Analysis and quantitative comparison of neural network dynamics on a neuron-wise and population level

Robin Gutzen

Information
Band / Volume 102
ISBN 978-3-95806-738-7

Mitglied der Helmholtz-Gemeinschaft

JÜLICH
Forschungszentrum

Schriften des Forschungszentrums Jülich
Reihe Information / Information Band / Volume 102

Forschungszentrum Jülich GmbH
Institut für Neurowissenschaften und Medizin (INM)
Computational and Systems Neuroscience (INM-6)

Analysis and quantitative comparison of neural network dynamics on a neuron-wise and population level

Robin Gutzen

Schriften des Forschungszentrums Jülich
Reihe Information / Information Band / Volume 102

ISSN 1866-1777 ISBN 978-3-95806-738-7

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 2024

Schriften des Forschungszentrums Jülich
Reihe Information / Information, Band / Volume 102

D 82 (Diss. RWTH Aachen University, 2023)

ISSN 1866-1777
ISBN 978-3-95806-738-7

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

The brain is the most elevated of all the viscera, and the nearest to the roof of the head ... The senses hold this organ as their citadel; it is in this that are centered all the veins which spring from the heart; it is here that they terminate; this is the very culminating point of all, the regulator of the understanding.

— Pliny the Elder (77 AD)

Nobody understands how the brain works, and everybody knows that.

— Anirban Bandyopadhyay (2022)

SUMMARY

Our goal is to better understand the working mechanisms of biological neural systems. To this end, describing neural systems as networks provides a powerful and widely-used analysis approach. The network syntax of interacting nodes exhibiting joint dynamics facilitates the quantitative characterizations of neural systems across scales. Moreover, this approach enables us to construct systematic comparisons of neural network descriptions across domains.

We aim to identify characterizations of neural network activity that reflect the underlying connectivity and relate to the network's ability to process information. In this context, we explore characteristic measures from experimental and simulated data sources. Concretely, we look at cortical activity data from mice and monkeys, from different recording techniques like implanted electrode arrays, laminar probes, ECoG, and calcium imaging, and further from simulations of stochastic processes, spiking, and mean-field network models. We investigate activity measures of different complexity, including measures on the level of individual neurons, higher-order measures of coordinated spiking activity, and population-level field potential measures describing spatial wave patterns. Such activity characterizations always represent an abstraction, and the right level of detail depends on the data type and the question of interest. For a given context, the appropriate abstraction level allows us to integrate and compare data and models from heterogeneous sources.

Evaluating the similarity between such different network descriptions is a common demand in computational neuroscience. Extending the concept of validation, we formalize and apply cross-domain comparisons in model vs. experiment, model vs. model, and experiment vs. experiment scenarios. In this framework, we further evaluate and extend existing statistical testing approaches and look at reproducibility, sources of variability, and technical limitations.

Through our exploration of network activity characterizations and their comparability, we evaluate the relationship between network connectivity, activity, and function. Concretely, over the course of five research projects, we implement and demonstrate systematic approaches to validate model simulators, statistically evaluate network organization, infer network connectivity from activity data, combine data sources of wave activity, and relate wave activity to external influences and behavior. With a focus on open and collaborative science practices, we implement our methodologies as reusable open-source tools while building upon existing open-source tools and standards.

ZUSAMMENFASSUNG

Das Ziel dieser Arbeit ist es, die Funktionsmechanismen biologischer neuronaler Systeme besser zu verstehen. Die Beschreibung neuronaler Systeme als Netzwerke bietet dafür einen idealen Analyseansatz. Die Netzwerk Syntax interagierender Knoten, die eine gemeinsame Dynamik aufweisen, erleichtert die quantitative Charakterisierung neuronaler Systeme auf verschiedenen Beschreibungsebenen. Darüber hinaus ermöglicht dies einen systematischen Vergleich von Netzwerken über verschiedene Anwendungsbereiche hinweg.

Wir streben an Charakterisierungen der Aktivität neuronaler Netze zu identifizieren, die die zugrunde liegende Konnektivität widerspiegeln und sich auf die Informationsverarbeitungsfähigkeit des Netzes beziehen. In diesem Zusammenhang untersuchen wir charakteristische Maße aus experimentellen und simulierten Datenquellen. Konkret betrachten wir kortikale Aktivitätsdaten von Mäusen und Affen, generiert mit verschiedene Aufzeichnungstechniken wie implantierten Elektrodenarrays, laminaren Sonden, ECoG, und Kalzium-Imaging, sowie Daten von Simulationen stochastischer Prozesse, Spiking und Mean-Field-Netzwerkmodellen. Wir untersuchen verschieden komplexe Aktivitätsmaße, einschließlich Maße auf der Einzelneurone Ebene, Maße höherer Ordnung der koordinierten Spiking-Aktivität, und Feldpotentialmaße auf Populationsebene. Solche Aktivitätscharakterisierungen stellen immer eine Abstraktion dar, wobei die richtige Abstraktionsebene es uns ermöglicht, Daten und Modelle aus heterogenen Quellen zu integrieren und zu vergleichen.

Die Bewertung der Ähnlichkeit zwischen solchen unterschiedlichen Netzwerkbeschreibungen ist eine gängige Herausforderung im Bereich der Computational Neuroscience. Indem wir das Konzept der Validierung erweitern, können wir bereichsübergreifende Vergleiche in den Szenarien Modell vs. Experiment, Modell vs. Modell und Experiment vs. Experiment formalisieren und anwenden.

Die Erforschung der Charakterisierung von Netzwerkaktivitäten und ihrer Vergleichbarkeit ermöglicht uns, die Beziehung zwischen Netzwerkkonnektivität, Aktivität, und Funktion zu bewerten. Dies stellen wir konkret anhand von fünf Forschungsprojekten dar: zur systematischen Ansätzen zur Validierung von Modellsimulatoren, statistischen Auswertung der Netzwerkorganisation, Ableitung von Netzwerkkonnektivität aus Aktivitätsdaten, Kombination von Datenquellen für Wellenaktivität, und Verknüpfung von Wellenaktivität mit externen Einflüssen und Verhalten. Im Sinne kollaborativer Wissenschaftspraktiken implementieren wir unsere Methoden als frei nutzbare Open-Source Tools und bauen auf bestehenden Open-Source Tools und Standards auf.

PUBLICATIONS

PAPER

- Trensch, Guido, Robin Gutzen, Inga Blundell, Michael Denker, and Abigail Morrison (2018). "Rigorous Neural Network Simulations: A Model Substantiation Methodology for Increasing the Correctness of Simulation Results in the Absence of Experimental Validation Data." In: *Frontiers in Neuroinformatics* 12, p. 81. DOI: 10.3389/fninf.2018.00081
- Gutzen, Robin, Michael von Papen, Guido Trensch, Pietro Quaglio, Sonja Grün, and Michael Denker (2018b). "Reproducible Neural Network Simulations: Statistical Methods for Model Validation on the Level of Network Activity Data." In: *Frontiers in Neuroinformatics* 12, p. 90. DOI: 10.3389/fninf.2018.00090
- Gutzen, Robin, Sonja Grün, and Michael Denker (2022). "Evaluating the Statistical Similarity of Neural Network Activity and Connectivity via Eigenvector Angles." In: *Biosystems* 223, p. 104813. DOI: 10.1016/j.biosystems.2022.104813
- Capone, Cristiano, Chiara De Luca, Giulia De Bonis, Robin Gutzen, Irene Bernava, Elena Pastorelli, Francesco Simula, Cosimo Lupo, Leonardo Tonielli, Francesco Resta, Anna Letizia Allegra Mascaro, Francesco Pavone, Michael Denker, and Pier Stanislao Paolucci (2023). "Simulations Approaching Data: Cortical Slow Waves in Inferred Models of the Whole Hemisphere of Mouse." In: *Communications Biology* 6.1, pp. 1–14. DOI: 10.1038/s42003-023-04580-0
- Gutzen, Robin, Giulia De Bonis, Chiara De Luca, Elena Pastorelli, Cristiano Capone, Anna Letizia Allegra Mascaro, Francesco Resta, Arnau Manasanch, Francesco Saverio Pavone, Maria V. Sanchez-Vives, Maurizio Mattia, Sonja Grün, Pier Stanislao Paolucci, and Michael Denker (2024). "A Modular and Adaptable Analysis Pipeline to Compare Slow Cerebral Rhythms across Heterogeneous Datasets." In: *Cell Reports Methods* 4.1, p. 100681. DOI: 10.1016/j.crmeth.2023.100681
- Morales-Gregorio, Aitor, Robin Gutzen, Paulina Dabrowska, Alper Yegenoglu, Sandra Diaz-Pier, Sarah Palmis, Alexendre René, Panos Sapountzis, Markus Diesmann, Sonja Grün, Johanna Senk, Georgia Gregoriou, Bjørg E. Kilavik, and Sacha J Van Albada (2022). "Activity-Driven Microconnectome Estimation of Macaque Visuomotor Cortices." In: *Preparation*

SOFTWARE

- Gutzen, Robin, Aitor Morales-Gregorio, Michael Von Papen, and Michael Denker (2018a). *NetworkUnit*. RRID:SCR_016543
- Gutzen, Robin, Chiara De Luca, and Giulia De Bonis (2023). *Collaborative Brain Wave Analyisis Pipeline (Cobrawap)*. RRID:SCR_022966. DOI: 10.5281/zenodo.10198748

CONTENTS

Part I

INTRODUCTION

1

BACKGROUND

1.1 SCALES OF NEURAL ORGANIZATION

Our brain has the relatively fixed and modest size of a clenched fist, with the additional appendices of the central nervous system reaching out through our body. This neural system is enormously complex and intricately structured (Koch and Laurent, 1999), representing aspects over many orders of magnitude in its neural organization, from molecular interactions to environmental influences. Although we use a mammal-centric view here, many of the discussed aspects also apply to other neural system such as typical model organisms like C.Elegans (Izquierdo and Beer, 2013) and, in some regards even artificial neural systems like neuromorphic architectures (Bartolozzi et al., 2016). All matter is built up from fundamental physical building blocks, and so are neural systems. However, the specificity to neural systems begins in most descriptions on the level of molecules and their interactions (except some hypotheses about the theory of mind, e.g., Penrose (1999)).

NANOMETERS On the scale of **nanometers**, there are a wide variety of *molecular processes* and interactions happening that are relevant to the brain and the nervous system. Genes are expressed, proteins are built acting in and in-between the brain's cells, e.g., as ion-channels, as catalysts for the biosynthesis of relevant molecules such as neurotransmitters, and as fluorescent markers in transgenic subjects (Barth et al., 2004). In contrast, the suppression of specific genes via knock-out can be used to model neurological disorders (Hall et al., 2009). Further molecular processes include, for example, the aggregation and folding of proteins (e.g., beta-amyloids which play a role in Alzheimer's disease (Sticht et al., 1995)), molecular transport (e.g., via diffusion, electrical and chemical gradients, and carrier proteins), and the binding to receptors (e.g., of neurotransmitters, hormones, caffeine, or certain trace elements) (Kandel et al., 2001). Many crucial processes in the nervous system take place around the *cell membranes* that act as a mechanical, chemical, and the electrical barrier between the inside and the outside of the cells. Therefore there is typically a concentration differential for many free molecules (including sodium, potassium, and calcium), resulting in a resting potential of about -70 mV across the membrane of a nerve cell, i.e., a neuron (Rubinstein, 1990). Facilitated by trans-membrane ion-pumps and channels, sufficiently depolarizing perturbations of this (Nernst) equilibrium trigger a stereotypical

response of a rapid depolarization and repolarization (within $\sim$ 2 ms) of the membrane, called an "action potential" (or "spike" if measured extracellularly), followed by a hyperpolarization before returning to the resting potential (after another $\sim$ 2 ms) (Hodgkin and Huxley, 1952).

MICROMETERS The action potentials are considered a fundamental unit in the information transport and processing in the nervous system. Action potentials propagate along the membranes of the *neuronal compartments*. Typically on a scale of **micrometers**, there are dendritic branches where the input of other neurons is received, and action potentials are evoked. The action potentials in the dendrites can interact and are aggregated in the soma where eventually, when a threshold potential is surpassed at the axon hillock, an outgoing action potential is created to propagate along the axon to be transmitted to other neighboring neurons via shared *synapses*, i.e., the neuron "fires" (Kandel et al., 2001). At the synapses, action potentials trigger the release of neurotransmitters from the axon of the pre-synaptic neuron, which are received by the dendrites of post-synaptic neurons, where then action potentials can again be elicited. The efficiency of this transmission, also referred to as synaptic weight, is dependent on multiple aspects such as the number of released neurotransmitters, the number of synaptic connections, the number of receptors, and the availability of intracellular ions (e.g., Ca^{+}). Furthermore it, and is subject to ongoing change as a function of the neuron's activities, i.e., "plasticity". While the soma of neurons has a diameter between a few and a hundred micrometers, the length of the neuron along its dendrites and axon can stretch up to meters (Niebur, 2008). The exact dimensions, morphology, electrophysiological properties, and transcriptomic profile of neurons depend strongly on the neuronal cell type and the species. Furthermore, there is also a variety of synaptic types, separating most broadly into electrical and chemical synapses and chemical synapses further into excitatory and inhibitory synapses. Transmissions via excitatory synapses stimulate the generation of a post-synaptic action potential, whereas transmissions via inhibitory synapses suppress it. According to Dale's law, each neuron produces the same types of neurotransmitters at all of its post-synaptic connections and can therefore be labeled as excitatory or inhibitory (Kandel, 1968). There is a diverse zoo of neuronal cell types that can also differ across species and brain regions (Tasic et al., 2016; Tasic et al., 2018; Boldog et al., 2018). This neuron zoo is further complemented by non-neural glial cells.

MILLIMETERS Neurons are highly interconnected, with each one typically having thousands of synapses, and are organized in various group structures, like *assemblies* (Harris, 2005), with their neighbors

on a scale of **millimeters**. This interconnection leads to a coordination of the neurons' activity and in many cases, a synchronization (Singer, 1999; Ben-Shaul et al., 2001; Sakurai and Takahashi, 2006). Coordinated firing activity is particularly relevant because synchronized input is much more likely to trigger a response in the post-synaptic neuron than asynchronous input (Abeles, 1982; Palm, 1990; König et al., 1996). The neuronal activity is not solely contained inside the cells. The depolarization of the cell membranes and the synaptic transmissions of the electric potentials within a local group of neurons also contribute to a field potential that propagates in the extracellular medium, the local field potential (LFP). Therefore, the LFP represents a summed population activity of a group of nearby neurons, for which the radius of influence depends on the brain area, species, and spectral frequency (Pesaran et al., 2018). The relationship between the LFP and the firing activity of neurons is not unidirectional as the LFP can influence a neuron's membrane potential and, therefore, its probability and timing to fire, in a process called "ephaptic coupling" (Aur and Jog, 2010).

CENTIMETERS Neuronal structures in the brain further organize in structurally and functionally differing areas on the scale of millimeters to **centimeter**, depending on the species. Besides many structurally distinct "deep" brain areas (such as the cerebellum, hippocampus, pons, or thalamus), the cerebral cortex represents an increasing fraction of the brain in higher mammals (Hofman, 1988). The cortex has a relatively consistent, layered structure (Douglas et al., 1989) but still separates into specific *cortical areas* that can be associated with specific functions based on their cytoarchitecture (Brodmann, 1909; Amunts and Zilles, 2015). The brain areas, and in particular the cortical areas, are further organized in hierarchical pathways structuring the flow of processing and integration of information. This complex organization of neural structures building up the *whole brain* further gives rise to characteristic oscillatory behaviors of the field potential, incorporating many different brain rhythms across a frequency range from < 0.1 Hz up to 600 Hz (Katz and Cracco, 1971). Different areas in the brain can differ in their spectral signature due to a multitude of mechanisms producing the various oscillatory bands (Buzsaki, 2006). Furthermore, local oscillations interact, resulting in traveling waves in the electrical potential field that can cover up to an entire hemisphere or the whole brain (Muller et al., 2016; Hindriks et al., 2014; Burkitt et al., 2000).

METERS The brain is not an isolated structure. It connects to the rest of the *central nervous system* that stretches out through the *whole body*, in some mammals for many **meters**. The body is an indispensable aspect in the organization of neural systems, as it represents the interface to perceive and interact with the environment ("embodiment"). Specialized neural structures process the input from the sensory or-

gans and integrate it with other inputs and existing dynamics. Further, the nervous system is specialized to command the body's activity for homeostasis and behavior (Friston, 2010).

KILOMETERS The influences on neural organization and function also extend beyond the body, on the scale of **kilometers**. Brain organization incorporates the laws and statistics of the environment where it develops, which is particularly prevalent for the visual cortex (Switkes et al., 1978; Beaulieu and Colonnier, 1989; Gomez et al., 2019). Furthermore, as many mammals live in some social structure, their brains are built to navigate the corresponding social relationships, including, for example, communication, cooperation, and competition. Moreover, social constructs like language and culture can shape perception (Winawer et al., 2007; Goldstein et al., 2009) and more complex social aspects, like socioeconomic status, can be reflected in the structure of the brain (Kweon et al., 2022). In fact, the brain is shaped by its entire surrounding *eco-system* (on a scale of **megameters**) by the adaptive processes of evolution. There is even evidence for coherence between brain rhythms and the electromagnetic resonance in the earth's ionosphere (Saroka et al., 2016).

SCALES IN THE FOCUS OF THIS WORK Consequently, neural systems incorporate into their organization influences from at least 15 spatial scales of magnitude, and there is active neuroscience research within and across all these different scales. Additionally, there is at least as wide of a span of temporal scales involved in the structuring and function of the brain. In the scope of the present work, we are particularly interested in the activity and function of neural structures on a scale reflecting "everyday experiences", agency (Hoel, 2018), and conscious behavior. This means, in space, we focus on groups of individual neurons up to cortical areas and hemispheres. In time, we focus on the scales of milliseconds (the timing resolution of an action potential) up to a few seconds and minutes (the timing of simple tasks). These spatiotemporal scales also roughly correspond to a theoretically postulated maximum of integrated information in the brain, i.e., the abstract description level with the maximal intrinsic causal power (Hoel et al., 2016).

1.2 A NETWORK VIEW ONTO NEURAL ORGANIZATION

Networks are an ensemble of connected nodes that interact via edges and exhibit joint dynamics determining the network activity. This syntax of networks is immensely powerful for representing a wide range of systems, from the trivially simple to the highly complex (see, e.g., Albert and Barabási, 2002). The syntax of networks also lends itself exceptionally well to the realm of neural systems (Bassett and

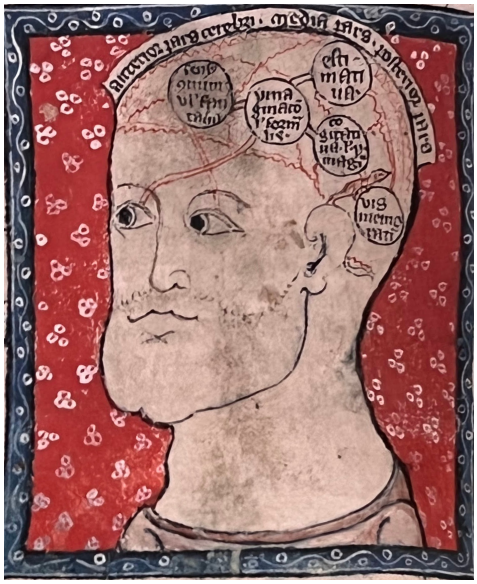

Figure 1.1: **Schematic illustration of the ventricular doctrine.** The schema shows five ventricles for the faculties of common sense, imagination, estimation, cogitation, and memory. The illustration is accompanied by a text titled "Qualiter caput hominis situatur" (How the human head is structured), from a trilingual compendium of texts on parchment from the first half of the 14th century, England *(Cambridge, Syndics of Cambridge University Library. Photo taken at the Fondazione Prada, Venice, Italy.)*

Sporns, 2017). One of the earliest descriptions of networks used to explain the brain is the so-called "ventricular doctrine" that originated from the works of early Christian philosophers Nemesis of Emesa (350-420) and Augustine of Hippo (354-430) and was further developed by Avicenna (980-1037), Thomas Aquinas (1225-1274), Albertus Magnus (1193-1280), and others (Manzoni, 1998; Rose, 2009). The ventricular doctrine already postulates the existence of functionally specific units in the brain where the higher faculties of humans reside, the interactions informing human behavior (Figure 1.1).

Since then, the network approach has only increased in popularity as a mental and mathematical tool to describe complex systems such as the brain. The approach offers qualitative and quantitative descriptions, as the edges can represent measures of the interaction between the nodes. Notably, as a tool, it can be applied in multiple different ways (Peel et al., 2022). The different ways of describing a neural system as a network differ in the scale of the description and the scale of its granularity: What is the scope of the represented system? What do the individual nodes represent? For example, networks

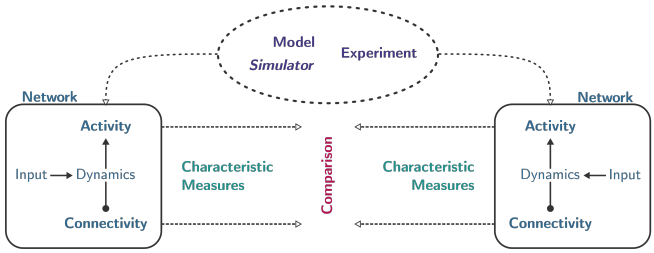

Figure 1.2: **Syntax of networks and their comparison.** Here, a *network* is an ensemble of neurons. The neurons interact via a *connectivity* exhibiting a joint *dynamics* that is entrained by an *input* and results in a corresponding *activity*. A network or network description can be derived from *experimental* measurements or a *model* implemented with a *simulator*. Abstractions of the network in the form of *characteristic measures* on the connectivity or activity level are the basis for *comparisons* between network descriptions.

can describe the interaction of the compartments of a single neuron, between individual neurons, between brain areas, or between entire brains. These scales can also be combined: networks of neurons that entail multi-compartment models (Schemmel et al., 2017; Yang et al., 2020b), multi-area networks interconnecting networks of individual neurons (Albada et al., 2019), or multi-brain networks incorporating networks of brain areas (Yang et al., 2020a).

Although there are multiple ways to describe a neural system in terms of a network, not all are equally useful. A useful network description should apply a node definition that reflects a reasonable level of biophysical segmentation. The abstraction of its scope and granularity should reflect the description level of interest. Therefore, given the scales of interest within this work identified previously in Section 1.1, we use networks to describe groups of neurons (up to cortical area) where nodes represent individual neurons. See Figure 1.2 for a schematic overview of the use of the network context within this thesis. Following the biophysical basis of neural systems, the edges between the nodes represent the strength of the synaptic connection between neurons. However, in many practical applications, information about physical connections may not be directly accessible. Therefore, one may instead quantify the interaction between the nodes based on measures derived from the neuron activity (e.g., correlations).

A benefit of describing neural systems as networks is that it enables systematic approaches to compare descriptions of the neural systems to each other regardless of whether they are based on experimental measurements of biological neural systems or computational models. Such comparisons can either focus on the connectivity level (aspects

of the network edges) or the activity level (dynamical aspects of the network nodes). Section 1.3 lays out the basis of measuring properties of the network activity. Section 1.4 elaborates on the need for comparing networks representing a neural system in contexts of experiments and models. Section 1.5 further addresses the connectivity level and activity level of network descriptions have a complex relationship.

1.3 QUANTITATIVE DESCRIPTION OF NEURAL DATA IN MEASURES AND MODELS

"To look, to see, to know" - Ludwik Fleck

The concept of a quantitative description of the world has a long and winding history. Aristotle distinguished distinctly between quantities and qualities. Quantities could only be compared via inequalities, whereas qualities could only be compared via degrees (Tal, 2020). Over the centuries, the popular philosophical perspective was developed towards establishing a relation between quantitative and qualitative descriptions. Quantity was rather understood as the intensity of a quality. This change of perspective may seem subtle, but this represents the logical foundation of the scientific use of measurements. This development enabled the popularization of the scientific approach of quantitative measurements and formulations of quantitative laws during the 16th and 17th centuries (Grant and Grant, 1996). A prominent example of this is Galileo Galilei, who employed and advanced this approach by a combination of experiment and mathematics (Sharratt, 1996).

The apparent conceptual discrepancies between qualities and quantities continued to be re-evaluated, guided by philosophical advances, as well as advances in measurement techniques and understanding of physical laws. For example, temperature slowly evolved from a quality to a measurable quantity (Sherry, 2011). Gottfried Wilhelm Leibniz further expanded the idea of quantity as the qualitative intensity, stating that all natural change is continuous, including the intensities of representational states of consciousness (Jorgensen, 2009). Indeed, the dichotomy of quality and quantity is still evaluated today. Not all things can be measured or should be described mathematically. For example, business math of the type "*Experience - Expectation = Satisfaction*" (Compeau, 2018) presents an ill-fated attempt to imprint aspects of quantities onto qualities (Diana, 2019). Besides the enormous effectiveness of quantitative and mathematical approaches (Wigner, 1960), there are still genuine limits. In the words of Bertrand Russel: "Physics is mathematical not because we know so much about the physical world, but because we know so little: it is only its mathematical properties that we can discover." (Russell, 2009).

Measurements are not a one-to-one mapping of reality (assuming an objective reality). Measurements are necessarily indirect and abstracting assessments of reality conceptually linked to theory. Measurement

both requires theory and is a prerequisite for theory. Therefore, one must avoid potential circular approaches where the theory being tested is already assumed in the method of measurement (e.g., measuring the expansion of objects due to temperature with a mercury thermometer (Franklin et al., 1989)). Consequently, measurements cannot test a theory in isolation but rather by applying a theoretic assumption that is only challenged when advancement in the measurement application unveils discrepancies that require alternative or more refined theories (Hacking and Hacking, 1983; Kuhn, 1961).

When applying this concept of measurement to the brain, we eventually want to measure cognitive processes (e.g., attention, conscious awareness, memory recall). However, without an underlying theory of these processes, we must suffice with accessible observables and proxy measures (cf. outside-in vs. inside-out strategy in Buzsáki, 2019). Here, we focus only on the network aspects of the neural system and on measurements of the network activity. Because of the brain's organization over multiple spatial and temporal scales (see Section 1.1), measurement techniques need to have corresponding resolutions. As mentioned above, measurements are inherently indirect. So, to represent neural activity, multiple alternative types of signals are used in different measurement techniques, including electric potential, magnetic field strength, fluctuation of calcium concentration, or radiant isotopes. The type of signal used for the activity measurement also informs the level of abstraction the measurement represents and which sort of biases the measurement technique might entail (Hong and Lieber, 2019; Sejnowski et al., 2014). Additionally, in the measurements of living organisms, there is often a trade-off between the quality of the measurement and the interference with the organism. On the one hand, acquiring accurate and prolonged recordings of cell activity in-vitro is much more tractable than in a living and behaving subject. Invasive recordings typically give a better resolution than non-invasive techniques. On the other hand, less interference with the organism allows for more behavioral flexibility and the natural function of the neural system. For neural systems, measuring with a "natural" environment and stimuli is more helpful in acquiring measurements that give insight into the functional mechanisms of "normal" activity. (Chamove, 1989; Felsen and Dan, 2005; Bartels and Zeki, 2005). Here, we primarily focus on the activity in the cerebral cortex, which is thought to be the leading actor in most higher cognitive functions in mammals, and conveniently is more accessible than deeper brain structures.

Measurements abstract the complexity of reality, in some aspects intentionally and in some aspects accidentally/circumstantial. For example, we abstract the activity to measure local electrical potentials with electrode recordings because we assume that they play a relevant role in the communication within and between neurons (e.g., via

voltage-gated membrane dynamics and the transmission of action potentials). However, electrode recordings and others types of recordings also abstract the network activity by measuring only a subset of neurons in place of the full network of interest. This is due to limitations of the measurement techniques in covering the space of the network, discerning the heterogeneous activity within that space (e.g., see the "dark matter" problem of neuroscience (Shoham et al., 2006)). This considerable degree of undersampling is a general challenge in modern neuroscience. Even with access to the activity of all neurons, it is not trivially evident that we could deduce how information processing in the brain works (Jonas and Kording, 2017).

One approach to reducing the complexity of the electric neural activity is to separate it into a pulsed binary signal and a continuous analog signal. The binary signal describes the action potentials of the neurons, extracellular measured as high-frequency pulses called *spikes*. The analog signal describes the lower-frequency components of the extracellular potential, called local field potential (LFP) (cf. Section 1.1). In Chapter 4, Chapter 5, and Chapter 6, we will focus on the spiking activity description, and in Chapter 7, and Chapter 8, we will focus on the LFP activity description. Although separating the types of electrical activity is helpful, the description level is typically still too complex to derive general mathematical activity descriptions. Therefore, we further reduce the activity signals to their characteristic properties. For the spiking activity, these can entail measures regarding the spike rate, its regularity (over time), its variability (over neurons), or its coordination (over time and neurons). The LFP standard measures include amplitude, oscillation phase, spectral power (w.r.t. to a frequency regime), or coordination (over time and locations). There is an extensive catalog of characteristic measures derived from various measurements. However, it is still unclear what precisely the uniquely relevant measure or combination of measures is to fully describe/explain/predict the functionality of a neural system.

The quantitative description of reality is not only the basis for experimental measurements but also for the creation of executable models. A conceptual model is a formal description of the system of interest. In contrast, an executable model is an operational implementation of the conceptual model that can be simulated and formulate predictions (Schlesinger, 1979; Thacker et al., 2004). Chapter 4 explores in more detail the concept of a distinct conceptual model and an executable model in the domain of computational neuroscience. Models do not suffer from the same issues of only being accessible via indirect measurements that provide a limited abstraction. All simulated activity and all network parameters are directly available. Instead, the model itself is an abstraction of the system they are describing. One might consider whether it is possible to gain new knowledge about the system of interest from its models when the model is just an abstraction

and inherently can not be entirely valid itself (Balci and Sargent, 1982; Sterman, 2000). This question is a matter of ongoing philosophical debate (Oreskes et al., 1994; David, 2009; Irobi et al., 2004). However, in practice, we are not concerned with a model's truth value but rather with its simulations' testable accuracy. Only the model's accuracy relative to the given system justifies its use as the basis for tractable analysis and prediction-making.

In particular, the field of computational neuroscience is driven by the development of models describing neuronal activity on different temporal and spatial scales, ranging from single cell dynamics (e.g., Koch and Segev, 2000; Izhikevich, 2004) to spiking activity in mesoscopic neural networks (e.g., Potjans and Diesmann, 2014; Markram et al., 2015), to whole-brain activity (e.g., Sanz Leon et al., 2013; Schmidt et al., 2018). Indeed, there are modeling approaches for most scopes and granularity levels of neural network descriptions (Section 1.2). Models employ various strategies in abstracting neural systems, for example, by reducing the variety of neural cell types, omitting physical extension of cells, and simplifying complex membrane dynamics to single differential equations (Teeter et al., 2018). However, even abstract models can give fairly realistic descriptions of neural systems and lead to new insights. For example, the Kuramoto model describes neural networks as a collection of coupled, self-sustained phase oscillators (Kuramoto, 1975). This abstract model can predict synchronization over large populations and phase transition of the network state at a critical coupling strength (Acebrón et al., 2005). In order to further leverage such insights, models of different description levels should be combinable, and models of different abstractions should be comparable (see Section 2.2).

The comparability and combinability of models entail multiple challenges. One is the heterogeneity of model descriptions, their implementations, and how they are made available. Although there is no consensus regarding this (Nordlie et al., 2009), several frameworks support researchers in documenting and implementing models beyond the level of custom-written code in standard high-level programming languages. These frameworks include guidelines for reproducible network model representations (Nordlie et al., 2009; McDougal et al., 2016), domain-specific model description languages (e.g., Plotnikov et al., 2016; Gleeson et al., 2010), modelling tool-kits (e.g BMTK[1], NetPyNE[2]), and generic network simulation frameworks (Davison et al., 2008). To share these models, but also data, with the community, several databases and repositories have emerged and are commonly used for this purpose, for example, GitHub[3], OpenSourceBrain[4], the

1 https://github.com/AllenInstitute/bmtk
2 https://www.netpyne.org
3 https://github.com
4 http://opensourcebrain.org

Neocortical Microcircuit Collaboration Portal[5] (Ramaswamy et al., 2015), the G-Node Infrastructure (GIN)[6], ModelDB[7], NeuroElectro[8] (Tripathy et al., 2014), or CRCNS [9] (Teeters et al., 2008).

In the following Section 1.4, we elaborate on the quantification of model accuracy and credibility.

1.4 COMPARABILITY OF MEASURED AND MODELED NEURAL ACTIVITY DATA

"What I cannot create I do not understand." - Richard Feynman

The abstraction of neural network descriptions from experimental measurements and computational models to characteristic measures provides the foundation for quantitative comparisons. The combination of quantitative descriptions of (neural)systems and their statistical evaluation of similarity is conceptually formalized as the process of validation (Schlesinger, 1979) (Figure 1.3). Validation is essential to modeling and has an extensive history of research and applications in many scientific and engineering domains. Many established processes, insights, and terminologies can be adapted to neuroscience. Although suggested by the term "validation", the goal of modeling and model validation is not "validity". Instead, modeling concerns the testable correctness of its simulated output relative to the system of interest. Only a model that can accurately predict the behavior of the *system of interest* is useful for inferring properties of that system. Validation, the process of evaluating this accuracy, is an ongoing task (Murray-Smith, 2015). There is not a single validation test that is sufficient to evaluate a model in its completeness (Forrester and Senge, 1980), and the outcome of a validation process is not a definite or binary verdict. The outcome of a validation process typically comes in the form of a quantitative score (or scores). The interpretation of whether this score indicates agreement or discrepancy between the model and the system of interest depends on the model's intended application and scope (Carnap, 1968).

In 1979, the Technical Committee on Model Credibility of the Society of Computer Simulation established a widely recognized description of model verification and validation environment. The validation setup is separated into three basic elements (see Figure 1.3). The *system of interest* can be defined as "an entity, situation, or system which has been selected for analysis" (Schlesinger, 1979), and constitutes the references against which validations are carried out. Originally referred to as *reality*, the term 'system of interest' better conveys that modeling efforts always operate within boundaries, for example, in terms of scope and granularity (Section 1.2). In this modeling environment, separating

5 https://bbp.epfl.ch/nmc-portal
6 https://gin.g-node.org
7 https://senselab.med.yale.edu/modeldb
8 https://neuroelectro.org
9 https://crcns.org

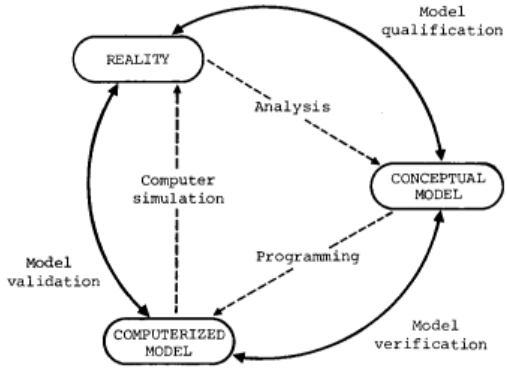

Figure 1.3: **Schematic view of the model simulation environment.** In a 1979 report (Schlesinger, 1979), the SCS Technical Committee on Model Credibility conceptually defined the relationships between the *conceptual model*, the executable *computerized model*, and *reality* (later also termed *system of interest* (Thacker et al., 2004)). These three entities describe each other via analysis, programming, and simulation activities (dashed lines). *Validation*, *verification*, and *model qualification* are defined as the assessment activities (solid lines) between the entities.

the conceptual model and its executable model implementation is crucial for the logic of validation and verification. The conceptual model represents the abstraction and analytic description of reality, but only the corresponding verified executable model can perform simulations that can be validated.

This separate consideration of the executable model corresponds to the increasing efforts to provide model description languages and simulator engines independent of any concrete conceptual model. For example, successful simulators with specific but overlapping application domains are NEST (Gewaltig and Diesmann, 2007), BRIAN (Goodman, 2009), and NEURON (Hines and Carnevale, 1997). Additional efforts like the modeling languages NESTML (Nagendra Babu et al., 2021) and PyNN (Davison et al., 2008) further generalize executable model descriptions so that they are independent of the simulator choice. These generalization efforts also tend to make the modeling process more reproducible, and reproducibility is necessary for comparability. Any model simulation must first reliably reproduce with the same outcome (preferably independent of the modeling language and simulator choice) before it can be reasonably compared to other data (Goodman et al., 2016).

Although the initial formalism by Schlesinger (1979) was adapted and refined by many others (e.g., AIAA, 1998; Giannasi et al., 2001; Thacker et al., 2004; Murray-Smith, 2015; Sargent et al., 2013), the

definition of the different concepts and terms, like 'verification' and 'validation', have been remarkably consistent. However, in practice, the conceptual steps are likely to be highly intertwined. For example, validation indicating a considerable discrepancy may cause a re-verification of the model implementation. Alternatively, the validation may provoke an iteration and further sophistication of the underlying conceptual model. When the validation result explicitly informs a change in the model description, e.g., the change of a parameter value, then this has to be considered a calibration. Calibration integrates information from the data, that the model simulation is compared against, back into the model description. This procedure naturally improves the agreement with the respective dataset. However, in contrast to a validation test, the agreement to the calibration data set is not informative about the model's predictive power. To combine calibration and validation in the model development process, they need to be strictly separated (Thacker et al., 2004). Validation can only be meaningful when done with previously unseen data. The independence and power of the validation can further be improved by comparing measures not used in the calibration.

Neuroscience works with models on different scales like single-neuron, microcircuit, or multi-area models. Further, these models can incorporate elements from different scales, for example, receptor models being part of neuron models or neuron models being part of microcircuit models (Section 1.2). This layered complexity allows for different strategies to model validation. Multifaceted validation approaches to complex systems are not exclusive to neuroscience (Forrester and Senge, 1980). A common approach is to start from the level of the smallest elements. Considering our network-level approach, this corresponds to validating single neuron responses or synapse behavior to experimental data (see, e.g., Markram et al., 2015; Reimann et al., 2015). This kind of single-cell validation suggests that when the basic building blocks of a system are validated, the entire system should perform appropriately. Validation of the entire system follows only when all the sub-elements have passed their validation with reasonable agreement (Thacker et al., 2004). However, the relation between the smallest elements' functionality and the composite systems' emergent dynamics is intrinsically complex. Indeed, the relations across scales in biological arrangements are often unknown and itself part of the modeling (Noble, 2008).

Nonlinear effects and sensitivity of individual network elements to minute changes (Marder and Taylor, 2011) often prevent extrapolating the behavior of the complete network. Moreover, when shifting the focus to the global features of the network dynamics, the individual neurons or sub-circuits in the model may be abstracted (e.g., the cortical microcircuit model (Potjans and Diesmann, 2014), multi-area model (Schmidt et al., 2018), or mean-field approaches (Barabási et al., 1999;

Helias and Dahmen, 2020; Layer et al., 2022)). Simplified neuronal dynamics models have the advantage that they can be mathematically approximated and facilitate an understanding of the governing mechanisms (Schuecker et al., 2015; Bos et al., 2016; Litwin-Kumar and Doiron, 2012; Renart et al., 2010; Ostojic et al., 2009). Despite abstract neuron models not sufficing to the same validated accuracy as more biologically realistic neuron models, the corresponding network models are able to reproduce many dynamical features observed in experimental data (Shadlen and Newsome, 1998; Renart et al., 2010; Potjans and Diesmann, 2014). In such cases, a bottom-up validation approach is not suitable. A top-down, network-level validation approach represents a complementary approach focusing on collective network dynamics. Network-level validation allows for a direct comparison of population activity features independent of the abstraction level of the neurons (e.g., spiking or rate-based neuron models).

Classical validation is concerned with comparing models to experimental observations. However, the same process can also be applied to compare model vs. model. These could be alternative model descriptions, versions, implementations, or simulations (see Chapter 4). Although there is no affirmation concerning the biological system of interest, model vs. model comparisons can also improve confidence in a model. Model vs. model comparisons can be validated by proxy when one of the models is already validated against experimental data (Martis, 2006). Further, they can evaluate the sensitivity with respect to input and parameter changes (De Schutter and Bower, 1994; Saltelli, 2002; Marino et al., 2008; Zi, 2011; Borgonovo and Plischke, 2016; Tennøe et al., 2018). Given the scarcity and specificity of available experimental data, model vs. model comparison also allows testing a more extensive dynamical regime that available data would not necessarily cover.

1.5 RELATIONSHIPS BETWEEN CONNECTIVITY, ACTIVITY, AND FUNCTION OF NEURAL NETWORKS

"Computation, intelligence and consciousness are patterns in the spacetime arrangement of particles that take on a life of their own, and it's not the particles but the patterns that really matter! Matter doesn't matter." - Max Tegmark

The function of neural networks is expressed via its activity, and network activity is mediated by network connectivity. Thus, there are functional aspects engrained in network connectivity. Korbinian Brodmann performed a prominent study of brain structures at the beginning of the 20th century. He used staining techniques to make neurons visible in human brain preparations. By distinguishing the cellular structures, he segmented the cortex into 52 distinct areas (Figure 1.4A). Prior works hypothesized the localization of cognitive functions in the brain based on the effects of localized brain lesions and external electrical stimulations (Ferrier, 1874; Swedenborg, 1882). However, Brodmann's works provided the basis for linking the local differences in brain anatomy to corresponding functionality. For ex-

ample, after his descriptions of the cortical cytoarchitecture, it was found that Brodmann area 4 is the motor cortex that controls voluntary movement, and Brodman area 17 is the primary visual cortex that receives and processes signals from the retina. However, it was, and still is, generally unclear which details of the cortical network connectivity provoke which activity signatures to encode the expression of a corresponding function.

Throughout history, theories explaining the function of the nervous system were often linked to the technologies and beliefs of the time. The Greek physician Galen (129 - 210) formulated the first prominent theory involving a sort of neural network. The so-called "balloonist" theory describes mechanisms in the body as hydraulic, where pressurized air or liquid transported within hollow neurons inflates the muscles (Pearn, 2002). This theory survived for about 1500 years, with famous supporters like Rene Descartes (Descartes, 1662), until it was finally completely refuted in the late 17th century (Cobb, 2002). About a century later, at the beginning of experimentation around electricity, Luigi Galvani found that electrical charges (the "animal spirits") transmitted in the body caused muscle movement (Galvani, 1792). With the technological advances over the next century (Figure 1.4B), in 1924, Hans Berger was able to make the first in-vivo electrocorticogram (EEG) recording of the ongoing electrical activity in the human cortex (Figure 1.4C, Berger (1929)). More recently, the brain's functionality is often conceptualized via the analogy to a computer (Chirimuuta, 2021; Brette, 2022). Correspondingly, the means of communication within the brain is referred to as *neural code* (Quiroga and Panzeri, 2013; Brette, 2019), which should represent the activity signatures that are responsible for the function of a neural system.

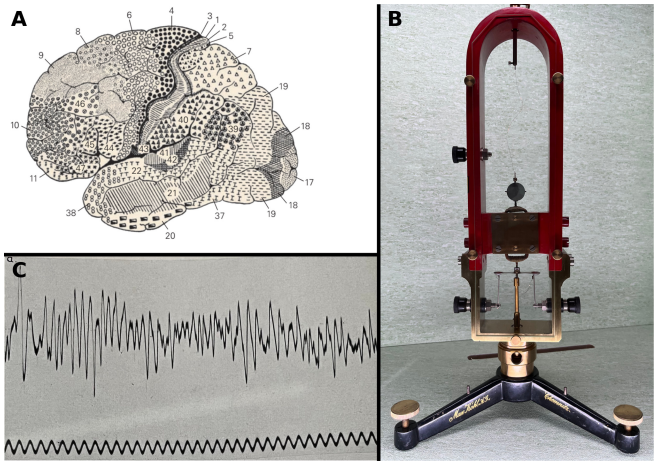

Figure 1.4: **A)** Segmentation of the human cerebral cortex into 52 (not all visible here) areas by Korbinian Brodmann in 1909. The area definitions are based on the cytoarchitectural organization of neurons made visible by cell staining. The figure shows a medial view of the human cerebral cortex (amended from Brodmann (1909) displayed in Kandel et al. (2001)). **B)** Max Kohl loop galvanometer for measuring brain waves. This type of recording device was used in the early 20th century for measuring small voltages *(Friedrich-Schiller-Universität Jena, Collection of scientific and technical devices for physics. Photo taken at the Fondazione Prada, Venice)*. **C)** Using a device similar to the one depicted in panel B, Hans Berger recorded the first EEG signal 1924. The figure shows an oscilloscope trace of a derivation from the cortex of the left frontal and right parietal lobes (within ~ 0.1 s) (Berger, 1929).

Even without the direct analogy to a conventional computer, we can examine our network view on the brain towards the eventual units of computation (MacLennan, 2004; Horsman et al., 2014; Wood, 2019; Papadimitriou et al., 2020). What are the network activity features used in the function of the brain? What are the strategies for how the digital and analog signals in biological neural networks represent and operate on information? Narrowing down the basic units of computations by identifying irrelevant network activity aspects is very difficult, as nearly every aspect seems to have at least some relevance. The most familiar abstraction of neural network activity represents the characteristic shape of spikes or action potentials only by their spike times. However, even this basic abstraction might discard relevant aspects of the network activity, as there are arguments that variations in the spike shapes play a role in computation (Juusola et al., 2007; Ramezani and Akan, 2018).

In contrast to trying to identify non-relevant aspects, machine learning demonstrates impressive results in replicating functionalities based

on model descriptions of very simplified abstraction levels. A multilayer perceptron architecture with binary neurons (Rosenblatt, 1958) and backpropagation (Rumelhart et al., 1985; LeCun et al., 1988) can already successfully perform a variety of classification and regression tasks (Murtagh, 1991; Verma and Kumar, 2020). Deep convolutional neural network models demonstrate a similar dynamic to the hierarchical processing in visual cortical areas (Yamins and DiCarlo, 2016). The machine learning approach explicitly optimizes only for functionality and typically not for a connectivity architecture or network activity that describes biological neural networks. Whereas neuroscience is concerned with identifying the attributes of the observed connectivity and activity that initiate a function. Albeit their different objectives, there is a long-lasting synergetic exchange between machine learning (or computer science in general) and neuroscience and continuous work to close the gap between biological realism and functional optimization. One might argue that in network architecture design, form follows function, and there are claims that the brain also performs function optimizations (e.g., for mutual information (Friston, 2010)). However, results from machine learning, for example, that rate-based neural networks currently still outperform spiking networks (Nunes et al., 2022), provide only circumstantial evidence for a purely rate-coding mechanism in biological neural networks. In Section 1.3, we already posed the question about functionally relevant aspects of network activity. More precisely, to link network activity to function, we are interested in the abstract activity features that contain the same or similar amount of information concerning a given function.

The relevance of action potentials is undisputed, as they represent the primary communication signal within and between neurons. However, different aspects of the spiking activity have been demonstrated to provide specific indications of neural function. The primary relevance of the firing rates is, for example, suggested by the experimental evidence of tuning curves (Seung and Sompolinsky, 1993), receptive fields (Hubel and Wiesel, 1959), and place cells (O'Keefe and Dostrovsky, 1971). Whereas the relevance of precise spike timing is, for example, suggested by the extremely fast observed reaction times (Kempter and Gerstner, 1995; Thorpe et al., 1996; Butts et al., 2007), and the observed prominent role of synchrony in neural networks (Grün et al., 2002; Abeles, 1991; Kreiter and Singer, 1996). Despite the apparent conflict between the spike-rate and spike-time coding paradigms, they are not necessarily mutually exclusive (Perkel and Bullock, 1968; Brette, 2015; Kiselev, 2016). Furthermore, there are aspects of the coordination of activity with a population of neurons that are sought to represent relevant abstraction levels. For example, the projection of activity onto low-dimensional manifolds (Semedo et al., 2019; Gallego et al., 2017), the organization of spiking activity into spatio-temporal patterns (Ikegaya et al., 2004; Quaglio et al., 2017;

Stella et al., 2022), or the selective locking of spike times to phases of the LFP (Vinck et al., 2010; Denker et al., 2010). The LFP itself is an abstraction of the combined activity of nearby neurons, mainly their postsynaptic potentials. However, many aspects contribute to the formation of the LFP, including the connectivity and morphologies of the neurons (Pesaran et al., 2018).

Analogous to the aspect of the spiking activity that it shows a spatial organization in the form of spatio-temporal patterns, the LFP also shows a spatial coordination of its local oscillation phases. This spatial coordination often forms patterns of traveling waves that can be observed across brain areas, brain states, measurement techniques, and frequency bands. It has been shown that the formation of waves can result from the underlying connectivity (Davis et al., 2021; Sanchez-Vives et al., 2017). Thinking of waves emerging from interacting nodes in a network and the generality of the network description (Section 1.2), it should not be surprising that we observe wave phenomena all around us. They are present in diverse physical contexts and are frequently associated with informative content: electromagnetic waves are the carriers of radiant energy and information, waves in compressible mediums enable sound transmission, gravitational waves in spacetime tell us about astronomical events , and waves in the oceans attract surfers. This ubiquity of waves stems from the fact that waves are, in principle, substrate-independent. They have distinct properties like frequency, amplitude, propagation velocity, and dispersion and can be mathematically described independently from the type of medium. Similarly, computation and even consciousness can be thought of as substrate-independent[10]. This analogy is not a direct argument for the relevance of waves in the brain, but it represents an intriguing perspective on the eventual role of waves in neural systems. Although their functional role is not yet fully understood (Wu et al., 2008), multiple works link properties of neural wave activity to function.

At the low end of the frequency spectrum < 1 Hz, oscillatory activity is often associated with rest and sleep (Amzica and Steriade, 1998). In particular, propagating transitions between low and high activity states called *slow waves* (Steriade et al., 1993; Contreras et al., 1996; Achermann and Borbély, 1997; Sanchez-Vives and McCormick, 2000) are consistently observed in states of NREM sleep and anesthesia. Slow waves were functionally associated with regulating or indicating the degrees of consciousness (Tononi and Massimini, 2008; Massimini et al., 2009; Massimini et al., 2012; Camassa et al., 2021) as well as memory formation and learning (Hanlon et al., 2009; Watrous et al., 2015; Wei et al., 2016; Schonhaut et al., 2020; Tukker et al., 2020; Capone et al., 2019a; Pazienti et al., 2022; Tonielli et al., 2022). In higher

10 see the response of Max Tegmark to the question "What scientific term or concept ought to be more widely known?" (2017) https://www.edge.org/response-detail/27126

frequencies, oscillatory wave activity can be observed in the alpha, beta, and gamma domains in various experiments of awake, behaving animals (Senseman and Robbins, 2002; Petersen et al., 2003; Wu et al., 2008; Keane and Gong, 2015; Greenwood et al., 2015; Townsend and Gong, 2018; Denker et al., 2018b; Roberts et al., 2019; Liang et al., 2021). There, the wave activity has been linked to functional aspects in visual performance (Freeman and Barrie, 2000; Gabriel and Eckhorn, 2003; Vinck et al., 2010; Davis et al., 2020; Townsend et al., 2017), motor tasks (Rubino et al., 2006; Balasubramanian et al., 2019; Heitmann et al., 2015), and has been the object of postulations of how computation may incorporate the wave activity (Gong and Leeuwen, 2009; Qi et al., 2015; Muller et al., 2018; Halgren et al., 2018; Heitmann et al., 2012).

2

THESIS STATEMENT

2.1 PRINCIPLES OF HOW CONNECTIVITY PROVOKES ACTIVITY PROVOKES FUNCTION

There are still many known unknowns about the specific working mechanisms of neural systems (see for example van Hemmen and Sejnowski, 2005). Given the complexity and lack of understanding of even small neural systems (e.g. Sarma et al., 2018), any single research project in unlikely to completely unravel the mysteries of the brain. Fundamental scientific progress requires insights from many disciplines addressing research questions on many scales, from molecular interactions to social and environmental interactions (Section 1.1). Such progress is likely the result of slow but steady, iterative research endeavors from many different approaches. There is a real benefit in the variance of approaches, measurement techniques, and analysis methods. On the one hand, this poses an excellent opportunity to create an evaluation of neural networks and their function from multiple viewpoints yielding a more comprehensive and robust understanding (see e.g. Botvinik-Nezer et al., 2020). On the other hand, additional challenges exist to put the individual pieces of insight together to build a cumulative understanding of biological neural network mechanisms. In order to effectively build on each other's work, we face non-trivial difficulties, for example, to find, access, reproduce, reuse, and combine the heterogeneous ensemble of previously published data, results, and methods from different sources.

The goal of this thesis is to propose, implement, and demonstrate systematic approaches to quantitatively evaluate neural network activity and measure its relation to the underlying connectivity and its eventual functional purpose. Concretely, we aim to characterize neural network activity on a meaningful description level to study how activity features, like coordinated spiking or traveling cortical waves, emerge from network connectivity and how they relate to network functionality (Section 1.5). We approach these research questions by performing comparisons of network descriptions from experimental recordings and computational models. Therefore, our analyses need to integrate data and models of heterogeneous sources and make them quantitatively comparable.

2.2 SYSTEMATIC QUANTITATIVE COMPARISONS OF NEURAL SYSTEMS

Today's research landscape excels in an unprecedented richness of experimental data and methodologies. However, different recording techniques differ considerably in how they capture neural activity. These differences include the type of signal (e.g., electric activity, magnetic fields, fluctuation of calcium concentration, or radiant isotopes) and the signal's scale in terms of temporal resolution (sub-milliseconds to seconds), spatial resolution (micrometer to centimeters), and spatial extent (single electrode to the whole brain) (Hong and Lieber, 2019; Sejnowski et al., 2014). The complementarity of the different experimental approaches, each focusing on specific aspects or features, enables a deeper understanding of the various neuronal phenomena. However, each recording technique has its trade-off (e.g., spatial vs. temporal resolution, latency, artifacts, and biases). Combining different measurements can also validate findings independently of a particular measurement device. It further informs the creation of more detailed brain dynamics and function models. To properly exploit the richness of the available data and translate it to the richness of analysis results, it becomes necessary to develop strategies to combine the heterogeneity in experimental measurements and model simulations. In this context, data integration and comparison are two sides of the same coin, as some data measures are to be grouped while others are to be contrasted, depending on the application. Applications include the calibration and validation of models, the quantification of biological variability for reaching a higher level of abstraction in describing phenomena, and measuring the effects of various influences on network activity, such as external stimuli in an experiment or parameter changes in a model description. However, performing such comparisons is not a trivial task and, so far, only rarely addressed.

Here, we formalize the approach of cross-domain comparisons based on the general concept of validation. In particular, we adapt the concept for comparison on the network-activity level in contrast to bottom-up validation. Besides the model vs. experiment validation, we extend the application to model vs. model and experiment vs. experiment applications (Section 1.4). For comparing two networks, this concept entails the quantitative characterizations of the network activities on a common description level, an assessment of similarity (in most cases based on statistical testing), and the interpretation of the level of resulting agreement in the context of the comparison. In this context, we evaluate and extend existing statistical testing approaches. Furthermore, we explore additional aspects of comparability, such as ensuring the reproducibility of results, identifying and separating sources of variability, evaluating technical limitations (e.g., given by simulators, measurements, or algorithms), and discussing the

convergence of network activity from different sources to a common description level.

2.3 CHARACTERIZATION OF NEURAL SYSTEMS ON A MEANINGFUL LEVEL OF ABSTRACTION

Arguably not every aspect of the activity of biological neural networks is of equal relevance for its function. In most cases, the abstraction of spike waveforms to spike times is useful. In some cases, further abstraction to rates can be useful (Mazurek and Shadlen, 2002). In other cases, the precise spike timing is relevant, but only for some (combination of) spikes and up to a certain precision (Reich et al., 1997; Gutkin et al., 2003; Gütig, 2014). Similarly, for the local field potential (LFP) it depends on the context which frequency components are regarded as relevant and which attributes are considered "noise" (Knoblauch and Palm, 2005). Furthermore, we can see neural network activity only through the lens of measurements that filter and distort the activity or via simplified recreations in the form of models (Section 1.3). On this basis, we aim to find meaningful quantitative characterizations of neural network activity that carry sufficient information to, for example, understand its functionality better, infer aspects of the connectivity, or distinguish it from another network activity. Besides meaningful characterizations, we are further concerned with the appropriate level of abstraction on which we can compare two network activity characterizations. What this constitutes depends on the respective data modality and the scientific questions to be addressed. Similar data modalities can have more immediate commonalities, whereas very different ones may only be compared via a more abstract description. Generally, comparing two networks benefits from having a common description level, corresponding characterizations, and equivalent or at least comparable methods for extracting these characterizations.

Having the above considerations in mind, we here explore different measures of neural network activity. On the level of the spiking activity neurons, we look at measures of different complexity in the sense of neuron-wise, pair-wise, and other higher-order measures. Toward population measures, we further investigate measures of continuous representations of network activity, such as the LFP. There, we particularly characterize the attributes of spatially propagating wave patterns. We derive these measures from different measurement techniques (multielectrode arrays, laminar probes, ECoG, and calcium imaging) and different model descriptions and implementations (e.g. spiking networks, mean-field networks, stochastic processes, network realizations on neuromorphic hardware).

2.4 THEMATIC THREADS

Based on the preceding overview in Chapter 1, we outlined challenges to the advancement of understanding of biological neural networks that we aim to address in with this work. We started from the scientific question of how the activity of biological neural networks is produced from its underlying connectivity and how it implements functions Section 2.1. We then traced back the requirements of how to get there via an approach of systematic network comparisons Section 2.2 and corresponding characterizations of network activity Section 2.3. From this assessment, we segment our approach of comparing neural networks to analyze their dynamics into three steps that illustrate the key aspects of the different comparison scenarios presented in the Results Chapters of this work. Therefore, these three steps represent the thematic threads that run through this thesis:

1 **Characterization of neural data**: which characteristic measures represent the relevant aspects of neural networks for a given application?

2 **Modality of comparison:** in which relation are the two networks to be compared?

3 **Interplay of connectivity, activity, function:** what insights can the systematic comparisons of networks provide regarding the networks' dynamics?

3

THEMATIC OUTLINE

The central part of the thesis progressively follows the three thematic threads that were introduced in the previous Chapter 2. In five research applications, we explore different aspects of neural network comparability, outlined in Figure 3.1.

NETWORK LEVEL VALIDATION & REPRODUCIBLE SIMULATIONS In Chapter 4, we set up a formal framework for performing validation tests of neural network models based on *univariate measures* on the network activity level. Building on the classical validation scenario of model vs. experiment, we introduce analogous practices of how models can be validated (or 'substantiated' to specify the terminology) without experimental data by comparing different model instances. We apply this framework to a particular comparison case with the two model instances that only differ in their respective simulation engine. The approach of comparing *model vs. model across simulators* presents a formal approach to evaluate the quantitative agreement between simulators. Furthermore, it evaluates the *consistency of the network activity w.r.t simulator differences*, helping to differentiate activity features that are artifactual from robust expressions of the network model.

STATISTICAL COMPARISON OF NETWORK ACTIVITY AND CONNECTIVITY VIA EIGENVECTOR ANGLES In Chapter 5, we construct a statistical test for comparing the *structure of pairwise measures of neural activity and connectivity* between two network realizations. In particular, it can evaluate the structural similarity between matrices containing the Pearson correlations between spike trains or matrices containing the synaptic connection strength between neurons. We demonstrate the test with stochastic, simulated, and experimental data scenarios. Furthermore, we create synaptically rewired variations of a network model and compare *model vs. model across connectivity realizations* w.r.t their connectivity structure and the activity correlation structure. This approach enables us to quantitatively *link changes of the connectivity to changes in the activity* in that network model.

ACTIVITY-DRIVEN MICROCONNECTOME ESTIMATION In Chapter 6, we approach the question of how much the columnar structure of the canonical cortical microcircuit varies across different cortical areas. To this end, we look at the estimated microscale connectivity within cortical areas. To *differentiate spiking activity across functional ar-*

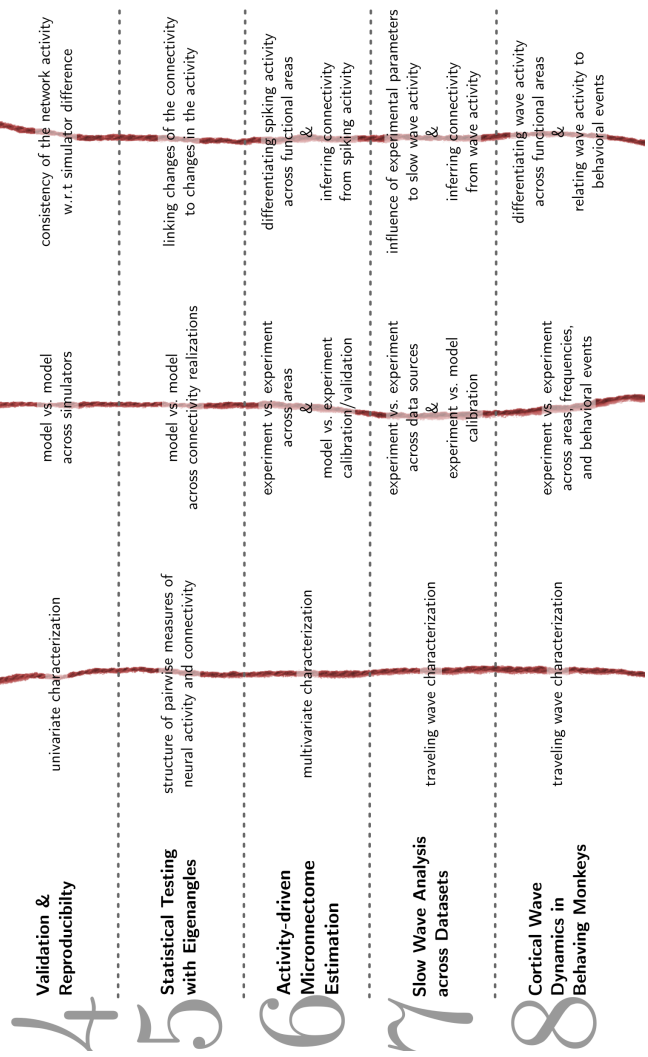

Figure 3.1: **Thematic outline of the thesis chapters.** Towards the goal of better understanding the networks of the brain, we identify three steps (columns) in our approach that run as common thematic threads through the five Results chapters (rows). Chapter 4 is based on the publications (Gutzen et al., 2018b; Trensch et al., 2018), Chapter 5 is based on (Gutzen et al., 2022), Chapter 6 is based on (Morales-Gregorio et al., 2022), Chapter 7 is based on (Gutzen et al., 2024; Capone et al., 2023), Chapter 8 is based on an ongoing research project.

eas, we create a *multivariate characterization* based on multiple network activity measures. With this characterization, we perform pairwise comparisons *experiment vs. experiment across areas* with spiking activity data from multiple datasets. Furthermore, we lay out and test a *model vs. experiment calibration/validation* workflow to optimize the connectivity parameters of a microcircuit model, effectively *inferring connectivity from activity*.

SLOW WAVE ANALYSIS ACROSS HETEROGENOUS DATASETS In Chapter 7, we introduce an adaptable analysis pipeline for analyzing cortical slow wave activity and realizing comparisons of *experiment vs. experiment across data sources* based on *traveling wave characterizations*. We demonstrate the pipeline implementation by performing a meta-study across multiple electrocorticography and wide-field calcium imaging datasets of anesthetized mice, evaluating the *influence of experimental parameters on slow wave activity*. Furthermore, by evaluating model simulation data with the analysis pipeline, we enable *experiment vs. model calibration* approaches and additionally support *inferring connectivity from wave activity*.

CORTICAL WAVE DYNAMICS IN BEHAVING MONKEYS In Chapter 8, we reuse our adaptable analysis pipeline for the *traveling wave characterization* of oscillatory LFP waves. The LFP activity is measured with multielectrode arrays in the visual cortex in awake monkeys performing a hand-eye coordination task. We analyze the wave activity from four different cortical areas along the visual pathway and four distinct frequency regimes stretching from 1 to 90 Hz. By comparing the recordings in a scheme *experiment vs. experiment across areas, frequencies, and behavioral events*, we identify wave attributes specific for the individual areas and frequency bands. Furthermore, by cutting, filtering, and aligning the trial structure of the behavioral task, we relate patterns of the *wave activity to behavioral events*.

Part II

RESULTS

4

NETWORK LEVEL VALIDATION & REPRODUCIBLE SIMULATIONS

This Chapter is based on the following publications:

- Gutzen, Robin, Michael von Papen, Guido Trensch, Pietro Quaglio, Sonja Grün, and Michael Denker (2018b). "Reproducible Neural Network Simulations: Statistical Methods for Model Validation on the Level of Network Activity Data." In: *Frontiers in Neuroinformatics* 12, p. 90. DOI: 10.3389/fninf.2018.00090
- Trensch, Guido, Robin Gutzen, Inga Blundell, Michael Denker, and Abigail Morrison (2018). "Rigorous Neural Network Simulations: A Model Substantiation Methodology for Increasing the Correctness of Simulation Results in the Absence of Experimental Validation Data." In: *Frontiers in Neuroinformatics* 12, p. 81. DOI: 10.3389/fninf.2018.00081

Author contributions:

- RG, MvP, GT, SG, MD designed the study.
- MD, SG, AM supervised the project.
- RG, MvP, MD conceptualized the validation framework.
- GT, RG, MD, AM established the terminology.
- GT implemented and verified the models.
- RG supported by MvP implemented the software for the validation testing.
- RG performed and visualized the analyses.
- PQ performed the SPADE analysis.
- RG wrote the initial draft of Gutzen et al. (2018b), contributed to the initial draft of Trensch et al. (2018), and wrote this chapter.
- All authors contributed to the writing of the respective published manuscripts.

4.1 INTRODUCTION

QUANTITATIVE AND REPRODUCIBLE VALIDATION PRACTICES

Building models of systems that we observe in the real world and predicting their behavior is a constructive approach to developing an understanding of these systems. However, measuring how well simulations describe the system of interest is generally not straightforward. Validation formalizes the evaluation of models regarding predictive accuracy. However, as outlined in Section 2.2, the statistical validation of models lacks a standardized approach and supporting software tools. Thus, it is usually up to the individual researcher to define how much the simulation outcome is supposed to match the experimental data. Consequently, we identify three difficulties with published model descriptions:

1. Models are only tested qualitatively instead of quantitatively. For example, the spike trains resulting from the simulation are visually classified (e.g., Voges and Perrinet, 2012), but without calculating specific statistics to quantify the features of the activity. This lack of concrete numbers and detailed records of how the numbers are calculated prevents a direct comparison to other models. This is a specific case of a general theme outlined in Section 2.3.

2. The information provided in a publication on the details of how the specific statistical analysis is performed and, thus, how a model is validated is insufficient to reproduce the validation scenario.

3. Models are only compared to a single experimental data set using a specific statistical measure. Moreover, the choices of data sets and measures are biased to address specifically the scientific aim of the publication. However, the absence of a standardized procedure to base the validation on a broad set of data sets and statistical measures limits the degree to which confidence in the model is quantified in a context detached from the research conducted in the publication. Moreover, it prevents the direct comparison between published models and their re-use in related studies.

SINGLE-CELL VS. NETWORK-LEVEL VALIDATION These issues are particularly prevalent for network models of neural systems. Validation in computational neuroscience is often just applied to the single neuron level (e.g. Markram et al., 2015). The advantage of single-cell validation is that the cellular activity, e.g., the response to current input, can be measured in different labs and even under different experimental conditions. Network-level validation, on the other hand, is more demanding. Experimentally, network dynamics can usually

only be measured *in-vivo*, involving more sophisticated experiments. Moreover, there can be a considerable variability between measured systems, e.g., different subjects. All these aspects need to be considered in a corresponding network-level validation approach. We further elaborate on the concept of validation, and its application to network-level and model vs. model approaches in Section 1.4.

PRIOR WORK AND OUTLINE Prior work on overcoming the listed issues and formalizing the validation process include the development of the Python module `SciUnit` (Omar et al., 2014; Sarma et al., 2016), and the description of workflows for the validation of models (Senk et al., 2017; van Albada et al., 2018; Kriegeskorte and Douglas, 2018). Furthermore, the work presented in this chapter builds on and progresses the work of Robin Gutzen's master thesis. There, we explored different statistical tests to compare network activity measures in a model vs. model approach. We applied this to the comparisons of activity correlations in a cortical microcircuit model (Potjans and Diesmann, 2014) implemented with the Nest simulator and on the neuromorphic hardware SpiNNaker (Furber et al., 2013). Furthermore, this previous work features an early version of the NetworkUnit Python module.

Here, we continue these efforts by refining and extending the validation test methodology, terminology, and implementation. Moreover, we investigate a new validation test use-case, developing and evaluating a model realization of the Izhikevich's polychronization model (Izhikevich, 2006) on the SpiNNaker system. We demonstrate a quantitative assessment of the SpiNNaker model via validation tests on a multiple network-level activity characterizations. Finally, we discuss the conditions under which the models are in acceptable agreement. We further discuss the applicability of the proposed workflow for other validation scenarios.

4.2 METHODS

4.2.1 *Terminology*

MODEL VERIFICATION AND VALIDATION Modeling is about representing a system with a desired level of accuracy. Evaluating the modeling effort is a non-trivial exercise that requires a rigorous validation process. The terms *verification* and *validation* may have different meanings in different contexts. For example, in software engineering verification and validation is the objective assessment of products and processes with the purpose is to support the quality of the developed product (Bourque and Fairley, 2014). In the development of computerized models, verification and validation are processes that accumulate evidence of a model's correctness regarding a specific application context (Thacker et al., 2004). The Society for Computer Simulation (SCS)

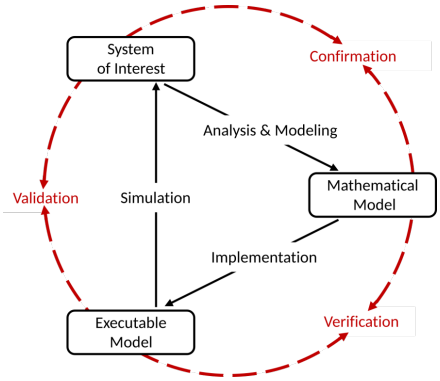

Figure 4.1: **Interrelationship of the basic elements for modeling and simulation.** The terminology introduced by Schlesinger (1979) (cf. Figure 1.3) describes the modeling and simulation processes in generic terms. For numerical models for neural network simulations, we propose the presented terminology, which we have adapted slightly from Thacker et al. (2004). The model distinguishes between modeling/simulation activities (solid black arrows), and assessment activities (dashed red arrows).

formulated a standard terminology to facilitate the communication between model builders and model users and establish credibility in computer simulations (Schlesinger, 1979). Their general definition may not do justice to a particular modeling domain. Thus, there are domain-specific adaptations that may specify the components by do not change the basic meanings. For neural network modeling, we propose the terminology displayed in Figure 4.1, amended from Thacker et al. (2004). The authors use the terms *reality of interest*, *conceptual model*, and *computerized model*. We prefer the terms *system of interest*, *mathematical model*, and *executable model*, as they are more explicit in expressing the underlying intent. Abstracting the neurobiological reality poses various empirical challenges (Section 1.3). Therefore, spiking neural network models are often not based on a specific biological ground truth from which a stereotypical behavior can be recorded. The term 'system of interest' recognizes that the modeling and simulation process can be applied to systems without concrete physical counterparts.

The **System of interest** is an *"entity, situation, or system which has been selected for analysis"*. The precise description of the modeler's intention is formulated via a **analysis and modeling** process in the **mathematical model** as a *"verbal description, equations, governing relationships, or natural laws that purport to describe reality or the system of interest"* (Schlesinger, 1979). The applicability of the model description is motivated by a **confirmation** process. However, the mathematical model alone cannot simulate the system of interest. By applying software

engineering efforts, it has to be implemented as an **executable model**. By separating the model into a mathematical and an executable entity, the terminology emphasizes the difference between verification and validation (Section 1.4). Model verification is the assessment of model implementation. Neural network models are mathematical models written down in source code as numerical algorithms. Therefore, the verification activities should be further specified:

- **Verification** describes the process of ensuring that the executable model appropriately represents the mathematical model and improving this fit.
- **Source code verification** tasks confirm that the functionality it implements works as intended.
- **Calculation verification** tasks assess the level of error that arises from various sources of error in numerical simulations as well as to identify and remove them (Thacker et al., 2004).

Only when the executable model is verified can it be reasonably validated.

- The **validation** process evaluates the consistency of the predictive simulation outcome with the system of interest.

This evaluation further considers the domain of the models' intended application and expected level of agreement. Since modeling abstracts the system of interest, it is only expected to match its reference to a certain degree and for certain prescribed conditions.

MODEL VERIFICATION AND SUBSTANTIATION For neural network simulations, the system of interest can be represented by empirical measurements, for example, multi-unit spiking activity acquired with electrophysiological recordings. However, this data may prove inadequate for validation. Depending on the specific system of interest, adequate data can be scarce and typically represents a substantial undersampling of neural network activity. Furthermore, the constraints of recording with biological organisms do not always allow for the required control to isolate the phenomenon that a model aims to describe. Consequently, in many neuronal network modeling applications, the available experimental data is just used to check for consistency instead of a complete validation process. However, the credibility and consistency of a model can be further increased by circumstantial evidence by comparing models against other model variations (e.g., a second implementation) (Thacker et al., 2004; Martis, 2006). Such a technique accumulate evidence for a model's plausibility even if the models are not formally validated. To avoid ambiguity with the existing model verification and validation terminology, we propose the term '***substantiation***'.

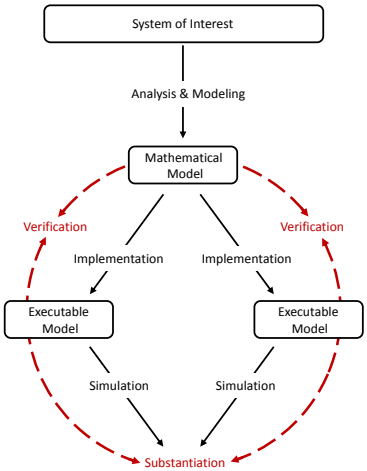

Figure 4.2: **Model verification and substantiation workflow.** The schematic workflow merges two separate model verification and validation processes (cf. Figure 4.1) without the backward reference to the system of interest. This concept evaluates the consistency of two executable models that share the same mathematical model. We term this assessment activity *'substantiation'*. Solid black arrows indicate modeling and simulation activities; dashed red arrows indicate assessment.

- **Substantiation** describes the process of evaluating and quantifying the level of agreement of two executable models.

Thus, *Model verification and substantiation* are processes that accumulate *circumstantial evidence* of a model's correctness or accuracy by a quantitative comparison of the simulation outcomes from validated or non-validated model implementations. The modeling, simulation, and assessment interrelationships are shown in Figure 4.2. In this chapter, we will demonstrate the usefulness of this approach.

REPRODUCIBILITY AND REPLICABILITY In good scientific practice, being able replicate and reproduce results is essential. However, the exact meaning of these two terms may differ between and even within disciplines. For example, in psychology *reproducibility* refers to repeating an experiment and *replicability* means that independent studies yield similar results (Patil et al., 2016). Whereas, for computational experiments, *reproducibility* can be understood as a different team with a different experimental setup as obtaining the same results (ACM, 2016). There are attempts to resolve this ambiguity across disciplines (Barba, 2018), multiple variations of the term's meaning persist Plesser, 2018. Since the terms meaning may also depend on aspects of the application, e.g. in computational applications to what degree outcomes are deterministic, it is advisable to specify the terms explicitly in the

context of their application. Here, we follow the definitions by the *Association for Computing Machinery* (ACM, 2016):

- **Replicability** *(Different team, same experimental setup)*: The measurement can be obtained with stated precision by a different team using the same measurement procedure, the same measuring system, under the same operating conditions, in the same or a different location on multiple trials. For computational experiments, this means that an independent group can obtain the same result using the author's own artifacts.
- **Reproducibility** *(Different team, different experimental setup)*: The measurement can be obtained with stated precision by a different team, a different measuring system, in a different location on multiple trials. For computational experiments, this means that an independent group can obtain the same result using artifacts which they develop completely independently.

To be more terminologically precise, we aim for *results reproducibility* (Goodman et al., 2016).

- **Results reproducibility**: Obtaining the same results from the conduct of an independent study whose procedures are as closely matched to the original experiment as possible.

A replication of an executable model should aim for bit-identicality. Although computers are deterministic, this is not always feasible. For example, the seed of the pseudorandom number generator may not be recorded, or the generated trajectory of pseudorandom numbers may be dependent on the software version or the underlying hardware. However, when a seed is set, successive simulations of replicable models should deliver the same result on the same hardware. A reproduction then constitutes a re-implementation of the model in a different framework, statistically getting the same results. In this terminology, we here replicate a published model and create a model reproduction on the SpiNNaker neuromorphic system. In an iterative substantiation process, we generate verified executable model implementations that are in reasonable agreement.

4.2.2 *Network activity characterization*

When validating neural network activity, one usually cannot expect a spike-to-spike equivalence between the simulated spiking activity and the experimental or simulated reference data. Even for different implementations of the same model, the computation depends on the capabilities and limitations of the computer hardware and the exact details of the computer environment (Glatard et al., 2015). Therefore, the simulation outcomes must be compared statistically to quantify

the similarity level. In the following, we outline several measures of increasing complexity that capture a broad range of network activity dynamics.

Neuron-wise, pairwise, and higher-order measures can, in a sense, be regarded in a hierarchical order. Neuron-wise statistics consider only the single unit activity, irrespective of other units' behavior. Pairwise and higher-order statistics consider how the activity is coordinated within the network. Notably, this conceptual hierarchy does not imply a hierarchy of failure, i.e., agreement on the highest order does not automatically imply agreement on lower-order measures. Therefore, each statistical property is to be evaluated.

4.2.2.1 *Neuron-wise measures*

We characterize the spiking activity of single neurons in the network via the distributions of several neuron-wise measures. The level of network activity can be estimated by the average firing rate

$$FR = n_{sp}/T, \tag{4.1}$$

n_{sp} denoting the number of spikes during an observation interval of length T. The inter-spike intervals are defined by

$$ISI_i = t_{i+1} - t_i, \tag{4.2}$$

where t_i denotes a neuron's ordered spike times. The distribution of ISI_i is used to characterize the temporal structure of the single spike trains. A measure particularly suited to analyze the regularity of the n spike intervals is the local coefficient of variation

$$LV = \frac{3}{n-1} \sum_{i=1}^{n-1} \frac{(ISI_i - ISI_{i+1})^2}{(ISI_i + ISI_{i+1})^2}, \tag{4.3}$$

which equals 1 for a Poisson process (Shinomoto et al., 2003).

4.2.2.2 *Pairwise measures*

We analyze the cross-correlation function

$$R_{xy}(\tau) = \langle x(t)y(t+\tau)\rangle = \frac{1}{N}\sum_{t=1}^{N} x(t)y(t+\tau)\ , \tag{4.4}$$

where $\langle\cdot\rangle$ denotes the temporal average (Tetzlaff and Diesmann, 2010). It quantifies correlations between spike counts of two binned spike trains, $x(t)$ and $y(t)$, for a range of lags τ given N bins. Subtracting the average spike counts $\mu_x = \langle x(t)\rangle$ and $\mu_y = \langle y(t)\rangle$ yields the covariance function

$$C_{xy}(\tau) = \big\langle (x(t) - \mu_x)\,\big(y(t+\tau) - \mu_y\big)\big\rangle = R_{xy}(\tau) - \mu_x\mu_y\ . \tag{4.5}$$

Normalizing the covariance function by the standard deviations $\sigma_x = \sqrt{C_{xx}(\tau = 0)}$ of the processes, one obtains the cross-correlation coefficient function

$$\rho_{xy}(\tau) = \frac{C_{xy}(\tau)}{\sigma_x \sigma_y} \,. \tag{4.6}$$

The Pearson correlation coefficient is given by $\rho_{xy}(\tau{=}0)$ (Perkel et al., 1967). The matrix of correlation coefficients, C, evaluates the non-delayed (i.e. zero-lag) correlation of spikes. The activity on different scales can be analyzed by applying different bin sizes. Here we use binned spike trains on a fine temporal scale (Pearson correlations denoted by CC, using a bin width of 2 ms) and on a coarse scale (Pearson correlations denoted by RC, using a bin width of 100 ms). The correlations on coarser scales are often referred to as rate correlation. In particular, RC is able to capture characteristic population-wide fluctuations of network activity that are often observed on the associated temporal scales (see e.g., the stripy asynchronous irregular state in Voges and Perrinet, 2012).

The model we use in this chapter was originally conceived to exhibit a spatiotemporal arrangement of the spiking activity (polychronous groups) (Izhikevich, 2006). Therefore, we further analyze potential time-lagged correlations. For the cross-correlation coefficient function $\rho_{xy}(\tau)$, we select a bin width of 2 ms and calculate its sum for lags up to 100 ms, corresponding to an interval of $[-\Delta; \Delta]$ bins around 0 with $\Delta = 50$:

$$P_{xy} = \sum_{\tau=-\Delta}^{\Delta} \rho_{xy}(\tau) \tag{4.7}$$

4.2.2.3 *Correlation structure*

Eigenvectors of the correlation matrix capture the correlation structure of network activity (Friston et al., 1993; Peyrache et al., 2010). Consider the eigendecomposition of the symmetric, zero-lag correlation matrix according to

$$C\mathbf{v_i} = \breve{}_\mathbf{i}\mathbf{v_i}, \tag{4.8}$$

where λ_i are eigenvalues and $\mathbf{v_i}$ are eigenvectors. Due to the symmetry of the real-valued matrix C it follows that $\lambda_i \geq 0$ and eigenvectors $\mathbf{v}_i$ are real and orthogonal to each other. A large eigenvalue corresponds to an intra-correlated group of neurons whose activity explains a large amount of variance in the system and relates to dominant features in the correlation structure. The loadings of the corresponding eigenvector $\mathbf{v}_i$ identify the neurons constituting such groups. Consequently, a suitable sorting algorithm, for example, hierarchical clustering, exposes intra-correlated groups as block-like features of the correlation matrix. Here, we use the scipy[1] implementa-

1 RRID:SCR_008058; v1.0.0

tion `scipy.cluster.hierarchy.linkage()` with `method='ward'` and otherwise default settings.

To quantify to which degree the correlation structure of two simulation outcomes (1 and 2) is similar, one may flatten the upper triangular matrices of the correlation matrices C_1 and C_2 into vectors $\mathbf{c}_1$ and $\mathbf{c}_2$, respectively. This omits duplicate entries due to symmetry and the unity auto-correlation on the diagonal. The normalized scalar product

$$0 \leq \frac{|\mathbf{c}_1 \cdot \mathbf{c}_2|}{\|\mathbf{c}_1\| \, \|\mathbf{c}_2\|} \leq 1 \tag{4.9}$$

then constitutes a measure of similarity. A value of 1 denotes two identical vectors and a value of 0 two perpendicular vectors. The order of pairwise correlation coefficients in the two vectors $\mathbf{c}_1$ and $\mathbf{c}_2$ needs to be identical, i.e., the similarity measure is sensitive to the labeling of the neurons. Therefore, it should only be applied to compare two network simulations of the same neuron population. Accordingly, reordering the neuron population of one network statistically decreases the similarity measure of any existing structured correlation matrices while preserving the value for non-structured, e.g., homogeneous, correlation matrices. As a test statistic, the distribution of the normalized scalar product is not known and depends on the distribution of cross-correlation coefficients in $\mathbf{c}_1$ and $\mathbf{c}_2$. The significance of the similarity measure is therefore estimated using surrogate data. The associated null distribution is computed by randomly shuffling the neuron order of one network 10000 times.

4.2.2.4 *Spatiotemporal patterns*

The evaluation of the correlation structure presented so far considers only pairwise measures. Nevertheless, the spiking activity of complex networks may include higher-order interactions. Several methods for the detection of higher-order correlation have been developed in recent years (for a review see Quaglio et al., 2018) that do not make any specific assumption about the underlying connectivity and are thus well suited as statistical measures for model validation. Here, we focus on the SPADE (Spike Pattern Detection and Evaluation) method (Torre et al., 2013; Quaglio et al., 2017). SPADE is a statistical method to detect spatiotemporal spike patterns, i.e. temporally precise spike sequences, including synchronous spiking activity. The method comprises two main steps: a) using Frequent-Itemset-Mining to detect repeated spike sequences in parallel spike trains, and b) selecting the sequences that occur often enough to be significant concerning the null hypothesis of independent firing. The features of the patterns (neurons forming the sequences, number and time of occurrences, lags between the spikes forming the sequence, statistical significance of the pattern) characterize the network activity in terms of higher-order statistics.

4.2.2.5 *Statistical comparison of distributions*

Consider two sample distributions with means μ_i and standard deviations σ_i. Here, such sample distributions represent the neuron-wise or pairwise evaluation of one of the measures described above. According to Hedges (1981), the effect size

$$d = \frac{\mu_1 - \mu_2}{\sigma}, \tag{4.10}$$

characterizes the difference in the mean values where

$$\sigma = \sqrt{\frac{(n_1 - 1)\sigma_1^2 + (n_2 - 1)\sigma_2^2}{(n_1 + n_2 - 2)}} \tag{4.11}$$

is the pooled standard deviation, and the n_i specifies the number of samples entering each distribution. In the case of equal sample sizes, the definition is equivalent to Cohen (1988, p. 67). In the case of multiple simulation runs, we calculate the average effect size of the respective measures, as the simulations are independent, and there is no systematic trend of the measures for the evolving network states. Calculating the effect size assumes that both distributions are Gaussian. Even though this assumption is not fulfilled for every measure, we calculate the effect size as a simple quantification of the difference between the non-normal distributions. Note that for non-normal distributions, a small effect size does not necessarily indicate similarity because there might still be a mismatch in the shape of the distribution. In these cases, additional tests are needed for a complete evaluation. Candidates are the scalar product measure to compare correlation structures and statistical hypothesis tests.

The present work employs hypothesis tests to assess the equality of the means (two-sample Student's t-test) and the equality of the distributions (Kolmogorov-Smirnov test, Mann-Whitney U test). This quantifies the discrepancy in the results by a p-value. The two-sample t-test only applies to normally distributed data, while the latter two tests are non-parametric and applicable to any form of distribution. The Kolmogorov-Smirnov test computes the supremum of the difference between the two cumulative distribution functions. In contrast, the Mann-Whitney U test compares the rank sums of the jointly sorted samples. When applying hypothesis tests, the interpretation of the p-values as a similarity assessment must also consider potential biases and dependencies, e.g., on the sample size and the simulation time (Cohen, 1994).

4.2.3 *Implementation of a validation test framework*

Rigorous validation testing requires that test results are not affected by details of the actual testing procedure. This translates to performing the extraction of test statistics and its evaluation with the same

methods for both data sources entering the test. In a more complex scenario, this also includes finding an appropriate mapping between the data sources, for instance, when comparing a large-scale simulation of spiking activity to experimental data taken from a few electrodes only. Ultimately, validation methodologies should be standardized within the neuroscientific community to ensure consistency of the validation scores across different validation cycles of related models or data sets. The starting point for drafting a joint base for validation testing is formalizing the validation workflow for the individual research domains. For network-level validation of spiking activity data, we created this formalization as the open-source Python module NetworkUnit[2]. All quantitative comparisons of statistical measures in this chapter are carried out in this framework and the workflow to reproduce the findings using NetworkUnit is available online as a Jupyter notebook[3].

NetworkUnit focuses on the statistical comparison of measures characterizing spiking neural network models. It is based on the Python package SciUnit [4] (Omar et al., 2014), which provides a generic basis for the testing of models, employing similar concepts to those of unit testing in software engineering. SciUnit consists of three base classes for models, tests, and scores. The model class defines the model to be validated and, if needed, handles its execution. The test defines which measure is extracted from the model and against which experimental data the model is validated. Finally, the score defines the validation method to be applied and quantifies the result of the validation cycle. Models and tests are connected via their capabilities, e.g., a definition of what types of data output a model provides and what type of data input the test requires to extract its measure. Figure 4.3 schematically depicts the interplay of these components and the class hierarchy for the cases of validation of a model against experimental data or substantiation against another model.

For the analysis presented in this chapter, the components in Figure 4.3 can be understood as follows: the underlying capability is the class `ProducesSpikeTrains` as all analyzed measures are based on the spike times. The SpiNNaker model is implemented as the `sim_model` that is to be validated. It could either be validated against experimental data (`exp_data`), or substantiated against another instance of the model (`sim_model_B`), e.g., the original implementation as illustrated in our worked example. The test statistics we use in `XYTest` are the distributions of the measures presented in Section 4.2.2, e.g., firing rate or correlation coefficient. All these tests involve the comparison of distributions, so they are derived from a corresponding `BaseTest` (and potentially additional base tests). Some statistics, e.g. the correlation coefficient, depend on additional parameters (controlled by

2 https://github.com/INM-6/NetworkUnit; RRID:SCR_016543
3 https://web.gin.g-node.org/INM-6/network_validation
4 RRID:SCR_014528

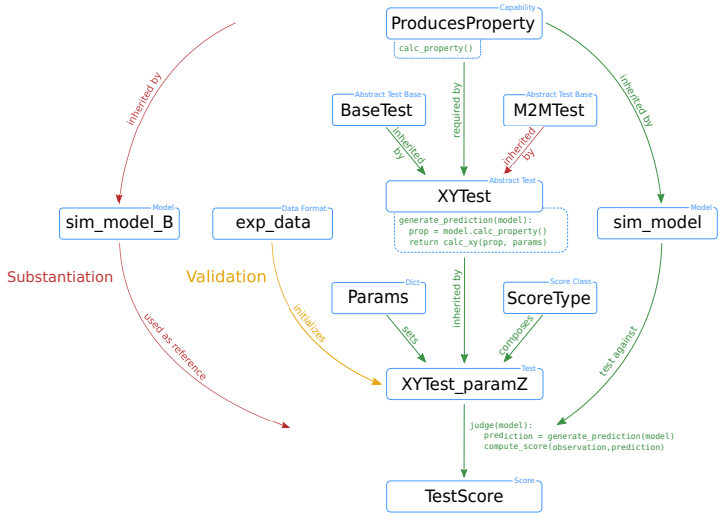

Figure 4.3: Illustration of a typical test design within `NetworkUnit`. The blue boxes indicate the components of the implementation of the validation test, i.e., classes, class instances, data sets, and parameters. Annotated arrows indicate the relation between the boxes. Green arrows show the basic functionality. The difference in the test design for comparing against experimental data (validation) and another simulation (substantiation) is indicated by yellow and red arrows, respectively. Some components' functionality for the test score computation is indicated by pseudo-code. The capability class `ProducesProperty` contains the function `calc_property()`. The test `XYTest` has a function `generate_prediction()`, which makes use of this capability, inherited by the model class, to generate a model prediction. The initialized test instance `XYTest_paramZ` makes use of its `judge()` function to evaluate this model prediction and compute the score `TestScore`. The `XYTest` can inherit from multiple abstract test classes (`BaseTest`), which is, for example, used with the `M2MTest` to add the functionality of evaluating multiple model classes. To make the test executable, it has to be linked to a `ScoreType`, and all free parameters need to be set (by a `Params` dict) to ensure a reproducible result.

`Params`) such as the bin size. In our case, the `ScoreType` are statistical hypothesis tests or the effect size.

The test instance uses spike trains from the model and the experimental data or, as in our case, from the reference model implementation to generate a 'prediction' and an 'observation', respectively. The calculation of features on activity data is performed using the Electrophysiology Analysis Toolkit[5]. Both observation and prediction are passed on to the score class, which evaluates their statistical congruence, e.g., in form of a two-sample t-test. Finally, the judge function of the test instance returns the results, for example, the p-value of the statistical hypothesis test. This design formalizes the generation of the results and makes them reproducible. The modular design of model and test classes enables the reuse of existing tests which facilitates the comparison of results of different models.

Performing a single test for validating a model does not sufficiently capture the model behavior to quantify it and comprehensively document its scientific applicability. Thus, a whole range of validation tests is usually performed, which may, in some cases, differ only in detail or may depend on a parameter. Instead of rewriting the test definition each time, it is more feasible to make use of class-based inheritance as indicated in Figure 4.3 (`BaseTest`→`XYTest`→`XYTest_paramZ`). All specific tests inherit from the `sciunit.Test` base class. They add and overwrite the required functionality, such as generating the prediction by calculating the correlation coefficients from spike trains. Because there may be a lot of different tests making use of correlation coefficients (for example, calculating correlations on different time scales), it is recommended to implement an abstract generic test class to handle correlations first. This abstract test class cannot be accessed explicitly by a user but only acts as a parent class for the actual executable test class, which, e.g., implements the test for a specific choice of bin size. This class-based inheritance guarantees that all tests build on the exact implementation and workflow. Here, we concentrate on model-to-model validation. In this scenario, the test instance compares the prediction of two model instances. For that scenario, we extended SciUnit with a `TestM2M` test class, in which the experimental data (`exp_data`) in Figure 4.3 are replaced by a second model class (`sim_model_B`).

4.2.4 *Simulation of the polychronization model*

To demonstrate the mode substantiation process, we use a minimal spiking network model that capable of exhibiting polychronous groups of spiking neurons, i.e., the polychronization model (Izhikevich, 2006). The model has several non-standard conceptual and implementation features that make it an illustrative example for reproduction on the

5 Elephant: `http://python-elephant.org`; RRID:SCR_003833

SpiNNaker neuromorphic system (Furber et al., 2013). The model exposes essential aspects in the formalization and simulation of neural networks as it produces a rich repertoire of network dynamics. However, we do not evaluate the emergence of polychronous groups, as it has been shown that they result from specific implementation choices and, therefore, are not robustly reproducible (Pauli et al., 2018). The original model is implemented in the C programming language and is available for download from the website of the author[6].

The polychronization model consists of 1000 neurons with four times more excitatory than inhibitory neurons. The corresponding neuron models are specified in (Izhikevich, 2003). Correspondingly, excitatory neurons are parameterized to exhibit regular spiking, and inhibitory neurons to exhibit fast spiking behavior. The neurons are connected randomly with a fixed out-degree of 100, where inhibitory neurons only form connections to the excitatory population. Each excitatory connection is assigned a fixed delay drawn from a discrete uniform distribution between 1 and 20 ms in intervals of 1 ms and all inhibitory connections are assigned a delay of 1 ms. Synaptic weights are initialized with an initial value of 6 for excitatory and -5 for inhibitory connections. The original model uses dimensionless variables, but currents can be interpreted in units of pA. The network is driven by random input realized by an external current pulse of 20 pA injected into one randomly chosen neuron in each time step. The simulation time step is 1 ms, within which multiple intermediate steps are calculated, depending on the implementation (Trensch et al., 2018). The simulated spiking activity in the network modifies the connection weights according to a spike-timing-dependent plasticity (STDP) rule. Synaptic weight changes are buffered for one biological second, then the weight matrix is updated for all plastic synapses simultaneously. We leave out a description of plasticity implementation here because it is irrelevant for the remainder of the analysis as we consider only the dynamics after freezing the learned connectivity matrix and refer to (Pauli et al., 2018) for a detailed description.

Figure 4.4 illustrates the setup of the simulations for the substantiation workflow. Analyzing 5 network states within one simulation process instead of the outcome of multiple different simulations with different random seeds is motivated by the findings of Pauli et al. (2018) showing that the model may converge into two distinctly different activity states. By analyzing the sample activity at different training times within one simulation, this ambiguity problem for the analysis can be bypassed. To generate the network activity data and to save the network states, the following three steps are performed:

1. Execute the C implementation with STDP for five hours of biological time. During this simulation run, save the network state at five points in time t_i, $i = (1, 2, ..., 5)$ after 1, 2, 3, 4, and 5

6 https://www.izhikevich.org/publications/spnet.htm

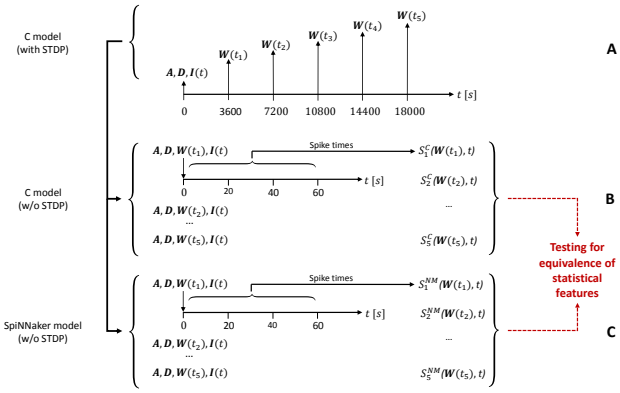

Figure 4.4: **Design of the simulation setup.** Three types of simulations are used in the substantiation scenario. Saving and loading variables is annotated at the corresponding points on respective simulation timelines. **A** Generation of the five initial network states used to initialize the following simulations. At the start ($t = 0$ s) of running the C implementation of the polychronization model (with STDP), the connectivity matrix A and delay matrix D are saved. At the following times t_i, the weight matrix $W(t_i)$ is saved. The random input stimulus to the network $I(t)$ is recorded for the simulation duration. **B** Generation of data from the five simulations of the C implementation (without STDP) based on the random input $I(t)$ and the five sets of initial conditions (A, D, $W(t_i)$) recorded in (A), respectively. The network spiking activity $S_i^{\mathrm{C}}(W(t_i), t)$ is recorded for 60 s. **C** Identical setup as in (B), but for the SpiNNaker implementation without STDP, where $S_i^{\mathrm{NM}}(W(t_i), t)$ denotes the simulation result. The data from (B) and (C) are subject to validation testing based on their statistical features (red dotted lines). Figure amended from Trensch et al. (2018).

hours. The network state is defined by the weight matrix $W(t_i)$ containing the current strength of each synapse, the connectivity matrix A, and the delay matrix D. Additionally, record the first 60 seconds of the random series of neurons to which the external stimulus is applied ($I(t)$, Figure 4.4A).

2. Switch off STDP in the C implementation. Re-initialize the network model with A, D, I, and the respective $W(t_i)$ for the five simulation runs $i = (1, 2, ..., 5)$. In each run record the network spiking data S_i^{C} over 60 seconds (illustrated in Figure 4.4B).

3. Repeat step (2) with the implementation on the SpiNNaker neuromorphic system (NM) of the polychronization model to obtain the spiking data S_i^{NM}.

The spiking data S_i^{C} and S_i^{NM} are used in the following statistical analyses. For the sake of simplicity, only the excitatory population is considered in the following validation. However, the results for the inhibitory population do not differ qualitatively.

4.2.5 *Iterative model improvements*

In the process of iterative verification and validation, the development of the model implementations can be segmented into three main iterations. For a more detailed account of the verification and model implementation, see (Trensch et al., 2018), whereas the statistical evaluation and comparison on each iteration step is presented in the Results of this chapter. First, the original C implementation of the polychronization model (Izhikevich, 2006) underwent a source code verification maintaining bit-wise replicability of the simulation outcome. The polychronization model was reproduced on the SpiNNaker neuromorphic system using the Izhikevich neuron model implementation provided by the SpiNNaker software stack, using the Explicit Solver Reduction (ESR) implementation of the dynamics described in Hopkins and Furber (2015).

The second iteration performed verification activities to resolve discrepancies in the simulated activity statistics. As a result, the ODE solver implementation for the SpiNNaker and the C model was replaced by a semi-implicit fixed-step size forward Euler scheme. Additionally, the revised implementations include a precise threshold detection and an optimized fixed-point representation for some critical calculations to improve the numerical precision of computations. The third iteration, initiated by a mismatch in the measured LV measures, corrects an implementation issue causing a small systematic lag in spike timing. Thus, each iteration constitutes a refinement of the implementation and subsequent verification and substantiation (using NetworkUnit). Table 4.1 summarizes the specific changes in each itera-

	C model	SpiNNaker model
iteration I	uses a semi-implicit fixed-step size forward Euler ODE-solver with step size 1 ms	(i) uses the SpiNNaker Explicit Solver Reduction (ESR) implementation of the Izhikevich neuron model
	remains unchanged	(ii) uses Izhikevich's algorithm for the neural dynamics
		(iii) uses a more exact fixed-step size forward Euler ODE-solver with step size 1 ms
iteration II	uses a 1/16 ms step size and more precise detection of threshold crossing	uses a 1/16 ms step size and more precise detection of threshold crossing, applies fixed-point conversion for critical calculations
iteration III	remains unchanged	resolves an implementation issue with the threshold detection

Table 4.1: Summary of the iterative development steps of the model implementations, based on a replication of Izhikevich's original implementation. Steps (ii) and (iii) represent incremental improvements between iterations I and II. (ODE = ordinary differential equation).

tion. The model source codes and simulation scripts are available on GitHub[7].

4.3 RESULTS

Although digital computer simulations are deterministic, given their input, achieving the equivalence of two neural network model implementations is not straightforward. This holds true for our application scenario, comparing implementations of the Izhikevich polychronization model on the SpiNNaker digital neuromorphic system and a cus-

7 https://github.com/gtrensch/RigorousNeuralNetworkSimulations (doi: 10.5281/zenodo.1435831)

tom C implementation. The SpiNNaker system uses 32-bit fixed-point numerics, while the C implementation employs floating-point numerics. Furthermore, Pauli et al. (2018) expose that the model dynamics are sensitive to small changes in model parameters and numerics. Therefore, we do not expect a spike-to-spike equivalence between the two model realizations. Instead, the model vs. model comparison rests on statistical characterization of the simulated network activity. In Section 4.3.1, we show that the application of validation tests guides the model development progress. Section 4.3.2 demonstrates the importance of incorporating multiple characteristic measures in a comparison application. Finally, Section 4.3.3 builds a comprehensive substantiation evaluation of the SpiNNaker model implementation against the C implementation.

4.3.1 *Network activity comparisons guide model development (iteration I)*

Applying continuous statistical comparisons during the iterative process of model implementation benefits the modeler. We illustrates the guided improvements in the development process of the SpiNNaker implementation in three iterative steps denoted by i-iii in Figure 4.5 (see Table 4.1.) The results are from 60 s of simulated data starting from the network state after 5 hours of biological time. Figure 4.5A shows a raster plot of the spiking data for the C implementation and the three consecutive steps of the SpiNNaker implementation. Step i represents the initial SpiNNaker implementation using an Explicit Solver Reduction (ESR) algorithm for the Izhikevich neuron dynamics. Step ii reimplements the neuron dynamics described in Izhikevich (2006). Step iii improves this algorithm by applying a fixed step size forward Euler method. The C simulation exhibits strong fluctuation in the population activity indicated as vertical stripes in the raster plot. The three steps of the SpiNNaker simulation successively show qualitatively more similar dynamics. As expected, there is no spike-to-spike equivalence between the implementations.

We evaluate the statistical agreement between the measure distributions (FRs, LVs, CCs) between the C and SpiNNaker simulations using the effect size defined in Section 4.2.2.5. The corresponding results in Figure 4.5B and C visually and quantitatively show an improving agreement with each step of the SpiNNaker implementation development. The effect size between the distributions declines consistently for all measures with each iteration step. Despite the improved agreements, the discrepancy in the distributions of firing rates is still considerable, and there remains a shape mismatch between the distributions of CCs (step iii, Figure 4.5B). The distribution of the sum of the cross-correlation coefficient of the SpiNNaker simulation (step iii, Figure 4.5D) is broader and features a longer tail than the distribution obtained from the C simulation. The corresponding correlation matrix

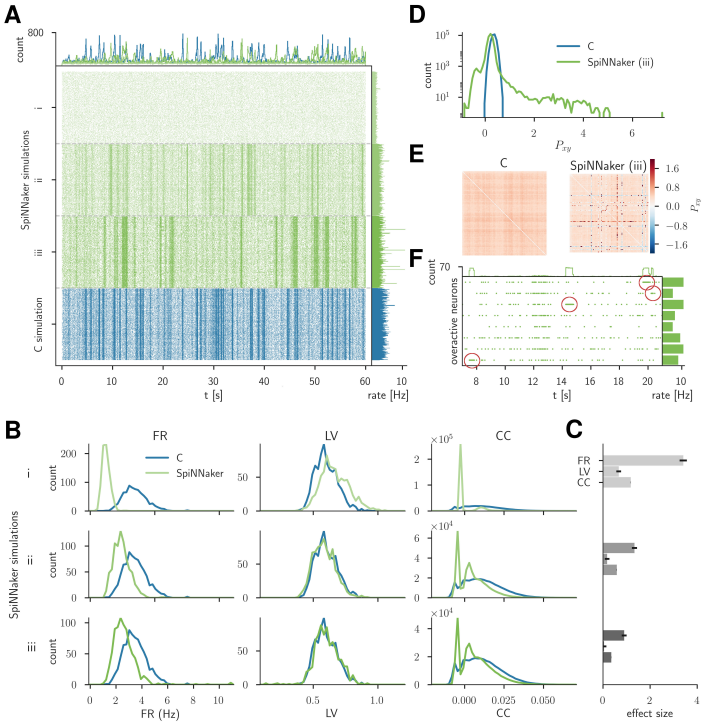

Figure 4.5: **Comparison of the C simulation with simulations of three consecutive stages of the SpiNNaker implementation. (A)** Raster plot of the spiking network activity (800 excitatory neurons) of the C simulation (bottom, blue) and three stages of the SpiNNaker implementation: i, ii, and iii (top, shades of green). The top and right histograms show the population spike counts in 60 ms bins and the mean firing rates, respectively. **(B)** Distributions of firing rates (FR, left), local coefficients of variation (LV, middle), and correlation coefficients (CC, right) for the C and SpiNNaker simulations. Each row (subsequent implementation steps: i, ii, iii) represents a specific SpiNNaker simulation (green) that differs in the underlying neuron model implementation. Data shown for the C simulation (blue) are identical in the three rows. **(**C) The difference between the distributions is quantified by the effect size with error bars indicating the 95% confidence interval. In step iii, the effect sizes for the FR, LV, and CC measures are 0.90, 0.05, and 0.36, respectively. **(D)** Distributions of the sum of the cross-correlation coefficient (P_{xy}, Eq. 4.7) in logarithmic representation for C and SpiNNaker (implementation step iii). **(E)** Color-coded correlation matrices for the sum of the cross-correlation coefficient in implementation step iii. The symmetric matrices display results for the subset of 100 excitatory neurons with the highest spike rates in the SpiNNaker simulation. **(F)** Raster plot of 8 overactive neurons in the SpiNNaker simulation (implementation step iii) showing episodes of 1 kHz spiking (emphasized by red markers). The top and right histograms show the population spike counts in 60 ms bins and the mean firing rates for the entire recording.

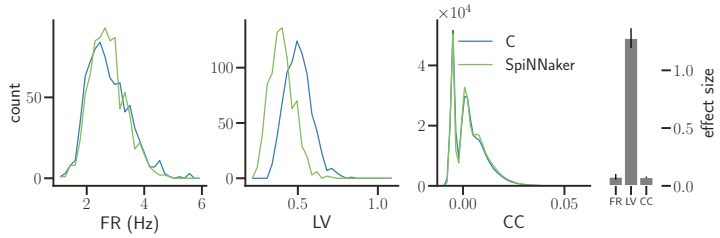

Figure 4.6: **Comparison of spiking activity measures after model refinement.** The panels show from left to right the distributions of FR, LV, and CC of the C and SpiNNaker simulation in iteration II Section 4.2.5, for the network state after $t_5 = 5$ h (same display as in Figure 4.5B). The histogram on the right visualizes the corresponding effect sizes (mean and standard deviation across all five network states $t_1, t_2, \ldots, t_5$). The numerical values are FR: 0.077 ± 0.025, LV: 1.28 ± 0.086, and CC: 0.074 ± 0.006 respectively.

(Figure 4.5E) for SpiNNaker reveals large and small outlier values that are arranged in horizontal and vertical lines. This visual line structure points towards individual neurons that are highly correlated or anti-correlated (within a ± 100 ms delay window) to a large number of other neurons.

Further investigation reveals 8 specific overactive neurons that exhibit larger firing rates than the rest of the population. Indeed, these neurons show occasional periods of several hundred milliseconds with firing rates of around 1 kHz (Figure 4.5F). In a subsequent analysis and source code verification, we trace the excessive firing rates to an overflow of a fixed-point variable in calculating the membrane potential. This example showcases how validation testing can reveal simulation mismatches and provides valuable information guiding further verification and development steps.

4.3.2 *Different characteristic measures provide distinct network assessments iteration II*

In iteration II, the SpiNNaker and the C code are refined with an improved forward Euler ODE solver, precise detection of threshold crossings, and a more accurate fixed-point representation on SpiNNaker. Based on these versions of the executable models, we compare the statistical properties of simulated spiking activity.

The distributions of the mean firing rates and correlation coefficients (Figure 4.6) show an improved agreement to the previous iteration (Figure 4.5B, bottom row). However, despite the similar firing rates and correlations, the spiking activity in the SpiNNaker simulations is considerably more regular, as indicated by a clear shift towards lower values in the LV distributions. Thus, the implementation refinements

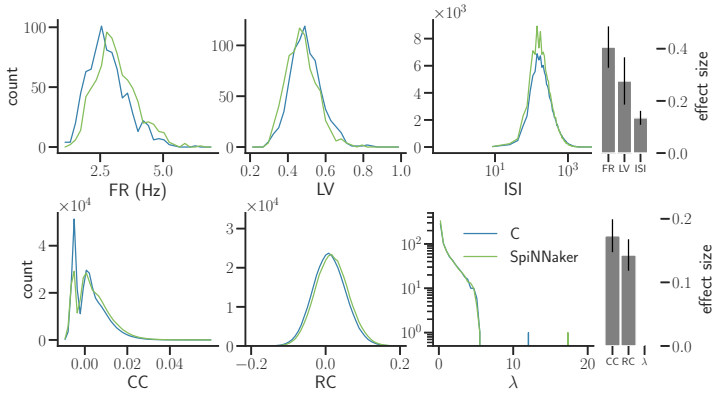

Figure 4.7: **Distributions of characteristic measures after the final step of model refinement.** The **top row** shows the single neuron statistics FR, LV, and the ISIs (same display and data specification as in Figure 4.5B). The histograms of ISIs are displayed using semi-logarithmic scaling. The **bottom row** shows pairwise statistics and network properties: the CC using 2 ms bins, the rate correlation (RC) using 100 ms bins, and the eigenvalues (λ) of the RC matrices using a logarithmic scaling of the vertical axis. **Right:** effect size using the same display as in Figure 4.6. The effect sizes are FR: 0.41 ± 0.08, LV: 0.28 ± 0.09, ISI: 0.14 ± 0.03, CC: 0.17 ± 0.03, RC: 0.14 ± 0.02, and λ: $< 8 \times 10^{-17}$, respectively.

improved two statistical measures while worsening the third. So evidently, any validation or substantiation process needs to consider multiple statistics.

4.3.3 *Higher-order network properties extend neuron-wise and pairwise evaluations iteration III*

In the third iteration of the model implementations, we correct the threshold detection algorithm of the SpiNNaker implementation while the C simulation remains unchanged. The agreement of the three activity measures FR, LV, CC improved substantially in this iteration, as measured by the corresponding effect sizes. Therefore, we include additional characteristic measures to extend the detail and scope of the comparison.

Figure 4.7 shows the previous and additional (ISIs, RC, eigenvalues λ of the RC matrix) distributions. According to the interpretation of Cohen (1988), the comparisons of all six measures exhibit effect sizes of small to medium size. With respect to the previous iteration (Section 4.3.2), the effect sizes for the FR and CC increases slightly. However, the overall agreement increases because of the improved match of the LV distributions. All measures show a relatively good

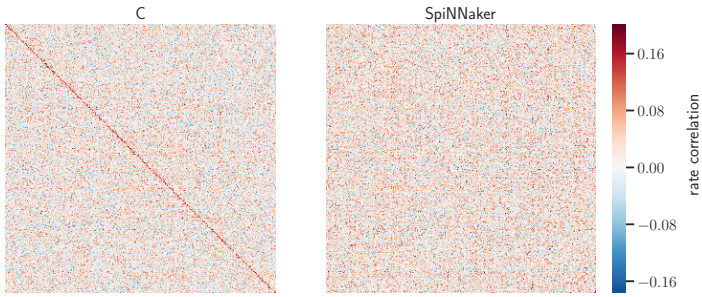

Figure 4.8: **Rate-correlation matrices for the C and SpiNNaker simulations**. Matrix elements show the RCs (color bar) of all pairs of the 800 excitatory neurons, computed from 60s of data with the network state after 5 hours. The order of the neurons in both symmetric matrices is determined by hierarchical clustering (Ward variance minimization algorithm, Section 4.2.2.3) of the C matrix. Auto-correlations are set to 0 not to disrupt the color scale.

agreement between the two model realizations. The remaining discrepancies are small but systematic shifts in the FR, CC, and RC distributions, indicating that SpiNNaker tends to exhibit slightly more correlated spikes. In accordance with this observation, the largest RC eigenvalue for the SpiNNaker model exceeds the one for the C model. A larger first eigenvalue indicates that more variance in the rate correlations is captured by one principle activity mode. Thus, the first eigenvalue difference suggests that, in SpiNNaker, the intermittent population increases are larger amplitude as compared to the C simulation (see, e.g., the oscillations described by Bos et al., 2016).

We test the six measure distribution pairs (for all 5 network states) for statistical equivalence using the non-parametric Kolmogorov-Smirnov test, the Mann-Whitney U test, parametric Student's-t test to the approximately normal distributed measures (FR, LV, RC, log(ISI)). All tests reject their null hypotheses with p-values below a 5% significance level (without correction for multiple comparisons), except the eigenvalue distributions with p-values between 0.17 and 0.96 for the 5 network states. Thus, all but the eigenvalue measure differ significantly between the SpiNNaker and C model implementation, although the overall effect sizes of these differences can be considered small.

Following the univariate comparison of the characteristic measures of the spiking network activity, we look closer at the organization of correlation in the network simulations. Figure 4.8 displays the rate-correlation matrices of all excitatory neurons for the C and SpiNNaker simulation. The C matrix is hierarchically clustered so that pairs and groups of highly correlated neurons are arranged next to each other, visible by high correlation values near the diagonal. The same neuron

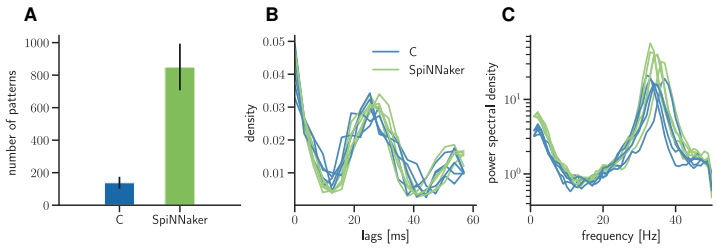

Figure 4.9: **Frequency and structure of spatiotemporal spike patterns. (A)** Bar diagram of the number of patterns detected using the SPADE method (Quaglio et al., 2017) in the two simulations. Displayed are the mean and standard deviation of the results for the 5 network states. The spike times of all 800 excitatory neurons are discretized by 3 ms bins. From the detected spike sequences, we only consider sequences that repeated $>= 3$ times, contain at least 5 spikes, and are shorter (first to last spike) than 60 ms. **(B)** Normalized distributions of the time lags between any two spikes involved in one of the patterns. The results for each of the 5 network states are displayed as a separate distribution. **(C)** Power spectra of the population activity in each network state. The spectra are calculated by Welch's method with a 100 Hz sampling frequency and a 1 Hz frequency resolution (window overlap: 50%). The C and SpiNNaker simulation data in all panels are indicated in blue and green, respectively.

ordering is applied to the SpiNNaker matrix, which does not show a similar diagonal line of high correlation values. Thus, by visual inspection, the two simulations seem not to share the same highly correlated neuron groups. However, we further quantify the similarity of the correlation structure using the normalized scalar products of the vectors containing the RC matrix (Section 4.2.2.3). The results for the 5 network states range from 0.176 to 0.209. A significance assessment with 10000 shuffled surrogate matrices results in corresponding scalar products between 0.081 and 0.108. Thus, the observed similarity score for of the two simulations is $>$ 43 standard deviations away from scores of independent surrogate matrices, indicating a similarity of the correlation structures beyond chance. In Chapter 5, we develop a corresponding null-hypothesis significance test following the idea of statistically comparing the structure of correlation matrices.

Furthermore, we analyze spatiotemporal spike patterns (STPs) as a higher-order statistic representing potential dynamic signatures of the underlying network connectivity. The STPs are detected with activities with SPADE method (Quaglio et al., 2017), considering all repeated spike sequences irrespective of their significance. Figure 4.9A&B summarizes the total number of patterns and the temporal lags between the spikes forming a specific STP. While ag distributions are qualitatively similar in both simulations, we detect considerably more patterns in

the SpiNNaker simulation. Furthermore, the power spectrum of the spiking activity pooled across all neurons shows a clear peak around 35 Hz for both simulations (Figure 4.9C). This peak corresponds to the dominant number of lags around 27 ms in the distribution of patterns' lags (Figure 4.9B). These observations provide a coherent picture of the SpiNNaker simulation behavior. The large number of spatiotemporal patterns can be explained by the simulation's larger average firing rate and correlation values, and a larger power around 35 Hz.

4.4 CONCLUSION

QUANTITATIVE CHARACTERIZATIONS OF NETWORK ACTIVITY In this chapter, we demonstrate how not only the explicit model parameters influence the spiking activity of a network model but also the details of the exact implementation and properties of the simulation engine affect the simulation outcome. Consequently, simulating the same model on two different simulation engines, a C-based simulator and the SpiNNaker neuromorphic hardware, results in simulated activities that are not identical in their spike times. Here, we investigate the mismatch in spike timings, going beyond the qualitative description, and formalize a systematic approach for the quantitative statistical evaluation of (dis)similarity (cf. Section 1.3). In the characterization of the spiking network activity, we incorporate measures of varying complexity: univariate measures (e.g., firing rates), pairwise measures (e.g., correlation coefficients), and higher-order measures (e.g., spike patterns). Each measure of the spiking statistics reflects only a limited aspect of network activity. Therefore, the comprehensive network activity characterization is enriched by including multiple measures to capture a broad range of network dynamics.

VALIDATION AND MODEL VS. MODEL COMPARISONS We built on a general framework of model validation to adapt the concept and terminology to network models of neural systems. Further, we extend its application beyond the statistical model vs. experiment comparisons to establish trust in the model to a more general framework of comparability (cf. Section 2.2). Direct model vs. model comparisons have the benefit of not being limited by data availability and can incorporate a broad dynamical range, different network states, and edge cases of the model parameterization. Such a comparison approach can serve many purposes depending on the models' relation and their differences. For example, it can measure the influence of changes to the connectivity (see Chapter 5), evaluate the robustness to variations of the parameters or input, or it could also establish trust in a model by proxy when the other model is already well validated against experimental data. Here, we use the model vs. model comparison approach to not directly evaluate the mathematical model description (using the polychronization

model by Izhikevich (2006)) but its implementation and underlying simulator engine. In order to not confuse the interpretation of this validation-like approach, we use the term *substantiation*.

TAKE-AWAYS FROM THE SIMULATOR SUBSTANTIATION The showcased application of the substantiation approach to comparing the C and SpiNNaker illustrates multiple aspects of network activity characterization, comparability, and model development that can be generalized to other comparison scenarios. The comparison of simulated population dynamics initially exposed an artifact in the numeric processing on SpiNNaker, causing unreasonably bursting neurons (Section 4.3.1). Iterative validation testing can guide the model development process by uncovering mismatches of the characteristic measures and can therefore complement the model verification. Even before fixing all the eventual issues in the model implementation (e.g., refinement of the threshold detection algorithm), the firing rates and correlation coefficients already showed a good agreement between the two simulations (Figure 4.6). However, the LV regularity measure still indicated a considerable mismatch. This example illustrates the need to have multiple different measures to have the most accurate representation of all aspects of the network activity. Furthermore, it shows that there is not necessarily a conditionality between higher and lower complexity measures. In this case, the statistics of the pairwise correlations are not informative about the statistics of the univariate LV measure. After the final adjustments, all our characteristic measures (firing rate, LV regularity, inter-spike intervals, correlation coefficients, rate correlation, and correlation eigenvalues) agree between the two simulations with effect sizes < 0.41 (Section 4.3.3). However, there are different levels of agreement. Additional analyses on spatio-temporal patterns in spiking activity reveal a considerable mismatch in the number of detected patterns (Figure 4.9). Furthermore, hierarchical clustering of the correlation matrices of the respective simulated spiking activities shows that most intra-correlated groups of neurons do not involve the same neurons in the two simulations (Figure 4.8). Even with a good agreement of multiple activity measures, there can still be additional aspects to the complex and non-linear neural network dynamics, with even minor details potentially causing significant discrepancies. For the remaining discrepancies in the SpiNNaker simulation, we suspect that the fixed-point arithmetic of the system (in contrast to the floating-point arithmetic in the C simulation) causes a slightly increased tendency for spike synchronization resulting in only minimal shifts of the correlation distributions but a more considerable difference regarding the occurrence of spatio-temporal patterns.

ACCEPTABLE LEVEL OF AGREEMENT As an outcome of the comparison approach, the statistical similarity evaluation of all character-

istic measures provides a quantitative assessment of the agreement between the compared entities. However, whether the resulting similarity quantification represents an acceptable agreement depends on the requirements of the intended application of the respective models. Conversely, the obtained similarity assessment defines the model's accuracy and therefore determines its reasonable applicability scope. Therefore, our study determines that the SpiNNaker implementation is sufficient for applications requiring an accurate representation of the spiking statistics in Izhikevich's polychronization model. However, it can not (yet) fulfill strong statistical requirements, including, for example, the equal number of patterns found with the SPADE method or the statistical equivalence of the calculated measure distributions (as assumed by standard null hypothesis significant tests). Whether a quantified level of agreement is acceptable can also be evaluated concerning the limitations and variability of the underlying system. In the biological reality, there are many sources of variability (see, e.g., Arieli et al., 1996; Mochizuki et al., 2016; Riehle et al., 2018, for trial-to-trial and subject-to-subject variability) so that we can typically not expect an exact agreement of the activity to either other experimental recordings or model simulations. Thus a model can be considered accurate in describing an activity feature when it reproduces the activity feature within the bounds of its observed variability. Similarly, we can use the intrinsic limitations of the simulation technology (e.g., the SpiNNaker neuromorphic hardware and its software stack) as a reference to judge the agreement level.

REPRODUCIBLE VALIDATION TESTS WITH NETWORKUNIT In this chapter, we outlined the concept of network-level validation and how the classical validation approach can be extended for other research applications employing a statistical comparison of network descriptions. Whether we perform a validation, calibration, or substantiation involving experimental data or model simulations, the practical steps for the statistical comparison are the same. Therefore, we developed the Python module NetworkUnit to provide the mechanics and method implementations to perform such quantitative comparisons. NetworkUnit is based on the SciUnit framework that is designed for general scientific model validation (Omar et al., 2014). NetworkUnit aims to provide a battery of tests applicable to compare network activity from spiking neural network models. All tests in this chapter have been implemented with NetworkUnit and are replicable with the provided open-access simulation data. Besides the example use-case presented here, another demonstration of how formalized validation testing can support model development tasks is described by van Albada et al. (2018) successful porting of the cortical microcircuit model (Potjans and Diesmann, 2014) to SpiNNaker. Formalizing validation tests and comparison processes with open-software tools and standard

tests is a pivotal step to increase the confidence in models developed by the neuroscience community. It ultimately leads not only to more replicability but also to true reproducibility of scientific findings.

FURTHER APPLICATION SCENARIOS The presented approach for the comparison of neural network activity can be adapted to a range of other scenarios, including: the calibration and validation of models against experimental data, which we address in Chapter 6; the evaluation of connectivity variations of the same model which we employ in Chapter 5; measuring the influence of the model parameters' numeric precision as investigated in the related study by (Dasbach et al., 2021); or the comparison of different experimental data sets as we present in Chapter 7 and Chapter 8. Notably, network-level comparison testing is not restricted to spiking activity. Models and experimental data can be evaluated based on continuous activity signals (e.g., the LFP) using the same approach and testing framework. The statistical testing structure is not dependent on the network description level (Section 1.2), so it can be applied to a network model of different scope and granularity and experimental measurements of different resolution (e.g., MEG, EEG, or ECoG).

5

NETWORK ACTIVITY AND CONNECTIVITY COMPARISONS VIA EIGENVECTOR ANGLES

This Chapter is based on the following publication:

- Gutzen, Robin, Sonja Grün, and Michael Denker (2022). "Evaluating the Statistical Similarity of Neural Network Activity and Connectivity via Eigenvector Angles." In: *Biosystems* 223, p. 104813. DOI: 10.1016/j.biosystems.2022.104813

Author contributions:

- RG, MD conceptualized the study.
- MD, SG supervised the project.
- RG established and implemented the methodology and models.
- RG performed and visualized the analyses.
- RG wrote the initial draft of (Gutzen et al., 2022) and wrote this chapter.
- All authors contributed to the writing of the published manuscript.

5.1 INTRODUCTION

CONNECTIVITY AND ACTIVITY DESCRIPTION LEVELS OF NETWORKS Neural networks, as introduced in Section 1.2, can be described on two levels of interaction. On the level of the synaptic connections, we can describe the physical coupling strength between pairs of neurons and the resulting efficiency with which they influence their respective activity, i.e., transmit information. On the activity level, we can describe the functional coupling between pairs of neurons via correlation-based evaluations of their respective (spiking) activity. There is a relation between these two levels of descriptions (Section 1.5). However, this relation is not trivial to determine (Sporns and Tononi, 2001), and object to change due to the external input (Stevenson et al., 2008; Pernice et al., 2011), network state (Olcese et al., 2016), neural dynamics (Ostojic et al., 2009; Aljadeff et al., 2015),

and network dynamics (Van Bussel, 2011; Curto and Morrison, 2019). Furthermore, in biological networks, it is technologically difficult to record in-vivo both the connectivity and the activity, so there is only little biological ground-truth data to determine their relationship (Gerhard et al., 2013, one exception is e.g.). Therefore, many efforts focus on evaluating computational models where we have full access to connectivity and activity.

RELEVANT CHARACTERIZATIONS ON THE DESCRIPTION LEVELS The analysis of networks on the connectivity level and the activity level typically requires quantitative measures that characterize the relevant aspects of the network (Section 1.3). Pairwise and higher-order measures that quantify the interactions between the nodes (i.e., neurons) of the network are particularly relevant for evaluating the relationship between the connectivity and activity description level. Examples include activity correlation measures, i.e., functional connectivity (Bullmore and Sporns, 2009; Eickhoff et al., 2010), the connectivity eigenspectrum (Zhou et al., 2009; Dahmen et al., 2019), and either connectivity or correlation-based graph measures (Pernice et al., 2011; Curto and Morrison, 2019; Haber and Schneidman, 2020).

COMPARABILITY OF THE NETWORK CHARACTERIZATIONS Based on such network characterizations, we can evaluate the similarity between realizations of networks (Section 1.4). We can compare two network realizations on the connectivity and activity description level. Moreover, by combining the two comparison levels approaches, we can investigate how differences in the network connectivities correspond to differences in their activities. Standard validation-type comparisons approaches are based on the univariate comparison of distributions of characteristic measures (as explored in Chapter 4). For single-neuron measures, this approach is straightforward, as the values are statistically independent and can be univariate compared by appropriate two-sample tests. Popular tests are, for example, the Student's t-test (Student, 1908) or the Kolmogorov-Smirnov test (Hodges, 1958), and, even though they are not formally statistical tests, comparative scores such as the Kullback-Leibler divergence (Kullback and Leibler, 1951) and the effect size (Cohen, 1988).

STATISTICAL COMPARISON OF PAIRWISE MEASURES However, as we are interested in the interactions between neurons, we aim to evaluate pairwise measures (i.e., correlations) for which a univariate representation is less suitable. Although pairwise measures may also be represented and compared via the distribution of the individual values (i.e., entries of the correlation matrix), the sample values are not statistically independent, and information is lost in the reduction process. Instead, comparing matrices representing the pairwise

measure in a meaningful way is not trivial, and only a few statistical tests are available (Flury, 1988; Krzanowski, 1990; Calsbeek and Goodnight, 2009; Box, 1949). Approaches to perform comparisons between matrices that do not involve a statistical test include using the correlations between matrix elements, the Euclidean distance, or the geodesic distance (Venkatesh et al., 2019).

CONSTRUCTION OF A STATISTICAL TEST FOR PAIRWISE MEASURES Here, we develop a statistical test to compare pairwise measures of two networks. In particular, we derive the method for symmetric matrices of Pearson correlation coefficients and asymmetric matrices of synaptic weights. Our approach uses the matrix's eigenspectrum to characterize the matrix's structure, including the composition of correlated or interconnected groups. The eigenvectors of a matrix span the space in which its element values are represented most naturally. The first eigenvector points along the direction of the largest variance in the data, which could represent, for example, a strongly correlated group of neurons. The second eigenvector points towards the largest variance within the orthogonal subspace, i.e., a separate structural aspects of the network. The corresponding eigenvalues quantify the variance along these axes.

Therefore, we evaluate the similarity between two matrices by quantifying the alignment of the respective eigenvectors with the angles between them, which we term "eigenangles". Small angles between pairs of eigenvectors indicate similar underlying network structures. We define a similarity score from the ensemble of eigenangles weighted with the respective eigenvalues and derive how the similarity behaves under the null hypothesis of independent random matrices. With the assumptions of large matrices that are defined in the same space (i.e., calculated from the same set of neurons), we can thus construct a statistical test and evaluate the similarity between two networks via a quantitative eigenangle score and a corresponding p-value.

CHAPTER OUTLINE In the following, we build up a theoretical basis for the statistical test by describing the behavior of angles for high dimensional random vectors and the special case of eigenvectors. We construct the "eigenangle test" by first considering symmetric correlation matrices for spiking network activity and then extending our approach to asymmetric connectivity matrices of specific network types. We characterize the statistical test by applying it to calibration scenarios of stochastic and simulated neural network activity, as well as experimental spiking activity data. Finally, we use the eigenangle test to explore the relationship between network connectivity and activity by comparing network model realizations with specific connectivity modifications on the connectivity and activity description level (Figure 5.1).

The idea and an initial version of the eigenangle test for correlation matrices were already presented in Robin Gutzen's Master thesis. There the test was applied to an example of stochastic spike data. The presented work extends this work by overhauling and refining the definition of the eigenangle score, the null hypothesis, and the analytic null distribution. Also, the corresponding test implementation was rewritten accordingly and made more numerically robust and considerably more efficient. This improved test version allows for the extension to asymmetric connectivity matrices and enables the efficient evaluation of the here presented application scenarios.

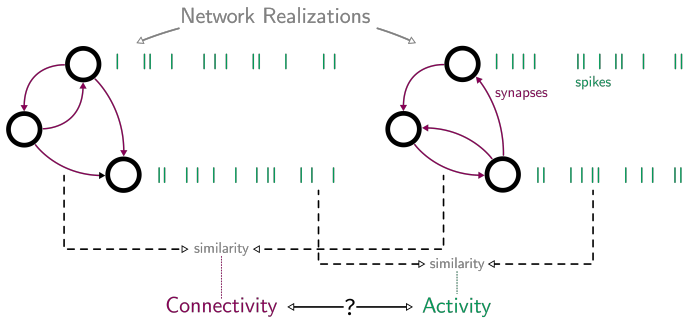

Figure 5.1: **Schematic approach for inferring connectivity-activity relation from comparisons of network realizations.** Two realizations of the same type of spiking network model may differ in their exact synaptic connectivity (magenta). For example, when rewiring some or all of the synapses. Consequently, the simulation of the two realizations may produce different sets of spiking activity (green). By evaluating the similarity of the respective activity dependent on the similarity of the respective connectivity, we can measure the impact of the specific rewiring on the connectivity-activity relationship.

DATA AND CODE AVAILABILITY All code and data for the simulations, analysis, and visualizations are stored in the gin repository `https://gin.g-node.org/INM-6/eigenangles`. The experimental data used in Figure 5.10 is available in the data publication of Brochier et al. (2018). The eigenangle test to compare correlation matrices will be provided as part of the v0.2 release of the Python validation test library NetworkUnit `https://github.com/INM-6/NetworkUnit` (RRID:SCR_016543).

5.2 METHODS

5.2.1 *Statistical eigenangle test for correlation matrices*

Our objective is to quantify the similarity of a pair of matrices using a statistical test based on the angles between their eigenvectors. Throughout this section, we will assume the case of comparing two symmetric correlation matrices as illustrated by the concrete use case of correlations between N neuronal spike trains. We generalize the concept to a more general class of matrices, such as graph asymmetric adjacency matrices, in Section 5.2.2.

Let $\{\mathbf{v}_i^A, i \in 1, \ldots, N\}$ denote the set of ordered, normalized eigenvectors of a matrix $\mathbf{A}$, such that $\lambda_i^A \geq \lambda_{i+1}^A$ for each corresponding eigenvalue λ_i^A. We define the i-th *eigenangle* ϕ_i as the $\phi_i = \measuredangle(\mathbf{v}_i^A, \mathbf{v}_i^B) = \arccos(\mathbf{v}_i^A \cdot \mathbf{v}_i^B) \in [0, \pi]$ between the eigenvectors $\mathbf{v}^A$ and $\mathbf{v}^B$ of matrices $\mathbf{A}$ and $\mathbf{B}$. Thus, we consider the angles between pairs of the i-th ordered eigenmode of each of the two matrices.

The underlying assumption for the eigenangle test is that a small angle between two eigenvectors indicates the similarity of the corresponding eigenmode, whereas a near-orthogonal angle indicates a discrepancy. Following this assumption, we will derive criteria to identify similar eigenmodes based on the eigenangles ϕ_i by calculating their expected distribution under the assumption of independent activity. Let us consider a set of N-dimensional normalized random vectors, i.e., vectors that are uniformly distributed over the N-dimensional unit sphere. Here, by assuming matrices $\mathbf{A}$ and $\mathbf{B}$ to be correlation matrices containing the correlation coefficients between individual neurons in the network, the dimensionality N of this space is equal to the number of recorded neurons. Contrary to the intuition suggested by a unit circle in $N = 2$, in higher dimensional spaces ($N > 2$), the probability distribution $f_\measuredangle(\phi)$ of angles between two random vectors is not uniform. To illustrate this dependency, let us imagine a point at the pole of the unit sphere (i.e., the intersection of the sphere with the first axis) and a randomly chosen second point on the unit sphere. For $N = 2$, each section of the circle $d\Theta$, located under angle θ from the pole, is equally likely to contain the second point. Thus, the probability distribution of the angle between two arbitrary vectors will be uniform. However, for $N = 3$, the ring section's surface $d\Theta$ under the polar angle θ increases with distance from the pole and is maximal at the equator. Therefore, the probability of observing an angle between two random vectors will be increased around $\pi/2$, i.e., the vectors tend to be perpendicular (cf., Figure 5.2A). The same effect is observed for larger dimensions $N > 3$. The probability distribution can be approximated numerically by sampling the angles within a set of random vectors, created by drawing the vector components independently from a normal distribution and then normalizing the vectors (Guhr

et al., 1998). Indeed, the probability distribution can also be calculated analytically (Cai et al., 2013):

$$f_{\measuredangle}(\phi) = \frac{\Gamma(\frac{N}{2})}{\sqrt{\pi}\Gamma(\frac{N-1}{2})} \sin(\phi)^{N-2} \qquad \phi \in [0, \pi] \tag{5.1}$$

However, here, we consider angles between eigenvectors that are not randomly distributed but instead are pairwise orthogonal due to correlation matrices being real and symmetric. We can numerically demonstrate that the effects of these additional constraints are only relevant for low-dimensional spaces and that the distribution of eigenangles can be approximated by $f_{\measuredangle}(\phi)$ for higher dimensions. To show this, we first define a random correlation matrix $\mathbf{A}$, which is positive definite, symmetric, whose elements are real-valued random variables $A_{j,i} \in [-1, 1]$, and diagonal elements are $A_{i,i} = 1$. A random correlation matrix can be created by calculating the Gram matrix from a set of normalized random vectors: $\mathbf{Y} = \mathbf{X}\mathbf{X}^*$, where $\mathbf{X}$ is a matrix with rows X_k being normalized vectors with elements drawn independently from a Gaussian distribution and $\mathbf{X}^*$ denoting the conjugate transpose of $\mathbf{X}$ (Holmes, 1991). The length of the row vectors X_k does not influence the distribution of the eigenvectors. Therefore, we describe this degree of freedom in the dimensionality of X as $\alpha \cdot N$. While α has no relevance for the distribution of eigenangles, it will influence the distribution of eigenvalues described below. Figure 5.2B demonstrates a representative example of how the distribution of angles between eigenvectors is well approximated by the analytic distribution $f_{\measuredangle}(\phi)$ of angles between random vectors for higher dimensions (about $N > 10$). Therefore, using $f_{\measuredangle}(\phi)$ as an analytical approximation is appropriate for our approach if we consider the analysis of correlations between large numbers of neurons (e.g., $N > 100$, which describes a common scenario in the analysis of electrophysiological data).

Since we motivated that small eigenangles indicate similarity, we define the deviation from orthogonality towards small angles $\Delta_i = 1 - \frac{\phi_i}{\pi/2}$, with $\phi \in [0, \pi]$, as the auxiliary variable *angle-smallness* quantifying the similarity of two vectors on a scale from -1 to 1. Performing a variable transformation on the random angle distribution (Eq. 5.1), we obtain the corresponding distribution for the angle-smallness:

$$\tilde{f}_{\measuredangle}(\Delta) \propto \cos^{N-2}(\Delta \cdot \pi/2), \qquad \Delta \in [-1, 1] \tag{5.2}$$

The eigenvalues of positive definite matrices extracted from neuronal activity provide a measure to describe the amount of variance captured by the relative contributions of individual neurons to the corresponding eigenvectors. For that reason, we argue that eigenvectors with larger eigenvalues have a more dominant role in defining the structure of the correlation matrix as opposed to eigenvectors with small eigenvalues. Thus, in designing a cumulative test score to quantify the similarity of two correlation matrices considering all

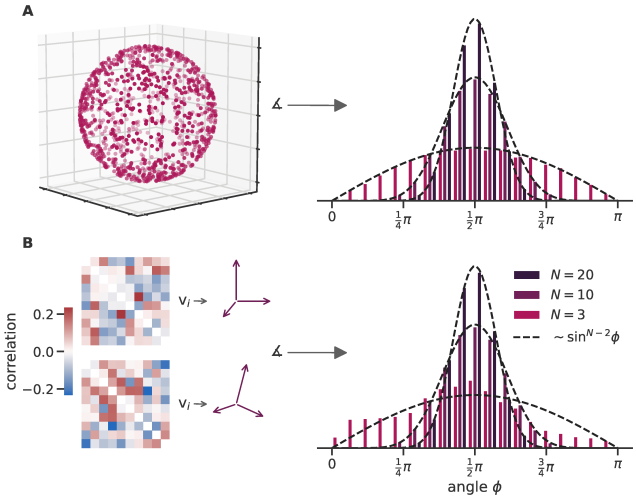

Figure 5.2: **Distributions of angles between vectors. A)** Examples of distributions of angles between normalized random vectors for dimensions $N = 3$, $N = 10$, and $N = 20$ (colors of the histogram). Dashed curve: analytical distribution proportional to $\sin(\phi)^{N-2}$. *Left:* example of random vectors on a sphere in $N = 3$. **B)** Examples of distributions of angles between the corresponding eigenvectors $\mathbf{v}_i$ of two random correlation matrices. Dashed curve: analytical prediction for random vectors as in panel A. *Left:* Example realization of two correlation matrices for $N = 10$. Each histogram contains 10^4 angles from the corresponding sampled vectors.

pairs of eigenvectors, those with the highest eigenvalues should be weighted stronger than those with low eigenvalues. To incorporate this aspect into the test, we weight the angle-smallness Δ_i between i-th eigenvectors of two matrices $\mathbf{A}$ and $\mathbf{B}$ with the quadratic mean of the corresponding eigenvalues $w_i = \sqrt{(\lambda_i^{A\,2} + \lambda_i^{B\,2})/2}$. Although a different average of the eigenvalues, e.g., the algebraic mean, could be chosen here, the quadratic mean better emphasizes large outlier values that correspond to the non-random structures in the matrices.

Next, we derive the analytic distribution of the weighted angle-smallness $w_i\Delta_i$ for pairs of independent random matrices. The Marchenko-Pastur distribution (Marčenko and Pastur, 1967) given by

$$h_\alpha(\lambda) = \frac{\alpha}{2\pi\lambda}\sqrt{(\lambda_+ - \lambda)\cdot(\lambda - \lambda_-)}\lambda_\pm = \left(1 \pm \sqrt{\frac{1}{\alpha}}\right)^2 \alpha > 1 \quad (5.3)$$

describes the distribution of eigenvalues λ for matrices of the type $\mathbf{Y}_N = \mathbf{X}\mathbf{X}^T$, where $\mathbf{X}$ is an $(\alpha N) \times N$ random matrix whose entries are

independent identically distributed random variables with mean 0 and variance $\sigma^2 < \infty$, and $\mathbf{X^T}$ its transpose. The distribution is asymptotic for $N \to \infty$, and can be considered a good estimate for the scenario of $N > 100$ considered in this study. The distribution $h_\alpha(\lambda)$ only depends on the parameter α, which we introduced above as the ratio between the length of the row vectors X_k and the dimensionality N. In the concrete application of $\mathbf{Y}_N$ representing a matrix of correlation coefficients between spike trains, the row vectors X_k correspond to the binned spike trains so that the parameter α is determined by the number of bins divided by the number of spike trains.

For two independent random matrices A and B with identically distributed eigenvalues, pairs of the sorted series of eigenvalues (λ_i^A, λ_i^B) are asymptotically equal for large N. Therefore, the above-defined weights w_i also follow the Marchenko-Pastur distribution. Combining the functions $\tilde{f_N}$ and h_α from Eq. 5.2 and Eq. 5.3, we can formulate the distribution of the angle-smallness Δ when weighted with the weights w as

$$g_{N,\alpha}(w\Delta) = \int_{\lambda_-}^{\lambda_+} \tilde{f_N}(\frac{w\Delta}{\lambda}) \cdot h_\alpha(\lambda) \cdot \frac{\mathrm{d}\lambda}{\lambda}. \tag{5.4}$$

In this formalism, alternative choices for the weights w are possible as long as there is a corresponding analytical description for their distribution.

Given the N individual values of the weighted angle-smallness obtained by comparing corresponding eigenvectors of the two matrices, we define the scalar *similarity score*

$$\eta = \frac{1}{N} \sum_i^N w_i \Delta_i \tag{5.5}$$

as their average. Therefore, a large positive η indicates that the angles between the eigenvectors (in particular, those corresponding to the largest eigenvalues) tend to be smaller than expected for independent random matrices. This indicates that the two matrices have common non-random structures. To interpret a given value of η for a sample with given N and α, we derive how η is distributed under the null hypothesis of random matrices. Assuming the N values of the weighted angle-smallness $w\Delta$ being independent random variables, the central limit theorem states that for large N their properly normalized sum will converge towards a normal distribution centered around the expected mean (here 0) with the standard deviation $\sigma = s/\sqrt{N}$, where s is the standard deviation of $g_{N,\alpha}(w\Delta)$. Expressing s as the integral over the product of the distribution with the squared distance to the mean, we obtain the final analytical description of the distribution of the similarity score

$$f(\eta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-\frac{\eta^2}{2\sigma^2}) \tag{5.6}$$

where

$$\sigma^2 = \frac{1}{N} \int x^2 \cdot g_{N,\alpha}(x)\, \mathrm{d}x.$$

Based on Eq. 5.6, we define the null distribution in the context of a Null Hypothesis Significance Test (NHST). As the construction of $f(\eta)$ was based on the assumption of independent random matrices, the null hypothesis could be best expressed as "the two matrices have no shared non-random structures". We assumed the matrices to be of the type $\mathbf{Y}_N = \mathbf{X}\mathbf{X}^T$. Moreover, we assumed that N is large to account for the constraint that eigenvectors are orthogonal in Eq. 5.1, to be able to apply the central limit theorem, and since the integrated Marchenko-Pastur distribution is asymptotic for $N \to \infty$. From numerical simulations, we found that $N > 10$ is sufficient so that the null distribution $f(\eta)$ reasonably represents the randomly sampled test data.

Violations of the null hypothesis, in particular, any substantial correlation between the matrices, resulting in a higher score that is unlikely to be explained just by the width of the null distribution. To evaluate a sample value of η concerning the null distribution, we assign it a one-sided p-value,

$$p_\eta = \int_\eta^\infty f(x)\, \mathrm{d}x \;=\; \frac{1}{2}\left(1 + \mathrm{erf}\left(-\frac{\eta}{\sigma\sqrt{2}}\right)\right). \tag{5.7}$$

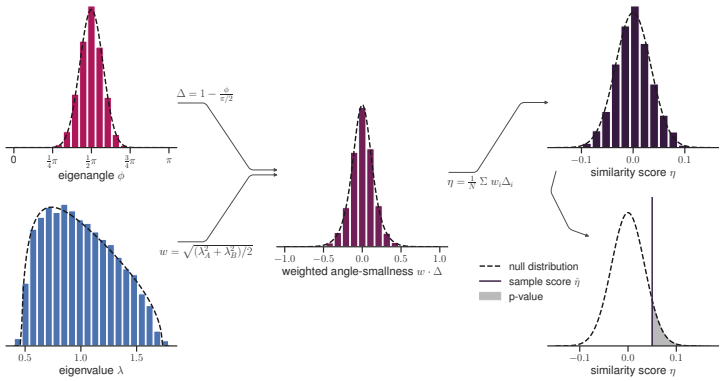

Figure 5.3: **Construction of the eigenangle test.** From left to right, the similarity score η and its corresponding null distribution is constructed by multiplying the angle-smallness Δ (derived from the eigenangles ϕ) with the corresponding eigenvalue weights w (derived from the corresponding eigenvalues λ) for each pair of matched eigenvectors, and then taking the average. p-values can be assigned to the similarity score η using the null distribution. The presented distributions are for dimension $N = 20$, and the histograms represent 10^4 samples.

Using the eigenangle NHST, we conclude the property of interest (here, similarity) as an alternative hypothesis when the null hypothesis (here, independence) is rejected in the case of a p-value smaller than a significance level α. Often, the significance level is arbitrarily set (e.g., $\alpha = 0.05$), which has implications discussed in detail in the literature (e.g., Nakagawa and Cuthill, 2007; Szucs and Ioannidis, 2017). Alternatively, the p-value could be regarded as another quantification, a random variable that should be calibrated with suitable reference scenarios. These scenarios should cover different ways the underlying null hypothesis could be violated to gauge the test's susceptibility. Therefore, to better assess the eigenangle test, in the following, we explore its behavior in use-cases of generated stochastic activity with inserted correlations (Section 5.3.1). Moreover, we extend our test to the case of asymmetric connectivity matrices (Section 5.2.2) and investigate the influence of synaptic rewiring on both the connectivity and activity for a simple balanced random network model (Section 5.3.3).

INTERACTIVE NOTEBOOK An interactive Jupyter notebook illustrates the construction of the eigenangle similarity measure and its statistical evaluation with a combination of text, code snippets, and interactive figures: `https://gin.g-node.org/INM-6/eigenangles/eigenangle_basics.ipynb`[1]

5.2.2 *Statistical eigenangle test for connectivity matrices*

In the following, we investigate under which conditions it is possible to go beyond the description of the eigenangle score for correlation matrices and adapt the statistical test to other kinds of matrices representing pairwise measures. The null distribution for the eigenangle score can be formulated for any kind of random matrix given two requirements: they have an analytic description for their eigenvalue distribution, and the angles between corresponding eigenvectors of two independent realizations of the matrix are distributed like random angles. This opens up additional applications of this approach, particularly the extension to connectivity matrices discussed in the following. Random connectivity matrices differ considerably from random correlation matrices, for example, in that they are sparse, may be inherently structured by the connectivity parameters of different subpopulations of neurons, and have complex-valued eigenvalues and eigenvectors since they are typically asymmetric.

We consider here a specific type of network that fulfills the above two assumptions. Rajan and Abbott (2006) present an analytic description for the absolute eigenvalue values of the connectivity matrices of a balanced excitatory-inhibitory network consisting of two (potentially

1 executable in the browser via `https://mybinder.org/v2/git/https%3A%2F%2Fgin.g-node.org%2FINM-6%2Feigenangles/HEAD?labpath=eigenangle_basics.ipynb`

sparsely connected) sub-populations with their respective synaptic weight distributions. The connections are drawn independently without adhering to a fixed in- or out-degree (i.e., 'pairwise Bernoulli'). In this study, for simplicity, we define these two sub-populations to consist of excitatory (E) and inhibitory (I) neurons, respectively. This implies two weight distributions for the E and I connections. To rank the eigenvectors and set the weights w_i for the eigenangles, we use only the real component $\lambda_{\mathbb{R}}$ of the eigenvalues. The corresponding distribution $h(\lambda_{\mathbb{R}})$ can be derived from the distribution of the absolute values of the eigenvalues $p(|\lambda|)$ by exploiting the point-symmetry of the spectrum for independently drawn matrix elements (Sommers et al., 1988; Girko, 1985),

$$h(\lambda_{\mathbb{R}}) = \int_{|\lambda_{\mathbb{R}}|}^{r_c} \frac{r}{\sqrt{r^2 - \lambda_{\mathbb{R}}^2}} p(r)\, \mathrm{d}r, \tag{5.8}$$

where r_c is the radius of the circle of complex eigenvalues. We are working with a specific network model incorporating the above properties. The exact configuration of the network model is documented in Table 5.1. Using random initializations of this network model, Figure 5.4A displays the sampled histogram of the real eigenvalues together with the analytical description.

We further observe that when the total weight variances of the two populations (taking into account also the variance caused by their relative size difference) are approximately the same, the eigenangles are well described by the distribution of random angles. As the eigenvectors ($\mathbf{v}$) are now complex-valued, we need to adapt the definition of angles to $\phi = \arccos(< \mathbf{v}_A, \mathbf{v}_B^* >_{\mathbb{R}})$ with $<,>$ representing the inner product and * the complex conjugate. Since the additional imaginary vector components double the degrees of freedom, the dimensionality factor N in Eq. 5.1 also needs to be adapted accordingly to $2N$ so that the random eigenangle distribution becomes

$$f_{\measuredangle}(\phi) = \frac{\Gamma(N)}{\sqrt{\pi}\Gamma(N - \frac{1}{2})} \sin(\phi)^{2N-2} \qquad \phi \in [0, \pi]. \tag{5.9}$$

Figure 5.4B displays this analytical eigenangle distribution and the histogram of sampled eigenangles for the comparison of re-initializations of our network model. Here, the weight variances in the network are not exactly equal, but $\frac{\mathrm{Var}\ J_{\mathrm{I}}}{\mathrm{Var}\ J_{\mathrm{E}}} = 6.16 > 1$. This deviation of the weight variance ratio from 1 causes a slight mismatch of the sampled eigenangle distribution with the analytic curve. However, further adjusting the weight variances would require the standard deviation of the underlying lognormal weight distribution to become considerably larger than its mean, leading to a worse estimation of the sample variance of the weights and, consequently, a worse fit of the eigenvalue distribution. Hence, the chosen configuration represents a compromise for this kind of E-I network with an E/I ratio of

80% to 20% to still allow for a reasonable fit of the eigenangle score distribution.

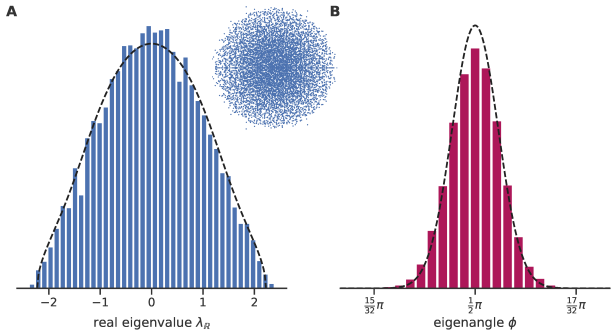

Figure 5.4: **Eigenvalue and eigenangle distribution of connectivity matrices.** For the connectivity matrix of a random balanced network (seed details in Table 5.1), panel **A** shows the sampled distribution of real eigenvalues (inset showing the scatter plot of complex eigenvalues) and the adapted Rajan-Abbott distribution. **B)** The sampled eigenangle distribution of the corresponding complex-valued eigenvectors and the prediction for random angles $\propto \sin \phi^{2N-2}$. The sampled histograms are based on eight re-initializations of the network's connectivity.

Using this extension to connectivity matrices, we can quantitatively compare the connectivity and the activity correlations of certain classes of neural networks with the same statistical method and are therefore able to directly relate the changes in one to the changes in the other. This provides a unique approach to investigating to what degree features of the network connectivity determine aspects of the neural activity.

5.2.3 *Neural network model*

To evaluate the extension of the eigenangle test to asymmetric connectivity matrices, we introduce a network model that adheres to the constraints of the assumed network type (Rajan and Abbott, 2006). The model describes a random balanced network with an excitatory (E) and an inhibitory (I) neuron population. Each population has randomly drawn connections to the other population and itself, with weights drawn from a respective lognormal distribution. Table 5.1 presents a detailed account of all the network model parameters. The model is implemented with the Nest simulator 3.1 [2] (Deepu et al., 2021) The simulated network activity is in an asynchronous state but exhibits isolated periods of synchronization between neurons (Figure 5.5).

2 RRID:SCR_002963

Table 5.1: **Network model configuration.**

parameter	value	description
N	1000	number of neurons
f	0.8	fraction of exc. neurons
J_{ex}	0.1	mean exc. strength [mV]
J_{in}	$-fJ_{ex}/(1-f)$	mean inh. strength [mV]
$P(J_i)$	lognormal	distribution to sample synaptic strengths
σ_{ex}	0.12	std of exc. weight distribution [mV]
σ_{in}	0.1	std of inh. weight distribution [mV]
T	60000	simulation time [ms]
T_0	1000	swinging-in time (additional to T) [ms]
ϵ	0.1	connection probability
η	0.9	external rate relative to threshold rate
delay	uniform(min=0.5, max=3.0)	distribution to sample synaptic delays [ms]
dt	0.1	time resolution [ms]
connection rule	pairwise Bernoulli	rule for connection neurons
synapse model	static	type of synaptic connection
neuron model	current-based, delta, leaky iaf	type of neuron model
τ_m	20.0	time constant of membrane potential [ms]
θ	20.0	membrane threshold potential [mV]
C_m	1.0	membrane capacitance [pF]
t_{ref}	2.0	duration of refractory period [ms]
E_L	0.0	resting membrane potential [mV]
V_{reset}	0.0	reset potential of the membrane [mV]
V_m	0.0	membrane potential [mV]
stimulus type	independent Poisson	driving stimulus to all neurons
stimulus rate	$1000 \cdot \eta \cdot \theta/(J_{ex}\tau_m)$	average rate of the Poisson generator [Hz]

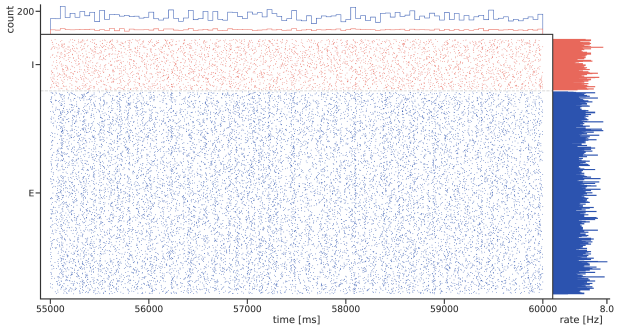

Figure 5.5: **Spiking activity of network model.** The raster plot shows the simulated spiking activity of all 1000 neurons, 800 excitatory (E, blue) and 200 inhibitory (I, red), of the network model (see Section 5.2.2 and Table 5.1) in a 5 s window. Top: population histogram for each population (50 ms bins). Right: mean firing rate for each neuron.

5.3 RESULTS

5.3.1 *Calibrating eigenangle similarity in stochastic activity data*

In correlation matrices, the eigenvectors associated with the largest eigenvalues will point to the dominant correlation structures, e.g., groups of correlated neurons. The eigenangle test formalizes this intuition by evaluating the angles between pairs of eigenvectors. In the following, we quantify the ability of the presented approach to properly detect correlation structures in neuronal data with three calibration scenarios.

In the first calibration scenario, we create two sets of independent stochastic spike trains with inserted correlations among the same sub-populations of neurons. Figure 5.6, bottom row (clustered correlation), shows an example spiking activity of such a correlated sub-population. The stronger the introduced correlations and the larger the sub-populations, the more similar the two correlation matrices and the smaller the test's p-values should be.

We create $N = 100$ spike trains of length $T = 30$ s, of which $N - n$ represent independent neurons modeled by a Poisson process (rate $\nu = 10$ Hz). A compound Poisson process (CPP) models the remaining sub-population (cluster) of size n (Staude et al., 2010). The CPP inserts synchronous spikes into the spike trains of the sub-population (with amplitude distribution $A(j) = \delta(j - n)$) while maintaining the average rate ν and realizing a pre-described average correlation coefficient. The correlation matrices are calculated with a bin size of 2 ms. Figure 5.7 shows that the eigenangle test indeed reflects the increasing similarity of the correlation structure with increasing cluster size and cluster cor-

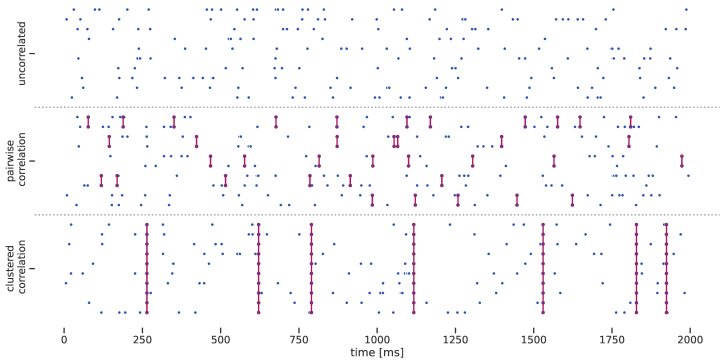

Figure 5.6: **Example stochastic spiking activity with inserted correlations**. The raster plot shows 2 s of stochastic Poisson spiking activity (rate 10 Hz) for three correlation scenarios (10 spike trains each) created by inserting synchronous spikes (red) with a compound Poisson process. The **top row** shows independent activity with only chance correlations. The **middle row** shows activity with synchronous spikes inserted in non-overlapping pairs. The **bottom row** shows activity with synchronous spikes inserted across multiple spike trains. The mean correlation coefficient for the pairwise (middle) and clustered (bottom) correlation neurons is 0.3.

relation. As expected, for zero correlation within the cluster, the null hypothesis that the correlation matrices are independent is plausible ($p \sim 50\%$). However, for a sufficiently large cluster size and cluster correlation (e.g., size$= 4$, correlation > 0.1), the test rejects the null hypothesis with a given significance level (e.g., $p = 5\%$), indicating a similarity of the correlation structures. This analysis confirms the intuition underlying the construction of the eigenangle test and provides an interpretation of its p-value. For example, a significance level of 0.05 for the eigenangle test corresponds to a shared correlation structure that is as least as dominant as 4% of neurons being correlated with an average coefficient of 0.1 with the other neurons showing independent activity.

We further characterize the features of the eigenangle test by comparing it to the well-established Kolmogorov-Smirnov (KS) test (Hodges, 1958). As for most common two-sample tests, the null hypothesis of the KS test is that both samples originate from the same underlying distribution, i.e., the test evaluates similarity, and a p-value smaller than the significance level rejects this similarity. In contrast, the eigenangle test has a null hypothesis regarding the independence of two samples (matrices). Thus, the eigenangle test evaluates difference, and a significant p-value indicates similarity.

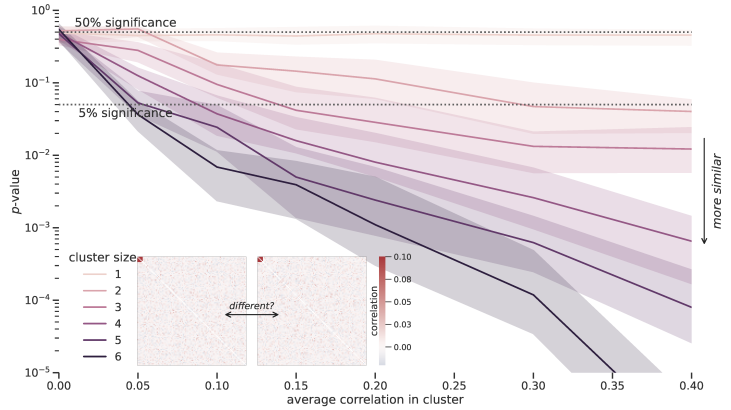

Figure 5.7: **Comparing sets of stochastic activity with similar correlation structures**. The eigenangle test p-values are presented for the comparison of two correlation matrices of stochastic Poisson activity (100 neurons) with a respective sub-population (cluster) of a given size (color hues) and prescribed average internal correlation. Median (curves) and 95% confidence intervals (shaded areas) of the p-value result from 105 repetitions. The matrices show two representative examples of correlation matrices for simulations with identical parameters.

Applying the KS test to the sampled correlation coefficient distributions, in the above parameter scan, there is no visible dependency of the p-value on either the cluster size or the correlation and no rejection of the KS null hypothesis (median p-value: 10.4%; 0.5% of p-values are smaller than a 1% significance level, Bonferroni corrected). This result agrees with the expectation because the correlation coefficients in the two matrices are sampled from the same distribution irrespective of the correlation structure. Further, when we look at the graph description where the edges represent the correlation strength between neurons and compare its local clustering coefficients (Onnela et al., 2005) with the KS test, there is also no dependency of the p-value on the cluster size or correlation.

We again generate correlation matrices for the second calibration scenario from stochastic Poisson activity with correlated sub-populations. Just one sample has two correlated clusters (Figure 5.6, bottom) of size 6 and 8 (average correlation coefficients: 0.3 and 0.1, respectively). In contrast, the other sample has the same "amount" of correlation which is, however, distributed among pairs of non-overlapping neurons. The middle row (pairwise correlation) in Figure 5.6 shows an exemplary spiking activity of such correlated neuron pairs. So, the correlation structure of the two samples is distinctly different, while the corresponding distributions of correlation coefficients are similar. Comparing the matrices with the eigenangle test and the distributions

with the KS test illustrates that information about the matrix structure is lost when only comparing the samples of correlation coefficients. Figure 5.8 shows that the eigenangle test results in no rejection of the null hypothesis and thus adequately indicates that the matrices are different (1.0% false positives, i.e., rejections of the null hypothesis for a 1% Bonferroni-corrected significance level, $N = 105$). In most cases (95.2%), the KS test does not reject its null hypothesis. It, therefore, indicates a principal similarity between the distributions as the difference between clustered and pairwise correlation is not represented in the distribution of correlation coefficients. Applying the KS test to the clustering coefficient distributions of the corresponding graph instead rejects the similarity null hypothesis, but only in about 71% of the cases.

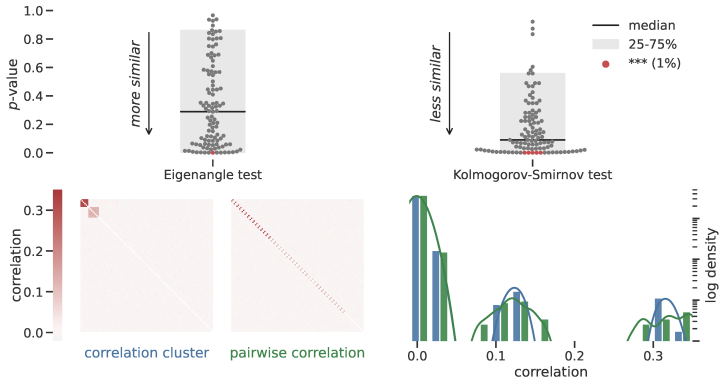

Figure 5.8: **Comparing clustered vs. distributed correlations.** Correlations of network activity are organized in distinct groups of multiple neurons or in independent pairs of two. **Top:** Swarm plots of p-values (dots) obtained from the eigenangle (*left*) and KS (*right*) tests for $N = 105$ random initializations of clustered and distributed correlations. Red dots: p-values significant at the 1% level. **Bottom:** Example realizations of correlation matrices (*left*) for the clustered (blue) and distributed (green) case, and corresponding distributions of correlation coefficients (*right*) used for the KS test.

In a third calibration scenario, we compare the correlation matrices from two simulations of the same network model implemented on two simulation engines. Given that the underlying structure of connectivity in the network model and the input is identical, we expect that both simulations exhibit activity with a matching correlation structure. Figure 5.9 shows the null hypothesis for the eigenangle test is indeed rejected, indicating that the two simulations produce a similar correlation structure. However, the null hypothesis for the KS test is rejected as well, indicating that there are different distributions of correlation coefficients. This discrepancy of yielding similar correlation structures but different "amounts" of correlations shows a tendency of one of

the simulator engines to exhibit more synchronous spikes. A more detailed comparison and discussion of these differences between the simulator outputs are presented in Gutzen et al. (2018b) and Trensch et al. (2018).

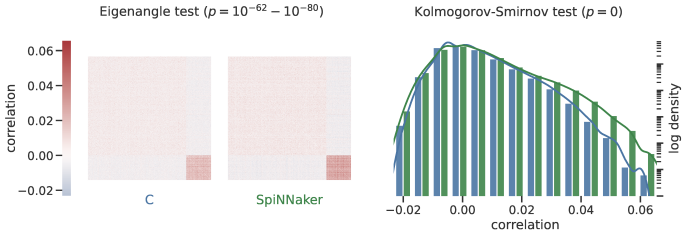

Figure 5.9: **Comparing the model activity from two simulators.** The Izhikevich polychronization model (Izhikevich, 2006) was simulated (Trensch et al., 2018; Gutzen et al., 2018b) with identical parameters and initial conditions on both a custom simulator written in C (blue) and the neuromorphic system SpiNNaker (green). Here, the correlation matrices (*left*) and distributions of correlation coefficient (*right*) from a 60 s recording are shown. Comparing 5 recordings, both the eigenangle test and the KS test clearly reject their respective null hypothesis, indicating a similarity of the correlation structure by the eigenangle test while the KS test indicates dissimilarity of the distributions of correlation coefficients.

The three calibration scenarios illustrate the proper behavior of the eigenangle test and how it is complementary to classical two-sample tests in explicitly evaluating structural matrix features instead of the distribution of values.

5.3.2 *Discriminating behavioral conditions in experimental activity data*

The eigenangle test can be applied to compare correlation matrices based on the same set of neurons. This includes experimental data acquired from the same recording setup. Hence, the test can quantify the similarity of the correlation structure across different task conditions. Here, we demonstrate the eigenangle test application to spiking activity data recorded in the motor cortex of a behaving monkey.

In the experiment, a macaque monkey (session l101210-001) performs an instructed delayed reach-to-grasp task (Brochier et al., 2018). During each trial, a visual cue (300 ms) informs which grip type to use, side-grip (SG) or precision-grip (PG), to pull an object. After a delay period of 1000 ms after cue offset, the GO signal appears, and a second cue informs about the force on the object, low-force (LF) or high-force (HF). This analysis considers each trial's delay and movement periods (1800 ms intervals from cue onset to 500 ms after GO). For each of the four different trial types (SGLF, SGHF, PGLF, PGHF), there are 30 trials in the session.

We construct the comparisons of the correlation matrices of each trial type against each other (including itself). To do this, we randomly split the 30 trials of each trial type into two sets of 15 and calculate the correlation matrices from the spiking activity (bin size 5 ms, 74 single-unit spike trains, respectively concatenated within each set). We then apply the eigenangle test to compare the first trial set of trial type A against the second trial set of trial type B. This way, there is no overlap in compared trials, even when we compare a trial type to itself ($A = B$). We resample the trial sets and repeat the comparisons 500 times to remove any eventual sampling bias and evaluate the statistical variance.

The distribution of the resulting test p-values shows considerable variance across the trial samples (Figure 5.10). Still, on average, the test indicates a greater similarity (smaller p-values) between trial types that share a grip type than trial types with different grip types. Furthermore, we see that the force type has seemingly no influence on the similarity of the respective correlation matrices. These results indicate that the spiking correlations during the behavioral task are specific for the instructed grip type and not for the instructed force type. However, we also learn that the eigenangle test should ideally not be applied to compare a single set of data but to analyze an ensemble, e.g., by resampling trials to obtain a more robust evaluation of the p-values.

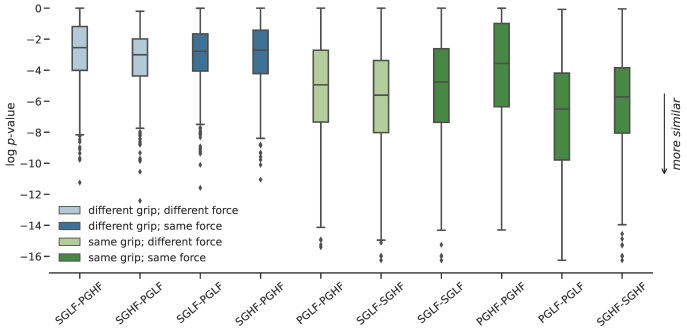

Figure 5.10: **Comparison of the spike correlation structure between 4 experimental conditions.** Each box plot represents the resulting p-values from comparing two trial types when resampling the used trial subset (500 repetitions). The comparisons are grouped by color, showing whether the trial conditions share the grip and/or force types. We observe a tendency that comparisons between conditions with the same grip type (green) yield smaller p-values, indicating more similarity, than those with a different grip type (blue). Comparisons with the same or different force types (light/dark) show similar p-values.

5.3.3 *Evaluating network rewiring effects in modeled connectivity and activity data*

Next, we apply the eigenangle test to evaluate the similarity between simulated networks of the random balanced network model type described in Section 5.2.3. We compare versions of this network model on the activity level via the correlation matrices calculated from the simulated spiking activity. Furthermore, we also perform comparisons on the connectivity level via the synaptic weight matrices using the extension of the eigenangle test described in Section 5.2.2. With the eigenangle test able to perform quantitative comparisons on both network description levels, we can systematically investigate the relationship between connectivity and activity (as schematically illustrated in Figure 5.1).

Specifically, we evaluate the effect that rewirings of the synaptic couplings have on the correlation structure. These simulation experiments are inspired by Mongillo et al. (2018), who evaluated the strategic shuffling and adding of synapses to explore the effects of synaptic volatility and learning. In their study, the authors quantified the influence of modifying connectivity on the activity by using the Pearson correlation between the firing rate vectors before and after rewiring and calculating the connectivity matrices' element-by-element correlation. In analogy to the work by Mongillo et al. (2018), we analyze three types of rewiring protocols.

Redraw: the entire network is re-initialized with a different random seed, i.e., a new connectivity matrix is drawn from the identical weight distribution as the original network.

Shuffle: keeping the exact weight values, the synaptic weights between all pairs of neurons in between a source and a target population are shuffled,

Add: new E-E synapses are drawn that target a sub-group constituting a fraction x of the excitatory population. Synaptic weights are sampled from the weight distribution of existing E-E connections. The number of new synapses is 20% of the number of existing synapses towards the target sub-group, i.e., $0.2x$ of all E-E synapses.

Figure 5.11 illustrates the corresponding test results for each of the comparisons (original vs. rewired network), repeated 100 times with different random initialization.

REDRAW In the redrawing protocol, we compare two independent random connectivity matrices. As expected, the corresponding test results for the synaptic weights (Figure 5.11A, top) are non-significant

p-values (with a median of $p = 0.52$). However, comparing the resulting correlation matrices indicates a certain degree of similarity (with a median of $p = 0.02$). The influence of the general network configuration and the state of the population dynamics can explain this initialization-independent similarity between the correlation structures of completely rewired networks. The network model exhibits fluctuations in the population activity, including times of synchronous spiking between many neurons of the network. This activity feature is represented by the first eigenvector containing similar vector loadings for all neurons. Therefore, synaptic redrawing or rearranging of neurons does not majorly change the direction of the first eigenvector. This characterization of the network's susceptibility to synaptic changes provides a reference to contextualize the effect of the other rewiring protocols (see dotted horizontal lines in Figure 5.11).

SHUFFLE Figure 5.11B illustrates the effects of shuffling the existing synaptic connections between and within the E and I populations (i.e., E-E, E-I, I-E, and I-I). The shuffling of the synaptic weights from E neurons leads to a larger change (i.e., less similarity) in the overall correlation structure than the shuffling of weights where the source population is I. However, the weight matrix structure changes are similar for all source and target population combinations. They only show slightly larger p-values for the E-E shuffling than the I-I shuffling, which likely corresponds to the larger number of shuffled E synapses. With the ratio between the p-value of correlation and weight comparison ($\log(p_{\text{correlations}}/p_{\text{weights}})$), we may gauge whether the changes to the correlation structure are under- or over-proportional as compared to the changes to the connectivity structure (Figure 5.11B, middle). In particular, we can compare this ratio concerning the reference scenario of a complete weight redrawing (Figure 5.11A, middle). This way, we identify an over-proportional influence of changing synapses originating from the E population on the correlation structure.

Additionally, Figure 5.11B-bottom shows the synaptic shuffling effects on the E and I firing rates by measuring the Pearson correlation coefficients between the respective rate vectors before and after rewiring. The shuffling only affects the target population's firing rate vector, whereas the source population's vector is nearly perfectly correlated (i.e., unchanged by the rewiring protocol). Further, the synapses originating from the E population have a larger influence on the target population's rates than those originating from the I population. However, this trend inverts when correcting for the different number of synapses in the E and I populations. When shuffling the same number of weights in each case, the synapses originating in the I population exhibit a larger influence on the rate vectors than those originating in the E population, particularly for I-I connections (Figure 5.12B, bottom). However, even when correcting for the number of synapses,

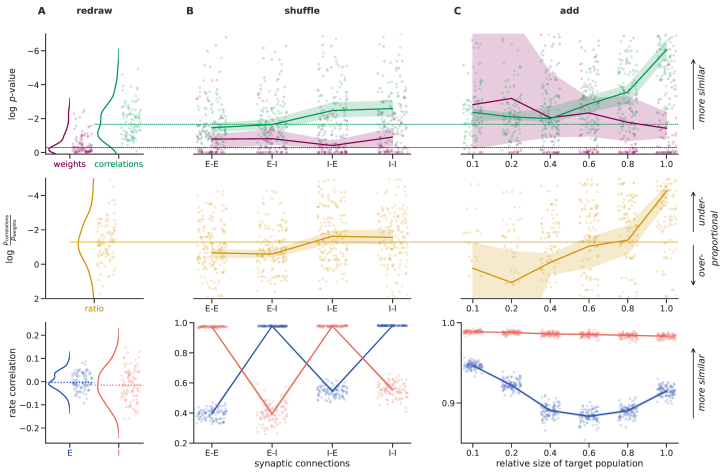

Figure 5.11: **Comparing the effects of rewiring protocols on the connectivity and activity of a random network model.** In each panel, the scattered points show the outcomes of 100 repeated comparisons between original and rewired networks realizations (some points lie outside the bounds of the plotted domain and are not shown), and the curve and shaded area show the median value and the bootstrapped 95% confidence interval. **Top row:** log p-values of the eigenangle test comparing the weight (magenta) and correlation (green) matrices of two network realizations. **Middle row:** corresponding ratios of the log p-values for the comparisons of correlations versus corresponding weights indicate if changes in activity correlation are over- or under-proportional given the change in the connectivity structure. **Bottom row:** correlation of the vector of firing rates, separated for excitatory (E, blue) and inhibitory (I, red) neurons. **A)** The "redraw" column shows the distribution of comparison results when comparing two random network initialization, i.e., different random seeds. **B)** The "shuffle" column compares network realizations with rewired versions where the synapses between the populations indicated on the axis are shuffled. **C)** The "add" column compares networks with rewired versions where the number of E-E synapses that targeted a sub-population constituting a relative fraction of x of the target E population is increased by 20%.

the trend that E synapses have a greater influence on the correlation structure remains (see Figure 5.12B, top and middle).

Mongillo et al. (2018) also find that synaptic shuffling mainly influences the rates of the targeted population. However, they also see a considerable influence on the non-targeted population. Furthermore, the study reports a greater influence of synapses originating in the I population even when not correcting for the number of shuffled synapses. These deviations emphasize the relevance of the underlying network configuration. The network model employed by Mongillo

et al. contains considerably higher connection probabilities (and thus, more indirect and recurrent connections) than our network model. Moreover, in their model, the weight distribution and firing rate variances are larger for the I population than for the E population. The authors highlight these two factors as the relevant attributes for the increased influence of I connections.

ADD When adding additional E-E synapses, we see that a few but targeted synapses (at $x = 0.1$, adding 1280 synapses) can cause a considerable change to the correlation structure (Figure 5.11C), comparable to a full redrawing. However, adding many synapses in a more distributed manner (up to $x = 1.0$, adding 12800 synapses) causes little change to the correlation structure, even though the additional synapses increase the overall synchronous activity and the average correlation. However, changes to the weight matrix similarity seem to be linked to the total number of added synapses, i.e., the more synapses are added, the more dissimilar the weight matrix becomes. Mongillo et al. (2018) show the same trend of the weight matrix similarity using the correlation coefficient between the original and rewired matrix elements.

The effect on the E firing rates is a trade-off between the absolute number of added synapses and the focus of their targets, resulting in a U-shape of the rate correlation curve (also reported in Mongillo et al. (2018)). The isolated effect of focusing the new connections becomes more visible when changing the protocol and adding the same number of synapses irrespective of the target population size (see Figure 5.12C).

5.4 CONCLUSION

SUMMARY OF THE EIGENANGLE TEST APPROACH Angles can serve as a measure of vector alignment. We demonstrate that a score based on the angles between eigenvectors can quantify the similarity between two symmetric correlation matrices (or two asymmetric connectivity matrices). With the analytical description of this similarity score under the null hypothesis of independent random matrices, we construct the statistical eigenangle test. The eigenangle test expands a niche of statistically comparing pairwise measures and is complementary to traditional two-sample tests (e.g. Kolmogorov-Smirnov test) that compare distributions of measures. We establish this statistical comparison approach's relevance for pairwise measures representing interactions between neurons in neural network descriptions.

TAKE-WAYS OF THE APPLICATION EXAMPLES In three calibration scenarios, we illustrate that the eigenangle test can accurately assess shared structural features (i.e., correlated sub-groups) in the compared

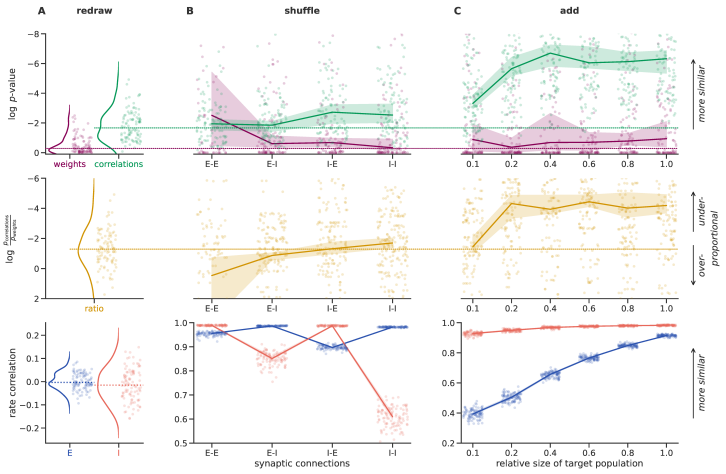

Figure 5.12: **Effects of network rewiring with same numbers of changes synapses.** Same as Figure 5.11, but for the *shuffle* protocol, the same number of synapses (3500) are rewired for E-E, E-I, I-E, and I-I. For the *add* protocol, the same number of synapses (12800, corresponds to 20% of E-E connections) are added for each target population size.

matrices and distinguish between different structural configurations even when the corresponding univariate distribution of values is comparable. Furthermore, applied to experimental spiking data of a behaving monkey, the eigenangle test can differentiate between the correlation structure in two distinct behavioral conditions.

For a balanced neural network model with an excitatory and inhibitory population, we apply the eigenangle test to compare both the synaptic weights and activity correlations between different realizations of the model. We evaluate the effects of synaptic rewiring protocols that redraw, shuffle, or add connections. We discover that the type and scope of change brought on by a particular synaptic rewiring in the network rely on the precise network configuration, making it challenging to draw general conclusions about the functional characteristics of random balanced networks. Using our specific network model type, we discover that the synapses from the E population are primarily responsible for determining the correlation structure. However, rewiring synapses that originate in the I population has a greater impact on the firing rates. Furthermore, the correlation structure is more influenced by how clustered synaptic connections are than by how many there are.

GENERALIZATION OF REWIRING EXPERIMENTS The network model in Mongillo et al. (2018) and our network model (following the conditions of Rajan and Abbott (2006)) are both relatively simple balanced

E-I networks. Mainly, they differ in their connection probabilities and synaptic weight variances. Still, most of the corresponding results of the rewiring experiments from Mongillo et al. (2018) can be qualitatively reproduced. The remaining discrepancies suggest a dependency between the impact of a specific synaptic rewiring and the configuration and state of the network. The rewiring (i.e., plasticity) effects' dependency on the network architecture would also be expected when assuming that different types of neural networks (e.g. different cortical areas) support different functionality, including mechanisms of learning, memory, or representational robustness. Our approach allows us to investigate such dependencies between the correlation structure and distinct network connectivity features. For example, this approach could measure the influences of the arrangement, composition, and interaction of cell assemblies which play a role in various theories of cortical information processing (Harris, 2005; Litwin-Kumar and Doiron, 2012; Aviel et al., 2003).

APPLICABILITY OF THE EIGENANGLE TEST The test depends on the assumption that the two compared networks have the same neuron identities to properly define an angle between corresponding eigenvectors. Thus, valid applications include model vs. model comparisons of two versions of the same model (Section 5.3.3) and experiment vs. experiment comparisons for data acquired with the same recording setup (Section 5.3.2). Notably, classical model vs. experiment validation is generally not possible unless there is a mapping between simulated and experimentally observed neurons. However, validation testing regarding a model's dependency on the input, its robustness to parameter variations, or its stability over time (Section 1.4) may be evaluated with the test. For any application, to guide the interpretation of the test results, it is generally advisable to probe the test's sensitivity modeled influences in the given context, as was down in Section 5.3.1.

The test assumes a sufficiently large network to apply certain statistical approximations. However, we estimate with numerical simulations that $N \geq 100$ represents a sufficient number of neurons for practical purposes. Furthermore, the test relies on ordering the eigenvalues to define the pairwise angles. When there are two equal eigenvalues (e.g. two identical clusters in a matrix), this ordering can become ambiguous. In corresponding numerical simulations, the test nevertheless correctly indicates similarity in most cases but with reduced accuracy. However, we expect such ambiguity in ordering (the largest) eigenvalues to be improbable in most experimental or simulated data.

SOURCES OF CORRELATION There can be different influences on the pairwise correlations in a neural network. A correlation between two neurons' activity may indicate a direct physical connection, an indirect connection (intermediary or common input), or a joint dynamic

(e.g. same response selectivity or non-stationarity). These mechanisms can occur on different time scales, and how much they are reflected in the correlation measure is a matter of the chosen bin size and recording length. An accurate estimation of the pairwise correlations requires a sufficiently long recording of the activity data. However, an increased recording length may entail an increased level of non-stationarity. For example, suppose the population firing rate is not stationary in time, as often in experimental data. This will induce rate correlations and influence the measured correlation values (the degree of influence depends on bin size). Besides a non-stationary population activity, the correlation structure itself can also change over time.

The eigenangle test evaluates the data in a non-time-resolved manner and only considers the measured correlation aggregated from various influences. Therefore, in more complex simulations or experiments, timescales for the correlation and the recording window should be carefully chosen, and the test result must be discussed concerning eventual non-stationary influences. Repeating an eigenangle test with different timescales or surrogate data can help disentangle these influences and interpret the test results. For example, to separate the influence of fine-temporal correlations and rate correlation, one could apply the test to spike-time dithered surrogate data (Grün, 2009; Stella et al., 2022) that preserves only the rate profile.

6

ACTIVITY-DRIVEN CALIBRATION OF NETWORK CONNECTIVITY

This Chapter is based on the following publications:

- Morales-Gregorio, Aitor, Robin Gutzen, Paulina Dabrowska, Alper Yegenoglu, Sandra Diaz-Pier, Sarah Palmis, Alexendre René, Panos Sapountzis, Markus Diesmann, Sonja Grün, Johanna Senk, Georgia Gregoriou, Bjørg E. Kilavik, and Sacha J Van Albada (2022). "Activity-Driven Microconnectome Estimation of Macaque Visuomotor Cortices." In: *Preparation*

Author contributions:

- AMG, BK, SvA conceptualized the study.
- SvA supervised the project.
- SaP, BK, SoP, PS, GG collected the data.
- AMG (supported by AR, RG, JSe) implemented the models.
- RG, AMG designed and implemented the calibration & validation workflow.
- AMG, AY, SD, PD implemented the optimization algorithm.
- AMG performed and visualized the optimization applications.
- AMG (with inputs from PD, RG, AR) wrote the initial draft of Morales-Gregorio et al. (2022).
- RG wrote this chapter.
- All authors are contributing to the writing of the manuscript.

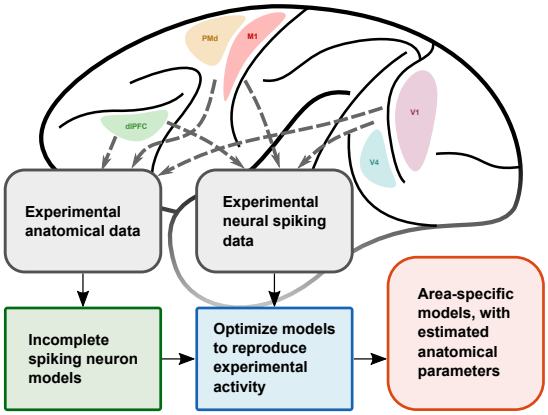

Figure 6.1: **Graphical overview.** We infer area-specific spiking neuron models based on spiking activity data from multiple cortical areas. We use an optimization algorithm to estimate the connectivity parameters of the models that lead to the maximum similarity between the simulated and experimental neuronal activity.

6.1 INTRODUCTION

FOCUS ON NETWORK CONNECTIVITY Networks are defined by their connections. So, in the network view of neural systems, we focus on how the network nodes (in our case, the neurons) interact. The physical connectivity in the form of synaptic adjacency thus represents the backbone of neuronal networks. These connections relay the information transport and processing to generate complex cognitive functions. In the previous Chapter 5, we measured the influence that specific changes to the connectivity structure have on the network activity in terms of its spiking correlations and firing rates. Our approach revealed how, in a given network model, some connectivity aspects and parameter statistics can be more critical in determining network activity than others. Here, we continue to explore how network connectivity determines network activity, but with the alternate approach of inferring connectivity parameters of a neural network model from characteristic measures of spiking network activity.

MEASURING CONNECTIVITY BIOLOGICAL NEURONAL NETWORKS We look, in particular, into the microscale connectivity in different cortical areas of macaque monkeys. While much is known about the cortico-cortical long-range connectivity in the macaque from tract-tracing studies (Felleman and Van Essen, 1991; Bakker et al., 2012; Markov et al., 2012), the connectivity within the cortical areas is less tractable. Currently, estimates of the within-area connectivity rely on

a combination of measurements from different animals and areas (Binzegger, 2004; Potjans and Diesmann, 2014). While cortical areas differ in their neural activity and functionality, they rely on a similar basic columnar architecture. This architecture is usually described as the cortical microcircuit and considered the canonical building block of the cortex (Douglas et al., 1989). However, there is a lack of detailed estimations of the microcircuit connectivity (microconnectome) outside the early sensory and motor areas. Estimating such local connectivity often relies on paired recordings (Thomson, 2002; Song et al., 2005; Perin et al., 2011; Kodandaramaiah et al., 2018) or glutamate uncaging (Kodandaramaiah et al., 2018; Nikolenko et al., 2007; Weiler et al., 2008; Anderson et al., 2010; Noguchi et al., 2011; Yamawaki et al., 2014; Yamawaki and Shepherd, 2015), which are both tedious processes and sample from relatively few neurons at a time.

Therefore, we propose an activity-driven method to estimate the microconnectome from *in vivo* data. Network activity data is used to calculate functional connectivity via various correlation-based methods (Berger et al., 2007; Dann et al., 2016; English et al., 2017; Pastore et al., 2018; Kobayashi et al., 2019). Under ideal stationary conditions, functional connectivity can approximate structural connectivity (Ostojic et al., 2009; Bullmore and Sporns, 2009; Eickhoff et al., 2010). However, simulation studies have shown that spurious correlations lead to systematic errors in the inference of structural connectivity from pairwise correlations (Das and Fiete, 2020).

INFERRING CONNECTIVITY VIA CALIBRATION Instead, we use a model calibration approach. We determine plausible connectivity parameters by simulating spiking neural network models and matching the joint single-neuron statistics between simulations and experimental activity data. There are several approaches to performing the parameter exploration to optimize the similarity of the simulated and target experimental activity, including brute force search (Prinz et al., 2004; Stringer et al., 2016), plasticity rules (Diaz-Pier et al., 2016), inference via analytic likelihood functions for tractable models (Paninski et al., 2004; Pillow, 2005; Ladenbauer et al., 2019; René et al., 2019), or machine learning approaches for analytically intractable models (Bittner et al., 2019; Gonçalves et al., 2020). However, here, we use an approach based on evolutionary optimization methods (Druckmann, 2007; Rossant, 2010; Carlson et al., 2014) because of their computational efficiency, little assumptions on the model, and broad applicability.

Generally, the calibration/optimization approach does not necessarily result in a unique set of parameters. Instead, it may result in one of the potentially many solutions because disparate network parameters can produce the same dynamics (Prinz et al., 2004). The degeneracy of the optimization problem depends on the model complexity and parameterization, as well as the similarity score. In turn, the successful

optimization of an adequate model consolidates that the chosen similarity score and its description level of the network activity capture the relevant features of the network dynamics (cf.. Chapter 1). We evaluate the unique convergence of our optimization setup via the application to synthetic data with known ground truth. Additionally, performing multiple optimization runs and applying a separate validation step further supports the plausibility of the optimization solution. For the calibration and validation steps, we build on the comparison test framework introduced in Chapter 4.

OUTLINE Robustly determining model parameters enables us to investigate the origins of the difference in neural dynamics across cortical areas. There are two options for how the variability of neuronal dynamics across cortical areas could emerge: 1) differences in the received activity from outside the cortical area, or 2) differences in the internal structure (connectivity and/or neuron properties). Therefore, we choose the external input and the internal connectivity probabilities as free variables in the optimization workflow to evaluate these two possibilities.

In the following, we use a method to estimate the parameters of neural network models based on single-neuron statistics. Based on electrophysiological data from different cortical areas and macaques ($N = 5$, *Macaca mulatta*), we explore the variability of cortical spiking activity. We then present a custom optimization algorithm to estimate anatomical parameters from the multi-dimensional activity statistics. We demonstrate the method's performance with the test case of a small balanced spiking neuron network. Further work will include adapting the methods to larger models and estimating the connectivity parameters from the experimental data.

6.2 METHODS

Workflow overview

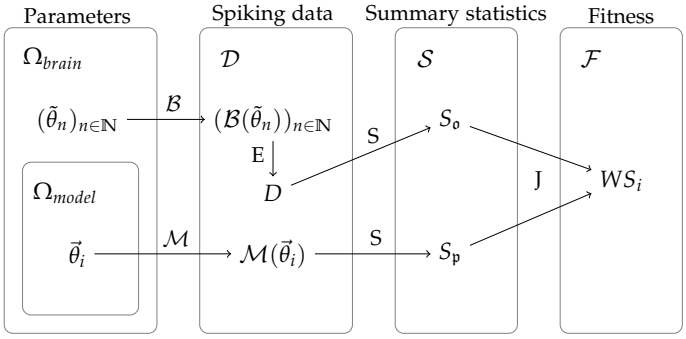

Figure 6.2: **Schematic illustration of the statistical value spaces on each level of abstraction.** From left to right, starting with the underlying parameter spaces in the model simulations (Ω_{model}) and the cortical areas (Ω_{brain}), which both produce spiking data (in $\mathcal{D}$). Deriving summary statistics (in $\mathcal{S}$) from the spiking data is the basis for constructing comparisons. Comparisons are performed via a difference measure (WS) in the space of a fitness function ($\mathcal{F}$).

Assessing the quantitative similarity between experimental and simulated data is a non-trivial procedure that can take many different forms (Chapter 4). Our approach uses a measure of the joint distance between the characteristic measures of the spiking activity data. Figure 6.2 shows the spaces between the abstraction levels of network parameters and the fitness function.

Let $(\tilde{\theta}_n)_{n\in\mathbb{N}}$ bet the set of all the possible parameters that describe the activity in the central nervous system. This includes anything that affects neural activity, such as neuroanatomy, chemical concentrations, or environment temperatures. The brain function $\mathcal{B}$ is a highly nonlinear function $\mathcal{B} : \Omega_{brain} \to \mathcal{D}$, mapping from the parameter space to the space of activity data, i.e., time series $(\mathcal{B}(\tilde{\theta}_n))_{n\in\mathbb{N}}$. Here, we focus on spiking activity, following the assumption that much of the relevant information is encoded within this signal. Current acquisition techniques do not enable simultaneous recording from all neurons in the brain. Thus we work with a subset $D \subset (\mathcal{B}(\tilde{\theta}_n))_{n\in\mathbb{N}}$ recorded with extracellular electrodes E. See a detailed description of the data D in the Methods Section 6.2.1.

On the modeling side, we constrain the range of all possible parameters that may describe the brain to the minimal subset $\theta_i \subset (\tilde{\theta}_n)_{n\in\mathbb{N}}$. Here, we focus on the neuroanatomy of the cortex at the level of the connectivity parameters. The choice of parameters is further deter-

mined by the chosen model $\mathcal{M}$. For a given set of parameters, the model will create a time series $\mathcal{M}(\theta_i)$ of spiking activity. The chosen models M are described in the Methods Section 6.2.1.

The study's objective is to assess whether a choice of model parameters generates a similar spiking activity as the recorded experimental activity, $\mathcal{M}(\theta_i) \approx D$. We extract summary statistics from the data as characteristic measures on the single-unit level. We call this function S, and it yields two clouds of points S_o (observation) and S_p (prediction), where each dot in the cloud corresponds to the summary statistics of a slice of recording from a single spike train. See Methods Section 6.2.2 for a description of the used measures.

Finally, we compare the multi-dimensional summary statistics to each other. Using a judge function J, we can estimate a fitness WS_i, which measures how similar the multi-dimensional summary statistics are. We used the Wasserstein distance (WS) as a fitness function (see Section 6.2.2 for details).

6.2.1 *Data sets*

All data was collected from macaque monkeys (*Macaca Mulatta*, $N = 5$) in the resting state. The macaques were sitting in a dimly lit room with no particular task while the continuous activity from the cortex was recorded. The relevant behavior in each experiment was also tracked using either videos or eye-tracking systems. The recordings were made using different devices, thus resulting in data for different layers. Here, we refer to the combined cortical layers L2 and L3 as L2/3. Consequently, L5 + L6 = L56, and L3 + L4 + L5 = L3/4/5.

Table 6.1: **Summary of subjects and recordings included in this study.** Name identifier of the subject, cortical area, layers from which the data were recorded, number of electrode contacts in each area, and number of recording sessions are indicated.

Name	Sessions	Areas	Layers	Contacts	Source
T	7	M1	All	24	Novel from
	11	PMd	All	24	Kilavik lab
F	59	dlPFC	L2/3	4	Novel from
	59	V4	L2/3	4	Gregoriou lab
L	1	V1	L5/6	896	Spike sorted
	1	V4	L5/6	128	Chen et al 2021
N	2	M1/PMd	L3/4/5	96	Brochier lab
E	2	M1/PMd	L3/4/5	96	Brochier lab

Data from macaque T

The data from macaque T was recorded from the premotor cortex (PMd) and primary motor cortex (M1) ($N = 1$ subject, N=11 $\approx$ 15 min sessions). Acute recordings were made with laminar probes (Plexon and Alpha Omega, 24 contacts, 100 and 200 μm pitch). The laminar probes enabled recording from across the cortical layers. The 200 μm pitch probes could record all layers simultaneously, while the 100 μm pitch probes did not span the entirety of the cortical gray matter. The motor cortex gray matter is known to be approximately 3.5 mm thick (Koo et al., 2012), with the superficial and deep layers roughly split in half. A guard zone was applied around the middle of the probes of 0.5 mm, where the contacts were excluded. The rest were identified as either superficial (L2/3) or deep layers (L5/6). A preliminary analysis of this data has already been published (Kilavik, 2018). The raw data was spike sorted offline. Spike sorting identified 5-13 clean single units per probe and session.

In addition to the spiking data, surface Electromyography (EMG) of the contralateral deltoid muscle, the heart rate with an ear clip, and a video of the macaque behavior were recorded. In all the behavioral videos, the screen LEDs were used to send a 1 s long blink every minute that can be used to realign video with the neural recordings. We performed a video-based segmentation into behavioral epochs: eyes-open, eyes-closed, and movement periods. We excluded the movement periods from our analysis since they are associated with high motor cortex activity and variability (Dąbrowska et al., 2021). Including the movement periods would bias the spiking statistics of the motor cortex to include non-resting state dynamics.

Data from macaque F

The data from macaque F was recorded from visual area V4 and dorsolateral prefrontal cortex (dlPFC) ($N = 1$ subject, $N = 59 \approx 5$ min sessions). Acute recordings were made with up to four simultaneous Plexon electrodes from the superficial layers (L2/3) during resting state. The eye pupil was tracked for behavioral segmentation into eyes-open and eyes-closed epochs. Spike sorting identified 4-10 clean single units per area and session.

Data from macaque L

The data from macaque L was recorded from visual areas V1 and V4 ($N = 1$ subject, $N = 1 \approx 20$ min session). Chronic recordings were made using 16 8x8 electrode Utah arrays (Blackrock microsystems), two of them in visual area V4 and the rest in the primary visual cortex (V1), with a total of 1024 electrodes. The electrodes were 1.5 mm long. Thus the recordings were made from the deep layers L5 and L6. A

full description of the experimental setup and the data collection and preprocessing has already been published (Chen et al., 2022).

Pupil position and diameter data were collected using an infrared camera to determine the macaques' gaze direction and eye closure. In addition to the resting state recordings, a visual response task was also performed. The visual response data were used to calculate each electrode's signal-to-noise ratio (SNR). All electrodes with an SNR lower than two were excluded from further analysis. Additionally, we excluded up to 100 electrodes that contributed to high-frequency cross-talk in each session, as reported in the original data publication (Chen et al., 2022).

The raw data were spike-sorted using a semi-automatic workflow with Spyking Circus—a free, open-source spike-sorting software written entirely in Python (Yger et al., 2018). An extensive method description can be found in their publication and the online documentation of Spyking Circus[1]. After the automatic sorting, the waveform clusters were manually merged and labeled as single-unit activity, multi-unit activity, or noise. Only single-unit activity (SUA) spike trains were included in this study. The waveform signal-to-noise ratio (wfSNR) was calculated for all SUA, and those with a wfSNR < 2 or electrode SNR < 2 (from the visual response task) were excluded from the analysis.

Data from macaques N & E

The data from macaques N & E was recorded with an implanted 10x10 electrode Utah array from the interface between premotor (PMd) and primary motor (M1) cortex ($N = 2$ subjects, $N = 2$ 15-20 min sessions per subject). A full description of the experimental setup, data collection, and preprocessing is described in (Brochier et al., 2018). An extensive analysis of the resting state data has also been published (Dąbrowska et al., 2021). In addition to the registration of brain activity, the monkey's behavior was video recorded and synchronized with the electrophysiology recording.

Spiking neuron models

In this study, we use two different spiking network models, a small balanced spiking network (Brunel and Brunel, 2000; Ostojic, 2014) and a cortical microcircuit model (Potjans and Diesmann, 2014). Both were simulated using NEST 3.2 (de Schepper et al., 2022).

Small balanced spiking neuron network model

The small balanced spiking network model consisted of one excitatory (E) and one inhibitory (I) population of leaky integrate-and-fire

1 spyking-circus.readthedocs.io

neurons. The neurons a randomly connected with a probability P for each possible connections (pairwise Bernoulli). The default connection probability P depends on the source and target population:

$$P_0 = \begin{bmatrix} P_{EE} & P_{EI} \\ P_{IE} & P_{II} \end{bmatrix} = \begin{bmatrix} 0.15 & 0.18 \\ 0.17 & 0.19 \end{bmatrix},$$

Each connection has a delay value that is drawn from a log-normal distribution. The network is driven by an input of independent spike trains simulated with an inhomogeneous Poisson process.

Microcircuit model

The cortical microcircuit model describes the neural populations under a square of 1 mm^2 cortical surface in the layers L2/3, L4, L5, and L6 (Potjans and Diesmann, 2014). Each layer consists of two populations of point neurons, one excitatory and one inhibitory, thus a total of eight populations. The total number of neurons was based on layer-resolved stereological neuron estimates. All layers receive a constant background input with Poisson noise.

For the synthetic data, we used the following connectivity matrix, with the elements $P_{\text{source,target}}$ with source and target being the populations [$L2/3_E$, $L2/3_I$, $L4_E$, $L4_I$, $L5_E$, $L5_I$, $L6_E$, $L6_I$]:

$$P_0 = \begin{bmatrix} 0.1009 & 0.1689 & 0.0437 & 0.0818 & 0.0323 & 0 & 0.0076 & 0 \\ 0.1346 & 0.1371 & 0.0316 & 0.0515 & 0.0755 & 0 & 0.0042 & 0 \\ 0.0077 & 0.0059 & 0.0497 & 0.1350 & 0.0067 & 0.0003 & 0.0453 & 0 \\ 0.0691 & 0.0029 & 0.0794 & 0.1597 & 0.0033 & 0 & 0.1057 & 0 \\ 0.1004 & 0.0622 & 0.0505 & 0.0057 & 0.0831 & 0.3726 & 0.0204 & 0 \\ 0.0548 & 0.0269 & 0.0257 & 0.0022 & 0.0600 & 0.3158 & 0.0086 & 0 \\ 0.0156 & 0.0066 & 0.0211 & 0.0166 & 0.0572 & 0.0197 & 0.0396 & 0.2252 \\ 0.0364 & 0.0010 & 0.0034 & 0.0005 & 0.0277 & 0.0080 & 0.0658 & 0.1443 \end{bmatrix}$$

which was previously derived from anatomical studies (Binzegger, 2004; Potjans and Diesmann, 2014).

6.2.2 *Characterization of spiking activity*

The experimental and simulated spike trains are sliced into 10 second samples. We compute the characteristic measures for each neuron in these data samples and construct a multi-dimensional cloud of spiking neuron statistics. These multi-dimensional summary statistics are used to characterize the activity in the cortical areas and to assess the similarities between experiments and simulations. We use the

following four characteristic measures to quantify multiple aspects of the spiking activity.

FIRING RATE The first order of neuronal activity is the firing rate, measured by the number of emitted spikes (N_{spikes}) per recording time T:

$$\text{FR} = \frac{N_{\text{spikes}}}{T}. \tag{6.1}$$

LOCAL COEFFICIENT OF VARIATION Looking at the intervals between n consecutive spikes of a neuron I, the revised local coefficient of variation LvR (Shinomoto et al., 2003) quantifies the instantaneous regularity of the spiking, also taking into account the refractory period $R = 5$ ms:

$$\text{LvR} = \frac{3}{n-1}\sum_{i=1}^{n-1}\left(1 - \frac{4I_iI_{i+1}}{(I_i + I_{i+1})^2}\right)\left(1 + \frac{4R}{I_i + I_{i+1}}\right). \tag{6.2}$$

Equivalent to other measures of spiking regularity, a value of 0 indicates perfectly rhythmic spiking, whereas a value of 1 corresponds to a Poisson point process.

AVERAGE CORRELATION COEFFICIENT The pairwise Pearson correlation coefficient of two spike trains quantifies the level of synchronized coordination between two neurons. It is defined as the covariance between two binned spike trains b with a mean number of spikes per bin μ, normalized by the roots of the individual variances:

$$\text{CC}_{ij} = \frac{\text{cov}_{ij}}{\sqrt{\text{cov}_{ii}\text{cov}_{jj}}}\text{cov}_{ij} = < b_i - \mu_i, b_j - \mu_j > \tag{6.3}$$

Since CC is a pairwise measure, we use the average correlation value of a neuron with all $N-1$ other simultaneously recorded neurons to have a neuron-wise characterization.

$$\text{CC}^i_{\text{avg}} = \frac{1}{(N-1)}\sum_{j\in[1,N]\setminus i}\text{CC}_{ij} \tag{6.4}$$

STANDARD DEVIATION OF CORRELATION COEFFICIENTS Using the above definition of correlation coefficient Eq. 6.3, we also use the standard deviation to measure the spread of correlation values of a neuron:

$$\text{CC}^i_{\text{std}} = \frac{1}{(N-2)}\sum_{j\in[1,N]\setminus i}(\text{CC}^i_{\text{avg}} - \text{CC}_j) \tag{6.5}$$

Additional measures were considered but not used in our analysis because they were strongly correlated to other measures. If strongly correlated measures were introduced, some properties would be overrepresented in the fitness function. For example, the coefficient of

variation (Cv, Cv2) (Holt et al., 1996) and local variability (Lv) (Shinomoto et al., 2003) of the inter-spike intervals are excluded due to their similarity to the LvR. Introducing Cv, Cv2, Lv and LvR would significantly increase the representation of the inter-spike interval variation without a large increase in the explained variance. We also considered the spike-triggered population response (stPR) (Okun et al., 2015), but did not include it due to its similarity with the CC_{avg}. As the LvR requires a minimum of three spikes to be computed, we set a threshold of FR $\geq$ 2 spikes/s for any given spike train to be included in the analysis. Since the considered measures assume stationarity, the length of the spike train slices must not be too long. We, therefore, choose a spike train sample length (t_{slice}) of 10 seconds. All metrics were computed using their implementation in the elephant toolbox (Denker et al., 2022), within the NetworkUnit reproducible testing framework (Gutzen et al., 2018b).

Multidimensional fitness function

Selection of a fitness function

To assess the fitness of a given parameter set, we need to compare the multi-dimensional single-neuron statistics to our target. Thus, we need to compute a distance between multi-dimensional probability distributions. Such problems have been explored, for example, in the context of GANs, where the Wasserstein distance (Arjovsky et al., 2017), also known as the earth mover's distance, was found to have many desirable properties:

1. agnostic about the underlying statistical distribution
2. jointly evaluates multivariate distributions, thus incorporating the covariance structure
3. can compare samples of different sizes
4. is a true distance: symmetric and positive definite
5. is extendable to higher dimensions
6. is numerically robust on point-distributions (no integration needed)

Not all common similarity measures share all these properties. For example, the Kullback-Leibler divergence violates properties 4 and 6, whereas the Jensen-Shannon entropy violates property 6.

Wasserstein distance

The Wasserstein distance (WS) is defined as an optimal transport problem between the observations (S_o, the target multi-dimensional

summary statistics) and the predictions ($S_{\mathfrak{p}}$, the multi-dimensional summary statistics from the candidate model). Each multi-dimensional summary statistics cloud is an $N \times M$ matrix of N number of samples (spike train slices) and M number of summary statistics. The observations and predictions must have the same number of summary statistics $M_{\mathfrak{o}} = M_{\mathfrak{p}}$ but can have different sample sizes $N_{\mathfrak{o}}$ and $N_{\mathfrak{p}}$. First, we normalized (z-scored) the statistics cloud concerning the observations $S_{\mathfrak{o}}$ across samples for both $S_{\mathfrak{o}}$ and $S_{\mathfrak{p}}$. This normalization ensures that the measured distances remain comparable across many different predictions $S_{\mathfrak{p}}$. The normalization step further ensures that all metrics M are equally weighted. Without normalization, the firing rates would affect the distance by several orders of magnitude more than the correlation statistics ($\mathrm{CC}_{\mathrm{avg}}$, $\mathrm{CC}_{\mathrm{std}}$). Second, we assigned equal weights (mass) $w_{\mathfrak{o}}$, $w_{\mathfrak{p}}$ to each neuron within the statistics clouds, such that $\sum^{N_{\mathfrak{o}}} w_{\mathfrak{o}} = \sum^{N_{\mathfrak{p}}} w_{\mathfrak{p}} = 1$. Thus, $S_{\mathfrak{o}}$ and $S_{\mathfrak{p}}$ have the same total mass, and the differently sized statistics clouds can be compared. Finally, we search for the optimal transport of mass between $S_{\mathfrak{o}}$ and $S_{\mathfrak{p}}$ by finding the graph configuration that minimizes the work required to transport all the weights:

$$\min \sum_{i}^{S_{\mathfrak{o}}} \sum_{j}^{S_{\mathfrak{p}}} w_{i,j} \cdot d_{i,j} \tag{6.6}$$

where $w_{i,j}$ is the weight transported between points i and j, and $d_{i,j}$ is the Euclidean distance between them. For finding the optimal graph configuration, we use the simplex algorithm implemented in OpenCV (Bradski, 2000). Then, the Wasserstein distance (WS) is the work normalized by the transported weight, which we defined to $\sum^{N} w = 1$. Thus $\sum_i^{S_{\mathfrak{o}}} \sum_j^{S_{\mathfrak{p}}} w_{i,j} = 1$, and the WS is therefore

$$\mathrm{WS} = \sum_{i}^{S_{\mathfrak{o}}} \sum_{j}^{S_{\mathfrak{p}}} w_{i,j} \cdot d_{i,j} \tag{6.7}$$

Multi-objective vs single-objective optimization

An alternative approach to a multi-dimensional distance is multi-objective optimization, where the multi-dimensional distribution is separated into univariate distributions. Each distribution is tested separately, and a Pareto front of all objectives is sought. However, toy tests with this approach struggled to fit all summary statistics simultaneously and would lead to inconsistent results. We encountered three problems with a multi-objective approach:

1. Logical "Or": different summary statistics fitted separately and not simultaneously.
2. Ignored covariance structure due to univariate testing.
3. Sensitive to the choice of statistical test or distance (e.g., Kolmogorov-Smirnov, t-test).

Ultimately we found that single-objective optimization with a multi-dimensional distance was more robust than a multi-objective optimization.

Optimization algorithm

We use an optimization algorithm to explore our models' parameter space, maximizing the similarity between the target observations and the model predictions. Analogous to evolutionary processes, each model parameterization is called an *individual*, and a *population* of individual is iteratively evaluated and mutated to generate the individual with the best fit. As described in the previous section (Section 6.2.2), we used the Wasserstein distance as the fitness function, which the optimizer has to minimize. The optimization algorithm consists of an iterative random search outlined in the pseudocode 1.

Algorithm 1 Pseudocode of random search algorithm

```
population ← generate random individuals
for generation ≤ N do
    population.fitness ← evaluate(population)
    survivors ← select(all individuals)
    mutants ← mutate(survivors)
    newcomers ← generate random individuals
    population ← newcomers and mutants
```

First, the optimizer generates an initial population of individuals $\vec{\theta}_0$, $\vec{\theta}_1$, ..., $\vec{\theta}_N$, of length $N_{\text{population}} = 128$. Then, the individuals are simulated and evaluated with the multi-dimensional summary statistics S_i of their simulated spiking activity. The fitness of is calculated by the Wasserstein distance between S_i and the pre-computed target multi-dimensional summary statistics S_0.

The evaluation step was parallelized in a high-performance computing (HPC) system. Thus, evaluating each generation step takes only as long as it's the slowest individual. It is to be noted that the computational cost of spiking model simulations scales up with the total number of synapses. Thus some parameter combinations with unrealistically many synapses would require a much longer simulation time. To avoid this problem, we set a time limit for the evaluation step (tuned to each model); if the time limit is reached, the individual will get the worst possible fitness. After the evaluation step, a fraction $p_{\text{survival}} = 0.5$ of the population with the best fitness (lowest Wasserstein distance) were selected, and the rest of the models were discarded. We call the remaining individuals the survivors. Additional to the survivors, we also include a small fraction of the all-over best individuals by selecting $p_{\text{from\_best}} = 0.1$ from the population with the

highest score of the whole optimization. This addition prevents the algorithm from drifting away from a good fitness area.

After selecting the survivors of a generation, we mutate them and the overall best individuals using a Gaussian mutation with the mean equal to the parameter values and the standard deviation fixed to a value $\sigma_{\text{mutation}} = 0.01$. In some cases, with a probability of $p_{gradient} = 0.5$, the mutation is not random but instead follows the natural gradient, calculated from the nearest individuals in parameter space (including previously discarded ones). We estimate the natural gradient as follows:

$$\vec{g} \approx \frac{1}{N_{nearest}} \sum_{j=1}^{N_{nearest}} \left(WS_j - WS_0\right) \frac{\vec{\theta}_j - \vec{\theta}_0}{\sqrt{\vec{\theta}_j \cdot \vec{\theta}_0}} \tag{6.8}$$

where $\vec{g}$ is the gradient vector, $\vec{\theta}_0$ is the current individual with fitness WS_0 (i.e., the point in parameter space for which the gradient is being estimated), and $\vec{\theta}_j$ are the nearest neighbors ($N_{\text{nearest}} = 1000$) to $\vec{\theta}_0$ in the parameter space.

Finally, the population is complemented with new random individuals, such that the new population has a length of $N_{\text{population}}$. The new population then undergoes the same evaluation, selection, and mutation steps until a maximum number of iterations $N_{iterations} = 250$ is reached. The optimization algorithm was implemented in the learning-to-learn (L2L) framework (Yegenoglu et al., 2022).

6.3 RESULTS

6.3.1 *Characterization of cortical activity across areas and subjects*

The first step is to acquire appropriate data sets for investigating the variation and specialization across cortical areas. Therefore, we sourced cortical spiking data measured in multiple cortical areas and layers from 5 macaque monkeys across 4 experimental setups in different labs (see Methods Section 6.2.1 for details). For a comparison between data sets and a relation to stationary models that is as fair as possible, all recordings were performed without any task behavior or stimulus in an idle resting state. Figure 6.3A shows an overview of the cortical locations and layers for each experimental recording. The spike train data was cut into 10 s slices for the following analysis. Exemplary samples are shown in Figure 6.3B. Since the direct comparison of spiking data is not feasible, we instead choose four single-neuron statistics to characterize the spiking activity per neuron, calculated within each 10 s slice, as the basis for quantitative comparisons: the firing rate (FR), the local variability (LvR), and the average of a neuron's cross-correlations (CC_{avg}), and its standard deviation (CC_{std}) (see Section 6.2.2 for details). Figure 6.3C, illustrates

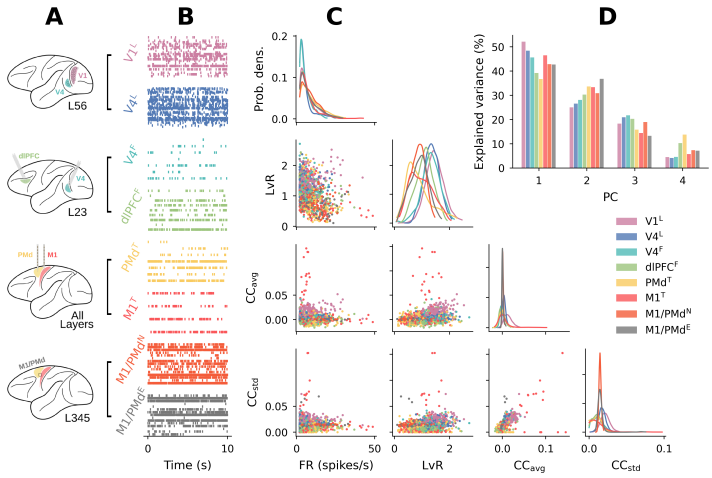

Figure 6.3: **Overview of experimental data sets and their summary statistics. A)** Schematic representation of the recording location. **B)** Sample recordings of simultaneous spike trains from each data set for a 10 s window, superscripts indicate the subject name. **C)** Summary statistics of the single unit spike trains. Each point in the scatter plots corresponds to a 10 s spike train of a single neuron. **D)** Variance explained by the first four principal components (PC) of the multi-dimensional summary statistics.

the univariate distributions of the neuron-wise measures (on the diagonal) and the relations between measures as point clouds (on the off-diagonal), separately for each data set. This choice of comparison measures informs about distinct aspects of the activity data without being redundant. Their individual relevance is underlined by a PCA analysis showing that they each explain at least 5-10% of the variance in each data set (Figure 6.3D).

We also measure the power spectral density from the spike time histograms (Figure 6.4) but do not include it in the summary statistics due to its different dimensionality.

Analogously to the summary statistics for the experimental data, we also extract those four characteristic measures from the two spiking network models we use in this study. A smaller balanced network consisting of 10000 excitatory (E) and 2500 inhibitory (I) neurons and a microcircuit model with E-I populations over four layers (total 61843 E and 15326 I neurons, see Methods Section 6.2.1). From simulations, with their "default" parameters, we already see that the multi-dimensional summary statistics for the small balanced network span a narrower range (Figure 6.5), whereas those of the microcircuit model span a broader range that is more similar to the experimental data (Figure 6.6).

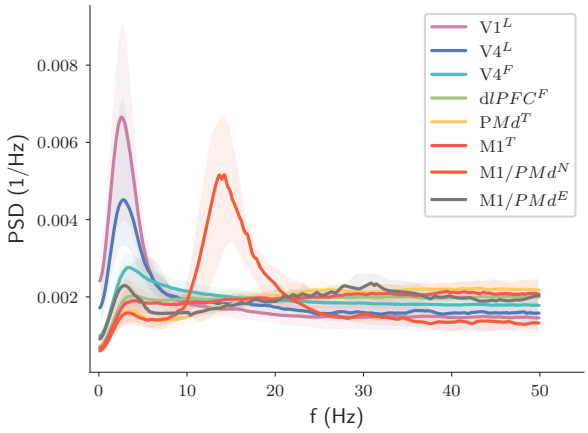

Figure 6.4: Power spectral density (PSD) of the spike time histogram (bin size 10 ms) within the groups for $t_{slice} = 20$ s data slices. The median for each area (solid line) and the 5-95% percentiles (shading) are shown.

6.3.2 *Comparison setup*

This study comprises comparisons between experimental data sets recorded from different subjects and different cortical areas, comparisons between models with different parameterizations, and comparisons of model vs. experiment. Although serving different purposes within the model optimization workflow and the notion of similarity is interpreted differently in each situation, the practical execution and implementation remain the same. In these comparisons, we are quantifying the similarity between the spiking activity of simultaneously recorded neurons. The setup to perform the comparison tests is built on the validation test library NetworkUnit (Chapter 4). We extend the package by implementing additional functionalities needed for our multi-dimensional comparisons.

For this comparison application with data from different areas, layers, and time slices, we need to be careful to separate and reasonably group the data for the extraction of measures and comparisons. With the spiking data being cut into 10 s slices, each data set is represented as a list (for each slice) of lists (for each neuron) of spike trains. The characteristic measures are only computed within each slice of simultaneously recorded neurons. This is particularly important for a valid definition of the correlation coefficients. Furthermore, each neuron is labeled with its attributed layer so that only the characteristics of neurons from the same layer are compared. A corresponding definition of comparison tests realizes this selective grouping. In NetworkUnit, tests are represented as 'class' objects so that variations of existing

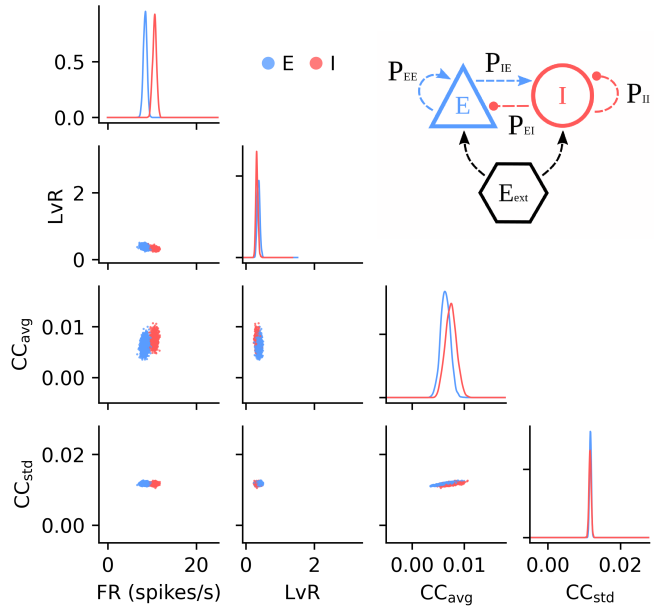

Figure 6.5: Multi-dimensional summary statistics of the small balanced spiking neuron network with the target parameters.

tests can be defined as their child classes, minimizing redundant code and computations. For example, the tests measuring CC_{avg} and CC_{std} are both derived from the general CC test. The four tests for our four characteristic measures (Section 6.2.2) are combined together in a "joint test". Since we chose only neuron-wise measures, the output of each test has a length equal to the number of tested neurons (i.e., spike trains), and the outputs of the four tests can be combined into a four-dimensional characterization of the spiking activity. So, instead of multiple univariate comparisons, we can perform a combined multi-dimensional comparison using the Wasserstein distance (Section 6.2.2). The Wasserstein distance evaluates the minimum distances that are required to move one multi-dimensional point cloud to another. Here, the dimensions correspond to the z-scored characteristic measures and the points to activity slices of individual neurons. This approach has the essential advantage that the information about the relations between the measures is retained and included in the comparison score. So, the outputs of our comparison setup are fitness scores, one for each layer label, that we average to a combined score to be used in the model optimization. It is to be noted that the fitness score does not have absolute interpretability, i.e., it is not a true metric. For the optimization, however, we only need its relative interpretability to

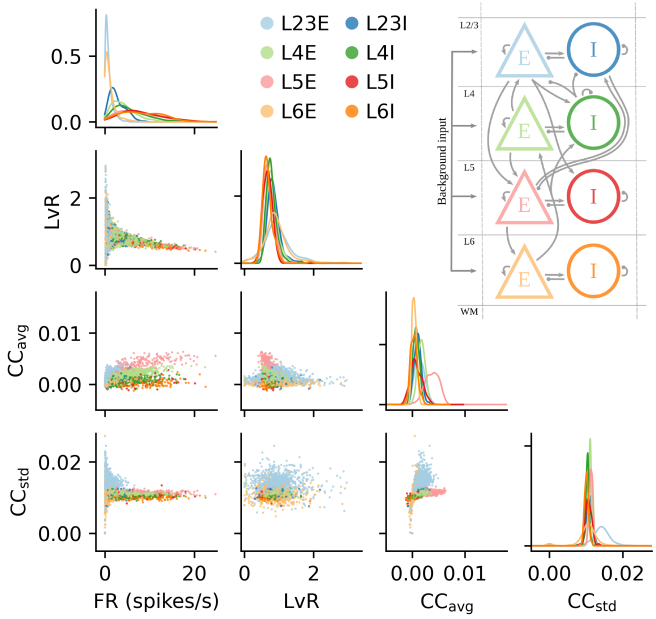

Figure 6.6: Multi-dimensional summary statistics of the microcircuit model with the target parameters.

maximize the model fitness, i.e., minimize the combined Wasserstein distance.

Due to the large amount of data from the experimental recordings and the model simulations and the many comparisons that need to be performed in the stepwise optimization process, there is a relevant concern for computational performance. Therefore, we improved the implementation of the comparisons with NetworkUnit by extending its caching functionality not to recalculate existing results between simulations and tests and introduced parallelization between tests and independent comparisons.

6.3.3 *Multidimensional activity characterization captures the distinct differences between cortical areas*

Before applying our comparison setup to evaluate spiking neural network models, we first gauge the variability across the datasets by applying the setup to experiment-to-experiment comparisons. Thus, we perform pairwise comparisons between all the individual experimental recordings across subjects, areas, and layers (Figure 6.7) using univariate comparisons (ANOVA) for each of the four characteristic measures (Panel A-D) and multivariate comparisons (MANOVA) where the four

measures are considered in combination (Panel E). For the individual tests, to not introduce sample-size biases, we subsample a fixed number N of points from the multi-dimensional summary statistics. We choose $N = 350$ (randomly sampled) since the tests are robust to the normality assumption if the samples are large (generally $N > 30$) and robust to the assumption of equal (co-)variance (homoscedasticity) if the sample sizes are equal (Lindman, 1992). The significance level $\alpha = 0.05$ for the ANOVA and MANOVA tests was Bonferroni corrected for multiple testing ($k = 47$ experiment pairs $\cdot$ 5 test types $= 235$), such that $\alpha = 0.05/k = 2.13 \cdot 10^{-4}$.

The univariate ANOVA tests (Figure 6.7A-D) reveal that no single characteristic measure can distinguish the activity from all areas, i.e., reject the hypothesis of similarity with $p < \alpha$. However, for comparing the combined measures with the MANOVA tests (Figure 6.7E), all recording pairs become statistically different from each other. Therefore, we can conclude that the multi-dimensional summary statistics are specific for each area and/or layer. This result motivates our approach to optimize neural network models separately for different cortical areas to estimate area-specific connectivity parameters.

Although all recordings are statistically different according to the MANOVA evaluations, some appear to be less different than others. Notably, the most similar recordings are for the same area and layer in different monkeys ($M1/PMd^{E}_{l3/4/5}$ vs. $M1/PMd^{N}_{l3/4/5}$), and for the same layer of adjacent areas ($M1^{T}_{l2/3}$ vs. $PMd^{T}_{l2/3}$), while recordings from different layers in the same area and subject tend to be more different ($V4^{L}$, $M1^{T}$, except PMd^{T}).

6.3.4 *Proof of concept: genetic optimization algorithms can infer connectivity parameters*

After demonstrating that our characterization and comparison setup for spiking network activity can differentiate recordings from different cortical areas(Section 6.3.3), we want to exploit this area specificity to estimate corresponding connectivity parameters for area-specific neural network models. To estimate the underlying model parameters from activity data, we search the parameter space for a configuration that minimizes the Wasserstein distance score between simulated and experimental recordings. We explore the parameter space with an evolutionary optimization algorithm (see Methods Section 6.2.2). However, we verify that the optimization method finds the correct parameter values. To this end, we test the optimization workflow with simulated data where the ground truth, i.e., the underlying model parameters, is known. So, we generate a set of spiking activity and calculate its summary statistics by simulating a small, balanced spiking neuron network ($N = 12500$, see Methods Section 6.2.1 with the target parameterization. By matching the summary statistics of the activity

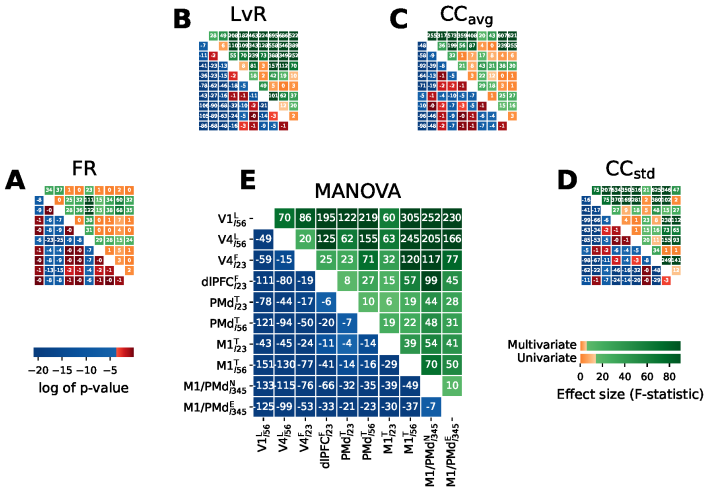

Figure 6.7: **Uni- and multivariate pairwise comparison tests of the four characteristic measures.** We test the null hypothesis that two or more groups have the same population mean. In all panels, lower triangular entries show the logarithm of the p-values, and the upper triangular part shows the F-statistic. **A-D)** shows the univariate pairwise analysis of variance (ANOVA) tests. **E)** shows the multivariate pairwise analysis of variance (MANOVA) tests. Significance levels ($\alpha = 0.05$) are corrected for multiple testing according to the Bonferroni correction. Note that the area ticklabels refer to the lower triangular entries, while the upper triangular part is mirrored along the diagonal.

of the target model with the simulated activity of an independent model realization, we aim to retrieve adequate estimations of the target connectivity parameters: (P_{EE}, P_{EI}, P_{IE} and P_{II}). See Figure 6.5 for the multi-dimensional summary statistics of the synthetic data.

Since the simulations are driven by randomized input, there are inevitable variations in the multi-dimensional summary statistics. We quantified this variability within the target model by simulating $N = 20$ different realizations of the network model, with all parameters equal except for the random seed. The Wasserstein distance between the simulations is presented in Figure 6.8A). Furthermore, we also evaluate the effect of subsampling on the variability and find only a slight increase in the variance of the Wasserstein distance score. Low sensitivity to subsampling is desirable since the recorded experimental data represents only a tiny fraction of the total number of neurons in a cortical area. Next, we ensure that the optimization landscape (parameter space) is smooth. A non-smooth parameter space indicates numerical instabilities in the model, the fitness function, or the summary statistics. Therefore, we perform a parameter scan in the ranges

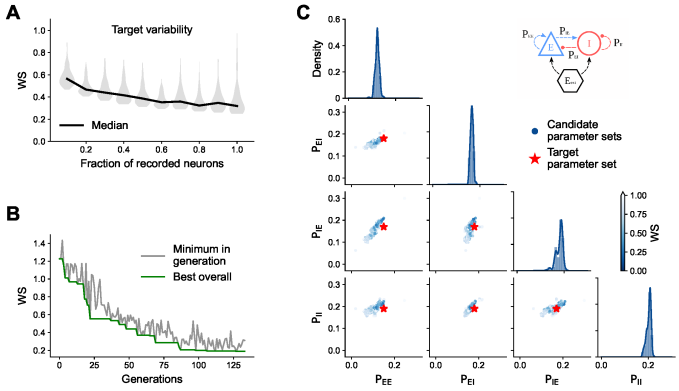

Figure 6.8: **Proof of concept for the optimization method using the Wasserstein distance (WS) between the dynamical ranges. A** WS variability of the target simulation due to different randomized inputs and randomized subsampling ($N = 20$). **B** Progress of the optimization algorithm, as the lowest WS overall (green) and within each generation (grey). **C** Pairplot of the estimated parameter sets. Off-diagonal plots show the parameters of the estimated models. Each point represents one model instance. The color indicates the WS for that model instance with respect to the target. Diagonal plots show the histogram of estimated parameters for models with $WS < 0.5$. The inset on top right shows a schematic representation of the network model.

$P_{xy} \in [0.05, 0.20]$ with steps of 0.01, confirming that the parameter space is indeed smooth. Thus, the optimization algorithm should be able to converge and find the global minimum consistently.

Indeed, the optimization converges in about 100 generations (Figure 6.8B), i.e., $N \approx 12800$ simulations were needed to find the solution since there were 128 simulations per generation. The estimated parameters with a Wasserstein distance < 1 were all found in the vicinity of the target parameters (Figure 6.8C). Considering the measured variability of the simulated activity, the multi-dimensional summary statistics of all those parameter combinations can be considered equivalent to the target simulation. Therefore, we can conclude that our optimization approach can successfully estimate the connectivity parameters (for an adequately simple model) based on multi-dimensional summary statistics of the spiking activity.

6.4 CONCLUSION

SUMMARY Different cortical areas show different functional specificity and neural activity. Given the notion of a canonical cortical microcircuit architecture, how are these differences between cortical ar-

eas manifested in the local network connectivity? Here, we presented a method to optimize the parameters of a neural network model so that the statistics of its spiking activity match the statistic of a target set of (experimentally recorded) spiking activity. The optimization method uses a multi-dimensional cloud of four characteristic single-neuron measures. We showed how the characteristic measures vary across the dataset and how their joint multi-dimensional statistics reflect distinct differences and uniquely characterize the cortical areas. Furthermore, we provided a proof of concept application of the optimization method using synthetic data. Further work will be needed to extend the methods to more biologically realistic models, which could then be used to estimate the connectivity parameters from the experimental data.

NETWORK ACTIVITY CHARACTERIZATIONS FOR CALIBRATION AND VALIDATION Our findings suggest rich and varied multi-dimensional summary spiking statistics across different cortical areas. In agreement with previous reports (Mochizuki et al., 2016), our findings show considerable variability in the spiking activity across the cortex. Furthermore, the characteristic measures (FR, LvR, CC_{avg}, CC_{std}) represent a reasonable description of the main statistics of the network activity. Each describes a considerable (min $5-10\%$) of activity variability, and while none can distinguish the data from different cortical areas alone, their joint statistics show a clear separation. Additionally, they are sufficient in describing the activity of a simple E-I network to determine its network configuration uniquely. Still, to further expand the comparison scope and make our approach more reliable, the calibration-type comparisons based on the four-dimensional summary statistics will be complemented by separate validation-type comparisons. Validation tests are used to evaluate the agreement between a model and an (experimental) target based on data features that are not yet integrated into the model via calibration. Therefore, we intend to compare the power spectral density of the population activities, the number of silent neurons, and the timescales of the spike train autocorrelation. The timescales are already well known to vary across the cortex at both the individual neuron and population levels (Chaudhuri et al., 2015) and therefore represent a relevant condition to check the plausibility of the model.

Our optimization approach was designed to balance exploration and exploitation. We initially used evolutionary algorithms akin to Rossant (2010) and Carlson et al. (2014). However, when fitting the connection probability between neuron populations, the crossover steps would move the optimizer away from the global optimum. Our interpretation was that the crossover shifting of one parameter pushed the model into a completely different regime. This is mostly not the case in genetics, where a single base pair swap often causes only marginal differences.

Thus, we removed the crossover step and substituted it with random sampling.

LIMITATIONS & OUTLOOK So far, this work has demonstrated a proof of concept with a small balanced E-I spiking neural network. We set only four free parameters, spanning a relatively constraint parameter space, that could even be scanned with a naive brute-force approach. The parameter space for the more biologically realistic microcircuit model is several orders of magnitude larger, containing (at least) 64 free parameters. Performing a brute-force parameter scan for the microcircuit model with the same resolution would require $N = 16^{64} \approx 10^{77}$ parameter configurations to evaluate (several orders of magnitude over the estimated total number of atoms of planet Earth, $\sim 10^{50}$). This combinatorial explosion drastically reduces the effectiveness of the random search algorithm. Therefore, the parameter space needs to be constrained, for example, by imposing a strict E-I balance or setting limits for the total number of synapses. Smarter search strategies and initial explorations could further support the exploration performance via computationally cheaper mean-field models (Layer et al., 2022).

Tackling these obstacles, future work will see this approach extended to larger and more biologically realistic models to estimate local cortical connectivity parameters. This will enable us to compare inferred parameters of local connectivity and network input across cortical areas, investigating the relationships between microscale connectivity, statistics of network activity, and specific functionality.

7

SLOW WAVE ANALYSIS ACROSS HETEROGENOUS DATASETS

This Chapter is based on the following publications:

- Gutzen, Robin, Giulia De Bonis, Chiara De Luca, Elena Pastorelli, Cristiano Capone, Anna Letizia Allegra Mascaro, Francesco Resta, Arnau Manasanch, Francesco Saverio Pavone, Maria V. Sanchez-Vives, Maurizio Mattia, Sonja Grün, Pier Stanislao Paolucci, and Michael Denker (2024). "A Modular and Adaptable Analysis Pipeline to Compare Slow Cerebral Rhythms across Heterogeneous Datasets." In: *Cell Reports Methods* 4.1, p. 100681. DOI: 10.1016/j.crmeth.2023.100681
- Capone, Cristiano, Chiara De Luca, Giulia De Bonis, Robin Gutzen, Irene Bernava, Elena Pastorelli, Francesco Simula, Cosimo Lupo, Leonardo Tonielli, Francesco Resta, Anna Letizia Allegra Mascaro, Francesco Pavone, Michael Denker, and Pier Stanislao Paolucci (2023). "Simulations Approaching Data: Cortical Slow Waves in Inferred Models of the Whole Hemisphere of Mouse." In: *Communications Biology* 6.1, pp. 1–14. DOI: 10.1038/s42003-023-04580-0

Author contributions:

- RG, GDB, MD, PP conceptualized the study.
- MD, PP supervised the project.
- RG, GDB, CDL, PP, MD, CC established the methodology.
- RG implemented the software, performed and visualized the analyses.
- AAM, AM, FR, FP, MS provided the experimental data.
- CC, CDL performed the mean-field model calibration.
- RG wrote the draft of Gutzen et al. (2024), and wrote this chapter. CC wrote the draft of Capone et al. (2023).
- All authors contributed to the respective manuscripts.

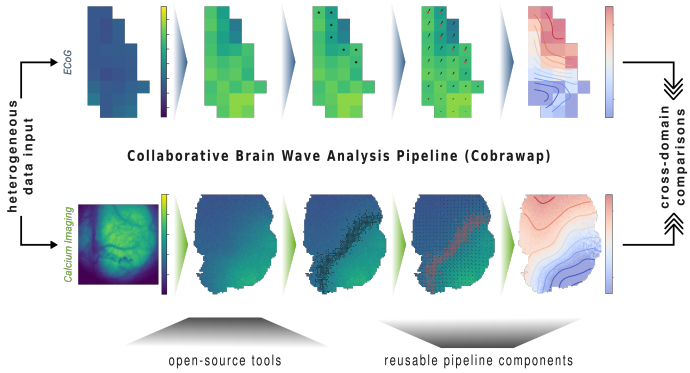

Figure 7.1: **Graphical overview.** The Collaborative Brain Wave Analysis Pipeline (Cobrawap) enables the analysis of slow wave activity from heterogeneous data, for example ECoG (top row) and wide-field calcium imaging (bottom row). In five successive stages the input data is processed and abstracted to a common wave description level, from which wave characterizations can be derived to serve as the basis for cross-domain comparisons. The pipeline is built with open-source tools and in a modular design that makes each component individually reusable.

7.1 INTRODUCTION

Traveling waves are a feature of neural network activity that can be observed in many brain areas, in particular in the cerebral cortex, and with many different measurement techniques. As introduced in Section 1.3, the prevalent phenomenon of cortical wave activity can on the one hand be linked to properties of the underlying connectivity architectures, and on the other has been associated with multiple functional properties (Wei et al., 2016; Davis et al., 2020; Balasubramanian et al., 2019; Heitmann et al., 2015; Townsend et al., 2017). However, their functional role in neural processes and computation remains unclear. Slow wave activity in the range of about $0.5 - 4$ Hz, defined by transitions between states of low activity (Down) and high activity (Up), represent a specific type of cortical waves. They, in particular, are reliably observed in mammals during NREM sleep and anesthesia (Figure 7.2) and are therefore associated with memory, consciousness, and the cognitive effects of sleep (Hanlon et al., 2009; Capone et al., 2019a; Golosio et al., 2021a; Shimaoka et al., 2017). Since the discovery and the early studies and the early studies on slow wave activity in the 1990s (Steriade et al., 1993; Contreras et al., 1996; Achermann and Borbély, 1997; Sanchez-Vives and McCormick, 2000), a lot of research has accumulated around this phenomenon. For example, findings include

that the transitions of cells between Down and Up states are synchronously coordinated over a wide cortical range implying a larger network mechanism (Volgushev et al., 2006), that coordinates with the synchronization of the astrocytic network (Szabó et al., 2017), thalamic activity (Szabó et al., 2017), and across the cortex as reoccurring slow wave patterns can appear over an entire hemisphere (Muller et al., 2016). Evidence from slice and in-vivo recordings further suggests that wave propagation is guided by excitability, i.e., predominantly resides in layers 4 and 5 (Capone et al., 2019b; Bharioke et al., 2022), and shows distinctly different oscillation characteristics across cortical regions (De Bonis et al., 2019). Another insight studying *in vitro* ferret slices is the influence of the temperature: the duration of Up states shortens with cooling while the Down states are shortest around the physiological temperature $36 - 37°C$ (Reig et al., 2009). Although slow wave activity is characteristic of sleep and anesthesia, it can even be observed in localized areas during wakefulness in EEG recordings of behaving mice (Vyazovskiy et al., 2011). Additionally, modeling approaches suggest the importance of long-range connections (Pastorelli et al., 2019), synchronous high-amplitude events (Jercog et al., 2017), and the correct E-I ratio (Keane and Gong, 2015) to exhibit propagating slow waves.

This richness of data and results from the literature, in principle, provides a promising basis to build a comprehensive understanding of the phenomenon. However, given this variety of studies, it is not surprising that they reveal also a very heterogeneous mosaic of datatypes, approaches, methods, metrics, and terminology. Due to this plurality, the relationships between the respective findings are rarely apparent and mostly qualitative, limiting the potential of cumulative discovery by the collection of studies (as introduced in Section 2.2).

Yet, even if authors adopt definitions and methods from other publications, it is not necessarily straightforward to compare their quantitative findings. For example, in one study Massimini et al. (2004) recorded slow traveling waves in sleeping subjects with HD EEG. They determined an average propagation velocity along a row of 20 electrodes of 2.7 ± 0.2 m/s. Botella-Soler et al. (2012) performed similar measurements in sleeping humans and adopted the wave definition and methodology from the previous study (e.g. settings of the bandpass filter and the thresholds for the negative peak detection). Interpolating from two electrodes at a time, they report an average propagation velocity of 1.0 ± 0.2 m/s. Although, the constructed similarity between the two studies makes it considerably easier to relate the results, a rigorous quantitative comparison is still difficult as there remain crucial differences in analysis implementations. For example, the first study uses an arithmetic mean to compute the average while the second study uses the mode peak of a kernel-density estimated velocity histogram. From a reader's perspective, it is difficult to im-

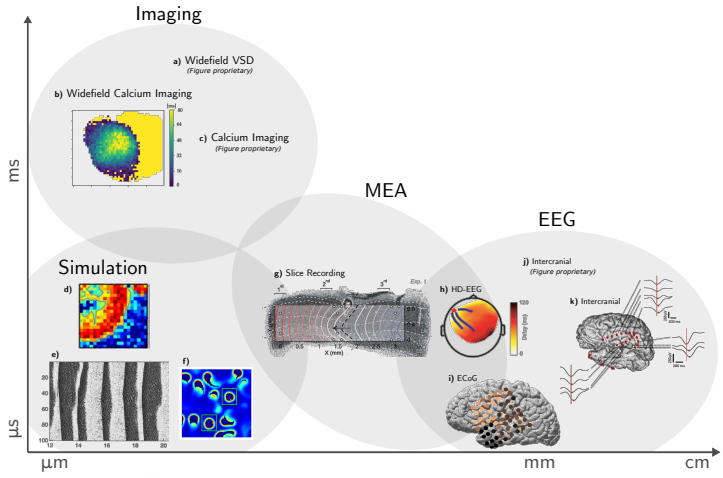

Figure 7.2: **Multiscale, Uniphenomenon: the many faces of slow waves.** **a)** wide-field voltage-sensitive dye imaging of awake C57BL6J mice (Chan et al., 2015), **b)** recorded anesthetized GCaMP6f mice with wide-field fluorescence microscopy (Celotto et al., 2020), **c)** optically evoked Ca2+ wave in anesthetized C57/BI6 mice (Stroh et al., 2013), **d**) distributed network of cortical columns of LIF with Spike Frequency Adaptation neurons (Pastorelli et al., 2019), **e)** one-dimensional multi-layer thalamo-cortical model with one- and two-compartment neuron models using Hodgkin-Huxley kinetics (Bazhenov et al., 2002), **f)** 2D balanced conductance-based spiking neural network model (Keane and Gong, 2015), **g)** multi-electrode recording in ferret cortical slices (Capone et al., 2019b), **h)** human HD-EEG during first sleep episode of the night (Massimini et al., 2004), **i)** human ECoG recording during sleep (Muller et al., 2016), **j)** intracranial depth EEG in sleeping human subjects (Nir et al., 2011), **k)** intracranial depth EEG in humans during sleep (Botella-Soler et al., 2012).

possible to retrace other differences or correct for confounds in their methodology and implementations, especially when the respective analysis code is not accessible or reusable. So, even when definitions, standards, and methods are shared, there can still be various potential undocumented differences in the analysis due to different algorithmic and implementation choices, which makes it increasingly ambitious to interpret and understand the differences in the quantitative results.

One part of the solution to this dilemma is that the analysis method (incl. the code) does not only need to be accessible enough to make the results reproducible, but it further needs to be general enough to be reusable. The other part revolves around finding a common analysis basis to further be able to integrate data from different sources for cross-domain comparisons (Section 1.4). However, defining an adequate description level representing all relevant properties of the phenomenon is typically rather ambiguous since it is generally unclear which observables or characteristic measures are essential. For slow wave activity, the properties that are typically reported are thus often heuristic and include, for example, transition slopes (Ruiz-Mejias et al., 2011), phase velocity (Massimini et al., 2004; Muller et al., 2016), wave type classification (Townsend et al., 2015; Denker et al., 2018b; Roberts et al., 2019), source/sink location and propagation patterns (Huang et al., 2010; Pazienti et al., 2022; Liang et al., 2021), excitability (De Bonis et al., 2019; Ruiz-Mejias et al., 2016; Mattia and Sanchez-Vives, 2012), event frequency (Capone et al., 2023), and others. Concluding from the insights of Chapter 4, there is no single characteristic measure that is sufficient on its own to comprehensively represent the network activity, therefore, any proper comparison approach must incorporate multiple characteristic measures. Thus, we here focus on the common observables planarity, inter-wave intervals, velocity, and direction.

Concluding, the comparability of analysis results requires common (or at least comparable) terminologies, methods, algorithms, and characteristic measures. Therefore, in this chapter we present an adaptable, reusable analysis pipeline for the evaluation of (slow) waves, that is able to integrate data from heterogeneous sources and builds on existing standards and open-source open-source tools (e.g. Sankemake (Mölder et al., 2021), Elephant (Denker et al., 2018a), Neo (Garcia et al., 2014), Nix (Stoewer et al., 2014), Pandas (McKinney, 2010), Scipy (Virtanen et al., 2020)). In the following, we first conceptually examine the development requirements of a flexible but formalized approach for constructing analysis pipelines (Section 7.2.1), that we then leverage to implement a modular pipeline to analyze cortical slow wave activity (Section 7.3.1): The Collaborative Brain-Wave Analysis Pipeline (Cobrawap). Then, we employ Cobrawap to perform a structured and partially automatized analysis of many heterogeneous datasets (Section 7.3.2), and to benchmark the Up state detection method by interchanging the corresponding method block (Section 7.3.3).

7.2 METHODS

7.2.1 *Implementation principles of constructing an adaptable analysis pipeline*

DESIGN APPROACH When constructing a rail system with only a few starting points and a few destination, it might be most efficient to just build the direct paths from between the combinations of starting and destination points. However, with increasing number of starting points and destinations (and their corresponding combinations), it quickly becomes more efficient to build common traffic nodes where tracks from similar domains come together and can then fan out to a multitude of destinations without requiring one specific track. We apply the same idea to the design of analysis pipelines to integrate data from different domains (e.g. measurement techniques) and enable their analysis for a multitude of applications (Figure 7.3). Aligning heterogeneous input data means that we need to find a common representation independent of whether it is obtained via EEG, implanted electrode arrays, imaging techniques, or even simulations. The input data may differ in terms of spatial- or temporal resolution, scale, or signal type. Still, the analysis pipeline should be able to process them appropriately to converge towards a common description of the phenomenon of interest (here, traveling wave activity). Using a common set of methods and algorithms the pipeline should enable a variety of specific analyses and derive characterization metrics that are agnostic about the data's origin. So that the resulting quantitative measures are comparable.

USING MODULAR ELEMENTS The key to making the pipeline adaptable to the different data processing requirements, analysis approaches, and scientific questions is *modularity*. Thus, in the first layer of structuring, we segment the analysis procedure into a series of sequential *stages* that are executed in order ("from left to right"). Each stage is a self-consistent logical step in an analysis workflow with a well-defined purpose, input, and output. A stage should be constructed general enough to be useable as a standalone or potentially reusable in other pipelines, and its output should be considered a reasonable intermediate result. Along the pipeline, the stages become necessarily more specific and tailored towards the scientific application, while the early stages cope with more general tasks such as data integration and preprocessing that are likely shared across different pipelines. Each stage is further segmented into *blocks*. A block defines a concrete action to be performed on the data, implementing a method. Similar to stages, blocks have a well-defined input and output by which they can be chained together. In contrast to stages, blocks are not necessarily executed in a predefined sequence. Instead, each stage

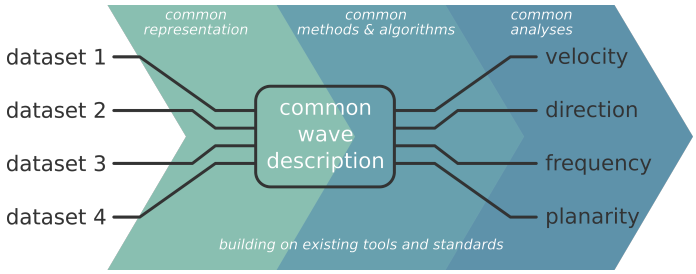

Figure 7.3: **Pipeline Approach.** The proposed pipeline design has the role of aligning methods and algorithms to create and operate on a common description of the phenomenon of interest: integrating data from heterogeneous sources on the one hand and being capable of extracting a variety of output metrics on the other, enabling a basis for rigorous comparison. The foundations of the pipeline are built on existing tools and standards, e.g., data and metadata representation, file formats, standard packages and implementations, environment handling, and workflow management. Within this framework, the catalog of applicable methods is flexibly extendable, making the analysis pipeline adaptable and reusable.

implements the mechanics of the block interactions and defines which block combinations and sequences can be chosen. Some blocks may need to be mandatory for the realization of the stage's purpose and have a *fixed* place in the execution order. Others may be optional and flexibly arrangeable. We identify two flavors of block selections: *choose one*, selecting one method block from multiple options; and *choose any*, selecting any number of method blocks in any order (see Figure 7.6).

Furthermore, the modular design of the pipeline is further supported by building on specialized solutions that address some general demands in research software development. Many of the challenges in constructing the analysis workflows are already formalized and being addressed by open-source software tools and standards, such as data- and metadata representation (Zehl et al., 2016; Sprenger et al., 2019; Rübel et al., 2021), provenance (Butt and Fitch, 2020), version control (Bell et al., 2017), standardized algorithms and frameworks (Virtanen et al., 2020; Denker et al., 2018a; Omar et al., 2014), and workflow management (Mölder et al., 2021; Crusoe et al., 2021; Garijo et al., 2017). Therefore, here, we can focus on making scientific progress by bringing the existing resources together.

REALIZATION WITH THE SNAKEMAKE WORKFLOW MANAGER The pipeline implementation uses the open-source language Python to ensure accessibility and reproducibility. Further, we designed the pipeline's architecture having in mind the features of the Python-

based workflow manager *snakemake* (Mölder et al., 2021), which employs input-to-output rules containing executable shell commands (e.g., Python scripts or bash commands). Snakemake structures the execution of the rules by building dependency trees from the final result file(s) back to the initial input, matching the input requirements to the outputs of preceding rules (see Figure 7.7). The rules are defined in script files called *snakefile* which also link to a config file. Thus, our pipeline structure is conveniently mappable onto the snakemake elements: blocks are represented by rules and stages by snakefiles. In addition, we use another top-level snakefile to combine the stages as snakemake *subworkflows* and make the pipeline executable as a whole. Furthermore, we expand the standard functionality of snakemake by three mechanics required by our pipeline design: 1) chaining the stages by linking the outputs and inputs of subworkflows, 2) manually selecting a specific block (i.e., method) or a sequence of blocks by choosing the desired methods in a config file, and 3) selecting and switching between sets of configs files (*"profiles"*) for all stages.

INTERFACING THE MODULAR ELEMENTS This degree of flexibility in the combination of stages and blocks and their reusability as standalone tools or in other workflows requires a detailed definition and standardizing of the input and output formats. By defining the input requirements for each stage and block, they can successfully interact while remaining interchangeable and thus reusable for other pipelines or applications. The agreement of any input data with the stage input definition checked by a fixed 'check_input' block. The output format of stages is similarly defined, because as intermediate results it should suffice to the same level of completeness and documentation as a final result. These definitions are collected in the stage's README file to guide developers of alternative pipelines as well as contributors of new blocks for the stage. Similarly, the individual blocks also need to clearly state the type and format of their in- and outputs. Other than for the stages, this is largely handled organically in form of the dependencies of the corresponding snakemake rule and the definition of the script's command line arguments and complemented by its docstring.

DATA AND METADATA REPRESENTATION The entire first stage is dedicated to being the interface between the pipeline and the data resource. For the analysis of slow waves with this pipeline, the data needs to be obtained from electrodes or pixels that are arranged on a rectangular grid (which may include empty sites), and that exhibit propagating Up states. The corresponding minimal set of metadata required for the pipeline to process the data are i) the sampling rate, ii) the distance between the electrodes/pixels, iii) and their relative spatial locations of the grid as integer x and y coordinates. Although

not explicitly used, it is strongly recommended to include more information such as the measured cortical location, the spatial scale of the grid, the units of the signal, the type and dosage of the anesthetic, an identifier of the dataset, etc. This additional metadata is propagated through the pipeline alongside the data in order to reasonably use and interpret the results. The consistent handling of data and metadata throughout the stages and block requires representation in a standardized format. For this, we chose the data format Neo (Garcia et al., 2014). Neo supports a variety of data types and reading and writing of various common file formats. This interoperability is, thus, ideal for aiding the flexible use of the pipeline. Since Neo itself is very versatile, there are multiple ways how to organize the data and metadata in the Neo structure, so we need to be even more precise in standardizing the data structure. That means that within the pipeline we store the data of all channels in one 'neo.AnalogSignal' object and the metadata in the corresponding annotations and array annotations for channel-wise metadata (like their x and y coordinates). Processing and transformation blocks overwrite the data in this Analogsignal object and add corresponding metadata. Additional 'neo.Event' objects may be added to represent transition times and wavefronts as well as an additional AnalogSignal object for derived vector fields (e.g., the optical flow). The file format to use for storing the intermediate results of blocks and stages can be format supported by Neo. We recommend Nix (Stoewer et al., 2014) for a robust file format, or the pickle or numpy for a less robust format that is, however, faster to read and write and produces smaller files.

LOGGING MECHANISMS The modular organization of the pipeline facilitates maintainability. Built-in means of the pipeline, such as automatic reports and storing intermediate results alongside their config files, further support reproducibility, and transparency. They contain, for example, logs, execution statistics, dependency trees, plots, and config settings.Additionally, we are currently working on integrating a formalized provenance tracking with *fairgraph*[1]. Furthermore, each block and each stage should also have corresponding plotting output. This helps as a sanity check and aid for potential debugging as well as enriching the results and execution logs.

CONFIGURABILITY The flip side of flexibility and adaptability is complexity and ambiguity. The many combinatorial possibilities need to be controlled by a user interface separate from the actual analysis scripts, e.g., what stages and blocks should be executed, in which order, and with which parameters. Config files (e.g., in csv, yaml, json format) offer human-readable access and control to a user to adapt and execute different variations of the pipeline. Furthermore, blocks

1 https://pypi.org/project/fairgraph/

need to be implemented having generality in mind. This means a given method should be presented in its most general form and any specification should be handled by corresponding parameters settings, given as command line arguments, i.e. within the pipeline via the config file. Even though this approach is initially more time-consuming, it does pay off in both the quality of the method implementation and its (re-)usability. Furthermore, the availability and aggregation of parameters allow for easier and more transparent calibration of the pipeline across blocks and stages. Each stage has one config file collecting the parameter settings for this stage and its blocks. Additionally, there is one config file for the entire pipeline that specifies the stages and their order and can define global parameters that may also overwrite stage parameters, e.g., for setting the file format or plotting parameters for all stages. Parameters in the config files are typically calibrated for a specific data type or experiment setup. To conveniently switch between calibration presets, the pipeline supports a hierarchical organization of config presets via *profiles*. By executing the pipeline with `PROFILE=data1`, for each stage the corresponding config file `config_data1.yaml` is used. For more versatility, profile names can use underscores to define subcategories and exceptions, e.g., `data1_subject3`. In this case, each stage first looks if a corresponding config file of the same name exists, and if not removes the subcategory with the last underscore from the name, and repeats this lookup until it finds the named config file or defaults to `config.yaml`. Furthermore, profiles can have variations indicated in the name with a '`|`', e.g., `data1_subject3|methodA`. This variation key is not removed when first looking up existing config files in the naming hierarchy, only when `config|methodA.yaml` doesn't exit it is removed and the lookup loop is repeated. In contrast to other typical analysis workflows, here, the construction of a specific workflow does not require the changing of any scripts but is rather like tracing a path along the selected stages and blocks within a larger framework offered by the pipeline.

VERSION CONTROLLED DEVELOPMENT The implementation of the general "collaborative brain-wave analysis pipeline" (Cobrawap) infrastructure is being developed on GitHub[2]. In the following, we present the application of the pipeline design for the specific context of slow wave activity and the corresponding configuration[3].

7.2.2 *LogMUA Estimation (in stage 2)*

The multi-unit activity (MUA) is an estimate of the local population firing rate, based on the relative spectral power in the high-frequency

2 https://github.com/NeuralEnsemble/cobrawap

3 https://gin.g-node.org/INM-6/cobrawap_publication_code

regime (200-1500 Hz) of the extracellular recordings. (Mattia et al., 2010; Reig et al., 2009; De Bonis et al., 2019) The corresponding algorithm first selects a moving window that samples the recording at a given rate. From these samples, the power spectral density (PSD) is calculated using the Welch algorithm. The MUA is defined as the average power in the defined frequency band divided by the average power of the full spectrum. Using the logarithm of the MUA helps to emphasize further the bi-modality of the distribution in the presence of slow oscillations. In the selection of the parameters for the algorithm, it is crucial to choose a moving window size large enough so that the chosen frequencies can be accurately estimated (window size $\leq \frac{1}{\text{highpass frequency}}$) and a corresponding MUA rate so that the full recording is sampled from (MUA rate $< \frac{1}{\text{window size}}$).

7.2.3 *Trigger Detection (in stage 3)*

The pipeline implementation provides multiple options to detect trigger events, i.e., transitions from a low activity state to a high activity state (Up).

threshold: The trigger events can either be defined by setting a threshold value for all the signals or by fitting a bimodal function to the amplitude distribution for each channel in order to set the threshold value. In the latter case, the fitting function is the sum of two Gaussians and the threshold value is set to the central minima. This option is applied to the ECoG datasets in this paper. As an alternative to a double Gaussian fit, there is also the option to only fit the first peak corresponding to the low activity state by only looking at the data left of the peak and defining the threshold as mean $+$ std $\cdot$ `SIGMA_FACTOR` with a user-defined `SIGMA_FACTOR`. Since the thresholding method detects also the corresponding downward transitions, this block is usually paired with an additional block that removes Up and Down states that are too short, given user-defined minimal Up and Down durations.

Hilbert phase: Instead of detecting threshold crossings on the actual signal, the upstrokes of the upward transitions can be detected by thresholding the phase signal of the corresponding analytic signal. An adequate threshold value is a matter of definition, here, we apply $-\pi/2$, which corresponds well to the beginning of the upstroke in the actual signal. To be more robust, the algorithm only selects time points where the threshold is crossed from smaller to larger values and where the crossing is followed by a peak (phase $= 0$). This option is applied to the calcium imaging datasets in this paper unless otherwise indicated.

minima: As a third option, we adapted and improved the minima detection method presented in (Celotto et al., 2020). This method relies on the assumption that in an adequately filtered signal that the existence of a local minimum followed by a peak of a certain height indicates the start of a upward transition. This is particularly suitable for recording techniques characterized by a fast characteristic rise time (i.e. comparable with the theoretical minimum time interval between the passage of two waves on a single channel, e.g. optical data). We improved this method by including some further refinement on trigger candidates. Under the assumption that only one minima candidate can lie between two "good" local maxima candidates, we impose that 1) local maxima candidates need to have a signal intensity higher than a relative threshold value, determined in a moving window; 2) local maxima candidates need to be separated by a minimum distance (associated with the characteristic frequency of the investigate phenomenon); 3) a local minima candidate needs to be followed by a monotonically rising signal for a defined time interval (also associated to the characteristic frequency of the investigated phenomenon). If more than one candidate minimum is found between two local maxima candidates, the last one before the following "good" maxima is selected.

7.2.4 *Trigger Clustering (in stage 4)*

Wavefronts are defined as clusters of trigger times in the three-dimensional space of the electrode arrangement (x,y) and samples in time (t). To run a clustering algorithm in this space, the units of the time dimension need to be translated to the units of the spatial dimensions. The ideal transformation factor (`TIME_SPACE_RATIO`) depends on the expected dynamic of the phenomena. A wave that propagates linearly with v_0 is best recognized in the cluster when the time dimension is transformed by a factor $v_0/(\text{sampling rate} \times \text{spatial scale})$. Thus, if we expect a propagation velocity roughly in the order of $\sim 10 - 20 \frac{mm}{s}$ then the transformation factor for the calcium imaging data with sampling rate 25 Hz and spatial scale 50μm is $\sim 8 - 16 \frac{pixel}{frame}$. Here, we choose a `TIME_SPACE_RATIO` of 11 for the calcium imaging data which scales according to the spatial resolution to a factor of 0.25 for the logMUA ECoG data with a sampling rate of 100 Hz. The clustering is performed by a density-based algorithm (`scipy.cluster.DBSCAN`), illustrated in Figure 7.4. The additional parameters for this algorithm are the minimum number of samples (`MIN_SAMPLES_PER_WAVE`) and the typical distance between neighboring sample points (`NEIGHBOUR_DISTANCE`) and were determined by calibrating test recordings from both calcium imaging and ECoG data and scaled consistently with the spatial resolution.

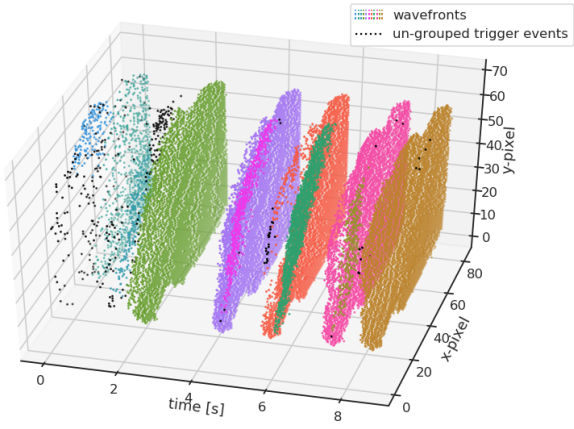

Figure 7.4: **Wavefront definition via trigger clustering.** Visualizing the clustering of detected transition times in the space-time domain for 10 s of an example calcium imaging recording. The trigger events are grouped based on their proximity in space and time using a density-based clustering algorithm (color coded).

7.2.5 *Optical Flow Estimation (in stage 4)*

The optical flow is the pattern of apparent motion in a visual scene, which here corresponds to the recorded signal on the recording grid. To estimate the optical flow of the spatial propagation of activation, we apply the Horn-Schunck algorithm with a quadratic penalty function and a 3x3 Scharr derivative filter on the phase of the signal (the alternative application using the signal's amplitude, as well as different derivative filters can be selected via the configuration). Although other penalty functions, i.e., the Charbonnier function, are more accurate, we found that here the simple quadratic function is sufficient. This observation is in agreement with Townsend and Gong (2018) who report good results for the near quadratic edge case of the penalty function. Their study also guided our choice of the parameter $\alpha = 1.5$, determining the weight of the smoothness constraint over the brightness constancy constraint. The resulting vector field is smoothed by a Gaussian kernel which reflects the dimensions of the expected wave activity with respect to the spatial and temporal scale of the data.

7.2.6 *Quantification and Statistical Analysis*

KERNEL ESTIMATION The kernel estimations for the plotted distributions in Figure 7.10 and Figure 7.11 use `scipy.gaussian_kde` with the default Scott's rule (Scott, 2015) as bandwidth method, except

for the distributions of inter-wave intervals which use 0.2 times the standard deviation as the kernel size.

VELOCITY FILTER Since the channel-wise velocity measure can produce unreasonably high values when there are near identical time delays between spatially distant triggers, we cap the presented distributions at 120 mm/s.

7.2.7 *Experimental model and subject details*

An overview of all the individual recordings is presented in Figure 7.5.

MOUSE ECOG RECORDINGS The three experimental ECoG datasets have been provided by IDIBAPS (Institut d'Investigacions Biomèdiques Agustí Pi i Sunyer): Wild-Type Williams Beuren Syndrome (WBS) 3-4 months old adult male mice, Wild-Type Fragile X Syndrome-Fmr1 knockout mice in FVB background (FXS) and Propagation Modes of Cortical Slow Waves across anesthesia levels in adult male C57BL/6J mice (PMSW) (Sanchez-Vives, 2019a; Sanchez-Vives, 2019b; Sanchez-Vives, 2020). All animals were bred in-house at the University of Barcelona and kept under a 12 h light/dark cycle with food and water *ad libitum*. All procedures were approved by the Ethics Committee at the Hospital Clínic of Barcelona and were carried out to the standards laid down in Spanish regulatory laws (BOE-A-2013-6271) and European Communities Directive (2010/63/EU).

For WBS subjects, an intraperitoneal injection of ketamine (100 mg/kg) and medetomidine (1.3 mg/kg) was administered to induce anesthesia. It was maintained by a constant administration of subcutaneous ketamine (37 mg/kg/h). For FXS subjects, anesthesia was induced by the inhalation of 4% isofluorane in 100% oxygen for induction and 1% for maintenance. Finally, for PMSW subjects, an intraperitoneal injection of ketamine (75 mg/kg) and medetomidine (1.3 mg/kg) and maintained by the inhalation of different concentrations of isoflurane in pure oxygen. In PMSW, three levels of anesthesia were reached that were classified according to the provided isoflurane concentrations: deep=1.16 $\pm$ 0.08% (s.e.m); medium=0.34 $\pm$ 0.06%; light=0.1 $\pm$ 0.0%. The volume delivered was 0.8 L/min.

In order to avoid respiratory secretions and edema, atropine (0.3 mg/kg), methylprednisolone (30 mg/kg), and mannitol (0.5 g/kg) were administered subcutaneously to all subjects. So as to aid breathing and once in the surgical plane of anesthesia, a tracheotomy was performed. The animal was then placed on a stereotaxic frame (SR-6M, Narishige, Japan) with constant body temperature monitoring maintained at 37°C by means of a thermal blanket (RWD Life Science, China). A wide craniotomy and durotomy were performed over the left or right

measurement technique	anesthetic	anesthetic dosage	disease model	subject	recording length
ECoG	ketamine & medetonidine	75 & 13 mg/kg	WBS KO	WBS_KO_6	170 s
ECoG	ketamine & medetonidine	75 & 13 mg/kg	WT	WBS_WT_2	122 s
ECoG	ketamine & medetonidine	75 & 13 mg/kg	WT	WBS_WT_5	140 s
ECoG	ketamine & medetonidine	100 & 13 mg/kg	WBS KO	WBS_KO_1	168 s
ECoG	ketamine & medetonidine	100 & 13 mg/kg	WBS KO	WBS_KO_2	132 s
ECoG	ketamine & medetonidine	100 & 13 mg/kg	WBS KO	WBS_KO_3	129 s
ECoG	ketamine & medetonidine	100 & 13 mg/kg	WBS KO	WBS_KO_4	172 s
ECoG	ketamine & medetonidine	100 & 13 mg/kg	WBS KO	WBS_KO_5	137 s
ECoG	ketamine & medetonidine	100 & 13 mg/kg	WBS KO	WBS_KO_6	170 s
ECoG	ketamine & medetonidine	100 & 13 mg/kg	WT	WBS_WT_1	124 s
ECoG	ketamine & medetonidine	100 & 13 mg/kg	WT	WBS_WT_3	164 s
ECoG	ketamine & medetonidine	100 & 13 mg/kg	WT	WBS_WT_4	138 s
ECoG	isoflurane	80 mg/kg	FXS KO	FXS_KO_1	197 s
ECoG	isoflurane	80 mg/kg	FXS KO	FXS_KO_2	106 s
ECoG	isoflurane	80 mg/kg	FXS KO	FXS_KO_3	243 s
ECoG	isoflurane	80 mg/kg	WT	FXS_WT_3	208 s
ECoG	isoflurane	90 mg/kg	WT	FXS_WT_4	151 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0%	WT	PM_WT_1	228 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0%	WT	PM_WT_2	198 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0%	WT	PM_WT_3	205 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0%	WT	PM_WT_4	205 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0%	WT	PM_WT_6	201 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0%	WT	PM_WT_7	199 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0%	WT	PM_WT_8	211 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0.2%	WT	PM_WT_2	202 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0.2%	WT	PM_WT_3	220 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0.2%	WT	PM_WT_6	216 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0.5%	WT	PM_WT_1	212 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0.5%	WT	PM_WT_4	204 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0.5%	WT	PM_WT_7	204 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0.5%	WT	PM_WT_8	210 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 0.75%	WT	PM_WT_3	213 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 1.0%	WT	PM_WT_1	198 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 1.0%	WT	PM_WT_2	213 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 1.0%	WT	PM_WT_4	210 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 1.0%	WT	PM_WT_6	209 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 1.25%	WT	PM_WT_7	221 s
ECoG	ketamine & medetonidine + iso	55 & 1.3 mg/kg + 1.25%	WT	PM_WT_8	211 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M2	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M2	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M2	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M2	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M2	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M2	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M2	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M2	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M3	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M3	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M3	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M3	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M3	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M3	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M3	40 s
calcium imaging	ketamine & xylazine	100 & 10 mg/kg	WT	M3	40 s
calcium imaging	isoflurane	1-2%	WT	190109	318 s
calcium imaging	isoflurane	1-2%	WT	190109	320 s
calcium imaging	isoflurane	1-2%	WT	190109	320 s
calcium imaging	isoflurane	1-2%	WT	191022	300 s
calcium imaging	isoflurane	1-2%	WT	191022	300 s
calcium imaging	isoflurane	1-2%	WT	191022	300 s

Figure 7.5: **Data overview.** Each row shows one of the 60 recordings used in this study. The columns show some of the attributes in which they can differ, and within each column, different values are colored differently.

(only left in FXS) hemisphere from -3.0 mm to +3.0 mm relative to the bregma and +3.0 mm relative to the midline. A 32-channel multielectrode array (550 μm spacing) covering a large part of the hemisphere's surface was used to record the extracellular micro-electrocorticogram (micro-ECoG) activity. For WBS and FXS datasets, recordings were acquired from spontaneous activity in the animal under anesthesia. Regarding the PMSW dataset, each anesthesia level was maintained for 20-30 minutes, and spontaneous recordings were consistently obtained in a stable slow oscillatory regime (approximately 10 minutes after the change in concentration). During the recording protocol, a precise visual inspection of all channels was made in order to ensure that all of them were properly acquiring the signal. The signals were amplified (Multichannel Systems, GmbH), digitized at 5 kHz, and fed into a computer via a digitizer interface (CED 1401 and Spike2 software, Cambridge Electronic Design, UK).

MOUSE WIDEFIELD CALCIUM IMAGING RECORDINGS Experimental data acquired from mice have been provided by LENS (European Laboratory for Non-Linear Spectroscopy[4]) (Resta et al., 2020a; Resta et al., 2020b). All procedures involving mice were performed in accordance with the rules of the Italian Minister of Health (Protocol Number 183/2016-PR). Mice were housed in clear plastic enriched cages under a 12 h light/dark cycle and were given *ad libitum* access to water and food.

Mouse Model: The transgenic mouse line used is the C57BL/6J-Tg(Thy1GCaMP6f)GP5.17Dkim/J (referred to as GCaMP6f mice[5]) from Jackson Laboratories (Bar Harbor, Maine USA). In this mouse model, the ultra-sensitive calcium indicator (GCaMP6f) is selectively expressed in excitatory neurons (Chen et al., 2013; Dana et al., 2014).

Surgery and wide-field imaging: Surgery procedures and imaging protocols were performed as described in (Celotto et al., 2020). Briefly, 6 months old male mice are anesthetized with either a mix of ketamine and Xylazine in doses of 100 mg/kg and 10 mg/kg respectively or isoflurane ($3-4\%$ induction and $1.5-2\%$ maintaining). To obtain optical access to neuronal activity over the right hemisphere, the local anesthetic lidocaine (20 mg/mL) was applied and the skin and the periosteum over the skull were removed. Wide-field imaging was performed right after the surgical procedure. GCaMP6f fluorescence imaging was performed with a 505 nm LED light (M505L3 Thorlabs, New Jersey, United States) deflected by a dichroic filter (DC FF 495-DIo2 Semrock, Rochester, New York, USA) on the objective (2.5x EC Plan Neofluar, NA 0.085, Carl Zeiss Microscopy, Oberkochen, Germany). The fluorescence signal was selected by a band-pass filter

4 LENS Home Page, http://www.lens.unifi.it (accessed on Nov. 2019)

5 For more details, see The Jackson Laboratory, Thy1-GCaMP6f, https://www.jax.org/strain/025393 (accessed on Nov. 2019).

(525/50 Semrock, Rochester, New York, USA) and collected on the sensor of a high-speed complementary metal-oxide semiconductor (CMOS) camera (Orca Flash 4.0 Hamamatsu Photonics, NJ, USA). A 3D motorized platform (M-229 for the xy plane, M-126 for the z-axis movement; Physik Instrumente, Karlsruhe, Germany) allowed the sample displacement.

7.3 RESULTS

7.3.1 *Construction of an adaptable analysis pipeline (Cobrawap) to enable slow wave comparisons*

Since there is no single fully comprehensive measure to characterize spatial activity patterns, we focus on identifying common observables and analysis metrics of slow wave activity that enable a reasonable comparison between datasets of different measurement types. By investigating the relations of the characteristics metrics with the comparison parameters (scale, brain state, anesthetic level, spatial/temporal resolution, etc.), we may evaluate the capabilities of measurement methods, identify biases, constrain theories, develop and benchmark analysis methods, contribute to defining standards, as well as aid the assessment of clinical data of, for example, in the case of coma patients.

Based on the conceptual framework illustrated previously (Section 7.2.1, we demonstrate the creation of a specific analysis pipeline application aimed at the analysis and comparison of slow wave activity across heterogeneous datasets. For this, we consider 5 publicly available datasets from the EBRAINS KnowledgeGraph platform[6] of electrocochleography (ECoG) and calcium imaging recordings of anesthetized mice. Besides the measurement technique, the 60 examined recordings vary in a range of factors such as experimental setup, the genetic strain of the mice, anesthetic, anesthesia level, temporal- and spatial resolution, and recording duration (see Section 7.2.7).

We organized the pipeline into 5 sequential stages, successively transforming the raw data, detecting wavefronts, and extracting slow wave characterizations, as illustrated in Figure 7.6. In the following, we describe the role of each stage in the analysis of the ECoG and calcium imaging data. The selection and execution order for each of the data types is illustrated for each stage in Figure 7.7. The general functionality of the stages and blocks is further described in the Methods (Section 7.2.1) and in the corresponding README files in the Github Repository.

stage 1 **Data Entry:** In the first stage the data is being prepared for the later stages by loading, structuring, and annotating the data

6 https://search.kg.ebrains.eu/

and metadata according to the defined representation scheme using the Neo data format (Garcia et al., 2014)). For loading the data and processing the often unique structure, typically, each data source requires a custom loading script to integrate the information from a data file and a corresponding config file. The loading script can be adapted from a template, making use of utility functions and, ideally, the IO functionality of Neo[7]. After loading, it is checked if the data object conforms with the requirements and an overview of a data sample is plotted.

stage 2 **Processing:** The stage offers a series of blocks implementing basic processing steps that can be arbitrarily combined o fit the data type and their analysis objectives. Both data types undergo a background subtraction, a normalization, and a detrending step to remove eventual recording artifacts. Considering the different measurement modalities and temporal resolutions of the data types, additionally, the calcium imaging recordings are cut to a region-of-interest and filtered from 0.1 to 5 Hz, while the raw ECoG signals are transformed to a logMUA signal (see Section 7.2.2) with a reduced sampling rate better suited to capture the slow oscillations and the transitions between Down and Up states (which are detected in the third stage). Where available, the blocks use standard function implementations by the Elephant Electrophysiology Analysis Toolkit (Denker et al., 2018a), the stack of scientific Python packages (i.e. scipy, scikit, etc.), or algorithms from the literature.

stage 3 **Trigger Detection:** This stage provides multiple options of trigger detection methods for identifying the time stamps of upward transitions (and downward if possible) in each channel as an indicator for the possible passage of a wavefront. Since the ECoG logMUA signal shows sharp state transitions, they are best detected by a threshold determined from a channel-wise fit of the amplitude distributions (as in De Bonis et al. (2019)). Conversely, the transitions in the imaging data are determined by the slow activation function of the fluorescent molecules. Therefore, they are better detected by identifying the upwards slopes by either the Hilbert phase signal crossing a specific value (here $-\pi/2$) or by the local minima preceding a dominant peak (see Section 7.2.3). In the following, we use the trigger detection via the Hilbert phase, however, we focus on the quantitatively comparison of the two methods in Section 7.3.3. Additionally, there are optional filter blocks that can be applied to clean the collection of detected triggers, like the removal of unreasonably short Up or Down states. In the following, only upward transitions are considered as triggers. The output of stage 3 is a series of trigger

7 https://neo.readthedocs.io/en/stable/io.html

events that is stored alongside the signals in the output data object. Although this stage is tailored towards the detection of state transitions it has also proven useful for other wave-like activity beyond slow waves, demonstrated in Chapter 8.

stage 4 **Wave Detection:** Within this stage, the channel-wise trigger times are grouped to define the individual waves (see Section 7.2.4), describing the wave activity on a abstract but common level. This wave representation is again stored as a series of grouped trigger events in the data object. The wave description is optionally enriched with additional descriptions such as the optical flow Section 7.2.5 and the critical points of the resulting vector field (Townsend and Gong, 2018) or an additional clustering of the waves into modes based on the spatial arrangement of the trigger delays (Ruiz-Mejias et al., 2011).

stage 5a **Wave-wise Characterization:** The final stage calculates one or multiple characteristic measure on the basis of the common and measurement-independent slow wave description generated by the previous stages. This entails scalar measures as but may also contain metadata information like analysis parameters or information about the dataset added in stage 1 (selected via the `annotations` block). The selection of characterization metrics can be tailored toward addressing specific scientific questions or research objectives. To have a consistent output format for the final stage, there are two distinct realizations: one for a characterization using wave-wise measures, e.g., determining one velocity value per wave; and another for a characterization using channel-wise measures, e.g., calculating local velocity values per channel and wave. For simplicity, these alternatives are presented together in Figure 7.6A and Figure 7.7.

stage 5b **Channel-wise Characterization:** This alternative final stage is equivalent to the 'Wave-wise Characterization' in its functionality, but its characteristic measures are calculated per wave and channel (i.e. electrode or pixel). The output in either case is a pandas dataframe (McKinney, 2010) where each row represents either one wave or one channel in a wave and each column an attribute/characteristic. Consequently, the output for one dataset can be directly merged or compared with the output for other datasets and serves as the basis for various cross-domain comparisons (e.g., data comparisons, model validation, method benchmarking).

To demonstrate the capabilities of the pipeline approach to generate meaningful quantification of slow-wave phenomena, we choose four metrics as the basis for cross-dataset comparisons: the local (i.e., channel-wise) inter-wave interval, velocity, and direction measures; and the global (i.e., wave-wise) planarity measure.

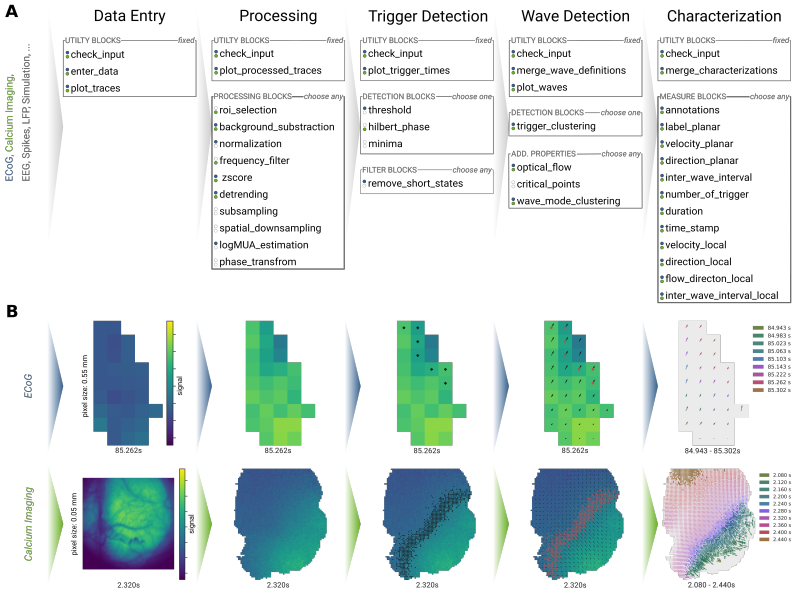

Figure 7.6: **Pipeline illustration.** This figure shows the progression of two test datasets (*ECoG* in blue, *calcium imaging* in green, see Section 7.2.7) through the analysis pipeline. **A:** The five successive stages contain a collection of modular blocks in three different selection modes (fixed, choose one, choose any). The analysis path is adaptable for specific datasets and analyses by selecting the desired blocks (indicated by colored dots for the test datasets) and by setting the corresponding parameters in the config file. Where reasonable, blocks either trigger the execution of a plotting block to illustrate the intermediate result or produce a plot as an additional output. **B:** The intermediate results after each stage are visualized for the two test datasets as color-coded signal on the electrode/pixel grid covering most of the right hemisphere of the mouse brain (up=anterior, right=lateral) with ECoG spanning 4.95×2.75 mm and calcium imaging 5×5 mm. From left to right, they show the analysis process from loading the raw data, processing the signal, detecting upward transitions (black marker), grouping them into wavefronts (red marker) together with their optical flow (arrows), to the example result of quantifying the linear flow alignment within the waves (i.e., planarity).

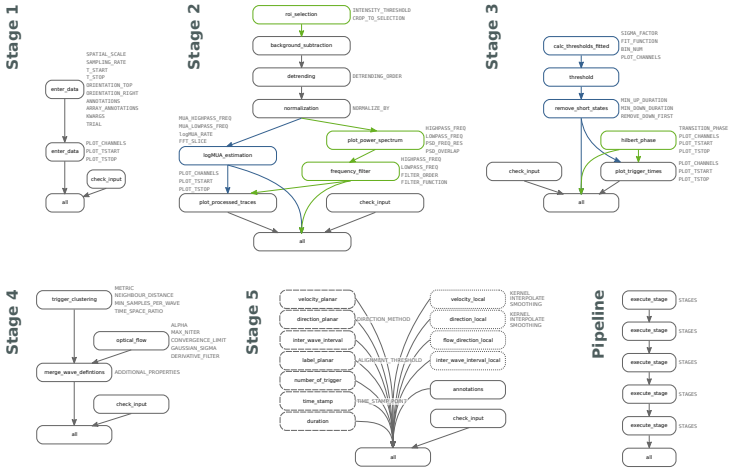

Figure 7.7: **Pipeline structure.** The diagrams shows for each stage of the pipeline illustrated in Figure 7.6 the execution order of the blocks, determined by the snakemake wokflow tool (Mölder et al., 2021). Blocks that are specific for the ECoG data are shown in blue, and blocks specific for wide-field calcium imaging in green, while common blocks are grey. In stage 5, dashed blocks are wave-wise measures and dotted blocks are channel-wise measures. Next to the blocks, the parameters are indicated that can be set in the corresponding config files.

The **inter-wave interval** is defined as the time delay between the occurrence of two consecutive waves at a recording site.

The channel-wise **velocity** is calculated from the derivatives of the delay function of a wave $T(x, y)$, which indicates when a wave has reached the position (x, y) in its propagation (Greenberg et al., 2018; Capone et al., 2019b; Pazienti et al., 2022):

$$v_{x,y} = \sqrt{\frac{1}{\partial_x T^2 + \partial_y T^2}} \tag{7.1}$$

The channel-wise **direction** of wave propagation can also be directly derived from the time delay function $T(x, y)$. However, in the following, we use the optical flow of the phase signal (see Section 7.2.5 for details). The optical flow is a continuous vector-valued signal for each position (x, y), indicating in which directions the contour lines of equal phase propagate. We define the channel-wise wave directions of a propagating wave as the optical flow vector directions at the time and position of its trigger events.

The **planarity** of a wave is also defined via the optical flow as the absolute value of the normalized channel-wise direction vectors at the times and positions of all trigger events that belong to a wave, quantifying their alignment on a scale from 0 to 1:

$$\text{planarity} = \frac{||\sum \vec{v}_i||}{\sum ||\vec{v}_i||} \tag{7.2}$$

For one of the calcium imaging recordings, Figure 7.8 presents some of the measures in the pipeline output, i.e., a table (pandas dataframe (McKinney, 2010)) containing all wave characteristics. The channel-wise and wave-wise measured directions and velocities, as well as the wave-wise planarities, are summarized for 4 groups of similar waves, i.e., "wave modes". The wave-mode clustering method (implemented as an optional block in stage 4 of the pipeline) applies a k-means clustering on the trigger delay matrix containing the relative trigger times for each channel in each wave (Ruiz-Mejias et al., 2011; Capone et al., 2019b; Pazienti et al., 2022). The number of modes was set by hand to reasonably represent the variability of wave types in the recording. Generally, the 'optimal' number of modes to set for the k-means algorithm depends on recording and the analysis application. For the presented recording, most waves are relatively planar and travel along the lateral-posterior-to-medial-anterior axis (modes #2 and #4), mode #1 is a variation of mode #4 with a lower average velocity, and mode #3 contains only one wave that seems to be a collision of two opposing wavefronts. Although the channel-wise and wave-wise measures for the direction and velocity (Figure 7.8B&D)

are defined and calculated in a completely different way, they agree considerably well for the modes #1, #2, and #4 when the wave pattern is predominantly planar. The complex wave pattern of mode #3 can not be accurately captured by a single wave-wise value for the direction and velocity, resulting in the presented deviations. Otherwise, the different measures provide a coherent characterization of each wave mode. In the following, we only use channel-wise measures for the analyses, with the exception of the wave-wise planarity measure, which has no channel-wise equivalent. An analog figure for a selected EcoG recording is presented in Figure 7.9.

7.3.2 *Evaluation of experimental influences via slow wave comparisons*

Based on the exemplary realization of the pipeline, we are now in a position to perform quantitative comparisons of slow wave dynamics across the described ECoG and calcium imaging datasets, contrasting various experimental parameters. In the following, we demonstrate the application of the pipeline to investigate the influences of the anesthetic type and dosage, the application of disease models via genetic knock-out, and the measurement technique itself, in particular, its spatial resolution.

To check the validity of the pipeline, we first qualitatively replicate some of the results that were previously published using the same datasets. It has been shown that the dynamics of slow waves crucially depend on the level of anesthesia. While the velocity of waves tends to decrease slightly in deeper anesthesia states (Pazienti et al., 2022), the inter-wave intervals become more prolonged, i.e., the frequency of waves decreases (Pazienti et al., 2022; Dasilva et al., 2021). The same trends are visible in the corresponding pipeline output presented in Figure 7.10A. The velocity and frequency of slow waves were also measured in the context of a disease model for Williams-Beuren Syndrome (WBS) in knock-out (KO) condition and wild-type (WT) (of the same genetic strain) (Davis et al., 2020; Sanchez-Vives, 2019b). In both the previous publication and the pipeline output (Figure 7.10B, there is no major effect on the wave characteristics visible except for a slight increase in the measures variances in the knock-out condition (Figure 7.10B). Including yet another dataset from a similar experiment that models the Fragile-X Syndrome (FXS) (Sanchez-Vives, 2019a) allows us to expand the analysis across experiments. Ignoring the genetic knock-out and only looking at the wild-type control subjects, we can compare the influence of other experimental parameters between the WBS and FXS experiments. A notable difference between the two experimental setups is that in the WBS experiment, ketamine was used as the anesthetic (100 mg/kg inducing + 37 mg/kg maintaining), while the FXS experiment used isoflurane (4 % inducing + 1 % maintaining). Comparing the distributions for the wild-type

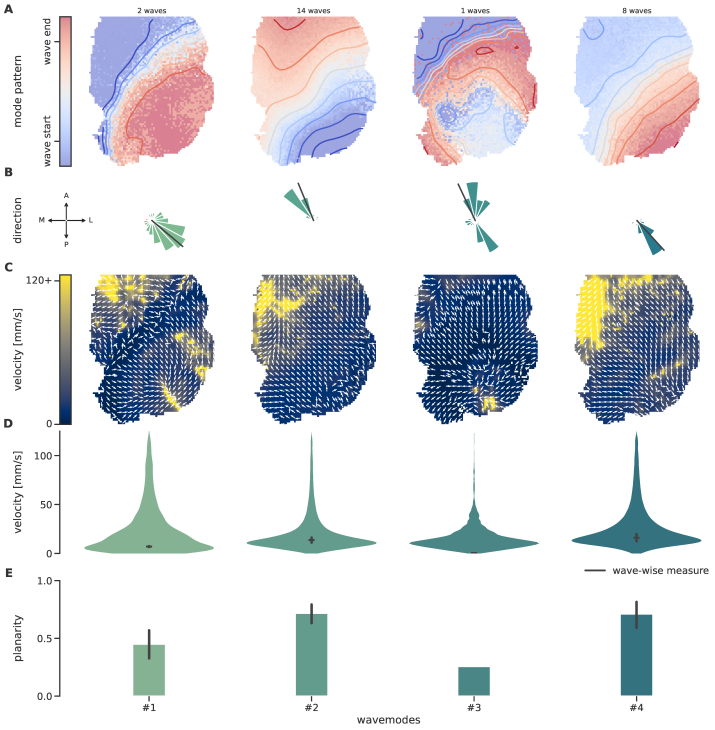

Figure 7.8: **Representation of the pipeline output for one exemplary calcium imaging recording.** Within the pipeline, the optional block wave_mode_clustering groups together similar waves into wave modes. The characterization of the waves in each of the 4 modes is shown in the corresponding columns (in no specific order). **A:** The average wave pattern (number of waves indicated on top) is illustrated as a time-delay heatmap with iso-delay contours. **B:** The aggregated histogram of channel-wise directions in waves of this mode. The black lines indicate the average wave-wise direction measure. **C:** Map of the average channel-wise velocities in waves of this mode, overlayed with the average channel-wise direction determined via the optical flow. **D:** The corresponding distributions of channel-wise velocities and, as black ticks and errorbars, the average and 95% CI of the corresponding wave-wise velocities. **E:** The average and 95% CI of the planarity values for the waves of this mode. An analogous example figure for an ECoG recording is shown in Figure 7.9.

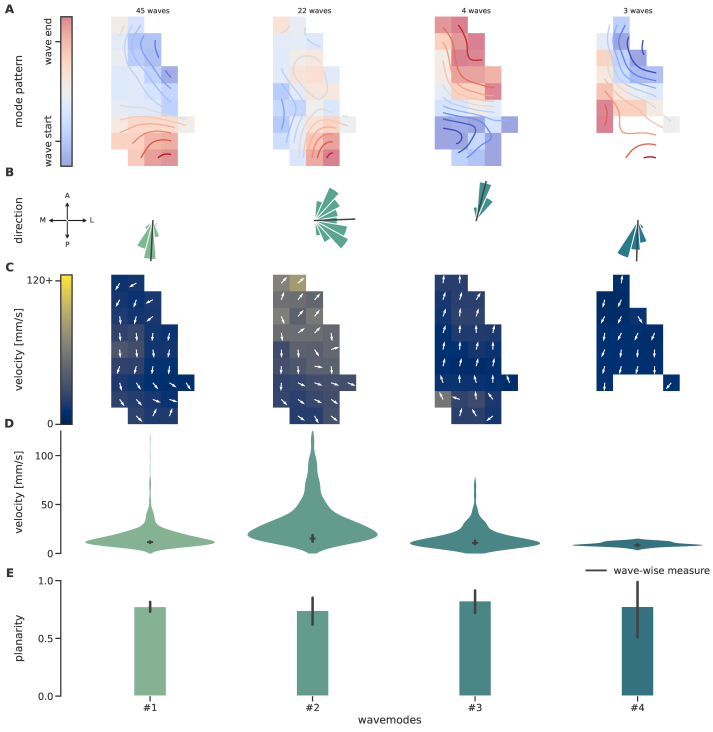

Figure 7.9: **Representation of the pipeline output for one exemplary ECoG recording.** The figure is analogous to Figure 7.8.

mice shows significantly larger velocities measured in the experiment that used isoflurane and a larger range of inter-wave intervals for the experiment that used ketamine (Figure 7.10B). In comparison to Figure 7.10A, where anesthesia was induced with ketamine (75 mg/kg) but maintained with isoflurane (0.1-1.16 %), we see a better agreement to the velocities in the WBS (ketamine) experiment than to the FXS (isoflurane) experiment. However, it is to be noted that since this is a meta-analysis, there is little control for confounding parameters between the different datasets. So, the attribution of the differences in wave characteristics to one parameter, here the anesthetic type, has to take this into account.

Next, we are expanding the scope of the analysis and contrast the ECoG recordings of ketamine- and isoflurane-anesthetized mice to analogous recordings of anesthetized mice that use calcium imaging to measure the cortex activity (Resta et al., 2020a; Resta et al., 2020b). Figure 7.10C illustrates the distributions of wave characteristics grouped by measurement technique and anesthetic type. A principal difference between the measurement techniques is their spatial resolution. The calcium imaging data has a resolution of 0.05 mm compared to 0.55 mm for the ECoG data. The finer resolution allows for a better distinction of complex non-planar wave patterns, as can be seen by the broader planarity distribution that is shifted towards smaller planarity values. Additionally, complex wave patterns with low planarity seem to be more prevalent under isoflurane-induced anesthesia than with ketamine-induced anesthesia, which can also be seen to a smaller extent in the ECoG recordings. Furthermore, the detected waves in the calcium imaging data are more frequent and regular, as shown by the inter-wave-interval distributions. Regarding the wave velocity distributions, there is a notable discrepancy between the measurement techniques for the isoflurane datasets, while the velocities for the ketamine dataset are quite similar. This considerably difference in wave velocities is likely related to a difference in the isoflurane concentration (1% in EcoG and $1.5-2\%$ in calcium imaging recordings), as even small differences in the concentration can have a considerable effect on the wave dynamics (cf. Figure 7.10A&B). The slow waves we detect with the pipeline tend to propagate along a preferred axis and primarily in one direction. This axis seems to be approximately consistent within the data of each measurement technique but not across. In the ECoG data, the preferred propagation axis spans from posterior-medial to anterior-lateral, with the preferred direction being different for the isoflurane and ketamine datasets. In the calcium imaging data, the preferred wave direction is from posterior-lateral to anterior-medial. Wave propagations being principally oriented in a back-to-front or front-to-back manner is also reported in previous studies (Massimini et al., 2004; Nir et al., 2011; Ruiz-Mejias et al., 2011; Sheroziya and Timofeev, 2014; Greenberg et al., 2018; Pazienti

et al., 2022). The spread of the wave direction histogram around the preferred directions can be either caused by a variance of channel-wise directions between waves or within waves, e.g., low-planarity waves have, per definition, a broader spread of channel-wise directions. To further investigate the observed differences in the wave characteristics between the ECoG and calcium imaging data, we eliminate the aspect of their different spatial resolution by spatially downsampling the calcium imaging data. By downsampling in a stepwise manner up to a downsampling factor of 11, for which the spatial resolution is then equal to the one of ECoG (0.55 mm), we can disentangle the effects the resolution has on the measured wave characteristics. Figure 7.10D shows how the distributions of wave characteristics change as a function of the downsampling factor. With a decreasing spatial resolution, fewer waves are detected, and they appear more planar as some complex local patterns are no longer detected. This effect is particularly obvious for the isoflurane datasets. A similar effect on the probability of detecting a planar wave as a function of ROI size has been previously shown by Liang et al. (2021). As mentioned before, there are relationships between the planarity of a wave and the corresponding measured channel-wise velocities and directions. The fact that the direction histogram of the downsampled calcium imaging data is more narrow indicates that the propagation directions are consistent across waves, and the variances observed in Figure 7.10C are caused mainly by non-planarity. With increasing planarity, in particular, the waves in the isoflurane datasets also exhibit faster channel-wise velocities that surpass the ketamine wave velocities, as is also visible in the ECoG data.

So, we demonstrate how the adaptable pipeline approach enables the comparison of slow wave characteristics across heterogenous datasets, including electrical and optical acquisition methods. This meta-analysis illustrates distinct differences within the aggregated data and potential dependencies on the experimental parameters to be further investigated.

7.3.3 *Method benchmarking via slow wave comparisons*

While applying the same analysis method to different data enables rigorous comparisons, applying different methods to the same data allows investigating the influence of the choice of the method itself. In the analysis of slow waves, the method for detecting the transitions from Down to Up states plays an important role. So far, in the calcium imaging data, we detected the trigger times at the upstroke of the transitions as the Hilbert phase of the signal crossing a threshold value of $-\frac{\pi}{2}$ (Section 7.2.3). However, there also exist alternative methods to define potential trigger times, such as using the local minima of the filtered signal, for example, done in Celotto

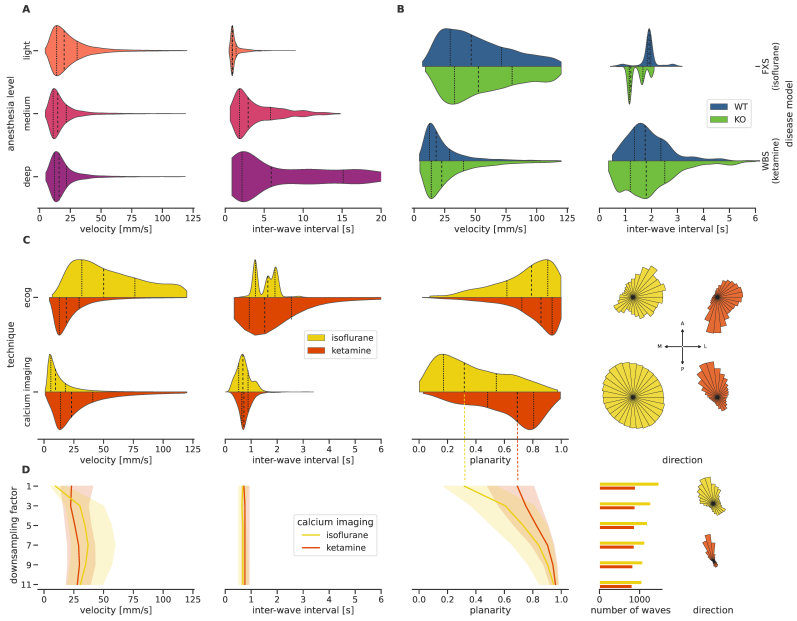

Figure 7.10: Quantitative comparison of slow waves across heterogeneous datasets. Violin plots show sample distributions with indications of the median (dashed line) and the quartiles (dotted lines). Line plots also show the median and quartiles (shaded areas). Polar plots show the distributions of wave directions in the right hemisphere so that up corresponds to an anterior direction and right to a lateral direction. **A**: Velocity and inter-wave intervals of slow waves in ECoG recordings as a function of the anesthesia level. **B**: Velocity and inter-wave intervals of slow waves in ECoG recordings of experiments modeling Willems-Beuren Syndrome (WBS) and Fragile-X Syndrome (FXS) split into wild-type (WT, blue) and knock-out (KO, green) subjects. **C**: The ECoG data from panel B is compared to calcium imaging data, split into anesthetic types, concerning wave velocities, inter-wave intervals, wave planarity, and wave direction. **D**: Effect of stepwise spatially downsampling the calcium imaging data from 0.05 mm (factor 1) to 0.55 mm (factor 11), the spatial resolution of the ECoG data. The wave direction histograms are only shown for the fully downsampled data (with factor 11).

et al. (2020). Figure 7.11 illustrates the influences the two different detection methods have on the resulting wave characteristics. Realizing this method-benchmarking workflow with our pipeline only requires selecting the corresponding detection block in the config file of the wave detection stage and to rerun the analysis on the calcium imaging data. As shown in Figure 7.11A, the resulting triggers differ clearly in number and exact timing, influencing the timing and number of detected waves. As a result, the representation of individual waves can differ to varying degrees (see an example in Figure 7.11B). Nevertheless, the overall distributions of wave characteristics remain very similar (Figure 7.11C). The slightly shorter inter-wave intervals for the minima-detected waves can be attributed to the increased number of detected waves.

This method comparison allows us to evaluate the strength of each way of detecting upward transitions and check for potential biases introduced in the wave characterization. The minima detection method is less strict and therefore detects more waves, including some smaller local ones. The Hilbert-phase method detects fewer waves, which are, however, better separated and more coherent across channels. Thus, we see that finding slow waves in different use cases is also a matter of applying the appropriate methods while being aware of their advantages and disadvantages. For a more extensive method comparison, including specific edge cases, this approach can be further combined with simulated data.

7.3.4 *Model calibration via slow wave comparisons*

Applying the Collaborative Brain-Wave Analysis Pipeline to experimental data enables the extraction of key spatio-temporal characteristics of slow waves acquired with multiple experimental methodologies at local and multi-areal spatial resolution. These insights are, of course, beneficial to modelers as they can inform the model development process and offer new approaches to exploit data to introduce complexity and biological realism in the models. However, the pipeline can also be directly applied to simulation outcomes to extract the same characteristics as from the experimental data to perform quantitative comparisons for model calibration and/or validation.

Here, we demonstrate how a mean-field neural network model can be analyzed with our pipeline and adapted to match the slow-wave characteristics of a calcium imaging dataset (Section 7.2.7). A detailed account of this modeling study is published in (Capone et al., 2023). The network model consists of approximately 1400 interacting populations, representing the pixels in the corresponding calcium imaging data. Each population entails 500 excitatory and inhibitory current-based AdEx (Adaptive Exponential integrate-and-fire) neurons and is modeled by a first-order meanfield equation (Capone et al.,

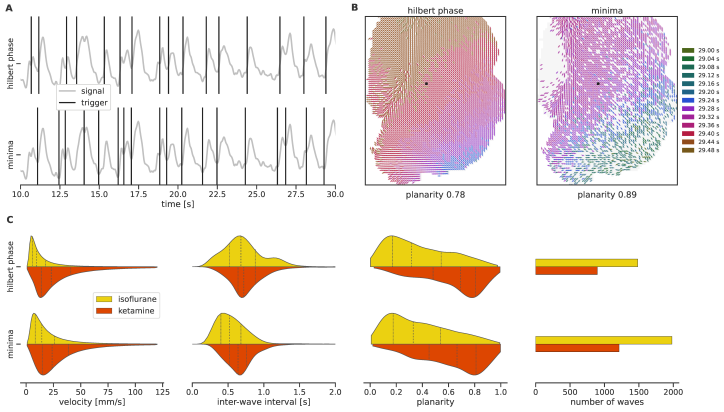

Figure 7.11: Comparison of trigger detection methods in calcium imaging data. **A**: Upward transitions (triggers, black vertical lines) found by two different detection algorithms, looking for either local minima or a crossing of the Hilbert phase $= -\pi/2$, in an example signal taken from the black-colored pixel indicated in panel B. **B**: The same exemplary wave (corresponding to the last trigger in panel A) illustrated over the recorded area as detected using the two trigger detection algorithms. The arrows indicate the local direction of the wavefront at the time of the trigger, which is also encoded as the color of the arrow. **C**: The distributions of the comparative characteristics of the resulting waves with the two different methods in the calcium imaging datasets.

2015; Capone et al., 2019a; El Boustani and Destexhe, 2009; Gigante et al., 2007), so that all the neurons in the population are described by their average firing rate. It has been shown that the spike-frequency adaptation of neurons is a crucial parameter in realizing a network state exhibiting slow-wave activity (Mattia and Sanchez-Vives, 2012; di Volo et al., 2019; Tort-Colet et al., 2021; Cattani et al., 2022). Therefore, a parameter for the adaption strength b is included in the model. Furthermore, we incorporate anatomical information to define connectivity priors (Schnepel et al., 2015), including the fact that long-range inter-areal connection are suppressed during the expression of slow waves (Olcese et al., 2016). This is represented by kernel functions for the connection probabilities between populations that have an elliptical shape on the cortical grid plane and are exponentially decaying with the distance. So, that the connectivity of each population is parameterized by the connectivity weight k_0, connectivity anisotropy a, the eccentricity of the elliptical kernel function e, and the spatial decay exponent λ. Each population is further driven by an external input I_{ext}.

The model is optimized in two steps, specifically for a 40 s recording of the experimental data. In the first step, the connection parameters for each model population are optimized by a likelihood maximization (using the gradient-based optimizer iR-prop (Igel and Hüsken, 2003)) with respect to the activity of a pixel in the recording. This parameter inference is performed on the 32 s of the recording. The remaining 8 s are used for a validation of the inference procedure to avoid overfitting. As the fluorescent indicator (Thy1-GCaMP6f) used in the experimental recordings targets specifically excitatory neurons, we can only consider the excitatory activity to constrain the model. The inhibitory activity is instead considered as hidden variable and determined via an adiabatic approximation (Mascaro and Amit, 1999). The resulting neural network model is able to produce simulated activity corresponding to a slow wave activity state (Figure 7.12B). However, the simulated slow wave dynamics are repeating nearly identically. Since the applied model constraints are local in time, they are not able to reproduce the variability over time that is found in the experimental data.

Therefore, the model optimization requires a second step, a calibration of a neuromodulatory input. This neuromodulation is expressed as an oscillation of the I_{ext} and b parameters with an amplitude A and period T around their previously determined channel-wise values (inspired by Goldman et al. (2019)). The calibration uses the characterization of the slow wave activity in the experimental and the simulated data resulting from our presented analysis pipeline. As illustrated in Figure 7.13, we compare the local velocity, direction, and inter-wave-interval measures individually using the Wasserstein distance (Arjovsky et al., 2017), and used the square root of their quadratic sum

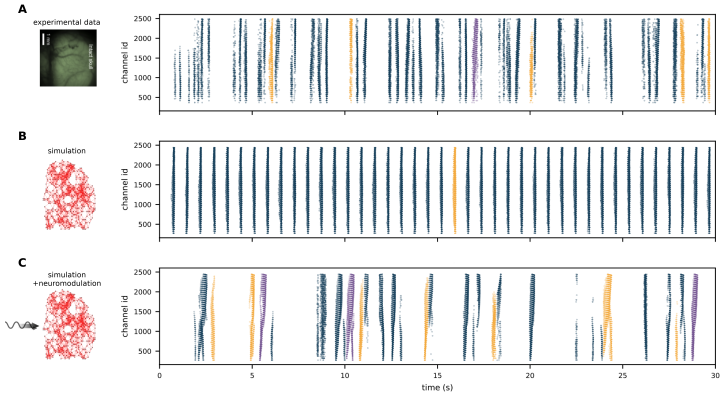

Figure 7.12: **Raster plots of experimental and simulated slow wave activity**. The raster plots on the right-hand side display the transition times from Down to Up states in each recording channel (pixel), so that a vertical group of trigger points indicates a slow wave passing over the recorded area. The colors of the waves indicate different wave modes (not discussed). **A)** Experimenal data from the right hemisphere of an anesthetized mouse, recorded via calcium imaging. **B)** Simulated data from a mean-field neural network model with connectivity parameters inferred from the experimental data via a likelihood optimization. **C)** Simulated data from the same model as in B, but with an additional oscillatory neuromodulation that is calibrated to match the slow wave characteristics (velocity, direction, inter-wave interval) of the experimental data. Figure adapted from Capone et al. (2023)

as a combined score. Using a grid search for A and T, we minimize the distance score to determine optimal neuromodulatory parameters. The resulting simulated slow wave activity has a much better resemblance of the target experimental data, reasonably reproducing most of the features of the non-stationary and non-linear dynamics of the slow wave activity (cf. Figure 7.12A&C).

7.4 CONCLUSION

7.4.1 *Comparability and reusability*

The use-case-specific slow wave analysis illustrates the benefits of a structured, modular pipeline approach. These benefits, however, are not limited to slow wave analyses but can contribute generally relevant aspects to rigorous data analysis procedures. So, the pipeline features built-in mechanisms (largely due to the use of snakemake) to retain information about the analysis scripts, their execution order, and parameter settings. The stage and block documentation and the visualization of their respective outcomes help to follow the workflow in detail to redo the analysis, but also to check results and increase confidence in the findings. By further aligning the workflows for different datasets within the pipeline using the same or analogous analysis methods (while still catering to their diverse processing demands), the individual results become more comparable. This does not mean that the results are expected to be the same, but by controlling the variations in the analysis workflow, the remaining differences (measurement techniques, species, anesthetics, etc.) can be reasonably evaluated and quantified. The flexibility to deal with the input of very different datasets also trivially extends to simulated data. Therefore, the pipeline design also provides a basis for the validation testing of models. The modularity that enables the pipeline to be adapted to different datasets also brings the additional advantage that the individual analysis steps become interchangeable and can thus be evaluated on their own. Analyzing the same dataset with such an adapted pipeline lets us test the influence of a method on the downstream results. Eventually, the gain of such a pipeline design extends even further as all its elements work as well as standalone components and can be reused for other pipeline or non-pipeline applications. Finally, the goal is not to fully automate data analyses but to provide a tool base to make the creation of analysis workflows more efficient, rigorous, and reproducible. No software can replace a thoughtful scientist, but the right software can give a scientist more room to think.

As the pipeline is built up in a modular fashion, there are multiple levels of reusability. The simplest is to use the adaptability via the config files to change the selection, order, or parameters of blocks to modify the pipeline execution. For larger, more specific changes to the

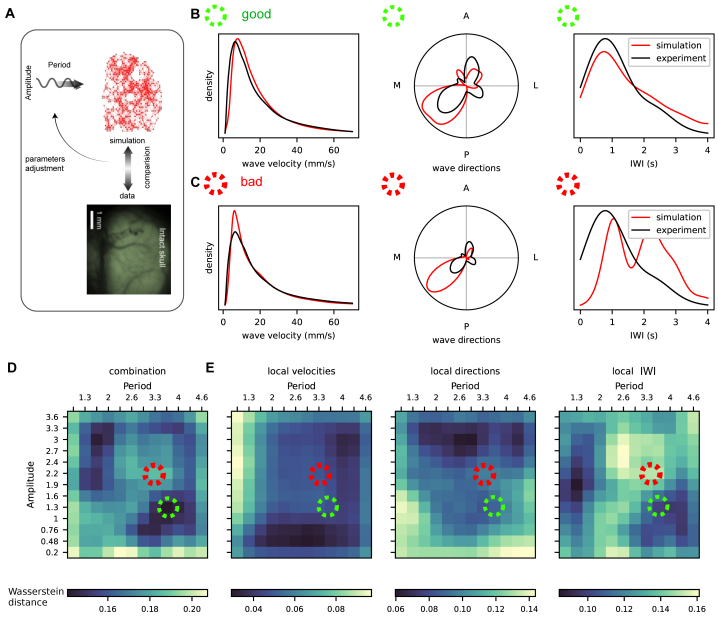

Figure 7.13: **Calibrating the oscillatory modulation of the mean-field model. A)** Schematic of the calibration step. The mean-field network model with inferred connectivity parameters is simulated with an additional oscillatory neuromodulation and the resulting slow-wave characteristics compared to the experimental data in order to determine the ideal oscillation paramters. **B,C)** The comparison is based on the local wave velocity, direction, and inter-wave interval distributions. The agreement of the simulation and experimental data depends on the parameter setting, panel B and C showing a good and a bad examplary parameter setting. **D,E** The agreement between the slow-wave characteristics is quantified by the Wasserstein distance (color scale). The heatmaps illustrate the results of a grid search for the amplitude and period of the neuromodulatory oscillations via the individual Wasserstein distances (panel E) and the combined distance (panel D). Figure adapted from Capone et al. (2023)

pipeline, new blocks can be added, for example, to integrate a new dataset or to carry out data analysis focused on other observables. For this, the pipeline repository provides corresponding templates and instructions. There is no principal restriction of types of data: as long as it conforms to the minimum requirements (e.g., an electrode/pixel layout on a grid), the pipeline may be applied on EEG, LFP, or even spiking data. While both blocks and stages are designed to also work as stand-alone analysis snippets in other workflows, most notably, the pipeline design allows deriving new analysis pipelines for other applications. Such a reuse is demonstrated in Chapter 8, to analyze the the cortical wave activity (e.g. in the alpha, beta, and gamma regime) in behaving monkeys.

7.4.2 *How a structured analysis pipeline can contribute to progressing the study of slow waves*

The presented meta-analysis across heterogenous datasets comprises ECoG and calcium imaging recordings (Figure 7.10). These measurement techniques are known to result in fundamentally different signals. While ECoG tends to record only the spiking activity of neurons in the superficial layer with a high firing rate and a high signal-to-noise ratio with high temporal and low spatial resolution, imaging Genetically Encoded Calcium Indicators (GECIs) such as GCaMP6s, measures population activity from specific cell types (i.e. excitatory neurons) in layers 2/3 and 5 as a delayed, low-pass filtered, non-linearly transformed fluorescence signal with low temporal and high spatial resolution (Siegle et al., 2021a; Siegle et al., 2021b; de Vries et al., 2020). Therefore, even elaborate models can't fully capture all aspects of their complex relationship, and there is generally no precise agreement between results derived from the two measurements beyond coarse qualitative measures (Wei et al., 2020; Chen et al., 2013; Stringer et al., 2019). In this context, our analysis pipeline can quantitatively illustrate the differences in the measurement types regarding the characteristics of slow wave activity. For analyzing wave characteristics, it is of particular interest from which cortical layer the measurement technique samples the recorded neurons since aspects like frequency power or propagation speed are known to vary considerably with cortical depth (Halgren et al., 2018; Capone et al., 2019b). Taking these considerations into account, an integrative approach of using multiple measurement techniques may benefit from the complementing viewpoints that, for example, ECoG and calcium imaging can provide.

Besides the biases of the measurement technique and its resolution, we further present the influences of the anesthetic type and dosage on the wave characteristics, showing in particular that ketamine tends to produce more planar waves than isoflurane, in turn also influencing the measured directions and velocities. This effect is likely linked to

the known attributes of the anesthetics, that ketamine is more effective in generating slow-wave activity as it increases LFP power in the delta frequency band while isoflurane rather enhances LFP activity in the theta band and above (Michelson and Kozai, 2018; Purdon et al., 2015).

The need to quantitatively relate results from the literature to each other becomes quite apparent when, for example, investigating the sources of variance of the velocity of slow waves, which can vary from a few mm/s in in-vitro measurements and in-vivo recordings of anesthetized rodents up to ~ 10 m/s in human sleep experiments (Massimini et al., 2004; Ruiz-Mejias et al., 2011; Muller et al., 2016). Studied influences to this variability include the extent of axonal projections (Massimini et al., 2004; Golomb et al., 1996), axonal conductances (Ruiz-Mejias et al., 2011), involved cell types (Bazhenov et al., 2002), neuronal excitability via the availability of neuromodulators (Destexhe et al., 1999a) and via cortico-cortical or cortico-thalamic loops (Mohajerani et al., 2013; Sanchez-Vives and McCormick, 2000). Furthermore, the velocity of a wave may depend on its direction, which in turn is influenced by an interplay of local and global connectivity properties and frequency effects (Galinsky and Frank, 2020; Mohajerani et al., 2013; Massimini et al., 2004). Comparison between data from different studies can help relate and discern such influences.

While the exploration of wave characteristics under different conditions can provide further insight into the understanding of the underlying processes, experimental data can also suffer from constraints in the data size, parameter regime, and uncontrolled confounds. Therefore, in many scenarios, it is crucial to be able also to include modeling data in the analysis. The following subsection discusses a use-case where the adaptability of the pipeline to simulation output is exploited in the context of model calibration.

7.4.3 *Conclusion*

In this paper, we show how formalizing analysis approaches can have miscellaneous advantages. By taking a data-science perspective, we work towards integrating heterogeneous insights from different data and analysis types. Furthermore, we experienced how structuring our methodology and implementation also contributed greatly to our structure of thought. We are confident that this work pushes forward the requirement of reusability for analysis resources and succeeds in promoting making use of and contributing to community software projects (like Elephant, Neo, Nix, Snakemake, and many more). Because in the end, the scientific community is more efficient when everyone can directly benefit and build on each other's work and can rigorously compare their results without re-implementing the wheel.

8

CORTICAL WAVE DYNAMICS IN BEHAVING MONKEYS

This Chapter is not based on a publication. The following people are/were involved:

- Robin Gutzen, Simon Essink, Alexander Kleinjohann, Cristiano Köhler, Regimantas Jurkus, Junji Ito, Julia Sprenger, Lucio Condro, Shrabasti Jana, Frédéric Barthélemy, Marcel De Haan, Thomas Brochier, Alexa Riehle, Sonja Grün, Michael Denker

Author contributions:

- RG, MD, SG conceptualized the study.
- MD, SG supervised the project.
- MDH, SG, AR, TB, FB designed and set up the experiment.
- LC, SH, FB, MDH acquired the experimental data.
- FB, JS, SE, AK, CK, RJ, JI, MD cleaned and formatted the experimental data.
- RG implemented, performed, and visualized the analyses.
- RG wrote this chapter.

8.1 INTRODUCTION

We are continuously objected to a ton of sensory input that is signaled to our brain. Additionally, there is neural activity from all kinds of internal processes, including learning, planning, attention, memory recall, emotion, homeostasis, daydreaming. Still, the immensely intricate information processing mechanisms manage to selectively filter the relevant sensory input in a given context and integrate it with existing activity and structures to generate corresponding reactions. These processing mechanisms feature both parallel and sequential processing (Parks et al., 1991; Musslick et al., 2016) as well as bottom-up and top-down signal integration (Buschman and Miller, 2007). All this happens continuously. Executing motor actions is simultaneous to perception, which again informs action. Examples of this enormous feat include

hand-eye coordination tasks and active sensing (Schroeder et al., 2010). Most of the exact working mechanisms that are involved are still unknown, while our approaches to probing the neural systems are limited to recording tiny fractions of the neural activity (Section 1.3). However, with carefully designed experiments and systematic analysis of behavioral and neurophysiological measurements, we can glimpse into the activity signatures of these underlying processes.

Here, we explore the sensory integration and cortical information processing in a hand-eye coordination task performed by a free-viewing monkey. The monkey is instructed to reach for illuminated targets that light up in different sequences. We measure the neural activity with implanted multielectrode arrays in the motor cortex and areas $V1$, $V2$, DP, and $7A$ in the visual cortex. We focus our analysis on the wave dynamics of the local field-potential (LFP) signal as a high-level representation of the neural network activity. Previously, it has been shown that the oscillatory LFP activity in the cortex is organized in traveling wave patterns which correlate to behavioral events and have been suggested to organize neural processes across space and time (Ermentrout and Kleinfeld, 2001; Rubino et al., 2006; Zhang et al., 2018; Muller et al., 2018).

We characterize the wave activity over the course of a recording session and compare it across the visual areas, frequency bands, and behaviorally relevant task events. We find distinct wave activity behaviors for the individual areas and frequency bands. Furthermore, we identify a wave pattern signature that is uniquely selective and directionally tuned to a relative transition from one visual target to its upper right neighbor.

8.2 METHODS

8.2.1 *Data acquisition*

The experimental data were acquired in a real-time visuomotor behavior and electrophysiology recording setup at the Institut de Neurosciences de la Timone in Marseille, France, and is described in detail in (de Haan et al., 2018). The setup is in a dark environment and includes a chair for a monkey with a projection surface in front of the chair (Figure 8.1A). The head of the subject is fixated with a custom mask to allow for an accurate recording of the free-viewing eye movements via an infrared LED and camera[1]. The right arm is attached to an exoskeleton[2] that allows for arm/hand movements in the horizontal plane while recording the hand position. The plane of the hand movement is underneath the non-transparent projection surface. A cursor (a round dot) indicates the hand's position on the projection

1 EyeLink 1000; SR Research; https//www.sr-research.com
2 KINARM; BKIN Technologies; www.bkintechnologies.com

surface. The cortical activity is recorded with implanted multielectrode arrays[3]. The simultaneous recordings of the hand movements, eye movements, trial events, and neural activity are collected from the individual recording systems, temporally aligned, and integrated into a common data object.

The experiment and neural recordings are performed with a macaque (*Macaca mulatta*) monkey (subject Y). The subject was implanted with multiple multielectrode arrays (inter-electrode distance 0.4 mm) in the left hemisphere of the cerebral cortex. A 10×10 array was placed in the motor/dorsal pre-motor cortex (M1_PMd), and four 6×6 arrays were placed along the visual pathway, covering the areas $V1$, $V2$, DP, and $7A$ (Figure 8.1B).

Here, we use a data recording acquired during the performance of a visually guided motor task. On the projection surface, six target locations are arranged in a hexagon, with an additional target in the center. In each trial, a sequence of four targets is presented. One target is illuminated at a time, and the next target lights up when the hand cursor reaches the target location (within a radius of 0.65 cm). A trial is successful if the subject reaches the sequences of illuminated targets with less than 1.5 s for each target. There are 12 different sequences, each starting with the central target position, that are presented randomly during a session. If a trial is unsuccessful, the same sequence is presented again. Here, we use a session (*y180306-land-001*) with 176 trials, of which 120 are successful, with a total duration of 658.12 s.

3 Utah Array; Cerebus, Blackrock Microsystems; http://blackrockmicro.com

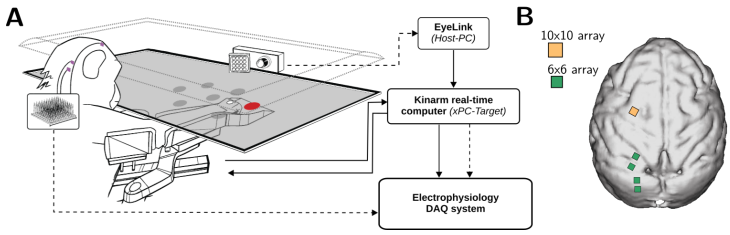

Figure 8.1: **Setup of the Vision-for-Action Experiment. A** schematically illustrates the components of the experimental setup and the neural and behavioral recording devices. The flowchart shows the flow of the corresponding analog (solid lines) and digital (dashed lines) signals. **B** shows the approximate placements of the Utah electrode arrays in the left hemisphere of the monkey. A 10×10 array (orange) is located in the border of the primary motor cortex, and dorsal premotor cortex (M1_PMd), and four 6×6 arrays (turquoise) are located in the visual cortex in the areas V1, V2, DP, and 7A. Adapted from de Haan et al. (2018).

8.2.2 *Data processing and analysis*

ANALYSIS PIPELINE The raw electrode signals (30 kHz) are downsampled to 1 kHz and with the behavioral data, trial events, and the various metadata stored as a combined object in the Neo format (Garcia et al., 2014). This data object is further processed and analyzed with the Collaborative Wave Analysis Pipeline (Cobrawap), which we introduced in Chapter 7. When entering the pipeline, the data is further annotated with information about the electrodes' spatial layout on the arrays and the cortical surface. In the pipeline's processing stage, the signal of each channel is centered around 0, 1st-order detrended, frequency filtered, and finally z-scored. The frequency filter (using a 4th-order 'sosfiltfilt' Butterworth filter) is applied to separate four frequency domains: *subalpha* ($1 - 10$ Hz), *lowbeta* ($10 - 23$ Hz), *highbeta* ($23 - 35$ Hz), and *gamma* ($60 - 90$ Hz), see Figure 8.2. The next stage detects the timing of potential wavefronts within each channel, defined by the oscillation phase (calculated with a Hilbert transform) crossing a value of $\phi = 0$. Subsequently, these trigger times in all channels are grouped into wavefronts via a density-based clustering

algorithm (using the DBSCAN function from the scikit-learn Python package (Pedregosa et al., 2011)). Further, we augment the filtered LFP signals with an optical flow vector-field signal, calculated with a Horn-Schunck algorithm, to indicate the instantaneous spatial propagation direction of the wavefronts (Townsend and Gong, 2018). Additionally, we group similar types of waves by applying a k-means clustering ($n = 4$) on the channel-wise trigger times, defining 4 wavemodes (Ruiz-Mejias et al., 2011).

CHARACTERISTIC WAVE MEASURES From the output of stage 4 of Cobrawap "Wave Detection," we extract continuous direction angle and planarity measures from the optical flow vector field. The **continuous direction angles** are defined per channel as the counter-clockwise angle from the posterior direction (positive horizontal axis in the here presented figures). The **continuous planarity** is defined per array as the normalized length of the sum of individual unit-length optical flow vectors, measuring the alignment of the vector field on a scale from 0 to 1. While these two measures do not depend on the detection of the individual wavefront, we additionally use the following wave-wise and channel-wise characteristic wave measures resulting from stage 5 of Cobrawap "Wave Characterization". Per wave, we also define a **wave-wise planarity** measure analogous to the continuous planarity but based on the optical flow vectors at the times and positions of wavefront triggers. We measure the **wave-wise inter-wave interval** as the time delay between successive waves averaged over channel locations. Further, we measure the **wave-wise velocity** by interpolating the horizontal and vertical locations of the wavefront trigger on the array as a function of time. Per channel per wave, we use the **channel-wise wave direction** and the **channel-wise wave velocity** that we calculate from the spatial gradient of the trigger delay function of a wave $T(x, y)$. Between the continuous optical flow directions and wave-wise gradient directions, we see a tendency that the optical flow estimation is slightly biased from the geometry of the electrode array, showing an artifactual preference for horizontal and vertical directions. The wave-wise velocity measure generally agrees with the channel-wise velocity, showing a somewhat higher variance due to the additional channel variability. The channel-wise measures are equivalent to the measures used in Chapter 7.

In the analysis of the LFP signals, we calculate the **power spectral density** and intra- and inter-area coherence, using the implementation of Welch's approximation method in the Elephant toolkit (Denker et al., 2018a) (using a "Hann" window with an overlap of 0.5). We calculate the **coherence** between two arrays by aggregating the coherence between pairs of individual channel signals from these arrays. 10 times, we sample the random pairs between all the channels, i.e., 320 channel pairs with each channel involved in 10 pairs.

8.3 RESULTS

8.3.1 *Frequency domains of the LFP activity*

The oscillatory LFP activity is commonly segmented into the characteristic frequency bands delta, theta, alpha, beta, and gamma (often splitting beta and gamma further into low and high parts) spanning the range from ~ 0.5 Hz up to ~ 150 Hz. These frequency bands were first defined as a pragmatic segmentation of the frequency range (Steriade et al., 1990). The exact generating mechanisms and their role in neural processing are still not fully understood. However, the different domains are associated with correlates of network state and behavior, and the their exact frequency span can vary with the subject's species and age, as well as the cortical location and depth (Buzsaki, 2006). Therefore, the definition of the LFP frequency bands are not universally fixed. While there are more detailed approaches to identifying frequency band boundaries in electrophysiology data (Shackman et al., 2010; Donoghue et al., 2020; Cohen, 2021), we here suffice with a less complex approach. To cover all potentially relevant frequencies, we define broad frequency ranges. We perform a qualitative assessment of the frequency power spectrum and the intra- and inter-area coherency spectra to separate different oscillatory mechanisms, i.e., peaks in the spectrum, without splitting them between two bands. Based on the spectra, shown in Figure 8.2, we thus define four frequency bands for our analysis: *subalpha* ($1-10$ Hz) combining most of the conventional delta, theta, and alpha domain; *lowbeta* ($10-23$ Hz); *highbeta* ($23-35$), and *gamma* ($60-90$). To avoid the influence of the 50 Hz line noise, and since there is no indication of spectral peaks above 35 Hz, we leave out the "low gamma" regime between our defined highbeta and gamma bands.

In the power spectrum (Figure 8.2A), we observe prominent peaks for the arrays in $7A$ and DP in the lowbeta band. Two distinct spectral peaks are visible in the lowbeta and highbeta bands for the motor array. We include, here, also the array placed in the motor area because, during the performance of the visuomotor task, we expect interactions between the motor cortex and the visual cortex that may involve the dominant frequencies from the motor array. The coherency analysis further depicts peaks in subalpha and lowbeta for the signals within areas $V1$, $7A$, and DP (Figure 8.2B). In between the different arrays, we observe strong coherency effects for the neighboring $7A$-DP areas in the subalpha and lowbeta band. Similarly, there is also a strong coherence between the neighboring $V1$-$V2$ areas in the lowbeta and highbeta bands.

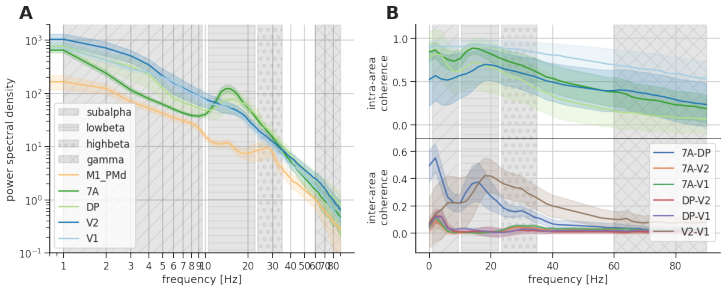

Figure 8.2: **Identifying the frequency domains of the LFP. A)** The spectral power density for multielectrode array (color coded) on a double-logarithmic scaling. The four frequency domains are displayed as gray-hatching patterns. **B)** The intra-area (top) and inter-area (bottom) coherence within and between the arrays. Values in the $45 - 55$ Hz range are removed to suppress the line noise.

8.3.2 *Characterization of traveling waves across areas and frequencies*

Visualizing the phase signal of the LFP activity in the space of the cortical surface for each of the frequency bands creates a fascinating display of wave activities. Waxing and waning patterns of changing directions, reorganizing in planar, sometimes spiral, sometimes more complex arrangements, travel over the electrode arrays. Regularly, the activity synchronizes over the area, and waves seem to travel from one area to the next. Figure 8.3A shows a snapshot of such a wave display over the four electrode arrays in the visual cortex in the lowbeta band. We find a first indication that the observed activity is not a random collection of wave-like patterns but that there are repeating characteristic patterns by applying a wavemode analysis (see Section 8.2.2). For example, looking at the $V2$ array in one of the first successful trials (trial 5), we see that there are two principle modes of wave occurring on that array (Figure 8.3B). Of the 65 waves detected in that trial, 43 waves propagate from the bottom to the top of the array (anterior-lateral to posterior-medial, mode A), 20 propagate from the top-right of the array to the bottom left (posterior to anterior, mode D). We set the number of wavemodes to discern by the clustering algorithm to 4. Therefore it necessarily produced two additional wavemodes. However, these two modes, B and C, only contain one wave, representing the furthest outlier patterns from the main mode clusters A and D.

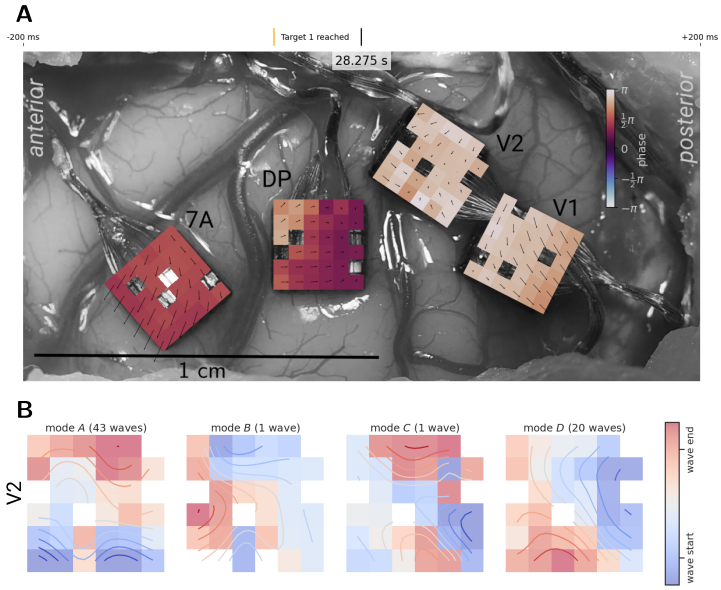

Figure 8.3: **Examplary wave patterns on the cortical surface.** For trial 5 of the visuomotor task, **A)** shows a snapshot shortly after the first target was reached with ongoing wave activity over the different arrays in their actual arrangement on the cortex. On the arrays, the color code displays the phase of the LFP activity filtered to the lowbeta band ($10 - 23$ Hz). The white arrows display the optical flow vector field as approximate wave directions. For the same trial, **B)** shows 4 wavemodes of typical wave pattern in the $V2$ array, with the most waves falling into the modes A and D. The array is oriented with the connector cable to the right side (turned $\sim 34^\circ$ counter-clockwise compared to panel A).

Our goal is to find correlates of the wave activity data that help us relate the activity in the cortical areas to each other and external stimuli and behavior. Due to the complexity of the wave activity, we need to abstract it to its characteristic features (as also motivated in Section 1.3). Thus we first derive quantitative measures to characterize the waves across the areas and frequencies. In Figure 8.4, we look at the planarity, inter-wave intervals, velocity, and direction of the detected waves (see details in Section 8.2.2). Generally, we see that the measure profiles for each area in each frequency band are distinctly different, except for the inter-wave intervals that become relatively similar for all areas towards higher frequencies.

A striking characteristic is that $V1$ exhibits considerably faster and more planar waves. With increasing frequency, the wave velocities are also higher. However, the velocity increase as a function of the frequency differs for the respective areas. $V1$ tends to have relatively

fast waves and DP rather slow waves in all frequency bands. However, $V2$ is on the slow end in subalpha and lowbeta but has a steep velocity increase between highbeta and gamma. In contrast, $7A$ has its steepest increase between subalpha and lowbeta with only slight increases in higher frequencies, having even slightly slower velocities than $V2$ in gamma. Furthermore, the wave planarity is also dependent on the frequency band. While the planarity distributions of $V2$ and DP remain unchanged throughout the frequencies, $V1$ and $7A$ both show the most planar waves in the lowbeta band and the least planar waves in the subalpha and gamma band.

Looking at the wave directions, we see a preferred propagation axis on each array. Wave propagation in DP and $7A$ tends to be oriented on an anterior/medial-to-posterior/lateral axis. Waves in $V1$ tend to be perpendicular to that on an anterior/lateral-to-posterior/medial axis, and $V2$ waves travel mostly on an anterior-to-posterior axis. There can also be a slight preference for one of the two directions on the preferred axes. This asymmetry also appears to change for some areas in some frequency bands. For example, for subalpha $V2$ and lowbeta DP, there is a switch of preferred direction along their axes compared to the other frequency bands. For all areas, the shape of the direction distribution also broadens in higher frequencies, which could be an effect of waves in higher frequency bands being associated with shorter wavelengths resulting in less homogenous and more complex local wave patterns. Correspondingly, there are observations in the literature that slow oscillations show coherence over a greater distance while fast oscillations are contained locally (e.g. von Stein and Sarnthein, 2000; Valderrama et al., 2012).

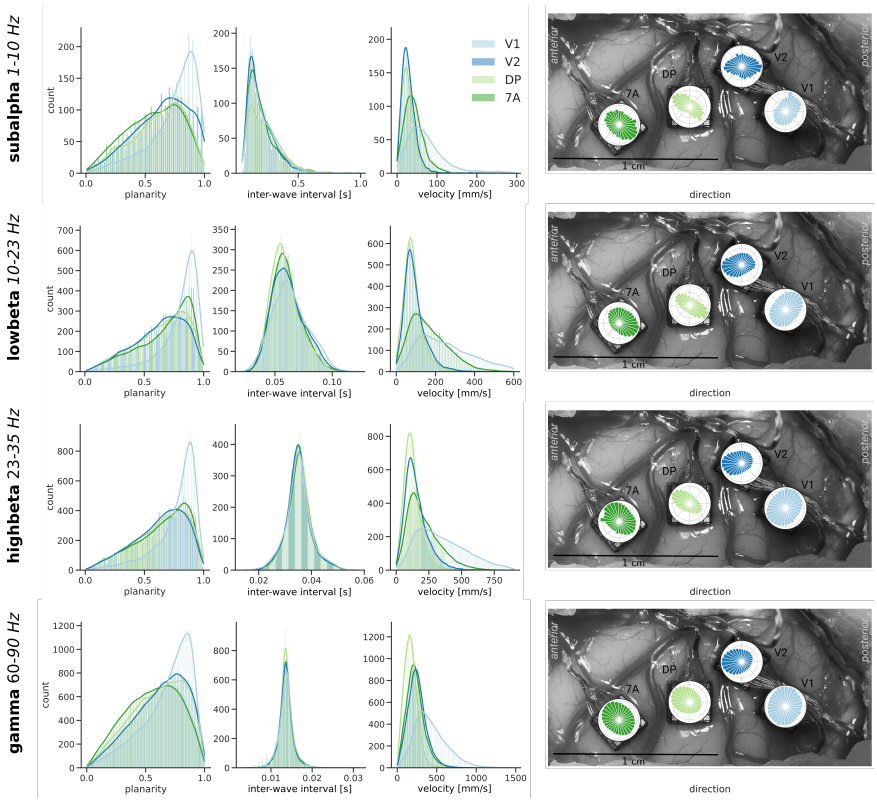

Figure 8.4: **Wave characterization across cortical areas and frequency bands.** The four rows present characterization of the wave activity in the four frequency bands defined in Figure 8.2. The columns present the summary distributions for the wave-wise measures planarity, inter-wave intervals, velocity, and channel-wise direction over all successful trials in the analyzed session, separated for the four areas in the visual cortex $V1$ (bright blue), $V2$ (dark blue), DP (bright green), $7A$ (dark green).

8.3.3 *Relation of wave characteristics to trial events*

There is still much to be analyzed and understood about the characteristics of the wave activity on their own. However, in order to come closer to identifying a potential functional role of the wave activity, we need to correlate the wave characteristic to behavior and external stimuli. To this end, we cut the data into individual trials, select only the successful ones, and align them to specific trial events (see the trial definition in Section 8.2.1) to see if any wave characteristic, on average, changes in response to the trial event. In particular, we

look at the number and channel-wise velocities of detected waves and the directions and the planarity of the continuous optical flow vector field signal (see Section 8.2.2 for details). The direction angles ϕ are circular values and, therefore, not straightforward to average or visualize as a function of time. Thus, we instead plot the horizontal (anterior→posterior) and vertical (lateral→medial) components of the direction as $\cos(\phi)$ and $\sin\phi$. In the trials, we select windows of ± 200 ms around the target-onset or target-reached events, either for one of the 4 targets in the sequence task (averaging over different sequences, i.e., positions) or for one of the 7 target positions (averaging over the order in the sequence).

Looking at the animated wave activity in a single trial, illustrated in Figure 8.5A, by visual inspection, one tends to see eventual relations between the wave dynamics and the occurrence of task events. However, in the trial averaged traces of the characteristic measures, no obvious relation appears for any alignment events, areas, and frequency bands. However, relations between trial events and wave activity become apparent when looking at specific sequences of target positions.

In the subalpha frequency band, the onset of the upper-right target position in the hexagonal layout after a prior onset and reaching of the central target position always triggers a specific wave pattern most robustly seen in $V2$ (Figure 8.5A). After the target onset, the wave direction shifts to an anterior direction ~ 50 ms after the onset event, independent of the wave direction before the onset event. Then the wave activity quickly changes the direction 180° propagating in a posterior direction at ~ 100 ms, which remains at least until 200 ms after onset. Within this time frame of $50 - 200$ ms after target onset, the direction-changing activity is accompanied by a steep increase in the vector fields' planarity in $V1$ and $V2$. Additionally, the number of detected waves in $V2$ increases considerably around ~ 100 ms after onset, and in $V1$ around ~ 200 ms after onset. Furthermore, we can also detect an increase in $V1$ wave velocities after the target onset.

However, this particular wave pattern that reoccurs in every trial with a center-out transition to the upper-right is not specific to the onset of the upper-right target position alone. Therefore, we look at instances where there is a similar relative target transition to the upper-right in the trials. Figure 8.5B and C show such instances for transitions from the left to the upper-left target position and the bottom-right to the right target position. Although the described measure trends are not as clear and with more variability as in Figure 8.5B, the same pattern can be identified in response to the target onset. Further testing the specificity of the pattern, we look at instances where there is a transition upwards to the upper-right target position from the lower-right target position (Figure 8.5D). One can identify aspects of the same pattern, but much less pronounced than in panels A-C, e.g., there

is no increased $V1$ wave velocity in panel D. Lastly, to check whether this pattern is indeed selective for a target transition approximately in the upper-right direction and not just for any transitions to one of the upper-right target positions, we look at instances where there is an upper-left transition from the right target position to the upper-right target position (Figure 8.5E). Indeed, there is no response in the form of the same wave pattern for an upper-left target transition.

8.4 CONCLUSION

SUMMARY In this chapter, we investigate the wave activity of the LFP signal in the cortical areas $V1$, $V2$, DP, and $7A$ along the visual pathway in awake macaque monkeys performing a visuomotor hand-eye coordination task. We analyze characteristic measures of the wave activity across cortical areas, frequency bands, and in relation to behaviorally relevant trial events. Figure 7.2 in Chapter 7 illustrates the issue of multiscale phenomena like wave activity that the assessment with a single measurement technique typically implies a compromise between temporal and spatial resolution, as well as a fine resolution or a larger scope. In this regard, spanning the analysis over multiple multielectrode arrays represents an approach to increase the scope of the observable network without reducing the granularity of the recording (cf. Section 1.2). To analyze the wave activity, we exploit the intentional reusability of the Cobrawap pipeline that we developed in the context of Chapter 7. We built the pipeline's modular processing and analysis blocks for analyzing slow wave activity (< 1 Hz) of propagating Up states recorded with ECoG and wide-field calcium imaging. Here, we demonstrate how the pipeline can be adapted to wave phenomena more generally, particularly for oscillatory LFP wave activity in higher frequencies ($1 - 90$ Hz).

AREA AND FREQUENCY SPECIFIC WAVE ACTIVITY With the detection and characterization of individual waves over the trials of a recording session, we show that the four neighboring visual areas exhibit wave activity with distinctly different (multivariate) property statistics (Section 8.3.2). This result is analogous to the multivariate characterization of the spiking activity across cortical areas presented in Chapter 6. In both cases, we see unique activity fingerprints of cortical areas that a combination of characteristic measures can differentiate (cf. neural characterization theme Section 2.3).

Furthermore, we observe that the wave activity statistics are specific to the cortical areas and vary across the frequency bands from sub-alpha to gamma. A difference between the frequency bands should be expected since they are associated with different oscillatory mechanisms, as introduced in the Section 8.1. However, many oscillatory mechanisms are not yet known or fully understood. Therefore, thor-

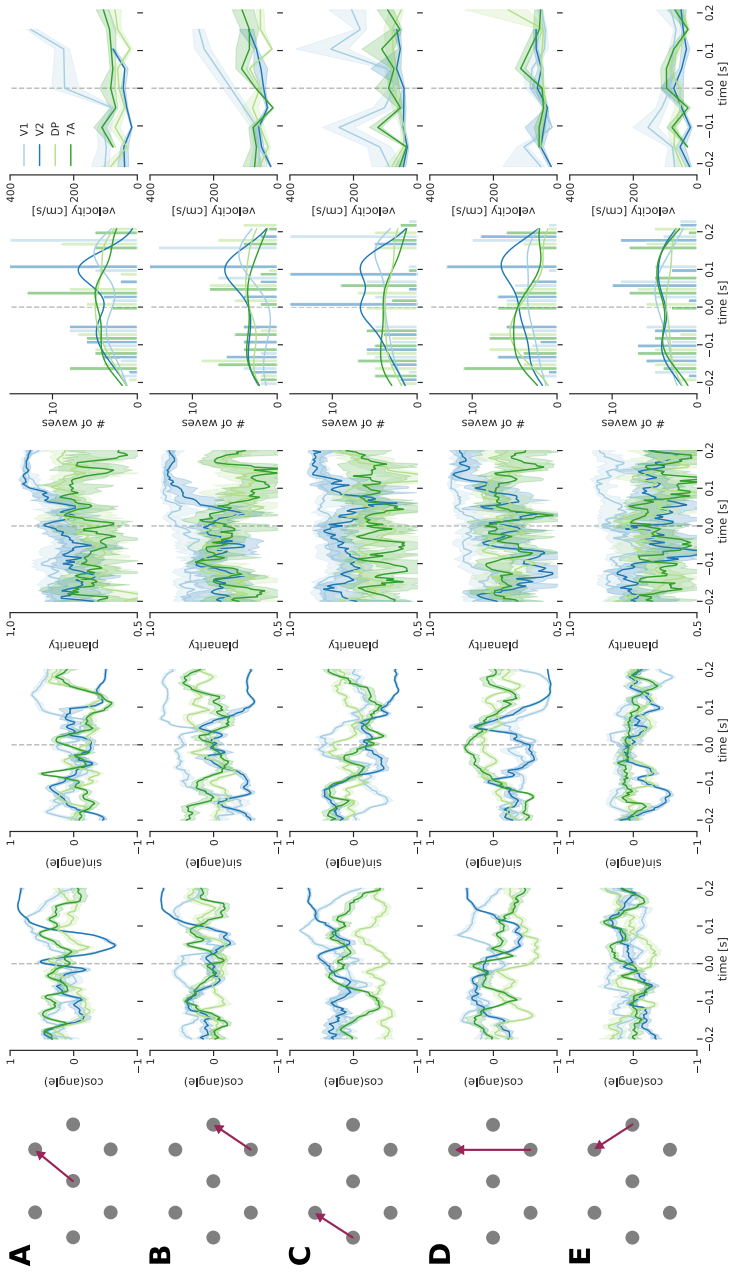

Figure 8.5: **Wave characteristics aligned to target onsets at specific positions.**
Caption on the next page.

Figure 8.5: **Wave characteristics aligned to target onsets at specific positions.** The *first column* shows the target onset events in the hexagonal target layout. The data is aligned to the target onset in the position marked by the arrow head for trials where the target at the position of the arrow base was on immediately prior. Therefore, **A**, **B**, and **C** correspond to a target position transition to the upper-right, **C** to a transition upwards, and **E** to a transition to the upper left. The individual plots in each *row* show the course of subalpha wave activity measures (window 400 ms) aligned to the target onsets and averaged over successful trials. The traces correspond to the median activity measures of the four cortical areas ($V1$-light blue, $V2$-dark blue, DP-light green, $7A$-dark green), aggregated over channels and trials, with the trace-surrounding shading representing the 99% confidence interval. The *columns* show the measures cos(angle) (i.e,anterior→posterior), sin(angle) (i.e., lateral→medial), planarity, number of waves, and velocity. The cosine and sine were calculated from the angle ϕ of the continuous channel-wise optical flow directions ($\phi = 0$ pointing posterior), and so is also the planarity measure. The number of waves and their velocity depends on the detection and separation into individual wavefronts.

ough characterizations of the oscillatory activity features, such as propagating waves, can contribute to building a broader knowledge collection to deduce the underlying mechanisms.

Furthermore, the different frequency bands are also assumed to have different functional roles. For example, in the context of predictive coding (Rao and Ballard, 1999), the feedforward and feedback signals traveling in the opposite direction along hierarchical pathways are thought to travel in different cortical layers and different frequency bands (Bastos et al., 2012; Bastos et al., 2015; Schneider et al., 2020; Vezoli et al., 2021). Accordingly, the feedforward signal should be more dominant in the low (subalpha) frequencies and the high (gamma) frequencies, while the feedback signal should be more dominant in intermediate (beta) frequencies. Suppose cortical waves are involved in such feedforward and feedback signals. In that case, one might, for example, expect waves to propagate along the visual pathway in different directions depending on the frequency band, not unlike the frequency-dependent preferred directions we see in our wave characterization in Figure 8.4.

FURTHER WAVE ACTIVITY COMPARISON APPLICATIONS To test the hypotheses of predictive coding and other theories, the presented characterizations of the wave activity may serve as a basis for statistical experiment vs. experiment comparison applications (Section 2.2). For example, such comparisons can evaluate wave activity differences between scenarios that are assumably either feedforward or feedback dominated: before vs. after movement onset, target onset vs. target

reached, or onset of earlier targets vs. later target (by the last target, the monkey has recognized the sequence). More forthright comparison scenarios could involve comparing perturbed and unperturbed trials. The experimental setup allows visual perturbations of the hand cursor or motor perturbations via the exoskeleton, similar to Tseng et al. (2007), Arce et al. (2010), and Torrecillos et al. (2015). However, it is to be noted that the corresponding analysis of wave directions across frequency bands relies on the electrode arrays' exact placement and orientation on the cortical surface. However, we know that implanted electrodes can drift and be moved out of place by scar tissue. Furthermore, we generally do not know the depth of the implanted electrodes (Mendoza-Halliday et al., 2022). Thus, the analysis cannot correctly incorporate effects specific to the measured cortical layer.

FUNCTIONAL RELEVANCE OF WAVE ACTIVITY In Section 8.3.3, we identify a neural correlate of the form of a characteristic wave activity pattern of at least 200 ms that is precisely triggered in response to a target lighting up to the upper-right relative to the last reached target position. This locking effect is most pronounced for the transition of the central location to the upper-right target position in the hexagonal layout. However, the effect is still clearly observable in relative transitions to the upper-right with other position pairs. However, the effect disappears when the relative orientation to the target position changes, even when involving the same absolute position. The clear and distinct specificity of a complex wave pattern that repeats reliably to the related event is genuinely astonishing.

The specificity to a relative target transition direction suggests that the response may be related to a corresponding eye movement and a focus shift. Indeed, traveling waves triggered by saccadic eye movement have been observed for several decades (Evans, 1953; Thickbroom et al., 1991; Zanos et al., 2015). These reports found that attributes of the triggered wave, such as the latency, amplitude, or direction, correlate with the saccade's duration, direction, and location in the visual field. While earlier works were limited to EEG recordings, Zanos et al. (2015) measured the saccade-triggered waves with an implanted multielectrode array in the visual cortex of a macaque monkey. They, however, report the wave activity in the $V4$ area (located next to $V2$) and in the beta band (defined as $20 - 40$ Hz). The earlier works assigned the triggered wave activity to the change in visual stimulation caused by the eye movement (Billings, 1989). However, Zanos et al. (2015) speculate further that the wave activity may also be related to transient activations of local inhibitory circuits involved in saccadic suppression mechanisms (Kleiser et al., 2004).

Our further research will include a correlation analysis of the trial events and the behavioral signals of eye and hand movement.

Part III

CONCLUSION

9 CONCLUSION

In this thesis, we set out to better understand neural network dynamics. Our approach consists of exploring meaningful characterizations of network activity, constructing network comparisons based on the resulting activity descriptions, and evaluating the similarity between network descriptions concerning their underlying structure and functional purpose. Each step of this three-step approach (i.e., thematic threads Figure 3.1) is represented in the five research projects that encompass this thesis. In the following, we outline how the work in Part ii contributes to each of the three themes and builds on one another.

9.1 CHARACTERIZATION OF NEURAL DATA

COMPLEXITY LEVELS OF SPIKING ACTIVITY CHARACTERISTICS Chapter 4 presents several measures that characterize different aspects of the spiking activity of neural networks. We classify the measures according to the level of interaction between the neurons they represent. Neuron-wise measures quantify isolated attributes of individual neurons, including the average firing rate and the spiking regularity. Pairwise measures quantify interactions between neurons, such as the Pearson correlation coefficients between spike trains. Higher-order measures represent more complex interactions between multiple neurons, for example, the eigenvalues of the correlation matrix or the statistics of spatio-temporal patterns.

UNIVARIATE DESCRIPTION Initially, we focus on univariate comparisons. This means we consider the individual measure values independent so they can be compared via their sample distribution, irrespective of whether the measure is neuron-wise, pairwise, or higher-order. Such a univariate description implicitly assumes that the measure values are identically distributed, which is a reasonable initial approximation. These assumptions allow us to apply two-sample null-hypothesis significance testing (NHST) as a statistical tool to perform a quantitative comparison between two sets of measured network activity characteristics. Despite these constraints (revised further down the thread), NHSTs present a well-established approach with many available and implemented tests (Nickerson, 2000).

TAKE-AWAYS FROM UNIVARIATE COMPARISONS Based on such univariate comparisons of network spiking activity, we evaluate the

similarity of two model implementations. These comparisons highlight two aspects of network characterization. First, no single measure is sufficient to represent the neural network comprehensively. Each measure is just one observable and can not capture all relevant information. Therefore, it is generally necessary to use many distinct measures in combination to construct a meaningful comparison. Second, although the characteristic measures can be classified in their complexity into neuron-wise, pairwise, and higher-order measures, there is no agreement hierarchy. For example, if two neural networks are similar in terms of their pairwise correlation coefficients, this does not necessarily imply that their neuron-wise spiking regularities are also similar. Conversely, an agreement in multiple neuron-wise and pairwise measures between two network realizations does not determine that, for example, the statistics of spatio-temporal patterns are also similar.

PAIRWISE CORRELATION MEASURES In Chapter 5, we focus on pairwise measures and review the assumption of the NHST approach that the measured values are independent. In particular, we look at correlations describing the interactions within a network. The correlations between neurons A and B and neurons A and C condition the correlation between B and C, resulting in a positive definite correlation matrix. This aspect of interdependence is not captured in a univariate distribution of correlation coefficients. Therefore, we introduce a new statistical test that can evaluate the similarity between correlation matrices. Applications to stochastic, simulated, and experimental spiking activity data demonstrate that the matrix representation of pairwise measures and its appropriate statistical comparison can capture specific aspects of the coordinated spiking activity that are not accessible by a univariate approach. For example, this pairwise characterization can distinguish the activity motifs associated with highly-connected neuron assemblies that are often proposed as vital elements in neural network architectures (Harris et al., 2003; Aviel et al., 2003; Litwin-Kumar and Doiron, 2012). Furthermore, we extend the considerations of (symmetric) pairwise activity measures to (asymmetric) pairwise measures of synaptic connectivity.

MULTIVARIATE DESCRIPTION The second assumption of the initial approach was that all measure values are identically distributed. This consideration needs to be revised, especially for network architectures consisting of distinct populations, for example, excitatory and inhibitory neurons across multiple layers in a cortical column model. Therefore, in Chapter 6, we extend our systematic approach by grouping neurons by population into sets of simultaneously recorded spike trains. So, in a comparison between two networks, only corresponding groups are compared, and their respective similarity scores

are combined into a joint score. Additionally, we elaborate on the insight that it takes multiple measures to evaluate a neural network reasonably. Each measure can represent a distinct network characterization. However, a measure may also correlate to other measures, which can itself be a characteristic attribute of the network. For example, it may be that neurons with a high firing rate also tend to spike more regularly. Such correlations between neuron-wise measures can be captured in a joint, multivariate description of the characteristic measures. Furthermore, they can be represented in multidimensional measure space and compared using the Wasserstein distance. In an evaluation of different sets of experimental data, we see that we can not distinguish all datasets with any single univariate description but only with the joint multivariate approach.

WAVES AS POPULATION ACTIVITY MEASURE Next, we focus on activity characteristics at the population level in the form of spatially propagating activity waves. Patterns of traveling waves are ubiquitous in cortical activity across species, frequency bands, and scales. Waves can represent propagating fronts of local transitions between activity states (slow waves) as well as spatial phase relationships of local field oscillations. We investigate their local (channel-wise, i.e., pixel or electrode-wise) and global (wave-wise) characteristics. In Chapter 7, we develop an adaptable analysis pipeline to describe slow wave activity data from heterogeneous sources on a common abstraction level. We extract characteristic measures of the wave activity from this abstract wave description, such as velocity, direction, planarity, and inter-wave intervals. By deriving comparable wave characteristics across data sources, we can relate the wave characteristic to data-specific parameters, such as the spatial resolution of the measurement technique. We reuse the same approach and measures in Chapter 8 for characterizing oscillatory phase waves in the local-field potential (LFP). The wave characterizations present distinct activity signatures for different visual areas and frequency bands. Furthermore, we see a correlation between wave characteristics and behaviorally relevant events.

9.2 MODALITIES OF COMPARISON

ADAPTING THE VALIDATION CONCEPT When comparing two network characterizations, we are interested in how the two networks are related. This comparison modality determines what kind of information we can gain and in what context it is useful. Networks can be described based on experimental recordings of biological tissue, computational model simulations, stochastically generated activity, and their many variations. The classic validation scenario compares model simulations against experimental measurements to establish

confidence in the accuracy of the model predictions. Validation has an established conceptual foundation and is applied in many scientific fields (Section 1.4). In Chapter 4, we adapt the validation approach and terminology to the network-level evaluation of the neural activity. Analog to model vs. experiment applications, we discuss the applicability of model vs. model and experiment vs. experiment comparison modalities. Model vs. model applications can not establish a model's confidence as validation testing does. However, they can, for example, evaluate plausibility, consistency, and sensitivity to input or parameter variations.

SIMULATOR EVALUATION Notably, validation qualitatively separates the mathematical model and the executable model description. For neural network models, this highlights the role that the simulator engine plays in evaluating a model. As a use case, we demonstrate a model vs. model comparison approach to evaluate the neuromorphic SpiNNaker system as a simulator. We develop a neural network model implementation on the SpiNNaker system and iteratively compare its simulation result to a reference simulation. This application showcases the benefits of validation testing for guiding the model development process. Furthermore, we discuss the interpretation of validation test results and emphasize a context-dependent level of agreement.

VALIDATION TEST FRAMEWORK The applications and interpretations of different comparison modalities can be very diverse. Still, the practical network comparison workflow is basically identical. Therefore, we formalized this process in the Python validation test library for network-level comparisons *"NetworkUnit"*. This framework implements validation testing in the form of interacting class instances. Network descriptions class A and B with the capability to simulate spiking activity are matched to a test class that extracts characteristic activity measures. In combination with a score class, the sample measures of A and B are statistically compared, resulting in a similarity score.

COMPARING NETWORK MODEL VARIATIONS In Chapter 5, we extend this framework by developing the statistical eigenangle test to evaluate the similarity of pairwise activity measures such as the Pearson correlation coefficient (and the synaptic weights as pairwise connectivity measure). We apply this test to explore further applications of model vs. model comparison scenarios (and an experiment vs. experiment scenario). In particular, we compare initializations of the same network model with a different random seed to measure the stochastic variability of the network connectivity and activity. Additionally, we apply specific changes to the connectivity of a model instance and compare the model before versus after to quantify the

impact of the applied changes. These applications serve to evaluate the model's sensitivity to variations. Moreover, evaluating such rewiring experiments allows us to gauge the impact of specific connectivity motifs on the activity and better understand the underlying network dynamics.

MODEL CALIBRATION Chapter 6 showcases how multiple comparison modalities can be combined to optimize network model parameters for experimental activity data from different cortical areas. Quantifying the variance of cortical activity across areas by experiment vs. experiment comparisons supports the need for area-specific model definitions. An analogous model vs. experiment workflow can be used to compare a model instance to a specific dataset to evaluate their similarity. To calibrate the model connectivity parameters, we adjust the model parameters accordingly in an iterative optimization loop. Calibration is conceptually different from validation because it explicitly incorporates the information from the experimental data into the model. After the optimization loop converges, the model vs. experiment validation testing thus needs to use a separate set of data (and ideally additional comparison metrics). We demonstrate the efficacy of our model optimization approach to converge to a unique solution of parameter settings by applying it to a test scenario, in which we calibrate a simple network model against simulated data with known ground truth. Chapter 7 shows another calibration and validation use-case, where a neuromodulatory input to a mean-field model is fitted so that the model exhibits accurate statistics of slow wave characteristics.

COMPARING VARIATION OF EXPERIMENTAL DATA Besides, in Chapter 7, we mainly focus on experiment vs. experiment comparisons across heterogenous datasets based on slow wave characteristics. In a meta-study, we aggregate datasets from different sources and compare them grouped by measurement technique, anesthetic type/dosage, and genetic strain. With this approach, we can evaluate the influence of these experimental parameters on the measured wave activity. In Chapter 8, we further compare experimental data of behaving macaque monkeys to evaluate the characteristics of cortical waves across (visual) areas, frequencies, and behaviorally relevant events.

9.3 INTERPLAY OF CONNECTIVITY, ACTIVITY, FUNCTION:

VARIABILITY OF SIMULATED NEURAL NETWORK ACTIVITY The exploration of different comparison metrics and modalities for neural network descriptions raises the question of what we can learn about a neural system from these comparisons. In particular, we are interested in the interplay between network connectivity, activity, and function.

The activity of neural networks is not a deterministic result of the underlying connectivity. The network dynamics are affected by additional influences such as the dynamic state of the network, type of external input, and internal and external sources of stochasticity. In neural network models, such influences are often explicitly included in the model description or the simulation environment. However, details of the model implementation and simulation environment also influence the simulated activity, for example, the type of random number generators, solver of ordinary differential equations, or floating point arithmetic. In Chapter 4, we quantitatively measure the variability of simulated spiking activity of a neural network model between two different computing architectures (classical Von-Neumann vs. digital neuromorphic).

MEASURING CONNECTIVITY → ACTIVITY EFFECTS In Chapter 5, we introduce a new methodology for analyzing the relationship between network connectivity and activity. The statistical test can (for a specific type of network) compare the structure of the synaptic connectivity and the activity correlation between two network realizations. Thus, we can quantitatively link connectivity changes to changes in the activity correlation. In our network model setup, we find that changes to specific synapse populations and motifs have an over-proportional effect on the activity correlations and firing rates while others are under-proportional. However, we also see that the susceptibility to specific synaptic changes depends on network type, state, and parameterization.

INFERRING ACTIVITY → CONNECTIVITY In Chapter 6, we investigate activity-to-connectivity inference by calibrating network model parameters. By matching the characteristics of the simulated spiking activity to experimental recordings, we can accurately determine the most likely parameter values for the model's connection probabilities and external input. The demonstrated effectiveness of this approach suggests that, within the constraints of the model parameterization and recorded network state, similar spiking activity is indicative of similar connectivity. We see distinctly different spiking activity in functionally separated cortical areas in the experimental data, suggesting different microscale connectivity in their network structure. Thus, a following investigation of the inferred structural differences may further indicate their specialized function.

WAVE ACTIVITY CAN INFORM CONNECTIVITY AND FUNCTION Finally, we focus on traveling wave activity, a prevalent neural network activity motif. On the one hand, it has been shown that traveling waves naturally emerge from delays in horizontal connectivity (Davis et al., 2021) and that slow waves constitute a "default mode" activity in the

cortex (Sanchez-Vives et al., 2017). In a model calibration application (Section 7.3.4), we also show that realistic wave activity can emerge from an anisotropic connection probability kernel and a rhythmic modulation of the external input and the spike-frequency adaptation. On the other hand, wave activity has been linked to functions such as memory consolidation (Wei et al., 2016; Capone et al., 2019a; Pazienti et al., 2022), movement preparation (Dean et al., 2012; Heitmann et al., 2015; Denker et al., 2018b), and visual perception (Muller et al., 2014; Zanos et al., 2015; Davis et al., 2020).

In Chapter 7, we integrate heterogeneous data sources into a joint meta-study to build towards a cumulative understanding of slow wave activity. In this meta-study, we evaluate the influence of dataset-specific influences on the characteristics of wave activity. For example, we measure the effects of different anesthesia states on wave dynamics, which may contribute to evaluating theories of unconscious brain states. Further, we compare the wave activity recorded with electrical and optical measurement techniques and isolate the influence of the spatial resolution on the wave measures.

In Chapter 8, we investigate the wave activity in four areas in the visual cortex that we know are connected within the hierarchical visual pathway. The (frequency-dependent) characterization of the wave activity in each area represents the basis for evaluating further how wave activity is mediated between these connected sub-networks. In a preliminary analysis, we find a specific wave pattern signature that exclusively locks to relative visual target transitions towards the upper-right. Based on this observation, we discuss a link between wave activity in the visual cortex and eye movement, which could further narrow down the functional role wave activity could play in visual processes.

10

DISCUSSION & OUTLOOK

10.1 OPEN AND COLLABORATIVE SCIENCE PRACTICES

GENERAL OPEN SCIENCE PRACTICES Practices of open and collaborative science are a prerequisite in modern science. The *FAIR* principles (Findable, Accessible, Interoperable, Reusable) formalize open science practices for research data management (Wilkinson et al., 2016) and research software development (Lamprecht et al., 2020; Barker et al., 2022). In this work, we explicitly address issues of reproducibility and reusability of data and code (discussed, for example, in Section 10.2 and Section 10.3). Moreover, our research applications benefit from open-access data (e.g., in the meta-study in Chapter 7) and code (e.g., in the model replication in Chapter 4). The data and code resources of our publications are published in a form supporting the replicability of results. We emphasize using open-source tools and standard data formats, such as the Python programming language[1], Neo data representation (Garcia et al., 2014), odML metadata representation (Grewe et al., 2011), Nest simulator (Deepu et al., 2021; de Schepper et al., 2022), Elephant analysis toolkit (Denker et al., 2018a), snakemake workflow management system (Mölder et al., 2021), and other Python modules. In our analysis-code development, we aim to contribute to such tools by providing application feedback and co-developing new utility features. We aim to make our methods and implementations generally reusable. Thus, in the scope of this thesis, we develop two open-source software products that we reuse in our projects and that are also being reused in external research projects.

NETWORKUNIT *NetworkUnit* (Gutzen et al., 2018a) is a validation test library to perform statistical network-level comparisons of network activity data. The module is based on the SciUnit validation framework (Omar et al., 2014), to which we, in turn, contribute by integrating new features. Furthermore, we align NetworkUnit with other SciUnit-based test libaries[2], and serving in the SciUnit executive committee[3]. Chapter 4, Chapter 5, and Chapter 6 use and develop NetworkUnit. Beyond our work, NetworkUnit is further applied in projects to evaluate simulators, i.e., comparing the accuracy of a GPU-accelerated Nest against conventional Nest (Tiddia et al., 2022) and comparing a native Nest model implementation against a PyNN-based

1 Python Software Foundation, https://www.python.org/
2 https://github.com/scidash/sciunit/network/dependents?dependent_type=REPOSITORY
3 https://github.com/scidash/sciunit/blob/master/COMMITTEE.md

(Davison et al., 2008) Nest implementation (Albers et al., 2022). Furthermore, NetworkUnit is used in projects evaluating the influence of numerical parameter resolution on the simulated outcome, i.e., measuring the robustness concerning the synaptic weight resolution (Dasbach et al., 2021) and synaptic delay resolution.

COBRAWAP *Cobrawap* (Gutzen et al., 2023) is a framework for constructing brain wave analysis pipelines in a modular manner. The tool is developed and applied in Chapter 7 and Chapter 8, including the related model calibration application described in Section 7.3.4 (Capone et al., 2023). Furthermore, Cobrawap is being applied in an external project using whole-brain calcium imaging studies of anesthetized and sleeping mice in the context of the works Shimaoka et al. (2017) and Zhang et al. (2022). Moreover, current work aims to extend the pipeline stages with additional methods (e.g., hierarchical subsampling (De Luca et al., 2022)) and to integrate the tool with high-performance computing resources and provide it as a workflow service on the EBRAINS platform[4] (Lupo et al., 2022).

10.2 CONTINUOUS VALIDATION TESTING

VALIDATION TESTING IMPROVES MODEL QUALITY AND REPRODUCIBILITY Standardized validation tests help establish an accurate confidence level in a model. Furthermore, in Chapter 4, we show how the repeated application of validation tests during the model development can support and improve the process. Validation tests during model development are distinctly different from calibration, where a bad comparison score triggers the adjustment of model parameters. Instead, validation tests guide the modelers in their model design and, in some cases, may trigger another pass of verification tests. In Chapter 4, we formalize the application of network-level validation tests (incl. calibration and substantiation). The standardized process and the corresponding implementation in the NetworkUnit module enhance reproducibility and ease of use.

VALIDATION TESTING FOR MODEL DEVELOPMENT AND PUBLISHING Elaborating on this idea, the next step would be to further formalize also the continuous application of validation tests, for example, in the context of model development. Instead of manually deciding when and which validation test to apply, predefined tests could be executed automatically with each iteration or adaption of the model. By applying predefined tests, the modeling process is evaluated for prior hypotheses, as opposed to applying suits of validation tests that may be adjusted ad-hoc and can obscure the model's explanatory power. In combination with an automatically triggered test

4 https://wiki.ebrains.eu/

execution, such a process would reduce the possibly arbitrary choices of a modeler and make the modeling process more reliable. Such a continuous validation process can also be applied to other application contexts. For example, performing automated tests on newly recorded data in an experiment by comparing it to sample data can quickly detect deviations from the expected activity regime. Such experimental data validation can indicate possibly faulty equipment and enrich the metadata in the form of a quality measure of a recording session. Furthermore, automated and standardized validation tests could benefit the publishing of models. Combining validation testing with the registration of models in model databases would provide an independent assessment in the form of scores on various standardized tests. The quantitative assessments would make the models comparable and help define their scope of application.

VALIDATION TESTING FOR INTEROPERABILITY Beyond evaluating how accurately a model describes a system of interest, validation tests can further evaluate the robustness of a model's technical realization. Computational results may be influenced by even minor details of the implementation and the computer system (Glatard et al. (2015), Chapter 4). Therefore, an evaluation of a model's confidence should include its robustness concerning different computational environments. This can include, for example, an evaluation of the influence of the precision of value representations (Dasbach et al., 2021)[1] or algorithm choice (e.g., for an ODE solver) (Trensch et al., 2018)[1]. Sensitivity checks can reveal artifactual network model behavior as demonstrated by Pauli et al. (2018)[1]. Furthermore, the model validation with respect to simulator engines can be combined with benchmarking tests of the simulation performance as done for the comparisons Nest vs. SpiNNaker (van Albada et al., 2018), Nest vs. Nest GPU (Golosio et al., 2021b; Tiddia et al., 2022), and Nest vs. PyNN Nest (Albers et al., 2022). The Nest vs. SpiNNaker benchmarking was performed parallel to the project described in Chapter 4, and the latter two examples are currently being adapted and expanded[1] with validation testing using NetworkUnit.

TECHNICAL REALIZATION OF A VALIDATION FRAMEWORK The process of continuous validation could be realized similarly to the practice of *continuous integration* in software development, which automatically executes unit tests that check the correct behavior of the code base whenever an update is pushed to the code repository. The aspect of storing validation test results and linking them to models and datasets is currently being developed by Appukuttan et al. (2022)

1 The author was involved in an advisory/supportive role in designing these validation workflows.

in the context of the EBRAINS platform[5]. This effort also includes the development of domain-specific validation test libraries based on the generic validation test module *SciUnit* (Omar et al., 2014). The NetworkUnit module for tests in the network-level domain that we develop in the context of Chapter 4, Chapter 5, and Chapter 6 contributes to this work.

10.3 REUSABILITY OF THE COLLABORATIVE BRAIN WAVE ANALYSIS PIPELINE

TYPES OF PIPELINE REUSE Not only physical but also digital resources should be recycled. Therefore, we explicitly designed the Collaborative Brain Wave Analysis Pipeline (Cobrawap) to be adaptable and reusable. Therefore, the continuous development and application of the pipeline beyond its initial application in Chapter 7 was expected. There are different types of possible pipeline reuse. In terms of its interfaces to larger workflow applications, the pipeline can be applied to other input types and/or for producing another type of output. In terms of the pipeline's modularity, the reuse types can be classified as adaption (e.g., new data, changed configuration, changed block selection), selection (e.g., extracting individual or a subset of blocks or stages), or extension (adding additional blocks and/or stages).

BRANCHING-OFF PIPELINE EXTENSIONS For example, current work entails the development of a pipeline realization that focuses on the local oscillations of multi-unit activity signals between Up and Down states and ignores the spatial propagation, similar to the work by De Bonis et al. (2019). The Slow Oscillation Analysis Pipeline (SOAP) reuses the first three stages of the Cobrawap and then branches off with specialized stages for more detailed local analysis. Another envisioned pipeline application will extend the analysis of spontaneous wave activity to wave activity triggered by external electrical stimulation. This extension requires storing information about the electrical stimulation, such as its timing, location, intensity, and pulse shape, in the standardized data representation of the pipeline (using Neo (Garcia et al., 2014)). This additional stimulus information is then available to new specialized blocks, for example, to calculate the stimulus-response complexity by a perturbation-complexity index (PCI) (Casali et al., 2013; Virmani and Nagaraj, 2019). Perturbational complexity is a clinical measure for evaluating levels of consciousness realized in multiple method definitions and implementations (Casali et al., 2013; Virmani and Nagaraj, 2019; Deco et al., 2018). Besides enabling PCI analyses on heterogeneous data, the corresponding pipeline integration could also be used to compare the different PCI methods, i.e., "method benchmarking" as demonstrated in Section 7.3.3.

5 https://wiki.ebrains.eu/bin/view/Collabs/model-validation

INPUT TYPE PIPELINE ADAPTIONS While the initial application of Cobrawap was tailored towards spatially coordinated propagating Up and Down states in the < 1 Hz regime (Chapter 7), we demonstrated in Chapter 8 how the pipeline can be adapted to oscillatory wave activity in higher frequency regimes. This adaption involves reinterpreting the detection of Down-to-Up transitions in slow waves as the detection of an equal-phase wavefront (e.g., defined via its peak $\phi = 0$) and calibration of the configuration parameters to the corresponding frequency regimes and spatial resolution. In this context, we integrated additional input datasets of oscillatory wave activity measured with implanted multielectrode arrays (Brochier et al., 2018; Chen et al., 2022). While traveling oscillatory waves are observed in many cortical regimes, frequency bands, and experimental contexts (e.g. Senseman and Robbins, 2002; Petersen et al., 2003; Wu et al., 2008; Townsend and Gong, 2018; Denker et al., 2018b; Balasubramanian et al., 2019), there are some discrepancies in the reported results and terminology. For example, both Townsend and Gong (2018) and Denker et al. (2018b) classify and investigate different wave types (including planar, spiral, and radial patterns) but apply different definitions. Integrating the corresponding datasets and methods into the pipeline could disentangle the different definitions and make the results comparable.

CONVERGENCE LEVEL OF HETEROGENOUS DATA When integrating heterogeneous data into a common analysis, the crucial question arises concerning the adequate comparison level at which all data abstractions should converge. For example, when we compare wave activity based on spiking and LFP data, we could either approximate an LFP signal from the spikes (e.g., via LFPy Hagen et al., 2018) and then process both data equivalently or derive wave measures directly from both signals. On the one hand, a direct calculation of wave measures from each data modality provides a more immediate wave description that may adjust for specific properties of the data modality. On the other hand, abstracting the more detailed data modality (spikes) to the more coarse data modality (LFP) allows a more immediate convergence of the data descriptions and equivalent calculations of wave measures. However, both approaches also have considerable drawbacks in terms of increasing complexity when adding new data modalities or wave measures. Transforming one data modality into another can get out of hand quickly when dealing with many different data sources and is not always easily possible. Similarly, developing methods to extract wave measures separately for each data modality increases the amount of methodological/coding effort and potential sources of errors and discrepancies many-fold. Therefore, in most cases, we would define the convergence level (i.e., abstract wave description) and derive wave measures independently of any data source. In Chapter 7, we motivate this approach with the

analogy of a rail system. When there are many starting points and many possible destinations, it is most effective to construct central traffic nodes instead of many direct connections. This option always involves an abstraction of the input data and may, in some cases, also omit to leverage the total informative value of the data. For example, the temporal precision and eventual phase-locking effects would be disregarded when abstracting spiking data via rate fluctuations to the abstract wave description defined in Cobrawap.

10.4 INFERRING CONNECTIVITY AND PREDICTING ACTIVITY

CONNECTIVITY-ACTIVITY RELATION A key aspect in network neuroscience is and remains the dynamic relationship between the structural connectivity of networks and the activity they exhibit. Due to their multifaceted relationship, it is not trivial to conclude when we can infer structure from the activity and when we can infer activity to structure. As the same computer hardware can run different kinds of software, a network with fixed connectivity can give rise to different types of activity. However, this analogy's extent is vague as the brain is not like a computer in many aspects (Brette, 2022). While connectivity shapes activity, the activity also shapes connectivity via multiple plasticity mechanisms on different timescales from milliseconds to hours (Gerstner et al., 2018).

FUNCTIONAL CONNECTIVITY Notably, the structural connectivity is more challenging to measure than the activity of a network, especially for in-vivo recordings. Therefore, investigating structure-activity relationships often involves the analysis of *functional connectivity*. Functional connectivity describes the interaction strength between network nodes only based on their activity correlation. Inferring causation (structural connectivity) from correlation (functional connectivity) is highly problematic, mainly due to correlation contributions via indirection connections and common inputs. While under ideal stationary conditions, functional connectivity can approximate structural connectivity (Bullmore and Sporns, 2009; Eickhoff et al., 2010), it has also been shown that spurious correlations induce systematic errors in the inference of structural connectivity (Das and Fiete, 2020). Nevertheless, there is continuous progress in inferring connectivity from activity, for example, by taking into account the sparseness level of the network (Pernice and Rotter, 2013), using the point-processes models of measured spike trains to guide the inference (Gerhard et al., 2013), or combined structural and functional analysis (Battiston et al., 2017). We know that the connectivity $\rightarrow$ activity relation is a one-to-many mapping that is influenced by additional factors such as the network input and its dynamical state (Sporns and Tononi, 2001). For the other way around, activity $\rightarrow$ connectivity, the ambiguity of functional

connectivity concerning the underlying connectivity also suggest a one-to-many (or at least one-to-few) mapping. However, evaluations of functional connectivity (e.g., in the above-cited works) further suggest that additional assumptions can reduce ambiguity and establish a one-to-one mapping from activity to connectivity.

ACTIVITY → CONNECTIVITY The consideration of a one-to-one mapping from activity to connectivity is used implicitly when calibrating connectivity parameters via comparisons of activity features. In order to make the unique inference of the connectivity tractable in a calibration application, the corresponding model needs to impose additional constraints for the connectivity, as we also do in Chapter 6. Therefore, we do not attempt to calibrate every aspect of the model connectivity but only specific aspects, such as the connectivity probabilities between neuron populations and the external input. Other connectivity aspects are fixed, such as the number of excitatory and inhibitory neurons, the total number of synapses, and their strength distribution. In particular, also any plasticity effects are ignored. Furthermore, these model constraints must also correspond to the experimental target system. Therefore, we focused in Chapter 6 on simple resting-state recordings that provide a certain degree of stationarity. Any remaining ambiguity in the connectivity inference would be indicated by repeated calibration steps not converging to a unique solution or by a subsequent validation procedure. We find in a proof-of-concept application with modeled data that our calibration workflow based on evolutionary optimization converges properly to the ground truth connectivity parameters based on a multivariate activity characterization. Furthermore, we find that our characterization approach for spiking activity can distinguish recordings from different cortical areas, layers, and subjects, with the similarity of characterizations generally corresponding to the assumed functional similarity of the recorded sets of neurons (e.g., between the same area and layer in different subjects). These first results, thus, provide good conditions to approximate connectivity parameters of cortical microcircuit models from spiking activity data, which we pursue in the next step. Similarly, in Chapter 7, we demonstrate a calibration application that infers connectivity and external input parameters for a mean-field network model via slow wave activity measured with wide-field calcium imaging. The successful modeling of realistic slow wave activity shows that many of the observed wave activity features can be accounted for by the interplay of spatially structured connectivity and fluctuating input (Capone and Mattia, 2017; Capone et al., 2023). This insight seems to also generalize for wave activity in higher frequency regimes (Kang et al., 2023).

CONNECTIVITY → ACTIVITY The activity of a network is constrained by its underlying connectivity. However, due to non-linear network dynamics and additional influences, the activity can typically not be determined precisely from a given connectivity (Albada et al., 2015). For example, in Chapter 4, we demonstrate how small numerical details in a model realization can substantially affect the exhibited spiking activity. However, multiple works draw relations between activity attributes (e.g., correlations) and specific connectivity aspects, such as recurrence and hub structures (Pernice et al., 2011), or background activity and synaptic properties (Ostojic et al., 2009). In Chapter 5, we also demonstrate a quantitative approach that can measure the influence of changes to a network model's connectivity on the network's activity correlations. The "eigenangle test" relies on the informative value of the eigenspectrum of a matrix-valued network measure. As, for example, demonstrated in Zhou et al. (2009) and Dahmen et al. (2019), the eigenspectrum of the connectivity matrix can inform the network state and dynamics. In our application, we use synaptic rewiring protocols to investigate the influences of the different connection changes between excitatory and inhibitory populations and the spread of additional excitatory connections. With this approach, we find that the correlation structure is more sensitive to excitatory connections, as previously suggested by Aertsen and Gerstein (1985), while firing rates are more sensitive to inhibitory connections, as also found by Mongillo et al. (2018). Our approach can be adapted to a range of other rewiring protocols, for example, to probe the effects of the arrangement, composition, and interaction of cell assemblies (Harris, 2005; Aviel et al., 2003; Litwin-Kumar and Doiron, 2012). Despite the complex and non-linear behavior of neural network dynamics (Dahmen et al., 2022; Wainrib and Touboul, 2013; García del Molino et al., 2013; Sompolinsky et al., 1988), we show that there are strategic approaches to evaluate its aspects for a given neural network type. Besides the type of network model that we used (Rajan and Abbott, 2006), the eigenangle test can, in principle, be applied to other network types with analytical eigenspectrum descriptions (e.g. Sommers et al., 1988; Kuczala and Sharpee, 2016; Ahmadian et al., 2015; Schuessler et al., 2020; Muir and Mrsic-Flogel, 2015), or numerically defined null hypothesis distributions via surrogate data (Grün, 2009; Stella et al., 2022).

10.5 RELATION BETWEEN SPIKE PATTERNS AND LFP WAVES

SEPARATION INTO SPIKES AND LFP Neural activity, measured by implanted electrodes, can be abstracted and separated into two types of activity: spikes, the binary pulse signals representing action potentials from nearby neurons; and the local field potential (LFP), the low-pass filtered analog signal component representing a population

activity from nearby neuron groups (Section 1.3). The link from intracellular action potentials to extra-cellular spikes is relatively clear (through the filter of spike sorting (Rey et al., 2015)). In contrast, there are many factors contributing to the LFP signal that may depend on the species, brain area, spectral band, and network state. However, the main contributors are the local postsynaptic potentials (Pesaran et al., 2018).

Evidently, there is an overlap in the respective underlying mechanisms, as spikes and LFP both represent aspects of electric neural activity. However, the unique contributions to the signals make them both relevant complementary components to understanding neural network dynamics. For example, in contrast to spikes, the LFP signal reflects subthreshold membrane dynamics and their correlation across neurons (Ness et al., 2016; Speed et al., 2020). Spikes, in contrast to LFP, can be attributed to individual neurons and can therefore resolve activations of specific combinations of neurons (Grün et al., 2002; Litwin-Kumar and Doiron, 2012). Both, spikes and LFP signal components are extracted from the same recorded signal and mainly distinguished by their frequency regime: LFP $\lessapprox$ 300 Hz, multi-unit activity (MUA) containing spikes $\gtrapprox$ 300 Hz. Chapter 4, Chapter 5, and Chapter 6 focus on the analysis of the spiking activity while Chapter 7 and Chapter 8 focus on wave dynamics in low-pass filtered (LFP) signals.

INTERACTION OF SPIKES AND LFP Even though the spikes and LFP are separated by frequency domain, the signals are not independent. In fact, there are remnants of the spiking activity in the higher frequencies of the LFP signal ($\gtrapprox$ 50 Hz), measured on the same electrode, that become visible in correlation analyses or spike-triggered averages (Ray and Maunsell, 2011; Zanos et al., 2011). Besides such artifactual relations, there are also meaningful relations between the signal components. Spikes not only influence the LFP at the same electrode but can create an outward propagating wave of LFP activity, that can be measured in nearby electrodes (Nauhaus et al., 2009). The other way around, the LFP can influence the membrane potential of a neuron and thus change its probability to emit an action potential (i.e., "ephaptic coupling") (Anastassiou et al., 2011). Additional to the exchange of action potentials via synapses, this type of interaction represents another communication mechanism between neurons. It is widely supported that rhythmic LFP activity plays a role in coordinating distant spiking activity and thus aligning different functional processes and neural circuits (Varela et al. (2001), Section 10.6).

One signature of this coordinating mechanism between LFP and spikes is that sometimes the timing of spikes is locked to the phase of the LFP or a frequency band of the LFP ("phase locking"). The degree of phase locking depends on the LFP oscillation strength (Denker

et al., 2007), surplus synchrony of the spikes (Denker et al., 2011), and can correlate with stimulus features (Engel et al., 1990; Fries et al., 2001; Vinck et al., 2010). These findings suggest a functional relevance of LPF-spike interactions, however, the exact mechanisms are not yet understood. Another indicator to an eventual functional mechanism is that patterns of coordinated spikes ("spatio-temporal patterns") tend to have a preferred spatial direction (Takahashi et al., 2015; Torre et al., 2016), just as wave activity shows a preferred directionality (Rubino et al. (2006) and Denker et al. (2018b), Chapter 7, Chapter 8). An ongoing research project under the author's supervision investigates the alignment of simultaneously co-occurring spike patterns and waves.

10.6 FUNCTIONAL ROLES OF CORTICAL WAVES

DIRECTION PREFERENCES OF TRAVELING CORTICAL WAVES Rubino et al. (2006) were the first to report the organization of local LFP oscillatory activity into propagating waves in the motor cortex. Among their results, they report the propagation of beta-band waves along preferred axes. However, the preferred axes are distinct for different functional areas in the motor cortex. In the primary motor cortex, waves travel on an anterior-posterior axis, whereas in the dorsal premotor area, they travel on a lateral-medial axis. Similarly, we observe preferred wave propagation axis in the visual cortex that differ in orientation (Section 8.3.2). Although there are also complex wave patterns, like spiral or radial patterns, that do not have a single defined direction, most observed waves tend to have an approximate planar directionality (Denker et al., 2018b; Townsend and Gong, 2018). The emergence of traveling cortical waves can be explained by the delays of horizontal fiber connections (Davis et al., 2021), and their direction preference by anisotropic connection probability kernels (Section 7.3.4, Capone et al. (2023)).

In Section 7.3.2, we further observe preferred direction axes on a hemisphere scale for slow wave activity. Notably, while the waves recorded with either ECoG or calcium imaging had preferred directions roughly on an anterior-posterior axis, the tilt of the axes in medial-lateral directions differed considerably between the two measurement techniques. A possible explanation for this discrepancy could be that the two techniques measure from different subsets of neurons in different cortical layers (ECoG: layer 1; calcium imaging: layer 2/3 & 5) (de Vries et al., 2020; Siegle et al., 2021b; Siegle et al., 2021a). Additionally, there is further evidence that cortical wave dynamics differ across the laminar layer structure of the cortex (Halgren et al., 2018; Capone et al., 2019b; Bharioke et al., 2022).

Furthermore, wave direction and direction variability seem to be a matter of spatial scope and frequency. Low-frequency oscillation patterns are coordinated over a longer distance (e.g., a hemisphere) than

the typically more localized high-frequency oscillations (von Stein and Sarnthein, 2000; Valderrama et al., 2012). This trend also corresponds to our observation of the direction variance as a function of frequency in recordings of implanted multielectrode arrays (Section 8.3.2). So, the directionality of cortical wave dynamics depends on the underlying connectivity, the specific area and layer, and frequency. Considering the popular view that synchronization of oscillations plays a role in the coordination of activity across local or distant neural population (Destexhe et al., 1999b; von Stein and Sarnthein, 2000; Halgren et al., 2018), wave direction is plausibly also of functional relevance for the directed flow of information (e.g., in feedback and feedforward processes) (Buffalo et al., 2011; van Kerkoerle et al., 2014; Michalareas et al., 2016).

COMMUNICATION WITHIN FREQUENCY BANDS When looking at the visualized LFP phase wave patterns across different frequency bands on multielectrode arrays, as in Chapter A, one can visually distinguish distinct patterns that occur on the same array at the same time, but in different frequency bands. In the context of information flow and inter/intra-areal communication, one might draw an analogy to frequency-separated communication channels as they are implemented, for example, in radio signals.

Indeed, similar concepts are proposed for cortical function and are widely studied. For example, synchronization of neural populations to different frequency bands can control different modes of attention (Fries et al., 2001; Buschman and Miller, 2007; Lakatos et al., 2008; Schroeder and Lakatos, 2009; Landau et al., 2015; Helfrich et al., 2018; Fiebelkorn et al., 2018); rhythmic synchronization modulates effective connectivity so that post-synaptic populations select inputs that are coherent to its dominant frequency ("communication-through-coherence") (Fries, 2005; Fries, 2015); feedback and feedforward signals (e.g., in a predictive coding schema) are separated into distinct frequency bands (Bastos et al., 2012; Bastos et al., 2015; Michalareas et al., 2016; Dann et al., 2016; Schneider et al., 2020; Vezoli et al., 2021; Vinck et al., 2022). Notably, these concepts are not mutually exclusive and share some common aspects. However, none of these studies explicitly relate the concept of frequency channels with traveling wave patterns.

Expanding these approaches by further considering the spatial organization of neural oscillations would be promising, as Davis et al. (2020) showed that the presence of waves can be more predictive of perceptual sensitivity than just the oscillation phase alone. Conversely, a detailed frequency consideration may be helpful to distinguish the feedforward and feedback attributes of stimulus-triggered and internally-generated (i.e., "spontaneous") waves. Attributes of stimulus-triggered waves in the visual cortex may depend, for example, on a stimulus's contrast (Nauhaus et al., 2009), or its position,

orientation, and frequency (Benucci et al., 2007). In contrast, there are also internally-generated waves that are instead indicative of attention (Davis et al., 2020) and general top-down communication (Roland et al., 2006). Furthermore, observed cross-frequency coupling effects (Lakatos et al., 2008; Schroeder and Lakatos, 2009; Helfrich et al., 2018; Greenwood and Ward, 2022) represent possible interaction modes between different frequency-specific mechanisms (Halgren et al., 2018).

In Section 8.3.2, we observe preliminary evidence that wave direction preferences change across frequencies. However, in order to substantiate this finding and narrow down its eventual relations to mechanisms suggested in the literature (e.g., feedback/feedforward separation), we need to further analyze the contextual details and reconcile previous evidence from the time, frequency, and spatial domains of oscillatory communication schemes and traveling wave dynamics.

COMPUTATIONAL MECHANISMS It has been suggested for a long time that additionally to the synaptic transmission of action potentials, oscillations of the local field potential (LFP) play a crucial role in the computational processes of neuro-electrical signals (Varela, 1995; Hopfield, 1995; Wilson et al., 2015). In Section 10.5, we discussed the unique contributions of spike and LFP signals. To narrow down the eventual computational mechanisms of neural circuits, one has to consider the various evidence for functionally relevant activity features (Section 1.5) and the complex biophysical spike-LFP and LFP-LFP interactions. Likely, there is not just a single way in which information is encoded and processed. Field potential oscillations (and their wave organizations) may be involved in various mechanisms in different brain regions, frequency bands, and spatial and temporal scales.

The LFP as a population activity measure is not only a helpful workaround for the grave undersampling limitations in recordings of simultaneous activity from single neurons. Even when neurons are in the direct vicinity of a recording electrode, many emit spikes extremely sparsely (Shoham et al., 2006) and are therefore not captured in the recorded spiking activity. Still, their subthreshold activity contributes to the LFP. The LFP, in turn, influences the membrane potential of neurons and, therefore, their spiking probability and timing (Fröhlich and McCormick, 2010; Anastassiou et al., 2011; Su et al., 2012; Han et al., 2018). The local LFP oscillations rhythmically hyperpolarize and depolarize the membrane potential of nearby neurons. During hyperpolarization, neurons are mostly unresponsive, while during depolarization, neurons are more sensitive to input, and spikes are shifted towards the peak of the LFP oscillation phase (Reyes and Fetz, 1993; Mellon and Wheeler, 1999).

As LFP oscillations modulate neuron sensitivity, the formation of cortical wave patterns means that within the extent of a wave, only a part of the cortical area is deactivated (hyperpolarized), while another part is excitable (depolarized). Therefore, wave activity in a cortical sensory field ensures that always a fraction of neurons is maximally sensitive (Ermentrout and Kleinfeld, 2001). This hypothesis agrees with reports of correlations between wave activity in the visual cortex and performance in a visual perception task (Davis et al., 2020). Sensory cortical areas are typically spatially organized in some map of the corresponding sensory field. So, different stimuli (e.g., different orientations of a visual stimulus) activate spatially separate parts in the sensory field. In the case of two such different stimuli occurring simultaneously, a coinciding wave could separate the two stimuli by their phase. This phase information could be used to categorize or selectively attend different inputs (Sompolinsky and Tsodyks, 1994; Ermentrout and Kleinfeld, 2001). Accordingly, there are observations that the spike phase of a neuron is selectively tuned to the orientation of a visual stimulus (Vinck et al., 2010). Furthermore, observations show that functional areas typically contain less than one wavelength of an LFP wave, i.e., a phase variation $< 2\pi$ (Roelfsema et al., 1997; Ribary et al., 1991; Prechtl et al., 1997; Lam et al., 2000; Rubino et al., 2006), so that each site in the area can be associated with a unique phase value. This notion of a relation between wavelength and area dimensions is further supported by observations of waves being reflected and compressed at the border between visual areas (Xu et al., 2007).

Besides within-area wave dynamics, the spatial spread of LFP rhythms is also postulated to support synchronization and modulation across spatially separated and functionally specialized areas (e.g. Varela et al., 2001; Shine et al., 2021). Although, such long-range interactions may be restricted to lower frequencies $\lessapprox$ 30 Hz as the range of gamma-modulated computation is limited (Ray and Maunsell, 2010). Propagating wave patterns in lower frequency regimes may reflect the communication between coordinated processes such as loops of action and perception (Jutras et al., 2013; Noel et al., 2022). In this regard, the rhythmic LFP oscillations are thought to segment the sensory input steam and integrate the information into a hierarchical structure of perception, for example, in active visual exploration (Jensen et al., 2021), or the understanding of language (Kazanina and Tavano, 2022).

WAVES IN VISUAL PROCESSING Notably, many of the observed LFP waves are recorded in cortical areas associated with visual processing and/or correlated to visual input. Even the observation of beta-band waves in the motor cortex by Rubino et al. (2006) were shown to not relate to movement planning or execution but instead hypothesized to

correspond to visual target input. In Chapter 8, we observe a distinct wave pattern that occurs very selectively only for a relative transition of a visual target to the upper right. In the 1-10 Hz (subalpha) regime, within 50-150 ms after target-onset to the upper right, we see an increase in detected waves and wave planarity in $V1$ and $V2$, a wave velocity increase in $V1$, and a wave propagation pattern in $V2$ that travels first in an anterior direction and then abruptly reverses into a posterior direction. The timing of this wave response to target onsets suggests a relation to a corresponding eye movement.

There have been many reports of changes in the oscillatory activity along the visual pathway that are related to eye movement (e.g. Lee and Malpeli, 1998; Reppas et al., 2002; Purpura et al., 2003; Sylvester et al., 2005; Rajkai et al., 2008; Barczak et al., 2019). Further, EEG recordings have shown that saccade-evoked potentials are organized in waves ("λ" waves) whose properties (e.g., latency, shape, amplitude) correlate with the saccade properties (e.g., duration, direction, position) (e.g. Evans, 1953; Billings, 1989; Thickbroom et al., 1991; Skrandies and Laschke, 1997). More recently, similar effects of saccade-triggered waves were also observed with implanted electrode arrays (Zanos et al., 2015). However, these various findings do not yet paint a coherent picture to which our observations could adhere. Consequently, we lack a comprehensive explanation of saccade-related potential oscillations, their wave properties, and their eventual functional role.

Initial theories focused on the fact that eye movements shift the visual scene over the retina (Thickbroom et al., 1991). However, the presence of saccade-triggered responses even in complete darkness point to additional mechanisms (Skrandies and Laschke, 1997). Ensuing hypotheses instead emphasize that the fixation at the end of a saccade is accompanied by increased excitability across the visual pathway (Lee and Malpeli, 1998; Nakamura and Colby, 2000; Reppas et al., 2002; Rajkai et al., 2008). So, it stands to reason that there is not only a visually evoked response but also a depolarizing preparation/expectation activity independent of retinal input and caused by a shifted gaze (Ito et al., 2011; Barczak et al., 2019). Such a preparation mechanism was postulated to integrate new observations with ongoing activity in higher-order regions (Sobotka and Ringo, 1997) by synchronizing activation onset across areas (Purpura et al., 2003; Jutras et al., 2013), mediated by phase-resets of neural oscillations (Lakatos et al., 2008; Rajkai et al., 2008; Barczak et al., 2019). Furthermore, multiple findings suggest that an inhibitory input suppression during the saccade precedes the excitability after fixation (Kleiser et al., 2004; Sylvester et al., 2005; Churan et al., 2012; Zanos et al., 2015; Barczak et al., 2019). In summary, there is accumulating evidence that active visual sensing (Schroeder et al., 2010; Barczak et al., 2019) is organized by oscillatory processes that modulate sensory-evoked responses (Davis et al., 2020) to enable the integration of visual inputs (Jensen et al., 2021).

These theories and results mostly align with our experimental observations (Section 8.3.3). Many of the cited studies report the corresponding oscillatory processes in low-frequency regimes (delta - alpha) (Purpura et al., 2003; Barczak et al., 2019; Jensen et al., 2021) (notably except (Zanos et al., 2015)), which have been shown to occur as traveling waves in the human neocortex (Zhang et al., 2018). Furthermore, there are also indications of directional tuning mechanisms in LFP saccade responses (Sobotka and Ringo, 1997; Purpura et al., 2003; Zanos et al., 2015). In the next steps of our analysis, we aim to incorporate behavioral hand and eye-movement data to better identify the correlates of the LFP wave dynamics. In this regard, we hope to integrate further aspects of the spatial organization of oscillations (i.e., waves) into the existing theories of visual processing and shed some light on the functional role(s) of traveling wave activity.

10.7 PERSONAL PERSPECTIVE ON THE FUTURE OF (COMPUTATIONAL) NEUROSCIENCE

There are still many fundamental questions in neuroscience that await to be answered. For example, in analogy to the Millennium Prize Problems in mathematics, van Hemmen and Sejnowski (2005) formulated "23 Problems in Systems Neuroscience". The issue with such fundamental questions is that they are unlikely to be answered by any single research project.

Looking back on the history of scientific breakthroughs, (Kuhn, 1962) describes a modest ascent of the general knowledge level intercepted with major jumps of scientific revolutions. Although sometimes attributed to a single research work, such revolutions or breakthroughs are, in most cases, an aggregation of many incremental research works that eventually culminate in a paradigm shift. However, it appears that the neuroscience research landscape is becoming increasingly complex, making it more difficult to aggregate many individual findings to improve our fundamental knowledge effectively. Fundamentally understanding neural systems and their host organisms involves many disciplines addressing research questions on many scales, from molecular to social and environmental interactions. This argument that research progress in the future is facing different obstacles than in the past may be biased by a hindsight view of the past and an uncertain view of the future. Nevertheless, there are real obstacles to scientific progress in our time, and although we are better positioned to advance scientific progress than ever in history, our advancements also bring their own challenges:

a) There are multiple known unknowns in our understanding of neural systems, for example: How do neurons encode information? What are the neural correlates of consciousness? Can we simulate an entire human brain? How do biological neural networks compute?

However, many of the questions still seem less tractable in the sense that it is already difficult to define the question. For example: "What do we mean by neural code?" (Brette, 2019), "Is there a hard problem of consciousness?" (Chalmers, 1995; O'Hara and Scutt, 1996), "Which whole brain simulations are actually useful?" (Destexhe, 2021). "What do we mean by biological computation?" (Wood, 2019). **b)** Furthermore, there have never been so many research resources and so many people working to advance science (Naujokaitytė, 2021). Still, in this huge assortment of scientific activity, we struggle to get an overview of the relevant works (Bornmann and Mutz, 2015), and therefore also to get research heard by the right people and allocate resources where they could be used most efficiently (Heinze, 2008). **c)** The technological resources available to science are incredible: Supercomputers (Alvarez, 2021), 7-Tesla+ MRI (Platt et al., 2021), Neuropixel probes (Steinmetz et al., 2018), and instant communication to any other researcher. But with great power comes great humility as many disciplines face an ongoing reproducibility crisis (Miłkowski et al., 2018). **d)** With advanced measurement technology and databases, there has never been so much detailed data on neural systems. Still, this only represents a tiny glimpse into the entire neural structure and activity ensemble. Instead, we face the issues of organizing, describing, representing, storing, sharing, processing, and analyzing the data at scale (Wilkinson et al., 2016).

Within this current state of neuroscience research, we argue that fundamental progress requires systematic approaches to aggregate the many heterogeneous advances into a collaborative effort to build a cumulative understanding of neural systems. The work presented in this thesis aims to contribute to this effort by developing systematic approaches for the quantitative comparability of neural network descriptions. In an ideal future, we would like to see the seamless reusability and combinability of research resources (including software, methods, results, data, and metadata) facilitated by open standards and interfaces. Furthermore, to tackle fundamental challenges in neuroscience, we hope to see a more systematic decoupling of technological, methodological, biological, and philosophical demands. For example: is the quest to determine the "neural code" limited by our measurement techniques? Is it limited because of biological stochasticity and variability? Or is the task limited by an insufficient definition of the question?

Part IV

APPENDIX

A

VISUALIZING CORTICAL WAVES

Neural oscillations have rich dynamics which are challenging to capture in a coherent picture. Complex interactions between signal components across space, time, and frequency exist. Such interactions give rise to various propagating wave patterns that are observable in the amplitude and phase of the neural activity. Figure A.1 is the result of an exploration of how to design data visualizations of cortical wave dynamics. The blog article, rgutzen.github.io/2020-06-25-visualizing_waves, provides a detailed description of the process of constructing the visualization. The visualization is also featured in the "SciPy John Hunter Excellence in Plotting Competition" and was awarded second place in 2020[1].

We use an open-access dataset of neural activity is analyzed that was recorded by a 10x10 electrode array in the motor cortex of a macaque monkey during a hand movement task (Brochier et al. (2018), Section 5.3.2). In this task, the monkey is instructed to reach for an object and to grasp it after receiving two visual cues specifying the grip type (precision grip or side grip) and the grip force (high force or low force).

The visualization shows four different frequency regimes, their spectral power in the frequency domain, their wave-like activity across the electrode array, and the signal evolution in time. The organization of the activity phase in space reveals spiral, circular, planar, and synchronous wave patterns. Distinct neural activity patterns appear simultaneously in the different frequency regimes. The power spectrum shows their varying contributions to the full signal. Further, the signal time course shows how the components align, especially around behaviorally relevant events. As a reaction to the 'Grip-Type Cue', the beta and alpha components increase in spectral power and show a more synchronized activity across the electrode array. However, they are reduced in favor of theta and delta contributions during the hand movement. Generally, there seems to be a correlation between the spectral power and the kind of wave pattern. This kind of visualized relations between the aspects of the neural signals and their correlations to behavioral events may serve as a basis to guide further investigations of the functional role of brain waves (for example as in Chapter 8).

1 https://jhepc.github.io/2020/entry_24

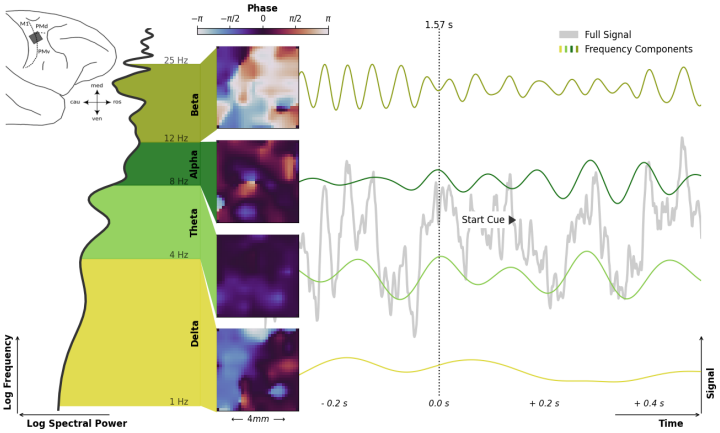

Figure A.1: Neural wave dynamics in the motor cortex of a macaque monkey during an instructed reach-to-grasp task. An animated version of the figure can be accessed at youtube.com/watch?v=keXdZirxxjQ. The figure presents the local field potential activity in four color-coded frequency domains (Delta, Theta, Alpha, Beta) with varying spectral power (*left*). The activity is recorded with an implanted 10x10 electrode array (*top-left inset*). The spatial organization of the oscillation phases over the area of the electrode array displays complex wave dynamics (*center*). The signal components evolve in time (*right*) and show a varying degree of coherence that seems to be related to trial events (labeled marker) of the task.

B

THE KINTSUGI BRAIN

Art and science can interact a symbiotic relationship. Scientific topics are often a sources for compelling art pieces; and art can as a creative approach to address scientific questions in thought provoking ways. There are many efforts that promote "SciArt", and "BrainArt" (or "NeuroArt") in particular. One of them is the OHBM BrainArt exhibition[1], that is organized as a part of the annual OHBM (*Organization for Human Brain Mapping*) conference. In this context, I presented the piece shown in Figure B.1.

> *This artwork is inspired by the brain's plasticity. The brain can rewire and repair its broken connections. Kintsugi is the Japanese art of repairing broken bowls or other pottery with gold, highlighting its cracks, and celebrating its imperfections as part of its history. Initially, plasticity allows for molding the clay and for molding the brain in its early development by the external influences imposed on it. Once in its apparent final shape, it is still subject to change and external forces. Bowls can break. Brains can break. Plasticity does not end. In many cases, they can still be repaired and become more valuable in the process. The philosophy of kintsugi, thus, raises the question: how much do we gain from our scars? And in which way is the brain entirely unlike a bowl?*

1 https://ohbm-brainart.github.io

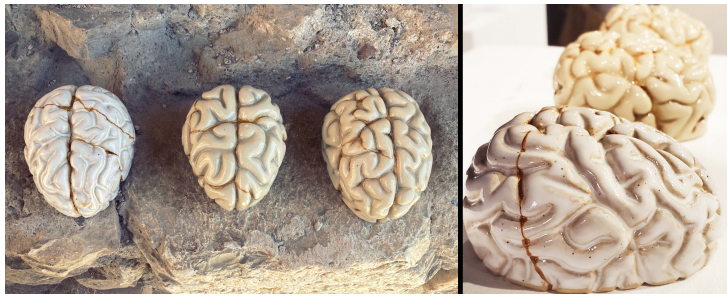

Figure B.1: **"The kintsugi brain"** is an artwork by Robin Gutzen, that consists of multiple hand-modeled clay figures of human brains, finished in different glazes. The brain figures were purposely broken and repaired with gold glue in the style of the Japanese art-form "kintsugi". (*The photo on the right was taken at the OHBM BrainArt exhibition 2022 in Glasgow by visitor Dr. Teodora Stoica*)

ACM (2016). *Artifact Review and Badging – Version 1.0 (Not Current)*. https://www.acm.org/publications/policies/artifact-review-badging. (Visited on 2022) (cit. on pp. 38, 39).

AIAA (1998). "Guide: Guide for the Verification and Validation of Computational Fluid Dynamics Simulations (AIAA G-077-1998(2002))." In: *Guide: Guide for the Verification and Validation of Computational Fluid Dynamics Simulations (AIAA G-077-1998(2002))*. AIAA Standards. Computational Fluid Dynamics Committee, American Institute of Aeronautics and Astronautics, Inc. DOI: 10.2514/4.472855.001. (Visited on 2022) (cit. on p. 14).

Abeles, M. (1982). "Role of the Cortical Neuron: Integrator or Coincidence Detector." In: *Israel journal of medical sciences* 18.1, pp. 83–92 (cit. on p. 5).

Abeles, Moshe (1991). *Corticonics* (cit. on p. 19).

Acebrón, Juan A., L. L. Bonilla, Conrad J. Pérez Vicente, Félix Ritort, and Renato Spigler (2005). "The Kuramoto Model: A Simple Paradigm for Synchronization Phenomena." In: *Reviews of Modern Physics* 77.1, pp. 137–185. DOI: 10.1103/RevModPhys.77.137. (Visited on 2022) (cit. on p. 12).

Achermann, P and A. A Borbély (1997). "Low-Frequency (<1Hz) Oscillations in the Human Sleep Electroencephalogram." In: *Neuroscience* 81.1, pp. 213–222. DOI: 10.1016/S0306-4522(97)00186-3. (Visited on 2022) (cit. on pp. 20, 112).

Aertsen, Ad M.H.J. and George L. Gerstein (1985). "Evaluation of Neuronal Connectivity: Sensitivity of Cross-Correlation." In: *Brain Research* 340.2, pp. 341–354. DOI: 10.1016/0006-8993(85)90931-X (cit. on p. 180).

Ahmadian, Yashar, Francesco Fumarola, and Kenneth D. Miller (2015). "Properties of Networks with Partially Structured and Partially Random Connectivity." In: *Physical Review E* 91.1, p. 012820. DOI: 10.1103/PhysRevE.91.012820. (Visited on 2019) (cit. on p. 180).

Albada, Sacha J van, Jari Pronold, Alexander van Meegen, and Markus Diesmann (2019). "Usage and Scaling of an Open-Source Spiking Multi-Area Model of Monkey Cortex." In: *International Workshop on Brain-Inspired Computing*. Springer, Cham, pp. 47–59 (cit. on p. 8).

Albada, Sacha Jennifer van, Moritz Helias, and Markus Diesmann (2015). "Scalability of Asynchronous Networks Is Limited by One-to-One Mapping between Effective Connectivity and Correlations." In: *PLOS Computational Biology* 11.9, e1004490. DOI: 10.1371/journal.pcbi.1004490. (Visited on 2022) (cit. on p. 180).

Albers, Jasper, Jari Pronold, Anno Christopher Kurth, Stine Brekke Vennemo, Kaveh Haghighi Mood, Alexander Patronis, Dennis Terhorst, Jakob Jordan, Susanne Kunkel, Tom Tetzlaff, Markus Diesmann, and Johanna Senk (2022). "A Modular Workflow for Performance Benchmarking of Neuronal Network Simulations." In: *Frontiers in Neuroinformatics* 16, p. 837549. DOI: 10.3389/fninf.2022.837549. (Visited on 2022) (cit. on pp. 174, 175).

Albert, Réka and Albert-László Barabási (2002). "Statistical Mechanics of Complex Networks." In: *Reviews of Modern Physics* 74.1, p. 47. DOI: 10.1103/RevModPhys.74.47 (cit. on p. 6).

Aljadeff, Johnatan, Merav Stern, and Tatyana Sharpee (2015). "Transition to Chaos in Random Networks with Cell-Type-Specific Connectivity." In: *Physical Review Letters* 114.8, p. 88101. DOI: 10.1103/PhysRevLett.114.088101 (cit. on p. 61).

Alvarez, Damian (2021). "JUWELS Cluster and Booster: Exascale Pathfinder with Modular Supercomputing Architecture at Juelich Supercomputing Centre." In: *Journal of large-scale research facilities JLSRF* 7, A183–A183. DOI: 10.17815/jlsrf-7-183. (Visited on 2022) (cit. on p. 188).

Amunts, Katrin and Karl Zilles (2015). "Architectonic Mapping of the Human Brain beyond Brodmann." In: *Neuron* 88.6, pp. 1086–1107. DOI: 10.1016/j.neuron.2015.12.001. (Visited on 2022) (cit. on p. 5).

Amzica, F. and M. Steriade (1998). "Electrophysiological Correlates of Sleep Delta waves1Correspondence and Reprint Requests May Be Addressed to Either Author.1." In: *Electroencephalography and Clinical Neurophysiology* 107.2, pp. 69–83. DOI: 10.1016/S0013-4694(98)00051-0. (Visited on 2022) (cit. on p. 20).

Anastassiou, Costas A., Rodrigo Perin, Henry Markram, and Christof Koch (2011). "Ephaptic Coupling of Cortical Neurons." In: *Nature Neuroscience* 14.2, pp. 217–223. DOI: 10.1038/nn.2727. (Visited on 2019) (cit. on pp. 181, 184).

Anderson, Charles T, Patrick L Sheets, Taro Kiritani, and Gordon M G Shepherd (2010). "Sublayer-Specific Microcircuits of Corticospinal and Corticostriatal Neurons in Motor Cortex." In: *Nature Neuroscience* 13.6, pp. 739–744. DOI: 10.1038/nn.2538. (Visited on 2022) (cit. on p. 89).

Appukuttan, Shailesh, Lungsi Sharma, Pedro Garcia-Rodriguez, and Andrew Davison (2022). "A Software Framework for Validating Neuroscience Models." (Visited on 2022) (cit. on p. 175).

Arce, Fritzie, Itai Novick, Yael Mandelblat-Cerf, Zvi Israel, Claude Ghez, and Eilon Vaadia (2010). "Combined Adaptiveness of Specific Motor Cortical Ensembles Underlies Learning." In: *The Journal of neuroscience* 30.15, pp. 5415–25. DOI: 10.1523/JNEUROSCI.0076-10.2010. PMID: 20392963 (cit. on p. 161).

Arieli, A, A Sterkin, A Grinvald, and A Aertsen (1996). "Dynamics of Ongoing Activity: Explanation of the Large Variability in Evoked Cortical Responses." In: *Science* 273.5283, pp. 1868–71. DOI: 10.1126/SCIENCE.273.5283.1868. PMID: 8791593 (cit. on p. 59).

Arjovsky, Martin, Soumith Chintala, and Léon Bottou (2017). "Wasserstein GAN." In: DOI: 10.48550/ARXIV.1701.07875 (cit. on pp. 97, 141).

Aur, Dorian and Mandar S. Jog (2010). *Neuroelectrodynamics: Understanding the Brain Language*. IOS Press (cit. on p. 5).

Aviel, Y., C. Mehring, M. Abeles, and D. Horn (2003). "On Embedding Synfire Chains in a Balanced Network." In: *Neural Computation* 15.6, pp. 1321–1340. DOI: 10.1162/089976603321780290 (cit. on pp. 85, 166, 180).

Bakker, Rembrandt, Thomas Wachtler, and Markus Diesmann (2012). "CoCoMac 2.0 and the Future of Tract-Tracing Databases." In: *Frontiers in Neuroinformatics* 6. DOI: 10.3389/fninf.2012.00030. (Visited on 2022) (cit. on p. 88).

Balasubramanian, Karthikeyan, Vasileios Papadourakis, Wei Liang, Kazutaka Takahashi, Matthew D Best, Aaron J Suminski, and Nicholas G Hatsopoulos (2019). "Propagating Patterns of Activity across Motor Cortex Facilitate Movement Initiation." In: *bioRxiv*, p. 549568. DOI: 10.1101/549568. (Visited on 2019) (cit. on pp. 21, 112, 177).

Balci, Osman and Robert G. Sargent (1982). "Validation of Multivariate Response Models Using Hotelling's Two-Sample T2 Test." In: *Simulation* 39.6, pp. 185–192. DOI: 10.1177/003754978203900602 (cit. on p. 12).

Barabási, Albert-László, Réka Albert, and Hawoong Jeong (1999). "Mean-Field Theory for Scale-Free Random Networks." In: *Physica A: Statistical Mechanics and its Applications* 272.1, pp. 173–187. DOI: 10.1016/S0378-4371(99)00291-5. (Visited on 2022) (cit. on p. 15).

Barba, Lorena A. (2018). *Terminologies for Reproducible Research*. DOI: 10.48550/arXiv.1802.03311. arXiv: 1802.03311 [cs]. (Visited on 2022) (cit. on p. 38).

Barczak, Annamaria, Saskia Haegens, Deborah A. Ross, Tammy McGinnis, Peter Lakatos, and Charles E. Schroeder (2019). "Dynamic Modulation of Cortical Excitability during Visual Active Sensing." In: *Cell Reports* 27.12, 3447–3459.e3. DOI: 10.1016/j.celrep.2019.05.072. (Visited on 2022) (cit. on pp. 186, 187).

Barker, Michelle, Neil P. Chue Hong, Daniel S. Katz, Anna-Lena Lamprecht, Carlos Martinez-Ortiz, Fotis Psomopoulos, Jennifer Harrow, Leyla Jael Castro, Morane Gruenpeter, Paula Andrea Martinez, and Tom Honeyman (2022). "Introducing the FAIR Principles for Research Software." In: *Scientific Data* 9.1, p. 622. DOI: 10.1038/s41597-022-01710-x. (Visited on 2022) (cit. on p. 173).

Bartels, Andreas and Semir Zeki (2005). "Brain Dynamics during Natural Viewing Conditions—A New Guide for Mapping Connectivity in

Vivo." In: *NeuroImage* 24.2, pp. 339–349. DOI: 10.1016/j.neuroimage.2004.08.044. (Visited on 2022) (cit. on p. 10).

Barth, Alison L., Richard C. Gerkin, and Kathleen L. Dean (2004). "Alteration of Neuronal Firing Properties after In Vivo Experience in a FosGFP Transgenic Mouse." In: *Journal of Neuroscience* 24.29, pp. 6466–6475. DOI: 10.1523/JNEUROSCI.4737-03.2004. (Visited on 2022) (cit. on p. 3).

Bartolozzi, C., R. Benosman, K. Boahen, G. Cauwenberghs, Tobi Delbrück, Giacomo Indiveri, Shih-Chii Liu, S. Furber, N. Imam, Bernabé Linares-Barranco, Teresa Serrano-Gotarredona, K. Meier, C. Posch, and M. Valle (2016). "Neuromorphic Systems." In: *Wiley Encyclopedia of Electrical and Electronics Engineering*. John Wiley & Sons, Ltd, pp. 1–22. DOI: 10.1002/047134608X.W8328. (Visited on 2022) (cit. on p. 3).

Bassett, Danielle S. and Olaf Sporns (2017). "Network Neuroscience." In: *Nature Neuroscience* 20.3, pp. 353–364. DOI: 10.1038/nn.4502. (Visited on 2022) (cit. on p. 6).

Bastos, Andre M., W. Martin Usrey, Rick A. Adams, George R. Mangun, Pascal Fries, and Karl J. Friston (2012). "Canonical Microcircuits for Predictive Coding." In: *Neuron* 76.4, pp. 695–711. DOI: 10.1016/J.NEURON.2012.10.038 (cit. on pp. 160, 183).

Bastos, André Moraes, Julien Vezoli, Conrado Arturo Bosman, Jan-Mathijs Schoffelen, Robert Oostenveld, Jarrod Robert Dowdall, Peter De Weerd, Henry Kennedy, and Pascal Fries (2015). "Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels." In: *Neuron* 85.2, pp. 390–401. DOI: 10.1016/j.neuron.2014.12.018. (Visited on 2022) (cit. on pp. 160, 183).

Battiston, Federico, Vincenzo Nicosia, Mario Chavez, and Vito Latora (2017). "Multilayer Motif Analysis of Brain Networks." In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 27.4, p. 047404. DOI: 10.1063/1.4979282. (Visited on 2022) (cit. on p. 178).

Bazhenov, Maxim, Igor Timofeev, Mircea Steriade, and Terrence J. Sejnowski (2002). "Model of Thalamocortical Slow-Wave Sleep Oscillations and Transitions to Activated States." In: *Journal of Neuroscience* 22.19, pp. 8691–8704. DOI: 10.1523/JNEUROSCI.22-19-08691.2002. (Visited on 2020) (cit. on pp. 114, 146).

Beaulieu, C. and M. Colonnier (1989). "Number and Size of Neurons and Synapses in the Motor Cortex of Cats Raised in Different Environmental Complexities." In: *The Journal of Comparative Neurology* 289.1, pp. 178–181. DOI: 10.1002/cne.902890115 (cit. on p. 6).

Bell, Jonathan, Thomas D. LaToza, Foteini Baldmitsi, and Angelos Stavrou (2017). "Advancing Open Science with Version Control and Blockchains." In: *2017 IEEE/ACM 12th International Workshop on Software Engineering for Science (SE4Science)*, pp. 13–14. DOI: 10.1109/SE4Science.2017.11 (cit. on p. 117).

Ben-Shaul, Yoram, Hagai Bergman, Ya'acov Ritov, and Moshe Abeles (2001). "Trial to Trial Variability in Either Stimulus or Action Causes Apparent Correlation and Synchrony in Neuronal Activity." In: *Journal of Neuroscience Methods* 111.2, pp. 99–110. DOI: 10.1016/S0165-0270(01)00389-2. (Visited on 2022) (cit. on p. 5).

Benucci, Andrea, Robert A. Frazor, and Matteo Carandini (2007). "Standing Waves and Traveling Waves Distinguish Two Circuits in Visual Cortex." In: *Neuron* 55.1, pp. 103–117. DOI: 10.1016/j.neuron.2007.06.017. (Visited on 2022) (cit. on p. 184).

Berger, Denise, David Warren, Richard Normann, Amos Arieli, and Sonja Grün (2007). "Spatially Organized Spike Correlation in Cat Visual Cortex." In: *Neurocomputing* 70.10-12, pp. 2112–2116. DOI: 10.1016/j.neucom.2006.10.141 (cit. on p. 89).

Berger, Hans (1929). "Über Das Elektroenkephalogramm Des Menschen." In: *Archiv für psychiatrie und nervenkrankheiten* 87.1, pp. 527–570 (cit. on pp. 17, 18).

Bharioke, Arjun, Martin Munz, Alexandra Brignall, Georg Kosche, Max Ferdinand Eizinger, Nicole Ledergerber, Daniel Hillier, Brigitte Gross-Scherf, Karl-Klaus Conzelmann, Emilie Macé, and Botond Roska (2022). "General Anesthesia Globally Synchronizes Activity Selectively in Layer 5 Cortical Pyramidal Neurons." In: *Neuron* 110.12, 2024–2040.e10. DOI: 10.1016/j.neuron.2022.03.032. (Visited on 2022) (cit. on pp. 113, 182).

Billings, R. J. (1989). "The Origin of the Occipital Lambda Wave in Man." In: *Electroencephalography and Clinical Neurophysiology* 72.2, pp. 95–113. DOI: 10.1016/0013-4694(89)90170-3. (Visited on 2022) (cit. on pp. 161, 186).

Binzegger, T. (2004). "A Quantitative Map of the Circuit of Cat Primary Visual Cortex." In: *Journal of Neuroscience* 24.39, pp. 8441–8453. DOI: 10.1523/JNEUROSCI.1400-04.2004. (Visited on 2022) (cit. on pp. 89, 95).

Bittner, Sean R., Agostina Palmigiano, Alex T. Piet, Chunyu A. Duan, Carlos D. Brody, Kenneth D. Miller, and John P. Cunningham (2019). "Interrogating Theoretical Models of Neural Computation with Emergent Property Inference." In: *bioRxiv : the preprint server for biology*. DOI: 10.1101/837567. (Visited on 2022) (cit. on p. 89).

Boldog, Eszter et al. (2018). "Transcriptomic and Morphophysiological Evidence for a Specialized Human Cortical GABAergic Cell Type." In: *Nature Neuroscience* 21.9, pp. 1185–1195. DOI: 10.1038/s41593-018-0205-2 (cit. on p. 4).

Borgonovo, Emanuele and Elmar Plischke (2016). "Sensitivity Analysis: A Review of Recent Advances." In: *European Journal of Operational Research* 248.3, pp. 869–887. DOI: 10.1016/j.ejor.2015.06.032. (Visited on 2022) (cit. on p. 16).

Bornmann, Lutz and Rüdiger Mutz (2015). "Growth Rates of Modern Science: A Bibliometric Analysis Based on the Number of Publica-

tions and Cited References." In: *Journal of the Association for Information Science and Technology* 66.11, pp. 2215–2222. DOI: 10.1002/asi.23329. (Visited on 2022) (cit. on p. 188).

Bos, Hannah, Markus Diesmann, and Moritz Helias (2016). "Identifying Anatomical Origins of Coexisting Oscillations in the Cortical Microcircuit." In: *PLOS Computational Biology* 12.10. Ed. by Peter E. Latham, e1005132. DOI: 10.1371/journal.pcbi.1005132. (Visited on 2022) (cit. on pp. 16, 55).

Botella-Soler, Vicente, Mario Valderrama, Benoît Crépon, Vincent Navarro, and Michel Le Van Quyen (2012). "Large-Scale Cortical Dynamics of Sleep Slow Waves." In: *PLoS ONE* 7.2. DOI: 10.1371/journal.pone.0030757. (Visited on 2020) (cit. on pp. 113, 114).

Botvinik-Nezer, Rotem et al. (2020). "Variability in the Analysis of a Single Neuroimaging Dataset by Many Teams." In: *Nature* 582.7810, pp. 84–88. DOI: 10.1038/s41586-020-2314-9. (Visited on 2022) (cit. on p. 23).

Bourque, Pierre and Richard E Fairley (2014). *SWEBOK: Guide to the Software Engineering Body of Knowledge*. IEEE Computer Society (cit. on p. 35).

Box, G. E. P. (1949). "A General Distribution Theory for a Class of Likelihood Criteria." In: *Biometrika* 36.3-4, pp. 317–346. DOI: 10.1093/biomet/36.3-4.317. (Visited on 2022) (cit. on p. 63).

Bradski, G. (2000). "The OpenCV Library." In: *Dr. Dobb's Journal of Software Tools* (cit. on p. 98).

Brette, Romain (2015). "Philosophy of the Spike: Rate-Based vs. Spike-Based Theories of the Brain." In: *Frontiers in Systems Neuroscience* 9, p. 151. DOI: 10.3389/fnsys.2015.00151. (Visited on 2019) (cit. on p. 19).

– (2019). "Is Coding a Relevant Metaphor for the Brain?" In: *Behavioral and Brain Sciences* 42, pp. 1–44. DOI: 10.1017/S0140525X19000049. (Visited on 2019) (cit. on pp. 17, 188).

– (2022). "Brains as Computers: Metaphor, Analogy, Theory or Fact?" In: *Frontiers in Ecology and Evolution* 10, p. 878729 (cit. on pp. 17, 178).

Brochier, Thomas, Lyuba Zehl, Yaoyao Hao, Margaux Duret, Julia Sprenger, Michael Denker, Sonja Grün, and Alexa Riehle (2018). "Massively Parallel Recordings in Macaque Motor Cortex during an Instructed Delayed Reach-to-Grasp Task." In: *Scientific Data* 5.1, pp. 1–23. DOI: 10.1038/sdata.2018.55. (Visited on 2020) (cit. on pp. 64, 78, 94, 177, 191).

Brodmann, Korbinian (1909). *Vergleichende Lokalisationslehre Der Grosshirnrinde in Ihren Prinzipien Dargestellt Auf Grund Des Zellenbaues*. Leibzig: Barth (cit. on pp. 5, 18).

Brunel, N and N Brunel (2000). "Dynamics of Sparsely Connected Networls of Excitatory and Inhibitory Neurons." In: *Computational*

Neuroscience 8, pp. 183–208. DOI: 10.1016/S0928-4257(00)01084-6 (cit. on p. 94).

Buffalo, Elizabeth A., Pascal Fries, Rogier Landman, Timothy J. Buschman, and Robert Desimone (2011). "Laminar Differences in Gamma and Alpha Coherence in the Ventral Stream." In: *Proceedings of the National Academy of Sciences* 108.27, pp. 11262–11267. DOI: 10.1073/pnas.1011284108. (Visited on 2022) (cit. on p. 183).

Bullmore, Ed and Olaf Sporns (2009). "Complex Brain Networks: Graph Theoretical Analysis of Structural and Functional Systems." In: *Nature Reviews Neuroscience* 10.3, pp. 186–198. DOI: 10.1038/nrn2575. (Visited on 2022) (cit. on pp. 62, 89, 178).

Burkitt, Guy R, Richard B Silberstein, Peter J Cadusch, and Andrew W Wood (2000). "Steady-State Visual Evoked Potentials and Travelling Waves." In: *Clinical Neurophysiology* 111.2, pp. 246–258. DOI: 10.1016/S1388-2457(99)00194-7. (Visited on 2022) (cit. on p. 5).

Buschman, Timothy J. and Earl K. Miller (2007). "Top-Down Versus Bottom-Up Control of Attention in the Prefrontal and Posterior Parietal Cortices." In: *Science* 315.5820, pp. 1860–1862. DOI: 10.1126/science.1138071. (Visited on 2022) (cit. on pp. 147, 183).

Butt, Anila Sahar and Peter Fitch (2020). "ProvONE+: A Provenance Model for Scientific Workflows." In: *Web Information Systems Engineering – WISE 2020*. Ed. by Zhisheng Huang, Wouter Beek, Hua Wang, Rui Zhou, and Yanchun Zhang. Lecture Notes in Computer Science. Cham: Springer International Publishing, pp. 431–444. DOI: 10.1007/978-3-030-62008-0_30 (cit. on p. 117).

Butts, Daniel A., Chong Weng, Jianzhong Jin, Chun-I. Yeh, Nicholas A. Lesica, Jose-Manuel Alonso, and Garrett B. Stanley (2007). "Temporal Precision in the Neural Code and the Timescales of Natural Vision." In: *Nature* 449.7158, pp. 92–95. DOI: 10.1038/nature06105. (Visited on 2022) (cit. on p. 19).

Buzsaki, Gyorgy (2006). *Rhythms of the Brain*. Oxford University Press (cit. on pp. 5, 152).

Buzsáki, György (2019). *The Brain from inside Out*. Oxford University Press (cit. on p. 10).

Cai, T Tony, Jianqing Fan, and Tiefeng Jiang (2013). "Distributions of Angles in Random Packing on Spheres." In: *Journal of Machine Learning Research* 14.1, pp. 1837–1864 (cit. on p. 66).

Calsbeek, Brittny and Charles J. Goodnight (2009). "Empirical Comparison of G Matrix Test Statistics: Finding Biologically Relevant Change." In: *Evolution* 63.10, pp. 2627–2635. DOI: 10.1111/j.1558-5646.2009.00735.x. (Visited on 2022) (cit. on p. 63).

Camassa, Alessandra, Maurizio Mattia, and Maria V. Sanchez-Vives (2021). "Energy-Based Hierarchical Clustering of Cortical Slow Waves in Multi-Electrode Recordings." In: *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society*

(EMBC), pp. 198–203. DOI: 10.1109/EMBC46164.2021.9630931 (cit. on p. 20).

Capone, Cristiano, Chiara De Luca, Giulia De Bonis, Robin Gutzen, Irene Bernava, Elena Pastorelli, Francesco Simula, Cosimo Lupo, Leonardo Tonielli, Francesco Resta, Anna Letizia Allegra Mascaro, Francesco Pavone, Michael Denker, and Pier Stanislao Paolucci (2023). "Simulations Approaching Data: Cortical Slow Waves in Inferred Models of the Whole Hemisphere of Mouse." In: *Communications Biology* 6.1, pp. 1–14. DOI: 10.1038/s42003-023-04580-0 (cit. on pp. 28, 111, 115, 139, 142, 144, 174, 179, 182).

Capone, Cristiano, Carla Filosa, Guido Gigante, Federico Ricci-Tersenghi, and Paolo Del Giudice (2015). "Inferring Synaptic Structure in Presence of Neural Interaction Time Scales." In: *PLOS ONE* 10.3, e0118412. DOI: 10.1371/journal.pone.0118412. (Visited on 2022) (cit. on p. 139).

Capone, Cristiano and Maurizio Mattia (2017). "Speed Hysteresis and Noise Shaping of Traveling Fronts in Neural Fields: Role of Local Circuitry and Nonlocal Connectivity." In: *Scientific Reports* 7.1, p. 39611. DOI: 10.1038/srep39611. (Visited on 2022) (cit. on p. 179).

Capone, Cristiano, Elena Pastorelli, Bruno Golosio, and Pier Stanislao Paolucci (2019a). "Sleep-like Slow Oscillations Improve Visual Classification through Synaptic Homeostasis and Memory Association in a Thalamo-Cortical Model." In: *Scientific Reports* 9.1, pp. 1–11. DOI: 10.1038/s41598-019-45525-0. (Visited on 2020) (cit. on pp. 20, 112, 141, 171).

Capone, Cristiano, Beatriz Rebollo, Alberto Muñoz, Xavi Illa, Paolo Del Giudice, Maria V Sanchez-Vives, and Maurizio Mattia (2019b). "Slow Waves in Cortical Slices: How Spontaneous Activity Is Shaped by Laminar Structure." In: *Cerebral Cortex* 29.1, pp. 319–335. DOI: 10.1093/cercor/bhx326. (Visited on 2021) (cit. on pp. 113, 114, 132, 145, 182).

Carlson, Kristofor D., Jayram Moorkanikara Nageswaran, Nikil Dutt, and Jeffrey L. Krichmar (2014). "An Efficient Automated Parameter Tuning Framework for Spiking Neural Networks." In: *Frontiers in Neuroscience* 8, p. 10. DOI: 10.3389/fnins.2014.00010. (Visited on 2022) (cit. on pp. 89, 108).

Carnap, Rudolf (1968). "Inductiove Logic and Inductive Intuition." In: 51, pp. 258–314. DOI: 10.1016/S0049-237X(08)71047-4 (cit. on p. 13).

Casali, Adenauer G., Olivia Gosseries, Mario Rosanova, Mélanie Boly, Simone Sarasso, Karina R. Casali, Silvia Casarotto, Marie-Aurélie Bruno, Steven Laureys, Giulio Tononi, and Marcello Massimini (2013). "A Theoretically Based Index of Consciousness Independent of Sensory Processing and Behavior." In: *Science Translational Medicine* 5.198, 198ra105–198ra105. DOI: 10.1126/scitranslmed.3006294. (Visited on 2022) (cit. on p. 176).

Cattani, Anna, Andrea Galluzzi, Matteo Fecchio, Andrea Pigorini, Maurizio Mattia, and Marcello Massimini (2022). *Adaptation Shapes Local Cortical Reactivity: From Bifurcation Diagram and Simulations to Human Physiological and Pathological Responses*. DOI: 10.1101/2022.06.11.493219. (Visited on 2022) (cit. on p. 141).

Celotto, Marco, Chiara De Luca, Paolo Muratore, Francesco Resta, Anna Letizia Allegra Mascaro, Francesco Saverio Pavone, Giulia De Bonis, and Pier Stanislao Paolucci (2020). "Analysis and Model of Cortical Slow Waves Acquired with Optical Techniques." In: *Methods and Protocols* 3.1, p. 14. DOI: 10.3390/mps3010014. (Visited on 2020) (cit. on pp. 114, 122, 126, 137).

Chalmers, D.J. (1995). "Facing up to the Problem of Consciousness." In: *Journal of Consciousness Studies* 2.3, pp. 200–219 (cit. on p. 188).

Chamove, A. S. (1989). "Cage Design Reduces Emotionality in Mice." In: *Laboratory Animals* 23.3, pp. 215–219. DOI: 10.1258/002367789780810608. (Visited on 2022) (cit. on p. 10).

Chan, Allen W., Majid H. Mohajerani, Jeffrey M. LeDue, Yu Tian Wang, and Timothy H. Murphy (2015). "Mesoscale Infraslow Spontaneous Membrane Potential Fluctuations Recapitulate High-Frequency Activity Cortical Motifs." In: *Nature Communications* 6.1, pp. 1–12. DOI: 10.1038/ncomms8738. (Visited on 2020) (cit. on p. 114).

Chaudhuri, Rishidev, Kenneth Knoblauch, Marie-Alice Gariel, Henry Kennedy, and Xiao-Jing Wang (2015). "A Large-Scale Circuit Mechanism for Hierarchical Dynamical Processing in the Primate Cortex." In: *Neuron* 88.2, pp. 419–431. DOI: 10.1016/j.neuron.2015.09.008. (Visited on 2022) (cit. on p. 108).

Chen, Tsai-Wen, Trevor J. Wardill, Yi Sun, Stefan R. Pulver, Sabine L. Renninger, Amy Baohan, Eric R. Schreiter, Rex A. Kerr, Michael B. Orger, Vivek Jayaraman, Loren L. Looger, Karel Svoboda, and Douglas S. Kim (2013). "Ultrasensitive Fluorescent Proteins for Imaging Neuronal Activity." In: *Nature* 499.7458, pp. 295–300. DOI: 10.1038/nature12354. (Visited on 2020) (cit. on pp. 126, 145).

Chen, Xing, Aitor Morales-Gregorio, Julia Sprenger, Alexander Kleinjohann, Shashwat Sridhar, Sacha J. van Albada, Sonja Grün, and Pieter R. Roelfsema (2022). "1024-Channel Electrophysiological Recordings in Macaque V1 and V4 during Resting State." In: *Scientific Data* 9.1, p. 77. DOI: 10.1038/s41597-022-01180-1. (Visited on 2022) (cit. on pp. 94, 177).

Chirimuuta, Mazviita (2021). "Your Brain Is Like a Computer: Function, Analogy, Simplification." In: *Neural Mechanisms: New Challenges in the Philosophy of Neuroscience*. Ed. by Fabrizio Calzavarini and Marco Viola. Studies in Brain and Mind. Cham: Springer International Publishing, pp. 235–261. DOI: 10.1007/978-3-030-54092-0_11. (Visited on 2022) (cit. on p. 17).

Churan, Jan, Daniel Guitton, and Christopher C. Pack (2012). "Spatiotemporal Structure of Visual Receptive Fields in Macaque Supe-

Contreras, D, I Timofeev, and M Steriade (1996). "Mechanisms of Long-Lasting Hyperpolarizations Underlying Slow Sleep Oscillations in Cat Corticothalamic Networks." In: *The Journal of Physiology* 494.1, pp. 251–264. DOI: 10.1113/jphysiol.1996.sp021488 (cit. on pp. 20, 112).

Crusoe, Michael R., Sanne Abeln, Alexandru Iosup, Peter Amstutz, John Chilton, Nebojša Tijanić, Hervé Ménager, Stian Soiland-Reyes, Bogdan Gavrilovic, and Carole Goble (2021). "Methods Included: Standardizing Computational Reuse and Portability with the Common Workflow Language." In: *arXiv:2105.07028 [cs]*. DOI: 10.1145/3486897. arXiv: 2105.07028 [cs] (cit. on p. 117).

Curto, Carina and Katherine Morrison (2019). "Relating Network Connectivity to Dynamics: Opportunities and Challenges for Theoretical Neuroscience." In: *Current Opinion in Neurobiology*. Computational Neuroscience 58, pp. 11–20. DOI: 10.1016/j.conb.2019.06.003 (cit. on p. 62).

Dąbrowska, Paulina Anna, Nicole Voges, Michael von Papen, Junji Ito, David Dahmen, Alexa Riehle, Thomas Brochier, and Sonja Grün (2021). "On the Complexity of Resting State Spiking Activity in Monkey Motor Cortex." In: *Cerebral Cortex Communications* 2.3, tgab033. DOI: 10.1093/texcom/tgab033 (cit. on pp. 93, 94).

Dahmen, David, Sonja Grün, Markus Diesmann, and Moritz Helias (2019). "Second Type of Criticality in the Brain Uncovers Rich Multiple-Neuron Dynamics." In: *Proceedings of the National Academy of Sciences* 116.26, pp. 13051–13060. DOI: 10.1073/pnas.1818972116 (cit. on pp. 62, 180).

Dahmen, David, Moritz Layer, Lukas Deutz, Paulina Anna Dąbrowska, Nicole Voges, Michael von Papen, Thomas Brochier, Alexa Riehle, Markus Diesmann, Sonja Grün, and Moritz Helias (2022). "Global Organization of Neuronal Activity Only Requires Unstructured Local Connectivity." In: *eLife* 11. Ed. by Stephanie E Palmer and Timothy E Behrens, e68422. DOI: 10.7554/eLife.68422 (cit. on p. 180).

Dana, Hod, Tsai-Wen Chen, Amy Hu, Brenda C. Shields, Caiying Guo, Loren L. Looger, Douglas S. Kim, and Karel Svoboda (2014). "Thy1-GCaMP6 Transgenic Mice for Neuronal Population Imaging In Vivo." In: *PLOS ONE* 9.9, e108697. DOI: 10.1371/journal.pone.0108697 (cit. on p. 126).

Dann, Benjamin, Jonathan A Michaels, Stefan Schaffelhofer, and Hansjörg Scherberger (2016). "Uniting Functional Network Topology and Oscillations in the Fronto-Parietal Single Unit Network of Behaving Primates." In: *eLife* 5. Ed. by Klaas Enno Stephan, e15719. DOI: 10.7554/eLife.15719 (cit. on pp. 89, 183).

Das, Abhranil and Ila R. Fiete (2020). "Systematic Errors in Connectivity Inferred from Activity in Strongly Recurrent Networks." In:

Organization of Neuronal Activity Only Requires Unstructured Local Connectivity." In: *eLife* 11. Ed. by Stephanie E Palmer and Timothy E Behrens, e68422. DOI: 10.7554/eLife.68422. (Visited on 2022) (cit. on p. 180).

Dana, Hod, Tsai-Wen Chen, Amy Hu, Brenda C. Shields, Caiying Guo, Loren L. Looger, Douglas S. Kim, and Karel Svoboda (2014). "Thy1-GCaMP6 Transgenic Mice for Neuronal Population Imaging In Vivo." In: *PLOS ONE* 9.9, e108697. DOI: 10.1371/journal.pone.0108697. (Visited on 2020) (cit. on p. 126).

Dann, Benjamin, Jonathan A Michaels, Stefan Schaffelhofer, and Hansjörg Scherberger (2016). "Uniting Functional Network Topology and Oscillations in the Fronto-Parietal Single Unit Network of Behaving Primates." In: *eLife* 5. Ed. by Klaas Enno Stephan, e15719. DOI: 10.7554/eLife.15719. (Visited on 2019) (cit. on pp. 89, 183).

Das, Abhranil and Ila R. Fiete (2020). "Systematic Errors in Connectivity Inferred from Activity in Strongly Recurrent Networks." In: *Nature Neuroscience* 23.10, pp. 1286–1296. DOI: 10.1038/s41593-020-0699-2. (Visited on 2022) (cit. on pp. 89, 178).

Dasbach, Stefan, Tom Tetzlaff, Markus Diesmann, and Johanna Senk (2021). "Prominent Characteristics of Recurrent Neuronal Networks Are Robust against Low Synaptic Weight Resolution." In: *Frontiers in Neuroscience* 15, p. 757790. DOI: 10.3389/fnins.2021.757790. arXiv: 2105.05002 [q-bio]. (Visited on 2022) (cit. on pp. 60, 174, 175).

Dasilva, Miguel, Alessandra Camassa, Alvaro Navarro-Guzman, Antonio Pazienti, Lorena Perez-Mendez, Gorka Zamora-López, Maurizio Mattia, and Maria V. Sanchez-Vives (2021). "Modulation of Cortical Slow Oscillations and Complexity across Anesthesia Levels." In: *NeuroImage* 224, p. 117415. DOI: 10.1016/j.neuroimage.2020.117415. (Visited on 2020) (cit. on p. 133).

David, Nuno (2009). "Validation and Verification in Social Simulation: Patterns and Clarification of Terminology." In: *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 5466 LNAI. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 117–129. DOI: 10.1007/978-3-642-01109-2_9 (cit. on p. 12).

Davis, Zachary W., Gabriel B. Benigno, Charlee Fletterman, Theo Desbordes, Christopher Steward, Terrence J. Sejnowski, John H. Reynolds, and Lyle Muller (2021). "Spontaneous Traveling Waves Naturally Emerge from Horizontal Fiber Time Delays and Travel through Locally Asynchronous-Irregular States." In: *Nature Communications* 12.1, p. 6057. DOI: 10.1038/s41467-021-26175-1. (Visited on 2022) (cit. on pp. 20, 170, 182).

Davis, Zachary W., Lyle Muller, Julio Martinez-Trujillo, Terrence Sejnowski, and John H. Reynolds (2020). "Spontaneous Travelling Cortical Waves Gate Perception in Behaving Primates." In: *Nature*,

pp. 1–5. DOI: 10.1038/s41586-020-2802-y. (Visited on 2020) (cit. on pp. 21, 112, 133, 171, 183–186).

Davison, Andrew P, Daniel Brüderle, Jochen M Eppler, Jens Kremkow, Eilif Muller, Dejan Pecevski, Laurent Perrinet, and Pierre Yger (2008). "PyNN: A Common Interface for Neuronal Network Simulators." In: *Frontiers in Neuroinformatics* 2, p. 11. DOI: 10.3389/neuro.11.011.2008 (cit. on pp. 12, 14, 174).

De Bonis, Giulia, Miguel Dasilva, Antonio Pazienti, Maria V. Sanchez-Vives, Maurizio Mattia, and Pier Stanislao Paolucci (2019). "Analysis Pipeline for Extracting Features of Cortical Slow Oscillations." In: *Frontiers in Systems Neuroscience* 13, p. 70. DOI: 10.3389/fnsys.2019.00070. (Visited on 2019) (cit. on pp. 113, 115, 121, 128, 176).

De Luca, Chiara, Irene Bernava, Giulia De Bonis, Francesco Simula, Francesco Resta, Elena Montagni, Anna Letizia Allegra Mascaro, Francesco Saverio Pavone, and Pier Stanislao Paolucci (2022). "Hierarchical Optimal Sampling (HOS): A Tool for Managing and Manipulating Wide-field Imaging Datasets." In: *Brain Activity across Scales and Species: Analysis of Experiments and Simulations (BASSES) Workshop* (cit. on p. 174).

De Schutter, E. and J. M. Bower (1994). "An Active Membrane Model of the Cerebellar Purkinje Cell. I. Simulation of Current Clamps in Slice." In: *Journal of Neurophysiology* 71.1, pp. 375–400. DOI: 10.1152/jn.1994.71.1.375. (Visited on 2022) (cit. on p. 16).

Dean, Heather L., Maureen A. Hagan, and Bijan Pesaran (2012). "Only Coherent Spiking in Posterior Parietal Cortex Coordinates Looking and Reaching." In: *Neuron* 73.4, pp. 829–841. DOI: 10.1016/j.neuron.2011.12.035. (Visited on 2022) (cit. on p. 171).

Deco, Gustavo, Joana Cabral, Victor M. Saenger, Melanie Boly, Enzo Tagliazucchi, Helmut Laufs, Eus Van Someren, Beatrice Jobst, Angus Stevner, and Morten L. Kringelbach (2018). "Perturbation of Whole-Brain Dynamics in Silico Reveals Mechanistic Differences between Brain States." In: *NeuroImage* 169, pp. 46–56. DOI: 10.1016/j.neuroimage.2017.12.009. (Visited on 2022) (cit. on p. 176).

Deepu, Rajalekshmi et al. (2021). *NEST 3.1*. Zenodo. DOI: 10.5281/zenodo.5508805. (Visited on 2022) (cit. on pp. 72, 173).

Denker, Michael, Cristiano Köhler, Aitor Morales-Gregorio, Alexander Kleinjohann, Peter Bouss, Alessandra Stella, Regimantas Jurkus, Maximilian Kramer, Paulina Dąbrowska, Moritz Kern, Anno Christopher Kurth, Robin Gutzen, Florian Porrmann, and Sarah Pilz (2022). *Elephant 0.11.1*. DOI: 10.5281/ZENODO.1186602. (Visited on 2022) (cit. on p. 97).

Denker, Michael, Alexa Riehle, Markus Diesmann, and Sonja Grün (2010). "Estimating the Contribution of Assembly Activity to Cortical Dynamics from Spike and Population Measures." In: *Journal of Computational Neuroscience* 29.3, pp. 599–613. DOI: 10.1007/s10827-010-0241-8 (cit. on p. 20).

Denker, Michael, Sébastien Roux, Henrik Lindén, Markus Diesmann, Alexa Riehle, and Sonja Grün (2011). "The Local Field Potential Reflects Surplus Spike Synchrony." In: *Cerebral Cortex* 21.12, pp. 2681–2695. DOI: 10.1093/cercor/bhr040 (cit. on p. 182).

Denker, Michael, Sébastien Roux, Marc Timme, Alexa Riehle, and Sonja Grün (2007). "Phase Synchronization between LFP and Spiking Activity in Motor Cortex during Movement Preparation." In: *Neurocomputing*. Computational Neuroscience: Trends in Research 2007 70.10, pp. 2096–2101. DOI: 10.1016/j.neucom.2006.10.088. (Visited on 2022) (cit. on p. 181).

Denker, Michael, Alper Yegenoglu, and Sonja Grün (2018a). "Collaborative HPC-enabled Workflows on the HBP Collaboratory Using the Elephant Framework." In: *Neuroinformatics*. DOI: 10.12751/incf.ni2018.0019. (Visited on 2021) (cit. on pp. 115, 117, 128, 151, 173).

Denker, Michael, Lyuba Zehl, Bjørg E. Kilavik, Markus Diesmann, Thomas Brochier, Alexa Riehle, and Sonja Grün (2018b). "LFP Beta Amplitude Is Linked to Mesoscopic Spatio-Temporal Phase Patterns." In: *Scientific Reports* 8.1, p. 5200. DOI: 10.1038/s41598-018-22990-7 (cit. on pp. 21, 115, 171, 177, 182).

Descartes, René (1662). *De homine*. Lugduni Batavorum : Apud Petrum Leffen & Franciscum Moyardum. (Visited on 2022) (cit. on p. 17).

Destexhe, A., D. Contreras, and M. Steriade (1999a). "Cortically-Induced Coherence of a Thalamic-Generated Oscillation." In: *Neuroscience* 92.2, pp. 427–443. DOI: 10.1016/S0306-4522(99)00024-X. (Visited on 2021) (cit. on p. 146).

Destexhe, Alain (2021). "In Silico, Computer Simulations from Neurons up to the Whole Brain." In: *eNeuro* 8.2, ENEURO.0124–21.2021. DOI: 10.1523/ENEURO.0124-21.2021. (Visited on 2022) (cit. on p. 188).

Destexhe, Alain, Diego Contreras, and Mircea Steriade (1999b). "Spatiotemporal Analysis of Local Field Potentials and Unit Discharges in Cat Cerebral Cortex during Natural Wake and Sleep States." In: *Journal of Neuroscience* 19.11, pp. 4595–4608. DOI: 10.1523/JNEUROSCI.19-11-04595.1999. (Visited on 2022) (cit. on p. 183).

Diana, Stefano (2019). *On the Barricades of the Incalculable: Against Algorithm Addiction*. (Visited on 2022) (cit. on p. 9).

Diaz-Pier, Sandra, Mikaël Naveau, Markus Butz-Ostendorf, and Abigail Morrison (2016). "Automatic Generation of Connectivity for Large-Scale Neuronal Network Models through Structural Plasticity." In: *Frontiers in Neuroanatomy* 10. DOI: 10.3389/fnana.2016.00057. (Visited on 2022) (cit. on p. 89).

Donoghue, Thomas, Matar Haller, Erik J. Peterson, Paroma Varma, Priyadarshini Sebastian, Richard Gao, Torben Noto, Antonio H. Lara, Joni D. Wallis, Robert T. Knight, Avgusta Shestyuk, and Bradley Voytek (2020). "Parameterizing Neural Power Spectra into Periodic and Aperiodic Components." In: *Nature Neuroscience* 23.12, pp. 1655–

1665. DOI: 10.1038/s41593-020-00744-x. (Visited on 2020) (cit. on p. 152).

Douglas, Rodney J., Kevan A.C. Martin, and David Whitteridge (1989). "A Canonical Microcircuit for Neocortex." In: *Neural Computation* 1.4, pp. 480–488. DOI: 10.1162/neco.1989.1.4.480. (Visited on 2022) (cit. on pp. 5, 89).

Druckmann, Shaul (2007). "A Novel Multiple Objective Optimization Framework for Constraining Conductance-Based Neuron Models by Experimental Data." In: *Frontiers in Neuroscience* 1.1, pp. 7–18. DOI: 10.3389/neuro.01.1.1.001.2007. (Visited on 2022) (cit. on p. 89).

Eickhoff, S. B., S. Jbabdi, S. Caspers, A. R. Laird, P. T. Fox, K. Zilles, and T. E. J. Behrens (2010). "Anatomical and Functional Connectivity of Cytoarchitectonic Areas within the Human Parietal Operculum." In: *Journal of Neuroscience* 30.18, pp. 6409–6421. DOI: 10.1523/JNEUROSCI.5664-09.2010. (Visited on 2022) (cit. on pp. 62, 89, 178).

El Boustani, Sami and Alain Destexhe (2009). "A Master Equation Formalism for Macroscopic Modeling of Asynchronous Irregular Activity States." In: *Neural Computation* 21.1, pp. 46–100. DOI: 10.1162/neco.2009.02-08-710. (Visited on 2022) (cit. on p. 141).

Engel, Andreas K., Peter König, Charles M. Gray, and Wolf Singer (1990). "Stimulus-Dependent Neuronal Oscillations in Cat Visual Cortex: Inter-Columnar Interaction as Determined by Cross-Correlation Analysis." In: *European Journal of Neuroscience* 2.7, pp. 588–606. DOI: 10.1111/j.1460-9568.1990.tb00449.x. (Visited on 2022) (cit. on p. 182).

English, Daniel Fine, Sam McKenzie, Talfan Evans, Kanghwan Kim, Euisik Yoon, and György Buzsáki (2017). "Pyramidal Cell-Interneuron Circuit Architecture and Dynamics in Hippocampal Networks." In: *Neuron* 96.2, 505–520.e7. DOI: 10.1016/j.neuron.2017.09.033. (Visited on 2022) (cit. on p. 89).

Ermentrout, G. Bard and David Kleinfeld (2001). "Traveling Electrical Waves in Cortex: Insights from Phase Dynamics and Speculation on a Computational Role." In: *Neuron* 29.1, pp. 33–44. DOI: 10.1016/S0896-6273(01)00178-7. (Visited on 2022) (cit. on pp. 148, 185).

Evans, C. C. (1953). "Spontaneous Excitation of the Visual Cortex and Association Areas — Lambda Waves." In: *Electroencephalography and Clinical Neurophysiology* 5.1, pp. 69–74. DOI: 10.1016/0013-4694(53)90054-6. (Visited on 2022) (cit. on pp. 161, 186).

Felleman, D J and D C Van Essen (1991). "Distributed Hierarchical Processing in the Primate Cerebral Cortex." In: *Cerebral cortex (New York, N.Y. 1991)* 1.1, pp. 1–47. DOI: 10.1093/cercor/1.1.1-a. (Visited on 2022) (cit. on p. 88).

Felsen, Gidon and Yang Dan (2005). "A Natural Approach to Studying Vision." In: *Nature Neuroscience* 8.12, pp. 1643–1646. DOI: 10.1038/nn1608. (Visited on 2022) (cit. on p. 10).

Ferrier, David (1874). "On the Localisation of the Functions of the Brain." In: *British Medical Journal* 2.729, pp. 766–767. (Visited on 2022) (cit. on p. 16).

Fiebelkorn, Ian C., Mark A. Pinsk, and Sabine Kastner (2018). "A Dynamic Interplay within the Frontoparietal Network Underlies Rhythmic Spatial Attention." In: *Neuron* 99.4, 842–853.e8. DOI: 10.1016/j.neuron.2018.07.038. (Visited on 2022) (cit. on p. 183).

Flury, Bernhard (1988). *Common Principal Components & Related Multivariate Models*. New York: John Wiley & Sons, Inc. (Visited on 2022) (cit. on p. 63).

Forrester, Jay W. and Peter M. Senge (1980). "Tests for Building Confidence in System Dynamics Models." In: *System dynamics, TIMS studies in management sciences* 14, pp. 209–228 (cit. on pp. 13, 15).

Franklin, A., M. Anderson, D. Brock, S. Coleman, J. Downing, A. Gruvander, J. Lilly, J. Neal, D. Peterson, M. Price, R. Rice, L. Smith, S. Speirer, and D. Toering (1989). "Can a Theory-Laden Observation Test the Theory?" In: *The British Journal for the Philosophy of Science* 40.2, pp. 229–231. JSTOR: 687514. (Visited on 2022) (cit. on p. 10).

Freeman, Walter J. and John M. Barrie (2000). "Analysis of Spatial Patterns of Phase in Neocortical Gamma EEGs in Rabbit." In: *Journal of Neurophysiology* 84.3, pp. 1266–1278. DOI: 10.1152/jn.2000.84.3.1266. (Visited on 2022) (cit. on p. 21).

Fries, Pascal (2005). "A Mechanism for Cognitive Dynamics: Neuronal Communication through Neuronal Coherence." In: *Trends in Cognitive Sciences* 9.10, pp. 474–480. DOI: 10.1016/j.tics.2005.08.011. (Visited on 2022) (cit. on p. 183).

– (2015). "Rhythms For Cognition: Communication Through Coherence." In: *Neuron* 88.1, pp. 220–235. DOI: 10.1016/j.neuron.2015.09.034. (Visited on 2022) (cit. on p. 183).

Fries, Pascal, John H. Reynolds, Alan E. Rorie, and Robert Desimone (2001). "Modulation of Oscillatory Neuronal Synchronization by Selective Visual Attention." In: *Science* 291.5508, pp. 1560–1563. DOI: 10.1126/science.1055465. (Visited on 2022) (cit. on pp. 182, 183).

Friston, K. J., C. D. Frith, P. F. Liddle, and R. S. J. Frackowiak (1993). "Functional Connectivity: The Principal-Component Analysis of Large (PET) Data Sets." In: *Journal of Cerebral Blood Flow & Metabolism* 13.1, pp. 5–14. DOI: 10.1038/jcbfm.1993.4. (Visited on 2022) (cit. on p. 41).

Friston, Karl (2010). "The Free-Energy Principle: A Unified Brain Theory?" In: *Nature Reviews Neuroscience* 11.2, pp. 127–138. DOI: 10.1038/nrn2787. (Visited on 2022) (cit. on pp. 6, 19).

Fröhlich, Flavio and David A. McCormick (2010). "Endogenous Electric Fields May Guide Neocortical Network Activity." In: *Neuron* 67.1, pp. 129–143. DOI: 10.1016/j.neuron.2010.06.005 (cit. on p. 184).

Furber, Steve B., David R. Lester, Luis A. Plana, Jim D. Garside, Eustace Painkras, Steve Temple, and Andrew D. Brown (2013). "Overview of the SpiNNaker System Architecture." In: *IEEE Transactions on Computers* 62.12, pp. 2454–2467. DOI: 10.1109/TC.2012.142. (Visited on 2022) (cit. on pp. 35, 47).

Gabriel, Andreas and Reinhard Eckhorn (2003). "A Multi-Channel Correlation Method Detects Traveling γ-Waves in Monkey Visual Cortex." In: *Journal of Neuroscience Methods* 131.1, pp. 171–184. DOI: 10.1016/j.jneumeth.2003.08.008. (Visited on 2022) (cit. on p. 21).

Galinsky, Vitaly L. and Lawrence R. Frank (2020). "Universal Theory of Brain Waves: From Linear Loops to Nonlinear Synchronized Spiking and Collective Brain Rhythms." In: *Physical Review Research* 2.2, p. 023061. DOI: 10.1103/PhysRevResearch.2.023061. (Visited on 2020) (cit. on p. 146).

Gallego, Juan A., Matthew G. Perich, Lee E. Miller, and Sara A. Solla (2017). "Neural Manifolds for the Control of Movement." In: *Neuron* 94.5, pp. 978–984. DOI: 10.1016/j.neuron.2017.05.025. (Visited on 2022) (cit. on p. 19).

Galvani, Luigi (1792). *De Viribus Electricitatis In Motu Musculari Commentarius*. Apud Societatem typographicam (cit. on p. 17).

García del Molino, Luis Carlos, Khashayar Pakdaman, Jonathan Touboul, and Gilles Wainrib (2013). "Synchronization in Random Balanced Networks." In: *Physical Review E* 88.4, p. 042824. DOI: 10.1103/PhysRevE.88.042824. (Visited on 2022) (cit. on p. 180).

Garcia, Samuel, Domenico Guarino, Florent Jaillet, Todd R. Jennings, Robert Pröpper, Philipp L. Rautenberg, Chris Rodgers, Andrey Sobolev, Thomas Wachtler, Pierre Yger, and Andrew P. Davison (2014). "Neo: An Object Model for Handling Electrophysiology Data in Multiple Formats." In: *Frontiers in Neuroinformatics* 8.February, p. 10. DOI: 10.3389/fninf.2014.00010. (Visited on 2020) (cit. on pp. 115, 119, 128, 150, 173, 176).

Garijo, Daniel, Yolanda Gil, and Oscar Corcho (2017). "Abstract, Link, Publish, Exploit: An End to End Framework for Workflow Sharing." In: *Future Generation Computer Systems* 75, pp. 271–283. DOI: 10.1016/j.future.2017.01.008. (Visited on 2021) (cit. on p. 117).

Gerhard, Felipe, Tilman Kispersky, Gabrielle J. Gutierrez, Eve Marder, Mark Kramer, and Uri Eden (2013). "Successful Reconstruction of a Physiological Circuit with Known Connectivity from Spiking Activity Alone." In: *PLoS Computational Biology* 9.7. Ed. by Olaf Sporns, e1003138. DOI: 10.1371/journal.pcbi.1003138. (Visited on 2022) (cit. on pp. 62, 178).

Gerstner, Wulfram, Marco Lehmann, Vasiliki Liakoni, Dane Corneil, and Johanni Brea (2018). "Eligibility Traces and Plasticity on Behavioral Time Scales: Experimental Support of NeoHebbian Three-Factor Learning Rules." In: *Frontiers in Neural Circuits* 12, p. 53. (Visited on 2022) (cit. on p. 178).

Gewaltig, Marc-Oliver and Markus Diesmann (2007). "NEST (NEural Simulation Tool)." In: *Scholarpedia* 2.4, p. 1430. DOI: 10.4249/scholarpedia.1430 (cit. on p. 14).

Giannasi, Frank, Philip Lovett, and Anthony N. Godwin (2001). "Enhancing Confidence in Discrete Event Simulations." In: *Computers in Industry* 44.2, pp. 141–157. DOI: 10.1016/S0166-3615(00)00084-1 (cit. on p. 14).

Gigante, Guido, Maurizio Mattia, and Paolo Del Giudice (2007). "Diverse Population-Bursting Modes of Adapting Spiking Neurons." In: *Physical Review Letters* 98.14, p. 148101. DOI: 10.1103/PhysRevLett.98.148101. (Visited on 2022) (cit. on p. 141).

Girko, V. L. (1985). "Circular Law." In: *Theory of Probability & Its Applications* 29.4, pp. 694–706. DOI: 10.1137/1129095. (Visited on 2022) (cit. on p. 71).

Glatard, Tristan, Lindsay B. Lewis, Rafael Ferreira da Silva, Reza Adalat, Natacha Beck, Claude Lepage, Pierre Rioux, Marc-Etienne Rousseau, Tarek Sherif, Ewa Deelman, Najmeh Khalili-Mahani, and Alan C. Evans (2015). "Reproducibility of Neuroimaging Analyses across Operating Systems." In: *Frontiers in Neuroinformatics* 9, p. 12. DOI: 10.3389/fninf.2015.00012 (cit. on pp. 39, 175).

Gleeson, Padraig, Sharon Crook, Robert C. Cannon, Michael L. Hines, Guy O. Billings, Matteo Farinella, Thomas M. Morse, Andrew P. Davison, Subhasis Ray, Upinder S. Bhalla, Simon R. Barnes, Yoana D. Dimitrova, and R. Angus Silver (2010). "NeuroML: A Language for Describing Data Driven Models of Neurons and Networks with a High Degree of Biological Detail." In: *PLoS Computational Biology* 6.6. Ed. by Karl J. Friston, e1000815. DOI: 10.1371/journal.pcbi.1000815 (cit. on p. 12).

Goldman, Jennifer S., Núria Tort-Colet, Matteo di Volo, Eduarda Susin, Jules Bouté, Melissa Dali, Mallory Carlu, Trang-Anh Nghiem, Tomasz Górski, and Alain Destexhe (2019). "Bridging Single Neuron Dynamics to Global Brain States." In: *Frontiers in Systems Neuroscience* 13, p. 75. (Visited on 2022) (cit. on p. 141).

Goldstein, Julie, Jules Davidoff, and Debi Roberson (2009). "Knowing Color Terms Enhances Recognition: Further Evidence from English and Himba." In: *Journal of Experimental Child Psychology* 102.2, pp. 219–238. DOI: 10.1016/j.jecp.2008.06.002. (Visited on 2021) (cit. on p. 6).

Golomb, D., X. J. Wang, and J. Rinzel (1996). "Propagation of Spindle Waves in a Thalamic Slice Model." In: *Journal of Neurophysiology* 75, pp. 750–769. DOI: 10.1152/jn.1996.75.2.750. (Visited on 2021) (cit. on p. 146).

Golosio, Bruno, Chiara De Luca, Cristiano Capone, Elena Pastorelli, Giovanni Stegel, Gianmarco Tiddia, Giulia De Bonis, and Pier Stanislao Paolucci (2021a). "Thalamo-Cortical Spiking Model of Incremental Learning Combining Perception, Context and NREM-sleep." In:

PLOS Computational Biology 17.6, e1009045. DOI: 10.1371/journal.pcbi.1009045. (Visited on 2021) (cit. on p. 112).

Golosio, Bruno, Gianmarco Tiddia, Chiara De Luca, Elena Pastorelli, Francesco Simula, and Pier Stanislao Paolucci (2021b). "Fast Simulations of Highly-Connected Spiking Cortical Models Using GPUs." In: *Frontiers in Computational Neuroscience* 15, p. 627620. (Visited on 2022) (cit. on p. 175).

Gomez, Jesse, Michael Barnett, and Kalanit Grill-Spector (2019). "Extensive Childhood Experience with Pokémon Suggests Eccentricity Drives Organization of Visual Cortex." In: *Nature Human Behaviour* 3.6, pp. 611–624. DOI: 10.1038/s41562-019-0592-8. (Visited on 2022) (cit. on p. 6).

Gonçalves, Pedro J, Jan-Matthis Lueckmann, Michael Deistler, Marcel Nonnenmacher, Kaan Öcal, Giacomo Bassetto, Chaitanya Chintaluri, William F Podlaski, Sara A Haddad, Tim P Vogels, David S Greenberg, and Jakob H Macke (2020). "Training Deep Neural Density Estimators to Identify Mechanistic Models of Neural Dynamics." In: *eLife* 9, e56261. DOI: 10.7554/eLife.56261. (Visited on 2022) (cit. on p. 89).

Gong, Pulin and Cees van Leeuwen (2009). "Distributed Dynamical Computation in Neural Circuits with Propagating Coherent Activity Patterns." In: *PLOS Computational Biology* 5.12, e1000611. DOI: 10.1371/journal.pcbi.1000611. (Visited on 2019) (cit. on p. 21).

Goodman, Dan F M. (2009). "The Brian Simulator." In: *Frontiers in Neuroscience* 3.2, pp. 192–197. DOI: 10.3389/neuro.01.026.2009. (Visited on 2022) (cit. on p. 14).

Goodman, Steven N., Daniele Fanelli, and John P. A. Ioannidis (2016). "What Does Research Reproducibility Mean?" In: *Science Translational Medicine* 8.341, 341ps12–341ps12. DOI: 10.1126/scitranslmed.aaf5027. (Visited on 2021) (cit. on pp. 14, 39).

Grant, Edward and Professor Emeritus Edward Grant (1996). *The Foundations of Modern Science in the Middle Ages: Their Religious, Institutional and Intellectual Contexts*. Cambridge University Press (cit. on p. 9).

Greenberg, Anastasia, Javad Karimi Abadchi, Clayton T. Dickson, and Majid H. Mohajerani (2018). "New Waves: Rhythmic Electrical Field Stimulation Systematically Alters Spontaneous Slow Dynamics across Mouse Neocortex." In: *NeuroImage* 174, pp. 328–339. DOI: 10.1016/j.neuroimage.2018.03.019. (Visited on 2022) (cit. on pp. 132, 136).

Greenwood, Priscilla E., Mark D. McDonnell, and Lawrence M. Ward (2015). "Dynamics of Gamma Bursts in Local Field Potentials." In: *Neural Computation* 27.1, pp. 74–103. DOI: 10.1162/NECO_a_00688 (cit. on p. 21).

Greenwood, Priscilla E. and Lawrence M. Ward (2022). "Phase Offset Determines Alpha Modulation of Gamma Phase Coherence and

Hence Signal Transmission." In: *Biosystems* 219, p. 104729. DOI: 10.1016/j.biosystems.2022.104729. (Visited on 2022) (cit. on p. 184).

Grewe, Jan, Thomas Wachtler, and Jan Benda (2011). "A Bottom-up Approach to Data Annotation in Neurophysiology." In: *Frontiers in Neuroinformatics* 5, p. 16. (Visited on 2022) (cit. on p. 173).

Grün, Sonja (2009). "Data-Driven Significance Estimation for Precise Spike Correlation." In: *Journal of Neurophysiology* 101.3, pp. 1126–1140. DOI: 10.1152/jn.00093.2008 (cit. on pp. 86, 180).

Grün, Sonja, Markus Diesmann, and Ad Aertsen (2002). "Unitary Events in Multiple Single-Neuron Spiking Activity: I. Detection and Significance." In: *Neural Computation* 14.1, pp. 43–80. DOI: 10.1162/089976602753284455 (cit. on pp. 19, 181).

Guhr, Thomas, Axel Müller–Groeling, and Hans A. Weidenmüller (1998). "Random-Matrix Theories in Quantum Physics: Common Concepts." In: *Physics Reports* 299.4, pp. 189–425. DOI: 10.1016/S0370-1573(97)00088-4. (Visited on 2022) (cit. on p. 65).

Gütig, Robert (2014). "To Spike, or When to Spike?" In: *Current Opinion in Neurobiology*. Theoretical and Computational Neuroscience 25, pp. 134–139. DOI: 10.1016/j.conb.2014.01.004. (Visited on 2022) (cit. on p. 25).

Gutkin, Boris, G. Bard Ermentrout, and Michael Rudolph (2003). "Spike Generating Dynamics and the Conditions for Spike-Time Precision in Cortical Neurons." In: *Journal of Computational Neuroscience* 15.1, pp. 91–103. DOI: 10.1023/A:1024426903582. (Visited on 2022) (cit. on p. 25).

Gutzen, Robin, Giulia De Bonis, Chiara De Luca, Elena Pastorelli, Cristiano Capone, Anna Letizia Allegra Mascaro, Francesco Resta, Arnau Manasanch, Francesco Saverio Pavone, Maria V. Sanchez-Vives, Maurizio Mattia, Sonja Grün, Pier Stanislao Paolucci, and Michael Denker (2024). "A Modular and Adaptable Analysis Pipeline to Compare Slow Cerebral Rhythms across Heterogeneous Datasets." In: *Cell Reports Methods* 4.1, p. 100681. DOI: 10.1016/j.crmeth.2023.100681 (cit. on pp. 28, 111).

Gutzen, Robin, Chiara De Luca, and Giulia De Bonis (2023). *Collaborative Brain Wave Analyisis Pipeline (Cobrawap)*. RRID:SCR_022966. DOI: 10.5281/zenodo.10198748 (cit. on p. 174).

Gutzen, Robin, Sonja Grün, and Michael Denker (2022). "Evaluating the Statistical Similarity of Neural Network Activity and Connectivity via Eigenvector Angles." In: *Biosystems* 223, p. 104813. DOI: 10.1016/j.biosystems.2022.104813 (cit. on pp. 28, 61).

Gutzen, Robin, Aitor Morales-Gregorio, Michael Von Papen, and Michael Denker (2018a). *NetworkUnit*. RRID:SCR_016543 (cit. on p. 173).

Gutzen, Robin, Michael von Papen, Guido Trensch, Pietro Quaglio, Sonja Grün, and Michael Denker (2018b). "Reproducible Neural Network Simulations: Statistical Methods for Model Validation on

the Level of Network Activity Data." In: *Frontiers in Neuroinformatics* 12, p. 90. DOI: 10.3389/fninf.2018.00090 (cit. on pp. 28, 33, 78, 97).

Haber, Adam and Elad Schneidman (2020). "Learning the Architectural Features That Predict Functional Similarity of Neural Networks." In: *bioRxiv*, p. 2020.04.27.057752. DOI: 10.1101/2020.04.27.057752. (Visited on 2020) (cit. on p. 62).

Hacking, Ian and Emeritus University Professor Ian Hacking (1983). *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge University Press (cit. on p. 10).

Hagen, Espen, Solveig Næss, Torbjørn V. Ness, and Gaute T. Einevoll (2018). "Multimodal Modeling of Neural Network Activity: Computing LFP, ECoG, EEG, and MEG Signals With LFPy 2.0." In: *Frontiers in Neuroinformatics* 12, p. 92. (Visited on 2022) (cit. on p. 177).

Halgren, Mila, Daniel Fabó, István Ulbert, Joseph R. Madsen, Lorand Erőss, Werner K. Doyle, Orrin Devinsky, Donald Schomer, Sydney S. Cash, and Eric Halgren (2018). "Superficial Slow Rhythms Integrate Cortical Processing in Humans." In: *Scientific Reports* 8.1, p. 2055. DOI: 10.1038/s41598-018-20662-0. (Visited on 2022) (cit. on pp. 21, 145, 182–184).

Hall, Bradford, Advait Limaye, and Ashok B. Kulkarni (2009). "Overview: Generation of Gene Knockout Mice." In: *Current Protocols in Cell Biology* 44.1, pp. 19.12.1–19.12.17. DOI: 10.1002/0471143030.cb1912s44. (Visited on 2022) (cit. on p. 3).

Han, Kyung-Seok, Chong Guo, Christopher H. Chen, Laurens Witter, Tomas Osorno, and Wade G. Regehr (2018). "Ephaptic Coupling Promotes Synchronous Firing of Cerebellar Purkinje Cells." In: *Neuron* 100.3, 564–578.e3. DOI: 10.1016/j.neuron.2018.09.018. (Visited on 2019) (cit. on p. 184).

Hanlon, Erin C., Ugo Faraguna, Vladyslav V. Vyazovskiy, Giulio Tononi, and Chiara Cirelli (2009). "Effects of Skilled Training on Sleep Slow Wave Activity and Cortical Gene Expression in the Rat." In: *Sleep* 32.6, pp. 719–729. DOI: 10.1093/sleep/32.6.719. (Visited on 2021) (cit. on pp. 20, 112).

Harris, Kd, Jozsef Csicsvari, Hajime Hirase, George Dragoi, and György Buzsáki (2003). "Organization of Cell Assemblies in the Hippocampus." In: *Nature* 424.July, pp. 552–556. DOI: 10.1038/nature01765.1. (cit. on p. 166).

Harris, Kenneth D (2005). "Neural Signatures of Cell Assembly Organization." In: *Nature Reviews Neuroscience* 6.5, pp. 399–407. DOI: 10.1038/nrn1669 (cit. on pp. 4, 85, 180).

Hedges, L. V. (1981). "Distribution Theory for Glass's Estimator of Effect Size and Related Estimators." In: *Journal of Educational and Behavioral Statistics* 6.2, pp. 107–128. DOI: 10.3102/10769986006002107 (cit. on p. 43).

Heinze, Thomas (2008). "How to Sponsor Ground-Breaking Research: A Comparison of Funding Schemes." In: *Science and Public Policy*

35.5, pp. 302–318. DOI: 10.3152/030234208X317151. (Visited on 2022) (cit. on p. 188).

Heitmann, Stewart, Tjeerd Boonstra, Pulin Gong, Michael Breakspear, and Bard Ermentrout (2015). "The Rhythms of Steady Posture: Motor Commands as Spatially Organized Oscillation Patterns." In: *Neurocomputing*. Advances on Biological Rhythmic Pattern Generation: Experiments, Algorithms and Applications 170, pp. 3–14. DOI: 10.1016/j.neucom.2015.01.088. (Visited on 2022) (cit. on pp. 21, 112, 171).

Heitmann, Stewart, Pulin Gong, and Michael Breakspear (2012). "A Computational Role for Bistability and Traveling Waves in Motor Cortex." In: *Frontiers in Computational Neuroscience* 6, p. 67. (Visited on 2022) (cit. on p. 21).

Helfrich, Randolph F., Ian C. Fiebelkorn, Sara M. Szczepanski, Jack J. Lin, Josef Parvizi, Robert T. Knight, and Sabine Kastner (2018). "Neural Mechanisms of Sustained Attention Are Rhythmic." In: *Neuron* 99.4, 854–865.e5. DOI: 10.1016/j.neuron.2018.07.032. (Visited on 2019) (cit. on pp. 183, 184).

Helias, Moritz and David Dahmen (2020). "Dynamic Mean-Field Theory for Random Networks." In: *Statistical Field Theory for Neural Networks*. Ed. by Moritz Helias and David Dahmen. Lecture Notes in Physics. Cham: Springer International Publishing, pp. 95–126. DOI: 10.1007/978-3-030-46444-8_10. (Visited on 2022) (cit. on p. 16).

Hindriks, Rikkert, Michel J. A. M. van Putten, and Gustavo Deco (2014). "Intra-Cortical Propagation of EEG Alpha Oscillations." In: *NeuroImage* 103, pp. 444–453. DOI: 10.1016/j.neuroimage.2014.08.027. (Visited on 2022) (cit. on p. 5).

Hines, M. L. and N. T. Carnevale (1997). "The NEURON Simulation Environment." In: *Neural Computation* 9.6, pp. 1179–1209. DOI: 10.1162/neco.1997.9.6.1179. (Visited on 2022) (cit. on p. 14).

Hodges, J. L. (1958). "The Significance Probability of the Smirnov Two-Sample Test." In: *Arkiv för Matematik* 3.5, pp. 469–486. DOI: 10.1007/BF02589501. (Visited on 2022) (cit. on pp. 62, 75).

Hodgkin, A. L. and A. F. Huxley (1952). "A Quantitative Description of Membrane Current and Its Application to Conduction and Excitation in Nerve." In: *The Journal of Physiology* 117.4, pp. 500–544. DOI: 10.1113/jphysiol.1952.sp004764 (cit. on p. 4).

Hoel, Erik P. (2018). "Agent Above, Atom Below: How Agents Causally Emerge from Their Underlying Microphysics." In: *Wandering Towards a Goal: How Can Mindless Mathematical Laws Give Rise to Aims and Intention?* Ed. by Anthony Aguirre, Brendan Foster, and Zeeya Merali. The Frontiers Collection. Cham: Springer International Publishing, pp. 63–76. DOI: 10.1007/978-3-319-75726-1_6. (Visited on 2022) (cit. on p. 6).

Hoel, Erik P., Larissa Albantakis, William Marshall, and Giulio Tononi (2016). "Can the Macro Beat the Micro? Integrated Information across Spatiotemporal Scales." In: *Neuroscience of Consciousness* 2016.1, niw012. DOI: 10.1093/nc/niw012. (Visited on 2022) (cit. on p. 6).

Hofman, Michel A. (1988). "Size and Shape of the Cerebral Cortex in Mammals." In: *Brain, Behavior and Evolution* 32.1, pp. 17–26. DOI: 10.1159/000116529. (Visited on 2022) (cit. on p. 5).

Holmes, Richard B (1991). "On Random Correlation Matrices." In: *SIAM journal on matrix analysis and applications* 12.2, pp. 239–272. DOI: 10.1137/0612019 (cit. on p. 66).

Holt, G. R., W. R. Softky, C. Koch, and R. J. Douglas (1996). "Comparison of Discharge Variability in Vitro and in Vivo in Cat Visual Cortex Neurons." In: *Journal of Neurophysiology* 75.5, pp. 1806–1814. DOI: 10.1152/jn.1996.75.5.1806. (Visited on 2022) (cit. on p. 97).

Hong, Guosong and Charles M. Lieber (2019). "Novel Electrode Technologies for Neural Recordings." In: *Nature Reviews Neuroscience* 20.6, pp. 330–345. DOI: 10.1038/s41583-019-0140-6. (Visited on 2021) (cit. on pp. 10, 24).

Hopfield, J. J. (1995). "Pattern Recognition Computation Using Action Potential Timing for Stimulus Representation." In: *Nature* 376.6535, pp. 33–36. DOI: 10.1038/376033a0. (Visited on 2022) (cit. on p. 184).

Hopkins, Michael and Steve Furber (2015). "Accuracy and Efficiency in Fixed-Point Neural ODE Solvers." In: *Neural Computation* 27.10, pp. 2148–2182. DOI: 10.1162/NECO_a_00772. (Visited on 2022) (cit. on p. 49).

Horsman, Clare, Susan Stepney, Rob C. Wagner, and Viv Kendon (2014). "When Does a Physical System Compute?" In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 470.2169, p. 20140182. DOI: 10.1098/rspa.2014.0182. arXiv: 1309.7979. (Visited on 2021) (cit. on p. 18).

Huang, Xiaoying, Weifeng Xu, Jianmin Liang, Kentaroh Takagaki, Xin Gao, and Jian-young Wu (2010). "Spiral Wave Dynamics in Neocortex." In: *Neuron* 68.5, pp. 978–990. DOI: 10.1016/j.neuron.2010.11.007. (Visited on 2018) (cit. on p. 115).

Hubel, D. H. and T. N. Wiesel (1959). "Receptive Fields of Single Neurones in the Cat's Striate Cortex." In: *The Journal of Physiology* 148.3, pp. 574–591. DOI: 10.1113/jphysiol.1959.sp006308 (cit. on p. 19).

Igel, Christian and Michael Hüsken (2003). "Empirical Evaluation of the Improved Rprop Learning Algorithms." In: *Neurocomputing* 50, pp. 105–123. DOI: 10.1016/S0925-2312(01)00700-7. (Visited on 2022) (cit. on p. 141).

Ikegaya, Yuji, Gloster Aaron, Rosa Cossart, Dmitriy Aronov, Ilan Lampl, David Ferster, and Rafael Yuste (2004). "Synfire Chains and Cortical Songs: Temporal Modules of Cortical Activity." In: *Science*

304.5670, pp. 559–564. DOI: 10.1126/science.1093173. (Visited on 2022) (cit. on p. 19).

Irobi, I S, Johan Andersson, and Anders Wall (2004). "Correctness Criteria for Models' Validation - A Philosophical Perspective." In: *Proceedings of the International Conference on Modeling, Simulation & Visualization Methods, MSV '04 & Proceedings of the International Conference on Algorithmic Mathematics & Computer Science, AMCS '04, June 21-24, 2004, Las Vegas, Nevada, USA*, pp. 475–476 (cit. on p. 12).

Ito, Junji, Pedro Maldonado, Wolf Singer, and Sonja Grün (2011). "Saccade-Related Modulations of Neuronal Excitability Support Synchrony of Visually Elicited Spikes." In: *Cerebral Cortex* 21.11, pp. 2482–2497. DOI: 10.1093/cercor/bhr020. (Visited on 2022) (cit. on p. 186).

Izhikevich, E.M. (2003). "Simple Model of Spiking Neurons." In: *IEEE Transactions on Neural Networks* 14.6, pp. 1569–1572. DOI: 10.1109/TNN.2003.820440. (Visited on 2022) (cit. on p. 47).

– (2004). "Which Model to Use for Cortical Spiking Neurons?" In: *IEEE Transactions on Neural Networks* 15.5, pp. 1063–1070. DOI: 10.1109/TNN.2004.832719. (Visited on 2022) (cit. on p. 12).

Izhikevich, Eugene M. (2006). "Polychronization: Computation with Spikes." In: *Neural Computation* 18.2, pp. 245–282. DOI: 10.1162/089976606775093882 (cit. on pp. 35, 41, 46, 49, 51, 58, 78).

Izquierdo, Eduardo J. and Randall D. Beer (2013). "Connecting a Connectome to Behavior: An Ensemble of Neuroanatomical Models of C. Elegans Klinotaxis." In: *PLOS Computational Biology* 9.2, e1002890. DOI: 10.1371/journal.pcbi.1002890. (Visited on 2022) (cit. on p. 3).

Jensen, Ole, Yali Pan, Steven Frisson, and Lin Wang (2021). "An Oscillatory Pipelining Mechanism Supporting Previewing during Visual Exploration and Reading." In: *Trends in Cognitive Sciences* 25.12, pp. 1033–1044. DOI: 10.1016/j.tics.2021.08.008. (Visited on 2022) (cit. on pp. 185–187).

Jercog, Daniel, Alex Roxin, Peter Bartho, Arthur Luczak, Albert Compte, and Jaime De La Rocha (2017). "UP-DOWN Cortical Dynamics Reflect State Transitions in a Bistable Network." In: *eLife* 6, e22425. DOI: 10.7554/eLife.22425.001. (Visited on 2020) (cit. on p. 113).

Jonas, Eric and Konrad Paul Kording (2017). "Could a Neuroscientist Understand a Microprocessor?" In: *PLOS Computational Biology* 13.1. Ed. by Jörn Diedrichsen, e1005268. DOI: 10.1371/journal.pcbi.1005268. (Visited on 2018) (cit. on p. 11).

Jorgensen, Larry M (2009). "The Principle of Continuity and Leibniz's Theory of Consciousness." In: *Journal of the History of Philosophy* 47.2, pp. 223–248. DOI: 10.1353/hph.0.0112 (cit. on p. 9).

Jutras, Michael J., Pascal Fries, and Elizabeth A. Buffalo (2013). "Oscillatory Activity in the Monkey Hippocampus during Visual Exploration and Memory Formation." In: *Proceedings of the National*

Academy of Sciences 110.32, pp. 13144–13149. DOI: 10.1073/pnas.1302351110. (Visited on 2022) (cit. on pp. 185, 186).

Juusola, Mikko, Hugh P.C. Robinson, and Gonzalo G. de Polavieja (2007). "Coding with Spike Shapes and Graded Potentials in Cortical Networks." In: *BioEssays* 29.2, pp. 178–187. DOI: 10.1002/bies.20532. (Visited on 2022) (cit. on p. 18).

Kandel, Eric R. (1968). "Dale's Principle and the Functional Specificity of Neurons." In: *Psychopharmacology; a Review of Progress, 1957-1967*, pp. 385–398 (cit. on p. 4).

Kandel, Eric R., James H. (James Harris) Schwartz, and Thomas M. Jessell (2001). *Principles of Neural Science*. Vol. 4. McGraw-Hill, Health Professions Division (cit. on pp. 3, 4, 18).

Kang, Ling, Jonas Ranft, and Vincent Hakim (2023). "Beta Oscillations and Waves in Motor Cortex Can Be Accounted for by the Interplay of Spatially-Structured Connectivity and Fluctuating Inputs." In: *eLife* 12. Ed. by Nicole C Swann, e81446. DOI: 10.7554/eLife.81446. (Visited on 2023) (cit. on p. 179).

Katz, Leonard and Roger Q. Cracco (1971). "A Review of Cerebral Rhythms in the Waking EEG." In: *American Journal of EEG Technology* 11.3, pp. 77–100. DOI: 10.1080/00029238.1971.11080839. (Visited on 2022) (cit. on p. 5).

Kazanina, Nina and Alessandro Tavano (2022). "What Neural Oscillations Can and Cannot Do for Syntactic Structure Building." In: *Nature Reviews Neuroscience*, pp. 1–16. DOI: 10.1038/s41583-022-00659-5. (Visited on 2022) (cit. on p. 185).

Keane, Adam and Pulin Gong (2015). "Propagating Waves Can Explain Irregular Neural Dynamics." In: *The Journal of Neuroscience* 35.4, pp. 1591–605. DOI: 10.1523/JNEUROSCI.1669-14.2015. PMID: 25632135 (cit. on pp. 21, 113, 114).

Kempter, Richard and Wulfram Gerstner (1995). "Temporal Coding in the Sub-Millisecond Range: Model of Barn Owl Auditory Pathway." In: *Advances in neural information processing systems* 8, p. 7 (cit. on p. 19).

Kilavik, B. E. (2018). *Directional Selectivity across Macaque Motor Cortical Layers during Reach Planning and Execution*. Tech. rep. (Visited on 2022) (cit. on p. 93).

Kiselev, Mikhail (2016). "Rate Coding vs. Temporal Coding - Is Optimum Between?" In: *2016 International Joint Conference on Neural Networks (IJCNN)*, pp. 1355–1359. DOI: 10.1109/IJCNN.2016.7727355 (cit. on p. 19).

Kleiser, Raimund, Rüdiger J Seitz, and Bart Krekelberg (2004). "Neural Correlates of Saccadic Suppression in Humans." In: *Current Biology* 14.5, pp. 386–390. DOI: 10.1016/j.cub.2004.02.036. (Visited on 2022) (cit. on pp. 161, 186).

Knoblauch, Andreas and Günther Palm (2005). "What Is Signal and What Is Noise in the Brain?" In: *Biosystems*. Selection of Papers from

the 5th International Workshop on Neural Coding 79.1, pp. 83–90. DOI: 10.1016/j.biosystems.2004.09.007. (Visited on 2022) (cit. on p. 25).

Kobayashi, Ryota, Shuhei Kurita, Anno Kurth, Katsunori Kitano, Kenji Mizuseki, Markus Diesmann, Barry J. Richmond, and Shigeru Shinomoto (2019). "Reconstructing Neuronal Circuitry from Parallel Spike Trains." In: *Nature Communications* 10.1, p. 4468. DOI: 10.1038/s41467-019-12225-2. (Visited on 2022) (cit. on p. 89).

Koch, Christof and Gilles Laurent (1999). "Complexity and the Nervous System." In: *Science* 284.5411, pp. 96–98. DOI: 10.1126/science.284.5411.96. (Visited on 2022) (cit. on p. 3).

Koch, Christof and Idan Segev (2000). "The Role of Single Neurons in Information Processing." In: *Nature Neuroscience* 3.S11, pp. 1171–1177. DOI: 10.1038/81444. (Visited on 2022) (cit. on p. 12).

Kodandaramaiah, Suhasa B, Francisco J Flores, Gregory L Holst, Annabelle C Singer, Xue Han, Emery N Brown, Edward S Boyden, and Craig R Forest (2018). "Multi-Neuron Intracellular Recording in Vivo via Interacting Autopatching Robots." In: *eLife* 7, e24656. DOI: 10.7554/eLife.24656. (Visited on 2022) (cit. on p. 89).

König, Peter, Andreas K. Engel, and Wolf Singer (1996). "Integrator or Coincidence Detector? The Role of the Cortical Neuron Revisited." In: *Trends in Neurosciences* 19.4, pp. 130–137. DOI: 10.1016/S0166-2236(96)80019-1. (Visited on 2022) (cit. on p. 5).

Koo, Bang-Bon, Steven P. Schettler, Donna E. Murray, Jong-Min Lee, Ronald J. Killiany, Douglas L. Rosene, Dae-Shik Kim, and Itamar Ronen (2012). "Age-Related Effects on Cortical Thickness Patterns of the Rhesus Monkey Brain." In: *Neurobiology of Aging* 33.1, 200.e23–200.e31. DOI: 10.1016/j.neurobiolaging.2010.07.010. (Visited on 2022) (cit. on p. 93).

Kreiter, Andreas K. and Wolf Singer (1996). "- On the Role of Neural Synchrony in the Primate Visual Cortex." In: *Brain Theory*. Ed. by A. Aertsen and V. Braitenberg. Amsterdam: Elsevier Science B.V., pp. 201–227. DOI: 10.1016/B978-044482046-4/50053-5. (Visited on 2022) (cit. on p. 19).

Kriegeskorte, Nikolaus and Pamela K. Douglas (2018). "Cognitive Computational Neuroscience." In: *Nature Neuroscience* 21.9, pp. 1148–1160. DOI: 10.1038/s41593-018-0210-5. (Visited on 2022) (cit. on p. 35).

Krzanowski, W. J. (1990). "Between-Group Analysis with Heterogeneous Covariance Matrices: The Common Principal Component Model." In: *Journal of Classification* 7.1, pp. 81–98. DOI: 10.1007/BF01889705. (Visited on 2022) (cit. on p. 63).

Kuczala, Alexander and Tatyana O. Sharpee (2016). "Eigenvalue Spectra of Large Correlated Random Matrices." In: *Physical Review E* 94.5, p. 050101. DOI: 10.1103/PhysRevE.94.050101. (Visited on 2019) (cit. on p. 180).

Kuhn, Thomas S. (1961). "The Function of Measurement in Modern Physical Science." In: *Isis* 52.2, pp. 161–193. DOI: 10.1086/349468. (Visited on 2022) (cit. on p. 10).

– (1962). *The Structure Of Scientific Revolutions*. University of Chicago Press. (Visited on 2022) (cit. on p. 187).

Kullback, S. and R. A. Leibler (1951). "On Information and Sufficiency." In: *The Annals of Mathematical Statistics* 22.1, pp. 79–86. DOI: 10.1214/aoms/1177729694. JSTOR: 2236703. (Visited on 2022) (cit. on p. 62).

Kuramoto, Yoshiki (1975). "Self-Entrainment of a Population of Coupled Non-Linear Oscillators." In: *International Symposium on Mathematical Problems in Theoretical Physics*. Ed. by Huzihiro Araki. Lecture Notes in Physics. Berlin, Heidelberg: Springer, pp. 420–422. DOI: 10.1007/BFb0013365 (cit. on p. 12).

Kweon, Hyeokmoon, Gökhan Aydogan, Alain Dagher, Danilo Bzdok, Christian C. Ruff, Gideon Nave, Martha J. Farah, and Philipp D. Koellinger (2022). "Human Brain Anatomy Reflects Separable Genetic and Environmental Components of Socioeconomic Status." In: *Science Advances* 8.20, eabm2923. DOI: 10.1126/sciadv.abm2923 (cit. on p. 6).

Ladenbauer, Josef, Sam McKenzie, Daniel Fine English, Olivier Hagens, and Srdjan Ostojic (2019). "Inferring and Validating Mechanistic Models of Neural Microcircuits Based on Spike-Train Data." In: *Nature Communications* 10.1, p. 4933. DOI: 10.1038/s41467-019-12572-0. (Visited on 2022) (cit. on p. 89).

Lakatos, Peter, George Karmos, Ashesh D. Mehta, Istvan Ulbert, and Charles E. Schroeder (2008). "Entrainment of Neuronal Oscillations as a Mechanism of Attentional Selection." In: *Science* 320.5872, pp. 110–113. DOI: 10.1126/science.1154735. (Visited on 2022) (cit. on pp. 183, 184, 186).

Lam, Ying-Wan, Lawrence B. Cohen, Matt Wachowiak, and Michal R. Zochowski (2000). "Odors Elicit Three Different Oscillations in the Turtle Olfactory Bulb." In: *Journal of Neuroscience* 20.2, pp. 749–762. DOI: 10.1523/JNEUROSCI.20-02-00749.2000. (Visited on 2022) (cit. on p. 185).

Lamprecht, Anna-Lena et al. (2020). "Towards FAIR Principles for Research Software." In: *Data Science* 3.1, pp. 37–59. DOI: 10.3233/DS-190026. (Visited on 2022) (cit. on p. 173).

Landau, Ayelet Nina, Helene Marianne Schreyer, Stan van Pelt, and Pascal Fries (2015). "Distributed Attention Is Implemented through Theta-Rhythmic Gamma Modulation." In: *Current Biology* 25.17, pp. 2332–2337. DOI: 10.1016/j.cub.2015.07.048. (Visited on 2022) (cit. on p. 183).

Layer, Moritz, Johanna Senk, Simon Essink, Alexander van Meegen, Hannah Bos, and Moritz Helias (2022). "NNMT: Mean-Field Based Analysis Tools for Neuronal Network Models." In: *Frontiers in Neu-*

roinformatics 16, p. 835657. DOI: 10.3389/fninf.2022.835657. (Visited on 2022) (cit. on pp. 16, 109).

LeCun, Yann, D Touresky, G Hinton, and T Sejnowski (1988). "A Theoretical Framework for Back-Propagation." In: *Proceedings of the 1988 Connectionist Models Summer School*. Vol. 1, pp. 21–28 (cit. on p. 19).

Lee, Daeyeol and Joseph G. Malpeli (1998). "Effects of Saccades on the Activity of Neurons in the Cat Lateral Geniculate Nucleus." In: *Journal of Neurophysiology* 79.2, pp. 922–936. DOI: 10.1152/jn.1998.79.2.922. (Visited on 2022) (cit. on p. 186).

Liang, Yuqi, Chenchen Song, Mianxin Liu, Pulin Gong, Changsong Zhou, and Thomas Knöpfel (2021). "Cortex-Wide Dynamics of Intrinsic Electrical Activities: Propagating Waves and Their Interactions." In: *Journal of Neuroscience* 41.16, pp. 3665–3678. DOI: 10.1523/JNEUROSCI.0623-20.2021. (Visited on 2021) (cit. on pp. 21, 115, 137).

Lindman, H. R. (1992). *Analysis of Variance in Experimental Design*. New York: Springer (cit. on p. 105).

Litwin-Kumar, Ashok and Brent Doiron (2012). "Slow Dynamics and High Variability in Balanced Cortical Networks with Clustered Connections." In: *Nature Neuroscience* 15.11, pp. 1498–1505. DOI: 10.1038/nn.3220. (Visited on 2022) (cit. on pp. 16, 85, 166, 180, 181).

Lupo, Cosimo et al. (2022). "Towards an EBRAINS Service for Brain Wave Analysis: Cobrawap." In: *School of Brain Cells & Circuits "Camillo Golgi"* (cit. on p. 174).

MacLennan, Bruce J. (2004). "Natural Computation and Non-Turing Models of Computation." In: *Theoretical Computer Science*. Super-Recursive Algorithms and Hypercomputation 317.1, pp. 115–145. DOI: 10.1016/j.tcs.2003.12.008. (Visited on 2021) (cit. on p. 18).

Manzoni, T. (1998). "The Cerebral Ventricles, the Animal Spirits and the Dawn of Brain Localization of Function." In: *Archives Italiennes De Biologie* 136.2, pp. 103–152 (cit. on p. 7).

Marčenko, V A and L A Pastur (1967). "Distribution of Eigenvalues for Some Sets of Random Matrices." In: *Mathematics of the USSR-Sbornik* 1.4, pp. 457–483. DOI: 10.1070/SM1967v001n04ABEH001994. (Visited on 2021) (cit. on p. 67).

Marder, Eve and Adam L. Taylor (2011). "Multiple Models to Capture the Variability in Biological Neurons and Networks." In: *Nature Neuroscience* 14.2, pp. 133–138. DOI: 10.1038/nn.2735. (Visited on 2019) (cit. on p. 15).

Marino, Simeone, Ian B. Hogue, Christian J. Ray, and Denise E. Kirschner (2008). "A Methodology for Performing Global Uncertainty and Sensitivity Analysis in Systems Biology." In: *Journal of Theoretical Biology* 254.1, pp. 178–196. DOI: 10.1016/j.jtbi.2008.04.011. (Visited on 2022) (cit. on p. 16).

Markov, Nikola T. et al. (2012). "A Weighted and Directed Interareal Connectivity Matrix for Macaque Cerebral Cortex." In: *Cerebral Cortex* 24.1, pp. 17–36. DOI: 10.1093/cercor/bhs270 (cit. on p. 88).

Markram, Henry et al. (2015). "Reconstruction and Simulation of Neocortical Microcircuitry." In: *Cell* 163.2, pp. 456–492. DOI: 10.1016/J.CELL.2015.09.029 (cit. on pp. 12, 15, 34).

Martis, Morvin Savio (2006). "Validation of Simulation Based Models: A Theoretical Outlook." In: *The Electronic Journal of Business Research Methods* 4.1, pp. 39–46 (cit. on pp. 16, 37).

Mascaro, Massimo and Daniel J Amit (1999). "Effective Neural Response Function for Collective Population States." In: *Network: Computation in Neural Systems* 10.4, pp. 351–373. DOI: 10.1088/0954-898X_10_4_305. (Visited on 2022) (cit. on p. 141).

Massimini, M., G. Tononi, and R. Huber (2009). "Slow Waves, Synaptic Plasticity and Information Processing: Insights from Transcranial Magnetic Stimulation and High-Density EEG Experiments." In: *European Journal of Neuroscience* 29.9, pp. 1761–1770. DOI: 10.1111/j.1460-9568.2009.06720.x. (Visited on 2022) (cit. on p. 20).

Massimini, Marcello, Fabio Ferrarelli, Simone Sarasso, and Giulio Tononi (2012). "Cortical Mechanisms of Loss of Consciousness: Insight from TMS/EEG Studies." In: *Archives Italiennes de Biologie* 150.2/3, pp. 44–55. DOI: 10.4449/aib.v150i2.1361. (Visited on 2022) (cit. on p. 20).

Massimini, Marcello, Reto Huber, Fabio Ferrarelli, Sean Hill, and Giulio Tononi (2004). "The Sleep Slow Oscillation as a Traveling Wave." In: *J. Neurosci.* 6.4, pp. 1160–1170. DOI: 10.1523/JNEUROSCI.1318-04.2004 (cit. on pp. 113–115, 136, 146).

Mattia, Maurizio, Stefano Ferraina, and Paolo Del Giudice (2010). "Dissociated Multi-Unit Activity and Local Field Potentials: A Theory Inspired Analysis of a Motor Decision Task." In: *NeuroImage*. Computational Models of the Brain 52.3, pp. 812–823. DOI: 10.1016/j.neuroimage.2010.01.063. (Visited on 2020) (cit. on p. 121).

Mattia, Maurizio and Maria V. Sanchez-Vives (2012). "Exploring the Spectrum of Dynamical Regimes and Timescales in Spontaneous Cortical Activity." In: *Cognitive Neurodynamics* 6.3, pp. 239–250. DOI: 10.1007/s11571-011-9179-4. (Visited on 2020) (cit. on pp. 115, 141).

Mazurek, Mark E. and Michael N. Shadlen (2002). "Limits to the Temporal Fidelity of Cortical Spike Rate Signals." In: *Nature Neuroscience* 5.5, pp. 463–471. DOI: 10.1038/nn836. (Visited on 2022) (cit. on p. 25).

McDougal, Robert A., Anna S. Bulanova, and William W. Lytton (2016). "Reproducibility in Computational Neuroscience Models and Simulations." In: *IEEE Transactions on Biomedical Engineering* 63.10, pp. 2021–2035. DOI: 10.1109/TBME.2016.2539602 (cit. on p. 12).

McKinney, Wes (2010). "Data Structures for Statistical Computing in Python." In: *Proceedings of the 9th Python in Science Conference*, pp. 56–61. DOI: 10.25080/Majora-92bf1922-00a. (Visited on 2022) (cit. on pp. 115, 129, 132).

Mellon, DeForest and Christopher J. Wheeler (1999). "Coherent Oscillations in Membrane Potential Synchronize Impulse Bursts in Central Olfactory Neurons of the Crayfish." In: *Journal of Neurophysiology* 81.3, pp. 1231–1241. DOI: 10.1152/jn.1999.81.3.1231. (Visited on 2022) (cit. on p. 184).

Mendoza-Halliday, Diego, Alex James Major, Noah Lee, Maxwell Lichtenfeld, Brock Carlson, Blake Mitchell, Patrich D. Meng, Yihan (Sophy) Xiong, Jacob A. Westerberg, Alexander Maier, Robert Desimone, Earl K. Miller, and Andre M. Bastos (2022). *A Ubiquitous Spectrolaminar Motif of Local Field Potential Power across the Primate Cortex*. DOI: 10.1101/2022.09.30.510398. (Visited on 2022) (cit. on p. 161).

Michalareas, Georgios, Julien Vezoli, Stan van Pelt, Jan-Mathijs Schoffelen, Henry Kennedy, and Pascal Fries (2016). "Alpha-Beta and Gamma Rhythms Subserve Feedback and Feedforward Influences among Human Visual Cortical Areas." In: *Neuron* 89.2, pp. 384–397. DOI: 10.1016/j.neuron.2015.12.018. (Visited on 2022) (cit. on p. 183).

Michelson, Nicholas J. and Takashi D. Y. Kozai (2018). "Isoflurane and Ketamine Differentially Influence Spontaneous and Evoked Laminar Electrophysiology in Mouse V1." In: *Journal of Neurophysiology* 120.5, pp. 2232–2245. DOI: 10.1152/jn.00299.2018. (Visited on 2021) (cit. on p. 146).

Miłkowski, Marcin, Witold M. Hensel, and Mateusz Hohol (2018). "Replicability or Reproducibility? On the Replication Crisis in Computational Neuroscience and Sharing Only Relevant Detail." In: *Journal of Computational Neuroscience* 45.3, pp. 163–172. DOI: 10.1007/s10827-018-0702-z. (Visited on 2022) (cit. on p. 188).

Mochizuki, Y. et al. (2016). "Similarity in Neuronal Firing Regimes across Mammalian Species." In: *Journal of Neuroscience* 36.21, pp. 5736–5747. DOI: 10.1523/JNEUROSCI.0230-16.2016. (Visited on 2020) (cit. on pp. 59, 108).

Mohajerani, Majid H., Allen W. Chan, Mostafa Mohsenvand, Jeffrey LeDue, Rui Liu, David A. McVea, Jamie D. Boyd, Yu Tian Wang, Mark Reimers, and Timothy H. Murphy (2013). "Spontaneous Cortical Activity Alternates between Motifs Defined by Regional Axonal Projections." In: *Nature Neuroscience* 16.10, pp. 1426–1435. DOI: 10.1038/nn.3499. (Visited on 2020) (cit. on p. 146).

Mölder, Felix, Kim Philipp Jablonski, Brice Letcher, Michael B. Hall, Christopher H. Tomkins-Tinch, Vanessa Sochat, Jan Forster, Soohyun Lee, Sven O. Twardziok, Alexander Kanitz, Andreas Wilm, Manuel Holtgrewe, Sven Rahmann, Sven Nahnsen, and Johannes Köster

(2021). "Sustainable Data Analysis with Snakemake." In: *F1000Research* 10, p. 33. DOI: 10.12688/f1000research.29032.2. (Visited on 2021) (cit. on pp. 115, 117, 118, 131, 173).

Mongillo, Gianluigi, Simon Rumpel, and Yonatan Loewenstein (2018). "Inhibitory Connectivity Defines the Realm of Excitatory Plasticity." In: *Nature Neuroscience* 21.10, p. 1463. DOI: 10.1038/s41593-018-0226-x. (Visited on 2019) (cit. on pp. 80, 82–85, 180).

Morales-Gregorio, Aitor, Robin Gutzen, Paulina Dabrowska, Alper Yegenoglu, Sandra Diaz-Pier, Sarah Palmis, Alexendre René, Panos Sapountzis, Markus Diesmann, Sonja Grün, Johanna Senk, Georgia Gregoriou, Bjørg E. Kilavik, and Sacha J Van Albada (2022). "Activity-Driven Microconnectome Estimation of Macaque Visuomotor Cortices." In: *Preparation* (cit. on pp. 28, 87).

Muir, Dylan R. and Thomas Mrsic-Flogel (2015). "Eigenspectrum Bounds for Semirandom Matrices with Modular and Spatial Structure for Neural Networks." In: *Physical Review E* 91.4, p. 042808. DOI: 10.1103/PhysRevE.91.042808. (Visited on 2022) (cit. on p. 180).

Muller, Lyle, Frédéric Chavane, John Reynolds, and Terrence J. Sejnowski (2018). "Cortical Travelling Waves: Mechanisms and Computational Principles." In: *Nature Reviews Neuroscience* 19.5, pp. 255–268. DOI: 10.1038/nrn.2018.20 (cit. on pp. 21, 148).

Muller, Lyle, Giovanni Piantoni, Dominik Koller, Sydney S Cash, Eric Halgren, and Terrence J Sejnowski (2016). "Rotating Waves during Human Sleep Spindles Organize Global Patterns of Activity That Repeat Precisely through the Night." In: *eLife* 5. Ed. by Frances K Skinner, e17267. DOI: 10.7554/eLife.17267. (Visited on 2020) (cit. on pp. 5, 113–115, 146).

Muller, Lyle, Alexandre Reynaud, Frédéric Chavane, and Alain Destexhe (2014). "The Stimulus-Evoked Population Response in Visual Cortex of Awake Monkey Is a Propagating Wave." In: *Nature Communications* 5.1, p. 3675. DOI: 10.1038/ncomms4675. (Visited on 2022) (cit. on p. 171).

Murray-Smith, David J. (David J.) (2015). *Testing and Validation of Computer Simulation Models*. Simulation Foundations, Methods and Applications. Springer International Publishing. DOI: 10.1007/978-3-319-15099-4 (cit. on pp. 13, 14).

Murtagh, Fionn (1991). "Multilayer Perceptrons for Classification and Regression." In: *Neurocomputing* 2.5, pp. 183–197. DOI: 10.1016/0925-2312(91)90023-5. (Visited on 2022) (cit. on p. 19).

Musslick, Sebastian, Biswadip Dey, Mostofa Ali Patwary, Theodore L Willke, and Jonathan D Cohen (2016). "Controlled vs. Automatic Processing: A Graph-Theoretic Approach to the Analysis of Serial vs. Parallel Processing in Neural Network Architectures." In: *CogSci*, p. 7 (cit. on p. 147).

Nagendra Babu, Pooja, Charl Linssen, Jochen Martin Eppler, Tobias Schulte to Brinke, Abolfazl Ziaeemehr, Tanguy Fardet, Younes

Bouhadjar, Renato Duarte, Bernhard Rumpe, and Abigail Morrison (2021). *NESTML 4.0*. Zenodo. DOI: 10.5281/zenodo.4740083. (Visited on 2022) (cit. on p. 14).

Nakagawa, Shinichi and Innes C Cuthill (2007). "Effect Size, Confidence Interval and Statistical Significance: A Practical Guide for Biologists." In: *Biological Reviews* 82.4, pp. 591–605. DOI: 10.1111/j.1469-185X.2007.00027.x (cit. on p. 70).

Nakamura, Kae and Carol L. Colby (2000). "Visual, Saccade-Related, and Cognitive Activation of Single Neurons in Monkey Extrastriate Area V3A." In: *Journal of Neurophysiology* 84.2, pp. 677–692. DOI: 10.1152/jn.2000.84.2.677. (Visited on 2022) (cit. on p. 186).

Nauhaus, Ian, Laura Busse, Matteo Carandini, and Dario L Ringach (2009). "Stimulus Contrast Modulates Functional Connectivity in Visual Cortex." In: *Nature Neuroscience* 12.1, pp. 70–76. DOI: 10.1038/nn.2232. (Visited on 2022) (cit. on pp. 181, 183).

Naujokaitytė, Goda (2021). "Number of Scientists Worldwide Reaches 8.8M, as Global Research Spending Grows Faster than the Economy." In: *Science | Business*. (Visited on 2022) (cit. on p. 188).

Ness, Torbjørn V., Michiel W. H. Remme, and Gaute T. Einevoll (2016). "Active Subthreshold Dendritic Conductances Shape the Local Field Potential." In: *The Journal of Physiology* 594.13, pp. 3809–3825. DOI: 10.1113/JP272022. (Visited on 2022) (cit. on p. 181).

Nickerson, Raymond (2000). "Null Hypothesis Significance Testing: A Review of an Old and Continuing Controversy." In: *Psychological methods* 5.2, pp. 241–301. DOI: 10.1037/1082-989X.5.2.241 (cit. on p. 165).

Niebur, Ernst (2008). "Neuronal Cable Theory." In: *Scholarpedia* 3.5, p. 2674. DOI: 10.4249/scholarpedia.2674. (Visited on 2022) (cit. on p. 4).

Nikolenko, Volodymyr, Kira E Poskanzer, and Rafael Yuste (2007). "Two-Photon Photostimulation and Imaging of Neural Circuits." In: *Nature Methods* 4.11, pp. 943–950. DOI: 10.1038/nmeth1105. (Visited on 2022) (cit. on p. 89).

Nir, Yuval, Richard J. Staba, Thomas Andrillon, Vladyslav V. Vyazovskiy, Chiara Cirelli, Itzhak Fried, and Giulio Tononi (2011). "Regional Slow Waves and Spindles in Human Sleep." In: *Neuron* 70.1, pp. 153–169. DOI: 10.1016/j.neuron.2011.02.043. (Visited on 2020) (cit. on pp. 114, 136).

Noble, Denis (2008). *The Music of Life: Biology Beyond Genes*. OUP Oxford (cit. on p. 15).

Noel, Jean-Paul, Edoardo Balzani, Eric Avila, Kaushik J Lakshminarasimhan, Stefania Bruni, Panos Alefantis, Cristina Savin, and Dora E Angelaki (2022). "Coding of Latent Variables in Sensory, Parietal, and Frontal Cortices during Closed-Loop Virtual Navigation." In: *eLife* 11. Ed. by Kristine Krug, Joshua I Gold, and Marieke

Schölvinck, e80280. DOI: 10.7554/eLife.80280. (Visited on 2022) (cit. on p. 185).

Noguchi, Jun, Akira Nagaoka, Satoshi Watanabe, Graham C. R. Ellis-Davies, Kazuo Kitamura, Masanobu Kano, Masanori Matsuzaki, and Haruo Kasai (2011). "*In vivo* Two-Photon Uncaging of Glutamate Revealing the Structure-Function Relationships of Dendritic Spines in the Neocortex of Adult Mice: *In vivo* Two-Photon Glutamate Uncaging." In: *The Journal of Physiology* 589.10, pp. 2447–2457. DOI: 10.1113/jphysiol.2011.207100. (Visited on 2022) (cit. on p. 89).

Nordlie, Eilen, Marc-Oliver Gewaltig, and Hans Ekkehard Plesser (2009). "Towards Reproducible Descriptions of Neuronal Network Models." In: *PLoS Computational Biology* 5.8. Ed. by Karl J. Friston, e1000456. DOI: 10.1371/journal.pcbi.1000456. (Visited on 2022) (cit. on p. 12).

Nunes, João D., Marcelo Carvalho, Diogo Carneiro, and Jaime S. Cardoso (2022). "Spiking Neural Networks: A Survey." In: *IEEE Access* 10, pp. 60738–60764. DOI: 10.1109/ACCESS.2022.3179968 (cit. on p. 19).

O'Hara, K. and T. Scutt (1996). "There Is No Hard Problem of Consciousness." In: *Journal of Consciousness Studies* 3.4, pp. 290–302 (cit. on p. 188).

O'Keefe, J. and J. Dostrovsky (1971). "The Hippocampus as a Spatial Map: Preliminary Evidence from Unit Activity in the Freely-Moving Rat." In: *Brain Research* 34, pp. 171–175. DOI: 10.1016/0006-8993(71)90358-1 (cit. on p. 19).

Okun, Michael, Nicholas A. Steinmetz, Lee Cossell, M. Florencia Iacaruso, Ho Ko, Péter Barthó, Tirin Moore, Sonja B. Hofer, Thomas D. Mrsic-Flogel, Matteo Carandini, and Kenneth D. Harris (2015). "Diverse Coupling of Neurons to Populations in Sensory Cortex." In: *Nature* 521.7553, pp. 511–515. DOI: 10.1038/nature14273. (Visited on 2022) (cit. on p. 97).

Olcese, Umberto, Jeroen J. Bos, Martin Vinck, Jan V. Lankelma, Laura B. van Mourik-Donga, Friederike Schlumm, and Cyriel M. A. Pennartz (2016). "Spike-Based Functional Connectivity in Cerebral Cortex and Hippocampus: Loss of Global Connectivity Is Coupled to Preservation of Local Connectivity During Non-REM Sleep." In: *Journal of Neuroscience* 36.29, pp. 7676–7692. DOI: 10.1523/JNEUROSCI.4201-15.2016. (Visited on 2022) (cit. on pp. 61, 141).

Omar, Cyrus, Jonathan Aldrich, and Richard C. Gerkin (2014). "Collaborative Infrastructure for Test-Driven Scientific Model Validation." In: *Companion Proceedings of the 36th International Conference on Software Engineering*. Hyderabad India: ACM, pp. 524–527. DOI: 10.1145/2591062.2591129. (Visited on 2022) (cit. on pp. 35, 44, 59, 117, 173, 176).

Onnela, Jukka Pekka, Jari Saramäki, János Kertész, and Kimmo Kaski (2005). "Intensity and Coherence of Motifs in Weighted Complex

Networks." In: *Physical Review E* 71.6, p. 065103. DOI: 10.1103/PhysRevE.71.065103 (cit. on p. 76).

Oreskes, Naomi, Kristin Shrader-Frechette, and Kenneth Belitz (1994). "Verification, Validation, and Confirmation of Numerical Models in the Earth Sciences." In: *Science* 263.February, pp. 641–646. DOI: 10.1126/science.263.5147.641. JSTOR: pdf/2883078.pdf (cit. on p. 12).

Ostojic, S., N. Brunel, and V. Hakim (2009). "How Connectivity, Background Activity, and Synaptic Properties Shape the Cross-Correlation between Spike Trains." In: *Journal of Neuroscience* 29.33, pp. 10234–10253. DOI: 10.1523/JNEUROSCI.1275-09.2009. (Visited on 2022) (cit. on pp. 16, 61, 89, 180).

Ostojic, Srdjan (2014). "Two Types of Asynchronous Activity in Networks of Excitatory and Inhibitory Spiking Neurons." In: *Nature Neuroscience* 17.4, pp. 594–600. DOI: 10.1038/nn.3658. (Visited on 2022) (cit. on p. 94).

Palm, Günther (1990). "Cell Assemblies as a Guideline for Brain Research." In: *Concepts in Neuroscience* 1, pp. 133–147 (cit. on p. 5).

Paninski, Liam, Jonathan W. Pillow, and Eero P. Simoncelli (2004). "Maximum Likelihood Estimation of a Stochastic Integrate-and-Fire Neural Encoding Model." In: *Neural Computation* 16.12, pp. 2533–2561. DOI: 10.1162/0899766042321797. (Visited on 2022) (cit. on p. 89).

Papadimitriou, Christos H., Santosh S. Vempala, Daniel Mitropolsky, Michael Collins, and Wolfgang Maass (2020). "Brain Computation by Assemblies of Neurons." In: *Proceedings of the National Academy of Sciences* 117.25, pp. 14464–14472. DOI: 10.1073/pnas.2001893117. (Visited on 2022) (cit. on p. 18).

Parks, Randolph W., Debra L. Long, Daniel S. Levine, David J. Crockett, Edith G. McGeer, Patrick L. McGeer, Irene E. Dalton, Ronald F. Zec, Robert E. Becker, Kerry L. Coburn, Gil Siler, Mark E. Nelson, and James M. Bower (1991). "Parallel Distributed Processing and Neural Networks: Origins, Methodology and Cognitive Functions." In: *International Journal of Neuroscience* 60.3-4, pp. 195–214. DOI: 10.3109/00207459109167033. (Visited on 2022) (cit. on p. 147).

Pastore, Vito Paolo, Paolo Massobrio, Aleksandar Godjoski, and Sergio Martinoia (2018). "Identification of Excitatory-Inhibitory Links and Network Topology in Large-Scale Neuronal Assemblies from Multi-Electrode Recordings." In: *PLOS Computational Biology* 14.8. Ed. by Daniele Marinazzo, e1006381. DOI: 10.1371/journal.pcbi.1006381. (Visited on 2022) (cit. on p. 89).

Pastorelli, Elena, Cristiano Capone, Francesco Simula, Maria V. Sanchez-Vives, Paolo Del Giudice, Maurizio Mattia, and Pier Stanislao Paolucci (2019). "Scaling of a Large-Scale Simulation of Synchronous Slow-Wave and Asynchronous Awake-Like Activity of a Cortical Model With Long-Range Interconnections." In: *Frontiers in Systems Neuro-*

science 13, p. 33. DOI: 10.3389/fnsys.2019.00033. (Visited on 2020) (cit. on pp. 113, 114).

Patil, Prasad, Roger D. Peng, and Jeffrey T. Leek (2016). *A Statistical Definition for Reproducibility and Replicability*. Preprint. DOI: 10.1101/066803. (Visited on 2022) (cit. on p. 38).

Pauli, Robin, Philipp Weidel, Susanne Kunkel, and Abigail Morrison (2018). "Reproducing Polychronization: A Guide to Maximizing the Reproducibility of Spiking Network Models." In: *Frontiers in Neuroinformatics* 12, p. 46. DOI: 10.3389/fninf.2018.00046. (Visited on 2022) (cit. on pp. 47, 51, 175).

Pazienti, Antonio, Andrea Galluzzi, Miguel Dasilva, Maria V. Sanchez-Vives, and Maurizio Mattia (2022). "Slow Waves Form Expanding, Memory-Rich Mesostates Steered by Local Excitability in Fading Anesthesia." In: *iScience* 25.3, p. 103918. DOI: 10.1016/j.isci.2022.103918. (Visited on 2022) (cit. on pp. 20, 115, 132, 133, 136, 171).

Pearn, John (2002). "A Curious Experiment: The Paradigm Switch from Observation and Speculation to Experimentation, in the Understanding of Neuromuscular Function and Disease." In: *Neuromuscular Disorders* 12.6, pp. 600–607. DOI: 10.1016/S0960-8966(01)00310-8. (Visited on 2022) (cit. on p. 17).

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay (2011). "Scikit-Learn: Machine Learning in Python." In: *Journal of Machine Learning Research* 12, pp. 2825–2830 (cit. on p. 151).

Peel, Leto, Tiago P. Peixoto, and Manlio De Domenico (2022). "Statistical Inference Links Data and Theory in Network Science." In: *Nature Communications* 13.1, p. 6794. DOI: 10.1038/s41467-022-34267-9. (Visited on 2022) (cit. on p. 7).

Penrose, Roger (1999). *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics*. Oxford [u.a.]: Oxford Univ. Press. (Visited on 2022) (cit. on p. 3).

Perin, Rodrigo, Thomas K. Berger, and Henry Markram (2011). "A Synaptic Organizing Principle for Cortical Neuronal Groups." In: *Proceedings of the National Academy of Sciences* 108.13, pp. 5419–5424. DOI: 10.1073/pnas.1016051108. (Visited on 2022) (cit. on p. 89).

Perkel, Donald H. and Theodore H. Bullock (1968). "Neural Coding." In: *Neurosciences Research Program Bulletin* 6, pp. 221–348 (cit. on p. 19).

Perkel, Donald H., George L. Gerstein, and George P. Moore (1967). "Neuronal Spike Trains and Stochastic Point Processes." In: *Biophysical Journal* 7.4, pp. 419–440. DOI: 10.1016/S0006-3495(67)86597-4 (cit. on p. 41).

Pernice, Volker and Stefan Rotter (2013). "Reconstruction of Sparse Connectivity in Neural Networks from Spike Train Covariances." In: *Journal of Statistical Mechanics: Theory and Experiment* 2013.03, P03008.

DOI: 10.1088/1742-5468/2013/03/P03008. (Visited on 2022) (cit. on p. 178).

Pernice, Volker, Benjamin Staude, Stefano Cardanobile, and Stefan Rotter (2011). "How Structure Determines Correlations in Neuronal Networks." In: *PLOS Computational Biology* 7.5, e1002059. DOI: 10.1371/journal.pcbi.1002059. (Visited on 2019) (cit. on pp. 61, 62, 180).

Pesaran, Bijan, Martin Vinck, Gaute T Einevoll, Anton Sirota, Pascal Fries, Markus Siegel, Wilson Truccolo, Charles E Schroeder, and Ramesh Srinivasan (2018). "Investigating Large-Scale Brain Dynamics Using Field Potential Recordings: Analysis and Interpretation - Supplementary Information." In: 21.7, pp. 903–919. DOI: 10.1038/s41593-018-0171-8 (cit. on pp. 5, 20, 181).

Petersen, Carl C. H., Amiram Grinvald, and Bert Sakmann (2003). "Spatiotemporal Dynamics of Sensory Responses in Layer 2/3 of Rat Barrel Cortex Measured In Vivo by Voltage-Sensitive Dye Imaging Combined with Whole-Cell Voltage Recordings and Neuron Reconstructions." In: *Journal of Neuroscience* 23.4, pp. 1298–1309. DOI: 10.1523/JNEUROSCI.23-04-01298.2003. (Visited on 2021) (cit. on pp. 21, 177).

Peyrache, Adrien, Karim Benchenane, Mehdi Khamassi, Sidney I Wiener, and Francesco P Battaglia (2010). "Principal Component Analysis of Ensemble Recordings Reveals Cell Assemblies at High Temporal Resolution." In: *Journal of Computational Neuroscience* 29.1-2, pp. 309–325. DOI: 10.1007/s10827-009-0154-6 (cit. on p. 41).

Pillow, J. W. (2005). "Prediction and Decoding of Retinal Ganglion Cell Responses with a Probabilistic Spiking Model." In: *Journal of Neuroscience* 25.47, pp. 11003–11013. DOI: 10.1523/JNEUROSCI.3305-05.2005. (Visited on 2022) (cit. on p. 89).

Platt, Tanja, Mark E. Ladd, and Daniel Paech (2021). "7 Tesla and Beyond." In: *Investigative Radiology* 56.11, pp. 705–725. DOI: 10.1097/RLI.0000000000000820. (Visited on 2022) (cit. on p. 188).

Plesser, Hans E. (2018). "Reproducibility vs. Replicability: A Brief History of a Confused Terminology." In: *Frontiers in Neuroinformatics* 11, p. 76. DOI: 10.3389/fninf.2017.00076. (Visited on 2020) (cit. on p. 38).

Plotnikov, Dimitri, Bernhard Rumpe, Inga Blundell, Tammo Ippen, Jochen Martin Eppler, and Abgail Morrison (2016). *NESTML: A Modeling Language for Spiking Neurons*. DOI: 10.48550/arXiv.1606.02882. arXiv: 1606.02882 [cs]. (Visited on 2022) (cit. on p. 12).

Potjans, Tobias C. and Markus Diesmann (2014). "The Cell-Type Specific Cortical Microcircuit: Relating Structure and Activity in a Full-Scale Spiking Network Model." In: *Cerebral Cortex* 24.3, pp. 785–806. DOI: 10.1093/cercor/bhs358 (cit. on pp. 12, 15, 16, 35, 59, 89, 94, 95).

Prechtl, J. C., L. B. Cohen, B. Pesaran, P. P. Mitra, and D. Kleinfeld (1997). "Visual Stimuli Induce Waves of Electrical Activity in Turtle Cortex." In: *Proceedings of the National Academy of Sciences* 94.14, pp. 7621–7626. DOI: 10.1073/pnas.94.14.7621. (Visited on 2022) (cit. on p. 185).

Prinz, Astrid A, Dirk Bucher, and Eve Marder (2004). "Similar Network Activity from Disparate Circuit Parameters." In: *Nature Neuroscience* 7.12, pp. 1345–1352. DOI: 10.1038/nn1352. (Visited on 2022) (cit. on p. 89).

Purdon, Patrick L., Aaron Sampson, Kara J. Pavone, and Emery N. Brown (2015). "Clinical Electroencephalography for Anesthesiologists: Part I: Background and Basic Signatures." In: *Anesthesiology* 123.4, pp. 937–960. DOI: 10.1097/ALN.0000000000000841 (cit. on p. 146).

Purpura, Keith P., Steven F. Kalik, and Nicholas D. Schiff (2003). "Analysis of Perisaccadic Field Potentials in the Occipitotemporal Pathway During Active Vision." In: *Journal of Neurophysiology* 90.5, pp. 3455–3478. DOI: 10.1152/jn.00011.2003. (Visited on 2022) (cit. on pp. 186, 187).

Qi, Yang, Michael Breakspear, and Pulin Gong (2015). "Subdiffusive Dynamics of Bump Attractors: Mechanisms and Functional Roles." In: *Neural Computation* 27.2, pp. 255–280. DOI: 10.1162/NECO_a_00698. (Visited on 2022) (cit. on p. 21).

Quaglio, Pietro, Vahid Rostami, Emiliano Torre, and Sonja Grün (2018). "Methods for Identification of Spike Patterns in Massively Parallel Spike Trains." In: *Biological Cybernetics* 112.1-2, pp. 57–80. DOI: 10.1007/s00422-018-0755-0. (Visited on 2022) (cit. on p. 42).

Quaglio, Pietro, Alper Yegenoglu, Emiliano Torre, Dominik M. Endres, and Sonja Grün (2017). "Detection and Evaluation of Spatio-Temporal Spike Patterns in Massively Parallel Spike Train Data with SPADE." In: *Frontiers in Computational Neuroscience* 11, p. 41. DOI: 10.3389/fncom.2017.00041. (Visited on 2022) (cit. on pp. 19, 42, 56).

Quiroga, Quian Rodrigo. and Stefano. Panzeri (2013). *Principles of Neural Coding* (cit. on p. 17).

Rajan, Kanaka and L. F. Abbott (2006). "Eigenvalue Spectra of Random Matrices for Neural Networks." In: *Physical Review Letters* 97.18, p. 188104. DOI: 10.1103/PhysRevLett.97.188104. (Visited on 2019) (cit. on pp. 70, 72, 84, 180).

Rajkai, Csaba, Peter Lakatos, Chi-Ming Chen, Zsuzsa Pincze, Gyorgy Karmos, and Charles E. Schroeder (2008). "Transient Cortical Excitation at the Onset of Visual Fixation." In: *Cerebral Cortex* 18.1, pp. 200–209. DOI: 10.1093/cercor/bhm046. (Visited on 2022) (cit. on p. 186).

Ramaswamy, Srikanth et al. (2015). "The Neocortical Microcircuit Collaboration Portal: A Resource for Rat Somatosensory Cortex."

In: *Frontiers in Neural Circuits* 9. DOI: 10.3389/fncir.2015.00044. (Visited on 2022) (cit. on p. 13).

Ramezani, Hamideh and Ozgur B. Akan (2018). "Impacts of Spike Shape Variations on Synaptic Communication." In: *IEEE Transactions on NanoBioscience* 17.3, pp. 260–271. DOI: 10.1109/TNB.2018.2838056 (cit. on p. 18).

Rao, Rajesh P. N. and Dana H. Ballard (1999). "Predictive Coding in the Visual Cortex: A Functional Interpretation of Some Extra-Classical Receptive-Field Effects." In: *Nature Neuroscience* 2.1, pp. 79–87. DOI: 10.1038/4580 (cit. on p. 160).

Ray, Supratim and John H. R. Maunsell (2010). "Differences in Gamma Frequencies across Visual Cortex Restrict Their Possible Use in Computation." In: *Neuron* 67.5, pp. 885–896. DOI: 10.1016/j.neuron.2010.08.004. (Visited on 2022) (cit. on p. 185).

– (2011). "Different Origins of Gamma Rhythm and High-Gamma Activity in Macaque Visual Cortex." In: *PLOS Biology* 9.4, e1000610. DOI: 10.1371/journal.pbio.1000610. (Visited on 2022) (cit. on p. 181).

Reich, Daniel S, Jonathan D Victor, Bruce W Knight, Tsuyoshi Ozaki, and Ehud Kaplan (1997). "Response Variability and Timing Precision of Neuronal Spike Trains in Vivo." In: *Journal of neurophysiology* 77.5, pp. 2836–2841 (cit. on p. 25).

Reig, R., M. Mattia, A. Compte, C. Belmonte, and M. V. Sanchez-Vives (2009). "Temperature Modulation of Slow and Fast Cortical Rhythms." In: *Journal of Neurophysiology* 103.3, pp. 1253–1261. DOI: 10.1152/jn.00890.2009. (Visited on 2020) (cit. on pp. 113, 121).

Reimann, Michael W., James G. King, Eilif B. Muller, Srikanth Ramaswamy, and Henry Markram (2015). "An Algorithm to Predict the Connectome of Neural Microcircuits." In: *Frontiers in Computational Neuroscience* 9. DOI: 10.3389/fncom.2015.00120. (Visited on 2022) (cit. on p. 15).

Renart, Alfonso, Jaime de la Rocha, Peter Bartho, Liad Hollender, Néstor Parga, Alex Reyes, and Kenneth D. Harris (2010). "The Asynchronous State in Cortical Circuits." In: *Science* 327.5965, pp. 587–590. DOI: 10.1126/science.1179850. (Visited on 2022) (cit. on p. 16).

René, Alexandre, André Longtin, and Jakob H. Macke (2019). "Inference of a Mesoscopic Population Model from Population Spike Trains." In: DOI: 10.48550/ARXIV.1910.01618. (Visited on 2022) (cit. on p. 89).

Reppas, John B, W. Martin Usrey, and R. Clay Reid (2002). "Saccadic Eye Movements Modulate Visual Responses in the Lateral Geniculate Nucleus." In: *Neuron* 35.5, pp. 961–974. DOI: 10.1016/S0896-6273(02)00823-1. (Visited on 2022) (cit. on p. 186).

Resta, Francesco, Anna Letizia Allegra Mascaro, and Francesco Pavone (2020a). "Study of Slow Waves (SWs) Propagation through Wide-Field Calcium Imaging of the Right Cortical Hemisphere of GCaMP6f

Mice (v2) [Data Set].” In: *EBRAINS*. DOI: 10.25493/QFZK-FXS. (Visited on 2021) (cit. on pp. 126, 136).

Resta, Francesco, Anna Letizia Allegra Mascaro, and Francesco Pavone (2020b). “Wide-Field Calcium Imaging of the Right Cortical Hemisphere of GCaMP6f Mice at Different Anesthesia Levels [Data Set].” In: *EBRAINS*. DOI: 10.25493/XJR8-QCA. (Visited on 2021) (cit. on pp. 126, 136).

Rey, Hernan Gonzalo, Carlos Pedreira, and Rodrigo Quian Quiroga (2015). “Past, Present and Future of Spike Sorting Techniques.” In: *Brain Research Bulletin*. Advances in Electrophysiological Data Analysis 119, pp. 106–117. DOI: 10.1016/j.brainresbull.2015.04.007. (Visited on 2022) (cit. on p. 181).

Reyes, A. D. and E. E. Fetz (1993). “Effects of Transient Depolarizing Potentials on the Firing Rate of Cat Neocortical Neurons.” In: *Journal of Neurophysiology* 69.5, pp. 1673–1683. DOI: 10.1152/jn.1993.69.5.1673. (Visited on 2022) (cit. on p. 184).

Ribary, U, A A Ioannides, K D Singh, R Hasson, J P Bolton, F Lado, A Mogilner, and R Llinás (1991). “Magnetic Field Tomography of Coherent Thalamocortical 40-Hz Oscillations in Humans.” In: *Proceedings of the National Academy of Sciences* 88.24, pp. 11037–11041. DOI: 10.1073/pnas.88.24.11037. (Visited on 2022) (cit. on p. 185).

Riehle, Alexa, Thomas Brochier, Martin Nawrot, and Sonja Grün (2018). “Behavioral Context Determines Network State and Variability Dynamics in Monkey Motor Cortex.” In: *Frontiers in Neural Circuits* 12, p. 52. DOI: 10.3389/fncir.2018.00052. (Visited on 2022) (cit. on p. 59).

Roberts, James A., Leonardo L. Gollo, Romesh G. Abeysuriya, Gloria Roberts, Philip B. Mitchell, Mark W. Woolrich, and Michael Breakspear (2019). “Metastable Brain Waves.” In: *Nature Communications* 10.1, pp. 1–17. DOI: 10.1038/s41467-019-08999-0. (Visited on 2020) (cit. on pp. 21, 115).

Roelfsema, Pieter R., Andreas K. Engel, Peter König, and Wolf Singer (1997). “Visuomotor Integration Is Associated with Zero Time-Lag Synchronization among Cortical Areas.” In: *Nature* 385.6612, pp. 157–161. DOI: 10.1038/385157a0. (Visited on 2022) (cit. on p. 185).

Roland, Per E., Akitoshi Hanazawa, Calle Undeman, David Eriksson, Tamas Tompa, Hiroyuki Nakamura, Sonata Valentiniene, and Bashir Ahmed (2006). “Cortical Feedback Depolarization Waves: A Mechanism of Top-down Influence on Early Visual Areas.” In: *Proceedings of the National Academy of Sciences* 103.33, pp. 12586–12591. DOI: 10.1073/pnas.0604925103. (Visited on 2022) (cit. on p. 184).

Rose, F. Clifford (2009). “Cerebral Localization in Antiquity.” In: *Journal of the History of the Neurosciences* 18.3, pp. 239–247. DOI: 10.1080/09647040802025052. (Visited on 2022) (cit. on p. 7).

Rosenblatt, F. (1958). "The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain." In: *Psychological Review* 65, pp. 386–408. DOI: 10.1037/h0042519 (cit. on p. 19).

Rossant, Cyrille (2010). "Automatic Fitting of Spiking Neuron Models to Electrophysiological Recordings." In: *Frontiers in Neuroinformatics* 4, p. 2. DOI: 10.3389/neuro.11.002.2010. (Visited on 2022) (cit. on pp. 89, 108).

Rübel, Oliver, Andrew Tritt, Ryan Ly, Benjamin K. Dichter, Satrajit Ghosh, Lawrence Niu, Ivan Soltesz, Karel Svoboda, Loren Frank, and Kristofer E. Bouchard (2021). "The Neurodata Without Borders Ecosystem for Neurophysiological Data Science." In: *bioRxiv*, p. 2021.03.13.435173. DOI: 10.1101/2021.03.13.435173. (Visited on 2021) (cit. on p. 117).

Rubino, Doug, Kay A. Robbins, and Nicholas G. Hatsopoulos (2006). "Propagating Waves Mediate Information Transfer in the Motor Cortex." In: *Nature Neuroscience* 9.12, pp. 1549–1557. DOI: 10.1038/nn1802. (Visited on 2019) (cit. on pp. 21, 148, 182, 185).

Rubinstein, Isaak (1990). *Electro-Diffusion of Ions*. SIAM (cit. on p. 3).

Ruiz-Mejias, Marcel, Laura Ciria-Suarez, Maurizio Mattia, and Maria V Sanchez-Vives (2011). "Slow and Fast Rhythms Generated in the Cerebral Cortex of the Anesthetized Mouse." In: *Journal of Neurophysiology* 106.6, pp. 2910–2921. DOI: 10.1152/jn.00440.2011 (cit. on pp. 115, 129, 132, 136, 146, 151).

Ruiz-Mejias, Marcel, Maria Martinez de Lagran, Maurizio Mattia, Patricia Castano-Prat, Lorena Perez-Mendez, Laura Ciria-Suarez, Thomas Gener, Belen Sancristobal, Jordi García-Ojalvo, Agnès Gruart, José M. Delgado-García, Maria V. Sanchez-Vives, and Mara Dierssen (2016). "Overexpression of Dyrk1A, a Down Syndrome Candidate, Decreases Excitability and Impairs Gamma Oscillations in the Prefrontal Cortex." In: *Journal of Neuroscience* 36.13, pp. 3648–3659. DOI: 10.1523/JNEUROSCI.2517-15.2016. (Visited on 2020) (cit. on p. 115).

Rumelhart, David E, Geoffrey E Hinton, and Ronald J Williams (1985). *Learning Internal Representations by Error Propagation*. Tech. rep. California Univ San Diego La Jolla Inst for Cognitive Science (cit. on p. 19).

Russell, Bertrand (2009). *An Outline of Philosophy*. London: Routledge. DOI: 10.4324/9780203875469 (cit. on p. 9).

Sakurai, Yoshio and Susumu Takahashi (2006). "Dynamic Synchrony of Firing in the Monkey Prefrontal Cortex during Working-Memory Tasks." In: *Journal of Neuroscience* 26.40, pp. 10141–10153. DOI: 10.1523/JNEUROSCI.2423-06.2006. (Visited on 2022) (cit. on p. 5).

Saltelli, Andrea (2002). "Sensitivity Analysis for Importance Assessment." In: *Risk Analysis* 22.3, pp. 579–590. DOI: 10.1111/0272-4332.00040. (Visited on 2022) (cit. on p. 16).

Sanchez-Vives, M. (2020). "Propagation Modes of Slow Waves in Mouse Cortex." In: *Human Brain Project Neuroinformatics Platform*. DOI: 10.25493/WKA8-Q4T (cit. on p. 124).

Sanchez-Vives, Maria V. (2019a). "Cortical Activity Features in Transgenic Mouse Models of Cognitive Deficits (Fragile X Syndrome) [Data Set]." In: *Human Brain Project Neuroinformatics Platform*. DOI: 10.25493/ANF9-EG3. (Visited on 2021) (cit. on pp. 124, 133).

– (2019b). "Cortical Activity Features in Transgenic Mouse Models of Cognitive Deficits (Williams Beuren Syndrome) [Data Set]." In: *Human Brain Project Neuroinformatics Platform*. DOI: DOI:10.25493/DZWT-1T8. (Visited on 2021) (cit. on pp. 124, 133).

Sanchez-Vives, Maria V., Marcello Massimini, and Maurizio Mattia (2017). "Shaping the Default Activity Pattern of the Cortical Network." In: *Neuron* 94.5, pp. 993–1001. DOI: 10.1016/j.neuron.2017.05.015. (Visited on 2020) (cit. on pp. 20, 171).

Sanchez-Vives, Maria V. and David A. McCormick (2000). "Cellular and Network Mechanisms of Rhythmic Recurrent Activity in Neocortex." In: *Nature Neuroscience* 3.10, pp. 1027–1034. DOI: 10.1038/79848. (Visited on 2021) (cit. on pp. 20, 112, 146).

Sanz Leon, Paula, Stuart A. Knock, M. Marmaduke Woodman, Lia Domide, Jochen Mersmann, Anthony R. McIntosh, and Viktor Jirsa (2013). "The Virtual Brain: A Simulator of Primate Brain Network Dynamics." In: *Frontiers in Neuroinformatics* 7. DOI: 10.3389/fninf.2013.00010. (Visited on 2022) (cit. on p. 12).

Sargent, R G, Jack Kleijnen, and R G Sargent (2013). "Verification and Validation of Simulation Models." In: *Journal of Simulation* 7.1, pp. 12–24. DOI: 10.1057/jos.2012.20 (cit. on p. 14).

Sarma, Gopal P., Travis W. Jacobs, Mark D. Watts, S. Vahid Ghayoomie, Stephen D. Larson, and Richard C. Gerkin (2016). "Unit Testing, Model Validation, and Biological Simulation." In: *F1000Research* 5, p. 1946. DOI: 10.12688/f1000research.9315.1. (Visited on 2022) (cit. on p. 35).

Sarma, Gopal P. et al. (2018). "OpenWorm: Overview and Recent Advances in Integrative Biological Simulation of Caenorhabditis Elegans." In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 373.1758, p. 20170382. DOI: 10.1098/rstb.2017.0382. (Visited on 2022) (cit. on p. 23).

Saroka, Kevin S., David E. Vares, and Michael A. Persinger (2016). "Similar Spectral Power Densities Within the Schumann Resonance and a Large Population of Quantitative Electroencephalographic Profiles: Supportive Evidence for Koenig and Pobachenko." In: *PLOS ONE* 11.1, e0146595. DOI: 10.1371/journal.pone.0146595. (Visited on 2022) (cit. on p. 6).

Schemmel, Johannes, Laura Kriener, Paul Müller, and Karlheinz Meier (2017). "An Accelerated Analog Neuromorphic Hardware System Emulating NMDA- and Calcium-Based Non-Linear Dendrites."

In: *2017 International Joint Conference on Neural Networks (IJCNN)*, pp. 2217–2226. DOI: 10.1109/IJCNN.2017.7966124 (cit. on p. 8).

Schlesinger, S (1979). "Terminology for Model Credibility." In: *Simulation* 32.3, pp. 103–104. DOI: 10.1177/003754977903200304 (cit. on pp. 11, 13, 14, 36).

Schmidt, Maximilian, Rembrandt Bakker, Claus C. Hilgetag, Markus Diesmann, and Sacha J. van Albada (2018). "Multi-Scale Account of the Network Structure of Macaque Visual Cortex." In: *Brain Structure and Function* 223.3, pp. 1409–1435. DOI: 10.1007/s00429-017-1554-4. (Visited on 2022) (cit. on pp. 12, 15).

Schneider, Marius, Benjamin Dann, Swathi Sheshadri, Hansjörg Scherberger, and Martin Vinck (2020). "A General Theory of Coherence between Brain Areas." In: *bioRxiv*, p. 2020.06.17.156190. DOI: 10.1101/2020.06.17.156190. (Visited on 2020) (cit. on pp. 160, 183).

Schnepel, Philipp, Arvind Kumar, Mihael Zohar, Ad Aertsen, and Clemens Boucsein (2015). "Physiology and Impact of Horizontal Connections in Rat Neocortex." In: *Cerebral Cortex* 25.10, pp. 3818–3835. DOI: 10.1093/cercor/bhu265. (Visited on 2022) (cit. on p. 141).

Schonhaut, Daniel R., Ashwin G. Ramayya, Ethan A. Solomon, Nora A. Herweg, Itzhak Fried, and Michael J. Kahana (2020). "Single Neurons throughout Human Memory Regions Phase-Lock to Hippocampal Theta." In: *bioRxiv*, p. 2020.06.30.180174. DOI: 10.1101/2020.06.30.180174. (Visited on 2020) (cit. on p. 20).

Schroeder, Charles E. and Peter Lakatos (2009). "Low-Frequency Neuronal Oscillations as Instruments of Sensory Selection." In: *Trends in Neurosciences* 32.1, pp. 9–18. DOI: 10.1016/j.tins.2008.09.012. (Visited on 2022) (cit. on pp. 183, 184).

Schroeder, Charles E, Donald A Wilson, Thomas Radman, Helen Scharfman, and Peter Lakatos (2010). "Dynamics of Active Sensing and Perceptual Selection." In: *Current Opinion in Neurobiology*. Cognitive Neuroscience 20.2, pp. 172–176. DOI: 10.1016/j.conb.2010.02.010. (Visited on 2022) (cit. on pp. 148, 186).

Schuecker, Jannis, Markus Diesmann, and Moritz Helias (2015). "Modulated Escape from a Metastable State Driven by Colored Noise." In: *Physical Review E* 92.5, p. 052119. DOI: 10.1103/PhysRevE.92.052119. (Visited on 2022) (cit. on p. 16).

Schuessler, Friedrich, Alexis Dubreuil, Francesca Mastrogiuseppe, Srdjan Ostojic, and Omri Barak (2020). "Dynamics of Random Recurrent Networks with Correlated Low-Rank Structure." In: *Physical Review Research* 2.1, p. 013111. DOI: 10.1103/PhysRevResearch.2.013111. (Visited on 2022) (cit. on p. 180).

Scott, David W. (2015). *Multivariate Density Estimation: Theory, Practice, and Visualization*. John Wiley & Sons (cit. on p. 123).

Sejnowski, Terrence J., Patricia S. Churchland, and J. Anthony Movshon (2014). "Putting Big Data to Good Use in Neuroscience." In: *Nature*

Neuroscience 17.11, pp. 1440–1441. DOI: 10.1038/nn.3839. (Visited on 2022) (cit. on pp. 10, 24).

Semedo, João D., Amin Zandvakili, Christian K. Machens, Byron M. Yu, and Adam Kohn (2019). "Cortical Areas Interact through a Communication Subspace." In: *Neuron* 102.1, 249–259.e4. DOI: 10.1016/j.neuron.2019.01.026. (Visited on 2022) (cit. on p. 19).

Senk, Johanna et al. (2017). "A Collaborative Simulation-Analysis Workflow for Computational Neuroscience Using HPC." In: *High-Performance Scientific Computing*. Ed. by Edoardo Di Napoli, Marc-André Hermanns, Hristo Iliev, Andreas Lintermann, and Alexander Peyser. Vol. 10164. Cham: Springer International Publishing, pp. 243–256. DOI: 10.1007/978-3-319-53862-4_21. (Visited on 2022) (cit. on p. 35).

Senseman, David M. and Kay A. Robbins (2002). "High-Speed VSD Imaging of Visually Evoked Cortical Waves: Decomposition Into Intra- and Intercortical Wave Motions." In: *Journal of Neurophysiology* 87.3, pp. 1499–1514. DOI: 10.1152/jn.00475.2001. (Visited on 2018) (cit. on pp. 21, 177).

Seung, H S and H Sompolinsky (1993). "Simple Models for Reading Neuronal Population Codes." In: *Proceedings of the National Academy of Sciences* 90.22, pp. 10749–10753. DOI: 10.1073/pnas.90.22.10749. (Visited on 2022) (cit. on p. 19).

Shackman, Alexander J., Brenton W. McMenamin, Jeffrey S. Maxwell, Lawrence L. Greischar, and Richard J. Davidson (2010). "Identifying Robust and Sensitive Frequency Bands for Interrogating Neural Oscillations." In: *NeuroImage* 51.4, pp. 1319–1333. DOI: 10.1016/j.neuroimage.2010.03.037. (Visited on 2022) (cit. on p. 152).

Shadlen, Michael N. and William T. Newsome (1998). "The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding." In: *The Journal of Neuroscience* 18.10, pp. 3870–3896. DOI: 10.1523/JNEUROSCI.18-10-03870.1998. (Visited on 2022) (cit. on p. 16).

Sharratt, Michael (1996). *Galileo: Decisive Innovator*. Cambridge University Press (cit. on p. 9).

Sheroziya, Maxim and Igor Timofeev (2014). "Global Intracellular Slow-Wave Dynamics of the Thalamocortical System." In: *Journal of Neuroscience* 34.26, pp. 8875–8893. DOI: 10.1523/JNEUROSCI.4460-13.2014. (Visited on 2022) (cit. on p. 136).

Sherry, David (2011). "Thermoscopes, Thermometers, and the Foundations of Measurement." In: *Studies in History and Philosophy of Science Part A* 42.4, pp. 509–524. DOI: 10.1016/j.shpsa.2011.07.001. (Visited on 2022) (cit. on p. 9).

Shimaoka, Daisuke, Chenchen Song, and Thomas Knöpfel (2017). "State-Dependent Modulation of Slow Wave Motifs towards Awakening." In: *Frontiers in Cellular Neuroscience* 11, p. 108. (Visited on 2022) (cit. on pp. 112, 174).

Shine, James M., Eli J. Müller, Brandon Munn, Joana Cabral, Rosalyn J. Moran, and Michael Breakspear (2021). "Computational Models Link Cellular Mechanisms of Neuromodulation to Large-Scale Neural Dynamics." In: *Nature Neuroscience* 24.6, pp. 765–776. DOI: 10.1038/s41593-021-00824-6. (Visited on 2022) (cit. on p. 185).

Shinomoto, Shigeru, Keisetsu Shima, and Jun Tanji (2003). "Differences in Spiking Patterns Among Cortical Neurons." In: *Neural Computation* 15.12, pp. 2823–2842. DOI: 10.1162/089976603322518759 (cit. on pp. 40, 96, 97).

Shoham, Shy, Daniel H. O'Connor, and Ronen Segev (2006). "How Silent Is the Brain: Is There a "Dark Matter" Problem in Neuroscience?" In: *Journal of Comparative Physiology A* 192.8, pp. 777–784. DOI: 10.1007/s00359-006-0117-6. (Visited on 2022) (cit. on pp. 11, 184).

Siegle, Joshua H. et al. (2021a). "Reconciling Functional Differences in Populations of Neurons Recorded with Two-Photon Imaging and Electrophysiology." In: *eLife* 10. Ed. by Ronald L Calabrese and Marius Pachitariu, e69068. DOI: 10.7554/eLife.69068. (Visited on 2022) (cit. on pp. 145, 182).

Siegle, Joshua H. et al. (2021b). "Survey of Spiking in the Mouse Visual System Reveals Functional Hierarchy." In: *Nature* 592.7852, pp. 86–92. DOI: 10.1038/s41586-020-03171-x. (Visited on 2022) (cit. on pp. 145, 182).

Singer, Wolf (1999). "Neuronal Synchrony: A Versatile Code for the Definition of Relations?" In: *Neuron* 24.1, pp. 49–65. DOI: 10.1016/S0896-6273(00)80821-1. (Visited on 2022) (cit. on p. 5).

Skrandies, W and K Laschke (1997). "Topography of Visually Evoked Brain Activity during Eye Movements: Lambda Waves, Saccadic Suppression, and Discrimination Performance." In: *International Journal of Psychophysiology* 27.1, pp. 15–27. DOI: 10.1016/S0167-8760(97)00749-6. (Visited on 2022) (cit. on p. 186).

Sobotka, Stanislaw and James L Ringo (1997). "Saccadic Eye Movements, Even in Darkness, Generate Event-Related Potentials Recorded in Medial Septum and Medial Temporal Cortex." In: *Brain Research* 756.1, pp. 168–173. DOI: 10.1016/S0006-8993(97)00145-5. (Visited on 2022) (cit. on pp. 186, 187).

Sommers, H. J., A. Crisanti, H. Sompolinsky, and Y. Stein (1988). "Spectrum of Large Random Asymmetric Matrices." In: *Physical Review Letters* 60.19, pp. 1895–1898. DOI: 10.1103/PhysRevLett.60.1895. (Visited on 2019) (cit. on pp. 71, 180).

Sompolinsky, H., A. Crisanti, and H. J. Sommers (1988). "Chaos in Random Neural Networks." In: *Physical Review Letters* 61.3, pp. 259–262. DOI: 10.1103/PhysRevLett.61.259. (Visited on 2022) (cit. on p. 180).

Sompolinsky, H. and M. Tsodyks (1994). "Segmentation by a Network of Oscillators with Stored Memories." In: *Neural Computation* 6.4,

pp. 642–657. DOI: 10.1162/neco.1994.6.4.642. (Visited on 2022) (cit. on p. 185).

Song, Sen, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B Chklovskii (2005). "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits." In: *PLoS Biology* 3.3. Ed. by Karl J. Friston, e68. DOI: 10.1371/journal.pbio.0030068. (Visited on 2022) (cit. on p. 89).

Speed, Anderson, Joseph Del Rosario, Navid Mikail, and Bilal Haider (2020). "Spatial Attention Enhances Network, Cellular and Subthreshold Responses in Mouse Visual Cortex." In: *Nature Communications* 11.1, p. 505. DOI: 10.1038/s41467-020-14355-4. (Visited on 2022) (cit. on p. 181).

Sporns, Olaf and Giulio Tononi (2001). "Classes of Network Connectivity and Dynamics." In: *Complexity* 7.1, pp. 28–38. DOI: 10.1002/cplx.10015. (Visited on 2022) (cit. on pp. 61, 178).

Sprenger, Julia, Lyuba Zehl, Jana Pick, Michael Sonntag, Jan Grewe, Thomas Wachtler, Sonja Grün, and Michael Denker (2019). "odMLtables: A User-Friendly Approach for Managing Metadata of Neurophysiological Experiments." In: *Frontiers in Neuroinformatics* 13, p. 62. DOI: 10.3389/fninf.2019.00062. (Visited on 2021) (cit. on p. 117).

Staude, Benjamin, Stefan Rotter, and Sonja Grün (2010). "CuBIC: Cumulant Based Inference of Higher-Order Correlations in Massively Parallel Spike Trains." In: *Journal of Computational Neuroscience* 29.1-2, pp. 327–350. DOI: 10.1007/s10827-009-0195-x (cit. on p. 74).

Steinmetz, Nicholas A, Christof Koch, Kenneth D Harris, and Matteo Carandini (2018). "Challenges and Opportunities for Large-Scale Electrophysiology with Neuropixels Probes." In: *Current Opinion in Neurobiology*. Neurotechnologies 50, pp. 92–100. DOI: 10.1016/j.conb.2018.01.009. (Visited on 2022) (cit. on p. 188).

Stella, Alessandra, Peter Bouss, Günther Palm, and Sonja Grün (2022). "Comparing Surrogates to Evaluate Precisely Timed Higher-Order Spike Correlations." In: *eNeuro* 9.3. DOI: 10.1523/ENEURO.0505-21.2022. (Visited on 2022) (cit. on pp. 20, 86, 180).

Steriade, M., P. Gloor, R. R. Llinás, F. H. Lopes da Silva, and M. M. Mesulam (1990). "Basic Mechanisms of Cerebral Rhythmic Activities." In: *Electroencephalography and Clinical Neurophysiology* 76.6, pp. 481–508. DOI: 10.1016/0013-4694(90)90001-Z. (Visited on 2022) (cit. on p. 152).

Steriade, M., A. Nunez, and F. Amzica (1993). "Intracellular Analysis of Relations between the Slow (<1 Hz) Neocortical Oscillation and Other Sleep Rhythms of the Electroencephalogram." In: *Journal of Neuroscience* 13.8, pp. 3266–3283. DOI: 10.1523/JNEUROSCI.13-08-03266.1993. (Visited on 2022) (cit. on pp. 20, 112).

Sterman, J D (2000). "Business Dynamics. System Thinking and Modeling for a Complex World." In: *McGraw-Hill Education. Boston:*

January 2000, 982 pp. DOI: 10.1016/S0022-3913(12)00047-9 (cit. on p. 12).

Stevenson, Ian H, James M Rebesco, Lee E Miller, and Konrad P Körding (2008). "Inferring Functional Connections between Neurons." In: *Current Opinion in Neurobiology* 18.6, pp. 582–588. DOI: 10.1016/j.conb.2008.11.005. (Visited on 2022) (cit. on p. 61).

Sticht, Heinrich, Peter Bayer, Dieter Willbold, Sonja Dames, Caroline Hilbich, Konrad Beyreuther, Rainer W. Frank, and Paul Rösch (1995). "Structure of Amyloid A4-(1–40)-Peptide of Alzheimer's Disease." In: *European Journal of Biochemistry* 233.1, pp. 293–298. DOI: 10.1111/j.1432-1033.1995.293_1.x. (Visited on 2022) (cit. on p. 3).

Stoewer, Adrian, Christian J. Kellner, Jan Benda, Thomas Wachtler, and Jan Grewe (2014). "File Format and Library for Neuroscience Data and Metadata." In: *Front. Neuroinform. Conference Abstract: Neuroinformatics*. DOI: 10.3389/conf.fninf.2014.18.00027. (Visited on 2021) (cit. on pp. 115, 119).

Stringer, Carsen, Marius Pachitariu, Nicholas A Steinmetz, Michael Okun, Peter Bartho, Kenneth D Harris, Maneesh Sahani, and Nicholas A Lesica (2016). "Inhibitory Control of Correlated Intrinsic Variability in Cortical Networks." In: *eLife* 5, e19695. DOI: 10.7554/eLife.19695. (Visited on 2022) (cit. on p. 89).

Stringer, Carsen, Marius Pachitariu, Nicholas Steinmetz, Matteo Carandini, and Kenneth D. Harris (2019). "High-Dimensional Geometry of Population Responses in Visual Cortex." In: *Nature* 571.7765, pp. 361–365. DOI: 10.1038/s41586-019-1346-5. (Visited on 2022) (cit. on p. 145).

Stroh, Albrecht, Helmuth Adelsberger, Alexander Groh, Charlotta Rühlmann, Sebastian Fischer, Anja Schierloh, Karl Deisseroth, and Arthur Konnerth (2013). "Making Waves: Initiation and Propagation of Corticothalamic Ca2+ Waves In Vivo." In: *Neuron* 77.6, pp. 1136–1150. DOI: 10.1016/j.neuron.2013.01.031. (Visited on 2020) (cit. on p. 114).

Student (1908). "The Probable Error of a Mean." In: *Biometrika* 6.1, pp. 1–25. DOI: 10.2307/2331554. JSTOR: 2331554. (Visited on 2022) (cit. on p. 62).

Su, Chih-Ying, Karen Menuz, Johannes Reisert, and John R. Carlson (2012). "Non-Synaptic Inhibition between Grouped Neurons in an Olfactory Circuit." In: *Nature* 492.7427, pp. 66–71. DOI: 10.1038/nature11712. (Visited on 2022) (cit. on p. 184).

Swedenborg, Emanuel (1882). *The Brain Considered Anatomically, Physiologically and Philosophically v. 1, 1882*. James Speirs (cit. on p. 16).

Switkes, Eugene, Melanie J. Mayer, and Jeffrey A. Sloan (1978). "Spatial Frequency Analysis of the Visual Environment: Anisotropy and the Carpentered Environment Hypothesis." In: *Vision Research* 18.10, pp. 1393–1399. DOI: 10.1016/0042-6989(78)90232-8 (cit. on p. 6).

Sylvester, Richard, John-Dylan Haynes, and Geraint Rees (2005). "Saccades Differentially Modulate Human LGN and V1 Responses in the Presence and Absence of Visual Stimulation." In: *Current Biology* 15.1, pp. 37–41. DOI: 10.1016/j.cub.2004.12.061. (Visited on 2022) (cit. on p. 186).

Szabó, Zsolt, László Héja, Gergely Szalay, Orsolya Kékesi, András Füredi, Kornélia Szebényi, Árpád Dobolyi, Tamás I. Orbán, Orsolya Kolacsek, Tamás Tompa, Zsombor Miskolczy, László Biczók, Balázs Rózsa, Balázs Sarkadi, and Julianna Kardos (2017). "Extensive Astrocyte Synchronization Advances Neuronal Coupling in Slow Wave Activity in Vivo." In: *Scientific Reports* 7.1, p. 6018. DOI: 10.1038/s41598-017-06073-7 (cit. on p. 113).

Szucs, Denes and John P. A. Ioannidis (2017). "When Null Hypothesis Significance Testing Is Unsuitable for Research: A Reassessment." In: *Frontiers in Human Neuroscience* 11, p. 390. DOI: 10.3389/fnhum.2017.00390 (cit. on p. 70).

Takahashi, Kazutaka, Sanggyun Kim, Todd P. Coleman, Kevin A. Brown, Aaron J. Suminski, Matthew D. Best, and Nicholas G. Hatsopoulos (2015). "Large-Scale Spatiotemporal Spike Patterning Consistent with Wave Propagation in Motor Cortex." In: *Nature Communications* 6.1, pp. 1–11. DOI: 10.1038/ncomms8169. (Visited on 2020) (cit. on p. 182).

Tal, Eran (2020). "Measurement in Science." In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2020. Metaphysics Research Lab, Stanford University. (Visited on 2022) (cit. on p. 9).

Tasic, Bosiljka et al. (2016). "Adult Mouse Cortical Cell Taxonomy Revealed by Single Cell Transcriptomics." In: *Nature Neuroscience* 19.2, pp. 335–346. DOI: 10.1038/nn.4216. (Visited on 2022) (cit. on p. 4).

Tasic, Bosiljka et al. (2018). "Shared and Distinct Transcriptomic Cell Types across Neocortical Areas." In: *Nature* 563.7729, pp. 72–78. DOI: 10.1038/s41586-018-0654-5. (Visited on 2022) (cit. on p. 4).

Teeter, Corinne, Ramakrishnan Iyer, Vilas Menon, Nathan Gouwens, David Feng, Jim Berg, Aaron Szafer, Nicholas Cain, Hongkui Zeng, Michael Hawrylycz, Christof Koch, and Stefan Mihalas (2018). "Generalized Leaky Integrate-and-Fire Models Classify Multiple Neuron Types." In: *Nature Communications* 9.1, p. 709. DOI: 10.1038/s41467-017-02717-4. (Visited on 2022) (cit. on p. 12).

Teeters, Jeffrey L., Kenneth D. Harris, K. Jarrod Millman, Bruno A. Olshausen, and Friedrich T. Sommer (2008). "Data Sharing for Computational Neuroscience." In: *Neuroinformatics* 6.1, pp. 47–55. DOI: 10.1007/s12021-008-9009-y. (Visited on 2022) (cit. on p. 13).

Tennøe, Simen, Geir Halnes, and Gaute T. Einevoll (2018). "Uncertainpy: A Python Toolbox for Uncertainty Quantification and Sensitivity Analysis in Computational Neuroscience." In: *Frontiers in*

Neuroinformatics 12, p. 49. DOI: 10.3389/fninf.2018.00049. (Visited on 2022) (cit. on p. 16).

Tetzlaff, Tom and Markus Diesmann (2010). "Dependence of Spike-Count Correlations on Spike-Train Statistics and Observation Time Scale." In: *Analysis of Parallel Spike Trains*. Ed. by Sonja Grün and Stefan Rotter. Springer Series in Computational Neuroscience. Boston, MA: Springer US, pp. 103–127. DOI: 10.1007/978-1-4419-5675-0_6. (Visited on 2023) (cit. on p. 40).

Thacker, Ben H, Scott W Doebling, Francois M Hemez, Mark C Anderson, Jason E Pepin, and Edward A Rodriguez (2004). *Concepts of Model Verification and Validation*. Tech. rep. Los Alamos National Lab., Los Alamos, NM (US) (cit. on pp. 11, 14, 15, 35–37).

Thickbroom, G. W., W. Knezevič, W. M. Carroll, and F. L. Mastaglia (1991). "Saccade Onset and Offset Lambda Waves: Relation to Pattern Movement Visually Evoked Potentials." In: *Brain Research* 551.1, pp. 150–156. DOI: 10.1016/0006-8993(91)90927-N. (Visited on 2022) (cit. on pp. 161, 186).

Thomson, A. M. (2002). "Synaptic Connections and Small Circuits Involving Excitatory and Inhibitory Neurons in Layers 2-5 of Adult Rat and Cat Neocortex: Triple Intracellular Recordings and Biocytin Labelling In Vitro." In: *Cerebral Cortex* 12.9, pp. 936–953. DOI: 10.1093/cercor/12.9.936. (Visited on 2022) (cit. on p. 89).

Thorpe, Simon, Denis Fize, and Catherine Marlot (1996). "Speed of Processing in the Human Visual System." In: *Nature* 381.6582, pp. 520–522. DOI: 10.1038/381520a0. (Visited on 2022) (cit. on p. 19).

Tiddia, Gianmarco, Bruno Golosio, Jasper Albers, Johanna Senk, Francesco Simula, Jari Pronold, Viviana Fanti, Elena Pastorelli, Pier Stanislao Paolucci, and Sacha J. van Albada (2022). "Fast Simulation of a Multi-Area Spiking Network Model of Macaque Cortex on an MPI-GPU Cluster." In: *Frontiers in Neuroinformatics* 16, p. 67. (Visited on 2022) (cit. on pp. 173, 175).

Tonielli, Leonardo, Chiara De Luca, Elena Pastorelli, Cristiano Capone, Francesco Simula, Cosimo Lupo, Irene Bernava, Giulia De Bonis, Gianmarco Tiddia, Bruno Golosio, and Pier Stanislao Paolucci (2022). *NREM and REM: Cognitive and Energetic Effects in Thalamo-Cortical Sleeping and Awake Spiking Model*. DOI: 10.48550/arXiv.2211.06889. arXiv: 2211.06889 [cs, q-bio]. (Visited on 2022) (cit. on p. 20).

Tononi, Giulio and Marcello Massimini (2008). "Why Does Consciousness Fade in Early Sleep?" In: *Annals of the New York Academy of Sciences* 1129.1, pp. 330–334. DOI: 10.1196/annals.1417.024. (Visited on 2022) (cit. on p. 20).

Torre, Emiliano, David Picado-Muiño, Michael Denker, Christian Borgelt, and Sonja Grün (2013). "Statistical Evaluation of Synchronous Spike Patterns Extracted by Frequent Item Set Mining." In: *Frontiers in Computational Neuroscience* 7. DOI: 10.3389/fncom.2013.00132. (Visited on 2022) (cit. on p. 42).

Torre, Emiliano, Pietro Quaglio, Michael Denker, Thomas Brochier, Alexa Riehle, and Sonja Grün (2016). "Synchronous Spike Patterns in Macaque Motor Cortex during an Instructed-Delay Reach-to-Grasp Task." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 36.32, pp. 8329–40. DOI: 10.1523/JNEUROSCI.4375-15.2016 (cit. on p. 182).

Torrecillos, Flavie, Julie Alayrangues, Bjørg Elisabeth Kilavik, and Nicole Malfait (2015). "Distinct Modulations in Sensorimotor Postmovement and Foreperiod β-Band Activities Related to Error Salience Processing and Sensorimotor Adaptation." In: *The Journal of neuroscience* 35.37, pp. 12753–65. DOI: 10.1523/JNEUROSCI.1090-15.2015. PMID: 26377464 (cit. on p. 161).

Tort-Colet, Núria, Cristiano Capone, Maria V. Sanchez-Vives, and Maurizio Mattia (2021). "Attractor Competition Enriches Cortical Dynamics during Awakening from Anesthesia." In: *Cell Reports* 35.12, p. 109270. DOI: 10.1016/j.celrep.2021.109270. (Visited on 2021) (cit. on p. 141).

Townsend, Rory G and Pulin Gong (2018). "Detection and Analysis of Spatiotemporal Patterns in Brain Activity." In: *PLOS Computational Biology* 14.12, e1006643. DOI: 10.1371/journal.pcbi.1006643. (Visited on 2020) (cit. on pp. 21, 123, 129, 151, 177, 182).

Townsend, Rory G, Selina S Solomon, Spencer C Chen, Alexander N J Pietersen, Paul R Martin, Samuel G Solomon, and Pulin Gong (2015). "Emergence of Complex Wave Patterns in Primate Cerebral Cortex." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 35.11, pp. 4657–62. DOI: 10.1523/JNEUROSCI.4509-14.2015 (cit. on p. 115).

Townsend, Rory G., Selina S. Solomon, Paul R. Martin, Samuel G. Solomon, and Pulin Gong (2017). "Visual Motion Discrimination by Propagating Patterns in Primate Cerebral Cortex." In: *Journal of Neuroscience* 37.42, pp. 10074–10084. DOI: 10.1523/JNEUROSCI.1538-17.2017. (Visited on 2022) (cit. on pp. 21, 112).

Trensch, Guido, Robin Gutzen, Inga Blundell, Michael Denker, and Abigail Morrison (2018). "Rigorous Neural Network Simulations: A Model Substantiation Methodology for Increasing the Correctness of Simulation Results in the Absence of Experimental Validation Data." In: *Frontiers in Neuroinformatics* 12, p. 81. DOI: 10.3389/fninf.2018.00081 (cit. on pp. 28, 33, 47–49, 78, 175).

Tripathy, Shreejoy J., Judith Savitskaya, Shawn D. Burton, Nathaniel N. Urban, and Richard C. Gerkin (2014). "NeuroElectro: A Window to the World's Neuron Electrophysiology Data." In: *Frontiers in Neuroinformatics* 8. DOI: 10.3389/fninf.2014.00040. (Visited on 2022) (cit. on p. 13).

Tseng, Ya-weng, Jörn Diedrichsen, John W. Krakauer, Reza Shadmehr, and Amy J. Bastian (2007). "Sensory Prediction Errors Drive Cerebellum-Dependent Adaptation of Reaching." In: *Journal of Neu-*

rophysiology 98.1, pp. 54–62. DOI: 10.1152/jn.00266.2007 (cit. on p. 161).

Tukker, John J., Prateep Beed, Dietmar Schmitz, Matthew E. Larkum, and Robert N. S. Sachdev (2020). "Up and Down States and Memory Consolidation Across Somatosensory, Entorhinal, and Hippocampal Cortices." In: *Frontiers in Systems Neuroscience* 14. DOI: 10.3389/fnsys.2020.00022. (Visited on 2020) (cit. on p. 20).

Valderrama, Mario, Benoît Crépon, Vicente Botella-Soler, Jacques Martinerie, Dominique Hasboun, Catalina Alvarado-Rojas, Michel Baulac, Claude Adam, Vincent Navarro, and Michel Le Van Quyen (2012). "Human Gamma Oscillations during Slow Wave Sleep." In: *PLOS ONE* 7.4, e33477. DOI: 10.1371/journal.pone.0033477. (Visited on 2022) (cit. on pp. 155, 183).

Van Bussel, Frank (2011). "Inferring Synaptic Connectivity from Spatio-Temporal Spike Patterns." In: *Frontiers in Computational Neuroscience*. DOI: 10.3389/fncom.2011.00003. (Visited on 2022) (cit. on p. 62).

Varela, Francisco J (1995). "Resonant Cell Assemblies: A New Approach to Cognitive Functions and Neuronal Synchrony." In: *Biological research* 28, pp. 81–81 (cit. on p. 184).

Varela, Francisco, Jean-Philippe Lachaux, Eugenio Rodriguez, and Jacques Martinerie (2001). "The Brainweb: Phase Synchronization and Large-Scale Integration." In: *Nature Reviews Neuroscience* 2.4, pp. 229–239. DOI: 10.1038/35067550. (Visited on 2022) (cit. on pp. 181, 185).

Venkatesh, Manasij, Joseph Jaja, and Luiz Pessoa (2019). "Comparing Functional Connectivity Matrices: A Geometry-Aware Approach Applied to Participant Identification." In: *NeuroImage*, p. 116398. DOI: 10.1016/j.neuroimage.2019.116398. (Visited on 2019) (cit. on p. 63).

Verma, Akanksha and Dr Kumar (2020). "Multilayer Perceptron Artificial Neural Network: A Review." In: pp. 1–38 (cit. on p. 19).

Vezoli, Julien, Martin Vinck, Conrado Arturo Bosman, André Moraes Bastos, Christopher Murphy Lewis, Henry Kennedy, and Pascal Fries (2021). "Brain Rhythms Define Distinct Interaction Networks with Differential Dependence on Anatomy." In: *Neuron* 109.23, 3862–3878.e5. DOI: 10.1016/j.neuron.2021.09.052. (Visited on 2022) (cit. on pp. 160, 183).

Vinck, Martin, Bruss Lima, Thilo Womelsdorf, Robert Oostenveld, Wolf Singer, Sergio Neuenschwander, and Pascal Fries (2010). "Gamma-Phase Shifting in Awake Monkey Visual Cortex." In: *Journal of Neuroscience* 30.4, pp. 1250–1257. DOI: 10.1523/JNEUROSCI.1623-09.2010. (Visited on 2022) (cit. on pp. 20, 21, 182, 185).

Vinck, Martin, Cem Uran, and Andrés Canales-Johnson (2022). *The Neural Dynamics of Feedforward and Feedback Interactions in Predictive Processing*. DOI: 10.31234/osf.io/n3afb. (Visited on 2022) (cit. on p. 183).

Virmani, Mohit and Nithin Nagaraj (2019). "A Novel Perturbation Based Compression Complexity Measure for Networks." In: *Heliyon* 5.2, e01181. DOI: 10.1016/j.heliyon.2019.e01181. (Visited on 2022) (cit. on p. 176).

Virtanen, Pauli et al. (2020). "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python." In: *Nature Methods* 17.3, pp. 261–272. DOI: 10.1038/s41592-019-0686-2. (Visited on 2021) (cit. on pp. 115, 117).

Voges, N. and L. Perrinet (2012). "Complex Dynamics in Recurrent Cortical Networks Based on Spatially Realistic Connectivities." In: *Frontiers in Computational Neuroscience* 6, p. 41. DOI: 10.3389/fncom.2012.00041 (cit. on pp. 34, 41).

Volgushev, Maxim, Sylvain Chauvette, Mikhail Mukovski, and Igor Timofeev (2006). "Precise Long-Range Synchronization of Activity and Silence in Neocortical Neurons during Slow-Wave Sleep." In: *Journal of Neuroscience* 26.21, pp. 5665–5672. DOI: 10.1523/JNEUROSCI.0279-06.2006. (Visited on 2020) (cit. on p. 113).

Vyazovskiy, Vladyslav V., Umberto Olcese, Erin C. Hanlon, Yuval Nir, Chiara Cirelli, and Giulio Tononi (2011). "Local Sleep in Awake Rats." In: *Nature* 472.7344, pp. 443–447. DOI: 10.1038/nature10009. (Visited on 2020) (cit. on p. 113).

Wainrib, Gilles and Jonathan Touboul (2013). "Topological and Dynamical Complexity of Random Neural Networks." In: *Physical Review Letters* 110.11, p. 118101. DOI: 10.1103/PhysRevLett.110.118101. (Visited on 2022) (cit. on p. 180).

Watrous, Andrew J., Juergen Fell, Arne D. Ekstrom, and Nikolai Axmacher (2015). "More than Spikes: Common Oscillatory Mechanisms for Content Specific Neural Representations during Perception and Memory." In: *Current opinion in neurobiology* 31, pp. 33–39. DOI: 10.1016/j.conb.2014.07.024. (Visited on 2019) (cit. on p. 20).

Wei, Yina, Giri P. Krishnan, and Maxim Bazhenov (2016). "Synaptic Mechanisms of Memory Consolidation during Sleep Slow Oscillations." In: *Journal of Neuroscience* 36.15, pp. 4231–4247. DOI: 10.1523/JNEUROSCI.3648-15.2016. (Visited on 2020) (cit. on pp. 20, 112, 171).

Wei, Ziqiang, Bei-Jung Lin, Tsai-Wen Chen, Kayvon Daie, Karel Svoboda, and Shaul Druckmann (2020). "A Comparison of Neuronal Population Dynamics Measured with Calcium Imaging and Electrophysiology." In: *PLOS Computational Biology* 16.9, e1008198. DOI: 10.1371/journal.pcbi.1008198. (Visited on 2020) (cit. on p. 145).

Weiler, Nicholas, Lydia Wood, Jianing Yu, Sara A Solla, and Gordon M G Shepherd (2008). "Top-down Laminar Organization of the Excitatory Network in Motor Cortex." In: *Nature Neuroscience* 11.3, pp. 360–366. DOI: 10.1038/nn2049. (Visited on 2022) (cit. on p. 89).

Wigner, Eugene P. (1960). "The Unreasonable Effectiveness of Mathematics in the Natural Sciences. Richard Courant Lecture in Mathematical Sciences Delivered at New York University, May 11, 1959."

In: *Communications on Pure and Applied Mathematics* 13.1, pp. 1–14. DOI: 10.1002/cpa.3160130102. (Visited on 2022) (cit. on p. 9).

Wilkinson, Mark D. et al. (2016). "The FAIR Guiding Principles for Scientific Data Management and Stewardship." In: *Scientific Data* 3.1, p. 160018. DOI: 10.1038/sdata.2016.18. (Visited on 2022) (cit. on pp. 173, 188).

Wilson, Matthew A, Carmen Varela, and Miguel Remondes (2015). "Phase Organization of Network Computations." In: *Current Opinion in Neurobiology*. SI: Brain Rhythms and Dynamic Coordination 31, pp. 250–253. DOI: 10.1016/j.conb.2014.12.011. (Visited on 2019) (cit. on p. 184).

Winawer, Jonathan, Nathan Witthoft, Michael C. Frank, Lisa Wu, Alex R. Wade, and Lera Boroditsky (2007). "Russian Blues Reveal Effects of Language on Color Discrimination." In: *Proceedings of the National Academy of Sciences* 104.19, pp. 7780–7785. DOI: 10.1073/pnas.0701644104. (Visited on 2021) (cit. on p. 6).

Wood, C. C. (2019). "The Computational Stance in Biology." In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 374.1774, p. 20180380. DOI: 10.1098/rstb.2018.0380. (Visited on 2021) (cit. on pp. 18, 188).

Wu, Jian-Young, Xiaoying Huang, and Chuan Zhang (2008). "Propagating Waves of Activity in the Neocortex: What They Are, What They Do." In: *The Neuroscientist* 14.5, pp. 487–502. DOI: 10.1177/1073858408317066. (Visited on 2019) (cit. on pp. 20, 21, 177).

Xu, Weifeng, Xiaoying Huang, Kentaroh Takagaki, and Jian-young Wu (2007). "Compression and Reflection of Visually Evoked Cortical Waves." In: *Neuron* 55.1, pp. 119–129. DOI: 10.1016/j.neuron.2007.06.016. (Visited on 2022) (cit. on p. 185).

Yamawaki, N. and G. M. G. Shepherd (2015). "Synaptic Circuit Organization of Motor Corticothalamic Neurons." In: *Journal of Neuroscience* 35.5, pp. 2293–2307. DOI: 10.1523/JNEUROSCI.4023-14.2015. (Visited on 2022) (cit. on p. 89).

Yamawaki, Naoki, Katharine Borges, Benjamin A Suter, Kenneth D Harris, and Gordon M G Shepherd (2014). "A Genuine Layer 4 in Motor Cortex with Prototypical Synaptic Circuit Connectivity." In: *eLife* 3, e05422. DOI: 10.7554/eLife.05422. (Visited on 2022) (cit. on p. 89).

Yamins, Daniel L. K. and James J. DiCarlo (2016). "Using Goal-Driven Deep Learning Models to Understand Sensory Cortex." In: *Nature Neuroscience* 19.3, pp. 356–365. DOI: 10.1038/nn.4244. (Visited on 2022) (cit. on p. 19).

Yang, Lifang, Mengmeng Li, Long Yang, Haofeng Wang, Hong Wan, and Zhigang Shang (2020a). "Functional Connectivity Changes in the Intra- and Inter-Brain during the Construction of the Multi-Brain Network of Pigeons." In: *Brain Research Bulletin* 161, pp. 147–157.

DOI: 10.1016/j.brainresbull.2020.04.015. (Visited on 2022) (cit. on p. 8).

Yang, Shuangming, Bin Deng, Jiang Wang, Huiyan Li, Meili Lu, Yanqiu Che, Xile Wei, and Kenneth A. Loparo (2020b). "Scalable Digital Neuromorphic Architecture for Large-Scale Biophysically Meaningful Neural Network With Multi-Compartment Neurons." In: *IEEE Transactions on Neural Networks and Learning Systems* 31.1, pp. 148–162. DOI: 10.1109/TNNLS.2019.2899936 (cit. on p. 8).

Yegenoglu, Alper, Anand Subramoney, Thorsten Hater, Cristian Jimenez-Romero, Wouter Klijn, Aarón Pérez Martín, Michiel van der Vlag, Michael Herty, Abigail Morrison, and Sandra Diaz-Pier (2022). "Exploring Parameter and Hyper-Parameter Spaces of Neuroscience Models on High Performance Computers With Learning to Learn." In: *Frontiers in Computational Neuroscience* 16, p. 885207. (Visited on 2022) (cit. on p. 100).

Yger, Pierre, Giulia LB Spampinato, Elric Esposito, Baptiste Lefebvre, Stéphane Deny, Christophe Gardella, Marcel Stimberg, Florian Jetter, Guenther Zeck, Serge Picaud, Jens Duebel, and Olivier Marre (2018). "A Spike Sorting Toolbox for up to Thousands of Electrodes Validated with Ground Truth Recordings in Vitro and in Vivo." In: *eLife* 7, e34518. DOI: 10.7554/eLife.34518. (Visited on 2022) (cit. on p. 94).

Zanos, Theodoros P., Patrick J. Mineault, Konstantinos T. Nasiotis, Daniel Guitton, and Christopher C. Pack (2015). "A Sensorimotor Role for Traveling Waves in Primate Visual Cortex." In: *Neuron* 85.3, pp. 615–627. DOI: 10.1016/j.neuron.2014.12.043. (Visited on 2022) (cit. on pp. 161, 171, 186, 187).

Zanos, Theodoros P., Patrick J. Mineault, and Christopher C. Pack (2011). "Removal of Spurious Correlations between Spikes and Local Field Potentials." In: *Journal of Neurophysiology* 105.1, pp. 474–486. DOI: 10.1152/jn.00642.2010 (cit. on p. 181).

Zehl, Lyuba, Florent Jaillet, Adrian Stoewer, Jan Grewe, Andrey Sobolev, Thomas Wachtler, Thomas G. Brochier, Alexa Riehle, Michael Denker, and Sonja Grün (2016). "Handling Metadata in a Neurophysiology Laboratory." In: *Frontiers in neuroinformatics* 10.July, p. 26. DOI: 10.3389/fninf.2016.00026 (cit. on p. 117).

Zhang, Honghui, Andrew J. Watrous, Ansh Patel, and Joshua Jacobs (2018). "Theta and Alpha Oscillations Are Traveling Waves in the Human Neocortex." In: *Neuron* 98.6, 1269–1281.e4. DOI: 10.1016/j.neuron.2018.05.019. (Visited on 2022) (cit. on pp. 148, 187).

Zhang, Xiaohui, Eric C. Landsness, Wei Chen, Hanyang Miao, Michelle Tang, Lindsey M. Brier, Joseph P. Culver, Jin-Moo Lee, and Mark A. Anastasio (2022). "Automated Sleep State Classification of Wide-Field Calcium Imaging Data via Multiplex Visibility Graphs and Deep Learning." In: *Journal of Neuroscience Methods* 366, p. 109421.

DOI: 10.1016/j.jneumeth.2021.109421. (Visited on 2022) (cit. on p. 174).

Zhou, Qingguo, Tao Jin, and Hong Zhao (2009). "Correlation Between Eigenvalue Spectra and Dynamics of Neural Networks." In: *Neural Computation* 21.10, pp. 2931–2941. DOI: 10.1162/neco.2009.12-07-671. (Visited on 2019) (cit. on pp. 62, 180).

Zi, Z. (2011). "Sensitivity Analysis Approaches Applied to Systems Biology Models." In: *IET Systems Biology* 5.6, pp. 336–346. DOI: 10.1049/iet-syb.2011.0015. (Visited on 2022) (cit. on p. 16).

de Haan, Marcel Jan, Thomas Brochier, Sonja Grün, Alexa Riehle, and Frédéric V Barthélemy (2018). "Real-Time Visuomotor Behavior and Electrophysiology Recording Setup for Use with Humans and Monkeys." In: *Journal of neurophysiology* 120.2, pp. 539–552. DOI: 10.1152/jn.00262.2017 (cit. on pp. 148, 150).

de Schepper, Robin et al. (2022). *NEST 3.2*. DOI: 10.5281/ZENODO.5886893. (Visited on 2022) (cit. on pp. 94, 173).

de Vries, Saskia E. J. et al. (2020). "A Large-Scale Standardized Physiological Survey Reveals Functional Organization of the Mouse Visual Cortex." In: *Nature Neuroscience* 23.1, pp. 138–151. DOI: 10.1038/s41593-019-0550-9. (Visited on 2022) (cit. on pp. 145, 182).

di Volo, Matteo, Alberto Romagnoni, Cristiano Capone, and Alain Destexhe (2019). "Biologically Realistic Mean-Field Models of Conductance-Based Networks of Spiking Neurons with Adaptation." In: *Neural Computation* 31.4, pp. 653–680. DOI: 10.1162/neco_a_01173. (Visited on 2022) (cit. on p. 141).

van Albada, Sacha J, Andrew G Rowley, Johanna Senk, Michael Hopkins, Maximilian Schmidt, Alan B Stokes, David R Lester, Markus Diesmann, and Steve B Furber (2018). "Performance Comparison of the Digital Neuromorphic Hardware SpiNNaker and the Neural Network Simulation Software NEST for a Full-Scale Cortical Microcircuit Model." In: *Frontiers in neuroscience* 12, p. 291. DOI: 10.3389/fnins.2018.00291 (cit. on pp. 35, 59, 175).

van Hemmen, J. Leo and Terrence J. Sejnowski (2005). *23 Problems in Systems Neuroscience*. Oxford University Press (cit. on pp. 23, 187).

van Kerkoerle, Timo, Matthew W. Self, Bruno Dagnino, Marie-Alice Gariel-Mathis, Jasper Poort, Chris van der Togt, and Pieter R. Roelfsema (2014). "Alpha and Gamma Oscillations Characterize Feedback and Feedforward Processing in Monkey Visual Cortex." In: *Proceedings of the National Academy of Sciences* 111.40, pp. 14332–14341. DOI: 10.1073/pnas.1402773111. (Visited on 2022) (cit. on p. 183).

von Stein, Astrid and Johannes Sarnthein (2000). "Different Frequencies for Different Scales of Cortical Integration: From Local Gamma to Long Range Alpha/Theta Synchronization." In: *International Journal of Psychophysiology* 38.3, pp. 301–313. DOI: 10.1016/S0167-8760(00)00172-0. (Visited on 2022) (cit. on pp. 155, 183).

"It's hard for students to follow the path of greater resistance. [..] But the paradox is to be really successful, to really make a difference, to change to course of scientific discovery you have to swim against the current. You have to pursue ideas that are not popular, that people might not agree with, that might not really make sense at first. You have to befriend the crackpots and read poetry, and take long walks, and do all the things which aren't extensively the right things to do, the productive things, the good uses of time. So, we have to be really attentive to the unappreciated importance of uselessness."

— Sam Gershman

ACKNOWLEDGMENTS

Sam Gershman's quote encapsulates much of what I valued in my scientific environment over the last years and in science in general. Especially interdisciplinary fields can be daunting as there are so many related topics to explore. However, going down the rabbit hole of any topic, even when seemingly "useless" and unrelated to my work, is usually beneficial in the end. It is always a great joy to hear the fun little 'click' sound that the brain makes when a new mental connection falls into place between previously apart domains. I'm very grateful for being given a lot of freedom to pursue my own ideas and shape my research projects. However, this would not have been as enriching without the support and creative input from the diverse environment of people with various backgrounds and passions.

First and foremost, I thank my supervisor Michael Denker for the endless support throughout the past years and for showing me a lot about diligence and clarity in good scientific analysis, data management, writing, and thinking. Also, I would like to thank Sonja Grün for her supervision, continuing guidance, and for instilling in me the reflex to question every detail in the data. Furthermore, I thank Björn Kampa for joining as second examiner and for nudging me to look beyond the scope of each individual research project.

I want to express my gratitude to the whole institute led by Markus Diesmann and Sonja Grün, the secretaries, scientific coordinators, IT administrators, and everyone working tirelessly to create the welcoming and supportive environment that lies at the heart of good research. This also extends to our collaborators at the INT in Marseille, including Alexa Riehle, Thomas Brochier, Frédéric Barthélemy, Marcel De Haan, and many others, for welcoming us in their lab once a year, providing data, metadata, and, most importantly, an experimentalist point of view.

I'd like to acknowledge the fortunate opportunity to work within the context of the Human Brain Project. I'm very thankful for experi-

encing the vibrant exchanges across labs, the various upsides (as well as downsides) of large-scale research projects, and the chance to work with collaborators from around Europe. In this regard, I thank Andrew Davison and his group, including Lungsi Sharma and Shailesh Appukuttan, for joint efforts in building open-source tools and detailed discussions about validation. Furthermore, I thank Pier Paolucci and his group, particularly Giulia De Bonis, Chiara De Luca, and Cristiano Capone, for the collaborative quest of taming slow wave activity data and for inviting me for a stay in their lab.

In the context of joint publications, I also would like to thank all the other co-authors I had the pleasure of working with, including (amongst others) Guido Trensch, Michael van Papen, Aitor Morales-Gregorio, Abigail Morrison, and Sacha van Albada.

I also have to express my enormous gratitude to all my friends and colleagues, current and former, that made these past years so exciting, including (as incomplete and unordered list) Alex van Meegen, Alessandra Stella, Aitor Morales-Gregorio, Anno Kurth, Simon Essink, Alex Kleinjohann, Michael Dick, Barna Zajzon, Jasper Albers, Tobias Schulte to Brinke, Johanna Senk, Julia Sprenger, Robin Pauli, Lyuba Zehl, Gina Schumacher, Philipp Weidel, Pietro Quaglio, Vahid Rostami, and Anna Jürgensen. I thank you all for endless discussions about food, philosophy, visualizations, fermentation, coding details, books, and much more. You are the fun in fundamental science. Furthermore, I thank my sister Kaja for always being supportive, even when it's just making photos of red twine for my outline figure. Finally, I'm deeply thankful to Greta for being there and keeping me sane in the last months of writing.

Band / Volume 89
Controlling the electrical properties of oxide heterointerfaces through their interface chemistry
M.-A. Rose (2022), vi, 162 pp
ISBN: 978-3-95806-667-0

Band / Volume 90
Modeling and Suppressing Unwanted Parasitic Interactions in Superconducting Circuits
X. Xu (2022), 123, XVIII pp
ISBN: 978-3-95806-671

Band / Volume 91
Activating molecular magnetism by controlled on-surface coordination
I. Cojocariu (2022), xi, 169 pp
ISBN: 978-3-95806-674-8

Band / Volume 92
Computational study of structural and optical properties of two-dimensional transition-metal dichalcogenides with implanted defects
S. H. Rost (2023), xviii, 198 pp
ISBN: 978-3-95806-682-3

Band / Volume 93
DC and RF characterization of bulk CMOS and FD-SOI devices at cryogenic temperatures with respect to quantum computing applications
A. Artanov (2023), xv, 80, xvii-liii pp
ISBN: 978-3-95806-687-8

Band / Volume 94
HAXPES study of interface and bulk chemistry of ferroelectric HfO_2 capacitors
T. Szyjka (2023), viii, 120 pp
ISBN: 978-3-95806-692-2

Band / Volume 95
A brain inspired sequence learning algorithm and foundations of a memristive hardware implementation
Y. Bouhadjar (2023), xii, 149 pp
ISBN: 978-3-95806-693-9

Band / Volume 96
Characterization and modeling of primate cortical anatomy and activity
A. Morales-Gregorio (2023), ca. 260 pp.
ISBN: 978-3-95806-698-4

Band / Volume 97
Hafnium oxide based memristive devices as functional elements of neuromorphic circuits
F. J. Cüppers (2023), vi, ii, 214 pp
ISBN: 978-3-95806-702-8

Band / Volume 98
Simulation and theory of large - scale cortical networks
A. van Meegen (2023), ca. 250 pp
ISBN: 978-3-95806-708-0

Band / Volume 99
Structure of two-dimensional multilayers and topological superconductors: surfactant mediated growth, intercalation, and doping
Y.-R. Lin (2023), x, 111 pp
ISBN: 978-3-95806-716-5

Band / Volume 100
Frequency mixing magnetic detection for characterization and multiplex detection of superparamagnetic nanoparticles
A. M. Pourshahidi (2023), X, 149 pp
ISBN: 978-3-95806-727-1

Band / Volume 101
Unveiling the relaxation dynamics of Ag/HfO_2 based diffusive memristors for use in neuromorphic computing
S. A. Chekol (2023), x, 185 pp
ISBN: 978-3-95806-729-5

Band / Volume 102
Analysis and quantitative comparison of neural network dynamics on a neuron-wise and population level
R. Gutzen (2024), xii, 252 pp
ISBN: 978-3-95806-738-7

Weitere ***Schriften des Verlags im Forschungszentrum Jülich*** unter
http://wwwzb1.fz-juelich.de/verlagextern1/index.asp

Information
Band / Volume 102
ISBN 978-3-95806-738-7

Mitglied der Helmholtz-Gemeinschaft

JÜLICH
Forschungszentrum