

Construction and analysis of a spatially organized cortical network model

Johanna Senk

Schlüsseltechnologien / Key Technologies

Band / Volume 194

ISBN 978-3-95806-390-7

Forschungszentrum Jülich GmbH
Institute of Neuroscience and Medicine
Computational and Systems Neuroscience (INM-6 / IAS-6)

Construction and analysis of a spatially organized cortical network model

Johanna Senk

Bibliografische Information der Deutschen Nationalbibliothek.
Die Deutsche Nationalbibliothek verzeichnet diese Publikation in der
Deutschen Nationalbibliografie; detaillierte Bibliografische Daten
sind im Internet über <http://dnb.d-nb.de> abrufbar.

Herausgeber und Vertrieb: Forschungszentrum Jülich GmbH
Zentralbibliothek, Verlag
52425 Jülich
Tel.: +49 2461 61-5368
Fax: +49 2461 61-6103
zb-publikation@fz-juelich.de
www.fz-juelich.de/zb

Umschlaggestaltung: Grafische Medien, Forschungszentrum Jülich GmbH

Druck: Grafische Medien, Forschungszentrum Jülich GmbH

Copyright: Forschungszentrum Jülich 2019

Schriften des Forschungszentrums Jülich
Reihe Schlüsseltechnologien / Key Technologies, Band / Volume 194

D 82 (Diss., RWTH Aachen University, 2018)

ISSN 1866-1807
ISBN 978-3-95806-390-7

Vollständig frei verfügbar über das Publikationsportal des Forschungszentrums Jülich (JuSER)
unter www.fz-juelich.de/zb/openaccess.



This is an Open Access publication distributed under the terms of the [Creative Commons Attribution License 4.0](https://creativecommons.org/licenses/by/4.0/),
which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

AUTHOR'S LIST OF PUBLICATIONS

The work presented in this thesis is in parts based on the following publications:

A collaborative simulation-analysis workflow for computational neuroscience using HPC

Johanna Senk^{*}, Alper Yegenoglu^{*}, Olivier Amblet, Yury Brukau, Andrew Davison, David R. Lester, Anna Lührs, Pietro Quaglio, Vahid Rostami, Andrew Rowley, Bernd Schuller, Alan B. Stokes, Sacha J. van Albada, Daniel Zielasko, Markus Diesmann, Benjamin Weyers, Michael Denker, Sonja Grün

Published in Lecture Notes in Computer Science, Springer: [Senk et al. \(2017b\)](#).

^{*} These authors contributed equally.

Parts of this publication enter [Chapter 2](#).

Performance comparison of the digital neuromorphic hardware SpiNNaker and the neural network simulation software NEST for a full-scale cortical microcircuit model

Sacha J. van Albada, Andrew G. Rowley, Johanna Senk, Michael Hopkins, Maximilian Schmidt, Alan B. Stokes, David R. Lester, Markus Diesmann, Steve B. Furber
Published in Frontiers in Neuroscience (Neuromorphic Engineering): [van Albada et al. \(2018\)](#).

Published, preliminary results: [van Albada et al. \(2016, 2017\)](#).

Parts of this publication enter [Chapter 3](#).

VIOLA - A multi-purpose and web-based visualization tool for neuronal-network simulation output

Johanna Senk, Corto Carde, Espen Hagen, Torsten W. Kuhlen, Markus Diesmann, Benjamin Weyers

Published in Frontiers in Neuroinformatics and on arXiv: [Senk et al. \(2018a\)](#) and [Senk et al. \(2018b\)](#).

Parts of this publication enter [Chapter 4](#).

Conditions for traveling waves in spiking neural networks

Johanna Senk, Karolína Korvasová, Jannis Schuecker, Espen Hagen, Tom Tetzlaff, Markus Diesmann, Moritz Helias

Published on arXiv: [Senk et al. \(2018d\)](#).

Published, preliminary results: [Senk et al. \(2017a\)](#).

Parts of this publication enter [Chapter 5](#).

Reconciliation of weak pairwise spike-train correlations and highly coherent local field potentials across space

Johanna Senk, Espen Hagen, Sacha J. van Albada, Markus Diesmann

Published on arXiv: [Senk et al. \(2018c\)](#)

Published, preliminary results: [Senk et al. \(2015\)](#); [Hagen et al. \(2016c\)](#).

Parts of this publication enter Chapter 6.

Author contributions are indicated in the beginning of the respective chapters.

SUMMARY

CONSTRUCTION AND ANALYSIS OF A SPATIALLY ORGANIZED CORTICAL NETWORK MODEL

The cerebral cortex is one of the most intricate natural systems known, due to the multitude and heterogeneity of interconnected cells and its operation on different temporal and spatial scales. Cortical activity on the mesoscopic scale, spanning square millimeters to centimeters of cortical surface area, can be recorded with multi-electrode arrays implanted in neural tissue. Such extracellular recordings provide simultaneous access to population signals like local field potentials (LFPs) as well as spiking activity of individual neurons, and expose spatiotemporal activity patterns emerging parallel to the cortical surface. Local neuronal connectivity is specific with respect to cortical layers and neuron types, and the probability that two neighboring neurons are connected decays with distance. Computational models of neuronal networks with corresponding spatial extents and signal predictions are needed to infer the relationship between connectivity structure and experimentally recorded activity. This thesis focuses on the development of mesoscopic spatially organized cortical network models with cellular resolution. We develop a multi-layer network model with realistic neuron density and distance-dependent connectivity covering $4 \times 4 \text{ mm}^2$, a similar area as covered by multi-electrode arrays in use today. The model comprises excitatory and inhibitory spiking neuron populations in four cortical layers, integrates experimentally obtained connectivity data, and reproduces features of observed in-vivo spiking statistics. As a finding, the model reconciles the seemingly contradictory experimental observations of weakly correlated spike trains and strong, distance-dependent correlations of LFPs. Experimental data on the structure and dynamics of cortical networks are only known within certain margins of error and severe simplifications need to be made. Therefore, mean-field theory is required to explore regimes of biologically realistic activity and uncover mechanisms governing the network dynamics. This thesis advances the theory of spatially organized networks to a point where predictions are in quantitative agreement with direct simulations of spiking neuronal networks. Since conventional high-performance computing architectures are not optimized for accelerated and massively parallel neuroscientific simulations, the community develops dedicated neuromorphic hardware. We compare the performance of the software simulator NEST to the neuromorphic hardware system SpiNNaker in terms of accuracy, runtime, and energy consumption. To capture spatiotemporal patterns in simulated activity data, we design concepts for visual data analysis and provide the interactive web-based tool VIOLA (VIsualization Of Layer Activity) as a reference implementation. Moreover, we assess the integration of collaborative and interdisciplinary simulation-analysis workflows into online platforms. This thesis discusses the foundations of a model platform for the stepwise refinement of mesoscopic spatially structured network models and paves the way towards tackling further questions on the brain's function, learning, and diseases.

ZUSAMMENFASSUNG

KONSTRUKTION UND ANALYSE EINES RÄUMLICH ORGANISIERTEN KORTIKALEN NETZWERKMODELLS

Die Großhirnrinde (Cortex cerebri) ist eines der kompliziertesten natürlich vorkommenden Systeme aufgrund der Vielzahl und Verschiedenartigkeit vernetzter Zellen und der Informationsverarbeitung auf unterschiedlichen zeitlichen und räumlichen Skalen. Kortikale Aktivität auf der mesoskopischen Skala, die Quadratmillimeter bis -zentimeter an kortikaler Oberfläche umfasst, kann mit in Nervengewebe implantierten Multi-Elektroden-Arrays aufgezeichnet werden. Derartige Messungen im extrazellulären Raum ermöglichen die gleichzeitige Erfassung sowohl von Populationssignalen wie lokalen Feldpotentialen (LFPs) als auch von Aktionspotenzial-Sequenzen einzelner Neuronen. Dabei können räumliche und zeitliche Aktivitätsmuster beobachtet werden, die sich parallel zur Oberfläche des Cortex ausbilden. Die lokale neuronale Verbindungsstruktur weist spezifische Charakteristika abhängig von kortikaler Schicht und Neuronentyp auf. Außerdem fällt die Wahrscheinlichkeit, dass zwei benachbarte Neuronen verbunden sind, mit ihrem Abstand ab. Um die Beziehung zwischen Verbindungsstruktur und experimentell aufgezeichneter Aktivität aufzudecken, werden Computermodelle neuronaler Netze mit entsprechender räumlicher Ausdehnung und Vorhersagefähigkeit von Messgrößen benötigt. Die vorliegende Arbeit behandelt die Entwicklung mesoskopischer, räumlich organisierter kortikaler Netzwerkmodelle mit zellulärer Auflösung. Wir entwickeln ein mehrschichtiges Netzwerkmodell mit realistischer Neuronendichte und abstandsabhängiger Verbindungsstruktur, das mit $4 \times 4 \text{ mm}^2$ eine ähnliche Fläche wie heutzutage verwendete Multi-Elektroden-Arrays einnimmt. Dieses Modell besteht aus vier kortikalen Schichten mit Populationen exzitatorischer und inhibitorischer Einzelneuronen, integriert experimentell ermittelte Daten zur Verbindungsstruktur und kann Merkmale von in-vivo aufgezeichnete Aktivität statistisch reproduzieren. Ein Ergebnis ist, dass das Modell die scheinbar widersprüchlichen experimentellen Beobachtungen schwach korrelierter Aktionspotenzial-Sequenzen und starker, abstandsabhängiger Korrelationen in LFPs in Einklang bringt. Experimentelle Daten zu Struktur und Dynamik kortikaler Netze sind nur innerhalb gewisser Fehlertoleranzen bekannt, die vereinfachende Annahmen erzwingen. Daher ist es erforderlich Molekularfeldtheorie anzuwenden, um Bereiche biologisch realistischer Aktivität zu erforschen und Mechanismen aufzuklären, die die Netzwerkdynamik dominieren. Diese Arbeit entwickelt die Theorie von räumlich organisierten Netzen weiter und erreicht quantitative Übereinstimmungen mit direkten Simulationen gepulster Netze. Da konventionelle High-Performance-Computing-Architekturen nicht für beschleunigte und massiv parallele neurowissenschaftliche Simulationen optimiert sind, wird zweckbestimmte neuromorphe Hardware entwickelt. Wir vergleichen die Performance der Simulations-Software NEST mit der des neuromorphen Hardware-Systems SpiN-Naker in Bezug auf Genauigkeit, Laufzeit und Energieverbrauch. Um raumzeitliche Muster in simulierten Aktivitätsdaten zu erfassen, entwerfen wir Konzep-

te zur visuellen Datenanalyse und entwickeln das interaktive, webbasierte Programm VIOLA (VIualization Of Layer Activity) als Referenzimplementierung. Des Weiteren untersuchen wir die Integration kollaborativer, interdisziplinärer Simulations- und Analyse-Workflows in Online-Plattformen. Diese Dissertation erarbeitet Grundlagen für eine Modell-Plattform zur schrittweisen Weiterentwicklung mesoskopischer, räumlich strukturierter Netzwerkmodelle und bereitet den Weg, um weiteren Fragestellungen zu Gehirnfunktion, Lernen und Krankheiten nachzugehen.

CONTENTS

1	INTRODUCTION	1	
1.1	Spatial dependencies in the cortex on the mesoscale	3	
1.2	Modeling of neuronal networks	12	
1.3	Simulation technology	17	
1.4	Visualization	20	
1.5	Scope of the thesis	23	
2	A COLLABORATIVE SIMULATION-ANALYSIS WORKFLOW	27	
2.1	Introduction	28	
2.2	Workflow	30	
2.3	Discussion	36	
3	NEUROMORPHIC HARDWARE FOR FULL-SCALE NETWORKS	39	
3.1	Introduction	40	
3.2	Methods	43	
3.2.1	The leaky integrate-and-fire neuron model	43	
3.2.2	Single-neuron tests	44	
3.2.3	Cortical microcircuit model	45	
3.2.4	Implementation	46	
3.2.5	Performance benchmarks	47	
3.3	Results	49	
3.3.1	Conceptual separation of biological time and wall-clock time	49	
3.3.2	Steps toward implementation on SpiNNaker	50	
3.3.3	Comparison of single-neuron results between NEST and SpiNNaker	54	
3.3.4	Comparison of network results between NEST and SpiNNaker	57	
3.3.5	Performance	61	
3.4	Discussion	65	
4	INTERACTIVE WEB-BASED DATA VISUALIZATION WITH VIOLA	71	
4.1	Introduction	72	
4.2	Results	80	
4.2.1	Views of VIOLA	80	
4.2.2	VIOLA use case	85	
4.3	Methods	90	
4.3.1	Data formats	90	
4.3.2	Reference implementation	91	
4.3.3	2D and 3D view implementations	93	
4.3.4	Network description	94	
4.3.5	LFP predictions	96	
4.3.6	Software summary	99	
4.4	Discussion	102	
5	CONDITIONS FOR TRAVELING WAVES IN SPIKING NEURAL NETWORKS	107	
5.1	Introduction	108	
5.2	Results	112	

5.2.1	Linear stability analysis of a neural-field model	114
5.2.2	Conditions for spatial and temporal oscillations, and traveling waves	116
5.2.3	Application to a network with excitation and inhibition	118
5.2.4	Network simulation with nonlinear rate neurons	119
5.2.5	Linearization of spiking network model	120
5.2.6	Comparison of neural-field and spiking models	123
5.2.7	Validation by simulation of spiking neural network	127
5.3	Methods	130
5.3.1	Linear stability analysis	130
5.3.2	Properties of the spatial profile	132
5.3.3	Transition curves for reduced profile	132
5.3.4	Linearization of the spiking model	135
5.3.5	Model comparison	136
5.3.6	Fixing the working point	137
5.3.7	Physical units	138
5.3.8	Network structure and parameters	138
5.3.9	Software and implementation	139
5.4	Discussion	143
6	A MESOSCOPIC, MULTI-LAYER, FULL-SCALE CORTICAL NETWORK MODEL	147
6.1	Introduction	148
6.2	Methods	149
6.2.1	Point-neuron networks	149
6.2.2	Forward modeling of extracellular potentials	160
6.2.3	Statistical analysis	163
6.2.4	Software accessibility	168
6.3	Results	169
6.3.1	Upscaling of a cortical microcircuit model using lateral distance-dependent connectivity	169
6.3.2	Spiking activity of the point-neuron networks	176
6.3.3	LFP predictions	183
6.4	Discussion	193
7	DISCUSSION	199
7.1	Conclusions	199
7.2	Outlook	202

Bibliography	205
--------------	-----

INTRODUCTION

The cerebral cortex is the few millimeter thick and folded outer region of the cerebrum that consists of gray matter, a neural tissue with a high density of neuronal cell bodies. The gray matter covers the subjacent white matter which mainly consists of long-range connections, glia, and subcortical nuclei. The well-orchestrated interplay of a large number of connected neurons within the cortex underlies cognitive functions and complex behavior. Human cortex, for example, contains about 16 billion neurons (Herculano-Houzel, 2009) and each neuron processes input from on the order of 10^4 other neurons (Abeles, 1991; Binzegger et al., 2004). Despite its ubiquity, the brain, and in particular the cortex, is one of the most intricate natural systems known and scientists are still far from a full understanding of it as a whole. The intrinsic complexity of the brain not only originates from the multitude and heterogeneity of cells contributing to the activity, but also from its operation on different temporal and spatial scales; from submilliseconds for neuronal processing to months or years for learning, and from submicrometer structures of cell components to the unfolded area of cortex covering about a quarter of a square meter (Sejnowski et al., 2014).

Brain research or neuroscience (Albright et al., 2000; Kandel et al., 2000) is not only driven forward by the pure desire to expand the territory of knowledge, but likewise by practical questions from different disciplines: how to treat neurodegenerative diseases, or how to exploit neural principles for efficient information processing to advance technology? Studies on neural systems become increasingly complex and heterogeneous, that is, endeavors that require cooperation and coordination of expertise from different disciplines (Kandel and Squire, 2000; Kandel et al., 2013). This development challenges the typical scientific procedure to collect experimental data, and to build theoretical models for explaining the data and making predictions, and calls for comprehensive and reproducible workflows (Denker and Grün, 2015; Bouchard et al., 2016; Zehl et al., 2016; Rougier et al., 2017; Antolík and Davison, 2018). A fundamental challenge in the interdisciplinary field of Computational Neuroscience (Sterratt et al., 2011; Kass et al., 2018) is to uncover the relationship between anatomical structure and emergent activity in the brain (Bouchard et al., 2018). Experimental measurements provide data of the biological object under investigation. Mathematical and physical methods are needed to analyze the data and to develop theoretical models and hypotheses. Techniques from computer science facilitate numerical model simulations as well as efficient data processing.

This thesis focuses on the construction and analysis of spatially structured cortical network models. In contrast to experimental recordings, computational models offer control over all parameters and do not suffer from undersampling (Levina and Priesemann, 2017). The two extreme ends of modeling strategies are bottom-up and top-down. While top-down models aim at finding a network architecture to realize a specific function or to reproduce an observation, bottom-up models

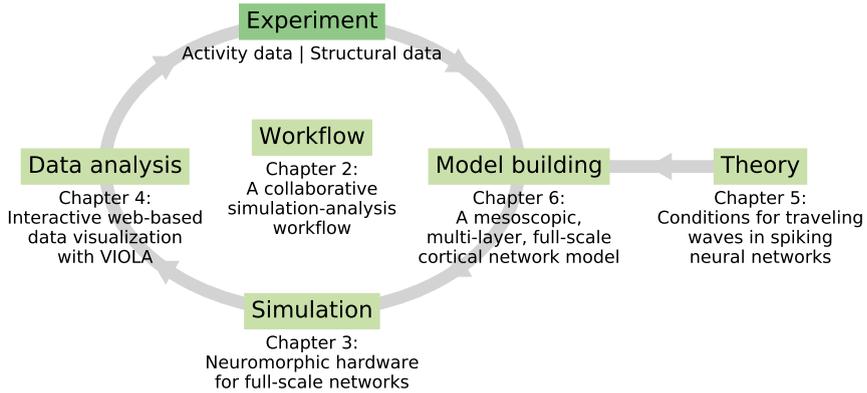


Figure 1.1: **Towards spatially structured cortical network models at cellular resolution.** Integrative loop for model development and structure of this thesis. **Experiment:** Structural data to be integrated into models; activity data for model validation. **Workflow (Chapter 2):** A collaborative simulation-analysis workflow using high-performance computing (HPC) for the comparison of data from NEST (software-based) and SpiNNaker (neuromorphic hardware) integrated into a web-based software platform. **Simulation (Chapter 3):** Performance comparison of the simulators NEST and SpiNNaker in terms of accuracy, runtime, and energy consumption for a full-scale cortical microcircuit model. **Data analysis (Chapter 4):** Visual inspection of spatially and temporally resolved neuronal activity data; visualization concepts and the corresponding interactive and web-based reference implementation VIOLA. **Theory (Chapter 5):** Analytical conditions for traveling waves in spiking networks derived via mean-field theory and parameter mapping to a neural-field model. **Model building (Chapter 6):** A mesoscopic, multi-layer, full-scale cortical network model with distance-dependent connectivity that accounts for spiking activity and local field potentials (LFPs).

are built up from elementary components. We here aim to approach the origin of spatiotemporal activity observed in cortex and we assume that such macroscopic collective dynamics rely on the microscopic interaction of connected single neurons. In line with the bottom-up modeling strategy, we integrate available structural data into models and analyze the emerging dynamics. However, our models incorporate also neuroscientific concepts such as cortical layers and simplified cell types. A model of a natural system requires both enough relevant detail to capture the phenomenon in study as well as an adequate level of abstraction achieved with simplifying assumptions to generate and test predictions. In the connectionist modeling approach (McLeod et al., 1998), complex dynamics can emerge from interconnected networks of simple units. Accordingly, the neuronal model networks considered here encompass large numbers of identical simple model neurons as basis elements (up to realistic neuron densities) and a complex connectivity structure.

The development of such large-scale neuronal network models can be split up into a number of individual components that together form the integrative loop shown in Figure 1.1. We identify the following components:

- The *Experiment* provides access to the natural system, the cerebral cortex, that we aim to understand. Structural data of the cortical connectivity are integrated into bottom-up models. Experimentally recorded activity data are compared to simulated model data for model validation.
- *Model building* deals with finding mathematical descriptions for cortical networks on the basis of structural experimental data. The level of description is a trade-off between tractability and predictive power.
- *Theory* contributes to model building. Predictions from reduced dynamics of the neuronal network as an interacting many-body system help to infer underlying mechanisms and to constrain model parameters.
- *Simulation* includes the efficient implementation of the model description and execution of simulation codes using high-performance computing (HPC) systems to generate model data.
- *Data analysis* covers processing of simulation output, in-depth statistical analysis, and visualization with the aim to interpret model data in relation to experimental activity data.
- The *Workflow* integrates all other components into a loop that facilitates the interoperability as well as the iterative refinement of individual components. Ideally, this loop enables repeatability of all steps and reproducibility of results.

These named components of model development require specific expertise and come with their own challenges. Section 1.1 establishes the biological detail of cortical structure and activity considered. Section 1.2 introduces mathematical and computational modeling of neuronal networks. Section 1.3 focuses on simulation technology for neuronal networks and distinguishes between software-based approaches and neuromorphic hardware. Section 1.4 deals with the visualization of experimental and model data. On the basis of these foundations and initial considerations, the scope of this thesis is finally specified in Section 1.5.

1.1 SPATIAL DEPENDENCIES IN THE CORTEX ON THE MESOSCALE

Horizontal and vertical organization of the cortex

The cerebral cortex has a distinct hierarchical spatial structure, both in parallel (horizontal) and perpendicular (vertical) to the cortical surface. To maximize the surface area, the cortex is folded and exhibits characteristic grooves called sulci and crests called gyri. The horizontal organization begins with a division into two hemispheres. Each hemisphere comprises four anatomically different lobes, illustrated in Figure 1.2A, with specialized functions: frontal (planning, movement control), parietal (somatic sensation), temporal (hearing) and occipital (vision). Further parcellation into brain areas is based on either the cytoarchitecture (Brodmann areas, Brodmann, 1909) or the functional involvement (Kandel et al., 2000). For example, Brodmann area 4 refers to the primary motor cortex (M1) in the precentral gyrus of the frontal lobe, and area 17 refers to the primary visual cortex (V1)

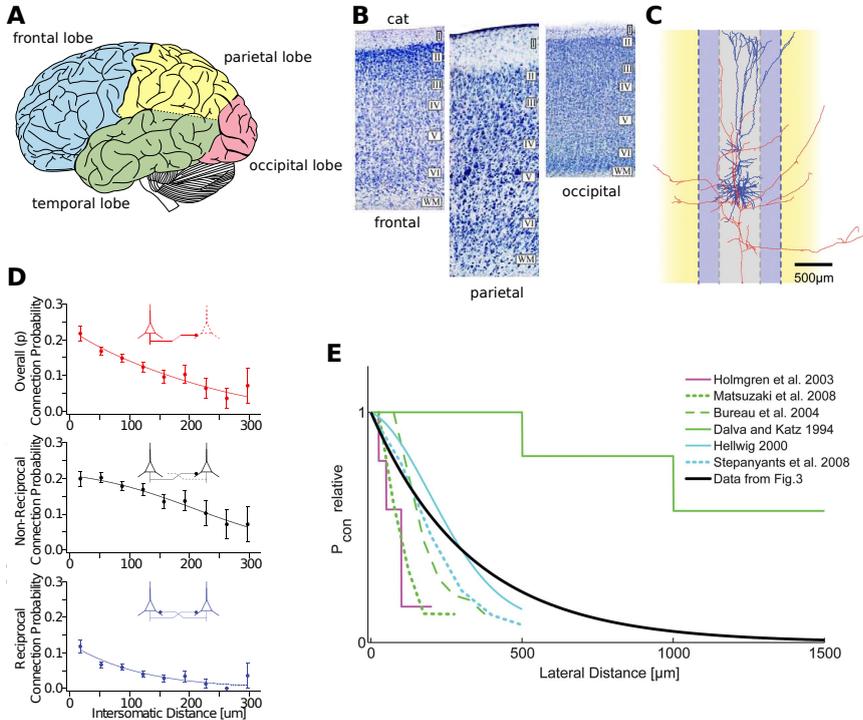


Figure 1.2: Structural organization of the cortex. **A** Cortical lobes; adapted from Gray (1918). **B** Cytoarchitectonic differences between frontal, parietal, and occipital areas of cat cortex; adapted from DeFelipe (2011, Figure 8). Photomicrographs from 100 μm thick Nissl-stained sections. Cortical layers 1–6 and white matter (WM) indicated. **C** Morphological reconstruction of a layer 5 pyramidal neuron from a rat somatosensory cortical slice of 300 μm thickness with dendritic (blue) and axonal (red) arborizations; adapted from Boucsein et al. (2011, Figure 4B). **D** Connection probability versus distance (error bars represent SEM); adapted from Perin et al. (2011, Figure 1 E–G) with permission. Simultaneous recordings from up to 12 thick-tufted layer 5 pyramidal neurons in rat somatosensory cortical slices of 300 μm thickness. **E** Connection probability versus distance; adapted from Boucsein et al. (2011, Figure 4A). Comparison of results from different studies with maximum connection probabilities normalized to unity. Methodological differences prevent direct numerical comparison. Data from Holmgren et al. (2003); Matsuzaki et al. (2008); Bureau et al. (2004); Dalva and Katz (1994); Hellwig (2000); Stepanyants et al. (2008), and black curve from Boucsein et al. (2011, Figure 3).

in the occipital lobe; note that cat V1 typically combines areas 17 and 18 (Payne and Peters, 2002). The surface area per hemisphere of area 17, also known as striate cortex, amounts to $\sim 1090\text{mm}^2$ in macaque monkey and $\sim 310\text{mm}^2$ in cat, as measured by van Essen and Maunsell (1980). This thesis mainly focuses on sensory cortices.

Vertically, the cortex is organized into layers, each with characteristic neuron compositions as already described by Ramon y Cajal (1899). The conventional division considers six layers with the topmost layer 1 containing only few neurons. Layers 2 and 3 are often combined into a layer 2/3, and deeper layers are sometimes subdivided even further (Abeles, 1991; Braitenberg and Schüz, 1998; Payne and Peters, 2002). Absolute and proportional layer thicknesses and neuron densities vary markedly between brain areas, as shown in Figure 1.2B for different sections of the cat cortex. Visual cortex, for instance, has a distinct layer 4, but this layer is hardly present in motor cortex; on the existence of layer 4, see Skoglund et al. (1997); García-Cabezas and Barbas (2014); Yamawaki et al. (2014); Barbas and García-Cabezas (2015). Motor cortex is one of the thickest regions with a thickness of more than 4 mm compared to other regions of about 1 mm in human cortex (Fischl and Dale, 2000). The neuron density in cortex is on the order of $10^4 - 10^5$ neurons/ mm^3 and at least twice as high in the thinner visual cortex compared to motor cortex (Beaulieu and Colonnier, 1983; Abeles, 1991; Herculano-Houzel et al., 2013). Within cortical areas, cylindrical volumes of roughly 1mm^2 surface area spanning all layers are associated with the concept of a cortical column (Mountcastle, 1957; Hubel and Wiesel, 1977). Columns do not have strict boundaries and are defined rather on the functional level as neurons within the same column are assumed to participate in the same function. An exception is barrel cortex in the rodent which exhibits clearly defined anatomical boundaries (see, for example, Petersen and Sakmann, 2001), but we here use the simplifying assumption that cortex is laterally shift-invariant. This thesis primarily considers the cortical mesoscopic scale spanning square millimeters to centimeters of cortical surface area (Muller et al., 2018), and we hereby refer to this scale as mesoscale.

Neurons and spiking activity

Having established the main cortical structures and relevant spatial scales, we next describe individual neurons and their interactions. The three main parts of a neuron are the cell body (soma), a branching tree-like structure around the soma serving as the main target for input (dendrites), and a nerve fiber transmitting output (axon), see for example Kandel et al. (2000). Neurons exchange signals across anatomical connections (synapses) that link an axon terminal of a presynaptic neuron to a dendrite of a postsynaptic neuron. The signals manifest in excursions of the membrane potential, that is an electrical potential difference due to relative differences in ion concentrations on each side of the resistive and capacitive cell membrane. A neuron continuously integrates its input currents and a sufficient depolarization causes the initiation of an action potential (spike) at the soma or axon hillock through activation of voltage-gated ion channels. This discrete electrical pulse propagates along the axon of the presynaptic neuron to a synapse causing the release of neurotransmitters into the synaptic cleft. The intake of neu-

rotransmitters by the postsynaptic neuron results again in a deflection of the membrane potential through opening of ion-specific channels that is propagated along the dendrites to the soma and is integrated. We here only consider interaction via such chemical synapses and not via electrical synapses, so-called gap junctions (Connors and Long, 2004). Further, we assume static synaptic connection strengths. That means that the effect of an incoming spike onto the postsynaptic membrane potential does not change over time, for example, dependent on pre- and postsynaptic neuronal activity (for synaptic plasticity, see Hebb, 1949; Abbott and Nelson, 2000; Zucker and Regehr, 2002).

Sequences of spikes (spike trains) are assumed to be the primary means of communication between neurons. Spiking activity in the neocortex is typically characterized by sparse and highly irregular spiking (Softky and Koch, 1993; Brunel and Hakim, 1999) with weak pairwise correlations (Ecker et al., 2010; Renart et al., 2010; Cohen and Kohn, 2011; Tetzlaff et al., 2012). Mochizuki et al. (2016) find that differences of the firing regularity across brain areas in one species are greater than the differences in similar areas across species. They observe that the spiking activity in visual areas is typically nearly random while motor areas exhibit most regular activity. The spike counts per unit time (firing rates) of cortical neurons range between < 1 and 30 spikes/s (Swadlow, 1988; de Kock and Sakmann, 2009). Firing rates of excitatory neurons are typically low in superficial and deep layers (layers 2/3 and 6, respectively), and higher in middle layers (layers 4 and 5), as reviewed in Potjans and Diesmann (2014, Table 6). Spike count correlations, however, appear to be more pronounced in superficial and deep layers than in middle layers (Smith et al., 2012). Although cortical neurons differ significantly in morphology, electrical properties, spiking behavior, and connectivity (Connors and Gutnick, 1990; Markram et al., 2015), Dale's principle (Eccles et al., 1954) simply classifies them based on the effect a spike occurrence has onto postsynaptically connected neurons: excitatory neurons increase the negative membrane potential of all postsynaptic neurons, thus facilitate their spiking, while inhibitory neurons decrease the potential further. Excitation and inhibition are typically balanced in cortical network activity (Haider et al., 2006; Dehghani et al., 2016). Although more than 80% of all cortical neurons are excitatory (Braitenberg and Schüz, 1998; Markram et al., 2015), many inhibitory neurons establish stronger connections and exhibit higher firing rates to maintain the balance.

We next introduce the most prominent neuron types in cortex (Abeles, 1991; Braitenberg and Schüz, 1998; Payne and Peters, 2002). Pyramidal neurons are the most frequently occurring type and they are excitatory. Their morphology, illustrated in Figure 1.2C, is characterized by a triangular-shaped soma, a thick apical dendrite extending towards the cortical surface, a dense tree of basal dendrites below the soma, and a long axon that often forms long-range connections. The apical dendrite grows more branches the closer it gets to the surface, while the basal dendrites spreads profusely in all directions within short range around the soma. The axon branches strongly close to the soma, but the main axon often leaves cortex and proceeds through the white matter. Pyramidal cells are encountered in all cortical layers with various deviations from the just described morphological features. Stellate cells are another common cell type in cortex, named after their star-shaped dendritic tree. Spiny stellate cells are excitatory and largely occur in

the middle region of cortex in primary sensory brain areas. For example, layer 4 of the primary visual cortex is termed granular layer due to its high density of spiny stellate cells; layers above and below are called supragranular and infragranular, respectively. Large spiny stellate cells can have axons extending into the white matter, while small cells remain local. A large number of subtypes of inhibitory neurons exists, but we here focus on inhibitory interneurons that release the neurotransmitter GABA. Such neurons are restricted to the gray matter, and they are commonly classified into basket and non-basket cells. The fast-spiking basket cell is the most-studied interneuron type with its compact shape due to a dense axonal arborization around the soma. Although our work only accounts for neuronal cells, nonneuronal cells like glial cells are also present in cortex. [Azevedo et al. \(2009\)](#) compute the ratio between nonneuronal and neuronal cells in the gray matter to 1.48 for the human cortex.

Distance-dependent connectivity

Already the composition of the cerebrum points towards the important role of neuronal connectivity: axons, dendrites and spines occupy more than 80% of the gray matter ([Schüz and Palm, 1989](#)), and the white matter is almost exclusively composed of myelinated axons. Cortical networks are hierarchically organized, from connections spanning the whole cortex to local microcircuits. An example for inter-area connectivity is the network linking the 32 visual or visual-association areas in the primate cortex that exhibits a distinct processing hierarchy ([Felleman and Van Essen, 1991](#); [Chaudhuri et al., 2015](#)). Within cortical areas, it is widely believed that canonical microcircuits are stereotypical circuits that are similar in connectivity and their ability to process information ([Douglas et al., 1989](#); [Douglas and Martin, 1991](#); [Thomson et al., 2002](#); [Binzegger et al., 2004](#); [Douglas and Martin, 2004](#); [Binzegger et al., 2009](#); [Harris and Shepherd, 2015](#)). A cortical microcircuit is loosely defined as the local circuit within a cortical column, thus connecting on the order of $10^4 - 10^5$ neurons ([Potjans and Diesmann, 2014](#); [Markram et al., 2015](#)). The characteristic flow of activity through the cortical layers of a microcircuit involves distinct feedforward and feedback projections. In primary visual cortex, for instance, thalamic afferents enter mainly in layer 4 (and partially also in layer 6), and excitatory signals are first transmitted upwards into superficial layers and then downwards into deeper layers ([Douglas and Martin, 2004](#); [Thomson, 2010](#); [Potjans and Diesmann, 2014](#)).

The connectivity within a microcircuit can be derived statistically in terms of connection probabilities between neurons of a certain type and within a certain layer, based on measurements obtained with different experimental techniques ([Thomson and Lamy, 2007](#)). To begin with, staining techniques render the otherwise translucent neuronal morphologies visible by injecting a tracer that is absorbed by the cells. Anterograde tracers (for example, biocytin, see [King et al., 1989](#)) are taken up by the soma and are transmitted along the axon to its terminals, while retrograde tracers (for example, horseradish peroxidase, see [Adams, 1977](#)) are taken up by the terminals and are transmitted along the axon back to the soma ([Abeles, 1991](#)). Staining and subsequent visualization with microscopes are often combined with a reconstruction of the morphologies. Statistical assumptions help further to

infer the connectivity from these reconstructed cells (see, for example, Binzegger et al., 2004; Hill et al., 2012; Helmstaedter, 2013; Berning et al., 2015; Kasthuri et al., 2015). Peter's rule (Peters and Feldman, 1976; Braitenberg and Schüz, 1998) is such a simplistic assumption, according to which the density of synaptic contacts is proportional to the overlap of the dendritic and axonal arborizations (Amirikian, 2005; Brown and Hestrin, 2009; Hill et al., 2012). Binzegger et al. (2004) compile a quantitative connectivity map of cat primary visual cortex based on an extension of Peter's rule.

While anatomical reconstructions already indicate whether connections are likely to exist or not, electrophysiological measurements can directly assess functional connectivity, that means, not only whether two neurons are connected but also how strong that connection is. The patch clamp technique (Sakmann and Neher, 1984) allows intracellular recordings by controlling either the current (current clamp) or the voltage (voltage clamp). Thomson et al. (2002) derive a comprehensive connectivity map with data from cat (visual cortex) and rat (somatosensory, motor and visual cortex) with pairwise recordings. They evoke spikes in a putative presynaptic neuron with an intracellular current-pulse injection and simultaneously record possible responses in putative postsynaptic neurons to compute connectivity ratios. Whole-cell recordings of the membrane potential of putative postsynaptic neurons can alternatively be combined with photostimulation of presynaptic neurons (Callaway and Katz, 1993; Boucsein et al., 2011; Schnepel et al., 2015). The technique used in these studies evokes presynaptic firing by applying short light pulses to uncage precursors of neurotransmitters like glutamate.

The probability for two neurons being connected depends on their relative spatial location, the morphologies of their dendrites and axons, their types, and their functional properties. Here, we are in particular interested in the distance dependency of connection probabilities. Experimental data indicates that the large part of local cortical connections are established within a distance of $\lesssim 500 \mu\text{m}$ from the source/target neuron (Voges et al., 2010), with probabilities that decay with distance according to a Gaussian or exponentially shaped profile (Hellwig, 2000; Budd and Kisvárdy, 2001; Boucsein et al., 2011; Packer and Yuste, 2011; Perin et al., 2011; Levy and Reyes, 2012; Reimann et al., 2013; Jiang et al., 2015; Schnepel et al., 2015; Reimann et al., 2017), as shown in Figure 1.2D–E. Excitation is typically assumed to be broader than inhibition (Budd and Kisvárdy, 2001; Binzegger et al., 2004; Buzás et al., 2006; Binzegger et al., 2007; Stepanyants et al., 2008, 2009; Ohana et al., 2012) but some interneuron types form longer-range lateral connections (McDonald and Burkhalter, 1993). The spatial reach of local connections depends not only on neuron types and layers (Reimann et al., 2017), but differs also between cortical areas in the same species (Kätzel et al., 2011a, Figure 4). Beside such local connectivity, in particular pyramidal neurons can develop longer horizontal axons spanning several millimeters. In cat and monkey visual cortex, these connections are often clustered or patchy and connect neurons with similar functions such as orientation tuning (Livingstone and Hubel, 1984; Gilbert and Wiesel, 1989; Bosking et al., 1997; Tanigawa et al., 2005; Buzás et al., 2006; Binzegger et al., 2007). The total number of incoming connections per neuron is on the order of 10^4 (Abeles, 1991). Stepanyants et al. (2009) find that a local cortical volume of radius $500 \mu\text{m}$, corresponding to the blue region in Figure 1.2C, comprises

only about 26% of all excitatory synapses, and that the remaining synapses originate from nonlocal neurons, while almost 90% of inhibitory synapses fall into the local volume.

Stimulation experiments reveal that connections are temporally delayed with a distance-dependent and a constant contribution. The distance-dependent contribution results from finite conduction speeds of action potentials along neuronal fiber. Long-range axons, connecting to distant or subcortical brain regions, are mostly covered by a myelin sheath that increases conduction speeds to typically several meters per second with maximum values of about 100 m/s (Debanne, 2004; Muller et al., 2018). Short-ranging axons within gray matter, as considered in this work, are typically not myelinated, and the conduction speed is with about 0.3 m/s much smaller (Andersen et al., 1978; Berg-Johnsen and Langmoen, 1992; Murakoshi et al., 1993; Kang et al., 1994; Lohmann and Rörig, 1994; Salin and Prince, 1996; Debanne, 2004; Muller et al., 2018). The constant contribution is estimated to 0.5 – 1 ms (Hirsch and Gilbert, 1991; Murakoshi et al., 1993; Kang et al., 1994) and depends on synaptic processing on the pre- and postsynaptic terminals, and includes neurotransmitter release, binding and spike initiation. Measured values for both conduction speed and synaptic delay rely on experimental conditions as a reduced ambient temperature results in smaller conduction speeds (Berg-Johnsen and Langmoen, 1992) and larger synaptic delays (Katz and Miledi, 1965; Sabatini and Regehr, 1996), reviewed by González-Burgos et al. (2000). Moreover, delays may change over time due to synaptic plasticity (Waldeck et al., 2000; Lin and Faber, 2002).

Cortical recordings of activity patterns in physical space and time

Cortical activity can exhibit inhomogeneities across the cortical surface with patterns evolving over time. Such spatiotemporal activity patterns considered here explicitly account for physical space, although the same term is often used for spike patterns of multiple neurons without any spatial information (see, for example, Abeles and Gerstein, 1988). Propagating or localized neuronal activity has been detected in many species and cortical areas (Wu et al., 2008; Muller and Destexhe, 2012; Sato et al., 2012; Muller et al., 2018). The spatiotemporal aspect of activity appears to be a general phenomenon of the cortex as various patterns are observed both in vitro (in brain slices) and in vivo (in the anesthetized or awake brain), and in the spontaneous activity of the resting state as well as under stimulation or during a task. Macroscopic traveling waves traverse across the whole cortex, while mesoscopic waves can be observed within single cortical areas (Muller et al., 2018). The two main experimental techniques that unveil spatiotemporal patterns emergent on the mesoscale with sufficient temporal and spatial resolution are optical imaging with voltage sensitive dyes (Chemla and Chavane, 2010) and electrophysiological recordings with multiple electrodes (for example, Maynard et al., 1997). This thesis concentrates on experiments performed with electrode arrays in vivo in the mammalian visual or motor cortex and activity propagating in parallel to the cortical surface.

Extracortical recording techniques like electroencephalography (EEG), magnetoencephalography (MEG) and electrocorticography (ECoG) can cover several corti-

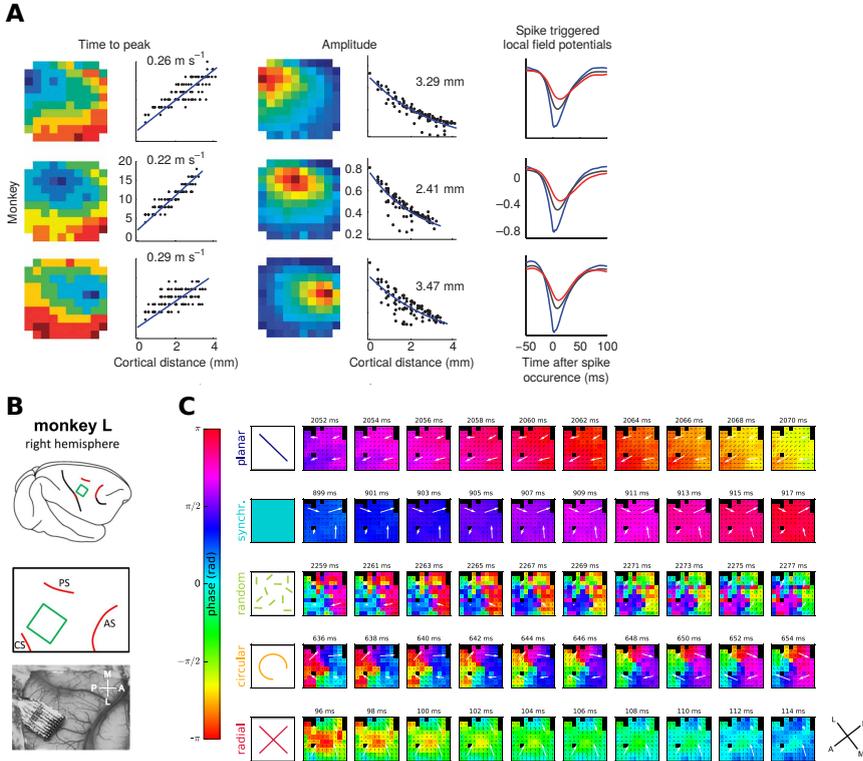


Figure 1.3: Spatiotemporal activity patterns recorded with multi-electrode arrays. **A** Spike-triggered traveling waves of local field potentials (LFPs) in monkey primary visual cortex (V1); adapted from [Nauhaus et al. \(2009, Figure 2\)](#) with permission from Springer Nature. Data from three example spike locations (rows). Left: Time to peak versus distance from triggering electrode, both as image and scatter plot, with estimated propagation speeds. Middle: Signal amplitudes, both as image (pseudo-color map: darkest red-saturated pixel at spike location) and scatter plot, with estimated space constants. Right: Average LFP waveforms at different distances r from spike location (blue: at spike location, gray: $r \in [400, 2,400] \mu\text{m}$, red: $r \in [2,400, 3,600] \mu\text{m}$). **B** Spatial location of the Utah multi-electrode array (green square) on the cortical surface in monkey L (motor cortex); adapted from [Denker et al. \(2018, Figure 1B–C\)](#). Anatomical features (red curves; CS: central sulcus, AS: arcuate sulcus, PS: precentral sulcus) estimated from photograph taken during surgery. **C** Automatically detected phase patterns of LFPs recorded from monkey L (see panel B); adapted from [Denker et al. \(2018, Figure 3B\)](#). Patterns: planar wave, synchronized, random, circular, and radial. Sequences (rows) show a total of 18 ms in steps of 2 ms. [Denker et al. \(2018\)](#) is licensed under a Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>).

cal areas. Chronic or acute implants of intracortical multi-electrode arrays can be used to record neuronal activity directly in neural tissue on the mesoscale (Buzsáki et al., 2012). The signal obtained with multi-electrode arrays has a submillisecond temporal resolution allowing the separation of the local field potential (LFP) as a mesoscopic population signal (low-frequency band, $\lesssim 100$ Hz) from the spiking activity of individual neurons (high-frequency band, $\gtrsim 100$ Hz), see Einevoll et al. (2013a). The spatial resolution is constrained by the electrode separation and typically on the order of $100 \mu\text{m}$ (Maynard et al., 1997). Despite the high neuron density of $10^4 - 10^5$ neurons/ mm^3 of cortical tissue (Herculano-Houzel, 2009), only a few neurons can be identified through spike sorting (Quiroga, 2007). Utah arrays (10×10 electrodes on $4 \times 4 \text{mm}^2$, Blackrock Microsystems, Salt Lake City, UT, USA), for instance, resolve little more than a hundred distinct neurons (Riehle et al., 2013).

Spatiotemporal patterns detected with multi-electrode arrays are categorized into different types of patterns and quantified. Muller and Destexhe (2012) distinguish globally propagating oscillatory waves from depolarizations spreading from a center of origin. For example, Rubino et al. (2006) show globally propagating LFP oscillations in the beta band ($10 - 45$ Hz), recorded in monkey motor cortex during an instructed-delay reaching task. Similar waves are observed in the human motor cortex (Takahashi et al., 2011). Nauhaus et al. (2009) demonstrate spreading depolarizations as LFP waves traveling outward from a spike occurrence in cat and monkey visual cortex (anesthetized), in spontaneous activity as well as under visual stimulation. Such spike-triggered LFP waves are shown in Figure 1.3A for the case of spontaneous activity in monkey cortex. The mechanism underlying these observations is debated (Ray and Maunsell, 2011b; Ian Nauhaus and Carandini, 2012). Muller et al. (2018, Figure 2) identify two different stimulus-evoked response types in the visual cortex: stationary bumps that are associated with relatively input-driven systems, and traveling waves that are associated with circuits dominated by recurrent activity. Riehle et al. (2013) observe movement-related spatiotemporal patterns in LFPs recorded in monkey motor cortex during a delayed reach-to-grasp task. Denker et al. (2018) later categorize such complex spatiotemporal phase patterns in LFP oscillations (beta band), here shown in Figure 1.3B. In visual cortex of anesthetized monkeys, Townsend et al. (2015) record complex patterns in low-frequency components of spontaneous LFP activity (delta band, $1 - 4$ Hz) and analyze them with methods from turbulence physics. Furthermore, Zanos et al. (2015) demonstrate that the execution of saccadic eye movements triggers traveling LFP waves in monkey visual cortex. The spatial reach of LFPs appears to be frequency dependent as shown by Łęski et al. (2013) and Dubey and Ray (2016). The latter study uses electrode stimulation in visual cortex of awake monkeys to expose that the spread is larger for intermediate frequencies (high-gamma band, $60 - 150$ Hz) than for lower or higher frequencies. Muller et al. (2018) summarize that traveling waves observed on the mesoscale exhibit propagation speeds of $0.1 - 0.8 \text{m/s}$ which is similar to axonal conduction speeds of unmyelinated long-range horizontal connections. Note, however, that such propagation also includes synaptic processing times and is affected by intrinsic dendritic filtering (Grinvald et al., 1994; Nauhaus et al., 2009; Takahashi et al., 2015; Zanos et al., 2015).

So far, we only discussed studies identifying spatiotemporal patterns in LFPs, a signal that includes the subthreshold activity of neuronal populations, but there is also evidence for spatial structure in spiking activity (Sato et al., 2012). Spike-count correlations decay with distance, measured in visual cortex of anesthetized monkeys (Smith and Kohn, 2008). Takahashi et al. (2015) find sequential firing that closely matches propagating LFP waves. With single-pulse intracortical microstimulations in motor cortex of awake monkeys, Hao et al. (2016) evoke activity spreading outward from a stimulating electrode, both in single-unit and multi-unit activity. In general, the detection of spatiotemporal patterns in population activity is more difficult on the level of spikes than from a continuous signal like the LFP. The two main reasons for that are sparse spiking (low firing rates) and sparse sampling of individual neurons inherent to the recording technique.

1.2 MODELING OF NEURONAL NETWORKS

Neural-field models (top-down)

The aim of top-down neuronal network models is to propose a network architecture that reproduces a desired observation, function, or computation of a biological neural system. Although guided by biophysical constraints, these models typically do not incorporate microscopic detail but rather provide phenomenological and coarse-grained descriptions using continuous variables (Schöner et al., 2015). Neural-field models (Wilson and Cowan, 1972a, 1973a; Amari, 1977), for example, describe the activity of neuronal populations on the meso- or macroscopic scale as a function of time and space with continuous nonlinear integro-differential equations and effective distance-dependent connectivity kernels. Such models provide explanations for diverse spatiotemporal activity patterns like traveling waves, wave fronts, bumps, pulses, and various periodic patterns (reviewed by Ermentrout, 1998; Coombes, 2005; Wyller et al., 2007a; Coombes, 2010; Bressloff, 2012; Coombes et al., 2014; Bressloff, 2014). The existence and uniqueness of activity patterns and their dependence on model parameters can be investigated either by constructing the patterns explicitly (Amari, 1977) or via bifurcation theory and linear stability analysis as described by Ermentrout (1998, Section 7) and Bressloff (2012, Sections 3-4). Neural fields also serve as models for hallucination patterns (Ermentrout and Cowan, 1979a; Bressloff et al., 2001), short-term memory (Laing et al., 2002), spatial working memory (Compte et al., 2000), movement preparation (Erlhagen and Schöner, 2002), saccade planning (Schneider and Erlhagen, 2002), visual perception (Erlhagen, 2003), motion perception (Giese, 2012), and binocular rivalry (Kilpatrick and Bressloff, 2010), to name some examples. While neural-field models already succeed in reproducing such observations, recent work strives to link these models to more realistic features like single-neuron properties (Hutt et al., 2015).

Spiking neuron and network models (bottom-up)

In contrast to top-down models, bottom-up models construct interconnected neuronal systems from its basic components to analyze the emerging dynamics. We

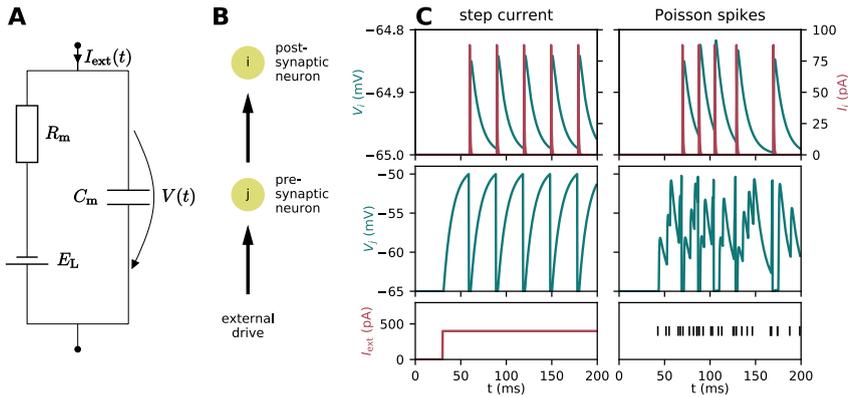


Figure 1.4: **The leaky integrate-and-fire (LIF) neuron model.** **A** RC circuit representing the subthreshold dynamics of the LIF neuron model (Equation 1.1). **B** Sketch of two connected neurons i (postsynaptic) and j (presynaptic) with external drive. **C** Simulated membrane potentials V and currents I for neurons i and j driven by a step current (left, $I_{\text{ext}} = 400$ pA) and by Poisson spike trains (right, rate $\nu_{\text{ext}} = I_{\text{ext}}/\tau_s/J_{\text{ext}}$ with a weight $J_{\text{ext}} = 4000$ pA), starting at $t = 30$ ms. Neuron and synapse parameters: $C_m = 250$ pF, $\tau_m = 10$ ms, $\tau_s = 0.5$ ms, $\tau_{\text{ref}} = 2$ ms, $E_L = V_{\text{reset}} = -65$ mV, $V_\theta = -50$ mV, $d = 1$ ms, $J = 87.8$ pA.

consider individual neurons as basic components and describe their dynamics and their interaction via spikes in a simplified yet biophysically inspired way. This level of detail is assumed to be enough to capture the main cortical processes (Hebb, 1949; Rumelhart et al., 1986; McLeod et al., 1998). A neuron is seen as a computational unit that receives and integrates input and generates output, yet there is a multitude of neuron models that differ in biological realism, capability to explain experimental observations, and analytical tractability (see Sterratt et al., 2011). The Hodgkin-Huxley model, for example, is a mathematical model that captures the time course of action potentials observed in voltage-clamp experiments including membrane currents and ion channels (Hodgkin and Huxley, 1952). Multicompartment models account for the spatial extent of the neuronal morphology with soma and axonal and dendritic structures (De Schutter and Van Geit, 2009a). More simplified neuron models are often used in large networks to study the effect of connectivity on the network dynamics, thereby disentangling this effects from biophysical detail on the single-cell level. In such models, the morphology is collapsed to a point, and the neuron dynamics are described by a single variable that often represents the neuron’s membrane potential. Rate-based neuron models, for instance, ignore spiking mechanisms and just convert their input by means of an activation (or characteristic) function to an output that can be interpreted as a firing rate or probability of firing; this level of description is comparable to phenomenological neural-field models. Binary neurons have only two possible states, on or off, and they can be modeled, for example, with a step function as activation function. Networks presented in this thesis are composed of leaky integrate-and-fire (LIF) neurons (Lapicque, 1907; Stein, 1967) with current-based synapses. LIF neuron models pose a compromise as they are analytically tractable but can still

reproduce in-vivo like spiking activity. A passive patch of the cell membrane is represented by an electrical RC circuit, illustrated in Figure 1.4A. The sub-threshold dynamics of a single LIF neuron i with exponentially decaying synaptic currents is described by a set of differential equations for the time evolution of the membrane potential V_i and its synaptic current I_i as

$$\begin{aligned}\tau_m \frac{dV_i}{dt} &= -(V_i - E_L) + R_m I_i(t), \\ \tau_s \frac{dI_i}{dt} &= -I_i + \tau_s \sum_j J_{ij} s_j(t - d_j),\end{aligned}\tag{1.1}$$

with the membrane time constant τ_m and the membrane resistance R_m (related by the membrane capacitance C_m as $\tau_m = R_m C_m$), the resistive leak reversal potential E_L , the synaptic time constant τ_s , the current jump J_{ij} due to a single spike from neuron j , incoming spike trains $s_j(t) = \sum_k \delta(t - t_k^j)$, and the transmission delay d_j . The current jump (also called connection strength or weight) is by convention positive for excitatory presynaptic neurons j and negative for inhibitory ones. Whenever V_i reaches the spike threshold V_θ , a spike is emitted and the membrane potential is reset to the spike reset potential V_{reset} and voltage-clamped for the refractory period τ_{ref} . Figure 1.4B illustrates an example of two synaptically connected neurons with the presynaptic neuron receiving external input. The neuronal responses are simulated in Figure 1.4C for the external drive being either a step current or a spike train drawn from a Poisson point process. The applied current input elicits regular spiking of the presynaptic neuron. Each spike is followed by a small deflection in the membrane current and potential of the postsynaptic neuron, that is, an excitatory postsynaptic current (EPSC) or potential (EPSP). Although with similar firing rate, the spiking activity caused by the noisy Poisson input is irregular; spikes occur only if the spike threshold is reached by sufficient summation of EPSPs.

Synaptic connectivity in spiking neuronal networks is typically described by statistical rules and realized at random. Such rules prescribe, for example, a connection probability, or a fixed number of incoming or outgoing connections per neuron, called in- or out-degree, respectively. In networks with spatial structure, connections can be restricted to nearest neighbors (Kriener et al., 2014b) or are established with a probability that falls off with distance between neurons (Mehring et al., 2003; Yger et al., 2011; Voges and Perrinet, 2012; Rosenbaum and Doiron, 2014; Keane and Gong, 2015; Schnepel et al., 2015; Pyle and Rosenbaum, 2017; Rosenbaum et al., 2017). Spatially organized networks also need to account for boundaries of the spatial domain modeled, and mostly assume periodic boundary conditions that are realized by a ring network in 1D (Roxin et al., 2005; Kriener et al., 2014b; Rosenbaum and Doiron, 2014) or with torus connectivity in 2D (Mehring et al., 2003; Yger et al., 2011; Voges and Perrinet, 2012; Rosenbaum and Doiron, 2014; Keane and Gong, 2015; Schnepel et al., 2015; Pyle and Rosenbaum, 2017; Rosenbaum et al., 2017).

Spiking neuronal networks without spatial structure already reproduce and explain features of in vivo activity like spike-train irregularity (Softky and Koch, 1993; van Vreeswijk and Sompolinsky, 1996; Amit and Brunel, 1997b; Shadlen

and Newsome, 1998), asynchronous firing (Ecker et al., 2010; Renart et al., 2010; Ostojic, 2014; Brunel, 2000), spike-train correlations (Gentet et al., 2010; Okun and Lampl, 2008; Helias et al., 2013), rate distributions across neurons (Griffith and Horn, 1966; Koch and Fuster, 1989; Roxin et al., 2011), and self-sustained activity (Ohbayashi et al., 2003; Kriener et al., 2014a). Brunel (2000), for example, finds that randomly and sparsely connected networks of excitatory and inhibitory neuron populations exhibit a rich space of distinct dynamical states classified by the synchrony and regularity of neuronal spiking activity. When inhibition dominates excitation, these models can reproduce asynchronous and irregular activity at low firing rates, as observed in vivo. Such a stable state is achieved for stronger inhibitory than excitatory connection weights while the number of excitatory connections equals four times the number of inhibitory connections, in line with experimental data (Braitenberg and Schüz, 1998). We refer to such models as “balanced random networks”. Models of this complexity can be treated analytically with mean-field models that reduce the single-neuron dynamics to an averaged population activity that resembles again rate-based equations (see, for example, Brunel, 2000; Deco et al., 2008; Tetzlaff et al., 2012; Helias et al., 2013; Schuecker et al., 2015). The approach introduced by Brunel (2000) accounts for both mean and variance of the input to a neuron in a self-consistent way. Balanced random networks serve as examples or building blocks for larger, more complex networks. Spatially structured network models that are constructed accordingly can be used to study the emergence of spatiotemporal activity patterns similarly to the approaches based on neural-field models (Mehring et al., 2003; Roxin et al., 2005; Yger et al., 2011; Voges and Perrinet, 2012; Kriener et al., 2014b; Rosenbaum and Doiron, 2014; Keane and Gong, 2015; Schnepel et al., 2015; Pyle and Rosenbaum, 2017; Rosenbaum et al., 2017).

Models that aim to represent the cortex on the mesoscale are typically diluted or encompass only one cortical layer. The comparability with biological cortical networks is thereby hampered since the preservation of neuronal dynamics including the correlation structure is severely limited upon downscaling (van Albada et al., 2015). Also, distance-dependent connectivity and emerging activity patterns rely on inter-neuron distances in comparison to the shape of connectivity kernels. In contrast, full-scale cortical network models aim to represent a patch of cortex at biologically realistic neuron and synapse density.

A full-scale cortical microcircuit model

The cortical microcircuit model by Potjans and Diesmann (2014), illustrated in Figure 1.5A, represents a 1 mm^2 patch of early sensory cortex at full density of neurons (approximately 80,000) and synapses (approximately 0.3 billion). The model is a minimal model with realistic in-degrees; larger networks are consequently less densely connected. Figure 1.5B shows spiking statistics of the simulated network activity. Biologically plausible firing rates of simple LIF model neurons (Equation 1.1) across four cortical layers with excitatory and inhibitory populations are explained by detailed connectivity derived from anatomical and electrophysiological data sets, that is, mainly from Binzegger et al. (2004) and Thomson and Lamy (2007). The model serves as a building block for a large multi-area model of the

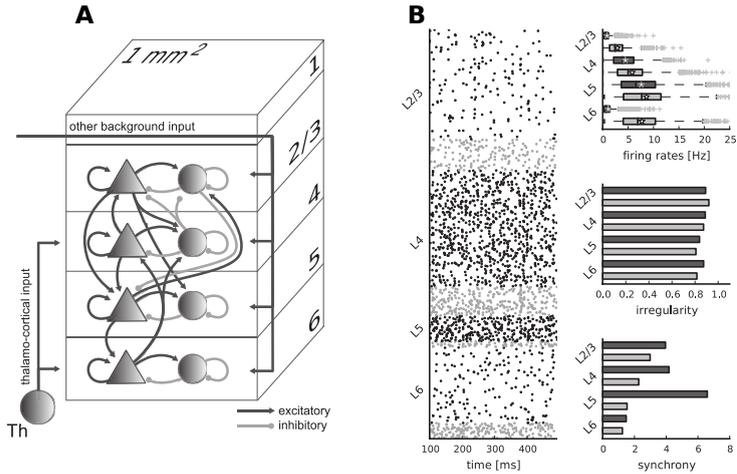


Figure 1.5: **A cortical microcircuit model with LIF neurons.** **A** Sketch of a microcircuit model; adapted from Potjans and Diesmann (2014, Figure 1) with permission from Oxford University Press. The model has four cortical layers with excitatory (triangles) and inhibitory (circles) populations of LIF neurons (same parameters as used in Figure 1.4C). **B** Simulated spiking activity of microcircuit model and statistics; adapted from Potjans and Diesmann (2014, Figure 6). Left: spike raster. Top right: boxplot of single-unit firing rates. Middle right: irregularity quantified by coefficient of variation of interspike intervals. Bottom right: synchrony quantified by variance of spike-count histogram divided by its mean.

visual system (Schuecker et al., 2017; Schmidt et al., 2018), and it is used as reference implementation for a forward-model based scheme to predict LFPs from spiking activity (Hagen et al., 2016a). Further studies analyze the role of different neuron populations in the model network with respect to attentional information (Wagatsuma et al., 2011), computational properties related to the input-output relationship (Cain et al., 2016), and network oscillations (Bos et al., 2016). In addition, analytical treatment of the model is feasible, for example by reducing the spiking dynamics via mean-field theory (Bos et al., 2016; Hahne et al., 2017; Schuecker et al., 2017) or a population density approach (Cain et al., 2016) to population rate dynamics, or by deriving stochastic population equations for the population activity (Schwalger et al., 2017). The source code of the model is publicly available at Open Source Brain¹ and is provided as an example for the simulator NEST².

1.3 SIMULATION TECHNOLOGY

In addition to experimental and theoretical approaches, numerical simulations play an important role for the investigation of the relationship between the structure of neuronal networks and emerging dynamics. Unlike earlier proof-of-concept simulations implemented from scratch, the emerging need for efficient and repeatable large-scale simulations calls for novel simulation technologies for the community (Bouchard et al., 2016; Eglén et al., 2017; Bouchard et al., 2018). The development of such technologies is driven by the requirements of neuroscientists and aims for both, performant exploitation of available hardware and ease of use for neuroscientific research. Code maintenance, continuous advancements and extensive testing by many scientists support the correctness of simulation results as well as the reproducibility of simulation studies (Rougier et al., 2017). Meanwhile, a variety of simulator systems are established with different design goals but also overlapping scopes, for example, in the biological detail represented. Many simulators have in common that they provide a Python interface (Muller et al., 2015) to ease the definition of simulation while hiding simulator infrastructure from the neuroscientist. Some simulator back ends can even be controlled with the same executable network model description if implemented in the simulator-independent language PyNN (Davison et al., 2009, 2010). We here primarily distinguish between software-based simulators (Brette et al., 2007) like NEST (NEural Simulation Tool, Gewaltig and Diesmann, 2007), and neuromorphic hardware (Furber, 2016a; Nawrocki et al., 2016) like SpiNNaker (SPIking Neural Network Architecture, Furber et al., 2013a). While software-based simulators focus on the optimal usage of conventional general-purpose hardware, the term neuromorphic hardware refers to dedicated hardware inspired by neural systems. Finally, we introduce a forward modeling scheme to predict LFPs from the simulated spiking activity of point-neuron networks.

¹ <http://opensourcebrain.org/projects/potjansdiesmann2014>

² <http://nest-simulator.org>

Software-based simulators

As it is a common hypothesis that essential aspects of brain function rest on neuronal interaction on the level of spikes and rates, a number of simulators focuses on the the description of individual neurons and their connectivity (Jordan et al., 2018). While other simulators even resolve fine structures like molecules (STEPS, Wils and De Schutter, 2009) or operate rather on the level of large groups of neurons or brain areas (MIIND, de Kamps et al., 2008; The Virtual Brain (TVB), Sanz Leon et al., 2013; Nengo, Bekolay et al., 2013), we here also consider this intermediate scale and focus on interactions via spikes. GENESIS (Bower and Beeman, 2007) and NEURON (Carnevale and Hines, 2006) allow for biological detail of individual neurons including complex morphologies (multicompartment neurons), whereas Brian (Goodman and Brette, 2013) and NEST (Gewaltig and Diesmann, 2007) rather aim for large networks of simple point-neurons, that is, for cellular resolution. These simulator systems have evolved over three decades of development (Brette et al., 2007) and advance further as network models are becoming ever more complex and general-purpose hardware progresses (Lytton et al., 2016; Jordan et al., 2018). Commonly used simulator code is mainly open source (Gleeson et al., 2017) and implementations of a large number of neuron and network models are hosted at publicly available resources (for example, ModelDB³, Hines et al., 2004, and Open Source Brain⁴, Gleeson et al., 2015).

The simulation code NEST⁵ (Gewaltig and Diesmann, 2007) is developed and maintained in a collaborative fashion by the NEST initiative⁶ under the GNU General Public License. Simulations are defined in terms of neuronal populations and connections employing the built-in scripting language SLI, the Python interface PyNEST (Eppler et al., 2009; Zaytsev and Morrison, 2014), or PyNN (Davison et al., 2009). The actual simulation kernel that executes network construction and state propagation of neuron dynamics is implemented in C++ to enable high computational performance. To optimize performance on available hardware, the simulation kernel is continually advanced with new data structures, algorithms and communication schemes (Morrison et al., 2005; Helias et al., 2012; Kunkel et al., 2014; Jordan et al., 2018). NEST uses a hybrid parallelization strategy combining MPI across compute nodes and multi-threading with OpenMP within each MPI process (Plesser et al., 2007). This strategy facilitates the simulation of large networks with neuron and synapse counts comparable to 10% of the human cortex (around 10^9 neurons and 10^{13} synapses) on contemporary supercomputers (RIKEN BSI, 2013; Kunkel et al., 2014; Jordan et al., 2018); smaller simulations can be run on laptops, workstations or moderately-sized clusters. NEST provides implementations of a number of common neuron and synapse models that are typically described by a system of a few delayed differential equations to be solved exactly, if possible, or numerically, otherwise (Rotter and Diesmann, 1999; Plesser and Diesmann, 2009). New models can be implemented either directly in C++ or in the domain-specific model description language NESTML (Plotnikov et al.,

³ <https://senselab.med.yale.edu/ModelDB>

⁴ <http://opensourcebrain.org>

⁵ <http://nest-simulator.org>

⁶ <http://nest-initiative.org>

2016). The Topology module allows the placement of neurons in 2D or 3D space and provides distance-dependent connectivity rules. Apart from spiking interactions via delayed static synapses, other connection schemes are also implemented in NEST such as plasticity (neuro-modulated, see Potjans et al., 2010, and structural, see Diaz-Pier et al., 2016), gap junctions (Hahne et al., 2016), and continuous (rate-based) interactions (Hahne et al., 2017).

Neuromorphic hardware

The term “neuromorphic hardware”, coined by Mead around 1990 (Mead, 1990), is an umbrella term for multi-disciplinary research that aims for novel brain-inspired hardware architectures and breaks with principles of the conventional von Neumann hardware (von Neumann, 1993). Expectations for such new systems are high and include massively parallel operation, asynchronous updates, low-power consumption, fault-tolerance, dimensionality reduction, speed-up, and so forth (Monroe, 2014; Ahmed and Sujatha, 2015; Furber, 2016b,a; Nawrocki et al., 2016; Vanarse et al., 2016). Possible fields of application are neuroscience research but also artificial intelligence, robotics, and machine learning in general. To date, a large number of different design approaches are followed that range from emulating single neurons or synapses to large-scale neuromorphic computing systems, use digital, analog or combined circuits, employ CMOS, organic electronic, memristive or hybrid devices, and tackle specific challenges posed by emerging limitations of current general-purpose hardware. A direct comparison between all these emerging systems is difficult due to different design goals coming with trade-offs between desirable objectives and missing standard benchmarks (Nawrocki et al., 2016; Furber, 2016a). To quantify energy consumption, a common measure is the energy per synaptic event, that is, the energy needed to transmit a spike through a synapse.

We here bring the attention to large-scale neuromorphic computing systems intended to simulate large neuronal networks with biological realism that aim to advance neuroscience research. Simulating a network with neuron and synapse counts comparable to 10% of the human cortex on the JUQUEEN supercomputer located in Jülich, Germany, with NEST demonstrates the need for both accelerated simulation times and a reduced power consumption: 1 s of biological times takes 40 min of wall-clock time and consumes about 2 MW of power (Kunkel et al., 2014), while the the brain’s power consumption is ~ 10 W (Herculano-Houzel, 2011). Such results render the investigation of slow processes like plasticity for similar network sizes unfeasible on contemporary HPC systems. Four large-scale systems are currently under consideration (Furber, 2016a): the IBM TrueNorth chip (Hsu, 2014; Merolla et al., 2014) is developed for real-time cognitive applications with digital neuron models; the Stanford Neurogrid (Benjamin et al., 2014) employs real-time sub-threshold analogue neural circuits; the Heidelberg Brain-ScaleS system (Schemmel et al., 2010; Scholze et al., 2012) uses waver-scale above threshold analogue neural circuits with a speed-up of 10^4 compared to biological time; the Manchester SpiNNaker (Furber et al., 2014) system is a real-time digital many-core system implementing neuron and synapse models in software running on small embedded processors.

SpiNNaker resembles a conventional supercomputer as it is a scalable, massively parallel, highly configurable digital system. It is in particular designed for the energy-efficient real-time simulation of large neural networks. To meet these requirements, SpiNNaker integrates a large number of ARM processor cores (typically used in mobile and embedded applications) with a communication infrastructure optimized for sending many small data packages (such as single spikes) to many destinations. Customized processor chips (Painkras et al., 2013) with on-board multicast packet routers are combined with standard memory chips to minimize the distances over which frequently accessed data is moved. Communication relies on routing tables using address-event-representation and data packets are exchanged via globally-asynchronous-locally-synchronous (Painkras et al., 2012) operations. Simulations are controlled with high-level neural description languages, for example, PyNN (Davison et al., 2009) or Nengo (Bekolay et al., 2013), on a host machine that is needed to map networks onto the SpiNNaker system and to retrieve simulation results. The largest simulations performed with SpiNNaker so far comprise up to 50 million (Sharp et al., 2014; Knight et al., 2016) and 86 million synapses (Stromatias et al., 2013).

Forward modeling of local field potentials

The main result of simulations implemented in software like NEST or run on neuromorphic hardware like SpiNNaker is the spiking activity of a large number of interconnected point neurons. Hagen et al. (2016a) propose a method for the forward-model prediction of extracellular potentials, that is, local field potentials, from such spiking activity. Forward modeling builds on the hypothesis that the complete knowledge of the spiking activity of all neurons is sufficient to predict other measures like LFPs (Einevoll et al., 2013a). In their framework, each point neuron has its equivalent multicompartment neuron, and spike trains are mapped to synapse activations that translate into a distribution of transmembrane currents (see, for example, De Schutter and Van Geit, 2009b). Volume conduction theory (Nunez and Srinivasan, 2006; Einevoll et al., 2013b) then relates these current sources to electric potentials in 3D space. Multicompartment neurons are modeled in NEURON⁷ (Carnevale and Hines, 2006; Hines et al., 2009) and forward-model computations rely on the tool LFPy⁸ (Lindén et al., 2014; Hagen et al., 2018).

1.4 VISUALIZATION

Visualization is the communication of information with graphical representations (Ward et al., 2010) that can enhance or even replace textual information. Graphical representations promote an efficient and effective understanding of relationships and processes. They are therefore an invaluable tool to gain and impart knowledge in scientific workflows, and can support all steps from the acquisition of raw data, across data processing and analysis, to the presentation of results. The optimization problem of visually encoding complex data to best convey information is non-trivial and constrained not only by features of the data itself, but also by

⁷ <https://neuron.yale.edu>

⁸ <https://lfp.readthedocs.io>

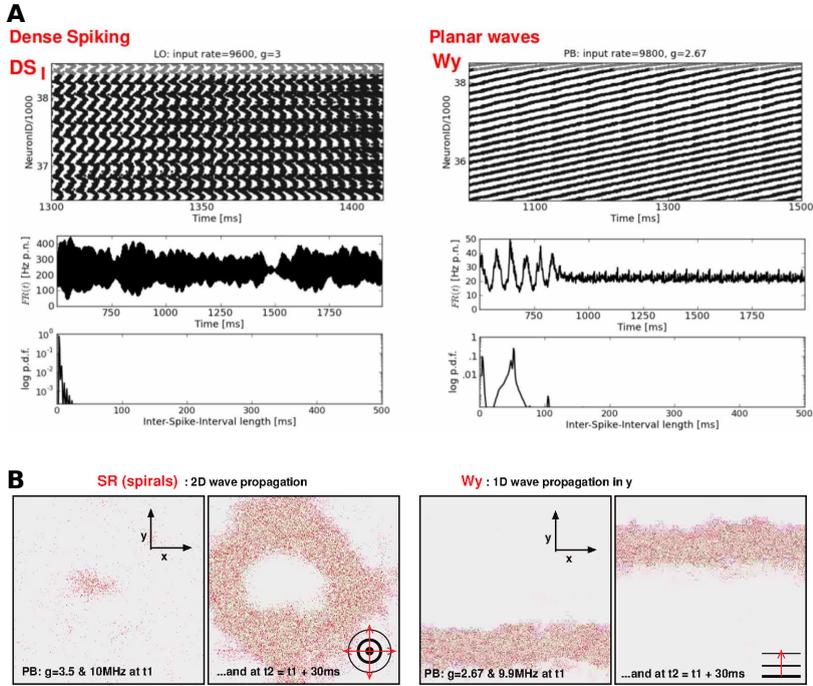


Figure 1.6: **Wave propagation in simulated spiking activity.** Spatially structured networks modeled as 2D cortical sheets with distance-dependent connection probabilities and conduction delays. **A** Spike raster (top, sorted according to neuron positions), average firing rate per neuron over time (middle), and distribution of interspike intervals (bottom) for two distinct network states; adapted from Voges and Perrinet (2012, Figure 4). Left: dense spiking, characterized by exceptionally high firing rates. Right: plane waves propagating in y -direction. **B** Two 2D cortical sheets showing dots at spike positions for two different time steps each; adapted from Voges and Perrinet (2012, Figure 6A–B). Left: circular spherical wave propagation. Right: plane wave of spikes propagating in y -direction (as in panel A, right).

available graphical representations or the choice of colors (Fairchild, 2013; Spence, 2014). For explorative and interactive visual analysis, as considered in this thesis, theoretical principles provide guidelines to approach multi-dimensional data sets. One example is the ‘information-seeking mantra’ introduced by Shneiderman (1996): overview first, zoom and filter, then details-on-demand. Another example is the concept of ‘coordinated multiple views’ established by Wang Baldonado et al. (2000) that suggests to study a complex conceptual entity with complementary visual representations in a coordinated way. The implementation of such visualizations can be realized with basic or advanced plotting functions of standalone programs (see Gnuplot⁹, Janert, 2010) or libraries associated with common programming languages (see Matplotlib for Python, Devert, 2014). An alternative is dedicated visualization tools that are optimized for a specific type of data (for example, VisNEST (Nowke et al., 2013, 2015) for the spiking activity of multi-area network models with the example model by Schmidt et al., 2018).

The data type most relevant for this thesis comprises spatially and temporally resolved neuronal activity resulting either from a network model simulation or an experimental recording. Such data may be multi-modal, meaning that the data can contain, for instance, discrete spike events of individual neurons, but also continuous population signals like LFPs as recorded with a multi-electrode array (Riehle et al., 2013). The data is high-dimensional as it accounts for up to three-dimensional spatial information of individual neurons or electrodes, and a temporal component. Activity can evolve on multiple scales: from single-neuron dynamics to slowly propagating features. In addition, the data may exhibit a high variability across space and time with locally confined activity and transient phenomena, for example, related to stimulation or behavior (Denker et al., 2018). Spike data is commonly visualized in spike raster diagrams or dot displays (Abeles, 1982) where rows show spike trains of different neurons and spike times are marked by dots; compare the spike raster shown in Figure 1.5A. Spatial resolution can be achieved by sorting neurons according to their location in physical space. In this way, Voges and Perrinet (2012) resolve different spatiotemporal activity patterns from a simulated spatially structured network such as dense spiking and planar waves shown in Figure 1.6A. Another common way to visualize spatially resolved spike data is to show sequences of temporal snapshots of the cortical sheet modeled as in Figure 1.6B. This method has proven beneficial also for analog data like the LFP phase patterns shown in Figure 1.3C; other examples are Mehring et al. (2003, Figure 5), Yger et al. (2011, Figures 2 and 12), and Keane and Gong (2015, Figure 1). Such routinely used graphic representation types are challenged by more and more complex data (Antolík and Davison, 2018). For example, advances in experimental recording techniques allow the simultaneous recording of an increasing number of neurons calling for novel analysis methods (Stevenson and Kording, 2011). Likewise, neuronal network models may comprise up to a realistic density of neurons and are organized in cortical layers (Potjans and Diesmann, 2014) or encompass even more biophysical detail (Markram et al., 2015).

⁹ <http://www.gnuplot.info>

1.5 SCOPE OF THE THESIS

This thesis intends to advance the development of spatially structured cortical network models in the framework of the integrative loop shown in Figure 1.1. For this purpose, we present five complementary studies that all encompass multiple aspects of the entire loop, but focus in particular on the progress of one component each: Workflow (Chapter 2), Simulation (Chapter 3), Data analysis (Chapter 4), Theory (Chapter 5), and Model building (Chapter 6). The common aim of these studies is to construct and analyze mesoscopic cortical network models accounting for similar spatial extents and signals as recorded using multi-electrode arrays like the Utah array with 10×10 electrodes on $4 \times 4 \text{ mm}^2$ (Section 1.1). Networks are composed of excitatory and inhibitory spiking point neurons, based on the concept of balanced random networks or the cortical microcircuit model by [Potjans and Diesmann \(2014\)](#) as a multi-layer extension (Section 1.2). We target full-scale neuronal networks, meaning, networks with realistic neuron and synapse densities, and assume specific experimentally derived connectivity that depends on neuron types, layers, and inter-neuron distances. Network simulations are implemented in NEST, and we compare the simulation performance to the neuromorphic hardware system SpiNNaker (Section 1.3). To capture spatiotemporal activity patterns eventually produced by such models, we further develop concepts for visual data analysis (Section 1.4). We also employ mean-field theory to reduce the dynamics of spiking neuronal networks to averaged population dynamics or neural fields to investigate the origin of observed spatiotemporal activity patterns (Section 1.2). In addition, we assess the integration into web-based platforms of collaborative, interdisciplinary workflows with similar complexity as illustrated by the loop in Figure 1.1. The following expands on the concrete objectives of the individual studies included.

Workflow

Chapter 2 addresses the challenges and demands for workflows for the acquisition and analysis of data in the interdisciplinary domain of Computational Neuroscience.

These workflows grow in complexity and heterogeneity, and are increasingly becoming large collaborative efforts that require diverse data generation and processing steps using different tools and HPC systems. We argue for the need of software platforms integrating HPC systems that allow scientists to construct, comprehend and execute workflows. As a use case we present a concrete implementation of such a complex workflow, covering diverse topics including HPC-based simulation using the NEST software, access to the SpiNNaker neuromorphic hardware platform, complex data analysis using the Elephant library¹⁰, and interactive visualization methods for facilitating further analysis. This workflow includes simulations of the full-scale cortical microcircuit model ([Potjans and Diesmann, 2014](#), see Section 1.2) both with NEST and SpiNNaker as simulator back end and a first comparison of the results. All tools used are embedded into a web-based software platform under development by the Human Brain Project, called the Collabora-

¹⁰ <http://elephant.readthedocs.io>

tory¹¹. On the basis of this prototype implementation, we discuss the state of the art and future challenges in constructing large, collaborative workflows with access to HPC resources.

Simulation

Chapter 3 provides a quantitative performance comparison of NEST and SpiNNaker simulations in terms of accuracy, time-to-solution, and energy-to-solution.

The study uses the full-density microcircuit model by [Potjans and Diesmann \(2014\)](#) as an example. The digital neuromorphic hardware SpiNNaker has been developed with the aim of enabling large-scale neural network simulations in real time and with low power consumption ([Furber et al., 2014](#)). Real-time performance is achieved with 1 ms integration time steps, and thus applies to neural networks for which faster time scales of the dynamics can be neglected. By slowing down the simulation, shorter integration time steps and hence faster time scales, which are often biologically relevant, can be incorporated. We here describe the first full-scale simulations of a cortical microcircuit with biological time scales on SpiNNaker. With approximately 80,000 neurons and 0.3 billion synapses, this model is among the largest networks simulated on SpiNNaker to date. The scale-up is enabled by recent developments in the SpiNNaker software stack that allow simulations to be spread across multiple boards. A comparison with simulations using the NEST software on a HPC cluster shows that both simulators can reach a similar accuracy, despite the fixed-point arithmetic of SpiNNaker. We thus demonstrate the usability of SpiNNaker for computational neuroscience applications with biological time scales and large network size.

Data analysis

Chapter 4 deals with the challenge to inspect spatially and temporally resolved activity data recorded from layered, spatially structured networks that resemble recordings with multi-electrode arrays.

To cover the surface area captured by today's experimental techniques and to achieve sufficient self-consistency, such models can contain millions of neurons. The interpretation of the resulting stream of multi-modal and multi-dimensional simulation data calls for integrating interactive visualization steps into existing simulation-analysis workflows. Here, we present a set of interactive visualization concepts, called views, for the visual analysis of activity data in topological network models, and a corresponding reference implementation VIOLA (Visualization Of Layer Activity). The software is a lightweight, open-source, web-based and platform-independent application combining and adapting modern interactive visualization paradigms, such as coordinated multiple views ([Wang Baldonado et al., 2000](#)), for massively parallel neurophysiological data ([Stevenson and Kording, 2011](#)). For a use-case demonstration we consider spiking activity data of a two-population, layered point-neuron network model incorporating distance-dependent connectivity subject to a spatially confined excitation originating from

¹¹ <http://collab.humanbrainproject.eu>

an external population. With the multiple coordinated views, an explorative and qualitative assessment of the spatiotemporal features of neuronal activity can be performed upfront of a detailed quantitative data analysis of specific aspects of the data.

Theory

Chapter 5 investigates the origin of spatiotemporal patterns like traveling waves in neural activity.

Although such patterns are frequently observed in experimental recordings, the mechanisms underlying their generation are largely unknown. Previous studies have investigated the existence and uniqueness of different types of waves or bumps of activity using neural-field models, phenomenological coarse-grained descriptions of neural-network dynamics (Coombes, 2005; Bressloff, 2012). But it remains unclear how these insights can be transferred to more biologically realistic networks of spiking neurons, where individual neurons fire irregularly (Softky and Koch, 1993). Here, we employ mean-field theory (Brunel and Hakim, 1999; Schuecker et al., 2015) to reduce a microscopic model of LIF neurons with distance-dependent connectivity to an effective neural-field model. In contrast to existing phenomenological descriptions, the dynamics in this neural-field model depends on the mean and the variance in the synaptic input, both determining the amplitude and the temporal structure of the resulting effective coupling kernel. For the neural-field model we derive conditions for the existence of spatial and temporal oscillations and periodic traveling waves using linear stability analysis; see Ermentrout (1998, Section 7) and Bressloff (2012, Sections 3–4). We demonstrate quantitative agreement between predictions of the analytically tractable neural-field model and numerical simulations of both networks of nonlinear rate-based units and networks of LIF neurons.

Model building

Chapter 6 finally presents the construction and analysis of a mesoscopic, multi-layer, full-scale, cortical network model.

The interpretation of data recorded with multi-electrode arrays, that is, extracellular potentials and spiking activity, calls for multiscale computational models with corresponding spatial dimensions and signal predictions. Such models can then facilitate the search of candidate mechanisms underlying experimentally observed spatiotemporal activity patterns in cortex. We here extend the microcircuit model by Potjans and Diesmann (2014) and the corresponding forward model for LFP predictions by Hagen et al. (2016a) from an area of 1 mm^2 to $4 \times 4 \text{ mm}^2$. The upscaling preserves the densities of neurons and local synapses, and introduces distance-dependent connection probabilities and conduction delays. As detailed experimental data on distance-dependent connectivity is partially lacking, we address this uncertainty in model parameters by testing different parameter combinations within biologically plausible bounds. Based on model predictions of spiking activity and LFPs, we find that the upscaling procedure preserves the overall spiking statistics of the original model and reproduces asynchronous irregular spiking

(Softky and Koch, 1993) across populations and weak pairwise spike-train correlations (Ecker et al., 2010) experimentally observed in sensory cortex. In contrast with the weak spike-train correlations, the correlation of LFP signals is strong and distance-dependent, compatible with experimental observations.

The last Chapter 7 summarizes the results of the individual studies, concludes on the combined advancement, and provides an outlook for the concrete network model and the more general problem of digitized modeling workflows in Computational Neuroscience.

This chapter is based on the following publication:

Johanna Senk, Alper Yegenoglu, Olivier Amblet, Yury Brukau, Andrew Davison, David R. Lester, Anna Lührs, Pietro Quaglio, Vahid Rostami, Andrew Rowley, Bernd Schuller, Alan B. Stokes, Sacha J. van Albada, Daniel Zielasko, Markus Diesmann, Benjamin Weyers, Michael Denker, Sonja Grün (2017), A collaborative simulation-analysis workflow for computational neuroscience using HPC, in Edoardo Di Napoli, Marc-André Hermanns, Hristo Iliev, Andreas Lintermann, Alexander Peyser, eds., High-Performance Scientific Computing. JHPCS 2016. Lecture Notes in Computer Science, vol 10164. Springer, Cham, 243–256, doi: 10.1007/978-3-319-53862-4_21

Author contributions:

Under the supervision of Michael Denker, Benjamin Weyers, Sonja Grün and Markus Diesmann, Alper Yegenoglu and the author equally share the main contribution to this publication, that involves conceptual work, coordination of the collaboration, code writing, development of Jupyter Notebooks and tasks, and the whole workflow integration into the HBP Collaboratory. While the author focused on the simulation of the microcircuit with the NEST back end, Alper Yegenoglu focused on data analysis with Elephant. All authors contributed to the writing of the manuscript.

2.1 INTRODUCTION

Workflows in the natural sciences that deal with the acquisition and analysis of experimental or simulated data often comprise an intricate sequence of processing steps, each of which requires the use of diverse software tools. The resulting heterogeneity in terms of both the composition of steps of the workflows and the diversity of tools employed generates a substantial degree of complexity that increases with the number of researchers involved. The situation is compounded if the tools themselves add an additional level of complexity, for instance requiring scientists to be trained in using the software. In particular, in interdisciplinary settings, users need to know how to integrate the various tools that may be unfamiliar to them in terms of practical usage and/or the scientific processing step they perform. A commonly encountered scenario where this holds are workflows and tools that rely on the capabilities of high-performance computing (HPC) systems, but where the access to and usage of such systems is complicated for less experienced users. Another problem dimension is added if requirements such as reproducibility or reusability are considered, for example in terms of version control of code and data or provenance tracking of the analysis. Ad hoc approaches are bound to fail as the complexity of the workflow increases. Instead, the heterogeneity and the emerging complexity of such workflows call for user-friendly standards and software tools that meet and integrate such requirements.

Interdisciplinary workflows in computational neuroscience are facing these problems. Computational neuroscience entails integrating and analyzing experimental data, building network models for brain simulations, and using theory to develop concepts concerning neuronal information processing. Datasets obtained from both experiments and simulations are highly diverse in their internal structure and content. Analysis tools are therefore often adapted to the specifics of the experiment or the simulation study. Moreover, the analysis tools employ methods with a different focus depending on the source of the data. For instance, while simulation studies often exploit the fact that data are controlled and can be acquired over long durations and large ensembles, experimentally obtained data are often analyzed with respect to the inherent non-stationarity of the data and the behavioral protocols. Thus, in the attempt to bring experiment and simulation closer together, a large array of heterogeneous data standards and tools exist. These need to be merged and linked into workflows for analysis, in particular for comparison of data from model and experiment. Moreover, workflows typically consist of complex chains of processing steps that often require the use of HPC systems for expensive computations, e.g., to run large-scale network simulations, to process extensive data records, or to perform parameter scans. At the same time, they necessitate the option of exploratory analysis in an interactive fashion (cf. [Denker and Grün, 2015](#)).

To propose a solution for this problem domain, this work presents an example of such an interdisciplinary and heterogeneous workflow in computational neuroscience. We describe in the following a realistic research question, derive its concrete challenges, and provide a possible approach to tackle them. The research question addressed is to what extent different simulators produce comparable results as they differ, for instance, in the biological detail they can represent, their

underlying architecture, performance, flexibility, or other design goals like application in robotics. We here compare two simulators, NEST and SpiNNaker, relying on two different types of digital hardware, and outline a workflow which is conceptually applicable for the comparison of other simulators as well. Both simulators aim at simulations of large networks of simple spiking neuron models which are currently gaining significant relevance in the field of computational neuroscience (van Albada et al., 2015). The simulator NEST¹ (NEural Simulation Tool, see Eppler et al., 2015; Gewaltig and Diesmann, 2007) is optimized to efficiently use existing HPC infrastructure and allows for exact and reproducible simulations. It combines ease of use (Python interface) and runtime performance (C++ kernel, multi-threading, and MPI-parallelism). Recent development of NEST has notably reduced the memory requirements (Kunkel et al., 2014), thus further facilitating large-scale simulations. In contrast, SpiNNaker² (Spiking Neural Network Architecture, see Furber et al., 2013b; Stokes et al., 2016) is a specific neuromorphic hardware designed for biological real-time operation, low power consumption, and scalability. The architecture of the SpiNNaker Neuromorphic Computing Platform itself is inspired by biological neuronal networks. A large number of low-powered and thus energy-efficient computation units are highly connected together by an asynchronous communication network.

Since there are by design major differences in how NEST and SpiNNaker operate, it remains to be evaluated to what extent simulation results are comparable at all. This suggests the refined research question (cf. above):

If a simulation of the same neural network model is run both on an HPC system using NEST and on the neuromorphic hardware system SpiNNaker, are the results the same?

When investigating the implications of this question in detail, the following distinct challenges emerge. First of all, access to HPC systems and the neuromorphic hardware is required. It must also be guaranteed that the same network model is simulated on both systems to enable a direct comparison of the simulation results; this calls for a model description valid for both systems. The model development itself needs to be performed under version control and the source code must be accessible to all scientists involved. Assuming that both simulations have been run successfully, the simulation output must be validated and compared using a suitable analysis tool. In order to evaluate the recorded series of spike times of each simulated neuron, i.e., the spike trains, and to assess whether the results from both simulations can be considered “the same”, a tool for statistical data analysis is needed. A prerequisite for the analysis is that the simulation output is readily accessible, for example after transferring it to the same data storage, and that the data is available in the same format. It is further desirable to have the possibility of an interactive data analysis and a more sophisticated one relying on HPC. Finally, visualization techniques are needed in order to convey a more intuitive understanding of the expectedly complex analysis results.

In summary, the initially posed research problem can be broken down into a list of separate demands which encompass the collaboration of experts from different scientific disciplines as well as a series of consecutive tasks that depend on

¹ <http://nest-simulator.org>

² <http://apt.cs.manchester.ac.uk/projects/SpiNNaker>

access to and usage of specific tools. On first sight, one could argue that solutions for the isolated problems already exist. To give an example, there are web-based repository hosting services like GitHub³ for source code management and version control which allow sharing repositories among researchers of different institutions. Sumatra⁴ allows for automated tracking of scientific computations. For the field of neuroscience, in particular, there are platforms to facilitate access, storage, analysis, and exchange of data, such as the G-Node Data Portal⁵. Resources for computational models of neural systems are Open Source Brain⁶ and ModelDB⁷. The Neuroscience Gateway⁸ provides an opportunity for neuroscientists to use HPC resources. When using such independent solutions, however, one faces major problems in terms of provenance tracking. If there are no links between the individual components, it will quickly become untraceable who did what, why, and when. Since scientific progress is rarely straightforward, but includes trial and error, repetitions, and iterative improvement, it is crucial to keep track of all steps involved and their history. Here, we suggest integrating such components into one collaboration platform in order to establish a stable and reproducible workflow.

In this study, we demonstrate how such a workflow can be implemented addressing the aforementioned problems by integrating established and emerging software tools using a web-based infrastructure. We will describe the workflow consisting of the following steps: (i) simulate the activity generated by a model of a cortical microcircuit (Potjans and Diesmann, 2014) using the NEST simulator, (ii) simulate the same network model with identical parameters using the SpiNNaker system, (iii) pool data on the centralized storage of the integrative software infrastructure, (iv) compare the resulting activity data using Elephant (Electrophysiology Analysis Toolkit, see Yegenoglu et al., 2015, 2016b), and (v) interactively visualize the analysis results. In the following, the individual steps comprising this workflow will be briefly highlighted, before we discuss the benefits and shortcomings of the currently available implementation of this workflow based on an integrative software architecture that is developed in the European FET Flagship “Human Brain Project” (HBP).

2.2 WORKFLOW

The principal layout of the workflow that we defined and implemented to compare the activity data coming from the classical NEST-based simulations, and the neuromorphic SpiNNaker-based simulations is depicted in Figure 2.1. The workflow comprises a collaboration of different laboratories with different expertises (marked by colored dots): HPC, neural network simulation, neuromorphic hardware, data analysis, and visualization. Each of the five steps in Figure 2.1 is defined by a set of methods and tools specific to these individual areas of expertise, and needs to be integrated into a common infrastructure that makes them accessible

3 <https://github.com>

4 <http://neuralensemble.org/sumatra>

5 <http://www.g-node.org>

6 <http://opensourcebrain.org>

7 <https://senselab.med.yale.edu/modeldb>

8 <https://www.nsgportal.org>

for cooperative work. All components integrated in the workflow are based on or accessible via Python, a programming language which is becoming commonly used in computational neuroscience for both simulation and data analysis (Davison, 2009). Furthermore, we incorporate tools like NEST and Elephant which follow standardized development processes in software engineering, e.g., test-driven development and continuous integration.

We embedded our workflow into an integrative software platform called the “HBP Collaboratory”⁹. The Collaboratory is a web-based portal which provides a common entry point to facilitate collaboration by providing a shared project space (termed the “Collab”) for groups of scientists. Specifically, for our project we created the “NEST SpiNNaker Elephant Demo” Collab¹⁰, which enables us to share simulation data and analysis results through the centralized “Collab storage”, to use all relevant applications including access to HPC infrastructures and neuromorphic hardware, and to document the workflow. For interactive Python programming, the Collaboratory provides Jupyter Notebooks¹¹ that run directly in the Collab and have NEST and Elephant preinstalled by default. The HPC resources required for simulations and analyses are launched via the Collaboratory’s internal task framework as jobs (“tasks”) which are sent to predefined compute clusters or supercomputers. Finally, the Collaboratory offers basic provenance tracking, providing the ability to reenact an already executed task.

The specific network we simulate is a full-scale neural network model of a cortical microcircuit (Potjans and Diesmann, 2014). Full-scale means that the natural density of neurons and synapses of the biological circuit is preserved. The microcircuit represents 1 mm² of cortex and contains around 80,000 spiking leaky integrate-and-fire point neurons connected by around 0.3 billion synapses in four cortical layers (L2/3, L4, L5, L6). Each layer comprises an excitatory and an inhibitory neuron population which are interconnected with cell-type- and layer-specific connection probabilities derived from experimental data on early sensory cortex. The model is well-suited for our workflow for two reasons: First, it is of neuroscientific interest since it is a minimal microcircuit that combines a realistic number of synapses per neuron with sparse network connectivity as found in cortex, exhibits realistic spiking activity, and serves as a prototype for larger networks (see Schmidt et al., 2017 for an example). Second, it fits onto both systems in terms of computational resources. More precisely, the network size indicates HPC for the NEST simulations, although it is still considered to present a small workload for HPC systems. For SpiNNaker, the model is an interesting use case because it requires the parallel use of multiple boards (van Albada et al., 2016). For comparability of the simulation results, we use a common model implementation based on PyNN and develop the source code using the version control system git¹². PyNN is a Python API for simulator-independent neuronal network model specification (Davison et al., 2009, 2013, 2015). The PyNN API enables writing generic code to control different simulators such as NEST, NEURON, Brian and also neuromorphic hardware (Brüderle et al., 2011), including the SpiNNaker platform. In the

⁹ <http://collab.humanbrainproject.eu>

¹⁰ <https://collab.humanbrainproject.eu/#/collab/507/nav/6326>

¹¹ <http://jupyter.org>

¹² <https://git-scm.com>

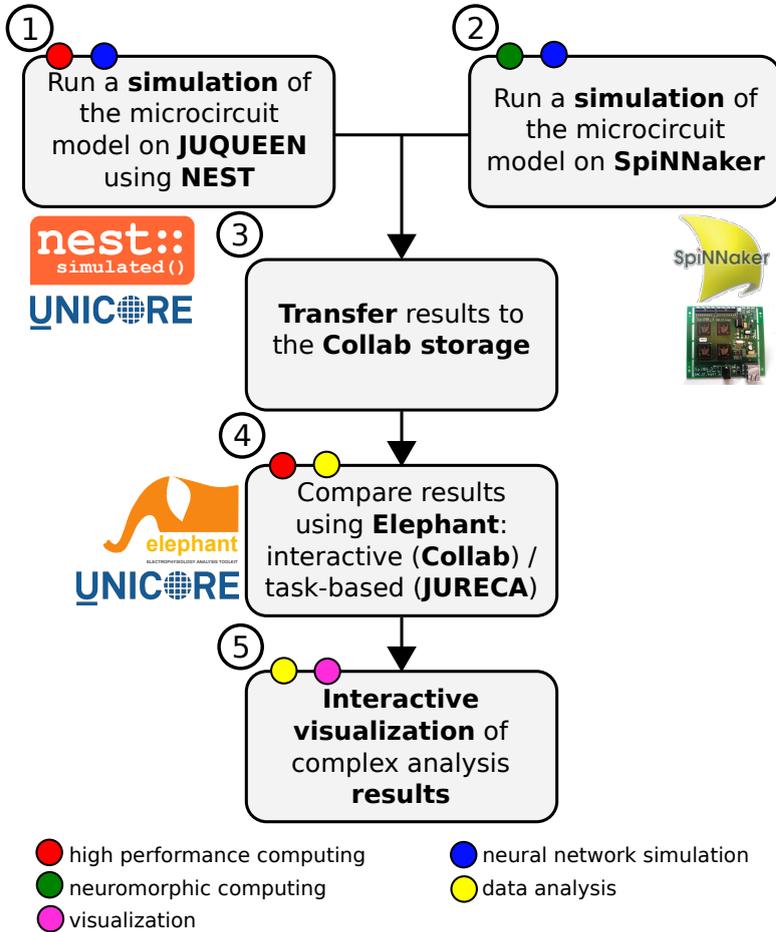


Figure 2.1: **Workflow overview.** A network simulation of a cortical microcircuit model is run using both NEST (1) and SpiNNaker (2). Simulation results are transferred to a common storage (3) and compared utilizing functionalities of the Elephant library (4). Complex analysis results are visualized to gain further insight (5). The middleware UNICORE is used to access HPC systems. Colored dots on top of each box indicate the disciplines involved. For example, supercomputers (red dots) are used to run a NEST simulation and to compare results using Elephant.

case of SpiNNaker, a software library is used to break down the Python network description into small chunks each of which can be run on a core, and to route communications between the parts of the network.

As the first step of the workflow (Figure 2.1, Step 1) we run the microcircuit simulation on an HPC system using NEST. We established two ways to access HPC resources from within the Collaboratory: using the task framework or interactively from a Jupyter Notebook running within the Collab. Specific simulation and network parameters such as the simulation duration (in our case: 10s) can be configured via a Jupyter Notebook before submitting the job. Both the task framework and the Jupyter Notebooks use UNICORE¹³ as middleware that yields secure and seamless access to supercomputing and data resources from a web-based environment such as the Collaboratory. UNICORE provides a wide range of features for HPC job submission and management as well as data transfer and handling. Concretely, the microcircuit simulation runs on JUQUEEN¹⁴, a supercomputer with an IBM BlueGene/Q architecture at the Jülich Supercomputing Centre (JSC), Forschungszentrum Jülich, Germany and is one of the fastest supercomputers in Europe and worldwide. Gathered simulation results, i.e., a down-sampled set of spike data from 100 excitatory and 100 inhibitory neurons from each of the four simulated cortical layers (800 neurons in total), are automatically copied from the supercomputer to the central Collab storage (Step 3).

In the second step of the workflow (Step 2), the simulation is run on (a part of) the half-million-core SpiNNaker machine located in Manchester, UK, with identical parameters. The Collaboratory integration in this case is implemented via the Neuromorphic Job Manager App. This allows users to submit PyNN scripts directly, through a git repository, or as a compressed archive of files using a webpage-based user interface. It is also possible to submit jobs directly from a Python script through the `hbp_neuromorphic_platform`¹⁵ library. Jobs for execution on SpiNNaker are periodically retrieved and run on a virtual machine cluster situated close to the SpiNNaker machine. Results are then retrieved from the machine, stored locally, and finally transferred to the Collab storage via the Job Manager App when requested (Step 3).

In order to compare the simulation results of the two systems and to characterize potential differences, we analyze the statistical features of the two datasets using the Elephant library (Step 4). Elephant is a community-centered, open-source Python library for analyzing multi-scale data on brain dynamics from experiments and simulations. The focus is on tools for the analysis of electrical activity, such as single-unit or massively parallel spike train data and local field potentials (LFPs). The scope of the library covers the analysis of analog signals (including time-domain and frequency-domain methods), spike-based analysis (e.g., spike train correlation, spike pattern analysis), and methods combining both signal types (e.g., spike-triggered averaging of an LFP signal). We first execute a task on the Collab to convert the data into the HDF5¹⁶ format and save the result. This data format is compatible with the Neo library (Garcia et al., 2016, 2014) which serves as a

¹³ <https://www.unicore.eu>

¹⁴ <http://www.fz-juelich.de/ias/jsc/EN>

¹⁵ https://pypi.python.org/pypi/hbp_neuromorphic_platform

¹⁶ <https://www.hdfgroup.org/hdf5>

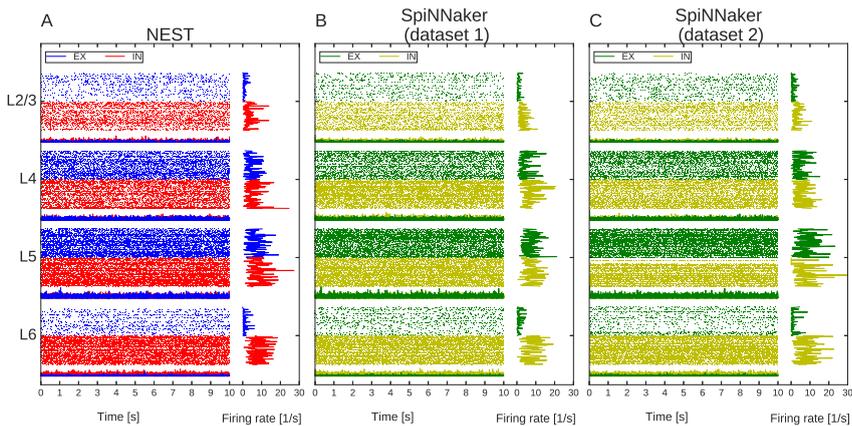


Figure 2.2: The data obtained from the NEST (A) and the first (B) and second (C) iteration of the SpiNNaker simulations are presented as raster displays. Each dot indicates a spike at its time of occurrence and each line represents the firing activity of a neuron. Neurons are grouped into different cortical layers (L2/3, L4, L5, L6) and sorted by neuron type, i.e., excitatory (EX) and inhibitory (IN). The histograms at the bottom and right side show the population activity and the firing rate of the neurons, respectively.

foundation of Elephant. The internal structure of the HDF5 file complies with the Neo¹⁷ architecture introduced in version 0.2 and is created using the Neo HDF5 I/O¹⁸.

Next, we compare the two simulation results in an interactive fashion using the Elephant library within a Jupyter Notebook. In particular, we consider features that are typically analyzed in neuroscience, such as the irregularity of the individual spike trains and correlations between pairs of neurons. A first visual impression (Step 5) of the firing behavior of individual neurons and the neuronal populations is provided by dot displays as shown for NEST in Figure 2.2A and for SpiNNaker in Figure 2.2B. Each spike is represented as a dot at the time of its occurrence. Multiple neurons are displayed below each other in different lines. The summed population activity is shown in the histogram below the dot display of the respective layer. On the right, the average firing rates of the individual neurons are depicted. The visualized spiking activity of NEST and SpiNNaker is qualitatively comparable. As we compare different network realizations on NEST and SpiNNaker, the neurons do not correspond one-to-one between the two systems, and hence, statistical measures for comparison are needed.

To capture properties of the coordination between individual neurons, we also computed Pearson correlation coefficients, i.e., the zero-delay correlation coefficients between all pairs of neurons in each population. Their distributions are visualized in Figure 2.3A,B. The shapes of the distributions agree between the two types of simulations—except for a remarkable difference for the layer four (L4) neurons (second row of Figure 2.3). In the process of tracking down the origin

¹⁷ <http://neo.readthedocs.io/en/0.4.1/core.html#grouping-objects>

¹⁸ <http://neo.readthedocs.io/en/0.4.1/io.html#neo.io.NeoHdf5IO>

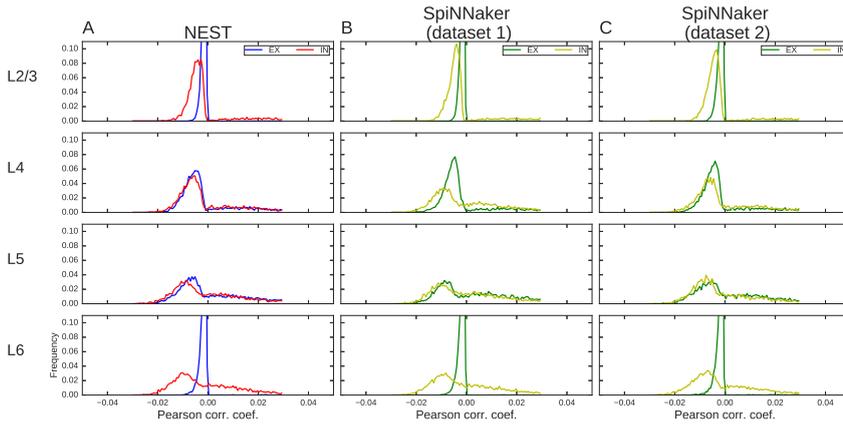


Figure 2.3: **The distributions of Pearson correlation coefficients are shown for the NEST (A) and the two SpiNNaker (B and C) simulations**. The correlation coefficients were computed for all pairs of recorded neurons of the same layer and neuron type. The y-axis is truncated for comparability between the different neuron populations.

of this variation, a slight improvement in the simulation and recording routine of SpiNNaker was made. With respect to the dot displays, the dataset resulting from this iteration step, shown in Figure 2.2C, does not exhibit striking differences to the initial results that are noticeable by eye. However, the distribution of correlation coefficients (Figure 2.3C) approaches the ones from the NEST simulation in Figure 2.3A in L4 much better.

We hence see that rather simple analysis methods already suffice to reveal prominent differences between the simulation results. To access the subtle differences, however, we aimed to uncover the full correlation structure of the activity using more sophisticated techniques. Since these demand HPC resources due to expensive surrogate generation and multiple hypothesis testing, we execute a task that uses UNICORE to send the data to the HPC system JURECA, located at the Forschungszentrum Jülich, and to parallelize the analysis (Step 4).

After completion, the results are again transferred to the Collab storage. Due to the complicated nature of the resulting data, we decided to visualize them using a special tool designed to interactively probe the correlation structure in order to obtain an overview and to gain further insight. To this end, the data are staged to a visualization server using dCache¹⁹, a distributed file system. This enables the use of high-fidelity visualization tools, e.g., based on the visualization toolkit ViSTA²⁰ (Virtual Reality Toolkit). Such a tool can interpret the data as a graph and render it as a node-link diagram, with the nodes representing the neurons and the weighted edges representing the correlations. The correlation value serves as attraction criterion in a force-directed layout algorithm that results in visual clusters of neurons where neurons are spatially close if they are strongly correlated.

¹⁹ <https://www.dcache.org>

²⁰ <http://www.itc.rwth-aachen.de/cms/IT-Center/Forschung-Projekte/Virtuelle-Realitaet/Infrastruktur/~fgmo/ViSTA-Virtual-Reality-Toolkit/?lidx=1>

This helps the analyst to identify and compare possible correlation patterns in the statistical data. Using a web-based streaming library, the visualization tool could deliver the rendered images to a website integrated into the Collab.

2.3 DISCUSSION

Our general aim is to map an interdisciplinary workflow involving multiple steps and tools to a common platform and to address major problems emerging from this setting: heterogeneous data, diverse knowledge of the participants in the workflow, the involvement of complex tools and infrastructures as well as aspects such as reproducibility, reusability and iterative refinement. The requirements for such a platform are to provide a collaborative environment which allows one to integrate and easily access software tools, libraries, and data, as well as HPC systems which are needed for demanding simulations and analyses. As an example, we demonstrated a concrete working solution implemented in the HBP Collaboratory and responded to the list of challenges identified in the introduction. From within the Collaboratory, we established access to the HPC systems JUQUEEN and JURECA by means of the middleware UNICORE. Likewise, a connection to a SpiNNaker machine was realized via the Neuromorphic Job Manager App. PyNN provides an interface to NEST and SpiNNaker and hence allows for a common model description, developed under version control with git. The simulation output from both systems was obtained in the HDF5 data format and, once transferred to the common Collab storage, it was read by the Elephant library which offers a variety of analysis methods. Jupyter Notebooks in the Collab were used for interactive Python programming and, as a last point, analysis results were visualized. In summary, our workflow comprises a variety of tools and resources which themselves are widely used within the communities involved. Our Collab is public within the Collaboratory, i.e., other users can inspect the developed tasks and Jupyter Notebooks as documented there or integrate them into their own Collabs.

Using a common framework usually restricts the user to available tools, but a versatile structure allows for adaption and extension if requested by the user. For example, a continuous exchange between us (computational neuroscientists as users and software engineers as developers of the Collaboratory) resulted in the integration of Jupyter Notebooks into the portal. This shows the importance of a bilateral communication between developers and users for a successful and ongoing development of a collaboration platform. Use cases based on the daily practice of the users are a main component of this development. During the implementation of the workflow into the Collaboratory, we not only aimed at the full integration of tools like Jupyter Notebooks, but also at establishing interfaces to tools outside the portal, for example by making supercomputers accessible via the middleware UNICORE. Furthermore, we accounted for different data types and formats with a conversion task to enable applying the same analysis functions to data obtained from different sources.

In addition to the inherent heterogeneity of workflow components, reproducibility and reusability are considered. For single workflow steps, we use tools that fulfill criteria of quality assurance, e.g., continuous integration and test-driven development. The task framework of the Collaboratory already allows provenance

tracking to some extent, but the whole workflow is not fully traceable, yet. Thus, we envision that all individual steps of the workflow can be tracked from the beginning of the simulations to the end of the visualization. The workflow can be improved through an iterative adjustment of single steps and parameters. Therefore, we need a flexible workflow implementation that allows for easy integration of individual parts.

Transparency of the workflow is an important property of a successful integration since it allows collaborators to comprehend and even carry out different steps of the workflow. The web interface of the Collaboratory serves as a common access point for collaboration where we collect documentation together with code, data, and results, as well as provenance information. However, a complete representation of how individual steps of the workflow are connected within the Collab is still ongoing work.

A possible next step is to set up a test battery to quantify detailed differences in results complementing the visual inspection in the interactive analysis. Furthermore, having common metadata (e.g., layer of neuron) for the data simulated by the two systems is important for the follow-up analysis tools. Here, we aim to make use of the odML metadata framework which is also used for experimental data (Grewe et al., 2011; Zehl et al., 2016). We seek a shared terminology which allows for easy handling and manipulation of the data, and avoids misinterpretations of vocabulary.

In contrast to well-established groupware solutions, such as BSCW (Appelt, 1999) for project management and file sharing or Moodle (Dougiamas and Taylor, 2003) used in the context of academic teaching, the Collaboratory offers a domain-specific integration of tools and middleware. It concentrates on content-based communication (e.g., sharing data and documentation) instead of direct communication using video or audio conference tools and therefore offers a single point of access for tools and data used in the neurosciences. The latter makes it especially useful for this scientific community and thus differentiates it from existing more general solutions.

Taken together, we believe that the workflow implementation presented in this work provides a promising vista of how a collaborative system such as the HBP Collaboratory, supported by a chain of compatible software tools, can help scientists to come together in large, interdisciplinary, and collaborative research endeavors. Indeed the availability of technologies that allow for large collaborative research endeavors is expected to become an indispensable asset as neuroscience moves towards questions that can no longer be handled by a single person. Considering the diversity of approaches and data types in the field of neuroscience, even the workflow presented here, consisting of a collaborative effort of eighteen researchers distributed over six institutes, may be considered a small collaboration in the future. These developments are expected to produce new challenges, e.g., the need for more dynamic ways of setting up workflows, better visualizations of the provenance information for generated data, or the ability to control more heterogeneous HPC environments required by the individual components of the workflows.

This chapter is based on the following publication:

Sacha J. van Albada, Andrew G. Rowley, Johanna Senk, Michael Hopkins, Maximilian Schmidt, Alan B. Stokes, David R. Lester, Markus Diesmann, Steve B. Furber (2018), Performance comparison of the digital neuromorphic hardware SpiNNaker and the neural network simulation software NEST for a full-scale cortical microcircuit model, *Frontiers in Neuroscience*, 12, 291, doi: 10.3389/fnins.2018.00291

Author contributions:

Under the supervision of Sacha van Albada and Markus Diesmann, the author designed and carried out the HPC performance benchmarks including power measurements. The author wrote simulation code, performed NEST simulations, and analyzed the data. All authors contributed to the conceptual work and the writing of the manuscript.

3.1 INTRODUCTION

Tools for simulating neural networks fall into two categories: simulation software and neuromorphic hardware. The available features, the speed at which the simulation engine arrives at the solution, and the power consumption differ between tools, but the tools are rarely systematically compared. To find out where we stand and to provide guidance for future research, we need to learn how to port network models discussed in the current literature from conventional software implementations to neuromorphic hardware and how to quantitatively compare performance.

The distinction between simulation software and neuromorphic hardware is not clear-cut. Next to the hardware, a neuromorphic system readily usable by neuroscientists requires a multi-level software stack engaging in tasks from the interpretation of a domain-specific model description language to the mapping of the neural network to the topology of the neuromorphic hardware. Reversely, simulation software profits from computer hardware adapted to the microscopic parallelism of neural networks with many computational cores and a tight integration of processing hardware and memory. For the purpose of the present study we refer to simulation software as a system that runs on conventional high-performance computing hardware without dedicated neuromorphic hardware.

The time as well as the energy required to arrive at the solution are becoming relevant as neuroscientists turn to supercomputers to simulate brain-scale neural networks at cellular resolution. Today's supercomputers require tens of minutes to simulate one second of biological time and consume megawatts of power (Kunkel et al., 2014; Jordan et al., 2018). This means that any studies on processes like plasticity, learning, and development exhibited over hours and days of biological time are outside our reach.

Although this is sometimes forgotten, not only speed and power consumption but also the accuracy of the simulation results is of importance: a highly inaccurate solution can be obtained arbitrarily fast. In other words, a statement on the wall-clock time required to arrive at the solution is meaningless without a statement on the achieved accuracy. Like runtime, energy consumption depends on the level of simulation accuracy. Low energy consumption is emphasized in the development of neuromorphic hardware, but accuracy is generally not explicitly taken into account when characterizing energy consumption. How one quantifies accuracy should be determined in the light of the trade-off between the combination of precision and flexibility on the one hand and the combination of speed and energy efficiency on the other hand which is the main idea behind dedicated hardware. If a dedicated hardware trades precision for speed and energy efficiency, for instance by having noisy components or not delivering every single spike, this is acceptable if the given precision still yields the desired network behavior. The relevant issue is then not whether but how to assess accuracy, that is, defining how the network should behave.

Here, we consider as a use case the digital neuromorphic hardware SpiNNaker (Furber et al., 2013b) and the neural network simulation software NEST (Gewaltig and Diesmann, 2007), both in use by the neuroscientific community and supporting the simulator-independent description language PyNN (Davison et al., 2009). Both NEST and SpiNNaker are designed to enable the simulation of large neural

network models. SpiNNaker enhances its efficiency through asynchronous update where spikes are processed as they come in and are dropped if the receiving process is busy over several delivery cycles. It is especially suited to robotic applications enabling the simulation to operate in real-time, but since it is general-purpose, in principle any type of neural network model can be simulated, including biological and artificial neural networks. In the context of the European Human Brain Project (HBP) a large system is under construction at the University of Manchester targeting brain-scale simulations. The networks in question may have static synapses or include plasticity. For simplicity, and since there is a close relationship between simulator performance with and without synaptic plasticity (e.g., [Knight and Furber, 2016](#)), we here focus on a non-plastic network: a spiking cortical microcircuit model ([Potjans and Diesmann, 2014](#)).

The microcircuit is regarded as unit cell of cortex repeated to cover larger areas of cortical surface and different cortical areas. The model represents the full density of connectivity in 1 mm² of the cortical sheet by about 80,000 leaky integrate-and-fire (LIF) model neurons and 0.3 billion synapses. This is the smallest network size where a realistic number of synapses and a realistic connection probability are simultaneously achieved. The capability to simulate this model constitutes a breakthrough as larger cortical models are necessarily less densely connected, with only a limited increase in the number of synapses per neuron for increased model size. Consequently, from this network size on, the computer memory required to store the synaptic parameters grows close to linearly with network size ([Lansner and Diesmann, 2012](#)). Further, a simulation technology can be devised such that the memory consumption of a compute node is independent of the total number of neurons in the network ([Jordan et al., 2018](#)). This renders the total memory consumption approximately directly proportional to the number of neuronal and synaptic elements in the model. The model already serves as a building block for a number of further studies and larger networks ([Wagatsuma et al., 2011](#); [Schmidt et al., 2016](#); [Cain et al., 2016](#); [Hagen et al., 2016a](#); [Schwalger et al., 2017](#)), and a first comparison of the simulation results of NEST and SpiNNaker for this model has served as a test case for a workflow implementation on the collaboration platform of the Human Brain Project ([Senk et al., 2017b](#)). The original implementation uses NEST, which can also handle much larger networks with trillions of synapses ([RIKEN BSI, 2013](#); [Kunkel et al., 2014](#); [Forschungszentrum Jülich, 2018](#); [Jordan et al., 2018](#)) under the increased memory consumption and run time costs indicated above. The previously largest simulations on SpiNNaker comprised about 50 million ([Sharp et al., 2014](#); [Knight et al., 2016](#)) and 86 million synapses ([Stromatias et al., 2013](#)). Thus, the present study describes the largest simulation on SpiNNaker to date, and also the first to implement the connectivity at full biological density.

SpiNNaker achieves real-time performance for an integration time step of 1 ms, which is suited to networks with dynamics on time scales sufficiently greater than 1 ms. While a resolution of 1 ms generally suffices for today's applications in robotics and artificial neural networks, a time step of 0.1 ms is typical for neuroscience applications due to the neurobiological time scales and the need to avoid artifacts of global synchronization ([Morrison et al., 2007b](#)). The model of [Potjans and Diesmann \(2014\)](#) has synaptic time constants of 0.5 ms, and therefore

requires integration time steps smaller than this. The current software controlling SpiNNaker enables using small time steps by slowing down the simulation. In the present work, we show how this feature in combination with further improvements of the software stack allows the cortical microcircuit model to be accurately integrated. This result demonstrates the usability of SpiNNaker for large-scale neural network simulations with biologically realistic time scales.

To assess accuracy, we compare simulation results with a reference solution obtained with an alternative solver (Hanuschkin et al., 2010; Morrison et al., 2007b) available in the NEST simulation code where spikes are not restricted to the grid spanned by the computation step size. The spike times from the cortical microcircuit model obtained with different simulation engines can only be compared in a statistical sense. Therefore we also look at single-neuron accuracy (Henker et al., 2012). Here, we consider both the 0.1 ms time step used in the microcircuit simulations, and 1 ms, the original design specification of SpiNNaker, and further, we investigate different spike rates to vary the relative contributions of subthreshold and spiking activity. This is relevant because NEST integrates the subthreshold dynamics exactly (Rotter and Diesmann, 1999), whereas SpiNNaker uses exponential integration (MacGregor, 1987), in which the synaptic currents are treated as piecewise constant. For the microcircuit model, we characterize accuracy based on distributions of spike rates, spike train irregularity, and correlations. Spike rates are chosen as a first-order measure of neural activity, and correlations together with spike train irregularity are relevant because cortical activity is known to be asynchronous irregular (van Vreeswijk and Sompolinsky, 1998); mesoscopic measures of brain activity like the local field potential (LFP) primarily reflect correlations in the microscopic dynamics (Hagen et al., 2016a); and correlations in spiking activity drive further aspects of network dynamics like spike-timing-dependent plasticity (STDP; Morrison et al., 2007a) underlying system-level learning. The three aforementioned measures of spiking activity are also the focus in the work of Potjans and Diesmann (2014).

Previous work has evaluated the energy consumption of various types of processors (Hasler and Marr, 2013) including SpiNNaker (Sharp et al., 2012; Stomatias et al., 2013) in relation to the number of operations performed. Here, we take a different approach, comparing the energy consumption of two simulation engines under the condition of comparable accuracy. This accuracy depends not only on the number of operations of a given precision, but also on the algorithms employed. For comparison with previous results (Sharp et al., 2012; Stomatias et al., 2013), we further derive the energy consumed per synaptic event.

In the following, we compare the accuracy of single-neuron LIF simulations between NEST and SpiNNaker, describe the adjustments made to SpiNNaker to enable the cortical microcircuit model to be implemented, and compare both simulators in terms of accuracy, runtime, and energy consumption. We also discuss the sources of differences in simulation results and performance between NEST and SpiNNaker. Thus, our study enables neuromorphic engineers to learn more about the internal workings of SpiNNaker and the implications for performance, and brings SpiNNaker closer to being a tool of choice for computational neuroscience use cases with large network size and short biological time scales.

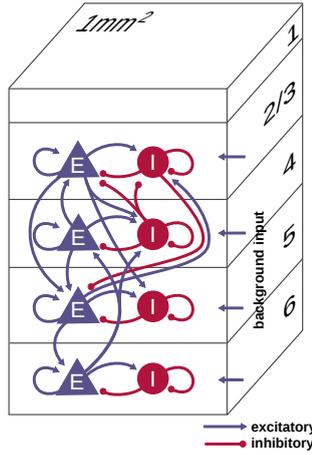


Figure 3.1: **Schematic illustration of the microcircuit model of early sensory cortex.** The model represents 1 mm^2 of cortex with the full density of neurons and synapses, for a total of 77,169 neurons and about 3×10^8 synapses. Each of the layers 2/3, 4, 5, and 6 contains an excitatory (E) and inhibitory (I) population of leaky integrate-and-fire model neurons. All neurons receive an external Poisson drive representing inputs from the rest of the brain. Figure adapted from Potjans and Diesmann (2014) with permission.

3.2 METHODS

3.2.1 The leaky integrate-and-fire neuron model

The cortical microcircuit model uses leaky integrate-and-fire (LIF) model neurons with synaptic currents modeled as jumps followed by an exponential decay. The subthreshold dynamics of each neuron is given by

$$\begin{aligned}\tau_m \frac{dV_i}{dt} &= -(V_i - E_L) + R_m I_i(t), \\ \tau_s \frac{dI_i}{dt} &= -I_i + \tau_s \sum_j J_{ij} s_j(t - d_j),\end{aligned}$$

where τ_m and τ_s are membrane and synaptic time constants, E_L is the leak or resting potential, R_m is the membrane resistance, V_i is the membrane potential of neuron i , I_i is the total synaptic current onto the neuron, J_{ij} is the jump in the synaptic current due to a single spike from neuron j , $s_j = \sum_k \delta(t - t_k^j)$ are the incoming spike trains, and d_j is the transmission delay. When V_i reaches a threshold θ , a spike is emitted, and the membrane potential is clamped to a level V_r for a refractory period τ_{ref} . Table 3.1 lists the single-neuron parameters.

NEST integrates this model using exact integration (Rotter and Diesmann, 1999), so that the subthreshold dynamics precisely follows the analytical solution. The spikes can either be constrained to the time grid or interpolated between grid points to yield precise spike times (Hanuschkin et al., 2010; Morrison et al., 2007b). In the present study, we consider both options, the latter providing a reference so-

membrane time constant	τ_m	10 ms
synaptic time constant	τ_s	0.5 ms
refractory period	τ_{ref}	2 ms
membrane resistance	R_m	40 M Ω
leak potential	E_L	-65 mV
threshold	θ	-50 mV
reset potential	V_r	-65 mV

Table 3.1: Parameters of the leaky integrate-and-fire model neurons used in the simulations.

lution. For reasons of modularity, SpiNNaker separates the neuron and synapse dynamics and uses exponential integration (MacGregor, 1987; reviewed in Rötter and Diesmann, 1999), in which the input current to the membrane potential equation is treated as piecewise constant. The synaptic currents are decayed over one time step before being added to the input, to ensure that the total charge transferred per synaptic event is $J\tau_s$, as in the exact solution.

3.2.2 Single-neuron tests

Simple systems such as single LIF model neurons allow a deterministic assessment of simulation accuracy (Henker et al., 2012). We assess the accuracy of NEST and SpiNNaker by comparing with precise solutions the subthreshold and spiking dynamics of single LIF model neurons receiving excitatory Poisson input with synaptic strength 87.8 pA, equal to the mean synaptic strength for the excitatory connections between most populations in the network model. The study considers both integration time steps of 0.1 ms to match the network simulations, and 1 ms, matching the primary design specification of SpiNNaker. Two input rates are investigated: 8,000 spikes/s (giving an output rate of around 17 spikes/s) and 10,000 spikes/s (giving an output rate of around 47 spikes/s), to study different proportions of subthreshold activity and spiking. As in the network simulations, the input spikes are constrained to the time grid. The simulations with the lower input rate are run for 16 s and those with the higher input rate for 4 s biological time to yield comparable total numbers of spikes for the low-rate and high-rate simulations, and a total of 10 simulations with different random seeds are performed for each setting. The other neuron parameters are as in the network model.

We characterize the accuracy of the single-neuron simulations in four ways, in each case comparing with NEST simulations with precise spike timing: 1) cross-correlation histograms of spike times with bin width equal to half the integration time step; 2) Pearson correlation coefficients between membrane potential traces recorded at the integration time steps; 3) the average percentage lead or lag in spike times; and 4) the root mean square error (RMSE) of the spike times after correcting for the average lead or lag. To determine the accumulated percentage lead or lag in spike times, we first find $N = \min(N_{spikes}[\text{precise}], N_{spikes}[\text{discrete}])$, where discrete refers to simulations with SpiNNaker or with NEST with spikes constrained to the grid and N_{spikes} is the total number of spikes in the respective simulation. The accumulated fractional lead or lag in spike times is then com-

puted as $[t_{\text{discrete}}(N) - t_{\text{precise}}(N)] / t_{\text{discrete}}(N)$, where $t_{\text{precise}}(N)$ and $t_{\text{discrete}}(N)$ refer to the time of the N^{th} spike in the respective simulation. The correction for the accumulated lead or lag before determining the RMSE is performed to obtain a measure of the variability of the spike times independent of overall rate differences between simulation methods. It consists of warping the spike times in the discrete simulation by the factor $t_{\text{precise}}(N) / t_{\text{discrete}}(N)$ such that the last spike times considered coincide. Denoting the resulting spike times as t_{discrete}^* , the RMSE is then determined as $\sqrt{\sum_{i=1}^N (t_{\text{discrete}}^*(i) - t_{\text{precise}}(i))^2 / N}$.

3.2.3 Cortical microcircuit model

The model, taken from [Potjans and Diesmann \(2014\)](#), represents the neurons under 1 mm^2 of surface of generic early sensory cortex, organized into layers 2/3, 4, 5, and 6 (see [Figure 3.1](#)). It comprises 77,169 neurons connected via approximately 3×10^8 synapses, with population-specific connection probabilities based on an extensive survey of the anatomical and physiological literature. The connectivity is otherwise random, drawing both source and target neurons with replacement. Each layer contains one excitatory and one inhibitory population of LIF model neurons. We denote the eight populations by 2/3E, 2/3I, 4E, 4I, 5E, 5I, 6E, and 6I. The synaptic strengths J_{ij} are normally distributed with mean \pm standard deviation of $351.2 \pm 35.21 \text{ pA}$ for inhibitory source neurons and $87.8 \pm 8.78 \text{ pA}$ for excitatory source neurons except for connections from 4E to 2/3E, which have weights $175.6 \pm 8.78 \text{ pA}$. Transmission delays are normally distributed with mean \pm standard deviation of $1.5 \pm 0.75 \text{ ms}$ for excitatory source neurons and $0.75 \pm 0.375 \text{ ms}$ for inhibitory source neurons, truncated at the simulation time step. All neurons receive independent Poisson inputs with population-specific rates reflecting connections from adjacent cortex, other cortical areas, and subcortical regions. For further details we refer to [Potjans and Diesmann \(2014\)](#).

We run the simulations over 10 s of biological time with a time step of 0.1 ms, the original time step used for simulating the model and one of the time steps for which SpiNNaker is designed. In one set of simulations, the Poisson input of the original model is replaced by a DC drive corresponding to its mean current. The second set of simulations uses Poisson input, drawn independently for each simulation. A 1 s transient is discarded before analysis. The accuracy of the network simulations is assessed by statistical comparisons with NEST simulations with precise spike timing, which avoid synchronization artifacts ([Hanuschkin et al., 2010](#); [Morrison et al., 2007b](#)). Specifically, we compute the Kullback-Leibler divergence D_{KL} between three sets of smoothed histograms for each neural population; 1) single-neuron firing rates averaged over the simulation duration; 2) single-neuron coefficients of variation of interspike intervals (CV ISI); 3) Pearson correlation coefficients between spike trains binned at 2 ms (corresponding to the refractory time) from all disjoint neuron pairs within a subpopulation of 200 neurons, which provides a trade-off between statistical precision and computation time. For each neural population and dynamic variable, the histogram bin sizes are determined using the Freedman-Diaconis rule ([Freedman and Diaconis, 1981](#)) on the histograms for NEST with precise spike timing, bin size $= 2 \frac{\text{IQR}(x)}{\sqrt[3]{n}}$ with IQR the

interquartile range and n the number of observations. For each population and variable, we determine $D_{\text{KL}}(P||Q)$ where Q represents the grid-based NEST or SpiNNaker data, and P represents the data from NEST with precise spike timing. The histograms are first smoothed via Gaussian kernel density estimation using the `scipy.stats.gaussian_kde` function with bandwidth 0.3 s^{-1} for the rates, 0.04 for the CV ISIs, and 0.002 for the correlations. To avoid excessive contributions of low-probability bins due to division by vanishingly small numbers, bins where the normalized histograms have values smaller than 10^{-15} are ignored. We do not perform significance tests on the results, because we know the ground truth: the simulation methods differ. The comparisons between the simulation methods therefore focus on the sizes of the differences between the outputs.

3.2.4 Implementation

The network model was originally implemented in the native simulation language interpreter (SLI) of NEST. To allow execution also on SpiNNaker and to unify the model description across back ends, we developed an alternative implementation in the simulator-independent language PyNN (version 0.7; Davison et al., 2009). On SpiNNaker, this works in conjunction with the sPyNNaker software (Rowley et al., 2015).

The NEST (version 2.8; Eppler et al., 2015) simulations are performed on a high-performance computing (HPC) cluster with 32 compute nodes. Each node is equipped with 2 Intel Xeon E5-2680v3 processors with a clock rate of 2.5 GHz, 128 GB RAM, 240 GB SSD local storage, and InfiniBand QDR (40 Gb/s). With 12 cores per processor and 2 hardware threads per core, the maximum number of threads per node using hyperthreading is 48. The cores can reduce and increase the clock rate (up to 3.3 GHz) in steps, depending on demand and thermal and power limits. Two Rack Power Distribution Units (PDUs) from Raritan (PX3-5530V) are used for power measurements. The HPC cluster uses the operating system CentOS 7.1 with Linux kernel 3.10.0. For memory allocation, we use jemalloc 4.1.0 in this study (see Ippen et al., 2017, for an analysis of memory allocation in multi-threaded simulations).

The SpiNNaker simulations are performed using the 4.0.0 release of the software stack. The microcircuit model is simulated on a machine consisting of 6 SpiNN-5 SpiNNaker boards, using a total of 217 chips and 1934 ARM9 cores. Each board consists of 48 chips and each chip of 18 cores, resulting in a total of 288 chips and 5174 cores available for use. Of these, two cores are used on each chip for loading, retrieving results and simulation control. Of the remaining cores, only 1934 are used, as this is all that is required to simulate the number of neurons in the network with 80 neurons on each of the neuron cores. Cores are also used for simulating delays of greater than 16 time steps using a “delay extension” implementation, and for simulating the Poisson input noise. Each of these cores also simulates 80 units per core, i.e., 80 sources in the case of the Poisson sources, and the extra delay for 80 neurons in the case of the delay extensions.

The given number of neurons per core was chosen as this is the smallest number of neurons that can be simulated on each core for this particular network, whilst still being able to allocate routing keys to the neurons and having the SpiNNaker

routing tables fit within the hardware constraints of the machine with the current software implementation. The routing tables grow as the problem is distributed across more chips on the machine, as this requires additional paths to be made to allow the cores to communicate. The fact that the number of neurons per core cannot be reduced further also restricts the maximum speed with which the network can be simulated given the network traffic rates of the microcircuit; we find that we need to slow the simulation down by a factor of 20 from real time to maintain a 0.1 ms time step and be able to process all the spikes without overrunning the time allocated for each step. The simulation could otherwise run faster by having fewer neurons per core, and so less work to do on each core.

Whereas NEST represents parameters and dynamic variables as double-precision floating-point numbers, SpiNNaker uses the ISO draft s16.15 fixed-point arithmetic type.

All analyses are carried out with Python 2.7.9, using the Elephant package (version 0.2.1; [Yeegenoglu et al., 2016a](#)) for computing spike train statistics.

3.2.5 Performance benchmarks

We compare the power efficiency and runtime of the microcircuit simulations with DC input between NEST and SpiNNaker. For NEST, the strong-scaling efficiency of the simulation is assessed on the HPC cluster. In line with NEST’s hybrid parallelization strategy, we use one MPI process per compute node and OpenMP-based multi-threading within each process. Since NEST internally treats threads like MPI processes, they are also referred to as “virtual processes”, and the total number of virtual processes vp equals the number of MPI processes times the number of threads per MPI process ([Plesser et al., 2007](#); [Kunkel et al., 2014](#)). For the benchmark simulations, vp is increased from one single thread on one compute node up to the saturation of the full cluster. One compute node is first filled with one thread per core, first on one and then on the second processor, before threads are also assigned to the second available hardware thread of each core (hyperthreading). During the benchmark simulations, the power consumption of the compute nodes under load is measured with the PDUs. The active power is read approximately once per second remotely from the PDUs using the Simple Network Management Protocol (SNMP). To account for additional contributions to the overall power consumption, we furthermore estimate the usage of service nodes and switches (2 Ethernet and 1 InfiniBand) based on PDU measurements and data sheets.

Timestamps in the simulation scripts allow the identification of different execution phases, such as “network construction” and “state propagation”, and to relate them directly to the temporally resolved power measurements. For the phase during which the dynamical state of the neural network is propagated, we compute the average power consumption, the energy consumption and the energy per synaptic event. The energy consumption is obtained by integrating the measured power. The energy per synaptic event is defined as the energy divided by the total number of transmitted spikes N_{tspikes} . N_{tspikes} is composed of all occurring spikes times the number of outgoing connections from the respective sending neurons.

On SpiNNaker, a maximum number of neurons to be simulated on each core can be specified in the current software implementation. Populations of neurons are specified in the PyNN script, and a core runs a subset of the neurons from at most a single population; neurons from several populations are not combined. Thus if the network specifies a population of 100 neurons and a second population of 50 neurons and requests 90 neurons on a core, three cores will be used split as 90 on the first core, 10 on the second core and 50 on the third, despite the fact that the last 10 neurons of the first population could be combined with the second population within the given constraints. This is purely due to software engineering decisions; it is easier to keep track of the neurons if a core can only contain part of a single population. This could change in a future version of the software.

SpiNNaker boards are either single boards or combined in units of 3 boards. This makes it easier to deal with the coordinate space on the boards. The boards are physically placed into subracks of 24 boards, where each subrack has a backplane providing power to the boards and a 48-port switch providing networking from the outside world to the boards, with one 100 Mb/s Ethernet connection to each board and a second Ethernet connection to the management processor on each board. This external network is used purely for I/O interactions with the boards; network traffic generated during the simulation is passed entirely via the SpiNNaker network on and between the boards. The management processors allow each of the boards to be powered on independently and the links between boards to be turned on and off; thus within a single 24-board rack, boards are allocated either individually or in groups of units of 3 boards. The software is capable of working out an approximate number of boards required for the simulation and then requesting this allocation; in the case of the cortical microcircuit model simulation, 6 boards are requested.

Once a SpiNNaker machine has been allocated, it interacts with a host computer, which reads the machine configuration information (including the layout of the machine and any hardware issues such as faulty cores, chips and links), and works out how the neural network is to be run on the machine. Once this has been determined, the network data is generated and loaded, and the network is run.

For estimates of power consumption we connected a single 24-board rack to a consumer power measurement device at the mains socket, and ensured that there were no other users; thus only 6 of the 24 boards were ever active, with the other 18 remaining switched off. The power measurement device integrates the power usage over time providing an energy consumption in kWh at chosen points in time. An estimate of baseline power results from a measurement with the power on but with all the boards powered off. This allows eliminating the power usage of the rack itself, including the power consumption of the network switch, though not the cooling system, which is activated dynamically. A webcam pointing at the meter takes snapshots of the device at appropriate moments in the simulation setup, loading, execution, and result extraction phases to obtain readings for these stages. These measurements yield the total energy consumption for each execution phase in steps of $\Delta E = 0.01$ kWh. When computing the power for a phase of duration T , we propagate this measurement inaccuracy according to $\Delta P = 1/T \cdot \Delta E$. The software of SpiNNaker presently does not allow turning off individual cores and

chips. Therefore it is not possible to subtract the power consumption of unused hardware components on each board.

For both NEST and SpiNNaker, the respective machines are exclusively used for the simulations under consideration. There are no contributions to the overall energy consumption from other jobs running.

3.3 RESULTS

3.3.1 *Conceptual separation of biological time and wall-clock time*

Model neurons in the SpiNNaker system are updated at regular intervals in wall-clock time; this allows the simulation to be divided between several CPUs with independent timers, and still maintain reasonable synchronization across the system. A typical setting in previous studies is $h_w = 1$ ms. As SpiNNaker was originally designed for real-time operation, the interpretation of the biological model of the time span between two update events was considered to be identical to the wall-clock time passing between two updates $h_b = h_w$. However, from the point of view of a general simulation engine the two quantities are conceptually not identical. If the equations of a neuron model require updates in intervals of 0.1 ms in order to achieve the desired numerical accuracy, h_b can be interpreted as $h_b = 0.1$ ms. Without further changes to the parameters of the SpiNNaker system this means that the dynamics of the neural network now evolves 10 times slower than wall-clock time. However, the number of spikes occurring per second of wall-clock time is now reduced by a factor of 10: if a neuron model emits a spike within 1 ms with probability 1, the probability to emit a spike within an interval of 0.1 ms is 0.1. Therefore, if the limiting factor for reliable operation of the hardware is the number of spikes per second of wall-clock time, it might be possible to increase the clock-speed of the system by a factor of 10 ($h_w = 0.1$ ms) and recover real-time performance while safely staying within the limits of the communication bandwidth.

However, the communication bandwidth is rarely the limiting factor in simulations on SpiNNaker; we must also consider the CPU cycles required to process each spike received, and each synapse the spike activates. The design specifications of SpiNNaker assume a connectivity of 1,000 incoming synapses per neuron; the cortical microcircuit model has a value closer to 10,000, which means that the simulation must be slowed down further to accommodate the extra computation this requires, otherwise the synchronization of the simulation is liable to drift between the cores, and the results will be unpredictable and unreliable. If we assume that at the design specifications of SpiNNaker, the computation is split roughly as 10% or 20,000 CPU cycles per time step for neural updates and 90% or 180,000 CPU cycles for synapse processing, setting the time step to 0.1 ms means that 10 times more work is required for neural processing, giving 200,000 CPU cycles per ms of biological time but the amount of work for synapses remains constant at 180,000 CPU cycles per ms as the number of synaptic events per time step is reduced by a factor of 10. Setting the number of synapses per neuron to 10,000 means 10 times more work, or 1,800,000 CPU cycles per ms, leading to a total of 2,000,000 CPU cycles per ms of biological time for all the computation required. This can be achieved by slowing down the simulation by a factor of 10. In practice, there are

additional overheads in these processes, and we achieve reliable operation when $h_w = 2$ ms, meaning a slow-down of the dynamics compared to real time by a factor of 20.

3.3.2 Steps toward implementation on SpiNNaker

We iteratively refined the SpiNNaker interface for PyNN to extend the range of functions covered, and to match their syntax and functionality. Furthermore, we enabled running long simulations, where it was previously only possible to have short runs due to the memory filling up with the recorded data. We also implemented the NEST connectivity routine used by [Potjans and Diesmann \(2014\)](#) on SpiNNaker. The representation of multapses (multiple synapses between a pair of neurons; [Crook et al., 2012](#)) was already supported in the software, and all that was required was to generate the connectivity data using the host Python software. The large number of synapses, however, were more of an issue; for previous models, the synaptic data for the entire network was generated in advance of execution. The representation of this data in Python required a large amount of RAM on the host PC. We therefore modified the software to perform initial estimates of resource usage on the SpiNNaker machine based on statistical information about the PyNN connectors, including the multapse connector created specifically for this network. The software now generates the actual connectivity data lazily for each core, one by one, just prior to loading onto the machine, reducing the RAM usage on the PC by orders of magnitude. This method is also faster, reducing the data generation time from more than 8 hours to around 1 hour. This process could be parallelized to further reduce the data generation time, but this is not done in the current software due the restrictions of Python running in parallel.

The limited resources and efficiently implemented data structures within the SpiNNaker simulation environment enforce a limit of 1.6 ms for connection delays when the biological time step is 0.1 ms, due to the use of 16-element ring buffers for synaptic inputs (for an explanation of ring buffers in neural network simulations, see for instance [Morrison et al., 2005](#)). The mean delay of excitatory connections in the microcircuit model is 1.5 ms, so it was not initially possible to draw delays from a normal distribution with reasonable width. To resolve this issue, we implemented a “delay extension” mechanism, whereby delays > 1.6 ms were split into a multiple of 1.6 ms steps plus a remainder:

$$\begin{aligned} \text{delay}_{\text{extended}} &= \left\lfloor \frac{\text{delay}_{\text{total}}}{1.6 \text{ ms}} \right\rfloor \times 1.6 \text{ ms}, \\ \text{delay}_{\text{remaining}} &= \text{delay}_{\text{total}} - \text{delay}_{\text{extended}}. \end{aligned}$$

The extended delay is handled by a separate core. Knowing that the delay is a multiple of 1.6 ms allows up to 8 such multiples, or 12.8 ms of delay, to be simulated within the limited resources of this core. These can in principle be chained together allowing any delay, but a single additional core combined with the maximum 1.6 ms in the neuron model itself (a total of 14.4 ms) was deemed sufficient in this model. Besides enabling longer delays to be represented, support for distributed delays was added.

The synaptic weights of the model are of the order of 10^2 pA. On SpiNNaker, a single synapse is represented by a 32-bit number consisting of 8 bits for the target neuron index (allowing up to 256 neurons per core—the identity of the core is not stored in the synaptic word but in the routing tables, so neurons can have more than 256 targets), 4 bits for the delay (allowing up to 16 values, as described above), 2 bits for the synapse type (excitatory or inhibitory), and 16 bits for the synaptic weights; 2 bits are reserved to allow for increasing the number of synapse types. The 16-bit weight values are stored as fixed-point values, but the position of the binary point is adjusted to ensure that also the largest summed synaptic inputs occurring in the simulation can be represented. The reason for this adjustment is that when a spike is received on a core, the weight from each synapse is added into one of the 16-delay ring-buffers, each of which is also 16 bits in size; thus ideally the combination of the additions of several weights should not overflow the buffer. Additionally, an appropriate degree of precision is required to represent the weight values given; for example, using 8 bits for the decimal part of the numbers would lead to a precision of $1 / (2^8)$ nA ≈ 4 pA. This would give a fairly large error as a fraction of the synaptic weights in our simulations. The calculation of the position of the binary point is done by finding the maximum value likely to be added to any single ring buffer element. In previous implementations, this was done by simply adding together all the weights incumbent on each of the neurons and taking the maximum. This guaranteed that the ring buffer elements would never overflow but tended not to leave enough precision in the weights for correct representation, especially not in the case of the cortical microcircuit where there are a large number of connections, but a relatively low firing rate. This calculation was therefore updated to combine the statistics of the connectivity to get an approximate upper bound on the sum of the weights in any ring buffer element. This is done by firstly assuming an average input spike rate and choosing a scale factor σ to use for the overhead in the calculation. We treat the ring buffer elements equally, since, although the delays are distributed, the ring buffer element that represents the given delay from the current time step is moving as the time steps progress. Also the number of ring buffer elements is unimportant, since, regardless of how many there are, delay values will appear which could place the weight in any one of the elements. Thus, the calculation concerns any delay ring buffer element. We then look to the distribution of the weights combined with the timing of the spikes, since it is the arrival of a spike that causes a weight to be added to the ring buffer. For the purpose of determining the maximal resolution of the synaptic weights that allows the summed inputs to the neurons to be represented, we assume a Poisson distribution in the number of spikes arriving. This does not mean that inputs in the model indeed need to be Poissonian; the estimated resolution will work under moderate deviations from Poisson statistics, and the resolution can be decreased in case of highly synchronous input. Under the Poisson assumption, we can expect this same distribution in the addition of the weights to the ring buffer elements. Taking the mean and standard deviation of the weights, we can then compute an expected mean and standard deviation in the sum of the weights in any ring buffer element. Thus, we can approximate the maximum weight as a number of standard deviations above this mean value. Applying these assumptions, the

following closed-form solution is derived from standard results on means and variances of products of independent variables:

$$\begin{aligned}
 v_r &= n w_{\text{mean}}^2 \\
 U &= \text{round}[n + 3\sqrt{n}] \\
 v_w &= \frac{e^{-n} n w_{\text{var}} (-n^U + e^n \Gamma_{\text{inc}}[1 + U, n])}{\Gamma[1 + U]}, \\
 M &= n w_{\text{mean}} + \sigma \sqrt{v_r + v_w}
 \end{aligned} \tag{3.1}$$

where M is the expected maximum value over time in any of the delay ring buffer elements, which is calculated using n , the average number of expected incoming spikes in a time step; w_{mean} , the mean of the incoming weights; σ , the number of standard deviations above the mean for safety overhead (set to 5 here); w_{var} , the variance of the incoming weights; the gamma function Γ , and the incomplete gamma function Γ_{inc} . In the cortical microcircuit simulation, we take the expected rates within the network to be 30 spikes/s and use the known rates of the Poisson generators for calculating the synaptic weight resolution. Requiring that M from Equation 3.1 can be represented leads to weights with 6 or 7 bits for the integer part (respectively allowing summed input values with integer parts up to $2^6 - 1 = 63$ nA and $2^7 - 1 = 127$ nA) and 10 or 9 bits respectively for the fractional part, depending on the total number of incoming synapses to the population in question, since it is the summed weights in the ring buffer elements that determine the necessary resolution. In terms of weights of single synapses, with 10 bits for the fractional part of the number, the weight of 0.0878 nA would be represented as 0.0869140625, and with 9 bits for the fractional part, the representation is 0.0859375. By comparison, the nearest double-precision (64-bit) floating point representation of 0.0878 is 0.0878000000000000304201108747292892076075077056884765625, the nearest single-precision (32-bit) representation is 0.087800003588199615478515625; using half-precision floating point (16-bit) for the value would result in the value 0.08779296875 being used. Thus, the precision of the weight values on SpiNNaker is reasonable given the 16 bits available for use. During the single-neuron tests, the whole 16 bits were used as the fractional part of the number, leading to 0.0878 being represented as 0.087799072265625. Note that even with this calculation, the chances of overflow of the buffer are non-zero, and overflows will still likely occur in long-running simulations. Thus the software counts the number of times an overflow occurs and reports this to the user at the end of the simulation.

The communication network of SpiNNaker can support up to 6 million spike packets per second, but it does not cope well with all the traffic occurring within a short time window within the time step. This is exacerbated by the initial synchronization of the simulation engine at the code level, making all cores likely to send spikes at the same time. The neuron cores pause the processing of neurons and thus the sending of spikes whilst they are processing incoming spikes, so after the initial spikes there is some spreading of the network traffic over the time step occurring naturally. However, the Poisson noise generating cores and the delay extension cores have little to no inputs, so they have no automatic spreading of the sending of the spikes over the time step due to time spent processing incoming spikes. Furthermore, in the microcircuit simulations, the spikes tended to be

concentrated within a small window within the time step despite the desynchronization due to the processing of inputs. Without any correction, it is likely that all the network traffic will therefore occur within a small window at the start of the time step, and all cores will send simultaneously. To overcome this issue, each core is firstly given a random wait time at the start of every time step. This gives a basic offset so that the first network packet sent by a core is unlikely to be synchronized with that of the other cores. Given the maximum number of network packets to be sent by the application within a single time step, it is then possible to work out an expected minimum number of CPU clock cycles between sending packets within a time step; in the case of neurons each neuron can only send at most one spike per time step, but with Poisson sources, the likely maximum number of spikes per time step has to be calculated statistically. For example, with a 0.1 ms biological time step in real time, a 200 MHz CPU and 100 neurons being executed on the core, there are at most 100 packets to be sent each time step, and 20,000 clock cycles in which to send them, so there should be 200 clock cycles between the sending of packets. In practice, we spread the packets over half the time step, to allow the spikes to be processed at the receiving end, so there would be 100 clock cycles between packets in this example. If the execution arrives at a point at which a packet is to be sent, but the expected number of CPU cycles since the start of the time step has not passed, the core is simply made to wait until this occurs. On neuron cores, the neurons continue processing spikes during this pause, whereas the delay extension and Poisson generator cores have little else to do when this occurs. Prior to this change, there were quite a few dropped packets in the simulation. With this change no packets were lost during the simulation.

The independent Poisson input sources of the cortical microcircuit model also required modifications of the SpiNNaker software stack. The software was designed for an input spike rate to each neuron around 10 spikes/s and assuming each source neuron to have synapses onto multiple target neurons within each population. This means that there is not too much network traffic, and that each Direct Memory Access (DMA) performed when a spike is received retrieves multiple synapses from the SDRAM, increasing the overall efficiency of the transfer by reducing the overheads of each transfer. The one-to-one connectivity of the Poisson sources coupled with their high firing rate breaks both these assumptions. The revised software contains a heuristic in the placement algorithm which attempts to place one-to-one connected populations on the same chip where possible. This reduces the communication overhead, since only the internal network-on-chip is used to transfer the spikes between the Poisson sources and the populations they feed. Furthermore, the synaptic connectivity data for the one-to-one connected populations are now stored in local Data Tightly Coupled Memory (DTCM); a DMA to transfer the data is no longer required. The high input rates also mean that multiple spikes often need to be sent in a single time step, for which support was added.

As SpiNNaker has limited SDRAM and no other backing store, the storage of recorded data can become an issue, even for short simulations. The improved software overcomes the problem by calculating the maximum duration of the simulation before the SDRAM is filled by the recorded data. The simulation runs

for this period, pauses whilst the data are extracted from the machine, and then resumes. This repeats until the simulation has covered the required duration.

3.3.3 Comparison of single-neuron results between NEST and SpiNNaker

Figure 3.2 shows the results of the single-neuron tests, comparing the simulation output of grid-based NEST and SpiNNaker with that of NEST with precise spike timing, which provides a near-exact reference solution (cf. Section 3.2.1). The example membrane potential traces in Figure 3.2A show that both simulators achieve a high accuracy. SpiNNaker displays a slight lead and grid-based NEST a slight lag with respect to the precise solution that is visible especially at 1 ms resolution. These deviations are also apparent in the cross-correlation histograms of the binned spike trains (Figure 3.2B; for a comparison of different numerical solvers, see Rotter and Diesmann, 1999). The histograms for time step 1 ms contain multiple peaks due to the 1000 Hz rhythm imposed by the grid-constrained input spikes. The lag of grid-based NEST is due to the fact that spike times are always rounded up to the nearest grid point, not down (Morrison and Diesmann, 2008; Krishnan et al., 2017). The early spiking of SpiNNaker is likely to be due to the use of fixed-point numerical representations and the separation of the exponential decay of the synaptic inputs from the integration of the membrane equation, as shown for a single input spike in Figure 3.3. Since the fixed-point synaptic weights in these simulations are slightly smaller than the floating-point values, the increased post-synaptic response with the fixed-point representation must be due to the limited resolution of the other neuron parameters and variables. The separated integration leads to consistently higher values of the membrane voltage in response to the incoming spike, but for 0.1 ms time steps, the numeric type appears more influential than this separation (insets of Figure 3.3). The deviations for the lower-rate simulations are slightly smaller than those for the high-rate simulations, because the limited memory of the dynamics causes subthreshold traces with different initial conditions to converge under identical inputs. In terms of membrane potential correlations, both simulators perform well at 0.1 ms resolution (Figure 3.2C). At 1 ms resolution and low rates, NEST outperforms SpiNNaker in terms of membrane potential correlations, accumulated spike lead or lag, and spike time precision (Figure 3.2C–E). This may be explained by the greater contribution of subthreshold dynamics, which NEST integrates exactly, as compared to spiking dynamics at low rates. At 1 ms resolution and high rates, SpiNNaker outperforms NEST on all three measures (Figure 3.2C–E).

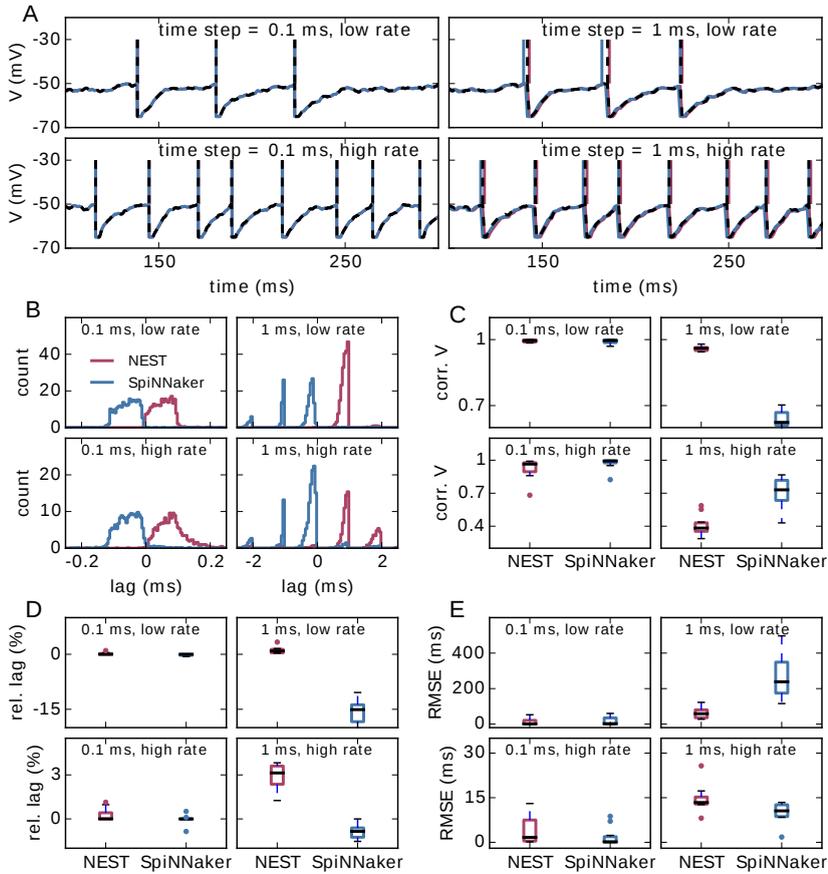


Figure 3.2: **Single-neuron tests for SpiNNaker and grid-based NEST simulations.** **A** Example membrane potential traces of leaky integrate-and-fire model neurons (parameters as in Table 3.1) receiving Poisson input with a “low rate” of 8,000 spikes/s and a “high rate” of 10,000 spikes/s (rows) for computation time steps 0.1 ms and 1 ms (columns). Red, NEST with spikes constrained to the grid; blue, SpiNNaker; black dashed curves, NEST with precise spike timing. Spike times are indicated by vertical lines. For time step 0.1 ms, all subthreshold traces overlap nearly precisely. **B** Average cross-correlation histograms over 10 simulations between binned spike trains from grid-based NEST (red) and SpiNNaker (blue) and those from NEST with precise spike timing. **C** Pearson correlation coefficients between membrane potential traces. **D** Accumulated fractional lead or lag in spike times. **E** Root mean square error of spike timing after correcting for accumulated lead or lag. All comparisons are with NEST with precise spike timing. Panels C–E: Thick black lines, median across 10 repeat simulations; boxes, interquartile range (IQR); whiskers extend to the most extreme observations within $1.5 \times \text{IQR}$ beyond the IQR.

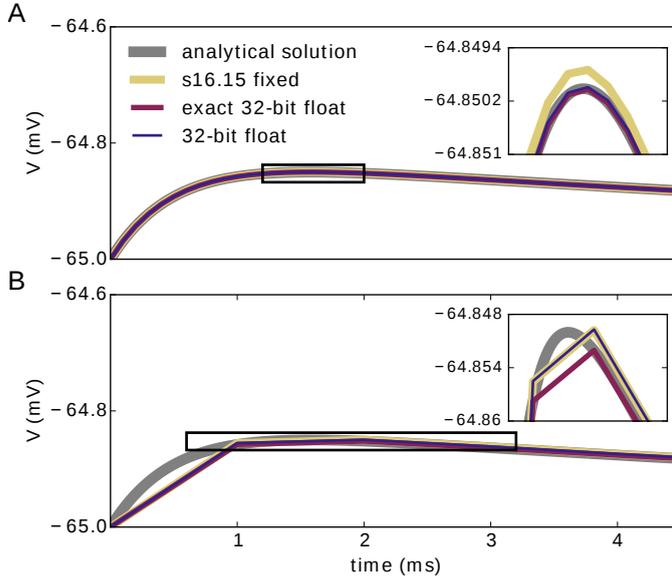


Figure 3.3: **Membrane potential excursion for different computation step sizes and numerical representations.** Panels show the peak region of the response of the model neuron on the SpiNNaker system to a single input spike (neuron parameters as in Table 3.1) for computation time steps 0.1 ms (A) and 1.0 ms (B). The ISO specification s16.15 fixed represents signed fixed-point numbers with 16 bits for the integer part and 15 bits for the fractional part; 32-bit float uses a single-precision ISO standard floating-point representation. Both s16.15 fixed and 32-bit float separate out the integration of the exponential decay of the synapses and the LIF neuron model (exponential integration), though each component uses a closed-form solution. Exact 32-bit float uses a single closed-form solution that encompasses both the exponential decay of the synapses and the LIF neuron model. This exact integration corresponds precisely to the analytical solution sampled at the integration time step. The s16.15 format leads to the same synaptic weights as with the 16-bit format in Figure 3.2. Double-precision (64-bit) floating-point numbers give membrane potential excursions that are visually indistinguishable from the single-precision results. Insets enlarge the membrane potential traces delineated by the black boxes.

3.3.4 Comparison of network results between NEST and SpiNNaker

Different slowdown factors were tested on SpiNNaker to determine the minimal slowdown factor at which no spike loss occurs in the simulation of the cortical microcircuit model. Based on the biological time step of 0.1 ms compared to the 1 ms design specification, this slowdown has to be at least a factor of 10. As explained in Section 3.3.1, additional slowdown is necessary to enable processing the high input rates to the neurons. At a slowdown factor of 20 with respect to real time SpiNNaker simulates this model without any spike loss. Also, the chosen precision for the synaptic weights prevents any overflows of the synaptic ring buffers from occurring in the SpiNNaker simulations. Therefore, differences between the NEST and SpiNNaker simulation results are only caused by floating-point vs. fixed-point numerical representations, exact subthreshold integration vs. separate integration of membrane voltage and synaptic inputs, and different random number generator seeds. Grid-based NEST, NEST with precise spike timing, and SpiNNaker produce closely similar spiking statistics, both for DC input (Figure 3.4) and for Poisson input (Figure 3.5). The raster plots of the spiking activity (Figure 3.4A–C; Figure 3.5A–C), a standard tool for the visual inspection of multi-channel spike data (Grün and Rotter, 2010), bring out the similarity. Despite the different initial conditions and different realizations of the connectivity, and also different input realizations in the case of Poisson input, distributions of average single-neuron firing rates (Figure 3.4D; Figure 3.5D), spiking irregularity (Figure 3.4E; Figure 3.5E), and correlation coefficients between binned spike trains (Figure 3.4F; Figure 3.5F) match closely between the three simulation methods and for all neural populations.

To assess how meaningful the differences between the simulation methods are, we compare these differences with those caused by the random number generator seeds alone. We perform three simulations of the microcircuit model with Poisson input for 10 s with NEST with precise spike timing with different random seeds for the connectivity, initial membrane potential distributions, and Poisson generators (Figure 3.6). In each case, we compare the distributions of rates, CV ISIs, and correlations, discarding a 1 s transient as before, in terms of the Kullback-Leibler divergence between the smoothed histograms. Since the simulations with the three methods (grid-based NEST, NEST with precise spike timing, and SpiNNaker) each use different random seeds, differences between the simulation results for these methods include the influence of the seeds, particularly in view of the finite length of the data. The results shown in Figure 3.6D–F indicate that the influence of the random seeds is comparable in size to the combined influence of the simulation method and the seeds. Thus, the simulation method itself contributes little to the variation in the dynamical properties of the microcircuit model, indicating in particular that SpiNNaker’s fixed-point numerics and approximations in the subthreshold integration do not compromise accuracy for networks of the given type.

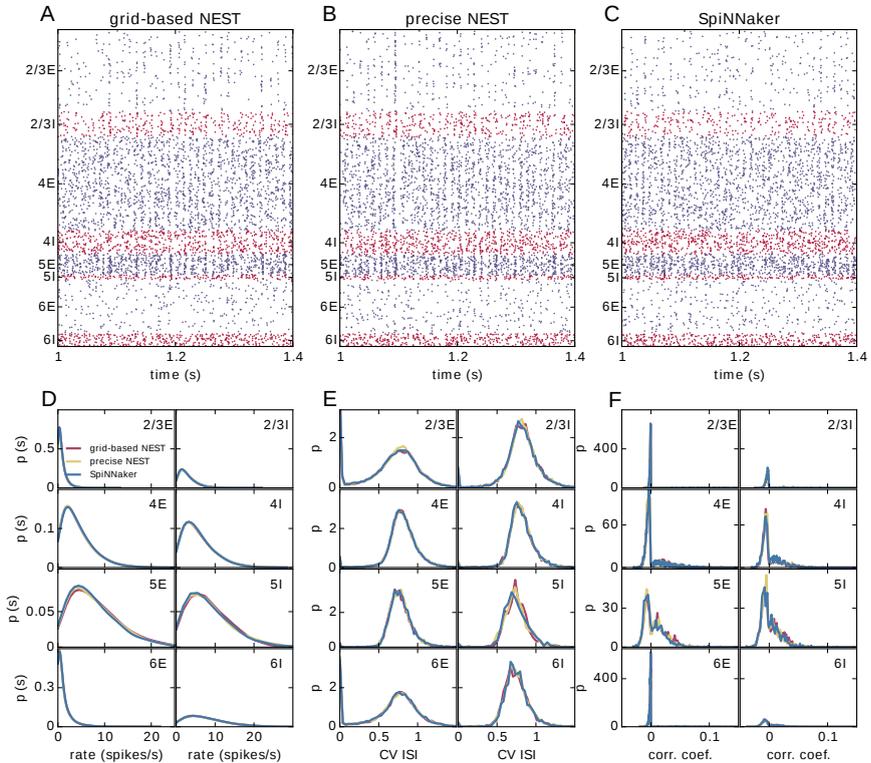


Figure 3.4: **Spiking output of the cortical microcircuit model with DC input.** **A–C** Raster plots showing spike times (dots) of excitatory neurons in blue and of inhibitory neurons in red. The spikes of 5% of all neurons (vertical) are displayed. **D–F** Distributions of spiking activity for each of the eight populations of the model. **D** Single-neuron firing rates of all neurons averaged over the last 9 s of the simulation. **E** CV ISI, a measure of irregularity of all neurons. **F** Correlation coefficients between binned spike trains for 200 neurons in each population. Histogram bin widths are determined by the Freedman-Diaconis rule.

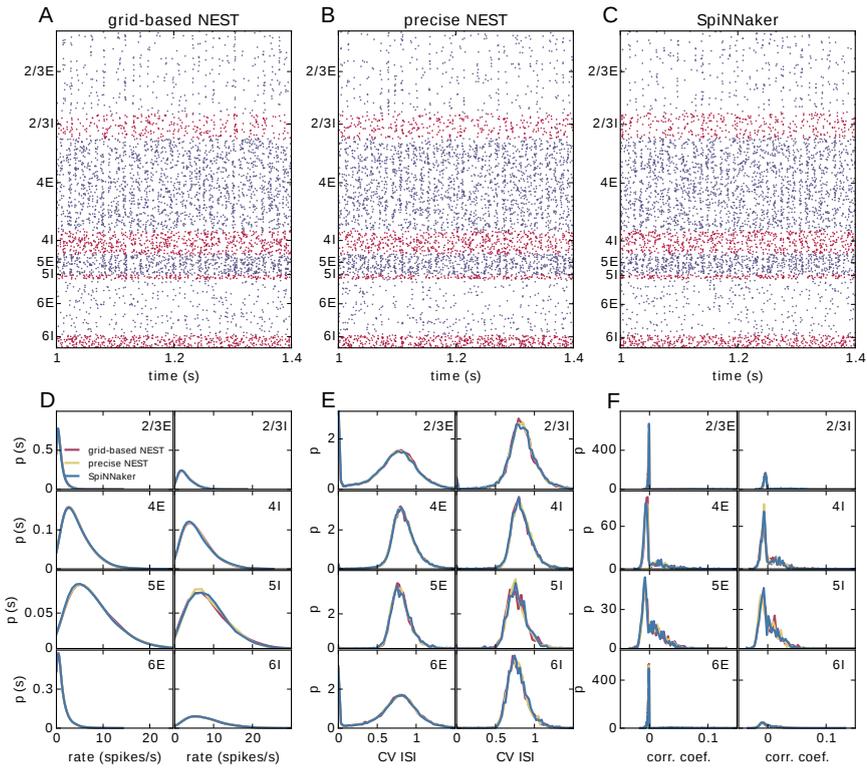


Figure 3.5: Spiking output of the cortical microcircuit model with Poisson input. Same display and parameters as in Figure 3.4.

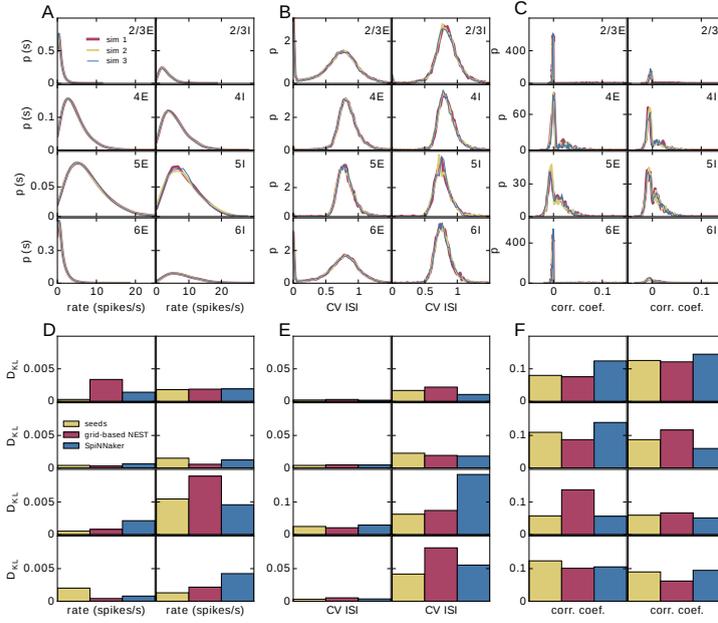


Figure 3.6: **Comparison of influence of random number generator seeds and simulation method.** **A–C** Distributions of dynamical properties of the microcircuit model for each of the 8 neural populations obtained using NEST with precise spike timing. Results for three simulations, each of 10 s duration discarding a 1 s transient, with different random seeds. **D–F** Kullback-Leibler (KL) divergences from NEST with precise spike timing as a reference, using results from NEST with precise spike timing and different random seeds (mean of KL divergences for two simulations with different seeds), and results from grid-based NEST and SpinNaker. **A,D** Time-averaged single-neuron firing rates. **B,E** Coefficient of variation of interspike intervals. **C,F** Pairwise correlations between binned spike trains.

3.3.5 Performance

Figure 3.7 shows results from measurements of the power consumption during benchmark simulations with NEST on one and two compute nodes of an HPC cluster. The simulations use an increasing number of threads on the single node, and all threads supported by the hardware on the two nodes (Figure 3.7A). The measured power consumption rises during script execution and we observe that it increases with the number of vp s whereas the required time decreases. In Figure 3.7B and C, we enlarge the traces for $vp = 48$ and $vp = 96$, respectively, and indicate the execution phases of the script. Prior to the execution of the script, the system exhibits a fluctuating baseline power consumption of the switched-on nodes; the baseline is higher for two nodes compared to one node. The phases “network construction” (red) and “state propagation” (blue) are the main phases as they refer to the setup of neurons and connections and the propagation of the dynamical state of the neural network, respectively. The color-coded areas for these phases have approximately the same size, indicating a similar energy consumption. The “writing output” phase transfers spike times from the simulation engine to file buffers after the dynamics has reached its final state. The corresponding PyNN function gathers data from all processes and uses only one thread per node for writing. Time spent otherwise during script execution is denoted in dark gray. These intervals correspond for instance to loading Python modules and setting simulation parameters before the network construction starts, and plotting the spiking activity after writing output. When the script has terminated, the timestamps are written to file, and after that, the power consumption returns to the baseline level.

We further spread the NEST simulations across up to all 32 compute nodes of the HPC cluster with 48 threads on each node ($vp = 1536$), shown in Figure 3.8. Panel A demonstrates the parallel scalability of network construction and propagation of the dynamics by showing the measured times together with the ideal linear expectations. The propagation time saturates at about three times the biological time. The jump in propagation time after $vp = 24$ coincides with the onset of hyperthreading. Network construction time continues to decrease over the full range of compute nodes but exhibits an intermediate increase starting at $vp = 15$ (see [Ippen et al., 2017](#) for a general discussion of network construction time). Figure 3.8B shows the power consumption averaged over the propagation phase as a total across all nodes used in the particular simulation. The change in slope at $vp = 48$ is due to the successive switching on of additional nodes. Integrating the power consumption traces over the propagation interval yields the energy consumption as depicted in Figure 3.8C. Due to the decrease in propagation time and the concomitant growth in power consumption with increasing vp , the energy consumption reaches a minimum at $vp = 96$, i.e., at two nodes. Thus, the hardware configuration requiring the minimal energy-to-solution is neither the one with the smallest number of hardware components involved nor the one with the shortest time-to-solution, but a system of intermediate size. The energy per synaptic event for the optimal configuration is $4.4 \mu\text{J}$.

Apart from the compute nodes, however, we also have to take other components of the cluster into account to estimate the total energy-to-solution. The HPC cluster

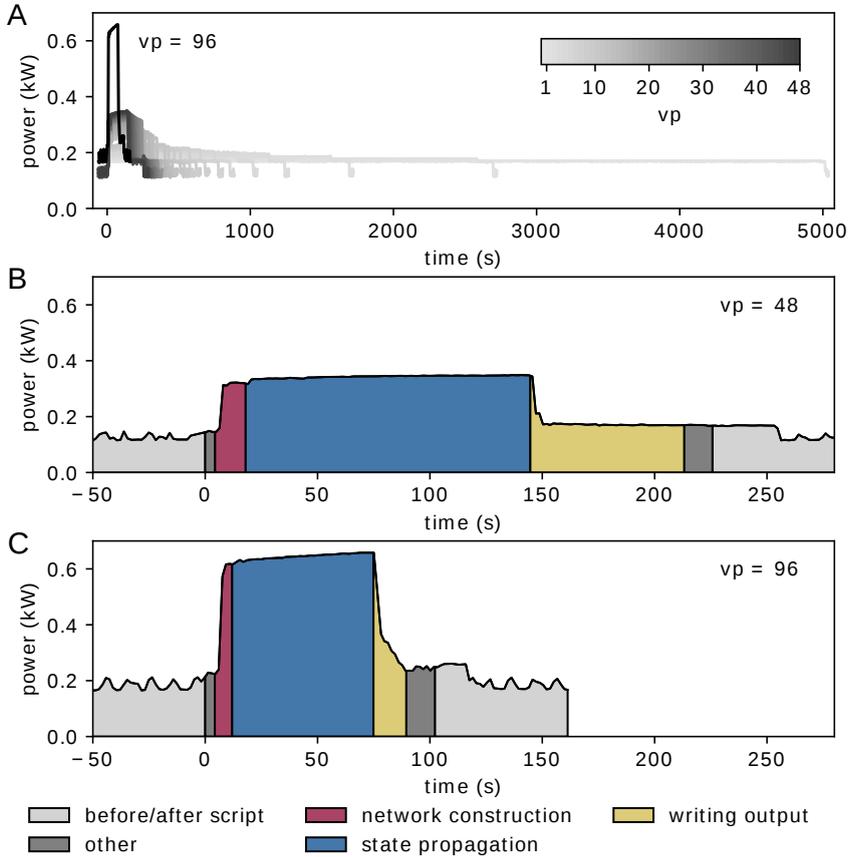


Figure 3.7: **Temporally resolved power consumption during microcircuit simulation with NEST on HPC cluster.** **A** Total power consumption as a function of time of a single compute node using 1–48 threads (gray code) and of two compute nodes with 48 threads per node ($vp = 96$, black). The curves terminate with the end of the simulations (for 10 s of biological time in all cases). **B** Power consumption in labeled execution phases of a simulation (legend) on a single compute node with 48 threads. **C** Execution phases for two compute nodes with 48 threads per node.

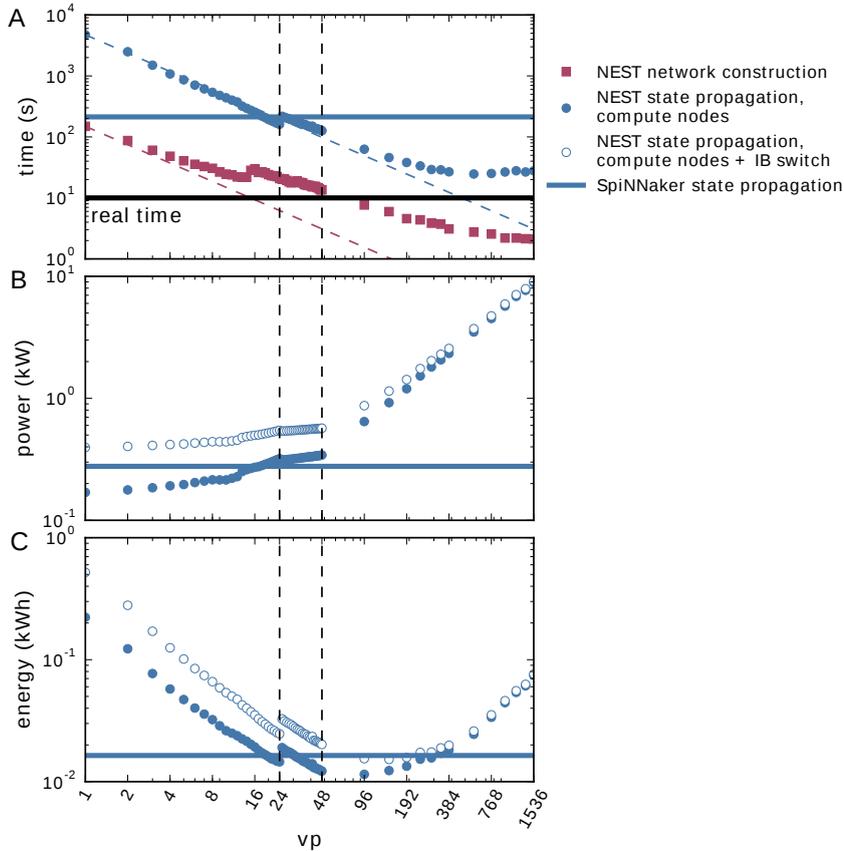


Figure 3.8: **Time to solution and energy consumption of NEST and SpiNNaker simulations.** **A** Duration of network construction (red square markers) and propagation of the dynamics (blue circular markers) vs. number of virtual processes for NEST simulations. Dashed lines represent ideal scaling. Black horizontal line indicates stretch of biological time simulated (10 s). **B** Mean power consumption during propagation for NEST simulations measured at involved compute nodes (filled markers) and with an additional power offset for the InfiniBand (IB) switch (open markers). **C** Energy consumption during propagation for NEST simulations measured at involved compute nodes (filled markers) and with an additional power offset for the IB switch (open markers). On a single node, vps bind initially to cores on one processor (up to $vp = 12$), then to cores on the second processor (up to $vp = 24$, left vertical dashed line), and finally to the second hardware thread on each core (up to $vp = 48$, right vertical dashed line). Blue horizontal lines in each panel indicate duration, power and energy of state propagation, respectively, for a SpiNNaker simulation. All panels in double-logarithmic representation.

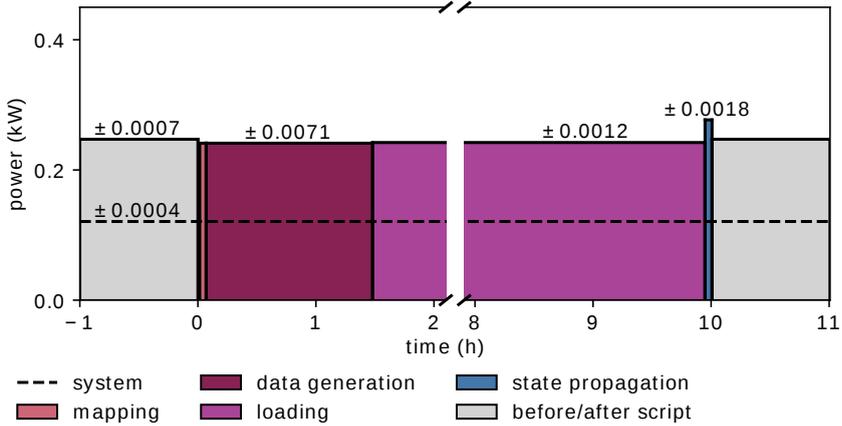


Figure 3.9: **Power consumption during microcircuit simulation with the SpiNNaker system.** The vertical axis shows the average power consumption during the color-coded execution phases. The power (black solid horizontal line segments) is computed from measurements of the energy consumption (colored areas) of the execution phases. The dashed horizontal line indicates the baseline measurement with the rack powered on but all 6 boards switched off. Numbers above black lines state the error in power estimation. Before and after the execution of the simulation script (light gray segments) boards are switched on and booted but idle. Power consumption during mapping (salmon) is set identical to consumption in data generation phase (raspberry). Propagation of the dynamical state by 10 s of biological time takes 200 s (blue segment). The respective power estimate results from the propagation by a longer stretch of biological time (1000 s) for increased accuracy.

requires two service nodes, with an estimated combined contribution of approximately 300 W, comparable to the base level of two compute nodes. Two Ethernet switches and one InfiniBand (IB) switch consume, based on their data sheets, a maximum of 64 W and 226 W, respectively. During the propagation phase, only the compute nodes and the IB switch are required. Figure 3.8B and C assess how an additional power offset accounting for the IB switch affects the power and the energy consumption as functions of vp . The increase in power consumption is crucial for small vp , but it is almost irrelevant for simulations across multiple compute nodes (large vp). We also observe that the minimum energy to solution shifts to a larger vp , and conclude that simulations become more efficient if distributed across more hardware. Including the contribution of the IB switch, the minimal energy per synaptic event is obtained at $vp = 144$ and equals $5.8 \mu\text{J}$. At this number of virtual processes, the simulation takes about 4.6 times real time.

Figure 3.9 illustrates the power consumption of the SpiNNaker system, derived from the measurements of the energy consumption for each execution phase (see Section 3.2.5). The background power usage caused by the network switch, the active cooling systems, and the power supply itself explains half of the total power consumption. As in the case of the HPC cluster we do not include any cooling of the room into the measure. The mapping phase is where the software of the host computer reads the machine configuration and then uses this description to

work out which parts of the neural network are to be executed on which chip, and the routes taken by network traffic that is to traverse the machine during simulation. Power consumption is mostly the same as in the idle phase, since the machine is only briefly contacted during this phase, with the rest of the work being done on the host computer. The data generation phase creates the data for each core; this includes the neuron parameters and synaptic matrices, as well as other SpiNNaker-specific data. Again, the machine is not in use during this phase, and hence could be turned off. The loading phase transfers the data generated on the host computer to the SpiNNaker machine. Although this requires communication with the machine, power consumption is still low, because only two cores on the machine are active at any time during this phase. These are the monitor core on an Ethernet-connected chip and the monitor core on the chip storing the data in memory. During the phase of state propagation, the power consumption increases significantly above the level of the idle state, reflecting the work done by the cores. The energy per synaptic event consumed during the propagation phase is $5.9 \mu\text{J}$.

Duration, power, and energy consumption of the propagation phase are included as horizontal lines in Figure 3.8 to facilitate a comparison with NEST. The energy consumed by the SpiNNaker simulation is close to the minimum energy of the NEST simulation for compute nodes and the IB switch, leading to a similar result for the energy per synaptic event on both systems.

Previous measurements of the SpiNNaker system indicated the approximate power usage of each chip to be 1 W when fully loaded, and the overhead for each board itself, excluding the chips, to be approximately 12 W. The 6 boards and 217 chips used in the present study thus predict a power consumption of $(6 \times 12 + 217) \text{ W} = 289 \text{ W}$. This is close to the 277 W measured during the state propagation phase, indicating that a calculation based on the number of chips and boards in use delivers a good estimate of the power consumption during this phase. With 48 chips per board, there are $48 \times 6 - 217 = 71$ unused chips, of which the power consumption is measured along with that of the active chips, but not taken into account in the back-of-the-envelope calculation. The fact that this calculation already gives a higher value than the measurement suggests that the power consumption of the unused chips is negligible.

3.4 DISCUSSION

On the example of a full-scale cortical microcircuit model (Potjans and Diesmann, 2014), the present work demonstrates the usability of SpiNNaker for large-scale neural network simulations with short neurobiological time scales and compares its performance in terms of accuracy, runtime, and power consumption with that of the simulation software NEST. With ~ 0.3 billion synapses, the model is the largest simulated on SpiNNaker to date, as enabled by the parallel use of multiple boards. The result constitutes a breakthrough: as the model already represents about half of the synapses impinging on the neurons, any larger cortical model will have only a limited increase in the number of synapses per neuron and can therefore be simulated by adding hardware resources. The synaptic time constants and delays of the model necessitate a shorter integration time step (here, 0.1 ms) than the original 1 ms design specification of SpiNNaker. The higher resolution is

achieved by a conceptual separation of biological time and wall-clock time. For the microcircuit model the current software stack of SpiNNaker requires the number of neurons per core to be set to exactly 80. This restriction is the result of the number of routing entries available on each router in the machine, combined with the current algorithm for assigning keys to the neurons. As a consequence of the combination of required computation step size and large numbers of inputs, the simulation has to be slowed down compared to real time. In future, we will investigate the possibility of adding support for real-time performance with 0.1 ms time steps. Reducing the number of neurons to be processed on each core, which we presently cannot set to fewer than 80, may contribute to faster simulation. More advanced software concepts using a synapse-centric approach (see [Knight and Furber, 2016](#)) open a new route for future work.

We assess accuracy by comparing grid-based NEST and SpiNNaker simulations with NEST simulations with precise spike timing, which provide a highly accurate reference solution. For the cortical microcircuit model, we consider firing rates, coefficients of variation of interspike interval distributions, and cross-correlations between binned spike trains. Which measures to use to quantify accuracy and which level of accuracy is considered to be acceptable has to be determined on an individual model basis from the acceptable range of desired model behaviors. For instance, for the microcircuit model, it is important that the simulations preserve asynchrony and irregularity of spiking, and differences in firing rates between the neural populations. Although models and their desired outcomes are diverse, these measures are also chosen because they characterize several fundamental aspects of single-neuron and population activity, and are therefore relevant to a wide range of models. Despite its fixed-point arithmetic, SpiNNaker is able to achieve comparable accuracy on these measures to that of NEST. Conversely, NEST executed on a high-performance cluster achieves speed and power efficiency comparable to the performance of SpiNNaker for some settings and in addition enabling a flexible trade-off between runtime and energy-to-solution. These results take into account that runtime and energy consumption should be assessed while controlling for simulation accuracy. For networks even larger than the cortical circuit considered, where runtime becomes strongly communication-dominated on traditional architectures, the asynchronous update of SpiNNaker may yet give it an advantage in terms of efficiency outside the scope of the present study due to an ability to simply expand the number of cores used by the simulation with minimal communication overhead. Thus, larger networks can be simulated in a weak scaling scenario where network size increases without increasing the rates of neuron state updates and synaptic events per neuron. For networks of any size, SpiNNaker is expected to yield accurate results as long as the simulation speed is chosen such that no spikes are lost and the resolution of the synaptic weights is sufficiently high. Further work is required to assess its scaling of runtime, memory, and energy consumption with network size.

The cortical microcircuit model consists of leaky integrate-and-fire (LIF) model neurons. To assess accuracy in a more controlled setting, we also consider single-neuron simulations. This reveals that grid-based NEST and SpiNNaker have similarly high accuracy at 0.1 ms time steps, with NEST slightly lagging behind and SpiNNaker slightly leading the precise solutions. This respective lag and lead can

be attributed to details of the neuron and synapse implementation, with NEST using exact integration (Rotter and Diesmann, 1999) for the subthreshold dynamics, whereas SpiNNaker uses fixed-point representations and a separation of the integration of the synaptic exponential decay and the neuron model. In terms of single-neuron dynamics, NEST performs relatively better at low spike rates, whereas SpiNNaker performs relatively better at high spike rates. The high accuracy of NEST for low rates may be due to the greater contribution of subthreshold dynamics in this condition.

Previous work has provided estimates of the power consumption of SpiNNaker executing spiking neural network models. Stomatias et al. (2013) instrument a 48-chip SpiNNaker circuit board to measure power consumption directly. They model locally and randomly connected networks of up to 200,000 Izhikevich model neurons and up to 250,000 LIF model neurons with a 1 ms time step and over a billion synaptic events per second, with a total board power consumption in the region of 30 W, arriving at a total energy per synaptic event of around 20 nJ. Subtracting baseline power, the incremental energy per synaptic event is found to be 8 nJ. Sharp et al. (2012) describe a small but detailed cortical model running with a 1 ms time step on a 4-chip SpiNNaker board instrumented to measure power. The model has 10,000 neurons and 4 million synapses, consuming just under 2 W, and the energy breakdown yields an incremental cost of 100 nJ per neuron per ms and 43 nJ per synaptic event, with a total energy per synaptic event of 110 nJ.

The present work measures the power consumption of the microcircuit model simulations on SpiNNaker and uses strong scaling with NEST on a high-performance compute cluster. At the optimal setting for NEST, with 144 virtual processes, the energy consumption of the compute nodes per synaptic event is $5.8 \mu\text{J}$, and for SpiNNaker the equivalent measurement is $5.9 \mu\text{J}$. There are several factors that contribute to the lower efficiency of SpiNNaker when running this model compared to the earlier studies, which mostly relate to the model being distributed sparsely over the SpiNNaker hardware, thereby causing baseline power to be amortized across many fewer synaptic events. The principal factors are: the use of a 0.1 ms time step, rather than the standard 1 ms; the biologically realistic number of $\sim 10,000$ synapses per neuron compared with the $\sim 1,000$ typical in neuromorphic models; and the highly distributed sparse connectivity of the biological model. With further software optimizations we expect such a network, with 80,000 neurons and 0.3 billion synapses, to map onto around 320 SpiNNaker cores—about half of a 48-chip board instead of the 6 boards used here—and to run in real time. With a 30 W power budget for half a board and 10 billion synaptic events over 10 seconds this yields 30 nJ per synaptic event, in line with the earlier total energy figures and two orders of magnitude below the present value. This ratio highlights the potential and the importance of further improvements of the software stack of the SpiNNaker system; this could include the use of the synapse-centric approach (Knight and Furber, 2016), which has been shown to accommodate the 0.1 ms time step and high synapse count better than the current mapping, but this is not yet available within the SpiNNaker tool flow. This would hopefully enable real-time operation of SpiNNaker during the network propagation phase, as well as reduce the number of cores and thus boards required for this simulation, and so result in a reduction in the power per synaptic event.

Mammalian brains consume about 6 kCal/day = 0.3 W per 1 billion neurons, of which roughly half is consumed by the cerebral cortex (Herculano-Houzel, 2011), and a substantial fraction is due to action potential signaling (Attwell and Laughlin, 2001; Lennie, 2003). In the human brain, with its 10^{11} neurons (Herculano-Houzel, 2012), cortex makes up close to 20% in terms of the number of neurons (Pakkenberg and Gundersen, 1997), so that we obtain 0.15 W per 2×10^8 cortical neurons. Assuming 10^4 synapses per neuron and an average spike rate of 4 spikes/s (Attwell and Laughlin, 2001), we arrive at an energy consumption of $0.15 \text{ W} / (2 \times 10^8 \times 10^4 \times 4 \text{ spikes/s}) = 19 \text{ fJ}$ per synaptic event. Since commonly used extracellular recording methods may miss a large fraction of neurons that are silent or nearly so, the average spike rate of cortex may actually be lower (Shoham et al., 2006). Taking an estimate of 0.1 spikes/s based on whole-cell recordings (Brecht et al., 2003; Margrie et al., 2002), we obtain $0.15 \text{ W} / (2 \times 10^8 \times 10^4 \times 0.1 \text{ spikes/s}) = 760 \text{ fJ}$ per synaptic event. These estimates indicate that with our computing systems and for the given model we are between about 7 and 9 orders of magnitude removed from the efficiency of mammalian cortex.

It is important that all contributing components are taken into account when comparing computing systems (noted by Hasler and Marr, 2013). In the present study, we have excluded the energy required for controlling the room temperature. In addition, we have ignored the contributions of the host computer and the Ethernet network to the energy consumption of the SpiNNaker simulations, and for the NEST simulations we have excluded the energy consumed by storage units, service nodes, and Ethernet switches in the derivation of the energy per synaptic event. The reason for ignoring these components is that they are in principle not needed during the phase in which the dynamics is propagated, except for data output, of which the contribution depends on the goal of the study. Similarly, it is in principle possible to power off the unused cores on the SpiNNaker boards, so the power usage of these cores could also be discounted. However, comparison of our measured power consumption, which includes both active and unused cores, with estimates based on active cores only, suggests that the contribution of the unused cores is negligible. Since we measured the power consumption of entire nodes on the HPC cluster, the measurements include cores not used for the NEST simulation up to the point where full nodes are assigned. This contribution is limited by the cores controlling their clock speed in steps depending on computational load, but could be discounted altogether.

Currently, both simulation engines require initial simulations to find the optimal setup in the first place, and these should be taken into account when evaluating their total energy consumption. In future, runtime models may be developed to estimate the optimal setup for a given network. This way, no additional simulations would be needed in order to determine at least a reasonable parallelization.

While the power consumption measurements described here only concern the phase in which the dynamics is propagated, on SpiNNaker the time taken to generate and load the network architecture is much longer, and needs to be addressed by future work. The present implementation already substantially reduces the time it takes to generate the connectivity. Work in progress includes developing the ability to generate the connectivity, which makes up the bulk of the data used

by the simulation, on the cores of SpiNNaker. This, as in the case of NEST (Morrison et al., 2005), has the potential to further reduce the network generation time through parallelization, as well as speeding up the loading of data by only transferring the parameters for the statistical generation of the synapses rather than the instantiated connections as is done now.

The current work considers networks of point neurons with static current-based synapses. In general, neural network models can contain more complex features, such as multi-compartment neuron models, conductance-based effects, and plasticity. Since such features increase the time required for neuron and synapse processing, they reduce the maximal rate at which the neurons on the SpiNNaker hardware can receive inputs and the number of neurons that can be mapped to a core while maintaining simulation speed. For instance, depending on the exact model and parameters chosen, simple pairwise spike-timing-dependent plasticity with additive weight dependence reduces both these quantities by a factor of 7 on SpiNNaker with the current software stack, and a factor of 2.5 with synapse-centric mapping of the network to the cores (Knight and Furber, 2016). More complex synaptic plasticity models with multiple dynamical variables like those described by Benna and Fusi (2016) can also be implemented but would further lower the number of neurons per core and their maximal input rates for a given simulation speed. One trend in computational neuroscience is toward ever larger-scale complex models (e.g., Traub et al., 2000; Lundqvist et al., 2006; Yu et al., 2013; Markram et al., 2015; Schmidt et al., 2016). Also such models can in principle be implemented on SpiNNaker; however, the scaling of the required resources and the corresponding simulation performance remain to be investigated.

Our comparison of SpiNNaker and NEST highlights concepts like accuracy, the influence of randomness, concreteness of use cases, and a common formal model specification that need to be considered when comparing systems of this sort. The concepts herein discussed facilitate the evaluation of other low-power platforms such as TrueNorth (Akopyan et al., 2015) and ROLLS (Qiao et al., 2015), and those that are similar to SpiNNaker but with other architectural features, such as described by Moradi et al. (2018).

Porting network models to dedicated hardware is a useful exercise to help identify requirements (the right product is built) and benchmark the results against existing simulation software (the product is built right). This gives us confidence that the co-design process in which we are engaged in the framework of a sequence of large-scale European consortia will continue to successfully guide us in the future. Close collaboration between hardware developers and computational neuroscientists ensures that the product can be used for realistic applications by its intended user community.

This chapter is based on the following publication:

Johanna Senk, Corto Carde, Espen Hagen, Torsten W. Kuhlen, Markus Diesmann, Benjamin Weyers (2018), VIOLA – A multi-purpose and web-based visualization tool for neuronal-network simulation output, *Frontiers in Neuroinformatics*, 12, 75, doi: 10.3389/fninf.2018.00075; *arXiv preprint arXiv:1803.10205v1*

Author contributions:

Under the supervision of Espen Hagen, Benjamin Weyers and Markus Diesmann, the author contributed to all parts of this publication including conceptual work and design of the visualization types. Corto Carde implemented the first version of VIOLA and the author co-developed the tool. Espen Hagen and the author jointly implemented simulation, preprocessing and analysis code for the example network and Espen Hagen incorporated LFP predictions. The author ran all simulations and created the figures. The author, Espen Hagen and Benjamin Weyers wrote the first draft of the manuscript; the other authors co-wrote the manuscript.

4.1 INTRODUCTION

One common technique to capture brain activity on the neuronal level is to record extracellular potentials in cortical tissue (Buzsáki et al., 2012; Einevoll et al., 2013a). The low frequency ($\lesssim 100$ Hz) part of the signal, often referred to as the local field potential (LFP), remains difficult to interpret as thousands to millions of proximal and distal neurons contribute to the signal (Kajikawa and Schroeder, 2011; Lindén et al., 2011; Łęski et al., 2013). From the high-frequency band ($\gtrsim 100$ Hz), however, one can detect sequences of spikes, the transient extracellular signatures of action potentials in single neurons nearby the recording electrode. The number of reliably identified neurons (through spike sorting, Quiroga, 2007) per recording session is low compared to the number of neurons in vicinity of the recording device, even if the experiment is performed with hundreds or more electrode contact points (Einevoll et al., 2012). The Utah array from Blackrock Microsystems¹, for example, resolves with 10×10 electrodes on 4×4 mm² little more than a hundred distinct neurons. Also optical methods for measuring neuronal activity have seen continuous improvements. As recently demonstrated, non-invasive three-photon fluorescence microscopy facilitates functional imaging at high optical resolution as deep as 1 mm (Ouzounov et al., 2017). While the method simultaneously images a comparably large number of neurons, the recordings lack the temporal resolution to reliably detect individual action potentials. Ouzounov et al. (2017) record from as many as 150 neurons in mouse hippocampal stratum pyramidale within a field of view of 200×200 μm^2 .

The rapidly improving parallel recording technology increases the need for suitable analysis methods for high-dimensional and dynamic data streams. Nevertheless, the recordings will remain to be characterized by a massive undersampling for some time. Therefore, detailed full scale models of the cortical tissue are required to understand the microscopic dynamics (van Albada et al., 2015) and to relate the microscopic activity to mesoscopic measures like the LFP. For this program to succeed, neuroscientists not only need to analyze model data in the same way as experimental data, but to explore data sets with orders of magnitude more channels than experimentally available.

Networks of model neurons incorporating varying levels of biophysical and anatomical detail reproduce a number of features of experimentally obtained spike trains. For networks of point- or one-compartment neuron models, this list of features includes irregular spike trains (Softky and Koch, 1993; van Vreeswijk and Sompolinsky, 1996; Amit and Brunel, 1997b; Shadlen and Newsome, 1998), asynchronous spiking (Ecker et al., 2010; Renart et al., 2010; Helias et al., 2014; Ostojic, 2014), correlation structure (Gentet et al., 2010; Okun and Lampl, 2008; Helias et al., 2013), self-sustained activity (Ohbayashi et al., 2003; Kriener et al., 2014b), realistic firing rates across cortical lamina (Potjans and Diesmann, 2014), single-neuron spiking activity of different cell types (Izhikevich, 2003; Kobayashi et al., 2009; Yamauchi et al., 2011) and responses under ‘in vivo’ conditions (Jolivet et al., 2008; Gerstner and Naud, 2009). Relating point-neuron network activity to population signals such as the LFP is, however, not straightforward. Approximations (see Mazzoni et al., 2015) or forward-model based schemes (Hagen et al., 2016a)

¹ <http://blackrockmicro.com>

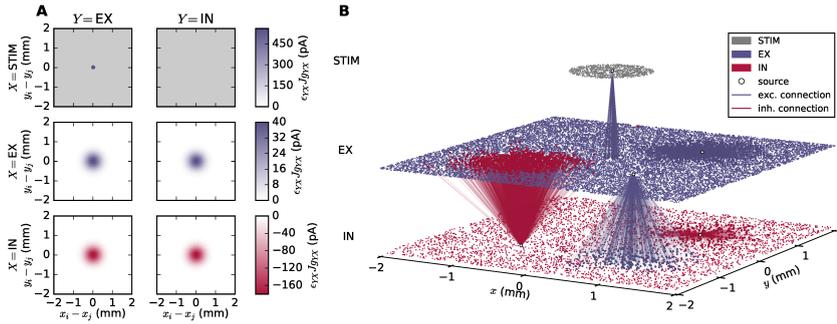


Figure 4.1: **Geometry and connectivity of a layered point-neuron network.** **A** Schematic illustration of distance-dependent network connectivity using connectivity pattern tables (Nordlie and Plesser, 2010). Each row represents source populations $X \in \{\text{STIM}, \text{EX}, \text{IN}\}$, and each column target populations $Y \in \{\text{EX}, \text{IN}\}$. The color coding in each image shows the connection intensity between presynaptic neurons j and postsynaptic neurons i located in (x_j, y_j) and (x_i, y_i) with origin $(0, 0)$ at the center. The connection intensities are defined as the product between pairwise connection probabilities $\epsilon_{YX}(r_{ij})$ and synapse strengths $g_{YX}J$ for each respective connection. Gray values denote connection intensities of zero. **B** Illustration of one network instantiation with randomly drawn neuron positions and outgoing connections from a subset of neuronal units. The colored dots represent individual units at their (x, y) -coordinates. Gray dots denote units in a stimulus (STIM) layer, blue dots excitatory (EX) units, and red dots inhibitory (IN) units. Blue and red lines denote excitatory and inhibitory connections respectively, from a source unit (white circles) onto neurons within the same or another layer.

are required to bridge the gap to experimental electrophysiological data which predominantly reflects population activity.

The focus of this study lies on visualization methods for activity of spatially extended neuronal network models. Incorporation of spatial structure is a prerequisite for models aiming to explain experimentally observed spatiotemporal patterns of activity (Rubino et al., 2006; Denker et al., 2011; Sato et al., 2012; Muller et al., 2014; Townsend et al., 2015). Such models have an arrangement of neurons in one-, two- or three-dimensional (1D, 2D or 3D) space and connection rules which typically depend on the distance between (parts of) the neurons (Mehring et al., 2003; Coombes, 2005; Yger et al., 2011; Bressloff, 2012; Voges and Perrinet, 2012; Kriener et al., 2014b; Keane and Gong, 2015; Rosenbaum et al., 2017). Although we primarily focus on model data, the same visualization methods can be applied with experimentally recorded data.

We here consider an example spiking point-neuron network consisting of an excitatory (EX), an inhibitory (IN) and an external stimulus (STIM) population. EX and IN units are positioned randomly within square domains while STIM units are randomly positioned within a circle at the center. A schematic representation of the network connectivity is shown in Figure 4.1A. We use connectivity pattern tables (Nordlie and Plesser, 2010) for source populations X (rows) and target populations Y (columns). The images indicate the ‘connection intensities’ for each connection,

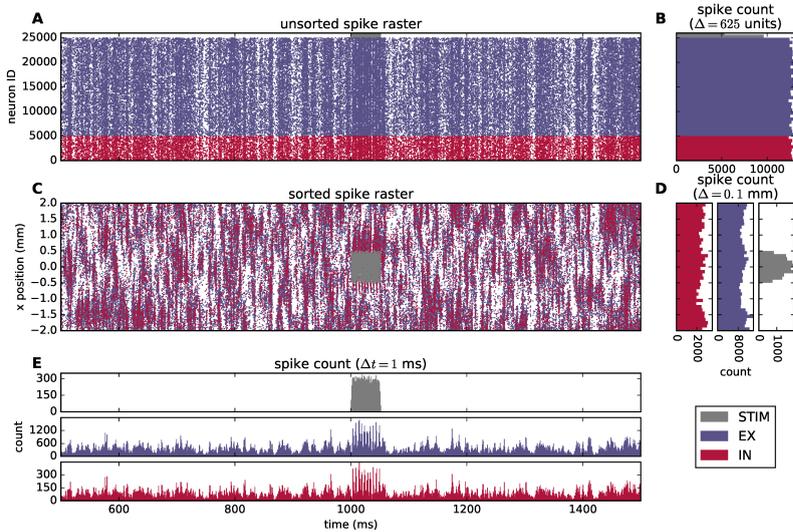


Figure 4.2: **Spiking activity of a layered point-neuron network model.** A Spike raster plot for STIM (gray dots), EX (blue dots) and IN (red dots) units from a simulation of the network instantiation depicted in Figure 4.1B. Each individual dot corresponds to a unit ID vs. spike time; only the spikes of every fifth neuron are shown in the raster. The color coding for each population is reused in the subsequent panels. B Spike count histogram across units in each population, calculated using a bin width of 625 units, sorted by neuron index j . C Sorted spike raster where dots correspond to the spatial location (projected onto the x -axis) and spike times of each unit. The raster-plot density is diluted as in panel A. D Spike count histogram across spatial bins with a width $\Delta l = 0.1$ mm. E Spike count histogram for each population across time, computed using a temporal bin width $\Delta t = 1$ ms.

defined as the product between averaged pairwise connection probabilities $\epsilon_{YX}(r_{ij})$ and synapse strengths $g_{YX}J$. The distance between a source and a target neuron is denoted by r_{ij} . Pairwise connection probabilities decay with horizontal distance between EX and IN units according to a Gaussian profile, while STIM units only connect locally to the EX population restricted by a cut-off radius. The geometry of one network instantiation is depicted in Figure 4.1B. EX (blue dots), IN (red dots) and STIM (gray dots) units are placed in separate layers. The distance dependency is illustrated by outgoing excitatory connections (blue lines) from single units in the STIM and EX populations and outgoing inhibitory connections (red lines) from single units in the IN population.

The visualization of neuronal activity data poses challenges due to the high dimensionality and time dependence of the data. Historically, electrophysiological data have been recorded from few electrodes or from many electrodes with undefined relative and absolute spatial coordinates (see, for example, the pioneering work of Krüger and Bach, 1981). This is not an essential limitation for recordings within the local cortical network where a neuron can form a synapse with any other neuron and there is little spatial organization. Furthermore, the fundamental interaction in a neuronal network is considered to be a dynamics on a graph; nodes solely interact via the edges of the graph. In this picture the spatial embedding of the graph is arbitrary as the dynamics is not constrained by the dimensions of physical space. Therefore, neuroscientists developed concepts for visualizing correlation structure and time dependence of neuronal activity in multi-channel recordings in ignorance of spatial properties. A temporal segment of activity of our example network is visualized in Figure 4.2. Panel A is the spike raster diagram or dot display in use for decades (explained in Abeles, 1982). Each row shows the spike train of one neuron where spike times are marked by dots. The rows either represent data of the same neuron in several trials or, as here, data of simultaneously recorded neurons in a single trial (Grün, 1996, Figure 6.2). The spike trains are vertically arranged by neuron ID and in addition color coded by population. The spike raster highlights global features of network activity and generations of neuroscientists have been trained to visually inspect these diagrams. For example, network synchrony appears as a stripy vertical pattern even if individual neurons only rarely participate in an individual synchronous event. The activation of the stimulus population is reflected in the other populations as an increased density of the dots. Epping et al. (1984) extend the concept of the raster diagram by assigning a unique color to the dots of a neuron. In this way multi-channel activity observed over multiple trials can be superimposed. Panel B shows spike counts along the temporal axis over neuronal units demonstrating that the per-neuron spike count is similar for the excitatory and the inhibitory populations. The spike count along the vertical axis in panel E is called the post-stimulus time histogram (PSTH, Perkel et al., 1967a) originally computed for an individual neuron observed over several trials. Later the display was also called peri-stimulus time histogram. Here the histogram is computed over simultaneously recorded neurons in a single trial. The display uncovers the fluctuations of population activity in time.

The development of adequate visualization concepts for multi-channel neuronal data is an ongoing endeavor (Allen et al., 2012). The cross-correlation function

(Perkel et al., 1967b) exposes the time-averaged relationship between the spike times of two neurons. The snowflake diagram generalizes the concept to three neurons (Perkel et al., 1975; Czanner et al., 2005). Gravitational clustering (Gerstein et al., 1985; Gerstein and Aertsen, 1985 and reviewed in Chapter 8 of Grün and Rotter, 2010) attempts to identify the emergence of correlated groups of neurons, so called cell-assemblies, and the temporal dynamics of the changing membership of individual neurons in such groups without averaging over trials. The joint peri-stimulus time histogram (JPSTH, Aertsen et al., 1989) generalizes the cross-correlation function to visualize the dynamics of the correlation between the spikes of two neurons in reference to a stimulus. Later, Prut et al. (1998) used the idea to investigate the occurrence of spatiotemporal patterns in the spike trains of three neurons, where “spatio” refers to the abstract space of neuron IDs not physical space. Because of the difficulties in determining statistical significance Grün et al. (2002) restricted the scope to patterns in the space of IDs and for visualization mapped significant events, so called unitary events, back into the spike raster diagram. Progress in the theory of neuronal networks showed that propagating spiking activity due to the stochastic nature of neuronal activity is likely to exhibit in each instance a random sub-pattern of spikes of some superset of neurons. Therefore, Schrader et al. (2008) designed a matrix spanned by binned ongoing time in both dimensions where matrix-elements represent the cardinality of the intersection set of the neurons spiking at the two respective time bins. With color-coded cardinality, in this matrix repeatedly occurring propagating spiking activity appears as a diagonal feature. Later an assessment of statistical significance was added (Torre et al., 2016a). Kemere et al. (2008) employ multi-channel recordings to construct the time course of a multi-dimensional vector of spike rates. A suitable projection to a lower dimensional space reveals differentiable trajectories of network activity depending on the experimental protocol (reviewed in Cunningham and Byron, 2014). Another line of work attempts to cope with the combinatorial explosion of patterns in multi-channel spike trains while maintaining sensitivity by the construction of a pattern spectrum: a two-dimensional histogram spanned by the number of spikes in a pattern, called pattern complexity, and the number of occurrences of the particular patterns (Gerstein et al., 2012; Torre et al., 2013).

Figure 4.2C modifies the spike raster diagram in panel A to arrange the spike trains on the vertical axis according to the x -coordinate of the position of the emitting neuron in physical space. In contrast to the regular spike raster, we observe inhomogeneous spatiotemporal features in network activity. The spatially binned spike counts along the temporal axis in panel D, however, do not reveal any unexpected structure. Thus, taking into account one coordinate of the neurons in physical space hints at some organization of neuronal activity. Nevertheless, a higher-dimensional analysis seems to be required to uncover its origin, as the features of spatiotemporal patterns can only be conjectured in 2D.

The emergence of planar wavelike spiking activity in 2D networks was shown by Voges and Perrinet (2012, Figures 3–5), but a 2D spatial visualization of the data could not faithfully capture intermediate mixed patterns such as rings and spiraling waves. Temporal snapshots of spatial activity show the evolution of patterns, as seen in Mehring et al. (2003, Figure 5), Yger et al. (2011, Figures 2,12), Voges and Perrinet (2012, Figure 6) and Keane and Gong (2015, Figure 1). Series

of such snapshots combined in an animation or movie can be informative, but require settings to be defined beforehand, leaving only little room for interactivity. With such non-interactive visualization methods, crucial decisions about a figure or an animation thus have to be made before a sufficient intuition about the data exists. Flexible, interactive visualization techniques identifying relevant dynamical features present in the data would have the potential to avoid the tedious and time-consuming loop of refining parameters and regenerating snapshots, animations or movies. In addition, high-dimensional and multi-modal data demand adequate workflows for analysis, from raw data to statistical measures, where interactive visual analysis methods can play a major role. It is for example essential to get a basic understanding of the datasets to better decide what statistical methods to use for more elaborate analysis. Furthermore, interactive visualization allows for explorative data analysis, including dimensionality reduction of complex datasets, highlighting of data points, and direct changes to visualization parameters.

For the development of supportive visual analytics tools, **Shneiderman (1996)** introduced the so-called ‘information-seeking mantra’. It describes the steps of common visual analysis workflows: “overview first, zoom and filter, details on demand”. The first step provides a superficial ‘overview’ of the data. In the second step, ‘zooming’ into the dataset allows the user to get a more detailed view on a chosen data subset. Application of ‘filters’ implies a change in dimensionality of the data or the extraction of particular features. Finally, **Shneiderman (1996)** proposes that visualization tools should enable the user to access all details of selected data points.

To not restrict the user to only one visual representation of the data, **Wang Baldonado et al. (2000)** established the concept of ‘coordinated multiple views’. Coordinated multiple views is a paradigm for the implementation of visual analysis applications that “use two or more distinct views to support the investigation of a single conceptual entity” (**Wang Baldonado et al., 2000**, page 110), and has been applied in various contexts (see for example **Roberts, 2007**). Basic coordination of views addresses selection operations (e.g., whether to display only a subset of the data) and also includes immediate control over animated frames (e.g., animation time step and playback speed for time-resolved data). In addition, each view may have an exclusive (view-specific) set of user controls and settings.

The activity exhibited by our example network is characterized by a non-trivial interplay between neuronal populations resulting in non-stationary activity in time and space. The neuroscientist needs to identify the propagation of spiking activity within and across individual layers over time and space, and simultaneously observe population activity measures such as the LFP. This is an opportunity to exercise the concepts by **Shneiderman (1996)** and **Wang Baldonado et al. (2000)**. Visualization in most cases focuses on a specific aspect or hypothesis to be tested by analyzing the corresponding data. Therefore, for each task the neuroscientist combines a different set of views. Sometimes particular views are not among the regularly used ones provided by the visualization framework but are created ad hoc specifically for the research question or the experimental protocol. Therefore the analysis software environment needs to facilitate fast prototyping of visualizations and an interface to a computing programming language used in the scientific domain. This focus on a specific aspect under investigation by the neuroscientist

necessarily entails an individual level of reduction or aggregation of the data. A particular visualization realizes this preprocessing of the data with methods like binning of data points in time or space, or by filtering out a certain subset of parameters of each data point. For instance, the spike raster plot in Figure 4.2A displays the individual spikes of all neurons whereas the bar chart in Figure 4.2E visualizes the total number of spikes per time step. The visualization abstracts from the spikes of individual neurons and turns the focus to the whole population. On the one hand the visualization simplifies interpretation by presenting less detail, on the other hand the reduction increases the chance of wrong or inaccurate conclusions. Historically, Vaadia et al. (1988, Figure 4 middle) illustrate a potential misinterpretation of the PSTH due to variability in the onset of the neuronal response: a neuron observed over multiple trials exhibits in the PSTH a smooth increase in spike rate, whereas the raster plots show in each trial an abrupt increase in spike density with a variable onset. Grün et al. (2002, Figure 8) demonstrate how such misalignments can propagate to measures of statistical significance: with respect to one trigger event the data show surplus spike synchrony simply due to non-stationarity of spike rate, whereas with respect to another trigger the rate is stationary and no excess synchrony is detected. However, if a multi-view approach is implemented that combines various visualizations, more than one aspect of the data (more than one visual representation of differently processed data) can be inspected simultaneously and can be put into relation. By interactive addition and removal of certain views, this process can be made flexible and thus address changes in the analysis goals or to consider findings during the analysis. Finally, we require a solution that allows for integration with platform-independent web-based technologies to keep the accessibility of the tool as high as possible.

A variety of coordinated multi-view applications for the interactive analysis of activity data has been described in literature, which generally follow the information-seeking mantra. For models of neuronal systems, the NEURON simulation environment (Carnevale and Hines, 2006) provides a graphical user interface based on a modified version of the discontinued InterViews library in addition to scripting in HOC and Python (Hines et al., 2009). The software itself offers the possibility of drawing multiple concurrent windows with dynamic and interactive plots of voltages, currents, morphology shapes and phase planes that are updated while simulations of single-neuron models or neuron networks are running. 3D visualization is not directly supported, but NEURON's Python bindings also allow running simulations to interact with modern visualization software, as for example incorporated by NeuronVisio (Mattioni et al., 2012) that relies on the OpenGL-accelerated Mayavi visualization toolset (Ramachandran and Varoquaux, 2011). The simulation software for large-scale neuronal network models NEST² (NEural Simulation Tool, Gewaltig and Diesmann, 2007) does not provide built-in interactive visualization. The original authors state this in their first report (Diesmann et al., 1995) as a design decision based on two considerations. First, in 1995 the life time of graphics frameworks and libraries appeared much shorter than the envisioned period of relevance of a simulation code. Thus, only a software stack with a strict separation of levels would ensure platform independence and sustainability of NEST. Second, a basic idea of the project is to contribute to a software

² <http://nest-simulator.org>

environment for ‘in virtu’ now often called ‘in silico’ experiments (restated in [Diesmann and Gewaltig, 2002](#)). In this concept, the authors state, simulated data and experimental data should be analyzed with the same analysis tools to maximize comparability and reproducibility. At the same time researchers at the department of Physiology and the Center for Neural Computation of the Hebrew University in Jerusalem started to work on an integrated analysis and visualization platform based on Open Inventor³ called Neural Data Analysis (NDA) but the project was abandoned with the advent of MATLAB ([Vaadia, 2017](#)). Recently [Nowke et al.](#) took on the challenge to develop a simulator independent visualization platform for brain-scale neuronal networks. The VisNEST ([Nowke et al., 2013, 2015](#)) framework visualizes the spiking activity of multi-area network models (using as an example [Schmidt et al., 2017](#)) in a virtual environment. The time-resolved activity data is mapped onto a 3D brain model. This enables the researcher to interact with the model in 3D to expose otherwise occluded parts of the brain and to relate brain activity to anatomy. In a different view, a dynamic 3D graph represents time course of spike exchange between different cortical areas. These representation of spatial information can be combined with classic charts such as spike raster plots. The tool does presently not account for the spatial organization of activity within brain areas. Apart from VisNEST, other standalone interactive multi-view applications have been developed for simulated spiking data, for instance SNN_{3D}Viewer ([Kasiński et al., 2009](#)) and ViSimpl ([Galindo et al., 2016](#)). SNN_{3D}Viewer focuses on 3D neuronal networks by visualizing individual neurons and their connections schematically, including interactive control over the 3D visualization (navigation, scale). ViSimpl combines a 3D particle-system-based visualization of the simulated neuronal network using color coding for the activity, supplemented by a set of data charts for single neurons and populations. Geppetto⁴ is a web-based modular platform for visualization and simulation of complex biological systems including spiking neuronal networks. Unlike the visualization concepts along which these tools have been developed, we here focus on concepts that expose the spatial organization of neuronal activity in layered networks and scale to signals from several square millimeters of brain surface.

Beside the aforementioned softwares applicable with spike data, general-purpose multi-view frameworks exist with different design goals and contexts of use (see [Roberts, 2007](#)). One generic high-level example is GLUE⁵, a Python and OpenGL-based multi-view framework. Another powerful framework is the now neglected OpenDX⁶. In contrast to the GLUE toolset, we here aim at web-based visualization. Easy access to libraries of common plotting functions and methods (scatter, line, surface plots etc.) is provided for most common programming languages (C++, Python, MATLAB, etc.). Nevertheless, a large amount of time and resources is still required to construct fully interactive visualization tools adhering to the principles outlined by [Shneiderman \(1996\)](#) and [Wang Baldonado et al. \(2000\)](#). Including interactivity and time synchronization between different visualizations may be

³ <https://www.openinventor.com>

⁴ <http://www.geppetto.org>

⁵ <http://glueviz.org>

⁶ <http://www.opendx.org>

demanding in terms of software design and development time, however, existing plotting libraries can be used to realize the individual visualizations.

As a reference implementation of our conceptual study, we introduce the interactive visualization tool VIOLA, an open-source, platform-independent and lightweight web-browser application. The tool is designed for initial visual inspection of massively parallel data generated primarily by simulations of spiking neuronal networks similar to the example network illustrated in Figure 4.1. VIOLA is designed around the information-seeking mantra and the concept of coordinated multiple views. 2D and 3D visualizations support the exploration of neuronal activity across space and time. The software can display raw spiking output as well as spatiotemporally binned data that may represent instantaneous spike counts gathered from nearby groups of neurons. Spike and LFP data can be displayed simultaneously, thus allowing for a multi-modal analysis.

The next sections are organized as follows: In Section 4.2 we present different visualization types and their application. Subsequently, in Section 4.3, we describe their implementation in the visualization tool VIOLA, the example network model and the phenomenological model for the LFP signal. Finally, in Section 4.4, we conclude our work and discuss general limitations of frameworks for explorative visualization and potential future developments.

4.2 RESULTS

For the analysis of data, static figures can help to highlight certain characteristics of the data or show results relevant for a particular hypothesis. However, static figures hamper an exploratory analysis of data as the adaption of data filters, visualization parameters, or changes in the perspective (in case of 3D visualization) require a re-rendering of the figure resulting in a very slow visual analysis process. Interactive visualization tools tackle these shortcomings by offering multiple views on the same data simultaneously, for example by projecting the data across different dimensions. This allows the user to investigate data at different levels of detail, and to adapt visualization parameters in a dynamic and explorative manner as the rendering of the visualization is continuously updated. Throughout this section, we use the spike output of the point-neuron network introduced in Section 4.1 as an example to demonstrate appropriate visualization types in an interactive and multi-view framework. Neurons in this network are placed in 2D sheets, and connections are drawn using distance-dependent probabilities between pairs of neurons. The model represents spatially heterogeneous neuronal activity across a $4 \times 4 \text{ mm}^2$ cortical sheet. As we here focus on visualization methods in VIOLA, we refer the reader to Sections 4.3.4 and 4.3.5 for the details on our network implementation in NEST (Kunkel et al., 2017), Python-based preprocessing steps and predictions of a mesoscopic population signal, the local field potential (LFP). We next describe in detail the different views of VIOLA and their use cases.

4.2.1 Views of VIOLA

VIOLA incorporates two conceptually different visualization types with two separate ‘views’ each. The first visualization type (view 1 and view 2 in Figure 4.3)

focuses on instantaneous snapshots of data across space. The second visualization type (view 3 and view 4 in Figure 4.4) shows time series of data. We first present the visualizations of preprocessed data described in Section 4.3.1. Views 1–3 may also be used to visualize raw data (non-preprocessed) as shown in Figure 4.5.

4.2.1.1 View 1: 2D spike-count rate

The 2D spike-count rate view (Figure 4.3A) shows instantaneous activity data in separate sub-panels for individual populations. The values in each panel correspond to the instantaneous spike-count rate v_β in one discrete spatiotemporal bin indexed by $\beta = (l_x, l_y, k)$ in our preprocessed data format (l_x and l_y denote spatial bin indices along the x - and y -axes, and k denotes a temporal bin index). In this format, each spike event is added to the corresponding spatiotemporal bin as described in detail in Section 4.3.1. The color bar denotes bin values in units of spike counts per second (spikes/s) and is shared between all sub-panels. This view provides a side-by-side comparison of the spatially resolved activity in each individual population. For a larger number of populations (than shown here), it is, however, difficult to relate the activity of one population to another population by visual observation. This problem can however be amended by combining multiple population activities in a single scene.

4.2.1.2 View 2: 3D layered spike-count rate

In the 3D layered spike-count rate view (Figure 4.3B), we combine the activity of all network layers in one 3D-animated scene. The view incorporates the possibility to show also other activity measures, for example the population LFP. The layers in view 2 correspond to the different sub-panels for each population in view 1. Different populations are here assigned unique colors. We chose to illustrate instantaneous spike-count rate v_β by dynamically sized cubic boxes. The box sizes are by default scaled such that their volumes are proportional to v_β at each time step, thus low-activity bins may still be visualized simultaneously with high-activity bins.

View 2 offers multiple possibilities for interactive adaptation of the visualization. As suggested by the information-seeking mantra (Shneiderman, 1996), the user can manually select which part of the data to show, for example by switching on or off individual layers that may occlude visibility of activity in other layers or by setting the horizontal x - and y -limits of the layers. It is also possible to reduce the opacity of the colored boxes or to scale their side length linearly. The camera can be set either to an orthographic or perspective-corrected projection mode. Dependent on the projection mode, the camera can be moved freely and allows for zooming, panning and rotating the scene. One can easily reset the camera to its default position by the click of a button or select different preset camera positions such as on top or to the side.

The major benefit provided by view 2 over view 1 is the possibility to visually relate the activity in one layer to other layers as all layers are drawn in the same 3D scene. As box volumes are computed from instantaneous spike-count rate values, this view brings the attention of the user to spatial regions of the network with high local activity. While views 1 and 2 offer flexible visualizations of instantaneous activity across space, we next consider scenes capable of showing time-series data.

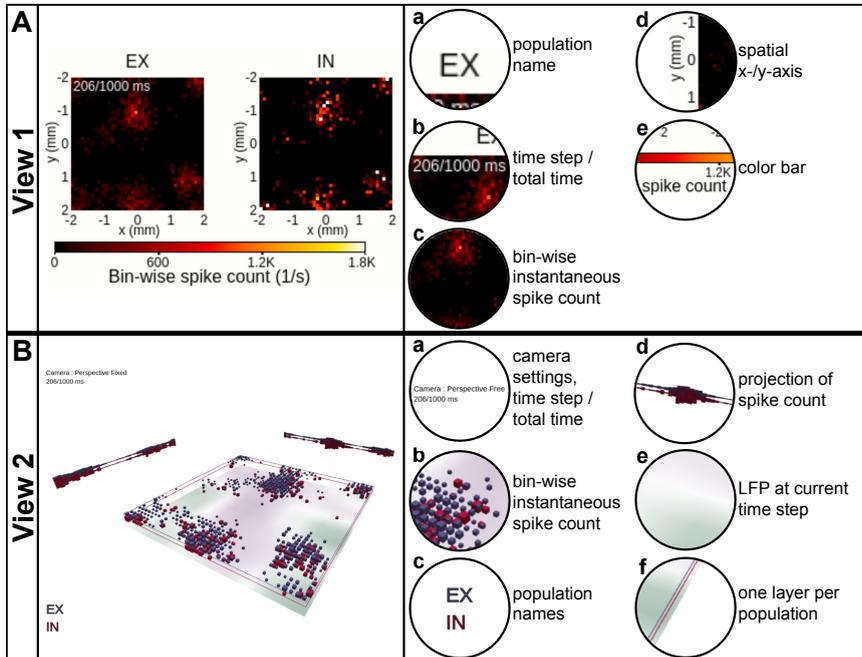


Figure 4.3: **View 1: 2D spike-count rate.** Panel A shows the instantaneous spike-count rates, defined as the number of spikes per second occurring within a spatiotemporal bin, using 2D image plots spanning the spatial x - and y -axes of the network layers. One separate image plot is created for each network population and denoted by the population name (EX, IN). The color map and corresponding color bar for instantaneous spike-count rate values are shared among all populations. In this and subsequent panels, we show the screen shot of the view itself to the left and highlight its components to the right. **View 2: 3D layered spike-count rate.** Panel B combines the data shown in Panel A in a single 3D scene by stacking the different population data on top of each other. The size of each cubic marker denotes the magnitude of the corresponding bin-wise instantaneous spike count, and its position corresponds to the spatial positions of the respective bin. Unique colors are assigned to each layer as indicated by the population names. Projections of the spike counts along the x - and y -axes are displayed towards the corresponding edges. The optional bottom image plot layer shows the spatial variation in an LFP-like signal at the present time step of the rendering loop.

4.2.1.3 View 3: Scrolling spike-count rate plot

The scrolling spike-count rate plot (Figure 4.4A) is a time-series representation of the data that neglects spatial features of the network and its activity. It shows the time evolution of the total spike-count rate v_k (black line), defined as the sum over the spike-count rates of all spatial bins and populations divided by the number of bins, together with the relative rate of each individual population (colored stacked plot). v_k is defined as $v_k \equiv 1/(L_x L_y) \cdot \sum_X \sum_{l_x} \sum_{l_y} v_{\beta}$ with $\beta = (l_x, l_y, k)$ and where L_x and L_y denote the number of bins along the x - and y -axes and neuronal populations are denoted X . The per-population spike-count rates can therefore be inferred by multiplication of the total rate with the fraction of spiking in individual sub-populations. The color coding of each population corresponds to the one used in view 2, but can also be read from the bar to the right. The plot is centered on the current time step (vertical white line indicator) when scrolling through data points in the animation. It allows for interactive change of the width of the visible time window and also permits to manually select or deselect individual populations to be displayed. View 3 provides a temporal overview of the data and allows to identify time intervals of interest, for example, due to an external perturbation.

4.2.1.4 View 4: Scrolling spike-count rate iso-surface plot

The instantaneous spike-count rates of our example network are time-series activity data with 2D spatial structure. In order to visualize such data without loss of dimensionality, a 3D representation is in general required (unlike for instance view 3). The scrolling spike-count rate iso-surface scene in Figure 4.4B simultaneously shows the evolution of network activity in space (as in views 1 and 2) and time (as in view 3). The rate iso-value surfaces of each population is rendered using the computer-graphics algorithm ‘marching cubes’ (Lorensen and Cline, 1987). The color coding of the individual populations matches the coding used in views 2 and 3. In terms of user interactivity, the user can set the threshold (isolation) for the surfaces. Furthermore, the user can select which populations to show, vary their opacity level, apply a temporal offset to individual populations, change the width of the time window, and take full control of the viewpoint in the 3D scene as in view 2.

4.2.1.5 Raw data views

In addition to visualizing spatiotemporally binned, preprocessed data, views 1–3 can also be used with raw simulation output files formatted according to the description in Section 4.3.1. With raw file output, view 1 (Section 4.2.1.1, Figure 4.5A) displays for each animation time step a square marker for each spike time t_j^k at the spatial location (x_j, y_j) of neuron j in population X . The square marker color is population specific. Likewise, view 2 (Section 4.2.1.2, Figure 4.5B) shows boxes of equal size for each spike event, colored according to population. The view allows, as with precomputed spike-count rates, to show spiking activity in each population in the same scene. We here show a snapshot of the spiking activity in perspective mode, and top-down. The main interactive feature of views 1 and 2 incorporated with raw data files is the option to reduce the neuron density to be displayed. As

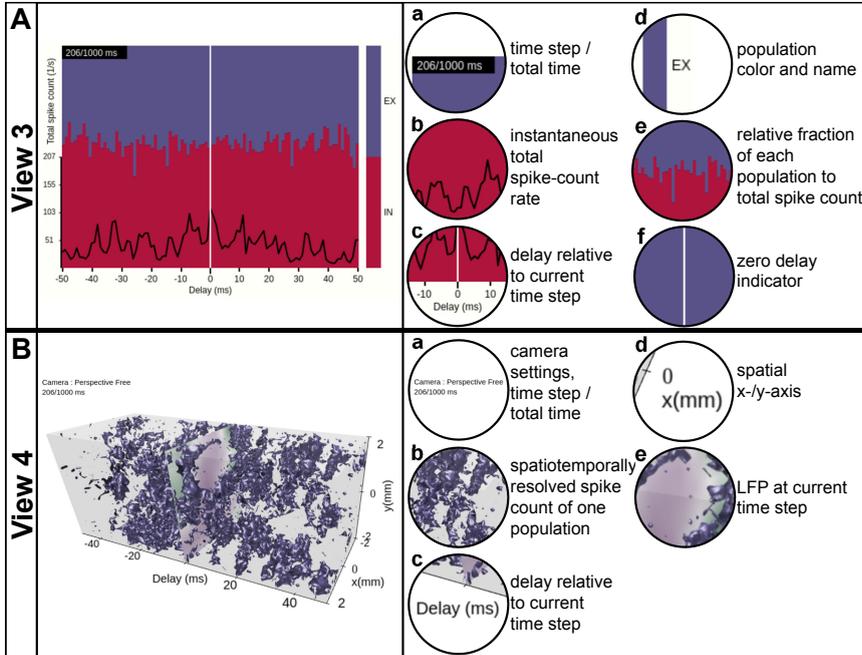


Figure 4.4: **View 3: Scrolling spike-count rate plot.** Panel A is a time-series representation of the data across a user-variable time interval around the present time step of VIOLA's rendering loop, indicated by the vertical white line. The instantaneous total spike-count rate summed over all populations is drawn using a black line. The relative fraction of spikes of each population to the total spike count is shown as a stacked, normalized histogram. The population outputs are color coded as in view 2. **View 4: Scrolling spike-count rate iso-surface plot.** Panel B provides a 3D representation of the spatiotemporally resolved spike-count rates of one selected network population across a user-variable time window. The spike-count rate is rendered as a closed iso-value surface in the color of the respective population and extends in both space (x - and y -axes) and time (delay axis). The present time step in the visualization is indicated by a time lag of zero on the time-delay axis. At zero time delay we also show the LFP signal corresponding to the present time step in the animation.

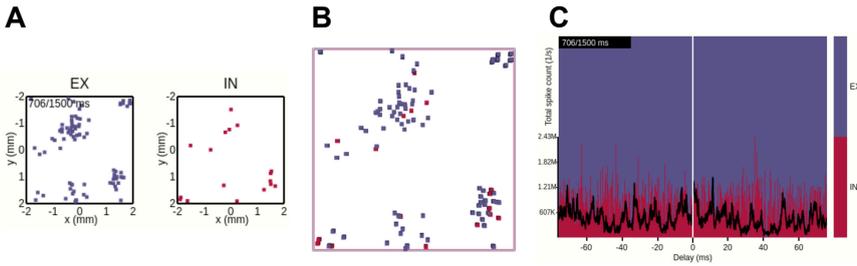


Figure 4.5: **Views 1–3 applied with raw data formats.** **A** View 1: Each dot corresponds to a single spike event of a neuronal unit at its spatial location in the network. **B** View 2: Perspective top-down view onto stacked layers of Panel A. **C** View 3: The stacked plot has the temporal bin-size of the simulation resolution. The black trace shows the spike count of all neurons per time bin (in units of spikes/s).

with preprocessed data visualization with view 2, individual layers can be activated/deactivated, one can switch between the orthographic and perspective viewing modes, and the camera can be positioned freely. View 3 (Section 4.2.1.3) applied to raw data is shown in Figure 4.5C. The temporal bin size of the animation is then equal to the simulation time step dt (one spike therefore results in the spike count rate $1/dt$ in units of spikes/s for that instant). The total spike count (black line) is summed over all neurons, in contrast to the mean over per-bin rates as in the case of the preprocessed data. The relative spike count per population is shown as a stacked plot normalized by the total amount of spikes in each temporal bin.

4.2.2 VIOLA use case

Numerical model development representing a physical system comprises implementation, simulation, analysis as well as comparison, validation and verification steps. Such model development is important for building hypotheses and aiding interpretation based on experimental data and observations. We here demonstrate how the views described above can be integrated with the development of a spiking point-neuron network model. For this purpose, we can hypothesize that transient external input to a layered spiking point-neuron network model with distance-dependent recurrent connections results in propagating spatiotemporal activity. We wish to assess the spatial extent and temporal duration of the network response to external perturbation and whether or not the unperturbed network state is recovered. In this use case we do this assessment by visual inspection prior to any detailed numerical analysis, focusing on the importance of coordinated multiple views (Wang Baldonado et al., 2000; Roberts, 2007) and Shneiderman’s information-seeking mantra (Shneiderman, 1996). For our hypothesis above we will therefore use our implementations of views 1–4 in VIOLA to rapidly analyze our network activity. We show that a combination of the different views is needed to assess the relevant aspects in the data, which is the evoked response to a network perturbation.

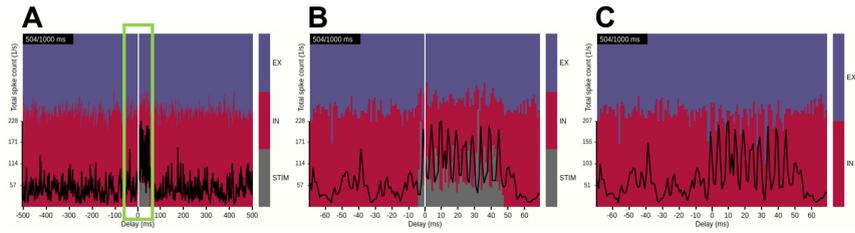


Figure 4.6: **Identifying a time interval of interest with view 3.** **A** The spike-count rate summed across spatial bins and the relative contributions by populations EX, IN and STIM shown for a time window of ± 500 ms around the current time step of the animation. **B** Same as Panel A, but with a narrower time window of ± 70 ms (indicated by the green frame in panel A), highlighting the activation of the STIM population and corresponding network response. **C** Same as panel B, but with the STIM contribution turned off.

The layered point-neuron network illustrated in Figure 4.1B consists of an excitatory (EX) and an inhibitory (IN) neuronal population plus one stimulus population (STIM). Each neuron is placed randomly within square sheets. EX, IN and STIM units are connected using distance-dependent rules as illustrated in Figure 4.1A. The connectivity is periodic across boundaries (torus connectivity). The detailed network description is given in Section 4.3.4. The main simulation output is spike times of individual neurons, neuron locations and a synthetic LFP signal (see Section 4.3.5 for details). Our initial preprocessing steps and corresponding data formats are described in Section 4.3.1.

4.2.2.1 Temporal features of evoked network activity

We first focus on ongoing activity of the network in the time domain, as provided by view 3. This view implements a scrolling spike-count rate plot which ignores spatial information. Interactive control of the view's time window allows for quick identification of events of interest from the full duration of the simulation (Figure 4.6A). One such event that is clearly differentiated from other ongoing activity is the activation of the external STIM population at the animation time step of 500 ms. Pausing the animation at 504 ms and zooming in onto the event (Figure 4.6B) allows for a detailed look on how the total spike-count rate (black trace) increases and oscillates while the stimulus is active, and confirms that the stimulus duration was 50 ms. The color-coded stacked histogram reveals that during stimulus activation a large relative fraction of spike events is contributed by the STIM units (gray), while the relative fraction generated by the recurrently connected EX (blue) and IN (red) units is reduced. We may also conclude that the transient onset of the stimulus results in temporally brief imbalances between excitatory and inhibitory populations in the network as the relative rate of the inhibitory population drops with regular intervals during the stimulation period. The imbalances occur at the stimulus onset and during each period of the resulting network oscillation (from recurrent interactions between excitatory and inhibitory neurons). This network spike-rate imbalance is even more pronounced when the STIM activity is hidden (Figure 4.6C). We note, however, that the rate balance averaged over the

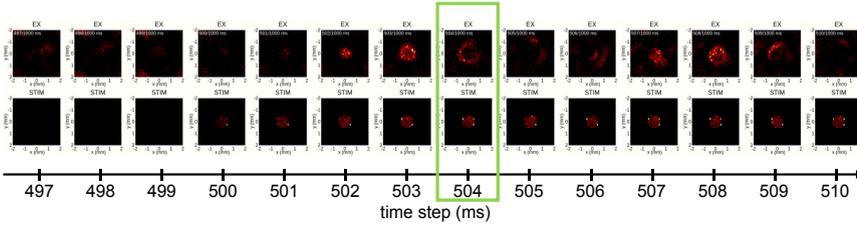


Figure 4.7: **From spontaneous to evoked activity, resolved in time and 2D space.** Time frames of populations EX (first row) and STIM (second row) are captured from view 1 every 1 ms. After three frames showing spontaneous activity of population EX, the STIM layer is activated (first visible in the fourth column, at 500 ms), resulting in a repeated pattern of outward spread of activity in the EX layer. The time step highlighted by a green outline (at 504 ms) corresponds to the animation time step in Figure 4.6.

stimulus duration is similar to time-averaged rate balance in the non-perturbed state.

From the visualization we can also infer that the external perturbation to the network does not result in a shifted network state after the stimulus is switched off. Overall rate fluctuations and relative fractions of spike-count rates appear comparable before and after the stimulus period, unlike networks that may display multi-stable patterns of activity (Litwin-Kumar and Doiron, 2012; Miller, 2016) wherein their state can shift from one attractor to another either spontaneously or due to a perturbation.

4.2.2.2 Spatial features of evoked network activity

Having identified a time segment of particular interest (the stimulus duration), we next exploit view 1, the 2D spike-count rate view, and focus on spatial aspects of the evoked network activity. Figure 4.7 shows a series of snapshots from the instantaneous spike-count rate animation across space for the EX (top row) and the STIM (bottom row) layers. Snapshots are shown for successive bins of width Δt . The first three columns in Figure 4.7 show spontaneous activity of the EX units. Thereafter the STIM population is switched on, as seen in the fourth column of the bottom row. The activity of the STIM layer is by construction confined to a circle at the center of the network. As the activity of the N_{STIM} units in the STIM layer is governed by Poisson processes with rate expectations ν_{STIM} , its spike intensity remains fairly constant (except for the bin at 500 ms as the time bin is centered on the time step). In layer EX, the stimulus elicits an increase in activity spreading outwards from the center. This response dies out after a few milliseconds due to recurrent inhibition, but reoccurs regularly as reflected by the oscillatory behavior observed in Figure 4.7B. The time step at 504 ms highlighted by the green outline is the same as in Figure 4.7 and latter Figures 4.8 and 4.9.

4.2.2.3 2D and 3D views of spatial activity

In order to relate the spatial relationship between activity in individual populations, we compare in Figure 4.8 three different layer-wise animations of neuronal

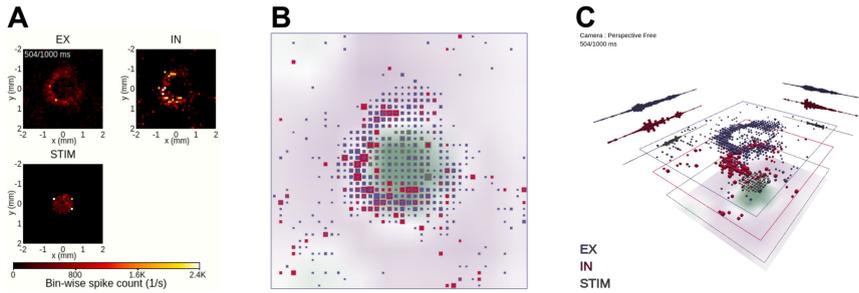


Figure 4.8: **Coordinated views on a temporal snapshot of the neuronal activity.** **A** Instantaneous spike-count rates in layers EX, IN and STIM using view 1. The animation time step of 504 ms is identical to the one in Figures 4.6 and 4.7 in this and subsequent panels. **B** Orthographic top-down view onto stacked population layers and LFP image plot with view 2. **C** Perspective view with large layer separation, including summed spike counts projected towards the layer edges in view 2.

activity. View 1 (Figure 4.8A) shows individual 2D image plots for the spike-count rates per population, with a shared color bar coding for instantaneous spike-count rate values. This view offers an accurate spatial representation of network activity in temporal bins of width Δt , showcasing the locality of the STIM layer activity and the wider spread of evoked activity in the EX and IN layers. This view does not, however, offer interactive features except time control of the animation (shared with views 2-4) and global scaling of the color-value mapping (sensitivity control, shared with views 2 and 3).

The 3D-scene provided by view 2 adds additional interactive features and incorporates the layer-resolved data of view 1 in one animation (Figure 4.8B,C). Panels B and C show the same temporal snapshot of activity as in Panel A. The view shows also the spatial variation of the LFP signal that we synthesized from network activity. The LFP signal, here shown as image plot with a color-coding reflecting its magnitude and sign, is more difficult to relate to the ongoing activity, as it is inherently a signal driven by past spiking activity (resulting of delayed synaptic activation on postsynaptic neurons from spiking activity in presynaptic neurons, cf. Section 4.3.5). We then compare rate values of one spatial bin and one population to other spatial locations and other populations through their different color codings and cube sizes. An observation is that activity in the EX and IN layers are typically confined within the same spatial region of the network, while a larger fraction of the network is quiescent at the time. This observation can for example explain high variability in interspike-intervals of individual neurons (Keane and Gong, 2015), as neurons may fire frequently while fronts of activity spread across the network and remain quiet until the next burst of activity.

In terms of using interactive features offered in view 2, we turn off the orthographic mode of panel B and go back to its default 3D perspective in panel C. We also rotate the viewpoint in order to directly focus on highly active parts of the network. Furthermore, the different layers of the network and LFP are offset vertically, with dynamic projections of the sum of spiking activity across each respective spatial axis for each network layer. From this setup of view 2, we can better infer the

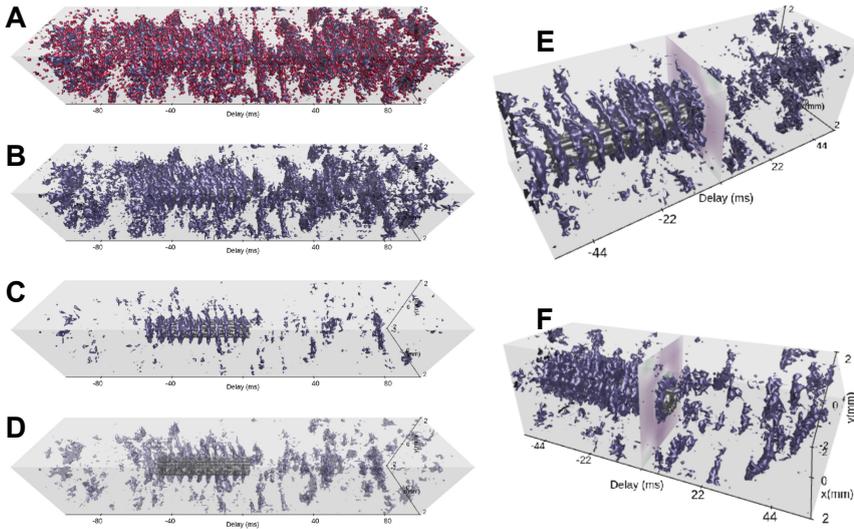


Figure 4.9: **Spatiotemporally resolved activity.** **A** Spike-count rates across time and space of populations EX (blue), IN (red) and STIM (gray) shown with view 4 for a time window of ± 100 ms around the present time step of the animation. The isolation threshold is set to a rate of 100 spikes/s. The animation time step is identical to the one in Figures 4.6–4.8 in this and subsequent panels. **B** Same as panel A, but with IN activity turned off. **C** Same as Panel B, but with an increased isolation threshold of 360 spikes per second. **D** Same as panel C, but with reduced opacity of layer EX activity and an isolation threshold of 195 spikes/s. **E** A narrower time window (± 55 ms) and shifted camera position (isolation threshold of 195 spikes/s). The image plot at a delay of 0 ms shows the synthesized LFP signal across space. **F** Same as panel E, but with the camera position rotated around the vertical z -axis.

activity in each individual layer, including that of the LFP layer, without switching off individual layers.

4.2.2.4 Spatiotemporally resolved network activity

We finally investigate network activity in space and time using the 3D scene provided by view 4. Similar to the scrolling spike-count rate plot of view 3, view 4 allows full control of the time axis. The activity of all populations EX, IN and STIM is displayed for a wide (200 ms) temporal segment using red, blue and gray iso-surfaces, respectively, in Figure 4.9A. We have centered the current time step (at 504 ms) on the evoked activity in the STIM layer (highlighted in Figure 4.7). It is already possible to identify activity patterns confined in space and time. However, it remains difficult to assess how spontaneous network activity changes in response to the stimulus due to occlusion of one surface by another, an inherent issue with multiple solid surfaces. In Figure 4.9B we therefore hide the activity of the IN layer and focus on the activity in the EX layer. The surfaces correspond to the bin-wise instantaneous spike-count rates at an isolation threshold of 100 spikes/s. Increasing this threshold to 360 spikes/s (Figure 4.9C) reveals that regular bursts of high

rates occur at the center of the layer, in the period when the STIM layer is activated. In the other views, these bursts may be seen as rate oscillations (Figure 4.6) or pulsating spatial activity (Figures 4.7 and 4.8). We here show that the attenuation of activity radiating outward from the center is rather strong.

Using view 4, both the oscillation frequency and the outward spread of activity in the EX population can be assessed. We highlight the STIM activity by reducing the opacity of the EX surfaces in Figure 4.9D. This reduces occlusion problems present with multiple overlapping opaque surfaces, and thus allows relating the activity in these two populations to one another. A smaller time segment of the scene is shown in panels E and F where we also demonstrate different camera positions. Rotating the camera allows us to observe the synthesized LFP signal at the current time step, and the corresponding network interactions resulting in a strong LFP fluctuation. We also observe the temporal offset between stimulus onset and a response in the EX activity as shown in Figure 4.9F.

In contrast to the previously discussed applications of views 1–3, the 3D-scene of view 4 allows to relate both temporal and spatial aspects of the spiking activity of different neuronal populations and the LFP signal to one another. With this view, we can get an overview of a large time segment and several populations and then use its incorporated interactive features in order to explore the network activity under influence of the stimulus. The focus of this view lies on highlighting qualitatively interesting features of the data on spatiotemporal scales such as the oscillating activity of EX population surrounding the STIM location (as in Figure 4.9C) or the temporal offset between STIM and EX seen in Figure 4.9F. Views 1 and 3, however, better resolve quantitative rate values or temporal offsets, respectively, than views 2 and 4.

4.3 METHODS

4.3.1 Data formats

The data we consider for visualization are sequences S_j of spike times, defined as $t_j = \sum_{s \in S_j} \delta(t_j^s)$, of a neuronal unit $j \in X$ located at coordinate (x_j, y_j) , where X denotes a neuronal population of size N_X . Individual spike times t_j^s are constrained to a discrete grid $n \cdot dt$ for $n \in \{0, 1, 2, \dots, n_{\text{steps}} - 1\}$, where dt is the time-resolution of spike acquisition and n_{steps} the number of time steps in the acquisition period T . We assume that the raw spike data to be visualized is available in two pure text files per population X . The first file contains two columns with values separated by a white space. Its first column contains integer numbers representing ‘global neuron identifiers’ (neuron IDs) j , while the second column contains corresponding spike times t_j^s in units of ms. This data format, first introduced for experimental data and reviewed in Rostami et al. (2017), is the default output format for spike data of the neuronal network simulator NEST (Kunkel et al., 2017). While the floating point data type is sufficient for displays and the computation of single-neuron and population spike rates, the format is only safe for correlation analysis if the time step is a power of two (Morrison et al., 2007c, A.2). The latter guarantees that spike times have a representation in the data type. An alternative is to use the original definition of the format and denote spike times by the integers n , thus

expressing time in units of the resolution of the grid. The second file contains three space-separated columns. Its first column contains unit IDs j , while columns two and three contain the corresponding coordinates x_j and y_j in units of mm. NEST internally represents networks as a graph where edges denote connections. Neurons cannot be interrogated for their location and are only identified by their ID, thus the information on the location must be defined and stored explicitly.

We consider another text-based data format for the visualization of spike data that are preprocessed by a temporal and spatial binning procedure. For the temporal binning we define a temporal bin size Δt as an integer multiple of the acquisition time resolution dt . For spatial binning of neuron positions along the x - and y -axes we define the bin widths Δl . The third spatial dimension (z -axis) is ignored. Assuming an acquisition period T and the side length L of the centered square network domain, the number of temporal bins is $K = T/\Delta t$ and the numbers of spatial bins along each axis $\{L_x, L_y\} = L/\Delta l$. A spatiotemporal bin is indexed by the length-three tuple of indices $\beta = (l_x \in \{0, 1, \dots, L_x - 1\}, l_y \in \{0, 1, \dots, L_y - 1\}, k \in \{0, 1, \dots, K - 1\})$, spanning $x \in [l_x \Delta l - L/2, (l_x + 1) \Delta l - L/2)$, $y \in [l_y \Delta l - L/2, (l_y + 1) \Delta l - L/2)$ and $t \in [k \Delta t, (k + 1) \Delta t)$. In each spatiotemporal bin, we sum for every population X the number of spike events and divide by the temporal bin size Δt . We refer to this measure as the instantaneous spike-count rate v_β in units of $1/s$. The preprocessed data is contained in one single file per population with four space-separated columns. Indices l_x, l_y and k for each spatiotemporal bin are put in columns 1, 2 and 3, respectively, while the 4th column contains the corresponding rate value. Rows are ordered in iteration running order according to $k \in [0, 1, \dots, K - 1]$ over all $l_x \in [0, 1, \dots, L_x - 1]$ and finally over all $l_y \in [0, 1, \dots, L_y - 1]$. Row entries where $v_\beta = 0$ is not written. The same data format is used to represent the evolution of spatially organized analog data with spatial resolution Δl^ϕ . The unit of the data depends on the actual measure, for example mV in case of the LFP.

4.3.2 Reference implementation

We have made implementations of the visualization types discussed throughout this chapter available in the tool VIOLA (VIsualization Of Layer Activity). Figure 4.10 illustrates the web-based JavaScript framework integrating the different visualizations we refer to as *views*. A central class named `Main` carries out the initialization and coordination of the views. The Graphical User Interface (GUI) is comprised of two main components, the `Setup Panel` and the `Main Panel`. The `Main Panel` also serves as a container for the views. The `Setup Panel` is the first entity presented to the user when the application is opened in a web browser. It serves mainly to specify the data types to be loaded and the basic data features (spatial dimensionality, time resolution) and visualization features such as the colors for each neuron population. Parameters can be set manually or be loaded from configuration files, for example specifying whether to load raw or preprocessed data. These configuration files are JavaScript Object Notation⁷ (JSON) files specifying the format of the loaded data, file names, and preset values for the visualizations. After confirming the entered information, the `Main.setup()` function

⁷ <http://www.json.org>

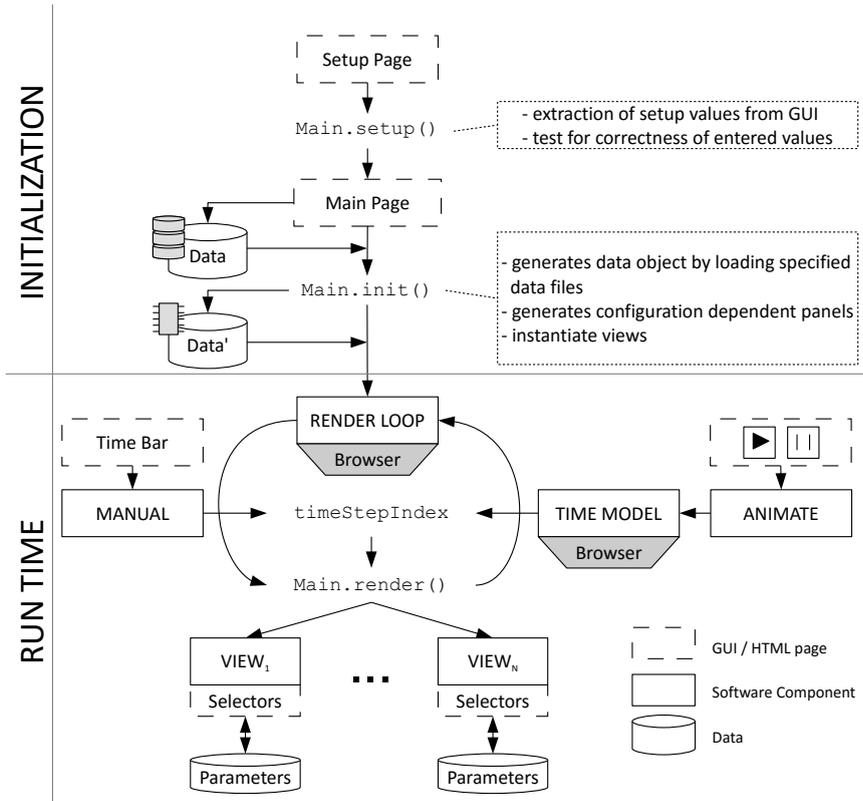


Figure 4.10: **Flow chart of VIOLA's components.** VIOLA incorporates two main parts: 'initialization' (top) and 'run time' (bottom). The initialization procedure defines a Setup and Main Panel in VIOLA's GUI and the corresponding `Main.setup()` and `Main.init()` functions. The `Main.setup()` function is used for setting initial values, while the `Main.init()` function allows for loading datasets with parameters that depend on setup values. The `Main.init()` function also instantiates the different views in the application. The run-time component of VIOLA uses a rendering loop and time model provided by the web browser. The time model is needed to synchronize the rendering of each view, such that at all times each view shows the same time step of the data and thus compensates for different rendering times of the views. By default, the `timeStepIndex` in the rendering loop is automatically updated by the browser (using the time model), but can also be set by the user, e.g. by a slider widget. The browser time model is controlled using the 'animate' widget, while a time bar is used for 'manual' time selection. Each update of the `timeStepIndex` triggers an execution of the `Main.render()` method and a corresponding update of all views. Parameters for the different views can be modified during run time as each view offers its individual input widgets.

extracts the entries and provides them globally to the other components. The `Main Panel` shown afterwards is used to load the input data files from the local file system using `JavaScript FileReader`, which includes setting up the internal data structure giving a coherent access to the data to be visualized. The `Main.init()` function then initializes the rendering loop, which is built into the web browser and controls the rendering of the various views. The `Main.render()` method is executed periodically by built-in functionality of the browser, which is further used to synchronize the rendering of all views. This is necessary as the rendering of an individual time step needs different amount of time per view. For example, rendering a complex 3D scene is slower (because it need more computational resources) than rendering a less complex 2D plot. All views must have finished rendering before the next rendering step is triggered. Each rendering call is triggered by updating the global `timeStepIndex` through the browser. The update of the `timeStepIndex` calls the rendering loop, which redirects the rendering call to all views. For rendering the data, all views access the loaded simulation data structure as part of the `Main` object.

The `timeStepIndex` variable can be controlled in two modes. First, it can be manually set by the user via a slider widget. The slider widget enables the user to scroll manually along the time axis, thus offering a simple navigation through the time series. Second, the user can start an animation of the loaded data set by pressing a start button. This triggers a periodic update of the `timeStepIndex`. As both manual and periodical updates of the `timeStepIndex` trigger the same downstream functionality for updating the views, manual navigation through the time series can be combined with automatic updates of the `timeStepIndex`. In case the animation is running, manual intervention by the user overrides the periodic update of the `timeStepIndex`, such that the shown data item (time step) corresponds to that manually selected one. The animation is continued from this manually selected time step.

All views offer view-specific selectors for visualization parameters. These parameters are read in and used in each render call.

4.3.3 2D and 3D view implementations

This section describes the visualization algorithms used with the different view implementations in VIOLA as introduced throughout [Section 4.2](#). As the output data produced by the simulation are scalars organized on a regular grid (related to the neuron's position), visualization methods applicable to scalar data are employed as well as standard chart types ([Hansen and Johnson, 2011](#)). `View 1` implements a standard image plot in which the color of each spatial bin represents a spike-count rate value. The display maps a binned measure of neuronal activity on a color by using a lookup table ranging from black to white over red and yellow (often referred to as 'red hot' or 'white hot' lookup table). We selected this lookup table as it is widely used and offers (if non-linear interpolated) equidistant colors according to the CIE $L^*a^*b^*$ color space [EN ISO 11664-4 from 1976]⁸. This color space is designed to represent equidistant colors according to human perception: a color twice as light in CIE $L^*a^*b^*$ space is also perceived twice as light by a human user

⁸ <https://www.vis4.net/blog/posts/mastering-multi-hued-color-scales>

(Fairchild, 2013). Such a heat map (Spence, 2014) facilitates the representation of the spatial structure of the data.

In the 3D visualization of view 2, VIOLA also implements a geometrical mapping of activity data to a cube's edge length, resulting in a cubic mapping of the single scalar value representing activity at each time step. The default scaling of the cube's edge length is such that the cube's volume is proportional to the data value in each bin.

In view 3, the concept of a stacked bar chart supports a global perspective on the simulated model (Spence, 2014). Individual populations in the model are color-coded to be separable in the bar chart. An additional line graph added on top of the bars shows the total spike-count rate as reference to the global activity.

For view 4, in order to support the visual interpretation of time series of spatially organized activity data, the 2D-organized activity data (considering the neurons in one layer) are extracted along the time axis resulting in a regularly structured 3D volume of scalar values. Through contouring, partial sub-volumes with a certain minimum threshold of activity get extracted and rendered as geometry. By means of a selected iso-value I_{th} such geometry gets extracted by applying the marching cubes algorithm for implicit volume rendering (Lorenson and Cline, 1987). For extraction of the geometry, the algorithm assumes that each data point of the data set is mapped onto a vertex (corner point) of a regular 3D grid, which can be subdivided into cells delimited by eight neighboring vertices each. Then, the algorithm calculates for each vertex of a cell whether the associated data value of the considered vertex lies inside or outside of the contour defined by the iso-value I_{th} by comparing the data value with the iso-value. If the data value of a vertex is smaller than the iso-value, the vertex is assumed to lay inside of the contour. For each possible combination of inside/outside states of the vertices of a cell, the topology of the contour for each cell gets extracted from a table by calculating a representative index. This table holds all possible topological states of a cell, which are constructed under the assumption that there are an infinite number of possibilities how a contour can pass a cell (for more details, please refer to, for example Hansen and Johnson (2011, Chapter 1). Finally, the exact position of the contour gets calculated by interpolation along the cell's edges.

Views 2, 3 and 4 all use the same color coding to identify the different neuronal populations. The implementation of the algorithms and views uses native JavaScript. The 2D rendering routine uses the HTML5 canvas element. The browser rendering engine supports HTML5 and especially the functionality of the canvas element, therefore no external libraries are required. 3D renderings relied on the three.js⁹ wrapper for WebGL content which is natively supported by the engine of modern browsers. Node.js¹⁰ facilitates the communication between views and the GUI for control.

4.3.4 Network description

The example network is based on an implementation of a random balanced network (Brunel, 2000) which is part of NEST as an example (`brunel_alpha_nest.py`

⁹ <https://threejs.org>

¹⁰ <https://nodejs.org>

in NEST 2.12.0 by Kunkel et al., 2017). The model is expressed using PyNEST (Eppler et al., 2009) in Python¹¹. The network consists of N_{EX} excitatory and N_{IN} inhibitory spiking point-neurons which are sparsely connected with connection probability c . Neurons have fixed in-degrees of cN_{EX} excitatory and cN_{IN} inhibitory incoming synapses with weights $g_{Y,\text{EX}}J$ and $g_{Y,\text{IN}}J$, respectively, with $Y \in \{\text{EX}, \text{IN}\}$. The integrate-and-fire model neurons are connected using static, current-based synapses with an alpha-shaped time course (NEST neuron model: `iaf_psc_alpha`). The intrinsic neuron parameters are identical for both neuron types. In addition to the recurrent connections, each neuron receives uncorrelated, external excitatory input from a Poisson process of a fixed rate $\nu_{\text{ext}} = \eta\nu_{\theta}$, where η denotes the external rate relative to the threshold rate ν_{θ} which is defined as $\nu_{\theta} = (V_{\theta} - E_L)C_m / (\exp(1)J\tau_m\tau_s)$. The threshold rate is the hypothetical external rate needed to bring the average membrane potential of a neuron to threshold V_{θ} (in the absence of an actual spiking mechanism). E_L denotes the resting potential, C_m the membrane capacitance, τ_m the membrane time constant and τ_s the postsynaptic current time constant.

Unlike the original network implementation which has no spatial information, we here place neurons randomly on a square 2D sheet with side lengths L . The connection probability between a presynaptic neuron j and postsynaptic neuron i decays with increasing horizontal distance r_{ij} (using periodic boundary conditions) while we preserve the in-degrees (number of incoming connections). A Gaussian-shaped profile $p_{YX}(r_{ij})$ is used with a standard deviation of σ_{YX} with $X, Y \in \{\text{EX}, \text{IN}\}$. We use $\epsilon_{YX}(r_{ij})$ to describe the distance-dependent connectivity profile assuming that the in-degree is preserved. The transmission delay function $d_{YX}(r_{ij})$ has a linear distance dependency with an offset d_{YX}^0 and a conduction velocity v_{YX} .

In addition to the stationary external input to each population, the network receives a spatially confined transient input with a duration t_{STIM} . The input is provided by a size N_{STIM} population of parrot neurons (NEST's `parrot_neuron` devices), positioned inside a circle of radius R_{STIM} around $(x, y) = (0, 0)$. Parrot neurons simply repeat input spike events as output spike events. Each parrot neuron receives input from a Poisson process with a rate expectation of ν_{STIM} and connected to K_{STIM} neurons of the EX population inside a connection mask radius R from the parrot-neuron location. The Poisson input starts at T_{STIM} and consequently the STIM units become active after a delay of d_{STIM} .

Table 4.1 summarizes the network description with model and simulation parameters listed in Table 4.2A and B. The original parameters for the EX-IN network in the NEST example are modified for this VIOLA use case demonstration bringing the network in a state with spatially confined network activity. For this we increase the network's neuron count, reduce the ratio of inhibitory to excitatory weights $g_{Y,\text{IN}}$ and the membrane capacitance C_m , while the postsynaptic amplitude J is increased. The parameter J is originally defined in units of mV, but is here re-defined in units of pA. Finally, the fixed conduction delay is replaced by a distance-dependent one.

The data sets result from simulations of duration T_{sim} with a temporal resolution of dt . We discard the startup transient period T_{trans} and record all spike times from all neurons. The unprocessed spike times together with the corresponding neuron

¹¹ <http://www.python.org>

positions are considered as raw output. The temporal and spatial bin sizes used for preprocessing, Δt and Δl respectively, are given in Table 4.2C.

4.3.5 LFP predictions

Generation of LFP-like data: The local field potential (LFP) is, due to its relative ease of measurement, a common measure of neuronal activity (Buzsáki et al., 2012; Einevoll et al., 2013a). The LFP is, in general, assumed to reflect synaptic activity and correlations of a large number of neurons in vicinity of the recording electrodes (Kajikawa and Schroeder, 2011; Lindén et al., 2011; Łeński et al., 2013). For the purpose of demonstrating VIOLA’s functionality, we synthesize LFP signals from network activity assuming a linear network-population spike to LFP relationship $H_X \equiv H_X(\vec{\Delta}, \tau)$ derived using a biophysical model. In this relationship, $\vec{\Delta}$ denotes the displacement between the center of a spatial bin and an electrode contact point γ at \mathbf{r}_γ , and τ the time relative to a presynaptic spike event (“lag”). Assuming linearity and homogeneous spike-LFP responses of individual presynaptic neurons located within the same bin of width Δl^ϕ indexed by $b = (l_x^\phi, l_y^\phi)$ (see Section 4.3.1), the signal ϕ_X at one contact γ of one population X is then given by

$$\phi_X(\mathbf{r}_\gamma, t) = \sum_b \left(\left(\sum_s \delta(t_b^s) \right) * H_X \right) (\mathbf{r}_\gamma, t). \quad (4.1)$$

Here, the term $\sum_s \delta(t_b^s)$ represents a series of spike times t_b^s of all presynaptic neurons in a bin b where δ denotes the Dirac delta function, and $*$ a convolution. As contributions of different populations X sum linearly, the total signal at each contact is

$$\phi(\mathbf{r}_\gamma, t) = \sum_X \phi_X(\mathbf{r}_\gamma, t). \quad (4.2)$$

Point-like neurons (as used in our network model) can not generate an extracellular potential, as all in- and outgoing currents sum to zero at the point’s location (due to conservation of charge). As in Hagen et al. (2016a) we assume that spatially extended (morphologically detailed) neurons and corresponding multicompartment models in combination with an electrostatic forward model are required to compute a biophysically meaningful LFP signal. To compute the LFP, we here derive for each presynaptic population $X \in \{\text{EX}, \text{IN}, \text{STIM}\}$ the phenomenological mapping $H_X(\vec{\Delta}, \tau)$ between a presynaptic spike event time t_b^s occurring in a spatial bin indexed by b to the extracellular potential.

Measurement sites: The electrode contact point locations are defined at the center of each spatial bin as $\mathbf{r}_\gamma = ((l_x^\phi + 1/2)\Delta l^\phi - L/2, (l_y^\phi + 1/2)\Delta l^\phi - L/2, 0)$.

Multicompartment model: We define a ball-and-stick type multicompartment model neuron with morphological features and passive parameters derived from the network’s LIF neuron description (membrane capacitance C_m , membrane time constant τ_m , passive leak reversal potential E_l). Assuming a homogeneous specific membrane capacitance c_m (capacitance per membrane area) and axial resistivity r_a (resistance times length unit), we choose the dendritic stick length L_{dend} and radius r_{dend} as follows: To preserve the total capacity of the point neuron (and equivalent surface area), we compute the corresponding soma radius as

$r_{\text{soma}} = \sqrt{\frac{C_m}{4\pi c_m} - \frac{r_{\text{dend}} L_{\text{dend}}}{2}}$. We define the passive leak conductivity as $g_L = c_m / \tau_m$ and leak reversal potential as E_L . For these calculations we choose c_m , r_a , L_{dend} and r_{dend} values as given in Table 4.3, resulting in $r_{\text{soma}} \approx 13.1 \mu\text{m}$. The compact ball-like soma is treated as a single segment, while the elongated dendrite is split into $n_{\text{dend}} = 11$ segments of equal length. The center of the soma segment is set to $\mathbf{r} = (0, 0, 0)$, and the dendritic stick is aligned in the positive direction along the vertical z -axis.

Synapse model: For LFP predictions we use the same current-based synapse model as in the network, defining the postsynaptic input current of a single presynaptic spike event as $I_{ij}(t) = J_{YX} \cdot (t - t_j^s - d_{ij}) / \tau_{\text{syn}} \exp(1 - (t - t_j^s - d_{ij}) / \tau_{\text{syn}}) \Theta(t - t_j^s - d_{ij})$, where J_{YX} denotes the connection-specific postsynaptic current amplitude as in the network, t_j^s the presynaptic spike time, $d_{ij} = d_{YX}(r_{ij})$ the conduction delay between presynaptic cell j and postsynaptic cell i and Θ the Heaviside step function. As we initially ignore delays and network spike times we set $d_{ij} = 0$ and $t_j^s = \tau^s$.

Synaptic connectivity: For outgoing connections of the excitatory populations $X \in \{\text{EX}, \text{STIM}\}$ we distribute synaptic input currents evenly along the entire length of the dendritic stick, while for outgoing connections of the inhibitory population $X = \text{IN}$ all synaptic input currents are assumed to be evenly distributed on the ball-like soma.

Electrostatic forward model: As described in detail in Lindén et al. (2014), we assume an extracellular conductive medium that is linear (frequency independent), isotropic (identical in all directions), homogeneous (identical in all positions) and ohmic (linear relationship between current density and electric potential), as represented by the scalar conductivity σ_e (cf. Table 4.3 for values). From the linearity of Maxwell's equations, contributions to the extracellular potential from different current sources sum linearly. Here, these current sources are transmembrane currents (summed over resistive, capacitive and synaptic currents). In the presently used volume conduction theory, the electric potential in location \mathbf{r}_γ from a point current with magnitude $I(t)$ in location \mathbf{r}_0 is

$$\phi_{\text{point}}(\mathbf{r}_\gamma, t) = \frac{1}{4\pi\sigma_e} \frac{I(t)}{|\mathbf{r}_\gamma - \mathbf{r}_0|}. \quad (4.3)$$

This relation is also valid for a sphere current source (i.e., our ball soma) centered at \mathbf{r}_0 with total transmembrane current $I_{\text{m,soma}}$ and radius r_{sphere} when $|\mathbf{r}_\gamma - \mathbf{r}_0| \geq r_{\text{sphere}}$. Thus

$$\phi_{\text{soma}}(\mathbf{r}_\gamma, t) = \frac{1}{4\pi\sigma_e} \frac{I_{\text{m,soma}}(t)}{|\mathbf{r}_\gamma - \mathbf{r}_{\text{soma}}|}. \quad (4.4)$$

The elongated dendritic segments are treated as 'line sources', obtained by integrating the point-source formula along the central axis of the segments (Holt and Koch, 1999; Lindén et al., 2014):

$$\phi_{\text{dend}}(\mathbf{r}_\gamma, t) = \frac{1}{4\pi\sigma_e} \sum_{u=1}^{n_{\text{dend}}} I_{\text{m},u}(t) \int \frac{d\mathbf{r}_u}{|\mathbf{r}_\gamma - \mathbf{r}_u|}. \quad (4.5)$$

The total extracellular potential from somatic and dendritic sources is then

$$\phi(\mathbf{r}_\gamma, t) = \phi_{\text{soma}}(\mathbf{r}_\gamma, t) + \phi_{\text{dend}}(\mathbf{r}_\gamma, t). \quad (4.6)$$

Our calculations of extracellular potentials rely on the Python package LFPy¹² (Lindén et al., 2014; Hagen et al., 2018). The tool implements the above forward-model formalism for extracellular potentials, and uses the NEURON simulation environment (Carnevale and Hines, 2006) to compute transmembrane currents $I_m(t)$ of multicompartment neuron models. As singularities may occur in the limit $|\mathbf{r}_\gamma - \mathbf{r}_u| \rightarrow 0$, the minimum distance between sources and measurement locations was set equal to the somatic or dendritic segment radius.

Prediction of spike-LFP relationship: We here describe the calculation of the linear spike-LFP relationships $H_X(\vec{\Delta}, \tau)$ which we use to construct an LFP-like signal from spatially binned network activity. While Hagen et al. (2016a) present a hybrid scheme to compute extracellular potentials from point-neuron network activity, and incorporated the biophysics-based forward model summarized above, this hybrid scheme is not adapted to laminar point-neuron networks with distance-dependent connections. We therefore construct a simpler and numerically much less demanding method inspired by the hybrid scheme, that still encompasses the governing biophysics underlying the generation of extracellular potentials and accounts for the laminar structure and distance-dependent connectivity of our network.

In this simplified model, we ignore heterogeneity in spike-LFP responses H_i , of individual presynaptic cells $i \in X$ located within a spatial bin b , i.e., $H_X \equiv \langle H_i \rangle$. H_i corresponds to the extracellular potential resulting of synaptic activation of postsynaptic populations of cells $j \in Y$ from a spike in cell i at time $\tau = 0$. We also assume that H_X is invariable across presynaptic bins, and encompasses the overall distance-dependent connection probabilities and connection delays in the network.

The calculation of $H_X(\vec{\Delta}, \tau)$ involves a number of steps. We first estimate the spatially averaged extracellular potential $\varphi_j(\vec{\Delta}, \tau)$ resulting from a single synapse activation at a time τ^s of the ball and stick neuron positioned at the center of a reference bin, for excitatory and inhibitory input. Electrode contact point locations \mathbf{r}_γ are defined at the centers of each square spatial bin indexed b (see above). With rotational symmetry around the z -axis and periodic (torus) connectivity of the network, we compute extracellular potentials at the unique subset of bin center-to-center distances $r \in \{|\vec{\Delta}|\}$ up to the maximum distance $\sqrt{2}L^2$, where L denotes the side length of the network layers, and $\{|\vec{\Delta}|\}$ the complete set of center-to-center displacement vector lengths from reference bin to all spatial bins. We utilize built-in functionality in LFPy to perform spatial averaging (cf. Equation 6 in Lindén et al. (2014)), assuming square contact points parallel to the horizontal xy -plane with side lengths equal to the bin width Δl^ϕ . In a following step we compute the average out-degree (number of outgoing connections of neuron i) $K_X = \sum_Y N_{Yc}$ for $X \in \{\text{EX}, \text{IN}\}$, where c denotes the overall connection probability between X and Y (cf. Table 4.2 which also gives K_{STIM} as a fixed parameter). With the distance-dependent connectivity $\epsilon_{YX}(r)$ used for each presynaptic population and out-degree K_X we compute the number of activated synapses (denoted by K_r) in each spatial bin at a distance r from the reference bin (including $r = 0$) by evaluating $p_{YX}(r)$ at the bin center points. The average connection delays from the reference bin to other bins are approximated as $d_{YX}(r) = d_{YX}^0 + r/v_{YX}$, where d_{YX}^0

¹² <http://lfp.py.github.io>

denotes a constant delay offset and v_{YX} the conduction speed of action potentials in the network, with values as in Table 4.2. With the elements of these steps in place (single-synapse LFP responses across bins, bin-wise number of activated synapses and delays), we construct $H_X(\vec{\delta}, \tau)$ as function of r as:

$$H_X(\vec{\Delta}, \tau) = \sum_{r \in \{|\vec{\Delta}|\}} K_r \cdot (\delta(d_{YX}(r)) * \varphi_j)(\vec{\Delta}, \tau), \quad (4.7)$$

where $\delta(\cdot)$ denotes the Dirac delta function. Note that we sum over all elements r in $\{|\vec{\Delta}|\}$.

LFP output: Each H_X is calculated at a spatial resolution Δl^ϕ and temporal resolution of dt (as in the network, cf. Table 4.2) for a total duration of $2\tau^s$, with synapse activation time at time τ^s . An identical spatial and temporal binning resolution is also used for spike events entering in Equation 4.1. The spike rates in each bin are filtered by a length Δt normalized boxcar filter using the `scipy.signal.lfilter` method prior to the convolution with the corresponding LFP kernel. Otherwise a temporal shift between the spatiotemporally binned spiking data (cf. Section 4.3.1) and the downsampled LFP in the visualization occurs. Discrete convolutions are incorporated using `numpy.convolve` and `scipy.signal.convolve2d` methods in Python. The final LFP signals are low-pass filtered and downsampled to the time resolution Δt of our preprocessed network output as described in Hagen et al. (2016a) in order to simultaneously show both datasets in VIOLA. Output is stored in a pure-text format as described in Section 4.3.1.

Table 4.3 summarizes the parameter values for the LFP predictions.

4.3.6 Software summary

All source codes of the tool VIOLA, the example network model and the processing of model output are hosted at <https://github.com/HBPVIS/VIOLA> (SHA:ca2f3c5).

We simulated the example network (`topo_brunel_alpha_nest.py`) with NEST v2.12.0 and Python v2.7.11. Further processing and plotting of Figures 4.1 and 4.2 (`nest_preprocessing.py`) also relied on Python with `numpy v1.10.4`, `SciPy v0.17.0`, and `matplotlib v1.5.1`. LFP signals (`fake_LFP_signal.py`) were computed using NEURON v7.5 and LFPy from <http://lfp.github.io> (SHA:5673a6). We visualized the neuronal activity with VIOLA using the Google Chrome browser, version 58.0.3029.110 (64-bit). VIOLA used JavaScript V8 5.8.283.38 with the 3D library `three.js` of revision 87, including WebGL and HTML5 build in the browser and `Node.js v4.8.3`. For colors, VIOLA used `Chroma.js` in the version 1.3.4.

Screenshots from VIOLA for the other figures were taken with Kazam-“NCC-80102” v1.4.5, and combined in Microsoft PowerPoint 2013.

A: Model summary	
Populations	Three: excitatory EX, inhibitory IN, external stimulus STIM
Topology	EX/IN: random neuron positions on square domain of size $L \times L$; STIM: random neuron positions inside a circle with radius R_{STIM} at the center of the domain; periodic boundary conditions
Connectivity	Random (EX/IN: convergent, fixed in-degree; STIM: divergent, fixed out-degree) connections described by distance-dependent probability kernels and cut-off masks
Neuron model	EX/IN: leaky integrate-and-fire (LIF), fixed threshold, absolute refractory time; STIM: parrot
Synapse model	Static weights, EX/IN: alpha-shaped postsynaptic currents, distance-dependent delays
Input	Independent fixed-rate Poisson spike trains to all neurons
Measurement	Spike activity
B: Network model	
Subthreshold dynamics	<p>EX/IN: If $t > t^* + \tau_{\text{ref}}$</p> $\frac{dV}{dt} = -\frac{V - E_L}{\tau_m} + \frac{I_{\text{syn}}(t)}{C_m}$ $I_{\text{syn}}(t) = \sum_j J_j \alpha(t - t_j^* - d_j)$ <p>with connection strength J_j, presynaptic spike time t_j^* and conduction delay d_j</p> $\alpha(t) = \frac{t}{\tau_s} e^{1-t/\tau_s} \Theta(t) \text{ with Heaviside function } \Theta$ <p>else</p> $V(t) = V_{\text{reset}}$
Spiking	<p>If $V(t-) < V_\theta \wedge V(t+) \geq V_\theta$</p> <ol style="list-style-type: none"> 1. set $t^* = t$ 2. emit spike with timestamp t^* 3. reset $V(t) = V_{\text{reset}}$
Distance-dependent connectivity	<p>Neuronal units $j \in X$ at location (x_j, y_j) and $i \in Y$ at (x_i, y_i) in pre- and postsynaptic populations X and Y, respectively. Distance between units i and j:</p> $r_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$ <p>Gaussian kernel for connection probability:</p> $p_{YX}(r_{ij}) = e^{-r_{ij}^2 / 2\sigma_{YX}^2}$ <p>R is the radius of a cut-off mask. Transmission delay function:</p> $d_{YX}(r_{ij}) = d_{YX}^0 + r_{ij} / v_{YX}$

Table 4.1: Description of the network model following the guidelines of Nordlie et al. (2009a).

A: Global simulation parameters		
Symbol	Value	Description
T_{sim}	1,500 ms	Simulation duration
dt	0.1 ms	Temporal resolution
T_{trans}	500 ms	Startup transient
T_{STIM}	999 ms	Start time of Poisson input to STIM
t_{STIM}	50 ms	Duration of STIM onset
B: Point-neuron network		
Populations and external input		
Symbol	Value	Description
X	EX, IN, STIM	Name
N_X	20,000	Population size: $X = \text{EX}$
	5,000	$X = \text{IN}$
	975	$X = \text{STIM}$
L	4 mm	Extent length
η	2	External rate relative to threshold rate for $X \in \{\text{EX}, \text{IN}\}$
R_{STIM}	0.5 mm	Radius of circle around (0,0) for locations of STIM
ν_{STIM}	300 Hz	External rate to each STIM neuron
Connection Parameters		
Symbol	Value	Description
c	0.1	Connection probability for recurrent connections between EX and IN
J	40 pA	Reference synaptic strength. All synapse weights are measured in units of J .
g_{YX}	1	Relative synaptic strengths: $X = \text{EX}, Y \in \{\text{EX}, \text{IN}\}$
	-4.5	$X = \text{IN}, Y \in \{\text{EX}, \text{IN}\}$
	1	$X = \text{STIM}, Y = \text{EX}$
R	0.1 mm	Radius of cut-off mask for $X = \text{STIM}, Y = \text{EX}$
K_{STIM}	300	Number of connections per STIM neuron
σ_{YX}	0.3 mm	Standard deviation of Gaussian kernel: $X, Y \in \{\text{EX}, \text{IN}\}$
d_{YX}^0	0.5 ms	Delay offset: $X, Y \in \{\text{EX}, \text{IN}\}$
	0.5 ms	$X = \text{STIM}, Y = \text{EX}$
v_{YX}	2 m/s	Conduction velocity: $X, Y \in \{\text{EX}, \text{IN}\}$
	-	$X = \text{STIM}, Y = \text{EX}$
d_{STIM}	0.5 ms	Delay from Poisson input to STIM
Neuron model		
Symbol	Value	Description
C_m	100 pF	Membrane capacitance
τ_m	20 ms	Membrane time constant
E_L	0 mV	Resting potential
V_θ	20 mV	Firing threshold
V_{reset}	0 mV	Reset potential
τ_{ref}	2 ms	Absolute refractory period
τ_s	0.5 ms	Postsynaptic current time constant
C: Preprocessing		
Symbol	Value	Description
Δt	1 ms	Temporal bin size
Δl	0.1 mm	Spatial bin size

Table 4.2: Simulation, network and preprocessing parameters.

Simplified LFP model parameters		
Symbol	Value	Description
c_m	$1 \mu\text{F}/\text{cm}^2$	Specific membrane capacitance
r_a	$150 \Omega\text{cm}$	Axial resistivity
L_{dend}	$500 \mu\text{m}$	Dendritic stick length
r_{dend}	$2.5 \mu\text{m}$	Dendritic stick radius
n_{dend}	11	Dendritic stick number of segments
r_{soma}	$13.1 \mu\text{m}$	Derived soma segment radius
τ^s	25 ms	Synapse activation time
Δl^ϕ	$400 \mu\text{m}$	Electrode separation, spatial bin width
σ_e	$0.3 \text{S}/\text{m}$	Extracellular conductivity

Table 4.3: Parameters for prediction of LFP signals.

4.4 DISCUSSION

The present study introduces four 2D and 3D visualization concepts, or views, for the interactive visual analysis of the activity of spiking neuronal network simulations, and a reference implementation for these views named VIOLA (Visualization of Layer Activity). VIOLA is an interactive web-technology based visualization tool designed to fit in between simulations and subsequent in-depth data analysis, and exemplifies key concepts of the information-seeking mantra by Shneiderman (1996) and the paradigm of coordinated multiple views (Wang Baladonado et al., 2000). The main application areas are the rapid validation of simulation results and the exploration of spatiotemporally resolved data prior to further quantitative analyses. As a use case, we demonstrate the usefulness of the tool with output from a simulation of a layered spiking point-neuron network model that incorporates distance-dependent connectivity. The use case shows that we can examine a perturbation of ongoing network activity caused by a temporally and spatially confined stimulus. The duration and the spatial spread of the event are quickly assessed with the help of multiple simultaneously displayed views.

In contrast to other visualization tools for simulated network output, for example VisNEST (Nowke et al., 2013, 2015), SNN3DViewer (Kasiński et al., 2009), ViSimpl (Galindo et al., 2016), and Geppetto or more generic multi-view tools like GLUE, the interactive JavaScript- and WebGLbased visualization integrates data analysis methods in a web application, thereby achieving mobility and deployability. Our approach builds on visualization concepts known from the literature for data of similar structure (reviewed in the introduction), but advances the concepts and adds interactivity and animation. For example, views 1 and 2 compare to series of snap shots (as in Mehring et al. 2003; Yger et al. 2011; Voges and Perrinet 2012; Keane and Gong 2015), but are here enhanced by the possibilities to show raw or preprocessed data, to specify visualization parameters interactively, and to provide a 3D and temporally animated view on the multi-dimensional data. View 4 presents a new concept combining 2D spatial and temporal resolution of multiple neuron populations, all shown simultaneously. This data representation delivers a wealth of information, but, to circumvent occlusion and instead expose interesting features of the data, it relies on interactive usage. The code of the reference implementation is open source and available in a pub-

lic repository (<https://github.com/HBPVIS/VIOLA>) together with the revision history and documentation. The present work uses the simulation code NEST to generate the data but the VIOLA implementation is completely independent of the former. The JavaScript code defines a standalone application (accessible at <http://hbpvis.github.io/VIOLA>) and interpretable by the browser running on the client device. In the last decade, JavaScript-based visualization got more and more versatile especially fostered by the introduction of HTML5 and its canvas environment. Furthermore, the development of WebGL enables the access to GPU-accelerated 3D rendering in the browser. Beside limitations regarding memory and access to low-level program control (as needed for controlled use of multi-threading), JavaScript offers the opportunity for simple deployment and handling of external libraries and dependencies. Unfortunately, JavaScript-based implementations tend to fail on one or the other browsers as browsers still differ in their interpretation of JavaScript and in the degree of following the HTML standard. Nevertheless, most browsers are free to use and usable on most operating systems. Therefore, this work explores the radical decision to use web-based technology to offer an easy-to-deploy tool for the visualization of dynamic simulation data. As a consequence, software development and deployment are integrated with minimal effort and no computational resources are required on the server: researchers immediately profit from progress on the development platform. Furthermore, due to the web-technology and the minimal requirements on the client, web portals can embed the application as a visualization backend; a prerequisite for the idea to create centralized ICT infrastructure for neuroscience. One such portal is currently being developed by the European Human Brain Project, named the HBP Collaboratory¹³. Another ongoing effort is the Neuroscience Gateway¹⁴ (Sivagnanam et al., 2013). Online embedding opens the possibility to accompany interactive visualization with server-side preprocessing steps and a database integration, in particular for simulation output being generated on the portal itself. This advances the goal of the HBP Collaboratory to provide a fully digitized workflow from data representation over model construction and simulation to model validation (Senk et al., 2017b). We argue that interactive visual analysis of simulated data is an obvious feature of a collaboratory, in addition to non-interactive script-based plotting relying for example on `matplotlib`.

The reference implementation accesses the file system of the host machine to load data. This is not recommended for web applications for security reasons. If data processing and storage were handled on the server-side, SQL-like database queries could restrict communication to only the data needed for the different view instances. Communication does not have to be limited to the raw data. Binning operations similar to those performed in our preprocessing steps can be handled by the database in a straightforward manner, and could also be performed in parallel. The data format HDF¹⁵ would also be an option to store and access large amounts of raw and preprocessed data with improved performance in terms of speed and compactness compared to the currently used text format.

¹³ <https://collab.humanbrainproject.eu>

¹⁴ <http://www.nsgportal.org>

¹⁵ <https://www.hdfgroup.org>

Inherent in interactive visualization is the problem of reproducibility. The raw data are insufficient to reproduce the visuals, only in combination with the full collection of GUI parameters adjusted by the researchers is the data set complete. In the same way as experimental and simulated data need to be enriched with metadata in order to uniquely specify their origin and enable reuse (Zehl et al., 2016), the visuals need to be enriched with the parameters of their creation. This new type of metadata could be stored in a database.

The JavaScript implementation imposes other shortcomings. Prominent is its limited capability for numerical analysis. While the `math.js`¹⁶ library provides a number of basic math functions and support for symbolic operations, complex numbers and arrays (matrices), the JavaScript libraries are not comparable to the Scientific Python stack (SciPy¹⁷) which provides an ecosystem of fundamental tools and methods encountered in mathematics, engineering and science. VIOLA implements the function computing the spatial correlation of neuronal activity from scratch (not shown). This approach has two conceptual weaknesses. First, the speed and accuracy of such functions are hampered by the fact that there is little native support for advanced mathematical operations, like the Fast Fourier Transform (FFT). Second, there is no separation between the code carrying out the statistical analysis and the code performing the visualization. This cuts visualization off from the rich set of analysis tools developed by the community and their reliable implementations, for example as collected in the Elephant package¹⁸. Future work needs to disentangle numerics from visualization code as separate building blocks in a visual analysis workflow.

As VIOLA's main focus lies on responsive interactive visualization, the reference implementation uses WebGL for all views. Prior tests exposed the low efficiency of the Document Object Model (DOM) as used in Scalable Vector Graphics (SVG) based visualization libraries such as `d3.js` as well as its high memory consumption. This led to the decision to the sole use of WebGL rendering, which has the limitation that external tools are required for generating screen shots and screen casts; vector graphics can neither be recorded nor exported. For the 2D views, an additional implementation based on the HTML 5 support of SVG graphics can be added and used for the export of vector-based image material. For extracting vector-based material from the 3D views, WebGL and its access to the underlying rendering pipeline can be facilitated. The 3D scene can be exported to be viewed in other 3D programs. To this end, `three.js` (as used in the reference implementation) offers export functionality for Wavefront OBJ file format, one standard for 3D content. The alternative is to extract the rendered scene prior to rasterization and use these data to generate a SVG or postscript-based representation similar to the operation of the C library `gl2ps`¹⁹. Nevertheless, any export mechanism needs to facilitate means of reproducibility. In particular metadata such as simulation and visualization parameters, time stamps, viewpoint angle and position etc. need to be bundled with the raw visualizations. As direct file writes may not be possible in a client side JavaScript application, one solution is server-based rendering and

¹⁶ <http://mathjs.org>

¹⁷ <https://www.scipy.org>

¹⁸ <http://elephant.readthedocs.io>

¹⁹ <http://www.geuz.org/gl2ps>

storage based on visualization parameters being communicated from the client back to the server. The resulting server-side rendered images are then stored as provenance information. Another option for reproducible visualization outcome is to only store the previously mentioned visualization parameters in the database, such that the client-side visualization application can be set back into the original captured state. If these parameters are captured over a longer period, the resulting data can ease the regeneration of content for demonstration purposes or a post-hoc video rendering.

While we here develop our arguments along model data, the different views and the reference implementation is equally suited for the exploration of experimental data. Our model network describes a neuronal layer covering a $4 \times 4 \text{ mm}^2$ patch of cortical tissue. Electrophysiological measurements of neuronal activity with the Utah multi-electrode array from Blackrock Microsystems sample both spiking activity of individual cells and population LFPs across near $4 \times 4 \text{ mm}^2$ of cortex (Milekovic et al., 2015; Torre et al., 2016b; Denker et al., 2018). Other multi-electrode arrays are used for in vitro experimentation on neural tissue or cell cultures (Massobrio et al., 2015). No changes to the reference implementation are required for the processing of these data. Other measurement modalities are of interest as well. One common experimental method is Ca^{2+} imaging which may infer changes in intracellular $[\text{Ca}^{2+}]$ of neurons in superficial (Grienberger and Konnerth, 2012) and deep layers (Ouzounov et al., 2017), while another method is voltage-sensitive dye imaging (VSDi) that measures membrane-voltage time-derivatives in surface-proximal tissues (Chemla and Chavane, 2010). With modifications to existing views or new view implementations, VIOLA can also represent this type of spatiotemporally resolved data. In particular the 3D visualization types incorporated in the present views 2 and 4 are well suited to represent the changes in intracellular Ca^{2+} ion concentrations across different cell bodies from 2- and 3-photon volumetric Ca^{2+} imaging in neural tissue. Within view 2 the visual representations of each cell's concentration can be set to a depth and position in the horizontal plane according to its image stack position in the raw imaging data. Units with baseline Ca^{2+} concentrations may then be hidden, and increasing levels can be visualized by scaling the box sizes as we have demonstrated with spike-rate data. A view similar to view 4 could show time-varying ion-concentrations of individual units as 3D tube plots where the tube diameter at a given time is proportional to a unit's Ca^{2+} concentration. As VSDi imaging data (typically) lack depth-information, color-image plotting can be applied similar to what we utilize here to show LFPs in our 3D view implementations. In addition, the multi-view aspect of visualization enables the combination of spatial representations with more abstract non-spatial representations of neuronal activity, as reviewed in the introduction.

The concepts developed here advance the visual exploration of data from cortical networks at cellular resolution. If the reference implementation finds more widespread interest it can be further developed by a community driven approach as all requirements like a proper licensing and a suitable development platform are in place, the primary purpose, however, is to serve as a living supplement to this publication. Creating a common web portal for the collaboration of neuroscientists is a central long-term goal of the Human Brain Project. In this endeavor our study

contributes knowledge on how a user interface for visual exploration needs to be designed and on the proper layout of the software stack at the troubled transition point between data processing and visualization.

CONDITIONS FOR TRAVELING WAVES IN SPIKING NEURAL NETWORKS

This chapter is based on the following publication:

Johanna Senk, Karolína Korvasová, Jannis Schuecker, Espen Hagen, Tom Tetzlaff, Markus Diesmann, Moritz Helias (2018), Conditions for traveling waves in spiking neural networks, *arXiv preprint arXiv:1801.06046v1*

Author contributions:

Under the supervision of Moritz Helias and Markus Diesmann, the author performed all parts of this publication. The sections about linear stability analysis and simulations with nonlinear rate neurons were developed jointly with Karolína Korvasová. The author wrote the first draft of the manuscript; the other authors co-wrote the manuscript.

5.1 INTRODUCTION

Experimental recordings of neural activity frequently reveal spatiotemporal patterns such as traveling waves propagating across the cortical surface (Rubino et al., 2006; Nauhaus et al., 2009; Muller and Destexhe, 2012; Sato et al., 2012; Muller et al., 2014; Townsend et al., 2015; Zanos et al., 2015; Denker et al., 2018) or within other brain regions such as the thalamus (Kim et al., 1995; Muller and Destexhe, 2012) or the hippocampus (Lubenov and Siapas, 2009). These large-scale dynamical phenomena are detected in local-field potentials (LFPs), see Riehle et al. (2013), and in the spiking activity (Takahashi et al., 2015) recorded with multi-electrode arrays, by voltage-sensitive dye imaging (Ferezou et al., 2006), or by two-photon imaging monitoring the intracellular calcium concentration (Garaschuk et al., 2000). They have been reported in in-vitro and in in-vivo experiments, in both anesthetized and awake states, and during spontaneous as well as stimulus-evoked activity (Muller and Destexhe, 2012).

Previous modeling studies have shown that networks of spiking neurons with distance-dependent connectivity, extending in one- or two-dimensional space, can exhibit a variety of such spatiotemporal patterns (Mehring et al., 2003; Yger et al., 2011; Voges and Perrinet, 2012; Keane and Gong, 2015). For illustration, consider the example in Figure 5.1. Depending on the choice of transmission delays, the spatial reach of connections and the strength of inhibition, a network of leaky integrate-and-fire (LIF) model neurons generates asynchronous-irregular activity (A), spatial patterns that are persistent in time (B), spatially uniform temporal oscillations (C), or propagating waves (D). Distance-dependent connectivity is a prominent feature of biological networks. In the neocortex, local connections are established within a radius of about $500\ \mu\text{m}$ around a neuron's cell body (Voges et al., 2010), and the probability of two neurons being connected decays with distance (Hellwig, 2000; Perin et al., 2011; Schnepel et al., 2015).

So far, the formation of spatiotemporal patterns in neural networks has mainly been studied by means of phenomenological neural-field models describing network dynamics at a macroscopic spatial scale (Wilson and Cowan, 1972b, 1973b; Amari, 1977). Such models can describe patterns in recorded brain activity that are related to movement (Erlhagen, 1997) or occur in response to a visual stimulus (Bressloff and Carroll, 2015). Neural-field models are formulated with continuous nonlinear integro-differential equations for a spatially and temporally resolved activity variable and usually possess an effective distance-dependent connectivity kernel. These models provide insights into the existence and uniqueness of diverse patterns which are stationary or non-stationary in space and time, such as waves, wave fronts, bumps, pulses, and periodic patterns (reviewed in Ermentrout, 1998; Coombes, 2005; Wyller et al., 2007a; Coombes, 2010; Bressloff, 2012, 2014; Coombes et al., 2014). There are two main techniques for analyzing spatiotemporal patterns in neural-field models (Bressloff, 2012): First, in the constructive approach introduced by Amari (1977), bump or wave solutions are explicitly constructed by relating the spatial and temporal coordinates of a nonlinear system (reviewed in Ermentrout, 1998, Section 7 and Bressloff, 2012, Sections 3–4). Second, the emergence of periodic patterns is studied with bifurcation theory as in the seminal works of Ermentrout and Cowan (1979a; 1979b; 1980a; 1980b). In

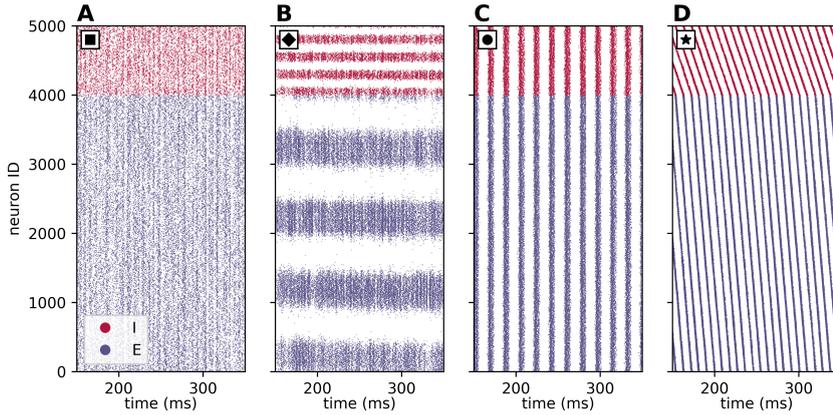


Figure 5.1: **Spatiotemporal patterns in a spiking neural network model.** Spiking activity of recurrently connected populations of excitatory (E, blue) and inhibitory (I, red) leaky integrate-and-fire neurons. Each dot represents the spike-emission time of a particular neuron. Neurons within each population are equally spaced on a ring with perimeter of 1 mm. Each neuron receives a fixed number of incoming connections from its excitatory (inhibitory) neighbors uniformly and randomly drawn within a distance of R_E (R_I). The spike-transmission delay d , the widths R_E and R_I of the spatial connectivity profiles, and the relative inhibitory synaptic weight g are varied. **A** Asynchronous-irregular activity ($d = 1$ ms, $R_E = R_I = 0.4$ mm, $g = 6$). **B** Oscillations in space ($d = 3$ ms, $R_E = 0.1$ mm, $R_I = 0.15$ mm, $g = 5$). **C** Oscillations in time ($d = 6$ ms, $R_E = R_I = 0.4$ mm, $g = 7$). **D** Propagating waves ($d = 3$ ms, $R_E = 0.2$ mm, $R_I = 0.07$ mm, $g = 5$). For remaining parameters, see Table 5.3.

this latter framework, linear stability analysis is often employed to detect pattern-forming instabilities and to derive conditions for the onset of pattern formation (see, for example, Bressloff, 1996 and Hutt et al., 2003 or the reviews Ermentrout, 1998, Section 8 and Bressloff, 2012, Section 5). There are four general classes of states that can linearly bifurcate from a homogeneous steady state: a new uniform stationary state, temporal oscillations (spatially uniform and periodic in time, also known as global ‘bulk oscillations’, see Bressloff and Kilpatrick, 2008), spatial oscillations (spatially periodic and stationary in time), and periodic traveling waves (spatially and temporally periodic), see Ermentrout (1998, Section 8) and also other studies (Roxin et al., 2005; Atay and Hutt, 2006; Venkov et al., 2007). The analysis of these states is often called ‘(linear) Turing instability analysis’ (Coombes, 2005; Coombes et al., 2007; Venkov et al., 2007) referring to the work of Turing on patterns in reaction-diffusion systems (Turing, 1952). The respective instabilities leading to these states are termed: a firing rate instability, Hopf instability (Kuramoto, 1984), Turing instability, and Turing-Hopf (Roxin et al., 2005) or ‘wave’ (Hutt et al., 2003) instability. The instabilities generating temporally periodic patterns (Hopf and Turing-Hopf instabilities) are known as ‘dynamic’ (Venkov et al., 2007) or ‘nonstationary’ (Hutt and Atay, 2005) instabilities, in contrast to ‘static’ (Venkov et al., 2007) or ‘stationary’ (Hutt and Atay, 2005) instabilities generating temporally stationary patterns. The emergence of pattern-forming instabilities has been investigated with respect to system parameters such as the spatial reach of excitation and inhibition in an effective connectivity profile (Ermentrout, 1998); specifically without transmission delays (Wyller et al., 2007b; Folias and Ermentrout, 2012), or with constant (Roxin et al., 2005, 2006), distance-dependent (Jirsa and Kelso, 2000; Hutt et al., 2003; Atay and Hutt, 2005, 2006; Coombes et al., 2007; Bressloff and Kilpatrick, 2008; Hutt, 2008; Bojak and Liley, 2010; Hutt and Rougier, 2010) or both types (Veltz, 2011, 2013) of delays.

Neural-field models treat neural tissue as a continuous excitable medium and describe neural activity in terms of a space and time dependent real-valued quantity. Throughout the current work the spatial coordinate refers to physical space, although in general it could also be interpreted as feature space. At the microscopic scale, in contrast, neural networks are composed of discrete units (neurons) – which interact via occasional short stereotypical pulses (spikes) rather than continuous quantities like firing rates. In the neocortex, spiking activity is typically highly irregular and sparse (Softky and Koch, 1993; Brunel and Hakim, 1999), with weak pairwise correlations (Ecker et al., 2010). To date, a rigorous link between this microscopic level and the macroscopic description by neural-field models is lacking (Coombes, 2010; Bressloff, 2014; Hutt et al., 2015; Montbrió et al., 2015). While randomly connected spiking networks have been extensively analyzed using mean-field approaches (Amit and Brunel, 1997a; Brunel and Hakim, 1999; Brunel, 2000; Lindner et al., 2005), the theoretical understanding of spatially structured spiking networks is still deficient. Hence, it remains unclear how to qualitatively transfer insights on the formation of spatiotemporal patterns from neural fields to networks of spiking neurons. Moreover, it is unknown how the multitude of neuron, synapse and connectivity parameters of spiking neural networks relates to the effective parameters in neural-field models. A quantitative link between the two levels of description is, for example, required for adjusting

parameters in a network of spiking neurons such that it generates a specific type of spatiotemporal pattern, and to enable model validation by comparison with experimental data.

Different efforts have already been undertaken to match spiking neuron models and time-continuous rate models with spatial structure. Certain assumptions and approximations allow the application of techniques for analyzing spatiotemporal patterns developed for neural-field models. The above mentioned constructive approach (Amari, 1977), for example, can be applied to networks of spiking neurons under the assumption that every neuron spikes at most once, thus ignoring the sustained spike generation and after-spike dynamics of biological neurons (Golomb and Ermentrout, 2001; Cremers and Herz, 2002; Osan and Ermentrout, 2002). A related simplification substitutes a spike train by an ansatz for a wave front. This leads to a mean-field description of single-spike activity often applied to a spike-response model (Fohlmeister et al., 1995; Kistler et al., 1998; Kistler, 2000; Bressloff, 2000). Traveling-wave solutions have also been proposed for a network of coupled oscillators and a corresponding continuum model (Crook et al., 1997). In the framework of bifurcation theory, Roxin et al. (2005; 2006) demonstrate a qualitative agreement between a neural-field model and a numerically simulated network of Hodgkin-Huxley-type neurons in terms of emerging spatiotemporal patterns. However, the authors do not observe stable traveling waves in the spiking network, even though the neural-field model predicts their occurrence. In the limit of slow synaptic interactions, spiking dynamics can be reduced to a mean-firing-rate model for studying bifurcations (Ermentrout, 1994; Bressloff and Coombes, 1998, 2000). An example is the lighthouse model (Haken, 2000a,b), defined as a hybrid between a phase oscillator and a firing-rate model, that reduces to a pure rate model for slow synapses (Chow and Coombes, 2006). Laing and Chow (2001) demonstrate a bump solution in a spiking network and discuss a corresponding rate model. Recently, the group around Doiron and Rosenbaum explored in a sequence of studies spatially structured networks of LIF neurons without transmission delays in the continuum limit with respect to the spatial widths of connectivity. The authors focus on the existence of the balanced state (Rosenbaum and Doiron, 2014), the structure of correlations in the spiking activity (Rosenbaum et al., 2017), and bifurcations in the linearized dynamics in relation to network computations (Pyle and Rosenbaum, 2017). Kriener et al. (2014b) employ static mean-field theory and extend the linearization of a network of LIF neurons with constant delays as described by Brunel (2000) to spatially structured networks. The work derives conditions for the appearance of spontaneous symmetry breaking that leads to stationary periodic bump solutions (spatial oscillations), and distinguishes between the mean-driven and the fluctuation-driven regime.

Despite these previous works on spatially structured network models of spiking neurons and attempts to link them with neural-field models, there still exists no systematic way of mapping parameters between these models. Furthermore, none of these studies focuses on uncovering the underlying mechanism of periodic traveling waves in spiking networks. In the present work, we establish the so far missing, quantitative link between a sparsely connected network of spiking LIF neurons with spatial structure and a typical neural-field model. An explicit parameter mapping between the two levels of description allows us to study the origin of

spatiotemporal patterns analytically in the neural-field model using linear stability analysis, and to reproduce the predicted patterns in spiking activity. We employ mean-field theory to derive the neural-field model as an effective rate model depending on the dynamical working point of the network that is characterized by both the mean and the variance of the synaptic input. The rate model accounts for biological constraints such as a static weight that is either positive (excitatory) or negative (inhibitory) and a spatial profile that can be interpreted as a distance-dependent connection probability. Given these constraints, we show that periodic traveling waves cannot occur in a single homogeneous population irrespective of the shape of distance-dependent connection probability. For two-population networks of excitatory and inhibitory neurons, in contrast, traveling waves emerge for specific types of spatial profiles and for sufficiently large delays, as shown in Figure 5.1D.

The remainder of this chapter is structured as follows: In Section 5.2 we derive the conditions for the existence of periodic traveling waves for a typical neural-field model by linear stability analysis, present an effective model corresponding to the microscopic description of spiking neurons, compare the two models, and show simulation results for validation. Section 5.3 contains details on our approach. Finally, in Section 5.4 we put our results in the context of previous literature.

5.2 RESULTS

We aim to establish a mapping between two different levels of description for spatially structured neural systems to which we refer as ‘neural-field model’ and ‘spiking model’ based on the initial model assumptions. While the neural-field model describes neural activity as a quantity that is continuous in space and time, the spiking model assumes a network of recurrently connected spiking model neurons in discrete space. Our methodological approach for mapping between these two models, as well as the structure of this section, are illustrated in Figure 5.2. (1) We start in Sections 5.2.1–5.2.3 with linear stability analysis of a typical neural-field model that is a well-known and analytically tractable rate equation. This approach builds on existing literature (cf. Ermentrout, 1998, Section 8 and Bressloff, 2012, Section 5) and introduces the concepts of our study with modest mathematical efforts. We analyze the neural-field model for one and two populations and derive conditions for the occurrence of periodic traveling waves based on spatial connectivity profiles and transmission delays. (2) In Section 5.2.4 we continue with simulations of a discrete version of the neural-field model, a network of nonlinear rate-based units, and show that the results from our linear analysis indeed accurately predict transitions between network states (homogeneously steady, spatial oscillations, temporal oscillations, waves). (3) Then, in Section 5.2.5 we linearize the population dynamics of networks of discrete spiking leaky integrate-and-fire (LIF) neurons using mean-field theory and derive expressions similar to the neural-field model. (4) Thus, both the linearized neural-field and spiking models can be treated in a conceptually similar manner, with the exception of an effective coupling kernel which is mathematically more involved for the spiking model. In Section 5.2.6 we perform a parameter mapping between the biophysically motivated parameters of the spiking model and the effective parameters of a neural-field model. (5) Finally,

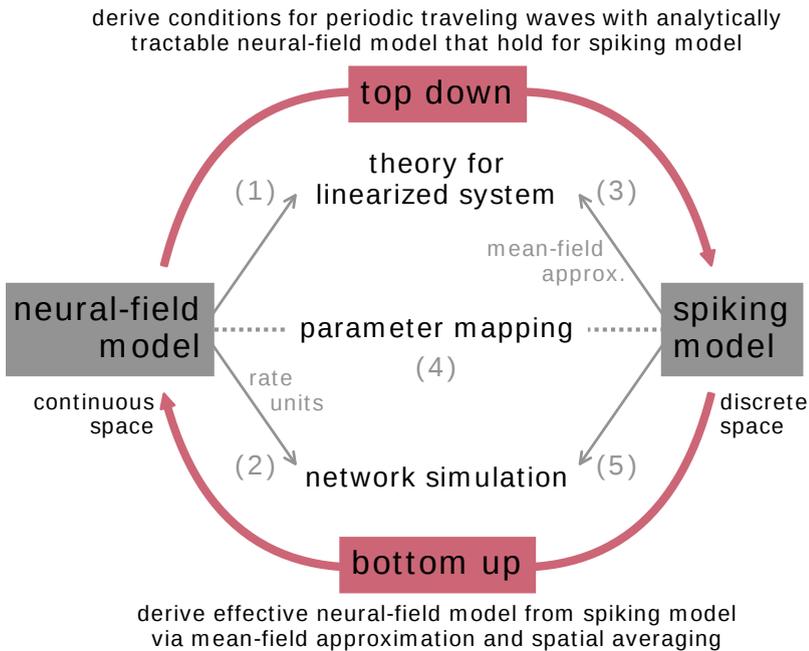


Figure 5.2: **Mapping microscopic single-neuron dynamics to spatially averaged population dynamics.** (1) Conditions for periodic traveling waves in a neural-field model. (2) Network simulation of discrete nonlinear rate neurons. (3) Mean-field approximation of the spiking model and spatial averaging lead to an effective linearized continuous system. (4) Parameter mapping between spiking and neural-field model. (5) Network simulation of spiking neurons and validation of analytical results.

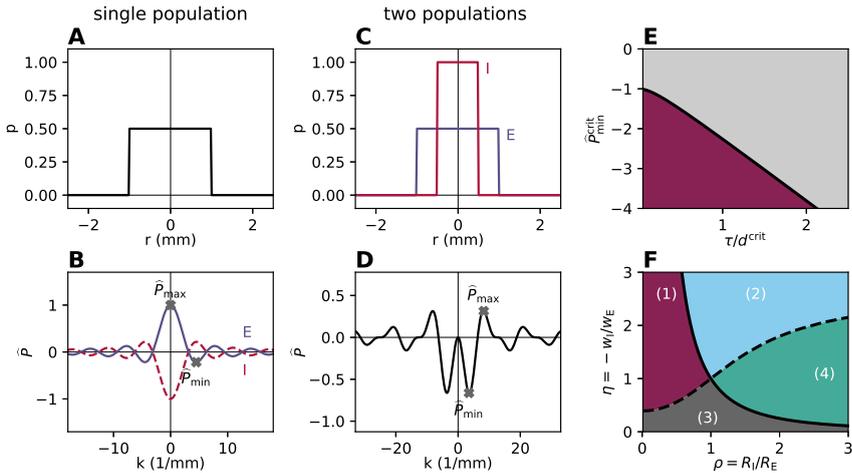


Figure 5.3: **Effective profile yields conditions for traveling waves.** **A** Boxcar-shaped spatial profile p of width $R = 1$ mm for a single population. **B** Effective profile \hat{P} (blue curve) denotes Fourier transform of spatial profile \hat{p} times positive weight $w_E = 1$. Gray crosses indicate maximum \hat{P}_{\max} and minimum \hat{P}_{\min} . Same spatial profile but with negative weight ($w_I = -w_E$) yields mirrored curve (red, dashed line). **C** Spatial profiles of different widths for two populations E ($R_E = 1$ mm, blue) and I ($R_I = 0.5$ mm, red). **D** Effective profile: $\hat{P}(k) = w_E \hat{p}_E(k) + w_I \hat{p}_I(k)$. **E** Transition curve $\hat{P}_{\min}^{\text{crit}}(\tau/d^{\text{crit}})$ given by Equation 5.10 for Hopf bifurcation indicating onset of delay-induced oscillations (appearing in purple region) with time constant τ and delay d . **F** Transition curves for relative width $\rho = R_I/R_E$ and relative weight $\eta = -w_I/w_E$. Colored regions indicate which extremum, the minimum \hat{P}_{\min} or the maximum \hat{P}_{\max} , has larger absolute value and if the dominant one occurs at $k = 0$ or at $k > 0$. (1, purple): \hat{P}_{\min} appears at $k_{\min} > 0$. (2, light blue): \hat{P}_{\min} appears at $k_{\min} = 0$. (3, dark gray): \hat{P}_{\max} appears at $k_{\max} = 0$. (4, green): \hat{P}_{\max} appears at $k_{\max} > 0$.

in Section 5.2.7 we demonstrate that the insights obtained in the analysis of the neural-field model apply to networks of simulated LIF neurons: The bifurcations indeed appear at the theoretically predicted parameter values.

In summary, the mapping of a microscopic spiking network model to a continuum neural-field model (bottom up) allows us to transfer analytically derived insights from the neural-field model directly to the spiking model (top down).

5.2.1 Linear stability analysis of a neural-field model

We first consider a neural-field model with a single population defined as a continuous excitable medium with a translation-invariant interaction kernel and delayed interaction in one spatial dimension. The dynamics follow an integro-differential equation

$$\tau \frac{du}{dt}(x, t) + u(x, t) = \int_{-\infty}^{\infty} P(x - y) \psi(u(y, t - d)) dy. \quad (5.1)$$

The variable u describes the activity of the neural population at position x and time t . Here $\tau > 0$ denotes a time constant and $d > 0$ a transmission delay. The function ψ describes the nonlinear transformation of the output activity u if considered as input to the neural field. The function P specifies the translation-invariant connectivity depending only on the displacement $r = x - y$ where x and y denote neuron positions. Earlier studies show that specific choices for connectivities P and nonlinear transformations ψ result in spatiotemporal patterns such as waves or bumps (Ermentrout, 1998; Coombes, 2005; Wyller et al., 2007a; Coombes, 2010; Bressloff, 2012, 2014; Coombes et al., 2014).

Here, we assume that the connectivity P is isotropic and define $P(r) := w p(r)$. The scalar weight w can either be positive (excitatory) or negative (inhibitory). The spatial profile $p(r)$ is a symmetric probability density function with the properties $p(r) = p(-r)$, $p(r) > 0$ for $r \in (-\infty, \infty)$ and $\int_{-\infty}^{\infty} p(r) dr = 1$. Figure 5.3A shows, as an example, a boxcar-shaped spatial profile with width R , defined by $p(r) = \frac{1}{2R} \Theta(R - |r|)$ where Θ denotes the Heaviside function.

Throughout this study we investigate bifurcations of the system in Equation 5.1 between a state of spatially and temporally homogeneous activity $u(x, t) = u_0$ to states where the activity shows structure in the temporal domain, in the spatial domain, or both. For this purpose we use Turing instability analysis (Bressloff, 1996; Hutt et al., 2003; Coombes, 2005). Initially we assume that the model parameters are chosen such that the homogeneous solution is locally asymptotically stable, implying that small perturbations away from u_0 will relax back to this baseline. We ask the question: In which regions of the parameter space (R, d, w, ψ) is the stability of the homogeneous solution lost? To this end we linearize around the steady state and denote deviations $\delta u(t) = u(t) - u_0$. Without loss of generality we assume the slope $\psi'(u_0)$ of the gain function to be unity; a non-zero slope can be absorbed into a redefinition of w . Because the resulting system is linear and invariant with respect to translations in time and space, its eigenmodes are Fourier-Laplace modes of the form

$$\delta u(x, t) = e^{ikx} e^{\lambda t}, \quad (5.2)$$

where the wave number $k \in \mathbb{R}$ is real and the temporal eigenvalue $\lambda \in \mathbb{C}$ is complex. Solutions constructed from these eigenmodes can oscillate in time and space, and exponentially grow or decay in time. The characteristic equation (see Equation 5.31 in Section 5.3)

$$(1 + \tau\lambda) e^{\lambda d} = \widehat{P}(k), \quad (5.3)$$

comprises the effective profile $\widehat{P}(k) := w \widehat{p}(k)$. The Fourier transform of the spatial profile is denoted by $\widehat{p}(k)$ which, by its definition as a probability density, is maximal at $k = 0$ with $\widehat{p}(0) = 1$ (see Equations 5.38 and 5.39 in Section 5.3). The effective profile for the boxcar-shaped spatial profile is shown in Figure 5.3B, for excitatory and inhibitory weights with absolute magnitudes of unity.

We next extend the system to two populations, an excitatory one denoted by E , and an inhibitory one denoted by I . Time constants τ and delays d are assumed to be equal for both populations, but u becomes a vector, $u = (u_E, u_I)^T$, and the connectivity $P(r)$ a matrix

$$P(r) = \begin{pmatrix} w_{EE} p_{EE}(r) & w_{EI} p_{EI}(r) \\ w_{IE} p_{IE}(r) & w_{II} p_{II}(r) \end{pmatrix}. \quad (5.4)$$

The linearized system again possesses the same symmetries as the counterpart for a single population so that the eigenmodes for the deviation from the stationary state are of the form $\delta u(x, t) = v e^{ikx} e^{\lambda t}$ with a constant vector v . Hence, we arrive at an auxiliary eigenvalue problem (see Equation 5.32 in Section 5.3) with the two eigenvalues

$$\hat{P}_{1,2}(k) = \frac{1}{2} \left(w_{EE} \hat{p}_{EE}(k) + w_{II} \hat{p}_{II}(k) \pm \sqrt{D} \right), \quad (5.5)$$

where

$$D = (w_{EE} \hat{p}_{EE}(k) - w_{II} \hat{p}_{II}(k))^2 + 4w_{EI} \hat{p}_{EI}(k) w_{IE} \hat{p}_{IE}(k). \quad (5.6)$$

These two eigenvalues play the same role as the effective profile \hat{P} in the one-population case above. As a consequence, the same characteristic equation (Equation 5.3) holds for both the one- and the two-population system.

In the following example, we restrict the weights and the spatial profiles to be uniquely determined by the source population alone, denoted by $w_{\alpha E} =: w_E, w_{\alpha I} =: w_I$ for $\alpha \in \{E, I\}$. An illustration of the two spatial profiles of different widths R_E and R_I is shown in Figure 5.3C. The respective effective profile (Equation 5.5) reducing to $\hat{P}(k) = w_E \hat{p}_E(k) + w_I \hat{p}_I(k)$ is shown in Figure 5.3D.

The characteristic equation (Equation 5.3) can be solved for the eigenvalues λ by using the Lambert W function defined as $z = W(z) e^{W(z)}$ for $z \in \mathbb{C}$ (Corless et al., 1996). The Lambert W function has infinitely many branches, indexed by b , and the branch with the largest real part is denoted the principle branch ($b = 0$), see Equations 5.34–5.36 in Section 5.3 for a proof. The characteristic equation determines the temporal eigenvalues (see Equation 5.37 in Section 5.3 and compare with Veltz, 2013)

$$\lambda_b(k) = -\frac{1}{\tau} + \frac{1}{d} W_b \left(\hat{P}(k) \frac{d}{\tau} e^{\frac{d}{\tau}} \right). \quad (5.7)$$

5.2.2 Conditions for spatial and temporal oscillations, and traveling waves

The homogeneous (steady) state of our system is locally asymptotically stable if the real parts of all eigenvalues λ_b are negative

$$\operatorname{Re} \left[W_b \left(\hat{P}(k) \frac{d}{\tau} e^{\frac{d}{\tau}} \right) \right] < \frac{d}{\tau}, \quad (5.8)$$

for all branches b of the Lambert W function. The system loses stability when the real part of the eigenvalue λ_0 on the principle branch becomes positive at a certain $k = k^*$. Such instabilities may occur either for a positive or a negative argument of the Lambert W function.

We denote the maximum of \hat{P} as \hat{P}_{\max} and the minimum as \hat{P}_{\min} occurring at k_{\max} and k_{\min} , respectively, as indicated in Figure 5.3B and D. The system becomes unstable for a positive argument of W if $\hat{P}_{\max} = 1$ where $\operatorname{Re} \left[W_0 \left(\frac{d}{\tau} e^{\frac{d}{\tau}} \right) \right] = \frac{d}{\tau}$

by the definition of the Lambert W function; so equality holds in Equation 5.8 independent of the values d and τ . The imaginary part of λ_0 is zero at such a transition. If the instability appears at a wave number $k^* = 0$, the population activity is collectively destabilized. This transition corresponds in networks of binary neurons and of spiking neurons to the transition between the asynchronous irregular (AI) state and the synchronous regular (SR) state, where the system ceases to be stabilized by negative feedback and leaves the balanced state (van Vreeswijk and Sompolinsky, 1996; Brunel, 2000). If this transition appears at a wave number $k^* > 0$, it follows from Equation 5.2 that the activity shows spatial oscillations that grow exponentially in time.

For a negative argument of W of less than $-1/e$, the eigenvalues (Equation 5.7) come in complex conjugate pairs. The real part of λ_0 becomes positive if the condition

$$\operatorname{Re} \left[W_0 \left(\hat{P}_{\min} \frac{d}{\tau} e^{\frac{d}{\tau}} \right) \right] = \frac{d}{\tau} \quad (5.9)$$

is fulfilled with a negative $\hat{P}_{\min} < -1$. Because the eigenvalues have non-zero imaginary parts, this transition corresponds to a Hopf bifurcation and the onset of temporal oscillations. The condition for this bifurcation has been derived earlier (Helias et al., 2013, Equation 10)

$$\frac{d^{\text{crit}}}{\tau} = \frac{\pi - \arctan \left(\sqrt{\hat{P}_{\min}^{\text{crit}^2} - 1} \right)}{\sqrt{\hat{P}_{\min}^{\text{crit}^2} - 1}}. \quad (5.10)$$

Here, d^{crit} denotes a critical delay and $\hat{P}_{\min}^{\text{crit}}$ a critical minimum of the effective profile for points on the transition curve. The system is stable for $\hat{P}_{\min} > -1$ for all delays. For larger absolute values of \hat{P}_{\min} , the bifurcation point is given by the critical value of the ratio between the time constant and the delay, shown in Figure 5.3E. If the transition occurs at $k^* = 0$, temporal oscillations emerge in which all neurons of the population oscillate in phase ('bulk oscillations', see Bressloff and Kilpatrick, 2008). In spiking networks this Hopf bifurcation corresponds to the transition from the AI regime to the state termed 'synchronous irregular fast (SI fast)', see Brunel and Hakim (1999). If the transition appears for $k^* > 0$, spatial and temporal oscillations occur simultaneously. This phenomenon is known as 'periodic traveling waves', see Ermentrout (1998, Section 8) and also other studies (Roxin et al., 2005; Atay and Hutt, 2006; Venkov et al., 2007). For the case that the system becomes unstable due to \hat{P}_{\max} reaching unity, the transition curve in Figure 5.3E also provides a lower bound $\hat{P}_{\min}^{\text{crit}}(\tau/d^{\text{crit}})$ above which temporal oscillations do not occur prior to the transition due to \hat{P}_{\max} .

In summary, the system is stable if $\hat{P}_{\max} < 1$ and $\hat{P}_{\min} > \hat{P}_{\min}^{\text{crit}}(\tau/d^{\text{crit}})$. For transitions occurring at either $\hat{P}_{\max} = 1$ or $\hat{P}_{\min} = \hat{P}_{\min}^{\text{crit}}(\tau/d^{\text{crit}})$ we distinguish between solutions with $k^* = 0$ or $k^* > 0$. In Figure 5.4 we provide an overview of the conditions for bifurcations leading to spatial, temporal, or spatiotemporal oscillatory states. These conditions imply that a one-population neural-field model does not permit traveling waves, which follows from the fact that the absolute value of \hat{p} is strictly maximal at $k = 0$ (see Equations 5.38–5.39 in Section 5.3). For a purely excitatory population ($w > 0$) the critical minimum $\hat{P}_{\min}^{\text{crit}}(\tau/d^{\text{crit}})$ therefore

	homogeneous	spatial oscillations	temporal oscillations	traveling waves
\hat{P}_{\max}	< 1	1	< 1	< 1
\hat{P}_{\min}	$> \hat{P}_{\min}^{\text{crit}}$	$> \hat{P}_{\min}^{\text{crit}}$	$\hat{P}_{\min}^{\text{crit}}$	$\hat{P}_{\min}^{\text{crit}}$
d	$< d^{\text{crit}}$	$< d^{\text{crit}}$	d^{crit}	d^{crit}
k^*	-	> 0	0	> 0

Figure 5.4: **Conditions for the onset of spatial and temporal oscillations, and traveling waves.** Gray cells in each column indicate the conditions required for the instability causing the bifurcation. White cells denote the conditions for the respective other bifurcation not to occur. Last row indicates whether the bifurcation happens for zero or nonzero wave number k^* . Here d^{crit} and $\hat{P}_{\min}^{\text{crit}}$, as defined in Equation 5.10 and shown in Figure 5.3E, denote the critical delay and the minimum of the effective profile on the transition curve for a Hopf bifurcation.

cannot be reached while keeping the maximum \hat{P}_{\max} stable as $\hat{P}_{\max} > |\hat{P}_{\min}|$. For a purely inhibitory population ($w < 0$), the condition $k_{\min} > 0$ is not fulfilled because \hat{P}_{\min} occurs at $k = 0$ as \hat{p} has its global maximum at the origin.

For a neural-field model accounting for both excitation and inhibition, however, we can select shapes and parameters of the spatial profiles, weights and the delay that fulfill the conditions for the onset of traveling waves as demonstrated by example in the next section.

5.2.3 Application to a network with excitation and inhibition

Based on the conditions derived in the previous section, the minimal network in which traveling waves can occur consists of one excitatory (E) and one inhibitory (I) population. As in the example in Figure 5.3, we assume that the connection weights and widths of boxcar-shaped spatial profiles only depend on the source population. The effective profile (Equation 5.5) in this case is

$$\hat{P}(k) = w_E \frac{\sin(R_E k)}{R_E k} + w_I \frac{\sin(R_I k)}{R_I k}, \quad (5.11)$$

and positive and negative peaks of the profile are responsible for bifurcations to spatial or temporal oscillations or wave solutions, respectively. The previous section derives that in particular the position and height of the minima and maxima of the effective profile are decisive. To assess parameter ranges in which the peaks of the effective profile in Equation 5.11 change qualitatively, we introduce the relative width $\rho := R_I/R_E > 0$ and the relative weight $\eta := -w_I/w_E > 0$, divide $\hat{P}(k)$ by w_E and introduce the rescaled wave number $\kappa = R_E k$ to arrive at the dimensionless reduced profile

$$\hat{B}(\kappa) = \frac{\sin(\kappa)}{\kappa} - \eta \frac{\sin(\rho\kappa)}{\rho\kappa}, \quad (5.12)$$

which simplifies the following analysis.

Our aim is to divide the parameter space (ρ, η) into regions that have qualitatively similar shapes of the effective profile. Section 5.3 describes the derivation of transition curves and Figure 5.3F illustrates the resulting parameter space. Above the first transition curve $\eta_{11}(\rho)$ (dashed curve, see Equation 5.46 in Section 5.3), the absolute value of \hat{B}_{\min} is larger than \hat{B}_{\max} (regions 1 and 2), and vice versa below this curve (regions 3 and 4). The second transition curve $\eta_{12}(\rho)$ (solid curve, see Equation 5.49 in Section 5.3) indicates whether the extremum with the largest absolute value occurs at $k = 0$ (regions 2 and 3) or at $k > 0$ (regions 1 and 4). The diagram provides the necessary conditions and corresponding parameter combinations required for both spatial and spatiotemporal patterns, purely based on the relative weights and the relative widths which determine the effective profile. The analysis shows that traveling waves require wider excitation than inhibition, $\rho < 1$, because only this relation simultaneously realizes a minimum at a non-zero wave number k^* and a maximum with a peak below unity (see Figure 5.4).

A neural-field model exhibiting traveling waves can therefore be constructed at will by first selecting a point within region 1 of Figure 5.3F where $\rho < 1$ and η ensures that $|\hat{B}_{\min}| > \hat{B}_{\max}$. Next, \hat{P} is fixed by scaling \hat{B} with the absolute weight w_E such that $\hat{P}_{\max} < 1$ for a stable bump solution and $\hat{P}_{\min} < -1$ for a Hopf bifurcation. Finally, a delay $d > d^{\text{crit}}$ specifies a point below the bifurcation curve shown in Figure 5.3E, given by the sufficient condition for the Hopf bifurcation in Equation 5.10. Likewise, solutions for purely temporal oscillations appear in region 2, where $\hat{P}_{\min} < -1$ is attained at a vanishing wave number k and a delay $d > d^{\text{crit}}$; in addition $\hat{P}_{\max} < 1$ ensures absence of the other bifurcation into spatial oscillations. For purely spatial oscillations, however, the comparison of the absolute values of \hat{B}_{\min} and \hat{B}_{\max} is not sufficient; it is hence not sufficient to rely on the dashed curve separating regions 2 and 4 in Figure 5.3F. A loss of stability due to $\hat{P}_{\max} > 1$ can emerge not only in region 4 but also in region 2, because even if $|\hat{P}_{\min}| > \hat{P}_{\max}$, stability of \hat{P}_{\min} can be ensured by a sufficiently short delay $d < d^{\text{crit}}$, as shown in Figure 5.4.

5.2.4 Network simulation with nonlinear rate neurons

We have so far only investigated the properties of an analytically tractable linear system that assumes time and space to be continuous variables. Next, we test the derived conditions for the onset of oscillations, summarized in Figure 5.4, for a nonlinear, discrete system in the continuum limit. We here consider a network of $N_E = 4,000$ excitatory (E) and $N_I = 1,000$ inhibitory (I) rate neurons described by a discrete version of the neural-field equation in Equation 5.1 (see Table 5.2 for details). The model neurons within each population are equally spaced on a ring of perimeter $L = 1$ mm. This rate-neuron network constitutes an intermediate step towards a network of spiking neurons. Each neuron has a fixed in-degree K_X (fixed number of incoming connections) per source population $X \in \{E, I\}$ with connections selected randomly within a distance R_X . A normalization of weights with the in-degree, $w'_X = w_X/K_X$, allows us to interpret p as a connection probability. The time constant τ and the delay d are the same as in the neural-field model. As nonlinear gain function in Equation 5.1 we choose $\psi(u) = \tanh(u)$.

The neuron activity of four rate-network simulations with different parameter combinations are shown in Figure 5.5A–D. The location of the specific parameter combinations is illustrated in Figure 5.5E–G with corresponding markers in the phase diagrams that visualize the stability conditions shown in Figure 5.3 derived with the neural-field model. Periodic traveling waves are possible if parameters are in the purple regions of the diagrams.

The system simulated in Figure 5.5A is stable according to the corresponding conditions. The square marker in the lower panels shows that $\hat{P}_{\max} < 1$ (panel E), and although $\hat{P}_{\min} < -1$, the delay is small such that the system is far away from the bifurcation (panel F). Indeed, the activity appears to not exhibit any spatial or temporal structure.

Figure 5.5B illustrates a case where $\hat{P}_{\max} > 1$ causes an instability (diamond marker in panel E). The Hopf bifurcation is remote in the parameter space (panel F) and panel G ensures $k_{\max} > 0$. A simulation of the corresponding rate-model network again confirms the predictions and exhibits stationary spatial oscillations (or periodic bumps) with a wave number of k_{\max} . In this finite-sized system with periodic boundary conditions, the bumps are homogeneously distributed across the domain and the wave numbers are integers.

Figure 5.5C demonstrates temporal oscillations at the parameter combination indicated by the circular marker. We here choose $\hat{P}_{\max} < 1$ and $\hat{P}_{\min} < -1$ (panel E). The latter condition leads to an entire range of delays that are beyond the bifurcation in panel F; we choose a delay slightly larger than the critical delay, lying to the left of the bifurcation curve. Inferred from panel G, $k_{\min} = 0$ and, as expected from the analytical prediction, the oscillations observed in simulations of the rate-neuron network are purely temporal.

Finally, Figure 5.5D depicts periodic traveling waves (denoted by star marker), as predicted by the analytically tractable neural-field model. The instability results from $\hat{P}_{\min} < \hat{P}_{\min}^{\text{crit}}$ (panel F) and occurs at $k_{\min} > 0$ (panel G) while \hat{P}_{\max} remains stable (panel E).

5.2.5 Linearization of spiking network model

To assess the validity of the predictions obtained from the analytical model for biologically more realistic spiking-neuron networks, we next linearize the dynamics of spiking leaky integrate-and-fire (LIF) neurons and derive a linear system similar to the neural-field model above. The sub-threshold dynamics of a single LIF neuron i with exponentially decaying synaptic currents is described by a set of differential equations for the time evolution of the membrane potential V_i and its synaptic current I_i as

$$\begin{aligned} \tau_m \frac{dV_i}{dt} &= -V_i + I_i(t), \\ \tau_s \frac{dI_i}{dt} &= -I_i + \tau_m \sum_j J_{ij} s_j(t-d), \end{aligned} \tag{5.13}$$

where we follow the convention of Fourcaud-Trocme and Brunel (2005), see Equation 5.61 in Section 5.3 for the relation to physical units. This definition, with

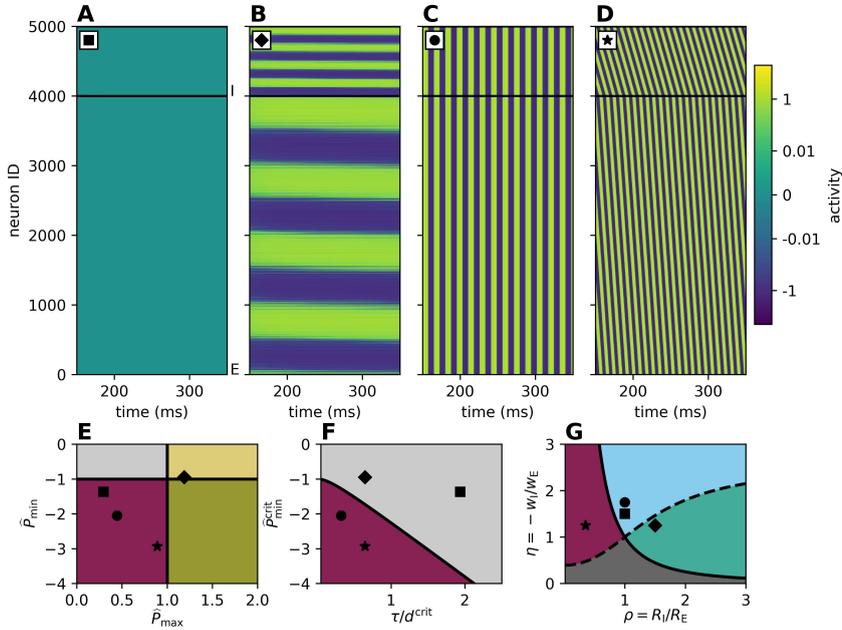


Figure 5.5: **Predictions from linear stability analysis lead to spatiotemporal patterns in simulated network of nonlinear rate neurons.** Different parameter combinations, selected according to stability conditions in Figure 5.4, cause pattern formation in rate-neuron network with tanh gain function. **A–D** Color-coded deviation from the activity δu per neuron over time. Neurons within each population are consecutively numbered with a ‘neuron ID’ according to their position on the ring, and neuron IDs of inhibitory neurons follow the ones of excitatory neurons. **E–F** Phase diagrams showing conditions and parameter choices indicated by corresponding markers. Purple regions indicate the possibility for periodic traveling waves. **A** Stable activity (square marker). **B** Spatial oscillations (diamond marker). **C** Temporal oscillations (circular marker). **D** Periodic traveling waves (star marker). Parameters: d , R_E and R_I as in Figure 5.1A–D, $w_E = 2.73$ in all panels. **A** $w_1 = -4.10$. **B** $w_1 = -3.42$. **C** $w_1 = -4.79$. **D** $w_1 = -3.42$.

both quantities V_i and I_i having the same unit, conserves the total integrated charge per impulse flowing into the membrane independent of the choice of the synaptic time constant τ_s . The membrane time constant, defined as $\tau_m = R_m C_m$ with membrane resistance R_m and membrane capacitance C_m , couples the current to the capacitance. We here assume τ_s to be much smaller than τ_m . The term $s_j(t) = \sum_k \delta(t - t_k^j)$ denotes a spike train of neuron j which is connected to neuron i with a constant connection strength J_{ij} and transmission delay d . Whenever V_i reaches the threshold V_θ , a spike is emitted and the membrane potential is reset to the resting potential V_r and voltage-clamped for the refractory period τ_{ref} .

Assuming that a neuron receives many uncorrelated and Poisson-distributed input spikes, and that amplitudes of postsynaptic potentials are small, we apply the diffusion approximation (Tuckwell, 1988; Amit and Brunel, 1997a; Ricciardi et al., 1999) and approximate the input to the neuron by a current with mean μ_i and variance σ_i as given by

$$\tau_m \sum_j J_{ij} s_j(t-d) \approx \mu_i(t) + \sqrt{\tau_m} \sigma_i(t) \xi(t). \quad (5.14)$$

The term $\xi(t)$ denotes a Gaussian white noise characterized by $\langle \xi(t) \rangle = 0$ and $\langle \xi(t) \xi(t') \rangle = \delta(t - t')$. Next, we introduce the instantaneous firing rate $v_j(t)$ of neuron j . If all presynaptic neurons j belong to a homogeneous population of identical neurons with uncorrelated activity and a postsynaptic neuron i receives input from K of these neurons, the mean and variance (first and second infinitesimal moments, Ricciardi et al., 1999) of the input current to neuron i are given by

$$\begin{aligned} \mu_i(t) &= \tau_m \sum_j J_{ij} v_j(t-d) = \tau_m J K v(t-d), \\ \sigma_i^2(t) &= \tau_m \sum_j J_{ij}^2 v_j(t-d) = \tau_m J^2 K v(t-d), \end{aligned} \quad (5.15)$$

where $v \equiv \langle v_j \rangle$ represents the population-averaged instantaneous firing rate.

Such a mean-field approach has been employed previously to study networks of spiking neurons without spatial structure (Amit and Brunel, 1997b; Brunel and Hakim, 1999; Brunel, 2000; Lindner et al., 2005). We extend on this approach by assuming that the neurons are placed on a discrete one-dimensional domain with an inter-neuron space constant Δx . In the continuum limit $\Delta x \ll 1$ we discard the single-neuron index i and use a continuous variable $x \in \mathbb{R}$ for space, replacing $\mu_i(t) \rightarrow \mu(x, t)$. We preserve the in-degree K and establish connections according to the symmetric and normalized connection probability $p(r)$. As before, p depends only on distance, the absolute value of the displacement $r = x - y$ for neurons at positions x and y .

In a spatially and temporally homogeneous state we may describe the state of the system by a stationary firing rate independent of time and space: $v(x, t) = v_0$ (see Equation 5.50 in Section 5.3). To investigate the stability of the spatially homogeneous stationary state, we consider a small excursion δv away from the stationary firing rate v_0 ,

$$v(x, t) = v_0 + \delta v(x, t), \quad \delta v \ll v_0, \quad (5.16)$$

caused by a perturbation of the synaptic input and study its effect on the dynamics of v . The response of the firing rate to the synaptic input can be approximated to linear order in δv by applying linear response theory to the Fokker-Planck equation (Risken, 1996), and expressing δv in terms of the temporal linear convolution kernels $h_\mu(t)$ and $h_{\sigma^2}(t)$ (Schuecker et al., 2015) as

$$\delta v(x, t) = [h_\mu * \delta\mu](x, t) + [h_{\sigma^2} * \delta\sigma^2](x, t). \quad (5.17)$$

The convolution operation is purely temporal, and the form of the response kernels is given in Equations 5.51–5.53 in Section 5.3.

As the network is recurrently connected, an excursion of the firing rate in turn leads to a perturbation of the mean of the synaptic input and its variance according to

$$\begin{aligned} \delta\mu(x, t) &= \tau_m JK \int_{-\infty}^{\infty} p(x-y) \delta v(y, t-d) dy \\ \delta\sigma^2(x, t) &= \tau_m J^2 K \int_{-\infty}^{\infty} p(x-y) \delta v(y, t-d) dy. \end{aligned} \quad (5.18)$$

In the following section, however, we ignore the h_{σ^2} terms because their contributions are usually small (Schuecker et al., 2015). The combination of Equations 5.17 and 5.18 provides a linearized system for the spiking model that is continuous in space and time and enables a direct comparison with the neural-field model in the following section.

5.2.6 Comparison of neural-field and spiking models

The linearization of the LIF model presented in the preceding section is the analogue to taking the derivative ψ' of the gain function in the linear stability analysis of the neural-field model in Section 5.2.1. Therefore, the results for the neural-field model carry over to the spiking case. To expose the similarities between the linearized systems of the spiking model and the neural-field model, we may bring the equations for the deviation from baseline activity

$$\delta o(x, t) = \begin{cases} \delta u(x, t) & \text{neural field} \\ \delta v(x, t) & \text{spiking} \end{cases} \quad (5.19)$$

to the form of the convolution equation

$$\begin{aligned} \delta o(x, t) &= [\tilde{h} * \delta i](x, t) \\ \delta i(x, t) &= \int_{-\infty}^{\infty} p(x-y) \delta o(y, t-d) dy, \end{aligned} \quad (5.20)$$

where the only difference is the convolution kernel relating the deviations from the input δi to those of the output δo defined as

$$\tilde{h}(t) := \begin{cases} \tilde{h}^{\text{nf}}(t) := \Theta(t) \frac{w}{\tau} e^{-\frac{t}{\tau}} & \text{neural field} \\ \tilde{h}^{\text{s}}(t) := \tau_m JK h_\mu(t) & \text{spiking.} \end{cases} \quad (5.21)$$

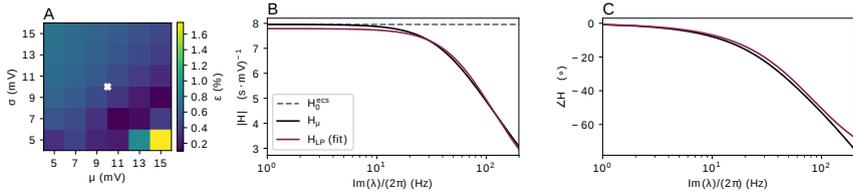


Figure 5.6: **Transfer function of spiking neuron model and its approximation.** **A** Fitting error of the low-pass filter approximation of the transfer function for LIF neurons derived in (Schuecker et al., 2015) over μ and σ (given relative to the reset potential). The fitting error $\epsilon = \sqrt{\epsilon_\tau^2 + \epsilon_{H_0}^2}$ is color-coded. **B** Amplitude of the transfer function and approximation (legend). Dashed line illustrates H_0 following from the analytically-determined effective coupling strength (see Equation 5.54 in Section 5.3). **C** Phase. The white cross in panel A indicates the working point (μ, σ) selected for the transfer function shown in panels B and C and used in the simulations throughout the study.

The kernel on the first line is the fundamental solution (Green's function) of the linear differential operator appearing on the left hand side of Equation 5.1, including the coupling weight w . As a consequence, the characteristic equations for both models result from the Fourier-Laplace ansatz $\delta o(x, t) = e^{ikx}e^{\lambda t}$ which relates the eigenvalues λ to the wave number k as

$$\tilde{H}(\lambda) \cdot e^{-\lambda d} \cdot \hat{p}(k) = 1. \quad (5.22)$$

The effective transfer function $\tilde{H}(\lambda)$ is defined as the Laplace transform of Equation 5.21 of the respective functions for the spiking model $\tilde{h}^s(t)$ and for the neural-field model $\tilde{h}^{\text{nf}}(t)$. As a result, we obtain the transfer function for the neural-field model

$$\tilde{H}^{\text{nf}}(\lambda) = \frac{1}{1 + \lambda\tau} w. \quad (5.23)$$

The corresponding expression for the effective spiking transfer function $\tilde{H}^s(\lambda)$ results from Equations 5.51–5.53 in Section 5.3.

5.2.6.1 Parameter mapping

So far the stability analysis shows that the characteristic equations for both the neural-field and the spiking model have the same form (Equation 5.22) given a proper definition of the respective transfer functions. The transfer function characterizes the transmission of a small fluctuation in the input to the output of the neuron model. Because these transfer functions differ between the two models, it is a priori unclear whether their characteristic equations have qualitatively similar solutions.

The transfer function of the LIF model in the fluctuation-driven regime investigated here can, however, be approximated by a first order low-pass (LP) filter (Lindner and Schimansky-Geier, 2001; Brunel et al., 2001; Helias et al., 2013)

$$H_\mu(\lambda) \approx H_{\text{LP}}(\lambda) = \frac{H_0}{1 + \lambda\tau}. \quad (5.24)$$

This simplified transfer function is of identical form as the transfer function (Equation 5.23) of the neural-field model, and thereby relates the phenomenological parameters w and τ of the neural-field model to the biophysically motivated parameters of the spiking model.

Fitting the absolute values of $H_{\text{LP}}(\lambda)$ to $H_\mu(\lambda)$ yields values for the parameters τ and H_0 . According to Equation 5.21, H_0 directly relates to w as

$$w = H_0 \tau_m J K. \quad (5.25)$$

The goodness of the fit of this transfer function to the first-order low pass filter depends on the mean μ and variance σ of the synaptic input, as shown in Figure 5.6A. The color-coded error of the fit combines the relative errors from both fitting parameters: $\epsilon = \sqrt{\epsilon_\tau^2 + \epsilon_{H_0}^2}$. For the majority of working points (μ, σ) the error is $< 1\%$ but the relative errors increase abruptly towards the mean-driven regime. In this regime input fluctuations are small and the mean input predominantly drives the membrane potential towards threshold, so that the model fires regularly and the transfer function exhibits a peak close to the firing frequency (Lindner and Schimansky-Geier, 2001; Brunel et al., 2001). We here fix the working point to the parameters indicated by the white cross (see Equation 5.59 in Section 5.3) for all populations, resulting in a common effective time constant τ . Here, we obtain a time constant $\tau = 1.94$ ms which thus lies in between the synaptic time constant, $\tau_s = 0.5$ ms, and the membrane time constant, $\tau_m = 5$ ms, of the LIF neuron model. For these parameters, Figure 5.6B shows the amplitude and Figure 5.6C the phase of the original transfer function $H_\mu(\lambda)$ in black and the fitted transfer function $H_{\text{LP}}(\lambda)$ in purple. The dashed gray line denotes H_0 obtained by computing the effective coupling strength from linear response theory, H_0^{ecs} , as a reference (see Equation 5.54 in Section 5.3).

5.2.6.2 Linear interpolation between the transfer functions

Evaluating the characteristic equation for the neural-field model yields an exact solution for each branch of the Lambert W function, given by Equation 5.7. For this model we already established that the principle branch is the most unstable one. An equivalent condition is not known for the general response kernel of the LIF neuron. To assess whether we may transfer the result for the neural-field model to the spiking case, we investigate the correspondence between the two characteristic equations that are both of the form Equation 5.22 but with different transfer functions. For this purpose, we define an effective transfer function

$$\tilde{H}_\alpha(\lambda) = \alpha \tilde{H}^s(\lambda) + (1 - \alpha) \tilde{H}^{\text{nf}}(\lambda), \quad (5.26)$$

with the parameter α that linearly interpolates between the effective transfer functions of the spiking and the neural-field model: $\tilde{H}_{\alpha=0}(\lambda) = \tilde{H}^{\text{nf}}(\lambda)$ and $\tilde{H}_{\alpha=1}(\lambda) = \tilde{H}^s(\lambda)$. Figure 5.7 illustrates two different ways for solving the combined characteristic equation

$$\tilde{H}_\alpha(\lambda) \cdot e^{-\lambda d} \cdot \hat{p}(k) = 1. \quad (5.27)$$

The first results from computing the derivative $d\lambda/d\alpha$ (see Equations 5.55–5.58 in Section 5.3) from the combined characteristic equation and integrating numerically

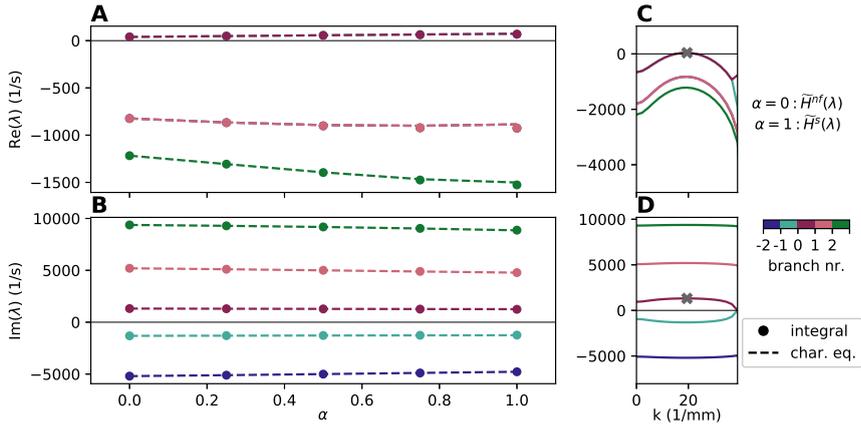


Figure 5.7: **Linear interpolation between neural-field ($\alpha = 0$) and spiking ($\alpha = 1$) model for eigenvalue close to bifurcation.** A Real and B imaginary part of the eigenvalue λ as a function of the linear interpolation parameter α for the characteristic equation in Equation 5.27. The solution at $\alpha = 0$ for the neural-field model is exact. C Real and D imaginary part of the eigenvalues (same units but different scaling as in A and B) with analytically exact solution (by Lambert W function, $\alpha = 0$) as functions of the wave number k . Different branches b are color-coded (legend); $b = 0$ corresponds to the principal branch with the maximum real eigenvalue (gray cross). Circular markers denote the linear interpolation according to the numerical integration of Equation 5.28. Dashed line segments for the linear interpolation are obtained by solving the characteristic equation (Equation 5.27) numerically. Both are evaluated at the same values for α . Parameters: $d = 1.5$ ms, $R_E = 0.2$ mm, $R_I = 0.07$ mm, $g = 5$.

with the exact solution of the neural-field model at $\alpha = 0$ for each branch b as initial condition:

$$\lambda(\alpha) = \int_0^\alpha \frac{d\lambda}{d\alpha'} d\alpha', \quad \lambda(0) = \lambda_b \quad (5.28)$$

with

$$\frac{d\lambda}{d\alpha} = -\frac{\tilde{H}^s(\lambda) - \tilde{H}^{nf}(\lambda)}{\alpha \frac{\partial \tilde{H}^s(\lambda)}{\partial \lambda} + (1 - \alpha) \frac{\partial \tilde{H}^{nf}(\lambda)}{\partial \lambda} - d \cdot \tilde{H}_\alpha(\lambda)}. \quad (5.29)$$

The spatial profile only enters the initial condition, and the derivative (Equation 5.29) is independent of the wave number k .

As an alternative approach, we directly solve the combined characteristic equation (Equation 5.27) numerically with the known initial condition. Figure 5.7A and B indicate that only the principle branch ($b = 0$) becomes positive while the other branches remain stable. The branches come in complex conjugate pairs. For the numerical solution of the characteristic equation, we fix the wave number to the value of k that corresponds to the maximum real eigenvalue.

The analysis shows that we may ignore the danger of branch crossing since different branches remain clearly separated in Figure 5.7A and B. In addition, the

eigenvalue on the principle branch is mostly independent of α , even if the system is close to the bifurcation (when the real part of λ_0 is close to zero). Thus for all values of α we expect qualitatively similar bifurcations, including $\alpha = 1$. This justification transfers the rigorous results from the bifurcation analysis of the neural-field model in Sections 5.2.2 and 5.2.3, and corresponding effective parameters, to the spiking model.

5.2.7 Validation by simulation of spiking neural network

Section 5.1 illustrates spatiotemporal patterns emerging in a spiking network simulation in Figure 5.1 and the subsequent sections derive a theory describing the mechanisms underlying such patterns. Finally, the parameter mapping between the spiking and the neural-field model explains the origin of the spike patterns by transferring the conditions found for the abstract neural-field model in Sections 5.2.2 and 5.2.3 to the spiking case. This section validates that the correspondence between network parameters in the two models is not incidental but covers the full phase diagram.

In the following, we simulate a network with the same neural populations and spatial connectivity used in the nonlinear rate-network in Figure 5.5, but replace the rate-model neurons by spiking neurons, and map the parameters as described in Section 5.2.6.1. The network model characterizes all neurons by the same working point (see Equation 5.59 in Section 5.3), which means that the connectivity matrix for the excitatory-inhibitory network has equal rows; entries in Equation 5.4 depend on the presynaptic population alone. Therefore the relative in-degree $\gamma = K_I/K_E$ and the relative synaptic strength $g = -J_I/J_E$ parametrize the spiking-network connectivity matrix as

$$P(r) = \tau_m J_E K_E \begin{pmatrix} p_E(r) & -\gamma g p_I(r) \\ p_E(r) & -\gamma g p_I(r) \end{pmatrix}. \quad (5.30)$$

The rightmost panels of Figure 5.8A–C show the same simulation results as Figure 5.1B–D; likewise the panels of Figure 5.1 have parameters that correspond to those of the rate-neuron network in Figure 5.5. The different patterns in Figure 5.1B–D emerge by gradually shifting a single network parameter that switches the system from a stable state (white filled markers in Figure 5.8D and E), across intermediate states (gray-scale filled markers) to the final states where stability is lost and the patterns have formed (black filled markers). Arrows visualize the sequences in the phase diagrams Figure 5.8D and E and the markers reappear in the upper left corners of the corresponding raster plots in Figure 5.8A–C.

The sequence of panels in Figure 5.8A illustrates a gradual transition from a stable (AI) state to spatial oscillations attained by increasing the amplitudes of excitatory postsynaptic current (PSC) amplitudes J'_E in the network. With J' we denote the weight as a jump in current while J denotes a jump in voltage in the physical sense, and the relationship is: $J' = C_m J / \tau_s$ (see Equation 5.61 in Section 5.3). The parameter variation thus homogeneously scales the effective profile \hat{P} but preserves the shape of the reduced profile \hat{B} (fixed position of diamond marker in panel F). Simultaneously an increasing rate of the external Poisson input compensates for the reduced PSC amplitudes to maintain the fixed working

point (μ, σ) of the neurons (see Equation 5.59 in Section 5.3). Diamond markers in Figure 5.8D show that along its path the system crosses the critical value $\hat{P}_{\max} = 1$, while $\hat{P}_{\min} > \hat{P}_{\min}^{\text{crit}}(\tau/d^{\text{crit}})$ stays in the stable regime, as shown in panel E. However, even for $\hat{P}_{\max} \lesssim 1$ (for $J'_E = 60$ pA) the network activity already exhibits weak spatial oscillations.

Choosing the synaptic delay d as a bifurcation parameter highlights the onset of temporal oscillations for the case $k = 0$ (panel B sequence, circular markers) and spatiotemporal oscillations for the case $k > 0$ (sequence in Figure 5.8C, star markers). In contrast to the case of purely spatial waves in panel A, the procedure preserves the effective spatial profile (fixed positions in panels D and F) and the system crosses the transition curve in panel E due to increasing delay alone, thus decreasing the ratio τ/d .

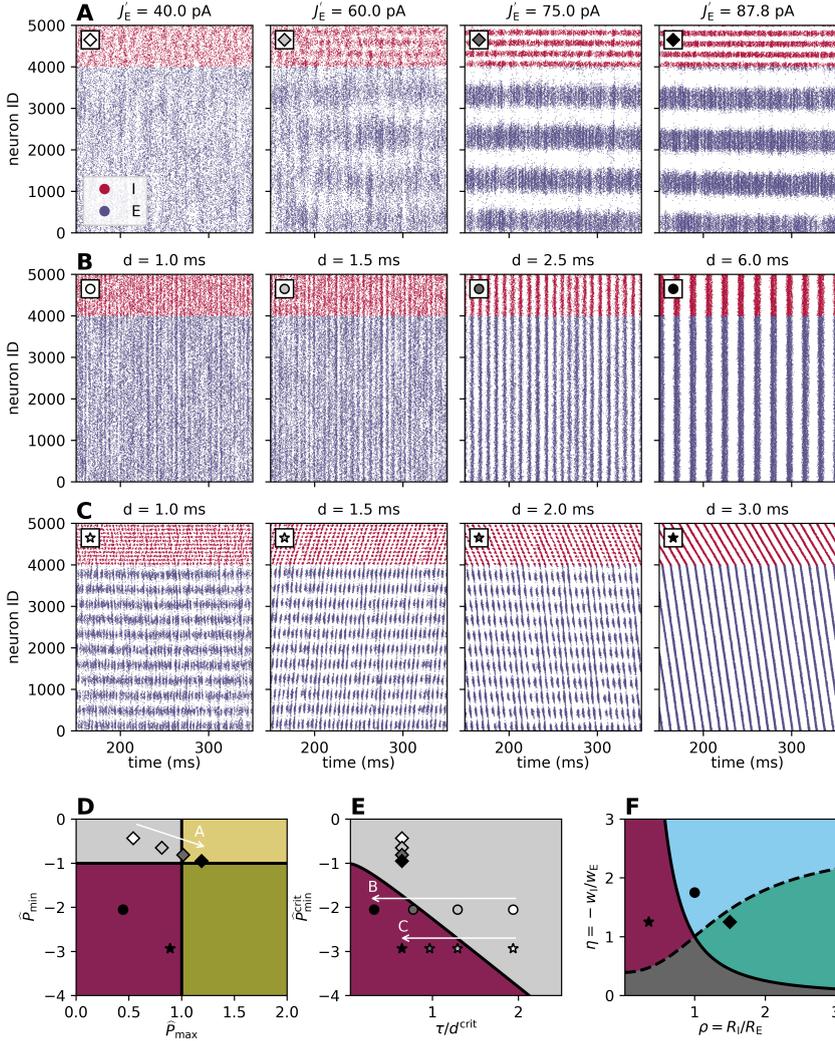


Figure 5.8: **Transitions from theoretically stable states to spatiotemporal patterns in spiking network simulation.** A–C Spike rasters showing transition to network states in Figure 5.1B–D (same markers, same parameter combinations). The changed parameter value is given on top of each raster plot. **A** Increasing recurrent weight J'_E leads to onset of spatial oscillations. **B** Increasing synaptic delay d leads to onset of temporal oscillations at $k = 0$. **C** Increasing delay d leads to onset of temporal oscillations at $k > 0$, i.e., periodic traveling waves. **D–E** Gray shaded markers and white arrows labeled according to respective panel A–C in phase diagrams indicate sequences of parameter combinations and breakdown of stability at $\hat{p}_{\max} = 1$ or at $\hat{p}_{\min} = \hat{p}_{\min}^{\text{crit}}$. For each sequence in panels A–C, delay d , excitatory profile width R_E , inhibitory profile width R_I , and the relative synaptic strength g correspond to the values given in Figure 5.1B–D with corresponding markers.

Figure 5.8C illustrates the gradual transition to traveling waves, where \hat{P}_{\max} remains in the theoretically stable regime at all times, but is close to the critical value of 1 (see the star marker in panel D). As a result, we observe spatial oscillations with a spatial frequency given by k_{\max} before and even after the Hopf bifurcation. For delays longer than the critical delay, mixed states occur in which different instabilities due to \hat{P}_{\max} and \hat{P}_{\min} compete. For delay values well past the bifurcation, this mixed state is lost resulting in a dependency only on \hat{P}_{\min} and periodic traveling waves with a spatial frequency that depends on k_{\min} .

5.3 METHODS

5.3.1 Linear stability analysis

5.3.1.1 Derivation of the characteristic equation

With the Fourier-Laplace ansatz $u(x, t) = e^{ikx}e^{\lambda t}$ for the integro-differential equation in Equation 5.1 linearized around u_0 and the choice to set the slope of the gain function to unity, the characteristic equation in Equation 5.3 results from

$$\begin{aligned} \tau\lambda e^{ikx}e^{\lambda t} &= -e^{ikx}e^{\lambda t} + \int_{-\infty}^{\infty} wp(x-y)e^{iky}e^{\lambda(t-d)} dy \\ \tau\lambda &= -1 + we^{-\lambda d} \int_{-\infty}^{\infty} p(x-y)e^{-ik(x-y)} dy \\ &= -1 - we^{-\lambda d} \int_{-\infty}^{\infty} p(r)e^{-ikr} dr, \quad r = x - y \quad (5.31) \\ &= -1 + we^{-\lambda d} \underbrace{\int_{-\infty}^{\infty} p(r)e^{-ikr} dr}_{\equiv \hat{p}(k)}. \end{aligned}$$

In the last row, we recognize the Fourier transform \hat{p} of the spatial profile p .

5.3.1.2 Effective connectivity profile for two populations

While the connectivity P is a scalar in the one-population model, it is a matrix in the case of two populations (given in Equation 5.4). The ansatz for deriving the characteristic equation in the latter case reads $\delta u(x, t) = ve^{ikx}e^{\lambda t}$, with v denoting a vector of constants. This leads to the auxiliary eigenvalue problem

$$\hat{P}(k)v = \tilde{P}(k)v, \quad (5.32)$$

where \hat{P} denotes an eigenvalue and \tilde{P} is an auxiliary matrix containing the Fourier transforms of the entries of P :

$$\tilde{P}(k) = \begin{pmatrix} w_{EE}\hat{p}_{EE}(k) & w_{EI}\hat{p}_{EI}(k) \\ w_{IE}\hat{p}_{IE}(k) & w_{II}\hat{p}_{II}(k) \end{pmatrix}. \quad (5.33)$$

Equation 5.32 possesses a nontrivial solution v if and only if $\det(\tilde{P}(k) - \hat{P}(k)\mathbb{1}) = 0$. Equation 5.5 explicitly states the two eigenvalues $\hat{P}_{1,2}$ solving this equation. These eigenvalues constitute the effective profile in the characteristic equation in Equation 5.3 that hence holds also for the two-population case.

5.3.1.3 Largest real part on principle branch of Lambert W function

The function $x(W) = We^W$ has a minimum at $W = -1$, no real solution for $x < -e^{-1}$, a single solution for $x > 0$, and two solutions for $x \in [-e^{-1}, 0)$. Typically, the term ‘principal branch’ of the Lambert W function with branch number $b = 0$ refers to the real branch defined on the interval $[-e^{-1}, \infty)$, where for negative arguments the larger solution is considered. Here we extend the definition to the whole real line by the complex branch with maximal real part and positive imaginary part on $(-\infty, -e^{-1})$.

We demonstrate that the branch of the Lambert W function with the largest real part is the principal branch. Considering only real-valued arguments $x \in \mathbb{R}$, we write $W(x) = |W(x)|e^{i\varphi} = \alpha + i\beta$ and

$$W(x)e^{W(x)} = |W(x)|e^{\alpha}e^{i(\varphi+\beta)} = x \in \mathbb{R} \quad (5.34)$$

$$\rightarrow e^{i(\varphi+\beta)} = \pm 1, \quad (5.35)$$

where $\varphi \in [-\pi, \pi]$ is the principal value. We index the branches by $q \in \mathbb{Z}$ according to the number of half-cycles of the exponential in Equation 5.35: $\varphi + \beta = q \cdot \pi$. The branch number is equal to $b = \lfloor \frac{q}{2} \rfloor$ with $\lfloor \cdot \rfloor$ denoting the floor function. The principle branch is therefore given by the index $q = 0$ for $x \geq 0$ and by $q = 1$ for $x < 0$.

Taking the absolute square of Equation 5.34 yields the real equation

$$x^2 e^{-2\alpha} = \alpha^2 + \beta^2. \quad (5.36)$$

Without loss of generality we may assume $\beta \geq 0$; this is certainly true for the real solutions with $\beta = 0$ and it also holds for one of the complex solutions for any complex pair. Complex solutions come in conjugate pairs due to the symmetry $(\varphi, \beta) \rightarrow (-\varphi, -\beta)$ exhibited by Equations 5.35 and 5.36. Since each member of a pair has by definition the same real part, it is sufficient to consider only the member with positive imaginary part $\beta > 0$.

To prove that the real part α of W is maximal for $b = 0$, we show that α is a decreasing function of β along the solutions of Equation 5.34. Investigating the intersections of the left-hand side and the right-hand side of Equation 5.36 as a function of α illustrates how increasing the imaginary part β affects the real part α . The left-hand side is a decaying function of α with an intercept of x^2 . The right-hand-side is a parabola with an offset of β^2 .

For $x \in (-\infty, -e^{-1}) \cup [0, \infty)$, an intersection occurs either at a positive real part $\alpha \geq 0$ if $x^2 \geq \beta^2$, or at a negative real part $\alpha < 0$ if $x^2 < \beta^2$. Increasing β moves the parabola upwards and therefore the intersection to the left, meaning that α decreases with increasing β .

For $x \in [-e^{-1}, 0)$, we distinguish the cases $\beta = 0$ and $\beta > 0$ which both have only solutions with $\alpha < 0$. First, the two real solutions ($q = \pm 1$) existing in this interval correspond to two simultaneously occurring intersections; in addition a third intersection is created by the squaring (Equation 5.36) but it is not an actual solution of Equation 5.34. The intersection at the larger real part per definition corresponds to the principal branch with index $q = 1$. Second, the complex solutions are indexed by odd numbers q with $|q| > 1$. Taking into account the interval

where φ is defined, the imaginary part is bounded from below such that $\beta \geq 2\pi$ for non-principal branches. Analogous to the previously discussed interval of x , there exists only one intersection between the exponential function and the parabola for large values of β (in particular: $x^2 < \beta^2$) that moves towards smaller values of α with increasing β .

So in summary we have shown that for real x , the principal branch harbors the solutions with maximal real part α .

5.3.1.4 Characteristic equation with Lambert W function

The characteristic equation in Equation 5.3 can be rewritten in terms of the Lambert W function to Equation 5.7 using the transformation:

$$\begin{aligned} (1 + \tau\lambda) e^{\lambda d} &= \widehat{P}(k) \cdot \frac{d}{\tau} e^{\frac{d}{\tau}} \\ \left(d\lambda + \frac{d}{\tau}\right) e^{d\lambda + \frac{d}{\tau}} &= \widehat{P}(k) \frac{d}{\tau} e^{\frac{d}{\tau}} \\ d\lambda + \frac{d}{\tau} &= W\left(\widehat{P}(k) \frac{d}{\tau} e^{\frac{d}{\tau}}\right). \end{aligned} \quad (5.37)$$

The last step collects terms using the definition of the Lambert W function, $z = W(z) e^{W(z)}$ with $z \in \mathbb{C}$.

5.3.2 Properties of the spatial profile

We assume that the spatial profile p is a symmetric probability density function, which implies that its Fourier transform \widehat{p} , also called the characteristic function, is real valued and even. Further, we can prove that $\widehat{p} \in (-1, 1]$ and that \widehat{p} attains 1 only at the origin in two steps:

- $|\widehat{p}(k)| \leq 1$ for all $k \in \mathbb{R}$:

$$\begin{aligned} |\widehat{p}(k)| &= \left| \int_{-\infty}^{\infty} p(r) e^{-ikr} dr \right| \leq \int_{-\infty}^{\infty} |p(r) e^{-ikr}| dr \\ &= \int_{-\infty}^{\infty} p(r) dr = 1 \quad \text{for all } k \in \mathbb{R}, \end{aligned} \quad (5.38)$$

- $|\widehat{p}(k)| < 1$ for all $k \neq 0$:

$$\begin{aligned} \left| \int_{-\infty}^{\infty} p(r) e^{-ikr} dr \right| &\leq \int_{-\infty}^{\infty} p(r) |\cos(kr)| dr \\ &< \int_{-\infty}^{\infty} p(r) dr = 1 \quad \text{for all } k \neq 0, \end{aligned} \quad (5.39)$$

because $|\cos(kr)| < 1$ almost everywhere in r if $k \neq 0$.

5.3.3 Transition curves for reduced profile

We here use a graphical approach to derive the transition curves shown first in Figure 5.3F. A necessary condition for an extreme value of the reduced profile $\widehat{B}(\kappa)$ from Equation 5.12 located at κ^* is: $\frac{\partial}{\partial \kappa} \widehat{B}(\kappa) |_{\kappa^*} = 0$. With the derivative

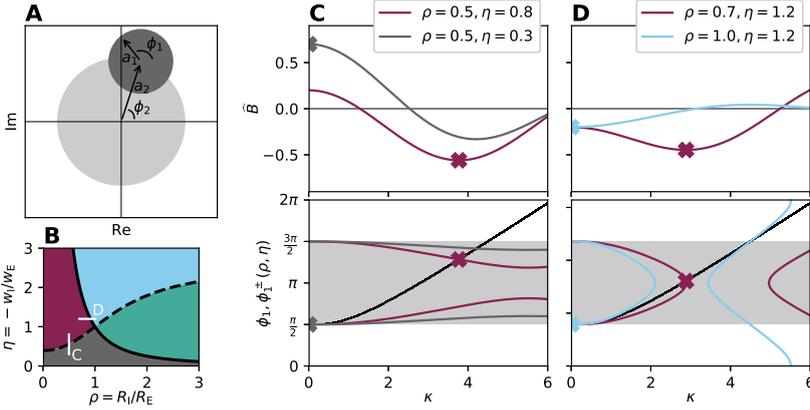


Figure 5.9: **Graphical analysis for extrema of reduced profile for derivation of transition curves.** **A** The condition for the extremum (Equation 5.41) amounts to the addition of two vectors in the complex plane whose sum is purely imaginary. The vectors have lengths a_1 and a_2 and angles ϕ_1 and ϕ_2 , defined in Equation 5.42. **B** Diagram of Figure 5.3F with indicated parameter combinations (ρ, η) as used in panels C and D. **C–D** Reduced profile \widehat{B} (top) and ϕ_1 and ϕ_1^\pm from Equation 5.43 vs. κ (bottom) for two different combinations of (ρ, η) with curve colors corresponding to regions in panel B. **C** $|\widehat{B}_{\min}| > \widehat{B}_{\max}$ in purple and vice versa in dark gray. **D** \widehat{B}_{\min} at $\kappa = 0$ in light blue and at $\kappa > 0$ in purple.

$$\frac{\partial}{\partial \kappa} \widehat{B}(\kappa) = \frac{\cos(\kappa)}{\kappa} - \frac{\sin(\kappa)}{\kappa^2} - \eta \frac{\cos(\rho\kappa)}{\kappa} + \eta \frac{\sin(\rho\kappa)}{\rho\kappa^2}, \quad (5.40)$$

this condition can be rewritten as

$$\begin{aligned} 0 &= \operatorname{Re} \left[(\kappa + i) e^{i\kappa} - \frac{\eta}{\rho} (\rho\kappa + i) e^{i\rho\kappa} \right] \\ &= \operatorname{Re} \left[a_1 e^{i\phi_1} + a_2 e^{i\phi_2} \right] \\ &= a_1 \cos(\phi_1) + a_2 \cos(\phi_2), \end{aligned} \quad (5.41)$$

where a_1 and a_2 are the absolute values of the complex numbers and ϕ_1 and ϕ_2 their phases, given by

$$\begin{aligned} a_1(\kappa) &= \sqrt{1 + \kappa^2} \\ \phi_1(\kappa) &= \kappa + \frac{\pi}{2} - \arctan(\kappa) \\ a_2(\kappa; \rho, \gamma) &= \frac{\eta}{\rho} \sqrt{1 + \rho^2 \kappa^2} \\ \phi_2(\kappa; \rho) &= \rho\kappa + \frac{3\pi}{2} - \arctan(\rho\kappa). \end{aligned} \quad (5.42)$$

The vanishing right-hand-side of Equation 5.41 implies that the term in the square brackets is purely imaginary. An example solution for the case $a_1 < a_2$ is illustrated in Figure 5.9A in the complex plane. Note that a_1 and ϕ_1 are independent

of the parameters ρ and η in this representation. In our graphical analysis, Equation 5.41 is interpreted as the sum of two vectors in the complex plane. As shown in Figure 5.9A, we determine ϕ_1 as the angle at which the tip of the second vector ends on the imaginary axis, which follows from elementary trigonometry as

$$\phi_1^\pm = \pi \pm \arccos \left(\frac{a_2}{a_1} \cos(\phi_2) \right). \quad (5.43)$$

The locations of extrema are then given by the intersections of ϕ_1^\pm with the second row of Equation 5.42. Here ϕ_2 is determined from the last equation in Equation 5.41.

Figure 5.9B reproduces Figure 5.3F. The white bars connect points given by parameter combinations (ρ, η) on both sides of the transition curves, and the parameters are specified in panels C and D. The first transition curve $\eta_{t1}(\rho)$ (dashed curve in Figure 5.9B) is determined by $\widehat{B}_{\max}(\kappa_{\max}) = \left| \widehat{B}_{\min}(\kappa_{\min}) \right|$, that means it is determined by parameters (ρ, η) for which the absolute values of the positive and negative extremum of the profile are equal. The top panel of Figure 5.9C compares two reduced profiles obtained for a fixed value for ρ and two values for η . The curve colors correspond to the colored regions in the diagram in Figure 5.9B for the respective parameter combination $\left| \widehat{B}_{\min} \right| > \widehat{B}_{\max}$ for the purple profile and vice versa for the dark gray profile. The point with the maximum absolute value of each profile is indicated with a cross. Exactly at the transition either κ_{\max} or κ_{\min} is zero (for example $\kappa_0 = 0$) and the other one is non-zero (for example $\kappa_1 > 0$). This condition, with Equation 5.12, yields the absolute value for both extrema at the transition, where they must be equal, thus $\left| \widehat{B}(\kappa_0) \right| = \left| \widehat{B}(\kappa_1) \right| = |1 - \eta|$. Any point on the transition curve is a unique triplet of parameters (ρ, η, κ_1) , and with the condition $\frac{\partial}{\partial \kappa} \widehat{B}(\kappa) \big|_{\kappa_1} = 0$ we obtain two equations that need to be fulfilled at each point for $\kappa = \kappa_1$:

$$\begin{aligned} 1 - \eta &= \frac{\sin(\kappa)}{\kappa} - \eta \frac{\sin(\rho\kappa)}{\rho\kappa} \\ 1 - \eta &= \cos(\kappa) - \eta \cos(\rho\kappa). \end{aligned} \quad (5.44)$$

The lower equation is obtained by identifying $\widehat{B}(\kappa)$ in its derivative in Equation 5.40. We solve both equations with respect to η and equate them to get

$$\frac{1}{\kappa} \sin(\kappa) [1 + \cos(\rho\kappa)] - \frac{1}{\rho\kappa} \sin(\rho\kappa) [1 + \cos(\kappa)] + \cos(\rho\kappa) - \cos(\kappa) = 0. \quad (5.45)$$

For a given value of ρ , we compute the roots of the left-hand-side expression, which defines $\kappa(\rho)$. The bottom panel of Figure 5.9C shows ϕ_1 from Equation 5.42 as a black curve and ϕ_1^\pm from Equation 5.43 for the parameters of the two effective profiles (same color coding as in the top panel). The intersections corresponding to the relevant extrema are highlighted by crosses. This visual analysis allows us to identify the interval for κ in which zero-crossings of the left-hand side of Equation 5.45 as a function of κ can correspond to the extrema, that is $\kappa \in (0, 4.49341)$ where the lower limit corresponds to $\phi_1 = \frac{\pi}{2}$ and the upper limit to $\phi_1 = \frac{3\pi}{2}$. The

zero-crossing at the smallest non-zero κ indicates the extremum at κ_1 . Finally, the transition curve is given by

$$\eta_{l1}(\rho) = \frac{1 + \cos(\kappa(\rho))}{1 + \cos(\rho\kappa(\rho))}, \quad (5.46)$$

where $\kappa(\rho)$ is given by the roots of Equation 5.45.

The second transition curve $\eta_{l2}(\rho)$ (solid curve in Figure 5.9B) indicates whether the extremum with the largest absolute value occurs at $\kappa = 0$ or at $\kappa > 0$. Figure 5.9D shows in the top panel two reduced profiles for a fixed value of η , but two values for ρ such that the \hat{B}_{\min} occurs once at $\kappa_{\min} = 0$ (light blue as in Figure 5.9B) and once at $\kappa_{\min} > 0$ (purple as in Figure 5.9B), indicated by cross markers.

Graphical analysis using the bottom panel of Figure 5.9D indicates that this transition happens when ϕ_1^- at $\kappa \gtrsim 0$ switches from lying slightly above (light blue curve) to below (purple curve) the parameter-independent function ϕ_1 (black curve). We observe that decreasing ρ moves the intersection point and with it the location of the extremum up the black line, starting from $\kappa = 0$ to larger values for κ .

Close to the transition, the intersection point comes arbitrarily close to $\kappa = 0$, which permits local analysis by a Taylor expansion of ϕ_1 for small κ :

$$\phi_1(\kappa) \approx \frac{\pi}{2} + \frac{\kappa^3}{3} + \mathcal{O}(\kappa^5) \quad (5.47)$$

$$\phi_1^-(\kappa; \rho, \eta) \approx \frac{\pi}{2} + \frac{\eta\rho\kappa^3}{3} + \mathcal{O}((\rho\kappa)^5). \quad (5.48)$$

A comparison of the coefficients of the third-order polynomials then gives the transition curve

$$\eta_{l2}(\rho) = \frac{1}{\rho^2}, \quad (5.49)$$

because this coefficient decides for small κ whether ϕ_1 (black curve) or ϕ_1^- as a function of the parameters (ρ, η) has a larger slope and lies on top.

5.3.4 Linearization of the spiking model

5.3.4.1 Stationary firing rate

The stationary firing rate ν_0 in the limit of short synaptic time constants ($\tau_s \ll \tau_m$) is given in Fourcaud and Brunel (2002) and Helias et al. (2013, Equation A.1):

$$\begin{aligned} \nu_0^{-1} &= \tau_r + \tau_m \sqrt{\pi} (F(y_\theta) - F(y_r)) \\ f(y) &= e^{y^2} (1 + \operatorname{erf}(y)), \quad F(y) = \int^y f(y') dy' \\ \text{with } y_{\{\theta,r\}} &= \frac{V_{\{\theta,r\}} - \mu}{\sigma} + \frac{\beta}{2} \sqrt{\frac{\tau_s}{\tau_m}}, \quad \beta = \sqrt{2} \left| \zeta\left(\frac{1}{2}\right) \right|, \end{aligned} \quad (5.50)$$

where ζ denotes the Riemann's zeta function (Abramowitz and Stegun, 1974).

5.3.4.2 Transfer function

The transfer function here denoted by H_μ is computed based on the first term of (Schuecker et al., 2015, Equation 29)

$$H_G(\omega) = \frac{v_0 \frac{\sqrt{2}}{\sigma}}{1 + i\omega\tau_m} \frac{\Phi'_\omega|_{x\theta}^{x_r}}{\Phi_\omega|_{x\theta}^{x_r}}, \quad (5.51)$$

for the oscillation frequency ω and $x_{\{r,\theta\}} = \sqrt{2}y_{\{\theta,r\}}$. The function $\Phi_\omega(x) = e^{\frac{1}{4}x^2} U(i\omega\tau_m - \frac{1}{2}, x)$ is defined by parabolic cylinder functions U (Abramowitz and Stegun, 1974; Lindner and Schimansky-Geier, 2001) and $\Phi'_\omega = \partial_x \Phi_\omega$. We need to multiply the transfer function with the transfer function of a first-order low-pass filter due to the exponential time course of our synaptic currents:

$$H_\mu(\omega) = H_G(\omega) \frac{1}{1 + i\omega\tau_s}. \quad (5.52)$$

We then obtain h_μ by an inverse Fourier transform and a Laplace transform because λ is a complex frequency and ω is real in the present context:

$$\begin{aligned} h_\mu(t) &= \mathcal{F}^{-1} [H_\mu](t) \\ H_\mu(\lambda) &= \mathcal{L} [h_\mu](\lambda). \end{aligned} \quad (5.53)$$

The latter relations imply a replacement $i\omega \rightarrow \lambda$ in Equation 5.51.

5.3.5 Model comparison

5.3.5.1 Effective coupling strength

For the numerical evaluation of the transfer function, we show $H_0^{\text{ecs}} = w^{\text{ecs}} / (\tau_m JK)$ as the dashed line in Figure 5.6B, obtained by calculating analytically the effective coupling strength w^{ecs} from linear-response theory. The effective coupling strength for a connection from neuron j with rate v_j to neuron i with rate v_i is defined as (Helias et al., 2013, Equations A.2 and A.3, correcting a typo in this previous work):

$$\begin{aligned} w_{ij}^{\text{ecs}} &= \frac{\partial v_i}{\partial v_j} \\ &= \tilde{\alpha} J_{ij} + \tilde{\beta} \tilde{J}_{ij}^2 \\ \text{with } \tilde{\alpha} &= \sqrt{\pi} (\tau_m v_i)^2 \frac{1}{\sigma_i} (f(y_\theta) - f(y_r)) \\ \text{and } \tilde{\beta} &= \sqrt{\pi} (\tau_m v_i)^2 \frac{1}{2\sigma_i^2} (f(y_\theta) y_\theta - f(y_r) y_r), \end{aligned} \quad (5.54)$$

where f and $y_{\{\theta,r\}}$ are defined as in Equation 5.50. The dashed line in Figure 5.6B is given by the term $\propto \tilde{\alpha}$ alone since we also ignore the small contribution of the variance to the transfer function of the LIF neuron (Schuecker et al., 2015).

5.3.5.2 Linear interpolation

To compute the derivative $d\lambda/d\alpha$ given in Equation 5.29, we use a method for computing the derivative of an implicit function: If $R(\alpha, \lambda) = 0$, it follows that the derivative

$$\frac{d\lambda}{d\alpha} = -\frac{\partial R/\partial\alpha}{\partial R/\partial\lambda} =: -\frac{R_\alpha}{R_\lambda}. \quad (5.55)$$

With the characteristic equation for the effective transfer function (Equation 5.27), we get

$$R(\alpha, \lambda) = \tilde{H}_\alpha(\lambda) \cdot e^{-\lambda d} \cdot \hat{p}(k) - 1 = 0. \quad (5.56)$$

The partial derivatives of R with respect to α and λ are

$$\begin{aligned} R_\alpha &= e^{-\lambda d} \cdot \hat{p}(k) \cdot \frac{\partial \tilde{H}_\alpha(\lambda)}{\partial \alpha} \\ &= e^{-\lambda d} \cdot \hat{p}(k) \cdot \left[\tilde{H}^s(\lambda) - \tilde{H}^{\text{nf}}(\lambda) \right], \end{aligned} \quad (5.57)$$

and

$$\begin{aligned} R_\lambda &= \hat{p}(k) \cdot \frac{\partial}{\partial \lambda} \left[\tilde{H}_\alpha(\lambda) \cdot e^{-\lambda d} \right] \\ &= e^{-\lambda d} \cdot \hat{p}(k) \cdot \left[\frac{\partial \tilde{H}_\alpha(\lambda)}{\partial \lambda} - d \cdot \tilde{H}_\alpha(\lambda) \right] \\ &= e^{-\lambda d} \cdot \hat{p}(k) \cdot \left[\alpha \frac{\partial \tilde{H}^s(\lambda)}{\partial \lambda} + (1 - \alpha) \frac{\partial \tilde{H}^{\text{nf}}(\lambda)}{\partial \lambda} - d \cdot \tilde{H}_\alpha(\lambda) \right] \\ &= e^{-\lambda d} \cdot \hat{p}(k) \cdot \left[\alpha \tilde{H}_\lambda^s(\lambda) + (1 - \alpha) \cdot \tilde{H}_\lambda^{\text{nf}}(\lambda) - d \cdot \tilde{H}_\alpha(\lambda) \right]. \end{aligned} \quad (5.58)$$

5.3.6 Fixing the working point

For the spiking model, we fix the mean input μ and its variance σ relative to their reset potential for both populations. Each neuron receives external excitatory and inhibitory input with Poisson-distributed interspike interval statistics (analogous to Helias et al., 2013, Equation E.1). The external input rates for excitatory neurons $\nu_{\text{E,ext}}$ and for inhibitory neurons $\nu_{\text{I,ext}}$ are

$$\nu_{\text{E,ext}} = \nu_{\text{E},0} + \nu_{\text{bal}}, \quad \nu_{\text{I,ext}} = \nu_{\text{bal}}/g, \quad (5.59)$$

$$\text{with } \nu_{\text{E},0} = \frac{\mu - \mu_{\text{loc}}}{J_{\text{E}} \tau_{\text{m}}} \quad \text{and} \quad \nu_{\text{bal}} = \frac{\sigma^2 - \sigma_{\text{loc}}^2 - \tau_{\text{m}} \nu_{\text{E},0} J_{\text{E}}^2}{\tau_{\text{m}} J_{\text{E}}^2 (1 + g^2)}, \quad (5.60)$$

where $\mu_{\text{loc}} = \tau_{\text{m}} \nu K J (1 - \gamma g)$ and $\sigma_{\text{loc}}^2 = \tau_{\text{m}} \nu K J^2 (1 + \gamma g^2)$ are the mean and variance due to local input from other neurons firing with the target rate ν , respectively. The rate ν_{bal} establishes a balance between excitation and inhibition, taking into account the variances, and the rate $\nu_{\text{E},0}$ only applies to the excitatory neurons with the aim to shift the mean.

5.3.7 Physical units

The sub-threshold dynamics of the LIF neuron in Equation 5.13 are, without loss of generality, given in scaled units. In this formulation, V , J and I are all quantities with unit Volt. For the parameter-wise comparison with numerical network simulation (for example using NEST, see [Gewaltig and Diesmann, 2007](#)), it is useful to consider a description where I' and J' represent electric currents in units of Ampere:

$$\begin{aligned}\tau_m \frac{dV'_i}{dt} &= - (V'_i - E_L) + R_m I'_i(t) \\ \tau_s \frac{dI'_i}{dt} &= -I'_i + \tau_s \sum_j J'_{ij} s_j(t-d).\end{aligned}\tag{5.61}$$

Here, we also introduce a resistive leak reversal potential E_L , and shift threshold and reset potentials $V'_\theta = V_\theta + E_L$ and $V'_r = V_r + E_L$, respectively. The membrane time constant $\tau_m = R_m C_m$ relates the membrane resistance R_m and capacitance C_m . In units of Ampere, the total current input $I' = I/R_m$ and the synaptic weight amplitude $J' = C_m J/\tau_s$.

5.3.8 Network structure and parameters

We simulate recurrently connected neural networks of one excitatory and one inhibitory populations each using the neural simulation software NEST ([Gewaltig and Diesmann, 2007](#)), using either spiking- or rate-neuron models. The support for rate neurons in NEST was recently added as described in [Hahne et al. \(2017\)](#). Tables 5.1 and 5.2 provide the complete neuron and network model descriptions and Table 5.3 summarizes all parameters as used for the network state showing periodic traveling waves (marked by black star in [Figure 5.1D](#), [Figure 5.5D](#) and [Figure 5.8C](#)). Other simulation parameters used to obtain other network states shown throughout this chapter are indicated with a \otimes marker in Table 5.3, and the changed parameters are given in the corresponding figures. The same marker always denotes the same parameter combination across figure panels. The tables distinguish between network properties and parameters valid for both spiking and rate neuron models and those specific to only one neuron model. Irrespective of the choice of neuron model (rate vs. spiking), the neuron parameters are shared between both neuron populations. The neurons of each population are positioned with equal spacing along a one-dimensional path of perimeter L and connections between neurons are drawn according to a distance-dependent rule with periodic boundary conditions (a “ring” network) using the NEST Topology module.

The number of excitatory neurons N_E in our network is four times larger than the number of inhibitory neurons N_I ([Braitenberg, 2001](#)). The number of incoming connections, the in-degree $K_{\{E,I\}}$, is proportional to the population size of the presynaptic population, assuming an overall connection probability of 10%. Around each postsynaptic neuron, the connection algorithm establishes connections from neighboring neurons within a distance of $R_{\{E,I\}}$. The width of the profile depends on the presynaptic population alone. Potentially presynaptic neurons within this

distance are picked at random and connections are established until the fixed in-degree is reached. Multiple connections between the same pair of neurons termed multapses are allowed, but self-connections (autapses) are prohibited.

The leaky integrate-and-fire model with exponential postsynaptic currents is implemented in NEST under the name `iaf_psc_exp`. The neuron parameters are the same as in the microcircuit model of [Potjans and Diesmann \(2014\)](#) with the difference that our membrane time constant τ_m is half of theirs and that we here omit the refractory period τ_{ref} , although our results generalize to a non-zero τ_{ref} . An excitatory and an inhibitory Poisson generator provide external input to all neurons. Their rates $\nu_{\{E,I\},ext}$ are determined according to [Equation 5.59](#) for fixing the working point (μ, σ) .

The dynamics of rate-based units in NEST is specified as stochastic differential equations using the Itô convention ([Hahne et al., 2017](#)), except that we here set the stochasticity (the variance of the input) to zero. We use the neuron model `tanh_ipn`, that employs a hyperbolic tangent as a gain function.

Simulations run for a simulation time T_{sim} with a temporal resolution of dt . During rate simulations, the instantaneous rate is recorded once at each time step dt . Our raster plots from simulations of the spiking model and the image plots from simulation of the rate model show the network activity from all simulated neurons after a start-up transient T_{trans} .

5.3.9 *Software and implementation*

Spiking- and rate-neuron network simulations were implemented in NEST v2.14.0 ([Peyser et al., 2017](#)), and Python v2.7.11. Post-processing and plotting relied on Python with numpy v1.10.4, SciPy v0.17.0, and matplotlib v2.0.2.

Model summary	
Populations	Excitatory (E), inhibitory (I)
Topology	Ring network: Neurons positioned equally spaced on one-dimensional domain of length L ; periodic boundary conditions
Connectivity	Random convergent connections with fixed in-degree, distance-dependent boxcar-shaped spatial profiles realized with cut-off masks
Spiking model	
Neuron model	Leaky integrate-and-fire (LIF), fixed threshold, absolute refractory time
Synapse model	Static weights and delays, exponentially shaped postsynaptic currents
Input	Independent fixed-rate Poisson spike trains to all neurons (excitatory and inhibitory Poisson sources)
Measurement	Spike activity
Rate model	
Neuron model	Rate neuron with tanh gain function
Synapse model	Delayed rate connection
Input	-
Measurement	Activity

Table 5.1: **Summary of network models following the guidelines of Nordlie et al. (2009b).** Separation between nonlinear spiking and rate neurons as used in NEST simulations.

Network models	
Distance-dependent connectivity	<p>Neural units $j \in X$ at location x_j and $i \in Y$ at x_i in pre- and postsynaptic populations X and Y, respectively. Displacement between units i and j:</p> $r_{ij} = x_i - x_j$ <p>Boxcar-shaped spatial profile with width R and Heaviside function Θ:</p> $p(r_{ij}) = \frac{1}{2R} \Theta(R - r_{ij})$
Spiking model	
Subthreshold dynamics	<p>If $t > t^* + \tau_{\text{ref}}$</p> $\frac{dV}{dt} = -\frac{V - E_L}{\tau_m} + \frac{I_{\text{syn}}(t)}{C_m}$ $I_{\text{syn}}(t) = \sum_j J_j I_{\text{PSC}}(t - t_j^* - d)$ <p>with connection strength J_j, presynaptic spike time t_j^* and conduction delay d $I_{\text{PSC}}(t) = e^{-t/\tau_s} \Theta(t)$ with Heaviside function Θ</p> <p>else</p> $V(t) = V_r$
Spiking	<p>If $V(t^-) < V_\theta \wedge V(t^+) \geq V_\theta$</p> <ol style="list-style-type: none"> 1. set $t^* = t$ 2. emit spike with timestamp t^* 3. reset $V(t) = V_r$
Rate model	
Differential equation	$\tau \frac{du}{dt} = -\delta u(t) + \sum_{j=1} w_j \psi(\delta u_j(t - d))$ $\psi(x) = \tanh(x)$

Table 5.2: **Description of network models.** Separation between nonlinear spiking and rate neurons as used in NEST simulations.

A: Global simulation parameters		
Symbol	Value	Description
T_{sim}	350 ms	Simulation duration
T_{trans}	150 ms	Start-up transient
dt	0.1 ms	Temporal resolution
B: Populations and external input		
Symbol	Value	Description
N_E	4,000	Population size of excitatory neurons
N_I	1,000	Population size of inhibitory neurons
L	1 mm	Domain length
Spiking model		
μ	10 mV	Mean input relative to reset potential
σ	10 mV	Variance of input relative to reset potential
$v_{E,\text{ext}}$	35085 Hz	⊗ Excitatory external rate (by fixing working point)
$v_{I,\text{ext}}$	3683 Hz	⊗ Inhibitory external rate (by fixing working point)
C: Connection parameters		
Symbol	Value	Description
R_E	0.2 mm	⊗ Profile width of excitatory neurons
R_I	0.07 mm	⊗ Profile width of inhibitory neurons
d	3 ms	⊗ Delay
Spiking model		
K_E	400	In-degree from excitatory neurons
γ	0.25	Relative in-degree, $\gamma = K_I/K_E$
J_E	87.8 pA	⊗ Reference synaptic strength
g	5	⊗ Relative synaptic strength, $g = -J_I/J_E$
Rate model		
w_E	2.73	⊗ Excitatory weight (by parameter mapping)
w_I	-3.42	⊗ Inhibitory weight (by parameter mapping)
D: Neuron model		
Symbol	Value	Description
Spiking model		
C_m	250 pF	Membrane capacitance
τ_m	5 ms	Membrane time constant
E_L	-65 mV	Leak potential
V_θ	-50 mV	Firing threshold
V_r	-65 mV	Reset potential
τ_{ref}	0 ms	Absolute refractory period
τ_s	0.5 ms	Postsynaptic current time constant
Rate model		
τ	1.94 ms	Time constant (by parameter mapping)

Table 5.3: **Simulation and network parameters.** Parameters according to setting for traveling waves as shown in Figure 5.1D, Figure 5.5D and Figure 5.8C (black star marker). Deviant parameters are given in the captions of the respective figures and indicated by different markers.

5.4 DISCUSSION

The present study employs mean-field theory (Brunel and Hakim, 1999) to rigorously map a spiking network model of leaky integrate-and-fire (LIF) neurons with constant transmission delay to a neural-field model. We use a conceptually similar linearization as Kriener et al. (2014b) combined with analytical expressions for the transfer function in the presence of colored synaptic noise (Schuecker et al., 2015). The insight that this transfer function in the fluctuation-driven regime resembles the one of a simple first-order low-pass filter facilitates the parameter mapping between the two models. The resulting analytically tractable effective rate model depends on the dynamical working point of the spiking network that is characterized by both the mean and the variance of the synaptic input. By means of bifurcation theory, in particular linear Turing instability analysis (Coombes, 2005; Coombes et al., 2007; Venkov et al., 2007), we investigate the origin of spatiotemporal patterns such as temporal and spatial oscillations and in particular periodic traveling waves emerging in spiking activity. The mechanism underlying these waves encompasses delay-induced fast global oscillations, as described by Brunel and Hakim (1999), with spatial oscillations due to a distance-dependent effective connectivity profile. We derive analytical conditions for pattern formation that are exclusively based on general characteristics of the effective connectivity profile and the delay. The profile is split into a static weight that is either excitatory or inhibitory for a given neural population, and a spatial modulation that can be interpreted as a distance-dependent connection probability. Given the biological constraint that connection probabilities depend on distance but weights do not, periodic traveling waves cannot occur in a single homogeneous population irrespective of the shape of distance-dependent connection probability. Only the effective connectivity profile of two populations (excitatory and inhibitory), permits solutions where a mode with finite non-zero wave number is the most unstable one, a prerequisite for the emergence of nontrivial spatial patterns such as traveling waves. We therefore establish a relation between the anatomically measurable connectivity structure and observable patterns in spiking activity. The predictions of the analytically tractable neural-field model are validated by means of simulations of nonlinear rate-unit networks (Hahne et al., 2017) and of networks composed of LIF-model neurons, both using the same simulation framework (Peyser et al., 2017). In our experience, the ability to switch from a model class with continuous real-valued interaction to a model class with pulse-coupling by changing a few lines in the formal high-level model description increases the efficiency and reliability of the research.

The here presented mathematical correspondence between these a priori distinct classes of models for neural activity has several implications. First, as demonstrated by the application in the current work, it facilitates the transfer of results from the well-studied domain of neural-field models to spiking models. The insight thus allows the community to arrive at a coherent view of network phenomena that appear robustly and independently of the chosen model. Second, the quantitative mapping of the spiking model to an effective rate model in particular reduces the parameters of the former to the set of fewer parameters of the latter; single-neuron and network parameters are reduced to just a weight and a time

constant. This dimensionality reduction of the parameter space conversely implies that entire manifolds of spiking models are equivalent with respect to their bifurcations. Such a reduction supports systematic data integration: Assume a researcher wants to construct a spiking model that reproduces a certain spatiotemporal pattern. The presented expressions permit the scientist to restrict further investigations to the manifold in parameter space in line with these observations. Variations of parameters within this manifold may lead to phenomena beyond the predictions of the initial bifurcation analysis. Additional constraints, such as firing rates, degree of irregularity, or correlations, can then further reduce the set of admissible parameters.

To keep the focus on the transferability of results from a neural-field to a spiking model, the study restricts the analysis to a rather simple network model. In many cases, extensions to more realistic settings are straight forward. As an example, we perform our analysis in one-dimensional space. In two dimensions, the wave number becomes a vector and bifurcations to periodic patterns in time and space can be constructed (see [Ermentrout, 1998](#), Section 8.4 and [Coombes, 2005](#)). Likewise, we restricted ourselves to a constant synaptic delay like [Roxin et al. \(2005; 2006\)](#) because it enables a separation of a spatial component, the shape of the spatial profile, and a temporal component, the delay. A natural next step is the inclusion of an axonal distance-dependent delay term as for instance in [Hutt et al. \(2003\)](#) to study the interplay of both delay contributions ([Veltz, 2013](#)). For simplification, we use here a boxcar-shaped spatial connectivity profile in the demonstrated application of our approach. For the emergence of spatiotemporal patterns, however, the same conditions on the connectivity structure and the delays hold for more realistic exponentially decaying or Gaussian-shaped profiles ([Hellwig, 2000](#); [Perin et al., 2011](#); [Schnepel et al., 2015](#)). If the spatial connectivity profiles are monotonically decaying in the Fourier domain (as it is the case for exponential or Gaussian shapes), the Fourier transform of the effective profile of a network composed of an excitatory and an inhibitory population exhibits at most one zero-crossing. Either the minimum or the maximum are attained at a non-zero and finite wave number k , but not both. With a cosine-shaped effective profile, only a single wave number dominates by construction ([Roxin et al., 2005, 2006](#)). Here, we decided for the boxcar shape because of its oscillating Fourier transform that allows us to study competition between two spatial frequencies corresponding to the two extrema.

Similar to our approach, previous neural-field studies describe the spatial connectivity profile as a symmetric probability density function (see, for example, [Wyller et al., 2007b](#)). For our aim, to establish a link to networks of discrete neurons, the interpretation as a connection probability and the separation from a weight are a crucial addition. This assumption enables us to distinguish between different neural populations, to analyze the shape of the profile based on parameters for the excitatory and the inhibitory contribution, and to introduce biophysically motivated parameters for the synaptic strength. Starting directly with an effective profile that includes both, excitation and inhibition, such as (inverse) Mexican hat connectivity, is mathematically equivalent and a common approach in the neural-fields literature ([Hutt et al., 2003](#); [Atay and Hutt, 2005](#); [Coombes, 2005](#); [Roxin et al., 2005](#)). But it neglects the biological separation of neurons into excitatory and inhibitory populations according to their effect on postsynaptic targets

(Dale's law, see Eccles et al., 1954) and their different spatial reach of connectivity (Stepanyants et al., 2009). A result of this simplification, these models can produce waves even with a single homogeneous population (Roxin et al., 2005; Atay and Hutt, 2006; Venkov et al., 2007), while with homogeneous stationary external drive we show that at least two populations are required.

Local excitation and distant inhibition are often used to support stationary patterns such as bumps, while local inhibition and distant excitation are associated with non-stationary patterns such as traveling waves (Cross and Hohenberg, 1993; Ermentrout, 1998; Hutt et al., 2003). For sufficiently long synaptic delays, we also observe periodic traveling waves with local inhibition and distant excitation, as often observed in cortex (Stepanyants et al., 2009). However, we show that the reason for this is the specific shape of the effective spatial profile, and not only the spatial reach itself. Our argumentation is therefore in line with Hutt et al. (2005; 2008) who demonstrate that wave instabilities can even occur with local excitation and distant inhibition for specific spatial interactions. The spatial connectivity structure and related possible activity states are in addition important factors for computational performance or function of model networks (Legenstein and Maass, 2007; Pyle and Rosenbaum, 2017).

The parameter mapping between a neural-field and a spiking model in this study relies on the insight that the transfer function of the LIF neuron in the fluctuation-driven regime resembles the one of a simple first-order low-pass filter. Since this approximation not only holds for LIF neurons, but also for other spiking neuron models, our results are transferable. A further candidate model with this property is the exponential integrate-and-fire model (Fermani and Richardson, 2015). Other examples include Nordlie et al. (2010) who characterize the firing-rate responses of LIF neurons with strong alpha-shaped synaptic currents and similarly Heiberg et al. (2013) for a LIF neuron model with conductance-based synapses and potassium-mediated after-hyperpolarization currents previously proposed (Casta et al., 2008).

In the literature, the time constant of neural-field models is often associated with the membrane or the synaptic time constant (Bressloff, 2000, 2014; Pyle and Rosenbaum, 2017). Here, we observe that the time constant of the neural-field model derived from the network of spiking neurons falls in between the two. In line with (Gerstner, 2000; Nordlie et al., 2010), we suggest to reconsider the meaning of the time constant in neural-field models.

A limitation of the approach employed here is that the linear theory is only exact at the onset of waves. Beyond the bifurcation, it is possible that nonlinearities in the spiking model govern the dynamics and lead to different prevailing wave numbers or wave frequencies than predicted. Roxin et al. (2006) report that the stability of traveling waves depends crucially on the nonlinearity. Nevertheless they do not observe traveling waves in their spiking-network simulations. In the present work, however, we identify biophysically motivated neuron and network parameters that allow traveling waves to establish in a spiking network. Still, we had to increase the delay beyond the predicted bifurcation point to obtain a stable wave pattern.

Furthermore, the theory underlying the mapping of the spiking network to the neural-field model is based on the diffusion approximation and therefore only applicable for sufficiently small synaptic weights. Widely distributed synaptic

weights, for example, may lead to larger deviations. We here primarily target a wave-generating mechanism for cortical networks. Since in other brain regions involved neuron types, connectivity structures and input characteristics are different, other mechanisms for pattern formation not covered in this work need to be taken into account (Muller and Destexhe, 2012).

The working-point dependence of the neural-field models derived here offers a new interpretation of propagating activity measured in vivo (Takahashi et al., 2015; Denker et al., 2018). Even if the anatomical connectivity remains unchanged during a period of observation, the stability of the neural system can be temporarily altered due to changes in activity. The transfer function of a LIF neuron depends on the mean and the variance of its input, and we have shown that stability is related to its parametrization. In particular, local changes of activity, for example due to a spatially confined external input, can affect stability and hence influence whether a signal remains rather local or travels across the cortical surface. That means, we would relate the tendency of a neural network to exhibit spatiotemporal patterns not only to its connectivity, but also to its activity state that can change over time.

A MESOSCOPIC, MULTI-LAYER, FULL-SCALE CORTICAL NETWORK MODEL

This chapter is based on the following publication:

Johanna Senk, Espen Hagen, Sacha J. van Albada, Markus Diesmann (2018), Reconciliation of weak pairwise spike-train correlations and highly coherent local field potentials across space, *arXiv preprint arXiv:1805.10235v1*

Author contributions:

Under the supervision of Espen Hagen and Markus Diesmann, the author performed all parts of this publication except for the LFP predictions incorporated by Espen Hagen. All authors contributed to the writing of the manuscript.

6.1 INTRODUCTION

Cortical activity on the mesoscopic scale (mesoscale), below a cortical surface area on the order of several square millimeters to centimeters (Muller et al., 2018), can be recorded extracellularly with chronic or acute implants of multi-electrode arrays (Maynard et al., 1997; Buzsáki et al., 2012; Einevoll et al., 2013a). The low-frequency part ($\lesssim 100$ Hz) of the measured extracellular potential, the local field potential (LFP), is a population signal with contributions from up to millions of local and remote neurons (Kajikawa and Schroeder, 2011; Lindén et al., 2011; Łęski et al., 2013). Spiking activity of individual neurons can be obtained from the high-frequency part ($\gtrsim 100$ Hz) of the signal through spike sorting (Quiroga, 2007). The number of reliably identified single neurons is on the order of 100 neurons for chronically implanted Utah arrays (10×10 electrodes on 4×4 mm², Blackrock microsystems¹) as in Riehle et al. (2013). The recordings expose LFP activity appearing to propagate across the cortex associated with distance dependency of statistical measures like correlations and coherences (Destexhe et al., 1999; Smith and Kohn, 2008; Wu et al., 2008; Muller and Destexhe, 2012; Sato et al., 2012; Dubey and Ray, 2016; Denker et al., 2018; Muller et al., 2018). The observation of coherent LFPs across space contrasts with the often reported low pairwise correlation in cortical spike trains obtained in asynchronous brain states (for example Ecker et al., 2010; Renart et al., 2010).

With the assumption of a neuron density of 10^5 neurons/mm² across the cortical surface (Herculano-Houzel, 2009), the number of neurons covered by a Utah array is more than a million. Every neuron receives up to 10^4 synapses from neighboring and distant neurons (Abeles, 1991). However, the local circuitry is highly specific with respect to cortical layers and neuron types (Douglas et al., 1989; Thomson et al., 2002; Binzegger et al., 2004). The majority of local cortical connections are established within a distance of $\lesssim 500$ μ m from the sending/receiving neuron (Voges et al., 2010), with probabilities that decay with distance according to a Gaussian or exponentially shaped profile (Hellwig, 2000; Boucsein et al., 2011; Packer and Yuste, 2011; Perin et al., 2011). Local connections are typically made by unmyelinated axons. Therefore, typical conduction delays between pre- and postsynaptic neurons are governed by propagation speeds estimated around 0.3 mm/ms (Hirsch and Gilbert, 1991; Murakoshi et al., 1993; Kang et al., 1994).

To date, the relationship between cortical connectivity structure and experimentally recorded activity of spikes and LFPs on the mesoscale remains poorly understood. Network models that encompass the relevant anatomical and physiological detail, spatial scales, and corresponding measurements can aid the interpretation of experimental observations and their underlying mechanisms. We here argue for full-scale models, in terms of realistic numbers of neurons and synapses: Down-scaled or diluted network models may not reproduce first- and second-order statistics (rates and correlations, respectively) of full-scale networks (van Albada et al., 2015). Also, Hagen et al. (2016a) demonstrate that biophysical forward-model predictions of LFP signals (and by extension electroencephalographic (EEG) and magnetoencephalographic (MEG) signals) must include the full density of cells and connections to account for network correlations. One such full-density model,

¹ <http://blackrockmicro.com>

the microcircuit model by Potjans and Diesmann (2014), represents a 1 mm^2 cortical patch of early sensory cortex with approximately 80,000 leaky integrate-and-fire neurons and about 0.3 billion synapses set up using neuron-type- and layer-specific connection probabilities derived from anatomical and electrophysiological data. This model produces biologically plausible firing rates across four cortical layers with one excitatory and inhibitory population per layer, is simple enough to allow for rigorous mathematical analysis, is publicly available, and has by now been used also in other studies (Wagatsuma et al., 2011; Bos et al., 2016; Cain et al., 2016; Hagen et al., 2016a; Hahne et al., 2017; Senk et al., 2017b; Schuecker et al., 2017; Schwalger et al., 2017; Schmidt et al., 2018; van Albada et al., 2018).

Here, we hypothesize that a version of this microcircuit model and corresponding LFP measurements upscaled laterally to an area of at least $4 \times 4\text{ mm}^2$ (similar to the Utah multi-electrode array), while accounting for distance-dependent connection probabilities, should not only preserve the main features of activity in the original model, but also explain features emerging on the mesoscale such as spatial propagation of evoked neuronal activity (Bringuier et al., 1999; Swadlow et al., 2002; Einevoll et al., 2007; Muller et al., 2014; Klein et al., 2016), and strong distance-dependent correlations and coherences in the measured LFP (Destexhe et al., 1999; Berens et al., 2008; Katzner et al., 2009; Nauhaus et al., 2009; Kajikawa and Schroeder, 2011; Jia et al., 2011; Srinath and Ray, 2014; Dubey and Ray, 2016) even for typically weak pairwise spike-train correlations in cortex (see, for example, Ecker et al., 2010; Renart et al., 2010). Furthermore, the upscaled model should serve as a test platform for parameters that are to date poorly constrained by available experimental data, and expose mechanisms underlying spatiotemporal pattern formation. Indeed, we find that the overall behavior of the original microcircuit is preserved when upscaled, and that the resulting model reconciles the observation of weak pairwise spike-train correlations in cortex with spatially correlated and coherent LFPs.

6.2 METHODS

6.2.1 Point-neuron networks

This section provides a compact description of the different network models considered in this study. The full network descriptions are given in Tables 6.1 to 6.3. Each network model represents a part of early sensory cortex with realistic densities of neurons and synapses. We first consider the original network model proposed by Potjans and Diesmann (2014) which describes a microcircuit under 1 mm^2 cortical surface, henceforth referred to as ‘reference model’. We then consider networks upscaled to greater surface areas, referred to as ‘upscaled models’. The eight neuron populations within each network are organized into four cortical layers, that is, layer 2/3 (L2/3), layer 4 (L4), layer 5 (L5) and layer 6 (L6), respectively. Each layer contains an excitatory (E) and an inhibitory (I) population of leaky integrate-and-fire (LIF) neurons, whose sub-threshold membrane dynamics are governed by Equation 6.13. The probabilities for two neurons to be connected are layer- and neuron-type-specific and derived from a number of anatomical and electrophysiological studies (Potjans and Diesmann, 2014). Post-

synaptic currents have static, normally distributed amplitudes at onset that decay exponentially (Equations 6.14 and 6.15). All neurons receive stationary external input in the form of Poisson spike trains with fixed rate parameters. In addition, one population of thalamocortical (TC) neurons targeting E and I neurons in both L4 and L6 can provide transient or stationary external input, for example to emulate stimuli of the sensory pathway.

6.2.1.1 Network model descriptions

We here describe the main differences between the original network model and upscaled models derived from it.

Reference model: Potjans and Diesmann (2014) parameterize the original micro-circuit model to cover a cortical column under a surface area of $A^r = 1 \text{ mm}^2$. The superscript r denotes ‘reference model’ here and throughout this chapter. The resulting network connects almost 80,000 neurons with approximately 0.3 billion synapses. The calculation of connection probabilities in the model assumes a Gaussian distance dependency of the form (see Potjans and Diesmann, 2014 for details)

$$c^r(r) = c_0 e^{-r^2/2\sigma_0^2}. \quad (6.1)$$

Here, r denotes the lateral distance between the two neurons. This distance dependency is introduced to reconcile connectivity measurements obtained using anatomical connectivity data (retrograde/anterograde staining, Binzegger et al., 2004) and electrophysiological data (in vitro, Thomson et al., 2002). The computed mean values averaged over all populations for zero-distance connection probability and standard deviation are $c_0 = 0.14$ and $\sigma_0 = 0.30 \text{ mm}$, respectively (Potjans and Diesmann, 2014, Equations 4–8, Figure 3). This spatial decay constant is large compared to the extent of a typical cortical column, which justifies their choice of a local network connectivity without distance dependency.

A neuron j in a source population X of size N_X^r connects at random to a neuron i in a target population Y of size N_Y^r with mean connection probability (Potjans and Diesmann, 2014, Equation 1)

$$C_{YX}^r = 1 - \left(1 - \frac{1}{N_X^r N_Y^r}\right)^{S_{YX}^r}, \quad (6.2)$$

where S_{YX}^r denotes the total number of synapses between these populations. The connection routine draws connections randomly between pairs of neurons i and j until the total number of synapses S_{YX}^r is reached. Multiple connections (multiples) between neuron pairs are allowed. The connection probability C_{YX}^r is here defined as the probability that a pair of neurons is connected via one or more synapses. Connection delays are normally distributed according to Equation 6.17 with different parameters for excitatory and inhibitory sources. The standard deviation of delays is 50% of the mean delay, and the excitatory mean delay is twice as long as the inhibitory one.

Upscaled models: We next consider cortical network models based on the reference network upscaled to cover an area of $A^u = L^2$. With square layers and a chosen side length $L = 4 \text{ mm}$ this area is similar to the area covered by the Utah array (10×10 electrodes, Blackrock Microsystems). The superscript u denotes ‘upscaled models’ here and throughout this chapter. In the upscaled models, neuron

positions are drawn randomly within a square domain of side length L with the origin $(0,0)$ at the center. We position neurons in the TC layer also within the area A^u , which facilitates the connectivity management between TC neurons and cortical neurons in the model. An analogy to the early visual pathway would be that the distance L in both thalamus and cortex corresponds to the same extent of the visual field. A source neuron $j \in X$ at location (x_j, y_j) connects to a target neuron $i \in Y$ at location (x_i, y_i) with a probability dependent on their distance r_{ij} given in Equation 6.11. This expression for distance accounts for periodic boundary conditions (torus connectivity). The distance-dependent connection probability is shaped as a two-dimensional (2D) Gaussian and cut off at a maximal radial distance R as defined in Equation 6.12. The zero-distance connection probability c_{YX} between populations X and Y is derived in Section 6.2.1.2. The corresponding standard deviation σ_X defines the spatial width of the profile and depends only on the source population X . Connection delays of the upscaled models are calculated using a linear distance dependency given by Equation 6.18 with a constant delay offset d_0 and a conduction speed v , plus a random offset drawn from a normal distribution with zero mean and standard deviation σ_d^u capped at values $\pm (d_0 - dt)$ in order to prevent delays smaller than the simulation time step dt . These values are the same for all cortical populations. For the external layer, TC neurons within a circle of adjustable radius $R_{\text{pulse}}^{\text{TC}}$ surrounding the center emit spikes in a synchronous and regular fashion (thalamic pulses) with time intervals Δt_{TC} .

These network model implementations rely on the neuronal network simulator NEST² (Gewaltig and Diesmann, 2007) and are set up such that the same code is used for both the reference and upscaled models, but with different parameters.

6.2.1.2 Upscaling procedure

We here describe the procedure used to derive parameters for the upscaled model(s) from the original reference network model description, in terms of neuron numbers, synapse numbers, distance-dependent connection probabilities, in-degrees of external input, and distance-dependent delays from available experimental data.

Neuron numbers: The upscaled networks preserve the neuron densities per square millimeter of the reference model. Assuming a homogeneous neuron density across space, the size of a population X in the upscaled networks is

$$N_X^u = N_X^r \frac{A^u}{A^r}. \quad (6.3)$$

Synapse numbers: With the aim to derive zero-distance connection probabilities c_{YX} for a Gaussian connectivity profile (Equation 6.12), we first compute average connection probabilities C_{YX}^u in the upscaled models similar to C_{YX}^r for the reference model (as in Equation 6.2, but with corresponding neuron and synapse numbers). We define this connection probability as

$$C_{YX}^u = C_{YX}^{\text{ui}} \cdot (1 - \delta C_{YX}). \quad (6.4)$$

The superscript *ui* denotes upscaled, intermediate connection probabilities. The term δC_{YX} is introduced to allow for selective modifications of the connection probabilities in the final upscaled network (for example to modify firing rate spectra,

² <http://www.nest-simulator.org>

see below). Thus, connections are unchanged for $\delta C_{YX} = 0$, meaning $C_{YX}^u = C_{YX}^i$, but a small positive or negative value results in an increase or decrease of a specific connection probability between populations X and Y , respectively. The connection probability C_{YX}^u depends linearly on the corresponding population-specific connection probability of the reference model, C_{YX}^r , and the ratio of mean connection probabilities from the upscaled and reference models (Schmidt et al., 2018, Equation 6)

$$C_{YX}^u = C_{YX}^r \frac{\bar{C}^u}{\bar{C}^r}. \quad (6.5)$$

Like Schmidt et al. (2018), we choose to use the average connection probability of the reference model $\bar{C}^r = 0.066$ as computed in (Potjans and Diesmann, 2014, Equation 9). To compute the average connection probability \bar{C}^u of the upscaled models, we integrate the Gaussian profile given in Equation 6.1 over all possible positions of a source neuron (x_1, y_1) and a target neuron (x_2, y_2) , located on a square domain of side length L . Accounting for the maximal radial distance of connections, set to $R = L/2$, and the periodic boundary conditions used for the upscaled model, we numerically solve

$$\bar{C}^u = \frac{1}{L^4} \int_{-L/2}^{L/2} \int_{-L/2}^{L/2} \int_{x_1-L/2}^{x_1+L/2} \int_{y_1-L/2}^{y_1+L/2} c^r(r_{21}) dy_2 dx_2 dy_1 dx_1 \quad (6.6)$$

where $r_{21} = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}$ with c^r as defined in Equation 6.1.

The total number of synapses S_{YX}^u follows from Equation 6.2, using connection probabilities and neuron numbers from the upscaled models. This in turn yields the average number of incoming connections to the target neurons, the synaptic in-degree, as $K_{YX}^u = S_{YX}^u / N_Y^u$. Connections in the upscaled model are drawn at random according to the spatial profile (Equation 6.12) and we fix only the zero-distance connection probability c_{YX} and the spatial width σ_X , such that the upscaled in-degree K_{YX}^u is achieved. Under the assumption of a homogeneous distribution of neurons and connections inside a disc with radius R around a target neuron, the local connection probability is then $c_{YX,R} = K_{YX}^u / N_{X,R}^u$, where $N_{X,R}^u$ denotes the number of potential source neurons. We eliminate $N_{X,R}^u$ from the expression for $c_{YX,R}$ by relating neuron numbers to surface areas: $N_{X,R}^u = N_X^u \cdot A_R / A^u$ with $A_R = \pi R^2$ and $A^u = L^2$. To achieve the same in-degree for the uniform connection probability $c(r) = c_{YX,R} \Theta(R - r)$ and the distance-dependent connection probability (Equation 6.12), the following volume integral in polar coordinates must be equal for both choices of $c(r)$: $\int_0^{2\pi} \int_0^\infty \int_0^\infty r dz dr d\varphi$. Due to isotropy, it is enough to equate $\int_0^\infty r c(r) dr$ for both connection probabilities to derive the zero-distance connection probability of the distance-dependent profile,

$$c_{YX} = \frac{K_{YX}^u L^2}{2\pi \sigma_X^2 N_X^u \left[1 - \exp\left(-\frac{R^2}{2\sigma_X^2}\right) \right]}. \quad (6.7)$$

The connection routine used for the upscaled models does not fix the total number of synapses, unlike the routine used for the reference model. Each pair of neurons is considered only once in contrast to the reference model which samples the neurons with replacement. If $c_{YX} > 1$, the routine is executed N_c times with

zero-distance connection probabilities c_{YX}/N_c where $N_c = \lceil c_{YX} \rceil$. In this case, a pair of neurons can be connected by up to N_c synapses.

Mean input: To preserve the mean input to each neuron of the reference network in the upscaled network, we adjust the in-degrees of the external stationary Poisson input to compensate for differences in internal in-degrees between the reference and the upscaled model that result from the above calculation of recurrent synaptic in-degrees. If the mean connection weight $g_{YX} \cdot J$ for internal connections, the weight for external input J , the population firing rates ν_X , and the external Poisson rate ν_{ext} are the same for both models, the external in-degrees $K_{Y,\text{ext}}^u$ per population Y of the upscaled model follow from the external in-degrees of the reference model $K_{Y,\text{ext}}^r$ and the difference in internal in-degrees:

$$\begin{aligned} \sum_X K_{YX}^u g_{YX} \nu_X + K_{Y,\text{ext}}^u \nu_{\text{ext}} &= \sum_X K_{YX}^r g_{YX} \nu_X + K_{Y,\text{ext}}^r \nu_{\text{ext}} \\ K_{Y,\text{ext}}^u &= K_{Y,\text{ext}}^r + \sum_X \frac{g_{YX} \nu_X}{\nu_{\text{ext}}} (K_{YX}^r - K_{YX}^u). \end{aligned} \quad (6.8)$$

This modification of external in-degrees in the upscaled network only preserves the mean of the spiking input (which is proportional to both in-degrees and weights), but not its variance (which is proportional to in-degrees and to weights squared); see, for example, [Brunel and Hakim \(1999\)](#); [van Albada et al. \(2015\)](#) for details.

Delays: To compare the mean delays of the reference model ([Equation 6.17](#)) and mean delays resulting from linear distance dependency in the upscaled model ([Equation 6.18](#)), we compute an effective delay for the upscaled model. The effective delay is computed as the average delay of the distance-dependent version evaluated on a disc of 1 mm^2 (with radius $Q = 1/\sqrt{\pi} \text{ mm}$), thus equalling the extent of the reference model. Accounting for all distances between random points on the disc, the effective delay in polar coordinates for a disc of radius Q is

$$\bar{d}_Q(\sigma_X) = \frac{1}{\pi^2 Q^4} \int_0^Q \int_0^{2\pi} \int_0^Q \int_0^{2\pi} \left(d_0 + \frac{r_{21}}{v} \right) \frac{1}{c_{\text{norm}}} e^{-\frac{r_{21}^2}{2\sigma_X^2}} r_1 r_2 d\varphi_1 dr_1 d\varphi_2 dr_2 \quad (6.9)$$

with $r_{21} = r_1^2 + r_2^2 - 2r_1 r_2 \cos(\varphi_1 - \varphi_2)$. We here account for the Gaussian distance dependency of the spatial profile ([Equation 6.12](#)) with spatial width σ_X but normalize the profile to unity for the integral over the disc by the factor c_{norm} , and ignore the Heaviside function because we only consider $Q < R$. The expression simplifies ([Sheng, 1985, Theorem 2.4](#)) to

$$\begin{aligned} \bar{d}_Q(\sigma_X) &= \frac{\int_0^{2Q} [d_0 + \frac{r}{v}] \exp\left(-\frac{r^2}{2\sigma_X^2}\right) r \left[4 \arctan\left(\sqrt{\frac{2Q-r}{2Q+r}}\right) - \sin\left(4 \arctan\left(\sqrt{\frac{2Q-r}{2Q+r}}\right)\right) \right] dr}{\int_0^{2Q} \exp\left(-\frac{r^2}{2\sigma_X^2}\right) r \left[4 \arctan\left(\sqrt{\frac{2Q-r}{2Q+r}}\right) - \sin\left(4 \arctan\left(\sqrt{\frac{2Q-r}{2Q+r}}\right)\right) \right] dr} \end{aligned} \quad (6.10)$$

which we evaluate numerically. Hence, the delay offset d_0 and conduction speed v can be set based on available experimental data, and the mean delays in the upscaled network can be compared with the corresponding excitatory and inhibitory mean delays of the reference model.

A: Model summary	
Structure	Multi-layer excitatory-inhibitory (E-I) network
Populations	8 cortical in 4 layers (L2/3, L4, L5, L6) and 1 thalamic (TC)
Input	Cortex: Independent fixed-rate Poisson spike trains to all neurons (population-specific in-degree)
Measurements	Spikes, LFP, CSD, MUA
Neuron model	Cortex: leaky integrate-and-fire (LIF); Thalamus: point process
Synapse model	Exponentially shaped postsynaptic currents with normally distributed static weights
Reference model	
Topology	None (no spatial information)
Delay model	Normally distributed delays
Connectivity	Random, independent, population-specific, fixed number of synapses
Upscaled models	
Topology	Random neuron positions on square domain of size $L \times L$; periodic boundary conditions
Delay model	Distributed distance-dependent delays
Connectivity	Random, distance-dependent connection probability, population-specific, number of synapses not fixed in advance
B: Network models	
Connectivity	Connection probabilities C_{YX} from population X to population Y with $\{X, Y\} \in \{L2/3, L4, L5, L6\} \times \{E, I\} \cup TC$, $C_{YX} = 0$ for $Y = TC$
	Reference model
	Fixed number of synapses S_{YX} between populations X and Y (see Equation 6.2), binomially distributed in-/out-degrees
	Upscaled models
	<ul style="list-style-type: none"> • Presynaptic neuron $j \in X$ at location (x_j, y_j) and postsynaptic neuron $i \in Y$ at (x_i, y_i) • Neuron inter-distance (periodic boundary conditions): $r_{ij} = \sqrt{\Delta x_{ij}^2 + \Delta y_{ij}^2} \quad (6.11)$ with $\Delta x_{ij} = x_i - x_j$ if $x_i - x_j \leq L/2$, otherwise $\Delta x_{ij} = L - x_i - x_j$ (same for Δy_{ij}) • Gaussian-shaped connection probability with maximal distance R, spatial width σ_X and zero-distance connection probability c_{YX} (see Equation 6.7): $c^u(r_{ij}) = c_{YX} e^{-r^2/2\sigma_X^2} \Theta(R - r_{ij}) \quad (6.12)$ Heaviside function $\Theta(t) = 1$ for $t \geq 0$, and 0 otherwise.

Table 6.1: Description of reference and upscaled network models following the guidelines of Nordlie et al. (2009a).

C: Neuron models	
Cortex	Leaky integrate-and-fire neuron (LIF) <ul style="list-style-type: none"> • Dynamics of membrane potential $V_i(t)$ for neuron i: <ul style="list-style-type: none"> – Spike emission at times t_s^i with $V_i(t_s^i) \geq V_\theta$ – Subthreshold dynamics: $\tau_m \dot{V}_i = -V_i + R_m I_i(t) \quad \text{if } \forall s : t \notin (t_s^i, t_s^i + \tau_{\text{ref}}] \quad (6.13)$ with $\tau_m = R_m C_m$ – Reset + refractoriness: $V_i(t) = V_{\text{reset}}$ if $\forall s : t \in (t_s^i, t_s^i + \tau_{\text{ref}}]$ • Exact integration with temporal resolution dt (Rotter and Diesmann, 1999) • Random, uniform distribution of membrane potentials at $t = 0$
Thalamus	Spontaneous activity: no thalamic input ($\nu_{\text{TC}} = 0$)
Upscaled models	
Thalamus	Thalamic pulses: coherent activation of all thalamic neurons inside a circle with radius $R_{\text{TC}}^{\text{pulse}}$ centered around $(0,0)$ at fixed time intervals Δt_{TC}

Table 6.2: Description of reference and upscaled network models (continuation of Table 6.1).

D: Synapse models	
Postsynaptic currents	<ul style="list-style-type: none"> • Instantaneous onset, exponentially decaying postsynaptic currents • Input current of neuron i from presynaptic neuron j: $I_i(t) = \sum_j J_{ij} \sum_s e^{-(t-t_s^j - d_{ij})/\tau_s} \Theta(t - t_s^j - d_{ij}) \quad (6.14)$
Weights	<ul style="list-style-type: none"> • Normal distribution with static weights, clipped to preserve sign: $J_{ij} \sim \mathcal{N}\{\mu = g_{YX} \cdot J, \sigma^2 = \sigma_{J,YX}^2\} \quad (6.15)$ <ul style="list-style-type: none"> • Probability density of normal distribution: $f(x \mu, \sigma^2) = \frac{1}{\sqrt{2\pi}\sigma^2} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (6.16)$
Reference model	
Delays	Normal distribution, left-clipped at dt : $d_{ij} = d_{ij}^r \sim \mathcal{N}\{\mu = \bar{d}_X, \sigma^2 = (\sigma_{d,X}^r)^2\} \quad (6.17)$
Upscaled models	
Delays	Linear distance dependency with delay offset d_0 and conduction speed v . Normally distributed additive noise, left-clipped at $-(d_0 - dt)$ and right-clipped at $d_0 - dt$: $d_{ij} = d_{ij}^u \sim d_0 + \frac{r_{ij}}{v} + \mathcal{N}\{\mu = 0, \sigma^2 = (\sigma_d^u)^2\} \quad (6.18)$

Table 6.3: Description of reference and upscaled network models (continuation of Tables 6.1 and 6.2).

A: Global simulation parameters		
Symbol	Value	Description
T_{sim}	5,000 ms	Simulation duration
dt	0.1 ms	Temporal resolution
T_{trans}	500 ms	Startup transient
B: Preprocessing		
Symbol	Value	Description
Δt	0.5 ms	Temporal bin size
Δl	0.1 mm	Spatial bin size
C: Global network parameters		
Connection parameters and external input		
Symbol	Value	Description
J	87.81 pA	Reference synaptic strength. All synapse weights are measured in units of J .
g_{YX}	1 -4 2	Relative synaptic strengths: $X \in \{\text{L2/3E, L4E, L5E, L6E, TC}\}$ $X \in \{\text{L2/3I, L4I, L5I, L6I}\}$, except for: $(X, Y) = (\text{L4E, L2/3E})$
$\sigma_{J,YX}$	$0.1 \cdot g_{YX} \cdot J$	Standard deviation of weight distribution
ν_{ext}	8 s^{-1}	Rate of external input with Poisson inter-spike interval statistics
LIF neuron model		
Symbol	Value	Description
C_m	250 pF	Membrane capacitance
τ_m	10 ms	Membrane time constant
E_L	-65 mV	Resistive leak reversal potential
V_θ	-50 mV	Spike detection threshold
V_{reset}	-65 mV	Spike reset potential
τ_{ref}	2 ms	Absolute refractory period after spikes
τ_s	0.5 ms	Postsynaptic current time constant

Table 6.4: Global simulation, preprocessing, and network parameters used for both reference and upscaled network models.

Additional network parameters for reference model													
Populations and external input													
Symbol	Value	L2/3I	L4E	L4I	L5E	L5I	L6E	L6I	TC	Description			
X	L2/3E	L2/3I	L4E	L4I	L5E	L5I	L6E	L6I	TC	Name			
N_X^I	20,683	5,834	21,915	5,479	4,850	1,065	14,395	2,948	902	Size			
$K_{X,ext}^I$	1,600	1,500	2,100	1,900	2,000	1,900	2,900	2,100	-	External in-degree			
Connection probabilities													
C_{YX}^I	from X												
	L2/3E	L2/3E	L2/3I	L4E	L4I	L5E	L5I	L6E	L6I	TC			
		0.1009	0.1689	0.0437	0.0818	0.0323	0.0	0.0076	0.0	0.0			
	L2/3I	0.1346	0.1371	0.0316	0.0515	0.0755	0.0	0.0042	0.0	0.0			
	L4E	0.0077	0.0059	0.0497	0.1350	0.0067	0.0003	0.0453	0.0	0.0983			
	L4I	0.0691	0.0029	0.0794	0.1597	0.0033	0.0	0.1057	0.0	0.0619			
	L5E	0.1004	0.0622	0.0505	0.0057	0.0831	0.3726	0.0204	0.0	0.0			
	L5I	0.0548	0.0269	0.0257	0.0022	0.0600	0.3158	0.0086	0.0	0.0			
	L6E	0.0156	0.0066	0.0211	0.0166	0.0572	0.0197	0.0396	0.2252	0.0512			
	L6I	0.0364	0.0010	0.0034	0.0005	0.0277	0.0080	0.0658	0.1443	0.0196			
	Connection Parameters												
	Symbol	Value	Description										
	\bar{d}_E	1.5 ms	Mean excitatory delay										
\bar{d}_I	0.75 ms	Mean inhibitory delay											
$\sigma_{d,X}^I$	$0.5 \cdot \bar{d}_X$	Standard deviation of delay distribution											

Table 6.5: Additional network parameters for the reference model.

Additional network parameters for the final upscaled model										
Populations and external input										
Symbol	Value	L2/3I	L4E	L4I	L5E	L5I	L6E	L6I	TC	Description
X										Name
N_{Y}^{u}	330,928	93,344	350,640	87,664	77,600	17,040	230,320	47,168	14,432	Size
$K_{X,ext}^{u}$	1,702	1,621	1,864	2,443	1,939	1,724	3,051	2,246	-	External in-degree
Connection probabilities										
C_{YX}^{u}										from X
		L2/3E	L2/3I	L4E	L4I	L5E	L5I	L6E	L6I	TC
L2/3E	0.007540	0.012622	0.003266	0.006113	0.002414	0.0	0.000568	0.0	0.0	0.0
L2/3I	0.010059	0.010245	0.002361	0.003849	0.005642	0.0	0.000314	0.0	0.0	0.0
L4E	0.000575	0.000441	0.003714	0.008575	0.000501	0.000022	0.003385	0.0	0.007346	0.0
L4I	0.005164	0.000217	0.005934	0.013725	0.000247	0.0	0.007899	0.0	0.004626	0.0
L5E	0.007503	0.004648	0.003774	0.000426	0.006210	0.029237	0.001524	0.0	0.0	0.0
L5I	0.004095	0.002010	0.001921	0.000164	0.003587	0.021240	0.000643	0.0	0.0	0.0
L6E	0.001166	0.000493	0.001577	0.001241	0.004275	0.001472	0.002959	0.016829	0.003826	0.0
L6I	0.002720	0.000075	0.000254	0.000037	0.002070	0.000598	0.004917	0.010784	0.001465	0.0

Table 6.6: Additional network parameters for the final upscaled model.

Connection probability modifications		
Symbol	Value	Description
δC_{YX}	0	$\{X, Y\} \in \{L2/3E, L2/3I, L4E, L4I, L5E, L5I, L6E, L6I, TC\}$, except for:
	-0.15	$(X, Y) = (L4I, L4E)$
	0.15	$(X, Y) = (L4I, L4I)$
	-0.2	$(X, Y) = (L5E, L5I)$
	0.05	$(X, Y) = (L5I, L5E)$
	-0.1	$(X, Y) = (L5I, L5I)$
Connection Parameters		
Symbol	Value	Description
d_0	0.5 ms	Delay offset
v	0.3 mm/ms	Conduction speed
σ_d^u	0.1 ms	Width of jitter distribution for delay
σ_E	0.35 mm	Excitatory spatial width
σ_I	0.1 mm	Inhibitory spatial width
Thalamus		
Symbol	Value	Description
R_{TC}^{pulse}	0.3 mm	TC neuron activation radius of disc around (0,0), all TC neurons in the disc are active during pulses
σ_{TC}	0.3 mm	Spatial width of TC neuron connections
Δt_{TC}	100 ms	Interval between thalamic pulses

Table 6.7: Additional network parameters for the final upscaled model (continuation of Table 6.6).

6.2.2 Forward modeling of extracellular potentials

In the present study we use a now well-established method to compute extracellular potentials from neuronal activity. The method relies on multicompartment neuron modeling to compute transmembrane currents (see, for example, De Schutter and Van Geit, 2009a) and volume conduction theory (Nunez and Srinivasan, 2006; Einevoll et al., 2013b) which relates current sources and electric potentials in space. Assuming a volume conductor model that is linear (frequency-independent), homogeneous (the same in all locations), isotropic (the same in all directions), and ohmic (currents depend linearly on the electric field \mathbf{E}), as represented by the scalar electric conductivity σ_e , the electric potential in location $\mathbf{r} \equiv (x, y, z)$ of a time-varying point current with magnitude $I(t)$ in location \mathbf{r}' is given by

$$\phi(\mathbf{r}, t) = \frac{I(t)}{4\pi\sigma_e|\mathbf{r} - \mathbf{r}'|}. \quad (6.19)$$

The potential is assumed to be measured relative to an ideal reference at infinite distance from the source. Consider a set of transmembrane currents of n_{comp} individual cylindrical compartments indexed by n in an N -sized population of cells indexed by j with time-varying magnitude $I_{jn}^m(t)$ embedded in a volume conductor representing the surrounding neural tissue. The extracellular electric potential is then calculated as the linear sum

$$\phi(\mathbf{r}, t) = \sum_{j=1}^N \sum_{n=1}^{n_{\text{comp}}} \frac{I_{jn}^m(t)}{4\pi\sigma_e} \int \frac{1}{|\mathbf{r} - \mathbf{r}_{jn}|} d\mathbf{r}_{jn}. \quad (6.20)$$

The integral term here enters as we utilize the *line-source* approximation (Holt and Koch, 1999) which amounts to assuming a homogeneous transmembrane current density per unit length and integrating Equation 6.19 along the center axis of each cylindrical compartment. The thick soma compartments (with $n = 1$) with magnitude $I_j^{\text{m,soma}}(t)$, however, are approximated as spherical current sources, which amounts to combining Equations 6.19 and 6.20 as Lindén et al. (2014)

$$\begin{aligned} \phi(\mathbf{r}, t) &= \sum_{j=1}^N \frac{1}{4\pi\sigma_e} \left(\frac{I_j^{\text{m,soma}}(t)}{|\mathbf{r} - \mathbf{r}_j^{\text{soma}}|} + \sum_{n=2}^{n_{\text{comp}}} \int \frac{I_{jn}^{\text{m}}(t)}{|\mathbf{r} - \mathbf{r}_{jn}|} d\mathbf{r}_{jn} \right) \\ &= \sum_{j=1}^N \frac{1}{4\pi\sigma_e} \left(\frac{I_j^{\text{m,soma}}(t)}{|\mathbf{r} - \mathbf{r}_j^{\text{soma}}|} + \sum_{n=2}^{n_{\text{comp}}} \frac{I_{jn}^{\text{m}}(t)}{\Delta s_{jn}} \ln \left| \frac{\sqrt{h_{jn}^2 + r_{\perp jn}^2} - h_{jn}}{\sqrt{l_{jn}^2 + r_{\perp jn}^2} - l_{jn}} \right| \right). \end{aligned} \quad (6.21)$$

Here, lengths of compartments n of cells j are denoted by Δs_{jn} , perpendicular distances from the electrode point contact to the axis of the line compartments by $r_{\perp jn}$, and longitudinal distances measured from the start of the compartment by h_{jn} . The distances $l_{jn} = \Delta s_{jn} + h_{jn}$ are measured longitudinally from the end of the compartment. As the above denominators can be arbitrarily small and cause singularities in the computed extracellular potential, we set the minimum separation $|\mathbf{r} - \mathbf{r}_j^{\text{soma}}|$ or $r_{\perp jn}$ equal to the radius of the corresponding compartment.

The above equations assume point electrode contacts, while real electrode contacts have finite extents. We employ the *disc-electrode* approximation (Camuñas Mesa and Quiroga, 2013; Lindén et al., 2014; Ness et al., 2015)

$$\phi_{\text{disc}}(\mathbf{u}, t) = \frac{1}{A_S} \iint_S \phi(\mathbf{u}, t) d^2r \approx \frac{1}{m} \sum_{h=1}^m \phi(\mathbf{u}_h, t) \quad (6.22)$$

to approximate the averaged potential across the uninsulated contact surface (Robinson, 1968; Nelson et al., 2008; Nelson and Pouget, 2010; Ness et al., 2015). We average the potential (Equation 6.21) in $m = 50$ randomized locations \mathbf{u}_h on each circular and flat contact surface S with surface area A_S and radius $5 \mu\text{m}$. The surface normal vector on the disc representing each contact is the unit vector along the vertical z -axis. All forward-model calculations are performed with the simulation tool LFPy³ (Lindén et al., 2014; Hagen et al., 2018), which uses the NEURON⁴ simulation software (Carnevale and Hines, 2006; Hines et al., 2009) to calculate transmembrane currents.

6.2.2.1 Modifications to the hybrid scheme

Extracellular potentials from the point-neuron network models are here calculated using a slightly modified version of the biophysics-based hybrid scheme introduced by Hagen et al. (2016a). The scheme combines forward modeling of extracellular potentials, or more specifically its low-frequency part termed the local field potential (LFP), from spatially extended multicompartment neuron models described above instead of point neurons. Point neurons cannot generate an extracellular potential, as the sum of all in- and outgoing currents vanishes in a point,

³ <https://lfp.readthedocs.io>

⁴ <https://neuron.yale.edu>

in contrast to multicompartment neuron models, which can account for in- and outgoing currents distributed in space. We refer the reader to the Methods of Hagen et al. (2016a) for an in-depth technical description of the implementation for randomly connected point-neuron network models. Here, we only summarize its main steps and list the main changes which allow accounting for extracellular potentials of networks with distance-dependent connectivity and periodic boundary conditions. This hybrid modeling scheme for extracellular potentials combines the simplicity and efficiency of point-neuron network models with multicompartment neuron models for LFP generation accounting for the biophysical origin of extracellular potentials. As in Hagen et al. (2016a), we assume that cortical network dynamics are well captured by the point-neuron network, and implement the hybrid scheme as follows:

- Spike trains of individual point neurons are mapped to synapse activation times on corresponding postsynaptic multicompartment neurons while overall connection parameters are preserved, that is, the distribution of delays, the mean postsynaptic currents, and the mean number of incoming connections onto individual cells (in-degree).
- Each multicompartment neuron has its equivalent in the point-neuron network and receives input spikes from presynaptic point neurons with the same distribution as in the point-neuron network (the mean in-degree of neurons in the network and the cell-type and layer specificity of connections is preserved, as in Hagen et al. (2016a)).
- The multicompartment neurons are mutually unconnected, and synaptic activations are translated into a distribution of transmembrane currents that contributes to the total LFP.
- Activity in multicompartment neuron models (and the corresponding LFP) does not interact with other multicompartment neurons or the activity in the point-neuron network model, that is, there are no ephaptic interactions.

The first version of the hybrid scheme implemented in hybridLFPy⁵ is developed for random networks such as the layered cortical microcircuit model of Potjans and Diesmann (2014) that is our reference network. In contrast to this reference model that contains no spatial information, the upscaled models described in Section 6.2.1 assign spatial coordinates to the neurons within each layer but ignore information about cortical depth, and draw connections between neurons with probabilities depending on lateral distance. Modifications to the hybrid scheme to account for upscaled networks thus include:

- We use the lateral locations of the point neurons also for the multicompartment neuron models, and assign population-dependent somatic depths as in Hagen et al. (2016a).
- We record the spiking activity from all neurons in the point-neuron network and associate each spike train to the corresponding neuron ID.

⁵ <https://INM-6.github.com/hybridLFPy>

- Presynaptic neuron IDs are drawn for each multicompartment neuron using the same distance-dependent probability rule as is used when constructing the point-neuron network (the connectivity is thus statistically reproduced). The same distance-dependent delay rule is also implemented in the hybrid scheme, and can be set separately for each pair of populations.
- We compute the extracellular potential at 100 contact sites arranged on a square regular grid with each contact separated by $400\ \mu\text{m}$, similar to the layout of the Blackrock ‘Utah’ multi-electrode array. The local field potential is computed at the center of layer 2/3 (L2/3).
- LFPy, which implements the above forward model and is used internally in the hybrid scheme, accounts for periodic boundary conditions.

6.2.2.2 Modifications to LFPy to account for periodic boundary conditions

As the upscaling procedure of the 1mm^2 reference point-neuron network model incorporates periodic boundary conditions, we modify the forward-model calculations in LFPy⁶ (Lindén et al., 2014; Hagen et al., 2018) to also account for such boundaries. The basic premise for this modification is that transmembrane currents of a neuron positioned near the network layer boundary should result in a fluctuation of the extracellular potential also due to sources across the boundary. This is analogous to input from network connections across the boundaries resulting from the distance-dependent connectivity rule. Thus, for a current source located in location $\mathbf{r}_{jn} = (x_{jn}, y_{jn}, z_{jn})$ the extracellular potential in location \mathbf{r} is computed as the sum

$$\phi(\mathbf{r}, t) = \sum_{p=-M}^M \sum_{q=-M}^M \phi_{pq}(\mathbf{r}, t), \quad (6.23)$$

where $\phi_{pq}(\mathbf{r}, t)$ corresponds to the extracellular potential with horizontally shifted source coordinates $(x_{jn} + pL, y_{jn} + qL, z_{jn})$, L the network layer side length and $M = 2$ a chosen integer setting the number of ‘mirror’ sources to either side.

6.2.3 Statistical analysis

As simulation output, we consider the spiking activity of the point-neuron networks (Section 6.2.1.1), and corresponding multi-unit activity (MUA), LFP (Section 6.2.2) and current-source density (CSD) estimates. We use simulated output data only after an initial time period of T_{trans} to avoid startup transients, and compute all measures for the whole time interval of the following simulation duration T_{sim} . Parameters are given in Tables 6.4 to 6.7.

6.2.3.1 Temporal binning of spike trains

Spike times t_i^s of the point-neuron networks simulated using temporal resolution dt are assigned to bins with width Δt . Temporally binned spike trains are used to compute pairwise spike-train correlations and population-rate power spectral densities, and to illustrate population-averaged rate histograms. The bin width Δt

⁶ <https://LFPy.readthedocs.io>

is an integer multiple of the simulation resolution dt . The simulation duration T_{sim} is an integer multiple of the bin width such that the number of bins is $K = T/\Delta t$. Time bins have indices $k \in \{0, 1, \dots, K-1\}$, spanning time points in $t \in [k\Delta t, (k+1)\Delta t)$.

6.2.3.2 Spatiotemporal binning of spike trains

In order to compute the propagation speed of evoked activity in the network, we perform a spatiotemporal binning operation of spiking activity in the network. As introduced in Section 6.2.1.1, neuron positions (x_i, y_i) of the point-neuron network are randomly drawn with $\{x_i, y_i\} \in [-L/2, L/2)$. We subdivide the spatial domain of each layer into square bins of side length Δl such that the integer numbers of bins along the x - and y -axis are $L_{\{x,y\}} = L/\Delta l$. The bin indices are $l_{\{x,y\}} \in \{0, 1, \dots, L_{\{x,y\}} - 1\}$, spanning $\{x, y\} \in [l_{\{x,y\}}\Delta l - L/2, (l_{\{x,y\}} + 1)\Delta l - L/2)$. Temporal bins of width Δt are defined as above. We compute for each population a spatially and temporally binned instantaneous spike-count rate in units of s^{-1} as the number of spike events from all neurons inside the spatial bin divided by Δt .

6.2.3.3 Current-source density (CSD) analysis

We estimate the current-source density (CSD) using the kernel CSD (kCSD) method introduced by Potworowski et al. (2012). The CSD is an estimate of the volume density of transmembrane currents nearby each LFP measurement site (in units of current per volume). Based on the Poisson equation in electrostatics,

$$\nabla(\sigma_e \nabla)\phi = -C, \quad (6.24)$$

which relates the electric potential $\phi \equiv \phi(\mathbf{r})$, conductivity $\sigma_e \equiv \sigma_e(\mathbf{r})$ (which is here assumed to be scalar as above), and current density $C \equiv C(\mathbf{r})$, one can make the assumption that the measured LFP at each electrode results from a sum of M current sources distributed across space. Similar to Łeński et al. (2011) and Potworowski et al. (2012), we consider the underlying CSD as a product

$$\tilde{f}(x, y, z) = \tilde{f}(x, y)H(z), \quad (6.25)$$

where the term $\tilde{f}(x, y)$ describes a spatial profile in the horizontal xy -plane and $H(z)$ the step function along the vertical z -axis,

$$H(z) = \begin{cases} 1 & -h \leq z \leq h, \\ 0 & \text{otherwise.} \end{cases} \quad (6.26)$$

The variable h denotes the half-thickness of the current-generating region. Under the assumption of a linear (frequency-independent) and homogeneous (equal in all locations) conductivity, it follows that the electric potential in a location $(x, y, 0)$ is

$$f(x, y, 0) = \frac{1}{2\pi\sigma_e} \int \text{arcsinh}\left(\frac{2h}{\sqrt{(x-x')^2 + (y-y')^2}}\right) \tilde{f}(x', y') \, dy' dx'. \quad (6.27)$$

We here choose to define $\tilde{f}(x, y)$ in terms of 2D Gaussians of the form

$$\tilde{b}_i(x, y) = \exp\left(-\frac{(x - x_i)^2 + (y - y_i)^2}{2\sigma_R^2}\right), \quad (6.28)$$

resulting in

$$b_i(x, y) = \frac{1}{2\pi\sigma_e} \int \operatorname{arcsinh}\left(\frac{2h}{\sqrt{(x - x')^2 + (y - y')^2}}\right) \tilde{b}_i(x, y) \, dy' dx'. \quad (6.29)$$

Introducing

$$\phi(x, y) = \mathcal{A}C(x, y) = \sum_{j=1}^M a_j b_j(x, y), \quad (6.30)$$

where $\mathcal{A} : \tilde{\mathcal{F}} \rightarrow \mathcal{F}$ is a linear operator connecting electric potentials and the underlying sources, the CSD is estimated as

$$C^*(x, y) = \tilde{\mathbf{K}}^T(x, y) \cdot \mathbf{K}^{-1} \cdot \mathbf{V}, \quad (6.31)$$

which minimizes the norm $\|\phi\|^2 = \sum_{i=1}^M |a_i|^2$. Here $\mathbf{V} = [\phi_1, \phi_2, \dots, \phi_N]^T$ is the observed LFP across channels,

$\tilde{\mathbf{K}}^T(x, y) = [\tilde{K}_1(x_1, y_1, x, y), \tilde{K}_2(x_2, y_2, x, y), \dots, \tilde{K}_N(x_N, y_N, x, y)]$ and

$$\mathbf{K} = \begin{bmatrix} K(x_1, y_1, x_1, y_1) & \cdots & K(x_1, y_1, x_N, y_N) \\ \vdots & \ddots & \vdots \\ K(x_N, y_N, x_1, y_1) & \cdots & K(x_N, y_N, x_N, y_N) \end{bmatrix}, \quad (6.32)$$

defined in terms of the kernel functions $K(x, y, x', y') = \sum_{i=1}^M b_i(x, y) b_i(x', y')$ and cross-kernel functions $\tilde{K}(x, y, x', y') = \sum_{i=1}^M b_i(x, y) \tilde{b}_i(x', y')$. See [Potworowski et al. \(2012\)](#) for details on the procedure. We use the implementation of the 2D kCSD method available in [Elephant⁷](#) (Electrophysiology Analysis Toolkit) with default parameters $\sigma_e = 0.3\text{S/m}$, $M = 1000$, $h = 1\text{mm}$, $\sigma_R^2 = 0.23\text{mm}^2$, and return the estimate at the space spanned by the LFP electrodes with resolution 0.4 mm.

6.2.3.4 Calculation of MUA signal

For each electrode contact point located in L2/3, we compute a signal representative of the so-called multi-unit activity (MUA) signal that can be obtained from recordings of extracellular potentials by high-pass filtering the signal ($\gtrsim 500\text{Hz}$), followed by signal rectification, temporal smoothing, and downsampling (see, for example, [Einevoll et al., 2007](#)). In a biophysical modeling study ([Pettersen et al., 2008](#)) it is shown that this signal is approximately linearly related to the firing rate of the local population of neurons in the vicinity of the measurement device. Neuron coordinates (x_i, y_i) of the upscaled point-neuron network are randomly drawn on the interval $\{x_i, y_i\} \in [-L/2, L/2]$. We subdivide the layers into square bins of side length $\Delta l_{\text{MUA}} = 0.4\text{mm}$ resulting in 10 bins along the x - and y -axis, respectively. Each electrode contact point is located at the center of the respective bin. We also define temporal bins of width Δt . We then compute for each population a spatially and temporally binned spike-count rate in units of s^{-1} by summing the number of spike events from all neurons inside the spatial bin and divide by the width of the temporal bin Δt . We then define the MUA signal as the sum of the per-bin contributions of the populations L2/3E and L2/3I.

⁷ <https://github.com/neuralensemble/elephant>

6.2.3.5 Visual analysis

The *spike raster* diagrams or dot displays show information on spiking activity. Each dot marks a spike event, and the dot position along the horizontal axis denotes the time of the event. Spike data of different neuron populations are stacked and the number of neurons shown is proportional to the population size. Within each population, neurons are sorted according to their lateral x -position and arranged accordingly on the vertical axis of the dot display.

We compute *population-averaged rate histograms* by deriving the per-neuron spike rates in time bins Δt and units of s^{-1} , averaged over all neurons per population within the center disc of 1 mm^2 . The corresponding histogram shows the rates in a time interval of $\pm 25\text{ ms}$ around the occurrence of a thalamic pulse. Such a display is comparable to the Peri-Stimulus Time Histogram (PSTH, [Perkel et al., 1967a](#)) that typically shows the spike count summed over different neurons or trials versus binned time.

Image plots with color bars can have a linear or a logarithmic scaling as specified in the respective captions. Since values of the distance-dependent cross-correlation functions can be positive or negative, we plot these with linear scaling up to a threshold, beyond which the scaling is logarithmic.

6.2.3.6 Statistical measures

Per-neuron spike rates v are defined as the number of spikes per neuron during each simulation divided by the simulation duration T_{sim} . Distributions of per-neuron spike rates are computed from all spike trains of each population separately for an interval from 0 to 30 s^{-1} using bins of width 1 s^{-1} . Histograms are normalized such that the cumulative sum over the histogram equals unity. We define the mean rate per population \bar{v} as the arithmetic mean of all per-neuron spike rates of each population.

The *coefficient of local variation* LV is a measure of spike-train irregularity computed from a sequence of length n of consecutive inter-spike intervals T_i ([Shinomoto et al., 2003](#), Equation 2.2), defined as

$$LV = \frac{1}{n-1} \sum_{i=1}^{n-1} \frac{3(T_i - T_{i+1})^2}{T_i + T_{i+1}}. \quad (6.33)$$

Like the conventional coefficient of variation CV ([Shinomoto et al., 2003](#), Equation 2.1), a sequence of intervals generated by a stationary Poisson process results in a value of unity, but the LV statistic is less affected by rate fluctuations compared to the CV ; thus, a non-stationary Poisson process should result in $LV \approx 1$. We compute the LV from the inter-spike intervals of the spike trains of all neurons within each population. Distributions of LV s are computed using bins of width 0.1, and histograms are normalized such that the cumulative sum over the histogram equals unity. We define the mean LV per population \overline{LV} as the arithmetic mean of all LV s of each population.

The *Pearson (product-moment) correlation coefficient* CC is a measure of synchrony that is defined for two signals u and v as

$$CC_{uv} = \frac{\text{cov}(u, v)}{\sqrt{\text{cov}(u, u) \text{cov}(v, v)}}, \quad (6.34)$$

with the covariance denoted by cov . The calculation is implemented using the function `numpy.corrcoef`. To compute distributions of correlation coefficients from spike trains, we randomly select 1000 neurons per population and assign their spike times to temporal bins with width $\Delta t_{CC} = 5$ ms (see Section 6.2.3.1). Then, we compute pairwise CCs for the spike counts $u = n_i$ and $v = n_j$ of selected neurons i from a population X and neurons j from a population Y (ignoring autocorrelations). Within each population, meaning $X = Y$, the CC is denoted by $E - E$ for an excitatory population or $I - I$ for an inhibitory population. Correlations between neurons from the excitatory and the inhibitory population in each layer are denoted by $E - I$. CC histograms have bins of width 0.003, are restricted to a range with a minimum and maximum CC of ± 0.08 , respectively, and are normalized such that the integral over the histogram equals unity. We also compute correlation coefficients for assessing the distance dependency of spikes, LFP, CSD, and MUA signals. In these cases, u and v are LFP, CSD, or MUA time series in different spatial locations. For spikes, we sample 40 excitatory and 10 inhibitory spike trains, bin them as above, compute their correlation coefficients (ignoring autocorrelations), and plot them according to distance between the pairs of neurons.

Coherences are computed as

$$\gamma_{uv}(f) = \frac{|\mathcal{S}_{uv}(f)|}{\sqrt{\mathcal{S}_{uu}(f)\mathcal{S}_{vv}(f)}}, \quad (6.35)$$

where $\mathcal{S}_{uv}(f)$ is the cross-spectral density between u and v , and $\mathcal{S}_{uu}(f)$ and $\mathcal{S}_{vv}(f)$ are the power spectral densities (PSDs) of each signal. The *cross-spectral density* and *power spectra* are computed using Welch's average periodogram method (Welch, 1967) as implemented by `matplotlib.mlab`'s `csd` and `psd` functions, respectively, with number of data points used in each block for the fast Fourier transform (FFT), that is, segment length $N_{\text{FFT}} = 256$, overlap between segments $N_{\text{overlap}} = 192$ and signal sampling frequency $F_s = 2$ kHz. To compute the *population-rate power spectral density*, we use the spike trains of all neurons per population (in Figure 6.3N and H only within the center disc of 1 mm^2), resampled into bins of size Δt , and with the arithmetic mean of the binned spike trains subtracted.

The effect of thalamic pulses is analyzed by means of *distance-dependent cross-correlation functions* $CC^v(\tau, r)$ evaluated for time lags τ . We discretize the network of size $L \times L$ into an even number of square bins of side length Δl . The spike trains from all neurons within each spatial bin are resampled into time bins of size Δt and averaged across neurons to obtain spatially and temporally resolved per-neuron spike rates. We select spatial bins on the diagonals of the network such that each distance to the center with coordinates $(0, 0)$ is represented by four bins. For 14 distances from consecutive spatial bins along the diagonal, we compute the temporal correlation function between the rates in the respective spatial bins with a binary vector containing ones at spike times of the thalamic pulses and zeros elsewhere, and then average over the four spatial bins at equal distance. The sequences are normalized by subtracting their mean and dividing by their standard deviation. Correlations between the sequences u and v with time steps k and the length of the sequences K are then computed as

$$CC_{u,v}(\tau) = \frac{1}{K} \sum_{k=1}^K u_{k+\tau} v_k \quad (6.36)$$

for $\tau \in [-25, 25]$ ms in steps of Δt . Finally, we subtract the baseline correlation value, obtained by averaging over all negative time lags (before thalamic activation at $\tau = 0$), and get $CC^v(\tau, r)$.

To estimate the *propagation speed* v_{prop} from the cross-correlation functions, we find for each distance the time lag corresponding to the largest CC^v . Values of CC^v smaller than 10 % of the maximum of all CC^v per population across distances and time lags are excluded. We further exclude distances smaller than the thalamic radius $R_{\text{TC}}^{\text{pulse}}$ plus the spatial width of thalamic connections σ_{TC} because a large part of neurons within this radius are simultaneously receiving spikes directly from thalamus upon thalamic pulses. A linear fit for the distance as function of time lag, $r_p(\tau) = r_{p,0} + v_p \cdot \tau$, yields the speed v_p and its fitting error, the standard deviation $\sigma_{v,p}$. We compute the speed for different populations p and obtain the propagation speed as weighted mean with its uncertainty:

$$v_{\text{prop}} = \frac{\sum_p v_p / \sigma_{v,p}^2}{\sum_p 1 / \sigma_{v,p}^2}, \quad \sigma_{v,\text{prop}} = \sqrt{\frac{1}{\sum_p 1 / \sigma_{v,p}^2}}. \quad (6.37)$$

6.2.3.7 Curve fitting

For certain measures, such as pairwise correlation coefficients computed for different distances between LFP electrode locations, we fit exponential functions of the form

$$y(r) = a \cdot e^{-r/b} + c, \quad (6.38)$$

where $\beta = (a, b, c)$ are constant parameters that minimize the sum $\sum_{i=1}^m |y_i(r_i) - y(r_i, \beta)|^2$ for the m data points y_i computed for distance r_i . The parameter fitting is implemented using the non-linear least squares function `curve_fit` provided by the `scipy.optimize` module, with initial guess $\beta = (0.1, 0.1, 0.1)$. Goodness of fit is quantified by the coefficient of determination, defined as

$$R^2 = 1 - \frac{\sum_{i=1}^m (y_i(r_i) - y(r_i, \beta))^2}{\sum_{i=1}^m (y_i(r_i) - \bar{y})^2}, \quad (6.39)$$

where \bar{y} is the mean of the observed data.

6.2.4 Software accessibility

We here summarize the details of software and hardware used to generate the results presented throughout this study. Point-neuron network simulations are implemented using the SLI interface of NEST v2.12.0 (Kunkel et al., 2017), and Python v2.7.11. We use the same network implementation for reference and all upscaled models and switch between them by adjusting parameters. Parameter scans rely on the parameters module of NeuroTools⁸. LFP signals are computed

⁸ <http://neuralensemble.org/NeuroTools>

using NEURON v7.5 and LFPy⁹ (branch ‘som_as_point_periodic’ at SHA:4cab667), hybridLFPy¹⁰ (branch ‘LFPy_dev’ at SHA:of1bf2). Analysis and plotting rely on Python with numpy v1.10.4, SciPy v0.17.0, and matplotlib v2.1.2. All simulations and analyses are conducted on the JURECA¹¹ supercomputer based on Intel Xeon E5-2680 v3 Haswell CPUs running the CentOS 7 Linux distribution. Simulations are run using 1152 and 2304 physical cores for the network and LFP simulations, respectively.

6.3 RESULTS

6.3.1 Upscaling of a cortical microcircuit model using lateral distance-dependent connectivity

Starting with a model of the cortical microcircuit (the reference model, see [Potjans and Diesmann, 2014](#)), we construct full-scale multi-layer neuronal network models with distance-dependent connectivity via the upscaling procedure described in [Section 6.2.1.2](#). The full network descriptions are provided in [Section 6.2.1.1](#) and in [Tables 6.1 to 6.3](#). Here, we point out similarities and differences between the reference model and an upscaled model with parameters set to the values given in [Tables 6.4 to 6.7](#). We refer to this parameterization as the ‘base parameters’. [Figure 6.1A](#) illustrates the reference model next to the laterally upscaled version. The reference model comprises almost 80,000 neurons under 1 mm² of cortical surface area, while the upscaled model consists of approximately 1.2 million neurons and covers an area of 4 × 4 mm², similar to the the area covered by the Utah multi-electrode array. To illustrate their connectivities, the figure shows in both network sketches incoming connections from population L5E to an example target neuron in population L4E. In the reference model without spatial structure, source neurons are picked randomly from the source population. In the upscaled model, source neurons are picked around the target neuron in layer 4 according to distance-dependent probabilities with Gaussian profiles of outgoing connections from layer 5 excitatory neurons. The width of the profile is 0.3 mm which is the average value σ_0 from the connectivity data underlying the reference model, see [Equation 6.1](#). A major fraction of source neurons falls into the center 1 mm², justifying the assumption of random connectivity in the reference model.

The population-specific connection probabilities in the reference model C_{YX}^r , shown in [Figure 6.1B](#), are equal to those in [Potjans and Diesmann \(2014, Table 5\)](#). The upscaling procedure yields connection probabilities C_{YX}^u that are decreased by approximately one order of magnitude in comparison to the reference model. The derived in-degrees K_{YX}^u , however, are slightly larger than K_{YX}^r for all population pairs. This is expected since the upscaling procedure adds connections at distances not accounted for within the limited extent of the reference model.

For the final upscaled model, we increase the excitatory and decrease the inhibitory spatial widths of the connection probability profiles ([Equation 6.12](#)) compared to the average value σ_0 of the reference model to $\sigma_E = 0.35$ mm and $\sigma_I =$

⁹ <http://lfp.py.github.io/>

¹⁰ <https://github.com/INM-6/hybridLFPy>

¹¹ http://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/JURECA/JURECA_node.html

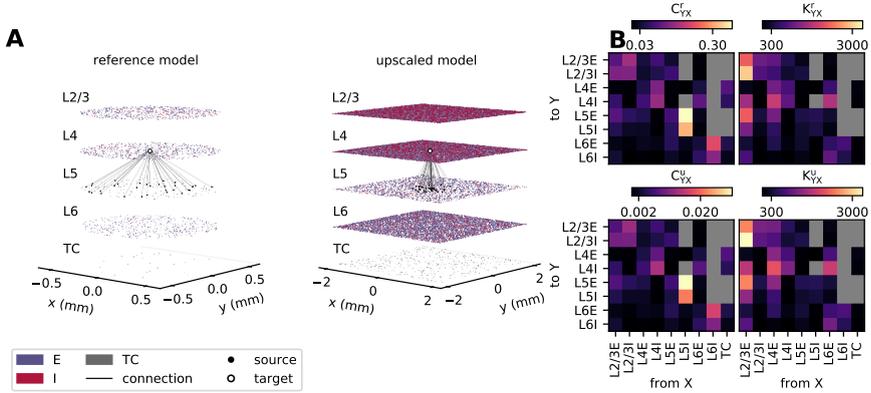


Figure 6.1: Layered cortical point-neuron network models. **A** Illustrations of the network geometry of the reference model (left, 1 mm^2 cortical microcircuit, introduced by Potjans and Diesmann, 2014) and an upscaled model (right, $4 \times 4 \text{ mm}^2$ cortical layers). Both models consist of four cortical layers (L2/3, L4, L5, L6) with an excitatory (E) and an inhibitory (I) population each, and an external thalamic population (TC). Colored dots represent individual neurons at their (x, y) -coordinates; excitatory neurons in blue, inhibitory neurons in red, and thalamic neurons in gray. The number of neurons shown per population is reduced by a factor 32 compared to the actual neuron number in each network to not saturate the illustrated layers. Black lines illustrate convergent connections from sources in L5E (black dots) to a target neuron in L4E (white dot). In-degrees correspond to the actual average in-degrees in both models rounded to the nearest integer: 33 in the reference model and 39 in the upscaled model. Sources are drawn at random in the reference model, but with lateral distance dependency (Gaussian-shaped profile) in the upscaled model. **B** Network connectivity of the reference model (top panels) and the upscaled model (bottom panels). Upscaled connection probabilities are computed as in Equation 6.4. Left panels show color-coded connection probabilities C_{YX}^r and C_{YX}^u (different color code) with the values given in Tables 6.5 and 6.6, and right panels show derived in-degrees K_{YX}^r and K_{YX}^u (same color code). Color maps have linear scaling with zero-values masked in gray.

0.1 mm, respectively. Accumulating experimental data indicate Gaussian or exponentially decaying connection probabilities with distance for both excitatory and inhibitory local connections; see, for example, the review by Boucsein et al. (2011), or Hellwig (2000) for pyramidal cells in layers 2 and 3 of rat visual cortex, Budd and Kisvárdy (2001) for clutch cells in layer 4 of cat visual cortex, Perin et al. (2011) for pyramidal cells in layer 5 of rat somatosensory cortex, Levy and Reyes (2012) for pyramidal cells and (non-)fast-spiking inhibitory cells in deep layer 2/3 and layer 4 of mouse auditory cortex, Schnepel et al. (2015) for excitatory input to pyramidal neurons in layer 5B of rat somatosensory cortex, Jiang et al. (2015) for pyramidal cells and different interneurons in layers 1, 2/3, and 5 of mouse visual cortex, Packer and Yuste (2011) for parvalbumin-positive cells connected to pyramidal cells in multiple layers of mouse neocortex, and Reimann et al. (2017) for morphologically classified cell types in an anatomical reconstruction and simulation of a rat hindlimb somatosensory cortex column (Markram et al., 2015). Such profiles result largely from the axo-dendritic overlap of the neuronal morphologies (Amirikian, 2005; Brown and Hestrin, 2009; Hill et al., 2012). Broader excitation than inhibition is in line with the experimental data since excitatory neurons, in particular pyramidal types, develop axons with larger horizontal reach compared to most inhibitory interneuron types (Budd and Kisvárdy, 2001; Binzegger et al., 2004; Buzás et al., 2006; Binzegger et al., 2007; Stepanyants et al., 2008, 2009; Ohana et al., 2012). Certain interneuron types may, however, have elaborate axons that span and form synapses across different layers within the cortical column (see, for example, Markram et al., 2015, Figure 2). Others may also form longer-range lateral connections (McDonald and Burkhalter, 1993).

The chosen value for the conduction speed $v = 0.3$ mm/ms is in the range of speeds reported for action potential propagation along unmyelinated nerve fibers in cortex. Conduction speeds can be measured, for example, in brain slices using electrical stimulation combined with electrophysiological recordings:

0.2 – 0.35 mm/ms in guinea pig hippocampus (Andersen et al., 1978), $1/(3.5 \text{ ms/mm}) \approx 0.29$ mm/ms at 34 – 35°C in cat visual cortex (Hirsch and Gilbert, 1991), 0.3 mm/ms at 35°C in rat hippocampus (Berg-Johnsen and Langmoen, 1992), 0.15 – 0.55 mm/ms at $31 \pm 0.5^\circ\text{C}$ in rat visual cortex (Murakoshi et al., 1993), 0.28 – 0.48 mm/ms (mean \pm standard deviation, 0.37 ± 0.37 mm/ms) at 35°C in cat motor cortex (Kang et al., 1994), 0.28 ± 0.19 mm/ms at 34°C in rat visual cortex (Lohmann and Rörig, 1994), 0.06 – 0.2 mm/ms at 34 – 35°C in rat somatosensory cortex (Salin and Prince, 1996), 0.508 mm/ms at 32 – 35°C in rat somatosensory cortex (Larkum et al., 2001, back-propagating action-potentials in dendrites), and 0.34 – 0.44 mm/ms at $34 \pm 1^\circ\text{C}$ in rat somatosensory cortex. Some of these values are likely underestimated because the separation of conduction speed from both the synaptic delay and spike initiation time is difficult (Hirsch and Gilbert, 1991). The bath temperature is provided if specified by the study because the conduction speed and the timing of synaptic processing depend strongly on environmental temperature (Katz and Miledi, 1965; Berg-Johnsen and Langmoen, 1992; Sabatini and Regehr, 1996; Hardingham and Larkman, 1998). We are here primarily interested in physiologically relevant body temperatures. Connections in the upscaled models have a delay offset $d_0 = 0.5$ ms comparable to the experimental estimates 0.5 – 1 ms (Murakoshi et al., 1993), 0.6 – 0.8 ms (Hirsch

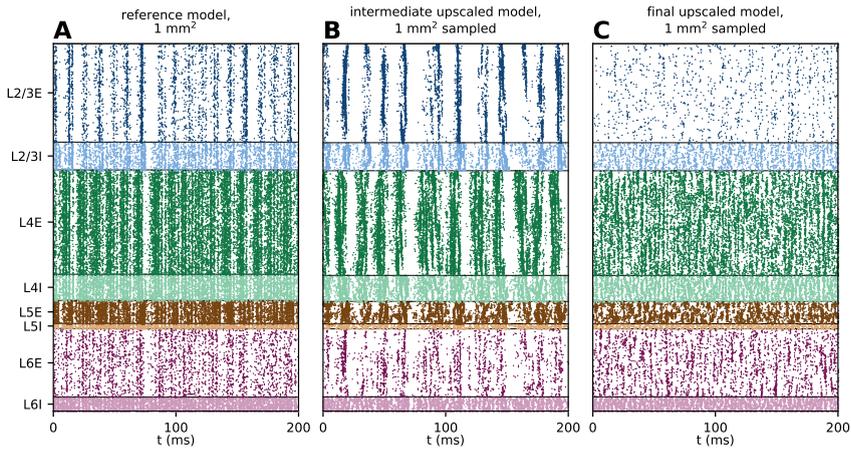


Figure 6.2: **Spiking activity of the reference, intermediate upscaled, and final upscaled models.** **A** Spike raster showing the spike times (horizontal) of all neurons of the reference model network (microcircuit below 1 mm^2 of cortical surface, no spatial connectivity structure) vertically organized according to layer (axes labeling and colors) and neuron type (lighter for inhibitory). **B** Spike raster of a model network upscaled to $4 \times 4 \text{ mm}^2$ with distance-dependent connectivity. The intermediate connection probabilities C_{YX}^{ui} resulting from the upscaling procedure are not modified ($\delta C_{YX} = 0$). Spike times of all neurons located inside a disc of 1 mm^2 shown (neurons are always sorted vertically according to their x -position). **C** Same as panel B, but with modified connection probabilities C_{YX}^{u} according to δC_{YX} given in Table 6.7. The parameters of the final upscaled model are referred to as ‘base parameters’ and given in Tables 6.4, 6.6 and 6.7.

and Gilbert, 1991) and 0.6 ms (Kang et al., 1994). To account for this variability in experimental data the delays have an additive normally distributed random component, see Equation 6.18. From a theoretical perspective, a wide delay distribution expands the region of stability in the phase space of stationary network activity (Brunel, 2000, Section 5.2).

Although delay offset and conduction speed have the same parameter values for excitatory and inhibitory connections in the upscaled model, the effective delays (Equation 6.10) within a given surface area differ due to the different space constants of the connectivity. Computing the mean delay for connections within a circle of 1 mm^2 with the respective spatial widths according to Equation 6.10 results in a shorter mean delay for inhibitory connections. The effective excitatory and inhibitory delays up to single decimal precision are 1.6 ms and 0.9 ms, respectively. Hence, a shorter inhibitory delay in a network model without distance dependence like the reference model is justified by a narrower inhibitory connectivity of the corresponding model with spatial structure.

The spike raster in Figure 6.2A shows that the reference model produces asynchronous irregular spiking with low firing rates (Softky and Koch, 1993; Brunel and Hakim, 1999; Brunel, 2000) across all populations. Network oscillations appear as weakly pronounced vertical stripes. The firing rates are on average higher for inhibitory populations than for excitatory populations within the same layer,

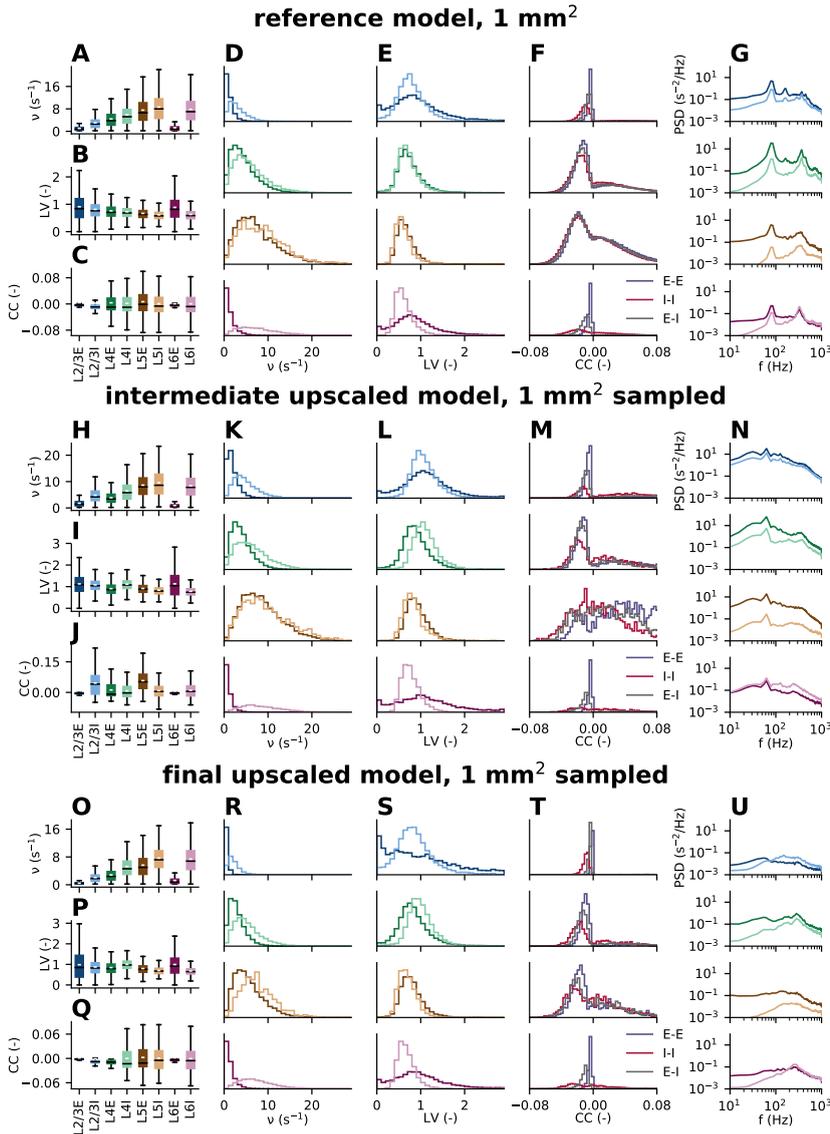


Figure 6.3: **Statistics of spiking activity of the reference, intermediate upscaled, and final upscaled models.** A–G Statistics of spiking activity of reference model shown in Figure 6.2A. A Heterogeneity of spike rates ν for each population (horizontal black lines: median, short white lines: mean, boxes in population-specific colors: lower and upper quartiles of the data, whiskers extend to most extreme observations within $1.5 \times \text{IQR}$ beyond the IQR (interquartile range) without outliers, see documentation of matplotlib.pyplot.boxplot). B Coefficients of local variation LV , see Equation 6.33. C Pearson correlation coefficients CC , see Equation 6.34. D Distributions of spike rates ν . E Distributions of coefficients of local variation LV . F Distributions of Pearson correlation coefficients CC . G Population-rate power spectral densities PSD . H–N Same as panels A–G for spiking activity of intermediate model shown in Figure 6.2B. O–U Same as panels A–G for spiking activity of upscaled model shown in Figure 6.2C.

see Figure 6.3A, and the mean illustrated in each box-chart is larger than the median. The latter corresponds to the long-tailed distributions of spike rates in Figure 6.3D with most neurons firing at lower rates, while few neurons have high ($> 20 \text{ s}^{-1}$) rates. This type of non-symmetric distribution of firing rates in the model resembles approximately lognormally distributed firing rates observed experimentally (reviewed in Buzsáki and Mizuseki, 2014). The mean values of the coefficients of local variation (Figure 6.3B,E) are slightly below unity, indicating more regular spike trains than events produced by a Poisson point process ($LV = 1$). The distributions are broad, that is, a fraction of neurons in each population has spike-train statistics with $LV > 1$. The mean LV values are comparable to values observed in visual cortex across different species (Mochizuki et al., 2016, Figure 5B). The box charts in Figure 6.3A,B are similar to (Potjans and Diesmann, 2014, Figure 6) showing firing rates and the conventional coefficient of variation (Shinomoto et al., 2003, Equation 2.1). The Pearson correlation coefficients (Figure 6.3C,F) are distributed and have a mean close to zero. Weak pairwise spike-train correlations (with mean values < 0.1 using 50 ms windows) are reported, for example, by Ecker et al. (2010) who record from nearby neurons in primary visual cortex of awake monkey under different stimulation conditions, and by Renart et al. (2010) in somatosensory and auditory cortex of anesthetized rats. The latter study finds that the mean correlations are not distance-dependent, but their standard deviations decay with distance (their Figure S11). The authors also include a theoretical analysis of this phenomenon for networks of infinite size and find that excitatory and inhibitory synaptic currents are anticorrelated, thereby leading to a suppression of shared-input correlations, and, hence, weak overall correlations in the asynchronous state. Tetzlaff et al. (2012) and Helias et al. (2014) identify the mechanism underlying the suppression of shared-input correlations for the realistic case of finite-sized networks, which differs from the mechanism in the infinite-size limit. They show that the decorrelation is due to dominant negative feedback, which leads to small correlations in both excitatory-inhibitory and purely inhibitory networks. However, depending on factors such as brain state and distance, stronger correlations are also detected in some cases (Smith and Kohn, 2008; Kriener et al., 2009; Peyrache et al., 2012; Smith et al., 2012; Doiron et al., 2016; Rosenbaum et al., 2017). The population-rate power spectral densities in Figure 6.3G show that the power tends to be higher in the activity of excitatory compared to inhibitory populations due to the overall larger density of excitatory neurons, except for layer 6, where the inhibitory rate is very high compared to the excitatory rate. Across layers the power is highest in layer 4, explained by the comparatively high spike rates and high cell densities. The power spectra reveal two dominant oscillation frequencies of the network in the low and high gamma ranges ($\sim 80 \text{ Hz}$ and $\sim 320 \text{ Hz}$). Recent theoretical work by Bos et al. (2016) provides insight into the main pathways between the recurrently connected populations involved in generating these high-frequency oscillations. The low-gamma peak is predominantly generated by a sub-circuit of layer 2/3 and layer 4 populations of excitatory and inhibitory neurons (pyramidal-interneuron gamma or “PING” mechanism (Leung, 1982; Börgers and Kopell, 2003, 2005), while the high-gamma peak results from interneuron-interneuron interactions (interneuron-interneuron gamma or “ING” mechanism, see Whittington et al., 1995; Wang and

Buzsáki, 1996; Chow et al., 1998; Whittington et al., 2000) within each layer. See Buzsáki and Wang (2012) for a review on the various mechanisms underlying gamma oscillations.

Before we discuss the final upscaled model in comparison to the reference model, we first introduce an intermediate model in order to differentiate between effects of pure upscaling and effects of modified connection probabilities on network activity. This intermediate model is upscaled as described in Section 6.2.1.2 resulting in connection probabilities C_{YX}^{ui} derived directly from C_{YX}^{r} (from Equation 6.5). No connection probabilities are otherwise perturbed ($\delta C_{YX} = 0$ for all X and Y). All model parameters are as specified in Tables 6.4 to 6.7 apart from the connection probabilities and the in-degrees of external input, which are derived as specified in Section 6.2.1.2. This intermediate model covers an area of $4 \times 4 \text{ mm}^2$, but we here choose to analyze only the spiking activity of neurons inside a disc of 1 mm^2 at the center to obtain a representative sample for comparison with the reference model in terms of neuron numbers and spatial scale. The spike raster of the intermediate model (Figure 6.2B) exhibits by visual inspection spatially inhomogeneous activity and network synchrony that are more pronounced than observed in the reference model. Compared to the reference model, spike-train correlations in this intermediate model are increased by approximately an order of magnitude (Figure 6.3J,M), the coefficients of local variation are slightly increased (Figure 6.3I), and finally the overall power in the rate spectra is increased across all frequencies (Figure 6.3N). The spectra also exhibit reduced low- and high-gamma peaks, and the activity is generally more broadband.

The high global synchrony observed in the spiking of the intermediate upscaled model is most likely exaggerated. There is accumulating evidence that the typical operating regime of sensory cortices is asynchronous and irregular in particular when no particular stimulus is present. Measures of LFP signals, which are assumed to mainly reflect synaptic activity, in for example visual cortex also do not show pronounced peaks in their spectra in the absence of stimuli (see, for example, Berens et al., 2008; Jia et al., 2011; Ray and Maunsell, 2011a; Jia et al., 2013a; van Kerkoerle et al., 2014). We therefore modify the network to suppress the amplitudes of the two dominant oscillations in the low- and high-gamma range, and reduce their frequencies to better resemble the low and high-gamma peaks more commonly reported in the literature. For the final upscaled model, we adapt connection probabilities by applying the modifications δC_{YX} given in Table 6.7. The connection probabilities in the reference model are estimated across different areas and species and are merely suggestive of typical cortical connectivity—we therefore consider small modifications to these values to be within the bounds of uncertainties of these probabilities. Our choices on which connections to perturb rely on the framework developed by Bos et al. (2016) who provide a ‘sensitivity measure’ that relates population rate spectra to the connectivity of the underlying neuron network in a systematic manner. With the example of our reference model, they expose which individual connections are crucial for peak amplitudes and frequencies of emerging oscillations, and demonstrate how modifications of these connections affect the power spectra. By applying this sensitivity measure to the intermediate upscaled network, we find that its rate spectra are primarily shaped by the same specific connections as in the reference network. To stabilize the circuit, Bos et al.

(2016) reduce the number of connections from L4I to L4E of the reference model for their analysis. With the same aim, we here reduce the connection probability from L4I to L4E and also increase that from L4I to L4I. Both of these modifications reduce amplitude and frequency of the low-gamma peak (Bos et al., 2016, Figure 8A for L4I-L4I). In addition, we increase the number of connections slightly from L5I to L5E and reduce the number of connections from L5I to L5I to further decrease the amplitude of this peak. A decrease of the number of connections from L5E to L5I amplifies low-frequency oscillations (Bos et al., 2016, Figure 8B). The resulting spike raster of the final upscaled model, similarly sampled in the center 1 mm^2 , exhibits temporally and spatially more homogeneous activity (Figure 6.2C) compared to the reference and intermediate networks. The mean spike-train correlations (Figure 6.3Q) are even lower than in the reference model. The power spectra have overall reduced power and its peaks are attenuated (Figure 6.3U). Most visible in populations L2/3E and L4E, a broad low-gamma peak spans roughly 40 – 60 Hz. Across all interneuron populations, a broad high-gamma peak above 100 Hz is present. The per-population spike rates of the reference model are now largely retained in the upscaled model (Figure 6.3O), as the upscaling procedure preserves the mean input of the neurons (see Section 6.2.1.2). The coefficients of local variation (Figure 6.3P) are similar to those of the reference model, although the *LV* of L4I is increased, which we also observe in the intermediate model (Figure 6.3H).

6.3.2 Spiking activity of the point-neuron networks

We have so far established an upscaling procedure of the reference network from an area of 1 mm^2 to an area of 16 mm^2 , which includes small perturbations to connection probabilities between key pre- and post-synaptic populations. The final upscaled network exhibits a stable network state that (1) is asynchronous and irregular across populations, (2) preserves the population rates, (3) preserves the distribution of firing rates, (4) preserves the variability of spike trains, (5) has very low average pairwise spike-train correlations, and (6) has rate spectra without pronounced peaks. We next investigate the spontaneous behavior around this network parameterization ('base parameters') by varying external input rates, inhibitory feedback weights, spatial connection widths, and the delay offset, which are all hard to constrain with available experimental data. We also study evoked thalamocortical activity in different network states in order to quantify the lateral propagation speed of the evoked network response, motivated by reports of propagating cortical activity.

6.3.2.1 Sensitivity to parameter perturbation during spontaneous activity

We here explore the state space of the upscaled network model by running parameter scans of both global network parameters (external input rate and inhibitory weights) and parameters governing distance-dependent connectivity (width of inhibition and delay offset). Theoretical work exposes a crucial sensitivity to network parameters studying the existence and stability of diverse dynamical states (Brunel, 2000; Roxin et al., 2005; Senk et al., 2018d). However, experimental data

from the literature are often sparse and disparate and the mapping of measured quantities to specific model parameters is not straightforward. Therefore, an exploration of the parameter space is necessary in order to characterize the range of possible model behaviors given the experimental constraints on the parameter values and also to obtain an intuition of the model behavior.

We first choose to vary the rate of the external Poisson input ν_{ext} and the relative inhibitory weight g . As shown in a simpler, analytically tractable case (Brunel, 2000), spatially unstructured networks of randomly and sparsely connected excitatory and inhibitory leaky integrate-and-fire neurons can transition between distinct activity states with respect to the regularity of individual neuron firing and the synchrony of population activity upon changing these two parameters. Jumps in $L\bar{V}$ (or the conventional coefficient of variation CV , see Shinomoto et al. 2003, Equation 2.1) and CC during parameter scans of comparable two-population networks typically indicate transitions between states. It is, however, not a priori clear whether or not this analytical insight obtained with a smaller random network generalizes to spatially extended networks incorporating multiple layers and realistic density of neurons and connections such as our upscaled network model. Mehring et al. (2003, Figure 2), Voges and Perrinet (2010, Figure 4) and Voges and Perrinet (2012, Figure 2) study the same parameter space with spatially organized network models; however, only in single-layer and diluted networks.

While the mean population rates $\bar{\nu}$ in a two-population network typically increase when increasing ν_{ext} or decreasing g (see Mehring et al. (2003, Figure 2D), Voges and Perrinet (2010, Figure 4) and Voges and Perrinet (2012, Figure 2) for examples), Figure 6.4A shows that a similar trend does not appear for all populations of our multi-layer upscaled model. Within the parameter range tested, the mean rate of L4E is nearly unaffected upon varying g , and varying ν_{ext} has little effect on the rate of L5E. For L6E, the trend is even reversed. Different responses in different populations is explained by the population-specific network connectivity and competing inhibition and excitation between the different populations. Both recurrent (excitatory and inhibitory) and external (only excitatory) in-degrees and corresponding presynaptic rates result in population-specific means and variances of synaptic inputs. Spike-train irregularity, here quantified by the mean coefficient of local variation $\bar{L}\bar{V}$ in Figure 6.4B, also shows different trends per population. For all populations, the $\bar{L}\bar{V}$ increases when increasing ν_{ext} . Increasing g results in an increased $\bar{L}\bar{V}$ only in layers 4 to 6, while the effect on L2/3 does not show a clear dependency on either parameter in the tested parameter range. The $\bar{L}\bar{V}$ remains below 1 across the whole parameter space for populations L4E and L6I, while the highest values (above 1.3) are observed in L4I and L6E. Mean pairwise spike train correlations $\bar{C}\bar{C}$ in Figure 6.4C, increase for all populations by increasing ν_{ext} and decreasing g .

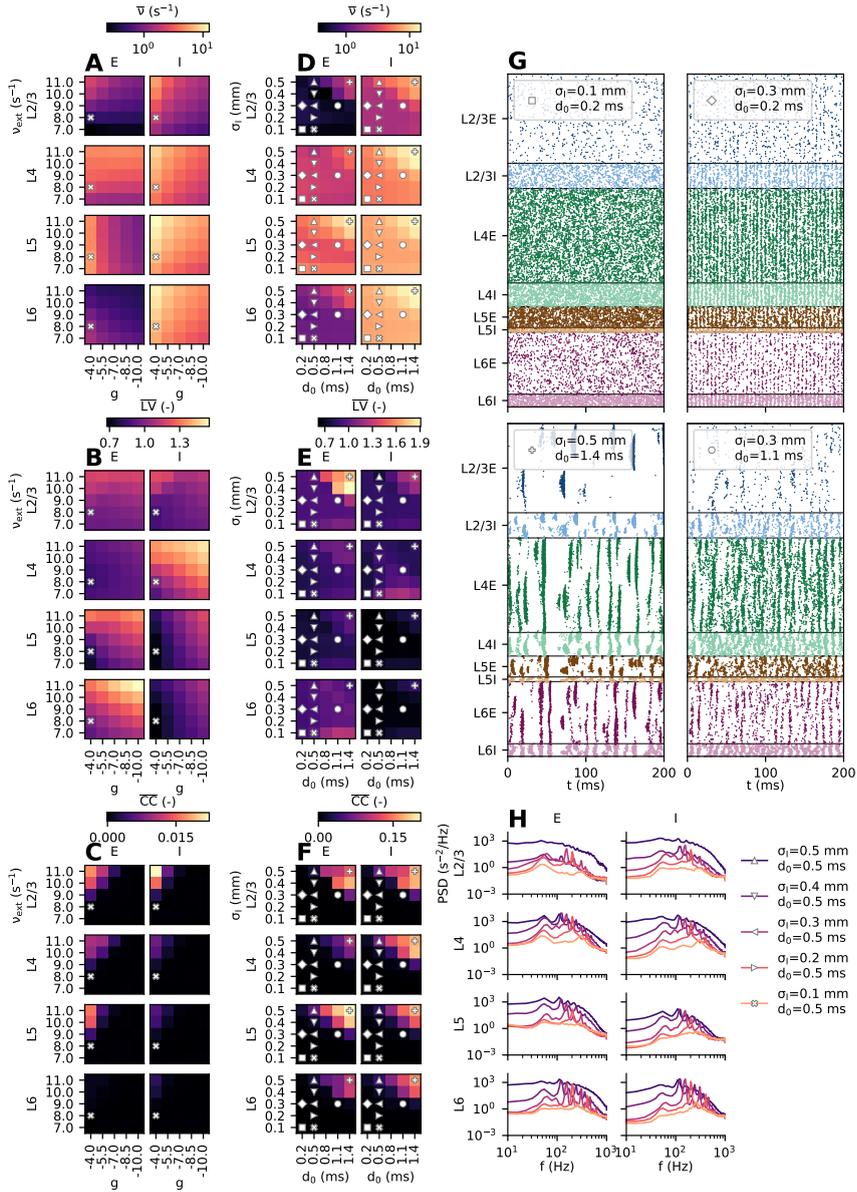


Figure 6.4: **Parameter sensitivity in the upscaled model.** **A–C** Dependency on external rate ν_{ext} and relative weight of inhibition g ($= g_{YX}$ with any inhibitory presynaptic population X). **A** Mean per-neuron spike rates $\bar{\nu}$ for each population (color map with logarithmic scaling). The cross marker denotes the default ‘base parameters’ in this and subsequent panels. **B** Mean coefficients of local variation $\bar{L}\bar{V}$ for each population, see Equation 6.33 (color map with linear scaling). **C** Mean Pearson correlation coefficients $\bar{C}\bar{C}$ between pairs of spike trains for each population, see Figure 6.3 (color map with linear scaling). **D–F** Same as panels A–C, but for dependency on inhibitory spatial width σ_1 and delay offset d_0 . Additional markers refer to parameter combinations used in panels G and H. **G** Spike rasters of selected parameter combinations (showing 3% of all neurons sampled from the full network of size $4 \times 4 \text{ mm}^2$, neurons are sorted as in Figure 6.2). The symbols in each raster plot legend mark the corresponding locations in the parameter space spanned by d_0 and σ_1 (panels D–F). **H** Population-rate power spectral densities (*PSD*) of selected parameter combinations. The markers correspond to the chosen parameter combinations in panels D–F.

Next, we vary the spatial width σ_1 of inhibitory connections and the delay offset d_0 , to assess the sensitivity of the upscaled network dynamics to variations in their chosen values. Although inhibitory spatial widths in terms of lateral axonal branching patterns are generally assumed to be shorter than excitatory widths (Stepanyants et al., 2009), estimates for the local excitatory and inhibitory decay of connection probabilities are broadly distributed and differ between brain areas, pre- and post-synaptic neuron types, and species (Hellwig, 2000; Budd and Kisvárdy, 2001; Boucsein et al., 2011; Kätzel et al., 2011b; Perin et al., 2011; Hill et al., 2012; Levy and Reyes, 2012; Jiang et al., 2015; Schnepel et al., 2015; Reimann et al., 2017). The reduction of multiple cell types and classes into only one excitatory and one inhibitory neuron type per layer in the reference model (Potjans and Diesmann, 2014) implicitly collapses the diversity of neuron morphologies (Amirikian, 2005; Brown and Hestrin, 2009; Hill et al., 2012) which have different spatial connectivity characteristics. Just as for the spatial widths of connections, experimental evidence on distance-dependent delay parameters is also sparse. As reviewed in Section 6.3.1, the estimates for the conduction speed in unmyelinated nerve fibers as well as for delay offsets are also widely distributed. In addition, experimentally obtained spiking statistics exhibit a high variability, even within the same brain area (Mochizuki et al., 2016). While available experimental data on the typical widths of connections of different types and corresponding conduction delays is inherently uncertain, theoretical neural-field model studies frequently investigate the strong influence of these parameters on the stability of the system (Ermentrout, 1998; Coombes, 2005; Roxin et al., 2005; Bressloff, 2012). In our upscaled model, broader inhibition and larger delays increase the mean per-neuron spike rates and the correlations in all populations, shown in Figure 6.4D and F. The effect of changing the parameters d_0 and σ_1 on $\bar{L}\bar{V}$ in Figure 6.4E is again population-specific. The highest $\bar{L}\bar{V}$ values (above 1.6) are obtained for long delays ($d_0 > 1$), and broader inhibition than excitation in L2/3E; the $\bar{L}\bar{V}$ remains low ($\lesssim 0.7$) in L5I and L6I across the whole parameter space. Figure 6.4G shows spike rasters of four distinct network states emerging from this parameter space. Short-range inhibition and short delays yield a spatially and temporally homogeneous state (square marker). Increasing the width of inhibition to an intermediate value results in fast global

oscillations (diamond marker). For broader inhibition than excitation, we observe localized activity spreading outwards (plus marker). Finally, we show an intermediate state (circular marker). These results are in line with predictions from neural-field studies, which indicate that long-range inhibition promotes localized states such as spatially periodic patterns. By contrast, long-range excitation promotes temporally periodic states that can also combine with spatial patterns; see [Ermentrout \(1998, Chapter 8\)](#) and [Senk et al. \(2018d\)](#). For a network of spiking neurons, [Rosenbaum and Doiron \(2014\)](#) show that a balanced state of excitation and inhibition requires broader excitation than inhibition. They demonstrate that the balanced state loses stability if excitation is too narrow compared to inhibition, leading to the emergence of spatial activity patterns.

Finally, [Figure 6.4H](#) shows population rate spectra (*PSD*) varying with the delay offset of the base parameters, $d_0 = 0.5$ ms, and different values for widths of inhibitory connections. While spatially inhomogeneous activity with localized patterns (large σ_1) are manifested as comparatively flat spectra with high power across all frequencies, reducing the spatial width also reduces the overall power, while peaks at the dominant oscillation frequencies emerge. Decreasing σ_1 not only reduces amplitudes of the power spectra, the frequency of the high-gamma peak is also gradually shifted to higher values. Both observations can be related to a reduction of the mean inhibitory delay averaged over all connections in the network due to the shorter-range connectivity. The faster inhibitory feedback results in a stronger decorrelation effect that reduces global oscillations ([Tetzlaff et al., 2012](#); [Helias et al., 2014](#)). The upward shift of the high-gamma frequency is explained by a shorter time period for the ING mechanism ([Bos et al., 2016](#)).

6.3.2.2 Sensitivity to perturbed parameters during evoked activity

We have so far only considered networks receiving external inputs with stationary rates. Cortical areas are, however, recurrently connected to other parts of cortex and subcortical structures, and receive inputs with large rate fluctuations. We here mimic a stimulation experiment, by activating all thalamic neurons inside a disc of radius R_{TC}^{pulse} around $(x, y) = (0, 0)$ once every time interval of Δt_{TC} (see [Table 6.7](#) for values). The activation could for example represent a visual stimulation experiment where activity in lateral geniculate nucleus (LGN, or visual thalamus) thalamocortical (TC) projection neurons is evoked by a brief flash stimulus to a part of the visual field ([Bringuier et al., 1999](#); [Muller et al., 2014](#)), air puffs or mechanical whisker deflections to stimulate whisker barrel cortex ([Swadlow et al., 2002](#); [Einevoll et al., 2007](#)), or direct electric or optogenetic stimulation of the thalamocortical pathway ([Klein et al., 2016](#)). In its population-specific responses to thalamic pulses, the reference model of [Potjans and Diesmann \(2014, page 802\)](#) exhibits a “handshake principle”, in which the receiving layer inhibits the sending layer as if to signal that it has received the message, so that the sending layer can stop transmitting. We test whether this effect and its strength are preserved in the upscaled model. Furthermore, we derive the propagation speed of evoked spiking activity spreading outward from the center of stimulation. Finally, we test the robustness of the propagation speed to parameter perturbations by varying the conduction speed and the delay offset.

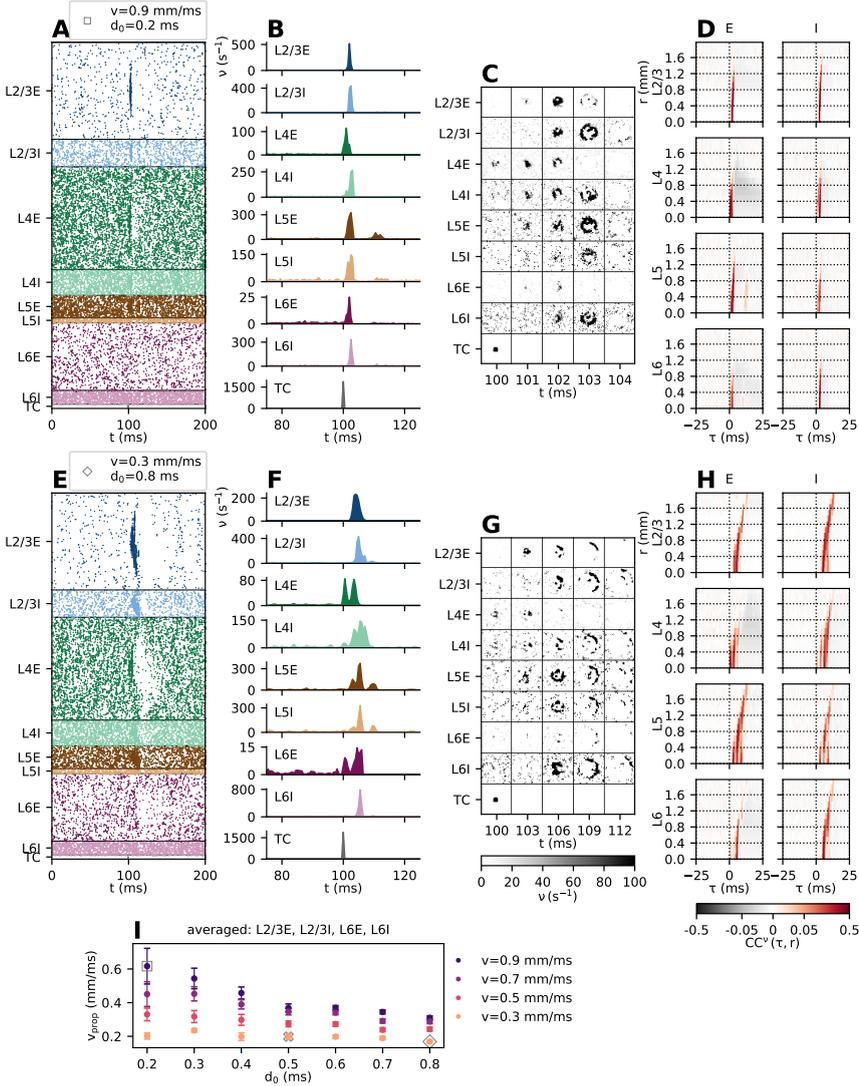


Figure 6.5: Activity evoked by thalamic pulses. **A** Spike raster (showing 3% of all neurons in $4 \times 4\text{ mm}^2$, neurons are sorted as in Figure 6.2). A single thalamic pulse occurs at $t = 100$ ms. **B** Population-averaged rate histogram for neurons within the center disc of 1 mm^2 with bin size Δt for a time interval around the thalamic pulse shown in panel A. **C** Series of snapshots of spatiotemporally binned activity per population over the whole $4 \times 4\text{ mm}^2$ network. **D** Distance-dependent cross-correlation functions between thalamic activation and spatially binned spiking activity $CC^v(\tau, r)$ where r is the distance to the center of the network and τ is the time lag. Color maps have a symmetric logarithmic scaling (linear up to threshold of ± 0.05 indicated by ticks in the color bar). Panels A–D are obtained in a network with conduction speed v and delay offset d_0 as indicated in the legend of panel A. **E–H** Same as panels A–D but with parameters as indicated in the legend of panel E. **I** Propagation speed v_{prop} estimated for parameter combinations of conduction speeds and delay offsets and averaged across populations named above the panel; error bars denote standard deviation $\sigma_{v,\text{prop}}$. The same markers correspond to the same parameter combinations throughout this figure. Base parameters are marked with a cross.

Panels A–D and E–H in Figure 6.5 show results for two different choices of conduction speed v and the delay offset d_0 . At times prior to a thalamic pulse at $t = 100$ ms, the spiking activities in Figure 6.5A and E are comparable, and both asynchronous and irregular, despite the different parameterization. However, the effect of the pulse on the network activity is more pronounced in panel E than in panel A according to visual inspection; the initial response lasts longer and the subsequent activity vanishes for tens of milliseconds in different populations. In Hao et al. (2016, Figure 3) a similar suppression period of tens of milliseconds is observed following a single-pulse electrical micro-stimulation in monkey motor cortex, often followed by a rebound of excitation. In panel E, the effective delay is larger due to the choice of a larger d_0 and a smaller v . In the population-averaged rate histograms of activity within 1 mm^2 in Figure 6.5B and F, corresponding to the spike rasters in panels A and E, respectively, we highlight the transient network responses by zooming into a smaller time window around the pulse. The strong initial response visible in populations L4E and L6E is expected since the thalamocortical input targets layers 4 and 6 directly (see Table 6.6). This evoked activity affects the other network populations via recurrent network connections across and within layers. The larger effective delay (panel H) here increases the response latency of the populations, and increases the duration of the responses while their maximum rates in some populations are reduced. The duration of the activation is overall similar to evoked multi-unit activity (MUA) following whisker stimulation as reported by Einevoll et al. (2007). The multiple peaks in the rate histograms in panel D, most prominent in populations L4E, L5E, L5I and L6E, are due to recurrent excitation and inhibition within and across layers. The overall increased delays expectedly break balance, that is, the high temporal correlation of excitatory and inhibitory spiking activity (see, for example, Renart et al., 2010). These results are comparable with Potjans and Diesmann (2014, Figure 10) and Hagen et al. (2016a, Figure 7), and we therefore conclude that the upscaling procedure does not fundamentally affect the response of the network to transient external input.

While the population-averaged rate histograms in Figure 6.5B and F expose the temporal effect of the perturbation of network activity, we next focus on the corresponding spatiotemporal responses. Figure 6.5C and G show series of snapshots of spatiotemporally binned activity of each population in the full network of size $4 \times 4 \text{ mm}^2$ (similar to Mehring et al., 2003; Yger et al., 2011, Figure 2). The temporal bin size is Δt as in the rate histograms, but we show snapshots only for selected time points as indicated below the frames. The thalamic pulse is visible only at $t = 100$ ms in the center of the network. The cortical populations respond with a ring-like outward spread of activity which can be described as a traveling wave in contrast to a stationary bump (Muller et al., 2018). The wave travels at a lower speed in the network with larger effective delay (compare selected time points in Figure 6.5C and G). In order to derive the radial propagation speed of activity evoked by thalamic pulses, we compute the distance-dependent cross-correlation functions (see Section 6.2.3.6) shown in Figure 6.5D and H. The maximum value of $CC'(\tau, r)$ shifts faster to larger time lags τ with increasing distances r in panel H compared to panel D, which indicates a lower propagation speed. Figure 6.5I summarizes the propagation speed estimates v_{prop} as a function of v and d_0 . The estimated propagation speeds increase with increasing conduc-

tion speed v and decreasing delay offset d_0 . Estimating the propagation speed in this way from spatially resolved spike trains can help to infer underlying network parameters from experimental data. It is to date difficult to observe wave-like activity on the spiking level (Takahashi et al., 2015). However, model predictions for spiking propagation speeds can be compared with population measures, keeping in mind potential differences between spiking activity and population measures such as the LFP. Both types of signals can reflect propagation along long-range horizontal connections which also includes synaptic processing times, but they are also affected by intrinsic dendritic filtering (Grinvald et al., 1994; Nauhaus et al., 2009; Takahashi et al., 2015; Zanos et al., 2015). Muller et al. (2018) remark that macroscopic waves traveling across the whole brain typically exhibit propagation speeds of 1 – 10 mm/ms similar to axonal conduction speeds of myelinated white matter fibers in cortex, while mesoscopic waves (as considered here) show propagation speeds of 0.1 – 0.8 mm/ms similar to axonal conduction speeds of unmyelinated long-range horizontal fibers within the superficial layers of cortex. For example, LFP ‘waves’ in visual cortex travel with such speeds. Nauhaus et al. (2009) study the propagation of spike-triggered LFPs both in spontaneous activity and with visual stimulation and derive speeds (mean \pm standard deviation) of 0.31 ± 0.23 mm/ms in cat and 0.24 ± 0.2 mm/ms in monkey (both anesthetized). Ian Nauhaus and Carandini (2012) reanalyze the data from Nauhaus et al. (2009) and further report a speed of 0.18 mm/ms in cat and 0.29 mm/ms in monkey for the impulse response of ongoing activity; for data from awake monkey (Ray and Maunsell, 2011b) they compute a speed of 0.13 mm/ms. Zanos et al. (2015) measure a speed of 0.31 ± 0.08 mm/ms triggered by saccades in monkey visual cortex. Propagation speeds obtained via voltage-sensitive dye imaging in visual cortex are comparable as well: an average speed of 0.28 mm/ms with a 75% confidence interval of 0.19 to 0.55 mm/ms in cat (Benucci et al., 2007), 0.1 – 0.25 mm/ms in monkey (Grinvald et al., 1994), and a range of 0.25 – 1.35 mm/ms with median \pm standard deviation of 0.57 ± 0.18 mm/ms in monkey (Muller et al., 2014). Estimates from monkey motor cortex are in the same range (Rubino et al., 2006; Takahashi et al., 2015; Denker et al., 2018). For the biologically plausible ranges of delay offsets and conduction speeds tested in the model, $d_0 \in [0.2, 0.8]$ ms and $v \in [0.3, 0.9]$ mm/ms, the resulting propagation speeds are mainly between 0.2 and 0.6 mm/ms. These derived propagation speeds are smaller than the corresponding conduction speeds because propagation through the network includes neuronal integration and the delay offsets. The values in the model cover the range of experimentally measured propagation speeds.

6.3.3 LFP predictions

We here summarize our findings for the predicted LFP signal across cortical space, with recording geometry similar to a 4×4 mm² Utah multi-electrode array. As in Hagen et al. (2016a), the eight cortical network populations spanning layers 2/3, 4, 5 and 6 are expanded into 16 different cell types in order to account for differences in layer specificity of synaptic connections among cell types in a single layer when predicting the LFP. While we here refrain from discussing the detailed derivation of these layer specificities (see Hagen et al., 2016a) from available anatomical data (i.e.,

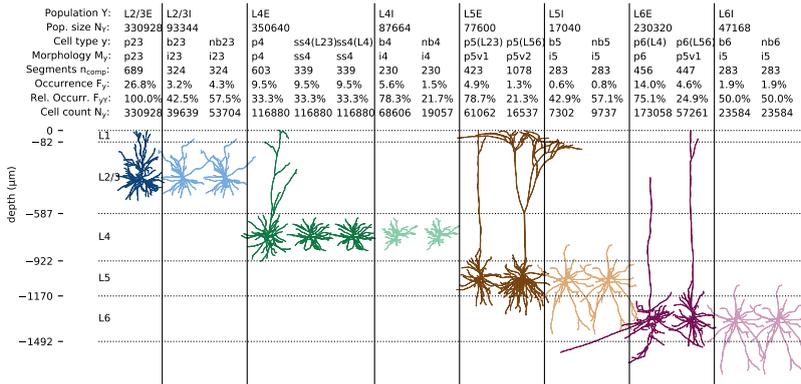


Figure 6.6: **Cell types and morphologies of the multicompartment-neuron populations.** The 8 cortical populations Y of size N_Y in the $4 \times 4 \text{ mm}^2$ network model are represented by 16 subpopulations of cell type y with detailed morphologies M_y (Binzegger et al., 2004; Izhikevich and Edelman, 2008). Neuron reconstructions are obtained from cat visual cortex and cat somatosensory cortex (source: NeuroMorpho.org by Kisvárday and Eysel (1992); Mainen and Sejnowski (1996); Contreras et al. (1997); Ascoli et al. (2007); Stepanyants et al. (2008), see Hagen et al., 2016a, Table 7). Each morphology M_y is here shown in relation to the layer boundaries (horizontal lines). Colors distinguish between network populations as in Figure 6.2. The number of compartments n_{comp} , frequencies of occurrence F_y , relative occurrence F_{yY} and cell count N_y are given for each cell type $y \in Y$.

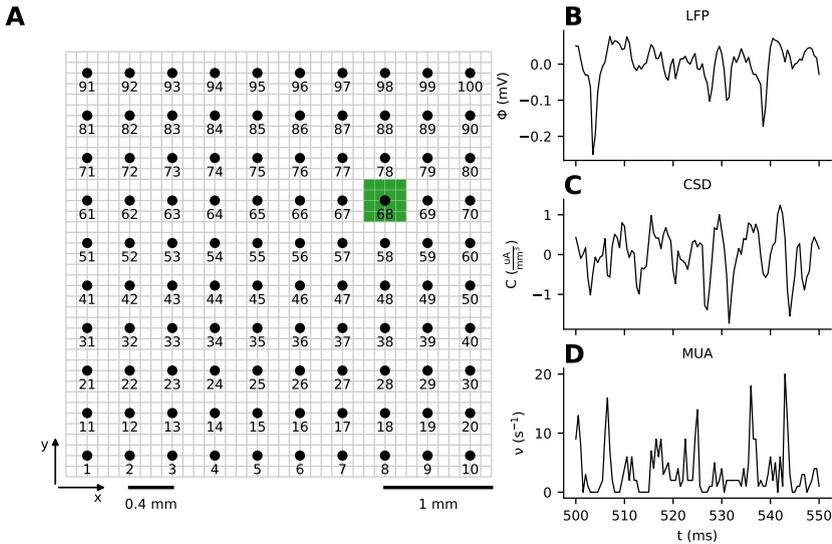


Figure 6.7: **Illustration of multi-electrode array geometry for LFP, CSD, and MUA predictions.** **A** Extracellular potentials are computed in 10×10 electrode locations denoted by circular markers at the depth corresponding to the center of layer 2/3. The electrode inter-contact distance is $400 \mu\text{m}$. The number under each circular marker denotes the channel number. **B–D** Example LFP, CSD, and MUA from one arbitrarily chosen contact (here channel number 68). The CSD is estimated from the LFP using an inverse method, and the MUA is calculated as the sum of excitatory and inhibitory spike events from layer 2/3 neurons in spatiotemporal bins of duration 0.5 ms and width $400 \mu\text{m}$ around each contact.

(Binzegger et al., 2004), in Figure 6.6 we show the reconstructed morphology used for each cell type y in population Y , with compartment counts and occurrences summarized in the table contained within the figure. The cortical layer boundaries and depths are also illustrated, and each morphology is positioned such that the soma is at the center of the corresponding layer. Different cell types belonging to the same population within a layer may have different geometries supporting different layer specificities of synaptic connections. This is the case for example for the p4 pyramidal cell type versus the ss4 spiny stellate cell types that both belong to population L4E of the point-neuron network. Previous modeling studies demonstrate the major effect of the geometry of the morphology on the measured extracellular potential due to intrinsic dendritic filtering of synaptic input (e.g., Lindén et al., 2010; Lindén et al., 2011; Łęski et al., 2013).

The geometry of the recording locations corresponding to the $4 \times 4 \text{ mm}^2$ Utah multi-electrode array is illustrated in Figure 6.7A. The 100 contact locations denoted by circular markers are positioned on a 10×10 grid with $400 \mu\text{m}$ separation between contact sites. LFPs are computed at the center of layer 2/3 (at $z = -334 \mu\text{m}$). An example LFP signal segment from one chosen channel (channel 68) is shown in panel B, corresponding to the spontaneous activity in our laminar,

upscaled point neuron network with ‘base parameters’ introduced above (in Figure 6.2C and corresponding text). The signal fluctuates with amplitudes similar to experimentally observed spontaneous potentials (0.1 – 1 mV, Maier et al., 2010; Hagen et al., 2015; Reyes-Puerta et al., 2016), with occasional larger transients. Further, we estimate from the LFP the underlying current source density (CSD) across space using the so-called kernel CSD method in two dimensions (2DkCSD, Potworowski et al., 2012). The CSD signal is expected to suppress correlations in the LFP resulting from volume conduction, and is therefore less correlated across space as it is taken to reflect the gross in- and outgoing transmembrane currents in vicinity to the recording device (Nicholson and Freeman, 1975; Mitzdorf, 1985; Pettersen et al., 2006, 2008; Potworowski et al., 2012). The LFP and corresponding CSD in general reflect correlations in synaptic input nearby the measurement site and therefore contain contributions from both local and remote neuronal activities. In contrast, the high-frequency ($\gtrsim 100$ Hz) part of experimentally obtained extracellular potentials contains information on spiking activity of local neurons. Activity of high-amplitude single neurons may be separated from the background based on classification of their extracellular action-potential waveforms (through ‘spike sorting’, Quiroga, 2007). Even if no units are clearly discernible in the high-frequency part of the signal, a previous biophysical forward-modeling study using biophysically detailed neuron models (Pettersen et al., 2008) shows that the envelope of the rectified high-pass filtered (750 Hz cutoff frequency) signal correlates with the spike rate in the local population of neurons. In this study, this rectified signal is referred to as the multi-unit activity (MUA), which we approximate by summing up all spiking activities of layer 2/3 neurons in $400 \times 400 \mu\text{m}^2$ spatial bins around each contact. The presently used LFP predictions rely on passive neuron models which do not generate spikes; spiking only occurs in the network. The contribution from excitatory and inhibitory spikes are weighted identically. One example MUA trace obtained at the same location as the LFP and CSD is shown in Figure 6.7D. A notable observation is that the MUA signal and its relations to the corresponding LFP and CSD signals are non-trivial.

6.3.3.1 Distance-dependent correlations of spike trains and LFPs

We next investigate the temporal correlation and coherence with distance for these measures of activity. The observation of weak pairwise spike-train correlations in cortical neuronal networks (for example, Ecker et al., 2010) is seemingly at odds with the typical observation of highly correlated LFPs across cortical space (for example, Nauhaus et al., 2009). We have so far established that the mean pairwise spike-train correlations within populations in our upscaled layered network are typically near zero (Section 6.3.1), and that the perturbation of key network parameters such as the external rate and delays affect the mean correlation (or ‘synchrony’) in the network (Section 6.3.2), as well as other measures like regularity (as measured by their mean coefficients of local variation \bar{LV}). It is, however, not clear how this weakly correlated network activity translates into population signals such as the LFP. Previous modeling studies of mechanisms of the spatial reach of the LFP highlight the crucial role of correlation in synaptic inputs to the LFP-generating neurons (Lindén et al., 2011; Łęski et al., 2013). In contrast to these studies, which use input spike trains with Poisson inter-spike statistics, we

here account for ongoing network interactions, and realistic numbers of neurons and connections under $4 \times 4 \text{ mm}^2$ of cortical surface using the methods to compute LFPs introduced by Hagen et al. (2016a). We thus extend our analysis to distance-dependent correlations in LFP, CSD, MUA, and pairs of spike trains.

For spontaneous spiking activity in the upscaled network (Figure 6.8A), we compute the LFP (panel C), reconstruct the underlying CSD from the LFP (panel F), and compute the MUA (panel I) across the 100 channel locations in layer 2/3 illustrated in Figure 6.7A. The network parameters and corresponding network state are those resulting from our upscaling procedure (see Section 6.2.1.2, base parameters given in Tables 6.4, 6.6 and 6.7). Visual inspection of panel C reveals that the LFP amplitude across channels is typically small ($\lesssim 0.5 \text{ mV}$) as highlighted in Figure 6.7B with occasional transients which may be seen also on neighboring channels. These transient events presumably result from spatially confined synchronization in the network, but are not seen across every LFP channel as would be the case with globally synchronous network events. The amplitudes observed here are similar to those from the forward-model predictions of LFPs from spontaneous activity in the original 1 mm^2 network model (Hagen et al., 2016a, Figure 8M), even if the total number of neurons in the upscaled model is increased by a factor of 16. These similar amplitudes are partially explained by the suppression of strong low-gamma oscillations in the upscaled network using modified connection probabilities. An increase in network synchrony (that is, increased correlations) can otherwise be expected to increase LFP amplitudes overall due to an increased pairwise cross-correlation between single-neuron contributions to the LFP (Hagen et al., 2016a). The network upscaling procedure does not obliterate the high- and low-gamma oscillations, which in the LFP spectra result in a large peak around 200 Hz and a small peak around 50 Hz. The network receives background input with a flat power spectrum (driven by a Poisson process with fixed rate) and has no internal sub-circuits capable of generating rate fluctuations or slow oscillations. Hence, the LFP in each channel contains little power towards small frequencies. Another factor explaining the lack of low frequency power is active decorrelation by inhibitory feedback, which is shown to suppress population-rate fluctuations (Tetzlaff et al., 2012; Helias et al., 2014).

We next compute the Pearson product-moment correlation coefficient between all possible pairs of LFP channels, and sort by inter-contact distance (panel E). The mean and standard deviation for each discrete contact separation are shown by the black line and corresponding error bars. Due to the periodic boundary conditions of the network, the longest possible inter-contact distance is $L/\sqrt{2} \approx 2.8 \text{ mm}$. The mean values are well fit by a simple exponential function (red line), with a spatial decay constant of $\sim 0.63 \text{ mm}$ and constant offset of ~ 0.3 . The histogram to the right is computed for all observed correlation coefficients. The correlations in the simulated LFP are lower compared to findings by Nauhaus et al. (2009, Fig. 8) during spontaneous activity in anesthetized macaque (approximately 0.95 at 0.4 mm and 0.75 at 2.4 mm electrode separation, respectively) and cat (approximately 0.93 at 0.4 mm and 0.83 at 2.4 mm electrode separation, respectively). With high-contrast drifting grating type stimuli, however, the correlations between pairs of LFP signals are shown to decrease to values around 0.5 at an electrode separation of 2.4 mm. Also Destexhe et al. (1999) analyze spatial correlations in the LFP

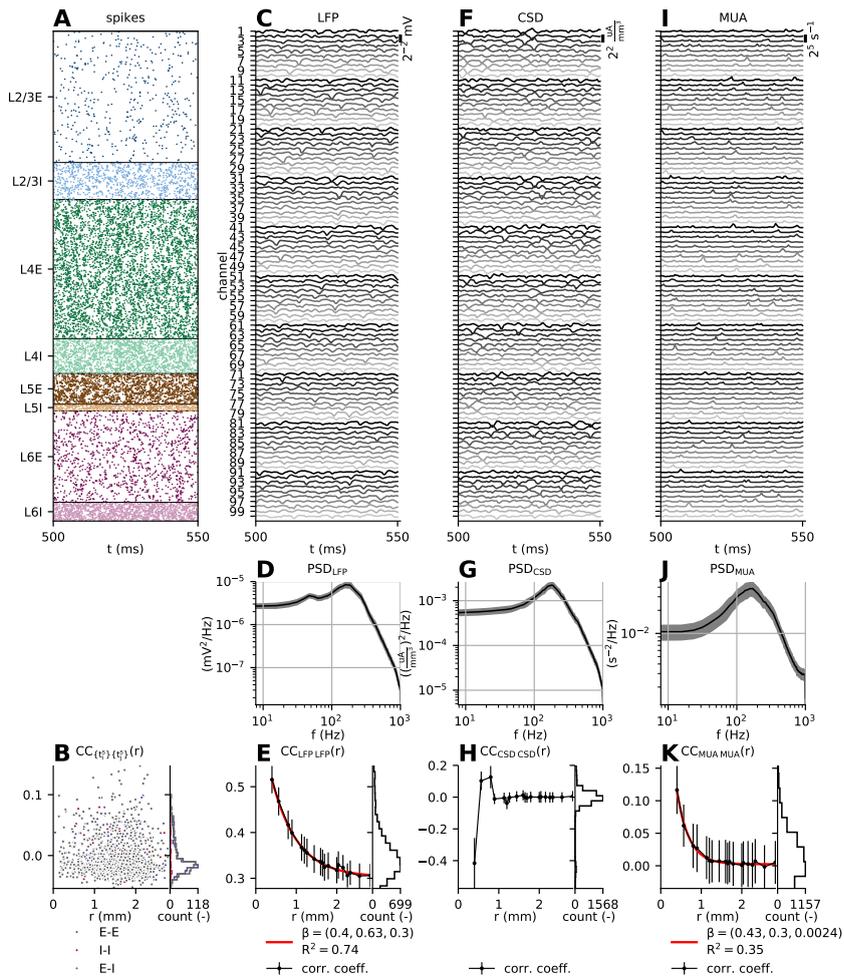


Figure 6.8: **Spikes, LFP, CSD, and MUA: Raw signals, power spectra, and distance-dependent correlations in L2/3.** **A** Spike raster (showing 10% of all neurons in $4 \times 4 \text{ mm}^2$, neurons are sorted as in Figure 6.2). **B** Pairwise spike-train correlations computed for pairs of excitatory (E-E, $n = 40$), inhibitory (I-I, $n = 10$) and excitatory and inhibitory (E-I) L2/3 neurons, sorted by inter-neuron distance r . **C** Local field potentials (LFP) across the 10×10 electrode contact points located at the center of layer 2/3, each separated by $400 \mu\text{m}$ in the lateral directions. **D** LFP power spectrum averaged over channels (black line). The gray area denotes the average spectrum plus/minus one standard deviation. **E** Pearson correlation coefficient between pairs of LFP signals as function of separation between channels. The black line shows the mean at each unique separation, whiskers denote one standard deviation. The red line shows the least-square fit of an exponential function to all values. The coefficient of determination (R^2) is given in the legend. **F** Current-source density (CSD) estimates from LFPs shown in panel C, calculated using the kCSD method in 2D. **G** CSD power spectrum (mean \pm one standard deviation). **H** Similar to panel E but for CSD signals, minus fit to exponential function. **I** Multi-unit activity (MUA) approximated as the bin-wise spike rates of layer 2/3 excitatory and inhibitory point neurons, calculated using a spatial bin width $\Delta h = 400 \mu\text{m}$. **J** MUA power spectrum (mean \pm one standard deviation). **K** Similar to panel E but for MUA signals.

of cat suprasylvian cortex during awake and different sleep states, and find mean correlations of approximately 0.6 at 2 mm contact separation in the awake state. These LFP correlations computed from experimental data are highly dependent on the choice of LFP reference which may introduce a shared signal component (which increases correlations), while the present model LFPs are computed with the assumption of an ideal reference electrode at infinite distance from the sources. The point neuron and corresponding LFP model also ignore rate fluctuations in their background input (here represented as Poisson generators with fixed rates) which is another source of spatial correlations. Global fluctuations or shared input correlations in the background input can be expected to increase pairwise LFP correlations (Lindén et al., 2011; Łęski et al., 2013; Hagen et al., 2016a).

We next bring our attention to the estimated CSD signal in panel F. By design the chosen CSD estimation method is expected to suppress correlations among channels due to volume conduction by reconstructing the sink/source pattern underlying the LFP (Nicholson and Freeman, 1975; Mitzdorf, 1985; Pettersen et al., 2006, 2008; Potworowski et al., 2012). This can, for example, allow the identification of loci of strong synaptic activity in experimental LFP data, which may be generated locally or due to some external drive. A brief inspection of the CSD traces computed from the model LFP reveals that ‘standout’ LFP events (e.g., in channel 31 at 525 ms) result in fluctuations in the corresponding CSD, but the traces appear overall more variable than the LFP. Just as for the LFP, we show the power spectra (panel G) and pairwise correlation coefficients with distance (panel H). In contrast to the LFP spectra, the low-gamma peak around 50 Hz is not present in the CSD spectra, but the high-frequency peak remains. The overall positive correlations observed for the LFP are largely canceled for the CSD. The CSD signals are typically anti-correlated with mean around -0.4 at the shortest electrode separations (0.4 mm), and then weakly correlated (~ 0.1) up to 1 mm. This CSD anti-correlation across proximal channels is expected, as a fraction of capacitive and resistive (‘leaky’) transmembrane return currents of synaptic input currents exits in vicinity to the synapse site and at the soma. The return currents are affected by intrinsic dendritic filtering (Lindén et al., 2010) throughout each individual LFP-generating neuron morphology. Our multicompartment cells are effectively treated as closed electric circuits and the basic principle of charge conservation must apply (see, for example, De Schutter and Van Geit, 2009a). The correlations between channels are negligible beyond 1 mm electrode separation. This negligible correlation at greater distances reflects in part that dendrites of each morphology (cf. Figure 6.6) are mostly confined within $\sim 300 \mu\text{m}$ in the lateral directions, and that local spontaneous network interactions for this particular network parameterization do not readily propagate across space. It is important to point out that the CSD estimate (cf. Section 6.2.3.3) is based on LFPs computed at a single depth only, and would change if LFPs across all depths were taken into account.

As an approximation to the so-called multi-unit activity (MUA) signal, we sum up spiking activity in layer 2/3 in the vicinity of each LFP contact point (cf. Section 6.2.3.4), resulting in the signals in Figure 6.8I. Similar to the computed LFP and CSD signals, we compute power spectra (panel J) and distance-dependent pairwise correlations among MUA signals (Figure 6.8K). In contrast to pairwise

spike-train correlations (Figure 6.8B), a sharply decaying distance dependency is observed, which is well fit by an exponential function with spatial decay constant of ~ 0.30 mm and vanishing offset from zero at greater distances. This sharp decay contrasts with the longer spatial decay constant observed for the LFP, and the anti-correlation between neighboring sites as observed for the CSD does not occur. These differences reflect that the LFP and CSD are measures resulting from synaptically driven transmembrane currents, while the MUA is a measure of the network spiking activity resulting from said synaptic input. Similar to the CSD spectra, the low-gamma oscillation around 50 Hz is not seen, while the high-gamma oscillation around 200 Hz is pronounced.

6.3.3.2 *Spatial coherence of local field potential is band-passed*

So far we have established that the model LFP is highly correlated with distance in qualitative agreement with experimental findings, while the corresponding CSD and MUA signals are hardly correlated beyond electrode separations of ~ 1 mm. We next extend this analysis to the frequency domain by considering distance-dependent coherences. This step is mainly motivated by two experimental observations: LFP coherence across channels depends on the inter-electrode distance as described by Jia et al. (2011) and Srinath and Ray (2014), and a recent study by Dubey and Ray (2016) shows that the ‘spatial spread’ of LFP has band-pass properties in the gamma range (60 – 150 Hz). Another modeling study (Łeński et al., 2013) extends the study of LFP ‘reach’ by Lindén et al. (2011) to distance-dependent coherences, showing that dendritic filtering (Lindén et al., 2010) introduces a low-pass effect on the LFP reach of uncorrelated synaptic input currents with an approximately white power spectrum. In contrast to these latter modeling studies, our combined point-neuron network and LFP-generating setup allows accounting for weakly correlated spiking activity in the network, at realistic density of neurons and connections.

From its spectra (Figure 6.8D) we infer that most of the variance in the spontaneous LFP data is due to a high-frequency gamma oscillation above 200 Hz in the network due to the ING mechanism present in each layer (Bos et al., 2016). In Figure 6.9A we show the mean coherences $\langle \gamma_{\text{LFP-LFP}} \rangle(f)$ between individual pairs of LFP signals from channels separated by a distance $r = \{0.4, 0.8, 1.2, 1.6, 2.0\}$ mm. While the coherence is, as expected, highest for the lowest frequencies ($\lesssim 10$ Hz) at all separations, it drops quickly for frequencies $f \approx 20$ Hz. For the shortest separation (0.4 mm), the coherence is around 0.35 at this frequency, and increases to ~ 0.5 in the low-gamma range (around 50 Hz). Broader peaks in the coherence with magnitudes around 0.3 and 2 are also seen for 100 – 150 Hz and 250 – 300 Hz, respectively. Beyond this range, the coherence drops to around 0.1. The coherence across all frequencies is further reduced for increased separations, but at 2 mm separation it still drops to the same value of ~ 0.1 at high frequencies. These model observations resemble coherences computed for experimental LFP data during stimulus conditions (Srinath and Ray, 2014, Figure 1A). There, a peak in low-gamma coherence around 40 Hz is seen for distances up to 4 mm in two different subjects. Also an increase in coherence is seen for frequencies around 80 Hz. The baseline coherence (no visual stimulus) shows no increase in the gamma range of frequencies, except for sharp peaks seen at 100 Hz due to the CRT display frequency and 120 Hz

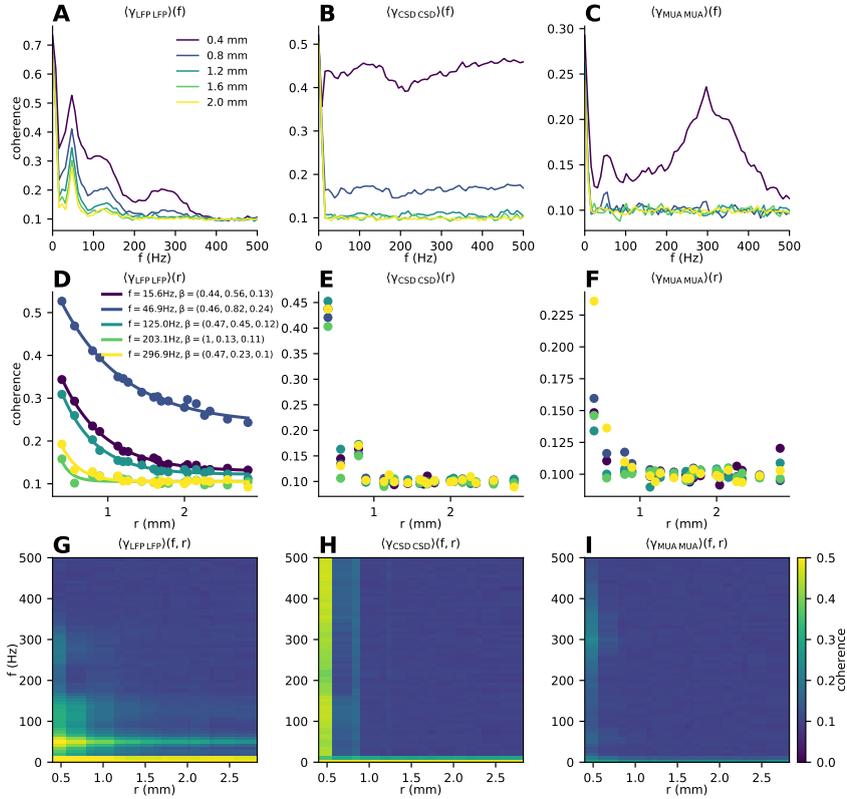


Figure 6.9: Distance dependency of LFP, CSD, and MUA coherences in L2/3. **A** Pairwise LFP coherences as function of frequency for different distances (color-coded) between electrode contacts r , averaged over pairs with identical electrode separation. **B** Similar to panel A but computed using the reconstructed CSD signal estimates at each electrode. **C** Similar to panel A and B but computed using the MUA signal at each electrode. **D** Mean LFP coherences as function of distance between electrode contacts for different frequencies (color-coded) with exponential fit to mean values ($R^2 = \{1, 0.99, 0.99, 0.70, 0.93\}$ for $f = \{15.6, 46.9, 125.0, 203.1, 296.9\}$ Hz, respectively) **E,F** Mean CSD, and MUA coherences as function of distance between electrode contacts for different frequencies (color-coded). **G–I** Color image plot of mean LFP, CSD, and MUA coherences as function of frequency and electrode separation.

due to the second harmonic of noise (Srinath and Ray, 2014). This lack of gamma peaks of physiological origin differs from our model predictions. We therefore conclude that the model LFP coherence more closely resembles the stimulus-driven LFP, but we note also that a baseline stationary thalamocortical activation level is assumed in the reference network (Potjans and Diesmann, 2014). This baseline activation enters the $K_{X,\text{ext}}^u$ parameter for populations $X \in \{L4, L6\} \times \{E, I\}$ also in the upscaled network. The corresponding mean-field theory (Bos et al., 2016) identifies sub-circuits located in and across layers 2/3 and 4 as the origins of the low-gamma oscillations. Therefore, a reduced external drive to layer 4 (by turning off the baseline thalamic activation altogether) should reduce the magnitude of this intrinsically generated oscillation and consequently reduce the corresponding spatial LFP coherence. An opposite effect on coherence can be expected by increasing the thalamocortical drive in the model. At present we do not pursue this possibility further. We also note that our LFP coherence is smaller than the comparable experimental values (Srinath and Ray, 2014), but see also Jia et al. (2011) and Dubey and Ray (2016). This smaller coherence underlies the reduced correlation with distance noted above which is likely due to the lack of temporally modulated input and intrinsically generated low-frequency fluctuations, and the use of an ideal reference. Some of these differences may also result from the fact that these experimental studies rely on the multitaper method (Thomson, 1982) in order to compute coherences while we use Welch’s average periodogram (see Section 6.2.3.6), and that the experimental data have longer durations than our chosen simulation period of $T_{\text{sim}} = 5$ s.

We next investigate the distance dependency of coherences for different frequencies. Dubey and Ray (2016) show an apparent band-pass effect in the LFP, in that the phase coherence across sites is increased and decays more slowly with distance in the gamma range compared to higher and lower frequencies. In Figure 6.9D we show the LFP-LFP coherences as functions of distance for different frequencies f , averaged over values computed for identical separation of channels. We also show the corresponding least-square fit to exponential functions, in order to investigate whether or not this model reproduces the experimentally observed band-pass effect. Indeed, we find for $f \approx 46.9$ Hz, which is at the center of the low-gamma peak in panel A, an elevated coherence with longer spatial decay constant ($\lambda = 0.82$ mm) than for frequencies where the overall coherence is reduced, such as $f = 15.6$ Hz ($\lambda = 0.56$ mm) and $f = 203.1$ Hz ($\lambda = 0.13$ mm). In the high-gamma band ($f = 296.9$ Hz) we again note a comparatively quick decay ($\lambda = 0.23$ mm) in coherence, which may reflect that the network interactions underlying the generation of this oscillation frequency remain local. In panel G we show the same data as displayed in panel A and D for all frequencies up to 500 Hz and average for each distance up to 2.8 mm. As implied by the above findings, the low-gamma peak in the coherence near 50 Hz is seen at all distances.

In a similar manner we compute distance-dependent coherences for the CSD (panels B,E,H) and MUA (panels C,F,I). The CSD shows only a weak frequency dependence in its coherence at all tested distances (panel B). The coherence is ~ 0.4 for a contact separation of 0.4 mm, and drops to levels below 0.2 at greater distances. The MUA coherence, however, is increased for the shortest distances (0.4 mm) around the high-frequency range of the high-gamma oscillation (~ 300 Hz)

as shown in panel C, but the coherence is at the baseline level at all greater distances.

6.4 DISCUSSION

The present work investigates a multi-layer point-neuron network model covering $4 \times 4 \text{ mm}^2$ of cortical surface at realistic neuron and connection density, amounting to $\sim 1.2 \cdot 10^6$ neurons and $\sim 5.5 \cdot 10^9$ synapses. The model accounts for spiking activity across excitatory and inhibitory neurons in layers 2/3, 4, 5, and 6 and one external thalamocortical population, as well as local field potentials (LFP). The $4 \times 4 \text{ mm}^2$ area covered by the model is similar to the one covered by a 10×10 Utah multi-electrode array commonly used for electrophysiological measurements in vivo in different cortical areas and species. The model is a laterally extended version of the cortical microcircuit under 1 mm^2 of cortical surface by [Potjans and Diesmann \(2014\)](#), but in contrast to this reference network the upscaled network accounts for distance-dependent connection probabilities and delays. The biophysics-based LFP predictions rely on the hybrid scheme for LFP predictions in point-neuron networks by [Hagen et al. \(2016a\)](#), which is here modified to account for spatially structured networks. Earlier work has shown that correlations are perturbed in downscaled networks ([van Albada et al., 2015](#)). The LFP reflects the fluctuations caused by network correlations and depends also on the spatial organization of networks (see, for example, [Hagen et al., 2016a](#)). Therefore, the development of biophysical network models that incorporate the full density of connections as well as the spatial organization of the observed system is crucial to aid the interpretation of the corresponding experimental data.

Our upscaling procedure preserves the overall features of activity in the reference network. This includes a stable network state with asynchronous and irregular spiking activity for the different neuron populations, distributed firing rates across neurons, spike trains with variability in agreement with observed activity in sensory cortex, weak pairwise spike-train correlations, and population firing rate spectra with peaks in the low-gamma range (40 – 80 Hz) and high-gamma range (200 – 300 Hz). Around this stable state, we investigate the effect of varying key network parameters, namely the weight of inhibitory connections and the external drive, as well as the width of inhibitory connection profiles and the minimum delays. We find that a strong external drive with reduced inhibitory feedback results in high synchrony, that conduction delays strongly affect the formation of temporal oscillations, and that wide inhibition results in spatial instabilities. Furthermore, the model exhibits spatially spreading activity evoked by thalamic pulses comparable to experiments with a brief flash stimulus to a part of the visual field in terms of the radial propagation speed of the evoked responses. Finally, the model accounts for spatially correlated and coherent LFPs even during spontaneous network activity when its pairwise spike-train correlations are low on average. LFP coherences are distance-dependent with a slower spatial decay around the frequency of the 50 Hz low-gamma oscillation compared to other frequencies, resulting in an apparent band-pass filter effect on the LFP coherence.

Comparison with other studies

To our knowledge, this computational study is the first to simultaneously account for both spiking activity and population activity measures such as the LFP in a layered network model that covers several square millimeters of cortical surface at the full density of neurons and synaptic connections. Compared to experimentally reported cortical neuron densities of $\sim 10^5$ neurons/mm² (see, for example, [Herculano-Houzel, 2009](#); [Ribeiro et al., 2013](#)), other studies of laminar point-neuron networks with distance-dependent connections ([Mehring et al., 2003](#); [Yger et al., 2011](#); [Voges and Perrinet, 2012](#); [Rosenbaum and Doiron, 2014](#); [Keane and Gong, 2015](#); [Schnepel et al., 2015](#); [Pyle and Rosenbaum, 2017](#); [Rosenbaum et al., 2017](#)) either rely on reducing the overall size of the network's geometry, reduce the neuron densities per cortical area, consider only one layer of excitatory and inhibitory neurons, or collapse all cortical layers into one. [Tomsett et al. \(2014\)](#) also incorporate LFP predictions from a recurrently connected network of $\sim 10^5$ multicompartment neurons, but consider only a thin cortical slice across layers similar to in vitro experiments. While reduced cell and connection counts speed up simulations, state-of-the-art point-neuron simulation software scales nearly linearly up to $\sim 10^9$ neurons ([Kunkel et al., 2017](#); [Jordan et al., 2018](#)). Hence, simulations of networks with $\sim 10^6$ neurons such as ours can be executed routinely on high-performance computing facilities.

We here choose to start from a previously published model of the cortical microcircuit by [Potjans and Diesmann \(2014\)](#). In increasing the model size, the choice of scaling procedure is critical. [van Albada et al. \(2015\)](#) show that the reducibility (downscaling) of randomly connected asynchronous networks is fundamentally limited if both spike rates and second-order statistics (correlations) are to be preserved. Their proposed scaling rules adjust the amplitudes of synaptic currents and mean and variance of noisy background input to the decreasing numbers of synapses. However, upscaling is different. In the microcircuit model each neuron receives a realistic number of synapses, originating either from within the circuit or attributed to the background. Increasing the network size necessarily decreases the probability for two neurons to be connected. The consideration of spatial organization, however, preserves a certain level of local recurrence while the total network size is growing. Consequently, our upscaling procedure works without the need to adjust the amplitudes of synaptic currents of the reference network. The distance-dependent connectivity results in modified in-degrees of recurrent network connections and noisy background input such that the mean input to each neuron is preserved, but not its variance. As demonstrated here, the activity statistics of neurons in a 1 mm² patch in the upscaled network is comparable to the statistics of the reference network. This retrospectively validates the decisions made in the construction of the microcircuit model by [Potjans and Diesmann \(2014\)](#).

The modeled LFP has amplitudes in agreement with spontaneous LFP amplitudes observed experimentally between 0.1 – 1 mV (see, for example, [Maier et al., 2010](#); [Hagen et al., 2015](#); [Reyes-Puerta et al., 2016](#)). The LFP spectra reveal a strong ongoing oscillation at high frequencies, in the 200 – 300 Hz range, and around 50 Hz. Spectra of spontaneous potentials in visual cortex do not typically reveal

strong oscillations at these frequencies, but elevated LFP gamma power in the 30 – 80 Hz range is frequently reported during stimuli (Jia et al., 2011; Ray and Maunsell, 2011a; Berens et al., 2008; Xing et al., 2012; Veit et al., 2017; Katzner et al., 2009; Jia et al., 2013b; Hadjipapas et al., 2015). A functional role in computation and synchronization between areas has therefore been hypothesized (Ray and Maunsell, 2010; Jia et al., 2013a; Buzsáki and Wang, 2012). The strong high-frequency oscillations here result from short interneuron conduction delays (Bos et al., 2016). Low frequencies are lacking in our spontaneous LFP as our network receives external drive with a stationary rate, does not intrinsically generate slow rate fluctuations, and is subject to active decorrelation (Tetzlaff et al., 2012), as well as due to the assumption of an ideal reference at infinite distance from the source. Nevertheless, the model produces highly correlated LFPs with a distance dependence compatible with experimental observations (Destexhe et al., 1999; Nauhaus et al., 2009). The model also reproduces elevated coherences in the low-gamma band as seen during visual stimulation (Jia et al., 2011; Srinath and Ray, 2014). The slower spatial decay for frequencies around 50 Hz in the model is consistent with a recent report of increased spatial LFP ‘reach’ analogous to a spatial band-pass filter effect in the low-gamma band (Dubey and Ray, 2016).

Possible model refinements

The upscaled model establishes local connections with a Gaussian decay of connection probabilities up to a radius of 2 mm. However, pyramidal neurons can develop long horizontal axons spanning several millimeters in addition to local axonal branching. In cat and monkey visual cortex, these connections are typically clustered or patchy and connect neurons with similar orientation tuning (Livingstone and Hubel, 1984; Gilbert and Wiesel, 1989; Bosking et al., 1997; Tanigawa et al., 2005; Buzás et al., 2006; Binzegger et al., 2007). In contrast, the visual cortex of rodents exhibits a salt-and-pepper layout without patchiness, but still some longer-distance connections (Ohki and Reid, 2007; Laramée and Boire, 2015). Although less common, subsets of inhibitory interneurons can also exhibit long-range connections (McDonald and Burkhalter, 1993). Voges and Perrinet (2012) assess the influence of different types of remote connections (none, random, or patchy) on the network activity of a 2D single-layer network, and conclude that the fraction of local versus remote connections is crucial for the resulting network dynamics, irrespective of the detailed spatial arrangement of remote connections.

For the type of model development conducted here, comprehensive datasets with detailed (distance-dependent) connection probabilities are mostly unavailable for all possible pairs of pre- and postsynaptic neuron types and different cortical layers. Some exceptions exist (for example, Binzegger et al., 2004), but most connectivity studies focus on specific connections, and due to differences in experimental methods, results may be difficult to compare and reconcile; see, for example, Schnepel et al. (2015, Supplementary Material) on the limitations of their photostimulation technique and Stepanyants et al. (2009) on truncated connections in brain slices. Neuron morphology appears to provide a valid first approximation for the distance dependency of connections (Amirikian, 2005; Brown and Hestrin, 2009; Hill et al., 2012; Rees et al., 2016), but the overlap between

dendrites and axons alone does not explain connectivity patterns, due to target neuron-type specificity (Potjans and Diesmann, 2014), specificity at the level of individual neurons (Kasthuri et al., 2015), and preferential locations of dendritic spines and synaptic boutons on connected neurons (Ohana et al., 2012). We make the conservative choice to let the spatial widths of connections and shape of post-synaptic potentials depend only on the presynaptic neuron type. Our hope is that the algorithmic approach pursued within consortia such as the Blue Brain Project (Reimann et al., 2015; Markram et al., 2015) and the Allen Brain Institute (Kandel et al., 2013) will provide more accurate neuronal connectomes of different brain regions across species in the future, including their distance dependencies (as in, for instance, Reimann et al., 2017 for rat somatosensory cortex).

Activity in finite-sized laminar networks is subject to effects that depend on the choice of boundary conditions. Periodic boundary conditions are frequently used in 1D networks (ring networks) (Roxin et al., 2005; Kriener et al., 2014b; Rosenbaum and Doiron, 2014) and in 2D networks with torus connectivity (Mehring et al., 2003; Yger et al., 2011; Voges and Perrinet, 2012; Rosenbaum and Doiron, 2014; Keane and Gong, 2015; Schnepel et al., 2015; Pyle and Rosenbaum, 2017; Rosenbaum et al., 2017) as also used here for the upscaled models. The model of a cortical slice by Tomsett et al. (2014) incorporates connections only within the confines of the modeled slice, but we consider networks that are part of a larger system (the intact brain). An advantage of periodic boundaries is the simplifying assumption that cortex is homogeneous and isotropic, that is, the connectivity of a neuron is independent of its location in the network. One disadvantage is that the maximal distance for connections is only $L/2$ for a ring domain with circumference L , or $L/\sqrt{2}$ for a square domain with side length L . Here, we restrict connections to a radius $R = L/2$. Another disadvantage is that propagating activity may travel across the boundary and directly influence its own propagation, resulting in for example wave-front annihilation (Muller et al., 2018). An option to suppress such effects would be to simulate a larger network and to sample only the activity of neurons across a smaller domain. In Figures 6.2 and 6.3, we extract activity of neurons within a center disc of 1 mm^2 . The network could be further upscaled, for example to cover a full cortical area. The lateral size of the unfolded cat striate cortex in one hemisphere is larger by a factor of almost 25 than the currently simulated upscaled network of 16 mm^2 , estimated in the range of $310 - 400 \mu\text{m}$ (Tusa et al., 1978; van Essen and Maunsell, 1980; Olavarria and Sluyters, 1985; Anderson et al., 1988). Striate cortex in macaque monkeys is even two to four times larger than in cats (van Essen and Maunsell, 1980). Networks of a full cortical area could also address the effects of borders to adjacent cortical areas. Anatomical borders between distinct areas are shown to affect wave propagation (Xu et al., 2007; Muller et al., 2014).

Spontaneous activity in our models is driven by uncorrelated external inputs with a fixed rate and Poisson statistics, to represent missing connections from remote and neighboring cortices, subcortical structures, and sensory inputs. Ongoing work aims to account for the structure of one hemisphere of macaque vision-related cortex in a spiking model (Schuecker et al., 2017; Schmidt et al., 2018). Mutual interactions between recurrently connected areas can be expected to profoundly affect their input statistics in terms of rates, spectra, and correlations. Fur-

thermore, we simulate evoked potentials by short thalamic pulses of activity, but sensory cortex receives continuously varying inputs. Ever more detailed models of, for example, the response properties of relay cells in visual thalamus are emerging (Martínez-Cañada et al., 2018), representing naturalistic image or movie stimuli to cortical models similar to ours.

Activity statistics such as distributions of correlations depend on simulation length (Tetzlaff et al., 2008). Here, we consider 5 s simulations, but experimental recordings are often longer (for example, Pan et al., 2013; Chu et al., 2014b,a). Future work can address how greater simulation durations affect the activity statistics, and their convergence across time.

In terms of signal predictions, the tool LFPy¹² embedded in the presently used hybrid scheme (Hagen et al., 2016a), facilitates the calculation of current dipole moments of individual neurons and associated contributions to electroencephalographic (EEG) signals and magnetoencephalographic (MEG) signals as recorded on the surface of the head (Hämäläinen et al., 1993; Nunez and Srinivasan, 2006; Hagen et al., 2018). Forward-model predictions of macroscopic signals like EEG and MEG are thus a tempting proposition, in particular under the consideration of mutual interactions between areas. Among other applications, this could provide an avenue towards a mechanistic model and understanding of visually evoked potentials (Sokol, 1976).

Significance of work

The present work represents a stepping-stone for understanding experimental data obtained by multi-electrode arrays that cover several square millimeters of cortical space. While the model description is highly reduced, it simultaneously accounts for spiking activity and LFPs and thereby enables a multi-scale comparison with corresponding experimental data. At the same time, its simplicity makes mathematical analysis in terms of mean-field and neural field theory viable (Bos et al., 2016; Senk et al., 2018d). Our hope is that the model facilitates a more principal understanding of the dependence of spike correlations on distance, spatially coherent and correlated LFPs, spike-LFP relationships, and emergent spatiotemporal patterns such as waves. The chapter describes not only a particular network model but a fully digitized “integrative loop”. We therefore envision the model as a starting point and building block for future work iteratively modifying parameters and adding further constraints to generate predictions for the activity of specific brain areas.

¹² <https://lfp.readthedocs.io>

DISCUSSION

This thesis contributes to understanding cortical networks on the mesoscale by constructing and analyzing spatially structured spiking neuronal network models. Chapter 6 presents a full-scale multi-layer network model that covers a similar area as multi-electrode arrays used today in experimental recordings, that is, $4 \times 4 \text{ mm}^2$. The preceding chapters address and advance relevant aspects of such large-scale model development, namely, the theory of neuronal networks, simulation technology, visual data analysis, and a workflow to coordinate complex scientific studies. Each chapter already contains a detailed discussion on its respective topic. Here, we collocate the main results in Section 7.1, and discuss in Section 7.2 their impact on the whole integrative loop illustrated in Figure 1.1 as well as for the field of Computational Neuroscience.

7.1 CONCLUSIONS

Workflow

Chapter 2 discusses the need for and the challenges of large, interdisciplinary, and collaborative research endeavors in Computational Neuroscience. We argue for the integration of a complex workflow involving a number of steps and tools into a common platform to facilitate its iterative construction, comprehension, and repeatable execution. The HBP Collaboratory¹ is one such web-based platform currently under development. On this collaboration platform, we implement a prototype workflow addressing the comparability of model data from simulations of the full-scale microcircuit model by Potjans and Diesmann (2014) using either the simulation software NEST or the neuromorphic hardware SpiNNaker as back end. Based on this workflow implementation, we identify, discuss and overcome difficulties like the inclusion of multiple simulation and analysis steps. These steps rely on community software tools and partially require access to high-performance computing facilities.

Simulation

Chapter 3 demonstrates the usability of SpiNNaker for large-scale neuronal network simulations with short neurobiological time scales, and compares its performance quantitatively in terms of accuracy, time-to-solution, and energy-to-solution with that of NEST running on an HPC cluster. The comparison is based on the microcircuit by Potjans and Diesmann (2014) which is with about 0.3 billion synapses the largest model simulated on SpiNNaker to date. To obtain an accuracy similar to that of NEST with 0.1 ms time steps, SpiNNaker requires a slowdown of around

¹ <https://collab.humanbrainproject.eu>

20 times real time. The runtime for NEST saturates around 3 times real time using hybrid parallelization with MPI and multi-threading. However, achieving this runtime comes at the cost of increased power and energy consumption. The lowest total energy consumption for NEST is reached at around 144 parallel threads and 4.6 times slowdown. At this setting, NEST and SpiNNaker have a comparable energy consumption per synaptic event. Our results widen the application domain of SpiNNaker and help guide its development, showing that further optimizations are necessary to enable real-time simulation of large biological neural networks.

Data analysis

Chapter 4 develops 2D and 3D visualization concepts for the interactive visual analysis of simulated spiking activity of spatially organized neuronal network models. These concepts, or views, aim to facilitate the rapid validation of simulation results and the exploration of spatiotemporally resolved data prior to in-depth quantitative analyses. As reference implementation for these concepts we develop the interactive, web-based and open-source tool VIOLA (VISualization of Layer Activity). The tool is designed around key principles for information visualization, that is, the information-seeking mantra (Shneiderman, 1996) and the paradigm of coordinated multiple views (Wang Baldonado et al., 2000). We construct and simulate a layered spiking network model covering a patch of $4 \times 4 \text{ mm}^2$. The model has an excitatory and an inhibitory neuron population connected according to a distance-dependent profile. Visual analysis of the spike data with VIOLA enables us to examine a perturbation of ongoing network activity caused by a temporally and spatially confined stimulus. As an example, the duration and the spatial spread of the event are quickly assessed with the help of multiple simultaneously displayed views. Since online platforms like the HBP Collaboratory require browser-compatible tools for integrated simulation-analysis workflows, with VIOLA we provide a first prototype for the interactive visual data analysis.

Theory

Chapter 5 identifies a mechanism underlying the generation of spatiotemporal patterns such as periodic traveling waves in the spiking activity of LIF-neuron networks. We employ mean-field theory (Brunel and Hakim, 1999) to rigorously map a spiking network model with distance-dependent connection probabilities and constant transmission delay to a neural-field model. Our approach combines a conceptually similar linearization in Kriener et al. (2014b) and an analytical expression for the transfer function in the presence of colored synaptic noise (Schuecker et al., 2015). This transfer function resembles the one of a simple first-order low pass filter in the fluctuation-driven regime. A parameter mapping based on this insight yields quantitative correspondence between a spiking network model and the corresponding neural-field model. The resulting analytically tractable effective rate model depends on the dynamical working point of the spiking network that is characterized by both the mean and the variance of the synaptic input. Using bifurcation theory, in particular linear Turing instability analysis (Coombes, 2005), we derive analytical conditions for pattern formation. These conditions depend

only on the delay (delay-induced fast global oscillations, see Brunel and Hakim, 1999) and on general characteristics of the effective connectivity profile. Spatial oscillations can occur if the most unstable mode has a finite non-zero wave number. We find that nontrivial patterns such as traveling waves require an effective connectivity profile of two populations (excitatory and inhibitory), and hence relate anatomically measurable connectivity structure and observable patterns in spiking activity. For a 1D ring network of an excitatory and inhibitory neuron population with boxcar-shaped connectivity profiles, we demonstrate quantitative agreement of analytical predictions obtained with the neural-field model and numerical simulations of nonlinear rate-unit networks (Hahne et al., 2017) and of networks composed of LIF-model neurons.

Model building

Chapter 6 presents a full-scale spiking neuronal network model covering $4 \times 4 \text{ mm}^2$ of cortical surface and extending across four cortical layers. With an excitatory and an inhibitory neuron population in each layer, the total number of neurons and synapses in the model amounts to $\sim 1.2 \cdot 10^6$ and $\sim 5.5 \cdot 10^9$, respectively. We construct this model by laterally extending the full-scale microcircuit model of 1 mm^2 by Potjans and Diesmann (2014) to cover a similar cortical surface area as a 10×10 Utah multi-electrode array commonly used for electrophysiological measurements in vivo in different cortical areas and species. Our upscaling to the mesoscale preserves realistic neuron and synapse densities and introduces a distance dependency of connection probabilities and delays. The spiking activity of the upscaled model is similar to the original model with asynchronous and irregular spiking activity for the different neuron populations, distributed firing rates across neurons, spike trains with variability in agreement with observed activity in sensory cortex, weak pairwise spike-train correlations, and population firing rate spectra with peaks in the low-gamma (40 – 80 Hz) and high-gamma range (200 – 300 Hz). Due to the partial lack of detailed experimental data on distance-dependent connectivity that leaves the model underconstrained, we test different parameter combinations around the stable network state. We show, for example, that a strong excitatory external drive results in high synchrony, that conduction delays strongly affect the formation of temporal oscillations, and that broad inhibition causes spatial instabilities. We further model thalamocortical input and demonstrate that thalamic pulses evoke spatially spreading activity with radial propagation speeds comparable to experimentally evoked responses to, for instance, a brief flash stimulus to a part of the visual field. The model also accounts for LFP predictions computed from the spiking activity (Hagen et al., 2016b), and achieves to reconcile the experimental observations of weak spike-train correlations (Ecker et al., 2010) and strong and distance-dependent correlations of LFP signals (Nauhaus et al., 2009). Enhanced spatial coherence in the low-gamma band around 50 Hz may explain the recent experimental report of an apparent band-pass filter effect in the spatial reach of the LFP (Dubey and Ray, 2016). This band-pass effect is neither explained by the frequency-independent volume conduction nor by cable properties of the neurons which effectively act as a low-pass filter (Lindén et al., 2010; Łęski et al.,

2013). Instead, the effect primarily results from correlations induced by network interactions.

7.2 OUTLOOK

In this thesis, the endeavor to develop cortical network models is separated into different components that are dealt with in isolation. We improve upon these components and establish a digitized workflow for the scientific loop of model construction and validation. Aiming for a full-scale, multi-layer spiking point-neuron network model representative of a cortical area on the mesoscale, we here expand on simplifying model assumptions made in this work and discuss next steps towards the intended level of description.

The workflow implementation into the HBP Collaboratory (Chapter 2) and the performance comparison of NEST and SpiNNaker (Chapter 3) both use the full-scale multi-layer microcircuit model by Potjans and Diesmann (2014). This microcircuit with about 80,000 neurons and 0.3 billion synapses serves as a prototype for even larger network models and corresponding simulations on HPC systems (for scaling performance, see the benchmark simulations in Section 3.3.5). Simulations of mesoscopic cortical networks can be achieved with state-of-the-art simulation technologies by simply adding hardware resources. Larger networks are less densely connected; increasing the number of neurons only entails a limited increase in the number of synapses per neuron. The mesoscopic model (Chapter 6) is an upscaled version of the microcircuit model and comprises about 1.2 million neurons and 5.5 billion synapses. This model can already be simulated routinely with NEST to test parameter combinations as shown in Section 6.3.2. A model of this size will be the next challenge for SpiNNaker but on the basis of the advancements of the SpiNNaker system, described in Section 3.3.2, it is presumably straight forward to implement by increasing the amount of boards used. In terms of memory consumption, current simulation technologies already facilitate large-scale network simulations with high accuracy. However, more work is needed to speed up these simulations to real time and even further.

The original microcircuit model does not incorporate a spatial structure. However, distance dependency of connections and conduction delays becomes relevant on the mesoscale (Voges et al., 2010). To investigate such distance dependencies and spatiotemporal patterns in network activity, we use simplified two-population networks (apart from Chapter 6). These networks comprise one excitatory and one inhibitory spiking neuron population, either in 2D (Chapter 4) for the visual analysis of layer activity, or in 1D (Chapter 5) for deriving a corresponding mean-field theory. The visualization concepts developed in Chapter 4 apply likewise to more neuronal populations or layers than exemplified with the reference implementation VIOLA. A model network that explicitly incorporates multiple cortical layers will make full use of the features of the tool to visualize 3D activity. The amount of data shown and its graphical representation can be interactively controlled for optimizing the visualization with respect to the given data set.

Extending the mean-field description of spatially organized spiking neuronal 1D-networks (Chapter 5) to 2D, will render the wave number a vector of wave numbers. Bifurcations to periodic patterns in time and space can then be constructed

(Coombes, 2005). Besides, more neuronal populations than the two considered in the study can be included directly. So far, the theory uses a boxcar-shaped connection profile as an example, but it is feasible to replace this profile by a biologically more realistic Gaussian or exponential shape. Since the theory is currently developed for constant delays, a next step is to introduce distance-dependent conduction delays (Hutt et al., 2003). A mean-field model with one-to-one correspondence to a mesoscopic full-scale spiking neuronal network model could not only be validated against network simulation results but also help to constrain model parameters.

Detailed experimental connectivity data for specific species and cortical areas are still partially unavailable, and the theory will generate predictions in addition to already doable parameter scans with the full model (Section 6.3.2). So far, the mesoscopic model distinguishes only between excitatory and inhibitory neurons with respect to the spatial width of connections, but more detailed profiles are conceivable (see Reimann et al., 2017 for rat somatosensory cortex). In Section 5.2.6 we establish a parameter mapping between mean-field models and analytically tractable neural-field models. This mapping can further contribute to transfer results from neural-field literature across mean-field approaches to spiking neuronal networks, that in turn allow for comparison with experimental activity data.

The integrative loop for model development shown in Figure 1.1 is a simplified representation and should in the ideal case exhibit even more connections and sub-loops. For instance, network simulations and predictions from reduced population dynamics generate hypotheses that influence the design of experiments. Also, the loop component data analysis is considered in this thesis primarily for simulated data but applies similarly to experimental data. We envision a direct comparison between experimental and simulated activity data that can be realized best by applying the same data analysis routine to both. VIOLA (Chapter 4) for visual data inspection and, for example, Elephant² for statistical data analysis are tools that can account for both data origins. In general, the community of scientists using certain software tools grows if the tools are flexible, reliable, and user-friendly. An active community in turn supports the development, maintenance, and correctness of the tools. In this thesis, we employ the same simulation software NEST for spiking LIF neurons and continuous rate-based units (Chapter 5). Similarly, the same PyNN implementation is used with NEST or SpiNNaker as simulator back end (Chapters 2 and 3). To provide access to different tools, promote collaborations, and keep track of the status and all steps of a complex workflow, Chapter 2 argues for the workflow integration into a common software platform like the HBP Collaboratory. The web-based visualization tool VIOLA (Chapter 4) is easily integrated into such a platform.

The endeavor to understand the brain and in particular the cerebral cortex brings together scientists from biology, physics, mathematics, and computer science. Each of these classical disciplines alone does not have sufficient means and methods to treat this intricate multi-scale natural system in detail without losing sight of the big picture. Combined in the field of Computational Neuroscience and with recently emerging opportunities of digitization to communicate, transfer knowledge and data, and collaborate, novel interdisciplinary research questions can be tackled

² <http://elephant.readthedocs.io>

that were unimaginable some decades ago. The development of cortical network models is one such challenge. Cortical network models serve the purpose to incrementally integrate experimental data from different sources and different scales and relate them to each other. Modeling allows us to assess consistency and detect gaps in the current knowledge. Direct simulations at the resolution of cells and synapses help us to validate mean-field theory and uncover discrepancies between model predictions and experimental activity data. This thesis develops a mesoscopic, multi-layer, full-scale spiking network model to be iteratively refined along the progress of experimental techniques, theoretical methods, and simulation technology. From the level of description considered here, that is, spatially structured spiking point-neuron networks representing the cortical mesoscale, research may continue to incorporate more biological detail (Markram et al., 2015) or to extend the model from the current patch of $4 \times 4 \text{ mm}^2$ to one full or even multiple cortical areas (Schmidt et al., 2018) towards brain-scale networks (Diesmann, 2013). Future work could make detailed predictions for a specific cortical area, for example, motor cortex studied intensely in our laboratory (Riehle et al., 2013; Denker et al., 2018). Further development of forward-modeling schemes would expand model predictions from spiking activity to other experimentally observable signals such as LFP (see Section 6.3.3), EEG, MEG, or ECoG (Buzsáki et al., 2012; Einevoll et al., 2013a; Hagen et al., 2016a). Similar to the microcircuit model (Potjans and Diesmann, 2014) this work is largely based on, it is our hope that this mesoscopic network model evolves to a platform facilitating further research on questions of network function and learning (Seger and Miller, 2010).

BIBLIOGRAPHY

- Abbott, L. F. and Nelson, S. B. (2000), Synaptic plasticity: taming the beast, *Nat. Neurosci.*, 3, 1178–1183, doi: 10.1038/81453
- Abeles, M. (1982), Local Cortical Circuits: An Electrophysiological Study, Studies of Brain Function (Springer-Verlag, Berlin, Heidelberg, New York)
- Abeles, M. (1991), Corticonics: Neural Circuits of the Cerebral Cortex (Cambridge University Press, Cambridge), 1st edition
- Abeles, M. and Gerstein, G. L. (1988), Detecting Spatiotemporal Firing Patterns Among Simultaneously Recorded Single Neurons, *Journal of Neurophysiology*, 60, 3, 909–924, doi: 10.1152/jn.1988.60.3.909
- Abramowitz, M. and Stegun, I. A. (1974), Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables (Dover Publications, New York), doi: 10.1119/1.15378
- Adams, J. (1977), Technical considerations on the use of horseradish peroxidase as a neuronal marker, *Neuroscience*, 2, 1, 141–145, doi: 10.1016/0306-4522(77)90074-4
- Aertsen, A., Gerstein, G., Habib, M., and Palm, G. (1989), Dynamics of neuronal firing correlation: modulation of "effective connectivity", *Journal of Neurophysiology*, 61, 5, 900–17, doi: 10.1152/jn.1989.61.5.900
- Ahmed, M. R. and Sujatha, B. (2015), A review on methods, issues and challenges in neuromorphic engineering, in 2015 International Conference on Communications and Signal Processing (ICCSP) (IEEE), doi: 10.1109/iccsp.2015.7322626
- Akopyan, F., Sawada, J., Cassidy, A., Alvarez-Icaza, R., Arthur, J., Merolla, P., et al. (2015), TrueNorth: Design and tool flow of a 65 mW 1 million neuron programmable neurosynaptic chip, *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 34, 10, 1537–1557, doi: 10.1109/tcad.2015.2474396
- Albright, T. D., Jessell, T. M., Kandel, E. R., and Posner, M. I. (2000), Neural science: A century of progress and the mysteries that remain, *Neuron*, 25, 1, S1–S55, doi: 10.1016/s0896-6273(00)80912-5
- Allen, E. A., Erhardt, E. B., and Calhoun, V. D. (2012), Data visualization in the neurosciences: Overcoming the curse of dimensionality, *Neuron*, 74, 4, 603–608, doi: 10.1016/j.neuron.2012.05.001
- Amari, S.-I. (1977), Dynamics of pattern formation in lateral-inhibition type neural fields, *Biol. Cybern.*, 27, 2, 77–87, doi: 10.1007/bf00337259
- Amirikian, B. (2005), A phenomenological theory of spatially structured local synaptic connectivity, *PLOS Comput. Biol.*, 1, 1, e11, doi: 10.1371/journal.pcbi.0010011

- Amit, D. J. and Brunel, N. (1997a), Dynamics of a recurrent network of spiking neurons before and following learning, *Network: Comp. Neural Sys.*, 8, 4, 373–404, doi: 10.1088/0954-898x_8_4_003
- Amit, D. J. and Brunel, N. (1997b), Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex, *Cereb. Cortex*, 7, 237–252, doi: 10.1093/cercor/7.3.237
- Andersen, P., Silfvenius, H., Sundberg, S. H., Sveen, O., and Wigström, H. (1978), Functional characteristics of unmyelinated fibres in the hippocampal cortex, *Brain Res.*, 144, 1, 11–18, doi: 10.1016/0006-8993(78)90431-6
- Anderson, P. A., Olavarria, J., and Sluyters, R. C. V. (1988), The overall pattern of ocular dominance bands in cat visual cortex, *Journal of Neuroscience*, 8, 6, 2183–2200, doi: 10.1523/JNEUROSCI.08-06-02183.1988
- Antolík, J. and Davison, A. P. (2018), Arkheia: Data management and communication for open computational neuroscience, *Front. Neuroinformatics*, 12, 6, doi: 10.3389/fninf.2018.00006
- Appelt, W. (1999), WWW based collaboration with the BSCW system, in SOFSEM'99: Theory and Practice of Informatics (Springer Berlin Heidelberg), 66–78, doi: 10.1007/3-540-47849-3_4
- Ascoli, G. A., Donohue, D. E., and Halavi, M. (2007), NeuroMorpho.org: A central resource for neuronal morphologies, *Journal of Neuroscience*, 27, 35, 9247–9251, doi: 10.1523/JNEUROSCI.2055-07.2007
- Atay, F. M. and Hutt, A. (2005), Stability and bifurcations in neural fields with finite propagation speed and general connectivity, *SIAM Journal on Applied Mathematics*, 65, 2, 664–666, doi: 10.1137/S0036139903430884
- Atay, F. M. and Hutt, A. (2006), Neural fields with distributed transmission speeds and long-range feedback delays, *SIAM Journal on Applied Dynamical Systems*, 5, 4, 670–698, doi: 10.1137/050629367
- Attwell, D. and Laughlin, S. B. (2001), An energy budget for signaling in the grey matter of the brain, *Journal of Cerebral Blood Flow & Metabolism*, 21, 1133–1145, doi: 10.1097/00004647-200110000-00001
- Azevedo, F. A. C., Carvalho, L. R. B., Grinberg, L. T., Farfel, J. M., Ferretti, R. E. L., Leite, R. E. P., et al. (2009), Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain, *J. Comp. Neurol.*, 513, 5, 532–541, doi: 10.1002/cne.21974
- Barbas, H. and García-Cabezas, M. Á. (2015), Motor cortex layer 4: less is more, *Trends Neurosci.*, 38, 5, 259–261, doi: 10.1016/j.tins.2015.03.005
- Beaulieu, C. and Colonnier, M. (1983), The number of neurons in the different laminae of the binocular and monocular regions of area 17 in the cat, *J. Comp. Neurol.*, 217, 3, 337–344, doi: 10.1002/cne.902170308

- Bekolay, T., Bergstra, J., Hunsberger, E., DeWolf, T., Stewart, T. C., Rasmussen, D., et al. (2013), Nengo: a python tool for building large-scale functional brain models, *Front. Neuroinformatics*, 7, doi: 10.3389/fninf.2013.00048
- Benjamin, B., Gao, P., McQuinn, E., Choudhary, S., Chandrasekaran, A., Bussat, J., et al. (2014), Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations, *Proc. IEEE*, 102, 5, 699–716, doi: 10.1109/JPROC.2014.2313565
- Benna, M. K. and Fusi, S. (2016), Computational principles of synaptic memory consolidation, *Nat. Neurosci.*, 19, 12, 1697–1706, doi: 10.1038/nn.4401
- Benucci, A., Frazor, R. A., and Carandini, M. (2007), Standing waves and traveling waves distinguish two circuits in visual cortex., *Neuron*, 55, 1, 103–117, doi: 10.1016/j.neuron.2007.06.017
- Berens, P., Keliris, G. A., Ecker, A. S., Logothetis, N. K., and Tolias, A. S. (2008), Comparing the feature selectivity of the gamma-band of the local field potential and the underlying spiking activity in primate visual cortex, *Front. Syst. Neurosci.*, 2, 2, 1–9, doi: 10.3389/neuro.06.002.2008
- Berg-Johnsen, J. and Langmoen, I. A. (1992), Temperature sensitivity of thin unmyelinated fibers in rat hippocampal cortex, *Brain Res.*, 576, 2, 319–321, doi: 10.1016/0006-8993(92)90696-7
- Berning, M., Boergens, K. M., and Helmstaedter, M. (2015), SegEM: Efficient image analysis for high-resolution connectomics, *Neuron*, 87, 6, 1193–1206, doi: 10.1016/j.neuron.2015.09.003
- Binzegger, T., Douglas, R. J., and Martin, K. A. C. (2004), A quantitative map of the circuit of cat primary visual cortex, *J. Neurosci.*, 39, 24, 8441–8453, doi: 10.1523/JNEUROSCI.1400-04.2004
- Binzegger, T., Douglas, R. J., and Martin, K. A. C. (2007), Stereotypical bouton clustering of individual neurons in cat primary visual cortex., *Journal of Neuroscience*, 27, 45, 12242–12254, doi: 10.1523/JNEUROSCI.3753-07.2007
- Binzegger, T., Douglas, R. J., and Martin, K. A. C. (2009), Topology and dynamics of the canonical circuit of cat V1, *Neural Networks*, 22, 8, 1071–1078, doi: 10.1016/j.neunet.2009.07.011
- Bojak, I. and Liley, D. T. J. (2010), Axonal velocity distributions in neural field equations, *PLOS Comput. Biol.*, 6, 1, e1000653, doi: 10.1371/journal.pcbi.1000653
- Börgers, C. and Kopell, N. (2003), Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity, *Neural Comput.*, 15, 3, 509–538, doi: 10.1162/089976603321192059
- Börgers, C. and Kopell, N. (2005), Effects of noisy drive on rhythms in networks of excitatory and inhibitory neurons, *Neural Comput.*, 17, 3, 557–608, doi: 10.1162/0899766053019908

- Bos, H., Diesmann, M., and Helias, M. (2016), Identifying anatomical origins of coexisting oscillations in the cortical microcircuit, *PLOS Comput. Biol.*, 12, 10, e1005132, doi: 10.1371/journal.pcbi.1005132
- Bosking, W. H., Zhang, Y., Schofield, B., and Fitzpatrick, D. (1997), Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex, *Journal of Neuroscience*, 17, 6, 2112–2127, doi: 10.1523/JNEUROSCI.17-06-02112.1997
- Bouchard, K. E., Aimone, J. B., Chun, M., Dean, T., Denker, M., Diesmann, M., et al. (2016), High-performance computing in neuroscience for data-driven discovery, integration, and dissemination, *Neuron*, 92, 3, 628–631, doi: 10.1016/j.neuron.2016.10.035
- Bouchard, K. E., Aimone, J. B., Chun, M., Dean, T., Denker, M., Diesmann, M., et al. (2018), International neuroscience initiatives through the lens of high-performance computing, *Computer*, 51, 4, 50–59, doi: 10.1109/mc.2018.2141039
- Boucsein, C., Nawrot, M., Schnepel, P., and Aertsen, A. (2011), Beyond the cortical column: abundance and physiology of horizontal connections imply a strong role for inputs from the surround, *Front. Neurosci.*, 5, 32, doi: 10.3389/fnins.2011.00032
- Bower, J. M. and Beeman, D. (2007), GENESIS (simulation environment), *Scholarpedia*, 2, 3, 1383
- Braitenberg, V. (2001), Brain size and number of neurons: an exercise in synthetic neuroanatomy, *Journal of Computational Neuroscience*, 10, 1, 71–77, doi: 10.1023/a:1008920127052
- Braitenberg, V. and Schüz, A. (1998), *Cortex: Statistics and Geometry of Neuronal Connectivity* (Springer-Verlag-Verlag, Berlin), 2nd edition
- Brecht, M., Roth, A., and Sakmann, B. (2003), Dynamic receptive fields of reconstructed pyramidal cells in layers 3 and 2 of rat somatosensory barrel cortex, *Journal of Physiology*, 553, 1, 243–265, doi: 10.1113/jphysiol.2003.044222
- Bressloff, P. C. (1996), New mechanism for neural pattern formation, *Phys. Rev. Lett.*, 76, 24, 4644–4647, doi: 10.1103/physrevlett.76.4644
- Bressloff, P. C. (2000), Traveling waves and pulses in a one-dimensional network of excitable integrate-and-fire neurons, *J. Math. Biol.*, 40, 2, 169–198, doi: 10.1007/s002850050008
- Bressloff, P. C. (2012), Spatiotemporal dynamics of continuum neural fields, *Journal of Physics A: Mathematical and Theoretical*, 45, 3, 033001, doi: 10.1088/1751-8113/45/3/033001
- Bressloff, P. C. (2014), *Waves in Neural Media* (Springer New York), doi: 10.1007/978-1-4614-8866-8

- Bressloff, P. C. and Carroll, S. R. (2015), Laminar neural field model of laterally propagating waves of orientation selectivity, *PLOS Comput. Biol.*, 11, 10, e1004545, doi: 10.1371/journal.pcbi.1004545
- Bressloff, P. C. and Coombes, S. (1998), Spike train dynamics underlying pattern formation in integrate-and-fire oscillator networks, *Phys. Rev. Lett.*, 81, 11, 2384–2387, doi: 10.1103/physrevlett.81.2384
- Bressloff, P. C. and Coombes, S. (2000), A dynamical theory of spike train transitions in networks of integrate-and-fire oscillators, *SIAM Journal on Applied Mathematics*, 60, 3, 820–841, doi: 10.1137/s0036139998339643
- Bressloff, P. C., Cowan, J. D., Golubitsky, M., Thomas, P. J., and Wiener, M. C. (2001), Geometric visual hallucinations, euclidean symmetry and the functional architecture of striate cortex, *Phil. Trans. R. Soc. B*, 356, 1407, 299–330, doi: 10.1098/rstb.2000.0769
- Bressloff, P. C. and Kilpatrick, Z. P. (2008), Nonlocal ginzburg-landau equation for cortical pattern formation, *Phys. Rev. E*, 78, 4, doi: 10.1103/physreve.78.041916
- Brette, R., Rudolph, M., Carnevale, T., Hines, M., Beeman, D., Bower, J. M., et al. (2007), Simulation of networks of spiking neurons: A review of tools and strategies, *Journal of Computational Neuroscience*, 23, 3, 349–398, doi: 10.1007/s10827-007-0038-6
- Bringuier, V., Chavane, F., Glaeser, L., and Frégnac, Y. (1999), Horizontal propagation of visual activity in the synaptic integration field of area 17 neurons, *Science*, 283, 5402, 695–699, doi: 10.1126/science.283.5402.695
- Brodmann, K. (1909), Vergleichende Lokalisationslehre der Großhirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues (Johann Ambrosius Barth, Leipzig)
- Brown, S. P. and Hestrin, S. (2009), Intracortical circuits of pyramidal neurons reflect their long-range axonal targets, *Nature*, 457, 7233, 1133–1136, doi: 10.1038/nature07658
- Brüderle, D., Petrovici, M., Vogginger, B., Ehrlich, M., Pfeil, T., Millner, S., et al. (2011), A comprehensive workflow for general-purpose neural modeling with highly configurable neuromorphic hardware systems, *Biol. Cybern.*, 104, 263–296, doi: 10.1007/s00422-011-0435-9
- Brunel, N. (2000), Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons, *Journal of Computational Neuroscience*, 8, 3, 183–208, doi: 10.1023/a:1008925309027
- Brunel, N., Chance, F. S., Fourcaud, N., and Abbott, L. F. (2001), Effects of synaptic noise and filtering on the frequency response of spiking neurons, *Phys. Rev. Lett.*, 86, 10, 2186–2189, doi: 10.1103/physrevlett.86.2186
- Brunel, N. and Hakim, V. (1999), Fast global oscillations in networks of integrate-and-fire neurons with low firing rates, *Neural Comput.*, 11, 7, 1621–1671, doi: 10.1162/089976699300016179

- Budd, J. M. and Kisvárday, Z. F. (2001), Local lateral connectivity of inhibitory clutch cells in layer 4 of cat visual cortex (area 17), *Exp. Brain Res.*, 140, 2, 245–250, doi: 10.1007/s002210100817
- Bureau, I., Shepherd, G. M. G., and Svoboda, K. (2004), Precise development of functional and anatomical columns in the neocortex, *Neuron*, 42, 5, 789–801, doi: 10.1016/j.neuron.2004.05.002
- Buzás, P., Kovács, K., Ferecskó, A. S., Budd, J. M. L., Eysel, U. T., and Kisvárday, Z. F. (2006), Model-based analysis of excitatory lateral connections in the visual cortex, *J. Comp. Neurol.*, 499, 861–881, doi: 10.1002/cne.21134
- Buzsáki, G., Anastassiou, C. A., and Koch, C. (2012), The origin of extracellular fields and currents – EEG, ECoG, LFP and spikes, *Nat. Rev. Neurosci.*, 13, 6, 407–427, doi: 10.1038/nrn3241
- Buzsáki, G. and Mizuseki, K. (2014), The log-dynamic brain: how skewed distributions affect network operations, *Nat. Rev. Neurosci.*, 15, 264–278, doi: 10.1038/nrn3687
- Buzsáki, G. and Wang, X. J. (2012), Mechanisms of gamma oscillations, *Annu. Rev. Neurosci.*, 35, 1, 203–225, doi: 10.1146/annurev-neuro-062111-150444
- Cain, N., Iyer, R., Koch, C., and Mihalas, S. (2016), The computational properties of a simplified cortical column model, *PLOS Comput. Biol.*, 12, 9, e1005045, doi: 10.1371/journal.pcbi.1005045
- Callaway, E. M. and Katz, L. C. (1993), Photostimulation using caged glutamate reveals functional circuitry in living brain slices, *Proceedings of the National Academy of Sciences*, 90, 16, 7661–7665, doi: 10.1073/pnas.90.16.7661
- Camuñas Mesa, L. A. and Quiroga, R. Q. (2013), A detailed and fast model of extracellular recordings, *Neural Comput.*, 25, 5, 1191–1212, doi: 10.1162/NECO_a_00433
- Carnevale, N. T. and Hines, M. L. (2006), *The NEURON Book* (Cambridge University Press, Cambridge)
- Casti, A., Hayot, F., Xiao, Y., and Kaplan, E. (2008), A simple model of retina-LGN transmission, *Journal of Computational Neuroscience*, 24, 2, 235–252, doi: 10.1007/s10827-007-0053-7
- Chaudhuri, R., Knoblauch, K., Gariel, M.-A., Kennedy, H., and Wang, X.-J. (2015), A large-scale circuit mechanism for hierarchical dynamical processing in the primate cortex, *Neuron*, 88, 2, 419–431, doi: 10.1016/j.neuron.2015.09.008
- Chemla, S. and Chavane, F. (2010), Voltage-sensitive dye imaging: Technique review and models, *J. Physiol. (Paris)*, 104, 1-2, 40–50, doi: 10.1016/j.jphysparis.2009.11.009
- Chow, C. and Coombes, S. (2006), Existence and wandering of bumps in a spiking neural network model, *SIAM Journal on Applied Dynamical Systems*, 5, 4, 552–574, doi: 10.1137/060654347

- Chow, C. C., White, J. A., Ritt, J., and Kopell, N. (1998), Frequency Control in Synchronized Networks of Inhibitory Neurons, *Journal of Computational Neuroscience*, 5, 4, 407–420, doi: 10.1023/a:1008889328787
- Chu, C. C. J., Chien, P. F., and Hung, C. P. (2014a), Multi-electrode recordings of ongoing activity and responses to parametric stimuli in macaque V1, *CRCNS.org*, doi: 10.6080/koj1012k
- Chu, C. C. J., Chien, P. F., and Hung, C. P. (2014b), Tuning dissimilarity explains short distance decline of spontaneous spike correlation in macaque V1, *Vision Res.*, 96, 113–132, doi: 10.1016/j.visres.2014.01.008
- Cohen, M. R. and Kohn, A. (2011), Measuring and interpreting neuronal correlations, *Nat. Rev. Neurosci.*, 14, 7, 811–819, doi: 10.1038/nrn.2842
- Compte, A., Brunel, N., Goldman-Rakic, P. S., and Wang, X. J. (2000), Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model, *Cereb. Cortex*, 10, 9, 910–923, doi: 10.1093/cercor/10.9.910
- Connors, B. W. and Gutnick, M. J. (1990), Intrinsic firing patterns of diverse neocortical neurons, *Trends in Neurosciences*, 13, 3, 99–104, doi: 10.1016/0166-2236(90)90185-d
- Connors, B. W. and Long, M. A. (2004), Electrical synapses in the mammalian brain, *Annu. Rev. Neurosci.*, 27, 1, 393–418, doi: 10.1146/annurev.neuro.26.041002.131128
- Contreras, D., Destexhe, A., and Steriade, M. (1997), Intracellular and computational characterization of the intracortical inhibitory control of synchronized thalamic inputs in vivo., *Journal of Neurophysiology*, 78, 1, 335–350, doi: 10.1152/jn.1997.78.1.335
- Coombes, S. (2005), Waves, bumps, and patterns in neural field theories, *Biol. Cybern.*, 93, 91–108, doi: 10.1007/s00422-005-0574-y
- Coombes, S. (2010), Large-scale neural dynamics: Simple and complex, *NeuroImage*, 52, 3, 731–739, doi: 10.1016/j.neuroimage.2010.01.045
- Coombes, S., bei Graben, P., Potthast, R., and Wright, J. (2014), *Neural Fields. Theory and Applications* (Springer-Verlag Berlin Heidelberg)
- Coombes, S., Venkov, N. A., Shiau, L., Bojak, I., Liley, D. T. J., and Laing, C. R. (2007), Modeling electrocortical activity through improved local approximations of integral neural field equations, *Phys. Rev. E*, 76, 5, doi: 10.1103/physreve.76.051901
- Corless, R. M., Gonnet, G. H., Hare, D. E. G., Jeffrey, D. J., and Knuth, D. E. (1996), On the lambert w function, *Advances in Computational Mathematics*, 5, 1, 329–359, doi: 10.1007/BF02124750

- Cremers, D. and Herz, A. V. M. (2002), Traveling waves of excitation in neural field models: Equivalence of rate descriptions and integrate-and-fire dynamics, *Neural Computation*, 14, 7, 1651–1667, doi: 10.1162/08997660260028656
- Crook, S., Bednar, J., Berger, S., Cannon, R., Davison, A., Djurfeldt, M., et al. (2012), Creating, documenting and sharing network models, *Network: Comp. Neural Sys.*, 23, 4, 131–149, doi: 10.3109/0954898x.2012.722743
- Crook, S. M., Ermentrout, G. B., Vanier, M. C., and Bower, J. M. (1997), The role of axonal delay in the synchronization of networks of coupled cortical oscillators, *Journal of Computational Neuroscience*, 4, 2, 161–172, doi: 10.1023/a:1008843412952
- Cross, M. C. and Hohenberg, P. C. (1993), Pattern formation outside of equilibrium, *Reviews of Modern Physics*, 65, 3, 851–1112, doi: 10.1103/revmodphys.65.851
- Cunningham, J. P. and Byron, M. Y. (2014), Dimensionality reduction for large-scale neural recordings, *Nat. Neurosci.*, 17, 11, 1500–1509, doi: 10.1038/nn.3776
- Czanner, G., Grün, S., and Iyengar, S. (2005), Theory of the snowflake plot and its relations to higher-order analysis methods, *Neural Comput.*, 17, 7, 1456–1479, doi: 10.1162/0899766053723041
- Dalva, M. B. and Katz, L. C. (1994), Rearrangements of synaptic connections in visual cortex revealed by laser photostimulation, *Science*, 265, 5169, 255–258, doi: 10.1126/science.7912852
- Davison, A., Brüderle, D., Eppler, J. M., Kremkow, J., Müller, E., Pevcevi, D., et al. (2009), PyNN: a common interface for neuronal network simulators, *Front. Neuroinformatics*, 2, 11, doi: 10.3389/neuro.11.011.2008
- Davison, A., Müller, E., Brüderle, D., and Kremkow, J. (2010), A common language for neuronal networks in software and hardware, *The Neuromorphic Engineer*, doi: 10.2417/1201001.1712
- Davison, A. P. (2009), Trends in programming languages for neuroscience simulations, *Front. Neurosci.*, 3, 3, 374–380, doi: 10.3389/neuro.01.036.2009
- Davison, A. P., Yger, P., Müller, E., Kremkow, J., Brüderle, D., Perrinet, L., et al. (2013), PyNN 0.7.5
- Davison, A. P., Yger, P., Müller, E., Kremkow, J., Brüderle, D., Perrinet, L., et al. (2015), PyNN 0.8.0
- de Kamps, M., Baier, V., Drever, J., Dietz, M., Mösenlechner, L., and van der Felde, F. (2008), The state of MIIND, *Neural Networks*, 21, 8, 1164–1181, doi: 10.1016/j.neunet.2008.07.006
- de Kock, C. P. J. and Sakmann, B. (2009), Spiking in primary somatosensory cortex during natural whisking in awake head-restrained rats is cell-type specific, *Proceedings of the National Academy of Sciences*, 106, 38, 16446–16450, doi: 10.1073/pnas.0904143106

- De Schutter, E. and Van Geit, W. (2009a), Modeling complex neurons, in E. De Schutter, ed., *Computational Modeling Methods for Neuroscientists* (MIT Press, Cambridge, MA), chapter 11, 1st edition, 260–283
- De Schutter, E. and Van Geit, W. (2009b), Modeling complex neurons, in E. De Schutter, ed., *Computational Modeling Methods for Neuroscientists* (MIT Press, Cambridge, MA), chapter 11, 1 edition, 260–283, doi: 10.7551/mitpress/9780262013277.003.0012
- Debanne, D. (2004), Information processing in the axon, *Nat. Rev. Neurosci.*, 5, 304–316, doi: 10.1038/nrn1397
- Deco, G., Jirsa, V. K., Robinson, P. A., Breakspear, M., and Friston, K. (2008), The dynamic brain: From spiking neurons to neural masses and cortical fields, *PLOS Comput. Biol.*, 4, 8, e1000092, doi: 10.1371/journal.pcbi.1000092
- DeFelipe, J. (2011), The evolution of the brain, the human nature of cortical circuits, and intellectual creativity, *Front. Neuroanatomy*, 5, 29, doi: 10.3389/fnana.2011.00029
- Dehghani, N., Peyrache, A., Telenczuk, B., Le Van Quyen, M., Halgren, E., Cash, S. S., et al. (2016), Dynamic balance of excitation and inhibition in human and monkey neocortex, *Scientific Reports*, 6, doi: 10.1038/srep23176
- Denker, M. and Grün, S. (2015), Designing workflows for the reproducible analysis of electrophysiological data, in International Workshop on Brain-Inspired Computing, volume 10087 LNCS, Springer (Springer, Cham), volume 10087 LNCS, 58–72, doi: 10.1007/978-3-319-50862-7_5
- Denker, M., Roux, S., Lindén, H., Diesmann, M., Riehle, A., and Grün, S. (2011), The local field potential reflects surplus spike synchrony, *Cereb. Cortex*, 21, 2681–2695, doi: 10.1093/cercor/bhr040
- Denker, M., Zehl, L., Kilavik, B. E., Diesmann, M., Brochier, T., Riehle, A., et al. (2018), LFP beta amplitude is linked to mesoscopic spatio-temporal phase patterns, *Scientific Reports*, 8, 1, 1–21, doi: 10.1038/s41598-018-22990-7
- 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, *Journal of Neuroscience*, 19, 11, 4595–4608, doi: 10.1523/jneurosci.19-11-04595.1999
- Devert, A. (2014), *Matplotlib Plotting Cookbook* (Packt Publishing Ltd)
- Diaz-Pier, S., Naveau, M., Butz-Ostendorf, M., and Morrison, A. (2016), Automatic generation of connectivity for large-scale neuronal network models through structural plasticity, *Front. Neuroanatomy*, 10, 57, doi: 10.3389/fnana.2016.00057
- Diesmann, M. (2013), The road to brain-scale simulations on K, *BioSupercomputing Newsletter*, 8, 8

- Diesmann, M. and Gewaltig, M.-O. (2002), NEST: An environment for neural systems simulations, in T. Plesser and V. Macho, eds., *Forschung und wissenschaftliches Rechnen, Beiträge zum Heinz-Billing-Preis 2001*, volume 58 of *GWDG-Bericht* (Ges. für Wiss. Datenverarbeitung, Göttingen), 43–70
- Diesmann, M., Gewaltig, M.-O., and Aertsen, A. (1995), SYNOD: an environment for neural systems simulations. Language interface and tutorial, Technical Report GC-AA-/95-3, Weizmann Institute of Science, The Grodetsky Center for Research of Higher Brain Functions, Israel
- Doiron, B., Litwin-Kumar, A., Rosenbaum, R., Ocker, G. K., and Josić, K. (2016), The mechanics of state-dependent neural correlations, *Nat. Neurosci.*, 19, 383–393, doi: 10.1038/nn.4242
- Dougiamas, M. and Taylor, P. (2003), Moodle: Using learning communities to create an open source course management system, in *World Conference on Educational Multimedia, Hypermedia and Telecommunications (EDMEDIA)*
- Douglas, R. J. and Martin, K. A. C. (1991), A functional microcircuit for cat visual cortex, *J. Physiol. (Lond)*, 440, 735–769, doi: 10.1113/jphysiol.1991.sp018733
- Douglas, R. J. and Martin, K. A. C. (2004), Neuronal circuits of the neocortex, *Annu. Rev. Neurosci.*, 27, 419–451, doi: 10.1146/annurev.neuro.27.070203.144152
- Douglas, R. J., Martin, K. A. C., and Whitteridge, D. (1989), A canonical microcircuit for neocortex, *Neural Comput.*, 1, 4, 480–488, doi: 10.1162/neco.1989.1.4.480
- Dubey, A. and Ray, S. (2016), Spatial spread of local field potential is band-pass in the primary visual cortex, *Journal of Neurophysiology*, 116, 4, 1986–1999, doi: 10.1152/jn.00443.2016
- Eccles, J. C., P, F., and Koketsu, K. (1954), Cholinergic and inhibitory synapses in a pathway from motor-axon collaterals to motoneurons, *Journal of Physiology*, 126, 3, 524–562, doi: 10.1113/jphysiol.1954.sp005226
- Ecker, A. S., Berens, P., Keliris, G. A., Bethge, M., and Logothetis, N. K. (2010), Decorrelated neuronal firing in cortical microcircuits, *Science*, 327, 5965, 584–587, doi: 10.1126/science.1179867
- Eglen, S. J., Marwick, B., Halchenko, Y. O., Hanke, M., Sufi, S., Gleeson, P., et al. (2017), Toward standard practices for sharing computer code and programs in neuroscience, *Nat. Neurosci.*, 20, 6, 770–773, doi: 10.1038/nn.4550
- Einevoll, G. T., Franke, F., Hagen, E., Pouzat, C., and Harris, K. D. (2012), Towards reliable spike-train recordings from thousands of neurons with multielectrodes, *Curr. Opin. Neurobiol.*, 22, 1, 11–17, doi: 10.1016/j.conb.2011.10.001
- Einevoll, G. T., Kayser, C., Logothetis, N. K., and Panzeri, S. (2013a), Modelling and analysis of local field potentials for studying the function of cortical circuits., *Nat. Rev. Neurosci.*, 14, 11, 770–785, doi: 10.1038/nrn3599

- Einevoll, G. T., Lindén, H., Tetzlaff, T., Łęski, S., and Pettersen, K. H. (2013b), Local field potentials: Biophysical origin and analysis, in R. Q. Quiroga and S. Panzeri, eds., *Principles of Neural Coding* (CRC Press), 37–59
- Einevoll, G. T., Pettersen, K. H., Devor, A., Ulbert, I., Halgren, E., and Dale, A. M. (2007), Laminar population analysis: Estimating firing rates and evoked synaptic activity from multielectrode recordings in rat barrel cortex, *Journal of Neurophysiology*, 97, 3, 2174–2190, doi: 10.1152/jn.00845.2006
- Epping, W., van den Boogaard, H., Aertsen, A., Eggermont, J., and Johannesma, P. (1984), The neurochrome. An identity preserving representation of activity patterns from neural populations, *Biol. Cybern.*, 50, 4, 235–240
- Eppler, J. M., Helias, M., Muller, E., Diesmann, M., and Gewaltig, M. (2009), PyNEST: a convenient interface to the NEST simulator, *Front. Neuroinformatics*, 2, 12, doi: 10.3389/neuro.11.012.2008
- Eppler, J. M., Pauli, R., Peyser, A., Ippen, T., Morrison, A., Senk, J., et al. (2015), NEST 2.8.0, doi: 10.5281/zenodo.32969
- Erlhagen, W. (1997), Lokalisierte, stationäre Verteilung in neuronalen Feldern: Modellierung experimenteller Befunde zur Planung und Kontrolle zielgerichteter Bewegungen (Verlag Harri Deutsch, Thun, Frankfurt/Main)
- Erlhagen, W. (2003), Internal models for visual perception, *Biol. Cybern.*, 88, 5, 409–417, doi: 10.1007/s00422-002-0387-1
- Erlhagen, W. and Schöner, G. (2002), Dynamic field theory of movement preparation, *Psychological review*, 109, 3, 545
- Ermentrout, B. (1994), Reduction of conductance-based models with slow synapses to neural nets, *Neural Comput.*, 6, 4, 679–695, doi: 10.1162/neco.1994.6.4.679
- Ermentrout, B. (1998), Neural networks as spatio-temporal pattern-forming systems, *Reports on Progress in Physics*, 61, 4, 353–430, doi: 10.1088/0034-4885/61/4/002
- Ermentrout, B. G. and Cowan, J. D. (1980a), Large scale spatially organized activity in neural nets, *SIAM J. Appl. Math.*, 38, 1, 1–21, doi: 10.1137/0138001
- Ermentrout, G. B. and Cowan, J. D. (1979a), A mathematical theory of visual hallucination patterns, *Biol. Cybern.*, 34, 3, 137–150, doi: 10.1007/bf00336965
- Ermentrout, G. B. and Cowan, J. D. (1979b), Temporal oscillations in neuronal nets, *J. Math. Biol.*, 7, 3, 265–280, doi: 10.1007/bf00275728
- Ermentrout, G. B. and Cowan, J. D. (1980b), Secondary bifurcation in neuronal nets, *SIAM Journal on Applied Mathematics*, 39, 2, 323–340, doi: 10.1137/0139028
- Fairchild, M. D. (2013), *Color appearance models* (John Wiley & Sons)
- Felleman, D. J. and Van Essen, D. C. (1991), Distributed hierarchical processing in the primate cerebral cortex, *Cereb. Cortex*, 1, 1–47, doi: 10.1093/cercor/1.1.1

- Ferezou, I., Bolea, S., and Petersen, C. C. (2006), Visualizing the cortical representation of whisker touch: Voltage-sensitive dye imaging in freely moving mice, *Neuron*, 50, 4, 617–629, doi: 10.1016/j.neuron.2006.03.043
- Fermani, F. and Richardson, M. J. E. (2015), Coarse-grained description of the spatio-temporal dynamics of network activity from experimentally verified single-neuron models and connectivity, in 24th Annual Computational Neuroscience Meeting: CNS*2015, BMC Neuroscience 2015, 16(Suppl 1):P206, doi: 10.1186/1471-2202-16-s1-p206
- Fischl, B. and Dale, A. M. (2000), Measuring the thickness of the human cerebral cortex from magnetic resonance images, *Proceedings of the National Academy of Sciences*, 97, 20, 11050–11055, doi: 10.1073/pnas.200033797
- Fohlmeister, C., Gerstner, W., Ritz, R., and van Hemmen, J. L. (1995), Spontaneous excitations in the visual cortex: Stripes, spirals, rings, and collective bursts, *Neural Comput.*, 7, 5, 905–914, doi: 10.1162/neco.1995.7.5.905
- Folias, S. E. and Ermentrout, G. B. (2012), Bifurcations of stationary solutions in an interacting pair of e-i neural fields, *SIAM Journal on Applied Dynamical Systems*, 11, 3, 895–938, doi: 10.1137/110860094
- Forschungszentrum Jülich (2018), Ready for exascale: researchers find algorithm for large-scale brain simulations on next-generation supercomputers, Press release
- Fourcaud, N. and Brunel, N. (2002), Dynamics of the firing probability of noisy integrate-and-fire neurons, *Neural Comput.*, 14, 9, 2057–2110, doi: 10.1162/089976602320264015
- Fourcaud-Trocmé, N. and Brunel, N. (2005), Dynamics of the instantaneous firing rate in response to changes in input statistics, *Journal of Computational Neuroscience*, 18, 3, 311–321, doi: 10.1007/s10827-005-0337-8
- Freedman, D. and Diaconis, P. (1981), On the histogram as a density estimator: L_2 theory, *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 57, 4, 453–476
- Furber, S. (2016a), Large-scale neuromorphic computing systems, *Journal of Neural Engineering*, 13, 5, 051001, doi: 10.1088/1741-2560/13/5/051001
- Furber, S., Galluppi, F., Temple, S., and Plana, L. (2014), The SpiNNaker Project, *Proc. IEEE*, 102, 5, 652–665, doi: 10.1109/JPROC.2014.2304638
- Furber, S., Lester, D., Plana, L., Garside, J., Painkras, E., Temple, S., et al. (2013a), Overview of the spinnaker system architecture, *IEEE Trans. Comp.*, 62, 12, 2454–2467, doi: 10.1109/tc.2012.142
- Furber, S. B. (2016b), Brain-inspired computing, *IET Computers & Digital Techniques*, 10, 6, 299–305, doi: 10.1049/iet-cdt.2015.0171

- Furber, S. B., Lester, D. R., Plana, L. A., Garside, J. D., Painkras, E., Temple, S., et al. (2013b), Overview of the SpiNNaker system architecture, *IEEE Transactions on Computers*, 62, 12, 2454–2467, doi: 10.1109/tc.2012.142
- Galindo, S. E., Toharia, P., Robles, O. D., and Pastor, L. (2016), ViSimpl: Multi-view visual analysis of brain simulation data, *Front. Neuroinformatics*, 10, 44, doi: 10.3389/fninf.2016.00044
- Garaschuk, O., Linn, J., Eilers, J., and Konnerth, A. (2000), Large-scale oscillatory calcium waves in the immature cortex, *Nat. Neurosci.*, 3, 5, 452–459, doi: 10.1038/74823
- Garcia, S., Davison, A. P., Rodgers, C., Yger, P., Mahnoun, Y., Estabanez, L., et al. (2016), Neo 0.4.1
- Garcia, S., Guarino, D., Jaillet, F., Jennings, T., Pröpper, R., Rautenberg, P. L., et al. (2014), Neo: an object model for handling electrophysiology data in multiple formats, *Front. Neuroinformatics*, 8, doi: 10.3389/fninf.2014.00010
- García-Cabezas, M. Á. and Barbas, H. (2014), Area 4 has layer IV in adult primates, *Europ. J. Neurosci.*, 39, 11, 1824–1834, doi: 10.1111/ejn.12585
- Gentet, L., Avermann, M., Matyas, F., Staiger, J. F., and Petersen, C. C. (2010), Membrane potential dynamics of GABAergic neurons in the barrel cortex of behaving mice, *Neuron*, 65, 422–435, doi: 10.1016/j.neuron.2010.01.006
- Gerstein, G. L. and Aertsen, A. M. H. J. (1985), Representation of cooperative firing activity among simultaneously recorded neurons, *Journal of Neurophysiology*, 54, 6, 1513–1528, doi: 10.1152/jn.1985.54.6.1513
- Gerstein, G. L., Perkel, D. H., and Dayhoff, J. E. (1985), Cooperative firing activity in simultaneously recorded populations of neurons: Detection and measurement, *Journal of Neuroscience*, 5, 4, 881–889, doi: 10.1523/jneurosci.05-04-00881.1985
- Gerstein, G. L., Williamns, E. R., Diesmann, M., Grün, S., and Trengove, C. (2012), Detecting synfire chains in parallel spike data, *J. Neurosci. Methods*, 206, 54–64, doi: 10.1016/j.jneumeth.2012.02.003
- Gerstner, W. (2000), Population dynamics of spiking neurons: fast transients, asynchronous states, and locking, *Neural Comput.*, 12, 1, 43–89, doi: 10.1162/089976600300015899
- Gerstner, W. and Naud, R. (2009), How good are neuron models?, *Science*, 326, 5951, 379–380, doi: 10.1126/science.1181936
- Gewaltig, M.-O. and Diesmann, M. (2007), NEST (NEural Simulation Tool), *Scholarpedia*, 2, 4, 1430, doi: 10.4249/scholarpedia.1430
- Giese, M. A. (2012), Dynamic neural field theory for motion perception, volume 469 (Springer Science & Business Media)

- Gilbert, C. D. and Wiesel, T. N. (1989), Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex, *Journal of Neuroscience*, 9, 7, 2432–2442, doi: 10.1523/jneurosci.09-07-02432.1989
- Gleeson, P., Davison, A. P., Silver, R. A., and Ascoli, G. A. (2017), A commitment to open source in neuroscience, *Neuron*, 96, 5, 964–965, doi: 10.1016/j.neuron.2017.10.013
- Gleeson, P., Silver, A., and Cantarelli, M. (2015), Open source brain, in *Encyclopedia of Computational Neuroscience* (Springer New York), 2153–2156, doi: 10.1007/978-1-4614-6675-8_595
- Golomb, D. and Ermentrout, G. B. (2001), Bistability in pulse propagation in networks of excitatory and inhibitory populations, *Phys. Rev. Lett.*, 86, 18, 4179–4182, doi: 10.1103/physrevlett.86.4179
- González-Burgos, G., Barrionuevo, G., and Lewis, D. A. (2000), Horizontal synaptic connections in monkey prefrontal cortex: An *in vitro* electrophysiological study, *Cereb. Cortex*, 10, 1, 82–92, doi: 10.1093/cercor/10.1.82
- Goodman, D. F. M. and Brette, R. (2013), Brian simulator, *Scholarpedia*, 8, 1, 10883, doi: 10.4249/scholarpedia.10883
- Gray, H. (1918), Anatomy of the human body (plate 728, labels removed), Wikimedia Commons, Public Domain (Retrieved from https://commons.wikimedia.org/wiki/File:Lobes_of_the_brain_NL.svg on June 3, 2018)
- Grewe, J., Wachtler, T., and Benda, J. (2011), A bottom-up approach to data annotation in neurophysiology, *Front. Neuroinformatics*, 5, 16, doi: 10.3389/fninf.2011.00016
- Grienberger, C. and Konnerth, A. (2012), Imaging calcium in neurons, *Neuron*, 73, 5, 862–885, doi: 10.1016/j.neuron.2012.02.011
- Griffith, J. S. and Horn, G. (1966), An analysis of spontaneous impulse activity of units in the striate cortex of unrestrained cats, *Journal of Physiology*, 186, 3, 516–534, doi: 10.1113/jphysiol.1966.sp008053
- Grinvald, A., Lieke, E. E., Frostig, R. D., and Hildesheim, R. (1994), Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex, *Journal of Neuroscience*, 14, 5, 2545–2568, doi: 10.1523/JNEUROSCI.14-05-02545.1994
- Grün, S. (1996), Unitary Joint-Events in Multiple-Neuron Spiking Activity: Detection, Significance, and Interpretation, *Reihe Physik*, Band 60 (Verlag Harri Deutsch, Thun, Frankfurt/Main), ISBN 3-8171-1506-7
- Grün, S., Diesmann, M., and Aertsen, A. (2002), ‘Unitary Events’ in multiple single-neuron spiking activity. II. Non-Stationary data, *Neural Comput.*, 14, 1, 81–119, doi: 10.1162/089976602753284464
- Grün, S. and Rotter, S., eds. (2010), *Analysis of Parallel Spike Trains* (Springer)

- Hadjipapas, A., Lowet, E., Roberts, M., Peter, A., and Weerd, P. D. (2015), Parametric variation of gamma frequency and power with luminance contrast: A comparative study of human MEG and monkey LFP and spike responses, *NeuroImage*, 112, 327–340, doi: 10.1016/j.neuroimage.2015.02.062
- Hagen, E., Dahmen, D., Stavrinou, M. L., Lindén, H., Tetzlaff, T., van Albada, S. J., et al. (2016a), Hybrid scheme for modeling local field potentials from point-neuron networks, *Cereb. Cortex*, 26, 12, 4461–4496, doi: 10.1093/cercor/bhw237
- Hagen, E., Dahmen, D., Stavrinou, M. L., Lindén, H., Tetzlaff, T., van Albada, S. J., et al. (2016b), Hybrid scheme for modeling local field potentials from point-neuron networks, *Cereb. Cortex*, bhw237, doi: 10.1093/cercor/bhw237
- Hagen, E., Næss, S., Ness, T. V., and Einevoll, G. T. (2018), Multimodal modeling of neural network activity: computing LFP, ECoG, EEG and MEG signals with LFPy2.0, *bioRxiv*, doi: 10.1101/281717
- Hagen, E., Ness, T. V., Khosrowshahi, A., Sørensen, C., Fyhn, M., Hafting, T., et al. (2015), ViSAPy: A python tool for biophysics-based generation of virtual spiking activity for evaluation of spike-sorting algorithms, *J. Neurosci. Meth.*, 245, 182–204, doi: 10.1016/j.jneumeth.2015.01.029
- Hagen, E., Senk, J., van Albada, S. J., and Diesmann, M. (2016c), Local field potentials in a 4x4mm² multi-layered network model, in 25th Annual Computational Neuroscience Meeting: CNS*2016, BMC Neuroscience 2016, 17(Suppl 1):P167, doi: 10.1186/s12868-016-0283-6
- Hahne, J., Dahmen, D., Schuecker, J., Frommer, A., Bolten, M., Helias, M., et al. (2017), Integration of continuous-time dynamics in a spiking neural network simulator, *Front. Neuroinformatics*, 11, 34, doi: 10.3389/fninf.2017.00034
- Hahne, J., Helias, M., Kunkel, S., Igarashi, J., Kitayama, I., Wylie, B., et al. (2016), Including gap junctions into distributed neuronal network simulations, in *Lecture Notes in Computer Science* (Springer International Publishing), 43–57, doi: 10.1007/978-3-319-50862-7_4
- Haider, B., Duque, A., Hasenstaub, A. R., and McCormick, D. A. (2006), Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition, *Journal of Neuroscience*, 26, 17, 4535–4545, doi: 10.1523/jneurosci.5297-05.2006
- Haken, H. (2000a), Phase locking in the lighthouse model of a neural net with several delay times, *Progress of Theoretical Physics Supplement*, 139, 96–111, doi: 10.1143/ptps.139.96
- Haken, H. (2000b), Quasi-discreted dynamics of a neural net: The lighthouse model, *Discrete Dynamics in Nature and Society*, 4, 3, 187–200, doi: 10.1155/S1026022600000182
- Hämäläinen, M., Haari, R., Ilmoniemi, R. J., Knuutila, J., and Lounasmaa, O. V. (1993), Magnetoencephalography — theory, instrumentation, and application to

- noninvasive studies of the working human brain, *Rev. Mod. Phys.*, 65, 2, 413–496, doi: 10.1103/revmodphys.65.413
- Hansen, C. D. and Johnson, C. R. (2011), *Visualization Handbook* (Academic Press)
- Hanuschkin, A., Kunkel, S., Helias, M., Morrison, A., and Diesmann, M. (2010), A general and efficient method for incorporating precise spike times in globally time-driven simulations, *Front. Neuroinformatics*, 4, 113, doi: 10.3389/fninf.2010.00113
- Hao, Y., Riehle, A., and Brochier, T. G. (2016), Mapping horizontal spread of activity in monkey motor cortex using single pulse microstimulation, *Frontiers in neural circuits*, 10, 104, doi: 10.3389/fncir.2016.00104
- Hardingham, N. R. and Larkman, A. U. (1998), The reliability of excitatory synaptic transmission in slices of rat visual cortex in vitro is temperature dependent, *Journal of Physiology*, 507, 1, 249–256, doi: 10.1111/j.1469-7793.1998.249bu.x
- Harris, K. D. and Shepherd, G. M. G. (2015), The neocortical circuit: themes and variations, *Nat. Neurosci.*, 18, 2, 170–181, doi: 10.1038/nn.3917
- Hasler, J. and Marr, H. B. (2013), Finding a roadmap to achieve large neuromorphic hardware systems, *Front. Neurosci.*, 7, 118, doi: 10.3389/fnins.2013.00118
- Hebb, D. O. (1949), *The organization of behavior: A neuropsychological theory* (John Wiley & Sons, New York)
- Heiberg, T., Kriener, B., Tetzlaff, T., Casti, A., Einevoll, G., and Plesser, H. (2013), Firing-rate models capture essential response dynamics of lgn relay cells, *Journal of Computational Neuroscience*, 35, 3, 359–375, doi: 10.1007/s10827-013-0456-6
- Helias, M., Kunkel, S., Masumoto, G., Igarashi, J., Eppler, J. M., Ishii, S., et al. (2012), Supercomputers ready for use as discovery machines for neuroscience, *Front. Neuroinformatics*, 6, 26, doi: 10.3389/fninf.2012.00026
- Helias, M., Tetzlaff, T., and Diesmann, M. (2013), Echoes in correlated neural systems, *New J Phys.*, 15, 023002, doi: 10.1088/1367-2630/15/2/023002
- Helias, M., Tetzlaff, T., and Diesmann, M. (2014), The correlation structure of local cortical networks intrinsically results from recurrent dynamics, *PLOS Comput. Biol.*, 10, 1, e1003428, doi: 10.1371/journal.pcbi.1003428
- Hellwig, B. (2000), A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex, *Biol. Cybern.*, 2, 82, 111–121, doi: 10.1007/pl00007964
- Helmstaedter, M. (2013), Cellular-resolution connectomics: challenges of dense neural circuit reconstruction, *Nat. Methods*, 10, 6, 501–507, doi: 10.1038/nmeth.2476
- Henker, S., Partzsch, J., and Schüffny, R. (2012), Accuracy evaluation of numerical methods used in state-of-the-art simulators for spiking neural networks, *Journal of Computational Neuroscience*, 32, 2, 309–326, doi: 10.1007/s10827-011-0353-9

- Herculano-Houzel, S. (2009), The human brain in numbers: a linearly scaled-up primate brain, *Front. Hum. Neurosci.*, 3, 31, doi: 10.3389/neuro.09.031.2009
- Herculano-Houzel, S. (2011), Scaling of brain metabolism with a fixed energy budget per neuron: Implications for neuronal activity, plasticity and evolution, *PLoS ONE*, 6, 3, e17514, doi: 10.1371/journal.pone.0017514
- Herculano-Houzel, S. (2012), The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost, *Proceedings of the National Academy of Sciences*, 109, Supplement 1, 10661–10668, doi: 10.1073/pnas.1201895109
- Herculano-Houzel, S., Watson, C., and Paxinos, G. (2013), Distribution of neurons in functional areas of the mouse cerebral cortex reveals quantitatively different cortical zones., *Front. Neuroanatomy*, 7, 35, doi: 10.3389/fnana.2013.00035
- Hill, S. L., Wang, Y., Riachi, I., Schürmann, F., and Markram, H. (2012), Statistical connectivity provides a sufficient foundation for specific functional connectivity in neocortical neural microcircuits, *Proceedings of the National Academy of Sciences*, 109, 42, E2885–E2894, doi: 10.1073/pnas.1202128109
- Hines, M. L., Davison, A. P., and Muller, E. (2009), NEURON and Python, *Front. Neuroinformatics*, 3, 1, doi: 10.3389/neuro.11.001.2009
- Hines, M. L., Morse, T., Migliore, M., Carnevale, N. T., and Hines, M. L. (2004), ModelDB: A database to support computational neuroscience, *Journal of Computational Neuroscience*, 17, 1, 7–11, doi: 10.1023/b:jcn.0000023869.22017.2e
- Hirsch, J. A. and Gilbert, C. D. (1991), Synaptic physiology of horizontal connections in the cat's visual cortex, *Journal of Neuroscience*, 11, 6, 1800–1809, doi: 10.1523/JNEUROSCI.11-06-01800.1991
- Hodgkin, A. L. and Huxley, A. F. (1952), A quantitative description of membrane current and its application to conduction and excitation in nerve, *Journal of Physiology*, 117, 500–544, doi: 10.1113/jphysiol.1952.sp004764
- Holmgren, C., Harkany, T., Svennenfors, B., and Zilberter, Y. (2003), Pyramidal cell communication within local networks in layer 2/3 of rat neocortex, *Journal of Physiology*, 1, 551, 139–153, doi: 10.1113/jphysiol.2003.044784
- Holt, G. R. and Koch, C. (1999), Electrical interactions via the extracellular potential near cell bodies, *Journal of Computational Neuroscience*, 6, 2, 169–184, doi: 10.1023/a:1008832702585
- Hsu, J. (2014), IBM's new brain [news], *IEEE Spectrum*, 51, 10, 17–19, doi: 10.1109/mspec.2014.6905473
- Hubel, D. H. and Wiesel, T. N. (1977), Ferrier lecture - functional architecture of macaque monkey visual cortex, *Proc. R. Soc. Lond. B.*, 198, 1130, 1–59, doi: 10.1098/rspb.1977.0085

- Hutt, A. (2008), Local excitation-lateral inhibition interaction yields oscillatory instabilities in nonlocally interacting systems involving finite propagation delay, *Physics Letters A*, 372, 5, 541–546, doi: 10.1016/j.physleta.2007.08.018
- Hutt, A. and Atay, F. M. (2005), Analysis of nonlocal neural fields for both general and gamma-distributed connectivities, *Physica D: Nonlinear Phenomena*, 203, 1-2, 30–54, doi: 10.1016/j.physd.2005.03.002
- Hutt, A., Bestehorn, M., and Wennekers, T. (2003), Pattern formation in intracortical neuronal fields, *Network: Computation in Neural Systems*, 14, 2, 351–368, doi: 10.1088/0954-898x_14_2_310
- Hutt, A., Hashemi, M., and beim Graben, P. (2015), How to render neural fields more realistic, in *Validating Neuro-Computational Models of Neurological and Psychiatric Disorders* (Springer International Publishing), 141–159, doi: 10.1007/978-3-319-20037-8_6
- Hutt, A. and Rougier, N. (2010), Activity spread and breathers induced by finite transmission speeds in two-dimensional neural fields, *Phys. Rev. E*, 82, 5, doi: 10.1103/physreve.82.055701
- Ian Nauhaus, D. L. R., Laura Busse and Carandini, M. (2012), Robustness of traveling waves in ongoing activity of visual cortex, *Journal of Neuroscience*, 32, 9, 3088–3094, doi: 10.1523/jneurosci.5827-11.2012
- Ippen, T., Eppler, J. M., Plesser, H. E., and Diesmann, M. (2017), Constructing neuronal network models in massively parallel environments, *Frontiers in Neuroinformatics*, 11, 30, doi: 10.3389/fninf.2017.00030
- Izhikevich, E. M. (2003), Simple model of spiking neurons, *IEEE Trans. Neural Netw.*, 14, 6, 1569–1572, doi: 10.1109/tnn.2003.820440
- Izhikevich, E. M. and Edelman, G. M. (2008), Large-scale model of mammalian thalamocortical systems, *Proceedings of the National Academy of Sciences*, 105, 9, 3593–3598, doi: 10.1073/pnas.0712231105
- Janert, P. K. (2010), *Gnuplot in action: understanding data with graphs* (Manning)
- Jia, X., Smith, M. A., and Kohn, A. (2011), Stimulus selectivity and spatial coherence of gamma components of the local field potential, *Journal of Neuroscience*, 31, 25, 9390–9403, doi: 10.1523/jneurosci.0645-11.2011
- Jia, X., Tanabe, S., and Kohn, A. (2013a), Gamma and the coordination of spiking activity in early visual cortex, *Neuron*, 77, 4, 762–774, doi: 10.1016/j.neuron.2012.12.036
- Jia, X., Xing, D., and Kohn, A. (2013b), No consistent relationship between gamma power and peak frequency in macaque primary visual cortex, *Journal of Neuroscience*, 33, 1, 17–25, doi: 10.1523/jneurosci.1687-12.2013
- Jiang, X., Shen, S., Cadwell, C. R., Berens, P., Sinz, F., Ecker, A. S., et al. (2015), Principles of connectivity among morphologically defined cell types in adult neocortex, *Science*, 350, 6264, aac9462–aac9462, doi: 10.1126/science.aac9462

- Jirsa, V. K. and Kelso, J. A. S. (2000), Spatiotemporal pattern formation in neural systems with heterogeneous connection topologies, *Phys. Rev. E*, 62, 6, 8462–8465, doi: 10.1103/physreve.62.8462
- Jolivet, R., Schürmann, F., Berger, T. K., Naud, R., Gerstner, W., and Roth, A. (2008), The quantitative single-neuron modeling competition, *Biol. Cybern.*, 99, 4-5, 417–426, doi: 10.1007/s00422-008-0261-x
- Jordan, J., Ippen, T., Helias, M., Kitayama, I., Sato, M., Igarashi, J., et al. (2018), Extremely scalable spiking neuronal network simulation code: From laptops to exascale computers, *Front. Neuroinformatics*, 12, 2, doi: 10.3389/fninf.2018.00002
- Kajikawa, Y. and Schroeder, C. E. (2011), How local is the local field potential?, *Neuron*, 72, 5, 847–858, doi: 10.1016/j.neuron.2011.09.029
- Kandel, E. R., Markram, H., Matthews, P. M., Yuste, R., and Koch, C. (2013), Neuroscience thinks big (and collaboratively), *Nat. Rev. Neurosci.*, 14, 9, 659–664, doi: 10.1038/nrn3578
- Kandel, E. R., Schwartz, J. H., and Jessel, T. M. (2000), Principles of Neural Science (McGraw-Hill, New York), 4th edition, ISBN 978-0838577011
- Kandel, E. R. and Squire, L. R. (2000), Neuroscience: Breaking down scientific barriers to the study of brain and mind, *Science*, 290, 5494, 1113–1120, doi: 10.1126/science.290.5494.1113
- Kang, Y., Kaneko, T., Ohishi, H., Endo, K., and Araki, T. (1994), Spatiotemporally differential inhibition of pyramidal cells in the cat motor cortex, *Journal of Neurophysiology*, 71, 1, 280–293, doi: 10.1152/jn.1994.71.1.280
- Kasiński, A., Pawłowski, J., and Ponulak, F. (2009), SNN3DViewer-3d visualization tool for spiking neural network analysis, in L. Bolc, J. L. Kulikowski, and K. Wojciechowski, eds., Computer Vision and Graphics, ICCVG 2008 (Springer-Verlag Berlin Heidelberg), 469–476, doi: 10.1007/978-3-642-02345-3
- Kass, R. E., Amari, S.-I., Arai, K., Brown, E. N., Diekman, C. O., Diesmann, M., et al. (2018), Computational neuroscience: Mathematical and statistical perspectives, *Annual Review of Statistics and Its Application*, 5, 1, 183–214, doi: 10.1146/annurev-statistics-041715-033733
- Kasthuri, N., Hayworth, K. J., Berger, D. R., Schalek, R. L., Conchello, J. A., Knowles-Barley, S., et al. (2015), Saturated reconstruction of a volume of neocortex, 162, 3, 648–661, doi: 10.1016/j.cell.2015.06.054
- Katz, B. and Miledi, R. (1965), The effect of temperature on the synaptic delay at the neuromuscular junction., *Journal of Physiology*, 181, 3, 656–670, doi: 10.1113/jphysiol.1965.sp007790
- Kätzel, D., Zemelman, B. V., Buetfering, C., Wölfel, M., and Miesenböck, G. (2011a), The columnar and laminar organization of inhibitory connections to neocortical excitatory cells, *Nat. Neurosci.*, 14, 1, 100–107, doi: 10.1038/nn.2687

- Kätzel, D., Zemelman, B. V., Buetfering, C., Wölfel, M., and Miesenböck, G. (2011b), The columnar and laminar organization of inhibitory connections to neocortical excitatory cells., *Nat. Neurosci.*, 14, 1, 100–107, doi: 10.1038/nn.2687
- Katzner, S., Nauhaus, I., Benucci, A., Bonin, V., Ringach, D. L., and Carandini, M. (2009), Local origin of field potentials in visual cortex, *Neuron*, 61, 1, 35–41, doi: 10.1016/j.neuron.2008.11.016
- Keane, A. and Gong, P. (2015), Propagating waves can explain irregular neural dynamics, *Journal of Neuroscience*, 35, 4, 1591–1605, doi: 10.1523/jneurosci.1669-14.2015
- Kemere, C., Santhanam, G., Byron, M. Y., Afshar, A., Ryu, S. I., Meng, T. H., et al. (2008), Detecting neural-state transitions using hidden markov models for motor cortical prostheses, *Journal of Neurophysiology*, 100, 4, 2441–2452, doi: 10.1152/jn.00924.2007
- Kilpatrick, Z. P. and Bressloff, P. C. (2010), Binocular rivalry in a competitive neural network with synaptic depression, *SIAM Journal on Applied Dynamical Systems*, 9, 4, 1303–1347, doi: 10.1137/100788872
- Kim, U., Bal, T., and McCormick, D. A. (1995), Spindle waves are propagating synchronized oscillations in the ferret LGNd in vitro, *Journal of Neurophysiology*, 74, 3, 1301–1323, doi: 10.1152/jn.1995.74.3.1301
- King, M. A., Louis, P. M., Hunter, B. E., and Walker, D. W. (1989), Biocytin: a versatile anterograde neuroanatomical tract-tracing alternative, *Brain Res.*, 497, 2, 361–367, doi: 10.1016/0006-8993(89)90281-3
- Kistler, W. M. (2000), Stability properties of solitary waves and periodic wave trains in a two-dimensional network of spiking neurons, *Phys. Rev. E*, 62, 6, 8834–8837, doi: 10.1103/physreve.62.8834
- Kistler, W. M., Seitz, R., and van Hemmen, J. (1998), Modeling collective excitations in cortical tissue, *Physica D: Nonlinear Phenomena*, 114, 3-4, 273–295, doi: 10.1016/s0167-2789(97)00195-4
- Kisvárdy, Z. F. and Eysel, U. T. (1992), Cellular organization of reciprocal patchy networks in layer III of cat visual cortex (area 17), *Neuroscience*, 46, 2, 275–286, doi: 10.1016/0306-4522(92)90050-c
- Klein, C., Evrard, H. C., Shapcott, K. A., Haverkamp, S., Logothetis, N. K., and Schmid, M. C. (2016), Cell-targeted optogenetics and electrical microstimulation reveal the primate koniocellular projection to supra-granular visual cortex, *Neuron*, 90, 1, 143–151, doi: 10.1016/j.neuron.2016.02.036
- Knight, J. C. and Furber, S. B. (2016), Synapse-centric mapping of cortical models to the SpiNNaker neuromorphic architecture, *Front. Neurosci.*, 10, 420, doi: 10.3389/fnins.2016.00420
- Knight, J. C., Tully, P. J., Kaplan, B. A., Lansner, A., and Furber, S. B. (2016), Large-scale simulations of plastic neural networks on neuromorphic hardware, *Front. Neuroanatomy*, 10, 37, doi: 10.3389/fnana.2016.00037

- Kobayashi, R., Tsubo, Y., and Shinomoto, S. (2009), Made-to-order spiking neuron model equipped with a multi-timescale adaptive threshold, *Front. Comput. Neurosci.*, 3, 9, doi: 10.3389/neuro.10.009.2009
- Koch, K. W. and Fuster, J. M. (1989), Unit activity in monkey parietal cortex related to haptic perception and temporary memory, *Exp. Brain Res.*, 76, 2, 292–306, doi: 10.1007/bf00247889
- Kriener, B., Enger, H., Tetzlaff, T., Plesser, H. E., Gewaltig, M.-O., and Einevoll, G. T. (2014a), Dynamics of self-sustained asynchronous-irregular activity in random networks of spiking neurons with strong synapses, *Front. Comput. Neurosci.*, 8, 136, doi: 10.3389/fncom.2014.00136
- Kriener, B., Helias, M., Aertsen, A., and Rotter, S. (2009), Correlations in spiking neuronal networks with distance dependent connections, *Journal of Computational Neuroscience*, 27, 2, 177–200, doi: 10.1007/s10827-008-0135-1
- Kriener, B., Helias, M., Rotter, S., Diesmann, M., and Einevoll, G. T. (2014b), How pattern formation in ring networks of excitatory and inhibitory spiking neurons depends on the input current regime, *Front. Comput. Neurosci.*, 7, 187, doi: 10.3389/fncom.2013.00187
- Krishnan, J., Porta Mana, P., Helias, M., Diesmann, M., and Di Napoli, E. A. (2017), Perfect detection of spikes in the linear sub-threshold dynamics of point neurons, *Front. Neuroinformatics*, 11, 75, doi: 10.3389/fninf.2017.00075
- Krüger, J. and Bach, M. (1981), Simultaneous recording with 30 microelectrodes in monkey visual cortex, *Exp. Brain Res.*, 41, 191–194, doi: 10.1007/bf00236609
- Kunkel, S., Morrison, A., Weidel, P., Eppler, J. M., Sinha, A., Schenck, W., et al. (2017), Nest 2.12.0, doi: 10.5281/zenodo.259534
- Kunkel, S., Schmidt, M., Eppler, J. M., Masumoto, G., Igarashi, J., Ishii, S., et al. (2014), Spiking network simulation code for petascale computers, *Front. Neuroinformatics*, 8, 78, doi: 10.3389/fninf.2014.00078
- Kuramoto, Y. (1984), *Chemical Oscillations, Waves, and Turbulence* (Springer Berlin Heidelberg), doi: 10.1007/978-3-642-69689-3
- Laing, C. R. and Chow, C. C. (2001), Stationary bumps in networks of spiking neurons, *Neural Comput.*, 13, 1473–1494, doi: 10.1162/089976601750264974
- Laing, C. R., Troy, W. C., Gutkin, B., and Ermentrout, B. G. (2002), Multiple bumps in a neuronal model of working memory, *SIAM J. Appl. Math.*, 63, 62–97, doi: 10.1137/S0036139901389495
- Lansner, A. and Diesmann, M. (2012), Virtues, pitfalls, and methodology of neuronal network modeling and simulations on supercomputers, in N. L. Novère, ed., *Computational Systems Neurobiology* (Springer, Dordrecht), chapter 10, 283–315, doi: 10.1007/978-94-007-3858-4_10
- Lapicque, L. (1907), Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation, *J. Physiol. Pathol. Gen.*, 9, 620–635

- Laramée, M.-E. and Boire, D. (2015), Visual cortical areas of the mouse: comparison of parcellation and network structure with primates, *Front. Neural Circuits*, 8, 149, doi: 10.3389/fncir.2014.00149
- Larkum, M. E., Zhu, J. J., and Sakmann, B. (2001), Dendritic mechanisms underlying the coupling of the dendritic with the axonal action potential initiation zone of adult rat layer 5 pyramidal neurons, *Journal of Physiology*, 533, 447–466, doi: 10.1111/j.1469-7793.2001.0447a.x
- Legenstein, R. and Maass, W. (2007), Edge of chaos and prediction of computational performance for neural circuit models, *Neural Networks*, 20, 3, 323–334, doi: 10.1016/j.neunet.2007.04.017
- Lennie, P. (2003), The cost of cortical computation, *Curr. Biol.*, 13, 6, 493–497, doi: 10.1016/s0960-9822(03)00135-0
- Łecki, S., Lindén, H., Tetzlaff, T., Pettersen, K. H., and Einevoll, G. T. (2013), Frequency dependence of signal power and spatial reach of the local field potential, *PLOS Comput. Biol.*, 9, e1003137, doi: 10.1371/journal.pcbi.1003137
- Łecki, S., Pettersen, K. H., Tunstall, B., Einevoll, G. T., Gigg, J., and Wójcik, D. K. (2011), Inverse current source density method in two dimensions: Inferring neural activation from multielectrode recordings, *Neuroinformatics*, 9, 4, 401–425, doi: 10.1007/s12021-011-9111-4
- Leung, L. S. (1982), Nonlinear feedback model of neuronal populations in hippocampal CA1 region, *Journal of Neurophysiology*, 47, 5, 845–868, doi: 10.1152/jn.1982.47.5.845
- Levina, A. and Priesemann, V. (2017), Subsampling scaling, *Nature Communications*, 8, 15140, doi: 10.1038/ncomms15140
- Levy, R. B. and Reyes, A. D. (2012), Spatial profile of excitatory and inhibitory synaptic connectivity in mouse primary auditory cortex, *Journal of Neuroscience*, 32, 16, 5609–5619, doi: 10.1523/jneurosci.5158-11.2012
- Lin, J.-W. and Faber, D. S. (2002), Modulation of synaptic delay during synaptic plasticity, *Trends Neurosci.*, 25, 9, 449–455, doi: 10.1016/s0166-2236(02)02212-9
- Lindén, H., Hagen, E., Łecki, S., Norheim, E. S., Pettersen, K. H., and Einevoll, G. T. (2014), LFPy: a tool for biophysical simulation of extracellular potentials generated by detailed model neurons, *Front. Neuroinformatics*, 7, 41, doi: 10.3389/fninf.2013.00041
- Lindén, H., Pettersen, K. H., and Einevoll, G. T. (2010), Intrinsic dendritic filtering gives low-pass power spectra of local field potentials, *Journal of Computational Neuroscience*, 29, 3, 423–444, doi: 10.1007/s10827-010-0245-4
- Lindén, H., Tetzlaff, T., Potjans, T. C., Pettersen, K. H., Grün, S., Diesmann, M., et al. (2011), Modeling the spatial reach of the LFP, *Neuron*, 72, 5, 859–872, doi: 10.1016/j.neuron.2011.11.006

- Lindner, B., Doiron, B., and Longtin, A. (2005), Theory of oscillatory firing induced by spatially correlated noise and delayed inhibitory feedback, *Phys. Rev. E*, 72, 6, 061919, doi: 10.1103/physreve.72.061919
- Lindner, B. and Schimansky-Geier, L. (2001), Transmission of noise coded versus additive signals through a neuronal ensemble, *Phys. Rev. Lett.*, 86, 2934–2937, doi: 10.1103/physrevlett.86.2934
- Litwin-Kumar, A. and Doiron, B. (2012), Slow dynamics and high variability in balanced cortical networks with clustered connections, *Nat. Neurosci.*, 15, 11, 1498–1505, doi: 10.1038/nn.3220
- Livingstone, M. S. and Hubel, D. H. (1984), Specificity of intrinsic connections in primate primary visual cortex, *Journal of Neuroscience*, 4, 11, 2830–2835, doi: 10.1523/jneurosci.04-11-02830.1984
- Lohmann, H. and Rörig, B. (1994), Long-range horizontal connections between supragranular pyramidal cells in the extrastriate visual cortex of the rat, *J. Comp. Neurol.*, 344, 543–558, doi: 10.1002/cne.903440405
- Lorensen, W. E. and Cline, H. E. (1987), Marching cubes: A high resolution 3D surface construction algorithm, in Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques (ACM, New York, NY, USA), SIGGRAPH '87, 163–169, doi: 10.1145/37401.37422
- Lubenov, E. V. and Siapas, A. G. (2009), Hippocampal theta oscillations are traveling waves, *Nature*, 459, 7246, 534–539, doi: 10.1038/nature08010
- Lundqvist, M., Rehn, M., Djurfeldt, M., and Lansner, A. (2006), Attractor dynamics in a modular network model of neocortex, *Network: Comp. Neural Sys.*, 17, 3, 253–276, doi: 10.1080/09548980600774619
- Lytton, W. W., Seidenstein, A. H., Dura-Bernal, S., McDougal, R. A., Schürmann, F., and Hines, M. L. (2016), Simulation neurotechnologies for advancing brain research: parallelizing large networks in NEURON, *Neural computation*, 28, 10, 2063–2090, doi: 10.1162/neco_a_00876
- MacGregor, R. J. (1987), *Neural and Brain Modeling* (Academic Press, San Diego)
- Maier, A., Adams, G. K., Aura, C., and Leopold, D. A. (2010), Distinct superficial and deep laminar domains of activity in the visual cortex during rest and stimulation, *Front. Syst. Neurosci.*, 4, 31, 31, doi: 10.3389/fnsys.2010.00031
- Mainen, Z. F. and Sejnowski, T. J. (1996), Influence of dendritic structure on firing pattern in model neocortical neurons, *Nature*, 382, 6589, 363–366, doi: 10.1038/382363a0
- Margrie, T., Brecht, M., and Sakmann, B. (2002), In vivo, low-resistance, whole-cell recordings from neurons in the anaesthetized and awake mammalian brain, *Pflügers Archiv European Journal of Physiology*, 444, 4, 491–498, doi: 10.1007/s00424-002-0831-z

- Markram, H., Muller, E., Ramaswamy, S., Reimann, M. W., Abdellah, M., Sanchez, C. A., et al. (2015), Reconstruction and simulation of neocortical microcircuitry, *Cell*, 163, 2, 456–492, doi: 10.1016/j.cell.2015.09.029
- Martínez-Cañada, P., Mobarhan, M. H., Halmes, G., Fyhn, M., Morillas, C., Pelayo, F., et al. (2018), Biophysical network modeling of the dLGN circuit: Effects of cortical feedback on spatial response properties of relay cells, *PLOS Comput. Biol.*, 14, 1, e1005930, doi: 10.1371/journal.pcbi.1005930
- Massobrio, P., Tessadori, J., Chiappalone, M., and Ghirardi, M. (2015), In Vitro Studies of neuronal networks and synaptic plasticity in invertebrates and in mammals using multielectrode arrays, *Neural Plasticity*, 2015, 1–18, doi: 10.1155/2015/196195
- Matsuzaki, M., Ellis-Davies, G. C. R., and Kasai, H. (2008), Three-dimensional mapping of unitary synaptic connections by two-photon macro photolysis of caged glutamate, *Journal of Neurophysiology*, 99, 3, 1535–1544, doi: 10.1152/jn.01127.2007
- Mattioni, M., Cohen, U., and Novère, N. L. (2012), Neuronvisio: A graphical user interface with 3d capabilities for NEURON, *Front. Neuroinformatics*, 6, 20, doi: 10.3389/fninf.2012.00020
- Maynard, E. M., Nordhausen, C. T., and Normann, R. A. (1997), The Utah intracortical electrode array: A recording structure for potential brain-computer interfaces, *EEG Clin. Neurophysiol.*, 102, 3, 228–239, doi: 10.1016/S0013-4694(96)95176-0
- Mazzoni, A., Lindén, H., Cuntz, H., Lansner, A., Panzeri, S., and Einevoll, G. T. (2015), Computing the local field potential (LFP) from integrate-and-fire network models, *PLOS Comput. Biol.*, 11, 12, e1004584, doi: 10.1371/journal.pcbi.1004584
- McDonald, C. T. and Burkhalter, A. (1993), Organisation of long-range inhibitory connections within rat visual cortex, *Journal of Neuroscience*, 13, 2, 768–781, doi: 10.1523/jneurosci.13-02-00768.1993
- McLeod, P., Plunkett, K., and Rolls, E. T. (1998), Introduction to Connectionist Modelling of Cognitive Processes (Oxford University Press, Oxford)
- Mead, C. (1990), Neuromorphic electronic systems, *Proc. IEEE*, 78, 10, 1629–1636, doi: 10.1109/5.58356
- Mehring, C., Hehl, U., Kubo, M., Diesmann, M., and Aertsen, A. (2003), Activity dynamics and propagation of synchronous spiking in locally connected random networks, *Biol. Cybern.*, 88, 5, 395–408, doi: 10.1007/s00422-002-0384-4
- Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., et al. (2014), A million spiking-neuron integrated circuit with a scalable communication network and interface, *Science*, 345, 6197, 668–673, doi: 10.1126/science.1254642

- Milekovic, T., Truccolo, W., Grün, S., Riehle, A., and Brochier, T. (2015), Local field potentials in primate motor cortex encode grasp kinetic parameters, *NeuroImage*, 114, 338–355, doi: 10.1016/j.neuroimage.2015.04.008
- Miller, P. (2016), Dynamical systems, attractors, and neural circuits, *F1000Research*, 5, 992, doi: 10.12688/f1000research.7698.1
- Mitzdorf, U. (1985), Current source-density method and application in cat cerebral cortex: Investigation of evoked potentials and EEG phenomena, *Physiol. Rev.*, 65, 1, 37–100, doi: 10.1152/physrev.1985.65.1.37
- Mochizuki, Y., Onaga, T., Shimazaki, H., Shimokawa, T., Tsubo, Y., Kimura, R., et al. (2016), Similarity in neuronal firing regimes across mammalian species, *Journal of Neuroscience*, 36, 21, 5736–5747, doi: 10.1523/JNEUROSCI.0230-16.2016
- Monroe, D. (2014), Neuromorphic computing gets ready for the (really) big time, *Communications of the ACM*, 57, 6, 13–15, doi: 10.1145/2601069
- Montbrío, E., Pazó, D., and Roxin, A. (2015), Macroscopic description for networks of spiking neurons, *Phys. Rev. X*, 5, 021028, doi: 10.1103/PhysRevX.5.021028
- Moradi, S., Qiao, N., Stefanini, F., and Indiveri, G. (2018), A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs), *IEEE Transactions on Biomedical Circuits and Systems*, 12, 1, 106–122, doi: 10.1109/tbcas.2017.2759700
- Morrison, A., Aertsen, A., and Diesmann, M. (2007a), Spike-timing dependent plasticity in balanced random networks, *Neural Comput.*, 19, 1437–1467, doi: 10.1162/neco.2007.19.6.1437
- Morrison, A. and Diesmann, M. (2008), Maintaining causality in discrete time neuronal network simulations, in P. b. Graben, C. Zhou, M. Thiel, and J. Kurths, eds., *Lectures in Supercomputational Neurosciences: Dynamics in Complex Brain Networks* (Springer, Berlin, Heidelberg), 267–278, doi: 10.1007/978-3-540-73159-7_10
- 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, doi: 10.1162/0899766054026648
- Morrison, A., Straube, S., Plesser, H. E., and Diesmann, M. (2007b), Exact sub-threshold integration with continuous spike times in discrete time neural network simulations, *Neural Comput.*, 19, 1, 47–79, doi: 10.1162/neco.2007.19.1.47
- Morrison, A., Straube, S., Plesser, H. E., and Diesmann, M. (2007c), Exact sub-threshold integration with continuous spike times in discrete-time neural network simulations, *Neural Comput.*, 19, 1, 47–79, doi: 10.1162/neco.2007.19.1.47
- Mountcastle, V. B. (1957), Modality and topographic properties of single neurons in cat's somatic sensory cortex, *Journal of Neurophysiology*, 20, 4, 408–434, doi: 10.1152/jn.1957.20.4.408

- Muller, E., Bednar, J. A., Diesmann, M., Gewaltig, M.-O., Hines, M., and Davison, A. P. (2015), Python in neuroscience, *Front. Neuroinformatics*, 9, 11, doi: 10.3389/fninf.2015.00011
- Muller, L., Chavane, F., Reynolds, J., and Sejnowski, T. J. (2018), Cortical travelling waves: mechanisms and computational principles, *Nat. Rev. Neurosci.*, 19, 5, 255–268, doi: 10.1038/nrn.2018.20
- Muller, L. and Destexhe, A. (2012), Propagating waves in thalamus, cortex and the thalamocortical system: Experiments and models, *J. Physiol. (Paris)*, 106, 5-6, 222–238, doi: 10.1016/j.jphysparis.2012.06.005
- Muller, L., Reynaud, A., Chavane, F., and Destexhe, A. (2014), The stimulus-evoked population response in visual cortex of awake monkey is a propagating wave, *Nature Communications*, 5, doi: 10.1038/ncomms4675
- Murakoshi, T., Guo, J.-Z., and Ichinose, T. (1993), Electrophysiological identification of horizontal synaptic connections in rat visual cortex in vitro, *Neuroscience Letters*, 163, 2, 211–214, doi: 10.1016/0304-3940(93)90385-x
- Nauhaus, I., Busse, L., Carandini, M., and Ringach, D. L. (2009), Stimulus contrast modulates functional connectivity in visual cortex, *Nat. Neurosci.*, 12, 70–76, doi: 10.1038/nn.2232
- Nawrocki, R. A., Voyles, R. M., and Shaheen, S. E. (2016), A mini review of neuromorphic architectures and implementations, *IEEE Transactions on Electron Devices*, 63, 10, 3819–3829, doi: 10.1109/ted.2016.2598413
- Nelson, M. J. and Pouget, P. (2010), Do electrode properties create a problem in interpreting local field potential recordings?, *Journal of Neurophysiology*, 103, 5, 2315–2317, doi: 10.1152/jn.00157.2010
- Nelson, M. J., Pouget, P., Nilsen, E. A., Patten, C. D., and Schall, J. D. (2008), Review of signal distortion through metal microelectrode recording circuits and filters, *Journal of Neuroscience Methods*, 169, 1, 141–157, doi: 10.1016/j.jneumeth.2007.12.010
- Ness, T. V., Chintaluri, H. C., Potworowski, J., Łęski, S., Głąbska, H., Wójcik, D. K., et al. (2015), Modelling and analysis of electrical potentials recorded in multi-electrode arrays (MEAs), *Neuroinformatics*, 13, 4, 403–426, doi: 10.1007/s12021-015-9265-6
- Nicholson, C. and Freeman, J. A. (1975), Theory of current source-density analysis and determination of conductivity tensor for anuran cerebellum., *Journal of Neurophysiology*, 2, 38, 356–368, doi: 10.1152/jn.1975.38.2.356
- Nordlie, E., Gewaltig, M.-O., and Plesser, H. E. (2009a), Towards reproducible descriptions of neuronal network models, *PLOS Comput. Biol.*, 5, 8, e1000456, doi: 10.1371/journal.pcbi.1000456
- Nordlie, E. and Plesser, H. E. (2010), Visualizing neuronal network connectivity with connectivity pattern tables, *Front. Neuroinformatics*, 3, 39, doi: 10.3389/neuro.11.039.2009

- Nordlie, E., Tetzlaff, T., and Einevoll, G. T. (2009b), Firing-rate response properties of spiking neurons: beyond the diffusion limit, *Front. Neuroinformatics*, doi: 10.3389/conf.neuro.11.2009.08.118
- Nordlie, E., Tetzlaff, T., and Einevoll, G. T. (2010), Rate dynamics of leaky integrate-and-fire neurons with strong synapses, *Front. Comput. Neurosci.*, 4, 149, doi: 10.3389/fncom.2010.00149
- Nowke, C., Schmidt, M., van Albada, S. J., Eppler, J. M., Bakker, R., Diesmann, M., et al. (2013), VisNEST: Interactive analysis of neural activity data, in IEEE Symposium on Biological Data Visualization 2013 (IEEE, Atlanta, Georgia, USA), 65–72, doi: 10.1109/BioVis.2013.6664348
- Nowke, C., Zielasko, D., Weyers, B., Peyser, A., Hentschel, B., and Kuhlen, T. W. (2015), Integrating visualizations into modeling NEST simulations, *Front. Neuroinformatics*, 9, 29, doi: 10.3389/fninf.2015.00029
- Nunez, P. L. and Srinivasan, R. (2006), *Electric Fields of the Brain, The Neurophysics of EEG* (Oxford University Press, Inc.), 2nd edition
- Ohana, O., Portner, H., and Martin, K. A. C. (2012), Fast recruitment of recurrent inhibition in the cat visual cortex, *PLoS ONE*, 7, 7, e40601, doi: 10.1371/journal.pone.0040601
- Ohbayashi, M., Ohki, K., and Miyashita, Y. (2003), Conversion of working memory to motor sequence in the monkey premotor cortex, *Science*, 301, 5630, 233–236, doi: 10.1126/science.1084884
- Ohki, K. and Reid, R. C. (2007), Specificity and randomness in the visual cortex, *Curr. Opin. Neurobiol.*, 17, 4, 401–407, doi: 10.1016/j.conb.2007.07.007
- Okun, M. and Lampl, I. (2008), Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities, *Nat. Neurosci.*, 11, 5, 535–537, doi: 10.1038/nn.2105
- Olavarria, J. and Sluyters, R. C. V. (1985), Unfolding and flattening the cortex of gyrencephalic brains, *J. Neurosci. Methods*, 15, 3, 191–202, doi: 10.1016/0165-0270(85)90098-6
- Osan, R. and Ermentrout, G. B. (2002), The evolution of synaptically generated waves in one- and two-dimensional domains, *Physica D*, 163, 217–235, doi: 10.1016/S0167-2789(02)00347-0
- Ostojic, S. (2014), Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons, *Nat. Neurosci.*, 17, 594–600, doi: 10.1038/nn.3658
- Ouzounov, D. G., Wang, T., Wang, M., Feng, D. D., Horton, N. G., Cruz-Hernández, J. C., et al. (2017), In vivo three-photon imaging of activity of GCaMP6-labeled neurons deep in intact mouse brain, *Nature Methods*, 14, 4, 388–390, doi: 10.1038/nmeth.4183

- Packer, A. M. and Yuste, R. (2011), Dense, unspecific connectivity of neocortical parvalbumin-positive interneurons: A canonical microcircuit for inhibition?, *Journal of Neuroscience*, 31, 37, 13260–13271, doi: 10.1523/JNEUROSCI.3131-11.2011
- Painkras, E., Plana, L. A., Garside, J., Temple, S., Davidson, S., Pepper, J., et al. (2012), SpiNNaker: A multi-core system-on-chip for massively-parallel neural net simulation, in Proceedings of the IEEE 2012 Custom Integrated Circuits Conference (IEEE), doi: 10.1109/cicc.2012.6330636
- Painkras, E., Plana, L. A., Garside, J., Temple, S., Galluppi, F., Patterson, C., et al. (2013), SpiNNaker: A 1-w 18-core system-on-chip for massively-parallel neural network simulation, *IEEE Journal of Solid-State Circuits*, 48, 8, 1943–1953, doi: 10.1109/jssc.2013.2259038
- Pakkenberg, B. and Gundersen, H. J. G. (1997), Neocortical neuron number in humans: Effect of sex and age, *J. Comp. Neurol.*, 384, 2, 312–320, doi: 10.1002/(sici)1096-9861(19970728)384:2<312::aid-cne10>3.0.co;2-k
- Pan, W.-J., Thompson, G. J., Magnuson, M. E., Jaeger, D., and Keilholz, S. (2013), Infralow LFP correlates to resting-state fMRI BOLD signals, *NeuroImage*, 74, 288–297, doi: 10.1016/j.neuroimage.2013.02.035
- Payne, B. R. and Peters, A., eds. (2002), *The Cat Primary Visual Cortex* (Elsevier), doi: 10.1016/b978-0-12-552104-8.x5000-7
- Perin, R., Berger, T. K., and Markram, H. (2011), A synaptic organizing principle for cortical neuronal groups, *Proceedings of the National Academy of Sciences*, 108, 13, 5419–5424, doi: 10.1073/pnas.1016051108
- Perkel, D. H., Gerstein, G. L., and Moore, G. P. (1967a), Neuronal spike trains and stochastic point processes. I. The single spike train, *Biophys. J.*, 7, 4, 391–418, doi: 10.1016/s0006-3495(67)86596-2
- Perkel, D. H., Gerstein, G. L., and Moore, G. P. (1967b), Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains, *Biophys. J.*, 7, 4, 419–440, doi: 10.1016/s0006-3495(67)86597-4
- Perkel, D. H., Gerstein, G. L., Smith, M. S., and Tatton, W. G. (1975), Nerve-impulse patterns: a quantitative display technique for three neurons, *Brain Res.*, 100, 271–296
- Peters, A. and Feldman, M. L. (1976), The projection of the lateral geniculate nucleus to area 17 of the rat cerebral cortex. I. General description., *J. Neurocytol.*, 5, 1, 63–84
- Petersen, C. C. H. and Sakmann, B. (2001), Functionally independent columns of rat somatosensory barrel cortex revealed with voltage-sensitive dye imaging, *Journal of Neuroscience*, 21, 21, 8435–8446, doi: 10.1523/jneurosci.21-21-08435.2001
- Petterson, K. H., Devor, A., Ulbert, I., Dale, A. M., and Einevoll, G. T. (2006), Current-source density estimation based on inversion of electrostatic forward

- solution: Effects of finite extent of neuronal activity and conductivity discontinuities., *J. Neurosci. Methods*, 154, 1-2, 116–133, doi: 10.1016/j.jneumeth.2005.12.005
- Pettersen, K. H., Hagen, E., and Einevoll, G. T. (2008), Estimation of population firing rates and current source densities from laminar electrode recordings., *Journal of Computational Neuroscience*, 24, 291–313, doi: 10.1007/s10827-007-0056-4
- Peyrache, A., Dehghani, N., Eskandar, E. N., Madsen, J. R., Anderson, W. S., Donoghue, J. A., et al. (2012), Spatiotemporal dynamics of neocortical excitation and inhibition during human sleep, *Proceedings of the National Academy of Sciences*, 109, 5, 1731–1736, doi: 10.1073/pnas.1109895109
- Peyser, A., Sinha, A., Vennemo, S. B., Ippen, T., Jordan, J., Graber, S., et al. (2017), *Nest* 2.14.0, doi: 10.5281/zenodo.882970
- Plesser, H. E. and Diesmann, M. (2009), Simplicity and efficiency of integrate-and-fire neuron models, *Neural Comput.*, 21, 353–359, doi: 10.1162/neco.2008.03-08-731
- Plesser, H. E., Eppler, J. M., Morrison, A., Diesmann, M., and Gewaltig, M.-O. (2007), Efficient parallel simulation of large-scale neuronal networks on clusters of multiprocessor computers, in A.-M. Kermarrec, L. Bougé, and T. Priol, eds., Euro-Par 2007: Parallel Processing, volume 4641 of *Lecture Notes in Computer Science* (Springer-Verlag, Berlin), volume 4641 of *Lecture Notes in Computer Science*, 672–681, doi: 10.1007/978-3-540-74466-5_71
- Plotnikov, D., Blundell, I., Ippen, T., Eppler, J. M., Rumpe, B., and Morrison, A. (2016), NESTML: a modeling language for spiking neurons, Modellierung 2016, Karlsruhe (Germany), 17 Mar 2016 - 19 Mar 2016 (Gesellschaft für Informatik e.V. (GI)), volume P-254 of *Lecture Notes in Informatics (LNI)*, 93–108
- Potjans, T. C. and Diesmann, M. (2014), The cell-type specific cortical microcircuit: Relating structure and activity in a full-scale spiking network model, *Cereb. Cortex*, 24, 3, 785–806, doi: 10.1093/cercor/bhs358
- Potjans, W., Morrison, A., and Diesmann, M. (2010), Enabling functional neural circuit simulations with distributed computing of neuromodulated plasticity, *Front. Comput. Neurosci.*, 4, 141, doi: 10.3389/fncom.2010.00141
- Potworowski, J., Jakuczun, W., Łęski, S., and Wójcik, D. (2012), Kernel current source density method, *NeuralComput*, 24, 2, 541–575, doi: 10.1162/neco_a_00236
- Prut, Y., Vaadia, E., Bergman, H., Haalman, I., Hamutal, S., and Abeles, M. (1998), Spatiotemporal structure of cortical activity: Properties and behavioral relevance, *Journal of Neurophysiology*, 79, 6, 2857–2874, doi: 10.1152/jn.1998.79.6.2857
- Pyle, R. and Rosenbaum, R. (2017), Spatiotemporal dynamics and reliable computations in recurrent spiking neural networks, *Physical Review Letters*, 118, 1, doi: 10.1103/physrevlett.118.018103

- Qiao, N., Mostafa, H., Corradi, F., Osswald, M., Stefanini, F., Sumislawska, D., et al. (2015), A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128k synapses, *Front. Neurosci.*, 9, 141, doi: 10.3389/fnins.2015.00141
- Quiroga, R. Q. (2007), Spike sorting, *Scholarpedia*, 2, 12, 3583, doi: 10.4249/scholarpedia.3583
- Ramachandran, P. and Varoquaux, G. (2011), Mayavi: 3d visualization of scientific data, *Computing in Science & Engineering*, 13, 2, 40–51, doi: 10.1109/mcse.2011.35
- Ramon y Cajal, S. (1899), *Textura del Systema Nervioso del Hombre y de los Vertebrados*. (Texture of the Nervous System of Man and the Vertebrates), volume 1 (New York: Springer)
- Ray, S. and Maunsell, J. H. R. (2010), Differences in gamma frequencies across visual cortex restrict their possible use in computation, *Neuron*, 67, 5, 885–896, doi: 10.1016/j.neuron.2010.08.004
- Ray, S. and Maunsell, J. H. R. (2011a), Different Origins of Gamma Rhythm and High-Gamma Activity in Macaque Visual Cortex, *PLOS Comput. Biol.*, 9, 4, e1000610, doi: 10.1371/journal.pbio.1000610
- Ray, S. and Maunsell, J. H. R. (2011b), Network rhythms influence the relationship between spike-triggered local field potential and functional connectivity, *Journal of Neuroscience*, 31, 35, 12674–12682, doi: 10.1523/jneurosci.1856-11.2011
- Rees, C. L., Moradi, K., and Ascoli, G. A. (2016), Weighing the evidence in Peters' rule: Does neuronal morphology predict connectivity?, *Trends Neurosci.*, 40, 2, 63–71, doi: 10.1016/j.tins.2016.11.007
- Reimann, M. W., Anastassiou, C. A., Perin, R., Hill, S. L., Markram, H., and Koch, C. (2013), A biophysically detailed model of neocortical local field potentials predicts the critical role of active membrane currents, *Neuron*, 79, 2, 375–390, doi: 10.1016/j.neuron.2013.05.023
- Reimann, M. W., Horlemann, A.-L., Ramaswamy, S., Muller, E. B., and Markram, H. (2017), Morphological diversity strongly constrains synaptic connectivity and plasticity, *Cereb. Cortex*, 27, 9, 4570–4585, doi: 10.1093/cercor/bhx150
- Reimann, M. W., King, J. G., Muller, E. B., Ramaswamy, S., and Markram, H. (2015), An algorithm to predict the connectome of neural microcircuits, *Front. Comput. Neurosci.*, 9, 120, doi: 10.3389/fncom.2015.00120
- Renart, A., De La Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., et al. (2010), The asynchronous state in cortical circuits, *Science*, 327, 587–590, doi: 10.1126/science.1179850
- Reyes-Puerta, V., Yang, J.-W., Siwek, M. E., Kilb, W., Sun, J.-J., and Luhmann, H. J. (2016), Propagation of spontaneous slow-wave activity across columns and layers of the adult rat barrel cortex in vivo, *Brain Structure Function*, 221, 9, 4429–4449, doi: 10.1007/s00429-015-1173-x

- Ribeiro, P. F. M., Ventura-Antunes, L., Gabi, M., Mota, B., Grinberg, L. T., Farfel, J. M., et al. (2013), The human cerebral cortex is neither one nor many: neuronal distribution reveals two quantitatively different zones in the gray matter, three in the white matter, and explains local variations in cortical folding, *Front. Neuroanatomy*, 7, 28, doi: 10.3389/fnana.2013.00028
- Ricciardi, L. M., Di Crescenzo, A., Giorno, V., and Nobile, A. G. (1999), An outline of theoretical and algorithmic approaches to first passage time problems with applications to biological modeling, *Math. Japonica*, 50, 2, 247–322
- Riehle, A., Wirtsohn, S., Grün, S., and Brochier, T. (2013), Mapping the spatio-temporal structure of motor cortical lfp and spiking activities during reach-to-grasp movements, *Frontiers in Neural Circuits*, 7, 48, doi: 10.3389/fncir.2013.00048
- RIKEN BSI (2013), Largest neuronal network simulation achieved using K computer, Press release
- Risken, H. (1996), *The Fokker-Planck Equation* (Springer Verlag Berlin Heidelberg), doi: 10.1007/978-3-642-61544-3_4
- Roberts, J. C. (2007), State of the art: Coordinated & multiple views in exploratory visualization, in *Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV 2007)* (IEEE Computer Society Press), doi: 10.1109/CMV.2007.20
- Robinson, D. A. (1968), The electrical properties of metal microelectrodes, *Proceedings of the IEEE*, 56, 6, 1065–1071, doi: 10.1109/PROC.1968.6458
- Rosenbaum, R. and Doiron, B. (2014), Balanced networks of spiking neurons with spatially dependent recurrent connections, *Physical Review X*, 4, 2, 021039, doi: 10.1103/PhysRevX.4.021039
- Rosenbaum, R., Smith, M. A., Kohn, A., Rubin, J. E., and Doiron, B. (2017), The spatial structure of correlated neuronal variability, *Nat. Neurosci.*, 20, 1, 107–114, doi: 10.1038/nn.4433
- Rostami, V., Ito, J., Denker, M., and Grün, S. (2017), [Re] Spike synchronization and rate modulation differentially involved in motor cortical function, *ReScience*, 3, 1, doi: 10.5281/zenodo.583814
- Rotter, S. and Diesmann, M. (1999), Exact digital simulation of time-invariant linear systems with applications to neuronal modeling, *Biol. Cybern.*, 81, 5-6, 381–402, doi: 10.1007/s004220050570
- Rougier, N. P., Hinsén, K., Alexandre, F., Arildsen, T., Barba, L. A., Benureau, F. C., et al. (2017), Sustainable computational science: the ReScience initiative, *PeerJ Computer Science*, 3, e142, doi: 10.7717/peerj-cs.142
- Rowley, A. G. D., Stokes, A. B., Knight, J., Lester, D. R., Hopkins, M., Davies, S., et al. (2015), *PyNN On Spinnaker Software 2015.004*, doi: 10.5281/zenodo.19230

- Roxin, A., Brunel, N., and Hansel, D. (2005), The role of delays in shaping spatio-temporal dynamics of neuronal activity in large networks, *Phys. Rev. Lett.*, 94, 23, 238103, doi: 10.1103/physrevlett.94.238103
- Roxin, A., Brunel, N., and Hansel, D. (2006), Rate Models with Delays and the Dynamics of Large Networks of Spiking Neurons, *Progress of Theoretical Physics Supplement*, 161, 68–85, doi: 10.1143/ptps.161.68
- Roxin, A., Brunel, N., Hansel, D., Mongillo, G., and van Vreeswijk, C. (2011), On the distribution of firing rates in networks of cortical neurons, *Journal of Neuroscience*, 31, 45, 16217–16226, doi: 10.1523/jneurosci.1677-11.2011
- Rubino, D., Robbins, K. A., and Hatsopoulos, N. G. (2006), Propagating waves mediate information transfer in the motor cortex, *Nat. Neurosci.*, 9, 12, 1549–1557, doi: 10.1038/nn1802
- Rumelhart, E., David, Hinton, E., Geoffrey, and Williams, J., Ronald (1986), Learning representations by back-propagating errors, *Nature*, 323, 533–536, doi: 10.1038/323533a0
- Sabatini, B. L. and Regehr, W. G. (1996), Timing of neurotransmission at fast synapses in the mammalian brain, *Nature*, 384, 6605, 170–172, doi: 10.1038/384170a0
- Sakmann, B. and Neher, E. (1984), Patch clamp techniques for studying ionic channels in excitable membranes, *Annu. Rev. Physiol.*, 46, 1, 455–472, doi: 10.1146/annurev.ph.46.030184.002323
- Salin, P. A. and Prince, D. A. (1996), Electrophysiological mapping of GABAA receptor-mediated inhibition in adult rat somatosensory cortex, *Journal of Neurophysiology*, 75, 4, 1589–1600, doi: 10.1152/jn.1996.75.4.1589
- Sanz Leon, P., Knock, S., Woodman, M., Domide, L., Mersmann, J., McIntosh, A., et al. (2013), The virtual brain: a simulator of primate brain network dynamics, *Front. Neuroinformatics*, 7, 10, doi: 10.3389/fninf.2013.00010
- Sato, T. K., Nauhaus, I., and Carandini, M. (2012), Traveling waves in visual cortex, *Neuron*, 75, 2, 218–229, doi: 10.1016/j.neuron.2012.06.029
- Schemmel, J., Brüderle, D., Grübl, A., Hock, M., Meier, K., and Millner, S. (2010), A wafer-scale neuromorphic hardware system for large-scale neural modeling, in Proceedings of the 2010 International Symposium on Circuits and Systems (ISCAS) (IEEE Press, Paris), 1947–1950, doi: 10.1109/iscas.2010.5536970
- Schmidt, M., Bakker, R., Hilgetag, C. C., Diesmann, M., and van Albada, S. J. (2017), Multi-scale account of the network structure of macaque visual cortex, *Brain Structure and Function*, doi: 10.1007/s00429-017-1554-4
- Schmidt, M., Bakker, R., Hilgetag, C. C., Diesmann, M., and van Albada, S. J. (2018), Multi-scale account of the network structure of macaque visual cortex, *Brain Structure & Function*, 223, 3, 1409–1435, doi: 10.1007/s00429-017-1554-4

- Schmidt, M., Bakker, R., Shen, K., Bezgin, G., Hilgetag, C.-C., Diesmann, M., et al. (2016), Full-density multi-scale account of structure and dynamics of macaque visual cortex, *arXiv preprint arXiv:1511.09364v4*
- Schneider, S. and Erlhagen, W. (2002), A neural field model for saccade planning in the superior colliculus: speed-accuracy tradeoff in the double-target paradigm, *Neurocomputing*, 44-46, 623–628, doi: 10.1016/S0925-2312(02)00449-6
- Schnepel, P., Kumar, A., Zohar, M., Aertsen, A., and Boucsein, C. (2015), Physiology and impact of horizontal connections in rat neocortex, *Cereb. Cortex*, 25, 10, 3818–3835, doi: 10.1093/cercor/bhu265
- Scholze, S., Eisenreich, H., Höppner, S., Ellguth, G., Henker, S., Ander, M., et al. (2012), A 32gbit/s communication SoC for a waferscale neuromorphic system, *Integration, the VLSI Journal*, 45, 1, 61–75, doi: 10.1016/j.vlsi.2011.05.003
- Schöner, G., Spencer, J., and Group, D. (2015), *Dynamic Thinking: A Primer on Dynamic Field Theory*, Oxford Series in Developmental Cognitive Neuroscience (Oxford University Press, Oxford)
- Schrader, S., Grün, S., Diesmann, M., and Gerstein, G. L. (2008), Detecting synfire chain activity using massively parallel spike train recording, *Journal of Neurophysiology*, 100, 2165–2176, doi: 10.1152/jn.01245.2007
- Schuecker, J., Diesmann, M., and Helias, M. (2015), Modulated escape from a metastable state driven by colored noise, *Phys. Rev. E*, 92, 052119, doi: 10.1103/PhysRevE.92.052119
- Schuecker, J., Schmidt, M., van Albada, S. J., Diesmann, M., and Helias, M. (2017), Fundamental activity constraints lead to specific interpretations of the connectome, *PLOS Comput. Biol.*, 13, 2, e1005179, doi: 10.1371/journal.pcbi.1005179
- Schüz, A. and Palm, G. (1989), Density of neurons and synapses in the cerebral cortex of the mouse, *J. Comp. Neurol.*, 286, 442–455, doi: 10.1002/cne.902860404
- Schwalger, T., Deger, M., and Gerstner, W. (2017), Towards a theory of cortical columns: From spiking neurons to interacting neural populations of finite size, *PLOS Comput. Biol.*, 13, 4, e1005507, doi: 10.1371/journal.pcbi.1005507
- Seeger, C. A. and Miller, E. K. (2010), Category learning in the brain, *Annual Review of Neuroscience*, 33, 1, 203–219, doi: 10.1146/annurev.neuro.051508.135546
- Sejnowski, T. J., Churchland, P. S., and Movshon, J. A. (2014), Putting big data to good use in neuroscience, *Nat. Neurosci.*, 17, 11, 1440–1441, doi: 10.1038/nn.3839
- Senk, J., Carde, C., Hagen, E., Kuhlen, T. W., Diesmann, M., and Weyers, B. (2018a), VIOLA - a multi-purpose and web-based visualization tool for neuronal-network simulation output, *Front. Neuroinformatics*, 12, 75, doi: 10.3389/fninf.2018.00075
- Senk, J., Carde, C., Hagen, E., Kuhlen, T. W., Diesmann, M., and Weyers, B. (2018b), VIOLA - a multi-purpose and web-based visualization tool for neuronal-network simulation output, *arXiv preprint arXiv:1803.10205v1*

- Senk, J., Hagen, E., van Albada, S. J., and Diesmann, M. (2015), From randomly connected to spatially organized multi-layered cortical network models, in 11th Göttingen Meeting of the German Neuroscience Society, T26-4B
- Senk, J., Hagen, E., van Albada, S. J., and Diesmann, M. (2018c), Reconciliation of weak pairwise spike-train correlations and highly coherent local field potentials across space, *arXiv preprint arXiv:1805.10235v1*
- Senk, J., Korvasová, K., Schuecker, J., Hagen, E., Tetzlaff, T., Diesmann, M., et al. (2018d), Conditions for traveling waves in spiking neural networks, *arXiv preprint arXiv:1801.06046v1*
- Senk, J., Korvasová, K., Schuecker, J., Hagen, E., Tetzlaff, T., Diesmann, M., et al. (2017a), Conditions for traveling waves in spiking neural networks obtained from a rigorous mapping to a neural-field model, in 26th Annual Computational Neuroscience Meeting (CNS*2017), BMC Neuroscience 2017, 18(Suppl 1):P94, doi: 10.1186/s12868-017-0371-2
- Senk, J., Yegenoglu, A., Amblet, O., Brukau, Y., Davison, A., Lester, D. R., et al. (2017b), A collaborative simulation-analysis workflow for computational neuroscience using HPC, in E. Di Napoli, M.-A. Hermanns, H. Iliev, A. Lintermann, and A. Peyser, eds., High-Performance Scientific Computing. JHPCS 2016. Lecture Notes in Computer Science, vol 10164. Springer, Cham, 243–256, doi: 10.1007/978-3-319-53862-4_21
- Shadlen, M. N. and Newsome, W. T. (1998), The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding, *Journal of Neuroscience*, 18, 10, 3870–3896, doi: 10.1523/jneurosci.18-10-03870.1998
- Sharp, T., Galluppi, F., Rast, A., and Furber, S. (2012), Power-efficient simulation of detailed cortical microcircuits on SpiNNaker, *J. Neurosci. Methods*, 210, 1, 110–118, doi: 10.1016/j.jneumeth.2012.03.001
- Sharp, T., Petersen, R., and Furber, S. (2014), Real-time million-synapse simulation of rat barrel cortex, *Front. Neurosci.*, 8, 131, doi: 10.3389/fnins.2014.00131
- Sheng, T. K. (1985), The distance between two random points in plane regions, *Adv. Appl. Prob.*, 17, 4, 748–773, doi: 10.2307/1427086
- Shinomoto, S., Shima, K., and Tanji, J. (2003), Differences in spiking patterns among cortical neurons, *Neural Comput.*, 15, 12, 2823–2842, doi: 10.1162/089976603322518759
- Shneiderman, B. (1996), The eyes have it: A task by data type taxonomy for information visualizations, in IEEE Symposium on Visual Languages, 1996 (IEEE), 336–343, doi: 10.1109/VL.1996.545307
- Shoham, S., O'Connor, D. H., and Segev, R. (2006), How silent is the brain: is there a “dark matter” problem in neuroscience?, *J. Comp. Phys.*, 192, 777–784, doi: 10.1007/s00359-006-0117-6
- Sivagnanam, S., Majumdar, A., Yoshimoto, K., Astakhov, V., Bandrowski, A., Martone, M. E., et al. (2013), Introducing the neuroscience gateway, in IWSG

- Skoglund, T. S., Pascher, R., and Berthold, C. H. (1997), The existence of a layer IV in the rat motor cortex, *Cereb. Cortex*, 7, 2, 178–180, doi: 10.1093/cercor/7.2.178
- Smith, M. A., Jia, X., Zandvakili, A., and Kohn, A. (2012), Laminal dependence of neuronal correlations in visual cortex, *Journal of Neurophysiology*, 940–947, doi: 10.1152/jn.00846.2012
- Smith, M. A. and Kohn, A. (2008), Spatial and temporal scales of neuronal correlation in primary visual cortex, *Journal of Neuroscience*, 28, 48, 12591–12603, doi: 10.1523/jneurosci.2929-08.2008
- Softky, W. R. and Koch, C. (1993), The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs, *Journal of Neuroscience*, 13, 1, 334–350, doi: 10.1523/jneurosci.13-01-00334.1993
- Sokol, S. (1976), Visually evoked potentials: Theory, techniques and clinical applications, *Survey of Ophthalmology*, 21, 1, 18–44, doi: 10.1016/0039-6257(76)90046-1
- Spence, R. (2014), Information visualization (Springer International Publishing Switzerland), 3rd edition
- Srinath, R. and Ray, S. (2014), Effect of amplitude correlations on coherence in the local field potential, *Journal of Neurophysiology*, 112, 4, 741–751, doi: 10.1152/jn.00851.2013
- Stein, R. B. (1967), Some models of neuronal variability, *Biophys. J.*, 7, 1, 37–68, doi: 10.1016/s0006-3495(67)86574-3
- Stepanyants, A., Hirsch, J. A., Martinez, L. M., Kisvárdy, Z. F., Ferecskó, A. S., and Chklovskii, D. B. (2008), Local potential connectivity in cat primary visual cortex, *Cereb. Cortex*, 18, 1, 13–28, doi: 10.1093/cercor/bhm027
- Stepanyants, A., Martinez, L. M., Ferecskó, A. S., and Kisvárdy, Z. F. (2009), The fractions of short- and long-range connections in the visual cortex, *Proceedings of the National Academy of Sciences*, 106, 9, 3555–3560, doi: 10.1073/pnas.0810390106
- Sterratt, D., Graham, B., Gillies, A., and Willshaw, D. (2011), Principles of Computational Modelling in Neuroscience (Cambridge University Press)
- Stevenson, I. H. and Kording, K. P. (2011), How advances in neural recording affect data analysis, *Nat. Neurosci.*, 14, 2, 139–142, doi: 10.1038/nn.2731
- Stokes, A. B., Rowley, A. G. D., Knight, J., Lester, D. R., Rast, A., Hopkins, M. W., et al. (2016), sPyNNaker 3.0.0
- Stromatias, E., Galluppi, F., Patterson, C., and Furber, S. (2013), Power analysis of large-scale, real-time neural networks on SpiNNaker, in Neural Networks (IJCNN), The 2013 International Joint Conference on (IEEE), 1–8, doi: 10.1109/ijcnn.2013.6706927
- Swadlow, H. A. (1988), Efferent neurons and suspected interneurons in binocular visual cortex of the awake rabbit: Receptive fields and binocular properties, *Journal of Neurophysiology*, 59, 4, 1162–1187, doi: 10.1152/jn.1988.59.4.1162

- Swadlow, H. A., Gusev, A. G., and Bezdudnaya, T. (2002), Activation of a cortical column by a thalamocortical impulse, *Journal of Neuroscience*, 22, 17, 7766–7773, doi: 10.1523/jneurosci.22-17-07766.2002
- Takahashi, K., Kim, S., Coleman, T. P., Brown, K. A., Suminski, A. J., Best, M. D., et al. (2015), Large-scale spatiotemporal spike patterning consistent with wave propagation in motor cortex, *Nature Communications*, 6, 7169, doi: 10.1038/ncomms8169
- Takahashi, K., Saleh, M., Penn, R. D., and Hatsopoulos, N. G. (2011), Propagating waves in human motor cortex, *Front. Hum. Neurosci.*, 5, 40, doi: 10.3389/fnhum.2011.00040
- Tanigawa, H., Wang, Q., and Fujita, I. (2005), Organization of horizontal axons in the inferior temporal cortex and primary visual cortex of the macaque monkey, *Cereb. Cortex*, 15, 12, 1887–1899, doi: 10.1093/cercor/bhi067
- Tetzlaff, T., Helias, M., Einevoll, G. T., and Diesmann, M. (2012), Decorrelation of neural-network activity by inhibitory feedback, *PLOS Comput. Biol.*, 8, 8, e1002596, doi: 10.1371/journal.pcbi.1002596
- Tetzlaff, T., Rotter, S., Stark, E., Abeles, M., Aertsen, A., and Diesmann, M. (2008), Dependence of neuronal correlations on filter characteristics and marginal spike-train statistics, *Neural Comput.*, 20, 9, 2133–2184, doi: 10.1162/neco.2008.05-07-525
- Thomson, A. M. (2010), Neocortical layer 6, a review, *Front. Neuroanatomy*, 4, 13, doi: 10.3389/fnana.2010.00013
- Thomson, A. M. and Lamy, C. (2007), Functional maps of neocortical local circuitry, *Front. Neurosci.*, 1, 1, 19–42, doi: 10.3389/neuro.01.1.1.002.2007
- Thomson, A. M., West, D. C., Wang, Y., and Bannister, A. P. (2002), Synaptic connections and small circuits involving excitatory and inhibitory neurons in layer 2-5 of adult rat and cat neocortex: Triple intracellular recordings and biocytin labelling in vitro, *Cereb. Cortex*, 12, 9, 936–953, doi: 10.1093/cercor/12.9.936
- Thomson, D. J. (1982), Spectrum estimation and harmonic analysis, *Proc. IEEE*, 70, 9, 1055–1096, doi: 10.1109/proc.1982.12433
- Tomsett, R. J., Ainsworth, M., Thiele, A., Sanayei, M., Chen, X., Gieselmann, M. A., et al. (2014), Virtual Electrode Recording Tool for EXtracellular potentials (VERTEX): comparing multi-electrode recordings from simulated and biological mammalian cortical tissue, *Brain Structure and Function*, 220, 4, 2333–2353, doi: 10.1007/s00429-014-0793-x
- Torre, E., Canova, C., Denker, M., Gerstein, G., Helias, M., and Grün, S. (2016a), ASSET: analysis of sequences of synchronous events in massively parallel spike trains, *PLOS Comput. Biol.*, 12, 7, e1004939, doi: 10.1371/journal.pcbi.1004939
- Torre, E., Picado-Muiño, D., Denker, M., Borgelt, C., and Grün, S. (2013), Statistical evaluation of synchronous spike patterns extracted by frequent item set mining, *Front. Comput. Neurosci.*, 7, 132, doi: 10.3389/fncom.2013.00132

- Torre, E., Quaglio, P., Denker, M., Brochier, T., Riehle, A., and Grün, S. (2016b), Synchronous spike patterns in macaque motor cortex during an instructed-delay reach-to-grasp task, *Journal of Neuroscience*, 36, 32, 8329–8340, doi: 10.1523/jneurosci.4375-15.2016
- Townsend, R. G., Solomon, S. S., Chen, S. C., Pietersen, A. N. J., Martin, P. R., Solomon, S. G., et al. (2015), Emergence of complex wave patterns in primate cerebral cortex, *Journal of Neuroscience*, 35, 11, 4657–4662, doi: 10.1523/jneurosci.4509-14.2015
- Traub, R. D., Bibbig, A., Fisahn, A., LeBeau, F. E. N., Whittington, M. A., and Buhl, E. H. (2000), A model of gamma-frequency network oscillations induced in the rat CA3 region by carbachol in vitro, *Europ. J. Neurosci.*, 12, 11, 4093–4106, doi: 10.1046/j.1460-9568.2000.00300.x
- Tuckwell, H. C. (1988), *Introduction to Theoretical Neurobiology*, volume 2 (Cambridge University Press, Cambridge)
- Turing, A. M. (1952), The chemical basis of morphogenesis, *Phil. Transact. Royal Soc.*, 237, 37–72, doi: 10.1098/rstb.1952.0012
- Tusa, R. J., Palmer, L. A., and Rosenquist, A. C. (1978), The retinotopic organization of area 17 (striate cortex) in the cat, *J. Comp. Neurol.*, 177, 2, 213–235, doi: 10.1002/cne.901770204
- Vaadia, E. (2017), private communication Sep 2017
- Vaadia, E., Kurata, K., and Wise, S. P. (1988), Neuronal activity preceding directional and nondirectional cues in the premotor cortex of rhesus monkeys, *Somatosensory & Motor Research*, 6, 2, 207–230, doi: 10.3109/08990228809144674
- van Albada, S. J., Helias, M., and Diesmann, M. (2015), Scalability of asynchronous networks is limited by one-to-one mapping between effective connectivity and correlations, *PLOS Comput. Biol.*, 11, 9, e1004490, doi: 10.1371/journal.pcbi.1004490
- van Albada, S. J., Rowley, A. G., Hopkins, M., Schmidt, M., Senk, J., Stokes, A. B., et al. (2016), Full-scale simulation of a cortical microcircuit on SpiNNaker, in *Front. Neuroinform. Conference Abstract: Neuroinformatics 2016*, doi: 10.3389/conf.fninf.2016.20.00029
- van Albada, S. J., Rowley, A. G., Senk, J., Hopkins, M., Schmidt, M., Stokes, A. B., et al. (2017), NEST-SpiNNaker comparison of large-scale network simulations, in *26th Annual Computational Neuroscience Meeting (CNS*2017)*, *BMC Neuroscience 2017*, 18(Suppl 1):P56, doi: 10.1186/s12868-017-0371-2
- van Albada, S. J., Rowley, A. G., Senk, J., Hopkins, M., Schmidt, M., Stokes, A. B., et al. (2018), Performance comparison of the digital neuromorphic hardware SpiNNaker and the neural network simulation software NEST for a full-scale cortical microcircuit model, *Front. Neurosci.*, 12, 291, doi: 10.3389/fnins.2018.00291
- van Essen, D. C. and Maunsell, J. H. R. (1980), Two-dimensional maps of the cerebral cortex, *J. Comp. Neurol.*, 191, 2, 255–281, doi: 10.1002/cne.901910208

- van Kerkoerle, T., Self, M. W., Dagnino, B., Gariel-Mathis, M.-A., Poort, J., van der Togt, C., et al. (2014), Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex, *Proceedings of the National Academy of Sciences*, 111, 40, 14332–14341, doi: 10.1073/pnas.1402773111
- van Vreeswijk, C. and Sompolinsky, H. (1996), Chaos in neuronal networks with balanced excitatory and inhibitory activity, *Science*, 274, 1724–1726, doi: 10.1126/science.274.5293.1724
- van Vreeswijk, C. and Sompolinsky, H. (1998), Chaotic balanced state in a model of cortical circuits, *Neural Comput.*, 10, 6, 1321–1371, doi: 10.1162/089976698300017214
- Vanarse, A., Osseiran, A., and Rassau, A. (2016), A review of current neuromorphic approaches for vision, auditory, and olfactory sensors, *Front. Neurosci.*, 10, 115, doi: 10.3389/fnins.2016.00115
- Veit, J., Hakim, R., Jadi, M. P., Sejnowski, T. J., and Adesnik, H. (2017), Cortical gamma band synchronization through somatostatin interneurons, *Nat. Neurosci.*, 20, 7, 951–959, doi: 10.1038/nn.4562
- Veltz, R. (2011), An analytical method for computing hopf bifurcation curves in neural field networks with space-dependent delays, *Comptes Rendus Mathématique*, 349, 13–14, 749–752, doi: 10.1016/j.crma.2011.06.014
- Veltz, R. (2013), Interplay between synaptic delays and propagation delays in neural field equations, *SIAM Journal on Applied Dynamical Systems*, 12, 3, 1566–1612, doi: 10.1137/120889253
- Venkov, N. A., Coombes, S., and Matthews, P. C. (2007), Dynamic instabilities in scalar neural field equations with space-dependent delays, *Physica D: Nonlinear Phenomena*, 232, 1, 1–15, doi: 10.1016/j.physd.2007.04.011
- Voges, N. and Perrinet, L. (2010), Phase space analysis of networks based on biologically realistic parameters, *J. Physiol. (Paris)*, 104, 1-2, 51–60, doi: 10.1016/j.jphysparis.2009.11.004
- Voges, N. and Perrinet, L. (2012), Complex dynamics in recurrent cortical networks based on spatially realistic connectivities, *Front. Comput. Neurosci.*, 6, 41, doi: 10.3389/fncom.2012.00041
- Voges, N., Schüz, A., Aertsen, A., and Rotter, S. (2010), A modeler’s view on the spatial structure of intrinsic horizontal connectivity in the neocortex, *Prog. Neurobiol.*, 92, 3, 277–292, doi: 10.1016/j.pneurobio.2010.05.001
- von Neumann, J. (1993), First draft of a report on the EDVAC, *Annals of the History of Computing, IEEE*, 15, 4, 27–75, doi: 10.1109/85.238389
- Wagatsuma, N., Potjans, T. C., Diesmann, M., and Fukai, T. (2011), Layer-dependent attentional processing by top-down signals in a visual cortical microcircuit model, *Front. Comput. Neurosci.*, 5, 31, doi: 10.3389/fncom.2011.00031

- Waldeck, R. F., Pereda, A., and Faber, D. S. (2000), Properties and plasticity of paired-pulse depression at a central synapse, *Journal of Neuroscience*, 20, 14, 5312–5320, doi: 10.1523/jneurosci.20-14-05312.2000
- Wang, X.-J. and Buzsáki, G. (1996), Gamma Oscillation by Synaptic Inhibition in a Hippocampal Interneuronal Network Model, *Journal of Neuroscience*, 16, 20, 6402–6413, doi: 10.1523/jneurosci.16-20-06402.1996
- Wang Baldonado, M. Q., Woodruff, A., and Kuchinsky, A. (2000), Guidelines for using multiple views in information visualization, in Proceedings of the Working Conference on Advanced Visual Interfaces (ACM, New York, NY, USA), AVI '00, 110–119, doi: 10.1145/345513.345271
- Ward, M., Grinstein, G., and Keim, D. (2010), Interactive data visualization: foundations, techniques, and applications (A. K. Peters, Ltd.)
- Welch, P. D. (1967), The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms, *IEEE Transactions on Audio Electroacoustics*, 15, 2, 70–73, doi: 10.1109/TAU.1967.1161901
- Whittington, M. A., Traub, R. D., and Jefferys, J. G. R. (1995), Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation, *Nature*, 373, 6515, 612–615, doi: 10.1038/373612a0
- Whittington, M. A., Traub, R. D., Kopell, N., Ermentrout, B., and Buhl, E. H. (2000), Inhibition-based rhythms: experimental and mathematical observations on network dynamics, *Int. J. Psychophysiol.*, 38, 3, 315–336, doi: 10.1016/S0167-8760(00)00173-2
- Wils, S. and De Schutter, E. (2009), STEPS: modeling and simulating complex reaction-diffusion systems with Python, *Front. Neuroinformatics*, 3, 15, doi: 10.3389/neuro.11.015.2009
- Wilson, H. R. and Cowan, J. D. (1972a), Excitatory and inhibitory interactions in localized populations of model neurons, *Biophys. J.*, 12, 1, 1–24, doi: 10.1016/S0006-3495(72)86068-5
- Wilson, H. R. and Cowan, J. D. (1972b), Excitatory and inhibitory interactions in localized populations of model neurons, *Biophysical Journal*, 12, 1, 1 – 24, doi: [http://dx.doi.org/10.1016/S0006-3495\(72\)86068-5](http://dx.doi.org/10.1016/S0006-3495(72)86068-5)
- Wilson, H. R. and Cowan, J. D. (1973a), A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue, *Kybernetik*, 13, 2, 55–80, doi: 10.1007/BF00288786
- Wilson, H. R. and Cowan, J. D. (1973b), A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue, *Kybernetik*, 13, 2, 55–80, doi: 10.1007/BF00288786
- Wu, J.-Y., Huang, X., and Zhang, C. (2008), Propagating waves of activity in the neocortex: What they are, what they do, *The Neuroscientist*, 14, 5, 487–502, doi: 10.1177/1073858408317066

- Wyller, J., Blomquist, P., and Einevoll, G. T. (2007a), On the origin and properties of two-population neural field models - a tutorial introduction, *Biophysical Reviews and Letters*, 02, 01, 79–98, doi: 10.1142/s1793048007000441
- Wyller, J., Blomquist, P., and Einevoll, G. T. (2007b), Turing instability and pattern formation in a two-population neuronal network model, *Physica D: Nonlinear Phenomena*, 225, 1, 75–93, doi: 10.1016/j.physd.2006.10.004
- Xing, D., Yeh, C. I., Burns, S., and Shapley, R. M. (2012), Laminar analysis of visually evoked activity in the primary visual cortex, *Proceedings of the National Academy of Sciences*, 109, 34, 13871–13876, doi: 10.1073/pnas.1201478109
- Xu, W., Huang, X., Takagaki, K., and young Wu, J. (2007), Compression and reflection of visually evoked cortical waves, *Neuron*, 55, 1, 119–129, doi: 10.1016/j.neuron.2007.06.016
- Yamauchi, S., Kim, H., and Shinomoto, S. (2011), Elemental spiking neuron model for reproducing diverse firing patterns and predicting precise firing times, *Front. Comput. Neurosci.*, 5, 42, doi: 10.3389/fncom.2011.00042
- Yamawaki, N., Borges, K., Suter, B. A., Harris, K. D., and Shepherd, G. M. G. (2014), A genuine layer 4 in motor cortex with prototypical synaptic circuit connectivity, *eLife*, 3, e05422, doi: 10.7554/elife.05422
- Yegenoglu, A., Davison, A., Holstein, D., Muller, E., Torre, E., Sprenger, J., et al. (2016a), Elephant 0.2.1. <https://github.com/neuralensemble/elephant/releases/tag/0.2.1>
- Yegenoglu, A., Davison, A. P., Holstein, D., Muller, E., Torre, E., Hagen, E., et al. (2016b), Elephant 0.3.0
- Yegenoglu, A., Denker, M., Phan, L., Holstein, D., Chorley, P., Ito, J., et al. (2015), Elephant – open-source tool for the analysis of electrophysiological data sets., in Bernstein Conference 2015: Abstract Book, W-05, doi: 10.12751/n-ncn.bc2015.0126
- Yger, P., El Boustani, S., Destexhe, A., and Frégnac, Y. (2011), Topologically invariant macroscopic statistics in balanced networks of conductance-based integrate-and-fire neurons, *Journal of Computational Neuroscience*, 31, 229–245, doi: 10.1007/s10827-010-0310-z
- Yu, Y., McTavish, T. S., Hines, M. L., Shepherd, G. M., Valenti, C., and Migliore, M. (2013), Sparse distributed representation of odors in a large-scale olfactory bulb circuit, *PLOS Comput. Biol.*, 9, 3, e1003014, doi: 10.1371/journal.pcbi.1003014
- Zanos, T. P., Mineault, P. J., Nasiotis, K. T., Guitton, D., and Pack, C. C. (2015), A sensorimotor role for traveling waves in primate visual cortex, *Neuron*, 85, 3, 615–627, doi: 10.1016/j.neuron.2014.12.043
- Zaytsev, Y. V. and Morrison, A. (2014), CyNEST: a maintainable Cython-based interface for the NEST simulator, *Front. Neuroinformatics*, 8, 23, doi: 10.3389/fninf.2014.00023

- Zehl, L., Jaillet, F., Stoewer, A., Grewe, J., Sobolev, A., Wachtler, T., et al. (2016), Handling metadata in a neurophysiology laboratory, *Front. Neuroinformatics*, 10, 26, doi: 10.3389/fninf.2016.00026
- Zucker, R. S. and Regehr, W. (2002), Short-term synaptic plasticity, *Annual Review of Physiology*, 64, 355–405, doi: 10.1146/annurev.physiol.64.092501.114547

FUNDING

This work received financial support from:

- Helmholtz Portfolio "Supercomputing and Modeling for the Human Brain" (SMHB)
- Helmholtz young investigator group VH-NG-1028
- European Union (BrainScaleS, grant 269921)
- European Union Seventh Framework Programme [FP7/2007-2013] under grant agreement No. 604102 (Human Brain Project, Ramp up phase)
- European Union's Horizon 2020 research and innovation programme under grant agreement No. 720270 (HBP SGA1) and No. 785907 (HBP SGA2)
- European Research Council under the European Union's Seventh Framework Program (FP7/2007-2013) / ERC grant agreement No. 320689
- ERS grant "Facing the multi-scale problem in neuroscience" of the RWTH Aachen University
- German Research Foundation (DFG, grant DI 1721/3-1 [KFO219-TP9])
- DFG SPP Priority Program 1665 (GR 1753/4-1 and DE 2175/1-1)
- Research Council of Norway (NFR through COBRA)
- UK Engineering and Physical Sciences Research Council (EPSRC) under grants EP/D07908X/1 and EP/G015740/1
- Jülich Aachen Research Alliance (JARA)

Computing time on the supercomputers JURECA and JUQUEEN at Forschungszentrum Jülich was granted by the JARA-HPC Vergabegremium (provided on the JARA-HPC partition, jinb33).

ACKNOWLEDGMENTS

I am very grateful to all the people who supported me during the last years and who contributed in different ways to this thesis. First and foremost, I would like to thank Markus Diesmann and Sonja Grün for giving me the opportunity to grow in an inspiring, multi-disciplinary, constructive and friendly atmosphere, both at INM-6 and beyond at conferences and courses. Many thanks go to Markus for his commitment in the doctoral supervision and for teaching me his view on both detail and greater vision of our research, and thanks to Torsten Kuhlen and Jörg Fitter for being the second and third reviewer of my thesis. I am very thankful to Espen Hagen who constantly supported and guided me in our joint projects and besides in all questions on the doctoral life and the world of science. Thank you so much for your availability over this whole time despite of geographical distances (Jülich, Ås / Oslo) and meetings hampered by problems with different video conferencing systems. Moritz Helias, Sacha van Albada, Michael Denker and Benjamin Weyers guided me through other projects and I would like to express my distinct gratitude to all of them for the opportunity to learn from their broad technical knowledge and their vast experiences on scientific processes. I thank Alper Yegenoglu for the great collaboration with (literally) equal shares in our project on collaboration; we even split a conference talk in two. Likewise, I thank Karolína Korvasová for our productive collaboration built around the births of her two children, with fruitful discussions enhanced by her rigorous mathematical perspective. Since a large part of this thesis deals with numerical simulations using NEST, I am furthermore thankful to the community developing and maintaining this simulator, and especially to Hans Ekkehard Plesser for his help with the topology module. A thanks for support in hardware and software questions goes to Bernd Wiebelt, Tiziano Zito and Michael Bontenackels. I owe many thanks to my office mates, Jakob Jordan, David Dahmen, Jonas Stapmanns, Hannah Bos and Lukas Deutz, as well as to my colleagues Julia Sprenger, Philipp Weidel, Tobias Kühn, Maximilian Schmidt, Jannis Schücker, Robin Pauli, Nicole Voges and others: for numerous discussions on science and more, and just for their company. Thanks, Jakob, for your enormous supply of optimism and soup. Thanks, David, for patiently answering all my questions and giving me feedback. Thanks also to the lunch group for sharing delicious and suspicious food. Another thanks goes to Martina Reske, Janine Lehm, Petra O'Brian, Stephanie Weitz and Verena Nelles for help with all bureaucratic issues. Finally, I would like to thank my family for their understanding and mental support. A special thanks goes to my boyfriend Markus Hrywniak for his patience, empathy and encouragement during this long-lasting stressful phase.

COLOPHON

This document was typeset using `classicthesis` style developed by André Miede. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". It is available for \LaTeX and $\text{L}\chi\text{X}$ at

<https://bitbucket.org/amiede/classicthesis/>

Band / Volume 181

Creating and characterizing a single molecule device for quantitative surface science

M. Green (2018), viii, 142 pp (untersch. Pag.)

ISBN: 978-3-95806-352-5

Band / Volume 182

8th Georgian-German School and Workshop in Basic Science

A. Kacharava (Ed.) erscheint nur als CD (2018)

ISBN: 978-3-95806-353-2

Band / Volume 183

Topological properties of complex magnets from an advanced *ab-initio* Wannier description

J.-P. Hanke (2018), xi, 173 pp

ISBN: 978-3-95806-357-0

Band / Volume 184

Translation Initiation with 70S Ribosomes: A Single Molecule Study

C. Remes (2018), iv, 113 pp

ISBN: 978-3-95806-358-7

Band / Volume 185

Scanning tunneling potentiometry at nanoscale defects in thin films

F. Lüpke (2018), iv, 144 pp (untersch. Pag.)

ISBN: 978-3-95806-361-7

Band / Volume 186

Inelastic neutron scattering on magnetocaloric compounds

N. Biniskos (2018), iii, 92 pp

ISBN: 978-3-95806-362-4

Band / Volume 187

Magnetic Order and Excitation in Frustrated Pyrochlore 5d - Transition Metal Oxides

E. Feng (2018), iv, 182 pp

ISBN: 978-3-95806-365-5

Band / Volume 188

Finite-Difference Time-Domain Simulations Assisting to Reconstruct the Brain's Nerve Fiber Architecture by 3D Polarized Light Imaging

M. Menzel (2018), ix, 296 pp

ISBN: 978-3-95806-368-6

Band / Volume 189

**Characterization of the cell-substrate interface
using surface plasmon resonance microscopy**

E. M. Kreysing (2018), xiii, 260 pp

ISBN: 978-3-95806-369-3

Band / Volume 190

Scattering! Soft, Functional and Quantum Materials

Lecture Notes of the 50th IFF Spring School 2019

11 – 22 March 2019, Jülich, Germany

ed. by M. Angst, T. Brückel, S. Förster, K. Friese, R. Zorn (2019),

ca 1000 pp

ISBN: 978-3-95806-380-8

Band / Volume 191

**Absolute scale off-axis electron holography of thin dichalcogenide
crystals at atomic resolution**

F. Winkler (2019), xxiii, 187 pp

ISBN: 978-3-95806-383-9

Band / Volume 192

**High-resolution genome and transcriptome analysis of *Gluconobacter
oxydans* 621H and growth-improved strains by next-generation
sequencing**

A. Kranz (2019), III, 182 pp

ISBN: 978-3-95806-385-3

Band / Volume 193

Group IV (Si)GeSn Light Emission and Lasing Studies

D. Stange (2019), vi, 151 pp

ISBN: 978-3-95806-389-1

Band / Volume 194

Construction and analysis of a spatially organized cortical network model

J. Senk (2019), 245 pp

ISBN: 978-3-95806-390-7

Weitere **Schriften des Verlags im Forschungszentrum Jülich** unter
<http://www.zb1.fz-juelich.de/verlagextern1/index.asp>

Schlüsseltechnologien / Key Technologies
Band / Volume 194
ISBN 978-3-95806-390-7