




Dynamical and statistical structure of spatially organized neuronal networks

Moritz Layer

Information

Band / Volume 85

ISBN 978-3-95806-651-9

DYNAMICAL AND STATISTICAL STRUCTURE OF
SPATIALLY ORGANIZED NEURONAL NETWORKS

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften der
RWTH Aachen University zur Erlangung des akademischen Grades eines
Doktors der Naturwissenschaften genehmigte Dissertation

vorgelegt von

Master of Advanced Study
Moritz Layer

aus Backnang

Berichter: Jun.-Prof. Dr. Moritz Helias
Univ.-Prof. Dr. Björn Kampa

Tag der mündlichen Prüfung: 29.08.2022

Diese Dissertation ist auf den Internetseiten der Universitätsbibliothek
online verfügbar.

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

Dynamical and statistical structure of spatially organized neuronal networks

Moritz Layer

Schriften des Forschungszentrums Jülich
Reihe Information / Information

Band / Volume 85

ISSN 1866-1777

ISBN 978-3-95806-651-9

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 2022

Schriften des Forschungszentrums Jülich
Reihe Information / Information, Band / Volume 85

D 82 (Diss. RWTH Aachen University, 2022)

ISSN 1866-1777
ISBN 978-3-95806-651-9

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.

EIDESSTATTLICHE ERKLÄRUNG

Ich, Moritz Layer, erkläre hiermit, dass diese Dissertation und die darin dargelegten Inhalte die eigenen sind und selbstständig, als Ergebnis der eigenen originären Forschung, generiert wurden.

Hiermit erkläre ich an Eides statt

1. Diese Arbeit wurde vollständig oder größtenteils in der Phase als Doktorand dieser Fakultät und Universität angefertigt;
2. Sofern irgendein Bestandteil dieser Dissertation zuvor für einen akademischen Abschluss oder eine andere Qualifikation an dieser oder einer anderen Institution verwendet wurde, wurde dies klar angezeigt;
3. Wenn immer andere eigene- oder Veröffentlichungen Dritter herangezogen wurden, wurden diese klar benannt;
4. Wenn aus anderen eigenen- oder Veröffentlichungen Dritter zitiert wurde, wurde stets die Quelle hierfür angegeben. Diese Dissertation ist vollständig meine eigene Arbeit, mit der Ausnahme solcher Zitate;
5. Alle wesentlichen Quellen von Unterstützung wurden benannt;
6. Wenn immer ein Teil dieser Dissertation auf der Zusammenarbeit mit anderen basiert, wurde von mir klar gekennzeichnet, was von anderen und was von mir selbst erarbeitet wurde;
7. Ein Teil oder Teile dieser Arbeit wurden zuvor veröffentlicht, siehe "Author's list of publications".

AUTHOR'S LIST OF PUBLICATIONS

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

NNMT: Mean-field based analysis tools for neuronal network models

Moritz Layer, Johanna Senk, Simon Essink, Alexander van Meegen, Hannah Bos, and Moritz Helias

Published in Frontiers in Neuroinformatics: Layer et al. (2022).

Parts of this publication enter Section 1.4 and Chapter 2.

Global organization of neuronal activity only requires unstructured local connectivity

David Dahmen*, Moritz Layer*, Lukas Deutz, Paulina Anna Dąbrowska, Nicole Voges, Michael von Papen, Thomas Brochier, Alexa Riehle, Markus Diesmann, Sonja Grün, and Moritz Helias

Published in eLife: Dahmen et al. (2022).

* These authors contributed equally.

Parts of this publication enter Chapter 3.

Author contributions are indicated at the beginning of the respective chapters and sections.

SUMMARY

DYNAMICAL AND STATISTICAL STRUCTURE OF SPATIALLY ORGANIZED NEURONAL NETWORKS

The cerebral cortex, the outer layer of mammalian brains, comprises a vast number of neurons arranged and connected in a highly organized fashion. The likelihood of neurons to be connected and how fast they may exchange signals depends, among other properties, on their spatial distance. Cortical networks may be well described as completely random networks on microscopic scales because cortical neurons have essentially uniform connection probabilities within a few tens of micrometers. However, the distance-dependence of neuronal connections certainly is important on mesoscopic scales spanning several millimeters, where many neurons are most likely unconnected. While the theory of random networks is already well-established, how such a spatial organization affects a network's activity is not yet fully understood. The objective of this thesis is to provide an overview of the current analytical understanding of spatially organized networks on a mesoscopic scale, as well as to advance this understanding with three studies covering complementary aspects of spatially organized network theory.

A variety of experimental recordings in cortex reveals that neuronal activity is coordinated across several millimeters: Multi-electrode-arrays covering a few square millimeters, for example, provide access to the local field potential, a measure of population activity, as well as single neuron spiking activity. While spiking activity exhibits distance-dependent correlation characteristics, population activity shows spatio-temporally coherent activity, like periodic patterns, waves, or bumps. In this thesis we employ a combination of network models, analytical tools, and simulations to gain an understanding of such findings. We particularly make use of mean-field theory, which is a viable tool for investigating statistical properties of populations made up of thousands of neurons, and it therefore may be utilized to gain a coarse-grained description of network activity at large scales. In the first main part, we present a Python package we developed to make previously developed analytical results from neuronal network mean-field theory applicable to concrete network models, giving access to estimates of model properties such as firing rates and power spectra, as well as more elaborate tools that can support network modeling. In the second study, we investigate how neurons may coordinate their activity dynamically across large distances, without the need for highly correlated input or long-range connections. In the third study, we explore how a temporal delay may affect pattern formation in planar networks.

As we demonstrate, spatial organization is a critical network feature that does not merely lead to obvious phenomena like spatially structured activity. On the contrary, as we show in this thesis, spatial organization leads to a variety of interesting, non-trivial effects, that on first sight might even seem counterintuitive, and this topic certainly provides a multitude of intriguing research questions for the near future.

ZUSAMMENFASSUNG

DYNAMISCHE UND STATISTISCHE STRUKTUR RÄUMLICH ORGANISierter NEURONALER NETZWERKE

Der Kortex, die äußere Schicht des Großhirns, besteht aus einer großen Anzahl Neuronen, die in einer hoch organisierten Weise angeordnet und verbunden sind. Die Wahrscheinlichkeit, mit der Neuronen miteinander verbunden sind und wie schnell sie Signale austauschen können, hängt unter anderem von ihrem räumlichen Abstand ab. Auf einer mikroskopischen Skala können kortikale Netzwerke gut als völlig ungeordnete, zufällige Netzwerke beschrieben werden, da kortikale Neuronen innerhalb einiger zehn Mikrometer im Wesentlichen einheitliche Verbindungswahrscheinlichkeiten aufweisen. Auf einer mesoskopischen Skala, die sich über mehrere Millimeter erstreckt, auf der viele Neuronen mit hoher Wahrscheinlichkeit nicht direkt verbunden sind, spielt die Abhängigkeit der neuronalen Verbindungen allerdings sicherlich eine wesentliche Rolle. Während die Theorie zufälliger neuronaler Netzwerke bereits gut entwickelt ist, ist noch nicht vollständig geklärt, wie sich eine solche räumliche Organisation auf die Aktivität eines Netzwerkes auswirkt. Ziel dieser Arbeit ist es, einen Überblick über das derzeitige analytische Verständnis räumlich organisierter neuronaler Netzwerke zu geben und dieses Verständnis durch drei Studien zu erweitern, die komplementäre Aspekte der Theorie räumlich organisierter Netzwerke abdecken.

Eine Vielzahl von Experimenten zeigt, dass kortikale neuronale Aktivität über mehrere Millimeter hinweg koordiniert ist: Multi-Elektroden-Arrays, die sich über einige Quadratmillimeter erstrecken, ermöglichen beispielsweise gleichzeitig sowohl das lokale Feldpotential, ein Maß für die neuronale Populationsaktivität, als auch die Aktionspotentiale einzelner Neurone aufzuzeichnen. Während die Aktionspotentiale entfernungsabhängige Korrelationsmuster aufweisen, zeigt die Populationsaktivität räumlich und zeitlich kohärente Aktivität in Form von periodischen Mustern, Wellen oder räumlich begrenzten Anregungen. In dieser Arbeit nutzen wir eine Kombination aus Netzwerkmodellen, analytischen Methoden und Simulationen, um ein besseres Verständnis für derartige Beobachtungen zu entwickeln. Wir nutzen dazu insbesondere Mean-Field Theorie (auch bekannt als Molekularfeldtheorie), die ein praktisches Werkzeug zum Verständnis statistischer Eigenschaften neuronaler Populationen, die aus Tausenden von Neuronen bestehen, darstellt und daher eine statistische Beschreibung der Netzwerkaktivität auf großen Skalen erlaubt. Im ersten Hauptteil stellen wir ein von uns entwickeltes Python-Paket, das es erlaubt, abstrakte analytische Ergebnisse aus der Mean-Field Theorie neuronaler Netzwerke auf konkrete Netzwerkmodelle anzuwenden. Das Paket enthält Methoden zu statistischen Schätzungen von Modelleigenschaften wie Feuerraten und Leistungsspektren, sowie komplexere Werkzeuge, die die Netzwerkmodellierung unterstützen und vereinfachen können. In einer zweiten Studie untersuchen wir, wie Neuronen ihre Aktivität dynamisch

über große Entfernungen koordinieren können, ohne dass ein stark korrelierter Input oder langreichweitige Verbindungen erforderlich wären, während wir in einer dritten Studie untersuchen, wie sich eine zeitliche Verzögerung bei der neuronalen Signalübermittlung auf die Bildung von Aktivitätsmustern in planaren Netzwerken auswirkt.

Wie wir im Verlauf dieser Arbeit zeigen, ist räumliche Organisation ein entscheidendes Netzwerkmerkmal, das nicht nur zu naheliegenden Phänomenen wie räumlich strukturierter Aktivität führt. Ganz im Gegenteil führt räumliche Organisation zu einer Vielzahl von interessanten, nicht-trivialen Effekten, die auf den ersten Blick vielleicht sogar kontraintuitiv erscheinen, und das Thema bietet sicherlich eine Reihe interessanter Fragestellungen für weitere Forschungsprojekte.

CONTENTS

1	INTRODUCTION	1
1.1	The spatio-temporal connectivity of neocortex	2
1.2	Experimentally observed mesoscopic in vivo activity in cortex	3
1.3	Neuron and network models utilized in this thesis	5
1.4	Mean-field theory of random networks	8
1.5	Theory of spatially organized networks	10
1.6	Scope of this thesis	11
2	NNMT: MEAN-FIELD BASED ANALYSIS TOOLS FOR NEURONAL NETWORK MODELS	13
2.1	Introduction	13
2.2	Workflows and architecture	14
2.2.1	Basic workflow	16
2.2.2	Model workflow	16
2.2.3	Structure of the toolbox	17
2.3	How to use the toolbox	18
2.3.1	Installation and setup	18
2.3.2	Usage example: analyzing a microcircuit model	18
2.4	Discussion	23
2.4.1	Comparison to other tools	23
2.4.2	Use cases	24
2.4.3	Limitations	25
2.4.4	How to contribute and outlook	26
3	CORRELATIONS IN TWO-DIMENSIONAL NETWORKS	29
3.1	Introduction	29
3.2	Macaque motor cortex shows long-range coordination patterns	30
3.3	Multi-synaptic connections determine correlations	32
3.4	Almost unstable networks show shallow exponential decay of covariances	34
3.5	Pairwise covariances in motor cortex decay on a millimeter scale	38
3.6	Firing rates alter connectivity-dependent covariance patterns	39
3.7	Coordination patterns in motor cortex depend on behavioral context	41
3.8	Discussion	44
3.9	Methods	48
4	ACTIVITY PATTERNS IN TWO-DIMENSIONAL NETWORKS	57
4.1	Introduction	57
4.2	The model	59
4.3	Conditions for patterns to emerge	63
4.3.1	Comparison of one- and two-dimensional networks	65
4.3.2	The effect of the neuron grid	65

4.3.3	Example: Gaussian connectivity profiles	68
4.4	Studying the interaction of Fourier modes	69
4.4.1	Center manifold reduction for ODEs	70
4.4.2	Center manifold reduction for ODEs with parameters	71
4.4.3	Center manifold reduction for DDEs	72
4.5	Simulation	73
4.6	Static patterns	74
4.6.1	Interaction of static spatial waves	75
4.6.2	Excitation of stable modes	78
4.7	Discussion	80
5	DISCUSSION	85
5.1	Conclusions	85
5.2	Sub-projects in the context of spatial organization	87
5.3	Overarching assumptions in this thesis	88
5.4	Relevance of spatial organization	90
5.5	Outlook	91
A	NNMT	93
A.1	Microcircuit parameters	93
B	LONG RANGE CORRELATIONS	95
B.1	Correlations and covariances	95
B.2	Robustness to E/I separation	97
B.3	Stationarity of behavioral data	98
B.4	Network model	103
B.5	Covariances	105
B.6	Cumulant generating function of connectivity matrix	106
B.7	Note on derivation of variance of covariances	107
B.8	Utilizing symmetries to reduce dimensionality	107
B.9	One-dimensional network with one population	109
B.9.1	Dimensionality reduction	110
B.9.2	Continuum limit	111
B.9.3	Prediction of exponential decay of covariance statistics	111
B.9.4	Long-range limit	112
B.10	One-dimensional network with two populations	114
B.10.1	Dimensionality reduction	115
B.10.2	General results	118
B.10.3	Long-range limit	119
B.11	Two-dimensional network with one population	121
B.11.1	Dimensionality reduction	122
B.11.2	Long-range limit	123
B.11.3	Note on higher order approximation	125
B.12	Two-dimensional network with two populations	126
B.12.1	Dimensionality reduction	126
B.12.2	Long-range limit	127
B.12.3	Higher order approximation	128

B.13	Validation of theory	129
B.14	Parameters of NEST simulation	131
B.15	Sources of heterogeneity	131
C	ACTIVITY PATTERNS IN TWO-DIMENSIONAL NETWORKS	133
C.1	Transition curves in network with Gaussian connectivity	133
C.2	NEST simulation parameters	135
C.3	Analysis of the convolutions in the non-linearity	136
C.4	Center manifold reduction for static patterns	137
	Bibliography	141

INTRODUCTION

The brain could just be a blob of neurons connected in a seemingly random fashion. Obviously, it is not. Even if one ignores the fact that the brain is a complex ensemble of numerous components and one only considers cortex, the specifics of its organization in space and time are overwhelming. But why is that? There are clearly various ways of approaching this issue, ranging from an evolutionary perspective to biological, space, and energy constraints, as well as functional considerations. In this thesis we are studying this question with a focus on network activity: How does spatial organization and a spatially structured network connectivity affect the network's activity? In particular, we have a look at cortical networks on the mesoscopic scale, which spans several millimeters and is now accessible to a variety of technologies enabling massively parallel recordings that can resolve single neuron activity. The availability of such experimental methods has given rise to a number of intriguing questions related to the spatial organization of network activity: How does the activity of neuronal populations consisting of thousands of neurons distributed across several millimeters arise from their microscopic connectivity properties? How do these neurons coordinate their activity across large distances? How does coherent population activity arise? These are the topics we address in the following chapters with the help of network models, analytical tools, and simulations.

To this end, we employ three different approaches: we introduce a Python package that implements a variety of analytical methods developed for estimating population activity statistics from underlying model parameters, allowing them to be applied for network model analysis (Chapter 2); we investigate how large-scale correlations can arise dynamically in a network with only unstructured local connectivity and without the need for strongly correlated input (Chapter 3); and we study how a constant communication delay and the neuronal input-output function influence the spontaneous emergence of spatio-temporally coherent population activity in planar networks (Chapter 4). We discuss these projects in the context of spatial organization and address why spatial structure is vital for understanding biological neuronal networks (Chapter 5).

The remainder of this chapter is dedicated to introducing the required background and concepts that occur throughout this thesis. Section 1.1 presents a brief overview of cortical connectivity structure, with a focus on distance-dependent connection probabilities on a local scale, as well as communication delays. Section 1.2 reviews some literature on experimental results on activity on the mesoscopic scale, particularly coherent population and spiking activity, to provide a notion of how cortical activity on this scale looks like. Section 1.3 introduces the neuron and network models used

in the main parts of this thesis. Section 1.4 discusses analytical results for random networks, which are well established, while in Section 1.5 we present some results of the far less developed theory of spatially organized networks. Finally, Section 1.6 explains the scope of this thesis in detail and gives a brief overview of the three main projects that constitute the core of this thesis and their different objectives.

1.1 THE SPATIO-TEMPORAL CONNECTIVITY OF NEOCORTEX

The mammalian brain consists of several hierarchically organized components that interact with each other to perform the manifold tasks of a brain. The organization of the central nervous system is described in detail in Kandel et al. (2013, Part IV, Chapter 15). In the following paragraphs, we only want to recapitulate the basics and concepts required to comprehend the work conducted within the scope of this thesis. Starting from the outside, the brain can be divided into three parts: the cerebrum, the largest component with its two hemispheres, a smaller part called cerebellum, and the brain stem, which connects the brain to the spinal cord. In this thesis we focus on the neocortex, which accounts for roughly 90 % of the cerebrum's outer layer of neuronal tissue, and is responsible for higher-order tasks like sensory perception, cognition, or the generation of motor commands. In rodents, the neocortex is smooth, but in primates and other large animals, it is folded, which increases its surface area. It can be subdivided into four lobes with differing functionalities: the occipital lobe (e.g., vision), the temporal lobe (e.g., hearing), the parietal lobe (e.g., somatosensation), and the frontal lobe (e.g., cognition and movement). A further division can be made on the basis of cytoarchitecture (Brodmann, 1909).

Apart from other cells, the neocortex is made up of approximately 80 % excitatory and 20 % inhibitory neurons, typically arranged in six horizontal layers. These differ in the contained neuron types as well as the connectivity between the neurons. Whereas here, we do not have a closer look at the different neuron types, the connectivity plays a major role in this thesis. But what exactly is connectivity? There are three aspects of neuronal interaction that fall into the category of connectivity: First, the synaptic or structural connectivity. If the axon of a neuron comes close to the dendrites of another neuron, a synapse may form, which determines the topography of the network. Second, the strength of the connectivity, or functional connectivity, which describes how strong the evoked post synaptic potential is, and this can be understood on top of the topography. Third, there is the temporal aspect of the connectivity, determining how long it takes for a neuron to communicate a signal to another neuron.

The connectivity of cortex is highly structured in space and time. We first consider the spatial component: Connectivity of cortex can be understood in terms of horizontal connectivity, parallel to the cortical surface, connecting nearby areas of cortex, vertical connectivity, perpendicular to the surface, connecting different layers, and connectivity through white matter, connecting distant areas and transmitting sensory input or motor output. In this thesis we focus on the horizontal connectivity, for which it is often useful to think in terms of local connectivity, within the range of a few hundred micrometers, and non-local connectivity going beyond this range (for a review see Voges et al., 2010; Binzegger et al., 2007). The non-local connectivity often is found to be patchy with neurons projecting to groups of neurons, typically on a millimeter scale

(Livingstone and Hubel, 1984; Bosking et al., 1997; Tanigawa et al., 2005; Buzás et al., 2006; Binzegger et al., 2007; Voges et al., 2010). In this thesis, we will investigate the effect of the local connectivity on network activity. On a local scale, neurons typically connect with a certain connection probability which decays with the inter-neuronal distance (Hellwig, 2000; Oswald et al., 2009; Boucsein et al., 2011; Markov et al., 2011; Levy and Reyes, 2012; Jiang et al., 2015; Schnepel et al., 2015; Horvát et al., 2016). The exact shape of the spatial connectivity profile is unclear: Markov et al. (2011) for example report an exponential decay, while Levy and Reyes (2012) find a Gaussian shape. The typical length scale on which the connection probability decays depends on neuron types and layers (Packer and Yuste, 2011; Jiang et al., 2015; Schnepel et al., 2015; Reimann et al., 2017), as well as on the brain area (Kätzel et al., 2011). Often, excitatory neurons are reported to have a broader profile than inhibitory neurons: Stepanyants et al. (2007) study the local potential connectivity in cat visual cortex by examining locations where the axon of a neuron is close to the dendrite of another neuron and find a broader distribution for excitatory than for inhibitory neurons. Stepanyants et al. (2009) show that in a section of $500\text{ }\mu\text{m}$ radius of cat visual cortex 74% of the excitatory synapses originate from neurons outside the section, while it is only around 10% for inhibitory neurons. Ohana et al. (2012) find that the mean distance between pre- and postsynaptic somata are larger for excitatory connections than for inhibitory ones in cat visual cortex. However, Levy and Reyes (2012) find comparable ranges for the connection profiles for excitatory and inhibitory neurons with only slightly longer ranges for excitatory neurons in mouse auditory cortex. Reimann et al. (2017) provide extensive data on connection probability profiles of different neuron types within and between different layers of cat visual cortex in their supplementary material.

The temporal connectivity structure is given by the signal transmission delays between two neurons, which comprise two components: a distance-dependent delay due to finite axonal conduction velocity, and a constant delay due to synaptic processing. Experimental findings regarding axonal conduction delays are reviewed in Swadlow and Waxman (2012). Axonal conduction velocities strongly depend on myelination and can range from about 0.3 m/s to over 100 m/s . Measurements of local horizontal axonal connections in cortex (Hirsch and Gilbert, 1991; Lohmann and Rörig, 1994) indicate that those axons are largely non-myelinated, such that for our purposes the lower end of conduction velocities may be more relevant (however, cf. discussion on Kisvárdy and Eysel (1992) in Swadlow and Waxman (2012)). Synaptic delays are caused by neurotransmitter release, binding, and spike initiation and are estimated to be in the range of $0.5 - 1.0\text{ ms}$ (Katz and Miledi, 1965; Murakoshi et al., 1993; Kang et al., 1994).

1.2 EXPERIMENTALLY OBSERVED MESOSCOPIC IN VIVO ACTIVITY IN CORTEX

Which experimental method is used to record neuronal activity in the neocortex depends on the spatial and temporal scales and resolutions of interest, as well as whether the activity is to be recorded in vivo or in vitro. In this thesis, we are interested in understanding in vivo activity on the mesoscopic scale, covering cortical areas of several square millimeters; both on a single neurons level as well as on

a population level. Such recordings may be made using a variety of approaches, including voltage sensitive dyes (Chemla and Chavane, 2010) and two-photon calcium imaging (Grienberger and Konnerth, 2012). Another tool that naturally lends itself for investigating these types of activity is multi-electrode arrays (Maynard et al., 1997), which combine multiple micro-electrodes on a grid that covers an area of several square millimeters. The Utah array (Blackrock Microsystems) that was used to record the data presented in Chapter 3 had 10×10 electrodes equally spaced on a grid of $4 \times 4 \text{ mm}^2$ with an inter-electrode distance of $400 \mu\text{m}$ (Riehle et al., 2013, 2018; Brochier et al., 2018; Dąbrowska et al., 2020). These devices can record with a high temporal resolution, resolving the single-neuron activity of about a hundred neurons at the same time. Electrode arrays can be implanted chronically in living animals and they can be used to record from several cortical areas while the animal is performing a task. The signal recorded is the extracellular potential that is caused by the electrical activity of neurons close to and between the electrodes (Buzsáki et al., 2012). For our purposes, there are especially two interesting signal components: the population activity represented by the local field potential and the single neuron spiking activity.

For studying the population activity, the recorded signal usually is low-pass filtered to remove the signal's high frequency spike components, yielding the local field potential (LFP), which is a composite signal including contributions from entire populations of neurons (Buzsáki et al., 2012; Pesaran et al., 2018). The LFP observed in cortex exhibits spatio-temporally coherent population activity across several square millimeters: Activity can be oscillatory, where the recorded LFP collectively rises and falls (Buzsáki and Draguhn, 2004; Wang, 2010), with frequencies covering several frequency bands; from slow up- and down-state like activity $\sim 1 \text{ Hz}$ to high-gamma oscillations $\geq 80 \text{ Hz}$. The LFP signal can show planar traveling waves (Wu et al., 2008; Nauhaus et al., 2009; Muller and Destexhe, 2012; Nauhaus et al., 2012; Sato et al., 2012; Muller et al., 2018), with velocities ranging from $0.1 - 0.8 \text{ m/s}$ (Muller et al., 2018). Even more complex forms have been observed, such as source- or sink-like radial activity (Denker et al., 2018), or spirals (Townsend et al., 2015; Denker et al., 2018). According to Muller et al. (2018, Figure 2), stimuli can also cause static bump-like activity.

Extracting the spiking activity from electrode-array recordings is a little more intricate and requires a procedure known as spike-sorting (Lewicki, 1998; Quiñ Quiroga, 2012). A typical recorded signal shows peaks associated with nearby spiking neurons. To extract the spike times, the straightforward way is to use a simple threshold, but there exist much more elaborate methods (Lewicki, 1998; Quiñ Quiroga, 2012). As a single electrode can record spiking signals from multiple neurons in its vicinity, each spike needs to be allocated to its source neuron. The most reliable way to do this is a sorting by the typical spike wave forms (Bartho et al., 2004). This also allows a categorization into putative excitatory and inhibitory neurons, as pyramidal neurons typically have a broader spike width than inhibitory neurons (Csicsvari et al., 1998). Cortical spiking activity is highly irregular (Softky and Koch, 1993) and usually assumed to be Poisson-like (Tolhurst et al., 1983; Lee et al., 1998; Shadlen and Newsome, 1998; Koch, 1999; Maynard et al., 1999). But, there are areas with other spiking statistics (Maimon and Assad, 2009); even binary (DeWeese et al., 2003). Often cortical spiking activity is reported to be more reliable than a Poisson process (Nawrot

et al., 2008; Amarasingham et al., 2006; Kara et al., 2000; Gur et al., 1997). The typical rate of spiking activity in cortex is layer-dependent and in the range of $\sim 0.3 - 8$ Hz (Potjans and Diesmann, 2014, Table 6). Spiking activity has been reported to show spatial organization as well: Sato et al. (2012) and Takahashi et al. (2015) find traveling waves of spiking activity in cat visual cortex and macaque motor cortex. Smith and Kohn (2008) investigate the spike count correlation in macaque visual cortex and find a decay with inter-neuron distance. Hao et al. (2016) show that a single pulse microstimulus can evoke spreading spiking activity across a multi-electrode array. Finally, grid cells within entorhinal cortex in many different species are well-known for their hexagonal spiking activity (Rowland et al., 2016; Nadasdy et al., 2017).

1.3 NEURON AND NETWORK MODELS UTILIZED IN THIS THESIS

For modeling neuronal behavior, or any system, there are essentially two opposing approaches: One might aim for the most complete description, incorporating as many details as possible. This may allow the model to fit real data with high accuracy, but it also makes the model prone to overfitting and usually necessitates the use of numerical methods. Alternatively, one may decide to develop a minimal model that includes only what is required to explain the phenomena of interest. This naturally limits the model's explanatory power and may necessitate greater abstraction, but it does allow for the use of analytical tools, which may result in a more intuitive understanding of the system.

Applying the first approach to neurons yields biologically realistic neuron models, like the Hodgkin-Huxley model, which gives an accurate description of the generation of action potentials by modeling the behavior of different ion channels and the trans-membrane voltage (Hodgkin and Huxley, 1952); or multi-compartment models that incorporate the complex morphology of axons and dendrites (Rall, 1964).

However, in this thesis we focus on an analytical understanding of the models utilized, therefore our models fall within the scope of the second approach. Here, we introduce the models we used in our research in increasing order of abstractness.

LEAKY INTEGRATE-AND-FIRE MODELS Simplifying the dynamical equations of the Hodgkin-Huxley model by employing simplifying assumptions regarding the ion channel dynamics, one can arrive at integrate-and-fire models (Abbott and Kepler, 1990). These models describe the temporal dynamics of the membrane potential V by ordinary differential equations like

$$\tau_m \frac{dV}{dt}(t) = -f(V(t)) + RI(t),$$

with the membrane time constant τ_m , a potentially non-linear function f , the synaptic input current I , and membrane resistance R (Stein, 1967; Tuckwell, 1988a). Such models describe the integration of inputs $RI(t)$ according to a dynamical rule determined by $-f(V(t))$. They cannot produce action potentials. Instead, once a threshold voltage V_{th} is reached, the voltage is reset to a fixed voltage V_0 , and the neuron is said to have spiked once. As refractory behavior is not captured by this model, sometimes a fixed refractory time τ_r is defined in which the membrane voltage is kept at the

reset voltage after spiking. This is one of the simplest models of spiking neurons including sub-threshold dynamics. There exist many versions of these types of neuron models (Brunel and Latham, 2003; Fourcaud-Trocmé et al., 2003; Richardson, 2007, 2008; Grabska-Barwinska and Latham, 2014; Montbrió et al., 2015; Gerstner et al., 2014, Part II), but we will focus on the *leaky integrate-and-fire (LIF) neuron* (Lapicque, 1907; Stein, 1967), with linear sub-threshold dynamics

$$\tau_m \frac{dV}{dt}(t) = -(V - V_0) + RI(t) .$$

If the membrane voltage is excited, it will decay exponentially to the resting potential V_0 .

As reviewed by Izhikevich (2004), due to its simplicity the LIF model only captures a minimal set of biological properties of cortical spiking neurons; in particular, it cannot emulate any kind of adaptive spiking activity, such as bursting behavior. However, it is one of the simplest ordinary differential equation models for spiking neurons, making it especially suitable for proving analytical results concerning network dynamics for which single neuron dynamics may be neglected. Thus, it has become a standard model in theoretical neuroscience, utilized to study a multitude of network properties, for example reviewed by Burkitt (2006).

In networks of LIF neurons, each neuron has its own membrane potential V_i , labeled by an index $i \in \{1, \dots, N\}$, with the total number of neurons N . The dynamical equation becomes

$$\tau_m \frac{dV_i}{dt}(t) = -(V_i - V_0) + RI_i(t) .$$

For *instantaneous synapses*, also known as *delta synapses*, the synaptic input current to neuron i has its simplest form, which is given by

$$I_i(t) = \sum_j J_{ij} \sum_k \delta(t - t_j^k - d_{ij}) ,$$

whereas for *exponential synapses* it is described by a differential equation in itself

$$\tau_s \frac{dI_i}{dt}(t) = -I_i(t) + \sum_j J_{ij} \sum_k \delta(t - t_j^k - d_{ij}) ,$$

with the synaptic time constant τ_s , the Dirac delta function δ , spike times t_j^k of neuron j , and delay d_{ij} from neuron j to neuron i . For the purposes of this thesis, one of the most important aspects is the *connectivity matrix* J . It describes whether and how strongly neuron j projects to neuron i . So when there is a non-zero entry J_{ij} , neuron j is connected to neuron i , and the post synaptic current is abstracted into the value of J_{ij} , called *synaptic weight*. The connectivity also determines the type of the neuron. A neuron j only having negative values J_{ij} for all values of i is inhibitory and vice versa.

Embedding spatial structure into these models requires adapting the connectivity matrix. For example one can assign each neuron a place on a two dimensional grid with positions $\mathbf{r}$ and assert a connection strength decaying with the inter-neuronal distance on a length scale R , e.g.,

$$J_{ij} = J(|\mathbf{r}_i - \mathbf{r}_j|) \propto \exp\left(-\frac{|\mathbf{r}_i - \mathbf{r}_j|}{R}\right) . \quad (1.1)$$

RATE MODELS If solely the statistics of the neuronal activity is of interest, and modeling the behavior of spiking neuronal networks on short time-scales is not relevant, one can go a step further and consider rate models (Wilson and Cowan, 1972; Dayan and Abbott, 2001; Gerstner et al., 2014, Chapter 15). Instead of modeling every single spike, these models aim at describing the spike rate ν and hence offer a statistical description of the neuron activity averaged across a short time interval. A simple rate network model is described by

$$\tau \frac{dv_i}{dt}(t) = -v_i(t) + \sum_j J_{ij} \psi(v_j(t - d_{ij})) ,$$

where v_i is the firing rate of neuron i , τ is the neuronal time constant, $\psi(v)$ is the neurons' non-linear input-output function. The first term on the right-hand side describes the single neuron dynamics, while the second term represents the non-linearly transformed recurrent input, capturing the fact that a neuron requires a minimal input to fire and that the output rate saturates for high inputs. To account for stochasticity, the model can be extended by adding an input or output noise term.

By construction, rate models cannot provide information on subthreshold dynamics or accurate spike timings. Remarkably, they can yield similar findings as more intricate models (Grytskyy et al., 2013; Heiberg et al., 2018), yet, because of their simplicity, rate models lend themselves even better to the application of analytical methods than LIF models.

NEURAL FIELD MODELS When dealing with spatially organized networks with distance dependent connectivities, such as in Eq. (1.1), a natural analytical step is to take the continuum limit of the neuronal grid (Wilson and Cowan, 1973). This means that the network is no longer assumed to be a grid of single neurons, but rather a continuous excitable medium, known as *neural field*. Mathematically, taking the continuum limit amounts to replacing the discrete neuron index i by a continuous variable $\mathbf{r} \in \mathbb{R}^n$, the position on an n -dimensional spatial manifold Ω , and replacing the sum over the synaptic inputs by an integral. The firing rate vector $\mathbf{v}(t)$ is replaced by a continuous function of space, the network's activity $\mathbf{X}(\mathbf{r}, t)$, as is the connectivity matrix $\mathbf{M}(\mathbf{r})$. An example of a neural field equation for an excitatory and an inhibitory population of neurons is

$$\tau \frac{\partial \mathbf{X}}{\partial t}(\mathbf{r}, t) = -\mathbf{X}(\mathbf{r}, t) + \int_{\Omega} d^n r' \mathbf{M}(\mathbf{r} - \mathbf{r}') \cdot \psi \{ \mathbf{X}[\mathbf{r}', t - d(\mathbf{r} - \mathbf{r}')] \}$$

$$\mathbf{X}(\mathbf{r}, t) = \begin{pmatrix} E(\mathbf{r}, t) \\ I(\mathbf{r}, t) \end{pmatrix}, \quad \mathbf{M}(\mathbf{r}) = \begin{pmatrix} M_{EE}(\mathbf{r}) & M_{EI}(\mathbf{r}) \\ M_{IE}(\mathbf{r}) & M_{II}(\mathbf{r}) \end{pmatrix},$$

with a non-linearity ψ , and a possibly distance dependent delay $d(\mathbf{r})$.

Neural field models do not describe single neuron properties, and, similar to rate models, they neglect all effects associated with spiking dynamics. They have, however, been employed successfully to analyze spatio-temporally coherent population activity in cortical networks (Coombes, 2005; Bressloff, 2012; Coombes et al., 2014). We discuss the literature in Section 1.5.

1.4 MEAN-FIELD THEORY OF RANDOM NETWORKS

This section is partly based on the following publication's introduction, for which the author wrote the original draft (see chapter Chapter 4 for the full contributions statement):

Moritz Layer, Johanna Senk, Simon Essink, Alexander van Meegen, Hannah Bos, & Moritz Helias (2022), NNMT: Mean-field based analysis tools for neuronal network models, *Frontiers in Neuroinformatics*, DOI: 10.3389/fninf.2022.835657

Biological neuronal networks are composed of large numbers of recurrently connected neurons, with a single cortical neuron typically receiving synaptic inputs from thousands of other neurons (Braitenberg and Schüz, 1998; DeFelipe et al., 2002). Although the inputs of distinct neurons are integrated in a complex fashion, such large numbers of weak synaptic inputs imply that average properties of entire populations of neurons do not depend strongly on the contributions of individual neurons (Amit and Tsodyks, 1991). Based on this observation, it is possible to develop analytically tractable theories of population properties, in which the effects of individual neurons are averaged out and the complex, recurrent input to individual neurons is replaced by a self-consistent effective input (reviewed, e.g., in Gerstner et al., 2014). In classical physics terms (e.g., Goldenfeld, 1992), this effective input is called *mean-field*, because it is the self-consistent mean of a field, which here is just another name for the input the neuron is receiving. The term self-consistent refers to the fact that the population of neurons that receives the effective input is the same that contributes to this very input in a recurrent fashion: the population's output determines its input and vice-versa. The stationary statistics of the effective input therefore can be found in a self-consistent manner: the input to a neuron must be set exactly such that the caused output leads to the respective input.

To state the basic principle of mean-field theory more precisely: A cortical neuron usually gets a large number of synaptic inputs. If those inputs are only weakly correlated, the central limit theorem implies that the synaptic inputs can be approximated as Gaussian white noise, with mean and noise intensity determined by the synaptic weights, the number of inputs, and the average firing rates. If a network of N neurons contains n populations of neurons with identical properties, a mean-field approximation reduces the original N dynamical equations to effectively n dynamical equations, significantly simplifying the problem. In particular, for some neuron models this simplification allows expressing the population firing rates as functions of the noise properties, such that average firing rates and noise properties may be determined self-consistently.

Mean-field theories have been developed for random networks, without any particular connectivity structure, for a variety of neuron models. They have been successfully applied to study average population firing rates (van Vreeswijk and Sompolinsky, 1996, 1998; Amit and Brunel, 1997b), and the various activity states a random network of spiking neurons can exhibit, depending on the network parameters (Amit and Brunel, 1997a; Brunel, 2000a; Ostojic, 2014), as well as the effects that different kinds of synapses have on firing rates (Fourcaud and Brunel, 2002; Lindner, 2004; Schuecker

et al., 2015; Schwalger et al., 2015; Mattia et al., 2019). In combination with linear response theory, they have been used to investigate how neuronal networks respond to external inputs (Lindner and Schimansky-Geier, 2001; Lindner et al., 2005), and they explain why neuronal networks can track external input on much faster time scales than a single neuron could (van Vreeswijk and Sompolinsky, 1996, 1998). Mean-field with linear-response theories allow studying correlations of neuronal activity (Sejnowski, 1976; Ginzburg and Sompolinsky, 1994; Lindner et al., 2005; Trousdale et al., 2012), which can be generalized to correlations of higher orders (Buice and Chow, 2013), and they were able to reveal why pairs of neurons in random networks, despite receiving a high proportion of common input, can show low output correlations (Hertz, 2010; Renart et al., 2010; Tetzlaff et al., 2012; Helias et al., 2014), which for example has important implications for information processing. Mean-field theories were utilized to show that neuronal networks can exhibit chaotic dynamics (Sompolinsky et al., 1988; van Vreeswijk and Sompolinsky, 1996, 1998), in which two slightly different initial states can lead to totally different network responses, which has been linked to the network's memory capacity (Toyoizumi and Abbott, 2011; Schuecker et al., 2018). Most of the results mentioned above have been derived for networks of either rate, binary, or spiking neurons of a linear integrate-and-fire type. But a multitude of other models has been investigated with similar tools as well; for example, just to mention a few, Hawkes processes, non-linear integrate-and-fire neurons (Brunel and Latham, 2003; Fourcaud-Trocmé et al., 2003; Richardson, 2007, 2008; Grabska-Barwinska and Latham, 2014; Montbrió et al., 2015), or Kuramoto-type models (Stiller and Radons, 1998; van Meegen and Lindner, 2018). Additionally, there is an ongoing effort showing that many of the results derived for distinct models are indeed equivalent and that those models can be mapped to each other under certain circumstances (Grytskyy et al., 2013; Ostojic and Brunel, 2011).

The mean-field approach is limited by the requirement for the central limit theorem to hold, namely a large number of uncorrelated inputs. Assuming neurons do have a large number of incoming connections, like in cortical networks, the validity of the mean-field approach is determined by the correlation of the inputs. If the network exhibits strongly correlated activity, for example caused by strong oscillations or external input, the mean-field approximation breaks down, whereas small oscillations of the population activity around a stationary state can be explained using linear-response theory on top of the mean-field approximation (Lindner and Schimansky-Geier, 2001; Lindner et al., 2005; Schuecker et al., 2015). The approximation particularly works well for asynchronous, irregular network states (Brunel, 2000b). van Vreeswijk and Sompolinsky (1996, 1998) demonstrate analytically that such activity arises dynamically in large, sparsely connected networks of excitatory and inhibitory neurons if the synaptic weights scale appropriately. Their argument is mainly based on an approximate balance of excitation and inhibition, and these dynamics, known as *balanced* state, are a very robust phenomenon that are found for a wide range of neuron models.

1.5 THEORY OF SPATIALLY ORGANIZED NETWORKS

The preceding section's literature review demonstrates that random unstructured network models have received a lot of attention, and their theory is well-established by now. In a largely independent stream of research, neuronal networks with spatially structured connectivities have been studied utilizing the continuous neural field models introduced in Section 1.4, which neglect single neuron and spiking properties. They have been used to investigate spatio-temporally coherent population activity in cortical networks with great success. Wilson and Cowan (1973) were the first to study global activity patterns, whereas Amari (1975, 1977) was the first to examine stable localized activity. Following that, neural field theories were able to explain hallucination patterns (Ermentrout and Cowan, 1979; Bressloff et al., 2001) and oscillatory activity observed in EEG and MEG recordings (Nunez, 1974; Jirsa and Haken, 1996, 1997). Working memory (Laing et al., 2002; Laing and Troy, 2003), motion perception (Giese, 2012), cognition (Schöner, 2008) are just a few examples of phenomena that have been modeled using neural field theories. The field is fairly extensive, and Chapter 4 discusses further neural field literature with a focus on delayed systems; for excellent reviews, we recommend Coombes (2005), Bressloff (2012), and Coombes et al. (2014).

The neural field approach describes the activity of a continuous excitable medium. Real neuronal networks, however, consist of distinct units, the neurons. The theory of discrete models of spatially organized networks has received far less attention than neural field theories or random network theory, and it has only lately seen some progress.

Simulations reveal that two-dimensional networks of spiking neurons exhibit complex, moving patches of activity, as well as regular periodic or wavelike activity, depending on whether excitation and inhibition are balanced or not (Usher et al., 1995; Keane and Gong, 2015). Bressloff and Coombes (1998, 2000) show that increasing the coupling strength in spiking networks with Mexican hat connectivity can lead to a desynchronization of firing times via a Turing-Hopf bifurcation. Murphy and Miller (2009) explore a mechanism explaining the selective amplification of activity patterns found in spontaneous visual cortex activity. Kriener et al. (2014) show that the critical coupling strength that determines pattern formation in spiking networks strongly depends on the input regime. Lim and Goldman (2014) investigate working memory in a LIF ring model with distance dependent connectivity, demonstrating that these networks can retain amplitude and spatial location of a stimulus. Pyle and Rosenbaum (2017) study pattern formation and reservoir computation in two-dimensional spiking networks and argue that, compared to random networks, including spatial organization improves the network's computational properties.

In addition to direct analytical analyses of discrete models, there is an ongoing effort to map the parameters of more complex spiking models to simpler matching rate or neural field models, allowing for the use of the well-developed neural field theory; Senk et al. (2020) provide a review of such attempts.

Another focus of the literature on spatially organized networks is correlations and their spatial configuration. Kriener et al. (2009) compare correlations in random, small world, and ring networks and show that the latter two exhibit heavy tailed correlation distributions. Rosenbaum and Doiron (2014) model the situation in sensory areas

and derive conditions on the spatial spread of recurrent and external connections for a stable balanced state in networks of integrate-and-fire neurons with distance-dependent connectivity and common external input. They continue their research in Rosenbaum et al. (2017), to study the distance-dependence of mean covariances in such networks. Darshan et al. (2018) explain how neuronal activity may be both temporally irregular and significantly spatially correlated by relating the strength of pair-wise correlations to the architecture of a network of binary neurons. Smith et al. (2018) demonstrate that anisotropy of connection profiles can result in long-range correlations, akin to those observed in visual cortex in the context of orientation tuning, from just short-range connections, and Huang et al. (2019) show that spatial organization is required for a spiking network model to show asynchronous activity if one incorporates that inhibitory synapses are slower than excitatory ones.

1.6 SCOPE OF THIS THESIS

This thesis intends to advance the understanding of spatially organized networks with a focus on the mesoscopic scale by means of analytical theory, its numerical implementation, and simulations. To this end, we present three complementary studies, each covering a different aspect of the theory of spatially organized networks. In Chapter 2, we present a Python package we developed for collecting numerical implementations of analytical tools based on mean-field theory (Section 1.4), which enables the computation of population properties such as firing rates or power spectra in concrete network model configurations. Continuing the investigation of the spatial structure of correlations reviewed in Section 1.5, we study pair-wise spike count correlations in Macaque motor cortex in Chapter 3 and present a network mechanism that explains how the observed long-range correlations can emerge in networks with short-range connectivity and uncorrelated input. In Chapter 4, we analytically investigate global spatio-temporally coherent activity pattern formation in planar two-dimensional networks with a constant synaptic delay.

These studies have the common objective of investigating network activity at the millimeter scale, as observable using multi-electrode arrays like the Utah array (Section 1.2). While the toolbox of the first study is applicable to a wide range of spatially organized, as well as random network models, the second and third study consider very similar spatially organized network models: both analyze the spontaneous activity of an excitatory and an inhibitory neuronal population with target-agnostic connectivity and population specific spatially extended connectivity profiles on a two-dimensional grid; the model in Chapter 4, however, includes a constant synaptic delay. While in Chapter 3 we study pair-wise correlations in a spatially homogeneous but almost critical balanced state, in Chapter 4 we investigate the spontaneous emergence of spatio-temporally coherent periodic, oscillatory, or wavelike activity in slightly super-critical networks. The two studies illustrate different network mechanisms that lead to coordinated neuronal activity across scales that are much larger than the range of typical direct connections between the neurons. The specific objectives of the three studies are presented in the following paragraphs.

NNMT: MEAN-FIELD BASED ANALYSIS TOOLS FOR NEURONAL NETWORK MODELS Chapter 2 presents an open-source, community centered Python package we developed for collecting reusable implementations of analytical methods for neuronal network model analysis based on mean-field theory. Because mean-field theories describe the mean activity of thousands of neurons and provide a coarse-grained perspective of brain dynamics, they are particularly important for understanding network activity statistics on the mesoscopic scale. Until previously, analytical methods for analyzing network models were implemented on a project-by-project basis, which meant that they were frequently quite specific, making them rarely reusable for other purposes, and they had not been collected and validated centrally. As we rely on correct implementations of existing methods for developing new ones and cross-checking them, we decided to introduce a well-tested reference collection. Therefore, we generalized existing implementations, collected them in one common framework, and developed a thorough testing procedure covering all included methods. We made this toolbox publicly available and invited other researchers to contribute their methods as well. This chapter describes the toolbox architecture, its features, and how to apply the tools. We provide an illustrative example of how to analyze a microcircuit model. We compare the package to other tools for network model analysis, explain limitations and use-cases, and discuss what and how new methods can be added.

CORRELATIONS IN TWO-DIMENSIONAL NETWORKS Chapter 3 investigates the effect of lateral connectivity on the coordination of neurons across a mesoscopic scale in motor cortex, leveraging a combination of electrode array recordings in Macaque monkey, network model simulation, and network theory. Concretely, we demonstrate that a heterogeneous non-patchy distance-dependent local connectivity on the range of a few hundred micrometers can explain the experimentally observed long-range coordination of neuronal activity across several millimeters. Furthermore, we explain how the coordination pattern can be altered during behavior, allowing the network to flexibly adapt which neurons are coordinated to respond to changing momentary demands.

ACTIVITY PATTERNS IN TWO-DIMENSIONAL NETWORKS Chapter 4 examines how a synaptic delay as well as the neurons' non-linear input-output functions influence spatio-temporal pattern formation in planar two-dimensional networks with distance-dependent connectivity. Spatio-temporal activity patterns, such as oscillations or waves, are ubiquitous in cortical recordings and can span large parts of cortex, often involving long time constants (e.g., slow oscillations). Analyzing the critical behavior of a respective neuronal network in two dimensions is a formidable task, due to accompanying mathematical intricacies, but some previous work exists for two-dimensional networks without a synaptic delay. A delay introduces additional mathematical difficulties, as it requires the analysis of functional differential equations, for which standard techniques are not applicable. However, we present a way to map the network with delay to a non-delayed network, allowing a direct comparison of the two systems and therefore clearly discerning the influence of the delay. We apply the analysis to static patterns and show that a delay leads to longer transient time-scales. We validate our analytical results using a simulation of rate neurons.

NNMT: MEAN-FIELD BASED ANALYSIS TOOLS FOR NEURONAL NETWORK MODELS

This chapter is based on the following publication:

Moritz Layer, Johanna Senk, Simon Essink, Alexander van Meegen, Hannah Bos, & Moritz Helias (2022), NNMT: Mean-field based analysis tools for neuronal network models, *Frontiers in Neuroinformatics*, DOI: 10.3389/fninf.2022.835657

Author contributions:

Under the supervision of Johanna Senk and Moritz Helias, the author developed and implemented the current version of the toolbox, including the test suite and the online documentation. The author wrote the original draft of all sections this chapter is based on (Sections 1, 2, 3.1, 3.2.2, and 4), except for the power spectra example, which is included in this thesis for completeness. The example is partly based on Section 3.3.2 of the aforementioned publication, for which Simon Essink wrote the original draft, and it was finalized in close collaboration with the author. Panel C of Figure 2.2 on page 19 was originally created by Hannah Bos in Bos et al. (2016) and reproduced using NNMT by Simon Essink for the aforementioned publication. All authors were involved in the design of the toolbox, made contributions to the toolbox, and participated in writing the manuscript.

2.1 INTRODUCTION

Section 1.4 demonstrates that analytical theories have clearly contributed to our understanding of neuronal networks and they offer a plethora of powerful and efficient methods for network model analysis. Comparing the predictions of analytical theories to simulations, experimental data, or other theories necessitates a numerical implementation applicable to various network models, depending on the research question. Such an implementation is often far from straightforward and at times requires investing substantial time and effort. Commonly, such tools are implemented as the need arises, and their reuse is not organized systematically and restricted to within a single lab. This way, not only are effort and costs spent by the neuroscientific community duplicated over and over again, but also are many scientists deterred from taking maximal advantage of those methods although they might open new avenues for investigating their research questions.

In order to make analytical tools for neuronal network model analysis accessible to a wider part of the neuroscientific community, and to create a platform for collecting well-tested and validated implementations of such tools, we have developed the Python toolbox NNMT (Layer et al., 2021), short for Neuronal Network Mean-field Toolbox. We would like to emphasize that NNMT is not a simulation tool; NNMT is a collection of numerically solved mean-field equations that directly relate the parameters of a microscopic network model to the statistics of its dynamics. NNMT has been designed to fit the diversity of mean-field theories, and the key features we are aiming for are modularity, extensibility, and a simple usability. Furthermore, it features an extensive test suite to ensure the validity of the implementations as well as a comprehensive user documentation. The current version of NNMT mainly comprises tools for investigating networks of leaky integrate-and-fire neurons as well as some methods for studying binary neurons and neural field models. The toolbox is open-source and publicly available on GitHub¹.

In the following, we present the design considerations that led to the structure and implementation of NNMT as well as an illustrative use case. Section 2.2 first introduces its architecture as well as the interfaces for accessing the tools. Section 2.3 then explains its usage by analyzing the cortical microcircuit model by Potjans and Diesmann (2014) and reproducing previously published results from Bos et al. (2016). Section 2.4 compares NNMT to other available toolboxes for neuronal network model analysis, discusses its use cases from a more general perspective, indicates current limitations and prospective advancements of NNMT, and explains how new tools can be contributed.

2.2 WORKFLOWS AND ARCHITECTURE

```

1 # basic workflow
  result = nnmt.<submodule>.<_tool>(*args, **kwargs)

# model workflow
my_model = nnmt.models.<model>(<network_params>, <analysis_params>)
6 result = nnmt.<submodule>.<tool>(my_model)

```

Listing 1: The two modes of using NNMT: In the basic workflow (top), quantities are calculated by passing all required arguments directly to the underscored tool functions available in the submodules of NNMT. In the model workflow (bottom), a model class is instantiated with parameter sets and the model instance is passed to the non-underscored tool functions which automatically extract the relevant parameters.

What are the requirements a package for collecting analytical methods for neuronal network model analysis needs to fulfill? To begin with, it should be adaptable and modular enough to accommodate many and diverse analytical methods while avoiding code repetition and a complex interdependency of package components. It should

¹ <https://github.com/INM-6/nnmt>




Figure 2.1: Structure and workflows of the Neuronal Network Mean-field Toolbox (NNMT).

(A) Basic workflow: individual mean-field based analysis methods are implemented as functions, called `_tools()`, that can be used directly by explicitly passing the required arguments. (B) Model workflow: to facilitate the handling of parameters and results, they can be stored in a model class instance, which can be passed to a `tool()`, which wraps the basic workflow of the respective `_tool()`. (C) Structure of the Python package. In addition to the tool collection (red frame), containing the `tools()` and the `_tools()`, and pre-defined model classes, the package provides utility functions for handling parameter files and unit conversions, as well as software aiding the implementation of new methods.

enable the application of the collected algorithms to various network models in a simple and transparent manner. It should make the tools easy to use for new users, while also providing experts with direct access to all parameters and options. Finally, the methods need to be thoroughly tested and well documented.

These are the main considerations that guided the development of NNMT. Panels A and B of Figure 2.1 on page 15 illustrate how the toolbox can be used in two different workflows, depending on the preferences and goals of the user. In the *basic workflow* the individual method implementations called *tools* are directly accessed, whereas the *model workflow* provides additional functionality for the handling of parameters and results.

2.2.1 Basic workflow

The core of NNMT is a collection of low-level functions that take specific parameters (or pre-computed results) as input arguments and return analytical results of network properties. In Panel A of Figure 2.1 on page 15, we refer to such basic functions as `_tools()`, as their names always start with an underscore. We term this lightweight approach of directly using these functions the basic workflow. The top part of Listing 1 demonstrates this usage; for example, the quantity to be computed could be the mean firing rate of a neuronal population and the arguments could be parameters which define neuron model and external drive. While the basic workflow gives full flexibility and direct access to every parameter of the calculation, it remains the user's responsibility to insert the arguments correctly, e.g., in the right units.

2.2.2 Model workflow

The model workflow is a convenient wrapper of the basic workflow (Panel B of Figure 2.1 on page 15). A *model* in this context is an object that stores a larger set of parameters and can be passed directly to a `tool()`, the non-underscored wrapper of the respective `_tool()`. The `tool()` automatically extracts the relevant parameters from the model, passes them as arguments to the corresponding core function `_tool()`, returns the results, and stores them in the model. The bottom part of Listing 1 shows how a model is initialized with parameters and then passed to a `tool()` function.

Models are implemented as Python classes and can be found in the submodule `nnmt.models`. We provide the class `nnmt.models.Network` as a parent class and a few child classes which inherit the generic methods and properties but are tailored to specific network models; custom models can be created straightforwardly. The parameters distinguish network parameters, which define neuron models and network connectivity, and analysis parameters; an example for an analysis parameter is a frequency range over which a function is evaluated. Upon model instantiation, parameter sets defining values and corresponding units are passed as Python dictionaries or `yaml` files. The model constructor takes care of reading in these parameters, computing dependent parameters from the imported parameters, and converting all units to SI units for internal computations. Consequently, the parameters passed as arguments and the functions for computing dependent parameters of a specific child class need to

be aligned. This design encourages a clear separation between a concise set of base parameters and functionality that transforms these parameters to the generic (vectorized) format that the tools work with. To illustrate this, consider the synaptic weight matrix of a network of excitatory and inhibitory neuron populations in which all excitatory connections have the same weight and all inhibitory ones another weight. As argument one could pass just a tuple of two different weight values and the corresponding model class would take care of constructing the full weight matrix. This happens in the example presented in Section 2.3.2: The parameter file `network_params_microcircuit.yaml` contains the excitatory synaptic weight and the ratio of inhibitory to excitatory weights. On instantiation, the full weight matrix is constructed from these two parameters, following the rules defined in `nnmt.models.Microcircuit`.

When a `tool()` is called, it checks whether the provided model object contains all required parameters and previously computed results. Then the `tool()` extracts the required arguments, calls the respective `_tool()`, and caches and returns the result. If the user attempts to compute the same property twice, using identical parameters, the `tool()` will retrieve the already computed result from the model's cache and return that value. Results can be exported to an HDF5 file and also loaded.

Using the model workflow instead of the basic workflow comes with the initial overhead of choosing a suitable combination of parameters and a model class, but has the advantages of a higher level of automation with built-in mechanisms for checking correctness of input (e.g., regarding units), reduced redundancy, and the options to store and load results. Both modes of using the toolbox can also be combined.

2.2.3 *Structure of the toolbox*

The structure of the Python package NNMT is depicted in Panel C of Figure 2.1 on page 15. It is divided into submodules containing the tools (e.g., `nnmt.lif.exp`, or `nnmt.binary`), the model classes (`nnmt.models`), helper routines for handling parameter files and unit conversions, as well as modules that collect reusable code employed in implementations for multiple neuron models (cf. Section 2.4.4). The tools are organized in a modular, extensible fashion with a streamlined hierarchy. To give an example, a large part of the currently implemented tools apply to networks of leaky integrate-and-fire (LIF) neurons, and they are located in the submodule `nnmt.lif`. The mean-field theory for networks of LIF neurons distinguishes between neurons with instantaneous synapses, also called delta synapses, and those with exponentially decaying post-synaptic currents. Similarly, the submodule for LIF neurons is split further into the two submodules `nnmt.lif.delta` and `nnmt.lif.exp`. NNMT also collects different implementations for computing the same quantity using different approximations or numerics, allowing for a comparison of different approaches.

Apart from the core package, NNMT comes with an extensive online documentation², including a quickstart tutorial, all examples presented in the original publication (Layer et al., 2022), a complete documentation of all tools, as well as a guide for contributors.

² <https://nnmt.readthedocs.io/>

Furthermore, we provide an extensive test suite that validates the tools by checking them against previously published results and alternative implementations where possible. This ensures that future improvements of the numerics do not break the tools.

2.3 HOW TO USE THE TOOLBOX

In this section, we demonstrate the practical use of NNMT by replicating previously published results. The example presented has been chosen to cover both stationary and dynamic network features, as mean-field theory is typically divided into two parts: stationary theory, which describes time-independent network properties of systems in a stationary state, and dynamical theory, which describes time-dependent network properties. The original paper (Layer et al., 2022) presents further examples, and all examples, including the used parameter files, are part of the online documentation².

2.3.1 Installation and setup

The toolbox can be either installed using pip:

```
pip install nnmt
```

or by installing it directly from the repository, which is described in detail in the online documentation. After the installation, the module can be imported:

```
import nnmt
```

2.3.2 Usage example: analyzing a microcircuit model

THE MICROCIRCUIT MODEL Here we show how to use the model workflow to calculate the firing rates and power spectra of the cortical microcircuit model by Potjans and Diesmann (2014). The circuit is a simplified point neuron network model with biologically plausible parameters, which has been recently used in a number of other works: for example, to study network properties such as layer-dependent attentional processing (Wagatsuma et al., 2011), connectivity structure with respect to oscillations (Bos et al., 2016), and the effect of synaptic weight resolution on activity statistics (Dasbach, Tetzlaff, Diesmann, and Senk, 2021); to assess the performance of different simulator technologies such as neuromorphic hardware (van Albada et al., 2018) and GPUs (Knight and Nowotny, 2018; Golosio et al., 2021); to demonstrate forward-model prediction of local-field potentials from spiking activity (Hagen et al., 2016); and to serve as a building block for large-scale models (Schmidt et al., 2018).

The model consists of eight populations of LIF neurons (Section 1.3), corresponding to the excitatory and inhibitory populations of four cortical layers: 2/3E, 2/3I, 4E, 4I, 5E, 5I, 6E, and 6I (see Panel A of Figure 2.2 on page 19). It defines the number of neurons in each population, the number of connections between the populations, the single neuron properties, and the external input. Simulations show that the model yields realistic firing rates for the different populations as observed in particular in the healthy resting-state of early sensory cortex (Potjans and Diesmann, 2014, Table 6).




Figure 2.2: Cortical microcircuit model by Potjans and Diesmann (2014). **(A)** Network diagram with nodes and edges according to the graphical notation proposed by Senk et al. (2021); only the strongest connections are shown as in Figure 1 of the original publication. **(B)** Simulation and mean-field estimate for average population firing rates using the parameters from Bos et al. (2016). **(C)** Power spectrum of the population spiking activity in population 2/3E of the microcircuit with the parameters from Bos et al. (2016). The spiking activity of each population in a 10 s simulation of the model is binned with 1 ms resolution and the power spectrum of the resulting histogram is calculated by a fast Fourier transform (FFT; light gray curves). In addition, the simulation is split into 500 ms windows, the power spectrum calculated for each window and averaged across windows (gray curves). Black curves correspond to analytical prediction obtained with NNMT.

THE IMPLEMENTED ANALYTICAL RESULTS As explained in Section 1.4, the mean-field approach approximates the input to a neuron of population a by Gaussian white noise with mean μ_a and noise intensity σ_a , which in turn are determined by the network properties. For LIF neurons one finds (Amit and Brunel, 1997b)

$$\mu_a = \tau_m \sum_b J_{ab} K_{ab} \nu_b, \quad \sigma_a^2 = \tau_m \sum_b J_{ab}^2 K_{ab} \nu_b,$$

where the sum runs over all populations, including external populations, and with the average number of inputs K_{ab} from population b to population a . The firing rate of a LIF neuron with instantaneous synapses driven by such Gaussian white noise is described by (Siegert, 1951; Tuckwell, 1988b; Amit and Brunel, 1997b)

$$\phi(\mu, \sigma) = \left(\tau_r + \tau_m \sqrt{\pi} \int_{\tilde{V}_0(\mu, \sigma)}^{\tilde{V}_{th}(\mu, \sigma)} e^{s^2} [1 + \operatorname{erf}(s)] ds \right)^{-1}, \quad (2.1)$$

with erf denoting the error function, and rescaled reset- and threshold-voltages

$$\tilde{V}_0(\mu, \sigma) = \frac{V_0 - \mu}{\sigma}, \quad \tilde{V}_{th}(\mu, \sigma) = \frac{V_{th} - \mu}{\sigma}. \quad (2.2)$$

The firing rates are then determined by the self-consistency equation

$$\nu = \phi(\mu, \sigma), \quad (2.3)$$

which may be solved numerically, for example using one of two algorithms included in our toolbox (Layer et al., 2022).

The LIF neurons in the microcircuit model, however, have exponentially shaped post-synaptic currents. Fourcaud and Brunel (2002) developed a method for calculating the firing rate for this synapse type. They have shown that, if the synaptic time constant τ_s is much smaller than the membrane time constant τ_m , the firing rate for LIF neurons with exponential synapses can be calculated using the results for delta neurons by adjusting the shifted integration boundaries

$$\tilde{V}_0(\mu, \sigma) = \frac{V_0 - \mu}{\sigma} + \frac{\alpha}{2} \sqrt{\frac{\tau_s}{\tau_m}}, \quad \tilde{V}_{th}(\mu, \sigma) = \frac{V_{th} - \mu}{\sigma} + \frac{\alpha}{2} \sqrt{\frac{\tau_s}{\tau_m}}, \quad (2.4)$$

with $\alpha = \sqrt{2} |\zeta(1/2)| \approx 2.07$, where $\zeta(x)$ denotes the Riemann zeta function.

Based on the mean-field approximation, linear response theory allows computing the power spectra depending on the network architecture, stationary firing rates, and the neurons' transfer function (Bos et al., 2016), which characterizes the network's response to oscillatory input. All of this can be readily computed using NNMT, which we will demonstrate in the following paragraphs.

COMPUTATION OF FIRING RATES The model `nnmt.models.Microcircuit` implements the microcircuit in NNMT. We here use the parameters of the circuit as published in Bos et al. (2016), which is slightly differently parameterized than the original model (Appendix Table A.1 on page 93). The parameters of the model are specified in a `yaml` file, which uses Python-like indentation and a dictionary-style syntax. List elements are indicated by hyphens, and arrays can be defined as nested lists. Parameters with

units can be defined by using the keys `val` and `unit`, whereas unitless variables can be defined without any keys. Listing 2 shows an example of how some of the microcircuit network parameters used here are defined. Which parameters need to be provided in the `yaml` file depends on the model used and is indicated in their respective docstrings.

```
# membrane time constant
tau_m:
  val: 10.0
  unit: ms

# neuron numbers
N:
  - 20683
  - 5834
  - 21915
```

Listing 2: Some microcircuit network parameters defined in a `yaml` file. A dictionary-like structure with the keys `val` (value) and `unit` is used to define the membrane time constant, which is the same across all populations. The numbers of neurons in each population are defined as a list. Only the numbers for the first three populations are displayed.

Once the parameters are defined, a microcircuit model is instantiated by passing the respective parameter file to the model constructor; the units are automatically converted to SI units. Then the firing rates are computed. For comparison, we finally load the simulated rates from Bos et al. (2016):

```
# create the network model using a network parameter yaml file
microcircuit = nnmt.models.Microcircuit('network_params.yaml',
                                         'analysis_params.yaml')

# calculate firing rates
firing_rates = nnmt.lif.exp.firing_rates(microcircuit)

# load simulated results
simulated_firing_rates = (
    nnmt.input_output.load_h5('Bos2016_results.h5')['rates'])
```

The simulated rates have been obtained by a numerical network simulation (for simulation details see Bos et al., 2016) in which the neuron populations are connected according to the model’s original connectivity rule: “random, fixed total number with multapses (autapses prohibited)”, see Senk et al. (2021) as a reference for connectivity concepts. The term *multapses* refers to multiple connections between the same pair of neurons and *autapses* are self-connections; with this connectivity rule multapses can occur in a network realization but autapses are not allowed. For simplicity, the theoretical predictions assume a connectivity with a fixed in-degree for each neuron. Dasbach et al. (2021) show that simulated spike activity data of networks with these two different connectivity rules are characterized by differently shaped rate distributions (“reference” in their Figures 3d and 4d). In addition, the weights in the

simulation are normally distributed while the theory replaces each distribution by its mean; this corresponds to the case $N_{\text{bins}} = 1$ in Dasbach et al. (2021). Nevertheless, our mean-field theoretical estimate of the average population firing rates is in good agreement with the simulated rates (Panel B of Figure 2.2 on page 19).

COMPUTATION OF POWER SPECTRUM Starting with the previously defined microcircuit, the toolbox’s modular structure in combination with the model workflow permits a step-by-step calculation of the power spectra:

```
# calculate required quantities
nnmt.lif.exp.working_point(microcircuit)
nnmt.lif.exp.transfer_function(microcircuit)
nnmt.network_properties.delay_dist_matrix(microcircuit)
nnmt.lif.exp.effective_connectivity(microcircuit)

# calculate power spectra
power_spectra = nnmt.lif.exp.power_spectra(microcircuit)
```

The inherent structure of the theory is emphasized in these steps: After instantiating the network model class with given network parameters, we determine the working point, which characterizes the statistics of the model’s stationary dynamics. It is defined by the population firing rates, the mean, and the standard deviation of the input to a neuron of the respective population. This is necessary for determining the transfer functions. The calculation of the delay distribution matrix is then required for calculating the effective connectivity and to finally get an estimate of the power spectra. Panel C of Figure 2.2 on page 19 reproduces the top left panel of Figure 1E in Bos et al. (2016) and shows the spectrum for population 2/3E of the adjusted version of the microcircuit model.

The numerical predictions obtained from the toolbox largely coincide with simulated data taken from the original publication (Bos et al., 2016) and reveal dominant oscillations of the population activities in the low- γ range around 63 Hz. Furthermore, faster oscillations with peak power around 300 Hz are predicted with higher magnitudes in the inhibitory populations 4I, 5I, and 6I.

The deviation between predicted and simulated power spectra seen at ~ 130 Hz in population 2/3E could be a harmonic of the correctly predicted, prominent 63 Hz peak; a non-linear effect not captured in linear response theory. Furthermore, the systematic overestimation of the power spectrum at large frequencies is explained by the limited validity of the analytical approximation of the transfer function for high frequencies.

FURTHER POSSIBILITIES NNMT offers a lot more options for analyzing and developing a network model, and we present a variety of them in the original publication Layer et al. (2022). For example, we show that using the included sensitivity measure (Bos et al., 2016), which provides an analytical relation between the network model’s connectivity and the peaks of the power spectra, we can systematically modify the frequency and amplitude of the low- γ peak of the microcircuit model. Furthermore, we demonstrate how the toolbox can be utilized to investigate a network model’s

non-linear response behavior to external input (Sanzeni et al., 2020), and — worth highlighting in the context of this thesis — we present a method for mapping a LIF network model to a simpler rate network model (Senk et al., 2020), which allows investigating spatio-temporal pattern formation properties with the help of neural field theory, as done in Chapter 4.

2.4 DISCUSSION

Mean-field theory grants important insights into the dynamics of neuronal networks. However, the lack of a publicly available numerical implementation for most methods entails a significant initial investment of time and effort prior to any scientific investigations. In this chapter, we present the open-source toolbox NNMT, which currently focuses on methods for LIF neurons but is intended as a platform for collecting standard implementations of various neuronal network model analyses based on mean-field theory that have been thoroughly tested and validated by the neuroscientific community (Riquelme and Gjorgjieva, 2021). As an illustrative use case, we examine the cortical microcircuit model of Potjans and Diesmann (2014), compute its firing rates and power spectra using NNMT, and compare the findings to known simulation results from the literature (Bos et al., 2016).

In the remainder of the discussion, we compare NNMT to other tools suited for network model analysis. We expand on the different use cases of NNMT and also point out the inherent limitations of analytical methods for neuronal network analysis. We conclude with suggestions on how new tools can be added to NNMT and how the toolbox may grow and develop in the future.

2.4.1 *Comparison to other tools*

There are various approaches and corresponding tools that can help to gain a better understanding of a neuronal network model. There are numerous simulators that numerically solve the dynamical equations for concrete realizations of a network model and all its stochastic components, often focusing either on the resolution of single-neurons, for example NEST (Gewaltig and Diesmann, 2007), Brian (Stimberg et al., 2019), or Neuron (Hines and Carnevale, 2001), or on the population level, for example TheVirtualBrain (Sanz Leon et al., 2013). Similarly, general-purpose dynamical system software like XPPAUT (Ermentrout, 2002) can be used. Simulation tools, like DynaSim (Sherfey et al., 2018), come with enhanced functionality for simplifying batch analysis and parameter explorations. All these tools yield time-series of activity, and some of them even provide the methods for analyzing the generated data. However, simulations only indirectly link a model’s parameters with its activity: to gain an understanding of how a model’s parameters influence the statistics of their activity, it is necessary to run many simulations with different parameters and analyze the generated data subsequently.

Other approaches provide a more direct insight into a model’s behavior on an abstract level: TheVirtualBrain and the Brain Dynamics Toolbox (Heitmann et al., 2018), for example, allow plotting a model’s phase space vector field while the

parameters can be changed interactively, allowing for exploration of low-dimensional systems defined by differential equations without the need for simulations. XPPAUT has an interface to AUTO-07P (Doedel and Oldeman, 1998), a software for performing numerical bifurcation and continuation analysis. It is worth noting that such tools are limited to models that are defined in terms of differential equations. Models specified in terms of update rules, such as binary neurons, need to be analyzed differently, for example using mean-field theory.

A third approach is to simplify the model analytically and simulate the simplified version. The simulation platform DiPDE³ utilizes the population density approach to simulate the statistical evolution of a network model's dynamics. Schwalger et al. (2017) start from a microscopic model of generalized integrate-and-fire neurons and derive mesoscopic mean-field population equations, which reproduce the statistical and qualitative behavior of the homogeneous neuronal sub-populations. Similarly, Montbrió et al. (2015) derive a set of non-linear differential equations describing the dynamics of the rate and mean membrane potentials of a population of quadratic integrate-and-fire (QIF) neurons. The simulation platform PyRates (Gast et al., 2019) provides an implementation of this QIF mean-field model, and allows simulating them to obtain the temporal evolution of the population activity measures.

However, mean-field and related theories can go beyond such reduced dynamical equations: they can directly link model parameters to activity statistics, and they can even provide access to informative network properties that might not be accessible otherwise. The spectral bound (Rajan and Abbott, 2006) of the effective connectivity matrix in linear response theory (Lindner et al., 2005; Pernice et al., 2011; Trousdale et al., 2012) is an example of such a property. It is a measure for the stability of the linearized system and determines, for example, the occurrence of slow dynamics and long-range correlations, as we demonstrate in Chapter 3. Another example is the sensitivity measure briefly mentioned at the end of Section 2.3.2, which directly links a network model's connectivity with the properties of its power spectrum. These measures are not accessible via simulations. They require analytical calculation.

Similarly, NNMT is not a simulator. NNMT is a collection of mean-field equation implementations that directly relate a model's parameters to the statistics of its dynamics or to other informative properties. It provides these implementations in a format that makes them applicable to as many network models as possible. This is not to say that NNMT does not involve numerical integration procedures; solving self-consistent equations, such as in the case of the firing rates calculations in Section 2.3.2 is a common task, and a collection of respective solvers is part of NNMT.

2.4.2 Use cases

In Section 2.3.2, we present a concrete application of some of the tools available in NNMT, and the original publication (Layer et al., 2022), as well as the online documentation², contain many more usage examples. Here, we consider such examples from a more general perspective to highlight the use cases NNMT lends itself to, as well as provide some ideas for how the toolbox could be utilized in future projects.

³ <http://alleninstitute.github.io/dipde>

Analytical methods have the advantage of being fast, and typically they only require a limited amount of computational resources. The computational costs for calculating analytical estimates of dynamical network properties like firing rates, as opposed to the costs of running simulations of a network model, are independent of the number of neurons the network is composed of. This is especially relevant for parameter space explorations, for which many simulations have to be performed. To speed up prototyping, a modeler can first perform a parameter scan using analytical tools from NNMT to get an estimate of the right parameter regimes and subsequently run simulations on this restricted set of parameters to arrive at the final model parameters. An example of such a parameter scan is given in Section 3.2.1 in Layer et al. (2022), where the firing rates of a network are studied as a function of the external input.

Additionally to speeding up parameter space explorations, analytical methods may guide parameter space explorations in another way: namely, by providing an analytical relation between network model parameters and network dynamics, which allows a targeted adjustment of specific parameters to achieve a desired network activity. The prime example implemented in NNMT is the sensitivity measure briefly mentioned in Section 2.3.2, which provides an intuitive relation between the network connectivity and the peaks of the power spectrum (Panel C of Figure 2.2 on page 19) corresponding to the dominant oscillation frequencies. As shown in the final part of Section 3.3.3 in (Layer et al., 2022), the sensitivity measure identifies the connections which need to be adjusted in order to modify the dominant oscillation mode in a desired manner. This illustrates a mean-field method that provides a modeler with additional information about the origin of a model's dynamics, such that a parameter space exploration can be restricted to the few identified crucial model parameters.

A modeler investigating which features of a network model are crucial for the emergence of certain activity characteristics observed in simulations might be interested in comparing models of differing complexity. The respective mappings can be derived in mean-field theory, and one variant included in NNMT, which is presented in Section 3.4 in Layer et al. (2022), allows mapping a LIF network to a simpler rate network. This is useful to investigate whether spiking dynamics is crucial for the observed phenomenon.

On a general note, which kind of questions researchers pursue is limited by and therefore depends on the tools they have at hand (Dyson, 2012). The availability of sophisticated neural network simulators for example has led to the development of conceptually new and more complex neural network models, precisely because their users could focus on actual research questions instead of implementations. We hope that collecting useful implementations of analytical tools for network model analysis will have a similar effect on the development of new tools and that it might lead to new, creative ways of applying them.

2.4.3 Limitations

As a collection of analytical methods, NNMT comes with inherent limitations that apply to any toolbox for analytical methods: it is restricted to network, neuron, and synapse models, as well as observables, for which a mean-field theory exists, and the tools are based on analytical assumptions, simplifications, and approximations,

restricting their valid parameter regimes and their explanatory power, which we expand upon in the next paragraphs. Analytical methods can provide good estimates of network model properties, but there are limitations that must be considered when interpreting results provided by NNMT: First of all, the employed numerical solvers introduce numerical inaccuracies, but they can be remedied by changing hyperparameters such as integration step sizes or iteration termination thresholds. More importantly, analytical methods almost always rely on approximations, which can only be justified if certain assumptions are fulfilled. Typical examples of such assumptions are fast or slow synapses, or a random connectivity. If such assumptions are not met, at least approximately, and the valid parameter regime of a tool is left, the corresponding method is not guaranteed to give reliable results. Hence, it is important to be aware of a tool's limitations, which we aim to document as thoroughly as possible.

An important assumption of mean-field theory is uncorrelated Poissonian inputs. As discussed in Section 1.4, asynchronous irregular activity is a robust feature of balanced networks, and mean-field theory is well-suited to describe the activity of such models. However, if a network model features highly correlated activity, or strong external input common to many neurons, approximating the input by uncorrelated noise no longer holds and mean-field estimates become unreliable.

In addition to the breakdown of such assumptions, some approaches, like linear response theory, rely on neglecting higher order terms. This restricts the tools' explanatory power, as they cannot predict higher order effects, such as the presence of higher harmonics in a network's power spectrum. Addressing these deficiencies necessitates using more elaborate analyses, and users should be aware of such limitations when interpreting the results.

Finally, a specific limitation of NNMT is that it currently only collects methods for LIF neurons. However, one of the aims of our publication (Layer et al., 2022) is to encourage other scientists to contribute to the collection. In the next section, we outline what sorts of contributions are suitable and how they can be made.

2.4.4 *How to contribute and outlook*

A toolbox like NNMT always is an ongoing project, and there are various aspects that can be improved. In this section, we briefly discuss how available methods could be improved, what and how new tools could be added, as well as the benefits of implementing a new method with the help of NNMT.

First of all, NNMT in its current state is partly vectorized but the included methods are not parallelized, e.g., using multiprocessing or MPI for Python (`mpi4py`). Vectorization relies on NumPy (Harris et al., 2020) and SciPy (Virtanen et al., 2020), which are thread-parallel for specific backends, e.g., IntelMKL. With the tools available in the toolbox at the moment, run-time only becomes an issue in extensive parameter scans, for instance, when the transfer function needs to be calculated for a large range of frequencies. To further reduce the runtime, the code could be made fully vectorized. Alternatively, parallelization of many tools in NNMT is straightforward, as many of them include for loops over the different populations of a network model and for loops over the different analysis frequencies. A third option is just-in-time compilation,

as provided by Numba (Lam et al., 2015). Another aspect to consider is the range of network models a tool can be applied to. Thus far, the toolbox primarily supports arbitrary block structured networks. Future developments could extend the class of networks to even more general models.

Due to the research focus at our lab, NNMT presently mainly contains tools for LIF neurons in the fast synaptic regime and networks with random connectivity. Nonetheless, the structure of NNMT allows for adding methods for different neuron types, like for example binary (Ginzburg and Sompolinsky, 1994) or conductance-based neurons (Izhikevich, 2007; Richardson, 2007), as well as more elaborate network models. Another way to improve the toolbox is adding tools that complement the existing ones: As discussed in Section 2.4.3, many mean-field methods only give valid results for certain parameter ranges. Sometimes, there exist different approximations for the same quantity, valid in complementary parameter regimes. A concrete example is the currently implemented version of the transfer function for leaky integrate-and-fire neurons, based on Schuecker et al. (2015), which gives a good estimate for small synaptic time constants compared to the membrane time constant $\tau_s/\tau_m \ll 1$. A complementary estimate for $\tau_s/\tau_m \gg 1$ has been developed by Moreno-Bote and Parga (2006). Similarly, the current implementation of the firing rates of leaky integrate-and-fire neurons, based on the work of Fourcaud and Brunel (2002), is valid for $\tau_s/\tau_m \ll 1$. Recently, van Vreeswijk and Farkhooi (2019) have developed a method accurate for all combinations of synaptic and membrane time constants.

In the following, we explain how such implementations can be added and how using NNMT helps implementing new methods. Clearly, the implementations of NNMT help implementing methods that build on already existing ones. An example is the firing rate for LIF neurons with exponential synapses `nnmt.lif.exp._firing_rates()` which wraps the calculation of firing rates for LIF neurons with delta synapses `nnmt.lif.delta._firing_rates()`. Additionally, the toolbox may support the implementation of tools for other neuron models. As an illustration, let us consider adding the computation of the mean activity for a network of binary neurons (included in NNMT 1.1.0). We start with the equations for the mean input μ_a , its variance σ_a^2 , and the vector of mean population activities $\mathbf{m}$ (Helias et al., 2014, Eqs. 4, 6, and 7)

$$\begin{aligned}\mu_a(\mathbf{m}) &= \sum_b K_{ab} J_{ab} m_b \quad , \\ \sigma_a^2(\mathbf{m}) &= \sum_b K_{ab} J_{ab}^2 m_b (1 - m_b) \quad , \\ m_a(\mu_a, \sigma_a) &= \frac{1}{2} \operatorname{erfc} \left(\frac{\Theta_a - \mu_a}{\sqrt{2}\sigma_a} \right) \quad ,\end{aligned}\tag{2.5}$$

with indegree matrix K_{ab} from population b to population a , synaptic weight matrix J_{ab} , and firing-threshold Θ_a . The sum $\sum_b$ may include an external population providing input to the model. This set of self-consistent equations has the same structure as the self-consistent equations for the firing rates of a network of LIF neurons Eq. (2.3): the input statistics are given as functions of the rate, and the rate is given as a function of the input statistics. Therefore, it is possible to reuse the firing rate integration procedure for LIF neurons, providing immediate access to the two different integration methods presented in Section 3.2.1 in Layer et al. (2022). Accordingly, it is sufficient to

implement Eq. (2.5) in a new submodule `nnmt.binary` and apply the solver provided by NNMT to extend the toolbox to binary neurons.

The above example demonstrates the benefits of collecting analytical tools for network model analysis in a common framework. The more methods and corresponding solvers the toolbox comprises, the easier implementing new methods becomes. Therefore, contributions to the toolbox are highly welcome; this can be done via the standard pull request workflow on GitHub (see the “Contributors guide” of the official documentation of NNMT²). We hope that in the future, many scientists will contribute to this collection of analytical methods for neuronal network model analysis, such that, at some point, we will have tools from all parts of mean-field theory of neuronal networks, made accessible in a usable format to all neuroscientists.

CORRELATIONS IN TWO-DIMENSIONAL NETWORKS

This chapter is based on the following publication:

David Dahmen*, Moritz Layer*, Lukas Deutz, Paulina Anna Dąbrowska, Nicole Voges, Michael von Papen, Thomas Brochier, Alexa Riehle, Markus Diesmann, Sonja Grün & Moritz Helias (2022), Global organization of neuronal activity only requires unstructured local connectivity, *eLife*, 11, e68422, DOI: 10.7554/eLife.68422

Author contributions:

The author shares first authorship of this publication with David Dahmen. Under the supervision of David Dahmen and Moritz Helias, the author carried out the analytical investigation and the corresponding numerical validation, developed the NEST simulation, and evaluated the simulated data, while David Dahmen focused on the conceptualization and, jointly with Paulina Anna Dąbrowska, experimental data analysis. All authors contributed to the conceptual work and the writing of the manuscript.

3.1 INTRODUCTION

As discussed in the introduction (Section 1.4), the theory of random networks is well established, and the coordination of neuronal activity in such networks has been explored extensively (Ginzburg and Sompolsky, 1994; Renart et al., 2010; Pernice et al., 2011; Tetzlaff et al., 2012; Grytskyy et al., 2013; Helias et al., 2014). In unstructured random networks each pair of neurons may have a direct connection. Introducing a distance-dependent connectivity, however, restricts the spatial range within which a neuron is likely to form direct connections; realistically, to a range of a few hundred micrometers (cf. Section 1.1). Nevertheless, neurons have to and clearly do coordinate their activity across whole cortical areas spanning several millimeters. The population coding paradigm (Georgopoulos et al., 1983), one of the early notions of coding in neuronal networks, postulates that entire populations of similarly tuned neurons behave coherently to encode information, necessitating a positive correlation between the participating neurons. Moreover, recently evidence accumulated that neuronal activity often evolves within complex low-dimensional manifolds, which implies even more involved ways of neuronal activity coordination (Gallego et al., 2017, 2018, 2020). The properties of such manifolds are determined by the structure of the correlations (Abbott et al., 2011; Mazzucato et al., 2016). Yet, it is unclear

how distant neurons that are unlikely to be connected may nonetheless be strongly coordinated.

To answer this question, we combine the analytical investigation of a spatially organized neuronal network model with the analysis of parallel spiking data from macaque motor cortex. We here quantify coordination by Pearson correlation coefficients and pair-wise covariances, which measure how temporal fluctuations of the neurons' activities around their mean firing rates are correlated. We show that even with only unstructured and short-range connections, strong covariances across distances of several millimeters emerge naturally in balanced networks if their dynamical state is within a critical regime which is caused by strongly heterogeneous local network connections, like in cortical brain networks. Intuitively, heterogeneous connectivity provides a large number of different indirect pathways linking two neurons, along which activity might propagate, determining the neurons' coordination. While mean covariances are readily accessible by mean-field techniques and have been shown to be small in balanced networks (Renart et al., 2010; Tetzlaff et al., 2012), explaining covariances on the level of individual pairs requires methods from statistical physics of disordered systems. With such a theory, here derived for spatially organized excitatory-inhibitory networks, we show that large pair-wise covariances arise at all distances if the network is close to the critical point. These predictions are confirmed by recordings of macaque motor cortex activity. The long-range coordination found in this study is not merely determined by the anatomical connectivity, but depends substantially on the network state, which is characterized by the individual neurons' mean firing rates. This allows the network to adjust the neuronal coordination pattern in a dynamic fashion, which we demonstrate through simulations and by comparing two behavioral epochs of a reach-to-grasp experiment.

In the remainder of this chapter, we first have a look at correlations in the macaque data in Section 3.2, and give an intuitive explanation for the observed coordination patterns in Section 3.3. In Section 3.4, we introduce our model and explain the results of our analysis, which yields predictions that are confirmed in the data in Section 3.5. The theory provides a second hypothesis, which is illustrated in Section 3.6 using a simulation of a simpler network model, and probed in the data in Section 3.7. Finally, we discuss our results in Section 3.8.

3.2 MACAQUE MOTOR CORTX SHOWS LONG-RANGE COORDINATION PATTERNS

We first analyze data from motor cortex of macaques during rest, recorded with $4 \times 4 \text{ mm}^2$, 100-electrode Utah arrays with $400 \mu\text{m}$ inter-electrode distance (Panel A of Figure 3.1 on page 31). The resting condition of motor cortex in monkeys is ideal to assess intrinsic coordination between neurons during ongoing activity. In particular, our analyses focus on true resting state data, devoid of movement-related transients in neuronal firing (see Methods). Parallel single-unit spiking activity of ≈ 130 neurons per recording session, sorted into putative excitatory and inhibitory cells, shows strong spike-count correlations across the entire Utah array, well beyond the typical scale of the underlying short-range connectivity profiles (Panels B and D of Figure 3.1 on page 31). Positive and negative correlations form patterns in space that are furthermore




Figure 3.1: Salt-and-pepper structure of covariances in motor cortex. (A) Sketch of 10×10 Utah electrode array recording in motor cortex of macaque monkey during rest. (B) Spikes are sorted into putative excitatory (blue triangles) and inhibitory (red circles) single units according to widths of spike waveforms (see Appendix Section B.2). Resulting spike trains are binned in 1 s bins to obtain spike counts. (C) Population-resolved distribution of pairwise spike-count Pearson correlation coefficients in session E2 (E-E: excitatory-excitatory, E-I: excitatory-inhibitory, I-I: inhibitory-inhibitory). (D) Pairwise spike-count correlation coefficients with respect to the neuron marked by black triangle in one recording (session E2, see Methods). Grid indicates electrodes of a Utah array, triangles and circles correspond to putative excitatory and inhibitory neurons, respectively. Size as well as color of markers represent correlation. Neurons within the same square were recorded on the same electrode.




Figure 3.2: **Correlations from direct and indirect connections.** (A) Positive correlation (green neuron i) follows from direct excitatory connection (top) or shared input (middle). (B) Negative correlation (magenta) between two excitatory neurons cannot be explained by direct connections: Neuronal interactions are not only mediated via direct connections ($n = 1$; sign uniquely determined by presynaptic neuron type) but also via indirect paths of different length $n > 1$. The latter may have any sign (positive: green; negative: purple) due to intermediate neurons of arbitrary type (triangle: excitatory, circle: inhibitory).

seemingly unrelated to the neuron types. All populations show a large dispersion of both positive and negative correlation values (Panel C of Figure 3.1 on page 31).

The classical view on pairwise correlations in balanced networks (Ginzburg and Sompolinsky, 1994; Renart et al., 2010; Pernice et al., 2011, 2012; Tetzlaff et al., 2012; Helias et al., 2014) focuses on averages across many pairs of cells: average correlations are small if the network dynamics is stabilized by an excess of inhibitory feedback; dynamics known as the ‘balanced state’ arise (van Vreeswijk and Sompolinsky, 1996; Amit and Brunel, 1997b; van Vreeswijk and Sompolinsky, 1998): Negative feedback counteracts any coherent increase or decrease of the population-averaged activity, preventing the neurons from fluctuating in unison (Tetzlaff et al., 2012). Breaking this balance in different ways leads to large correlations (Rosenbaum and Doiron, 2014; Darshan et al., 2018; Baker et al., 2019). Can the observation of significant correlations between individual cells across large distances be reconciled with the balanced state? In the following, we provide an intuitive mechanistic explanation.

3.3 MULTI-SYNAPTIC CONNECTIONS DETERMINE CORRELATIONS

Connections mediate interactions between neurons. Many studies therefore directly relate connectivity and correlations (Pernice et al., 2011, 2012; Trousdale et al., 2012; Brinkman et al., 2018; Kobayashi et al., 2019). From exclusively direct connections, one would expect positive correlations between excitatory neurons, negative correlations between inhibitory neurons, and a mix of negative and positive correlations only for excitatory-inhibitory pairs. Likewise, a shared input from inside or outside the network only imposes positive correlations between any two neurons (Panel A of

Figure 3.2 on page 32). The observations that excitatory neurons may have negative correlations (Panel D of Figure 3.1 on page 31), as well as the broad distribution of correlations covering both positive and negative values (Panel C of Figure 3.1 on page 31), are not compatible with this view. In fact, the sign of correlations appears to be independent of the neuron types. So how do negative correlations between excitatory neurons arise?

The view that equates connectivity with correlation implicitly assumes that the effect of a single synapse on the receiving neuron is weak. This view, however, regards each synapse in isolation. Could there be states in the network where, collectively, many weak synapses cooperate, as perhaps required to form low-dimensional neuronal manifolds? In such a state, interactions may not only be mediated via direct connections but also via indirect paths through the network (Panel B of Figure 3.2 on page 32). Such effective multi-synaptic connections may explain our observation that far apart neurons that are most likely not connected directly display considerable correlation of arbitrary sign.

Let us here illustrate the ideas first and corroborate them in subsequent sections. Direct connections yield correlations of a predefined sign, leading to correlation distributions with multiple peaks, e.g. a positive peak for excitatory neurons that are connected and a peak at zero for neurons that are not connected. Multi-synaptic paths, however, involve both excitatory and inhibitory intermediate neurons, which contribute to the interaction with different signs (Panel B of Figure 3.2 on page 32). Hence, a single indirect path can contribute to the total interaction with arbitrary sign (Pernice et al., 2011). If indirect paths dominate the interaction between two neurons, the sign of the resulting correlation becomes independent of their type. Given that the connecting paths in the network are different for any two neurons, the resulting correlations can fall in a wide range of both positive and negative values, giving rise to the broad distributions for all combinations of neuron types in Panel C of Figure 3.1 on page 31. This provides a hypothesis why there may be no qualitative difference between the distribution of correlations for excitatory and inhibitory neurons. In fact, their widths are similar and their mean is close to zero (see Methods for exact values); the latter being the hallmark of the negative feedback that characterizes the balanced state. The subsequent model-based analysis will substantiate this idea and show that it also holds for networks with spatially organized heterogeneous connectivity.

To play this hypothesis further, an important consequence of the dominance of multi-synaptic connections could be that correlations are not restricted to the spatial range of direct connectivity. Through interactions via indirect paths the reach of a single neuron could effectively be increased. But the details of the spatial profile of the correlations in principle could be highly complex as it depends on the interplay of two antagonistic effects: On the one hand, signal propagation becomes weaker with distance, as the signal has to pass several synaptic connections. Along these paths mean firing rates of neurons are typically diverse, and so are their signal transmission properties (De la Rocha et al., 2007). On the other hand, the number of contributing indirect paths between any pair of neurons proliferates with their distance. With single neurons typically projecting to thousands of other neurons in cortex, this leads to involved combinatorics; intuition here ceases to provide a sensible hypothesis on what is the spatial profile and range of coordination between neurons. Also it is

unclear which parameters these coordination patterns depend on. The model-driven and analytical approach presented in the next section will provide such a hypothesis.

3.4 ALMOST UNSTABLE NETWORKS SHOW SHALLOW EXPONENTIAL DECAY OF COVARIANCES

We first note that the large magnitude and dispersion of pair-wise correlations in the data and their spatial structure primarily stem from features in the underlying covariances between neuron pairs (Appendix Figure B.1 on page 96). Given the close relationship between correlations and covariances (Panels D and E of Appendix Figure B.1 on page 96), we analyze covariances in the following, as these are less dependent on single neuron properties and thus simpler to treat analytically.

To gain an understanding of the spatial features of intrinsically generated covariances in balanced critical networks, we investigate a network of N excitatory and inhibitory rate neurons on a two-dimensional sheet (Panel A of Figure 3.3 on page 35), whose activity $\mathbf{x} \in \mathbb{R}^N$ is described by

$$\tau \frac{d}{dt} \mathbf{x} = -\mathbf{x} + \mathbf{J} \cdot \boldsymbol{\psi}(\mathbf{x}) + \boldsymbol{\xi},$$

such that each neuron receives external uncorrelated Gaussian white noise input $\boldsymbol{\xi}$, $\langle \xi_i(t) \rangle = \mu$, $\langle \xi_i(s) \xi_j(t) \rangle = D_i \delta_{ij} \delta(s - t)$. The neuronal integration time constant is τ , and $\boldsymbol{\psi}$ is the neurons' non-linear input-output function. To allow for multapses, multiple connections between pairs of neurons, the connections J_{ij} between the neurons are drawn from a binomial distribution, and the connection probability decays with inter-neuronal distance on a characteristic length scale R (for more details see Methods or Appendix Section B.4). We investigate the covariance statistics in this model by help of linear-response theory, which has been shown to approximate spiking neuron models well (Pernice et al., 2012; Trousdale et al., 2012; Tetzlaff et al., 2012; Helias et al., 2013; Grytskyy et al., 2013; Dahmen et al., 2019). To that end, we linearize the equation around the network's *working point* $\mathbf{x}_0$, its stationary state, which yields a description of the fluctuations $\delta \mathbf{x} = \mathbf{x} - \mathbf{x}_0$ around the working point

$$\tau \frac{d}{dt} \delta \mathbf{x} = -\delta \mathbf{x} + \mathbf{W} \cdot \delta \mathbf{x} + \boldsymbol{\xi}',$$

with $\langle \xi'_i(t) \rangle = 0$, $\langle \xi'_i(s) \xi'_j(t) \rangle = D_i \delta_{ij} \delta(s - t)$. Here we introduced the matrix $W_{ij} = J_{ij} \boldsymbol{\psi}'(\mathbf{x}_{0,i})$, in the following referred to as *effective connectivity*.

Previous studies have used linear-response theory in combination with methods from statistical physics and field theory to gain analytic insights into both mean covariances (Ginzburg and Sompolsky, 1994; Lindner et al., 2005; Pernice et al., 2011; Tetzlaff et al., 2012) and the width of the distribution of covariances (Dahmen et al., 2019). Field-theoretic approaches, however, were so far restricted to purely random networks devoid of any network structure and thus not suitable to study spatial features of covariances. To analytically quantify the relation between the spatial ranges of covariances and connections, we therefore here develop a theory for spatially organized random networks with multiple populations. The randomness in our model is based on the sparseness of connections, which is one of the main sources




Figure 3.3: **Spatially organized E-I network model.** (A) Network model: space is divided into cells with four excitatory (triangles) and one inhibitory (circle) neuron each. Distance-dependent connection probabilities (shaded areas) are defined with respect to cell locations. (B) Eigenvalues λ of effective connectivity matrix for network in dynamically balanced critical state. Each dot shows the real part $\text{Re}(\lambda)$ and imaginary part $\text{Im}(\lambda)$ of one complex eigenvalue. The spectral bound (dashed vertical line) denotes the right-most edge of the eigenvalue spectrum. (C) Simulation: covariances of excitatory neurons over distance x between cells (blue dots: individual pairs; cyan: mean; orange: standard deviation; sample of 150 covariances at each of 200 chosen distances). (D) Theory: Variance of covariance distribution as a function of distance x for different spectral bounds of the effective connectivity matrix. *Inset*: effective decay constant of variances diverges as the spectral bound approaches one. (E) For large spectral bounds, the variances of EE, EI, and II covariances decay on a similar length scale. Spectral bound $\rho = 0.95$. Other parameters see Appendix Table B.3 on page 131.

of heterogeneity in cortical networks in that it contributes strongly to the variance of connections (see Appendix Section B.15).

Our aim is to find an expression for mean and variance of the covariances as functions of distance between two neurons. The detailed derivation is presented in Appendix Section B.4–Section B.12; a summary of the extensive derivation is provided in the Methods. Here we simply state the main results: Field-theoretic methods, combined with linear-response theory, allow us to obtain expressions for the mean covariance $\bar{c}$ and variance of covariance $\overline{\delta c^2}$

$$\begin{aligned}\bar{c} &= [\mathbf{1} - \mathbf{M}]^{-1} \frac{D}{1 - \rho^2} [\mathbf{1} - \mathbf{M}]^{-T}, \\ \overline{\delta c^2} &= [\mathbf{1} - \mathbf{S}]^{-1} \left(\frac{D}{1 - \rho^2} \right)^2 [\mathbf{1} - \mathbf{S}]^{-T},\end{aligned}\tag{3.1}$$

with identity matrix $\mathbf{1}$, input noise strength D , and the effective connectivity's *spectral bound* ρ , the largest real part among all eigenvalues of $\mathbf{W}$ (Panel B of Figure 3.3 on page 35). The matrices $\mathbf{M}$ and $\mathbf{S}$ represent the mean and variance of the random effective connectivity matrix $\mathbf{W}$, and hence contain the information about the distance-dependent connectivity statistics. Because the analysis of mean and variance of the covariances is identical, we just look at the variances here. Leveraging the symmetries $\mathbf{S}$ inherits from the connectivity, taking the continuum limit, and using approximations for large spectral radii $\rho \lesssim 1$, allows us to diagonalize the expressions. This reveals the spatial dependency of the variances of the covariance

$$\overline{\delta c^2}(\mathbf{x}) = \left(\frac{D}{1 - \rho^2} \right)^2 [\delta(|\mathbf{x}|) + B(\mathbf{x}) + (B * B)(\mathbf{x})],$$

with

$$B(\mathbf{x}) \propto K_0 \left(-\frac{|\mathbf{x}|}{R_{\text{eff}}} \right) \underset{|\mathbf{x}| \rightarrow \infty}{\propto} \frac{\exp(-|\mathbf{x}|/R_{\text{eff}})}{\sqrt{|\mathbf{x}|}},$$

where $K_0(x)$ denotes the modified Bessel function of second kind and zeroth order (Olver et al., 2010). The effective decay constant R_{eff} is given by

$$\frac{R_{\text{eff}}}{R} \sim \sqrt{\frac{\rho^2}{1 - \rho^2}} + \text{const.} \gg 1.\tag{3.2}$$

The equivalent results for the mean covariances come with different constants and prefactors, and the spectral bound replaced by the *population eigenvalue*, which corresponds to the all-ones eigenvector (Panel B of Figure 3.3 on page 35).

The analytical results show that despite the complexity of the various indirect interactions, both the mean and the variance of covariances follow simple exponential laws in the long-distance limit. These laws are universal in that they do not depend on details of the spatial profile of connections, as long as the profile decays with distance. Our theory shows that the associated length scales are strikingly different for means and variances of covariances. They each depend on the reach of direct connections and on specific eigenvalues of the effective connectivity matrix. These eigenvalues summarize various aspects of network connectivity and signal transmission into a single number: Each eigenvalue belongs to a “mode”, a combination of neurons that

act collaboratively, rather than independently, coordinating neuronal activity within a one-dimensional subspace. To start with, there are as many such subspaces as there are neurons. But if the spectral bound in Panel B of Figure 3.3 on page 35 is close to one, only a relatively small fraction of them, namely those close to the spectral bound, dominate the dynamics; the dynamics is then effectively low-dimensional. Additionally, the eigenvalue quantifies how fast a mode decays when transmitted through a network. The eigenvalues of the dominating modes are close to one, which implies a long lifetime. The corresponding fluctuations thus still contribute significantly to the overall signal, even if they passed by many synaptic connections. Therefore, indirect multi-synaptic connections contribute significantly to covariances if the spectral bound is close to one, and in that case we expect to see long-range covariances.

To quantify this idea, for the mean covariance $\bar{c}$ we find that the dominant behavior is an exponential decay $\bar{c} \sim \exp(-x/\bar{R})$ on a length scale $\bar{R}$. This length scale is determined by the population eigenvalue, corresponding to the mode in which all neurons are excited simultaneously. Its position solely depends on the ratio between excitation and inhibition in the network and becomes more negative in more strongly inhibition-dominated networks (Panel B of Figure 3.3 on page 35). We show in Appendix Section B.9.4 that this leads to a steep decay of mean covariances with distance. The variance of covariances, however, predominantly decays exponentially on a length scale R_{eff} that is determined by the spectral bound ρ (Panels B and D of Figure 3.3 on page 35). In inhibition-dominated networks, ρ is determined by the heterogeneity of connections. For $\rho \lesssim 1$ we obtain the effective length scale Eq. (3.2). What this means is that precisely at the point where ρ is close to one, when neural activity occupies a low-dimensional manifold, the length scale R_{eff} on which covariances decay exceeds the reach of direct connections by a large factor (Panel D of Figure 3.3 on page 35). As the network approaches instability, which corresponds to the spectral bound ρ going to one, the effective decay constant diverges (Panel D inset of Figure 3.3 on page 35) and so does the range of covariances. A distance-resolved histogram of the covariances in a concrete instance of the spatially organized E-I network close to the critical point confirms this prediction: Panel C of Figure 3.3 on page 35 shows that the mean covariance is close to zero but the width or variance of the covariance distribution stays large, even for large distances.

Our population-resolved theoretical analysis, furthermore, shows that the larger the spectral bound the more similar the decay constants between different populations, with only marginal differences for $\rho \lesssim 1$ (Panel E of Figure 3.3 on page 35). This holds strictly for target-agnostic connectivity, in which the connection weights only depend on the type of the presynaptic neuron but not on the type of the postsynaptic neuron. Moreover, we find a relation between the squared effective decay constants and the squared anatomical decay constants of the form

$$R_{\text{eff,E}}^2 - R_{\text{eff,I}}^2 = \text{const.} \cdot (R_E^2 - R_I^2) . \quad (3.3)$$

This relation is independent of the eigenvalues of the effective connectivity matrix, as the constant of order $\mathcal{O}(1)$ does only depend on the choice of the connectivity profile. For $\rho \simeq 1$, this means that even though the absolute value of both effective length




Figure 3.4: **Long-range covariances in macaque motor cortex.** Variance of covariances as a function of distance. (A) Population-specific exponential fits (lines) to variances of covariances (dots) in session E2, with fitted decay constants indicated in the legend (I-I: putative inhibitory neuron pairs, I-E: inhibitory-excitatory, E-E: excitatory pairs). Dots show the empirical estimate of the variance of the covariance distribution for each distance. Size of the dots represents relative count of pairs per distance and was used as weighting factor for the fits to compensate for uncertainty at large distances, where variance estimates are based on fewer samples. Mean squared error 2.918. (B) Population-specific exponential fits (lines) analogous to (A), with slopes constrained to be identical. This procedure yields a single fitted decay constant of 1.029 mm. Mean squared error 2.934. (C) Table listing decay constants fitted as in (B) for all recording sessions and the ratios between mean squared errors of the fits obtained in procedures B and A.

scales on the left hand side is large, their relative difference is small because it equals the small difference of anatomical length scales on the right hand side.

3.5 PAIRWISE COVARIANCES IN MOTOR CORTEX DECAY ON A MILLIMETER SCALE

To check if these predictions are confirmed by the data from macaque motor cortex, we first observe that, indeed, covariances in the resting state show a large dispersion over almost all distances on the Utah array (Figure 3.4 on page 38). Moreover, the variance of covariances agrees well with the predicted exponential law: Performing an exponential fit reveals length constants above one millimeter. These large length constants have to be compared to the spatial reach of direct connections, which is about an order of magnitude shorter, in the range of 100 – 400 μm (Schnepel et al., 2015), so below the 400 μm inter-electrode distance of the Utah array (Section 1.2). The shallow decay of the variance of covariances is, next to the broad distribution of covariances, a second indication that the network is in the dynamically balanced critical regime, in line with the prediction by Eq. (3.2).

The population-resolved fits to the data show a larger length constant for excitatory covariances than for inhibitory ones (Panel A of Figure 3.4 on page 38). This is qualitatively in line with the prediction of Eq. (3.3) given the — by tendency — longer reach of excitatory connections compared to inhibitory ones, as derived from

morphological constraints (Section 1.1). In the dynamically balanced critical regime, however, the predicted difference in slope for all three fits is practically negligible. Therefore, we performed a second fit where the slope of the three exponentials is constrained to be identical (Panel B of Figure 3.4 on page 38). The error of this fit is only marginally larger than the ones of fitting individual slopes (Panel C of Figure 3.4 on page 38). This shows that differences in slopes are hardly detectable given the empirical evidence, thus confirming the predictions of the theory given by Eq. (3.2) and Eq. (3.3).

3.6 FIRING RATES ALTER CONNECTIVITY-DEPENDENT COVARIANCE PATTERNS

Since covariances measure the coordination of temporal fluctuations around the individual neurons' mean firing rates, they are determined by how strong a neuron transmits such fluctuations from input to output (Abeles, 1991). To leading order this is explained by linear-response theory (Ginzburg and Sompolinsky, 1994; Lindner et al., 2005; Pernice et al., 2011; Tetzlaff et al., 2012): How strongly a neuron reacts to a small change in its input depends on its dynamical state, foremost the mean and variance of its total input, which determine the network's working point. If a neuron receives almost no input, a small perturbation in the input will not be able to make the neuron fire. If the neuron receives a large input, a small perturbation will not change the firing rate either, as the neuron is already saturated. Only in the intermediate regime the neuron is susceptible to small deviations of the input. Mathematically, this behavior is described by the gain of the neuron, which is the derivative of the neuron's input-output function $\psi'(x_0)$ at the working point (Abeles, 1991). Due to the non-linearity of the input-output function, the gain is vanishing for very small and very large inputs and non-zero in the intermediate regime. How strongly a perturbation in the input to one neuron affects one of the subsequent neurons therefore not only depends on the synaptic weight J but also on the gain $\psi'(x_0)$ and thereby the working point. This relation is captured by the effective connectivity $W_{ij} = J_{ij}\psi'(x_{0,j})$. What is the consequence of the dynamical interaction among neurons depending on the working point? Can it be used to reshape the low-dimensional manifold, the collective coordination between neurons?

The first part of this study reveals that long-range coordination can be achieved in a network with short-range random connections if effective connections are sufficiently strong. Alteration of the working point, for example by a different external input level, can affect the covariance structure: The pattern of coordination between individual neurons can change, even though the anatomical connectivity remains the same. In this way, routing of information through the network can be adapted dynamically on a mesoscopic scale. This is a crucial difference of such coordination as opposed to coordination imprinted by complex but static connection patterns.

Here we first illustrate this concept by simulations of a network of 2000 sparsely connected threshold-linear (ReLU) rate neuron models that receive Gaussian white noise inputs centered around neuron-specific non-zero mean values (see Methods and Appendix Section B.14 for more details). The ReLU activation function thereby acts as a simple model for the vanishing gain for neurons with too low input levels. Note that in cortical-like scenarios with low firing rates, neuronal working points are far




Figure 3.5: Changes in effective connectivity modify coordination patterns. (A) Visualization of effective connectivity: A sparse random network with given structural connectivity (left network sketch) is simulated with two different input levels for each neuron (depicted by insets), resulting in different firing rates (grayscale in right network sketches) and therefore different effective connectivities (thickness of connections) in the two simulations. Parameters can be found in Appendix Table B.4 on page 131. (B1) Histogram of input currents across neurons for the two simulations (N1 and N2). (B2) Scatter plot of inputs to subset of 1500 corresponding neurons in the first and the second simulation (Pearson correlation coefficient $r = 0.90$). (C) Correlation coefficients of rates and of covariances between the two simulations (b, black) and within two epochs of the same simulation (w, gray). (D) Correlation coefficient of rates (gray) and covariances (black) between the two simulations as a function of the spectral bound ρ . (E1) Distribution of rates in the two simulations (excluding silent neurons with $|\text{rate}| < 10^{-3}$). (E2) Scatter plot of rates in the first compared to the second simulation (Pearson correlation coefficient $r = 0.81$). (F1) Distribution of covariances in the two simulation (excluding silent neurons). (F2) Scatter plot of sample of 5000 covariances in first compared to the second simulation (Pearson correlation coefficient $r = 0.40$). Here silent neurons are included (accumulation of markers on the axes). Other parameters: number of neurons $N = 2000$, connection probability $p = 0.1$, spectral bound for panels B, C, E, F is $\rho = 0.8$.

away from the high-input saturation discussed above, which is therefore neglected by the choice of the ReLU activation function. For independent and stationary external inputs covariances between neurons are solely generated inside the network via the sparse and random recurrent connectivity. External inputs only have an indirect impact on the covariance structure by setting the working point of the neurons.

We simulate two networks with identical structural connectivity and identical external input fluctuations, but small differences in mean external inputs between corresponding neurons in the two simulations (Panel A of Figure 3.5 on page 40). These small differences in mean external inputs create different gains and firing rates and thereby differences in effective connectivity and covariances. Since mean external inputs are drawn from the same distribution in both simulations (Panel B of Figure 3.5 on page 40), the overall distributions of firing rates and covariances across all neurons are very similar (Panels E1 and F2 of Figure 3.5 on page 40). But individual neurons' firing rates do differ (Panel E2 of Figure 3.5 on page 40). For the simple ReLU activation used here, we in particular observe neurons with vanishing firing rate in one simulation and non-zero firing rate in the other simulation. This resulting change of working points substantially affects the covariance patterns (Panel F2 of Figure 3.5 on page 40): Differences in firing rates and covariances between the two simulations are significantly larger than the differences across two different epochs of the same simulation (Panel C of Figure 3.5 on page 40). The larger the spectral bound, the more sensitive are the intrinsically generated covariances to the changes in firing rates (Panel D of Figure 3.5 on page 40). Thus, a small offset of individual firing rates is an effective parameter to control network-wide coordination among neurons. As the input to the local network can be changed momentarily, we predict that in the dynamically balanced critical regime coordination patterns should be highly dynamic.

3.7 COORDINATION PATTERNS IN MOTOR CORTEX DEPEND ON BEHAVIORAL CONTEXT

In order to test the theoretical prediction in experimental data, we analyze parallel spiking activity from macaque motor cortex, recorded during a reach-to-grasp experiment (Brochier et al., 2018; Riehle et al., 2013). In contrast to the resting state, where the animal was in an idling state, here the animal is involved in a complex task with periods of different cognitive and behavioral conditions (Panel A of Figure 3.6 on page 42). We compare two epochs in which the animal is requested to wait and is sitting still but which differ in cognitive conditions. The first epoch is a starting period (S), where the monkey has self-initiated the behavioral trial and is attentive because it is expecting a cue. The second epoch is a preparatory period (P), where the animal has just received partial information about the upcoming trial and is waiting for the missing information and the GO signal to initiate the movement.

Within each epoch, S or P, the neuronal firing rates are mostly stationary, likely due to the absence of arm movements which create relatively large transient activities in later epochs of the task, which are not analyzed here (see Appendix Section B.3). The overall distributions of the firing rates are comparable for epochs S and P, but the firing rates are distributed differently across the individual neurons: Panel C of Figure 3.6 on page 42 shows one example session of monkey N, where the changes in




Figure 3.6: **Behavioral condition reshapes mesoscopic neuronal coordination.** (A) Trial structure of the reach-to-grasp experiment (Brochier et al., 2018). Blue segments above the time axis indicate data pieces at trial start (dark blue: S (S_1+S_2)) and during the preparatory period (light blue: P (P_1+P_2)). (B) Salt-and-pepper structure of covariance during two different epochs (S_1 and P_1) of one recording session of monkey N (151 trials, 106 single units, cf. Figure 3.1 on page 31 for recording setup). For some neurons the covariance completely reverses, while in the others it does not change. Inhibitory reference neuron indicated by black circle. (C1) Distributions of firing rates during S_1 and P_1 . (C2) Scatter plot comparing firing rates in S_1 and P_1 (Pearson correlation coefficient $r = 0.69$). (D1/D2) Same as panels C1/C2, but for covariances (Pearson correlation coefficient $r = 0.40$). (E) Correlation coefficient of firing rates across neurons in different epochs of a trial for eight recorded sessions. Correlations between sub-periods of the same epoch (S_1 - S_2 , P_1 - P_2 ; within-epoch, gray) and between sub-periods of different epochs (S_x - P_y ; between-epochs, black). Data shown in panels B-D is from session 8. Box plots to the right of the black dashed line show distributions obtained after pooling across all analyzed recording sessions per monkey. The line in the center of each box represents the median, box's area represents the interquartile range, and the whiskers indicate minimum and maximum of the distribution (outliers excluded). Those distributions differ significantly (Student t-test, two-sided, $p \ll 0.001$). (F) Correlation coefficient of covariances, analogous to panel E. The distributions of values pooled across sessions also differ significantly (Student t-test, two-sided, $p \ll 0.001$). For details of the statistical tests, see Methods. Details on number of trials and units in each recording session are provided in Appendix Table B.1 on page 99.

firing rates between the two epochs are visible in the spread of markers around the diagonal line in panel C2. To assess the extent of these changes, we split each epoch, S and P, into two disjoint sub-periods, S₁/S₂ and P₁/P₂ (Panel A of Figure 3.6 on page 42). We compute the correlation coefficient between the firing rate vectors of two sub-periods of different epochs (“between” markers in Panel E of Figure 3.6 on page 42) and compare it to the correlation coefficient between the firing rate vectors of two sub-periods of the same epoch (“within” markers): Firing rate vectors in S₁ are almost perfectly correlated with firing rate vectors in S₂ ($r \approx 1$ for all of the five/eight different recording sessions from different recording days for monkey E/N, similarly for P₁ and P₂), confirming stationarity investigated in Appendix Section B.3. Firing rate vectors in S₁ or S₂, however, show significantly lower correlation to firing rate vectors in P₁ and P₂, confirming a significant change in network state between epochs S and P (Panel E of Figure 3.6 on page 42).

The mechanistic model in the previous section shows a qualitatively similar scenario (Panels C and E of Figure 3.5 on page 40). By construction it produces different firing rate patterns in the two simulations. While the model is simplistic and in particular not adapted to quantitatively reproduce the experimentally observed activity statistics, its simulations and our underlying theory make a general prediction: Differences in firing rates impact the effective connectivity between neurons and thereby evoke even larger differences in their coordination if the network is operating in the dynamically balanced critical regime (Panel D of Figure 3.5 on page 40).

To check this prediction, we repeat the correlation analysis between the two epochs, which we described above for the firing rates, but this time for the covariance patterns. Despite similar overall distributions of covariances in S and P (Panel D₁ of Figure 3.6 on page 42), covariances between individual neuron pairs are clearly different between S and P: Panel B of Figure 3.6 on page 42 shows the covariance pattern for one representative reference neuron in one example recording session of monkey N. In both epochs, this covariance pattern has a salt-and-pepper structure, similar to the resting state data in Panel D of Figure 3.1 on page 31. Yet, neurons change their individual coordination: a large number of neuron pairs even changes from positive covariance values to negative ones and vice versa. These neurons fire cooperatively in one epoch of the task while they show antagonistic firing in the other epoch. The covariances of all neuron pairs of that particular recording session are shown in Panel D₂ of Figure 3.6 on page 42. Markers in the upper left and lower right quadrant show neuron pairs that switch the sign of their coordination (45 % of all neuron pairs). The extent of covariance changes between epochs is again quantified by correlation coefficients between the covariance patterns of two sub-periods (Panel F of Figure 3.6 on page 42). As for the firing rates, we find rather large correlations between covariance patterns in S₁ and S₂ as well as between covariance patterns in P₁ and P₂. Note, however, that correlation coefficients are around 0.8 rather than 1, presumably since covariance estimates from 200 ms periods are noisier than firing rate estimates. The covariance patterns in S₁ or S₂ are, however, significantly more distinct from covariance patterns in P₁ and P₂, with correlation coefficients around 0.5 (Panel F of Figure 3.6 on page 42).

This more pronounced change of covariances compared to firing rates is predicted by a network whose effective connectivity has a large spectral bound, in the dynamically

balanced critical state. In particular, the theory provides a mechanistic explanation for the different coordination patterns between neurons on the mesoscopic scale (range of a Utah array), which are observed in the two states S and P (Panel B of Figure 3.6 on page 42). The coordination between neurons is thus considerably reshaped by the behavioral condition.

3.8 DISCUSSION

In this chapter, we investigate covariance patterns of many neurons across mesoscopic distances analytically and in macaque motor cortex. We show that these patterns have a salt-and-pepper structure, which can be explained by a network model with a spatially dependent random connectivity operating in a dynamically balanced critical state. In this state, cross-covariances are shaped by a large number of parallel, multi-synaptic pathways, leading to interactions reaching far beyond the range of direct connections. Strikingly, this coordination on the millimeter scale is only visible if covariances are resolved on the level of individual neurons; the population mean of covariances quickly decays with distance and is overall very small. In contrast, the variance of covariances is large and predominantly decreases exponentially on length scales of up to several millimeters, even though direct connections typically only reach a few hundred micrometers.

Since the observed coordination patterns are determined by the effective connectivity of the network, they are dynamically controllable by the network state; for example, due to modulations of neuronal firing rates. Parallel recordings in macaque motor cortex during resting state and in different epochs of a reach-to-grasp task confirm this prediction. Simulations indeed exhibit a high sensitivity of coordination patterns to weak modulations of the individual neurons' firing rates, providing a plausible mechanism for these dynamic changes.

THE DYNAMICALLY BALANCED CRITICAL STATE Models of balanced networks have been investigated before (van Vreeswijk and Sompolinsky, 1996; Brunel, 2000a; Renart et al., 2010; Tetzlaff et al., 2012) and experimental evidence for cortical networks operating in the balanced state is overwhelming (Okun and Lampl, 2008; Reinhold et al., 2015; Dehghani et al., 2016). Excess of inhibition in such networks yields stable and balanced population-averaged activities as well as low average covariances (Tetzlaff et al., 2012). Recently the notion of balance has been combined with criticality in the dynamically balanced critical state that results from large heterogeneity in the network connectivity (Dahmen et al., 2019). Here we focus on the spatial organization of cortical networks and study the interplay between balance, criticality, and spatial connectivity in networks of excitatory and inhibitory neurons. We show that in such networks, heterogeneity generates disperse covariance structures between individual neurons on large length-scales with a salt-and-pepper structure.

RELATION TO PREVIOUS STUDIES Spatially organized balanced network models have been investigated before in the limit of infinite network size, as well as under strong and potentially correlated external drive, as is the case, for example, in primary sensory areas of the brain (Rosenbaum et al., 2017; Baker et al., 2019). In this scenario,

intrinsically generated contributions to covariances are much smaller than external ones. Population-averaged covariances then fulfill a linear equation, called the “balance condition” (van Vreeswijk and Sompolinsky, 1996; Hertz, 2010; Renart et al., 2010; Rosenbaum and Doiron, 2014) that predicts a non-monotonous change of population-averaged covariances with distance (Rosenbaum et al., 2017). In contrast, we here consider covariances on the level of individual cells in finite-size networks receiving only weak inputs. While we cannot strictly rule out that the observed covariance patterns in motor cortex are a result of very specific external inputs to the recorded local network, we believe that the scenario of weak external drive is more suitable for non-sensory brain areas, such as, for example, the motor cortex in the resting state conditions studied here. Under such conditions, covariances have been shown to be predominantly generated locally rather than from external inputs: Helias et al. (2014) investigated intrinsic and extrinsic sources of covariances in ongoing activity of balanced networks and found that for realistic sizes of correlated external populations the major contribution to covariances is generated from local network interactions (Fig. 7a in Helias et al., 2014). Dahmen et al. (2019) investigated the extreme case, where the correlated external population is of the same size as the local population (Fig. S6 in Dahmen et al., 2019). Despite sizable external input correlations projected onto the local circuit via potentially strong afferent connections, the dependence of the statistics of covariances on the spectral bound of the local recurrent connectivity is predicted well by the theory that neglects correlated external inputs (see supplement section 3 in Dahmen et al., 2019).

Our analysis of covariances on the single-neuron level goes beyond the balance condition and requires the use of field-theoretical techniques to capture the heterogeneity in the network (Dahmen et al., 2019; Helias and Dahmen, 2020). It relies on linear-response theory, which has previously been shown to faithfully describe correlations in balanced networks of nonlinear (spiking) units (Tetzlaff et al., 2012; Trousdale et al., 2012; Pernice et al., 2012; Grytskyy et al., 2013; Helias et al., 2013; Dahmen et al., 2019). These studies mainly investigated population-averaged correlations with small spectral bounds of the effective connectivity. Subsequently, Dahmen et al. (2019) showed the quantitative agreement of this linear-response theory for covariances between individual neurons in networks of spiking neurons for the whole range of spectral bounds, including the dynamically balanced critical regime. The long-range coordination studied in the current chapter requires the inclusion of spatially non-homogeneous coupling to analyze excitatory-inhibitory random networks on a two-dimensional sheet with spatially decaying connection probabilities. This new theory allows us to derive expressions for the spatial decay of the variance of covariances. We primarily evaluate these expressions in the long-range limit, which agrees well with simulations for distances $r > 2R \sim \mathcal{O}(1 \text{ mm})$, which is fulfilled for most distances on the Utah array (cf. Figure 3.3 on page 35, Appendix Figure B.7 on page 129). For these distances we find that the decay of covariances is dominated by a simple exponential law. Unexpectedly, its decay constant is essentially determined by only two measures, the spectral bound of the effective connectivity, and the length scale of direct connections. The length scale of covariances diverges when approaching the breakdown of linear stability. In this regime, differences in covariances induced by differences in length scales of excitatory and inhibitory connections become negligible.

The predicted emergence of a single length scale of covariances is consistent with our data.

OTHER POSSIBLE MECHANISMS FOR LONG-RANGE COORDINATION This study focuses on local and isotropic connection profiles to show that long-range coordination does not rely on specific connection patterns but can result from the network state alone. Alternative explanations for long-range coordination are based on specifically imprinted network structures: Anisotropic local connection profiles have been studied and shown to create spatio-temporal sequences (Spreizer et al., 2019). Likewise, embedded excitatory feed-forward motifs and cell assemblies via excitatory long-range patchy connections DeFelipe et al. (1986) can create positive covariances at long distances (Diesmann et al., 1999; Litwin-Kumar and Doiron, 2012). Yet these connections cannot provide an explanation for the large negative covariances between excitatory neurons at long distances (see e.g. Panel D of Figure 3.1 on page 31). Long-range connectivity, for example arising from a salt-and-pepper organization of neuronal selectivity with connections preferentially targeting neurons with equal selectivity (Ben-Yishai et al., 1995; Hansel and Sompolinsky, 1998; Roxin et al., 2005; Blumenfeld et al., 2006), would produce salt-and-pepper covariance patterns even in networks with small spectral bounds where interactions are only mediated via direct connections. However, in this scenario one would expect that neurons which have similar selectivity would throughout show positive covariance due to their mutual excitatory connections and due to the correlated input they receive. Yet, when analyzing two different epochs of the reach-to-grasp task, we find that a large fraction of neuron pairs actually switches from being significantly positively correlated to negatively correlated and vice versa (see Panel D2 of Figure 3.6 on page 42, upper left and lower right quadrant). This state-dependence of covariances is in line with the here suggested mechanism of long-range coordination by indirect interactions: Such indirect interactions depend on the effective strengths of various connections and can therefore change considerably with network state. In contrast, correlations due to imprinted network structures are static, so that a change in gain of the neurons will either strengthen or weaken the specific activity propagation, but it will not lead to a change of the sign of covariances that we see in our data. The static impact of these connectivity structures on covariances could nevertheless in principle be included in the presented formalism.

Long-range coordination can also be created from short-range connections with random orientations of anisotropic local connection profiles (Smith et al., 2018). This finding can be linked to the emergence of tuning maps in the visual cortex. The mechanism is similar to ours in that it uses nearly linearly unstable modes that are determined by spatial connectivity structures and heterogeneity. Given the different source of heterogeneity, the modes and corresponding covariance patterns are different from the ones discussed here: Starting from fully symmetric networks with corresponding symmetric covariance patterns, Smith et al. (2018) found that increasing heterogeneity (anisotropy) yields more randomized, but still patchy regions of positive and negative covariances that are in line with low-dimensional activity patterns found in visual cortex. In motor cortex we instead find salt-and-pepper patterns that can be explained in terms of heterogeneity through sparsity. We provide the theoretical

basis and explicit link between connectivity eigenspectra and covariances and show that heterogeneity through sparsity is sufficient to generate the dynamically balanced critical state as a simple explanation for the broad distribution of covariances in motor cortex, the salt-and-pepper structure of coordination, its long spatial range, and its sensitive dependence on the network state. Note that both mechanisms of long-range coordination, the one studied in Smith et al. (2018) and the one presented here, rely on the effective connectivity for the network to reside in the dynamically balanced critical regime. The latter regime is, however, not just one single point in parameter space, but an extended region that can be reached via a multitude of control mechanisms for the effective connectivity, for example by changing neuronal gains (Salinas and Sejnowski, 2001a,b), synaptic strengths (Sompolinsky et al., 1988), and network microcircuitry (Dahmen et al., 2020).

FUNCTIONAL IMPLICATIONS What are possible functional implications of the coordination on mesoscopic scales? Recent work demonstrated activity in motor cortex to be organized in low-dimensional manifolds (Gallego et al., 2017, 2018, 2020). Dimensionality reduction techniques, such as PCA or GPFA (Yu et al., 2009), employ covariances to expose a dynamical repertoire of motor cortex that is comprised of neuronal modes. Previous work started to analyze the relation between the dimensionality of activity and connectivity (Aljadeff et al., 2015, 2016; Mastrogiuseppe and Ostojic, 2018; Dahmen et al., 2019, 2020; Hu and Sompolinsky, 2020), but only in spatially unstructured networks, where each neuron can potentially be connected to any other neuron. The majority of connections within cortical areas, however, stems from local axonal arborizations (Schnepel et al., 2015). Here we add this biological constraint and demonstrate that these networks, too, support a dynamically balanced critical state. This state in particular exhibits neural modes which are spanned by neurons spread across the experimentally observed large distances. In this state a small subset of modes that are close to the point of instability dominates the variability of the network activity and thus spans a low-dimensional neuronal manifold. As opposed to specifically designed connectivity spectra via plasticity mechanisms (Hennequin et al., 2014) or low-rank structures embedded into the connectivity (Mastrogiuseppe and Ostojic, 2018), the dynamically balanced critical state is a mechanism that only relies on the heterogeneity which is inherent to sparse connectivity and abundant across all brain areas.

While we here focus on covariance patterns in stationary activity periods, the majority of recent works studied transient activity during motor behavior (Gallego et al., 2017). How are stationary and transient activities related? During stationary ongoing activity states, covariances are predominantly generated intrinsically (Helias et al., 2014). Changes in covariance patterns therefore arise from changes in the effective connectivity via changes in neuronal gains, as demonstrated here in the two periods of the reach-to-grasp experiment and in our simulations for networks close to criticality (Panel D of Figure 3.5 on page 40). During transient activity, on top of gain changes, correlated external inputs may directly drive specific neural modes to create different motor outputs, thereby restricting the dynamics to certain subspaces of the manifold. In fact, Elsayed et al. (2016) reported that the covariance structures during movement preparation and movement execution are unrelated and

corresponding to orthogonal spaces within a larger manifold. Also Luczak et al. (2009) studied auditory and somatosensory cortices of awake and anesthetized rats during spontaneous and stimulus-evoked conditions and found that neural modes of stimulus-evoked activity lie in subspaces of the neural manifold spanned by the spontaneous activity. Similarly, visual areas V1 and V2 seem to exploit distinct subspaces for processing and communication (Semedo et al., 2019), and motor cortex uses orthogonal subspaces capturing communication with somatosensory cortex or behavior-generating dynamics (Perich et al., 2020). Gallego et al. (2018) further showed that manifolds are not identical, but to a large extent preserved across different motor tasks due to a number of task-independent modes. This leads to the hypothesis that the here described mechanism for long-range cooperation in the dynamically balanced critical state provides the basis for low-dimensional activity by creating such spatially extended neural modes, whereas transient correlated inputs lead to their differential activation for the respective target outputs. The spatial spread of the neural modes thereby leads to a distributed representation of information that may be beneficial to integrate information into different computations that take place in parallel at various locations. Further investigation of these hypotheses is an exciting endeavor for the years to come.

3.9 METHODS

Experimental design and statistical analysis

Two adult macaque monkeys (monkey E - female, and monkey N - male) are recorded in behavioral experiments of two types: resting state and reach-to-grasp. The recordings of neuronal activity in motor and pre-motor cortex (hand/arm region) are performed with a chronically implanted $4 \times 4 \text{ mm}^2$ Utah array (Blackrock Microsystems). Details on surgery, recordings, spike sorting and classification of behavioral states can be found in (Riehle et al., 2013, 2018; Brochier et al., 2018; Dąbrowska et al., 2020). All animal procedures were approved by the local ethical committee (C2EA 71; authorization A1/10/12) and conformed to the European and French government regulations.

Resting state data

During the resting state experiment, the monkey is seated in a primate chair without any task or stimulation. Registration of electrophysiological activity is synchronized with a video recording of the monkey's behavior. Based on this, periods of "true resting state" (RS), defined as no movements and eyes open, are chosen for the analysis. Eye movements and minor head movements are included. Each monkey is recorded twice, with a session lasting approximately 15 and 20 min for monkeys E (sessions E1 and E2) and N (sessions N1 and N2), respectively, and the behavior is classified by visual inspection with single second precision, resulting in 643 and 652 s of RS data for monkey E and 493 and 502 s of RS data for monkey N.

Reach-to-grasp data

In the reach-to-grasp experiment, the monkeys are trained to perform an instructed delayed reach-to-grasp task to obtain a reward. Trials are initiated by a monkey closing a switch (TS, trial start). After 400 ms a diode is illuminated (WS, warning signal), followed by a cue after another 400 ms (CUE-ON), which provides partial information about the upcoming trial. The cue lasts 300 ms and its removal (CUE-OFF) initiates a 1 s preparatory period, followed by a second cue, which also serves as GO signal. Two epochs, divided into 200 ms sub-periods, within such defined trials are chosen for analysis: the first 400 ms after TS (starting period, S1 and S2), and the 400 ms directly following CUE-OFF (preparatory period, P1 and P2) (cf. Panel A of Figure 3.6 on page 42). Five selected sessions for monkey E and eight for monkey N provide a total of 510 and 1111 correct trials, respectively. For detailed numbers of trials and single units per recording session see Appendix Table B.1 on page 99.

Separation of putative excitatory and inhibitory neurons

Offline spike-sorted single units (SUs) are separated into putative excitatory (broad-spiking) and putative inhibitory (narrow-spiking) based on their spike waveform width (Bartho et al., 2004; Kaufman et al., 2010, 2013; Peyrache et al., 2012; Peyrache and Destexhe, 2019). The width is defined as the time (number of data samples) between the trough and peak of the waveform. Widths of all average waveforms from all selected sessions (both resting state and reach-to-grasp) per monkey are collected. Thresholds for “broadness” and “narrowness” are chosen based on the monkey-specific distribution of widths, such that intermediate values stay unclassified. For monkey E the thresholds are 0.33 ms and 0.34 ms, and for monkey N 0.40 ms and 0.41 ms. Next, a two step classification is performed session by session. Firstly, the thresholds are applied to average SU waveforms. Secondly, the thresholds are applied to SU single waveforms and a percentage of single waveforms pre-classified as the same type as the average waveform is calculated. SU for which this percentage is high enough are marked classified. All remaining SUs are grouped as unclassified. We verify the robustness of our results with respect to changes in the spike sorting procedure in Appendix Section B.2.

Synchrofacts, i.e., spike-like synchronous events across multiple electrodes at the sampling resolution of the recording system (1/30 ms) (Torre et al., 2016), are removed. In addition, only SUs with a signal-to-noise ratio (Hatsopoulos et al., 2004) of at least 2.5 and a minimal average firing rate of 1 Hz are considered for the analysis, to ensure enough and clean data for valid statistics.

Statistical analysis

All RS periods per resting state recording are concatenated and binned into 1 s bins. Next, pairwise covariances of all pairs of SUs are calculated according to the following formula:

$$\text{COV}(i, j) = \frac{\langle b_i - \mu_i, b_j - \mu_j \rangle}{l - 1} , \quad (3.4)$$

with b_i, b_j - binned spike trains, μ_i, μ_j being their mean values, l the number of bins, and $\langle x, y \rangle$ the scalar product of vectors x and y . Obtained values are broadly distributed, but low on average in every recorded session: in session E1 E-E pairs: 0.19 ± 1.10 (M $\pm$ SD), E-I: 0.24 ± 2.31 , I-I: 0.90 ± 4.19 , in session E2 E-E: 0.060 ± 1.332 , E-I: 0.30 ± 2.35 , I-I: 1.0 ± 4.5 , in session N1 E-E: 0.24 ± 1.13 , E-I: 0.66 ± 2.26 , I-I: 2.4 ± 4.9 , in session N2 E-E: 0.41 ± 1.47 , E-I: 1.0 ± 3.1 , I-I: 3.9 ± 7.3 .

To explore the dependence of covariance on the distance between the considered neurons, the obtained values are grouped according to distances between electrodes on which the neurons are recorded. For each distance the average and variance of the obtained distribution of cross-covariances is calculated. The variance is additionally corrected for bias due to a finite number of measurements (Dahmen et al., 2019). In most of cases, the correction does not exceed 0.01%.

In the following step, exponential functions $y = a e^{-\frac{x}{R}}$ are fitted to the obtained distance-resolved variances of cross-covariances (y corresponding to the variance and x to distance between neurons), which yields a pair of values (a, R). The least squares method implemented in the Python `scipy.optimize` module (SciPy v.1.4.1) is used. Firstly, three independent fits are performed to the data for excitatory-excitatory, excitatory-inhibitory, and inhibitory-inhibitory pairs. Secondly, analogous fits are performed, with the constraint that the decay constant R should be the same for all three curves.

Covariances in the reach-to-grasp data are calculated analogously but with different time resolution. For each chosen sub-period of a trial, data are concatenated and binned into 200 ms bins, meaning that the number of spikes in a single bin corresponds to a single trial. The mean of these counts normalized to the bin width gives the average firing rate per SU and sub-period. The pairwise covariances are calculated according to Eq. (3.4). To assess the similarity of neuronal activity in different periods of a trial, Pearson product-moment correlation coefficients are calculated on vectors of SU-resolved rates and pair-resolved covariances. Correlation coefficients from all recording sessions per monkey are separated into two groups: using sub-periods of the same epoch (*within*-epoch), and using sub-periods of different epochs of a trial (*between*-epochs). These groups are tested for differences with significance level $\alpha = 0.05$. Firstly, to check if the assumptions for parametric tests are met, the normality of each obtained distribution is assessed with a Shapiro-Wilk test, and the equality of variances with an F-test. Secondly, a t-test is applied to compare *within*- and *between*-epochs correlations of rates or covariances. Since there are two *within* and four *between* correlation values per recording session, the number of degrees of freedom equals: $df = (N_{\text{sessions}} \cdot 2 - 1) + (N_{\text{sessions}} \cdot 4 - 1)$, which is 28 for monkey E and 46 for monkey N. To estimate the confidence intervals for obtained differences, the mean difference between groups m and their pooled standard deviation s are calculated for each comparison

$$m = m_{\text{within}} - m_{\text{between}},$$

$$s = \sqrt{\frac{(N_{\text{within}} - 1)s_{\text{within}}^2 + (N_{\text{between}} - 1)s_{\text{between}}^2}{N_{\text{within}} + N_{\text{between}} - 2}},$$

with m_{within} and m_{between} being the mean, s_{within} and s_{between} the standard deviation and N_{within} and N_{between} the number of *within*- and *between*-epoch correlation coefficient values, respectively.

This results in 95% confidence intervals $m \pm t(df) \cdot s$ of 0.192 ± 0.093 for rates and 0.32 ± 0.14 for covariances in monkey E and 0.19 ± 0.14 for rates and 0.26 ± 0.17 for covariances in monkey N.

For both monkeys the *within*-epoch rate-correlations distribution does not fulfill the normality assumption of the t-test. We therefore perform an additional non-parametric Kolmogorov-Smirnov test for the rate comparison. The differences are again significant; for monkey E $D = 1.00$, $p = 6.66 \cdot 10^{-8}$; for monkey N $D = 1.00$, $p = 8.87 \cdot 10^{-13}$.

For all tests we use the implementations from the Python `scipy.stats` module (SciPy v.1.4.1).

Mean and variance of covariances for a two-dimensional network model with excitatory and inhibitory populations

The mean and variance of covariances are calculated for a two-dimensional network consisting of one excitatory and one inhibitory population of neurons. The connectivity profile $p(x)$, describing the probability of a neuron having a connection to another neuron at distance x , decays with distance. We assume periodic boundary conditions and place the neurons on a regular grid (Panel A of Figure 3.3 on page 35), which imposes translation and permutation symmetries that enable the derivation of closed-form solutions for the distance-dependent mean and variance of the covariance distribution. These simplifying assumptions are common practice and simulations show that they do not alter the results qualitatively.

As mentioned in Section 3.4, our aim is to find an expression for the mean and variance of covariances as functions of distance between two neurons. While the theory in Dahmen et al. (2019) is restricted to homogeneous connections, understanding the spatial structure of covariances here requires us to take into account the spatial structure of connectivity. We start with the equations for mean and variance of the covariances of the network model Eq. (3.1)

$$\begin{aligned}\bar{c} &= [\mathbf{1} - \mathbf{M}]^{-1} \frac{D}{1 - \rho^2} [\mathbf{1} - \mathbf{M}]^{-\text{T}}, \\ \overline{\delta c^2} &= [\mathbf{1} - \mathbf{S}]^{-1} \left(\frac{D}{1 - \rho^2} \right)^2 [\mathbf{1} - \mathbf{S}]^{-\text{T}},\end{aligned}$$

and use that $\mathbf{M}$ and $\mathbf{S}$ have a similar structure, such that mean and variance can be derived in the same way, which is why we only consider variances in the following.

To simplify Eq. (3.1), we need to find a basis in which $\mathbf{S}$, and therefore also $\mathbf{A} = \mathbf{1} - \mathbf{S}$, is diagonal. Due to invariance under translation, the translation operators $\mathbf{T}$ and the matrix $\mathbf{S}$ have common eigenvectors, which can be derived using that translation operators satisfy $\mathbf{T}^N = \mathbf{1}$, where N is the number of lattice sites in x - or y -direction (see Appendix Section B.8). Projecting onto a basis of these eigenvectors shows that

the eigenvalues s_k of S are given by a discrete two-dimensional Fourier transform of the connectivity profile

$$s_k \propto \sum_{\mathbf{x}} p(\mathbf{x}) e^{-i\mathbf{k}\mathbf{x}} .$$

Expressing A^{-1} in the eigenvector basis yields $A^{-1}(\mathbf{x}) = \mathbf{1} + \mathbf{B}(\mathbf{x})$, where $\mathbf{B}(\mathbf{x})$ is a discrete inverse Fourier transform of the kernel $s_k / (1 - s_k)$. Assuming a large network with respect to the connectivity profiles allows us to take the continuum limit

$$B(\mathbf{x}) = \frac{1}{(2\pi)^2} \int d^2k \frac{s(\mathbf{k})}{1 - s(\mathbf{k})} e^{i\mathbf{k}\mathbf{x}} .$$

As we are only interested in the long-range behavior, which corresponds to $|\mathbf{x}| \rightarrow \infty$, or $|\mathbf{k}| \rightarrow 0$, respectively, we can approximate the Fourier kernel around $|\mathbf{k}| \approx 0$ by a rational function, quadratic in the denominator, using a Padé approximation. This allows us to calculate the integral which yields

$$B(\mathbf{x}) \propto K_0(-|\mathbf{x}|/R_{\text{eff}}) ,$$

where $K_0(x)$ denotes the modified Bessel function of second kind and zeroth order (Olver et al., 2010), and the effective decay constant R_{eff} is given by Eq. (3.2). In the long-range limit the modified Bessel function behaves like

$$B(\mathbf{x}) \stackrel{|\mathbf{x}| \rightarrow \infty}{\propto} \frac{\exp(-|\mathbf{x}|/R_{\text{eff}})}{\sqrt{|\mathbf{x}|}} .$$

Writing Eq. (3.1) in terms of $B(\mathbf{x})$ gives

$$\overline{\delta c^2}(\mathbf{x}) = \left(\frac{D}{1 - \rho^2} \right)^2 [\delta(|\mathbf{x}|) + B(\mathbf{x}) + (B * B)(\mathbf{x})] ,$$

with the double asterisk denoting a two-dimensional convolution. $(B * B)(\mathbf{x})$ is a function proportional to the modified Bessel function of second kind and first order (Olver et al., 2010), which has the long-range limit

$$(B * B)(\mathbf{x}) \stackrel{|\mathbf{x}| \rightarrow \infty}{\propto} \sqrt{|\mathbf{x}|} \exp(-|\mathbf{x}|/R_{\text{eff}}) .$$

Hence, the effective decay constant of the variances is given by R_{eff} . Note that further details of the above derivation can be found in the Appendix Section B.4–Section B.12.

Network model simulation

The explanation of the network state dependence of covariance patterns presented in the main text is based on linear-response theory, which has been shown to yield results quantitatively in line with non-linear network models, in particular networks of spiking leaky integrate-and-fire neuron models (Tetzlaff et al., 2012; Trousdale et al., 2012; Pernice et al., 2012; Grytskyy et al., 2013; Helias et al., 2013; Dahmen et al., 2019). The derived mechanism is thus largely model independent. We here chose to illustrate it with a particularly simple non-linear input-output model, the rectified linear unit (ReLU). In this model, a shift of the network’s working point can turn some neurons

completely off, while activating others, thereby leading to changes in the effective connectivity of the network. In the following, we describe the details of the network model simulation.

We performed a simulation with the neural simulation tool NEST (Jordan et al., 2019) using the parameters listed in Appendix Table B.4 on page 131. We simulated a network of N inhibitory neurons (threshold_lin_rate_ipn, Hahne et al., 2017), which follow the dynamical equation

$$\tau \frac{dz_i}{dt} = -z_i + \sum_j J_{ij} v_j + \mu_{\text{ext},i} + \xi_i \sqrt{\tau} \sigma_{\text{noise},i}, \quad (3.5)$$

where z_i is the input to neuron i , v the output firing rate with (threshold linear activation function)

$$v = \phi(z) = \begin{cases} 0 & \text{for } z \leq 0 \\ z & \text{for } z > 0 \end{cases},$$

time constant τ , connectivity matrix J , a constant external input $\mu_{\text{ext},i}$, and uncorrelated Gaussian white noise $\langle \xi_i(t) \rangle = 0$, $\langle \xi_i(s) \xi_j(t) \rangle = \delta_{ij} \delta(s - t)$, with noise strength $\sqrt{\tau} \sigma_{\text{noise},i}$. The neurons were connected using the fixed_indegree connection rule, with connection probability p , indegree $K = p \cdot N$, and delta-synapses (rate_connection_instantaneous) of weight w .

The constant external input $\mu_{\text{ext},i}$ to each neuron was normally distributed, with mean μ_{ext} , and standard deviation σ_{ext} . It was used to set the firing rates of neurons, which, via the effective connectivity, influence the intrinsically generated covariances in the network. The two parameters μ_{ext} and σ_{ext} were chosen such that, in the stationary state, half of the neurons were expected to be above threshold. Which neurons are active depends on the realization of $\mu_{\text{ext},i}$ and is therefore different for different networks.

To assess the distribution of firing rates, we first considered the static variability of the network and studied the stationary solution of the noise-averaged input $\langle z \rangle_{\text{noise}}$, which follows from Eq. (3.5) as

$$\langle z_i \rangle_{\text{noise}} = \sum_j J_{ij} \langle v_j \rangle_{\text{noise}} + \mu_{\text{ext},i}. \quad (3.6)$$

Note that $\langle v_j \rangle_{\text{noise}} = \langle \phi(z_j) \rangle_{\text{noise}}$, through the nonlinearity ϕ , in principle depends on fluctuations of the system. This dependence is, however, small for the chosen threshold linear ϕ , which is only nonlinear in the point $z = 0$.

The derivation of μ_{ext} is based on the following mean-field considerations: according to Eq. (3.6) the mean input to a neuron in the network is given by the sum of external input and recurrent input

$$\mu = \mu_{\text{ext}} + \mu_{\text{recurrent}} = \mu_{\text{ext}} + Kw \text{Mean}(v).$$

The variance of the input is given by

$$\sigma^2 = \sigma_{\text{ext}}^2 + \sigma_{\text{recurrent}}^2 = \sigma_{\text{ext}}^2 + Kw^2 \text{Var}(v).$$

The mean firing rate can be calculated using the diffusion approximation (Tuckwell, 1988a; Amit and Tsodyks, 1991), which is assuming a normal distribution of inputs due to the central-limit theorem, and the fact that a linear threshold neuron only fires if its input is positive

$$\begin{aligned}
 \text{Mean}(v) &= \int_{-\infty}^{\infty} dv \mathcal{P}(\mu, \sigma^2, v) v \\
 &= \int_{-\infty}^{\infty} dz \mathcal{N}(\mu, \sigma^2, z) \phi(z) \\
 &= \int_0^{\infty} dz \mathcal{N}(\mu, \sigma^2, z) z \\
 &= \frac{\sigma}{\sqrt{2\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right) + \frac{\mu}{2} \left[1 + \text{erf}\left(\frac{\mu}{\sqrt{2}\sigma}\right)\right],
 \end{aligned}$$

where $\mathcal{P}$ denotes the probability density of the firing rate v . The variance of the firing rates is given by

$$\begin{aligned}
 \text{Var}(v) &= \text{Mean}(v^2) - \text{Mean}(v)^2 \\
 &= \frac{\mu^2}{4} \left[1 - \text{erf}^2\left(\frac{\mu}{\sqrt{2}\sigma}\right)\right] \\
 &\quad + \frac{\sigma^2}{2} \left[1 - \frac{1}{\pi} \exp\left(-\frac{\mu^2}{\sigma^2}\right) + \text{erf}\left(\frac{\mu}{\sqrt{2}\sigma}\right)\right] \\
 &\quad + \frac{\mu\sigma}{\sqrt{2\pi}} \text{erf}\left(\frac{\mu}{\sqrt{2}\sigma}\right).
 \end{aligned}$$

The number of active neurons is the number of neurons with a positive input, which we set to be equal to $N/2$

$$\begin{aligned}
 \frac{N}{2} &\stackrel{!}{=} N \int_0^{\infty} dz \mathcal{N}(\mu, \sigma^2, z) \\
 &= \frac{N}{2} \left[1 + \text{erf}\left(\frac{\mu}{\sqrt{2}\sigma}\right)\right],
 \end{aligned}$$

which is only fulfilled for $\mu = 0$. Inserting this condition simplifies the equations above and leads to

$$\mu_{\text{ext}} = -\frac{Kw\sigma}{\sqrt{2\pi}}.$$

For the purpose of relating synaptic weight w and spectral bound ρ , we can view the nonlinear network as an effective linear network with half the population size (only the active neurons). In the latter case, we obtain

$$w = -\frac{\rho}{\sqrt{\frac{N}{2}p(1-p)}}.$$

For a given spectral bound ρ , this relation allows us to derive the value

$$\mu_{\text{ext}} = \sqrt{\frac{Np}{\pi(1-p) - (\pi-1)\rho^2}} \sqrt{\sigma_{\text{ext}}^2} \rho, \quad (3.7)$$

that, for a arbitrarily fixed σ_{ext} (here $\sigma_{\text{ext}} = 1$), makes half of the population being active. We were aiming for an effective connectivity with only weak fluctuations in the stationary state. Therefore, we fixed the noise strength for all neurons to the small value $\sigma_{\text{noise}} = 0.1 \ll \sigma_{\text{ext}}$ compared to the external input, such that the noise fluctuations did not have a large influence on the calculation above that determines which neurons were active.

To show the effect of a change in the effective connectivity on the covariances, we simulated two networks with identical connectivity, but supplied them with slightly different external inputs. This was realized by choosing

$$\mu_{\text{ext},i}^{(\alpha)} = \mu_{\text{ext},i} + \mu_{\text{ext},i}^{(\alpha)},$$

with

$$\begin{aligned} \mu_{\text{ext},i} &\sim \mathcal{N}(\mu_{\text{ext}}, [1 - \epsilon] \sigma_{\text{ext}}^2), \\ \mu_{\text{ext},i}^{(\alpha)} &\sim \mathcal{N}(0, \epsilon \sigma_{\text{ext}}^2), \end{aligned}$$

$\epsilon \ll 1$, and $\alpha \in \{1, 2\}$ indexing the two networks. The main component $\mu_{\text{ext},i}$ of the external input was the same for both networks. But, the small component $\mu_{\text{ext},i}^{(\alpha)}$ was drawn independently for the two networks. This choice ensures that the two networks have a similar external input distribution (Panel B1 of Figure 3.5 on page 40), but with the external inputs distributed differently across the single neurons (Panel B2 of Figure 3.5 on page 40). How similar the external inputs are distributed across the single neurons is determined by ϵ .

The two networks have a very similar firing rate distribution (Panel E1 of Figure 3.5 on page 40), but, akin to the external inputs, the way the firing rates are distributed across the single neurons differs between the two networks (Panel E2 of Figure 3.5 on page 40). As the effective connectivity depends on the firing rates

$$W_{ij} = J_{ij} \phi'(v_j),$$

this leads to a difference in the effective connectivities of the two networks and therefore to different covariance patterns, as discussed in Figure 3.5 on page 40.

We performed the simulation for spectral bounds ranging from 0.1 to 0.9 in increments of 0.1. We calculated the correlation coefficient of firing rates and the correlation coefficient of time-lag integrated covariances between N_{sample} neurons in the two networks (Panel D of Figure 3.5 on page 40) and studied the dependence on the spectral bound.

To check whether the simulation was long enough to yield a reliable estimate of the rates and covariances, we split each simulation into two halves, and calculate the correlation coefficient between the rates and covariances from the first half of the simulation with the rates and covariances from the second half. They were almost perfectly correlated (Panel C of Figure 3.5 on page 40). Then, we calculated the correlation coefficients comparing all halves of the first simulation with all halves of the second simulation, showing that the covariance patterns changed much more than the rate patterns (Panel C of Figure 3.5 on page 40).

ACTIVITY PATTERNS IN TWO-DIMENSIONAL NETWORKS

Author contributions:

Under the supervision of Johanna Senk and Moritz Helias, the author carried out all the work presented in this chapter.

4.1 INTRODUCTION

As discussed in Section 1.2, experimental recordings in cortex show a wide range of spatio-temporally coherent population activity patterns across mesoscopic scales, including oscillations (Buzsáki and Draguhn, 2004), standing and traveling waves (Wu et al., 2008; Nauhaus et al., 2009; Muller and Destexhe, 2012; Nauhaus et al., 2012; Sato et al., 2012; Muller et al., 2018), hexagonal activity (Rowland et al., 2016; Nadasdy et al., 2017), spirals, or more complex patterns (Townsend et al., 2015; Denker et al., 2018). Analytically, pattern formation in neuronal networks has often been investigated in the neural field framework, introduced in Section 1.3. A neural field is an approximation of the activity of large populations of neurons that neglects the discrete nature of the neuronal network and assumes the network to be a continuous excitable medium instead. In contrast to theories on a discrete grid, describing the network's activity as a function of a continuous space variable has the advantage of making the equations amenable to standard analysis methods like differentiation or Fourier transforms, and hence they frequently allow deriving analytical results explicitly.

Nevertheless, analytical investigation typically entails simplifying assumptions. In the case of pattern formation in neuronal networks, common simplifications include: studying one-dimensional ring networks rather than sheetlike or three-dimensional networks; only including a single neuronal population, while representing the effects of excitation and inhibition via an effective connectivity; selecting a simple non-linearity for representing the neurons' input-output curves; or assuming instantaneous recurrent input, ignoring any signal transmission delays.

Clearly, neurons do not communicate instantaneously. Actually, neuronal networks have multiple sources of signal transmission delays, and from a modeling perspective they are often separated into two categories: a distance dependent delay due to finite axonal signal propagation speed, and a constant delay representing finite synaptic processing times. It is well known that delays can cause complex behavior in dynamical systems. In particular, delays can induce oscillations and the emergence of waves.

Therefore, neglecting delays appears to be a simplification that might lead to major limitations of a network model.

However, studying delayed systems analytically is much more difficult than studying systems with instantaneous interactions. As a results, pattern formation was initially studied using neural field theories without delays in one-dimensional networks, e.g. Wyller et al. (2007); Folias and Ermentrout (2012), but also in two-dimensional networks: In their seminal paper, Ermentrout and Cowan (1979) present an analysis of lattice periodic patterns, providing an explanation for different types of reported visual hallucinations. Tass (1995) performs a weakly non-linear analysis to derive amplitude equations, explaining which of multiple possible static patterns develop in a network on a square, and expands the theory in Tass (1997) to study stationary oscillatory patterns. While Tass restricts his analysis to non-linearities with only cubic contributions, Ermentrout (1998) reports the amplitude equations for non-linearities with quadratic and cubic contributions. Bressloff et al. (2001) extend the work of Ermentrout and Cowan by incorporating orientation tuning into their model of visual cortex, and Rule et al. (2011) investigate the influence of periodic stimuli on pattern formation.

Nunez (1974) was the first to include a distance-dependent delay in an attempt to model activity observed in EEG recordings. However, only in the last two decades has there been a surge of interest in investigating delayed neural field theories. Roxin et al. (2005) conduct a comprehensive bifurcation analysis of a network model with a single constant delay and extend their work in Roxin and Montbrió (2011), focusing on the influence of the non-linear input-output function. Venkov et al. (2007) analyze the selection of standing over traveling waves in a one-dimensional network with finite transmission velocity using a weakly non-linear analysis approach. In a series of papers (Faye and Faugeras, 2010; Veltz and Faugeras, 2011, 2013), Faugeras and colleagues establish the groundwork for applying methods from functional differential equation theory (Hale and Lunel, 1993; Diekmann et al., 1995) to delayed neural field equations. Similarly, van Gils et al. (2012) show that delayed neural fields can be studied as abstract delay differential equations using sun-star calculus. Together these works offer the foundation for applying standard methods like linear stability analysis, center manifold reduction and normal form calculation to delayed neural fields, which we make use of in this chapter as well. Applying these tools to a scalar neural field model, Veltz (2013) studies the relative roles of constant delays and space dependent delays on neural field dynamics, and Dijkstra et al. (2015) investigate a Pitchfork-Hopf bifurcation in a network with transmission delays analytically. There is much more work on one-dimensional delayed neural fields that we cannot discuss here in detail (Liley et al., 2002; Hutt et al., 2003; Hutt and Atay, 2005, 2007; Coombes et al., 2007; Hutt, 2008; Coombes and Laing, 2009; Visser et al., 2012; Fang and Faye, 2016; Senk et al., 2020; Spek et al., 2020); instead, we would like to refer to Ermentrout (1998); Coombes (2005, 2010); Bressloff (2012); Coombes et al. (2014) for excellent reviews of the neural field literature.

There are also some papers studying pattern formation in two-dimensional delayed networks: Visser et al. (2017) use linear stability analysis and center manifold reduction via sun star calculus to investigate a neural field with sigmoidal non-linearity, wizard hat connectivity, and distance dependent delay on a sphere. Daini et al. (2020) study

a comparable network with homogeneous and heterogeneous connectivity, and use spherical harmonics and perturbation theory to derive amplitudes equations, which they then integrate numerically. Spek et al. (2022) examine a scalar neural field with distance-dependent delay, sum of exponentials connectivity, and bounded non-linearity on a rectangle by transforming the delay differential equation's characteristic integral equation into a linear partial differential equation with boundary conditions.

A natural next step is combining the work of Ermentrout and Cowan (1979) and Tass (1995, 1997) with the tools established for assessing delayed neural field theories (Faye and Augeras, 2010; Veltz and Augeras, 2011; van Gils et al., 2012; Veltz and Augeras, 2013) to examine pattern formation in two-dimensional planar networks, a setting that is particularly relevant for understanding experimentally accessible patterns on mesoscopic scales (cf. Section 1.2). Therefore, in this chapter, we study spontaneous global pattern formation in a two-dimensional sheetlike E-I network with target-agnostic isotropic connectivity, constant synaptic delay, and arbitrary non-linearity analytically and with simulations. We focus on the analytical similarities and differences between one-dimensional and two-dimensional networks, as well as the influence of delay and the neurons' input-output curve on the selection of stationary patterns.

The chapter is structured as follows: First, we describe the network model, as well as the concepts required to comprehend the analysis. Then, we derive the distinct conditions for different types of patterns to emerge, discuss the key difference between one-dimensional and two-dimensional networks, consider the effects of the neuron grid, and provide an in depth examination of pattern formation conditions in networks with Gaussian connectivity profiles. Next, we introduce the analytical methods required to examine the interaction of distinct pattern components, followed by our simulation approach, and a detailed analysis of static pattern formation. Finally, we discuss the assumptions of the presented analysis and provide an outlook.

4.2 THE MODEL

As motivated in the introduction, we are analyzing a network model with an excitatory and an inhibitory population of neurons on a two-dimensional plane, in which the populations receive non-linearly transformed recurrent input with a fixed delay. To account for the distance-dependence of the connectivity, we assume that a neuron's number of outgoing connections decays with the inter-neuron distance and that neurons connect equally in all directions. The model's connectivity is a statistical representation of this connectivity profile. That is, the model is not described in terms of single neurons and a corresponding matrix defining the connections between all neuron pairs. Instead, the connectivity is the synaptic input strength multiplied with the probability of two neurons being connected given their distance. This is equivalent to a network in which all neurons are connected to one another and the synaptic input strength decays with inter-neuron distance. We assume that the spatial spread of the connectivity profiles is solely determined by the source population.

Such a network will have a certain working point, a homogeneous stationary activity state exhibited by the network in the absence of external influences. In the following sections, we study how and what types of spatially coherent activity patterns can

emerge spontaneously around this working point, without any external input. Here, we define the model mathematically, and we perform a Fourier transform of the activity to get a description in terms of a convenient base for characterizing activity patterns.

We start with the equation describing the dynamics of the activity

$$\mathbf{X}(\mathbf{r}, t) = \begin{pmatrix} E(\mathbf{r}, t) \\ I(\mathbf{r}, t) \end{pmatrix}$$

of an excitatory population $E(\mathbf{r}, t)$ and an inhibitory population $I(\mathbf{r}, t)$ at position $\mathbf{r} \in \mathbb{R}^2$ and time t around the working point

$$\begin{aligned} \tau \frac{\partial}{\partial t} \mathbf{X}(\mathbf{r}, t) &= -\mathbf{X}(\mathbf{r}, t) + [\mathbf{M} ** \boldsymbol{\psi}(\mathbf{X})](\mathbf{r}, t - d) \\ &= -\mathbf{X}(\mathbf{r}, t) + \int d^2 r' \mathbf{M}(\mathbf{r} - \mathbf{r}') \cdot \boldsymbol{\psi}(\mathbf{X}(\mathbf{r}', t - d)) \\ &= \begin{pmatrix} -E(\mathbf{r}, t) + \int d^2 r' \{M_E(\mathbf{r} - \mathbf{r}') \psi[E(\mathbf{r}', t - d)] + M_I(\mathbf{r} - \mathbf{r}') \psi[I(\mathbf{r}', t - d)]\} \\ -I(\mathbf{r}, t) + \int d^2 r' \{M_E(\mathbf{r} - \mathbf{r}') \psi[E(\mathbf{r}', t - d)] + M_I(\mathbf{r} - \mathbf{r}') \psi[I(\mathbf{r}', t - d)]\} \end{pmatrix}, \end{aligned} \quad (4.1)$$

$$(4.2)$$

where the first term on the right side forces fluctuations of the population activity around the working point to decay exponentially, while the second term represents the recurrent input to each population, which arrives with the constant delay d . The symbol $**$ denotes the two-dimensional convolution (Baddour, 2011). The target-agnostic but source-specific connectivity matrix is given by

$$\mathbf{M}(\mathbf{r}) = \begin{pmatrix} M_E(\mathbf{r}) & M_I(\mathbf{r}) \\ M_E(\mathbf{r}) & M_I(\mathbf{r}) \end{pmatrix},$$

and we assume the connectivity of the source populations $a \in \{E, I\}$ to be isotropic and to decay with the inter-neuron distance $r = |\mathbf{r}|$

$$\begin{aligned} M_a(\mathbf{r}) &= M_a(r), \\ \lim_{r \rightarrow \infty} M_a(r) &= 0. \end{aligned}$$

The function ψ is represents the neurons' non-linear input-output function, often simply referred to as non-linearity, and we assume $\psi(0) = 0$ because we are studying the activity around the working point.

Clearly, the homogeneous steady state $\mathbf{X}_0 = 0$ is a fixed-point solution of Eq. (4.2), and we can expand the non-linearity around the homogeneous solution in a Taylor series

$$\begin{aligned} \boldsymbol{\psi}(\mathbf{X}) &= \begin{pmatrix} \psi'(0)E(\mathbf{r}, t) \\ \psi'(0)I(\mathbf{r}, t) \end{pmatrix} + \frac{1}{2} \begin{pmatrix} \psi''(0)E^2(\mathbf{r}, t) \\ \psi''(0)I^2(\mathbf{r}, t) \end{pmatrix} + \frac{1}{6} \begin{pmatrix} \psi'''(0)E^3(\mathbf{r}, t) \\ \psi'''(0)I^3(\mathbf{r}, t) \end{pmatrix} + \dots \\ &\equiv \beta_1 \begin{pmatrix} E(\mathbf{r}, t) \\ I(\mathbf{r}, t) \end{pmatrix} + \beta_2 \begin{pmatrix} E^2(\mathbf{r}, t) \\ I^2(\mathbf{r}, t) \end{pmatrix} + \beta_3 \begin{pmatrix} E^3(\mathbf{r}, t) \\ I^3(\mathbf{r}, t) \end{pmatrix} + \dots \\ &\equiv \beta_1 \mathbf{X} + \beta_2 \mathbf{X}^{\odot 2} + \beta_3 \mathbf{X}^{\odot 3} + \dots, \end{aligned}$$

where $X^{\odot n}$ denotes the elementwise n -th power of X . This yields the dynamical equation

$$\begin{aligned} \tau \frac{\partial}{\partial t} X(r, t) = & -X(r, t) + \beta_1 [M ** X](r, t - d) \\ & + \beta_2 [M ** X^{\odot 2}](r, t - d) + \beta_3 [M ** X^{\odot 3}](r, t - d). \end{aligned}$$

We perform a two-dimensional Fourier transform in the space variable to get rid of the convolution in the term linear in the fields

$$\begin{aligned} \tau \frac{\partial}{\partial t} X(k, t) = & \underbrace{-X(k, t) + \beta_1 \widetilde{M}(k) \cdot X(k, t - d)}_{\text{linear part}} \\ & \underbrace{+ \beta_2 \widetilde{M}(k) \cdot [X \otimes \otimes X](k, t - d) + \beta_3 \widetilde{M}(k) \cdot [X \otimes \otimes X \otimes \otimes X](k, t - d)}_{\text{non-linear part}}, \end{aligned} \quad (4.3)$$

with $k = |k|$, and $\otimes \otimes$ denoting the two-dimensional, elementwise convolution

$$[X \otimes \otimes X](k, t) = \int d^2 k' \begin{pmatrix} E(k, t) E(k' - k, t) \\ I(k, t) I(k' - k, t) \end{pmatrix} = \int d^2 k' X(k, t) \odot X(k' - k, t),$$

and the elementwise product $\odot$.

How can Eq. (4.3) be understood? It is a differential equation describing the temporal behavior of the model's Fourier modes $X(k, t)$, which are illustrated in Panel A of Figure 4.1 on page 62. Two-dimensional Fourier modes are planar waves of wavelength $2\pi/|k|$, along the direction of wavevector k . The neural field activity $X(r, t)$ is a sum of the Fourier modes weighted by their *amplitudes*, which describe how strongly a Fourier mode contributes to the activity (Panel B of Figure 4.1 on page 62). Eq. (4.3) defines how the Fourier mode amplitudes change over time. The equation can be understood intuitively by considering its two parts separately:

The linear part describes how each Fourier mode develops in the absence of all other Fourier modes and therefore determines the mode's stability. If the amplitude of a mode decays to zero after it has been excited, for example by an initial stimulus, the mode is said to be stable. If its amplitude grows after being excited, it is called unstable. The mode is said to be critical if it is on the cusp of being stable or unstable, in which case the linear part is no longer sufficient to explain the mode's behavior and the non-linear part has to be studied instead.

The non-linear part describes how distinct Fourier modes interact with one another. Distinct modes can excite or suppress one another, and the non-linear interaction determines under which conditions different modes compete or coexist.

Because Eq. (4.3) comprises terms that enter the equation with a constant delay d , it is a delayed differential equation, which means that the dynamics of the Fourier mode amplitudes at time t depend on the activity at time $t - d$. In the following, we pay special attention to how this delayed feedback affects the dynamics of pattern formation.




Figure 4.1: Fourier modes: a convenient base for describing activity patterns. (A) Different Fourier modes on a square of side length 2π . A mode's wavevector $\mathbf{k}$ determines its properties: the spatial frequency and the direction. (B) Different activity patterns are composed of several Fourier modes. Hexagonal: $\mathbf{k}_1 = 4(1, 0)$, $\mathbf{k}_2 = 4(1/2, \sqrt{3}/2)$, $\mathbf{k}_3 = 4(-1/2, \sqrt{3}/2)$. Square: $\mathbf{k}_1 = (3, 0)$, $\mathbf{k}_2 = (0, 3)$. Rhombic: $\mathbf{k}_1 = (2, 5)$, $\mathbf{k}_2 = (-2, 5)$.

4.3 CONDITIONS FOR PATTERNS TO EMERGE

We would like to understand what types of activity patterns the introduced network model supports depending on its parameters: the connectivity, the non-linearity, and the delay. Therefore, we need to analyze the model's bifurcations: abrupt qualitative changes of the model's activity caused by small changes in the model's parameters, such as the emergence of oscillations or waves from a homogeneous steady state following a modification of the model's delay. To identify possible bifurcation points, we investigate for which model configurations the homogeneous steady state $\mathbf{X}(\mathbf{r}, t) = 0$ is a stable solution of Eq. (4.2). This is the case if none of the Fourier modes $\mathbf{X}(\mathbf{k}, t)$ is unstable.

Therefore, we perform a linear stability analysis, investigating the dynamics of the Fourier modes determined by the linear part of Eq. (4.3)

$$\tau \frac{\partial}{\partial t} \mathbf{X}(\mathbf{k}, t) = -\mathbf{X}(\mathbf{k}, t) + \beta_1 \widetilde{\mathbf{M}}(\mathbf{k}) \cdot \mathbf{X}(\mathbf{k}, t - d) .$$

The procedure is the same as in one-dimensional networks (Senk et al., 2020), with the only difference that one-dimensional Fourier transforms are substituted by their two-dimensional counterparts. We insert a Laplace ansatz $\mathbf{X}(\mathbf{k}, t) = e^{\lambda t} \mathbf{X}(\mathbf{k})$ and find

$$\begin{aligned} 0 &= \Delta \cdot \mathbf{X}(\mathbf{k}) , \\ &\equiv \left(\lambda \tau + 1 - \beta_1 \widetilde{\mathbf{M}}(\mathbf{k}) e^{-\lambda d} \right) \cdot \mathbf{X}(\mathbf{k}) , \end{aligned} \quad (4.4)$$

which only has non-trivial solutions for $\mathbf{X}(\mathbf{k})$ if $\det \Delta = 0$, yielding the characteristic equation

$$0 = (\lambda \tau + 1) e^{\lambda d} \left[(\lambda \tau + 1) e^{\lambda d} - \beta_1 c(\mathbf{k}) \right] , \quad (4.5)$$

with the eigenvalue of the Fourier transformed connectivity matrix, in the following referred to as effective profile,

$$c(\mathbf{k}) = \widetilde{\mathbf{M}}_{\text{E}}(\mathbf{k}) + \widetilde{\mathbf{M}}_{\text{I}}(\mathbf{k}) . \quad (4.6)$$

This gives two types of temporal characteristic values for each Fourier mode:

$$\lambda_{\text{s}} = -\frac{1}{\tau} , \quad (4.7)$$

and an infinite number of solutions

$$\lambda_b(\mathbf{k}) = -\frac{1}{\tau} + \frac{1}{d} W_b \left[\beta_1 c(\mathbf{k}) \frac{d}{\tau} e^{d/\tau} \right] , \quad (4.8)$$

where W_b denotes branch b of the multi-valued Lambert-W function (Corless et al., 1996). Demanding mutual orthonormality, the characteristic vector and its adjoint corresponding to λ_{s} are

$$\mathbf{v}_{\text{s}}(\mathbf{k}) = \begin{pmatrix} -\frac{\widetilde{\mathbf{M}}_{\text{I}}(\mathbf{k})}{\widetilde{\mathbf{M}}_{\text{E}}(\mathbf{k})} \\ 1 \end{pmatrix} , \quad \mathbf{w}_{\text{s}}(\mathbf{k}) = \frac{1}{1 + \frac{\widetilde{\mathbf{M}}_{\text{I}}(\mathbf{k})}{\widetilde{\mathbf{M}}_{\text{E}}(\mathbf{k})}} \begin{pmatrix} -1 & 1 \end{pmatrix} , \quad (4.9)$$

and for $\lambda_b(k)$ we find

$$v_c = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad w_c(k) = \frac{1}{1 + \frac{\tilde{M}_E(k)}{\tilde{M}_I(k)}} \begin{pmatrix} \frac{\tilde{M}_E(k)}{\tilde{M}_I(k)} & 1 \end{pmatrix}. \quad (4.10)$$

The temporal behavior and therefore the stability of the system is determined by the characteristic values λ_s and $\lambda_b(k)$. As long as the real parts of these are negative for all wavevectors k , the homogeneous activity state is stable. λ_s is always negative, hence the subscript “s” for stable. However, depending on the effective profile $c(k)$, the real part of $\lambda_b(k)$ might vanish or be positive at certain values k , rendering the corresponding Fourier modes critical or unstable. Therefore, the structure of the network’s connectivity determines whether the network activity displays patterns and which frequency components contribute to the patterns.

Furthermore, the analysis reveals that each Fourier mode can be understood in terms of two independent components, an always stable component corresponding to v_s and a potentially critical component corresponding to v_c

$$X(k, t) = X_c(k, t) v_c + X_s(k, t) v_s(k).$$

These components determine how the activity of the two neuronal populations is correlated. The stable vector v_s is part of the null-space of the Fourier transformed connectivity $\tilde{M}(k)$, implying that the feedback via the recurrent connectivity vanishes for the respective component, reducing the amplitude $X_s(k, t)$ to decay exponentially. Accordingly, the observable activity is described by v_c , and the excitatory and the inhibitory population will display identical dynamics; a result of the target-agnostic structure of the connectivity, because of which the two populations receive precisely the same input.

In a network in which multiple Fourier modes are unstable simultaneously, the non-linear part of Eq. (4.3) determines how the distinct modes interact and therefore what a possible activity pattern can look like. To improve our understanding of this interaction, we study the model configuration in which the smallest possible number of Fourier modes becomes unstable simultaneously. This corresponds to the critical case, in which the real part of $\lambda_b(k)$ on the largest branch of the Lambert-W function, the principal branch $b = 0$ (Senk et al., 2020, Appendix A3), vanishes for only a single value of k . There are two different scenarios for the real part of $\lambda_0(k)$ to vanish at a single critical value of $k = k_c$:

1. The static case or Pitchfork bifurcation: The effective profile $c(k)$ has a maximum at k_c such that

$$\beta_1 c_{\max} \equiv \beta_1 c(k_c) \stackrel{!}{=} 1. \quad (4.11)$$

The modes with $|k| = k_c$ are critical, and the temporal characteristic value $\lambda(k_c) = 0$ vanishes, which implies that activity patterns emerging via this bifurcation do not oscillate but are static once the network reaches a stationary state.

2. The dynamic case or Hopf bifurcation: The effective profile has a minimum at k_c , and the delay d assumes a critical value d_c such that (Helias et al., 2013, Eq. (10))

$$\beta_1 c_{\min} \equiv \beta_1 c(k_c) \stackrel{!}{\leq} -1, \quad (4.12)$$

$$d \stackrel{!}{=} d_c \equiv \tau \frac{\pi - \arctan\left(\sqrt{\beta_1^2 c_{\min}^2 - 1}\right)}{\sqrt{\beta_1^2 c_{\min}^2 - 1}}. \quad (4.13)$$

The temporal characteristic value is non-zero and given by the imaginary pair $\lambda(k_c) = \pm i\omega$, which implies that activity patterns emerging via this bifurcation oscillate with a temporal frequency ω once the network reaches a stationary state.

Clearly, it is possible that an effective profile fulfills both conditions simultaneously (Pitchfork-Hopf bifurcation, e.g. investigated in Veltz (2013); Dijkstra et al. (2015)), but we will restrict our analysis to a scenario in which only one of the two conditions is met.

4.3.1 Comparison of one- and two-dimensional networks

With the help of the analysis so far, we can understand why patterns formation in two-dimensional networks is much richer than pattern formation in one-dimensional systems. For simplicity, we consider the static case, even though the following arguments similarly hold in the dynamic case.

The situation in a one-dimensional system is shown in Panel A1 of Figure 4.2 on page 66. The effective profile $\beta_1 c(k)$ has maxima that just touch the critical value 1. Below the curve the critical wavevectors $k_{\max}$ are marked. Due to the one-dimensional nature of the system, only two Fourier modes corresponding to $k_{\max}$ and $-k_{\max}$ are critical. As the activity has to be real, the Fourier amplitudes corresponding to the two critical k -values have to be complex conjugates. Hence, the problem reduces to studying the behavior of one critical mode.

The situation for two-dimensional systems is shown in Panel A2 of Figure 4.2 on page 66. The two-dimensional effective profile corresponds to the one-dimensional version, rotated around the origin. Accordingly, not only two Fourier modes are critical simultaneously, but a continuum of modes with k -values on the circle $|k| = k_{\max}$. Fully understanding the behavior of the system therefore requires studying an infinite number of critical modes interacting non-linearly, which cannot be done using standard methods.

4.3.2 The effect of the neuron grid

The origin of this problem is the assumption of a homogeneous continuous neural field that is invariant with respect to arbitrary rotations and shifts. Neuronal networks, however, are made up of separate neurons distributed in space. Introducing a grid of neurons breaks the symmetries of the continuous system. In natural neuronal networks, neurons might be distributed arbitrarily, but we here assume the neurons




Figure 4.2: Comparison of effective profile in one- and two-dimensional, continuous and discrete networks for the static case. **(A)** Continuous network: **(A1)** Effective profile (blue), supported wavevectors (black), and critical wavevectors (orange) in one dimension. **(A2)** Effective profile (background) and critical wavevectors (orange) in two dimensions. **(A3)** One-dimensional projection of the peak of the effective profile in a weakly super-critical state, showing stable (blue) and unstable (orange) modes. **(B)** Discrete network: **(B1)** Effective profile (blue), supported wavevectors (black), and critical wavevectors (orange) in one dimension. **(B2)** Effective profile (background), supported wavevectors (black), and critical wavevectors (orange) in two dimensions. **(B3)** One-dimensional projection of the peak of the effective profile in a weakly super-critical state, showing stable (blue) and unstable (orange) modes.

to be arranged on a regular grid. A grid of neurons given by $\mathbf{R} = n_1\mathbf{a}_1 + n_2\mathbf{a}_2$, with lattice vectors $\mathbf{a}_1, \mathbf{a}_2$, only supports Fourier modes with wavevectors on the respective grid in Fourier space $\mathbf{K} = k_1\mathbf{b}_1 + k_2\mathbf{b}_2$, with $\mathbf{a}_i \cdot \mathbf{b}_j = 2\pi\delta_{ij}$, and Kronecker delta δ_{ij} . A square grid of neurons is therefore mapped to a square grid in Fourier space, a hexagonal grid to a hexagonal grid, and a rhombic grid to a rhombic grid. For simplicity, we assume the neurons to be on a square grid, although other grid geometries work similarly.

Mathematically, introducing a grid of neurons corresponds to replacing the Fourier transform and convolutions in Eq. (4.3) by their discrete counterparts. Clearly, this also affects the shape of the effective profile, which contains the Fourier transformed connectivity profiles. Yet, as long as the neuron grid is dense enough and therefore samples the continuous connectivity profile sufficiently well, the continuous effective profile provides a decent approximation of the discrete version.

Introducing a neuronal grid in a one-dimensional system has no effect on the number of critical modes, as illustrated in Panel B1 of Figure 4.2 on page 66. In two-dimensional networks, however, the continuum of critical k -modes is reduced to a discrete subset of modes, depicted in Panel B2 of Figure 4.2 on page 66. The symmetry of the lattice allows only fixed numbers of k -modes to become critical simultaneously. Assuming a grid spacing such that all k 's are given by integers, we distinguish the following possibilities:

1. The zero-mode at the origin $k_c = (0,0)$ becoming unstable corresponds to a global instability. In the static case, this leads to the emergence of a global up- or down-state, whereas in the dynamic case the populations show oscillating activity.
2. If the critical wavevector k_c is non-zero, the number of solutions to $k_c^2 = k_x^2 + k_y^2$ determines how many modes are critical simultaneously. The critical modes always occur in opposing pairs (k_i, k_j) , with $k_i = -k_j$. A four-mode instability is illustrated in Panel B2 of Figure 4.2 on page 66. The number of solutions can be four (e.g., $(\pm 2, 0)$ and $(0, \pm 2)$), eight (e.g. $(\pm 1, \pm 3)$ and $(\pm 3, \pm 1)$), twelve (e.g. $(\pm 5, 0)$, $(0, \pm 5)$, $(\pm 3, \pm 4)$, and $(\pm 4, \pm 3)$), or more. The critical modes represent planar waves of activity with similar spatial frequency $k_c/2\pi$, but different directions. In the static case, these waves are standing, whereas in the dynamic case they are oscillating with a temporal frequency $\omega/2\pi$, moving along the direction of their wavevector with velocity ω/k_c . As explained in 4.2, the emerging activity pattern is given by a combination of these critical modes with different amplitudes (and phases in the dynamic case), which are determined by the way the modes interact with one another.

Introducing a grid is not the only solution to the issue of infinitely many critical modes: Ermentrout and Cowan (1979) circumvent the problem, by restricting their analysis to doubly-periodic solutions that tessellate the two-dimensional plane. The fact that all interacting modes have the same wavelength, as $|k_i| = k_c$, restricts the solutions to either hexagonal, square, or rhombic symmetries (Panel B of Figure 4.1 on page 62). Tass (1995, 1997) instead study the system on a square and assume periodic boundary conditions. This breaks the symmetry of the system because the continuous two-dimensional Fourier transforms of the connectivity profiles in $c(k)$ are replaced by




Figure 4.3: Gaussian bifurcation profile. (A) Possible types of patterns for given peak values $c_{\min}$ and $c_{\max}$ of the effective profile 4.6. (B) Transition curve determining at which network parameters τ and d delay-induced oscillations occur for a given value of $c_{\min}$. (C) Transition curves indicating which types of patterns can develop when the homogeneous solution loses its stability given the connectivity weights and spatial spreads.

the Fourier transform over the square, which is not invariant with respect to rotations, leading to a finite number of simultaneously critical modes, similar to our ansatz.

Introducing a grid of neurons, however, has another advantage: We are interested in studying the behavior of the network at the critical point. In a simulation, however, it is impossible to target this point precisely. Instead, the network parameters have to be adjusted to bring the network close to the critical point such that only a small number of modes is slightly super-critical, meaning that $c(k)$ is slightly above the critical value of one. An illustration of what happens in this weakly super-critical state in a continuous and a discrete system is shown in Panels A3 and B3 of Figure 4.2 on page 66. Whereas in the continuous case an infinite number of modes around the critical mode k_c becomes critical as well, in the discrete case the vertical distance of the discrete values of $c(k)$ allows systematically adjusting the network parameters such that the maximum of $c(k)$ is lowered until only the critical mode k_c remains critical. As a result, studying the network on a grid allows us to investigate the weakly-super critical state without the additional complication of unstable neighboring modes.

4.3.3 Example: Gaussian connectivity profiles

For understanding the results of the analysis so far, it is helpful to consider a concrete example for the network connectivity: a Gaussian connectivity profile

$$M_a(\mathbf{r}) = \frac{w_a}{2\pi R_a^2} \exp\left(-\frac{r^2}{2R_a^2}\right),$$

with connectivity weight w_a and the spatial spread of the connectivity R_a for each pre-synaptic population $a \in \{E, I\}$. This yields the effective profile

$$c(k) = w_E \exp\left(-\frac{R_E^2 k^2}{2}\right) + w_I \exp\left(-\frac{R_I^2 k^2}{2}\right), \quad (4.14)$$

given by the sum of two Gaussians with standard deviations $1/R_a$. One example of this connectivity profile is shown in Panel A2 of Figure 4.2 on page 66, or Panel A1 for its one-dimensional projection. As explained above, the network's bifurcations are governed by the position of the extrema of the effective profile, which are determined by the spatial spread and connectivity weights. Studying the conditions determining the peak positions and their relative values, we can derive a bifurcation profile, shown in Figure 4.3 on page 68, similar to FIG. 4(e)–(g) in Senk et al. (2020). The details are presented in Appendix Section C.1.

What types of patterns may develop depending on the values of the effective profile peaks $c_{\min}$ and $c_{\max}$ is determined by the stability conditions Eq. (4.11) and Eq. (4.12), which are illustrated in Panel A of Figure 4.3 on page 68. For $c_{\max} < 1$ and $c_{\min} > -1$, the network exhibits homogeneous steady activity. Static patterns may emerge if $c_{\max} > 1$ and $c_{\min} > -1$, and dynamic patterns require $c_{\max} < 1$ and $c_{\min} < -1$. If both, $c_{\max} > 1$ and $c_{\min} < -1$, a Pitchfork-Hopf bifurcation may occur, which we will not discuss here (Veltz, 2013; Dijkstra et al., 2015). The minimum delay required for dynamic patterns to occur for a given $c_{\min}$, described by Eq. (4.13), can be read off Panel B of Figure 4.3 on page 68. Finally, which weights and spatial spreads of the connectivities lead to which types of patterns when the homogeneous solution loses its stability is depicted in Panel C of Figure 4.3 on page 68. It shows that static patterns with multiple interacting Fourier modes need a broader inhibition than excitation, whereas dynamic patterns with multiple unstable modes require a wider excitation. Furthermore, global up- or down-states necessitate a dominant excitation; global oscillations a dominant inhibition. Note that, contrary to Panels A and B of Figure 4.3 on page 68, the results in Panel C of Figure 4.3 on page 68 depend on the details of the connectivity profile. Senk et al. (2020), e.g., used a boxcar-shaped connectivity profile, which leads to qualitatively similar, but slightly different transition curves in their FIG. 4(g).

4.4 STUDYING THE INTERACTION OF FOURIER MODES

Our aim is to figure out what kind of stationary patterns emerge after a possible transitory period, which is governed by how unstable Fourier modes interact with one another. To that end, we have to investigate the non-linear part of Eq. (4.3) in the discretized system, in which the convolutions and Fourier transforms are substituted by their discrete counterparts. The difficulty in understanding the interaction stems from the fact that the convolutions

$$[X \circledast \circledast X](k, t) = \sum_{k'} X(k, t) \odot X(k' - k, t)$$

result in interactions between all Fourier modes, critical and stable. Furthermore, because Eq. (4.3) is a delayed differential equation, standard methods for studying

differential equations are not applicable. To proceed with the analysis, we apply a center manifold reduction (Hale and Lunel, 1993; Carr, 2006; Campbell, 2009) to the dynamical equation Eq. (4.3). In a nutshell, a center manifold reduction allows expressing the stable mode amplitudes as functions of the critical mode amplitudes, resulting in a dynamical equation of the critical modes only. Additionally, the center manifold reduction for delayed differential equations maps the system of delayed differential equations onto a system of ordinary differential equations with topologically equivalent bifurcation behavior, in which the delay is a simple parameter. As we want to study the behavior of the network around the critical point, not just at the critical point, we apply the center manifold reduction for systems with parameters (Qesmi et al., 2006, 2007; Guo and Man, 2011).

To develop an intuition for the center manifold reduction for delay differential equations (DDEs) with parameters, we discuss the center manifold reduction for ordinary differential equations (ODEs), then explain how to deal with parameters, and finally consider DDEs.

4.4.1 Center manifold reduction for ODEs

Performing a center manifold reduction around the fixed point $\mathbf{z} = 0$ of ordinary differential equations of the form

$$\frac{d\mathbf{z}}{dt}(t) = \mathbf{L} \cdot \mathbf{z}(t) + \mathbf{F}(\mathbf{z}(t)) ,$$

with $\mathbf{z} \in \mathbb{R}^n$, matrix $\mathbf{L}$, and non-linearity $\mathbf{F}$ with $\mathbf{F}(0) = 0$ and Jacobian $D\mathbf{F}(0) = 0$, requires determining the eigenvalues Λ of $\mathbf{L}$ and an eigenbasis $(\phi_1, \dots, \phi_n)$ such that

$$\mathbf{L} \cdot \phi_i = \lambda_i \phi_i ,$$

and finding an orthonormal adjoint basis $(\psi_1, \dots, \psi_n)$. We assume, the set of eigenvalues comprises c critical eigenvalues, which are defined to lie on the imaginary axis; the other s eigenvalues having negative real parts. In the following we will distinguish the matrices $\Phi = (\phi_1, \dots, \phi_c)$ and $\Psi = (\psi_1, \dots, \psi_c)^T$, collecting the eigenvectors and adjoint eigenvectors corresponding to the critical eigenvalues, and their stable counterparts Φ_s and Ψ_s . Projecting the differential equation onto the adjoint basis yields two separate equations for the critical modes $\mathbf{x} = \Psi \cdot \mathbf{z} \in \mathbb{R}^c$ and the stable modes $\mathbf{y} = \Psi_s \cdot \mathbf{z} \in \mathbb{R}^s$

$$\begin{aligned} \dot{\mathbf{x}}(t) &= \mathbf{L}_c \cdot \mathbf{x}(t) + \mathbf{f}(\mathbf{x}(t), \mathbf{y}(t)) , \\ \dot{\mathbf{y}}(t) &= \mathbf{L}_s \cdot \mathbf{y}(t) + \mathbf{g}(\mathbf{x}(t), \mathbf{y}(t)) , \end{aligned} \tag{4.15}$$

where $\mathbf{L}_c$, the first $c \times c$ -block of the projection of $\mathbf{L}$ onto the adjoint basis, only has eigenvalues on the imaginary axis, and $\mathbf{L}_s$, the last $s \times s$ -block of the projection of $\mathbf{L}$ onto the adjoint basis, has only eigenvalues with negative real parts.

Given that the growth or decay of the modes near the origin is primarily determined by the real parts of the eigenvalues of the linear operators, we can see that the critical modes have slow dynamics around the origin, whereas the stable modes show a relatively fast decay towards a value determined by the stable mode equation.

As a result, the interesting dynamics near the origin are essentially imposed by the critical modes.

This is formalized by the center manifold theorem, which states that there exists a local center manifold $W^c(0) = \{ (x, y) \in \mathbb{R}^c \times \mathbb{R}^s \mid y = h(x), h(0) = 0, Dh(0) = 0 \}$ around $(x, y) = (0, 0)$, which is tangential to the critical subspace of L and whose shape is determined by the non-linearity F . The dynamics restricted to the center manifold is described by

$$\dot{u}(t) = B \cdot u(t) + f[u(t), h(u(t))] , \quad (4.16)$$

with $u \in \mathbb{R}^c$, as long as u is sufficiently small; for ODEs $B = L_c$ and $f = \Psi \cdot F$. Furthermore, the center manifold theorem states that the stability properties of the zero solution of the differential equation reduced to the center manifold are the same as for the full differential equation, and if $z(t) = (x(t), y(t))$ is a solution of the full ODE, with $z(0) = (x(0), y(0))$ being sufficiently small, there exists a solution $u(t)$, such that as $t \rightarrow \infty$

$$\begin{aligned} x(t) &= u(t) + O(e^{-\gamma t}) , \\ y(t) &= h(u(t)) + O(e^{-\gamma t}) , \end{aligned}$$

with some constant γ . In the literature two different ways of defining h are found: the definition above, which highlights that the center manifold expresses the stable modes in terms of the critical ones, and a convenient definition that allows a more compact representation of the reduced dynamical equations, which we will use in the following. Therefore, we redefine $h \in \mathbb{R}^s$ as $h \in \mathbb{R}^n$, such that we can write

$$z(t) = \Phi \cdot u(t) + h(u(t)) + O(e^{-\gamma t}) , \quad (4.17)$$

with

$$\Psi \cdot h(u(t)) = 0 ,$$

which does not affect the center manifold reduction. Inserting a polynomial ansatz, $h(u(t)) = \sum_{ij} a_{ij} u_i(t) u_j(t) + \mathcal{O}(|u|^3)$ with $a_{ij} \in \mathbb{R}^n$, into Eq. (4.15) and applying the chain rule using Eq. (4.16) yields the center manifold at the critical point to lowest order.

4.4.2 Center manifold reduction for ODEs with parameters

If the linear part L of the differential equation depends on some parameters α , such that the system's eigenvalues are critical for vanishing α , the above procedure only holds for $\alpha = 0$. However, we would like to study interacting Fourier modes in the weakly super-critical regime, which corresponds to non-vanishing α . To generalize the center manifold reduction to differential equations with parameters, the differential equation is rewritten

$$\frac{dz}{dt}(t) = \underbrace{L(0) \cdot z(t)}_{\text{new linear part}} + \underbrace{[L(\alpha) - L(0)] \cdot z(t) + F(z(t))}_{\text{new non-linear part}} ,$$

and the linear and non-linear parts are reinterpreted. Consequently, z needs to be redefined

$$\tilde{z}(t) = \begin{pmatrix} z(t) \\ \alpha(t) \end{pmatrix},$$

with

$$\tilde{L} = \begin{pmatrix} L(0) \\ 0 \end{pmatrix}, \quad \tilde{F}(\tilde{z}) = \begin{pmatrix} [L(\alpha) - L(0)] \cdot z(t) + F(z(t)) \\ 0 \end{pmatrix},$$

where we used that parameters do not change with time, $\frac{d\alpha}{dt} = 0$. Similarly, the polynomial ansatz for the center manifold must be modified to include the parameters α . With these definitions, the center manifold reduction for ODEs can be performed, and the reduced dynamical equation for the critical mode amplitudes u becomes

$$\begin{aligned} \frac{du}{dt}(t) = & B \cdot u(t) + \Psi(0) \cdot (L(\alpha) - L(0)) \cdot [\Phi \cdot u + h(u(t))] \\ & + \Psi(0) \cdot F[\Phi \cdot u + h(u(t))] . \end{aligned}$$

4.4.3 Center manifold reduction for DDEs

A delayed differential equation of the form

$$\frac{d}{dt}z(t) = L_0 z(t) + \sum_{i=1}^M L_i z(t - d_i) + F[z(t), z(t - d_1), \dots, z(t - d_M)] ,$$

with matrices $L_i, i \in \{0, \dots, M\}$, a non-linear vector function F , and temporal delays $0 < d_1 < \dots < d_M \equiv d$, maps an initial condition defined on an initial interval $[-d, 0]$ to the dynamics at a future time t , in contrast to ODEs which describe a system's evolution given an initial condition at a single time point t_0 . As the linear part in such a case may no longer be described by a single matrix, but may instead contain multiple terms for several delays, the standard procedure for center manifold reduction no longer can be applied. Instead, the trick here is to introduce

$$z_t(\theta) = z(t + \theta) , \quad -d \leq \theta \leq 0 ,$$

and to interpret the DDE as a functional differential equation

$$\frac{d}{dt}z_t(\theta) = \begin{cases} \frac{d}{d\theta}z_t(\theta) & , \quad -d \leq \theta < 0 , \\ Lz_t + F(z_t) & , \quad \theta = 0 , \end{cases} \quad (4.18)$$

with linear functional operator L , and non-linear functional operator F , and solutions z_t in the Banach space of continuous mappings $C([-d, 0], \mathbb{R}^n)$. Note that the first part of the functional differential equation merely represents the correct evolution on the initial time interval. In our model the linear operator has the form (cf. Eq. (4.3))

$$Lz_t = L_0 z_t(0) + L_d z_t(-d) .$$

The characteristic equation that can be derived after making a Laplace ansatz (cf. Eq. (4.4)), $\Delta(\lambda_c) \phi = 0$, plays the role of the eigenvalue equation in the ODE case.

The important difference here is that the eigenvectors $\Phi(\theta)$ are no longer vectors in $\mathbb{R}^n$, but vectors in the Banach space of continuous mappings $C([-d, 0], \mathbb{R}^n)$, and thus functions of θ . The choice of eigenvectors depends on the critical characteristic values and therefore the studied bifurcation type (Hale and Lunel, 1993, Chapter 7.4). Similarly to the ODE procedure, one then defines a bilinear form

$$\langle \psi, \phi \rangle = \psi \cdot \phi + \int_{-d}^0 \psi(\sigma + d) L_d \phi(\sigma) d\sigma, \quad (4.19)$$

analogous to the standard scalar product, finds respective adjoint eigenvectors $\Psi(\theta)$, and uses them to project the DDE into a critical and a stable subspace. Finally, a corresponding center manifold theorem guarantees the existence of a center manifold (Hale and Lunel, 1993, Chapter 10.2), coming with a reduced dynamical equation for the critical modes u . Determining the center manifold itself is a little different, as the polynomial ansatz is replaced by

$$h(u(t), \theta) = \sum_{ij} a_{ij}(\theta) u_i(t) u_j(t) + \mathcal{O}(|u|^3), \quad (4.20)$$

with functions $a_{ij}(\theta) \in \mathbb{R}^n$. Inserting the ansatz into the first line of Eq. (4.18) determines a set of linear ODEs, which determines the functions $a_{ij}(\theta)$ up to integration constants; the second line of Eq. (4.18) is then used to fix the free constants.

In order to apply the center manifold reduction to our network model Eq. (4.3), we combine the presented approach for ODEs with parameters with the approach for DDEs, in which case the reduced dynamical equation is given by

$$\begin{aligned} \frac{du}{dt}(t) = & B \cdot u(t) + \Psi(0) \cdot [L(\alpha) - L(0)] (\Phi \cdot u(t) + h(u(t), \alpha)) \\ & + \Psi(0) \cdot F(\Phi \cdot u(t) + h(u(t), \alpha)). \end{aligned} \quad (4.21)$$

More details can be found in Appendix Section C.4.

4.5 SIMULATION

To validate our results, we numerically simulate an E-I network using the simulator NEST (Hahne et al., 2021). For this purpose, we implemented a custom neuron type `polynomial_rate_ipn`, which is based on the provided rate neuron implementation (Hahne et al., 2017), but allows setting the polynomial coefficients of the non-linearity explicitly. We set the neuron model's external noise input to zero (`mu=0.`, `sigma=0.`), and `linear_summation=FALSE`, and choose the synapse type `rate_connection_delayed`, such the dynamic equation for the rate X_i of neuron i is given by

$$\tau \frac{dX_i}{dt}(t) = -X_i(t) + \sum_j J_{ij} \psi[X_j(t-d)],$$

with the custom non-linearity

$$\psi(X) = \beta_1 X + \beta_2 X^2 + \beta_3 X^3.$$

We define N_I inhibitory neurons $\mathcal{I}$ and $N_E = nN_I$ excitatory neurons $\mathcal{E}$. The inhibitory neurons are distributed across a two-dimensional regular grid, and n

excitatory neurons are placed on the same grid point as one inhibitory neuron, respectively. We define the connectivity to be all-to-all, including self-connections (autapses), but no multiple connections between the same neurons (multapses), with Gaussian weights defined by

$$J_{ij} = \begin{cases} p_{ij}w_E & , \text{ if } j \in \mathcal{E} \\ p_{ij}w_I & , \text{ if } j \in \mathcal{I} \end{cases}, \quad p_{ij} = \begin{cases} \frac{\tilde{p}_{ij}}{\sum_{j \in \mathcal{E}} \tilde{p}_{ij}} & , \text{ if } j \in \mathcal{E} \\ \frac{\tilde{p}_{ij}}{\sum_{j \in \mathcal{I}} \tilde{p}_{ij}} & , \text{ if } j \in \mathcal{I} \end{cases}$$

$$\tilde{p}_{ij} = \begin{cases} \exp\left(-\frac{(r_i - r_j)^2}{2R_E^2}\right) & , \text{ if } j \in \mathcal{E} \\ \exp\left(-\frac{(r_i - r_j)^2}{2R_I^2}\right) & , \text{ if } j \in \mathcal{I} \end{cases}$$

taking into account the doubly periodic boundary conditions in the inter-neuronal distance $r_i - r_j$. The weights are normalized such that each neuron receives synapses with a total weight $\sum_j p_{ij}w_j = w_E + w_I$. At the beginning of the simulation, we set the initial mode amplitudes a_j corresponding to wavevector k_j by fixing the activity during an initial period to

$$X(r) = \sum_j a_j \cdot \sin(k_j \cdot r) .$$

All simulation parameters and their chosen values are listed in Appendix Section C.2. The results of the simulations are presented in the following section.

4.6 STATIC PATTERNS

In this section, we finally investigate which static patterns the model can display close to the critical point $\beta_{1c}(k_{\max}) = \beta_{1c_{\max}} = 1$. We introduce the parameter $\epsilon = \beta_{1c_{\max}} - 1$, which is determined by the network's connectivity, and measures the network's distance to the critical point: $\epsilon > 0$ corresponds to the super-critical state, $\epsilon = 0$ to the critical, and $\epsilon < 0$ to the sub-critical state.

For simplicity, we restrict our analysis to the case of four simultaneously unstable modes. The investigation of a scenario with more critical modes does not change the analysis, it merely makes it more laborious. To prepare the center manifold reduction (Section 4.4) of our network model Eq. (4.3), we first analyze the non-linear interaction using the results of the center manifold theorem to identify the dynamical equations of the critical and the involved stable modes.

One of the consequences of the center manifold theorem is that stable mode amplitudes are much smaller than critical mode amplitudes, as long as the critical mode amplitudes are not too large. This allows us to simplify the dynamical equations to identify the modes which we need to calculate the center manifold for. A detailed analysis (Appendix Section C.3) shows that the cubic non-linearity primarily introduces couplings between the critical modes, whereas the quadratic non-linearity introduces

a coupling between critical and stable modes. The simplified dynamical equations for the critical mode amplitudes $x_i(t) = \mathbf{v}_c \cdot \mathbf{u}(\mathbf{k}_i, t)$ are given by

$$\begin{aligned} \tau \frac{\partial}{\partial t} x_i(t) = & -x_i(t) + \beta_1 c(\mathbf{k}_c) x_i(t-d) \\ & + 2\beta_2 c(\mathbf{k}_c) \left[x_i(t-d) y_0(t-d) + \sum_{j \neq i} x_j(t-d) y_{ij}(t-d) \right] \\ & + 3\beta_3 c(\mathbf{k}_c) x_i(t-d) [2x_1(t-d)x_3(t-d) + x_2(t-d)x_4(t-d)] + \mathcal{O}(|x|^4), \end{aligned} \quad (4.22)$$

where we introduced the short-hand notations for the zero-mode amplitude $y_0(t) = \mathbf{v}_c \cdot \mathbf{u}(0, t)$ and the difference mode amplitudes $y_{ij}(t) = \mathbf{v}_c \cdot \mathbf{u}(\mathbf{k}_i - \mathbf{k}_j, t)$. The dynamical equations for the stable modes appearing in Eq. (4.22) are

$$\begin{aligned} \tau \frac{\partial}{\partial t} y_0(t) = & -y_0(t) + \beta_1 c(0) y_0(t-d) \\ & + 2\beta_2 c(0) [x_1(t-d)x_3(t-d) + x_2(t-d)x_4(t-d)] + \mathcal{O}(|x|^3), \end{aligned} \quad (4.23)$$

$$\begin{aligned} \tau \frac{\partial}{\partial t} y_{i(i+1)}(t) = & -y_{i(i+1)}(t) + \beta_1 c(\sqrt{2}k_c) y_{i(i+1)}(t-d) \\ & + 2\beta_2 c(\sqrt{2}k_c) x_i(t-d)x_{i+1}(t-d) + \mathcal{O}(|x|^3), \end{aligned} \quad (4.24)$$

$$\begin{aligned} \tau \frac{\partial}{\partial t} y_{i(i+2)}(t) = & -y_{i(i+2)}(t) + \beta_1 c(2k_c) y_{i(i+2)}(t-d) \\ & + \beta_2 c(2k_c) x_i^2(t-d) + \mathcal{O}(|x|^3), \end{aligned} \quad (4.25)$$

where we assume periodic indices $i+4 \equiv i$. As discussed in Section 4.3, the stable components $z(\mathbf{k}, t) = \mathbf{v}_s(\mathbf{k}) \cdot \mathbf{u}(\mathbf{k}, t)$ decay exponentially

$$\tau \frac{\partial}{\partial t} z(\mathbf{k}, t) = -z(\mathbf{k}, t). \quad (4.26)$$

4.6.1 Interaction of static spatial waves

The center manifold reduction presented in Section 4.4 can now be applied to Eq. (4.22)–Eq. (4.26). The calculation is fairly extensive, which is why we make use of the symbolic mathematics Python library SymPy (Meurer et al., 2017). The details are presented in Appendix Section C.4.

In the four-mode instability scenario the center manifold reduction results in four reduced dynamic equations. However, because the activity of the populations must be a real quantity, we know that the amplitudes of the critical mode pairs $\mathbf{k}_i$ and $\mathbf{k}_{i+2} = -\mathbf{k}_i$ must be complex conjugates. Consequently, it is sufficient to study the dynamics of two of the four modes. The reduced dynamic equations for the two critical mode amplitudes are

$$(\tau + \beta_1 d) \dot{u}_1 = u_1 \left[\epsilon + \left(A |u_1|^2 + B |u_2|^2 \right) \right], \quad (4.27)$$

$$(\tau + \beta_1 d) \dot{u}_2 = u_2 \left[\epsilon + \left(B |u_1|^2 + A |u_2|^2 \right) \right], \quad (4.28)$$




Figure 4.4: Static patterns for four-mode instability: Numerically integrated reduced dynamic equations and illustrative snapshots of the inhibitory activity at three different time points of a corresponding NEST simulation using the same network configuration (squares). The fixed parameters are listed in Appendix Section C.2. (A) Competing mode configuration with $\beta_2 = 0.075$ and initial mode amplitudes $2.03e-3$ and $1.98e-3$. Simulation at 640 ms, 2000 ms, and 3000 ms. (B) Coexisting modes configuration with $\beta_2 = 0.15$ and initial mode amplitudes $4.00e-3$ and $0.40e-3$. Simulation at 640 ms, 850 ms, and 3000 ms. (C) Coexisting mode configuration for three different delays $d = 0, d = 3, d = 10$.

with

$$A = \beta_2^2 [2\Lambda(0) + \Lambda(2k_c)] + 3\beta_3, \quad (4.29)$$

$$B = \beta_2^2 [2\Lambda(0) + 4\Lambda(\sqrt{2}k_c)] + 6\beta_3, \quad (4.30)$$

$$\Lambda(k) = \frac{c(k)}{1 - \beta_1 c(k)}, \quad (4.31)$$

where u_1 and u_2 represent the amplitudes of two spatial waves orthogonal to each other.

The reduced dynamical equations have three types of stationary solutions:

1. The homogeneous zero-solution $u_1(t) = u_2(t) = 0$.
2. A static wave along one direction: one mode amplitude is non-zero $u_i(t) = \pm \sqrt{-\frac{\epsilon}{A}}$, and the other one vanishes (Panel A of Figure 4.4 on page 76).
3. A static checkerboard pattern: both mode amplitudes are non-zero with the same amplitudes $u_1(t) = u_2(t) = \pm \sqrt{-\frac{\epsilon}{A+B}}$ (Panel B of Figure 4.4 on page 76).

In the sub-critical state, $\epsilon < 0$, the zero-solution is stable. In the super-critical state, $\epsilon > 0$, the zero-solution is unstable, and one of the other solutions describes the stationary pattern — which one is determined by A and B : $|A| < |B|$ results in a winner-takes-all system, in which only one mode survives; for $|A| > |B|$ the modes coexist.

In Figure 4.4 on page 76 we integrated the reduced dynamic equations numerically for different network parameters and show corresponding snapshots of the inhibitory activity of a network with identical parameters simulated with NEST (Section 4.5). The excitatory population shows similar behavior. We set the initial amplitude of one mode to be slightly higher than the amplitude of the second one. In Panel A of Figure 4.4 on page 76 the network shows winner-takes-all dynamics and only the strongest mode survives, leading to the emergence of vertical stripes. In Panel B of Figure 4.4 on page 76 the network is in the cooperating state, the two modes coexist, and a checkerboard pattern develops.

EFFECT OF THE DELAY As seen on the left hand side of Eq. (4.27) and Eq. (4.28), the delay merely enters the reduced dynamic equations as a modification of the network's integration time constant. Therefore, in the static case, its sole effect is a slowing down of the transient dynamics, as illustrated in Panel C of Figure 4.4 on page 76, where we integrate the reduced dynamical equations for three different delays. This is consistent with findings in one-dimensional systems (Veltz, 2013), and it is expected intuitively for static patterns, as time-independent activity patterns should not depend on the delay once the network has reached the stationary state.

Setting $d = 0$ reproduces the results found for non-delayed two-dimensional systems (Tass, 1995; Ermentrout, 1998).

EFFECT OF THE NON-LINEARITY The reduced dynamic equations unravel the influence of the individual components of the non-linearity on the network dynamics.

The linear slope β_1 occurs in a few places: Most significantly, it enters ϵ and therefore directly controls the network's criticality, that is whether activity patterns emerge or the network remains in a homogeneous state. Similarly, it determines the absolute value of the stationary amplitudes. Furthermore, β_1 enters Eq. (4.27) and Eq. (4.28) in combination with the delay and therefore enhances or decreases the effect of the delay. Finally, as a part of $\Lambda(k)$, β_1 has a non-trivial influence on the interaction of the critical modes.

If the non-linearity only has a cubic part β_3 , the network always exhibits winner-takes-all dynamics, implying that the cubic part promotes a competitive interaction of the modes (Panel A of Figure 4.4 on page 76). The network can only produce patterns with multiple coexisting modes when the non-linearity features a quadratic part β_2 (Panel B of Figure 4.4 on page 76). Eq. (4.23)–Eq. (4.25) demonstrate that, unlike the cubic component, the quadratic part of the non-linearity leads to an excitation of stable modes, which in turn excite the critical modes. Thus, β_2 introduces an indirect coupling of the critical modes via the stable modes. The indirect quadratic interaction may be able to overcome the competitive cubic interaction if it supports a cooperation of the critical modes; this is determined by the connectivity.

EFFECT OF THE CONNECTIVITY First and foremost, the connectivity determines which modes are unstable and therefore the spatial frequency of the static patterns, as discussed in Section 4.3.3. However, Eq. (4.29)–Eq. (4.31) show that the connectivity, or rather its Fourier transforms in the effective profile $c(k)$, also determine the nature of the indirect interaction induced by the quadratic non-linearity. Depending on the values of the effective profile at 0, $\sqrt{2}k_c$, and $2k_c$, the indirect interaction is either competitive or cooperative. Thus, the connectivity is critical in determining which static patterns the network can support.

EFFECT OF HOMOGENEOUS INPUT Adding a homogeneous input to the network amounts to assigning a finite value to the non-linearity at zero input $\psi(0) \neq 0$, which modifies the working point around which the network dynamics are linearized and therefore affects the values of β_1 , β_2 , and β_3 . We have demonstrated that the stability of the homogeneous stationary solution is determined by the effective profile and the slope β_1 of the neurons' input-output function at the working point. As a result, a homogeneous rise of the external input may push a network across a critical point, such that patterns may form; yet, if the external input drives the network into a regime in which the neuron activity saturates, it may also stabilize the homogeneous steady state. Depending on the particular shape of the neurons' input-output functions, a working point change may even affect the way unstable modes interact with each other by modifying the relative strength of the quadratic and cubic contributions, β_2 and β_3 , at the working point.

4.6.2 *Excitation of stable modes*

As a part of the center manifold reduction (Section 4.4), the stable mode amplitudes are expressed in terms of the unstable mode amplitudes. We find the amplitudes for the zero mode and the difference modes as functions of the unstable mode amplitudes




Figure 4.5: Critical and excited stable modes for four-mode instability: Effective profile (background), supported wavevectors (black), critical wavevectors (orange), and excited wavevectors (red). The red dot corresponds to $k = 0$, the red squares to $k = 2k_c$, and the red diamonds to $k = \sqrt{2}k_c$.

u_i , $i \in \{1, 2, 3, 4\}$, to lowest order in the quadratic non-linearity β_2 , the distance to criticality ϵ , and the critical modes u_i

$$h_0(u_1, u_2, u_3, u_4, \epsilon) = \frac{2\beta_{2c}(0)}{1 - \beta_{1c}(0)} (u_1 u_3 + u_2 u_4), \quad (4.32)$$

$$h_{i(i+2)}(u_1, u_2, u_3, u_4, \epsilon) = \frac{\beta_{2c}(2k_c)}{1 - \beta_{1c}(2k_c)} u_i^2, \quad (4.33)$$

$$h_{i(i+1)}(u_1, u_2, u_3, u_4, \epsilon) = \frac{2\beta_{2c}(\sqrt{2}k_c)}{1 - \beta_{1c}(\sqrt{2}k_c)} u_i u_{i+1}. \quad (4.34)$$

The results reveal that the critical modes excite certain stable modes due to the quadratic non-linearity (illustrated in Figure 4.5 on page 79). Each critical mode pair excites the zero-mode proportionally to the strength of the quadratic non-linearity and depending on the value of the effective profile at $k = 0$. Similarly, the critical modes excite the difference modes, which correspond to the wavevectors $k = 2k_c$ and $k = \sqrt{2}k_c$, proportionally to the quadratic non-linearity and dependent on the effective profile at the respective k -values.

Eq. (4.32) indicates that when a static spatial wave develops from a homogeneous state, the overall activity level increases or decreases a little bit, depending on the sign of the prefactor in Eq. (4.32). The other two equations describe the two-dimensional counterparts to higher harmonics in simple oscillators: Eq. (4.33) shows that static spatial waves always excite corresponding spatial waves with twice the spatial frequency, and Eq. (4.34) demonstrates the emergence of spatial waves diagonal to the two critical mode directions.

The center manifold reduction provides equal results as an adiabatic approximation: Around the critical point, the dynamics of the critical modes will be very slow — a universal phenomenon known as critical slowing down. As a result, we can assume that the much faster stable modes always relax to their equilibrium value. Assuming temporarily fixed values for the critical modes, solving Eq. (4.23)–Eq. (4.25) for their stationary solutions, by setting $\partial/\partial t y(t) = 0$ and ignoring the delay, yields solutions that are identical to those obtained by applying the center manifold reduction Eq. (4.32)–Eq. (4.34). However, an adiabatic approximation cannot explain the role of the delay, which does not appear in the amplitudes for the stable modes but does occur in the reduced dynamical equation for the critical modes Eq. (4.27) and Eq. (4.28).

4.7 DISCUSSION

In this chapter, we have shown how to investigate global pattern formation in a two-dimensional sheetlike E-I network with target-agnostic isotropic connectivity, constant synaptic delay, and arbitrary non-linearity analytically and with simulations. We demonstrated that the conditions for pattern formation are identical in one-dimensional and two-dimensional networks, and we explored the key difference between the two, namely the inevitable occurrence of multiple, simultaneously unstable, interacting Fourier modes in two-dimensional networks. We discussed how a neuron grid limits the number of simultaneously unstable modes, and we presented a detailed analysis of the conditions for pattern formation in a network with Gaussian connectivity profiles. Then we studied the non-linear interaction of simultaneously unstable modes for static patterns, and discussed how the network parameters influence the network's pattern formation behavior. We could show that the quadratic part of the non-linearity leads to an excitation of stable modes, which mediates an indirect coupling between the critical modes, determining whether the network exhibits winner-takes-all dynamics or cooperative dynamics. Additionally, we demonstrated that the delay merely slows down the transient activity. Finally, we compared our findings to a numerical network simulation using NEST in combination with a newly implemented rate neuron non-linearity.

In the remainder of this section, we elaborate on the assumptions of the used model and how some of them might be overcome. Finally, we discuss how the presented analysis can be used to study dynamic patterns.

FOUR MODE INSTABILITY For the sake of simplicity, we have only considered the case in which four Fourier modes become unstable simultaneously. Nonetheless, the presented analysis works similarly for scenarios involving more unstable modes. The next higher number of simultaneously unstable modes for a square grid would be eight. The cubic non-linearity leads to an equal interaction between all Fourier mode pairs (cf. Tass, 1995, eq.5.3 for case without delay), as does the indirect coupling via the stable zero mode. But the indirect coupling via the difference modes introduces terms including the effective profile at all difference modes $k_i - k_j$. So, while networks with a purely cubic non-linearity will still show winner-takes-all dynamics, the quadratic non-linearity might introduce more interesting interactions, like groups of cooperating modes.

GLOBAL PATTERNS While the presented analysis can explain lattice periodic patterns, it cannot deal with localized activity such as bumps, breathers, or spirals. Analyzing them necessitates explicitly constructing a localized solution and subsequently evaluating its stability properties. Often these types of patterns are explored in network models that include synaptic adaption or external input breaking the translation symmetry of the model. Bressloff (2012, Section 4) provides an excellent introduction to the subject.

REGULAR GRID & HOMOGENEOUS CONNECTIVITY When developing the simulation for this chapter, we observed that the simulations do not display the anticipated patterns as long as the network has any random component that introduces heterogeneity. Instead, the observed activity exhibited small patches of wavelike activity with the predicted spatial frequencies but varying sizes and orientations. This might suggest that the theory holds locally but fails to capture the global behavior due to the occurrence of small patches with preferred directions that violate the assumption of isotropy. However, this requires further investigation.

TARGET-AGNOSTIC CONNECTIVITY Our model assumes a target-agnostic connectivity, which means that neurons do not distinguish the types of neurons they are projecting to. As a result, both the excitatory and the inhibitory population receive the same input, and the neural field is effectively scalar. This assumption was introduced to allow us to focus solely on the effect of a synaptic delay. As demonstrated in Section 4.3, the delay can induce oscillatory activity. However, oscillations can also be caused by the network's connectivity, namely by I-I and E-I loops, two mechanisms known as Interneuron Network Gamma (ING) and Pyramidal Interneuron Network Gamma (PING) (Whittington et al., 2010). Pattern formation in two-dimensional networks with source- and target-specific connectivity has been studied without taking into account the synaptic delay (Tass, 1995, 1997). Including such a connectivity and a delay results in an intricate interaction of the two oscillation mechanisms. This is illustrated mathematically by the fact that in such a model, the effective profile $c(k)$ in the condition for the emergence of dynamic patterns Eq. (4.12) is no longer a real function, but a complex function. So, in this case, analyzing the conditions for pattern formation necessitates a complex analysis of the Lambert W function, and performing a bifurcation analysis as in Section 4.3.3 requires a categorization of the complex effective profile in terms of all parameters w_{EE} , w_{EI} , w_{IE} , w_{IL} , R_{EE} , R_{EI} , R_{IE} , R_{IL} , and d . This is an interesting approach that might be investigated in future studies.

SHEETLIKE NETWORK TOPOLOGY In the presented analysis, we assume a sheetlike network topology, which means that we assume an infinitely large flat network. This allows us to use standard two-dimensional Fourier modes as base functions for categorizing the patterns. In the course of our analysis, we introduce a neuron grid, and the minimum distance between the neurons puts a high frequency limit on the supported Fourier modes. Investigating a finite patch of network models like in Tass (1995, 1997) works analogously. However, the finite size of the patch puts a low frequency limit on the Fourier modes. As we can only simulate a finite network patch containing a finite number of neurons, we see both limitations in the simulation.

Aside from network size, boundary conditions can play an important role in shaping the details of stationary patterns beyond their frequency and type. We neglect the boundary conditions in our analysis, although we use periodic boundary conditions in the simulation. The analyzed network topology might be a good approximation of pattern formation on the mesoscopic range of a few square millimeters, as real neuronal networks are almost sheetlike on this scale and typically extend much further.

However, the large scale topology of cortex can be essential, especially when investigating large scale activity, as seen in EEG recordings. In this context, we would like to highlight Visser et al. (2017), who put special emphasis on the role of a spherical topology in pattern formation. They explore a scalar neural field with delay on the sphere using methods very similar to those presented in this chapter: they use a combination of linear stability analysis, center manifold reduction, albeit using the slightly different sun-star calculus (Diekmann et al., 1995), and simulations of the neural field model on a mesh, opposed to the simulation of discrete neurons performed here. On the sphere, spherical harmonics, rather than two-dimensional Fourier modes, are the natural option for describing patterns (Daini et al., 2020).

SINGLE DELAY The network model we examined has only a single constant delay. Yet, realistic neuronal networks have at least two sources of delays: a finite synaptic processing time and a distance dependent delay due to a finite signal transmission speed. However, because neurons within a cortical area are fairly close, the delay due to limited transmission speed may be neglected or absorbed into the constant delay for analyzing pattern formation on a mesoscopic scale. Using a single constant delay has the advantage of permitting an analytical derivation that begins with the network model equations and ends with the reduced dynamical equations that clearly illustrate the influence of the delay.

Nevertheless, the key steps of the presented analysis should hold for multiple and even distance dependent delays. In this case, the characteristic equation Eq. (4.5) gets more complicated and can only be solved numerically for the characteristic value determining the stability of the homogeneous steady state. The conditions for bifurcations are defined implicitly by inserting either $\lambda = 0$ or $\lambda = \pm i\omega$ into the characteristic equation Eq. (4.5). This strategy is pursued by Veltz (2013) and Visser et al. (2017). Finally, for static patterns, the delay enters the reduced dynamical equations only through the normalization of the adjoint eigenbasis, where it appears as a numerically solvable integral. Therefore, as expected intuitively, the finding that a constant delay only slows down the transient phase of the formation of static patterns should hold true for multiple and even for distance dependent delays, which confirms earlier findings in one-dimensional networks (Veltz, 2013).

NEXT STEP: INVESTIGATING DYNAMICAL PATTERNS The logical next step in developing our analysis is to apply the developed procedure for investigating interactions of simultaneously unstable modes to examine dynamic patterns. This requires choosing a different eigenbasis for the center manifold reduction, but the analysis remains the same. As dynamic modes have two free parameters, amplitude and phase, the analysis provides twice the number of reduced dynamic equations. The interplay of amplitudes and phases complicates the investigation of the corresponding unstable

modes' interactions. However, employing the developed network simulation and configuring the parameters in such a way that the criterion for a dynamic bifurcation is met, we observed the expected global oscillations, traveling waves, or blinking rolls, as reported by Tass (1997) for delay free networks. What stationary patterns the network exhibits depending on the network model parameters will be investigated in a future project.

DISCUSSION

The previous three chapters present the different studies that constitute the heart of this thesis. In this final chapter, we conclude the thesis by summarizing the results of each study, putting the studies in the context of the topic of this thesis — spatial organization — and discussing the assumptions the presented analyses rely on. Finally, we address the relevance of spatial organization to neuronal network theory and provide an outlook.

5.1 CONCLUSIONS

A TOOLBOX FOR MEAN-FIELD BASED NETWORK MODEL ANALYSIS TOOLS Chapter 2 presents a Python package we developed for collecting and providing easily-usable, well-tested, and well-documented implementations of mean-field results that can be applied to analyze a variety of network models. Mean-field theory and related methods like linear-response theory describe the activity statistics of entire populations of neurons (Section 1.4). While the respective analytical results themselves are insightful about underlying network mechanisms that shape the activity dynamics, these tools may be leveraged even further by numerically solving corresponding model-specific equations. Developing the required implementations, however, can be challenging and time-consuming. Thus, the aim of the presented toolbox is to collect such numerical implementations and to make them accessible to the neuroscientific community. In Chapter 2, we present the modular and extensible architecture of the toolbox and the different workflows for using the tools, showcase how to compute the mean-field estimates of a microcircuit model’s firing rates and power spectra, compare the toolbox to other tools for network model analysis, and discuss use-cases, limitations, and how to contribute new tools.

We argue that employing mean-field techniques for network model analysis comes with considerable advantages, especially for network model development. Performing parameter explorations with simulations requires running multiple simulations with different network model instances to infer the network activity statistics and to find a suitable model configuration. Depending on the network size, this might take a long time and demand a lot of computer power. The numerical computation of analytical estimates, on the other hand, is independent of the number of neurons as long as the theories’ assumptions are met. Furthermore, they can guide parameter explorations, as they directly link a model’s parameters to its activity statistics, allowing for systematic adjustments of network parameters to adapt a network’s behavior to one’s

demands. That being said, applying mean-field tools always requires considering their limitations, since they provide unreliable estimates outside of their valid parameter regimes. Currently, the applicability of the toolbox is limited by the fact that so far primarily methods for the commonly used LIF neuron model have been implemented. Nevertheless, having such a freely accessible package in the first place may foster the development of new methods that can take advantage of the toolbox's solvers and use existing tools as building blocks.

COVARIANCES IN TWO-DIMENSIONAL NETWORKS Chapter 3 studies the spatial structure of spike count covariances in networks with distance-dependent connection probability profiles. Resting state data in macaque motor cortex reveal a salt-and-pepper organization of spike count correlations and covariances with strong positive as well as negative covariances between pairs of neurons separated by several millimeters, independent of the putative neuron types. We demonstrate analytically that such long-range covariance patterns with quickly decaying mean covariances and slowly decaying variance of covariances across several millimeters can be explained by a model with heterogeneous, short-range connectivity. The heterogeneity puts the network into an almost critical regime in which a large number of indirect paths between pairs of neurons effectively influences their coordination. We predict that in this regime the anatomical differences of the spread of excitatory and inhibitory connections should be irrelevant for the spread of the covariances and verify this in the macaque data. Furthermore, we predict that the described dynamically emerging covariance patterns may be adjusted effectively by a shift of the mean network input, which we illustrate using a simple network model simulation. We observe such a significant modification of the covariance structure in two distinct recording epochs that correspond to different attentive states of the monkey.

The proposed coordination mechanism does not require any specific imprinted network connection structure, which would yield less flexible covariance patterns, or any correlated shared external input, which would only result in positive correlations. It is worth noting that this is only one possible explanation for strong correlations. The theory assumes uncorrelated input, making the network mechanism less suitable for understanding sensory areas, which receive strongly correlated external inputs. Rosenbaum et al. (2017), for example, have shown that such networks may exhibit non-monotonously decaying correlations. Smith et al. (2018) demonstrate that a different type of heterogeneity can lead to long range patchy correlations, as seen in visual cortex, whereas the mechanism presented here can explain the salt-and-pepper correlation structure seen in motor cortex and is particularly applicable to resting state activity.

ACTIVITY PATTERNS IN TWO-DIMENSIONAL NETWORKS Chapter 4 studies the effects of a constant synaptic delay and the different neuronal input-output function components on pattern formation and selection in a planar two-dimensional E-I network. A linear stability analysis confirms that the conditions on the network parameters for spatio-temporally structured activity to emerge from homogeneous activity are the same as in the one-dimensional case, and we perform a detailed analysis of the conditions in networks with Gaussian connectivity profiles. However, due to

the continuous rotation symmetry in two-dimensional neural fields, a continuum of Fourier modes becomes critical simultaneously when the homogeneous state loses its stability, and we discuss that this number is finite for discrete networks. We present a center manifold reduction that allows studying the non-linear interaction of the critical modes despite the temporal delay and apply the method to study stationary pattern formation. This results in a set of simplified ordinary differential equations for the mode amplitudes, which clearly discern how the different model components affect pattern formation: the delay primarily slows down the transient dynamics, and the input-output function determines whether the critical modes show winner-takes-all or cooperative dynamics. We integrate the equations numerically and show that a matching network simulation displays the predicted patterns. Finally, we discuss non-linear phenomena related to stable mode excitation.

5.2 SUB-PROJECTS IN THE CONTEXT OF SPATIAL ORGANIZATION

The three sub-projects of this thesis cover complementary aspects of spatially organized neuronal networks:

The package introduced in Chapter 2 provides tools for computing network properties that only indirectly depend on spatial organization, although some methods required for studying pattern formation are included as well. Mean-field estimates of properties such as average firing rates for asynchronous irregular network states are dependent on the number of inputs but otherwise are exclusively determined by single neuron properties; very similar to the transfer function. However, such population properties are crucial for understanding spatially organized networks: as shown in Chapter 3, the spatial covariance structure is dictated by the effective connectivity, which in turn is determined by the neurons' average firing rates and transfer functions. Apart from that, mean-field theory describes network activity averaged over thousands of neurons, providing a coarse-grained perspective of brain dynamics; particularly relevant on the millimeter scale. Furthermore, as mentioned in Section 2.4.2, the toolbox provides a method for mapping a spiking network model to a corresponding rate network model. Such mappings enable the use of results from neural field theory, such as those presented in Chapter 4, to support the analysis of pattern formation in spatially organized spiking networks.

Clearly, the other two sub-projects are related to spatially organized networks, since they study spatial phenomena. They both study spontaneous activity in the absence of any structured external input with a focus on the mesoscopic scale. They do, however, address complementary aspects, namely the spatial structure of covariances in spiking activity in an almost critical regime and coherent population activity in slightly super-critical networks. While the covariance project is clearly tied to the spiking component of multi-electrode recordings, the pattern formation project might be seen in the context of the LFP component. However, many of the ideas of the pattern formation analysis performed in Chapter 4 certainly are transferable to pattern formations in spiking activity. Just to give an example, hexagonal spiking patterns (Rowland et al., 2016; Nadasdy et al., 2017) similarly might be understood in terms of three simultaneously unstable interacting Fourier modes. That is also why we place a major emphasis on the effect of the different components of the non-linearity on

the dynamics: spiking neurons in a realistic regime show an expansive non-linearity, as demonstrated by Roxin et al. (2011), implying a strong influence of the quadratic component on pattern formation in realistic spiking networks.

5.3 OVERARCHING ASSUMPTIONS IN THIS THESIS

Here we discuss assumptions that were introduced to simplify the analyses or render them possible in the first place.

MEAN-FIELD ASSUMPTIONS The mean-field approach clearly plays a major role in the first two sub-projects. The core idea of mean-field analyses is based on the central limit theorem, which allows approximating the recurrent synaptic input by Gaussian noise as long as a neuron receives a large number of uncorrelated inputs. Hence, the mean-field approach is particularly suitable to describe asynchronous irregular activity, which is often observed in cortex (Softky and Koch, 1993; Shadlen and Newsome, 1998) and naturally emerges in networks with a balance between excitation and inhibition (van Vreeswijk and Sompolinsky, 1996, 1998). Small fluctuations and oscillations of network activity around an asynchronous irregular stationary state can be understood in terms of linear-response theory (Lindner and Schimansky-Geier, 2001; Lindner et al., 2005). Strong input correlations, as occurring in sensory areas, and other strongly correlated network states, such as synchronized activity, however, violate the assumptions of mean-field theory and require using different approaches (see e.g., Rosenbaum et al., 2017). A uniform increase of the external input to an entire network, on the other hand, can be understood within the described frameworks, as already discussed in the corresponding chapters. A homogeneous external input does affect the network's working point and therefore its dynamics: It may alter the spatial coordination pattern in the sub-critical regime described in Section 3.6, and it can have multiple effects on spontaneous pattern formation as discussed in Section 4.6.1.

MONOTONOUSLY DECAYING ISOTROPIC HORIZONTAL CONNECTIVITY In this thesis we focus on the effects of an isotropic local horizontal connectivity. On the range of a few hundred micrometers the connection probability decays monotonously (see discussion of literature in Section 1.1), and we demonstrate in this thesis that even such a simple spatial structure leads to interesting non-trivial dynamics. That said, biological networks feature much more elaborate horizontal connectivity, including anisotropic patchy and long range connections (McDonald and Burkhalter, 1993; Livingstone and Hubel, 1984; Bosking et al., 1997; Budd and Kisvárdy, 2001; Lewis et al., 2002; Tanigawa et al., 2005; Buzás et al., 2006; Voges et al., 2010; Boucsein et al., 2011). Such long range connectivity leads to increased coupling across larger distances and can induce a non-monotonous distance dependence of correlations, as seen in distributed orientation specific networks in visual cortex (Hubel and Wiesel, 1959). However, such specific connectivity is not required for long-range correlations as shown in Chapter 3; not even for patchy correlations as demonstrated by Smith et al. (2018), who study the effect of an anisotropic spatial connectivity profile. Incorporating a more complex connectivity in the pattern formation analysis of Chapter 4 is straightforward. This will only affect the concrete shape of the effective

profile, possibly resulting in simultaneously unstable Fourier modes with various frequencies. Breaking the isotropy amounts to breaking the rotation symmetry, which may result in preferred directions of stationary periodic patterns and traveling waves. Note that we implicitly neglect the vertical connectivity between different cortical layers by studying two-dimensional networks and limiting our analyses to only two populations.

TARGET-AGNOSTIC CONNECTIVITY Both Chapter 3 and Chapter 4 at some point make the assumption of a target-agnostic but source-specific connectivity, which means that we assume neurons do not differentiate between the types of neurons to which they connect, but excitatory neurons' connectivity can differ from inhibitory neurons' connectivity. Stepanyants et al. (2004, 2008) indeed report that excitatory neuron axons show no structural specificity towards their targets; inhibitory neurons, however, seem to connect target-specific. Assuming spatial spreads to only depend on the source neuron is in line with experimental results (Levy and Reyes, 2012, Table 3) and (Ohana et al., 2012, Table 1). A source- and target-specific connectivity does not introduce any fundamental difficulties in Chapter 3, since we first obtain the full results for a source- and target-specific connectivity and make the simplifying assumption solely to be able to study the long-range limit analytically. In Chapter 4, however, a source- and target-specific connectivity introduces another potential origin of oscillations, additional to the delay. We discuss a possible approach for studying such systems in Section 4.7. Here, we make this simplification to be able to clearly discern the effects of the delay.

MESOSCOPIC SCALE In this thesis, we focus on the mesoscopic scale covering a few millimeters. The specifics of this scale are highlighted by comparing it to the characteristics of other scales: On the microscopic scale of few hundred micrometers, the distance dependence of the connection probability is irrelevant, and a microscopic network may be well described by a model with random connectivity with uniform connection probabilities. The distance-dependent axonal delay should be insignificant due to the short distances and might be combined with the constant synaptic delay. On the macroscopic scale, which might cover several centimeters up to the entire brain, long-range connectivity and the brain's overall topology play a much more important role. Visser et al. (2017), for example, argue for the importance of suitable boundary conditions for understanding EEG, which might be taken into account by considering a spherical network topology. Furthermore, due to large distances between different parts of the brain, finite axonal transmission speeds may be the dominant contributors to communication delays. The mesoscopic scale is somewhere in the middle of these two extremes. Mesoscopic networks are large enough that the distance-dependence of connection probabilities has to be considered; however, long-range connections may be negligible on this scale. As a local approximation, a planar topology is appropriate for describing a few square millimeters of cortical networks. Whether or not the axonal delay has to be taken into account in addition to the synaptic delay is determined by the local axonal propagation velocities and hence the local myelination (Swadlow and Waxman, 2012).

5.4 RELEVANCE OF SPATIAL ORGANIZATION

In this section, we discuss why spatially organized connectivity is an important property worth investigating.

Obviously, including spatial organization is essential for explaining spatial phenomena like traveling waves or spatially periodic activity, reviewed in Section 1.2; only spatially structured networks can exhibit spatially organized activity.

However, spatial organization is more than merely a feature that may be considered for studying spatial phenomena and ignored otherwise: Keane and Gong (2015), for example, show that propagating waves exhibited by spatially extended balanced spiking neuronal networks result in individual neurons receiving highly synchronized synaptic inputs when the waves travel across the network; they discuss that this has been proposed as an alternative explanation for highly variable spiking activity (Stevens and Zador, 1998) that is consistent with non-Gaussian input distributions as, for example, observed in auditory cortex (DeWeese and Zador, 2006), and that can explain experimentally recorded cross-correlations of nearby pairs of cortical neurons (Okun and Lampl, 2008). These two observations are inconsistent with the mean-field assumption of uncorrelated, Gaussian input distributions, but can be understood by incorporating spatial organization. In a similar vein, Huang et al. (2019) demonstrate that a model with slower inhibition than excitation, as observed in experiments (Salin and Prince, 1996; Geiger et al., 1997; Xiang et al., 1998; Angulo et al., 1999), displays pathologic levels of rhythmic synchronization unless a spatially organized connectivity is incorporated, resulting in asynchronous spiking activity. The presented examples, such as non-Gaussian input distributions and asynchronous activity despite slower inhibition than excitation, highlight experimental findings that do not have a clear link to spatial organization but may be explained by taking it into account.

Networks with spatially structured connectivity are strongly related to networks with functionally structured connectivity, such as visual cortex (Hubel and Wiesel, 1959). In models of orientation selectivity, for example, the connectivity depends on the distance in a functional space, e.g., the orientation ranging from $-\pi/2$ to $\pi/2$, rather than the spatial distance (Ben-Yishai et al., 1995; Hansel and Sompolinsky, 1998; Blumenfeld et al., 2006); they even can be combined, as in the model of visual cortex by Bressloff et al. (2001). For studying such networks the same approaches can be used. Incorporating organization in (functional) space, however, is not the only way of structuring networks: networks can be clustered (Litwin-Kumar and Doiron, 2012), feature feed-forward structures (Diesmann et al., 1999), or have low-rank connectivity (Mastrogiuseppe and Ostojic, 2018), to name a few examples.

In general, connectivity — and especially structured connectivity — leads to a certain redundancy in the activity and hence reduces the activity's dimensionality, which indicates the minimal number of independent latent variables, or neuronal modes, that describe most of the activity's variability. Due to the availability of massively parallel recording methods, there has been a recent surge of interest in the dimensionality of brain activity (Yu et al., 2009; Cunningham and Byron, 2014), and cortical activity has been observed to be on low-dimensional activity manifolds (Gallego et al., 2017, 2018, 2020) or high-dimensional manifolds (Stringer et al., 2019a,b) depending on brain area, task, and stimulus. The occurrence of such activity manifolds requires the structured

coordination of neurons, and as already discussed in Section 3.8, investigating spatially structured networks is crucial for understanding how neurons can coordinate their activity across large distances in the first place. Recent studies have investigated the connection between connectivity and dimensionality in recurrent networks (Recanatesi et al., 2019; Dahmen et al., 2020; Hu and Sompolinsky, 2020) in more detail. Together with other types of structured connectivity, the spatial organization determines the global structure of the activity manifold. Different submanifolds may then be explored in the course of different tasks: Elsayed et al. (2016), for example, find that motor cortex uses orthogonal subspaces for movement preparation and movement execution; Luczak et al. (2009) demonstrate that stimulus evoked activity lies on submanifolds of the manifold explored by spontaneous activity; and Stringer et al. (2019a) find that the dimensionality in visual cortex depends on the dimensionality of the input signal. So, while the global structure of the manifold may be determined by the connectivity, the activity may be restricted dynamically to lower-dimensional subspaces to perform different tasks. However, how the connectivity shapes the structure of the global manifold is an open question that should be explored in future studies.

5.5 OUTLOOK

This thesis explores the theory of spatially organized networks. We demonstrate that spatial structure is a critical feature of biological neuronal networks that needs to be taken into account for understanding a wide range of phenomena observed in real brains, including large-scale coordination, spatio-temporal activity patterns, experimentally observed variability of spiking activity, and network activity dimensionality. We introduce a Python package that allows utilizing the coarse-grained perspective of analytical mean-field results to investigate neuronal network models, we explain how long-range covariances on a mesoscopic scale naturally arise from short-range connectivity, and we investigate the effects of synaptic delays as well as the neuronal input-output function on pattern formation in two-dimensional networks.

As highlighted in the preceding discussions, spatially organized networks, and particularly the theory of spatially organized spiking networks, constitute rich fields of research, the latter of which only recently has gained traction. Future research will certainly establish closer ties between the neural field and the spiking literature, and we expect that the relation between spatial organization and properties of spiking networks, as well as how spatially organized connectivity affects the properties of activity manifolds, will be explored in greater depth. Additionally, we hope to see studies of the effects of more intricate and long-range connectivities on network activity, which would allow studying networks on an intermediate scale at which brain topology may start playing a role. To conclude, the spatial organization of networks and its impact on network activity pose a variety of intriguing questions that will provide plenty of research opportunities for the foreseeable future.

NNMT

A.1 MICROCIRCUIT PARAMETERS

Symbol	Value PD14	Value B16	Description
$K_{4E,4I}$	795	675	In-degree from 4I to 4E
$K_{4E,ext}$	2100	1780	External in-degree to 4E
$D(\omega)$	none	truncated Gaussian	Delay distribution
$d_e \pm \delta d_e$	1.5 ± 0.75 ms	1.5 ± 1.5 ms	Mean and standard deviation of excitatory delay
$d_i \pm \delta d_i$	0.75 ± 0.375 ms	0.75 ± 0.75 ms	Mean and standard deviation of inhibitory delay

Table A.1: Parameter adaptations used here are introduced by Bos et al. (2016) (B16) compared to original microcircuit model. K_{ij} denotes the in-degrees from population j to population i . The delays in the simulated networks were drawn from a truncated Gaussian distribution with the given mean and standard deviation. The mean-field approximation of the microcircuit by Potjans and Diesmann (2014) (PD14) assumes the delay to be fixed at the mean value, which is specified in the toolbox by setting the parameter `delay_dist` to none.

LONG RANGE CORRELATIONS

B.1 CORRELATIONS AND COVARIANCES

A typical measure for the strength of neuronal coordination is the Pearson correlation coefficient, here applied to spike counts in 1 s bins. Correlation coefficients, however, comprise features of both auto- and cross-covariances. From a theoretical point of view, it is simpler to study cross-covariances separately. Indeed, linear-response theory has been shown to faithfully predict cross-covariances in spiking leaky integrate-and-fire networks (Tetzlaff et al., 2012; Trousdale et al., 2012; Pernice et al., 2012; Grytskyy et al., 2013; Helias et al., 2013; Dahmen et al., 2019). Figure B.1 on page 96 justifies the investigation of cross-covariances instead of correlation coefficients for the purpose of this study. It shows that the spatial organization of correlations closely matches the spatial organization of cross-covariances.




Figure B.1: **Correlations and covariances.** The shown data is taken from session E2. (E-E: excitatory-excitatory, E-I: excitatory-inhibitory, I-I: inhibitory-inhibitory). (A) Population-resolved distribution of pairwise spike-count Pearson correlation coefficients. Same data as in Figure 3.1 on page 31C. (B) Population-resolved distribution of pairwise spike-count covariances. (C) Population-resolved distribution of variances. (D) Pairwise spike-count correlation coefficients with respect to the neuron marked by black triangle. Grid indicates electrodes of a Utah array, triangles and circles correspond to putative excitatory and inhibitory neurons, respectively. Size as well as color of markers represent correlation. Neurons within the same square were recorded on the same electrode. Same data as in Figure 3.1 on page 31D. (E) Pairwise spike-count covariances with respect to the neuron marked by black triangle.

B.2 ROBUSTNESS TO E/I SEPARATION




Figure B.2: **Distance-resolved variance of covariance: robustness of decay constant estimation.** Exponential fits (lines) to variances of covariances (dots) analogous to Figure 3.4 on page 38A,B (columns 1&3 and 2&4, respectively) for all analyzed resting state sessions. The two sets of plots differ in E/I separation consistency values chosen during data preprocessing. Panels a-h: default (lowest) required consistency (~ 0.6), used throughout the main analysis; panels i-p: $c = 0.75$. The values of the obtained decay constants are listed in Appendix Table B.1 on page 99.

The analysis of the experimental data involves a number of preprocessing steps, which may affect the resulting statistics. In our study one such critical step is the separation of putative excitatory and inhibitory units, which is partially based on

setting thresholds on the widths of spike waveform, as described in the Methods section. We tested the robustness of our conclusions with respect to these thresholds.

As mentioned in Methods, two thresholds for the width of a spike waveform are chosen, based on all SU average waveforms: A width larger than the “broadness” threshold indicates a putative excitatory neuron. A width lower than the “narrowness” threshold indicates a putative inhibitory neuron. Units with intermediate widths are unclassified. Additionally, to increase the reliability of the classification, we perform it in two steps: first on the SU’s average waveform, and second on all its single waveforms. We calculate the percentage of single waveforms classified as either type. Finally, only SUs showing a high enough percentage of single waveforms classified the same as the average waveform are sorted as the respective type. The minimal percentage required, referred to as consistency c , is initially set to the lowest value which ensures no contradictions between average- and single-waveform thresholding results. While the “broadness” and “narrowness” thresholds are chosen based on all available data for a given monkey, the required consistency is determined separately for each recording session. For monkey N c is set to 0.6 in all but one sessions: In resting state session N1 it is increased to 0.62. For monkey E the values of c equals 0.6 in the resting state recordings and take the following values in five analyzed reach-to-grasp sessions: 0.6, 0.89, 0.65, 0.61, 0.64.

The only step of our analysis for which the separation of putative excitatory and inhibitory neurons is crucial is the fitting of exponentials to the distance-resolved covariances. This step only involves resting state data. To test the robustness of our conclusions, we manipulate the required consistency value for sessions E1, E2, N1, and N2 by setting it to 0.75. Figure B.2 on page 97 and Table B.1 on page 99 summarize the resulting fits.

It turns out that increasing c to 0.75, which implies disregarding about 20 – 25 percent of all data, does not have a strong effect on the fitting results. The obtained decay constants are smaller than for a lower c value, but they stay in a range about an order of magnitude larger than the anatomical connectivity. We furthermore see that fitting individual slopes to different populations in some sessions leads to unreliable results (cf. yellow lines in Panel A and I of Figure B.2 on page 97 and blue lines in Panels C, D, K, and L of Figure B.2 on page 97). Therefore, the data is not sufficient to detect differences in decay constants for different neuronal populations. Fitting instead a single decay constant yields trustworthy results (cf. yellow lines in Panels E and M of Figure B.2 on page 97 and blue lines in Panels G, H, O, and P of Figure B.2 on page 97). Our data thus clearly expose that decay constants of covariances are in the millimeter range.

B.3 STATIONARITY OF BEHAVIORAL DATA

The linear-response theory, with the aid of which we develop our predictions about the covariance structure in the network, assumes that the processes under examination are stationary in time. However, this assumption is not necessarily met in experimental data, especially in motor cortex during active behavioral tasks. For this reason we analyzed the stationarity of average single unit firing rate and pairwise zero time-lag covariance throughout a reach-to-grasp trial, similarly to Dahmen et al. (2019).

c		E1	E2	N1	N2
0.6 (default)	#exc/#inh	56/50	67/56	76/45	78/62
	unclassified	0.078	0.075	0.069	0.091
	relative error	1.1157	1.0055	1.0097	1.0049
	1-slope fit	1.674	1.029	1.676	4.273
	I-I	1.919	0.996	1.647	4.156
	I-E	0.537	1.206	2.738	96100.688
	E-E	1.642	1.308	80308.482	94096.871
0.75	#exc/#inh	45/42	47/48	70/36	74/48
	unclassified	0.24	0.28	0.18	0.21
	relative error	1.1778	1.0141	1.0102	1.0090
	1-slope fit	1.357	0.874	1.420	2.587
	I-I	1.794	0.809	1.394	2.550
	I-E	0.496	1.123	3.682	40.852
	E-E	1.390	1.199	80548.500	10310.780

Table B.1: **Summary of exponential fits to distance-resolved variance of covariance.** For each value of E/I separation consistency c the numbers of sorted putative neurons and the percentages of unclassified units, and therefore not considered for fitting SUs, are listed per resting state session, along with the resulting fits (cf. Figure 3.4 on page 38)




Figure B.3: **Rate and covariance stationarity during a reach-to-grasp trial: monkey E.** Black line indicates population mean and gray area ± 1 population standard deviation of single unit firing rate (left column) and pairwise zero time-lag covariance (right column) during trial of a given session (row). Blue bars indicate starting (S) and preparatory (P) periods used in the analysis (cf. Figure 3.6 on page 42). First, second and fourth dashed lines indicate visual signals lighting up and the third dashed line indicates the removal of a visual cue and beginning of a waiting period.




Figure B.4: **Rate and covariance stationarity during a reach-to-grasp trial: monkey N.** Analogous to Figure B.3 on page 99.

Although the spiking activity becomes highly non-stationary during the movement, those epochs that are chosen for the analysis in our study (S and P) show only moderate variability in time (Figure B.3 on page 99). An analysis on the level of single-unit resolved activity also shows that the majority of neurons has stationary activity statistics within the relevant epochs S and P, especially when comparing to their whole dynamic range that is explored during movement transients towards the end of the task (Figure B.5 on page 101). Figure B.6 on page 102 shows that there are, however, a few exceptions (e.g. units 11, 84 in this session) that show moderate transients also within an epoch. Nevertheless, these transients are small compared to changes between the two epochs S and P.

Thus both the population-level and single-unit level analyses are in line with the second test for stationarity that we show in Panels E and F of Figure 3.6 on page 42. There we compare the firing rate and covariance changes between two 200 ms segments of the same epoch to the firing rate and covariance changes between two 200 ms segments of different epochs. If the neural activity was not stationary within an epoch then we would not obtain correlation coefficients of almost 1 between firing rates in Panel E of Figure 3.6 on page 42 and correlation coefficients up to 0.9 between covariance patterns within one epoch in Panel F of Figure 3.6 on page 42. In summary, the analyses together make us confident that assuming stationarity within an epoch is a good approximation to show that there are significant behaviorally related changes in covariances across epochs of the reach-to-grasp experiment.




Figure B.5: **Stationarity of single-unit activity during a reach-to-grasp trial (monkey N, session i140627-001).** Black lines indicate mean and gray areas ± 1 standard deviation across trials of single unit activity in each panel (sliding window analysis with 5 ms step size and 100 ms window length). Blue bars indicate starting (S) and preparatory (P) periods used in the analysis (cf. Figure 3.6 on page 42). First, second and fourth dashed lines (marked with stars) indicate visual signals lighting up and the third dashed line indicates the removal of a visual cue and beginning of a waiting period.




Figure B.6: **Stationarity of single-unit activity during a reach-to-grasp trial (monkey N, session i140704-001).** Analogous to Figure B.5 on page 101.

session	N_{trials}	$N_{\text{single units}}$
e161212-002	108	129
e161214-001	99	118
e161222-002	102	118
e170105-002	101	116
e170106-001	100	113
i140613-001	93	137
i140617-001	129	155
i140627-001	138	145
i140702-001	157	134
i140703-001	142	142
i140704-001	141	124
i140721-002	160	96
i140725-002	151	106

Table B.2: **Numbers of trials and single units per reach-to-grasp recording session.** Session names starting with “e” correspond to monkey E and session names starting with “i” to monkey N.

B.4 NETWORK MODEL

We are considering neuronal network models with isotropic and distance-dependent connection profiles. Ultimately, we are interested in describing cortical networks with two-dimensional sheet-like structure. But, for developing the theory, we first consider the simpler case of a one-dimensional ring and subsequently develop the theory on a two-dimensional torus, ensuring periodic boundary conditions in both cases. N equidistantly distributed neurons form a grid on these manifolds. The position of neuron $i \in \{1, \dots, N\}$ is described by the vector $\mathbf{r}_i \in \mathbb{R}^D$, $D \in \{1, 2\}$. The connections W_{ij} from neuron j to neuron i are drawn randomly with a connection probability that decays with distance between neurons $|\mathbf{r}_i - \mathbf{r}_j|$, described by the normalized connectivity profile $p(r)$, $\int p(r) d^D r = 1$, which we assume to obey radial symmetry. The connection probability decays on a characteristic length scale R . As we are working on discrete lattices, we introduce the probability of two neurons being connected p_{ij} , which is defined by the relation $p(\mathbf{r}_i - \mathbf{r}_j) = \lim_{a \rightarrow 0} p_{ij}/a$, with lattice spacing a . We set the synaptic weights for connections of a single type to a fixed value w , but allow for multiple connections between neurons, that is $W_{ij} \in \{0, w, 2w, \dots\} = n_{ij} \cdot w$ for all sending neurons j of a given type, where n_{ij} is binomially distributed. Such multapses are required to simultaneously meet biological constraints on neuronal indegrees, neuron densities, and spatial ranges of connections. If instead one assumed Bernoulli connectivity, an analysis analogous to Senk et al. (2018, see Eq. 7) would yield a connection probability exceeding unity.

We introduce two populations of neurons, excitatory (E) and inhibitory (I) neurons. The number of neurons of a given population $a \in \{E, I\}$ is N_a , and their ratio is

$q = N_E/N_I$, which, for convenience, we assume to be an even number (see permutation symmetry below). The connection from population b to population a has the synaptic weight w_{ab} and characteristic decay length of the connectivity profile R_{ab} . The average number of inputs drawn per neuron is fixed to K_{ab} . In order to preserve translation symmetry, q excitatory neurons and one inhibitory neuron are put onto the same lattice point, as shown in Panel A of Figure 3.3 on page 35.

Linear-response theory has been shown to faithfully capture the statistics of fluctuations in asynchronous irregular network states (Lindner et al., 2005). Here we follow Grytskyy et al. (2013), who show that different types of neuronal network models can be mapped to an Ornstein-Uhlenbeck process and that the low-frequency limit of this simple rate model describes spike count covariances of spiking models well (Tetzlaff et al., 2012). In particular, Dahmen et al. (2019) showed quantitative agreement of linear-response predictions for the statistics of spike-count covariances in leaky integrate-and-fire networks for the full range of spectral bounds $\rho \in [0, 1)$. Therefore, we consider a network of linear rate neurons, whose activity $\mathbf{x} \in \mathbb{R}^N$ is described by

$$\tau \frac{d}{dt} \mathbf{x} = -\mathbf{x} + \mathbf{W}\mathbf{x} + \boldsymbol{\zeta} \quad ,$$

with uncorrelated Gaussian white noise $\boldsymbol{\zeta}$, $\langle \zeta_i(t) \rangle = 0$, $\langle \zeta_i(s) \zeta_j(t) \rangle = D_i \delta_{ij} \delta(s - t)$. The solution to this differential equation can be found by multiplying the whole equation with the left eigenvectors $\mathbf{u}_\alpha$ of $\mathbf{W}$

$$\tau \frac{d}{dt} y_\alpha = -y_\alpha + \lambda_\alpha y_\alpha + \tilde{\zeta}_\alpha \quad , \quad (\text{B.1})$$

where $y_\alpha = \mathbf{u}_\alpha \cdot \mathbf{x}$, $\tilde{\zeta}_\alpha = \mathbf{u}_\alpha \cdot \boldsymbol{\zeta}$ and λ_α is denoting the corresponding eigenvalue of $\mathbf{W}$. Neglecting the noise term, the solutions are given by

$$y_\alpha(t) \propto \Theta(t) \exp \left[-\frac{t}{\tau} (1 - \lambda_\alpha) \right] \quad , \quad (\text{B.2})$$

with Heaviside function $\Theta(t)$. These are the eigenmodes of the linear system and they are linear combinations of the individual neuronal rates

$$y_\alpha = \sum_{i=1}^N (u_\alpha)_i x_i \quad .$$

Note that the weights $(u_\alpha)_i$ of these linear combinations depend on the details of the effective connectivity matrix $\mathbf{W}$. The stability of an eigenmode is determined by the corresponding eigenvalue λ_α . If $\text{Re}(\lambda_\alpha) < 1$, the eigenmode is stable and decays exponentially. If $\text{Re}(\lambda_\alpha) > 1$, the eigenmode is unstable and grows exponentially. If $\text{Im}(\lambda_\alpha) \neq 0$, the eigenmode is oscillatory with an exponential envelope. $\text{Re}(\lambda_\alpha) = 1$ is here referred to as the critical point. This type of stability is also called linear stability to stress that these considerations are only valid in the linear approximation. Realistic neurons have a saturation at high rates, which prevents activity from diverging indefinitely. A network is called linearly stable if all modes are stable. This is determined by the real part of the largest eigenvalue of $\mathbf{W}$, called spectral bound ρ . In inhibition-dominated networks, the spectral bound is determined by the heterogeneity

in connections and $\rho \lesssim 1$ defines the dynamically balanced critical state (Dahmen et al., 2019).

The different noise components ξ_α excite the corresponding eigenmodes of the system and act as a driving force. A noise vector ξ that is not parallel to a single eigenvector u_α excites several eigenmodes, each with the corresponding strength ξ_α .

Note that the different eigenmodes do not interact, which is why the total activity x is given by a linear combination, or superposition, of the eigenmodes

$$x = \sum_{\alpha=1}^N y_\alpha v_\alpha \quad ,$$

where v_α denotes the α -th right eigenvector of the connectivity matrix W .

B.5 COVARIANCES

Time-lag integrated covariances $c_{ij} = \int d\tau \langle x_i(t)x_j(t+\tau) \rangle - \langle x_i(t) \rangle \langle x_j(t+\tau) \rangle$ can be computed analytically for the linear dynamics Eq. (B.1). They follow from the connectivity W and the noise strength D as (Lindner et al., 2005; Pernice et al., 2011; Trousdale et al., 2012; Grytskyy et al., 2013)

$$c = [1 - W]^{-1} D [1 - W]^{-T}, \quad (B.3)$$

with identity matrix 1 . These covariances are equivalent to covariances of spike counts in large time windows, given by the zero-frequency component of the Fourier transform of x (sometimes referred to as Wiener-Khinchin theorem (Gardiner, 1985); even though the theorem proper applies in cases where the Fourier transforms of the signals x do not exist). Spike count covariances (cf. Panel B of Figure 3.1 on page 31) can be computed from trial-resolved spiking data (Dahmen et al., 2019). This equivalence allow us to directly relate theoretical predictions for covariances to the experimentally observed ones.

While equation (B.3) provides the full information on covariances between any two neurons in the network, this information is not available in the experimental data. Only a small subset of neuronal activities can be recorded such that inference of connectivity parameters from Eq. (B.3) is unfeasible. Dahmen et al. (2019) recently proposed to instead consider the statistics of covariances as the basis for comparison between models and data. Using Eq. (B.1) and Eq. (B.3) as a starting point, field theoretical techniques allow the derivation of equations for the mean $\bar{c}$ and variance $\overline{\delta c^2}$ of cross-covariances in relation to the mean M and variance S of the connectivity matrix W (Dahmen et al., 2019):

$$\bar{c} = [1 - M]^{-1} D_r [1 - M]^{-T}, \quad (B.4)$$

$$\overline{\delta c^2} = [1 - S]^{-1} D_r^2 [1 - S]^{-T}. \quad (B.5)$$

M and S are defined in the subsequent section. The renormalized input noise strength is given by

$$D_r = D + \text{diag} \left[D (1 - S)^{-1} S \cdot I \right] \quad , \quad (B.6)$$

with input noise covariance $\mathbf{D}$, and the all-ones vector $\mathbf{I} = (1, \dots, 1)^T \in \mathbb{R}^N$. Note that Eq. (B.5) only holds for cross-covariances ($i \neq j$). The diagonal terms $[\overline{\delta c^2}]_{ii}$, i.e. the variance of auto-covariances, do get a second contribution, which is negligible for the cross-covariances considered here.

B.6 CUMULANT GENERATING FUNCTION OF CONNECTIVITY MATRIX

For calculating the mean and variance of the covariances of the network activity (Eq. (B.4) and Eq. (B.5)) we need mean $\mathbf{M}$ and variance $\mathbf{S}$ of connectivity $\mathbf{W}$. In the following, we derive the cumulant generating function (Gardiner, 1983) of W_{ij} .

The number of connections n from neuron j to neuron i is a binomial random variable with K trials with the probability of success given by p_{ij} (in the following, for brevity, we ignore the index i , $p_{ij} \equiv p_j$)

$$\text{Prob}_j(n) = \binom{K}{n} p_j^n (1 - p_j)^{K-n} .$$

The average number of connections from neuron j to neuron i is $K_j = p_j K$, which assures the correct average total indegree

$$\sum_j K_j = K \sum_j p_j = K .$$

The moment generating function of a connectivity matrix element $W_j \equiv W_{ij} \in \{0, w, 2w, \dots\}$ is given by

$$Z_{W_j}(k) = \sum_{n=0}^K \binom{K}{n} p_j^n (1 - p_j)^{K-n} e^{nwk} .$$

In a realistic network, K is very large. In the limit $K \rightarrow \infty$, while keeping $Kp = \text{const.}$, the binomial distribution converges to a Poisson distribution and we can write

$$\begin{aligned} Z_{W_j}(k) &\approx \sum_{n=0}^K \frac{K^n}{n!} e^{-K} e^{nwk} \\ &= \sum_{n=0}^K \frac{(K e^{wk})^n}{n!} e^{-K} \\ &\xrightarrow{K \rightarrow \infty} \exp \left[K (e^{wk} - 1) \right] . \end{aligned}$$

Taking the logarithm leads to the cumulant generating function

$$G_{W_j}(k) \approx p_j K (e^{wk} - 1) ,$$

and the first two cumulants

$$\begin{aligned} M_{ij} &= \left. \frac{\partial}{\partial k} G_{W_j}(k) \right|_{k=0} = p_j K w = p (|x_i - x_j|) K w , \\ S_{ij} &= \left. \frac{\partial^2}{\partial k^2} G_{W_j}(k) \right|_{k=0} = p_j K w^2 = p (|x_i - x_j|) K w^2 . \end{aligned}$$

B.7 NOTE ON DERIVATION OF VARIANCE OF COVARIANCES

Note that $\mathbf{M}$ and $\mathbf{S}$ have an identical structure determined by the connectivity profile and the structure of the covariance equation is identical for the mean Eq. (B.4) and variance Eq. (B.5) as well. This is why in the following we only derive the results for the mean of covariances. The results for the variance of covariances is obtained by substituting w by w^2 and D_r by D_r^2 . As we show, divergences in expressions related to the mean covariances arise if the population eigenvalue λ_0 of the effective connectivity matrix approaches one. In expressions related to the variance of covariances, the divergences are caused by the squared spectral bound ρ^2 being close to one. In general expressions, we sometimes write ζ in order to denote either the population eigenvalue or the spectral bound, corresponding to the context of mean or variance of covariances.

B.8 UTILIZING SYMMETRIES TO REDUCE DIMENSIONALITY

For real neuronal networks, the anatomical connectivity is never known completely, let alone the effective connectivity. This is why we are considering disorder-averaged systems. They are described by the mean $\mathbf{M}$ and variance $\mathbf{S}$ of the connectivity. The latter inherit the underlying symmetries of the network, like for example the same radially symmetric connectivity profile for all neurons of one type. As neuronal networks are high dimensional systems, calculating covariances from Eq. (B.4) and Eq. (B.5) first seems like a daunting task. But, leveraging the aforementioned symmetries similarly as in Kriener et al. (2014) allows for an effective reduction of the dimensionality of the system, thereby rendering the problem manageable.

As a demonstrative example of how this is done, consider a random network of N neurons on a one-dimensional ring, in which a neuron can form a connection with weight w to any other neuron with probability p_0 . In that case, $\mathbf{M}$ is a homogeneous matrix, with all entries given by the same average connectivity weight

$$\mathbf{M} = \begin{pmatrix} p_0 w & p_0 w & \dots & p_0 w \\ p_0 w & p_0 w & \dots & p_0 w \\ \vdots & \vdots & \ddots & \vdots \\ p_0 w & p_0 w & \dots & p_0 w \end{pmatrix}.$$

This corresponds to an all-to-all connected ring network. Due to the symmetry of the system, moving all neurons by one lattice constant does not change the system. The translation operator T , representing this operation mathematically, is defined via its effect on the vector of neuron activity $\mathbf{x}$

$$T\mathbf{x} = T \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{pmatrix} = T \begin{pmatrix} x_N \\ x_1 \\ \vdots \\ x_{N-1} \end{pmatrix}.$$

Applying T N -times yields the identity operation

$$T^N = \mathbf{1}.$$

Hence, its eigenvalues are given by complex roots of one

$$e^{-i2\pi l/N} = e^{-i2\pi la/L} = e^{-ik_l a} \quad , \quad l \in \{0, 1, \dots, N-1\} \quad ,$$

with $L = Na$ denoting the circumference of the ring. This shows that T has N one-dimensional eigenspaces. Since the system is invariant under translation, M is invariant under the transformation

$$TMT^{-1} = M \quad ,$$

and thus M and T commute. As M leaves eigenspaces of T invariant (if v is an eigenvector of T , Mv is an eigenvector with the same eigenvalue, so they need to be multiples of each other), all eigenvectors of T must be eigenvectors of M . Accordingly, knowing the eigenvectors of T allows diagonalizing M . The normalized (left and right) eigenvectors of T are given by

$$v_{k_l} = \frac{1}{\sqrt{N}} \begin{pmatrix} 1 \\ e^{ik_l a} \\ e^{i2k_l a} \\ \vdots \\ e^{i(N-1)k_l a} \end{pmatrix} \quad .$$

We get the eigenvalues of M by multiplying it with the eigenvectors of T

$$\begin{aligned} m_{k_l} &= v_{k_l}^\dagger M v_{k_l} \\ &= \frac{1}{N} p_0 w v_{k_l}^\dagger \cdot \begin{pmatrix} \sum_{j=0}^{N-1} e^{ik_l a j} \\ \sum_{j=0}^{N-1} e^{ik_l a j} \\ \vdots \\ \sum_{j=0}^{N-1} e^{ik_l a j} \end{pmatrix} \quad , \end{aligned}$$

which is always zero, except for $l = 0$, which corresponds to the population eigenvalue $\lambda_0 := m_{k_0} = N p_0 w$ of W (Panel C of Figure 3.3 on page 35). Now, we can simply write down the diagonalized form of M

$$\begin{pmatrix} \lambda_0 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 \end{pmatrix} \quad ,$$

and we effectively reduced the N -dimensional to a one dimensional problem. Inverting $A := \mathbf{1} - M$ in Eq. (B.4) is straightforward now, since it is diagonal in the new basis. Its eigenvalues can be written as

$$a_k = 1 - m_k \quad ,$$

where we suppressed the index l . Therefore its inverse is given by

$$\begin{aligned}
 A_{ij}^{-1} &= \sum_k a_k^{-1} (v_k)_i (v_k^\dagger)_j \\
 &= \frac{1}{N} \sum_k \frac{1}{1 - m_k} e^{ik(x_i - x_j)} \\
 &= \frac{1}{N} \sum_k \left(1 + \frac{m_k}{1 - m_k} \right) e^{ik(x_i - x_j)} \\
 &= \delta_{ij} + \frac{1}{N} \frac{\lambda_0}{1 - \lambda_0}.
 \end{aligned}$$

The renormalized noise can be evaluated using that the all-ones vector occurring in equation Eq. (B.6) is the eigenvector v_0 of S . After identifying the eigenvalue s_0 with the squared spectral bound ρ^2 , we find

$$D_r = \text{diag} \left(\frac{D}{1 - \rho^2} \right) ,$$

which allows us to express the mean cross-covariances $\bar{c}$ (see Eq. (B.4)) and the variance of cross-covariances $\overline{\delta c^2}$ (see Eq. (B.5)) in terms of the eigenvectors of M and S respectively

$$\begin{aligned}
 \bar{c} &= \frac{D}{1 - \rho^2} \left\{ \frac{2\lambda_0}{N(1 - \lambda_0)} + \left[\frac{\lambda_0}{N(1 - \lambda_0)} \right]^2 N \right\} , \\
 \overline{\delta c^2} &= \left(\frac{D}{1 - \rho^2} \right)^2 \left\{ \frac{2\rho^2}{N(1 - \rho^2)} + \left[\frac{\rho^2}{N(1 - \rho^2)} \right]^2 N \right\} .
 \end{aligned}$$

B.9 ONE-DIMENSIONAL NETWORK WITH ONE POPULATION

The simplest network with spatial connectivity is a one-dimensional ring of neurons with one population of neurons. Following Section B.6, the mean connectivity matrix has the form

$$M = Kw \begin{pmatrix} p_{11} & p_{12} & \dots & p_{1N} \\ p_{21} & p_{22} & \dots & p_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ p_{N1} & p_{N2} & \dots & p_{NN} \end{pmatrix} .$$

As p_{ij} only depends on the distance of two neurons, the rows in M are identical, but shifted by one index.

B.9.1 Dimensionality reduction

We follow the procedure developed in Section B.8, as the system is invariant under translation as well. Suppressing the subscripts of k , we get the eigenvalues of $\mathbf{M}$

$$\begin{aligned}
 m_k &= \frac{1}{N} K w \left(1, e^{-ika}, \dots, e^{-i(N-1)ka} \right) \begin{pmatrix} \sum_{j=0}^{N-1} p_{1(j+1)} e^{ikaj} \\ \sum_{j=0}^{N-1} p_{2(j+1)} e^{ikaj} \\ \vdots \\ \sum_{j=0}^{N-1} p_{N(j+1)} e^{ikaj} \end{pmatrix} \\
 &= K w \sum_{j=0}^{N-1} p_{(j+1)} e^{ikaj} \\
 &= K w a \sum_x p(x) e^{-ikx} \quad ,
 \end{aligned}$$

where the sum over x denotes a sum over all lattice sites. We used the translational symmetry from the first to the second line. The change of sign in the exponential from line two to three is due to the fact that we are summing over the second index of p_{ij} . Thus, the eigenvalues are effectively given by the discrete Fourier transform of the connectivity profile. Expressing $\mathbf{A}^{-1}$ using the eigenvectors $\mathbf{v}_k$ of $\mathbf{M}$ leads to

$$\begin{aligned}
 A_{ij}^{-1} &= \frac{1}{N} \sum_k \frac{1}{1 - m_k} e^{ik(x_i - x_j)} \\
 &= \frac{1}{N} \sum_k \left(1 + \frac{m_k}{1 - m_k} \right) e^{ik(x_i - x_j)} \\
 &= \delta_{ij} + \frac{1}{N} \sum_k \frac{m_k}{1 - m_k} e^{ik(x_i - x_j)} \\
 &\equiv \delta_{ij} + \mu_{ij} \quad ,
 \end{aligned} \tag{B.7}$$

where we extracted an identity for later convenience, and we defined μ_{ij} .

Next, we consider the renormalized noise, which is given by Eq. (B.6). Using that the all-ones vector $\mathbf{I}$ in the second term is the eigenvector of $\mathbf{S}$ corresponding to $k = 0$, we get

$$\mathbf{D} (\mathbf{I} - \mathbf{S})^{-1} \mathbf{S} \cdot \mathbf{v}_0 = D \frac{s_0}{1 - s_0} \quad .$$

Again, we identify s_0 with the spectral bound ρ^2 , and find

$$D_r = D + D \frac{\rho^2}{1 - \rho^2} = \frac{D}{1 - \rho^2} \quad . \tag{B.8}$$

Inserting Eq. (B.7) and Eq. (B.8) into Eq. (B.4) yields

$$\bar{c}_{ij} = \frac{D}{1 - \rho^2} \left(\delta_{ij} + 2\mu_{ij} + \sum_k \mu_{ik} \mu_{kj} \right) \quad .$$

B.9.2 Continuum limit

As we assume the lattice constant to be small, we know that the connectivity profile is sampled densely, and we are allowed to take the continuum limit. Therefore, we write

$$\begin{aligned}
 m_k &= Kw \sum_j p_{(j+1)} e^{ika_j} \\
 &= Kw \sum_j a \frac{p_{(j+1)}}{a} e^{ika_j} \\
 &\xrightarrow{a \rightarrow 0} Kw \int_{-L/2}^{L/2} dx p(-x) e^{ikx} \\
 &= Kw \int_{-L/2}^{L/2} dx p(x) e^{-ikx} .
 \end{aligned}$$

Note that $\lim_{a \rightarrow 0} \sum_j p_j/a = \lim_{a \rightarrow 0} \sum_j p(x_i - x_j)/a = \int dx p(-x)$, because we are summing over the second index j . If the decay constant R of the connectivity profile is small compared to the size of the network L , we can take L to infinity and finally end up with

$$m(k) = Kw \int dx p(x) e^{-ikx} . \quad (\text{B.9})$$

Analogously, we find

$$A^{-1}(x) = \delta(x) + \frac{1}{2\pi} \int dk \frac{m(k)}{1 - m(k)} e^{ikx} \equiv \delta(x) + \mu(x) , \quad (\text{B.10})$$

where we defined

$$\mu(x) = \frac{1}{2\pi} \int dk \mu(k) e^{ikx} , \quad (\text{B.11})$$

with

$$\mu(k) = \frac{m(k)}{1 - m(k)} . \quad (\text{B.12})$$

Finally, we get

$$\bar{c}(x) = \frac{D}{1 - \rho^2} [\delta(x) + 2\mu(x) + (\mu * \mu)(x)] , \quad (\text{B.13})$$

where the asterisk denotes the convolution.

B.9.3 Prediction of exponential decay of covariance statistics

Note that the integral in equation Eq. (B.11) can be interpreted as an integral in the complex plane. According to the residue theorem, the solution to this integral is a weighted sum of exponentials, evaluated at the poles of $[1 - m(k)]^{-1}$. As $\mu(x)$ appears in the equation for the mean covariances, and the convolution of two exponentials is an exponential with the prefactor $(\text{const.} + |x|)$, we expect the dominant behavior to be an exponential decay in the long-range limit, with decay constants given by the inverse imaginary part of the poles. The poles which are closest to zero are the ones which lead to the most shallow and thereby dominant decay. A real part of the poles leads to oscillations in $\mu(x)$.

B.9.4 Long-range limit

We cannot expect to solve the integral in Eq. (B.10) for arbitrary connectivity profiles. To continue our analysis, we make use of the Padé method, which approximates arbitrary functions as rational functions (Basdevant, 1972). We approximate $\mu(k)$ around $k = 0$ using a Padé approximation of order (0,2)

$$\mu(k) \approx \frac{m(0)}{1 - m(0) - \frac{m''(0)}{2m(0)}k^2} \quad ,$$

with

$$\begin{aligned} m(0) &= Kw \int dx p(x) = Kw = \lambda_0 \quad , \\ m''(0) &= -Kw \int dx x^2 p(x) = -Kw \langle x^2 \rangle \quad . \end{aligned} \quad (\text{B.14})$$

This allows us to calculate the approximate poles of $\mu(k)$

$$k_0 = \pm \sqrt{\frac{2m(0)}{m''(0)} [1 - m(0)]} \quad . \quad (\text{B.15})$$

As $2m(0)/m''(0)$ will be negative, due to factor i^2 from the second derivative of the Fourier integral, we write

$$k_0 = \pm i \sqrt{-\frac{2m(0)}{m''(0)} [1 - m(0)]} \quad .$$

Closing the integral contour in Eq. (B.11) in the upper half plane for $x > 0$, and in the lower half plane for $x < 0$, we get

$$\begin{aligned} \mu(x) &= -\frac{m(0)^2}{m''(0)} \sqrt{-\frac{m''(0)}{2m(0)} \frac{1}{1 - m(0)}} \exp \left(-\frac{|x|}{\sqrt{-\frac{m''(0)}{2m(0)} \frac{1}{1 - m(0)}}} \right) \\ &\equiv -\frac{m(0)^2}{m''(0)} \bar{R} \exp \left(-\frac{|x|}{\bar{R}} \right) \quad , \end{aligned}$$

where we defined the effective decay constant for the mean covariances

$$\bar{R} = \sqrt{-\frac{m''(0)}{2m(0)} \frac{1}{1 - m(0)}} = \sqrt{\frac{\langle x^2 \rangle}{2} \frac{1}{1 - \lambda_0}} \quad ,$$

with $m(0) = \lambda_0$ and $m''(0) = \lambda_0 \langle x^2 \rangle$, since $m(k)$ is the Fourier transform of the connectivity profile Eq. (B.9). Note that $\lambda_0 = Kw$ again is the population eigenvalue of the effective connectivity matrix $\mathbf{W}$. For evaluating Eq. (B.4) and Eq. (B.5), we need to calculate the convolution of μ with itself

$$(\mu * \mu)(x) = \int dy \mu(x - y) \mu(y) = \frac{m(0)^4}{m''(0)^2} \bar{R}^2 (\bar{R} + |x|) \exp \left(-\frac{|x|}{\bar{R}} \right) \quad .$$

The final expression for the mean covariances is

$$\bar{c}(x) = \frac{D}{1-\rho^2} \left\{ \delta(x) + \left[\left(\frac{m(0)^4}{m''(0)^2} \bar{R}^2 - 2 \frac{m(0)^2}{m''(0)} \right) \bar{R} + \frac{m(0)^4}{m''(0)^2} \bar{R}^2 |x| \right] \exp \left(-\frac{|x|}{\bar{R}} \right) \right\} .$$

Equivalently, for the variance of covariances we obtain the final result

$$\overline{\delta c^2}(x) = \left(\frac{D}{1-\rho^2} \right)^2 \left\{ \delta(x) + \left[\left(\frac{s(0)^4}{s''(0)^2} R_{\text{eff}}^2 - 2 \frac{s(0)^2}{s''(0)} \right) R_{\text{eff}} + \frac{s(0)^4}{s''(0)^2} R_{\text{eff}}^2 |x| \right] \exp \left(-\frac{|x|}{R_{\text{eff}}} \right) \right\} ,$$

where

$$s(k) = Kw^2 \int dx p(x) e^{-ikx} .$$

Note that the quality of the Padé approximation depends on the outlier eigenvalue and the spectral bound. For the variances, the approximation works best for spectral bounds ρ close to 1. The reason for this is that we are approximating the position of the poles in the complex integral Eq. (B.11). We make an approximation around $k = 0$ and Eq. (B.15) shows that the position of the complex poles moves closer to $k = 0$ as $s(0) \equiv \rho^2 \rightarrow 1$.

GENERAL RESULTS: Using Eq. (B.14)

$$m(0) = Kw = \lambda_0 , \quad m''(0) = -Kw \langle x^2 \rangle ,$$

we find

$$\bar{c}(x) = \frac{D}{1-\rho^2} \left\{ \delta(x) + \left[\frac{Kw(1-3Kw)}{2 \langle x^2 \rangle (1-Kw)} \bar{R} + \frac{(Kw)^2}{\langle x^2 \rangle^2} \bar{R}^2 |x| \right] \exp \left(-\frac{|x|}{\bar{R}} \right) \right\} ,$$

with

$$\bar{R} = \sqrt{\left| \frac{\langle x^2 \rangle}{2} \frac{1}{1-\lambda_0} \right|} .$$

For the variance we use

$$s(0) = Kw^2 = \rho^2 , \quad s''(0) = -Kw^2 \langle x^2 \rangle ,$$

to get

$$\overline{\delta c^2}(x) = \frac{D^2}{(1-\rho^2)^2} \left\{ \delta(x) + \left[\frac{Kw^2(1-3Kw^2)}{2 \langle x^2 \rangle (1-Kw^2)} R_{\text{eff}} + \frac{(Kw^2)^2}{\langle x^2 \rangle^2} R_{\text{eff}}^2 |x| \right] \exp \left(-\frac{|x|}{R_{\text{eff}}} \right) \right\} ,$$

with

$$R_{\text{eff}} = \sqrt{\left| \frac{\langle x^2 \rangle}{2} \frac{1}{1-\rho^2} \right|} .$$

EXPONENTIAL CONNECTIVITY PROFILE: Using an exponential connectivity profile given by

$$p(x) = \frac{1}{2R} e^{-|x|/R} ,$$

we find $\langle x^2 \rangle = 2R^2$ and

$$\bar{R} = \sqrt{\left| \frac{1}{1 - \lambda_0} \right|} R , \quad R_{\text{eff}} = \sqrt{\left| \frac{1}{1 - \rho^2} \right|} R ,$$

with $\lambda_0 = Kw$ for the mean, and $\rho^2 = Kw^2$ for the variance.

GAUSSIAN CONNECTIVITY PROFILE: Analogously, using a Gaussian connectivity profile given by

$$p(x) = \frac{1}{\sqrt{2\pi R^2}} e^{-x^2/(2R^2)} ,$$

we find $\langle x^2 \rangle = R^2$, and get

$$\bar{R} = \sqrt{\left| \frac{1}{2} \frac{1}{1 - \lambda_0} \right|} R , \quad R_{\text{eff}} = \sqrt{\left| \frac{1}{2} \frac{1}{1 - \rho^2} \right|} R . \quad (\text{B.16})$$

B.10 ONE-DIMENSIONAL NETWORK WITH TWO POPULATIONS

Realistic neuronal network consist of excitatory and inhibitory neurons. So we need to introduce a second population to our network. Typically, there are more excitatory than inhibitory neurons in the brain. Therefore, we introduce q excitatory neurons for each inhibitory neuron. We place q excitatory neurons and one inhibitory neuron together in one cell. The cells are distributed equally along the ring. For convenience, we define $N \equiv N_I$.

The structure of the connectivity matrix depends on the choice of the activity vector $\mathbf{x}$. For later convenience we choose

$$\mathbf{x} = \begin{pmatrix} \mathbf{x}_1^{(\text{E})} \\ \mathbf{x}_1^{(\text{I})} \\ \mathbf{x}_2^{(\text{E})} \\ \mathbf{x}_2^{(\text{I})} \\ \vdots \\ \mathbf{x}_N^{(\text{E})} \\ \mathbf{x}_N^{(\text{I})} \end{pmatrix} ,$$

where $\mathbf{x}_i^{(E)}$ is a q -dimensional vector denoting the activity of the q excitatory neurons in cell i . $\mathbf{M}$ is a $(q+1)N \times (q+1)N$ -matrix, which qualitatively has the structure

$$\mathbf{M} = \begin{pmatrix} \text{EE}_{11} & \text{EI}_{11} & \text{EE}_{12} & \text{EI}_{12} & \cdots & \text{EE}_{1N} & \text{EI}_{1N} \\ \text{IE}_{11} & \text{II}_{11} & \text{IE}_{12} & \text{II}_{12} & \cdots & \text{IE}_{1N} & \text{II}_{1N} \\ \text{EE}_{21} & \text{EI}_{21} & \text{EE}_{22} & \text{EI}_{22} & \cdots & \text{EE}_{2N} & \text{EI}_{2N} \\ \text{IE}_{21} & \text{II}_{21} & \text{IE}_{22} & \text{II}_{22} & \cdots & \text{IE}_{2N} & \text{II}_{2N} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \text{EE}_{N1} & \text{EI}_{N1} & \text{EE}_{N2} & \text{EI}_{N2} & \cdots & \text{EE}_{NN} & \text{EI}_{NN} \\ \text{IE}_{N1} & \text{II}_{N1} & \text{IE}_{N2} & \text{II}_{N2} & \cdots & \text{IE}_{NN} & \text{II}_{NN} \end{pmatrix}. \quad (\text{B.17})$$

Note that EE_{ij} are $q \times q$ matrices, EI_{ij} are $q \times 1$ matrices, IE_{ij} are $1 \times q$ matrices and II_{ij} are 1×1 matrices. The entries ab_{ij} describe the connectivities from population b in cell j to population a in cell i . The entries are given by

$$ab_{ij} = \begin{cases} \frac{1}{q} w_{ab} K_{ab} (p_{ab})_{ij} & \text{if } b = E \\ w_{ab} K_{ab} (p_{ab})_{ij} & \text{if } b = I \end{cases}.$$

The difference stems from the fact that we have q times as many excitatory neurons. As the total number of indegrees from excitatory neurons should be given by K_{aE} , we need to introduce a reducing factor of $1/q$, as the connection probability is normalized to one.

B.10.1 Dimensionality reduction

In the following, we will reduce the dimensionality of $\mathbf{M}$ as done before in the case with one population. First, we make use of the symmetry within the cells. All entries in $\mathbf{M}$ corresponding to connections coming from excitatory neurons of the same cell need to be the same. For that reason, we change the basis to

$$e_i^{(E)} = \frac{1}{\sqrt{q}} \begin{pmatrix} 0 \\ 0 \\ \vdots \\ \mathbf{I} \\ \vdots \\ 0 \end{pmatrix}, \quad e_i^{(I)} = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{pmatrix}, \quad (\text{B.18})$$

where $\mathbf{I}$ denotes a q -dimensional vector containing only ones. For a full basis, we need to include all the vectors with $\mathbf{I}$ being replaced by a vector containing all possible permutations of equal numbers of ± 1 . In this basis $\mathbf{M}$ is block diagonal

$$\begin{pmatrix} \mathbf{M}' & 0 \\ 0 & 0 \end{pmatrix},$$

and M' is an $2N \times 2N$ matrix, which has the same qualitative structure as shown in Eq. (B.17), but the submatrices $(ab)_{ij}$ are replaced by

$$ab_{ij} = \begin{cases} w_{EE} K_{EE} (p_{EE})_{ij} & \text{if } ab = EE \\ \sqrt{q} w_{EI} K_{EI} (p_{EI})_{ij} & \text{if } ab = EI \\ w_{IE} K_{IE} (p_{IE})_{ij} / \sqrt{q} & \text{if } ab = IE \\ w_{II} K_{II} (p_{II})_{ij} & \text{if } ab = II \end{cases}.$$

Next, we use translational symmetry of the cells. The translation operator is defined by

$$T\mathbf{x} = T \begin{pmatrix} \mathbf{x}_1^{(E)} \\ \mathbf{x}_1^{(I)} \\ \mathbf{x}_2^{(E)} \\ \mathbf{x}_2^{(I)} \\ \vdots \\ \mathbf{x}_N^{(E)} \\ \mathbf{x}_N^{(I)} \end{pmatrix} = \begin{pmatrix} \mathbf{x}_N^{(E)} \\ \mathbf{x}_N^{(I)} \\ \mathbf{x}_1^{(E)} \\ \mathbf{x}_1^{(I)} \\ \vdots \\ \mathbf{x}_{N-1}^{(E)} \\ \mathbf{x}_{N-1}^{(I)} \end{pmatrix}.$$

As the system is invariant under moving each cell to the next lattice site, M' is invariant under the transformation

$$TM'T^{-1} = M'.$$

Again, the eigenvalues of T can be determined using $T^N = 1$ and they are the same as in the case of one population. But, note that here the eigenspaces corresponding to the single eigenvalues are two dimensional. The eigenvectors

$$\mathbf{v}_k^{(E)} = \frac{1}{\sqrt{N}} \begin{pmatrix} 1 \\ 0 \\ e^{ika} \\ 0 \\ \vdots \\ e^{i(N-1)ka} \\ 0 \end{pmatrix}, \quad \mathbf{v}_k^{(I)} = \frac{1}{\sqrt{N}} \begin{pmatrix} 0 \\ 1 \\ 0 \\ e^{ika} \\ \vdots \\ 0 \\ e^{i(N-1)ka} \end{pmatrix},$$

belong to the same eigenvalue. In this basis, M' is block diagonal, with each block consisting of a 2×2 matrix, corresponding to one value of $k_l = \frac{2\pi l}{L}$, $l \in \{0, \dots, N-1\}$

$$M' = \begin{pmatrix} \mathbf{M}_{k_0} & 0 & \cdots & 0 \\ 0 & \mathbf{M}_{k_1} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{M}_{k_{N-1}} \end{pmatrix}.$$

Since all block matrices can be treated equally, we further reduced the problem to diagonalizing a 2×2 matrix. The submatrices take the form

$$\mathbf{M}_k = \begin{pmatrix} m_{\text{EE}}(k) & \sqrt{q}m_{\text{EI}}(k) \\ m_{\text{IE}}(k)/\sqrt{q} & m_{\text{II}}(k) \end{pmatrix} ,$$

with the discrete Fourier transform

$$m_{ab}(k) = K_{ab}w_{ab} \sum_{x=-Na/2}^{Na/2} p_{ab}(x)e^{-ikx} . \quad (\text{B.19})$$

Note that x and k are still discrete here, but we could take the continuum limit at this point. The eigenvalues of $\mathbf{M}_k$ are given by

$$\begin{aligned} m_{\pm}(k) &= \frac{1}{2}(m_{\text{EE}}(k) + m_{\text{II}}(k)) \\ &\pm \frac{1}{2}\sqrt{m_{\text{EE}}(k)^2 + m_{\text{II}}(k)^2 - 2m_{\text{EE}}(k)m_{\text{II}}(k) + 4m_{\text{EI}}(k)m_{\text{IE}}(k)} . \end{aligned} \quad (\text{B.20})$$

The corresponding eigenvectors are

$$\mathbf{v}_{1,2}(k) = \mathcal{N}_{\pm} \begin{pmatrix} \sqrt{q}m_{\text{EI}}(k) \\ m_{\pm}(k) - m_{\text{EE}}(k) \end{pmatrix} , \quad (\text{B.21})$$

with normalization $\mathcal{N}_{\pm}$. The eigenvectors written in the Fourier basis are given by

$$\mathbf{v}_{\pm}(k) = \mathcal{N}_{\pm} \left[\sqrt{q}m_{\text{EI}}(k)\mathbf{v}_k^{(\text{E})} + (m_{\pm}(k) - m_{\text{EE}}(k))\mathbf{v}_k^{(\text{I})} \right] , \quad (\text{B.22})$$

and we can get the eigenvectors $\tilde{\mathbf{v}}_{\pm}(k)$ in the basis we started with by extending $\mathbf{v}_k^{(\text{E})}$ and $\mathbf{v}_k^{(\text{I})}$ to vectors similar to Eq. (B.18), where the elements corresponding to excitatory neurons are repeated q -times. Note that the normalization of the original basis leads to an additional factor $1/\sqrt{q}$ in the first term of Eq. (B.22).

Analogously, we can find the left eigenvectors of $\mathbf{M}$ by conducting the same steps with the transpose of $\mathbf{M}$

$$\mathbf{u}_{\pm}(k) = \mathcal{N}_{\pm} \left[m_{\text{IE}}(k)\mathbf{v}_k^{(\text{E})\dagger} / \sqrt{q} + (m_{\pm}(k) - m_{\text{EE}}(k))\mathbf{v}_k^{(\text{I})\dagger} \right] , \quad (\text{B.23})$$

and the vectors in the original basis $\tilde{\mathbf{u}}_{\pm}(k)$ are obtained similarly to the right eigenvectors. The normalization $\mathcal{N}_{\pm}$ is chosen such that

$$\begin{aligned} \tilde{\mathbf{u}}_+(k) \cdot \tilde{\mathbf{v}}_+(k) &= 1 , \\ \tilde{\mathbf{u}}_+(k) \cdot \tilde{\mathbf{v}}_-(k) &= 0 , \\ \tilde{\mathbf{u}}_-(k) \cdot \tilde{\mathbf{v}}_+(k) &= 0 , \\ \tilde{\mathbf{u}}_-(k) \cdot \tilde{\mathbf{v}}_-(k) &= 1 , \end{aligned}$$

which leads to

$$\mathcal{N}_{\pm} = \sqrt{m_{\text{EI}}(k)m_{\text{IE}}(k) - (m_{\pm}(k) - m_{\text{EE}}(k))^2} .$$

Now, we can express A^{-1} in terms of the eigenvalues and eigenvectors of M

$$A^{-1} = 1 + \sum_k \left(\frac{m_+(k)}{1 - m_+(k)} \tilde{v}_+(k) \cdot \tilde{u}_+(k) + \frac{m_-(k)}{1 - m_-(k)} \tilde{v}_-(k) \cdot \tilde{u}_-(k) \right) , \quad (\text{B.24})$$

which leads to

$$A_{ij}^{-1} = \delta_{ij} + \frac{1}{N} \sum_k \mu_{ij}(k) e^{ik|x_i - x_j|} , \quad (\text{B.25})$$

where we defined $\mu(k)$ similar to Eq. (B.12). Let E and I be the sets of indices referring to excitatory or inhibitory neurons respectively. We find

$$\mu_{ij}(k) \equiv \begin{cases} \mu_{EE}(k) & \text{for } i, j \in E \\ \mu_{EI}(k) & \text{for } i \in E, j \in I \\ \mu_{IE}(k) & \text{for } i \in I, j \in E \\ \mu_{II}(k) & \text{for } i, j \in I \end{cases} ,$$

with

$$\begin{aligned} \mu_{EE}(k) &= \frac{1}{q} \frac{m_{EE}(k) + m_{IE}(k)m_{EI}(k) - m_{EE}(k)m_{II}(k)}{1 - \zeta(k)} , \\ \mu_{EI}(k) &= \frac{m_{EI}(k)}{1 - \zeta(k)} , \\ \mu_{IE}(k) &= \frac{1}{q} \frac{m_{IE}(k)}{1 - \zeta(k)} , \\ \mu_{II}(k) &= \frac{m_{II}(k) + m_{IE}(k)m_{EI}(k) - m_{EE}(k)m_{II}(k)}{1 - \zeta(k)} , \end{aligned} \quad (\text{B.26})$$

and

$$\zeta(k) = m_{EE}(k) + m_{II}(k) + m_{EI}(k)m_{IE}(k) - m_{EE}(k)m_{II}(k) .$$

B.10.2 General results

The renormalized noise is evaluated using the same trick as in the one population case. We express the all-ones vector using eigenvectors of the variance matrix S

$$I = a\tilde{v}_+(0) + b\tilde{v}_-(0) .$$

Evaluating the coefficients a and b and inserting the corresponding solutions into Eq. (B.6) yields

$$D_r = \text{diag} \left(\underbrace{\underbrace{D_r^{(E)}, \dots, D_r^{(E)}}_{q\text{-times}}, D_r^{(I)}, \underbrace{D_r^{(E)}, \dots, D_r^{(E)}}_{q\text{-times}}, D_r^{(I)}, \dots, \underbrace{D_r^{(E)}, \dots, D_r^{(E)}}_{q\text{-times}}, D_r^{(I)}}_{N(q+1)\text{-entries}} \right) , \quad (\text{B.27})$$

with

$$\begin{aligned} D_r^{(E)} &= D \left[1 + \frac{s_{EE}(0) + s_{EI}(0) + s_{EI}(0)s_{IE}(0) - s_{EE}(0)s_{II}(0)}{1 - \rho^2} \right] , \\ D_r^{(I)} &= D \left[1 + \frac{s_{IE}(0) + s_{II}(0) + s_{EI}(0)s_{IE}(0) - s_{EE}(0)s_{II}(0)}{1 - \rho^2} \right] , \end{aligned}$$

with the eigenvalues $s_{ab}(k)$ of $\mathbf{S}$. We again identified the spectral bound

$$\rho^2 = s_{EE}(0) + s_{II}(0) + s_{EI}(0)s_{IE}(0) - s_{EE}(0)s_{II}(0) . \quad (\text{B.28})$$

The mean covariances can be written as

$$\bar{\mathbf{c}} = \mathbf{D}_r + \boldsymbol{\mu} \mathbf{D}_r + \mathbf{D}_r \boldsymbol{\mu}^T + \boldsymbol{\mu} \mathbf{D}_r \boldsymbol{\mu}^T ,$$

where $\boldsymbol{\mu} = \boldsymbol{\mu}(x)$. We can distinguish three different kinds of covariances depending on the type of neurons involved

$$\bar{c}_{ij} \equiv \begin{cases} \bar{c}_{EE}(x) & \text{for } i, j \in E \\ \bar{c}_{EI}(x) & \text{for } i \in E, j \in I \text{ or } i \in E, j \in I \\ \bar{c}_{II}(x) & \text{for } i, j \in I \end{cases} .$$

with

$$\begin{aligned} \bar{c}_{EE}(x) &= D_r^{(E)} \delta(x) + 2D_r^{(E)} \mu_{EE}(x) + D_r^{(E)} q (\mu_{EE} * \mu_{EE})(x) + D_r^{(I)} (\mu_{EI} * \mu_{EI})(x) , \\ \bar{c}_{EI}(x) &= D_r^{(E)} \mu_{IE}(x) + D_r^{(I)} \mu_{EI}(x) + D_r^{(E)} q (\mu_{EE} * \mu_{IE})(x) + D_r^{(I)} (\mu_{II} * \mu_{EI})(x) , \\ \bar{c}_{II}(x) &= D_r^{(I)} \delta(x) + 2D_r^{(I)} \mu_{II}(x) + D_r^{(E)} q (\mu_{IE} * \mu_{IE})(x) + D_r^{(I)} (\mu_{II} * \mu_{II})(x) . \end{aligned}$$

B.10.3 Long-range limit

From here on, we consider the special case in which the synaptic connections only depend on the type of the presynaptic neuron and not on the type of the postsynaptic neuron. This is in agreement with network parameters used in established cortical network models (Potjans and Diesmann, 2014; Senk et al., 2018), in which the connection probabilities to both types of target neurons in the same layer are usually of the same order of magnitude. In that case, all expressions become independent of the first population index $A_{ab} \equiv A_b$, and the only expressions we need to evaluate become

$$\mu_a(k) = \gamma_a \frac{m_a(k)}{1 - \zeta(k)} ,$$

with

$$\zeta(k) = m_E(k) + m_I(k) ,$$

and

$$\gamma_a = \begin{cases} 1 & \text{if } a = I \\ 1/q & \text{if } a = E \end{cases} . \quad (\text{B.29})$$

After taking the continuum limit, we can make a (0,2)-Padé approximation again

$$\mu_a(k) \approx \frac{\gamma_a m_a(0)}{1 - \zeta(0) - \left[\frac{\zeta''(0)}{2} + (1 - \zeta(0)) \frac{m_a''(0)}{2m_a(0)} \right] k^2} ,$$

which leads to the poles

$$k_0 = \pm \sqrt{\left[\frac{\zeta''(0)}{2\zeta(0)} \frac{\zeta(0)}{1 - \zeta(0)} + \frac{m_a''(0)}{2m_a(0)} \right]^{-1}} ,$$

or the effective decay constant of the mean covariances

$$\bar{R}_a = \text{Im}(k_0)^{-1} = \sqrt{-\frac{\zeta''(0)}{2\zeta(0)} \frac{\zeta(0)}{1 - \zeta(0)} - \frac{m_a''(0)}{2m_a(0)}} .$$

Using

$$\begin{aligned} \zeta &\equiv \zeta(0) = w_E K_E + w_I K_I , \\ \zeta'' &\equiv \zeta''(0) = -w_E K_E \langle x^2 \rangle_E - w_I K_I \langle x^2 \rangle_I , \\ m_a &\equiv m_a(0) = w_a K_a , \\ m_a'' &\equiv m_a''(0) = -w_a K_a \langle x^2 \rangle_a , \end{aligned}$$

we get

$$\begin{aligned} \bar{R}_a &= \sqrt{\frac{w_E K_E \langle x^2 \rangle_E + w_I K_I \langle x^2 \rangle_I}{w_E K_E + w_I K_I} \frac{\zeta}{1 - \zeta} + \frac{\langle x^2 \rangle_a}{2}} \\ &= \sqrt{\frac{(\omega \kappa \tilde{\eta}^2 + 1)}{\omega \kappa + 1} \frac{\zeta}{1 - \zeta} \frac{\langle x^2 \rangle_I}{2} + \frac{\langle x^2 \rangle_a}{2}} , \end{aligned}$$

after introducing relative parameters

$$\omega = \frac{w_E}{w_I}, \quad \kappa = \frac{K_E}{K_I}, \quad \tilde{\eta}^2 = \frac{\langle x^2 \rangle_E}{\langle x^2 \rangle_I}, \quad \eta = \frac{\lambda_E}{\lambda_I} .$$

The renormalized noise Eq. (B.6) reduces to

$$D_r = \frac{D}{1 - \rho^2} . \quad (\text{B.30})$$

The mean covariances are

$$\begin{aligned} \overline{c_{EE}}(x) &= D_r [\delta(x) + 2\mu_E(x) + q (\mu_E * \mu_E)(x) + (\mu_I * \mu_I)(x)] , \\ \overline{c_{EI}}(x) &= D_r [\mu_E(x) + \mu_I(x) + q (\mu_E * \mu_E)(x) + (\mu_I * \mu_I)(x)] , \\ \overline{c_{II}}(x) &= D_r [\delta(x) + 2\mu_I(x) + q (\mu_E * \mu_E)(x) + (\mu_I * \mu_I)(x)] , \end{aligned}$$

with

$$\mu_a(x) = \gamma_a \frac{m_a(0)}{2(1 - \zeta)\bar{R}_a} \exp\left(-\frac{|x|}{\bar{R}_a}\right) ,$$

and

$$(\mu_a * \mu_a)(x) = \left(\gamma_a \frac{m(0)}{2(1-\zeta)\bar{R}_a} \right)^2 (\bar{R}_a + |x|) \exp\left(-\frac{|x|}{\bar{R}_a}\right) .$$

Note that expressions coming from both populations contribute to each kind of covariance. Therefore, all mean covariances contain a part that decays with either of the decay constants we just determined. If, for example, the inhibitory decay constant is much larger than the excitatory one, $\bar{c}_{EI}(x)$ will decay with the largest decay constant in the long-range limit

EXPONENTIAL CONNECTIVITY PROFILE: Just as in Section B.9.4 we get

$$\bar{R}_a = \sqrt{\frac{(\omega\kappa\eta^2 + 1)}{\omega\kappa + 1} \frac{\lambda_0}{1 - \lambda_0} R_I^2 + R_a^2} , R_{\text{eff},a} = \sqrt{\frac{(\omega^2\kappa\eta^2 + 1)}{\omega^2\kappa + 1} \frac{\rho^2}{1 - \rho^2} R_I^2 + R_a^2} ,$$

with $\lambda_0 = w_E K_E + w_I K_I$ for the decay constant of the mean covariances, and $\rho^2 = w_E^2 K_E + w_I^2 K_I$ for the decay constant of the variances.

GAUSSIAN CONNECTIVITY PROFILE: And similar to Section B.9.4 we get

$$\bar{R}_a = \sqrt{\frac{(\omega\kappa\eta^2 + 1)}{\omega\kappa + 1} \frac{\lambda_0}{1 - \lambda_0} \frac{R_I^2}{2} + \frac{R_a^2}{2}} , R_{\text{eff},a} = \sqrt{\frac{(\omega\kappa\eta^2 + 1)}{\omega\kappa + 1} \frac{\rho^2}{1 - \rho^2} \frac{R_I^2}{2} + \frac{R_a^2}{2}} .$$

B.11 TWO-DIMENSIONAL NETWORK WITH ONE POPULATION

In the following, we are considering two-dimensional networks, which are supposed to mimic a single-layered cortical network. Neurons are positioned on a two-dimensional lattice ($N_x \times N_y$ grid) with periodic boundary conditions in both dimensions (a torus). We define the activity vector to be of the form

$$\mathbf{x} = \begin{pmatrix} x_{1,1} \\ x_{1,2} \\ \vdots \\ x_{1,N_y} \\ x_{2,1} \\ \vdots \\ x_{2,N_y} \\ \vdots \\ x_{N_x,1} \\ \vdots \\ x_{N_x,N_y} \end{pmatrix} .$$

The connectivity matrix is defined correspondingly.

B.11.1 Dimensionality reduction

In two dimensions we have to define two translation operators that move all neurons either one step in the x -direction, or the y -direction, respectively. They are defined via their action on $\mathbf{x}$

$$T_x \mathbf{x} = \begin{pmatrix} x_{N_x,1} \\ x_{N_x,2} \\ \vdots \\ x_{N_x,N_y} \\ x_{1,1} \\ \vdots \\ x_{1,N_y} \\ \vdots \\ x_{N_x-1,1} \\ \vdots \\ x_{N_x-1,N_y} \end{pmatrix}, \quad T_y \mathbf{x} = \begin{pmatrix} x_{1,N_y} \\ x_{1,1} \\ \vdots \\ x_{1,N_y-1} \\ x_{2,N_y} \\ \vdots \\ x_{2,N_y-1} \\ \vdots \\ x_{N_x,N_y} \\ \vdots \\ x_{N_x,N_y-1} \end{pmatrix}. \quad (\text{B.31})$$

Similar reasoning as in one dimension leads to the eigenvalues

$$e^{-ik_l^{(x)}a}, \quad k_l^{(x)} = \frac{2\pi}{L_x}l, \quad l \in \{0, 1, \dots, N_x - 1\},$$

and similar for the y -direction. The eigenvectors can be inferred from the recursion relations

$$\begin{aligned} T_x \mathbf{v} = e^{-ik_l^{(x)}a} \mathbf{v} &\Rightarrow v_{(\alpha+1)\beta} = e^{ik_l^{(x)}a} v_{\alpha\beta}, \\ T_y \mathbf{v} = e^{-ik_l^{(y)}a} \mathbf{v} &\Rightarrow v_{\alpha(\beta+1)} = e^{ik_l^{(y)}a} v_{\alpha\beta}, \end{aligned}$$

where entries $v_{\alpha\beta}$ of the vector $\mathbf{v}$ are defined analogously to Eq. (B.31). The eigenvectors are given by

$$\mathbf{v}_k = \frac{1}{\sqrt{N_x N_y}} \begin{pmatrix} \mathbf{v}^{(x)} \\ e^{ik^{(y)}a} \mathbf{v}^{(x)} \\ \vdots \\ e^{-ik^{(y)}a} \mathbf{v}^{(x)} \end{pmatrix}, \quad \mathbf{v}^{(x)} = \begin{pmatrix} 1 \\ e^{ik^{(x)}a} \\ \vdots \\ e^{i\frac{N_x-1}{2}k^{(x)}a} \\ e^{-i\frac{N_x-1}{2}k^{(x)}a} \\ \vdots \\ e^{-ik^{(x)}a} \end{pmatrix},$$

where we suppressed the subscripts of $k^{(x)}$ and $k^{(y)}$ again. Using that these eigenvectors are eigenvectors of $\mathbf{M}$ as well, yields the eigenvalues of $\mathbf{M}$

$$m_k = \mathbf{v}_k^\dagger \mathbf{M} \mathbf{v}_k = Kw \sum_x \sum_y p(|\mathbf{x}|) e^{-ik \cdot \mathbf{x}}.$$

In the continuum limit, this becomes the two-dimensional Fourier transform

$$m(\mathbf{k}) = K w \int d^2x p(\mathbf{x}) e^{-i\mathbf{k} \cdot \mathbf{x}} \quad . \quad (\text{B.32})$$

The inverse of A is given by

$$A^{-1}(\mathbf{x}) = \delta(\mathbf{x}) + \mu(\mathbf{x}) \quad , \quad (\text{B.33})$$

with the inverse two-dimensional Fourier transform

$$\mu(\mathbf{x}) = \frac{1}{(2\pi)^2} \int d^2k \frac{m(\mathbf{k})}{1 - m(\mathbf{k})} e^{i\mathbf{k} \cdot \mathbf{x}} \quad . \quad (\text{B.34})$$

The expression for the renormalized noise is the same as in the one-dimensional case with one population. Hence, the mean covariances are given by

$$\bar{c}(\mathbf{x}) = \frac{D}{1 - \rho^2} [\delta(\mathbf{x}) + 2\mu(\mathbf{x}) + (\mu * * \mu)(\mathbf{x})] \quad , \quad (\text{B.35})$$

which is the one-dimensional expression, except for the convolution, which is replaced by its two-dimensional analogon denoted here by the double asterisk.

B.11.2 Long-range limit

Employing the symmetry of the connectivity kernel, we rewrite the integral in $\mu(\mathbf{x})$ using polar coordinates

$$\mu(\mathbf{x}) = \frac{1}{(2\pi)^2} \int_0^\infty dk \int_0^{2\pi} d\varphi k \frac{m(k)}{1 - m(k)} e^{ikr \cos(\varphi)} \quad , \quad (\text{B.36})$$

with $r = |\mathbf{x}|$, and make a Padé approximation of order (0,2) of the integration kernel

$$\mu(\mathbf{x}) = \frac{1}{(2\pi)^2} \int_0^\infty dk \int_0^{2\pi} d\varphi k \frac{m(0)}{1 - m(0) - \frac{m''(0)}{2m(0)} k^2} e^{ikr \cos(\varphi)} \quad . \quad (\text{B.37})$$

Following Goldenfeld (1992, p.16of), we can interpret this as calculating the Green's function of the heat equation

$$\left[1 - m(0) + \frac{m''(0)}{2m(0)} \nabla^2 \right] \mu(\mathbf{x}) = m(0) \delta(r) \quad , \quad (\text{B.38})$$

which can be solved, using the fact that $\mu(\mathbf{x})$ can only be a function of the radial distance r , due to the given symmetry of the kernel. Rewriting leads to

$$\left[-\frac{1}{r} \frac{d}{dr} \left(r \frac{d}{dr} \right) + \bar{R}^{-2} \right] \mu(r) = \Gamma \delta(r) \quad ,$$

with the effective decay constant

$$\bar{R} = \sqrt{-\frac{m''(0)}{2m(0)} \frac{1}{1 - m(0)}} \quad , \quad (\text{B.39})$$

and $\Gamma = -2m(0)^2/m''(0)$. Defining $\tilde{r} \equiv r/\bar{R}$, $\tilde{\mu}(\tilde{r}) \equiv \mu(\tilde{r}/\bar{R})$, and using $\delta(\tilde{r}/\bar{R}) = \bar{R}^{-2}\delta(\tilde{r})$, we get

$$\left[-\frac{1}{\tilde{r}} \frac{d}{d\tilde{r}} \left(\tilde{r} \frac{d}{d\tilde{r}} \right) + 1 \right] \tilde{\mu}(\tilde{r}) = \Gamma \delta(\tilde{r}) \quad .$$

The solution to this equation is given by the modified Bessel function of second kind and zeroth order K_0

$$\tilde{\mu}(\tilde{r}) = \frac{\Gamma}{2\pi} K_0(\tilde{r}) \quad .$$

Reinserting the defined variables yields

$$\mu(r) = -\frac{m(0)^2}{\pi m''(0)} K_0\left(\frac{r}{\bar{R}}\right) \quad . \quad (\text{B.40})$$

Note that the modified Bessel functions of second kind decay exponentially for long distances

$$K_i\left(\frac{r}{\bar{R}}\right) \xrightarrow{r \rightarrow \infty} \sqrt{\frac{\pi \bar{R}}{2r}} e^{-r/\bar{R}} \quad . \quad (\text{B.41})$$

But, consider that the inverse square root of the distance appears in front of the exponential. Formally, this is the one-dimensional result. The only difference here is, that $m(k)$ is a two-dimensional Fourier transform instead of a one-dimensional one and $\mu(r)$ contains modified Bessel functions of second kind instead of exponentials.

In order to evaluate the expression for the mean covariances Eq. (B.35), one needs to calculate the two-dimensional convolution of a modified Bessel function of second kind with itself, for which we use the following trick

$$\begin{aligned} (K_0 * K_0)\left(\frac{r}{\bar{R}}\right) &= \mathcal{F}^{-1}[\tilde{K}_0 \cdot \tilde{K}_0]\left(\frac{r}{\bar{R}}\right) \\ &= \frac{1}{2\pi} \mathcal{H}^{-1}\left[\frac{1}{(\beta + k^2)^2}\right](\sqrt{\beta}r) \\ &= -\frac{1}{2\pi} \frac{d}{d\beta} \mathcal{H}^{-1}\left[\frac{1}{\beta + k^2}\right](\sqrt{\beta}r) \\ &= -\frac{1}{2\pi} \frac{d}{d\beta} K_0(\sqrt{\beta}r) \\ &= \frac{R_{\text{eff},\mu} r}{4\pi} K_1\left(\frac{r}{\bar{R}}\right) \quad , \end{aligned}$$

where $\mathcal{F}$ denotes the Fourier transform, $\mathcal{H}$ denotes the Hankel transform, and $\beta = \bar{R}^{-2}$. The last step can be found in Abramowitz and Stegun (1964, 9.6.27).

The mean covariances are given by

$$\begin{aligned} \bar{c}(r) &= \frac{D}{1-\rho^2} \left[\delta(r) - 2\frac{m(0)^2}{\pi m''(0)} K_0\left(\frac{r}{\bar{R}}\right) + \frac{m(0)^4}{m''(0)^2} \frac{\bar{R}r}{4\pi^3} K_1\left(\frac{r}{\bar{R}}\right) \right] \\ &\xrightarrow{r \rightarrow \infty} \frac{D}{1-\rho^2} \left[\delta(r) - \frac{m(0)^2}{m''(0)} \sqrt{\frac{2\bar{R}}{\pi r}} e^{-r/\bar{R}} + \frac{m(0)^4}{m''(0)^2} \sqrt{\frac{\bar{R}^3 r}{32\pi^5}} e^{-r/\bar{R}} \right] \quad . \end{aligned}$$

Using

$$m(0) = Kw \equiv \zeta \quad , \quad m''(0) = -Kw \langle r^2 \rangle \quad , \quad (\text{B.42})$$

we get the effective decay constant

$$\bar{R} = \sqrt{\frac{1}{1-\zeta} \frac{\langle r^2 \rangle}{2}} R \quad . \quad (\text{B.43})$$

EXPONENTIAL CONNECTIVITY PROFILE: Using a two-dimensional exponential connectivity profile

$$p(x) = \frac{1}{2\pi R^2} e^{-|x|/R} \quad ,$$

leads to $\langle r^2 \rangle = 6R^2$, and we get

$$\bar{R} = \sqrt{\frac{3}{1-\lambda_0}} R \quad , \quad R_{\text{eff}} = \sqrt{\frac{3}{1-\rho^2}} R \quad ,$$

with $\lambda_0 = Kw$, and $\rho^2 = Kw^2$.

GAUSSIAN CONNECTIVITY PROFILE: Using a two-dimensional Gaussian connectivity profile

$$p(x) = \frac{1}{2\pi R^2} e^{-x^2/(2R^2)} \quad ,$$

leads to $\langle r^2 \rangle = 2R^2$, and we get

$$\bar{R} = \sqrt{\frac{1}{1-\lambda_0}} R \quad , \quad R_{\text{eff}} = \sqrt{\frac{1}{1-\rho^2}} R \quad .$$

B.11.3 Note on higher order approximation

While the (0,2)-Padé approximation seems to yield good results for the one-dimensional cases, in two dimensions the results only coincide for large spectral radii (Figure B.7 on page 129). One can extract a higher order approximation of the poles of the integration kernel of $\mu(x)$ and thereby the effective decay constant R_{eff} using the DLog-Padé-method, for which one calculates an $(n, n+1)$ -Padé approximation of the logarithmic derivative of the integration kernel around zero (Pelizzola, 1994). Using a (1,2)Padé approximation leads to

$$\bar{R} = \sqrt{-\frac{3(2m(0)-1)m''(0)^2 + (1-m(0))m(0)m''''(0)}{6m''(0)m(0)(1-m(0))}} \quad ,$$

which coincides with our previous results in the limit $m(0) \rightarrow 1$, and thus for large spectral radii. Note that this expression contains the fourth moment of the connectivity kernel $m''''(0) = wK \langle x^4 \rangle$.

B.12 TWO-DIMENSIONAL NETWORK WITH TWO POPULATIONS

Finally, we consider a two-dimensional network with two populations of neurons. As in the one dimensional case, the neurons are gathered in cells, which contain one inhibitory and q excitatory neurons. Again, they are placed on a two-dimensional lattice with periodic boundary conditions. The activity vector takes the form

$$\mathbf{x} = \begin{pmatrix} \mathbf{x}_{1,1}^{(E)} \\ \mathbf{x}_{1,1}^{(I)} \\ \mathbf{x}_{1,2}^{(E)} \\ \mathbf{x}_{1,2}^{(I)} \\ \vdots \\ \mathbf{x}_{1,N_y}^{(E)} \\ \mathbf{x}_{1,N_y}^{(I)} \\ \mathbf{x}_{2,1}^{(E)} \\ \mathbf{x}_{2,1}^{(I)} \\ \vdots \\ \mathbf{x}_{N_x,N_y}^{(E)} \\ \mathbf{x}_{N_x,N_y}^{(I)} \end{pmatrix}, \quad (\text{B.44})$$

where $\mathbf{x}_{i,j}^{(E)}$ denotes a q -dimensional vector.

B.12.1 Dimensionality reduction

We apply the procedure developed so far, which leads to the results we found in the one-dimensional case with two populations, with Fourier transforms and convolutions replaced by their two-dimensional analogons and modified Bessel functions of second kind instead of exponentials. So, we end up with

$$\begin{aligned} \overline{c_{EE}}(x) &= D_r^{(E)} \delta(x) + 2D_r^{(E)} \mu_{EE}(x) + D_r^{(E)} q (\mu_{EE} * \mu_{EE})(x) + D_r^{(I)} (\mu_{EI} * \mu_{EI})(x) , \\ \overline{c_{EI}}(x) &= D_r^{(E)} \mu_{IE}(x) + D_r^{(I)} \mu_{EI}(x) + D_r^{(E)} q (\mu_{EE} * \mu_{IE})(x) + D_r^{(I)} (\mu_{II} * \mu_{EI})(x) , \\ \overline{c_{II}}(x) &= D_r^{(I)} \delta(x) + 2D_r^{(I)} \mu_{II}(x) + D_r^{(E)} q (\mu_{IE} * \mu_{IE})(x) + D_r^{(I)} (\mu_{II} * \mu_{II})(x) , \end{aligned}$$

and $\mu_{ab}(x)$ given by Eq. (B.26) and the two-dimensional Fourier transform

$$m_{ab}(\mathbf{k}) = K_{ab} w_{ab} \int d^2x p_{ab}(x) e^{-i\mathbf{k} \cdot \mathbf{x}} .$$

The renormalized noise is given by Eq. (B.27) with spectral bound Eq. (B.28), with the eigenvalues $s_{ab}(k)$ replaced by the two-dimensional Fourier transforms $s_{ab}(\mathbf{k})$.

B.12.2 Long-range limit

Again, considering the special case in which the synaptic connections only depend on the type of the presynaptic neuron and not on the type of the postsynaptic neuron, the expressions simplify to

$$\mu_a(k) = \frac{m_a(k)}{1 - \zeta(k)} \quad , \quad (\text{B.45})$$

with

$$\zeta(k) = m_E(k) + m_I(k) \quad .$$

Padé approximation of the Fourier kernel, integration using Goldenfeld (1992, p.160f) and suppressing the zero arguments of ζ and m_a leads to

$$\begin{aligned} \mu_a(r) &= -\frac{\gamma_a w_a K_a}{2\pi(1-\zeta)} \bar{R}_a^2 K_0\left(\frac{r}{\bar{R}_a}\right) \\ &\xrightarrow{r \rightarrow \infty} -\frac{\gamma_a w_a K_a}{(1-\zeta)} \sqrt{\frac{1}{8\pi r \bar{R}_a^3}} e^{-r/\bar{R}_a} \quad , \end{aligned} \quad (\text{B.46})$$

with

$$\bar{R}_a = \sqrt{-\frac{\zeta''}{2\zeta} \frac{\zeta}{1-\zeta} - \frac{m_a''}{2m_a}} \quad .$$

After introducing the same relative parameters as in Section B.10.3, we find

$$\bar{R}_a = \sqrt{\frac{(\omega\kappa\tilde{\eta}^2 + 1)}{\omega\kappa + 1} \frac{\zeta}{1-\zeta} \frac{\langle x^2 \rangle_I}{2} + \frac{\langle x^2 \rangle_a}{2}} \quad . \quad (\text{B.47})$$

The two-dimensional convolutions are given by

$$\begin{aligned} (\mu_a * \mu_a)(r) &= \left[\frac{\gamma_a w_a K_a}{4(1-\zeta)} \right]^2 \frac{1}{\pi^3 \bar{R}_a^3} K_1\left(\frac{r}{\bar{R}_a}\right) \\ &\xrightarrow{r \rightarrow \infty} \left[\frac{\gamma_a w_a K_a}{4(1-\zeta)} \right]^2 \sqrt{\frac{1}{2\pi^5 \bar{R}_a^5 r}} e^{-r/\bar{R}_a} \quad . \end{aligned} \quad (\text{B.48})$$

The renormalized noise simplifies to Eq. (B.30). The mean covariances are given by

$$\begin{aligned} \overline{c_{EE}}(x) &= D_r [\delta(x) + 2\mu_E(x) + q(\mu_E * \mu_E)(x) + (\mu_I * \mu_I)(x)] \quad , \\ \overline{c_{EI}}(x) &= D_r [\mu_E(x) + \mu_I(x) + q(\mu_E * \mu_E)(x) + (\mu_I * \mu_I)(x)] \quad , \\ \overline{c_{II}}(x) &= D_r [\delta(x) + 2\mu_I(x) + q(\mu_E * \mu_E)(x) + (\mu_I * \mu_I)(x)] \quad . \end{aligned} \quad (\text{B.49})$$

Remember that the result for the variances of the covariances is obtained by substituting D_r by its square, and w_a , or ω respectively, by its square and setting $\zeta = \rho^2$.

Eq. (3.3) can be proven by inserting the result for $\bar{R}_a$

$$\bar{R}_E^2 - \bar{R}_I^2 = \frac{\langle x^2 \rangle_E}{2} - \frac{\langle x^2 \rangle_I}{2} = \text{const.} \cdot (R_E^2 - R_I^2) \quad .$$

Using an exponential connectivity profile yields $\text{const.} = 3$, a Gaussian connectivity profile yields $\text{const.} = 1$.

EXPONENTIAL CONNECTIVITY PROFILE: Using the results from B.11.2, we find

$$\bar{R}_a = \sqrt{3 \left[\frac{(\omega\kappa\eta^2 + 1)}{\omega\kappa + 1} \frac{\lambda_0}{1 - \lambda_0} R_I + R_a \right]}, \quad R_{\text{eff},a} = \sqrt{3 \left[\frac{(\omega^2\kappa\eta^2 + 1)}{\omega^2\kappa + 1} \frac{\rho^2}{1 - \rho^2} R_I + R_a \right]},$$

with $\lambda_0 = w_E K_E + w_I K_I$, and $R^2 = w_E^2 K_E + w_I^2 K_I$.

GAUSSIAN CONNECTIVITY PROFILE: Using the results from B.11.2, we find

$$\bar{R}_a = \sqrt{\frac{(\omega\kappa\eta^2 + 1)}{\omega\kappa + 1} \frac{\lambda_0}{1 - \lambda_0} R_I + R_a}, \quad R_{\text{eff},a} = \sqrt{\frac{(\omega^2\kappa\eta^2 + 1)}{\omega^2\kappa + 1} \frac{\rho^2}{1 - \rho^2} R_I + R_a}.$$

B.12.3 Higher order approximation

Using a (1,2)-DLog-Padé method as in B.11.3 yields

$$\bar{R}_a = \sqrt{-\frac{(1 - \zeta)^2 (m_a m_a'''' - 3m_a''^2) + m_a^2 [(1 - \zeta)\zeta'''' + 3\zeta''^2]}{6m_a(1 - \zeta) [(1 - \zeta)m_a'' + m\zeta'']}}, \quad (\text{B.50})$$

which again contains the fourth moments of the connectivity kernels.

B.13 VALIDATION OF THEORY




Figure B.7: **Comparison of simulation and theory.** (A) Variance of EE, EI, and II covariances as a function of distance. Darker dots are the results of the simulation. Lighter ones are the prediction of the discrete theory. (B) Variance of EE covariances as a function of distance (Eq. (B.49) for variances). The lightest blue dots are the predictions of the discrete theory (μ_a replaced by the discrete Fourier transform of Eq. (B.45), taking into account Section B.7), the medium blue line is the (0,2)-Padé prediction (μ_a replaced by its Padé approximation Eq. (B.46), taking into account Section B.7), and the dark blue line is the higher order (1,2)-DLog-Padé prediction (μ_a replaced by its Padé approximation Eq. (B.46), using Eq. (B.50), and taking into account Section B.7). (C) Fitted slope of linear regions in panel B for different spectral bounds ρ (light blue: discrete theory, medium blue: Padé approximation, dark blue: higher order Padé approximation).

In order to validate our results, we performed simulations, in which an effective connectivity matrix W of a two-dimensional network was drawn randomly, and covariances were calculated using the result from Lindner et al. (2005); Pernice et al. (2011); Trousdale et al. (2012)

$$c(W) = (1 - W)^{-1} D (1 - W)^{-T}.$$

The elements of the different components W_{ab} of the effective connectivity matrix, similar to Eq. (B.17), were drawn from a binomial distribution with K_b trials and a success probability of $\gamma_b p_b(|x|)$, with γ_b given by Eq. (B.29) and $|x|$ denoting the distance between the neurons.

We compared the results to the predictions by our discrete theory, continuum theory, and the long-range limit. We did this for all cases presented above: one dimension with one population, one dimension with two populations, two dimensions with one population, and two dimensions with two populations. In the cases of two populations we solely considered the special case of synaptic connections only depending on the type of the presynaptic neuron. The first three cases are not reported here. We simulated several sets of parameters, varying the number of neurons, the number of inputs, the decay constants and the spectral bound, of which we only report the one

using the parameters listed in Table B.3 on page 131, because the results do not differ qualitatively. Using

$$\rho^2 = s(0) = K_E w_E^2 + K_I w_I^2 \quad ,$$

and choosing

$$\frac{w_I}{w_E} = -\frac{N_E}{N_I} = -q \quad ,$$

we calculated the synaptic weights

$$w_E = \frac{\rho}{\sqrt{K_E + q^2 K_I}} \quad , \quad w_I = -\frac{q\rho}{\sqrt{K_E + q^2 K_I}} \quad .$$

The comparison of simulation and discrete theory is shown in Panel A of Figure B.7 on page 129. Simulation and theory match almost perfectly. The continuum theory, which is shown in Panels D and E of Figure 3.3 on page 35, matches as well as the discrete theory (not shown here). The slight shift in y-direction in Panel A of Figure B.7 on page 129 is due to the fact that in the random realization of the network the spectral bound is not exactly matching the desired value, but is slightly different for each realization and distributed around the chosen value. This jittering around the real spectral bound is more pronounced as $\rho \rightarrow 1$. Note that the simulated networks were small compared to the decay constant of the connectivity profile, in order to keep simulation times reasonable. This is why the variances do not fall off linearly in the semi-log plot. The kink and the related spreading starting around $x/R = 1.5$ is a finite size effect due to periodic boundary conditions: The maximal distance of two neurons along the axes in units of spatial decay constants is $(N/2)/R \approx 1.5$. Because of the periodic boundary conditions, the covariances between two neurons increases once the distance between them exceeds the maximal distance along an axis. This, together with the fact that the curve is the result of the discrete Fourier transform of Eq. (B.45), implies a zero slope at the boundary. This holds for any direction in the two dimensional plane, but the maximal distances between two neurons is longer for directions not aligned with any axis and depends on the precise direction, which explains the observed spreading.

In order to validate the long-range limit, we compared our discrete theory with the result from the Padé approximation at large distances (Panel B of Figure B.7 on page 129). We do not expect the Padé approximation to hold at small distances. We are mainly interested in the slope of the variance of covariances, because the slope determines how fast typical pairwise covariances decay with increasing inter-neuronal distance. The slope at large distances for the (0,2)-Padé approximation is smaller than the prediction by our theory, but the higher order approximation matches our theory very well (Panel C of Figure B.7 on page 129). In the limit $\rho \rightarrow 1$ both Padé predictions yield similar results. The absolute value of the covariances in the Padé approximation can be obtained from a residue analysis. The (0,2)-Padé approximation yields absolute values with a small offset, analogous to the slope results. Calculating the residues for the (1,2)-DLog Padé approximation would lead to a better approximation. Note that for plotting the higher order prediction in Panel B of Figure B.7 on page 129, we just inserted Eq. (B.50) into Eq. (B.46) and Eq. (B.48).

	Fig. 3.3B,C	Fig. 3.3D,E	Fig. B.7A	Fig. B.7B	
N_x	61	201	61	1001	# neurons in x-direction
N_y	61	201	61	1001	# neurons in y-direction
q	4	4	4	4	Ratio of E to I neurons
K_E	100	100	100	100	# E inputs per neuron
K_I	50	50	50	50	# I inputs per neuron
R_E	20	20	20	20	Spatial spread of E conn.
R_I	10	10	10	10	Spatial spread of I conn.
D	1	1	1	1	Squared noise amplitude
ρ	0.95	0.95	0.8	0.95	Spectral bound
	exp	exp	exp	exp	Connectivity kernel

Table B.3: Parameters used to create theory figures. Decay consts. in units of lattice const. a .

B.14 PARAMETERS OF NEST SIMULATION

Network Parameters		
N	2000	Number of neurons
p	0.1	Connection probability
τ	1 ms	Time constant
σ_μ	1 Hz	Standard deviation of external input
σ_{noise}	0.1 Hz	Standard deviation of noise
ρ	[0.1, 0.2, ..., 0.9]	Spectral bound
ϵ	0.1	Parameter controlling difference of two simulations
Simulation Parameters		
dt	0.1 ms	Simulation step size
t_{init}	100 ms	Initialization time
t_{sim}	2000000 ms	Simulation time without initialization time
t_{sample}	1 ms	Sample resolution at which rates where recorded
Analysis Parameters		
N_{sample}	200	Sample size
T	100 ms	Correlation time window

Table B.4: Parameters used for NEST simulation and subsequent analysis.

B.15 SOURCES OF HETEROGENEITY

Sparseness of connections is a large source of heterogeneity in cortical networks. It contributes strongly to the variance of effective connection weights that determines the spectral bound, the quantity that controls stability of balanced networks (Sompolinsky

et al., 1988; Dahmen et al., 2019): Consider the following simple model $W_{ij} = \mathcal{W}_{ij}\zeta_{ij}$ for the effective connection weights W_{ij} , where $\zeta_{ij} \in \{0, 1\}$ are independent Bernoulli numbers, which are 1 with probability p and 0 with probability $1 - p$, and $\mathcal{W}_{ij}$ are independently distributed amplitudes. The ζ_{ij} encode the sparseness of connections and the $\mathcal{W}_{ij}$ encode the experimentally observed distributions of synaptic amplitudes and single neuron heterogeneities that lead to different neuronal gains. Since $\mathcal{W}_{ij}$ and ζ_{ij} are independent, the variance of W_{ij} is

$$\text{Var}(W_{ij}) = p \cdot \text{Var}(\mathcal{W}_{ij}) + p(1 - p) \cdot \text{Mean}(\mathcal{W}_{ij})^2 .$$

For low connection probabilities as observed in cortex ($p(1 - p) \approx p$), assessing the different contributions to the variance thus amounts to comparing the mean and standard deviation of $\mathcal{W}_{ij}$. Even though synaptic amplitudes are broadly distributed in cortical networks, one typically finds that their mean and standard deviation are of the same magnitude (see e.g. Sayer et al., 1990, Tab.1; Feldmeyer et al., 1999, Tab. 1; Song et al., 2005, Fig.1; Lefort et al., 2009, Tab.2; Ikegaya et al., 2013, Fig.1; Loewenstein et al., 2011, Fig. 2). Sparseness of connections (second term on the right hand side) is thus one of the dominant contributors to the variance of connections. For simplicity, the other sources, in particular the distribution of synaptic amplitudes, are left out in this study. They can, however, be straight-forwardly added in the model and the theoretical formalism, because it only depends on $\text{Var}(\mathcal{W}_{ij})$.

ACTIVITY PATTERNS IN TWO-DIMENSIONAL NETWORKS

C.1 TRANSITION CURVES IN NETWORK WITH GAUSSIAN CONNECTIVITY

Here, we derive the transition curves plotted in Panel C of Figure 4.3 on page 68. We start with the effective profile for a network with Gaussian connectivity profiles Eq. (4.14). Introducing the dimensionless parameters

$$\kappa = \rho_E k, \quad \chi = \frac{R_I}{R_E}, \quad \eta = -\frac{w_I}{w_E},$$

we find the dimensionless effective profile

$$\tilde{c}(\kappa) = \frac{c(\kappa/R_E)}{w_E} = \exp\left(-\frac{\kappa^2}{2}\right) - \eta \exp\left(-\frac{\chi^2 \kappa^2}{2}\right).$$

The bifurcations of the system are determined by the position and value of the extrema of the effective profile, which are at

$$\kappa_0 = 0, \quad \kappa_{\pm} = \pm \sqrt{\frac{2 \ln(\eta \chi^2)}{\chi^2 - 1}}.$$

The transition curves indicate which types of patterns emerge when the network parameters are modified such that the homogeneous steady state loses its stability. The first transition curve separates the parameter subspaces into static and dynamic patterns. The second separates the subspaces for an unstable zero-mode $k_c = 0$ from multiple unstable modes $k_c \neq 0$.

Eq. (4.11) and Eq. (4.12) show that the homogeneous steady state might become unstable if either $c_{\max} = 1$, or $c_{\min} = -1$. Hence,

$$|\tilde{c}_{\min}| = |\tilde{c}_{\max}|$$

determines the first transition curve, as the system undergoes a Hopf bifurcation for $|\tilde{c}_{\min}| > |\tilde{c}_{\max}|$, and a Pitchfork bifurcation otherwise. Using the positions of the extrema, we find

$$\begin{aligned} |\tilde{c}(\kappa_0)| &= |\tilde{c}(\kappa_{\pm})|, \\ \Rightarrow |1 - \eta| &= \eta^{-1/(\chi^2 - 1)} \chi^{-2\chi^2/(\chi^2 - 1)} |1 - \chi^2|, \end{aligned}$$

which implicitly determines $\eta(\chi)$. We solve this equation numerically using `scipy.optimize.fsolve` (Virtanen et al., 2020).

The first transition curve determines whether the instability is caused by the maximum or minimum of the effective profile, whereas the second transition curve determines whether or not the respective extremum is at the origin. Because the effective profile is a sum of two Gaussian functions, we know that one of the extrema will be a minimum while the other will be a maximum as long as the non-zero extremum $\kappa_{\pm}$ exists. Thus, the second transition curve is determined by

$$\left. \frac{d^2}{d\kappa^2} \tilde{c}(\kappa) \right|_{\kappa=0} \stackrel{!}{=} 0 .$$

For $\left. \frac{d^2}{d\kappa^2} \tilde{c}(\kappa) \right|_{\kappa=0} < 0$ the maximum will be at the origin and vice versa. This yields

$$\eta(\chi) = \frac{1}{\chi^2} .$$

As one can show, the conditions for the non-zero extremum $\kappa_{\pm}$ to exist are always fulfilled in the parameter subspaces in which the relevant extremum is not at the origin.

C.2 NEST SIMULATION PARAMETERS

Symbol	Value	Description
$L_x \times L_y$	$2\pi \times 2\pi$	Area of simulated patch
N_I	28×28	Number of inhibitory neurons
N_E	28×28	Number of excitatory neurons
K_I	400	Number of inhibitory inputs
K_E	400	Number of excitatory inputs
τ	1.0	Model time constant
d	1.0	Delay
R_I	0.9	Inhibitory spatial spread
R_E	0.5	Excitatory spatial spread
w_I	-2.37	Inhibitory input weight
w_E	2.45	Excitatory input weight
β_0	0	Non-linearity offset $\psi(0)$
β_1	1	Non-linearity slope $\psi'(0)$
β_3	-0.083	Cubic non-linearity
T	3000	Simulation time
dt	0.1	Simulation time step
use_wfr	True	Use wavefront relaxation method
wfr_comm_interval	1.0	Wavefront communication interval
k_1, k_2	$(2, 0), (0, 2)$	Initial mode wavevectors

Table C.1: Parameters that are fix across all simulations.

C.3 ANALYSIS OF THE CONVOLUTIONS IN THE NON-LINEARITY

Assuming a four mode instability with critical modes $k_1 = -k_3$ and $k_2 = -k_4$, $k_c \equiv |k_1| = |k_2| = |k_3| = |k_4|$, $\mathcal{K} = \{k_1, k_2, k_3, k_4\}$, we write the activity using the characteristic vectors (Eq. (4.9) and Eq. (4.10))

$$\mathbf{X}(k, t) = \begin{cases} x(k, t) \mathbf{v}_c + z(k, t) \mathbf{v}_s(k) , & k \in \mathcal{K} , \\ y(k, t) \mathbf{v}_c + z(k, t) \mathbf{v}_s(k) , & k \notin \mathcal{K} . \end{cases} \quad (\text{C.1})$$

According to the center manifold theorem, y and z are at least of order $\mathcal{O}(x^2)$. We want to derive the dynamical equation for the critical x modes up to lowest order in the non-linearity. Therefore, we insert Eq. (C.1) into the dynamical equation Eq. (4.3) and project the result onto the adjoint characteristic vector $w_c(k_c)$, which yields

$$\begin{aligned} \tau \frac{\partial}{\partial t} x_i(t) = & -x_i(t) + \beta_1 c(k_c) x_i(t-d) \\ & + \beta_2 c(k_c) w_c(k_c) \cdot [\mathbf{u} \otimes \otimes \mathbf{u}](k_i, t-d) \\ & + \beta_3 c(k_c) w_c(k_c) \cdot [\mathbf{u} \otimes \otimes \mathbf{u} \otimes \otimes \mathbf{u}](k_i, t-d) , \end{aligned}$$

with

$$c(k) = w_c(k) \cdot \widetilde{\mathbf{M}}(k) = \widetilde{M}_E(k) + \widetilde{M}_I(k) .$$

Next, we need to consider the convolutions to lowest order in x :

$$\begin{aligned} & [\mathbf{X} \otimes \otimes \mathbf{X}](k_i, t) \\ &= \int d^2 k' \mathbf{X}(k_i - k', t) \odot \mathbf{X}(k', t) \\ &= 2 \sum_{j=1}^4 \mathbf{X}(k_i - k_j, t) \odot \mathbf{X}(k_j, t) \\ &= 2 \sum_{j=1}^4 [y(k_i - k_j, t) \mathbf{v}_c + z(k_i - k_j, t) \mathbf{v}_s(k_i - k_j)] \odot [x(k_j, t) \mathbf{v}_c + z(k_j, t) \mathbf{v}_s(k_j)] \\ &= 2 \sum_{j=1}^4 [x(k_j, t) y(k_i - k_j, t) \mathbf{v}_c + x(k_j, t) z(k_i - k_j, t) \mathbf{v}_s(k_i - k_j)] + \mathcal{O}(x^4) . \end{aligned}$$

Projection onto $w_c(k_c)$ yields

$$\begin{aligned} & w_c(k_c) \cdot [\mathbf{X} \otimes \otimes \mathbf{X}](k_i, t) \\ &= 2 \sum_{j=1}^4 [x(k_j, t) y(k_i - k_j, t)] + \mathcal{O}(x^4) \\ &= 2 \left[x(k_i, t) y(0, t) + \sum_{j \neq i} x(k_j, t) y(k_i - k_j, t) \right] + \mathcal{O}(x^4) \\ &\equiv 2 \left[x_i(t) y_0(t) + \sum_{j \neq 1} x_j(t) y_{ij}(t) \right] + \mathcal{O}(x^4) , \end{aligned}$$

where we introduced the shorthand notation

$$x(k_i, t) \equiv x_i(t) , \quad y(0, t) \equiv y_0(t) , \quad y(k_i - k_j, t) \equiv y_{ij}(t) .$$

Note that some of these modes are equivalent: for our analysis, we only use

$$y_0, y_{13}, y_{24}, y_{31}, y_{42}, y_{12}, y_{23}, y_{34}, y_{41} ,$$

as

$$y_{14} = y_{23}, y_{43} = y_{12}, y_{32} = y_{41}, y_{21} = y_{34} .$$

The convolution coming from the third order non-linearity yields

$$\begin{aligned} w_c(k_c) \cdot [X \circledast \circledast X \circledast \circledast X](k_i, t - d) \\ = w_c(k_c) \cdot \int d^2 k' d^2 k'' X(k_i - k', t) \odot X(k' - k'', t) \odot X(k'', t) \\ = 3x_i(t) x_i(t) x_{i+2}(t) + 6x_i(t) x_{i+1}(t) x_{i+3}(t) + \mathcal{O}(x^4) , \end{aligned}$$

where we assumed periodic indices $i + 4 \equiv i$. The convolutions in the dynamical equations of the zero- and difference-modes can be analyzed similarly and lead to Eq. (4.23), Eq. (4.24), and Eq. (4.25).

C.4 CENTER MANIFOLD REDUCTION FOR STATIC PATTERNS

For performing the center manifold reduction, we take the critical modes, including the distance to criticality, to be

$$\tilde{\mathbf{u}} = (u_1, u_2, u_3, u_4, \epsilon) .$$

The tilde serves as a reminder that all operators and vectors take into account the critical parameter ϵ in addition to the critical modes. We define the linear operator components and their actions on $\boldsymbol{\phi}(\theta) \in C([-d, 0], \mathbb{R}^N)$ with the number of neuron grid points N

$$\begin{aligned} \tilde{L}_0 \boldsymbol{\phi} &= -\frac{1}{\tau} \boldsymbol{\phi}(0) , \\ \tilde{L}_d(\epsilon) \boldsymbol{\phi} &= \frac{\beta_1}{\tau} \text{diag} \left[1 + \epsilon, 1 + \epsilon, 1 + \epsilon, 1 + \epsilon, \right. \\ &\quad c(0), c(2k_c), c(2k_c), c(2k_c), c(2k_c), \\ &\quad c(\sqrt{2}k_c), c(\sqrt{2}k_c), c(\sqrt{2}k_c), c(\sqrt{2}k_c), \\ &\quad \left. 0, \dots, 0 \right] \cdot \boldsymbol{\phi}(-d) . \end{aligned}$$

Using the results Eq. (4.22)–Eq. (4.26) and $c(k_c) = 1 + \epsilon$, we can define the non-linear operator to lowest order

$$\tilde{F}(\mathbf{u}, \mathbf{h}) = \frac{\beta_2}{\tau} \begin{pmatrix} u_1 h_0(-d) + u_2 h_{12}(-d) + u_3 h_{13}(-d) + u_4 h_{23}(-d) \\ u_1 h_{34}(-d) + u_2 h_0(-d) + u_3 h_{23}(-d) + u_4 h_{24}(-d) \\ u_1 h_{31}(-d) + u_2 h_{41}(-d) + u_3 h_0(-d) + u_4 h_{12}(-d) \\ u_1 h_{41}(-d) + u_2 h_{42}(-d) + u_3 h_{12}(-d) + u_4 h_0(-d) \\ 2c(0)(u_1 u_3 + u_2 u_4) \\ c(2k_c)u_1^2 \\ c(2k_c)u_2^2 \\ c(2k_c)u_3^2 \\ c(2k_c)u_4^2 \\ 2c\left(\sqrt{2}k_c\right)u_1 u_2 \\ 2c\left(\sqrt{2}k_c\right)u_2 u_3 \\ 2c\left(\sqrt{2}k_c\right)u_3 u_4 \\ 2c\left(\sqrt{2}k_c\right)u_1 u_4 \\ 0 \\ \vdots \\ 0 \end{pmatrix},$$

where we suppressed the $\tilde{\mathbf{u}}$ dependence of $\tilde{\mathbf{h}}(\theta, \tilde{\mathbf{u}})$ and used the same index notation as for the zero and difference modes in 4.6. As the characteristic value for static bifurcations is $\lambda = 0$, we find

$$\tilde{\mathbf{B}} = \text{diag}(0, 0, 0, 0, 0).$$

which we get from inserting Eq. (4.17) into the second line of Eq. (4.18) and applying the chain rule in the derivative of h using Eq. (4.21). This yields the ODEs that determine the functions in the center manifold ansatz up to integration constants. For fixing the constants, we use the second line of Eq. (4.18) (cf. Campbell (2009, Eq. 32), Qesmi et al. (2006, Eq. 2.8), or Qesmi et al. (2007, Eq. 3.6))

$$\begin{aligned} L(h(\theta, u)) + F(\Phi \cdot u) = & \left\{ \frac{\partial h}{\partial u}(\theta, u) \cdot B \cdot u + \frac{\partial h}{\partial u}(\theta, u) \Psi(0) \cdot (L(\epsilon) - L(0)) \Phi \cdot u \right. \\ & + \Phi(\theta) \Psi(0) \cdot (L(\epsilon) - L(0)) [\Phi \cdot u + h(\theta, u(t))] \\ & \left. + \Phi(\theta) \Psi(0) \cdot F(\Phi \cdot u) \right\} \Big|_{\theta=0}. \end{aligned}$$

The calculations are performed with SymPy (Meurer et al., 2017). The results of the analysis are presented in Section 4.6.1 and 4.6.2.

BIBLIOGRAPHY

- L F Abbott, K Rajan, and H Sompolinsky. Interactions between intrinsic and stimulus-evoked activity in recurrent neural networks. In *The Dynamic Brain: An Exploration of Neuronal Variability and its Functional Significance*, pages 65–82, 2011.
- LF Abbott and Thomas B Kepler. Model neurons: from Hodgkin-Huxley to Hopfield. In *Statistical mechanics of neural networks*, pages 5–18. Springer, 1990.
- Moshe Abeles. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, Cambridge, 1st edition, 1991.
- M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions*, volume 55 of *Applied Mathematics Series*. National Bureau of Standards, Washington, 1964.
- Johnatan Aljadeff, David Renfrew, Marina Vugué, and Tatyana O Sharpee. Low-dimensional dynamics of structured random networks. *Phys. Rev. E*, 93(2):022302, 2016.
- Johnatan Aljadeff, Merav Stern, and Tatyana Sharpee. Transition to chaos in random networks with cell-type-specific connectivity. *Phys. Rev. Lett.*, 114:088101, Feb 2015. URL <http://link.aps.org/doi/10.1103/PhysRevLett.114.088101>.
- Asohan Amarasingham, Ting-Li Chen, Matthew T Harrison, and David L Sheinberg. Spike count reliability and the poisson hypothesis. *Journal of Neuroscience*, 26(4): 801–809, January 2006.
- Shun-Ichi Amari. Homogeneous nets of neuron-like elements. *Biol. Cybern.*, 17(4): 211–220, 1975. ISSN 1432-0770.
- Shun-Ichi Amari. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol. Cybern.*, 27(2):77–87, 1977. URL <https://doi.org/10.1007/bf00337259>.
- D. J. Amit and M. V. Tsodyks. Quantitative study of attractor neural network retrieving at low spike rates I: Substrate – spikes, rates and neuronal gain. *Network*, 2:259, 1991.
- Daniel J. Amit and Nicolas Brunel. Dynamics of a recurrent network of spiking neurons before and following learning. *Network: Comp. Neural Sys.*, 8(4):373–404, 1997a. URL https://doi.org/10.1088/0954-898x_8_4_003.
- Daniel J. Amit and Nicolas Brunel. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex*, 7: 237–252, Apr. 1997b. URL <https://doi.org/10.1093/cercor/7.3.237>.
- MC Angulo, J Rossier, and E Audinat. Postsynaptic glutamate receptors and integrative properties of fast-spiking interneurons in the rat neocortex. *Journal of Neurophysiology*, 82(3):1295–302, Sep 1999.

- Natalie Baddour. Two-dimensional fourier transforms in polar coordinates. In *Advances in Imaging and Electron Physics*, pages 1–45. Elsevier, 2011. URL <https://doi.org/10.1016/b978-0-12-385861-0.00001-4>.
- Cody Baker, Christopher Ebsch, Ilan Lampl, and Robert Rosenbaum. Correlated states in balanced neuronal networks. *Phys. Rev. E*, 99:052414, May 2019. URL <https://link.aps.org/doi/10.1103/PhysRevE.99.052414>.
- Peter Bartho, Hajime Hirase, Lenaic Monconduit, Michael Zugaro, Kenneth D. Harris, and György Buzsáki. Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *Journal of Neurophysiology*, 92: 600–608, 2004.
- JL Basdevant. The padé approximation and its physical applications. *Fortschritte der Physik*, 20(5):283–331, 1972.
- R. Ben-Yishai, R.L. Bar-Or, and H. Sompolinsky. Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences*, 92:3844, 1995.
- Tom Binzegger, Rodney J. Douglas, and Kevan A. C. Martin. Stereotypical bouton clustering of individual neurons in cat primary visual cortex. *Journal of Neuroscience*, 27(45):12242–12254, Nov. 2007. URL <https://doi.org/10.1523/jneurosci.3753-07.2007>.
- Barak Blumenfeld, Dmitri Bibitchkov, and Misha Tsodyks. Neural network model of the primary visual cortex: From functional architecture to lateral connectivity and back. *Journal of computational neuroscience*, 20(2):219, 2006.
- Hannah Bos, Markus Diesmann, and Moritz Helias. Identifying anatomical origins of coexisting oscillations in the cortical microcircuit. *PLOS Comput. Biol.*, 12(10): e1005132, Oct. 2016. URL <http://doi.org/10.1371/journal.pcbi.1005132>.
- William H. Bosking, Yiang Zhang, Brett Schofield, and David Fitzpatrick. Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *Journal of Neuroscience*, 17(6):2112–2127, 1997. URL <https://doi.org/10.1523/jneurosci.17-06-02112.1997>.
- Clemens Boucsein, Martin Nawrot, Philipp Schnepel, and Ad Aertsen. Beyond the cortical column: abundance and physiology of horizontal connections imply a strong role for inputs from the surround. *Front. Neurosci.*, 5:32, Apr 2011. URL <https://doi.org/10.3389/fnins.2011.00032>.
- Valentin Braitenberg and Almut Schüz. *Cortex: Statistics and Geometry of Neuronal Connectivity*. Springer-Verlag, Berlin, 2nd edition, 1998. ISBN 3-540-63816-4.
- P. C. Bressloff and S. Coombes. Spike train dynamics underlying pattern formation in integrate-and-fire oscillator networks. *Phys. Rev. Lett.*, 81(11):2384–2387, Sep. 1998. URL <https://doi.org/10.1103/physrevlett.81.2384>.
- P. C. Bressloff and S. Coombes. A dynamical theory of spike train transitions in networks of integrate-and-fire oscillators. *SIAM Journal on Applied Mathematics*, 60(3):820–841, Jan. 2000. URL <https://doi.org/10.1137/s0036139998339643>.

- Paul C. Bressloff. Spatiotemporal dynamics of continuum neural fields. *Journal of Physics A: Mathematical and Theoretical*, 45(3):033001, 2012. URL <https://doi.org/10.1088/1751-8113/45/3/033001>.
- Paul C. Bressloff, Jack D. Cowan, Martin Golubitsky, Peter J. Thomas, and Matthew C. Wiener. Geometric visual hallucinations, euclidean symmetry and the functional architecture of striate cortex. *Phil. Trans. R. Soc. B*, 356(1407):299–330, Mar. 2001. URL <https://doi.org/10.1098/rstb.2000.0769>.
- Braden A. W. Brinkman, Fred Rieke, Eric Shea-Brown, and Michael A. Buice. Predicting how and when hidden neurons skew measured synaptic interactions. *PLOS Comput. Biol.*, 14(10):e1006490, oct 2018. URL <https://doi.org/10.1371/journal.pcbi.1006490>.
- Thomas Brochier, Lyuba Zehl, Yaoyao Hao, Margaux Duret, Julia Sprenger, Michael Denker, Sonja Grün, and Alexa Riehle. Massively parallel recordings in macaque motor cortex during an instructed delayed reach-to-grasp task. *Scientific Data*, 5: 180055, Apr. 2018. URL <https://doi.org/10.1038/sdata.2018.55>.
- Korbinian Brodmann. *Vergleichende Lokalisationslehre der Großhirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Johann Ambrosius Barth, Leipzig, 1909.
- Nicolas Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8(3):183–208, 2000a. URL <https://doi.org/10.1023/a:1008925309027>.
- Nicolas Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8(3):183–208, 2000b. URL <https://doi.org/10.1023/a:1008925309027>.
- Nicolas Brunel and Peter Latham. Firing rate of the noisy quadratic integrate-and-fire neuron. *Neural Comput.*, 15(10):2281–2306, 2003.
- Julian M. Budd and Zoltán F. Kisvárdy. Local lateral connectivity of inhibitory clutch cells in layer 4 of cat visual cortex (area 17). *Exp. Brain Res.*, 140(2):245–250, Jul. 2001. URL <https://doi.org/10.1007/s002210100817>.
- Michael A Buice and Carson C Chow. Beyond mean field theory: statistical field theory for neural networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(03):P03003, 2013.
- A. N. Burkitt. A review on the integrate-and-fire neuron model: I. homogenous synaptic input. *Biol. Cybern.*, 95(1):1–19, 2006.
- P. Buzás, K. Kovács, A. S. Ferecskó, J. M. L. Budd, U. T. Eysel, and Z. F. Kisvárdy. Model-based analysis of excitatory lateral connections in the visual cortex. *J. Comp. Neurol.*, 499:861–881, 2006. URL <https://doi.org/10.1002/cne.21134>.
- György Buzsáki, Costas A. Anastassiou, and Christof Koch. The origin of extracellular fields and currents – EEG, ECoG, LFP and spikes. *Nat. Rev. Neurosci.*, 13(6):407–427, June 2012. URL <https://doi.org/10.1038/nrn3241>.

- György Buzsáki and Andreas Draguhn. Neuronal oscillations in cortical networks. *Science*, 304:1926–1929, 2004.
- Sue Ann Campbell. Calculating centre manifolds for delay differential equations using mapleTM. In *Delay Differential Equations*, pages 1–24. Springer US, 2009. URL https://doi.org/10.1007/978-0-387-85595-0_8.
- Jack Carr. Center manifold. *Scholarpedia*, 1(12):1826, 2006. URL <https://doi.org/10.4249/scholarpedia.1826>.
- S. Chemla and F. Chavane. Voltage-sensitive dye imaging: Technique review and models. *J. Physiol. (Paris)*, 104(1-2):40–50, Jan. 2010. URL <https://doi.org/10.1016/j.jphysparis.2009.11.009>.
- S. Coombes. Waves, bumps, and patterns in neural field theories. *Biological Cybernetics*, 93:91–108, Jul. 2005. URL <https://doi.org/10.1007/s00422-005-0574-y>.
- S. Coombes. Large-scale neural dynamics: Simple and complex. *NeuroImage*, 52(3):731–739, Sep. 2010. URL <https://doi.org/10.1016/j.neuroimage.2010.01.045>.
- S. Coombes, N. A. Venkov, L. Shiau, I. Bojak, D. T. J. Liley, and C. R. Laing. Modeling electrocortical activity through improved local approximations of integral neural field equations. *Phys. Rev. E*, 76(5), Nov. 2007. URL <https://doi.org/10.1103/physreve.76.051901>.
- Stephen Coombes, Peter bei Graben, Roland Potthast, and James Wright. *Neural Fields. Theory and Applications*. Springer-Verlag Berlin Heidelberg, 2014.
- Stephen Coombes and Carlo Laing. Delays in activity-based neural networks. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1891):1117–1129, Feb. 2009. URL <https://doi.org/10.1098/rsta.2008.0256>.
- R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth. On the lambert w function. *Advances in Computational Mathematics*, 5(1):329–359, 1996. URL <https://doi.org/10.1007/BF02124750>.
- Jozsef Csicsvari, Hajime Hirase, Andras Czurko, and György Buzsáki. Reliability and state dependence of pyramidal cell–interneuron synapses in the hippocampus: an ensemble approach in the behaving rat. *Neuron*, 21(1):179–189, 1998.
- John P. Cunningham and M. Yu Byron. Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.*, 17(11):1500–1509, Aug. 2014. URL <https://doi.org/10.1038/nn.3776>.
- Paulina Anna Dąbrowska, Nicole Voges, Michael von Papen, Junji Ito, David Dahmen, Alexa Riehle, Thomas Brochier, and Sonja Grün. On the complexity of resting state spiking activity in monkey motor cortex. *bioRxiv*, May 2020. URL <https://doi.org/10.1101/2020.05.28.121095>.
- David Dahmen, Sonja Grün, Markus Diesmann, and Moritz Helias. Second type of criticality in the brain uncovers rich multiple-neuron dynamics. *Proceedings of*

- the National Academy of Sciences*, 116:13051–13060, 2019. URL <https://doi.org/10.1073/pnas.1818972116>.
- David Dahmen, Moritz Layer, Lukas Deutz, Paulina Anna Dąbrowska, Nicole Voges, Michael von Papen, Thomas Brochier, Alexa Riehle, Markus Diesmann, Sonja Grün, et al. Global organization of neuronal activity only requires unstructured local connectivity. *eLife*, 11:e68422, 2022. URL <https://doi.org/10.7554/elife.68422>.
- David Dahmen, Stefano Recanatesi, Gabriel Koch Ocker, Xiaoxuan Jia, Moritz Helias, and Eric Shea-Brown. Strong coupling and local control of dimensionality across brain areas. *bioRxiv*, 2020.
- Daniele Daini, Giacomo Ceccarelli, Enrico Cataldo, and Viktor Jirsa. Spherical-harmonics mode decomposition of neural field equations. *Physical Review E*, 101(1), Jan. 2020. URL <https://doi.org/10.1103/physreve.101.012202>.
- Ran Darshan, Carl van Vreeswijk, and David Hansel. Strength of correlations in strongly recurrent neuronal networks. *Phys. Rev. X*, 8:031072, Sep 2018. URL <https://link.aps.org/doi/10.1103/PhysRevX.8.031072>.
- Stefan Dasbach, Tom Tetzlaff, Markus Diesmann, and Johanna Senk. Dynamical characteristics of recurrent neuronal networks are robust against low synaptic weight resolution. *Front. Neurosci.*, 15, 2021. URL <http://doi.org/10.3389/fnins.2021.757790>.
- P. Dayan and L. F. Abbott. *Theoretical Neuroscience*. MIT Press, Cambridge, 2001.
- Jaime De la Rocha, Brent Doiron, Eric Shea-Brown, Josic Kresimir, and Alex Reyes. Correlation between neural spike trains increases with firing rate. *Nature*, 448(16): 802–807, august 2007.
- J. DeFelipe, L. Alonso-Nanclares, and J.I. Arellano. Microstructure of the neocortex: comparative aspects. *J. Neurocytol.*, 31:299–316, 2002. URL <https://doi.org/10.1023/A:1024130211265>.
- J. DeFelipe, M. Conley, and E. G. Jones. Long-range focal collateralization of axons arising from corticocortical cells in monkey sensory-motor cortex. *Journal of Neuroscience*, 6(12):3749–3766, 1986.
- Nima Dehghani, Adrien Peyrache, Bartosz Telenczuk, Michel Le Van Quyen, Eric Halgren, Sydney S. Cash, Nicholas G. Hatsopoulos, and Alain Destexhe. Dynamic balance of excitation and inhibition in human and monkey neocortex. *Scientific Reports*, 6, Mar. 2016. URL <https://doi.org/10.1038/srep23176>.
- Michael Denker, Lyuba Zehl, Bjørge E. Kilavik, Markus Diesmann, Thomas Brochier, Alexa Riehle, and Sonja Grün. LFP beta amplitude is linked to mesoscopic spatio-temporal phase patterns. *Scientific Reports*, 8(1):1–21, Mar. 2018. URL <https://doi.org/10.1038/s41598-018-22990-7>.
- Michael R DeWeese, Michael Wehr, and Anthony M Zador. Binary spiking in auditory cortex. *J. Neurosci.*, 23(21):7940–7949, 2003.

- Michael R. DeWeese and Anthony M. Zador. Non-gaussian membrane potential dynamics imply sparse, synchronous activity in auditory cortex. *Journal of Neuroscience*, 26(47):12206–12218, 2006.
- Odo Diekmann, Sjoerd M. Verduyn Lunel, Stephan A. van Gils, and Hanns-Otto Walther. *Delay Equations*. Springer New York, 1995. URL <https://doi.org/10.1007/978-1-4612-4206-2>.
- Markus Diesmann, Marc-Oliver Gewaltig, and Ad Aertsen. Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402(6761):529–533, 1999.
- K. Dijkstra, S.A. van Gils, S.G. Janssens, Yu.A. Kuznetsov, and S. Visser. Pitchfork–hopf bifurcations in 1d neural field models with transmission delays. *Physica D*, 297: 88–101, Mar. 2015. URL <https://doi.org/10.1016/j.physd.2015.01.004>.
- Eusebius J Doedel and B Oldeman. Auto-07p: continuation and bifurcation software. *Concordia Univ. Canada*, 1998.
- Freeman J Dyson. Is science mostly driven by ideas or by tools? *Science*, 338(6113): 1426–1427, 2012.
- Gamaleldin F Elsayed, Antonio H Lara, Matthew T Kaufman, Mark M Churchland, and John P Cunningham. Reorganization between preparatory and movement population responses in motor cortex. *Nat. Commun.*, 7(1):1–15, 2016.
- Bard Ermentrout. Neural networks as spatio-temporal pattern-forming systems. *Reports on Progress in Physics*, 61(4):353–430, 1998. URL <https://doi.org/10.1088/0034-4885/61/4/002>.
- Bard Ermentrout. *Simulating, Analyzing, and Animating Dynamical Systems: A Guide to Xppaut for Researchers and Students (Software, Environments, Tools)*. Soc for Industrial & Applied Math, 2002. ISBN 898715067.
- G. B. Ermentrout and J. D. Cowan. A mathematical theory of visual hallucination patterns. *Biol. Cybern.*, 34:137–150, 1979.
- Jian Fang and Grégory Faye. Monotone traveling waves for delayed neural field equations. *Mathematical Models and Methods in Applied Sciences*, 26(10):1919–1954, Aug. 2016. URL <https://doi.org/10.1142/s0218202516500482>.
- Grégory Faye and Olivier Faugeras. Some theoretical and numerical results for delayed neural field equations. *Physica D*, 239(9):561–578, May 2010. URL <https://doi.org/10.1016/j.physd.2010.01.010>.
- Dirk Feldmeyer, Veronica Egger, Joachim Lübke, and Bert Sakmann. Reliable synaptic connections between pairs of excitatory layer 4 neurones within a single "barrel" of developing rat somatosensory cortex. *Journal of Physiology*, 521(1):169–190, 1999. URL <http://jp.physoc.org/content/521/1/169.abstract>.
- Stefanos E. Folias and G. Bard Ermentrout. Bifurcations of stationary solutions in an interacting pair of e-i neural fields. *SIAM Journal on Applied Dynamical Systems*, 11(3):895–938, Jan. 2012. URL <https://doi.org/10.1137/110860094>.

- Nicolas Fourcaud and Nicolas Brunel. Dynamics of the firing probability of noisy integrate-and-fire neurons. *Neural Comput.*, 14:2057–2110, 2002. URL <https://doi.org/10.1162/089976602320264015>.
- N. Fourcaud-Trocmé, D. Hansel, C. van Vreeswijk, and N. Brunel. How spike generation mechanisms determine the neuronal response to fluctuating inputs. *Journal of Neuroscience*, 23:11628–11640, 2003.
- Juan A Gallego, Matthew G Perich, Rameez H Chowdhury, Sara A Solla, and Lee E Miller. Long-term stability of cortical population dynamics underlying consistent behavior. *Nat. Neurosci.*, 23(2):260–270, 2020.
- Juan A Gallego, Matthew G Perich, Lee E Miller, and Sara A Solla. Neural manifolds for the control of movement. *Neuron*, 94(5):978–984, 2017.
- Juan A Gallego, Matthew G Perich, Stephanie N Naufel, Christian Ethier, Sara A Solla, and Lee E Miller. Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nat. Commun.*, 9(1):1–13, 2018.
- Crispin W. Gardiner. *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*. Number 13 in Springer Series in Synergetics. Springer-Verlag, Berlin, 1983.
- Crispin W. Gardiner. *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*. Number 13 in Springer Series in Synergetics. Springer-Verlag, Berlin, 2nd edition, 1985. ISBN 3-540-61634-9, 3-540-15607-0.
- Richard Gast, Daniel Rose, Christoph Salomon, Harald E Möller, Nikolaus Weiskopf, and Thomas R Knösche. Pyrates - a python framework for rate-based neural simulations. *PLOS ONE*, 14(12):e0225900, 2019.
- Jörg RP Geiger, Joachim Lübke, Arnd Roth, Michael Frotscher, and Peter Jonas. Submillisecond ampa receptor-mediated signaling at a principal neuron-interneuron synapse. *Neuron*, 18:1009–1023, 1997.
- A. P. Georgopoulos, R. Caminiti, J. F. Kalaska, and J. T. Massey. Spatial coding of movement: A hypothesis concerning the coding of movement direction by motor cortical populations. *Experimental Brain Research [Suppl.]*, 7:327–336, 1983.
- Wulfram Gerstner, Werner M. Kistler, Richard Naud, and Liam Paninski. *Neuronal Dynamics. From Single Neurons to Networks and Models of Cognition*. Cambridge University Press, Cambridge, 2014.
- Marc-Oliver Gewaltig and Markus Diesmann. NEST (NEural Simulation Tool). *Scholarpedia*, 2(4):1430, 2007. URL <https://doi.org/10.4249/scholarpedia.1430>.
- Martin A. Giese. *Dynamic neural field theory for motion perception*, volume 469. Springer Science & Business Media, 2012.
- Iris Ginzburg and Haim Sompolinsky. Theory of correlations in stochastic neural networks. *Phys. Rev. E*, 50:3171–3191, Oct 1994. URL <https://link.aps.org/doi/10.1103/PhysRevE.50.3171>.

- Nigel Goldenfeld. *Lectures on phase transitions and the renormalization group*. Perseus books, Reading, Massachusetts, 1992.
- Bruno Golosio, Gianmarco Tiddia, Chiara De Luca, Elena Pastorelli, Francesco Simula, and Pier Stanislao Paolucci. Fast simulations of highly-connected spiking cortical models using GPUs. *Front. Comput. Neurosci.*, 15, Feb. 2021. URL <https://doi.org/10.3389/fncom.2021.627620>.
- A. Grabska-Barwinska and P. Latham. How well do mean field theories of spiking quadratic-integrate-and-fire networks work in realistic parameter regimes? *Journal of Computational Neuroscience*, 36(3):469–81, 2014.
- Christine Grienberger and Arthur Konnerth. Imaging calcium in neurons. *Neuron*, 73(5):862–885, Mar. 2012. URL <https://doi.org/10.1016/j.neuron.2012.02.011>.
- D. Grytskyy, T. Tetzlaff, M. Diesmann, and M. Helias. A unified view on weakly correlated recurrent networks. *Front. Comput. Neurosci.*, 7:131, 2013.
- Shangjiang Guo and Juanjuan Man. Center manifolds theorem for parameterized delay differential equations with applications to zero singularities. *Nonlinear Analysis: Theory, Methods & Applications*, 74(13):4418–4432, Sep. 2011. URL <https://doi.org/10.1016/j.na.2011.04.003>.
- M. Gur, A. Beylin, and D.M. Snodderly. Response variability of neurons in primary visual cortex (v1) of alert monkeys. *Journal of Neuroscience*, 8(17):2914–2920, apr 15th 1997.
- Espen Hagen, David Dahmen, Maria L. Stavrinou, Henrik Lindén, Tom Tetzlaff, Sacha J. van Albada, Sonja Grün, Markus Diesmann, and Gaute T. Einevoll. Hybrid scheme for modeling local field potentials from point-neuron networks. *Cereb. Cortex*, 26(12):4461–4496, Oct. 2016. URL <https://doi.org/10.1093/cercor/bhw237>.
- Jan Hahne, David Dahmen, Jannis Schuecker, Andreas Frommer, Matthias Bolten, Moritz Helias, and Markus Diesmann. Integration of continuous-time dynamics in a spiking neural network simulator. *Front. Neuroinformatics*, 11:34, May 2017. ISSN 1662-5196. URL <https://doi.org/10.3389/fninf.2017.00034>.
- Jan Hahne, Sandra Diaz, Alexander Patronis, Wolfram Schenck, Alexander Peyser, Steffen Graber, Sebastian Spreizer, Stine Brekke Vennemo, Tammo Ippen, Håkon Mørk, Jakob Jordan, Johanna Senk, Sara Konradi, Philipp Weidel, Tanguy Fardet, David Dahmen, Dennis Terhorst, Jonas Stapmanns, Guido Trensche, Alexander van Meegen, Jari Pronold, Jochen Martin Eppler, Charl Linssen, Abigail Morrison, Ankur Sinha, Jessica Mitchell, Susanne Kunkel, Rajalekshmi Deepu, Espen Hagen, Tom Vierjahn, Nilton Liuji Kamiji, Robin de Schepper, Pedro Machado, Jasper Albers, Wouter Klijn, Alex Myczko, William Mayner, Pooja Nagendra Babu, Hanjia Jiang, Sebastian Billaudelle, Benedikt S. Vogler, Guilherme Miotto, Lionel Kusch, Alberto Antonietti, Aitor Morales-Gregorio, Joris Dolderer, Younes Bouhadjar, and Hans Ekkehard Plesser. Nest 3.0, Jun. 2021. URL <https://doi.org/10.5281/zenodo.4739103>.

- Jack K. Hale and Sjoerd M. Verduyn Lunel. *Introduction to Functional Differential Equations*. Springer New York, 1993. URL <https://doi.org/10.1007/978-1-4612-4342-7>.
- David Hansel and Haim Sompolinsky. Modeling feature selectivity in local cortical circuits. In Christof Koch and Idan Segev, editors, *Methods in Neuronal Modeling*, pages 499–567. MIT Press, Cambridge, Massachusetts, 2 edition, 1998.
- Yaoyao Hao, Alexa Riehle, and Thomas G. Brochier. Mapping horizontal spread of activity in monkey motor cortex using single pulse microstimulation. *Frontiers in neural circuits*, 10:104, Dec. 2016. URL <https://doi.org/10.3389/fncir.2016.00104>.
- Charles R. Harris, K. Jarrod Millman, Stéfan J van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585:357–362, 2020.
- Nicholas Hatsopoulos, Jignesh Joshi, and John G. O’Leary. Decoding continuous and discrete motor behaviors using motor and premotor cortical ensembles. *Journal of Neurophysiology*, 92(2):1165–1174, Aug. 2004. URL <https://doi.org/10.1152/jn.01245.2003>.
- Thomas Heiberg, Birgit Kriener, Tom Tetzlaff, Gaute T. Einevoll, and Hans E. Plesser. Firing-rate models for neurons with a broad repertoire of spiking behaviors. *Journal of Computational Neuroscience*, 2018. URL <https://doi.org/10.1007/s10827-018-0693-9>.
- Stewart Heitmann, Matthew J Aburn, and Michael Breakspear. The brain dynamics toolbox for matlab. *Neurocomputing*, 315:82–88, 2018.
- Moritz Helias and David Dahmen. *Statistical Field Theory for Neural Networks*, volume 970. Springer International Publishing, 2020.
- Moritz Helias, Tom Tetzlaff, and Markus Diesmann. Echoes in correlated neural systems. *New J Phys.*, 15:023002, 2013.
- Moritz Helias, Tom Tetzlaff, and Markus Diesmann. The correlation structure of local cortical networks intrinsically results from recurrent dynamics. *PLOS Comput. Biol.*, 10(1):e1003428, 2014. URL <https://doi.org/10.1371/journal.pcbi.1003428>.
- Bernhard Hellwig. A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biological Cybernetics*, 2(82):111–121, Feb. 2000. URL <https://doi.org/10.1007/pl00007964>.
- G. Hennequin, T.P. Vogels, and W. Gerstner. Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron*, 82:1394–1406, 2014.

- John Hertz. Cross-correlations in high-conductance states of a model cortical network. *Neural Comput.*, 22:427–447, 2010.
- Michael L. Hines and Nicholas T. Carnevale. NEURON: a tool for neuroscientists. *Neuroscientist*, 7(2):123–135, Apr 2001.
- Judith A. Hirsch and Charles D. Gilbert. Synaptic physiology of horizontal connections in the cat’s visual cortex. *Journal of Neuroscience*, 11(6):1800–1809, Jun. 1991. URL <https://doi.org/10.1523/jneurosci.11-06-01800.1991>.
- A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 117: 500–544, 1952.
- Szabolcs Horvát, Răzvan Gămănuț, Mária Ercsey-Ravasz, Loïc Magrou, Bianca Gămănuț, David C. Van Essen, Andreas Burkhalter, Kenneth Knoblauch, Zoltán Toroczka, and Henry Kennedy. Spatial embedding and wiring cost constrain the functional layout of the cortical network of rodents and primates. *PLOS Biology*, 14(7):1–30, 07 2016. URL <https://doi.org/10.1371/journal.pbio.1002512>.
- Yu Hu and Haim Sompolinsky. The spectrum of covariance matrices of randomly connected recurrent neuronal networks. *bioRxiv*, 2020.
- Chengcheng Huang, Douglas A. Ruff, Ryan Pyle, Robert Rosenbaum, Marlene R. Cohen, and Brent Doiron. Circuit models of low-dimensional shared variability in cortical networks. *Neuron*, 101(2):337 – 348.e4, 2019. ISSN 0896-6273. URL <http://www.sciencedirect.com/science/article/pii/S0896627318310432>.
- D. H. Hubel and T. N. Wiesel. Receptive fields of single neurones in the cat’s striate cortex. *J Physiol*, 148:574–591, 1959.
- A. Hutt. Local excitation-lateral inhibition interaction yields oscillatory instabilities in nonlocally interacting systems involving finite propagation delay. *Physics Letters A*, 372(5):541–546, Jan. 2008. URL <https://doi.org/10.1016/j.physleta.2007.08.018>.
- Axel Hutt and Fatihcan M. Atay. Analysis of nonlocal neural fields for both general and gamma-distributed connectivities. *Physica D: Nonlinear Phenomena*, 203(1-2): 30–54, Apr. 2005. URL <https://doi.org/10.1016/j.physd.2005.03.002>.
- Axel Hutt and Fatihcan M. Atay. Spontaneous and evoked activity in extended neural populations with gamma-distributed spatial interactions and transmission delay. *Chaos, Solitons & Fractals*, 32(2):547–560, Apr. 2007. URL <https://doi.org/10.1016/j.chaos.2005.10.091>.
- Axel Hutt, Michael Bestehorn, and Thomas Wennekers. Pattern formation in intracortical neuronal fields. *Network: Computation in Neural Systems*, 14(2):351–368, Jan. 2003. URL https://doi.org/10.1088/0954-898x_14_2_310.
- Yuji Ikegaya, Takuya Sasaki, Daisuke Ishikawa, Naoko Honma, Kentaro Tao, Naoya Takahashi, Genki Minamisawa, Sakiko Ujita, and Norio Matsuki. Interpyramidal spike transmission stabilizes the sparseness of recurrent network activity. *Cereb. Cortex*, 23(2):293–304, 2013.

- E.M. Izhikevich. Which model to use for cortical spiking neurons? *IEEE Trans. Neural Netw.*, 5(15):1063–1070, Sep. 2004.
- Eugene M Izhikevich. *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*. MIT Press, 2007.
- Xiaolong Jiang, Shan Shen, Cathryn R. Cadwell, Philipp Berens, Fabian Sinz, Alexander S. Ecker, Saumil Patel, and Andreas S. Tolias. Principles of connectivity among morphologically defined cell types in adult neocortex. *Science*, 350(6264):aac9462–aac9462, Nov. 2015. URL <https://doi.org/10.1126/science.aac9462>.
- Viktor K Jirsa and Hermann Haken. Field theory of electromagnetic brain activity. *Phys. Rev. Lett.*, 77(5):960, 1996.
- Viktor K Jirsa and Hermann Haken. A derivation of a macroscopic field theory of the brain from the quasi-microscopic neural dynamics. *Physica D: Nonlinear Phenomena*, 99(4):503–526, 1997.
- Jakob Jordan, Håkon Mørk, Stine Brekke Vennemo, Dennis Terhorst, Alexander Peyser, Tammo Ippen, Rajalekshmi Deepu, Jochen Martin Eppler, Alexander van Meegen, Susanne Kunkel, Ankur Sinha, Tanguy Fardet, Sandra Diaz, Abigail Morrison, Wolfram Schenck, David Dahmen, Jari Pronold, Jonas Stapmanns, Guido Trensche, Sebastian Spreizer, Jessica Mitchell, Steffen Graber, Johanna Senk, Charl Linssen, Jan Hahne, Alexey Serenko, Daniel Naoumenko, Eric Thomson, Itaru Kitayama, Sebastian Berns, and Hans Ekkehard Plesser. *Nest 2.18.0*, Jun. 2019. URL <https://doi.org/10.5281/zenodo.2605422>.
- Eric R. Kandel, James H. Schwartz, Thomas M. Jessell, Steven A. Siegelbaum, A.J. Hudspeth, and Sarah Mack. *Principles of Neural Science*. McGraw-Hill, New York, 5 edition, 2013. ISBN 978-0-071-39011-8.
- Youngnam Kang, Takeshi Kaneko, Hitoshi Ohishi, Katsuaki Endo, and Tatsunosuke Araki. Spatiotemporally differential inhibition of pyramidal cells in the cat motor cortex. *Journal of Neurophysiology*, 71(1):280–293, Jan. 1994. URL <https://doi.org/10.1152/jn.1994.71.1.280>.
- P. Kara, P. Reinagel, and R. C. Reid. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron*, 3(27):635–646, Sep. 2000.
- B. Katz and R. Miledi. The measurement of synaptic delay, and the time course of acetylcholine release at the neuromuscular junction. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 161(985):483–495, Feb. 1965. URL <https://doi.org/10.1098/rspb.1965.0016>.
- Dennis Kätzel, Boris V. Zemelman, Christina Buetfering, Markus Wölfel, and Gero Miesenböck. The columnar and laminar organization of inhibitory connections to neocortical excitatory cells. *Nat. Neurosci.*, 14(1):100–107, Nov. 2011. URL <https://doi.org/10.1038/nn.2687>.

- Matthew T. Kaufman, Mark M. Churchland, Gopal Santhanam, Byron M. Yu, Afsheen Afshar, Stephen I. Ryu, and Krishna V. Shenoy. Roles of monkey premotor neuron classes in movement preparation and execution. *Journal of Neurophysiology*, 104(2): 799–810, Aug. 2010. URL <https://doi.org/10.1152/jn.00231.2009>.
- Matthew T. Kaufman, Mark M. Churchland, and Krishna V. Shenoy. The roles of monkey m1 neuron classes in movement preparation and execution. *Journal of Neurophysiology*, 110(4):817–825, Aug. 2013. URL <https://doi.org/10.1152/jn.00892.2011>.
- A. Keane and P. Gong. Propagating waves can explain irregular neural dynamics. *Journal of Neuroscience*, 35(4):1591–1605, Jan. 2015. URL <https://doi.org/10.1523/jneurosci.1669-14.2015>.
- Z. F. Kisvárdy and U. T. Eysel. Cellular organization of reciprocal patchy networks in layer III of cat visual cortex (area 17). *Neuroscience*, 46(2):275–286, 1992. URL [https://doi.org/10.1016/0306-4522\(92\)90050-c](https://doi.org/10.1016/0306-4522(92)90050-c).
- James C. Knight and Thomas Nowotny. GPUs outperform current HPC and neuromorphic solutions in terms of speed and energy when simulating a highly-connected cortical model. *Front. Neurosci.*, 12:1–19, Dec. 2018. URL <https://doi.org/10.3389/fnins.2018.00941>.
- Ryota Kobayashi, Shuhei Kurita, Anno Kurth, Katsunori Kitano, Kenji Mizuseki, Markus Diesmann, Barry J Richmond, and Shigeru Shinomoto. Reconstructing neuronal circuitry from parallel spike trains. *Nat. Commun.*, 10(1):1–13, 2019.
- Christof Koch. *Biophysics of Computation. Information Processing in Single Neurons*. Oxford University Press, New York, Oxford, 1999. ISBN 0-19-510491-9.
- Birgit Kriener, Moritz Helias, Ad Aertsen, and Stefan Rotter. Correlations in spiking neuronal networks with distance dependent connections. *Journal of Computational Neuroscience*, 27(2):177–200, Jul. 2009. URL <https://doi.org/10.1007/s10827-008-0135-1>.
- Birgit Kriener, Moritz Helias, Stefan Rotter, Markus Diesmann, and Gaute T. Einevoll. How pattern formation in ring networks of excitatory and inhibitory spiking neurons depends on the input current regime. *Front. Comput. Neurosci.*, 7(187):1–21, 2014. URL <https://doi.org/10.3389/fncom.2013.00187>.
- Carlo R Laing and William C Troy. Two-bump solutions of amari-type models of neuronal pattern formation. *Physica D: Nonlinear Phenomena*, 178(3-4):190–218, 2003.
- Carlo R. Laing, William C. Troy, Boris Gutkin, and Bard G. Ermentrout. Multiple bumps in a neuronal model of working memory. *SIAM J. Appl. Math.*, 63:62–97, 2002. URL <https://doi.org/10.1137/s0036139901389495>.
- Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. Numba: A llvm-based python jit compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pages 1–6, 2015.

- L. Lapicque. Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation. *J. Physiol. Pathol. Gen.*, 9:620–635, 1907.
- Moritz Layer, Johanna Senk, Simon Essink, Alexander van Meegen, Hannah Bos, and Moritz Helias. NNMT 1.0.0, Dec. 2021. URL <https://doi.org/10.5281/zenodo.5779549>.
- Moritz Layer, Johanna Senk, Simon Essink, Alexander van Meegen, Hannah Bos, and Moritz Helias. NNMT: Mean-field based analysis tools for neuronal network models. *Front. Neuroinformatics*, 16, 2022. URL <https://doi.org/10.3389/fninf.2022.835657>.
- Daeyeol Lee, Nicholas L. Port, Wolfgang Kruse, and Apostolos P. Georgopoulos. Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *Journal of Neuroscience*, 18(3):1161–1170, 1998.
- Sandrine Lefort, Christian Tamm, J.-C. Floyd Sarria, and Carl C. H. Petersen. The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex. *Neuron*, 61:301–316, 2009.
- Robert B. Levy and Alex D. Reyes. Spatial profile of excitatory and inhibitory synaptic connectivity in mouse primary auditory cortex. *Journal of Neuroscience*, 32(16):5609–5619, Apr. 2012. URL <https://doi.org/10.1523/jneurosci.5158-11.2012>.
- M. S. Lewicki. A review of methods for spike sorting: the detection and classification of neural action potentials. *Network*, 9(4):R53–78, Nov 1998.
- D.A. Lewis, D.S. Melchitzky, and G.-G. Burgos. Specificity in the functional architecture of primate prefrontal cortex. *J. Neurocytol.*, 31:265–276, 2002.
- D.T.J. Liley, P.J. Cadusch, and M.P. Dafilis. A spatially continuous mean field theory of electrocortical activity. *Network: Computation in Neural Systems*, 13(1):67–113, Jan. 2002. URL <https://doi.org/10.1080/net.13.1.67.113>.
- S. Lim and M. S. Goldman. Balanced cortical microcircuitry for spatial working memory based on corrective feedback control. *Journal of Neuroscience*, 34(20):6790–6806, May 2014. URL <https://doi.org/10.1523/jneurosci.4602-13.2014>.
- Benjamin Lindner. Interspike interval statistics of neurons driven by colored noise. *Phys. Rev. E*, 69:0229011–0229014, 2004.
- Benjamin Lindner, Brent Doiron, and André Longtin. Theory of oscillatory firing induced by spatially correlated noise and delayed inhibitory feedback. *Phys. Rev. E*, 72(6):061919, Dec. 2005. URL <https://doi.org/10.1103/physreve.72.061919>.
- Benjamin Lindner and Lutz Schimansky-Geier. Transmission of noise coded versus additive signals through a neuronal ensemble. *Phys. Rev. Lett.*, 86:2934–2937, Apr. 2001. URL <https://doi.org/10.1103/physrevlett.86.2934>.
- Ashok Litwin-Kumar and Brent Doiron. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat. Neurosci.*, 15(11):1498–1505, Sep. 2012. URL <https://doi.org/10.1038/nn.3220>.

- M. S. Livingstone and D. H. Hubel. Specificity of intrinsic connections in primate primary visual cortex. *Journal of Neuroscience*, 4(11):2830–2835, Nov. 1984. URL <https://doi.org/10.1523/jneurosci.04-11-02830.1984>.
- Yonatan Loewenstein, Annerose Kuras, and Simon Rumpel. Multiplicative dynamics underlie the emergence of the log-normal distribution of spine sizes in the neocortex in vivo. *Journal of Neuroscience*, 31(26):9481–9488, 2011.
- H. Lohmann and B. Rörig. Long-range horizontal connections between supragranular pyramidal cells in the extrastriate visual cortex of the rat. *J. Comp. Neurol.*, 344: 543–558, Jun. 1994. URL <https://doi.org/10.1002/cne.903440405>.
- Artur Luczak, Peter Barthó, and Kenneth D. Harris. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, 62(3):413–425, 2009. ISSN 0896-6273. URL <http://www.sciencedirect.com/science/article/pii/S0896627309002372>.
- Gaby Maimon and John A. Assad. Beyond poisson: Increased spike-time regularity across primate parietal cortex. *Neuron*, 62:426–440, May 2009.
- N. T. Markov, P. Misery, A. Falchier, C. Lamy, J. Vezoli, R. Quilodran, M. A. Gariel, P. Giroud, M. Ercsey-Ravasz, L. J. Pilaz, C. Huisoud, P. Barone, C. Dehay, Z. Toroczkai, D. C. Van Essen, H. Kennedy, and K. Knoblauch. Weight consistency specifies regularities of macaque cortical networks. *Cereb. Cortex*, 21(6):1254–1272, 2011.
- Francesca Mastrogiuseppe and Srdjan Ostojic. Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron*, 99(3):609–623, 2018.
- Maurizio Mattia, Matteo Biggio, Andrea Galluzzi, and Marco Storace. Dimensional reduction in networks of non-markovian spiking neurons: Equivalence of synaptic filtering and heterogeneous propagation delays. *PLOS Comput. Biol.*, 15(10): e1007404, Oct. 2019. ISSN 1553-7358. URL <https://dx.plos.org/10.1371/journal.pcbi.1007404>.
- E. M. Maynard, N. G. Hatsopoulos, C. L. Ojakangas, B. D. Acuna, J. N. Sanes, R. A. Normann, and J. P. Donoghue. Neuronal interactions improve cortical population coding of movement direction. *Journal of Neuroscience*, 19(18):8083–8093, 1999.
- Edwin M. Maynard, Craig T. Nordhausen, and Richard A. Normann. The Utah intracortical electrode array: A recording structure for potential brain-computer interfaces. *EEG Clin. Neurophysiol.*, 102(3):228–239, Mar. 1997. URL [https://doi.org/10.1016/s0013-4694\(96\)95176-0](https://doi.org/10.1016/s0013-4694(96)95176-0).
- Luca Mazzucato, Alfredo Fontanini, and Giancarlo La Camera. Stimuli reduce the dimensionality of cortical activity. *Frontiers in systems neuroscience*, 10:11, 2016.
- Colin T. McDonald and Andreas Burkhalter. Organisation of long-range inhibitory connections within rat visual cortex. *Journal of Neuroscience*, 13(2):768–781, Feb. 1993. URL <https://doi.org/10.1523/jneurosci.13-02-00768.1993>.

- Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, Ondřej Čertík, Sergey B. Kirpichev, Matthew Rocklin, AMiT Kumar, Sergiu Ivanov, Jason K. Moore, Sartaj Singh, Thilina Rathnayake, Sean Vig, Brian E. Granger, Richard P. Muller, Francesco Bonazzi, Harsh Gupta, Shivam Vats, Fredrik Johansson, Fabian Pedregosa, Matthew J. Curry, Andy R. Terrel, Štěpán Roučka, Ashutosh Saboo, Isuru Fernando, Sumith Kulal, Robert Cimrman, and Anthony Scopatz. SymPy: symbolic computing in python. *PeerJ Computer Science*, 3:e103, Jan. 2017. URL <https://doi.org/10.7717/peerj-cs.103>.
- Ernest Montbrió, Diego Pazó, and Alex Roxin. Macroscopic description for networks of spiking neurons. *Phys Rev X*, 5:021028, Jun 2015. URL <http://link.aps.org/doi/10.1103/PhysRevX.5.021028>.
- Ruben Moreno-Bote and Nestor Parga. Auto- and crosscorrelograms for the spike response of leaky integrate-and-fire neurons with slow synapses. *Phys. Rev. Lett.*, 96: 028101, 2006.
- Lyle Muller, Frédéric Chavane, John Reynolds, and Terrence J. Sejnowski. Cortical travelling waves: mechanisms and computational principles. *Nat. Rev. Neurosci.*, 19 (5):255–268, Mar. 2018. URL <https://doi.org/10.1038/nrn.2018.20>.
- Lyle Muller and Alain Destexhe. Propagating waves in thalamus, cortex and the thalamocortical system: Experiments and models. *J. Physiol. (Paris)*, 106(5-6):222–238, Sep. 2012. URL <https://doi.org/10.1016/j.jphysparis.2012.06.005>.
- Takayuki Murakoshi, Jian-Zhong Guo, and Tomomi Ichinose. Electrophysiological identification of horizontal synaptic connections in rat visual cortex in vitro. *Neuroscience Letters*, 163(2):211–214, Dec. 1993. URL [https://doi.org/10.1016/0304-3940\(93\)90385-x](https://doi.org/10.1016/0304-3940(93)90385-x).
- Brendan K. Murphy and Kenneth D. Miller. Balanced amplification: A new mechanism of selective amplification of neural activity patterns. *Neuron*, 61(4):635–648, 2009.
- Zoltan Nadasdy, T. Peter Nguyen, Ágoston Török, Jason Y. Shen, Deborah E. Briggs, Pradeep N. Modur, and Robert J. Buchanan. Context-dependent spatially periodic activity in the human entorhinal cortex. *Proceedings of the National Academy of Sciences*, 114(17), Apr. 2017. URL <https://doi.org/10.1073/pnas.1701352114>.
- Ian Nauhaus, Laura Busse, Matteo Carandini, and Dario L. Ringach. Stimulus contrast modulates functional connectivity in visual cortex. *Nat. Neurosci.*, 12:70–76, 2009. URL <https://doi.org/10.1038/nn.2232>.
- Ian Nauhaus, Laura Busse, Dario L. Ringach, and Matteo Carandini. Robustness of traveling waves in ongoing activity of visual cortex. *Journal of Neuroscience*, 32(9): 3088–3094, 2012. URL <https://doi.org/10.1523/jneurosci.5827-11.2012>.
- Martin P Nawrot, C Boucsein, V Rodriguez Molina, A Riehle, A Aertsen, and S Rotter. Measurement of variability dynamics in cortical spike trains. *Journal of Neuroscience Methods*, 169:374–390, 2008.

- Paul L Nunez. The brain wave equation: a model for the eeg. *Mathematical Biosciences*, 21(3-4):279–297, 1974.
- Ora Ohana, Hanspeter Portner, and Kevan A. C. Martin. Fast recruitment of recurrent inhibition in the cat visual cortex. *PLoS ONE*, 7(7):e40601, Jul. 2012. URL <https://doi.org/10.1371/journal.pone.0040601>.
- Michael Okun and Ilan Lampl. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat. Neurosci.*, 11(5):535–537, Mar. 2008. URL <https://doi.org/10.1038/nn.2105>.
- F. W. J. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark, editors. *NIST Handbook of Mathematical Functions*. Cambridge University Press, New York, NY, 2010.
- Srdjan Ostojic. Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons. *Nat. Neurosci.*, 17:594–600, Feb. 2014. URL <https://doi.org/10.1038/nn.3658>.
- Srdjan Ostojic and Nicolas Brunel. From spiking neuron models to linear-nonlinear models. *PLOS Comput. Biol.*, 7(1):e1001056, 2011.
- A.-M. M. Oswald, B. Doiron, J. Rinzel, and A. D. Reyes. Spatial profile and differential recruitment of GABAB modulate oscillatory activity in auditory cortex. *Journal of Neuroscience*, 29(33):10321–10334, 2009. URL <https://doi.org/10.1523/jneurosci.1703-09.2009>.
- Adam M. Packer and Rafael Yuste. Dense, unspecific connectivity of neocortical parvalbumin-positive interneurons: A canonical microcircuit for inhibition? *Journal of Neuroscience*, 31(37):13260–13271, Sep. 2011. URL <https://doi.org/10.1523/jneurosci.3131-11.2011>.
- Alessandro Pelizzola. Cluster variation method, padé approximants, and critical behavior. *Phys. Rev. E*, 49(4):R2503, 1994.
- Matthew G Perich, Sara Conti, Marion Badi, Andrew Bogaard, Beatrice Barra, Sophie Wurth, Jocelyne Bloch, Gregoire Courtine, Silvestro Micera, Marco Capogrosso, et al. Motor cortical dynamics are shaped by multiple distinct subspaces during naturalistic behavior. *BioRxiv*, 2020.
- Volker Pernice, Benjamin Staude, Stefano Cardanobile, and Stefan Rotter. How structure determines correlations in neuronal networks. *PLOS Comput. Biol.*, 7(5):e1002059, May 2011.
- Volker Pernice, Benjamin Staude, Stefano Cardanobile, and Stefan Rotter. Recurrent interactions in spiking networks with arbitrary topology. *Phys. Rev. E*, 85(3):031916, 2012.
- Bijan Pesaran, Martin Vinck, Gaute T. Einevoll, Anton Sirota, Pascal Fries, Markus Siegel, Wilson Truccolo, Charles E. Schroeder, and Ramesh Srinivasan. Investigating large-scale brain dynamics using field potential recordings: analysis and interpretation. *Nat. Neurosci.*, 21(7):903–919, Jun. 2018. URL <https://doi.org/10.1038/s41593-018-0171-8>.

- Adrien Peyrache, Nima Dehghani, Emad N. Eskandar, Joseph R. Madsen, William S. Anderson, Jacob A. Donoghue, Leigh R. Hochberg, Eric Halgren, Sydney S. Cash, and Alain Destexhe. Spatiotemporal dynamics of neocortical excitation and inhibition during human sleep. *Proceedings of the National Academy of Sciences*, 109(5): 1731–1736, Jan. 2012. URL <https://doi.org/10.1073/pnas.1109895109>.
- Adrien Peyrache and Alain Destexhe. Electrophysiological monitoring of inhibition in mammalian species, from rodents to humans. *Neurobiology of Disease*, 130:104500, Oct. 2019. URL <https://doi.org/10.1016/j.nbd.2019.104500>.
- Tobias C. Potjans and Markus Diesmann. The cell-type specific cortical microcircuit: Relating structure and activity in a full-scale spiking network model. *Cereb. Cortex*, 24(3):785–806, Dec. 2014. URL <https://doi.org/10.1093/cercor/bhs358>.
- Ryan Pyle and Robert Rosenbaum. Spatiotemporal dynamics and reliable computations in recurrent spiking neural networks. *Physical Review Letters*, 118(1), Jan. 2017. URL <https://doi.org/10.1103/physrevlett.118.018103>.
- R. Qesmi, M. Ait Babram, and M.L. Hbid. Center manifolds and normal forms for a class of retarded functional differential equations with parameter associated with fold-hopf singularity. *Applied Mathematics and Computation*, 181(1):220–246, Oct. 2006. URL <https://doi.org/10.1016/j.amc.2006.01.030>.
- R. Qesmi, M. Ait Babram, and M.L. Hbid. Symbolic computation for center manifolds and normal forms of bogdanov bifurcation in retarded functional differential equations. *Nonlinear Analysis: Theory, Methods & Applications*, 66(12):2833–2851, Jun. 2007. URL <https://doi.org/10.1016/j.na.2006.04.010>.
- Rodrigo Quian Quiroga. Spike sorting. *Current Biology*, 22(2):R45–R46, Jan. 2012. URL <https://doi.org/10.1016/j.cub.2011.11.005>.
- K. Rajan and L. F. Abbott. Eigenvalue spectra of random matrices for neural networks. *Phys. Rev. Lett.*, 97:188104, 2006.
- W. Rall. Theoretical significance of dendritic trees for neuronal input-output relations. In R. F. Reiss, editor, *Neural Theory and Modeling*. Stanford University Press, Palo Alto, 1964.
- Stefano Recanatesi, Gabriel Koch Ocker, Michael A. Buice, and Eric Shea-Brown. Dimensionality in recurrent spiking networks: Global trends in activity and local origins in connectivity. *PLOS Computational Biology*, 15(7):1–29, 07 2019. URL <https://doi.org/10.1371/journal.pcbi.1006446>.
- Michael W. Reimann, Anna-Lena Horlemann, Srikanth Ramaswamy, Eilif B. Muller, and Henry Markram. Morphological diversity strongly constrains synaptic connectivity and plasticity. *Cereb. Cortex*, 27(9):4570–4585, Jun. 2017. URL <https://doi.org/10.1093/cercor/bhx150>.
- Kimberly Reinhold, Anthony D Lien, and Massimo Scanziani. Distinct recurrent versus afferent dynamics in cortical visual processing. *Nat. Neurosci.*, 18:1789–1797, 2015.

- Alfonso Renart, Jaime De La Rocha, Peter Bartho, Liad Hollender, N'ector Parga, Alex Reyes, and Kenneth D. Harris. The asynchronous state in cortical circuits. *Science*, 327:587–590, Jan. 2010. URL <https://doi.org/10.1126/science.1179850>.
- Magnus J. E. Richardson. Firing-rate response of linear and nonlinear integrate-and-fire neurons to modulated current-based and conductance-based synaptic drive. *Phys. Rev. E*, 76(021919):1–15, August 2007.
- Magnus J. E. Richardson. Spike-train spectra and network response functions for non-linear integrate-and-fire neurons. *Biol. Cybern.*, 99:381–392, 2008.
- A. Riehle, T. Brochier, M. Nawrot, and S. Grün. Behavioral context determines network state and variability dynamics in monkey motor cortex. *Front. Neural Circuits*, 12(52), 2018. ISSN 1662-5110.
- Alexa Riehle, Sarah Wirtsohn, Sonja Grün, and Thomas Brochier. Mapping the spatio-temporal structure of motor cortical LFP and spiking activities during reach-to-grasp movements. *Frontiers in Neural Circuits*, 7:48, 2013. URL <https://doi.org/10.3389/fncir.2013.00048>.
- Juan Luis Riquelme and Julijana Gjorgjieva. Towards readable code in neuroscience. *Nat. Rev. Neurosci.*, 22(5):257–258, 2021.
- Robert Rosenbaum and Brent Doiron. Balanced networks of spiking neurons with spatially dependent recurrent connections. *Phys. Rev. X*, 4(2):021039, may 2014. URL <https://doi.org/10.1103/physrevx.4.021039>.
- Robert Rosenbaum, Mathew A. Smith, Adam Kohn, Jonathan E. Rubin, and Brent Doiron. The spatial structure of correlated neuronal variability. *Nat. Neurosci.*, 20(1):107–114, Oct. 2017. URL <https://doi.org/10.1038/nn.4433>.
- David C. Rowland, Yasser Roudi, May-Britt Moser, and Edvard I. Moser. Ten years of grid cells. *Annual Review of Neuroscience*, 39(1):19–40, Jul. 2016. URL <https://doi.org/10.1146/annurev-neuro-070815-013824>.
- Alex Roxin, Nicolas Brunel, and David Hansel. The role of delays in shaping spatio-temporal dynamics of neuronal activity in large networks. *Phys. Rev. Lett.*, 94(23):238103, Jun. 2005. URL <https://doi.org/10.1103/physrevlett.94.238103>.
- Alex Roxin, Nicolas Brunel, David Hansel, Gianluigi Mongillo, and Carl van Vreeswijk. On the distribution of firing rates in networks of cortical neurons. *Journal of Neuroscience*, 31(45):16217–16226, Nov. 2011. URL <https://doi.org/10.1523/jneurosci.1677-11.2011>.
- Alex Roxin and Ernest Montbrió. How effective delays shape oscillatory dynamics in neuronal networks. *Physica D*, 240(3):323–345, Feb. 2011. URL <https://doi.org/10.1016/j.physd.2010.09.009>.
- Michael Rule, Matthew Stoffregen, and Bard Ermentrout. A model for the origin and properties of flicker-induced geometric phosphenes. *PLOS Comput. Biol.*, 7(9):e1002158, Sep. 2011. URL <https://doi.org/10.1371/journal.pcbi.1002158>.

- P. A. Salin and D. A. Prince. Spontaneous GABAA receptor-mediated inhibitory currents in adult rat somatosensory cortex. *Journal of Neurophysiology*, 75(4):1573–1588, Apr. 1996. URL <https://doi.org/10.1152/jn.1996.75.4.1573>.
- Emilio Salinas and Terrence J. Sejnowski. Correlated neuronal activity and the flow of neural information. *Nat. Rev. Neurosci.*, 2(8):539–550, 2001a.
- Emilio Salinas and Terrence J. Sejnowski. Gain modulation in the central nervous system: Where behavior, neurophysiology, and computation meet. *Neuroscientist*, 7: 430–440, 2001b.
- Paula Sanz Leon, Stuart Knock, M. Woodman, Lia Domide, Jochen Mersmann, Anthony McIntosh, and Viktor Jirsa. The virtual brain: a simulator of primate brain network dynamics. *Front. Neuroinform.*, 7:10, 2013. URL <https://doi.org/10.3389/fninf.2013.00010>.
- Alessandro Sanzeni, Mark H Histed, and Nicolas Brunel. Response nonlinearities in networks of spiking neurons. *PLOS Comput. Biol.*, 16(9):e1008165, 2020.
- Tatsuo K. Sato, Ian Nauhaus, and Matteo Carandini. Traveling waves in visual cortex. *Neuron*, 75(2):218–229, Jul. 2012. URL <https://doi.org/10.1016/j.neuron.2012.06.029>.
- RJ Sayer, MJ Friedlander, and SJ Redman. The time course and amplitude of epsps evoked at synapses between pairs of ca3/ca1 neurons in the hippocampal slice. *Journal of Neuroscience*, 10(3):826–836, 1990.
- Maximilian Schmidt, Rembrandt Bakker, Claus C. Hilgetag, Markus Diesmann, and Sacha J. van Albada. Multi-scale account of the network structure of macaque visual cortex. *Brain Struct. Func.*, 223(3):1409–1435, Apr. 2018. ISSN 1863-2661. URL <https://doi.org/10.1007/s00429-017-1554-4>.
- Philipp Schnepel, Arvind Kumar, Mihael Zohar, Ad Aertsen, and Clemens Boucsein. Physiology and impact of horizontal connections in rat neocortex. *Cerebral Cortex*, 25(10):3818–3835, Oct. 2015. URL <https://doi.org/10.1093/cercor/bhu265>.
- Gregor Schöner. Dynamical systems approaches to cognition. *Cambridge handbook of computational cognitive modeling*, pages 101–126, 2008.
- Jannis Schuecker, Markus Diesmann, and Moritz Helias. Modulated escape from a metastable state driven by colored noise. *Phys. Rev. E*, 92:052119, Nov. 2015. URL <https://doi.org/10.1103/physreve.92.052119>.
- Jannis Schuecker, Sven Goedeke, and Moritz Helias. Optimal sequence memory in driven random networks. *Phys. Rev. X*, 8:041029, 2018. URL <https://link.aps.org/doi/10.1103/PhysRevX.8.041029>.
- Tilo Schwalger, Moritz Deger, and Wulfram Gerstner. Towards a theory of cortical columns: From spiking neurons to interacting neural populations of finite size. *PLOS Comput. Biol.*, 13(4):e1005507, Apr. 2017. URL <https://doi.org/10.1371/journal.pcbi.1005507>.

- Tilo Schwalger, Felix Droste, and Benjamin Lindner. Statistical structure of neural spiking under non-poissonian or other non-white stimulation. *Journal of computational neuroscience*, 39:29, 2015.
- TJ Sejnowski. On the stochastic dynamics of neuronal interaction. *Biological cybernetics*, 22(4):203–211, 1976.
- João D Semedo, Amin Zandvakili, Christian K Machens, M Yu Byron, and Adam Kohn. Cortical areas interact through a communication subspace. *Neuron*, 102(1):249–259, 2019.
- Johanna Senk, Espen Hagen, Sacha J. van Albada, and Markus Diesmann. Reconciliation of weak pairwise spike-train correlations and highly coherent local field potentials across space. *arXiv*, pages 1805.10235 [q-bio.NC], 2018. URL <https://arxiv.org/abs/1805.10235>.
- Johanna Senk, Karolína Korvasová, Jannis Schuecker, Espen Hagen, Tom Tetzlaff, Markus Diesmann, and Moritz Helias. Conditions for wave trains in spiking neural networks. *Phys. Rev. Res.*, 2(2), May 2020. URL <https://doi.org/10.1103/physrevresearch.2.023174>.
- Johanna Senk, Birgit Kriener, Mikael Djurfeldt, Nicole Voges, Han-Jia Jiang, Lisa Schüttler, Gabriele Gramelsberger, Markus Diesmann, Hans E. Plesser, and Sacha J. van Albada. Connectivity concepts in neuronal network modeling. *arXiv*, pages 2110.02883 [q-bio.NC], 2021.
- Michael N. Shadlen and William T. Newsome. The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *Journal of Neuroscience*, 18(10):3870–3896, May 1998. URL <https://doi.org/10.1523/jneurosci.18-10-03870.1998>.
- Jason S Sherfey, Austin E Soplata, Salva Ardid, Erik A Roberts, David A Stanley, Benjamin R Pittman-Polletta, and Nancy J Kopell. Dynasim: a matlab toolbox for neural modeling and simulation. *Front. Neuroinformatics*, 12:10, 2018.
- Arnold J.F. Siegert. On the first passage time probability problem. *Physical Review*, 81(4):617–623, 1951.
- Gordon B Smith, Bettina Hein, David E Whitney, David Fitzpatrick, and Matthias Kaschube. Distributed network interactions and their emergence in developing neocortex. *Nat. Neurosci.*, 21(11):1600, 2018.
- Matthew A. Smith and Adam Kohn. Spatial and temporal scales of neuronal correlation in primary visual cortex. *Journal of Neuroscience*, 28(48):12591–12603, Nov. 2008. URL <https://doi.org/10.1523/jneurosci.2929-08.2008>.
- William R. Softky and Christof Koch. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *Journal of Neuroscience*, 13(1):334–350, 1993. URL <https://doi.org/10.1523/jneurosci.13-01-00334.1993>.

- H. Sompolinsky, A. Crisanti, and H. J. Sommers. Chaos in random neural networks. *Phys. Rev. Lett.*, 61:259–262, Jul 1988. URL <http://link.aps.org/doi/10.1103/PhysRevLett.61.259>.
- Sen Song, Per Sjöström, Markus Reigl, Sacha Nelson, and Dmitri Chklovskii. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol.*, 3(3):e68, 2005.
- L. Spek, K. Dijkstra, S.A. van Gils, and M. Polner. Dynamics of delayed neural field models in two-dimensional spatial domains. *Journal of Differential Equations*, 317:439–473, Apr. 2022. URL <https://doi.org/10.1016/j.jde.2022.02.002>.
- Len Spek, Yuri A. Kuznetsov, and Stephan A. van Gils. Neural field models with transmission delays and diffusion. *The Journal of Mathematical Neuroscience*, 10(1), Dec. 2020. URL <https://doi.org/10.1186/s13408-020-00098-5>.
- Sebastian Spreizer, Ad Aertsen, and Arvind Kumar. From space to time: Spatial inhomogeneities lead to the emergence of spatiotemporal sequences in spiking neuronal networks. *PLOS Comput. Biol.*, 15(10):e1007432, 2019.
- Richard B. Stein. Some models of neuronal variability. *Biophys. J.*, 7(1):37–68, Jan. 1967. URL [https://doi.org/10.1016/s0006-3495\(67\)86574-3](https://doi.org/10.1016/s0006-3495(67)86574-3).
- A. Stepanyants, J. Hirsch, L. M. Martinez, Z. F. Kisvárdy, A. S. Ferecsko, and D. B. Chklovskii. Local potential connectivity in cat primary visual cortex. *Cereb. Cortex*, 18(1):13–28, 2008.
- Armen Stepanyants, Judith A. Hirsch, Luis M. Martinez, Zoltán F. Kisvárdy, Alex S. Ferecskó, and Dmitri B. Chklovskii. Local potential connectivity in cat primary visual cortex. *Cerebral Cortex*, 18(1):13–28, Apr. 2007. URL <https://doi.org/10.1093/cercor/bhm027>.
- Armen Stepanyants, Luis M. Martinez, Alex S. Ferecskó, and Zoltán F. Kisvárdy. The fractions of short- and long-range connections in the visual cortex. *Proceedings of the National Academy of Sciences*, 106(9):3555–3560, Feb. 2009. URL <https://doi.org/10.1073/pnas.0810390106>.
- Armen Stepanyants, Gábor Tamás, and Dmitri B Chklovskii. Class-specific features of neuronal wiring. *Neuron*, 43(2):251–259, Jul. 2004. URL <https://doi.org/10.1016/j.neuron.2004.06.013>.
- Charles F. Stevens and Anthony M. Zador. Input synchrony and the irregular firing of cortical neurons. *Nat. Neurosci.*, 1(3):210–217, July 1998.
- JC Stiller and G Radons. Dynamics of nonlinear oscillators with random interactions. *Phys. Rev. E*, 58(2):1789, 1998.
- Marcel Stimberg, Romain Brette, and Dan FM Goodman. Brian 2, an intuitive and efficient neural simulator. *eLife*, 8, Aug. 2019. URL <https://doi.org/10.7554/elife.47314>.

- Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Matteo Carandini, and Kenneth D Harris. High-dimensional geometry of population responses in visual cortex. *Nature*, 571(7765):361–365, 2019a.
- Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D Harris. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437):eaav7893, 2019b.
- Harvey A. Swadlow and Stephen G. Waxman. Axonal conduction delays. *Scholarpedia*, 7(6):1451, 2012. URL <https://doi.org/10.4249/scholarpedia.1451>.
- Kazutaka Takahashi, Sanggyun Kim, Todd P. Coleman, Kevin A. Brown, Aaron J. Suminski, Matthew D. Best, and Nicholas G. Hatsopoulos. Large-scale spatiotemporal spike patterning consistent with wave propagation in motor cortex. *Nature Communications*, 6(7169), May 2015. URL <https://doi.org/10.1038/ncomms8169>.
- Hisashi Tanigawa, QuanXin Wang, and Ichiro Fujita. Organization of horizontal axons in the inferior temporal cortex and primary visual cortex of the macaque monkey. *Cereb. Cortex*, 15(12):1887–1899, Dec. 2005. URL <https://doi.org/10.1093/cercor/bhi067>.
- Peter Tass. Cortical pattern formation during visual hallucinations. *Journal of Biological Physics*, 21(3):177–210, 1995. URL <https://doi.org/10.1007/bf00712345>.
- Peter Tass. Oscillatory cortical activity during visual hallucinations. *Journal of Biological Physics*, 23(1):21–66, 1997. URL <https://doi.org/10.1023/a:1004990707739>.
- Tom Tetzlaff, Moritz Helias, Gaute T. Einevoll, and Markus Diesmann. Decorrelation of neural-network activity by inhibitory feedback. *PLOS Comput. Biol.*, 8(8):e1002596, aug 2012. URL <https://doi.org/10.1371/journal.pcbi.1002596>.
- D.J. Tolhurst, J.A. Movshon, and A.F. Dean. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.*, 23(8):775–785, Jan. 1983. URL [https://doi.org/10.1016/0042-6989\(83\)90200-6](https://doi.org/10.1016/0042-6989(83)90200-6).
- Emiliano Torre, Pietro Quaglio, Michael Denker, Thomas Brochier, Alexa Riehle, and Sonja Grün. Synchronous spike patterns in macaque motor cortex during an instructed-delay reach-to-grasp task. *J. Neurosci.*, 36(32):8329–8340, 2016.
- R. G. Townsend, S. S. Solomon, S. C. Chen, A. N. J. Pietersen, P. R. Martin, S. G. Solomon, and P. Gong. Emergence of complex wave patterns in primate cerebral cortex. *Journal of Neuroscience*, 35(11):4657–4662, Mar. 2015. URL <https://doi.org/10.1523/jneurosci.4509-14.2015>.
- Taro Toyozumi and Larry F Abbott. Beyond the edge of chaos: Amplification and temporal integration by recurrent networks in the chaotic regime. *Phys. Rev. E*, 84(5):051908, 2011.
- J Trousdale, Y Hu, E Shea-Brown, and K Josic. Impact of network structure and cellular response on spike time correlations. *PLOS Comput. Biol.*, 8(3):e1002408, 2012.

- Henry C. Tuckwell. *Introduction to Theoretical Neurobiology*, volume 1, chapter 3, The Lapique model of the nerve cell, pages 85–123. Cambridge University Press, Cambridge, 1988a. ISBN 0-521-35096-4.
- Henry C. Tuckwell. *Introduction to Theoretical Neurobiology*, volume 2. Cambridge University Press, Cambridge, 1988b.
- M. Usher, M. Stemmler, and Z. Olami. Dynamic pattern formation leads to $1/f$ noise in neural populations. *Phys. Rev. Lett.*, 74(2):326–329, 1995.
- Sacha J. van Albada, Andrew G. Rowley, Johanna Senk, Michael Hopkins, Maximilian Schmidt, Alan B. Stokes, David R. Lester, Markus Diesmann, and Steve B. Furber. 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, May 2018. URL <https://doi.org/10.3389/fnins.2018.00291>.
- S. A. van Gils, S. G. Janssens, Yu. A. Kuznetsov, and S. Visser. On local bifurcations in neural field models with transmission delays. *J. Math. Biol.*, 66(4-5):837–887, Nov. 2012. URL <https://doi.org/10.1007/s00285-012-0598-6>.
- Alexander van Meegen and Benjamin Lindner. Self-consistent correlations of randomly coupled rotators in the asynchronous state. *Phys. Rev. Lett.*, 121(25):258302, 2018.
- C. van Vreeswijk and H. Sompolinsky. Chaotic balanced state in a model of cortical circuits. *Neural computation*, 10(6):1321–1371, Aug. 1998. URL <https://doi.org/10.1162/089976698300017214>.
- Carl van Vreeswijk and Farzad Farkhooi. Fredholm theory for the mean first-passage time of integrate-and-fire oscillators with colored noise input. *Phys. Rev. E*, 100(6):060402, 2019.
- Carl van Vreeswijk and Haim Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274:1724–1726, December, 6 1996. URL <https://doi.org/10.1126/science.274.5293.1724>.
- Romain Veltz. Interplay between synaptic delays and propagation delays in neural field equations. *SIAM Journal on Applied Dynamical Systems*, 12(3):1566–1612, Jan. 2013. URL <https://doi.org/10.1137/120889253>.
- Romain Veltz and Olivier Faugeras. Stability of the stationary solutions of neural field equations with propagation delays. *The Journal of Mathematical Neuroscience*, 1(1):1, 2011. URL <https://doi.org/10.1186/2190-8567-1-1>.
- Romain Veltz and Olivier Faugeras. A center manifold result for delayed neural fields equations. *SIAM Journal on Mathematical Analysis*, 45(3):1527–1562, Jan. 2013. URL <https://doi.org/10.1137/110856162>.
- N. A. Venkov, S. Coombes, and P. C. Matthews. Dynamic instabilities in scalar neural field equations with space-dependent delays. *Physica D: Nonlinear Phenomena*, 232(1):1–15, Aug. 2007. URL <https://doi.org/10.1016/j.physd.2007.04.011>.

- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods*, 17:261–272, 2020.
- S. Visser, R. Nicks, O. Faugeras, and S. Coombes. Standing and travelling waves in a spherical brain model: The nunez model revisited. *Physica D: Nonlinear Phenomena*, 349:27–45, Jun. 2017. URL <https://doi.org/10.1016/j.physd.2017.02.017>.
- Sid Visser, Hil GE Meijer, Michel JAM van Putten, and Stephan A van Gils. Analysis of stability and bifurcations of fixed points and periodic solutions of a lumped model of neocortex with two delays. *The Journal of Mathematical Neuroscience*, 2(1):8, 2012. URL <https://doi.org/10.1186/2190-8567-2-8>.
- Nicole Voges, Almut Schüz, Ad Aertsen, and Stefan Rotter. A modeler’s view on the spatial structure of intrinsic horizontal connectivity in the neocortex. *Prog. Neurobiol.*, 92(3):277–292, Nov. 2010. URL <https://doi.org/10.1016/j.pneurobio.2010.05.001>.
- Nobuhiko Wagatsuma, Tobias C. Potjans, Markus Diesmann, and Tomoki Fukai. Layer-dependent attentional processing by top-down signals in a visual cortical microcircuit model. *Front. Comput. Neurosci.*, 5:31, 2011. URL <https://doi.org/10.3389/fncom.2011.00031>.
- Xiao-Jing Wang. Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological Reviews*, 90(3):1195–1268, 2010. URL <https://doi.org/10.1152/physrev.00035.2008>.
- Miles A. Whittington, Mark O. Cunningham, Fiona E.N. LeBeau, Claudia Racca, and Roger D. Traub. Multiple origins of the cortical gamma rhythm. *Developmental Neurobiology*, 71(1):92–106, Dec. 2010. URL <https://doi.org/10.1002/dneu.20814>.
- H. R. Wilson and J. D. Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys. J.*, 12(1):1–24, Jan 1972. URL [https://doi.org/10.1016/s0006-3495\(72\)86068-5](https://doi.org/10.1016/s0006-3495(72)86068-5).
- H. R. Wilson and J. D. Cowan. A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13(2):55–80, Sep. 1973. URL <https://doi.org/10.1007/bf00288786>.
- Jian-Young Wu, Xiaoying Huang, and Chuan Zhang. Propagating waves of activity in the neocortex: What they are, what they do. *The Neuroscientist*, 14(5):487–502, 2008. URL <https://doi.org/10.1177/1073858408317066>.

- John Wyller, Patrick Blomquist, and Gaute T. Einevoll. Turing instability and pattern formation in a two-population neuronal network model. *Physica D: Nonlinear Phenomena*, 225(1):75–93, Jan. 2007. URL <https://doi.org/10.1016/j.physd.2006.10.004>.
- Zixiu Xiang, John R. Huguenard, and David A. Prince. GABAAreceptor-mediated currents in interneurons and pyramidal cells of rat visual cortex. *Journal of Physiology*, 506(3):715–730, Feb. 1998. URL <https://doi.org/10.1111/j.1469-7793.1998.715bv.x>.
- Byron M. Yu, John P. Cunningham, Gopal Santhanam, Stephen I. Ryu, Krishna V. Shenoy, and Maneesh Sahani. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, 102(1): 614–635, jul 2009. URL <https://doi.org/10.1152/jn.90941.2008>.

FUNDING

This work has received funding from the European Union's Q17 Horizon 2020 Framework Programme for Research and Innovation under Specific Grant Agreement No. 720270 (HBP SGA1), 785907 (HBP SGA2), and 945539 (HBP SGA3), ANR grant GRASP, was partially supported by HGF young investigator's group VH-NG-1028, and has been partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)–368482240/GRK2416. This research was supported by the Joint Lab Supercomputing and Modeling for the Human Brain.

ACKNOWLEDGEMENTS

I am deeply grateful to all people who supported me during the last four years, which included two years of the pandemic that probably were difficult for everyone. I want to express special gratitude towards my supervisors: Moritz Helias, from whom I definitely learned the most, for many helpful discussions and his inexhaustible knowledge, patience, and optimism; Sonja Grün for providing me with invaluable insight into the fascinating world of data analysis and the plethora of issues related to reproducibility; and Björn Kampa for allowing me to observe actual experiments in his lab, which may be necessary from time to time to bring a physicist who likes to float in the spheres of abstract equations and theories back into the real world. Similarly, I would like to thank the RTG2416 for giving me the opportunity to meet with a wide range of people working on really interesting problems, as well as for delicious dinners. I am particularly grateful to Johanna Senk and David Dahmen, my excellent direct supervisors, for guiding me through all stages of our projects with a lot of wit and empathy. My colleagues and friends at the INM-6 deserve special recognition for interesting, sometimes highly questionable, lunch time conversations. And I would like to thank Maren Maus for many amusing hours spent as office neighbors.

Outside of the world of my work, I would like to thank my entire family for just being the best. Thanks to Heike Denninger, in particular, for extended phone calls when none of us could find the drive to work; to my brother, Tom Layer, for organizing numerous hilarious virtual game sessions during the pandemic; and to my mother, Meike Unruh, for always providing me the recipes I need. Special thanks go to my sister, Maxi Layer, for sharing seemingly endless hours of virtual synchronized-music-and-work-sessions. Finally, I want to express my deepest gratitude to my wife, Natalie Layer, who has to put up with me every day. Without her we would miss out on so many funny noises, and life would never be as beautiful.

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 L^AT_EX and L_YX at

<https://bitbucket.org/amiede/classicthesis/>

Band / Volume 71

Investigating the Interaction between π -Conjugated Organic Molecules and Metal Surfaces with Photoemission Tomography

X. Yang (2021), xviii, 173 pp

ISBN: 978-3-95806-584-0

Band / Volume 72

Three-Dimensional Polymeric Topographies for Neural Interfaces

F. Milos (2021), 133 pp

ISBN: 978-3-95806-586-4

Band / Volume 73

Development, characterization, and application of compliant intracortical implants

K. Srikantharajah (2021), xiv, 155, xv-xvii pp

ISBN: 978-3-95806-587-1

Band / Volume 74

Modelling, implementation and characterization of a Bias-DAC in CMOS as a building block for scalable cryogenic control electronics for future quantum computers

P. N. Vliex (2021), xiv, 107, xv-xxviii pp

ISBN: 978-3-95806-588-8

Band / Volume 75

Development of Electrochemical Aptasensors for the Highly Sensitive, Selective, and Discriminatory Detection of Malaria Biomarkers

G. Figueroa Miranda (2021), XI, 135 pp

ISBN: 978-3-95806-589-5

Band / Volume 76

Nanostraw- Nanocavity MEAs as a new tool for long-term and high sensitive recording of neuronal signals

P. Shokoohimehr (2021), xi, 136 pp

ISBN: 978-3-95806-593-2

Band / Volume 77

Surface plasmon-enhanced molecular switching for optoelectronic applications

B. Lenyk (2021), x, 129 pp

ISBN: 978-3-95806-595-6

Band / Volume 78

Engineering neuronal networks in vitro: From single cells to population connectivity

I. Tihaa (2021), viii, 242 pp

ISBN: 978-3-95806-597-0

Band / Volume 79

Spectromicroscopic investigation of local redox processes in resistive switching transition metal oxides

T. Heisig (2022), vi, 186 pp

ISBN: 978-3-95806-609-0

Band / Volume 80

Integrated Control Electronics for Qubits at Ultra Low Temperature

D. Nielinger (2022), xviii, 94, xix-xxvi

ISBN: 978-3-95806-631-1

Band / Volume 81

Higher-order correlation analysis in massively parallel recordings in behaving monkey

A. Stella (2022), xiv, 184 pp

ISBN: 978-3-95806-640-3

Band / Volume 82

Denoising with Quantum Machine Learning

J. Pazem (2022), 106 pp

ISBN: 978-3-95806-641-0

Band / Volume 83

Hybrid hydrogels promote physiological functionality of long-term cultured primary neuronal cells in vitro

C. Meeßen (2022), x, 120 pp

ISBN: 978-3-95806-643-4

Band / Volume 84

Surface states and Fermi-level pinning on non-polar binary and ternary (Al,Ga)N surfaces

L. Freter (2022), 137 pp

ISBN: 978-3-95806-644-1

Band / Volume 85

Dynamical and statistical structure of spatially organized neuronal networks

M. Layer (2022), xiii, 165 pp

ISBN: 978-3-95806-651-9

Weitere **Schriften des Verlags im Forschungszentrum Jülich** unter
<http://www.zb1.fz-juelich.de/verlagextern1/index.asp>

Information
Band / Volume 85
ISBN 978-3-95806-651-9