Technische Universität München
TUM School of Life Sciences

Dynamic propagation of neuronal activity in the reptilian brain

Juan Luis Riquelme

Vollständiger Abdruck der von der TUM School of Life Sciences der Technischen Universität München zur Erlangung eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitz:

Prof. Dr. Frank Johannes

Prüfende der Dissertation:

1. Prof. Dr. Julijana Gjorgjieva
2. Prof. Dr. Gilles Laurent
3. Prof. Dr. Anton Sirota

Die Dissertation wurde am 08.05.2023 bei der Technischen Universität München eingereicht und durch die TUM School of Life Sciences am 17.07.2023 angenommen.

Acknowledgements

First and foremost, I would like to thank my supervisor Julijana Gjorgjieva for her invaluable support and guidance throughout my PhD. I will forever be grateful for her mentoring. I have been so lucky to have had her in my corner all these years.

I want to especially thank all the co-authors from my publications. Mike Hemberger for his support from day one, Lorenz Fenk for the passionate long talks, and Gilles Laurent for his profound insights and exceptional mentorship.

In addition, I want to thank the previous and current members of the Gjorgjieva lab, the Laurent lab, and my colleagues at the Max Planck Institute for Brain Research. Together they have been an essential part of my personal and scientific development.

I want to thank the members of my thesis advisory committee, Gilles Laurent, Hiroshi Ito, and Matthias Kaschube, for their guidance has led to the completion of this thesis, as well as the members of my examining committee.

Finally, I am grateful to my friends, family, and especially Eloïse for their support.

Prior publication of parts of this thesis

Some of the results and figures presented in Chapters 4 and 5 are part of an article entitled *Single spikes drive sequential propagation and routing of activity in a cortical network* which I wrote together with my supervisor Prof. Dr. Julijana Gjorgjieva, and feedback from my collaborators Dr. Mike Hemberger and Prof. Dr. Gilles Laurent. The article has been peer-reviewed and published in the journal *eLife* [1].

Some of the results and figures presented in Chapter 6 are part of an article entitled *Interhemispheric competition during sleep* which I wrote with my collaborators Dr. Lorenz Fenk and Prof. Dr. Gilles Laurent. The article has been peer-reviewed and published in the journal *Nature* [2].

All methods, results, figures, and tables from these articles which are part of this thesis were my contribution to the articles unless specifically mentioned otherwise.

Summary

The brain is a remarkable biological machine made of complex networks of neurons that generate behavior by producing and processing electrical signals. Understanding how these signals propagate and change through the brain is crucial for uncovering the mechanisms underlying cognitive processes. Reliable and flexible propagation effectively produces a mapping between inputs and outputs and is, therefore, a key ingredient for implementing brain computations. Network mechanisms shape propagation at multiple temporal and spatial scales, resulting in intricate dynamics such as competition or coordination of multiple signals. This thesis investigates the propagation of neuronal activity in the reptilian brain, using computational methods to analyze the routing of single spikes within the turtle dorsal cortex and the coordination of neuronal populations across multiple areas in the sleeping lizard. By studying the neuronal activity in reptiles, we can gain insights into the conserved principles and evolutionary modifications of propagation mechanisms while taking advantage of the unique physiological and behavioral characteristics of these animals.

Recent evidence from the turtle dorsal cortex shows that the activation of a single neuron can trigger a sequence of spiking activity in the surrounding network. However, it is still unclear how these sequences propagate reliably within the network. To investigate the propagation of spiking activity in the turtle dorsal cortex, I used a computational model of a randomly connected network constrained by experimental measurements. This model includes a long-tailed distribution of synaptic strengths, which allows the generation of reliable sequences from single spikes without structured connectivity. Manipulating these networks through an additional input spike could flexibly alter the particular sequence of neurons that activity traverses. Studying the propagation of multiple parallel running sequences in the recurrent network revealed a highly combinatorial repertoire of highly specific and reliable interactions between sequences. Therefore, a sparse network of strong connections is a simple yet powerful mechanism for precisely routing spike propagation within the network and enabling reliable and flexible cortical computations.

I developed computational analysis methods to analyze multi-area Local Field Potential (LFP) recordings from the sleeping lizard during the Rapid Eye Movement (REM) phase of sleep. These methods involved algorithms that detected and quantified LFP events, overcoming technical challenges such as processing large data resulting from high sampling rates and the analysis of millisecond-scale events over hours-long recordings. As a result of this quantification, my collaborators and I found that REM LFP activity in the claustrum, a sub-cortical area, showed tightly coordinated bilateral activity with a fixed lag that alternated during the night,

reflecting competitive dynamics between multiple populations of midbrain nuclei. These results highlight the importance of computational methods to study LFP recordings and how these can inform us about the propagation of activity across multiple areas in the lizard brain.

In conclusion, understanding the propagation of neuronal activity is crucial to uncover the mechanisms of brain computation. This thesis employs computational methods to investigate signal propagation within the turtle dorsal cortex and across multiple areas in the sleeping lizard. My research contributes to a better understanding of brain function, showing how a single spike can trigger reliable and flexible firing sequences in cortical networks and describing complex dynamics across sub-cortical populations. These insights, obtained from the brains of reptiles, may inform us of general principles of brain function and their evolution. Furthermore, I show how computational methods can help us better understand the mechanisms of signal propagation in the brain. Overall, my findings provide insights into the mechanisms of signal propagation in the brain and how networks of neurons might implement reliable and flexible computations.

Zusammenfassung

Das Gehirn ist eine bemerkenswerte biologische Maschine, die aus komplexen Netzwerken von Neuronen besteht, welche Verhalten erzeugen, indem sie elektrische Signale produzieren und verarbeiten. Es ist entscheidend zu verstehen, wie sich diese Signale im Gehirn ausbreiten und verändern, um die Mechanismen zu entschlüsseln, die kognitiven Prozessen zugrunde liegen. Eine verlässliche und flexible Signalverbreitung führt zu einer effektiven Zuordnung zwischen Eingängen und Ausgängen und ist daher eine Schlüsselkomponente für die Umsetzung von Gehirnberechnungen. Netzwerkmechanismen formen die Ausbreitung auf verschiedenen zeitlichen und räumlichen Skalen, was zu komplizierten Dynamiken wie Wettbewerb oder Koordination mehrerer Signale führt. In dieser Arbeit wird die Ausbreitung von neuronaler Aktivität im Reptiliengehirn untersucht. Dabei werden Berechnungsmethoden eingesetzt, um die Weiterleitung einzelner Spikes im dorsalen Kortex der Schildkröte und die Koordination von Neuronenpopulationen in mehreren Hirnarealen der schlafenden Eidechse zu analysieren. Durch die Untersuchung der neuronalen Aktivität in Reptilien können wir Einblicke in konservierte Prinzipien und evolutionäre Veränderungen der Ausbreitungsmechanismen gewinnen und gleichzeitig die einzigartigen physiologischen und verhaltensbezogenen Merkmale dieser Tiere nutzen.

Jüngste Studien des dorsalen Kortex von Schildkröten zeigen, dass die Aktivierung eines einzelnen Neurons eine Folge von Spike-Aktivitäten im umgebenden Netzwerk auslösen kann. Es ist jedoch noch unklar, wie sich diese Sequenzen zuverlässig innerhalb des Netzwerks ausbreiten. Um die Ausbreitung der Spike-Aktivität im dorsalen Kortex der Schildkröte zu untersuchen, habe ich ein Computermodell eines zufällig verbundenen Netzwerks verwendet, das durch experimentelle Messungen eingeschränkt wurde. Dieses Modell beinhaltet eine endlastige Verteilung der Synapsenstärken, die es ermöglicht, zuverlässige Sequenzen aus einzelnen Spikes ohne strukturierte Konnektivität zu erzeugen. Die Manipulation dieser Netzwerke durch einen zusätzlichen Eingangsspike konnte die bestimmte Abfolge der Neuronen, die aktiviert werden, flexibel verändern. Die Untersuchung der Ausbreitung mehrerer parallel laufender Sequenzen im rekurrenten Netzwerk ergab ein hochkombinatorisches Repertoire an sehr spezifischen und zuverlässigen Interaktionen zwischen den Sequenzen. Ein dünnbesetztes Netzwerk aus starken Verbindungen ist daher ein einfacher, aber leistungsfähiger Mechanismus, um die Ausbreitung von Spikes innerhalb des Netzwerks präzise zu steuern und zuverlässige und flexible kortikale Berechnungen zu ermöglichen.

Ich habe computergestützte Analysemethoden entwickelt, um die Aufzeichnungen der lokalen Feldpotentiale (LFP) von schlafenden Eidechsen während der REM-Phase

(Rapid Eye Movement) zu analysieren. Diese Methoden umfassten Algorithmen, die LFP-Ereignisse erkannten und quantifizierten und dabei technische Herausforderungen wie die Verarbeitung großer Datenmengen aufgrund hoher Samplingraten und die Analyse von Ereignissen im Millisekundenbereich bei stundenlangen Aufzeichnungen meisterten. Als Ergebnis dieser Quantifizierung fanden meine Kollegen und ich heraus, dass die REM-LFP-Aktivität im Klaustrum, einem subkortikalen Bereich, eine eng koordinierte bilaterale Aktivität mit einer festen Verzögerung aufweist, die sich während der Nacht abwechselt und die Konkurrenzdynamik zwischen mehreren Populationen von Mittelhirnkernen widerspiegelt. Diese Ergebnisse zeigen, wie wichtig computergestützte Methoden für die Untersuchung von LFP-Aufzeichnungen sind und wie sie uns Aufschluss über die Ausbreitung der Aktivität in verschiedenen Bereichen des Eidechsengehirns geben können.

Zusammenfassend ist das Verständnis der Ausbreitung neuronaler Aktivität entscheidend, um die Mechanismen der Gehirnberechnung aufzudecken. In dieser Arbeit wurden computergestützte Methoden angewendet, um die Signalausbreitung im dorsalen Kortex der Schildkröte und in mehreren Hirnarealen der schlafenden Eidechse zu untersuchen. Meine Forschung trägt zu einem besseren Verständnis der Hirnfunktionen bei, indem sie zeigt, wie ein einzelner Spike zuverlässige und flexible Zündsequenzen in kortikalen Netzwerken auslösen kann, und beschreibt die komplexe Dynamik in subkortikalen Populationen. Diese Erkenntnisse, die aus den Gehirnen von Reptilien gewonnen wurden, können uns Aufschluss über allgemeine Prinzipien der Gehirnfunktion und ihrer Evolution geben. Außerdem zeige ich, wie wir mit Hilfe von Berechnungsmethoden die Mechanismen der Signalausbreitung im Gehirn besser verstehen können. Insgesamt geben meine Ergebnisse Aufschluss über die Mechanismen der Signalausbreitung im Gehirn und darüber, wie Netzwerke von Neuronen zuverlässige und flexible Berechnungen durchführen können.

Contents

Acknowledgements i

Prior publication of parts of this thesis iii

Summary v

German summary vii

List of Figures xi

Acronyms and abbreviations xiii

1 Introduction 1
- 1.1 Thesis overview . . . 1
- 1.2 Why activity propagation? . . . 3
- 1.3 Why reptiles? . . . 10
- 1.4 Why a computational approach? . . . 16

2 Materials and methods 21
- 2.1 Electrophysiological data . . . 21
- 2.2 Code availability . . . 22
- 2.3 Simulations . . . 23

3 Biological and theoretical background 27
- 3.1 The reptilian cortex . . . 27
  - 3.1.1 Circuitry of the turtle dorsal cortex . . . 31
    - 3.1.1.1 The turtle cortex ex vivo . . . 32
    - 3.1.1.2 Lessons from mammalian cortical connectivity . . . 34
  - 3.1.2 Vision in the turtle cortex . . . 37
- 3.2 Two-stage sleep . . . 40
  - 3.2.1 Sleep and the reptilian brain . . . 43
- 3.3 Computational models of spike propagation . . . 45
  - 3.3.1 Synfire chains . . . 45
  - 3.3.2 Polychronous chains . . . 51
  - 3.3.3 Rate propagation . . . 53

4 A model of the turtle cortex that produces repeatable spiking sequences 57
- 4.1 Overview . . . 57
  - 4.1.1 Experimental background and modeling questions . . . 58

4.2 Experimentally defined constraints on a random network model . . . 60
4.3 Single spikes trigger reliable sequential activity 66
4.4 Strong connections provide reliability and weak connections modulate it 77
4.5 Discussion 82

5 Flexible routing of single spikes in a random network model of the turtle cortex 85
5.1 Overview 85
5.1.1 Experimental background and modeling questions 86
5.2 Sequences are composed of sub-sequences that correspond to sub-networks of strong connections 88
5.3 Sparse external input can halt or facilitate the propagation of activity 92
5.4 Sequences from multiple triggers reliably activate combinations of followers 96
5.5 Discussion 101

6 Inter-hemispheric competition during REM sleep in the lizard brain 107
6.1 Overview 107
6.1.1 Experimental observations and computational challenges . . . 108
6.2 Sharp Negative deflections in claustral activity during REM 111
6.3 Coordination of activity between the left and right claustra during NREM and REM 113
6.4 Dynamics of left and right claustral leadership. 122
6.5 Coordination between REM activity in the claustrum and an ipsilateral midbrain nucleus 126
6.6 Discussion 134

7 Discussion and outlook 139
7.1 Modeling propagation in the turtle cortex 139
7.1.1 Routing and cortical function. 140
7.1.2 Outlook 143
7.1.2.1 Beyond turtles 143
7.1.2.2 Plasticity 145
7.2 Analyzing propagation in the lizard brain 148
7.2.1 Outlook 149
7.3 Final remarks 152

A Appendix: Multiple network instantiations 155

Bibliography 161

List of Figures

1.1 The action potential . . . 3
1.2 An internal frame to study brain computation . . . 4
1.3 Neuronal tuning . . . 5
1.4 Implementing logic functions with networks of neurons . . . 8
1.5 Evolution of circuits for avoidance and approach behaviors . . . 11
1.6 Phylogeny of lizards and turtles . . . 13
1.7 Computational tools in neuroscience . . . 16

3.1 Phylogenetic tree of vertebrate brains . . . 28
3.2 Subdivisions of the developing and adult vertebrate brain . . . 30
3.3 Circuitry of the turtle cortex . . . 33
3.4 Oscillatory signals during sleep . . . 41
3.5 Synfire chains . . . 48
3.6 Polychronous groups . . . 52
3.7 Alternative definitions of firing rate . . . 54

4.1 Sequence reactivation in the turtle cortex . . . 59
4.2 Network model . . . 61
4.3 Long-tailed distribution of connection strengths . . . 65
4.4 Network spontaneous activity . . . 67
4.5 Definition of followers . . . 69
4.6 Spiking sequences in the model . . . 70
4.7 Properties of model sequences . . . 73
4.8 Sequences extent . . . 74
4.9 Evolution of sequences at rates above turtle in vivo estimates . . . 76
4.10 Connectivity underlying activity propagation in the network . . . 78
4.11 Motifs of connectivity leading to postsynaptic spikes . . . 79
4.12 Connectivity underlying spiking sequences in the network . . . 81

5.1 Routing of firing sequences in the turtle cortex . . . 87
5.2 Clustering of followers by their activation . . . 89
5.3 Connectivity behind the activation of sub-networks . . . 91
5.4 Gating sub-network activations . . . 93
5.5 Effect of timing and strength on sub-network activations . . . 95
5.6 Coactivation of multiple triggers . . . 97
5.7 Properties of combination-specific and core followers . . . 98
5.8 Contextual activation of sub-networks . . . 99

6.1 Sharp negative field-potential events populate Claustrum REM activity 109
6.2 Sharp negative detection and properties 112
6.3 Sharp negatives make up beta spectral power during REM sleep . . 114
6.4 Spiking aligned to the downward phase of sharp negatives 115
6.5 Absence of bilateral coordination between claustral SWR during NREM sleep 116
6.6 Cross-correlation analysis of short-timescale bilateral coordination during REM sleep 118
6.7 Bilateral SN matching 120
6.8 Claustral spiking is most common during ipsilateral SN leadership . 121
6.9 Dominance switching between and within REM periods 124
6.10 Slow dynamics of leadership across REM periods 125
6.11 Changing dynamics of leadership throughout the night 127
6.12 Recordings in Imc reveal close coordination with ipsilateral claustral SNs 128
6.13 Coordination of bilateral claustra and Imc 130
6.14 Mutual exclusion between bilateral Imc nuclei in their correlation to a single claustrum. 131
6.15 Antagonistic spiking of single Imc units during REM 133
6.16 Effects of Imc unilateral lesion 135

7.1 Routing using sparse strong connectivity 140
7.2 Plastic effects on sequential propagation 147
7.3 Modeling inter-hemispheric competition 150

A.1 Clustering of followers in multiple networks 156
A.2 Gating through external input 157
A.3 Gating through external input (cont.) 158
A.4 Interactions of multiple trigger neurons 159

Acronyms and abbreviations

Name	Meaning
E/I	Excitatory/Inhibitory
MEA	Microelectrode Array
STDP	Spike Timing-Dependent Plasticity
SWR	Sharp Wave Ripple
IEI	Inter-Event Interval
LFP	Local Field Potential
REM	Rapid Eye Movement
SWS	Slow Wave Sleep
Imc	Nucleus isthmi pars magnocellularis (in reptiles and birds)
Ipc	Nucleus isthmi pars parvocellularis (in reptiles and birds)
OT	Optic Tectum
EEG	Electroencephalography
EMG	electromyography
PCA	Principal component analysis
LIF	Leaky integrate-and-fire
AdEx	Adaptive exponential integrate-and-fire
AMPA	α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid
NMDA	N-methyl-D-aspartate
GABA	γ-aminobutyric acid
EPSP	Excitatory postsynaptic potential
IPSP	Inhibitory postsynaptic potential
CPU	Central processing unit
HVC	High vocal center (in songbirds)
SST	Somatostatin
PV	Parvalbumin
VIP	Vasoactive intestinal peptide
V1	Primary visual cortex
LGN	Lateral geniculate nucleus in the thalamus
DVR	Dorsoventricular ridge (in reptiles)
CA1	First region of the cornu ammonis (hippocampus)
PWS	Propagating waves sleep (in zebrafish)
SBS	Slow bursting sleep (in zebrafish)
kHz	Kilohertz
ms	Milliseconds
µm	Micrometers

1 Introduction

1.1 Thesis overview

The brain is one of the most complex machines in the universe. It is responsible for generating actions, behaviors, and thoughts. Advances in experimental technology and computational techniques have enabled scientists to study the brain in ever greater depth, providing new insights into how neuronal systems process information and generate behavior. Computational neuroscience, in particular, investigates the activity of neuronal systems using computational tools such as simulations, mathematical theories, and computational analysis.

This thesis uses computational techniques to study how neuronal signals propagate and are processed within the reptilian brain. I combined computational modeling and analysis of electrophysiological data to examine the activity of the brains of turtles and lizards to shed light on the mechanisms behind the propagation of neuronal activity in these systems. The resulting insights, obtained from the brains of reptiles, may inform us of principles of general brain function and their evolution.

Most of this thesis is dedicated to understanding the fundamental principles underpinning the propagation of single action potentials in the visual cortex of turtles. I describe my results in Chapters 4 and 5. I took a bottom-up modeling approach to this question, which begins with the basic biophysical characteristics of networks of neurons constrained by experimental recordings (Chapter 4) and then progresses to increasingly complex predictions about the emergent behaviors of that system (Chapter 5).

In Chapter 6, I zoom out from single spikes in the cortex to study inter-area communication. I examine the activity of sub-cortical regions of the lizard brain during sleep. The data obtained by my experimental collaborators are local field potential (LFP) recordings, which capture the collective electrical activity of entire populations of neurons. I developed computational methods of analysis to examine this electrophysiological data recorded across multiple regions in the two hemispheres of the lizard brain.

Overall, my findings offer insights into some of the processes and mechanisms driving the transmission and communication of neuronal activity within the reptilian brain.

In the remaining sections of this chapter, I provide an introduction and motivation to the three key elements that form the subject of this thesis:

- why we should study the mechanisms of propagation of neuronal activity (section 1.2),
- what is interesting about the reptilian brain (section 1.3), and
- how computational techniques are useful in neuroscience (section 1.4).

For each question, I provide a broad view of the field and delineate the particular focus of my research. I further detail the background for these topics in Chapter 3, where I describe relevant literature about the reptilian brain, summarize our current understanding of its physiology and evolutionary context (section 3.1), and outline the main theories and models of signal propagation in the brain (section 3.3). The main results of this thesis are presented in Chapters 4, 5, and 6. Each results chapter begins with a brief introduction that outlines the specific research question addressed, followed by a detailed discussion of the results obtained. Finally, in Chapter 7, I summarize all of my results, put them in the context of the ideas that I expose in this introduction, and discuss future directions.

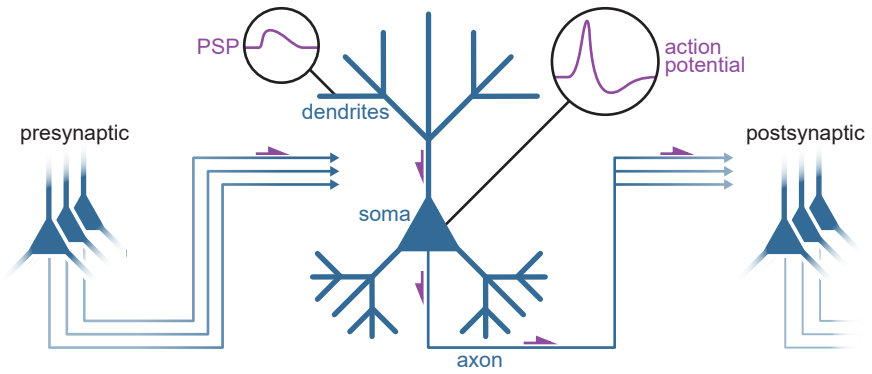

Figure 1.1: The action potential. Simplified schematic of the fundamental unit of communication between neurons. Single neurons (middle) receive synaptic connections from multiple presynaptic partners and project to multiple postsynaptic partners. Electrical impulses (purple arrows) travel through the axons of presynaptic partners and arrive at the dendrites of the target neuron, where they elicit deflections in the membrane voltage potential (a postsynaptic potential, PSP). When sufficient inputs have been integrated in the membrane of the neuron, they will trigger an "action potential", also called a "spike", an all-or-nothing electrical impulse that will then travel through the axon of the neuron to reach a new set of target neurons.

1.2 Why activity propagation? An internal frame to study brain computation

Studying the mechanisms underlying the generation and processing of electrical activity in the brain is a fundamental aspect of neuroscience and can help us understand the neuronal basis of perception, thought, and action. In the brain, there are many kinds of electrical signals, with the most well-known type being the "action potential", which is a fundamental unit of neuronal activity. Action potentials are thought to underlie mental representations and cognitive processes, making them essential for understanding how the brain works. An action potential is a brief electrical impulse that travels along the membrane of a neuron, initiated by a threshold depolarization event in response to a series of synaptic inputs (Fig. 1.1). Indeed, neurons are connected in complex and recurrent networks, and the activation of a single neuron is often the consequence of the concerted activity from multiple presynaptic partners. As such, an action potential is also a fundamental unit of neuronal communication [3]. Thus, we can assign two forms of meaning to the action potential: as a unit of representation and as a unit of communication.

In neuroscience, there are different approaches to studying the brain, each stemming from different perspectives. This thesis focuses on studying neuronal activity in relation to the internal mechanics of the brain, which I will call in this section an "internal" frame, rather than in relation to its environment, which I will call an "external" frame (Fig. 1.2). An external frame investigates how the brain processes sensory information or controls movement and behavior and is often the focus of

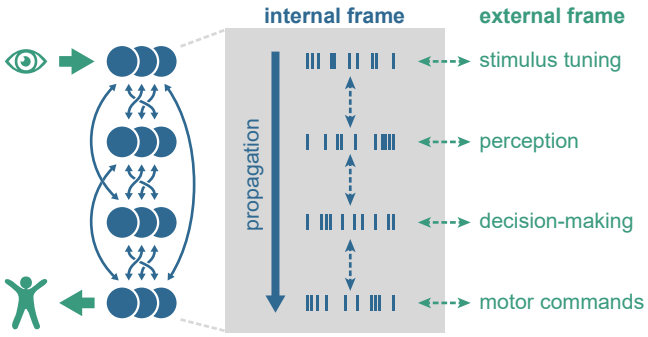

Figure 1.2: An internal frame to study brain computation. Simplified schematic of a highly recurrent system of neurons (blue circles) receiving sensory inputs and producing motor commands. In the process, the system internally generates a series of patterns of activity (blue ticks). One may approach understanding the brain by detecting correlates (dashed green arrows) between some of those neuronal activity patterns and external elements such as environmental measurements. However, only a tiny fraction of all neurons interact directly with the environment. Instead, the majority of neurons that make up the brain propagate activity to and from other neurons with varying degrees of reliability. The recurrent connectivity and properties of different groups of neurons implement mappings from an input pattern of activity to an output pattern of activity (dashed blue arrows). The approach of this thesis (gray shading) is to study brain computation from the perspective of the mechanics of signal propagation. Put simply, how does a network of neurons reliably transform an activity pattern into a different one?

modern neuroscience research. On the other hand, an internal frame focuses on the intrinsic and spontaneous patterns of activity and how these are transformed within the brain by examining how neurons and networks of neurons communicate with each other. Although I will here contrast these frames, they are not mutually exclusive. Most neuroscience research mixes them with different degrees of emphasis, as they offer complementary understandings of the functioning of the brain.

An external frame: the brain and its environment

An external frame to studying the brain focuses on the relationship between brain activity and its environment. This includes examining how the activity of individual neurons or networks of neurons changes in response to different stimuli, such as sight, sound, or touch. By studying the response of neurons to external stimuli, this frame has produced valuable insights into the elements underlying perception, cognition, and behavior. For instance, the classical work by Hubel and Wiesel eminently addressed the external frame of neuroscience [4]. In their experiments, they focused on the relationship between brain activity and external visual stimuli by studying the response of single neurons in the cat visual system to different stimuli types, such as lines, edges, and colors (Fig. 1.3A). They were able to gain insights into how the visual cortex might encode visual information. These experiments have led to an understanding that the visual system is organized

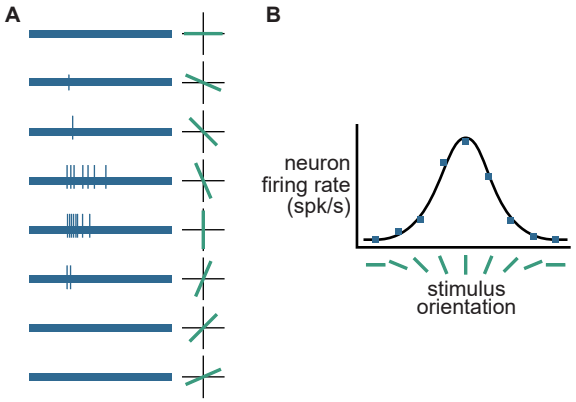

Figure 1.3: Neuronal tuning. **A.** Single unit responses (left, blue) to shining a rectangular light spot at various angles (right, green) to the contralateral retina of an anesthetized cat. Modified from [4]. **B.** Schematized tuning curve for a neuron that codes for stimulus orientation.

hierarchically, with lower-level neurons responding to simple features such as edges and angles and higher-level neurons responding to more complex features such as objects and faces. This external frame is most clearly exemplified by the concept of neuronal tuning curves (Fig. 1.3B): graphs that show how the firing rate of a neuron (the number of times it fires an action potential per unit of time) varies as a function of some property of the external stimulus that is presented to the animal. Tuning curves are among the most common tools for dissecting brain activity in modern neuroscience since they can capture the relationship between brain activity and externally measurable quantities, like the angle of a line in the visual field or an animal's location in a maze.

Note that studying brain activity in relation to external objects easily extends to involve concepts at the cognitive level. For instance, a neuron reliably activating upon the presentation of a particular visual stimulus may suggest that this neuron is part of the network representing such stimulus and that it may underlie the subjective experience of the animal perceiving the stimulus [4]. On the other hand, a neuron ramping up its firing rate before one particular motor movement but not another in a two-choice task may suggest that the neuron is involved in representing a choice and that its surrounding network may be involved in the process of decision-making [5]. This extension into cognition gives us a chance to provide a neurobiological grounding to these concepts.

The external frame is tightly linked to the idea of a neural code, which is some form of cipher that the brain uses to represent the external world. Although popular, the neural code has been criticized as misleading and potentially limiting the type of questions that we ask [6]. Most importantly, the idea of neural code fully ignores the dynamic, recurrent, and distributed causal structure of the brain by focusing on individual electrical signals and not considering how they relate to each other. This is particularly important when studying neurons embedded in recurrent networks

without understanding the origin of their inputs or the causal role of their outputs. Indeed, a neuron in the visual cortex may activate reliably for one particular visual stimulus, but it is never truly exposed to the direct action of light. Instead, such a cortical neuron receives inputs only from other neurons in the retina, thalamus, or other brain regions, so its exposure to the original stimulus is only indirect. If this cortical neuron reliably activates for a particular stimulus orientation, what does that tell us about the neurons that provided its inputs and together led it to activate? In summary, an external frame trades studying causality and recurrency within the brain for a simplified view of the indirect relationships between neurons and the environment.

An internal frame: interactions within the brain

A complementary perspective to studying the brain, which this thesis focuses on, addresses the internal workings of each of its parts and how they interact with one another, independently of elements external to the brain itself. This internal frame can involve looking at how neurons communicate with each other, the dynamics of that communication, the different types of signals that propagate within the brain, intrinsic and spontaneous patterns of activity, and the mechanisms by which the brain self-organizes. For example, the study of the self-organization of networks of neurons can provide insights into how network structures form and the types of neuronal activity that they can generate [7]. In self-organized networks, individual neurons are able to adjust their connections with one another based on the inputs they receive, allowing the network as a whole to develop and adapt to changing conditions [8].

The internal frame that I describe here overlaps with but is not strictly interchangeable with that of an "inside-out" perspective on the brain as introduced by György Buzsáki [9]. An "inside-out" perspective generally focuses on the internal workings of the brain, independently of external stimuli, and understanding the brain's intrinsic properties and patterns of activity. Furthermore, it describes the brain as a self-organized system that is not passive when processing external information but instead generates constant internal patterns of activity and predictions for its actions. I use the term "internal" to underscore the understanding of the causal relationships and processes through which different elements in the brain interact. That is, I am interested in the processes that relate different elements within the brain, while the "inside-out" perspective takes a broader view and emphasizes the role of spontaneous activity.

Marr's levels of analysis of information processing systems is another popular framework to classify neuroscientific questions. The external frame that I described above most strongly addresses the computational and algorithmic levels, where we try to understand what the goal of the system is and what representations it relies on. On the other hand, the internal frame addresses questions mainly at the implementational and algorithmic level: it focuses on what neuronal structures

and activities form the system and how these activities are modulated through the system dynamics and processes.

Differentiating between an internal and external framing of neuroscientific questions is not a purely philosophical or theoretical endeavor. It impacts the methods that we develop and, thus, the type of insights that are within our reach. As an example, demixed-PCA (Principal Component Analysis) is a method to analyze neuronal signals that tries to find the projection of the data that best represents external parameters such as stimuli or rewards [10]. That is, in demixed-PCA, elements external to the brain are used to algorithmically construct an interpretation of the activity inside the brain. Therefore, demixed-PCA enables us to find potential representations of external variables in high-dimensional neural data that could not be otherwise inspected. By contrast, methods of analysis might use brain activity as the reference for interpreting other brain activity. For instance, time-warping algorithms can adjust multiple brain signals to one another in order to correct for uncontrolled temporal variability in experiments [11]. As a result, time-wrapping algorithms are able to reveal precise spike-time patterns even when these were decoupled from behavior and thus invisible to an external observer. In Chapter 6, I describe the development of methods of analysis that also use a temporal comparison of brain signals recorded in multiple areas and how these methods helped me and my experimental collaborators detect and quantify the striking temporal precision with which those areas coordinate.

Reliable propagation underlies computation

The brain is not just one of the most complex systems in nature; it is also the most powerful computing device we know of. Understanding how electrical signals are propagated within the brain is a fundamental question of the internal frame that can shed light on the mechanisms underlying brain computation. When a network propagates activity in a reliable way, neuronal communication essentially performs a mapping from one pattern of neuronal activity onto another pattern of neuronal activity. If the same network can flexibly and reliably switch between multiple of these mappings, the network of neurons can implement "functions" in the mathematical sense: mappings from input variables to output variables. These mathematical functions can then act as building blocks that, when combined with different types of dynamics, such as cooperation or competition, implement complex computations. For instance, theoretical studies have leveraged reliable signal propagation to implement logical operations using groups of neurons. Using firing rates and altering the balance of excitation and inhibition, models have shown how multiple interacting pathways of propagation may implement logic gates, switches, and even memory units [12, 13] (Fig. 1.4). Thus, the reliable and flexible propagation of neuronal signals is a crucial element for carrying out brain computations and understanding this process can provide insights into how the brain processes information.

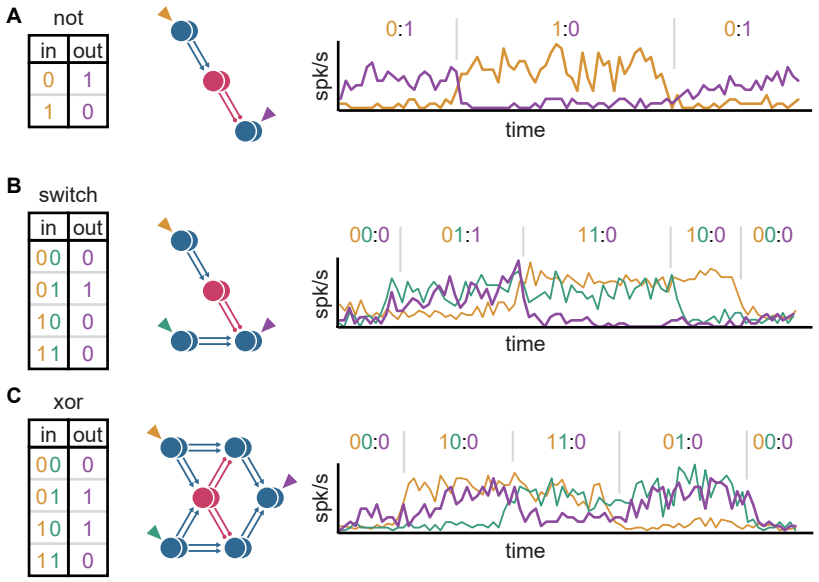

Figure 1.4: Implementing logic functions with networks of neurons. **A.** Left: Truth table of a standard "not" logic function. Each row is a specific mapping from an input pattern to an output pattern. Middle: connectivity motif implementing the logic function. Blue (red) circles indicate an excitatory (inhibitory) population. Colored arrowheads mark the populations whose activations represent the input and output variables of the function. Right: Mean firing rates of the input and output populations over time as the activation of the network express each of the mappings of the truth table (indicated above). A high firing rate encodes a 1, and a low a 0. **B.** Same as **A** but for a "switch gate". Activating the yellow population gates on or off the signal from the green population. **C.** Same as **A** but for an "xor" logic function. The output of the network is 1 only if both inputs differ. Note that network motifs are built by combining simpler ones (**C** from **B**, **B** from **A**). Proper functioning of these logic functions requires the reliable propagation of activity between each pair of populations. Modified from [13].

One example of research that addresses the propagation of action potentials is the theory of synfire chains, which I describe in more detail in section 3.3. The key issue that synfire chains address is that individual synaptic inputs are typically very weak, meaning that they can not cause a neuron to fire an action potential on their own. In order to overcome the weak individual inputs, the synfire chain theory proposes that many of them converge onto a single neuron, summing together until the postsynaptic neuron reaches its firing threshold [14, 15]. Thus, it makes the specific prediction that, for information to be transmitted from one group of neurons to the next, synapses must be organized in a convergent-divergent fashion and that the sequential activation of groups of neurons must occur in a synchronous manner, with the neurons in each group firing at nearly the same time. This synchronous firing pattern ensures that the signal is propagated accurately along the chain, allowing for the transmission of information in a highly reliable manner.

Other theoretical work has shown that, with synfire chains as the core mechanism of propagation, networks of neurons can also implement flexible gating [15]. Therefore, synfire chains are, in principle, capable of flexible propagation and thus of implementing computations such as those that I describe above (Fig. 1.4). Related theoretical work has taken the idea of synfire chains from the internal into the external frame by showing how synfire chains may encode concepts and how their synchronous activations may be used to bind multiple concepts together in working memory [14, 16].

In summary, since neurons are rarely directly exposed to external stimuli and fundamentally talk to each other, we may develop a more comprehensive understanding of how the brain works by studying how this internal communication happens. This thesis takes an internal frame to neuroscientific questions, aiming to understand how circuits of neurons work by focusing on how neuronal signals propagate within the brain. Complementing an external approach that describes brain activity in terms of its relationship to the external environment, the internal frame might inform us of how computations might be implemented.

1.3 Why reptiles? An evolutionary perspective in neuroscience

The process of evolution, through random mutation and natural selection, has produced a diversity of biological systems. The evolution of metazoans over 600 million years has produced a vast number of adaptations to the dynamic environment of Earth. It has led to an immense richness and variety of neuronal systems. Taking an evolutionary approach to study the brain involves taking a comparative view across a variety of species [17]. This view can inform our understanding of brain function and structure in several ways: by differentiating implementation from function and by providing logic to the organization of the brain and behavior [18].

Comparing the mechanisms of neuronal function across multiple organisms allows us to identify those that are the most consistently conserved. In other words, by studying shared properties across species, we can examine the essential, and therefore conserved, principles for the proper operation of neuronal systems. Alternatively, the comparative method can also identify principles at the algorithmic level. Throughout millions of years of evolution, nature has repeatedly explored multiple solutions to the same set of problems, from navigation to perception. We can only identify different mechanisms that perform the same function by examining mechanisms in diverse biological systems. This makes it possible to study algorithmic solutions rather than implementations of those solutions. Importantly, it can reveal which aspects of biological processes are necessary for function.

By studying the neuronal systems of a wide range of species, we might gain insight into the common features shared across different animals and how these features have been modified and adapted to fit the specific needs of each species. For instance, different species, such as bats and finches, may have evolved comparable solutions for tracing three-dimensional flight routes, but these solutions might be implemented differently in the brains of each species. An evolutionary approach thus helps us to identify the general principles that govern the organization and function of neuronal systems. Overall, it is a framework for understanding the basic principles that underlie the functioning of neuronal systems, as well as how these principles might have been implemented differently.

Understanding brain evolution can also help us define a conceptual taxonomy of mechanisms that can explain not only how a function is implemented but also what functions are being implemented. Although from the perspective of the researcher, defining "what" must come before investigating "how", from an evolutionary perspective, both questions are addressed at once. The ultimate purpose of neuronal systems is to direct animal behavior in their environment. Several tools of comparison, from transcriptomics to anatomy, may help us identify what components of these systems were added, removed, and modified throughout evolution [20, 21]. If we can infer those changes, we gain insight into the new behavioral problems that need to be solved and the spectrum of possible solutions.

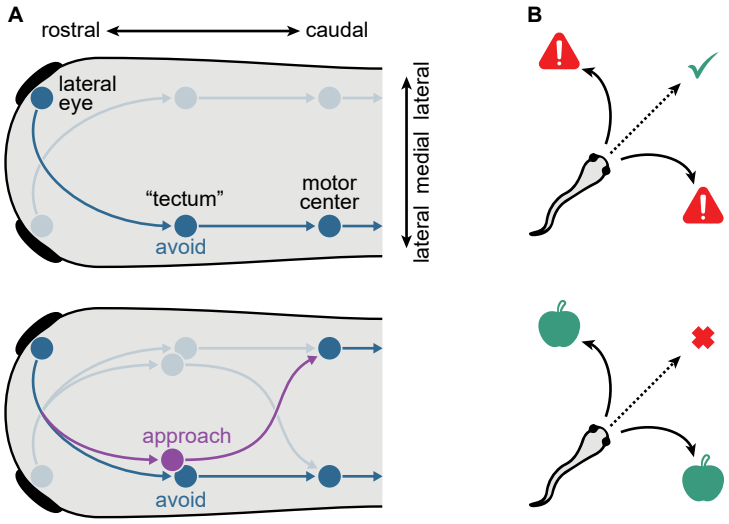

Figure 1.5: Evolution of circuits for avoidance and approach behaviors. **A.** Schematic of hypothetical evolution of circuits for avoidance and approach behaviors. Top: Avoidance in a hypothetical chordate ancestor. Activity from photosensitive cells in the lateral eye would propagate contralaterally to an early version of the tectum, which then propagated ipsilaterally to neurons controlling locomotion, resulting in a turn away from aversive visual stimuli. Bottom: Visually-guided approach appears in early vertebrates. The tectum expanded in parallel to retinal evolution, differentiating to include a rostral region that projected contralaterally, resulting in the possibility of turning towards attractive visual stimuli. **B.** Top: Averaging the direction of multiple threats (solid line) is an effective strategy to avoid all of them. Bottom: By contrast, averaging over multiple rewards is undesirable, so an approach behavior requires a form of attention: winner-take-all-like dynamics that select one stimulus and suppress others. Modified from [19].

Modern systems neuroscience has sought to identify neural correlates of cognitive processes as one of its primary objectives. With evolution as the focal point, we can reevaluate cognitive concepts such as attention and decision-making and redefine them in the natural context that gave rise to them, with explicit links to neuroanatomical and neurophysiological data [18, 19]. For instance, we might come to the conclusion that multiple brain areas evolved to produce diverse but related behaviors that we have classically combined into a single cognitive concept such as "attention" [19, 22]. Indeed, studies from the tectum of lampreys, an extant link to early vertebrates, revealed the presence of multiple sub-circuits that project ipsi- and contralaterally, likely producing avoidance and approach behaviors Fig. 1.5A [19, 21]. Contrary to avoidance behaviors, approach behaviors require the predominance of one stimulus over others through competing winner-take-all dynamics, likely involving contralateral inhibition Fig. 1.5B [22]. Thus, by studying the circuitry of this and other species, we may arrive at a mechanistic definition of different sub-types of attentive behavior and how they were built on top of existing circuitry. In a way, the evolutionary approach illuminates the relationship between the brain and behavior in a bidirectional manner: we may use behavior to assign meaning to circuits and use circuits to define how behavior is organized.

In summary, an evolutionary approach to neuroscience can guide our understanding of both the structure and function of neuronal systems. It is thus a pre-requisite for the comparative approach in neuroscience that the neuroscientific community studies the function of neuronal circuits in a diverse set of species [17]. This thesis focuses on reptiles, specifically turtles and lizards. While mammals, particularly rodents and monkeys, are more commonly used in modern neuroscience, reptiles offer a number of unique advantages that make them interesting subjects for research.

Advantages of reptiles as model organisms in neuroscience

In terms of their evolutionary history, reptiles diverged from modern mammals over 300 million years ago (Fig. 1.6). Around 320 million years ago emerged the first amniotes, with the capacity to produce amniotic eggs, resulting in the invasion of land [21]. Amniotes went on to become one of the most successful groups of animals on Earth, eventually giving rise to mammals, non-avian reptiles, and birds. Although birds descended from ancient reptiles, they developed physiological characteristics (such as endothermy) similar to modern mammals, filled similar behavioral niches, and, as increasing literature suggests, developed comparably complex cognitive functions ([23]). This puts extant non-avian reptiles, such as lizards and turtles, at a very interesting cross-point from the point of view of evolutionary neuroscience. They provide us with a point of reference that connects extinct ancient amniotes, modern mammals, and birds. Studying reptiles thus can provide insights into the evolutionary origins, conserved principles, and divergences between these lineages.

Differences in the architectural complexity of brain areas can also be an advantage when addressing basic neuronal mechanisms of information processing. Evolutionary

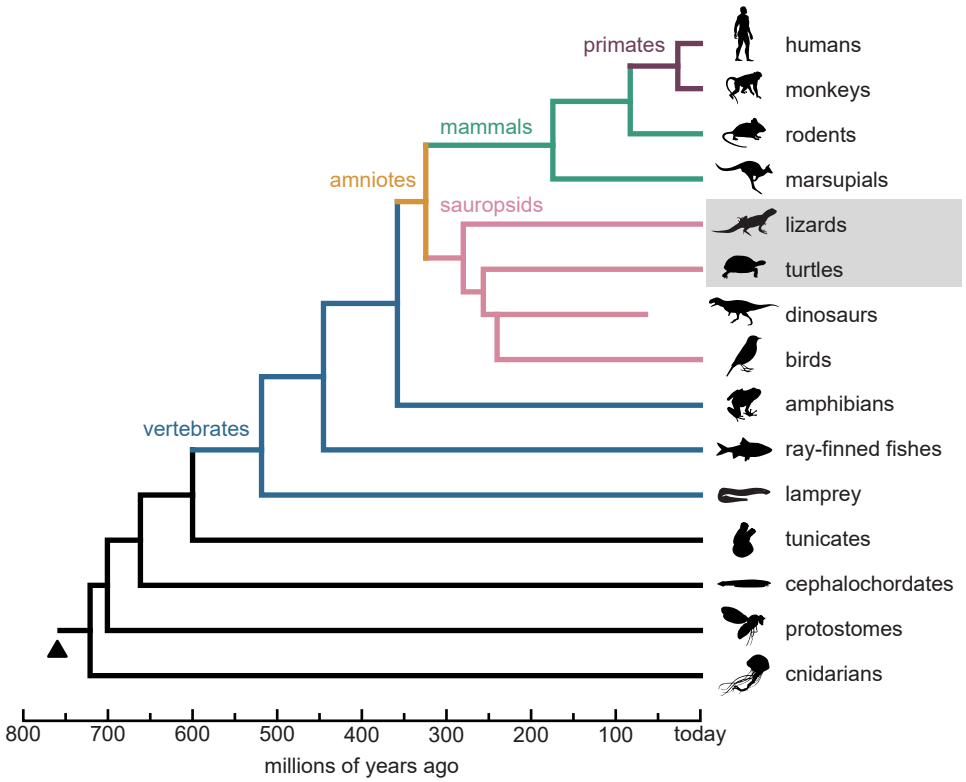

Figure 1.6: Simplified phylogeny of lizards and turtles. Simplified phylogenetic tree of the species studied in this thesis (gray shading). Branches represent the approximate divergence of lineages. Please note many branches are omitted for clarity and that the vertical ordering of the branches is arbitrary. The arrowhead at the tree root indicates the earliest estimate of the first neuronal systems over 700 million years ago. Early amniotes appeared approximately 320 million years ago, giving rise to mammals, reptiles, and birds. Dinosaurs spanned between 230 and 65 million years ago. Aspects of convergent evolution between mammals and birds have resulted in common niches of physiology and complex cognitive capacity [21, 23, 24]. Modern lizards and turtles sit at a cross-point, giving us a reference that bridges ancient amniotes, modern mammals, and birds. Data from [18, 21, 25]. The silhouettes of example species are from `http://phylopic.org` under Creative Commons CC0 1.0 licenses.

pressure has been expressed differently in the brain of reptiles compared to mammals, resulting in certain architectural aspects that are, at least at first sight, simpler or more complex across the different lineages. For instance, while both turtles and rodents possess visual cortices, the latter have developed specialized structures, with up to 6 differentiated layers of neurons and specialized organizational principles, such as retinotopy. On the other hand, the visual cortex of turtles is made of 3 layers and lacks some of these organizational features [25, 26]. Perhaps, part of the reason for this difference in visual cortices can be explained by the expansion of the optic tectum in reptiles. The optic tectum and its mammalian homolog, the superior colliculus, are important non-cortical brain regions involved in the brain response to visual stimuli. It is found in all vertebrates but is more complex in the reptilian and bird lineages compared to mammals. For example, the optic tectum in reptiles is usually recognized to have 14 layers of cells, and the avian optic tectum has 15 layers, while the superior colliculus of mammals has only 8 layers [27]. The differences in architectonic features between the reptilian and mammalian optic tecta and cortices thus suggest different evolutionary specializations of common solutions to visual processing.

Beyond the comparative interest of reptiles, certain species also display a number of unique characteristics that make them interesting subjects for neuroscientific study. These characteristics include experimental accessibility in terms of physiology and behavior.

Reptiles are cold-blooded animals, which means their physiology is able to tolerate a wide range of body temperatures while still maintaining function. Additionally, certain species, such as turtles, are aquatic animals and have developed robust mechanisms to support reduced oxygen levels. This combination makes their neural tissue especially amenable to certain experimental setups such as ex vivo preparations [28, 29] (see section 3.1.1.1). Indeed, electrophysiological studies often require working on thin slices of brain tissue, where local and mid-range connectivity might be severely damaged. By contrast, ex vivo preparations preserve this connectivity enabling different types of experiments, such as those that I introduce in Chapter 4. Experimental studies can also take advantage of anatomical differences. For instance, the claustrum, a thin sub-cortical region of difficult access in rodents, presents a different shape and location in the lizard brain that enabled the targeted electrophysiological recordings that I describe and analyze in Chapter 6 [30, 31].

Certain reptiles can provide experimental accessibility to interesting or specialist behaviors. For instance, lizards, like mammals, display two alternating phases of sleep in terms of their electrophysiological and behavioral signatures: rapid eye movement sleep (REM) and slow-wave sleep (SWS) (see section 3.2) [32]. In Chapter 6, I discuss several findings on the nature of neuronal activity in the sleeping lizard *Pogona vitticeps* during REM sleep. These discoveries were possible thanks to these lizards showing extremely regular sleep cycles, where REM to SWS transitions happen approximately every 80 seconds [32]. In addition, these lizards present a sustained sleep state for 9-11 hours with little movement or wakeful interruptions

that could affect the recordings and analysis. This regularity and the high number of REM/SWS cycles per recorded night contrasts with the sleeping patterns of humans, where a full cycle can last over 90 minutes, or mice, who typically display irregular sleep bouts that go from a few minutes to an hour, interleaved with frequent wakeful periods [33]. In summary, by focusing on a sleeping specialist, we were able to study a behavior common to many species in detail and with abundant data.

In conclusion, reptiles are remarkable neuroscientific research models. Their evolutionary history, simple neuronal systems, unique physiological characteristics, and interesting behaviors make them valuable study subjects for neural function and behavior mechanisms. By studying reptiles in greater depth, we can gain valuable insights into the evolution and operation of neuronal systems. By examining the brains and behaviors of numerous species, we can gain a deeper understanding of the principles that govern the functioning of neuronal systems and the ways in which these principles have been adapted and modified throughout natural history.

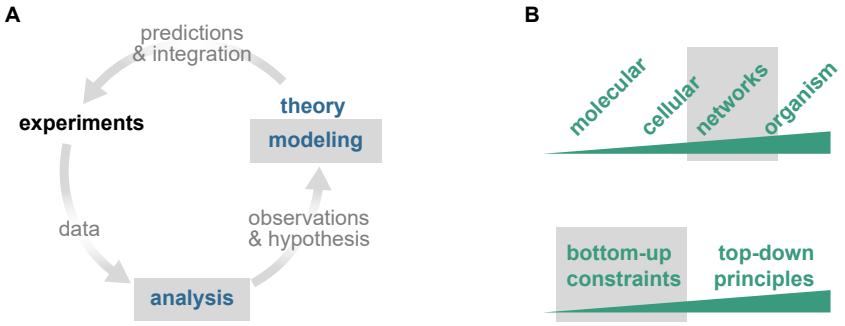

Figure 1.7: Computational tools in neuroscience. **A.** Computational neuroscience offers three broad categories of tools (blue): methods of data analysis, mathematical theories, and computational modeling. Analysis methods allow us to digest complex experimental data into observations and formulate hypotheses. These become formally described in the form of mathematical theories and computational models. Through these formalisms, we integrate knowledge and produce new predictions and refined hypotheses that can guide further experiments. **B.** Defining a model requires fixing an abstraction level (top) and determining the core observations that the model embodies (bottom). These observations might take the form of measurable biological constraints or abstract principles of how the brain might work. All models require both elements but, depending on how these are emphasized, results in bottom-up or top-down models. The present thesis covers (gray shading): biologically-constrained modeling of networks of neurons (Chapters 4 and 5) and methods of data analysis (Chapter 6).

1.4 Why a computational approach? Modeling and data analysis in neuroscience

The word "computational" in computational neuroscience refers to the use of computational and mathematical techniques to study the brain and its functions (Fig. 1.7A). This involves using mathematical theories and computer models to describe or simulate the complex interactions between neurons and the networks they form. Computational neuroscience also involves the development of algorithms and analytical tools to analyze and interpret large amounts of neural data. This thesis deals with both aspects: modeling in Chapters 4 and 5, and data analysis in Chapter 6. Here I briefly describe the benefits of using computational tools in neuroscience and the particular approach of this thesis.

Modeling: goals

Modeling plays a critical role in neuroscience as it allows us to construct simplified representations of complex biological systems, which can aid in understanding the underlying mechanisms and behaviors of these systems [34]. Computational models can be used to test hypotheses and make predictions about complex neural processes that would otherwise be difficult to investigate experimentally. For

instance, we may test the effects of different parameters or variables on neuronal systems for which experimental tools like pharmacology or genetics do not exist yet. By systematically varying different factors within a model, we can gain insights into how the system responds to different conditions, which can help identify key mechanisms underlying neural function.

Beyond predictions, models provide another valuable purpose in neuroscience: they can be used to integrate knowledge. Experimental research typically addresses single questions in highly controlled settings. By using models, we can integrate multiple experimentally observed phenomena into a cohesive picture. In this aspect, mathematical theories and computational models are a form of formalizing our understanding of brain function to detect inconsistencies or highlight gaps in our knowledge.

In Chapter 4, I use a network model of the turtle dorsal cortex to first recapitulate experimental findings and then explore the system beyond what was experimentally accessible. Physical limitations on the placement of electrodes and the properties of the tissue make the monitoring of the entire biological network activity experimentally unfeasible. In a computational model, on the other hand, I could study signal propagation considering the entire population. In addition, I was able to simulate network states that might be present in the behaving animal, but that could not be easily explored in an experimental setup.

Modeling: levels of abstraction

Selecting the appropriate level of abstraction for a model is a critical step, as it determines the level of detail and complexity that the model will capture. The level of abstraction may range from detailed biophysical models that incorporate molecular, cellular, and network properties to simplified models that focus on a few key features or behaviors (Fig. 1.7B top). The main challenge in selecting the level of abstraction is striking a balance between complexity, traceability, and interpretability. On the one hand, detailed biophysical models can capture the nuances of neural processing and provide mechanistic insights into how the brain works. On the other hand, these models may require large amounts of data to parameterize and are often computationally expensive to simulate. Furthermore, even when fitting and simulation are possible, large detailed models risk becoming highly complex, with emergent properties growing unmanageable, which might make it challenging to extract meaningful insights from them.

As described in section 1.2, my modeling work focuses on how groups of neurons interact with one another at the network level. That is why I employ networks of leaky integrate-and-fire neurons, which are neuron models that ignore sub-cellular structures and instead treat each neuron as a single compartment [35] (see Chapter 2). This trade-off in single-neuron complexity allowed me to simulate how a large number of relatively simple neurons interact together. I thus chose to focus on the network properties and mechanisms that are involved in reliable signal propagation.

Modeling: bottom-up and top-down

Another related critical aspect in modeling is that of the choice of assumptions. In this sense, the most common classification of modeling approaches differentiates between "bottom-up" and "top-down" approaches [34] (Fig. 1.7B bottom).

A top-down approach is driven by principles of brain function or organization and typically involves starting with a general description of the behavior of a system and then constructing a model that can reproduce that behavior. This approach is often used when the underlying components are poorly understood, but the overall behavior or principle of function is. For instance, a top-down approach to sensory modeling may start by determining that sensory neurons must maximize information transmission to downstream areas and then using that principle to derive the properties that those neurons must follow.

A bottom-up approach to modeling is driven by biological constraints and typically involves constructing models by starting with individual components and building upward to more complex systems. This approach relies on understanding the properties and interactions of the underlying components and then using that knowledge to build models that accurately represent the behavior of the system as a whole. In other words, the focus is on constructing models from the ground up, exploring the emerging interactions between multiple components. A classic example of a bottom-up modeling approach is the Hodgkin-Huxley model, which explains how individual neurons generate and propagate electrical signals. This model is based on a detailed understanding of the biophysical properties of the neuronal membrane and simulates the behavior of ion channels to explain how an action potential works at the level of a single neuron (Fig. 1.1).

In practice, all models include elements from both top-down and bottom-up approaches, although they may be emphasized differently. Top-down approaches describe their modeling assumptions as "principles" of neuronal organization, while bottom-up approaches refer to biological "constraints". However, determining what a constraint or a principle is may not be straightforward when one considers partial experimental evidence. For instance, there is experimental evidence for the existence of assemblies of neurons and for activity-dependent plasticity in the cortex, but their exact properties are still objects of active research [36], meaning models that rely on them must make assumptions based on an incomplete picture. The extent to which these assumptions can be called constraints or principles is often non-obvious.

The model that takes the focus of the majority of this thesis (Chapter 4) follows a predominantly bottom-up approach. It is quantitatively constrained by experimental measurements made by collaborators from the turtle dorsal cortex in terms of neuronal populations, densities, connection strengths, and single-cell properties. By modifying some of these elements and studying alternative models, I was able to pinpoint which element played a pivotal role in the reliable propagation of cortical spikes.

Data analysis

Modern neuroscience experiments generate large amounts of data, often with a high degree of variability and noise. Computational methods of data analysis allow us to extract meaningful patterns and relationships from this data, which can help us to draw conclusions about how the brain functions. This is a critical component of neuroscience research, as it enables us to digest data into observations that can guide theory as well as to search for the predictions resulting from those theories.

In Chapter 4, I rely on simple methods of data analysis to quantify several properties of the turtle dorsal cortex in order to constrain a computational model. I then repeated the experimental protocols in my model and employed analysis methods to detect and quantify the reliable activation of neurons in the simulation results. These analyses allowed me to validate my model, showing that it could recapitulate the key experimental observations before exploring potential predictions under different manipulations.

Finally, in Chapter 6, I detail the new analysis methods that I developed to process, quantify and visualize recordings from multiple sub-cortical areas of the lizard brain. Relying on these advanced computational methods, I was able to process large data sets of recorded local field potential (LFP) and find the fine temporal coordination of neuronal signals across multiple areas of the lizard during sleep.

In conclusion, computational neuroscience plays a crucial role in understanding the complex mechanisms of the brain. It includes theory, modeling, and analysis tools that provide insights into neuronal function. Models can integrate knowledge, make predictions, and explore these systems beyond experimental accessibility. Methods of data analysis can extract meaningful patterns and relationships from large amounts of data. This thesis exemplifies both techniques, with a network model of the turtle dorsal cortex and data analysis methods of cross-area activity in the lizard brain. Neuroscience is an interdisciplinary field, and computational techniques will continue to make significant contributions to our understanding of the brain and its function.

2 Materials and methods

The target of this thesis is to investigate the network mechanisms of signal propagation in the brains of reptiles using computational methods. These methods include modeling through computer simulations and analysis techniques to process electrophysiological data collected by collaborators.

In Chapter 4, I construct a large-scale network model of the visual cortex of turtles. This model is constrained by quantification from electrophysiological data. I then simulate the resulting network model repeatedly under different conditions and analyze its output. In Chapter 5, I generate and analyze a new set of simulations of the same model with additional manipulations.

In Chapter 6, I describe several algorithms that I developed to analyze electrophysiological data recorded from sub-cortical areas of the lizard brain. Chapter 6 contains all the methodological descriptions of these computational methods which were my main contribution to the study published in [2].

In this chapter, I add the technical details for those results, including a description of the data, code, and mathematical model that I used.

2.1 Electrophysiological data

All the data described in this thesis were collected by my collaborators. All the analyses described in this thesis were performed by me. The analyses and numerical simulations were executed on one or multiple compute nodes with processor Intel(R) Xeon(R) Gold 6152 CPU @ 2.10 GHz. Any partial results or cached results were stored using the HDF5 format through the Python package *pandas*.

Turtle cortex data

The data (described in Chapter 4) were originally collected by my collaborator Dr. Mike Hemberger and published in his academic article [37]. From this data, I used quantifications already performed by Dr. Mike Hemberger in the form of tables of EPSP amplitudes, tested connections between neuron pairs, and estimates of adaptation indices. In addition, I performed quantifications from membrane voltage traces that had been recorded under current injections in order to estimate the membrane time constant and membrane capacitance of neurons in the turtle cortex (4.2B), as well as membrane voltage traces under spontaneous activity to estimate membrane variance and mean (4.4A). All of these were the result of whole-cell

patch-clamp recordings performed by Dr. Mike Hemberger and deposited in a data repository: https://doi.org/10.6084/m9.figshare.19763017.v1.

In addition, I used several other estimates of neuronal densities, connection probabilities, and mean firing rates that had been originally reported in [37].

All fitting procedures to obtain model parameters were performed using the standard methods of maximum likelihood estimation and least squares as implemented in the Python package *scipy*.

Lizard subcortical LFP

The data (described in Chapter 6) were collected by my collaborator Dr. Lorenz Fenk and published as part of an academic article of which I am also an author [2]. From this data, I used the raw LFP recorded by Dr. Lorenz Fenk using either 32-channel silicon NeuroNexus probes sampled at 32 kHz or Neuropixels 1.0 probes sampled at 30 kHz. Dr. Lorenz Fenk performed histological revisions to confirm the location of the probes in the claustrum and Imc of *Pogona vitticeps* lizards.

In addition, I used spike times resulting from spike sorting using IronClust or Kilosort2 and manual curation performed by Dr. Lorenz Fenk.

The processing of this data presented several challenges that required novel computational methods. I describe these challenges and the solutions that I developed in Chapter 6.

As part of my quantifications of lizard subcortical LFP, I performed several standard statistical tests. These statistical tests are implemented in the standard Python package *scipy* (v.1.6.2). Often, due to the long duration of the recordings and large resulting datasets, P values resulted in zero values. These zeroed P values correspond to values too small to be computed with a standard float64 ($< 510^{-324}$) in our computing machines.

2.2 Code availability

All the code that went into the making of this thesis was cleaned up, organized, and packaged into pip-installable Python packages. These packages were archived with the corresponding publications and are publicly available in an effort to increase code quality, reproducibility, and sharing in neuroscience [38]:

- **Model simulation and analysis code** (described in Chapters 4 and 5) was deposited at https://github.com/comp-neural-circuits/tctx.
- **LFP analysis code** (described in Chapter 6) was deposited at https://gitlab.mpcdf.mpg.de/mpibr/laur/inter-hemispheric-rem.

In both cases, I used Python version 3.8 and several standard packages: *scipy* (version 1.6.2), *numpy* (version 1.20.3), *pandas* (version 1.3.0), and *xarray* (version 0.18.2).

In addition, as part of my research, I found limitations in the NEST software package that I used to simulate my model (see section 2.3). These were technical issues related to the capacity of this software to generate exponential distributions of synaptic weights of inhibitory neurons and uniform distributions of synaptic delays. In an effort to contribute to the computational neuroscience community, I developed fixes for these issues using C++ and contributed them publicly. The fixes can be found in NEST version 2.16.0 [39].

2.3 Simulations

For the simulation of the large networks described in Chapters 4 and 5, I relied on an existing software package called NEST.

Single neuron

I modeled single neurons as point neurons using a previously developed model: the adaptive exponential integrate-and-fire (AdEx) neuronal model [35]. In AdEx, the membrane potential is defined as:

$$\begin{aligned} C\frac{dV}{dt} =& - g_L(V - E_L) + g_L\Delta_T \exp\frac{V - V_T}{\Delta_T} \\ & - g_e(V - E_e) - g_i(V - E_i) - w + I_e \\ \tau_w\frac{dw}{dt} =& a(V - E_L) - w \end{aligned}$$

The variable w accounts for the adaptation current. After each spike, it is increased by b, the spike-triggered adaptation. The values used for the single neuron parameters, either resulting from the literature or from fitting to electrophysiological data, can be found in Table 2.1.

NEST implements the AdEx model as the model "aeif_cond_exp", combined with exponential decay kernels for the excitatory and inhibitory synaptic conductances (g_e and g_i) [39].

Network

All simulated networks contained 100k single-neuron models. The integration time for the simulations was 0.1 ms. All simulations contained 100 trials where one or multiple trigger neurons were forced to spike at 400 ms intervals. Simulations also included an initial period of 1000 ms to account for the development of self-sustained activity (see section 4.3). A single simulation containing all trials for a single network took between 30 and 45 minutes of real-time on a single compute node with processor Intel(R) Xeon(R) Gold 6152 CPU @ 2.10 GHz.

All simulated networks were subject to a random spontaneous activity modeled as a random current independently sampled for every neuron in the network every

Variable	**Value**
Excitatory reversal potential (E_e)	10 mV
Inhibitory reversal potential (E_i)	-75 mV
Synaptic conductance time constant*	1.103681 ms
Excitatory synaptic conductance (mean)*	3.73 nS
Excitatory synaptic conductance (std)*	6.51 nS
Excitatory Synaptic conductance (max)*	67.8 nS
Leak reversal potential (E_L)	-70.6 mV
Spike detection threshold	0 mV
Membrane reset potential	-60 mV
Spike initiation threshold (V_T)	-50.4 mV
Membrane capacitance* (C)	239.8 pF
Leak conductance* (g_L)	4.2 nS
Subthreshold adaptation (a)	4 nS
Spike-triggered adaptation (b)	80.5 pA
Adaptation time constant (w)	144 ms
Slope factor (T)	2 mV
Refractory period	2 ms

Table 2.1: Neuron and synapse model parameters. Asterisk (*) indicates parameters fitted from experimental data [37]. The parameters corresponding to the distribution of synaptic conductance of excitatory connections (mean, std) refer to a lognormal distribution that will have its tail truncated at the maximum experimentally-observed connection (max).

1 millisecond. These currents were sampled from a Gaussian distribution

$$N(\mu_{in}, \sigma_{in})$$

the parameters μ_{in} and σ_{in} (mean and standard deviation) were fixed for the whole network. Unless otherwise specified, these parameters were sampled from uniform distributions (see Fig. 4.4B). Whenever a simulation was "re-ran", I used the same connectivity and trigger neuron as well as the same parameters μ_{in} and σ_{in} for the spontaneous drive, but not the exact same current values, which were resampled.

I initially generated an original set of 6000 simulations where I identified sequences of spikes (Fig. 4.6) and sub-sequences (Fig. 5.2). From this original set of simulations, I re-ran several random selections for different analyses:

- I analyzed a random selection of 900 simulations to quantify the presence of motifs of spike transfers (Fig. 4.11).
- I re-ran a random selection of 2000 simulations where I first removed connections based on their strengths (Fig. 4.12).
- I re-ran a random selection of 2000 simulations where a random gate neuron received an additional external spike to investigate how sub-sequence activation can be reliably controlled (Fig. 5.4).
- I selected 8 simulations and re-ran each one while varying the strength and timing of the external spike. Each of these 8 simulations was re-run between 1142 and 3655 times.
- I selected 1 representative simulation (same as in Fig. 5.2B) to study the coactivation of multiple triggers. I re-ran this simulation 7000 times: 2000 with single, randomly chosen triggers, and 5000 times with random pairs of those triggers (Fig. 5.6).
- I repeated the above procedure involving multiple triggers with 50 other representative simulations in order to verify the generality of my results. In that case, for each of those 50 simulations, I randomly selected 20 triggers and 80 pairs of triggers, resulting in $50 \times (20 + 80) = 5000$ new simulations (Fig. A.4).

Note that one of the goals of my model was to provide predictions beyond what could be recorded experimentally (see section 4.1.1), which required very large networks (100k neurons), which in turn resulted in very computationally heavy and time-consuming simulations and analysis. For this reason and for clarity of presentation, in some of the results presented in Chapters 4 and 5, I focus on representative simulations. Nonetheless, whenever I focused on single simulations, I then verified the generality of my conclusions through additional simulations and re-instantiations of the network. The summaries of these additional simulations can be found in Appendix A.

3 Biological and theoretical background

This thesis investigates how neuronal activity propagates within the reptilian brain, with a focus on the propagation of spikes in the cortex of turtles (Chapters 4 and 5) and the interaction of sub-cortical populations in the sleeping lizard (Chapter 6). In this thesis, I first summarize previous experimental and theoretical work and introduce several concepts required to interpret the novelty and context of my results. I split this literature review into three sections:

- In section 3.1, I will present the physiology and evolutionary context of the reptilian brain, with an emphasis on the circuitry and properties of the three-layered turtle cortex.
- In section 3.2, I discuss our current understanding of the nature of sleep as a repeating state characterized by distinct brain activity patterns that are observed across various animal species.
- Finally, in section 3.3, I will outline the main theories and models of propagation of neuronal activity in the brain, with a focus on computational models of spike propagation.

To complement the general view of the field that I summarize here, each result chapter contains a brief introduction that introduces the specific research question addressed (sections 4.1, 5.1, and 6.1), followed by a discussion of the results in relation to some of the previous work presented here (sections 4.5, 5.5, and 6.6).

3.1 The reptilian cortex

The cortices of all currently existing reptiles and mammals likely descended from a single primordial cortex, which suggests that they have preserved beneficial aspects of their shared heritage in terms of circuit structure and dynamics. The existence of a cerebral cortex likely predates the divide of the amniotes between the sauropsids and the synapsids some 320 million years ago [40, 41]. The sauropsids, or "reptilian lineage", gave rise to avians (birds) and modern non-avian reptiles, while the synapsids were the ancestors of modern mammals. The key characteristic that gives name to the amniotes is the amnion, a membrane that protects their embryos and enables them to survive in non-aquatic habitats, leading to the eventual invasion of land [21].

Modern "reptiles" comprise a diverse group of thousands of amniote species, which is paraphyletic, that is, defined by exclusion, with respect to the subclades of

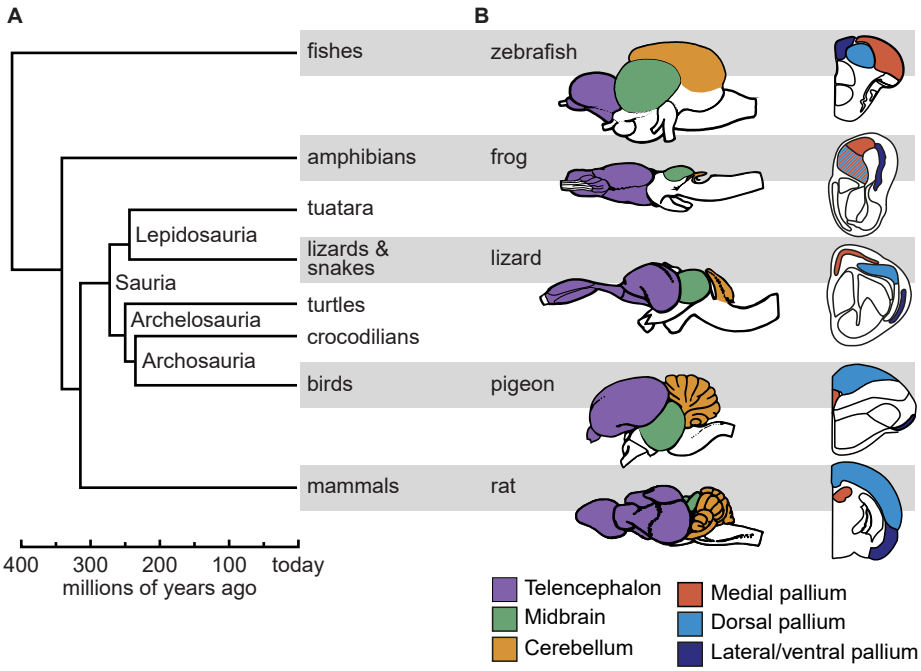

Figure 3.1: Phylogenetic tree of vertebrate brains. **A.** A simplified phylogenetic tree of vertebrates. Modern "reptiles" include all Sauria, excluding Avians (birds). **B.** Left: Schematized lateral view of illustrative species for groups in grey. The colored areas correspond to subdivisions common to all the species. Right: Schematized coronal section of the right hemisphere with pallial subdivisions in different colors. Modified from [25].

birds and mammals. This group includes the tuatara, snakes, lizards, crocodilians, and turtles, which all display common or similar features in terms of development, physiology, and function [42,43]. Defining similarities across species has traditionally involved examining their developmental period, especially the phylotypic stage during mid-embryogenesis, where related species typically express the most similar morphological and molecular features [44,45]. In spite of the diversity of adult brains, the basic division of the vertebrate brain during development is preserved, according to modern transcriptomics, which uses the co-expression of transcription factors to identify homologous brain areas during embryogenesis [20, 25, 45] (Fig. 3.1). Within the telencephalon, the cerebral cortex originates from a subregion of progenitor cells within the dorsal pallium. Modern molecular techniques based on the expression of transcription factors show us that the cerebral cortex contains homologous subdivisions common to all vertebrates [46] (Fig. 3.2).

The cerebral cortex developed as a part of the pallium in the telencephalon, and in its ancestral form, likely had three layers, an architecture that has been preserved in extant non-avian reptiles and in allocortical structures, namely the hippocampus and piriform cortex of modern mammals [47]. In avians, the pallium underwent a drastic reorganization, losing this layered architecture, which makes it harder to establish precise comparisons between cortical regions and structures in birds and other corticated species [48]. The six-layered neocortex that we observe in the mammalian lineage likely evolved from an anterior-dorsal region of the ancestral three-layered cortex located between the hippocampus and piriform cortex [25, 41, 49]. Altogether, the study of the reptilian cortex gives a point of reference that connects multiple lineages (ancestral amniotes, modern mammals, and birds) and potentially can illustrate evolutionarily retained properties and elemental function across multiple cortical structures [44].

The reptilian cortex contains at least three key subdivisions: the medial, lateral, and dorsal cortices. The medial cortex is estimated to be the homolog to the mammalian hippocampus and the lateral cortex to the mammalian piriform cortex. Between these two areas, the dorsal cortex seems to receive multisensorily and has been established through transcriptomic studies as corresponding to the mammalian neocortex [20, 41]. In turtles, the dorsal cortex receives prominent visual input, but its structure and physiology suggest functions beyond early sensory processing, with likely associative characteristics [26, 50]. The existence of somatosensory and motor cortices is not supported by any strong evidence.

The six-layered mammalian isocortex and the three-layered reptilian dorsal cortex develop from the dorsal pallium [51]. The middle layer (layer 2) of the reptilian cortex contains densely packed principal cells, mainly excitatory pyramidal cells, and is surrounded by two layers (layers 1 and 3) containing mainly GABAergic interneurons and dense neuropil [50]. Inputs to the reptilian cortex arrive to layer 1, which seems to act as well as an associative layer [50, 51]. In the mammalian 6-layered neocortex, the deep layers (layers 5 and 6) are considered output layers with corticofugal projection fibers, while the upper layers (layers 1, 2, 3, and

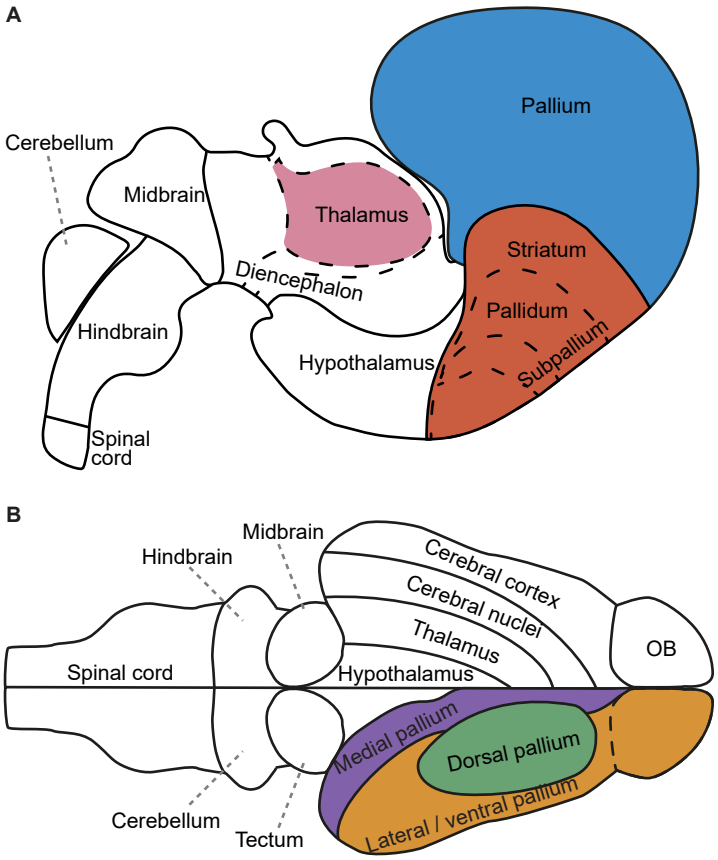

Figure 3.2: Subdivisions of the developing and adult vertebrate brain. **A.** Schematic of the main sections of an idealized developing vertebrate brain. Lateral section; anterior is to the right. Colors indicate subdivisions that can be identified with transcription factors, which become refined through development. **B.** Schematized top view of an adult turtle brain indicating the major subdivisions common across vertebrates. Colors indicate subdivisions between the lateral and the medial pallium that can be identified through protein expression during development. Modified from [25].

4) are input or associative layers, receiving many corticopetal fibers. The single principal cell layer found in the reptilian cortex has been hypothesized to have evolved into the deep layers of the mammalian neocortex, with the upper layers of the mammalian neocortex being an evolutionary addition in the mammalian lineage [52]. However, molecular studies have revealed that layer 2 of the anterior dorsal cortices of turtles contain a mixture of glutamatergic neurons that express either upper or deep layer markers [41]. Although there seems to be no one-to-one homology between turtle neurons and individual mammalian glutamatergic neurons, this molecular evidence now points to an evolutionary reorganization of the transcription factors in place during development [41, 45]. In the mammalian hippocampus, which is an archicortex and architectonically comparable to the three-layered turtle cortex, electrophysiological studies have shown differentiated roles and input/output patterns of glutamatergic cells during sharp-wave events between the deeper (closer to stratum oriens) and more superficial (closer to stratum radiatum) cells. Unsupervised clustering of the electrophysiological properties of principal neurons in the turtle cortex also reveals two broad groups [53]. Interneuron subtypes in the turtle cortex are very conserved in reptiles and mammals, with the presence of somatostatin (SST) expressing neurons as well as transcriptomically identified groups matching the mammalian parvalbumin-expressing (PV) and vasoactive intestinal peptide (VIP) expressing neurons. Both glutamatergic and GABAergic interneurons of the reptilian cortex follow the same developmental process as in the mammalian cortex, where glutamatergic cells derive from multipotent cortical cells, and GABAergic cells initiate in the subpallium and then migrate to the cortex [41].

Altogether, the architectonic, transcriptomic, and developmental similarities across the reptilian dorsal cortex and the mammalian iso-, paleo-, and archi-cortices suggest that the core components of cortical circuitry likely appeared already in their amniote common predecessor. These components may have undergone recombination and specialization across species, but their core elements seem to have remained mostly preserved over the last 320 million years.

3.1.1 Circuitry of the turtle dorsal cortex

Next, I briefly describe the details of the architecture and circuitry of the dorsal cortex of turtles as relevant to Chapters 4 and 5.

Single-cell transcriptomics supports the classical anatomical division of the turtle cortex into lateral, dorsal, dorsomedial, and medial [41, 54] (Fig. 3.3A). The dorsal cortex displays a reduction of cell density in layer 2 and a thickening of layer 3, which can be used to establish a further subdivision into the lateral dorsal cortex and the medial dorsal cortex [54].

The molecular layer (layer 1) is the most superficial one and contains the apical dendrites of principal cells. Incoming thalamo-cortical and cortico-cortical axons establish input connections in layer 1 with principal cells as well as interneurons. In the neuropil of layer 1, there are the bodies of a few GABAergic interneurons

with comparable transcriptomic signatures to those of neurogliaform interneurons typically found in layer one of the mammalian isocortex [41]. Layer 1 interneurons have been sub-classified into subpial cells, with dendrites extending parallel to the pial surface, and stellate cells, located in the deeper parts of layer 1 and displaying sparsely spiny dendrites that extend across all layers [55]. Subpial GABAergic interneurons display spike-rate adaptation, while stellate cells do not.

The cellular layer (layer 2) is located in the middle of the 3-layered cortex and contains densely packed somata of excitatory pyramidal neurons. Their density is estimated at 25000 per mm^2. These principal neurons display spiny dendrites that extend towards the pia, resulting in a morphology similar to the granule cells in the dentate gyrus of the mammalian hippocampus. These pyramidal neurons possess multiple apical dendrites and a small or no basal dendrite. Note that mammalian neocortical pyramidal cells typically possess a prominent single apical dendrite and multiple basal ones [56] (Fig. 3.3B). Reconstructions and analysis from electrophysiological signatures reveal broad axonal arborizations [28, 37].

In terms of their electrical properties, turtle cortical principal neurons display complex behaviors. Their dendrites can display dendritic spikes as a result of voltage-dependent non-linear conductances, and their action potentials can backpropagate into their dendrites, triggering the voltage-dependent release of calcium [56, 57].

Finally, the subcellular layer (layer 3) is the deepest layer, sitting, in the dorsal cortex, above a ventricle. Layer 3 is composed mainly of basal dendrites and axons from layer 2 pyramidal neurons, as well as a larger population of GABAergic interneurons when compared to layer 1. These interneurons display aspiny dendrites that target all layers, producing local inhibition through GABAa (fast) and GABAb (slow) receptors. As layer 1 stellate cells, these interneurons show little spike-frequency adaptation [57, 58].

Putting the circuit together, it seems that both pyramidal neurons and layer 1 interneurons receive excitatory thalamic inputs, which trigger recurrent spiking, activating strong, non-adaptive, local inhibition that may act feedforward or feedback [57] (Fig. 3.3C). Indeed, a short electrical stimulation of afferent fibers to the dorsal cortex triggers a volley of excitatory postsynaptic potentials in both the layer 1 stellate interneurons and the layer 2 pyramidal cells [57]. The excitatory input on stellate cells is long-lasting and leads to copious spiking, while excitation is only brief on pyramidal neurons and is followed by strong and long-lasting inhibitory input, likely masking any additional excitation. Using light stimulation in an eye-attached whole-brain ex vivo preparation results in similar results, with pyramidal neurons receiving inputs for 50–100 ms and being dominated by inhibition [59]. Note that these responses are not unlike those seen in the layer 2/3 pyramidal neurons of awake mice during visual stimulation [60].

3.1.1.1 The turtle cortex ex vivo

Some of the computational results that this thesis develops (Chapters 4 and 5) are grounded on measurements and phenomena that had been observed in the ex

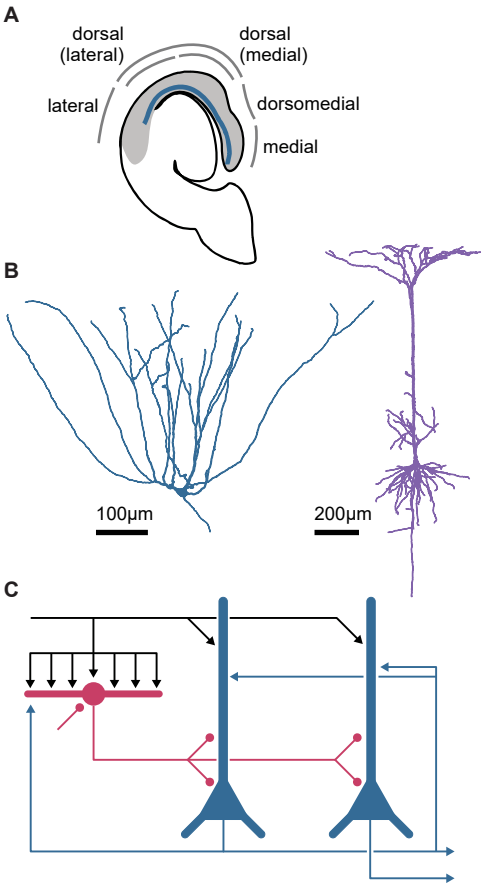

Figure 3.3: Circuitry of the turtle cortex. **A.** Simplified schema of a coronal cortical slice of the turtle brain with the main divisions of the cortical sheet. Blue indicates the layer of principal cell bodies (layer 2). Note dorsal (visual) cortex rests atop a ventricle. **B.** Outline of filled pyramidal neurons from turtle and rat (layer 5). **C.** Simplified schema of intracortical connections of the turtle visual cortex. Volleys from thalamocortical afferent fibers (black) directly excite the apical dendrites of pyramidal cells (blue) and inhibitory cells (pink). Activation of interneurons causes feedforward inhibition onto pyramidal cells. Activation of pyramidal cells triggers reciprocal excitation and local feedback inhibition. Axons from pyramidal neurons project basally, providing excitatory output from the cortex. **A** and **B** modified from [56]; **C** modified from [57].

vivo turtle cortex. I then use simulations to produce quantitative predictions for an in vivo scenario. Next, I discuss the particularities of the turtle cortex ex vivo preparation and why it is an experimentally advantageous preparation for studying cortical connectivity.

Fresh-water turtles are ectotherms (cold-blooded) animals that regularly retreat to the low-oxygen muddy bottom of ponds in the winter, where their metabolism slows down while tucked in the mud. As a result, pond turtles may go for months without eating and with little oxygen. Fresh-water turtles thus have a natural resistance to low (hypoxia) or no (anoxia) oxygen levels [61]. Their neural tissue thus possesses several physiological adaptations to tolerate anoxic conditions, such as reduced ion permeability or the blockage of certain channels and receptors to save energy. As a result, the brain of turtles can be preserved healthy and spontaneously active and even be stored at very low temperatures overnight without incurring tissue damage.

The robustness of turtle neural tissue has led to the development of several ex vivo preparations. These can be more experimentally advantageous than in vivo because of the absence of movement or respiratory artifacts, highly stable recording setups, and more accessible areas. Compared to sliced in vitro preparations, ex vivo preparations may preserve not only close-range but also mid- and long-range connectivity across the whole brain and even include some sensory organs. For instance, the brain may be prepared as an eye-attached explant, allowing the study of cortical activity triggered by the presentation of visual stimuli onto the retina [26, 28, 62].

The model that I present in Chapter 4 recapitulates and builds upon an experimental study where the authors employed an ex vivo cortical "slab" preparation [37]. This slab constituted a piece of about 3 mm^2 of the dorsal cortical sheet, with the pial surface at the top and the surface corresponding to the ventricle directly below. Thanks to its size and anatomy, a large section of long- and mid-range cortico-cortical connectivity was preserved. Additionally, the flat organization of the slab allowed the use of high-density flat Multi-Electrode Arrays (MEA) to record neuronal activity from below while simultaneously using multiple patch-clamps from above. This enabled the recording of large-scale populations with well-preserved connectivity and controlled inputs [28]. I summarize the results from these experiments as well as some of the open questions in section 4.1.

3.1.1.2 Lessons from mammalian cortical connectivity

In spite of turtles being an interesting (section 1.3) and useful (section 3.1.1.1) animal model for neuroscience, it is still heavily understudied when compared to mammals, and in particular, rodents. For this reason, and considering the evolutionary links between turtles, early amniotes, and mammals, as well as the architectonic similarities between turtle cortex and mammalian paleo- and archicortices, it makes sense to turn to mammalian literature on cortical architecture. In

this section, I briefly review some insights into cortical connectivity from mammalian studies, with an emphasis on those relevant to the results described in Chapter 4: patterns of excitatory and inhibitory connectivity and the presence of powerful excitatory connections.

Our knowledge of the connectivity of neural circuits is ultimately essential to our comprehension of how they process sensory information. Neurons in circuits interact by means of synaptic contacts resulting in complex recurrent networks. The synaptic connections of neuronal circuits are subject to plastic modifications, which alter, over time, the strength of the postsynaptic potential that they elicit after a presynaptic activation [8, 63, 64]. The majority of our understanding of cortical function comes from observing individual neurons or small groups of neurons [65–68]. Whole-cell patch-clamp recording techniques have been one of the primary methods that have been used to study synaptic connectivity between pairs or small groups of neurons [69–74]. Studies using multi-patch clamp setups, for example, have been conducted in an effort to map out the statistical features of the somatosensory cortex of rats [68, 72, 75]. Computational modeling that is heavily constrained by these experimental measurements can then attempt to produce various in sillico reconstructions of this cortical tissue in order to produce a common framework for study [76–78]. More recent research has shifted the focus of the field to the study of larger groups of neurons of multiple populations, primarily through the use of high-density extracellular electrodes or imaging tools [79, 80]. Despite all these efforts, we do not yet have a complete picture of the connectivity profiles of primary sensory cortices.

Inhibition

Neurotransmitter GABA, which inhibits activity in cortical circuits by acting on postsynaptic GABAa or GABAb receptors, is released by local cortical interneurons. Each of these receptor types elicits responses in the postsynaptic neuron that vary in ionic properties, resulting in different time courses, from fast GABAa receptors to slow GABAb receptors [81]. A similar phenomenon has been found in the turtle dorsal cortex [57, 58]. When compared to excitatory connections, the connectivity between inhibitory interneurons is understudied. This can be partially attributed to the fact that inhibitory interneurons are smaller in size and have a relatively low number in the cortex. However, estimates of inhibitory connectivity may not lack too much accuracy when compared to estimates of excitatory connectivity due to many studies relying on tissue slicing, which can truncate long-range axonal projections but preserves axonal arbors that are mainly local, such as those of interneurons [82].

Interneurons are so diverse that researchers have made several efforts to categorize them into separate cell types, which makes it hard to pool together data for a coherent picture of inhibitory-to-inhibitory connectivity [65, 83]. For instance, recent work in the mouse primary visual cortex distinguishes 15 sub-types of interneurons with particular connectivity patterns for each pair [74]. The use of molecular markers has

become the most commonly accepted method of classification of cortical interneurons into a small number of non-overlapping classes [84]. It typically distinguishes between those interneurons expressing parvalbumin (PV), somatostatin (SST), and a serotonin receptor known as 5HT3aR. It is common practice to further subdivide the 5HT3aR type into vasoactive-intestinal peptide (VIP) positive interneurons and non-VIP interneurons [83]. Note that these major interneuron sub-types have also been identified in the reptilian cortex through transcriptomic profiles [41]. In terms of general connectivity, there are three main forms of inhibitory connectivity [74]: interneurons primarily inhibiting local pyramidal cells, interneurons that primarily inhibit other interneurons, and interneurons that project to all other neuron types within their same layer.

Excitation

Excitatory neurons release glutamate as their main neurotransmitter, resulting in a depolarization of their postsynaptic partner called an Excitatory Post Synaptic Potential (EPSP).

Compared to excitatory-to-excitatory, excitatory-to-inhibitory synaptic connections are more frequently observed, and they tend to produce stronger depolarizations with cell-type-specific postsynaptic dynamics [85, 86]. This strong depolarization, resulting from a strong synaptic gain, has been speculated to be behind frequently observed spike-to-spike transfers in vitro and in vivo from excitatory to inhibitory neurons [87, 88]. Short-term synaptic dynamics are typically classified as depressing or facilitating. These have been demonstrated to translate into early- versus late-onset inhibitory feedback phenomena, which are able to influence activity at various timescales, possibly involving morphologically distinct interneuron types, according to studies conducted in the hippocampus and the neocortex [87, 89]. In addition to the difference in timescales, these different interneuron types may target specific areas of other excitatory cells, such as soma or dendritic domains. Therefore, excitatory-to-inhibitory spike transfers may cause different forms of feedback inhibition in the circuit: early somatic feedback or late dendritic feedback [90]. Indeed, both the hippocampus and the neocortex have been demonstrated to include a di-synaptic inhibitory microcircuit with an excitatory-to-inhibitory-to-excitatory pattern [87, 89]. For instance, the bulk of connections between excitatory stellate neurons in the entorhinal cortex might be, in fact, a form of lateral inhibition [91]. It is possible that this frequent circuit motif of di-synaptic lateral inhibition may be responsible for neuronal competition between excitable neurons and for activity normalization at the scale of populations [87, 92, 93].

Paired recordings from layer 4 neurons in the barrel cortex of rats have revealed the presence of very strong excitatory connections [68, 94]. However, these are rare connections. The overall distribution of EPSP amplitudes of excitatory connections follows a long-tailed distribution, where the great majority of connections elicit very weak EPSPs (in the order of less than 1 mV), and few rare connections elicit very strong EPSPs (several mV). This skewed distribution is usually well described by

the log of a normal distribution (log-normal) [68, 95, 96]. These rare and powerful connections have been observed between spiny stellate neurons in layer 4, where a single action potential from the presynaptic neuron can elicit a spike in the postsynaptic one [75]. In mammals, similar reliable activation of other neurons from a single spike of a pyramidal neuron has been observed in slices of the human prefrontal cortex, but in those cases, the activations traversed a reliable poly-synaptic chain, possibly involving rebound inhibitory [97].

Powerful excitatory synapses are, in fact, common in the mammalian hippocampus (a three-layered archicortex) in the form of the Mossy Fiber Synapse within the tri-synaptic pathway [98, 99]. The tri-synaptic circuit involves connections from the entorhinal cortex to the dentate gyrus, from the dentate gyrus to pyramidal cells in CA3 through mossy fibers, and from CA3 to CA1 through Schaffer collaterals. The Mossy Fiber Synapse is a very strong synapse, with each one composed of several active zones on the same postsynaptic spine [100]. A postsynaptic CA3 pyramidal neuron can be depolarized by as much as 20-30 mV from a single presynaptic activation [99]. These strong depolarizations typically lead to spike-to-spike transfers, observed in both in vitro and in vivo setups, which in turn has granted the Mossy Fiber Synapse the name "detonator" synapse [98].

Recent studies using whole-cell patch-clamp measurements have also shown strong synapses in recurrent connections between pyramidal CA3 neurons and across different layers in the barrel cortex [101, 102]. However, how rare strong excitatory connections are across all cortical structures remains an unsolved question that will require higher-throughput methods than pairwise patch-clamp. For instance, a recent study using Electron Microscopy reconstructions of a piece of rodent cortex showed long tails of spine sizes in the synapses between L2/3 pyramidal neurons, with an explicit binarization into weak and strong connections [96].

Overall, the presence of excitatory synaptic strengths that are orders of magnitude apart suggests the existence of differentiated roles between weak and strong connections. Their existence across multiple cortical areas and species may point to these long-tailed distributions as a core element of cortical circuit function [95].

3.1.2 Vision in the turtle cortex

In both turtles and mammals, visual signals travel from the retina to the cortex through the thalamus [103]. Contrary to mammals, the retina of turtles does not contain a fovea, which was likely lost with the invasion of land and re-developed in the mammalian lineage with primates [21, 104]. Instead, the turtle retina contains a "visual streak", maybe a residue of the fish fovea, in the form of an elongated area with densely packed ganglion neurons oriented horizontally [105].

Axons from retinal ganglion cells bundle together to form the optic nerve and target the telencephalon. Most of the fibers cross the optic chiasm and take one of two main pathways: the retino-geniculo-cortical pathway and the retino-tecto-fugal pathway. The first reaches the dorsolateral geniculate nucleus (dLGN) of the thalamus, which

forwards activity to the dorsal cortex. The second pathway is more indirect and includes several stages: to the optic tectum (a homolog of the mammalian superior colliculus), the nucleus rotundus of the thalamus, the dorsoventricular ridge (DVR, containing a homolog of the mammalian claustrum [30]), and eventually, the dorsal cortex [106, 107].

The dLGN of each hemisphere receives inputs from retinal ganglion cells of both the contralateral and ipsilateral retinas. The fibers from dLGN then arrive laterally and project dorsally onto the dorsal cortex. Earlier studies have described these fibers as following an order matching the iso-azimuth axis of the visual field [108]. However, more recent tracing evidence argues that these projections are non-topographic [26]. The dLGN axons then produce synaptic connections in the upper region of layer 1 of the dorsal cortex. As described in section 3.1.1, these thalamic inputs impinge onto the spiny apical dendrites of pyramidal neurons and the smooth dendrites of subpial glutamatergic interneurons. Experimental evidence suggests that a single interneuron may receive over six times more thalamic fiber synapses than a single pyramidal cell [109].

The primary visual cortex of mammals has a retinotopic mapping of visual space, where single neurons activate reliably to the presentation of visual stimulus in a limited area of the visual field [4]. Individual neurons in the dorsal cortex of turtles, on the other hand, have very large receptive fields that frequently include the entire visual field [26].

When retinal spots are stimulated in turtles who are awake and immobilized, activity is evoked in large areas of their pallium. [110, 111]. Single neurons in the dorsal cortex react more favorably to novel and moving stimuli [111]. Initial reports suggest that the nasotemporal axis of the visual field correlates to the nasotemporal axis of the turtle dorsal brain, whereas the dorsoventral axis of the visual field correlates with the lateromedial dorsal cortex axis [111]. However, more recent evidence supports an absence of retinotopy in the turtle dorsal cortex [26]. Nonetheless, the receptive fields of single neurons, even though covering most of the visual field, are inhomogeneous. The composition of these inhomogeneities enables decoding the position of a stimulus in the visual field through a population decoding approach. In summary, it seems likely that visual processing in the dorsal cortex is not specialized in local computations of small parts of the visual field but rather that it involves computations that are visually global, population-based, and associative.

The generation of activity waves that can propagate through neural networks is an emergent trait observable in multiple animal models and brain areas, from invertebrate sensory areas to the human cortex [112, 113]. Waves are not exclusively but often associated with sleep (see section 3.2). This form of propagation has also been reported as a form of spontaneous and evoked activity in the turtle dorsal cortex [114]. Using a whole-brain eye-attached ex vivo preparation (see section 3.1.1.1), it is possible to project light stimuli onto the turtle retina while recording brain activity [28]. The presentation of visual stimulation triggers traveling waves in the dorsal cortex of turtles that can be imaged through voltage-sensitive

dyes [114, 115]. The waves are depolarizing, with a supra-threshold and a sub-threshold component. These waves expand from anterior to posterior and from lateral to medial following visual stimulation, with a speed of propagation of around 50–90 µm/ms [114]. The traveling waves may be studied as neuronal avalanches in the context of critical dynamics in the turtle brain, a state that ensures maximal transmission of information [29].

Intracranial local field potential in the dorsal cortex of awake turtles responds with oscillatory activity to visual stimulation [26, 116, 117]. Visual stimuli trigger slow oscillations in the 1–5 Hz range, plus a sustained oscillatory activity in the 15–35 Hz (beta) range. These beta oscillations in the dorsal cortex show a rostrocaudal phase gradient and are linked to intracortical current loops that occur from the phase locking of the activation of pyramidal neurons at this frequency [26]. The slow oscillatory component also presents a phase gradient that moves from lateral to medial, as do the traveling waves, and parallel to the thalamocortical afferent fibers.

In conclusion, the visual processing pathways in turtles and mammals present some differences but also significant similarities. The visually responding cortical neurons of turtles display receptive fields that cover the entire visual field, suggesting that visual processing in turtles is population-based and globally associative. Additionally, visual stimulation triggers traveling waves in the turtle dorsal cortex, with oscillations linked to intracortical current loops. Overall, studying the visual system of turtles provides insights into the evolution and diversity of visual processing strategies across different species.

3.2 Two-stage sleep

Sleep is a ubiquitous behavior observed across the animal kingdom that involves distinct brain states characterized by unique electrical signatures. Yet, despite the intensive study of sleep, we still lack a mechanistic understanding of how many of the electrical signals that underlie sleep initiate and propagate within the brain. In Chapter 6, I describe a series of findings by my experimental collaborators and me on the dynamics of propagation of neuronal activity across several sub-cortical areas of the *Pogona vitticeps* lizard. In this section, I briefly introduce sleep and how it relates to evolution and reptiles.

Sleep is an essential and repeating state that can be seen in most animals. Several behavioral characteristics, such as decreased reactivity to stimuli, behavioral quiescence, quick reversibility, and homeostatic adjustment following deprivation, have been used to define sleep behaviorally [33]. Non-invasive recording methods have also been used to measure and define sleep states in humans. These methods include electroencephalography (EEG) and electromyography (EMG). EEG measures the electrical activity of the brain by placing small electrodes on the scalp, while EMG measures the electrical activity of muscles by placing small needle electrodes into the muscle tissue. Combined data from these methods during sleep in humans and rodents has led to the detection of distinct patterns of electrophysiological activity, such as alpha waves, spindles, and slow oscillations, which in turn break down the sleep state into multiple stages (Fig. 3.4). Yet, the underlying processes of many of the electrical signals connected with sleep are still poorly understood, despite the fact that sleep is a ubiquitous behavior and that distinct sleep phases seem to perform essential tasks across a variety of animal species.

A characteristic phase in human sleep involves periods lasting 10 to 20 minutes with bursts of abrupt, fast eye movements. Consequently, this phase has been termed "rapid eye movement" (REM) sleep. REM sleep is also accompanied by changes in physiology, such as a rise in heart rate and blood pressure, as well as in the temperature of the brain and the blood flow to the cerebral cortex [33]. The electrophysiological signatures of REM sleep include high-frequency oscillations in the bands called "theta" (6–9 Hz in rodents, 4–6 Hz in humans) and "gamma" (30–90 Hz). The EEG activity in the brain during REM sleep is very similar to that during wakefulness which has led to REM sleep also being referred to as "paradoxical" or "desynchronized" sleep.

On the other hand, non-REM sleep corresponds to prolonged periods without movement from the eyes that separate REM periods. A full cycle containing all phases of non-REM and REM sleep in humans is about 90 minutes [33]. This phase of sleep is sometimes referred to as slow-wave (SW) sleep as it is distinguished by EEG signals that have a low frequency and a big amplitude. In humans, non-REM sleep is typically considered to have three different "depth" stages according to the properties of EEG signals, resulting in a progression of stages: strong theta activity (roughly 4–8 Hz; "N1" stage), presence of spindles and K-complexes ("N2" stage), and delta

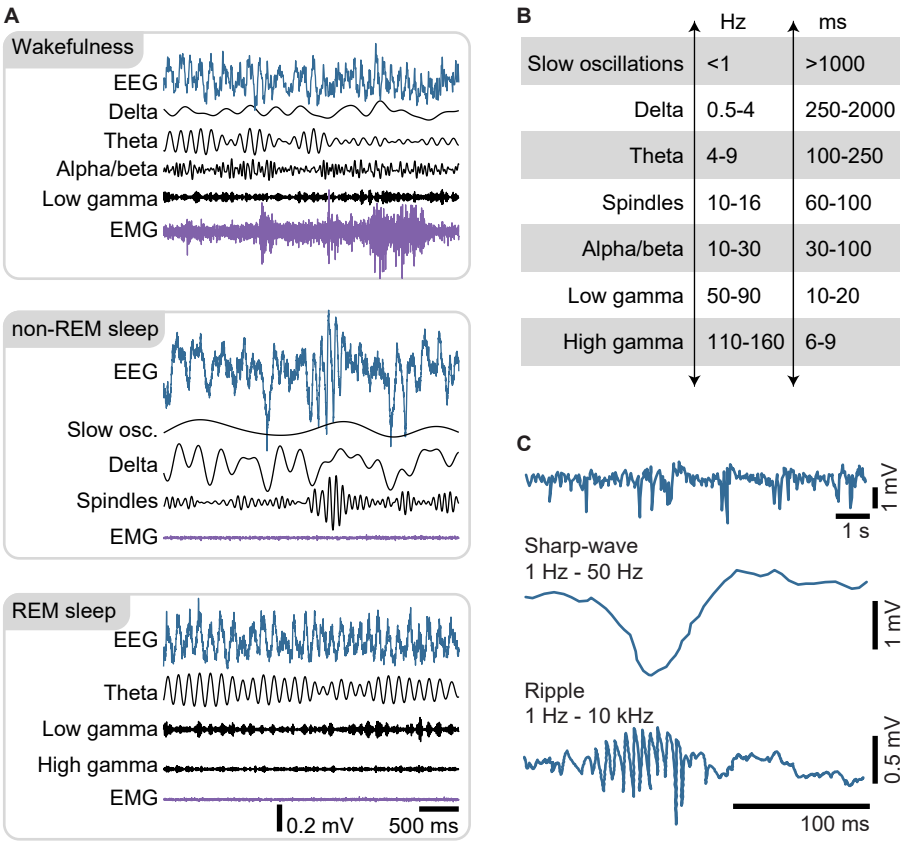

Figure 3.4: Oscillatory signals during sleep. **A.** Example electroencephalographies (EEG) from freely moving mice during wakefulness and the two main sleep states: rapid eye movement sleep (REM) and non-REM. EEG filtered at different frequency bands reveals different spectral profiles of each of the states. Electromyography (EMG) indicates movement. **B.** Summary of the main spectral bands used in analyzing electrophysiological signals during sleep. **C.** Top: Rat hippocampal CA1 LFP recorded during quiet wakefulness. Middle and bottom: zoom-in on a Sharp-wave Ripple (SWR). **A** and **B** modified from [33]. **C** modified from [118].

waves ("N3" stage) [33]. Compared to REM, electrophysiological signatures of non-REM sleep have been studied in more detail and typically involve [33, 119] (Fig. 3.4):

- Slow oscillations and delta waves are caused by changes in the resting membrane potentials of thalamic and cortical neurons, which alternate between depolarized (UP) and hyperpolarized (DOWN) states [120]. These oscillations occur at frequencies lower than 1 Hz.
- Delta waves, which have a frequency range of approximately 1 to 4 Hz and are frequently observed during the deeper stages of non-REM sleep. These waves are thought to have two possible mechanisms: hyperpolarization in the thalamus and intrinsic oscillations in layer 5 cortical neurons.
- K-complexes, which occur as spontaneous events during stage N2 sleep in humans and consist of a slow component followed by a regular 14 Hz rhythm.
- Spindles, which are transient oscillations of 11–15 Hz. Spindles are frequently, but not exclusively, associated with the UP state of a slow wave and are temporally associated with distinct blood flow changes in the lateral and posterior thalamus. K-complexes and spindles occurring spontaneously during stage N2 sleep in humans reflect inhibitory and excitatory microstates [33].

We still do not completely understand why animals need to sleep for a significant portion of their lives. Theories of the function of sleep can be broadly summarized in two general directions: those suggesting that sleep serves restorative and housekeeping functions and those suggesting that sleep is involved in cognitive tasks such as brain plasticity and may underlie learning and memory consolidation [121–124].

An important electrical signal often associated with quiet wakefulness, non-REM sleep, and plasticity is the sharp-wave ripple (SWR). SWR are relatively brief, high-frequency electrical signals that can be recorded in the extracellular local field potential (LFP). They are characterized by large-amplitude negative deflections (the sharp-wave) and are typically accompanied by fast oscillations in the local field potential (the ripple) [118] (Fig. 3.4C). While SWRs have been intensely studied in the CA1 region of the hippocampus, they can also occur in other hippocampal regions, such as CA2 and CA3, as well as in other brain regions, such as the entorhinal cortex, or in the claustrum [30, 32, 118]. However, the characteristics of SWRs may vary across different brain regions and species, and their functional roles may also differ. SWRs in the CA1 of rats during quiet wakefulness and non-REM sleep have been implicated in memory consolidation and synaptic plasticity [119, 125]. For instance, following maze exploration, hippocampal CA1 "replay" spiking sequences may occur during a SWR and recapitulate some of the patterns of activation experienced by the animal when it navigated the maze [126]. SWRs are thought to play a role in the transfer of information from the hippocampus to other brain regions and may be involved in the integration of new memories with existing knowledge.

In summary, while the precise function of sleep remains elusive, current evidence indicates a role in the consolidation of memories and learning through synaptic plasticity. Further research is required to arrive at a full view of the biological functions of sleep.

3.2.1 Sleep and the reptilian brain

The question of whether the REM and non-REM stages of sleep existed in the common ancestor of amniotes has been debated until recently [24]. All mammals, birds, and reptiles are descended from amniotes, the earliest vertebrates whose life cycle was capable of occurring outside of water. Despite anatomical variations across all these lineages, many brain regions are preserved and continue to operate in a similar fashion [21] (see section 3.1).

Although sleep is experienced by all vertebrates, research on sleep has primarily focused on mammals and birds. It has been suggested that the evolution of REM and non-REM sleep in mammals and birds was driven by common constraints and selective pressures related to homeothermy [24]. However, there is now evidence to suggest that these sleep stages could be ancestral to both mammals and birds. A combination of behavioral tracking and electrophysiological recordings was used to investigate cortical and subcortical brain activity in the Australian *Pogona vitticeps* lizard [28]. Later, the same approach was applied to the South American *Salvator merianae* lizard [24]. This research confirmed the existence of behavioral and electrophysiological two-phase sleep in these lizards. Note that *Pogona vitticeps* belongs to the earliest subclass to branch out from the sauropsid line, resulting in a reptilian species that is among the furthest from avians, which strongly suggests a common origin of sleep in amniotes.

Interestingly, a recent study has shown two electrophysiologically distinct stages of sleep in zebrafish [127]. The two identified stages are characterized by either propagating waves (PWS) or synchronous slow bursting oscillations (SBS) in the dorsal pallium. Although PWS and SBS display electrical dynamics similar to, respectively, REM and non-REM, there are also differences between these states. For instance, PWS lacks the rapid eye movements that characterize mammalian and reptilian REM. Nonetheless, the observation of two distinct sleep states in fish has brought up the possibility of common roots of two-stage sleep that could go back beyond amniotes to the first vertebrates 500 million years ago.

In conclusion, sleep is a widespread phenomenon in the animal kingdom, yet its functions and mechanisms remain largely unclear. Sleep involves unique brain states marked by forms of neuronal activity that develop in distinct stages. Although the functions of sleep are not yet fully understood, recent research on *Pogona vitticeps* lizards and fish suggests that two-stage sleep may have originated from a common ancestor in amniotes and maybe in early vertebrates, respectively. Despite intensive research, many of the underlying processes of sleep remain a mystery, but studying

the electrophysiological mechanisms of sleep in different animal species could offer valuable insights into its evolution and purposes.

3.3 Computational models of spike propagation

Action potentials, or spikes, are the main form of communication between neurons in cortical networks [3] (Fig. 1.1). Spikes are all-or-none electrical impulses generated by individual neurons that will travel through the axon and trigger the synaptic release of neurotransmitters. These synaptic activations will cause voltage changes in the postsynaptic neurons. The generation and propagation of each action potential happen on a scale of just a few milliseconds [128, 129] and may have an effect on the membrane potential of thousands of synaptic targets [130]. In addition, cortical synaptic connectivity is highly recurrent and often reciprocal [68]. If a single excitatory spike of one neuron is expected to produce a spike in more than one of its postsynaptic partners, the system could easily degenerate in a recurrent amplification loop leading to runaway excitation and the synchronization of the entire network. On the other hand, synaptic connections are, in their majority, extremely weak in terms of the postsynaptic potential (PSP) that they elicit and may often run into a synaptic failure where no PSP is evoked [37, 68, 101, 102]. In that case, the propagation of action potentials can easily die out after very few synaptic jumps. If the reliable propagation of neuronal activity is the foundation for computations (section 1.2), how do cortical networks solve these problems?

Theoretical and computational neuroscientists have addressed the question of reliable propagation by proposing different models in which cortical neurons might be organized. These models make predictions on the type of connectivity structures shared by cortical neurons, as well as the spatiotemporal patterns of neuronal activity that result from them. In this chapter, I review the ideas and properties of some of these models as relevant to the results described in Chapters 4 and 5: synfire chains, polychronous chains, and rate propagation.

3.3.1 Synfire chains

Early theoretical studies have already suggested that feedforward structures of fully connected neurons should allow for the reliable transmission of excitatory signals in a "complete transmission line" [131]. This idea became formalized and further developed by Moshe Abeles and others throughout the 80s and 90s as the "synfire chain" model and remains an area of research today [14, 132, 133]. This model is one of the most popular theories of spike propagation as it is able to account for reliable transmission and repeating temporally-precise spatiotemporal patterns of action potentials. In this section, I describe the synfire chain theory and its relationship to spatiotemporal spike patterns.

The fundamental membrane dynamics of a neuron state that synaptic inputs are integrated up over time, and once the membrane voltage crosses a certain threshold, the neuron emits a spike. This has led to the description of neurons as integrators. Yet, an alternative interpretation of these membrane dynamics is possible, where neurons are described as coincidence detectors [14, 132, 134]. If single connections are of small amplitude, and spikes arrive with some temporal variability, the membrane

potential will display a certain variance, which will reduce the probability of a spike transfer, that is, that a given presynaptic spike leads to a postsynaptic spike.

How important is synchrony in ensuring a spike transfer? We can examine the effect of synchrony in a simplified example where a neuron receives a volley of temporally-spread input spikes (Fig. 3.5A). Let's imagine that the distance from the resting potential to the threshold is approximately equivalent to the combined amplitude of k EPSPs, all with the same amplitude. For simplicity, let's assume time can be divided into small bins that approximately correspond to the duration of a small EPSP and that an asynchronous volley of n input spikes is randomly spread over t time bins so that each spike may fall in any of the time bins with equal probability $1/t$. The probability that a given input spike of the volley causes a postsynaptic spike, i.e., that the membrane potential of the postsynaptic neuron reaches the spiking threshold k, is equal to the probability that at least $k-1$ of the remaining $n-1$ spikes fall in the same time bin. Under this scenario, the number S of spikes that do fall in a given bin is a random variable following a binomial distribution:

$$S \sim B(n-1, p = 1/t)$$

The cumulative distribution function F of a binomial random variable S gives the probability that it takes on a value less than or equal to a given value $F(x) = P(S \leq x)$. Thus, the probability that a given input spike transfers into a postsynaptic spike is

$$\begin{aligned} P(\text{transf.}) &= P(k-1 \leq S) \\ &= 1 - P(S < k-1) \\ &= 1 - P(S \leq k-2) \\ &= 1 - F(k-2) = 1 - \sum_{i=0}^{k-2} \binom{n-1}{i} p^i (1-p)^{n-1-i} \end{aligned}$$

Fig. 3.5A shows how this probability is affected by the total number of spikes in the volley and how temporally spread they are. For a volley of $n = 100$ spikes spread over $t = 40$ time bins and a threshold that is $k = 10$ times the amplitude of a single EPSP, the probability that a given presynaptic spike produces a postsynaptic spike results in $P(\text{transf.}) \approx 0.001$. If the volley was compressed to just $t = 4$ time bins (a 10-fold reduction), then the probability of a spike transfer rises to $P(\text{transf.})_{t=4} \approx 0.999$ (a 1000-fold increase) and becomes almost guaranteed. This simple example illustrates how synchronous inputs are radically more effective and efficient than asynchronous inputs in producing a spike transfer.

As described above, the most effective way to get a neuron to produce a response spike is to provide it with a sufficient number of input spikes that arrive almost synchronously. These kinds of inputs may be supplied by a pool of neurons that are projecting convergently onto the neuron of interest. On the other hand, if two neurons receive the same inputs, they have a greater likelihood of producing spikes

at the same time. In other words, synchrony has a tendency to preserve itself: it is effective at driving a single postsynaptic neuron, and it is the result of multiple neurons receiving a common input.

These observations led Abeles [14, 132] to describe a pattern of convergent and divergent connectivity (Fig. 3.5B). In this pattern, neurons are organized in groups (also called pools) which are placed sequentially in a feedforward fashion. Each neuron in a given group receives inputs from most neurons in the previous group (convergence), and it projects to most of the neurons in the next group (divergence). Fully all-to-all connectivity is not strictly required, but connections between groups need to be sufficiently dense to accommodate for the expected membrane potential distance between resting and threshold, as well as the probability of failure to spike of some neurons in the previous group. Using this pattern of connectivity, it is possible to repeatedly arrange groups and create extensive chains of neurons. This architecture circumvents the problem of weak one-to-one synaptic connections by grouping them and creating an overall strong connection between each pair of consecutive groups. Note that the activation of the chain requires that neurons in the first group start with near-synchronous activity since a sparse or temporally dispersed activation will fail to produce robust propagation (Fig. 3.5C). At each level, the synchronous activity present in one group replicates itself in the next group so that a volley of spikes will move across the network in steps, activating one group after the other. This feedforward structure of neurons that **fire syn**chronously in sequence gave rise to the term "**synfire** chain" [14, 132].

Spatiotemporal patterns

The delay of a single spike transfer, that is, the time between a presynaptic somatic spike and the postsynaptic somatic spike, is estimated in the scale of < 10 milliseconds [37, 132]. If a sequence of multiple spikes lasts for hundreds of milliseconds, and each spike requires the arrival of tens of presynaptic inputs, then the flow of activity must involve a large number of sequential synaptic activations. In addition, hundreds of neurons linked by single synapses will often fail to propagate single spikes due to the weak amplitude and frequent failure of single synapses. Thus, a chain of neurons linked by single synapses seems like an unlikely structure to be underlying long spiking patterns.

On the other hand, neurons in synfire chains activate reliably and at precise times relative to one another. Each neuron activates in synchrony with respect to its own group and at a fixed delay with respect to other groups, which results from the expected monosynaptic delay scaled by the number of groups between them. Long or short delays between spikes in a spatiotemporal pattern can thus be explained by a larger or smaller number of groups in the chain. Thus the synfire model can provide an explanation for the occurrence of spatiotemporal spike patterns by assigning the neurons that contribute to a pattern a particular position in the chain. Note that synaptic delays are assumed to be short, but not all neurons may be experimentally

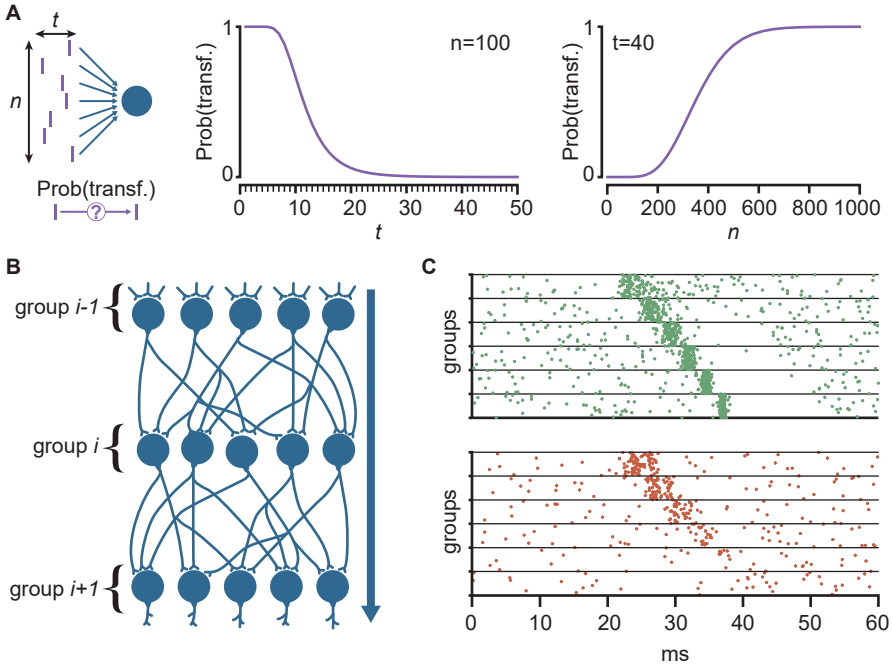

Figure 3.5: Synfire chains. **A.** Left: schematic of a neuron receiving a volley of input spikes with a fixed number of spikes (n) and spread over a period of time (t). Probability of a spike transfer (*Prob(transf.)*) is the probability that a given presynaptic spike within the volley triggers a postsynaptic spike. Middle: Probability under volleys of 100 spikes with different temporal spreads. Note t is a discrete variable representing the number of time bins. A time bin is idealized to correspond to the duration of a single EPSP duration (see text). A spike transfer is almost guaranteed if the volley is concentrated in 4 time-bins, but it has a probability of 1 in 1000 if it is spread over 40 time-bins. Right: Probability under volleys of different sizes in a fixed time window of 40 time-bins. Note that at this temporal spread, compared to middle, we need 6 times more spikes to guarantee a transfer. **B.** Schematic of synfire chain with a convergent-divergent feedforward architecture resulting in the reliable propagation of neuronal activity. All neurons are excitatory. Neurons in each horizontal group receive inputs from most neurons in the previous group (convergence) and project to most of the neurons in the next group (divergence). **C.** Example of successful (top) and unsuccessful (bottom) propagation in a synfire chain, depending on the properties of the initial signal. Top: a volley of spikes forced onto most of the neurons in the first group is able to cause activation in the next group. Because the volley is sufficiently big and temporally concentrated, synchrony and neuronal participation increase as the volley propagates in the chain. Bottom: the initial volley is too sparse and temporally dispersed, so propagation of spikes becomes unstable and dies out. **B** and **C** modified from [15].

observed, which gives this explanation a high descriptive power to accommodate spatiotemporal patterns with a wide range of delays.

In addition, neurons may belong to multiple chains. Since a single neuronal activation does not trigger postsynaptic spikes, each activation needs to be contextualized by the spikes of the rest of its group. Consequently, a neuron may belong to multiple groups, each part of a different chain, as long as the next steps in each of those chains do not overlap in excess. Through this mechanism, the synfire model can explain the possibility that a neuron participates in multiple, different spatiotemporal patterns.

It is important to note that synfire activations are built on a cooperative effect. Indeed, the spike train of a single neuron cannot be used to infer the membership of a neuron to any synfire chain. On the other hand, since experimental techniques can only sub-sample the activity of a population, the observation of a single spike absent of population synchrony or external cues does not exclude that this spike belongs to a synfire chain. Note that since a neuron can belong to multiple groups and the number of neurons can be very high, the total number of synfire combinations becomes extremely large, which places an extra burden on the experimental sampling of neuronal activity.

The synfire chain model is fundamentally designed not to need that the synapses within the chain be more powerful than the synapses found elsewhere in the network. In addition, it can operate with partial activation of each group. Thus, if we were to examine the connection between two neurons in two consecutive groups under forced stimulation of the presynaptic one, we would only observe a small EPSP and a weak cross-correlation. If the entire presynaptic group activates, however, the accumulated EPSP will be large and the cross-correlation strong.

Early experimental studies of spatiotemporal spike patterns of simultaneously recorded neurons showed that it is common for a single neuron to contribute many spikes to a given pattern [135, 136]. This behavior may be explained, according to the synfire model, by the involvement of a neuron in numerous different neuron groups in succession [132]. The transmission of synchronous activity does not get disrupted at any point along the chain as a result of the infrequent reappearance of neurons in multiple groups. One can imagine that synfire chains wrap around over themselves, reusing different combinations of the same neurons to design groups that are minimally overlapping, resulting in a structure that is feedforward in practice but that appears recurrent when measuring connectivity. This structure is called "synfire reverberation" and has been shown to be plausible through computational simulations [136].

The sequential schematic of a synfire chain (Fig. 3.5B) does not correspond to the anatomical placement of the neurons. Indeed, single neurons may reappear in multiple groups, even within the same chain. In this model, neurons may hold completely arbitrary locations. Classical synfire work refers to the groups of neurons as "groups", "nodes", or "pools", instead of layers to avoid confusion with the cortical layers and the layers of artificial neural networks ([14, 137]. The expected

size of these groups is called the "width" of the chain, which, in most modeling studies, is around 100-300 neurons [14, 15, 138]. Although most models explore fully connected ("complete") chains, these are improbable in biology. Incomplete ("diluted") chains with only a small percent of all possible connections (around 1 in 3) can still produce reliable propagation, but they require that the width of the chain is scaled up in order to keep the expected net input onto individual cells above the spiking threshold [14].

Finally, some theoretical studies have shown that synfire-like structures can also occur as a result of plasticity rules. Certain learning rules, such as spike-timing-dependent plasticity (STDP), have a tendency to develop strong unidirectional connections on random recurrently connected networks [139]. Combining this rule with homeostatic or heterosynaptic mechanisms and random spontaneous activity can turn the recurrent network into a mainly-feedforward structure that behaves similarly to classical synfire chains [133, 140, 141].

Limitations

Although simulations and theoretical studies have shown that synfire chains are able to propagate neuronal activity reliably, they are highly organized structures for which we lack anatomical evidence. These structures typically require highly dense or even all-to-all connectivity [12, 138, 142–144]. This requirement might be relaxed by using stronger synaptic efficacies or by increasing the width of the chain, but for synchrony to propagate, the underlying structure must remain largely feedforward [15,144]. Evidence on cortical connectivity from multi-patch experiments suggests a highly recurrent network with frequent bidirectional connectivity [68]. Electron microscopy studies that are able to reconstruct the dense connectivity in a small volume also suggest high degrees of recurrency [96]. Because the dense sampling of connectivity at a large scale remains an open problem, most attempts at detecting synfire chains have focused on detecting synchrony above chance levels in the activation of neuronal populations [135, 136, 145–148].

Could synfire chains hide within the recurrent connectivity of the cortex? Computational simulations where feedforward chain structures are embedded into recurrent networks have highlighted severe limitations for the reliability of propagation in that case. On the one hand, the synchronous activation of a large group of neurons tends to ripple within the recurrent structure, creating a spreading wave of activity that may result in a pathological state of "synfire explosion" [142, 143]. On the other hand, spurious correlations in the background activity of the recurrent network may cause some of the pools of neurons of the embedded chain to activate in sufficient amounts to ignite the chain or, at least, create echoes that die out [12, 143]. To circumvent these issues, models typically resort to using different synaptic efficacies for the recurrent and feedforward part of the model so as to enhance chain propagation and reduce its attachment to the background network [142, 143]. That is, the synfire chain model proposed synchrony as a way to solve the issue that connections are too weak to ensure propagation, yet attempts to model synfire chains under

more realistic connectivity required the presence of a heterogeneous distribution of synaptic strengths where some must be stronger than others.

In summary, the synfire model models where neurons are connected in a feedforward structure have shown that reliable propagation of neuronal activity through excitatory connections is possible [131, 132]. These transmission lines are composed of groups of neurons that fire synchronously, and where each pair of consecutive groups share feedforward divergent-convergent connectivity. This model can account for the millisecond accuracy of some of the neuronal firing patterns that have been observed experimentally [14,128,146]. However, even though this model has gathered a lot of attention from theoretical and computational neuroscientists, there is still very little or inconclusive experimental data to support it [135, 136, 147]. Furthermore, modeling studies that attempt to use synfire chains under more realistic features of connectivity, such as the strong recurrency and frequent bidirectionality of cortical networks, find severe limitations in the capacity of the model to operate as theoretically desired.

3.3.2 Polychronous chains

Classical synfire chains make two core assumptions on how homogeneous synaptic properties are: all synaptic connections are very weak, and synaptic delays are negligible (or constant). As described in the section above, embedding synfire chains into recurrent networks typically leads to a relaxation of the first assumption in order to segregate feedforward and recurrent connectivity. Studies that relaxed the second assumption led to the development of a different, closely related model: polychronous chains [7].

In polychronous chains, the synaptic delays between multiple neurons in a feedforward network are assumed to be heterogeneous. As a consequence, even if neurons in a group fire synchronously, the spikes received by a common postsynaptic neuron will look asynchronous and may not add up to produce a spike transfer. Conversely, if synaptic delays are heterogenous but fixed, there must exist a particular temporal order in which the presynaptic group can fire so that all spikes reach their common postsynaptic target at the same time (Fig. 3.6).

Interestingly, the heterogeneity of synaptic delays opens a new possibility: a given group may have multiple common postsynaptic targets but a different set of delays for each one. Thus, although the spikes are propagated to all of the targets, only one of them will receive them synchronously, depending on the particular spatiotemporal pattern of activation. This feedforward structure is called a "polychronous chain" since the groups of neurons may activate with multiple (*poly*) temporal delays (*chronous*) to generate reliable propagation of activity. Note that this form of activation is different from what is often described as "asynchronous": polychronization assumes diverse temporal patterns that are somehow fixed, while asynchronous activation often refers to variable, noisy, delays [7].

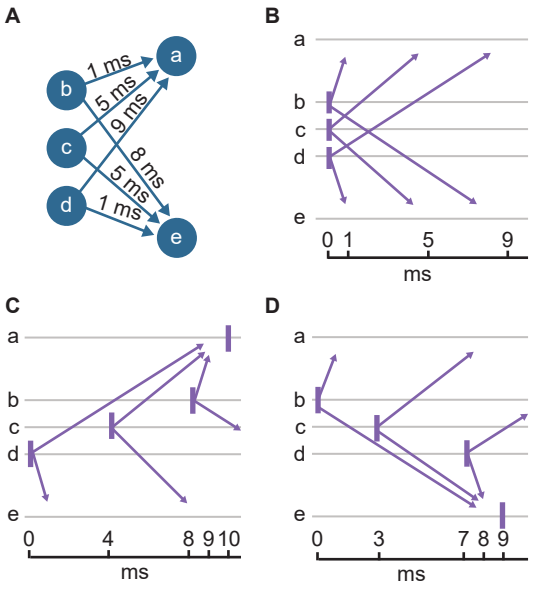

Figure 3.6: Polychronous groups. **A.** Schematic of a hypothetical polychronous group (neurons b, c, and d) that project with a diverse but fixed set of synaptic delays onto two different partners (neurons a and e). **B.** The synchronous firing of the polychronous group is unable to cause activity propagation on either postsynaptic target due to the synaptic delays. **C-D.** Only firing in a particular temporal pattern can the polychronous group activate one of its postsynaptic partners, but not both, resulting in the routing of activity in two mutually exclusive paths. Modified from [7].

Note that all the benefits from the synfire model in terms of explanatory power also apply to polychronous chains. That includes explaining long and temporally precise spatiotemporal patterns where neurons may appear more than once in the same pattern or participate in multiple different patterns (see section 3.3.1).

In addition, the original study of polychrony included the development of polychronous groups through spike-timing-dependent plasticity (STDP) to produce feedforward structures in a recurrent network [7]. Depending on the temporal duration of EPSPs, the number of delay combinations becomes huge, exceeding by several orders of magnitude the number of neurons in the entire network. In addition, the network developed gamma-like rhythms and a form of excitatory/inhibitory balance.

Although the polychronous model captured several interesting features of cortical dynamics, it remains poorly explored in comparison to the synfire chain field. The initial model was built over a simplified recurrent network with powerful connections and a large range of synaptic delays (0.1–44 ms). Interestingly, very recent work combining modeling and experimental measurements has suggested that polychrony may underlie sequence propagation within the HVC area of the songbird brain [129]. Local axonal collaterals in HVC are unmyelinated, resulting in slower activity propagation over long distances. Consequently, a series of Electron Microscopy reconstructions, together with compartmental modeling, suggested a long-tailed distribution of synaptic delays. Using those slow and heterogeneous delays in network models with a spatial component, the authors found polychrony as the most consistent model for spatiotemporal patterns of activity in the songbird HVC.

3.3.3 Rate propagation

The models described in the previous sections deal with precise spike times and can thus account for spatiotemporal patterns as observed in many cortical areas [126, 148–156], and that are the main focus of Chapters 4 and 5 of this thesis. However, a very different form of propagation is often considered in neuroscientific literature: firing rates. Although there is no unified definition of firing rates, at least three versions are often used in describing experimental and modeling results (Fig. 3.7) [3, 157]:

(1). An average of the number of spikes over a certain period of time.

(2). An average over a population of neurons.

(3). An average over several repetitions of an experimental trial.

Note that all three definitions may potentially be applied in a sliding window to include a temporal component and result in a time series.

If we consider that the membrane potential of a neuron acts primarily as an integrator of inputs [134], in the cases of definitions (1) and (2), a firing rate time series might be interpreted as a proxy for the dynamics of an ideal postsynaptic target.

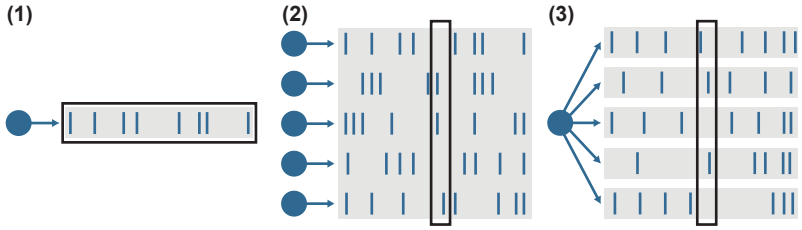

Figure 3.7: Alternative definitions of firing rate. Circles indicate neurons. Vertical ticks indicate spikes. Gray shadings indicate trials. Black boxes indicate what spikes are counted to estimate the firing rate. Modified from [3]. The three definitions are **(1).** Average over a time period. **(2).** Average over a population. **(3).** Average over repetitions.

Nonetheless, firing rates as such do not physically exist: they are averages, that is, the mathematical quantification of a set of real physical phenomena (spikes) [3]. This makes firing rates a concept that belongs predominantly to the "external" frame of studying the brain (see section 1.2) since they do not exist without an external observer, that is, the experimenter or modeler quantifying the spike train. This is particularly important in the case of definition (3), where one might take firing rates as the estimation of the probability that a spike occurs in any one trial. Indeed, the postsynaptic recipients of a hypothetical spike train do not have access to the probability across trials but rather only to the specific train in each trial independently.

Maybe because of this reliance on an external frame, discussions of propagation through firing rates are often tightly linked to the question of neuronal coding, that is, what type of cipher the brain uses to represent the external world [15,158]. These discussions often place the precise timing of spikes and firing rates as opposing strategies for coding.

Modeling work has suggested that rate-coded signals can be reliably transmitted over multiple feedforward layers embedded in recurrent networks and that this reliable encoding model can be used to construct networks that compute logic gates [12] (Fig. 1.4). On the other hand, synfire-like structures failed to propagate reliably, running into the synfire explosions, reverberations, and die-out issues that I described in section 3.3.1. Further work by the same authors then showed that if these networks are constructed so that excitatory signals are canceled by inhibitory feedback (a setup known as excitatory-inhibitory detailed balance), then signals can also be reliably gated on or off by altering the strength of local inhibition [13].

Although the gating through balance is able to select among different rate inputs, it can not effectively gate short transient signals in the scale of tens of milliseconds [13]. Later work showed that synfire structures could instead solve this issue [159]. Additional modeling work suggested that feedforward structures may act in both modes of operation: combining synchronous (synfire-like) propagation with asynchronous (rate) propagation [15].

In line with the idea of combining precise timing and rate coding, experimental evidence for the presence of precise spike patterns at the onset of up-states but not later has led to the hypothesis of "packet-based" communication. In this hypothesis, information may be organized in discrete units that begin encoding using fine spike times and later transition to a rate-based encoding [156]. However, modeling suggests that switching between these modes requires the fine-tuning of multiple elements, most notably the connection probabilities, and synaptic strengths [15]. The way in which such quick switching of fine-tuning may be achieved remains unclear.

Although rate propagation fails to handle short timescales, it can be very effective at handling noise. For instance, the combination of experimental recordings and modeling suggests that a single excitatory cortical spike may result in 28 additional spikes in the network [158]. These spikes will, in turn, trigger other spikes, overall resulting in a reverberating amplification that, even if contained through local feedback inhibition, will lead to the variability of the membrane potential of neurons in the network. If there is even a small probability that single spikes are spontaneously emitted, with no relevance for information processing, then the resulting membrane variance of all neurons in the network contains no information, and any coding mechanism must be able to handle it. This variance will affect the capacity of the network to produce precise spikes, which makes rate codes a particularly appealing solution for propagating information. Alternatively, propagation must rely on strong depolarizing events of amplitude higher than this hypothetical non-informative membrane variance. The alternative of large depolarizations does not seem implausible given that current injections with random currents, including these large depolarizations, show that single neurons can produce temporally precise spike trains [160] and since long-tailed distributions of synaptic strengths are common across cortical areas [95]. In Chapters 4 and 5, I explore a model that shows rare strong connections as a reliable mechanism of propagation.

In summary, firing rates are an important framework for understanding the propagation of neuronal activity, but they are very tightly linked to the question of the neural code. This reliance on an external frame has led to confronting precise spike time models of propagation with models based on firing rates. Networks, feedforward or recurrent, can be set up to reliably propagate and gate firing rate signals. Firing rates are well suited to handle spontaneous spiking noise, but they can not account for the precise spiking patterns often reported in cortical networks or handle operations such as gating in short timescales. Overall, while precise timing and firing rates may have their respective advantages and disadvantages, understanding both concepts is critical for a comprehensive understanding of the propagation of neuronal activity in the brain.

4 A model of the turtle cortex that produces repeatable spiking sequences

Remark: Some of the methods, results, and figures in this chapter are part of an article entitled *Single spikes drive sequential propagation and routing of activity in a cortical network*, which was written together with my supervisor Prof. Dr. Julijana Gjorgjieva and feedback from my collaborators Dr. Mike Hemberger and Prof. Dr. Gilles Laurent. The article has been reviewed and published in the journal *eLife* [1]. All methods, results, and figures from that article that are part of this chapter were my contribution to the article unless specifically mentioned otherwise.

All experimental data described and analyzed in this chapter was originally collected by my collaborator Dr. Mike Hemberger and published in [37].

4.1 Overview

Both experimental and modeling methods have suggested that, in the face of noise and irregular activity, cortical circuits may depend on firing rates to transmit information reliably [158, 161, 162]. This implies that individual spikes may not be essential for computation. However, it has been observed that even a single spike can significantly increase network firing rates in the rat barrel cortex [158] and reliably trigger sequences of spikes in the human cortex [97]. This suggests that the precise timing of spikes may play a crucial role in cortical computations [15, 148, 163–167]. Since spikes are the main form of communication between neurons, understanding how they propagate within a network is crucial for defining the fundamental elements of cortical computation [3, 6]. How relevant are single spikes for cortical function?

The effect of individual neurons on network activity, brain state, and behavior has been observed in rat and mouse cortices through single-cell stimulation [168–172]. Likewise, recent experiments by my collaborator, Dr. Mike Hemberger, have demonstrated that even one to three spikes of a single neuron in the turtle cortex can reliably trigger activity sequences in the surrounding network [37]. Nonetheless, it remains unclear how these spiking patterns propagate, particularly in the presence of irregular and apparently noisy activity.

In this chapter, I describe my computational attempt at describing how these sequences are triggered and sustained and to provide experimentally measurable predictions about their properties. To explore the mechanism of sequence generation from single spikes, I constructed and studied a model network using experimental measurements as constraints. By replicating the experimental protocol in simulations,

I found that the model network is capable of producing repeatable sequences in response to the activation of individual neurons without the need for parameter searches and under the assumption of random connectivity. I could then analyze the properties of sequences as a function of model parameters, including the mean level of network activity that is uncorrelated to the trigger spikes. Finally, by analyzing activity propagation and building alternative models, I found a plausible mechanistic description of sequence propagation: few strong connections support sequence propagation, while many weak connections provide sequence flexibility.

4.1.1 Experimental background and modeling questions

The turtle visual cortex, which has only three layers, shares architectural similarities with the mammalian olfactory cortex or hippocampus and is evolutionarily connected to the six-layered mammalian neocortex [41,50] (Fig. 4.1A). Additionally, this cortex is suitable for long ex vivo experiments, in which local connectivity remains intact (section 3.1.1.1). My collaborators employed a multi-electrode array (MEA) under the ventricular surface of a slab of turtle dorsal cortex while simultaneously patching a single pyramidal cell (Fig. 4.1B) [28,37]. Applying short step currents in whole-cell patch-clamp mode, they could cause the patched pyramidal neuron to trigger one to three spikes while simultaneously monitoring neuronal activity with the multi-electrode array. They found that the network reacted reliably to these trigger spikes with a sequence of activations of other neurons. The triggered network activity was reliable in three ways: responses were repeatable across trials, the responses involved the same neurons, and their activations respected the same temporal order [37] (Fig. 4.1C).

While these experiments revealed striking reliability in the response of an ex vivo network to just one spike, they opened three additional questions that I computationally address in the results of this chapter (Fig. 4.1D). First, whether such reliability would be affected under more in vivo-like conditions. If the membrane potential of single neurons shows a high variance in vivo [161], it seems reasonable to assume limits to the reliability of response under those conditions. Second, how far the influence of the trigger neuron extended. The limited size of the MEA, smaller than the cortical slab, which was, in turn, smaller than the full turtle cortex, spatially restricted the monitoring of the network. Lastly and most importantly, how did activity propagate within the excitatory layer? The physical distance between the pyramidal cell layer and the MEA limited the observation of sequence propagation within the excitatory neuron population, with a majority of sorted units being inhibitory interneurons [26, 28]. Given that network responses included the activation of some of these interneurons with a delay of up to 200 ms, it seems reasonable to assume that spike propagation was not direct from the initial pyramidal neuron to each of the responding interneurons, but rather that activity first propagated over several pyramidal neurons within the excitatory layer. The question of reliable propagation within an excitatory cortical population is still

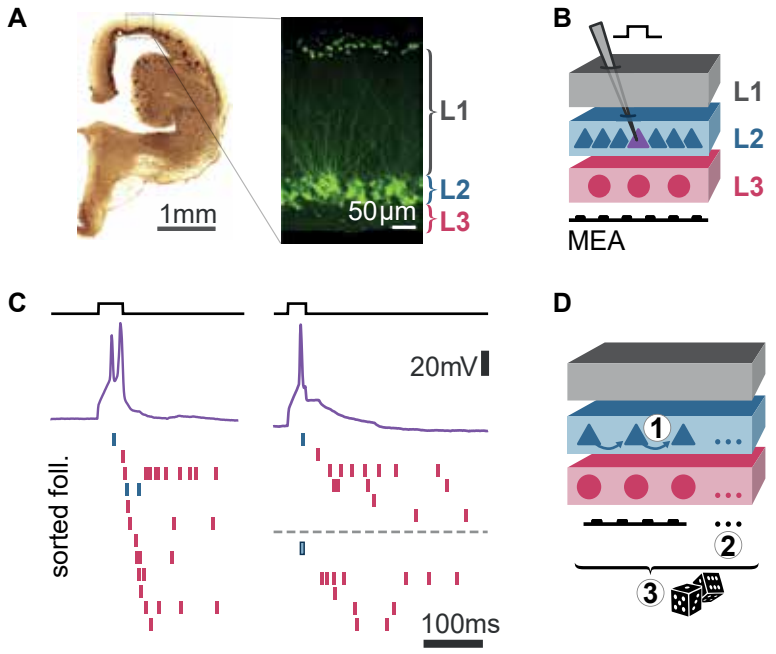

Figure 4.1: Sequence reactivation in the turtle cortex. **A.** Transverse section of turtle forebrain with zoom-in of the three-layered dorsal cortex. Modified from [28]. **B.** Schematic indicating the recording setup on an ex vivo slab of turtle dorsal cortex used in [37] (see section 3.1.1.1). A rectangular multi-electrode array (MEA) was placed under the ventricular surface, while a single pyramidal neuron was stimulated with square pulses under whole-cell patch-clamp mode. **C.** Examples of triggered sequences. Left: membrane voltage potential (purple) of a patched pyramidal neuron under a square current pulse (black). Ticks below correspond to spikes sorted from other neurons in the surrounding network, detected using the MEA. Blue ticks correspond to putative excitatory neurons, and red to putative inhibitory neurons. Right: two different trials (separated by a dashed line) where a patched pyramidal neuron was brought to spike once. Note the same network neurons react to this spike, and they preserve the same temporal order of activation. Modified from [37]. **D.** Schematic of the main questions addressed computationally in this chapter: (1) what is the network mechanism of excitatory propagation of activity from a single spike, (2) what are the properties of the resulting spiking sequences beyond what can be measured experimentally, and (3) how do these properties change in relation to uncorrelated ongoing network activity.

an experimental incognita and has been the subject of several theoretical models (section 3.3).

Reliable sequences of spikes have been observed in various in vivo systems, including replay or pre-play in the rat hippocampus [126, 149, 150], rat auditory and somatosensory cortex [155, 156], mouse visual and auditory cortex [153, 154], and human middle temporal lobe [152]. While spiking sequences are frequently associated with behavioral or sensory cues, the specific network mechanisms responsible for them are unclear. Various theoretical frameworks based on structured connectivity have been proposed to account for the consistent transmission of activity of these sequences. A classic mechanistic framework is that of synfire chains, which relies on divergent-convergent connectivity to link groups of neurons that fire simultaneously [14, 15, 138] (section 3.3.1). Another example is polychronous chains, which depend on precise time-locked patterns of neuron firing to account for transmission delays [7] (section 3.3.2). Yet, there is insufficient direct experimental evidence to support the existence of these particular connectivity structures [75, 146, 173]. Instead, it has been suggested that structured connectivity resulting in sequences can emerge through connection training [140, 174, 175]. Lastly, models have been developed based on the cortex of turtles, which explore the spread of waves throughout the network or the statistical characteristics of population activity, such as neuronal avalanches. [62, 176]. On the whole, these cortical propagation models center predominantly on organized population activity and have not explored how the activation of just a single neuron can initiate a repeatable sequence. This will be the focus of the present chapter.

4.2 Experimentally defined constraints on a random network model

I created a recurrent network model with single-cell features and connectivity constrained by previously acquired experimental data from the turtle visual cortex [37] in order to investigate the network dynamics that could result in reliable activity propagation from single spikes. According to experimental estimates of neuronal density, the network model's 100,000 neurons (93 percent excitatory and 7 percent inhibitory) correspond to a 2 by 2 mm slice of turtle visual cortex (Fig. 4.2A). The model network extended over a square with a side of 2000 µm where neurons were randomly placed and connected.

Based on experimental data that pyramidal (excitatory) neurons exhibit adaptive spiking, I modeled neurons as adaptive exponential integrate-and-fire point neurons (AdEx) (see Table 2.1 for parameters and Chapter 2 for the mathematical definition) [35, 37].

For the adaptation currents, I employed parameter values that were already discovered and discussed in the literature [35]. These captured well the median Adaptation index (0.3) of pyramidal neurons in the turtle brain that had been obtained experimentally. For both the experimental data and the model, the adaptation index was calculated using current injections into single neurons that

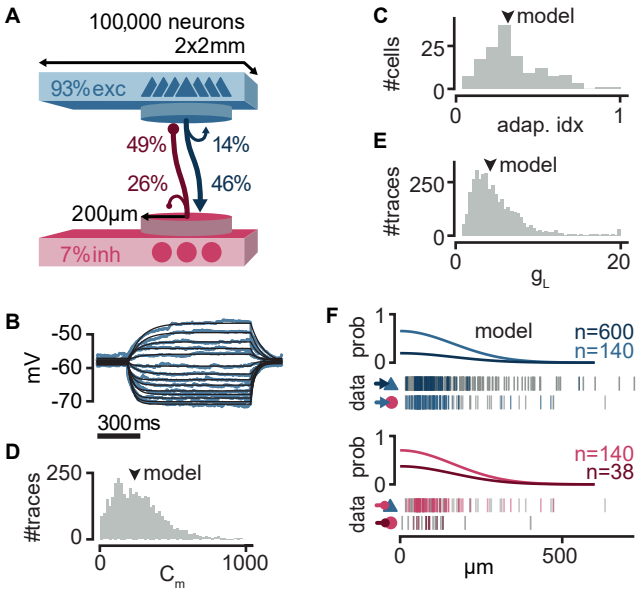

Figure 4.2: Network model. **A.** Diagram of the network, including the size of the network, the ratio of excitatory and inhibitory neurons, and the probabilities of connections within a 200 µm radius disc. The excitatory neurons are indicated by blue triangles, while the inhibitory neurons are indicated by red circles. The mean number of outgoing connections per cell for the different connection types are as follows: 750 for E-to-E, 110 for I-to-I, 190 for E-to-I, and 2690 for I-to-E. **B.** Model membrane potential responses (black) were fitted to those recorded experimentally (blue) in a neuron subjected to current injection. **C.** Distribution of adaptive indices for 145 recorded neurons. Arrowhead indicates the parameters of model neurons (median). **D.** Distribution of fitted membrane capacitance (C_m) ($n = 3886$ traces). **E.** Distribution of leak conductance (g_L) ($n = 3886$ traces). **F.** Gaussian profiles of the probability of connection as a function of the distance between neurons for different connection types. To fit the profiles, I extracted the ratio of pairs of neurons that displayed a connection relative to all tested pairs. Ticks below the profiles indicate which pairs of neurons were connected (colored) or not connected (gray) for each connection type that was tested at different distances. Blue: excitatory. Red: inhibitory. From [1].

lasted 1 second. From the resulting spike trains, the ratio of the time between the final two spikes and the time between the first two spikes was calculated. The strength of the current was about 2–5 times stronger than the minimum value required to trigger just one spike (the rheobase current).

Using prior experimental recordings of membrane potential made in current-clamp mode, I fitted the membrane capacitance and leak conductance parameters of the model (Fig. 4.2B-E). I used least-squares for fitting these parameters to current-clamp traces acquired under whole-cell clamp with varying amounts of current injection [37]. All neurons in the model shared the same values, the medians of the distributions. For a single isolated model neuron, the resulting membrane properties produced a 150 pA rheobase current.

I linked the model neurons using probabilities that decreased with their Euclidean distance. This distance-dependent connectivity was suggested by reconstructions of axonal arbors of principal cells in the turtle cortex [37]. I used periodic boundary conditions for this connectivity rule. In order to estimate the decay of the probabilities with distance, I used data from 918 pairs of neurons recorded under a whole-cell patch clamp, and I evaluated the connection probabilities at regular binned distances. After that, I used the method of least squares to fit the standard deviation of a Gaussian profile to those probability estimations (Fig. 4.2F). I adjusted the height of the Gaussian profile so that any particular disc with a 200 µm radius matched estimates of population-specific probabilities (Fig. 4.2A). This radio was chosen because most experimentally recorded pairs of cells fell within that distance. Because this entire process was random, from the placement of neurons to connections, and I did not enforce fixed degrees of connectivity, the resulting connectome presented stochastic inhomogeneities with some variance in the total number of connections that each neuron made. I avoided any autapses (self-connections) in the network.

There is both direct and indirect evidence from paired-patch studies in the turtle cortex that points to the existence of infrequent but strong excitatory synapses. This evidence comes from paired clamped cells that show either direct strong excitatory postsynaptic potentials (EPSPs) or poly-synaptic ones. This results in a distribution of amplitudes of EPSPs with a long tail [37]. In the AdEx model that I used, the synaptic conductance value and the synaptic time constant have an effect on the amplitude of EPSPs. I fitted both of these values using experimental data. In the first place, I used least-squares regression to fit single synaptic time constants to the rising timings of EPSPs acquired in paired patch recordings [37], and then I took the median of the distribution for my model synapses. Having fixed the time constant, I used least squares to estimate the synaptic conductances from experimentally recorded EPSP amplitudes. I used maximum likelihood estimation to fit a lognormal distribution to the resulting distribution of conductances (Fig. 4.3A). For each network instantiation, I sampled the strength of single model connections from this long-tailed lognormal distribution (Fig. 4.3B). If any sampled conductance resulted in a value greater than the maximum value fitted from EPSP amplitudes, I

re-sampled it. This limit was 67.8 nS, which corresponded to a 21 mV EPSP. Note that the main body of the distribution is still heavily placed in the low conductance values. As a result of this long-tailed distribution, each model excitatory neuron projects onto other neurons a great number of weak synapses and very few strong ones. However, since the model connectivity is dense, even though strong connections were relatively rare, any two excitatory neurons still shared, on average, approximately two connections that were so strong that they could push the postsynaptic neuron to near its spiking threshold (Fig. 4.3A inset, Fig. 4.3C). The real effectiveness of these infrequent but strong connections is contingent on the conductance state of the postsynaptic cell. Consequently, the level of activity of the network as well as the status of adaption currents may play a key role in whether these connections trigger a postsynaptic spike. In summary, the network connectivity was heavily based on experimental constraints and suggested dense weak excitatory connectivity overlapping with a sparse strong one.

Inhibitory connections were obtained by first making a copy of the distribution of excitatory conductances that was scaled up. Experimental estimates under voltage-clamp experiments suggest that inhibitory currents were 2-3 times greater than excitatory ones. Consequently, I scaled the inhibitory distribution by 8, which resulted in 2.5 stronger currents. Note that because synapses in the model are conductance-based, these currents depend on the distance between membrane voltage and the reversal inhibitory or excitatory potential. In an isolated model neuron that has been depolarized to near their threshold (at 50 mV), the scaled inhibitory synaptic conductances resulted in the strongest IPSPs of -21 mV, resulting in a symmetrical picture to the excitatory conductances, which, from resting, push the membrane potential by +21 mV. The skewed distribution of inhibitory synapses has been reported in other animal cortices [177–179].

The sampling of connection strengths was independent of distance for both excitatory and inhibitory connections.

The model assumed fixed but heterogeneous synaptic delays that were sampled from a uniform distribution in the range of 0.5 to 2 milliseconds. This accounted for the experimentally observed short and predictable lag between presynaptic action potentials and monosynaptic EPSPs.

In summary, the model displayed heterogeneous connectivity, with probabilities dependent on distance, variable number of connections, inhibitory and excitatory strengths following lognormal distributions, and delays following a narrow uniform distribution. Apart from these distributions, connections were completely random in each model instantiation.

Lognormal bootstrap

Under the connection rules introduced above, how realistic is it that a pair of neurons share a very strong connection? Might this model be strongly biased by the particular experimental dataset that I was using? I addressed these questions by bootstrapping the likelihood that an excitatory neuron sends one or more

strong connections onto another excitatory neuron (Fig. 4.3D). A connection was considered to be strong if its conductance value was within the range of 50.6 to 67.8 nS. This range corresponds, respectively, to the value in the top 99.7 percent of the originally fitted distribution (represented by the orange triangle in Fig. 4.3B) and the maximum strength that can be achieved by a connection, according to the experimental dataset. In the bootstrap step j, the count of connections that classify as strong s_{ij} of a model neuron i is distributed according to a Binomial distribution

$$s_{ij} \sim B(n_i, p_j)$$

where n_i is the total count of connections from neuron i and p_j is the probability that a connection belongs to the range of strong connections. I began with the experimental set of 122 EPSP amplitudes and then sampled from that set 122 times using replacement. I then fit a new lognormal to the new re-sampled data and extracted from it the corresponding p_j. My model neurons were each given a different number of connections, n_i, and these connections followed approximately a discretized Gaussian distribution with the parameters:

$$n_i \sim N(\mu = 745, \sigma = 27)$$

Thus, I re-sampled n_i 1,000 times for each bootstrapped step j. For each sample, I drew a corresponding s_{ij} from each $B(n_i, p_j)$. Finally, I approximated the probability of a strong connection in step j as the ratio of all s_{ij} that were greater than or equal to 1.

In total, I performed this process over 50,000 bootstrapped samples, generating a distribution of the chance that a neuron had at least one very strong excitatory-to-excitatory link. I discovered that this distribution was strongly skewed, with 39% of the bootstrapped fits giving a probability that was even greater than the model I had originally developed. Consequently, at least under the connectivity assumptions of my model, there is a high chance that any given excitatory neuron has at least one very strong connection.

Spontaneous activity

Neurons that were patch-clamped and recorded ex vivo showed a noisy resting potential and generated spontaneous spikes even in the absence of external stimuli [37]. In the model, the source of this spontaneous firing and noisy membrane dynamics was modeled as a white noise current that was sampled individually for each neuron, but that had identical statistics (mean and standard deviation) for all of them (Fig. 4.4A).

The model was instantiated in 300 different networks with unique random connectivity in each one. Each of these configurations was run 20 different times with unique, spontaneous activity, which resulted in a total of 6,000 single simulations (Fig. 4.4A).

Every simulation had unique, fixed statistics of the white noise current. These were sampled at random from uniform distributions. The mean current parameter

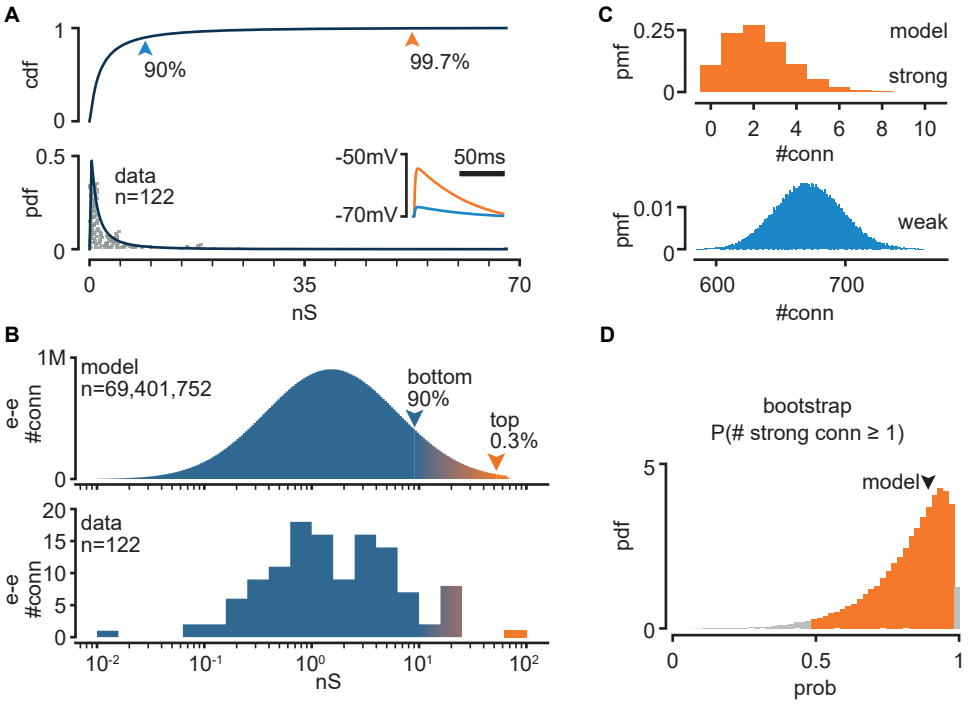

Figure 4.3: Long-tailed distribution of connection strengths. **A.** Lognormal fit to estimate the distribution of peak excitatory synaptic conductances. The top graph represents the cumulative distribution function (cdf), while the bottom graph represents the probability density function (pdf). The estimated conductances were calculated from experimentally obtained EPSP amplitudes recorded from pairs of connected neurons, shown as gray dots. The inset shows an example of modeled EPSPs at resting potential for different synaptic weights (top arrowheads) in an isolated postsynaptic neuron. **B.** Top: Connection strength distribution from an example model network. Bottom: Connection strengths estimated from data in log-scale (same data as **A**). Note both distributions are well approximated by a Gaussian distribution in log-scale. **C.** Probability mass function (pmf) of the count of weak (bottom 90%) or strong (top 0.3%) connections per excitatory model neuron onto other excitatory neurons. **D.** This graph shows the bootstrap estimate of the likelihood that an excitatory neuron will have at least one strong excitatory-to-excitatory connection. The colored area represents the 95% confidence interval, and the arrowhead indicates the fit with the original dataset shown in **A**. From [1].

was sampled in the range 50–110 pA (μ_{in}), and the standard deviation parameter in the range 0–110 pA (σ_{in}) (Fig. 4.4B). Every neuron in every simulation was injected with a random current that was drawn from a Gaussian distribution $N(\mu_{in}, \sigma_{in})$ and re-sampled independently once per millisecond.

The injected spontaneous currents caused many levels of membrane potential fluctuations and spontaneous firing, which were comparable to what was seen in the experiments (Fig. 4.4C). On the membrane potential variation of model neurons, the influence of the current mean was much more significant than the effect of the variance (Fig. 4.4D). This was the outcome of self-sustained firing in the network at sufficiently depolarizing currents. For instance, see the spikes in the bottom left corner of Fig. 4.4C. With just a few spikes in the network, under certain mean current values (μ_{in}), the network could sometimes display self-sustaining spontaneous activity. Note that this activity saturated quickly, and the network never displayed runaway excitation. However, when exploring extremely low mean firing rates, this self-sustained activity could result in lower activity early on and higher activity later on in the same simulation. To prevent a bias from this effect, the simulations started with a burst of spikes that were sent to a randomly selected group of 500 excitatory neurons in the first 100 milliseconds. The first 1,000 milliseconds of the simulation were ignored.

The random spontaneous drive caused in the model low firing rates on average, similar to those of the turtle brain. Note that in the experiment, these are likely consequences of extremely adaptable network and membrane dynamics combined with recordings that last for many hours [26, 37]. The model also generates bouts of greater instantaneous firing rates, which, as a result of the spatial component of the connectivity rules, take the form of waves of activity that spread over the plane (Fig. 4.4E). Similar fluctuations have been repeatedly reported in the ex vivo turtle cortex [114, 180].

In summary, the model that I developed contains properties of individual neurons and connectivity distributions that were restricted by biological data. The resulting network showed random and heterogeneous connectivity. The model was able to recreate real cortical waves of activity as observed experimentally. I next used this network model as a testing ground for investigating the mechanistic foundations of sequence triggering and propagation seen in the ex vivo turtle cortex under minimal stimulation.

4.3 Single spikes trigger reliable sequential activity in a biologically-constrained network model

Having constructed a bottom-up network model of the turtle cortex, I next evaluated if the biological restrictions were sufficient to reliably activate repeating sequences of firing neurons from a single spike in a single pyramidal neuron picked at random.

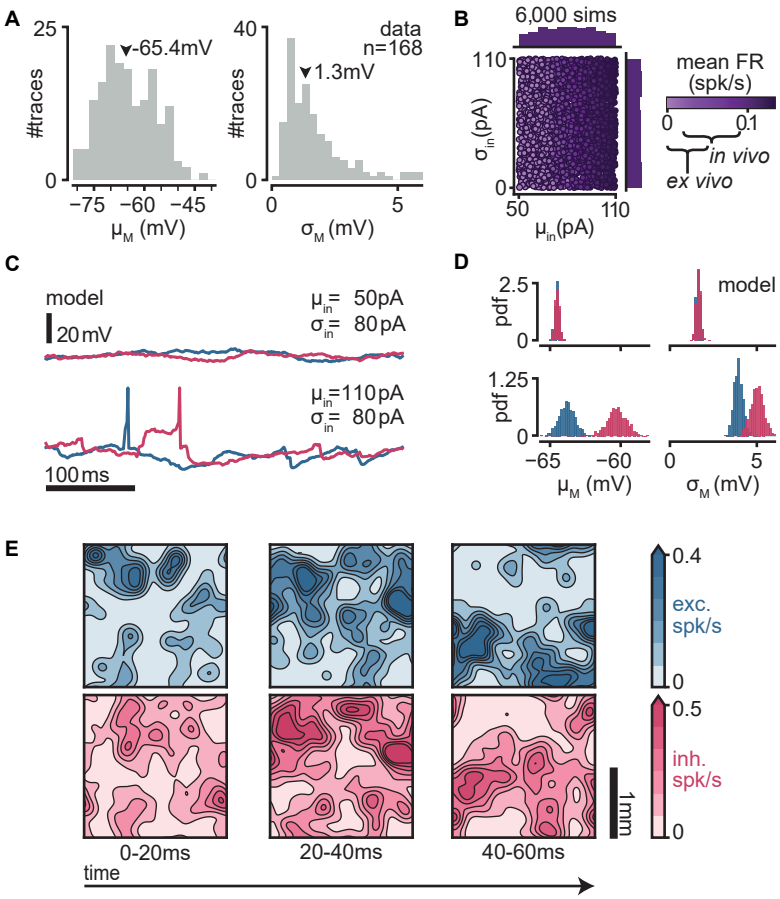

Figure 4.4: Network spontaneous activity. **A.** Standard deviation and mean of membrane potentials for 168 patched neurons in ex vivo preparations from experimental measurements. The arrowheads on the plot indicate the medians. **B.** Each dot is a single simulation with the given input statistics (μ_{in} and σ_{in}). Input statistics were sampled from uniform priors (marginal histograms), resulting in a wide range of network mean firing rates (color), including regimes compatible with estimates for the in vivo and ex vivo turtle dorsal cortex [37]. **C** The membrane potential of two types of model neurons, excitatory (blue) and inhibitory (red), under different white noise current conditions. When the mean current (μ_{in}) is high, action potentials and excitatory postsynaptic potentials (EPSPs) are present in the voltage traces. **D.** Distributions of statistics (mean and and standard deviation) of the membrane potential of model neurons under the same white noise current parameters in **C** (blue: excitatory; red: inhibitory). **E.** Spatial distribution of population firing rates for three consecutive time bins in the middle of an example simulation. Spikes are counted in a 2D grid of 1 µm. Counts are convolved with a 100 µm Gaussian filter and divided by the duration of the time bins (20 ms). Blue: excitatory. Red: inhibitory. From [1].

I chose one excitatory neuron at random from each of the 300 random networks mentioned above (Fig. 4.4B). I called this the "trigger" neuron. During each simulation, I forced the trigger neuron to fire 100 spikes, one at a time, at 400 ms intervals. I generated an instantaneous spike by setting the voltage of the trigger model neuron above its threshold. I define a "trial" as the time between each trigger neuron action potential. For each trial, I assessed the modulation of firing rate (ΔFR, Fig. 4.5A) displayed by each of the remaining network neurons in response to this spike. I split trials in a baseline time window of 100 ms before the spike and a post-spike time window of 300 ms after the injection of the trigger spike. The difference in firing rate between these two windows is the firing rate modulation (ΔFR), which I took on average across all trials for each of the other neurons in the network. To normalize ΔFR, I divided it by the value that would be obtained by an ideal neuron that never spikes before the trigger and always spikes exactly once after the trigger.

Followers

I carried out a statistical test on the distribution of ΔFR in order to locate consistently activated network neurons. The test relied on a null distribution for ΔFR, which was defined as the difference between two samples drawn from two Poisson random variables that had the same underlying rate. The two samples corresponded to the expected number of spikes before and after the trigger spike injection. The underlying rate for both random variables was determined by taking the mean firing rate of all of the model neurons. Because of the significant differences in the firing rates of the inhibitory and excitatory populations, I generated a different null distribution with a different mean for each population (Fig. 4.5B). Note that due to the difference in the length of the time windows before and after the spike injection, the rates of the two Poisson random variables have to be modified. The null distribution that was produced by this subtraction had a mean of zero but a standard deviation that was not zero, as seen in Fig. 4.5C. For the purpose of detecting spurious increases in firing rate from those caused by reliably responding neurons, I set the upper threshold of the null distribution at $p = 10^{-7}$. Neurons displaying a ΔFR above this threshold had a statistically significant increase in spiking as a consequence of the activation of the trigger neuron. I refer to those neurons as "followers". Since the model is composed of 10^5 neurons and I test all of them individually (with the exception of the trigger), one expects one false positive for every 100 simulations.

If each simulation had been run with a greater number of trials, a more accurate calculation of ΔFR and the identification of more followers would have been possible. Nevertheless, I discovered that 100 trials were sufficient to find the most reliable followers while restricting simulation periods (between 30 and 60 minutes for each simulation) and maintaining the experimental relevance of both my predictions and protocol.

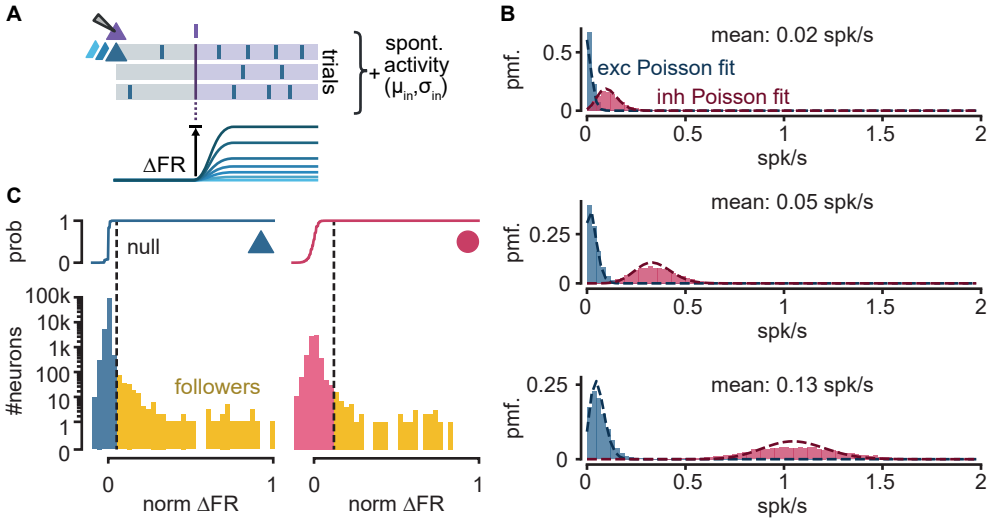

Figure 4.5: Definition of followers. **A.** Diagram of the stimulation protocol used in the model. The top neuron is stimulated to produce one action potential, indicated by a purple tick. The response of all other neurons in the network is then observed for changes in their firing rate. The entire simulation is carried out with random spontaneous activity present in the network with fixed standard deviation (σ_{in}) and mean (μ_{in}). **B.** Distributions of mean firing rates per cell for three example simulations at increasing firing rates. Top to bottom: 0.017 spk/s, 0.051 spk/s, and 0.13 spk/s. Dashed lines indicate Poisson fits. Blue: exc; red: inh. **C** Distribution of single-cell firing rate modulation (ΔFR) for a single simulation. The top panel shows the cumulative mass function of the null distribution, and the dashed line indicates the threshold for follower classification, which is set at $p = 10^{-7}$. The bottom panel shows the ΔFR for followers (shown in yellow) and other neurons in the network (shown in blue for excitatory neurons and red for inhibitory neurons). The ΔFR is normalized to the modulation of a perfect follower, which is exactly one spike per trial. From [1].

Because of the very low average firing rates that I explored (see next section), I was unable to identify neurons that lowered their firing rate in response to the trigger spike induction. Those neurons would be the reverse of followers, which increase their firing rate. On the other hand, a later set of simulations under a modified protocol (described in Chapter 5) revealed that neurons of this kind do, in fact, exist.

In 94.6 percent of my 6,000 simulations, there was at least one follower detected per trigger neuron. Note that this includes simulations spanning two orders of magnitude in their output firing rates (between 0.01 and 0.13 spk/s, Fig. 4.6A) and with a wide range of input current standard deviation (between 0 and 110 pA, σ_{in}). When I ordered the followers by their activation latency with respect to the trigger neuron spike, I found that not only these neurons responded reliably but that their spikes followed a sequence that was consistent from trial to trial, just as they were in the original experiments (Fig. 4.6B).

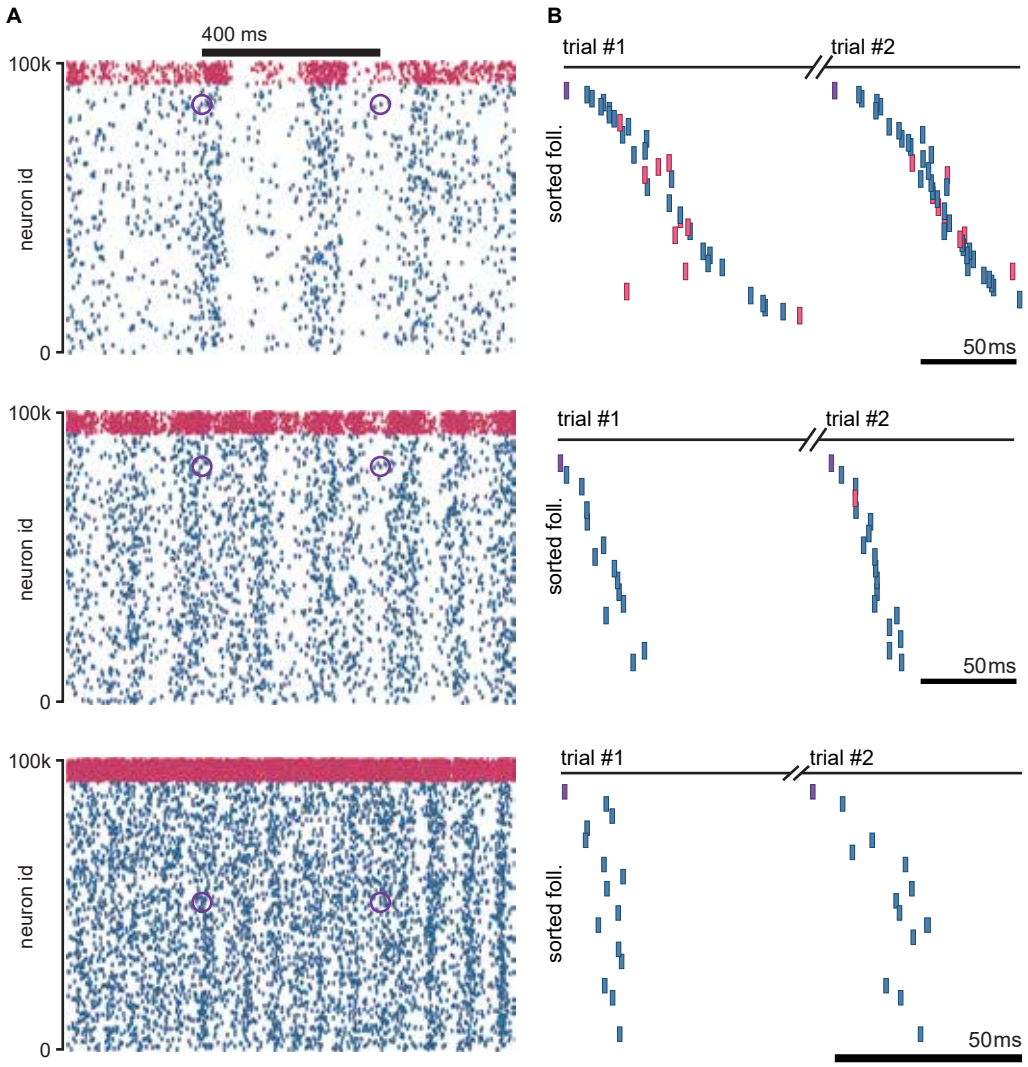

Figure 4.6: Spiking sequences in the model. **A.** Spike raster of three example simulations at increasing firing rates. Top to bottom: 0.017 spk/s, 0.051 spk/s, and 0.13 spk/s. Raster plots consisting of all spikes that were fired by each neuron in the network during an interval of 1 second in the middle of the simulation. Each raster contains two consecutive trials. Purple circles highlight the injected spikes of the trigger neuron. Blue: exc; red: inh. **B.** Top to bottom: Example sequences of spikes from followers in two consecutive trials in three example simulations (same as **A**). Y-axis: followers sorted by median spike-time across all trials. Same sorting in both trials. Spikes from non-followers are not shown. From [1].

Follower sequences

How reliable is the temporal order in sequences generated by my model? I employed the same entropy-based method that was used in the turtle experiments in order to quantify the variability of follower identification per rank in the sequences generated by my model. This allowed me to determine how much variation there was. Followers were ranked according to the time at which they had their first spike after the trigger spike. For each rank, I computed the entropy of follower identity as:

$$H_k = -\sum_{i}^{n} p_{ik} \log_2 p_i k$$

where k is the rank, n is the total number of followers in the given sequence, and p_{ik} is the estimated probability (frequency) that a particular follower i fired in that rank. The resulting value was then normalized by the maximal entropy H_k that would be expected from a uniform distribution:

$$H_u = -\sum_{i}^{n} \frac{1}{n} \log_2 \frac{1}{n}$$

If one observes the result $H_k = 0$, it means that the exact same neuron fired in the rank k during each and every trial.

It is important to keep in mind that this method allows me to compare to the values obtained experimentally but that it was designed for the extremely robust followers that could be detected experimentally, and it does not always behave well in a scenario where the network is fully observable. For instance, this metric is particularly sensitive to the activations or failures of followers that occur early within the sequence since these events have an effect on the rankings of all later followers who participate in the same sequence. Both the random exploration of varying fire rate levels (see next section) and random connection instantiations resulted in occasional sequence failure and total numbers of followers that were considerably larger than the number of trials (100). This discrepancy hampered my ability to correctly estimate order entropy in all scenarios. To reduce this effect, I restricted my estimates to simulations that included at least as many trials as followers, as well as simulations that included at least 25% of the followers in each trial.

The resultant spike-rank entropy was very similar to what was seen in the experimental findings [37]: the predictability of the follower ordering was highest for the first four ranks in the sequence (Fig. 4.7A).

Next, I repeated the study of the center-of-mass of follower activations in the same manner as it was done experimentally (see Figure 5A of [37]). This analysis demonstrates that, both in the model and the experiments, follower activations spread away from the source neuron. I estimated the center-of-mass of follower activations over time so that I could examine the spatial evolution of sequences as they develop. I determined the average X and Y position for those followers whose activity occurred during a sliding window of 5 milliseconds. The resultant X and Y

time series were then smoothed down with the use of a Gaussian window with a duration of 15 milliseconds. In accordance with what was found in the experiments, I observed that the sequences in my model develop over time and expand outward from the trigger neuron (Fig. 4.7B). In addition, approximately 18 percent of all followers in the model fall outside the area that could be accessed in the experimental recordings using a multi-electrode array of size 1.3x1.3mm (dashed line in Fig. 4.7C). The findings indicate that the initial experiments may have underestimated both the number of followers and the length of the sequence.

Effect of spontaneous activity

Experimentally detecting excitatory followers was hindered by the electrical distance that separated the multi-electrode array and the pyramidal middle layer; nonetheless, my model predicts that the quantity of excitatory followers far outnumbers that of inhibitory followers. As described in the previous section, I had already seen strong effects of the spontaneous random drive on the variance and mean of the membrane potential for both excitatory and inhibitory neurons in the model (Fig. 4.4). These altered membrane properties could, in turn, alter their capability to function as reliably responding followers. In order to evaluate how the amount of spontaneous activity in the network affected the number of followers, I examined the numbers for all simulations as a function of their mean firing rates (Fig. 4.7D). I found that when spontaneous firing rates were very close to zero, the model almost completely failed to generate any followers. However, the number of followers in the model increased very rapidly with mean firing rates and quickly peaked at rather low firing rates. The mean firing rate for the top 1% of simulations ranked by follower count was 0.007 spk/s. The number of followers then gradually dropped with further increases in spontaneous firing rates.

Next, I categorized my simulations into those with low activity and those with high activity. I defined these two regimes by the mean firing rates that had been reported in experiments for ex vivo ([0, 0.05] spk/s; [25th, 75th] percentiles) and in vivo ([0.02, 0.09] spk/s). From the random sampling of spontaneous activity parameters, I ended up with 41% of my simulations classified as low and 33% as high. Note that both ranges slightly overlap, which translated in some simulations classifying as both, while the model could generate firing rates even higher than in vivo, which translated in some simulations not being used for the quantifications below.

Sequences were present in both groups but exhibited properties that were noticeably distinct between excitatory and inhibitory model followers. The chance of a low-activity simulation exhibiting at least one or ten inhibitory followers (74 percent and 19 percent, respectively) was in very good agreement with actual experimental data (77 percent and 20 percent, respectively, Fig. 4.7E). To my surprise, I found that running simulations with higher levels of activity resulted in a decreased likelihood of obtaining inhibitory followers (4% probability of having 10 or more inhibitory followers). Excitatory neurons, on the other hand, show the reverse trend:

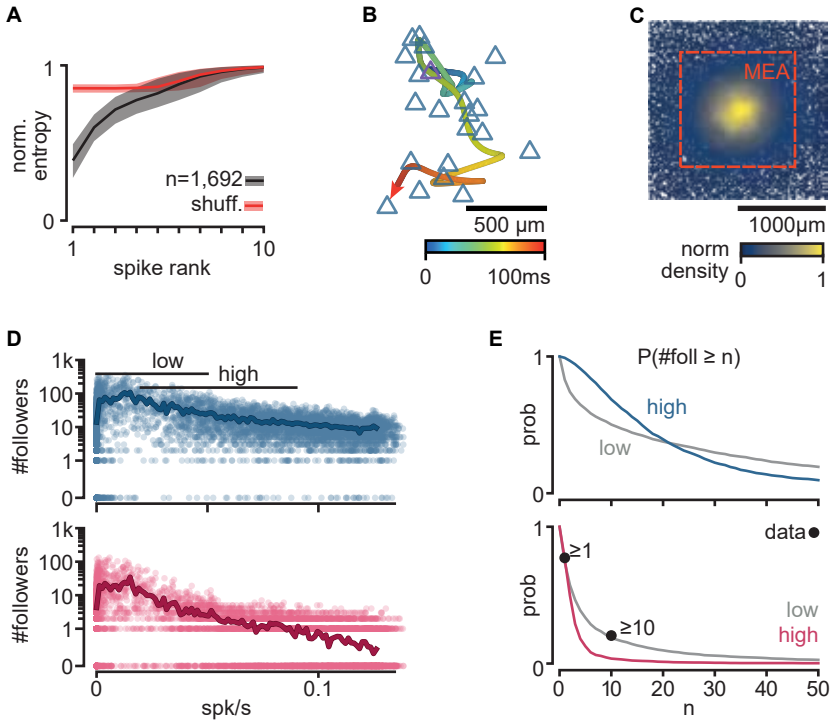

Figure 4.7: Properties of model sequences. **A.** Normalized spike rank entropy of simulated sequences (mean in black, standard deviation in gray). Normalized spike rank entropy of shuffled sequences is shown in red. **B.** Spatial evolution of the center-of-mass of follower activations during the first 100 ms of an example sequence. Trigger neuron in purple outline. Exc. followers in blue outlines. **C.** The location of excitatory followers relative to the trigger pooled from all simulations and colored by local density (n=117,427). Stippled square: MEA used in experiments. **D.** The number of identified followers in each simulation is plotted against the mean level of activity in the network. The excitatory neurons are in blue, and the inhibitory neurons are in red. The thick line represents the moving average **E.** The likelihood of producing a specified minimum number of followers in both excitatory and inhibitory groups in simulations under high and low levels of activity. Dots show the estimates from an experimental ex vivo setup [37] From [1].

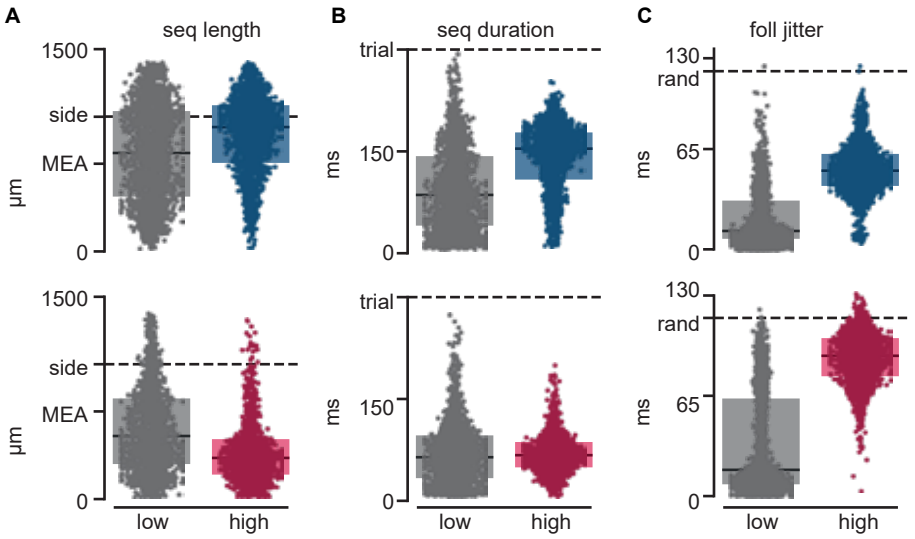

Figure 4.8: Sequences extent. Colored: high-activity simulations; gray: low. Blue: exc; red: inh. Boxes: median and [25th, 75th] percentiles. **A.** Distance from trigger neuron to the farthest detected follower. "Side" indicates the half-width of the area covered by the model; MEA: half-width of experimental MEA. **B.** Delay to median spike-time of the last activated follower. "Trial" indicates the maximum value imposed by the protocol. **C.** The average standard deviation of follower spike times. "Rand" indicates the standard deviation that is expected if spike times were randomly distributed within the trial. From [1].

the chance of having 10 or more excitatory followers was greater in high-activity simulations (69%) than it was in low-activity simulations (50%).

In both low- and high-activity simulations, the sequences of activations of excitatory followers often lasted longer than 150 milliseconds. They also reached the spatial limits of my model circuit, which extended on a plane of 2 mm by 2 mm, containing 10^5 neurons (Fig. 4.8A). On the other hand, inhibitory followers in the model network tended to fire early in the sequence and were located physically closer to the trigger neuron in comparison to excitatory followers (Fig. 4.8B). These disparities are even stronger when looking at high-activity simulations as opposed to low-activity ones. Under high-activity regimes, inhibitory followers exhibited near-random amounts of temporal jitter (Fig. 4.8C). which I quantified as the standard deviation of their spike delays with respect to the trigger spike. As a consequence of this, I hypothesize that we should be able to identify followers in trials conducted in vivo but almost exclusively among the excitatory neurons in the pyramidal layer.

Super-high firing rates

I finally investigated sequence propagation in simulations at rates that were far higher than the estimates of the spontaneous activity of in vivo turtles. Indeed, while the expected number of followers decays under increased rates (Fig. 4.7D),

the number of excitatory followers remains on the order of tens of neurons (see the spiking sequence in Fig. 4.6 bottom, with mean rate 0.13 spk/s).

I re-ran the three sample simulations from Fig. 4.6 under even larger mean currents (200–2000 pA, Fig. 4.9AB) in order to get an understanding of whether or not the existence of followers continued at even higher levels. An increase in the input variance did not have a significant impact on the firing rates. Mean inputs that were greater than the ones I present here resulted in highly unnatural synchronous firing. In this extended regime, inhibitory neurons fired at rates that were many orders of magnitude greater than the rates of excitatory units (mean rate: 27.75 inh spk/s; 0.36 exc spk/s, mean peak rate: 33.71 inh. spk/s, 2.51 exc. spk/s). The network-wide mean firing rate went up to 2.28 spk/s (example raster in Fig. 4.9C). At these rates, the Poisson that makes up the null model for detecting followers can no longer capture the shape of the distribution of excitatory firing rates, which becomes highly skewed and with high variance (Fig. 4.9D). This high variance results in a very wide distribution of ΔFR (Fig. 4.9E). Compare this to the distribution of firing rate modulation in Fig. 4.5C. As a direct result of this wide distribution, the simple statistical test was unable to discern between followers and non-followers.

To circumvent this limitation in this extended firing rate regime, I examined the neurons that were designated as followers in the original simulations. Interestingly, followers, although indistinguishable from background activity, often activated after the injected spike, even under the greatest drive. However, they activated in a very sparse manner, with each follower activating, on average, in just 15 percent of all trials. As a direct consequence of this sparse activation of single followers, the initial sequences were never displayed in their entirety but rather resulted in several partial failures in all trials. In Chapter 5, I examine in detail the mechanisms behind these types of sequence failures and propose them as a hallmark of flexible routing.

According to this exploration of super-high firing rates, discovering sequences within these regimes would be made possible by first identifying followers under lower firing rates. This conclusion is particularly important in the context of attempts to experimentally discover firing sequences. For instance, the turtle cortex may be particularly well suited for the study of such sequences precisely because it displays low spontaneous firing rates and strongly adaptive properties. Similarly, in the rat somatosensory cortex, repeatable sequences have been observed in vivo just after the transition from a state of low firing to a state of high firing states, but not later during the state of high firing itself [181].

Based on these results, it appears that my bottom-up model is still capable of producing partial sequences under spontaneous rates that are significantly higher than those of the in vivo turtle cortex. Nevertheless, this does bring up certain limits for follower identification that would require further work if the system being modeled had an operational firing rate much higher than the one studied here.

In summary, I developed a model that was biologically constrained and absent of structured connectivity but that still generates repeatable firing sequences of neurons in response to the injection of a single spike in most randomly chosen trigger neurons.

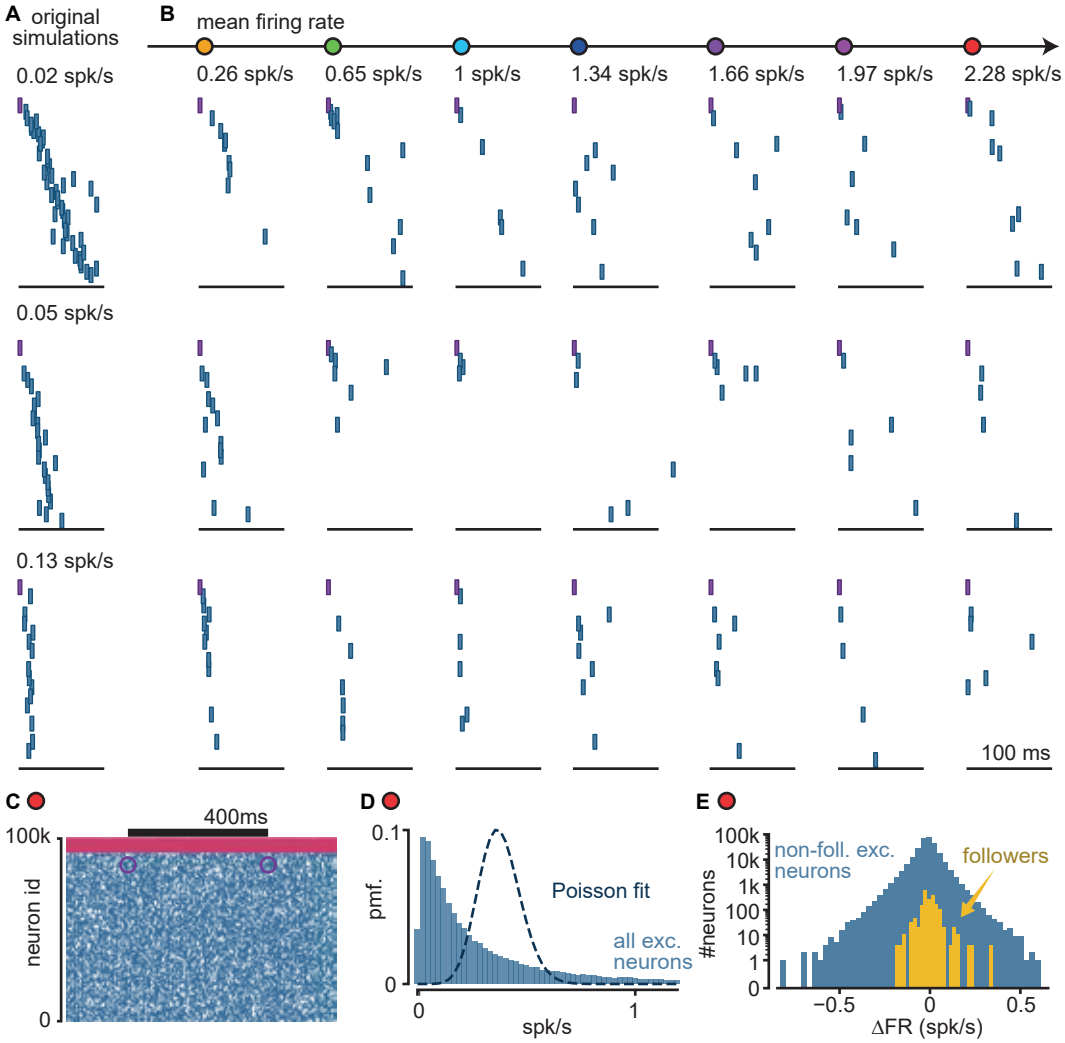

Figure 4.9: Evolution of sequences at rates above turtle in vivo estimates.
A. Example trials from original simulations at low, high, and super-high firing rates. **B.** Examples from each simulation after re-running them with a higher mean firing rate (specified at the top). Each row represents one of the three initial simulations in **A**. Only spikes from excitatory followers are shown. Purple: trigger spike. Within each row, the same neurons and the same vertical sorting. **C.** Spike raster of a simulation with 2.28 spk/s mean rate (red dot in **B**). The spikes of all neurons in the network are shown for an interval of 1 second. Purple circles highlight injected spikes. Blue: exc.; red: inh. **D.** Excitatory single-cell mean firing rate distribution is shown for one simulation with a mean network rate of 2.28 spk/s (represented by a red dot in panel A, where the mean excitatory rate is 0.36 spk/s). The dashed line represents the best Poisson fit. **E.** Distribution corresponding to the firing rate modulation (ΔFR) of excitatory single-cells for an example simulation. Mean rate is 2.28 spk/s (red dot in A). Yellow indicates the ΔFR for followers as identified in original simulations under lower activity. Blue indicates non-follower neurons in the network.

The model network was constrained by measurements obtained in an ex vivo setup where firing rates are low. However, by randomly sampling parameters of spontaneous activity, my model was able to make quantitative predictions on the properties sequences; namely, that sequences can occur even when the spontaneous firing rate is at in vivo levels and that, in that circumstance of greater baseline activity, the sequences are mostly made up of excitatory followers, while the responses produced by inhibitory followers are less reliable in terms of repeatable activation and of temporal ordering.

4.4 Strong connections provide the reliability of activity propagation, while weak connections modulate it

Next, I investigated the process through which spikes travel through my model networks to better understand the mechanics that underpin reliable follower activation. I found that the presence of converging weak connections alters the spontaneous activity level of the network and induces inhibition, whereas single strong connections are responsible for driving reliable excitatory responses.

Spike transfers

As a first step, I searched for instances in which a single spike successfully traversed a connection in a random subset of 900 out of my total of 6,000 original simulations. A successful traversal of a connection is defined as the firing of the postsynaptic neuron within 100 milliseconds of the firing of the presynaptic neuron. I refer to this event as a "spike transfer" from the presynaptic neuron to the postsynaptic one. I did this by mapping the spiking activity of the network onto the recurrent connectivity of the network. The result was a directed acyclic graph where single nodes are spikes, and edges correspond to connections that lead from spike to spike (Fig. 4.10A). Please take note that I investigated spike transfers among all of the neurons in the network. Note that a single connection may be represented as multiple edges in the directed graph if a connection is repeatedly traversed. This graph thus may characterize the properties of the synaptic connections that are traversed the most frequently.

The time between the pre and postsynaptic spikes corresponds to the delay of the spike transfer. The majority of spike transfers between excitatory neurons in low-activity simulations exhibited a delay of 6-8 milliseconds (Fig. 4.10B), which is in very good agreement with the delays reported in the turtle brain [37]. It is interesting to note that excitatory-to-inhibitory spike transfers display consistently shorter delays. This is also true at higher firing rates (Fig. 4.10B inset). This may be because inhibitory neurons are typically in a more depolarized state (Fig. 4.10D), which would cause them to require less time to go from resting to their firing threshold.

The following step was to extract the strength of the connections underlying spike transfers (Fig. 4.10C). As expected, the general shape of the distribution of strengths

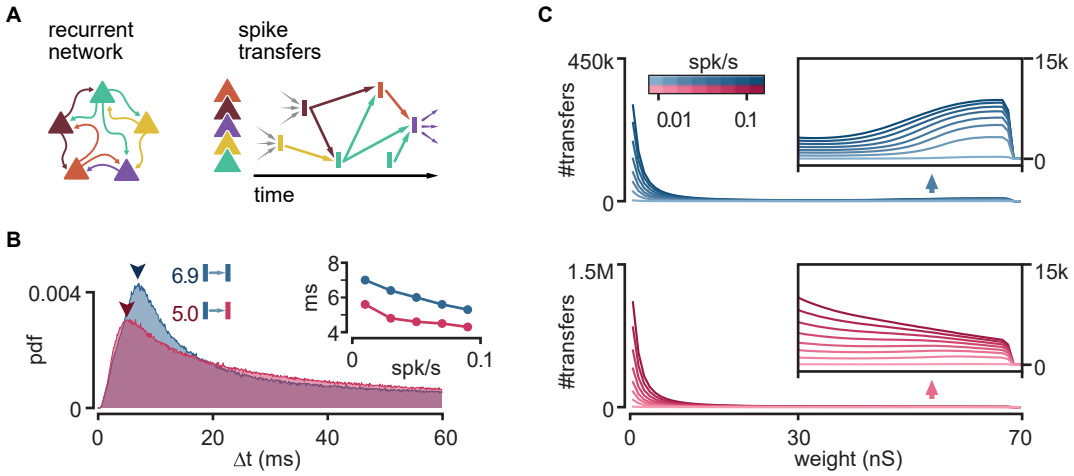

Figure 4.10: Connectivity underlying activity propagation in the network. **A.** Diagram of how the spiking activity and recurrent network structure are integrated to create a directed acyclic graph that contains the transfer of spikes between neurons. **B.** Delay distributions of spike transfers from excitatory to excitatory neurons (blue) or to inhibitory neurons(red). The data represents the average of all low-activity simulations with a mean firing rate between 0 and 0.05 spk/s. The arrowheads indicate the mode of each distribution. The inset graph shows the relationship between the most common delay and the mean firing rate. **C.** Histograms showing the strength of connections involved in excitatory-to-excitatory (blue) or excitatory-to-inhibitory (red) spike transfers. The inset shows the tails with a zoom-in of the abscissa axis. The lines are averages across multiple simulations after grouping them by their mean network firing rate (color bar).

of spike transfers mirrors the shape of the distribution of the full connectivity (Fig. 4.3A). It has a main body of very weak connections followed by a long tail. Surprisingly, the distribution of spike transfers between excitatory model neurons exhibits a secondary peak, which indicates an over-representation of connections within the tail of the distribution (inset in Fig. 4.10C top). On the other hand, the distribution for excitatory-to-inhibitory spike transfers lacks this secondary peak and shows a more populated body of weak connections (1.5M). This suggests that inhibitory spikes are mainly driven by the convergence of many weak connections (Fig. 4.10C bottom). At higher levels of activity, the difference between the distributions for each of these populations becomes even more pronounced.

I investigated the connectivity motifs present in the directed graph of spike transfers with the goal of gaining an understanding of the connectivity structure that underlies them. The greater the number of neurons that are engaged, the greater the potential number of motifs. Because of the combinatorial nature of the problem, I was only able to extract low-order patterns that included less than four spikes and had depths of up to two spike transfers (Fig. 4.11A). I found that excitatory spikes in low-activity simulations are unlikely to be triggered by fan or convergence motifs (Fig. 4.11B top, conv.). Instead, excitatory spikes are most often the result

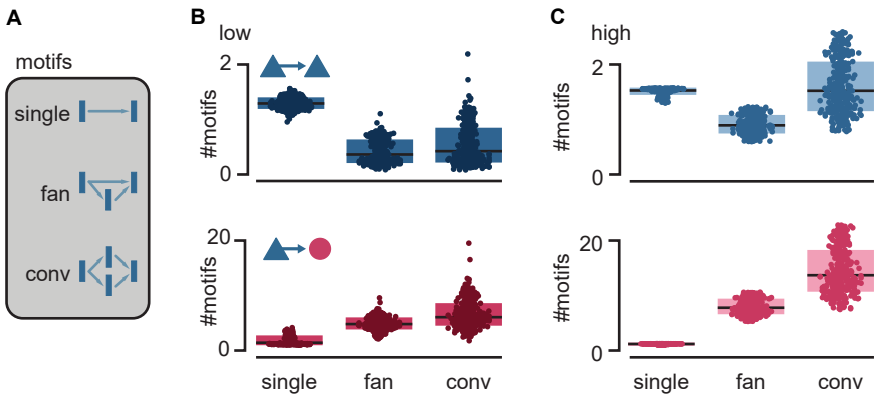

Figure 4.11: Motifs of connectivity leading to postsynaptic spikes. **A.** Schematic of motifs detected in the graph of spike transfers (conv: convergence). **B.** The number of motifs leading to an excitatory (blue) or inhibitory (red) neuron spike. Data points correspond to the mean value across all spikes in each simulation. Low-activity simulations. Boxes: median and [25th, 75th] percentiles. **C.** Same as **B** for high-activity simulations.

of single one-to-one spike transfers (Fig. 4.11B top, single). In contrast, I observed that inhibitory spikes are the consequence of interactions between multiple spike transfers, with a dominance of convergence motifs (Fig. 4.11B bottom, conv.). In high-activity simulations, spike transfers through convergence are more prevalent for both populations (Fig. 4.11C), although the increase is much higher for motifs that lead to inhibitory spikes. Single spike transfers producing inhibitory spikes in high-activity simulations are extremely uncommon (Fig. 4.11C bottom, single). These results show that the activity propagation in these networks is driven by a variety of distinct excitatory and inhibitory connection patterns.

In conclusion, this motif analysis gave me the ability to make a number of important hypotheses regarding the ways in which excitatory and inhibitory neurons participate in the propagation of sequences. This propagation involves differentiated types of motifs and strengths of connections. On the one hand, the activation of inhibitory neurons is often caused by the convergent activation of a group of weak connections. Since there is a relatively small number of inhibitory neurons and a high possibility that excitatory neurons would synaptically target inhibitory neurons (Fig. 4.2A), inhibitory neurons are very responsive to spontaneous network activity. If weak connections are as effective as driving inhibitory spikes as strong connections are, then it is more likely that inhibition is indeed driven by weak connections since these are several times more common. This effect is made more pronounced by adaptation properties that would prevent inhibitory neurons from firing too frequently. As a consequence of this, I conclude that inhibitory neurons can not be selective to their excitatory trigger, which reduces their potential to become followers. On the other hand, I discovered that excitatory neurons are less vulnerable to the constant activity of a network than inhibitory neurons. As a result, single

strong excitatory inputs constitute the predominant force behind excitatory spikes and, as a result, the most probable mean of sequence propagation.

Alternative models

In order to determine whether or not strong connections have a role in the propagation of sequences, I developed many variants of my model in which I eliminated excitatory-to-excitatory connections based on the strengths of those connections (Fig. 4.12A).

The previous analysis showed that there are two main modes in the distribution of spike transfer strengths for excitatory spikes. I used these modes to define "weak" and "strong" connections as, respectively, the bottom 90 percent and the top 0.3 percent of all connections. Removing all connections, except the strong ones, resulted in a very sparse excitatory network; specifically, 61.4% of excitatory neurons get fewer than 2 excitatory connections (Fig. 4.3C). On the other hand, removing all connections except weak ones resulted in a well-preserved network. In these two alternative models, I did not change the connection from or to inhibitory neurons in any way. After modifying the connectivity, I re-ran a randomized selection of 2000 simulations from the original set of 6000 simulations and extracted the number of followers in each one for each of the two new versions of the model. I found that networks with only weak connections had to be driven with a much stronger spontaneous mean current to produce any spiking activity (Fig. 4.12B). Yet, they produced almost no excitatory followers, with not even one of the 2000 simulations generating more than 10 followers (Fig. 4.12C). By contrast, networks with only strong connections showed a number of followers comparable to the original set of simulations with full connectivity (Fig. 4.12D). As a result, I have come to the conclusion that the presence of sparse connections that are powerful enough to elicit a postsynaptic spike is both essential and sufficient for the production of repeating sequences in this model.

As compared to the whole model, it is interesting to note that the model that only contains strong connections often generates a greater number of followers. This suggests that weak connections not only do not contribute to the presence of followers but actually reduce the reliability of network neurons under typical conditions. Getting rid of weak connections will make the input/output curve of the network slightly flatter (Fig. 4.12B). As a consequence of this change in slope, networks consisting of solely strong connections continue to maintain low firing rates for a wider range of inputs, in which a greater number of neurons respond reliably to trigger spikes (Fig. 4.7D). In addition to the slope change, removing weak connections also produces a shift of the input/output curve towards higher inputs, which increases the overlap between the range of sampled mean currents and the range of inputs where output activity is almost zero. Note that at these values, the network is almost completely silent and barely generates any followers. Thus, this shift defines a narrow range of mean currents under which the presence of weak connections can actually increase the number of followers rather than decrease it (Fig. 4.12E). Thus, weak

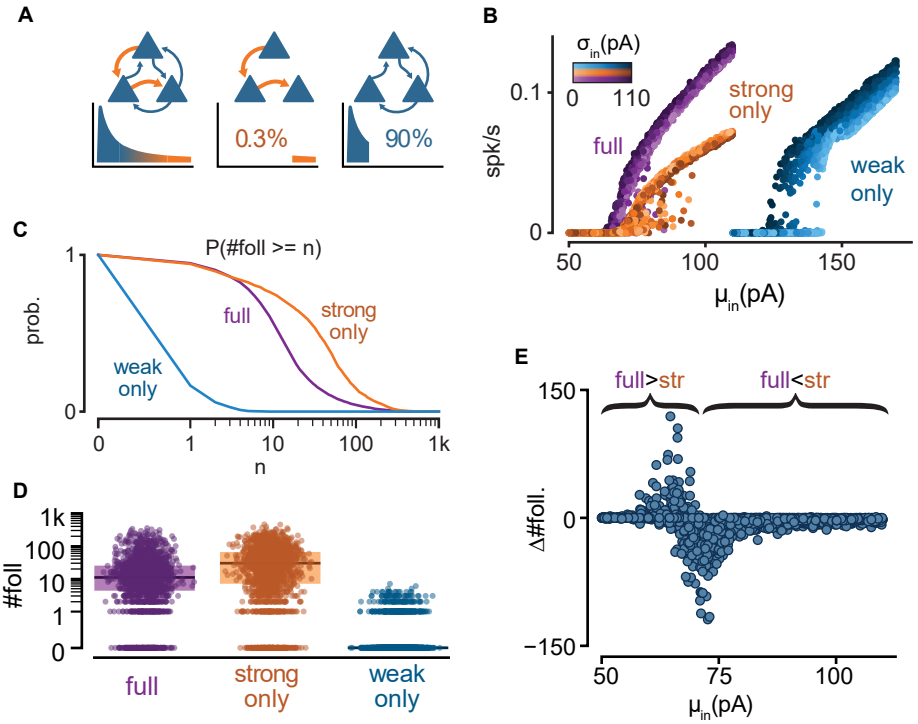

Figure 4.12: Connectivity underlying spiking sequences in the network. **A.** Schematic of alternative network models and their truncated distribution of synaptic strengths. **B.** Input-output curves for the full (purple), strong-only (orange), and weak-only (blue) models. **C.** The number of detected excitatory followers per simulation (n=2,000 each). Boxes: median and [25th, 75th] percentiles **D.** Probability of obtaining at least n followers for each model. **E.** The difference in the number of followers detected in 2,000 simulations of full and strong-only models under the same spontaneous drive. The brackets highlight ranges of inputs where the absence or existence of weak connections decreases (full < str) or increases (full > str) the number of detected followers.

connections have the capacity to both increase and decrease the total number of followers, depending on the specific input value presented to the network.

In summary, the propagation of activity in the presented model of the turtle cortex may be affected differently by connections that are frequent but weak and by connections that are rare but strong. Weak connections are able to amplify ongoing spontaneous network activity and drive recurrent inhibition. These effects effectively modulate the reliability of follower responses to single spikes. Strong connections, on the other hand, drive postsynaptic excitatory spikes in isolation and are thus the main conduit of reliable propagation. In conclusion, these two kinds of connectivity play distinct but complementary functions in the reliable propagation of activity throughout the network.

4.5 Discussion

In this chapter, I discussed my computational approach to investigate the potential mechanisms through which reliable firing sequences are initiated from a single cortical spike. The goal of my modeling was to determine the potential mechanisms behind this reliability, as well as to provide predictions for the presence and properties of sequences beyond the limits of experimental accessibility.

I used a computational model of a randomly connected network that I then constrained by experimental measurements taken from the visual cortex of the turtle. I found that thanks to an experimentally-observed long-tailed distribution of synaptic strengths, this model could generate reliable sequences from single spikes even without structured connections. This model provided predictions about the properties of spiking sequences, including nontrivial differences between the populations of excitatory and inhibitory neurons. In addition, my model suggested that sequences are robust enough to be present and detectable under the higher variance of membrane potentials that we might expect from in vivo baseline firing rates. Through the analysis of spike transfers and alternative model networks with truncated connectivity profiles, I showed that strong but infrequent connections might establish a substrate for reliable propagation. On the other hand, the weak but numerous connections might either enhance or disturb the process of propagation.

In conclusion, taking a bottom-up modeling approach that is strongly grounded on experimental data, I developed a model that explains the plausible mechanisms of reliable propagation of spiking sequences, with insights into how cortical networks may transfer streams of information and key experimentally-testable predictions.

The findings of my modeling indicate that infrequent but strong connections are essential to the production of sequences of cortical spikes from a single initial one. Repeatable sequences of action potentials have been reported in multiple mammalian cortical structures, from the hippocampus of rats to the primary somatosensory cortex of mice [126, 149, 150, 152–156]. Interestingly, although the model introduced here was contained by turtle data, these strong and sparse connections are a typical characteristic of the connectivity in the cortex of multiple species. For instance, heavy-tailed distributions of synaptic connection strengths have been discovered in the mouse barrel and visual cortices [101, 182], the rat visual cortex [68], the guinea pig and rat hippocampus [128, 183], and the human cortex [184]. In the hippocampus and the cortex, modeling and in vitro experiments have led researchers to the conclusion that those strong connections not only impact their postsynaptic partners but may alter the dynamics of the network [128, 185, 186]. Given the connection between strong synapses and sequences in the turtle cortex, it is possible that strong synapses lay at the heart of cortical sequences more broadly. If so, could sequences through strong connections be a preserved mechanism of signal propagation across multiple corticated species? We should take into account that the model that I have presented is founded on the data of the turtle cortex, while the cortices of other species may exhibit significant differences. These differences

may include variations in single-cell features, specialized connectivity, or operational firing rates. For instance, it may be possible that the level of reliability might be decreased or increased depending on some species-specific properties of connectivity that I did not investigate here. In order to arrive at a thorough comparative view, it will be necessary to further investigate, through models or experiments, the cortical responses of a variety of species to single spikes [17].

I found separate roles for strong and weak connections while analyzing the connectivity motifs and strengths that underlie spiking within the model in general (Fig. 4.11) as well as their impact on the number of reliably activating neurons (Fig. 4.12). While strong connections seem to provide reliability, weak ones provide flexibility. Experiments in the cortices of rats and mice have shown that trial-to-trial variability of EPSP sizes is reduced for large EPSPs when compared to small ones, supporting the idea that strong connections are more generally reliable [95, 101, 128]. Interestingly, the combination of two lognormal distributions is the model that best describes the distribution of synapse sizes between L2/3 pyramidal neurons, according to new data from electron imaging of the primary visual cortex of mice [96]. This potential categorization of excitatory synapses in two types might be the morphological counterpart of the predictions made by my model.

In my modeling approach, I tried to constrain as many model parameters as possible, which included connection strengths, neuronal densities, single-neuron properties, and connection probabilities. However, the connections in the model were still assigned randomly across network instantiations, resulting in a high degree of heterogeneity. The model produced sequences from single spikes without any fine-tuning of the connectivity, showing that structured connectivity might not be a requirement for propagation and providing potentially more general predictions about spiking sequences. However, certain aspects of cortical architecture, such as neuronal diversity, dendritic interactions, and specialized connectivity, might have an effect on the way in which activity is transmitted and directed.

Experimental evidence from morphology and transcriptomics suggests that the turtle brain has a variety of neuronal types, similar to what is seen in the cortices of mammals [41, 176] (see section 3.1.1). In terms of signal propagation, it is possible that the various kinds of neurons play different roles. For instance, parvalbumin- and somatostatin-expressing neurons in mouse V1 may play distinct but complementary functions in modulating reliability across multiple trials [187].

The current model depicts neurons as single compartments, yet, complex input interactions and non-linearities may take place at the level of dendritic arbors, which might act as different compartments with particular electrophysiological properties and micro-structures of organization. For example, synapses may cluster in short dendritic segments to successfully drive postsynaptic spiking [8, 188]; inhibition may locally balance these excitatory inputs [177]; and sequential activation within clusters may be particularly effective at triggering postsynaptic spiking [189].

Finally, as described in section 3.1.1, some nonrandom characteristics are likely to be present in the connectivity of the turtle dorsal cortex. These include things

like axonal projection biases [28], gradients of connectivity throughout the cortical surface, and differences in connections of the basal and apical dendritic trees [50]. Furthermore, connectivity structures derived from plastic rules guided by sensory input are likely to be present in the visual cortex.

The effects of all of these elements of cortical architecture on the way single spikes propagate in the cortex are still open questions in need of further research.

In conclusion, I developed a bottom-up model of the turtle cortical network that provides insights into the type of connectivity that might lead to spiking sequences from single spikes. This modeling approach starts with the most basic components as measured experimentally, yet it produces emergent network behavior that can be quantified to guide future experiments in the turtle cortex. Furthermore, the model shows that structured connectivity is not necessary for reliable propagation and establishes a link between long-tailed distributions of synaptic strengths and spiking sequences, both phenomena observed in many species and brain areas beyond the turtle dorsal cortex.

5 Flexible routing of single spikes in a random network model of the turtle cortex

Remark: Some of the methods, results, and figures in this chapter are part of an article entitled *Single spikes drive sequential propagation and routing of activity in a cortical network*, which was written together with my supervisor Prof. Dr. Julijana Gjorgjieva and feedback from my collaborators Dr. Mike Hemberger and Prof. Dr. Gilles Laurent. The article has been reviewed and published in the journal *eLife* [1]. All methods, results, and figures from that article that are part of this chapter were my contribution to the article unless specifically mentioned otherwise.

5.1 Overview

In the brain, information is encoded in the activation of neurons, which communicate with each other through synaptic connections [3]. However, not all information is equally relevant or important for the particular function that a neuronal circuit needs to implement. The ability to process and integrate information in a selective and flexible manner is a fundamental feature of brain function. This feature relies fundamentally on gating mechanisms that regulate the flow of neuronal activity through cortical circuits by amplifying or suppressing neural activity in a highly precise and context-dependent manner. Understanding the gating of neuronal activity is key to understanding the implementation of cortical computations, that is, the fundamental operations that transform input information into output information [13, 15]. How does the brain regulate the flow of information within itself?

Evidence from the turtle cortex has now shown that even a single spike from a neuron can reliably trigger a repeatable response [37]. In Chapter 4, I build upon these experimental results using a computational model and demonstrate that activity propagation from a single neuron in a recurrent network could be made reliable thanks to the presence of rare but powerful connections. However, a circuit that inevitably responds with the same pattern of activations to all spikes of a given single neuron is not capable of flexibility. Instead, to implement flexible computations, cortical propagation must be malleable through context-dependent mechanisms.

In this chapter, I study the variability of sequences of spikes in my experimentally-constrained model of the turtle cortex. I find that responding neurons can be

grouped together in atomic units that activate in full or not at all. These groups of neurons are connected to one another through sparse strong connections, which create the opportunity to gate the propagation of activity between them selectively. I find that because these strong connections are rare, manipulating just one neuron through a single external input is enough to halt or promote propagation. Finally, I explore how multiple spikes interact in the recurrent network and find that these interactions are reliable. These interactions enable a highly combinatorial repertoire of computations, understood as mappings of patterns of initial activations (triggers) to different patterns of network activations (followers).

5.1.1 Experimental background and modeling questions

The traditional view of cortical networks has been that excitatory neurons display a high variance in their membrane potential, resulting in highly irregular firing [161, 162, 190]. Additionally, single synaptic connections are considered to be weak, requiring the activation of many such connections in order to produce a postsynaptic spike [128]. Together, the unreliability of single units and the weakness of single connections have led to theoretical schemes of propagation that require population activity in the form of firing rate codes or synchronous firing [12, 15] (see section 3.3). In consequence, the models of gating of activity have typically required the concerted manipulation of a large number of neurons, either in their conductance state through a change of the ratio of excitation and inhibition [13] or in their precise spike timing [15]. Overall, this form of population gating is coarse by definition and does not explain how reliable gating and routing of single spikes might be implemented on a cortical network.

Experimental evidence from the turtle cortex suggests that reliable routing of spikes is possible. Having established that single neuron activations could lead to the reliable triggering of spiking sequences in an ex vivo slab of turtle cortex, my collaborators performed an additional set of experiments to investigate how multiple sequences might interact. While monitory network activity with a multi-electrode array (MEA), they simultaneously patched two pyramidal neurons and injected current pulses in one or both of them (Fig. 5.1A) [28,37]. As in previous experiments, they found neurons ("followers") in the surrounding network that activated reliably to the activation of these trigger neurons (see examples in Chapter 4, Fig. 4.1). Interestingly, they found that the activation of multiple of these triggers could result in the modulation of follower activation in diverse ways (Fig. 5.1B). Compared to the single-trigger setup, some followers responded to two-trigger activations by increasing their likelihood of activating, while other followers decreased it.

The diversity of follower responses in the turtle cortex to multi-trigger activations meant that, on a trial-by-trial basis, an initial spike from the same trigger neuron could propagate to some followers but not to others. Most interestingly, the end recipient of these spikes might be determined by very few additional inputs, in this case, the activation of a single additional neuron in the network.

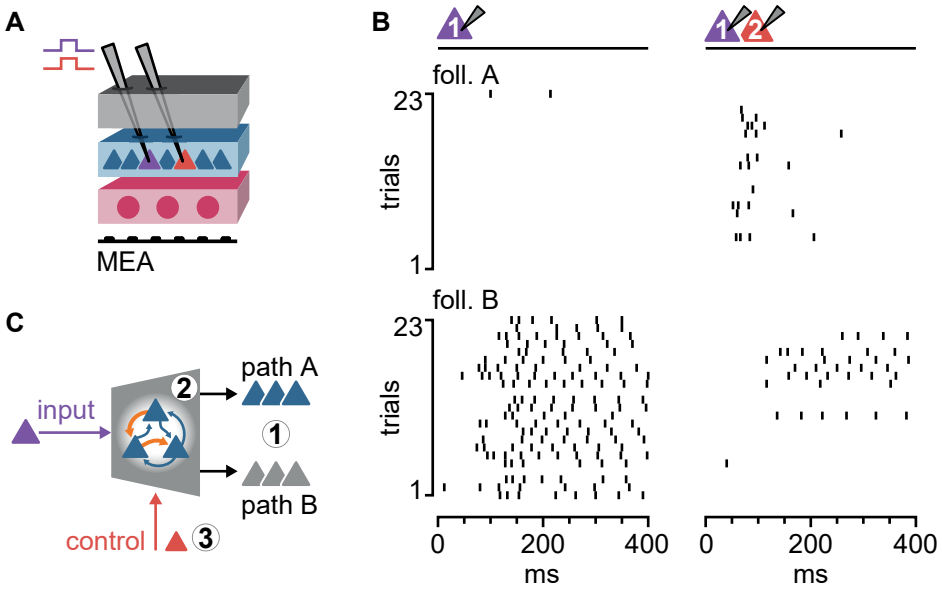

Figure 5.1: Routing of firing sequences in the turtle cortex. **A.** Schematic indicating the recording setup on an ex vivo slab of turtle dorsal cortex used in [37] (see section 3.1.1.1). A rectangular multi-electrode array (MEA) was placed under the ventricular surface, while two pyramidal neurons were stimulated with square pulses under whole-cell patch-clamp mode. **B.** (adapted from [37]) Examples of follower modulation by the activation of an additional trigger neuron in the ex vivo turtle cortex. Each tick corresponds to an action potential of follower A (top) or follower B (bottom) over 46 trials. In 23 of these trials (left), only trigger neuron 1 was forced to spike. In the remaining 23 trials (right), both trigger neurons 1 and 2 were forced to spike. Note both follower neurons are strongly and differently modulated by the activation of trigger neuron 2: follower A reliably increases its firing rate while follower B decreases it. **C.** Schematic of the main questions addressed computationally in this chapter (see text): (1) can different groups of followers of the same trigger neuron be activated independently, (2) what are the network mechanisms that determine the path of sequential activations that a single initial spike takes, and (3) what signals are available to reliably control this path of propagation.

In this chapter, I study what network mechanisms might route activity on top of the sparse network of strong connections in the turtle cortex (Fig. 5.1C). I first focus on the cases where my model showed partial failures of propagation to establish that gating is possible in this experimentally-constrained model at the level of individual spikes. I find that the sparsity of strong connections creates "gate" neurons that act as a mechanism to determine the path of activation. I then study two different types of control signals to reliably gate propagation: the conductance state of single neurons, manipulated by external temporally-constrained input spikes, and the interaction of multiple spikes within the recurrent circuit.

5.2 Sequences are composed of sub-sequences that correspond to sub-networks of strong connections

In order to get a better understanding of the ways in which sequential neuronal activations might be modulated, I first focused on the circumstances under which sequences fail to spread and the mechanisms by which this occurs. I found that, as a consequence of activity spreading via parallel sub-networks of strong connections that are sparsely connected, followers might activate or fail to activate in groups.

First, I returned to the original set of 6,000 simulations that I described in Chapter 4 (Fig. 4.4B) and selected a representative simulation in a firing rate regime compatible with the mean firing rates of both the ex vivo and in vivo turtle cortex. The simulation was picked so that its follower count was closest to the average count of all simulations in this firing rate range. This simulation had an average fire rate of 0.034 spk/s and produced 25 followers. Repeated trials with the same trigger neuron (Fig. 5.2A columns) can result in the failure of followers to activate (Fig. 5.2A rows). As a consequence, the specific spiking composition of a sequence might change from one trial to the next. Remember that I defined followers through a statistical test (section 4.3), resulting in neurons that are active in a large proportion of, but not necessarily all, trials. In addition, follower activations in this model rely mostly on rare but very strong connections (Fig. 4.10 and Fig. 4.11). Consequently, it appeared reasonable that the activation or failure of any follower may alter the unfolding of a sequence.

In order to examine the possibility of follower-to-follower dependence of activation, I evaluated what followers were active in which trials and carried out k-modes clustering of the neurons based on their activation across trials (Fig. 5.2B left). I compiled a summary of the activity of each follower as a single vector consisting of one hundred binary elements, each of which represented the presence (defined as at least one spike) or absence of that follower over all of the 100 trials in the simulation. After that, I performed a k-modes clustering analysis on these vectors. K-modes is an unsupervised technique that may be used to allocate data samples to K different clusters. It is very similar to k-means clustering, but instead of using the average values across sample vectors to identify the centroid of each cluster, it uses the number of matching values between sample vectors. This allows it to work very

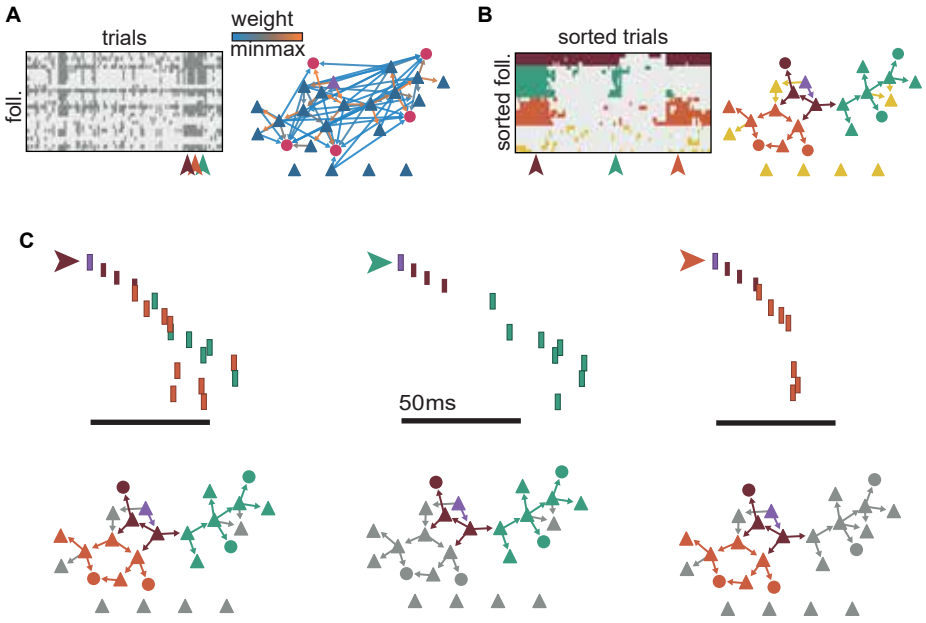

Figure 5.2: Clustering of followers by their activation. **A.** Left: Matrix of a representative simulation indicating what followers activate. Followers are in rows, and trials are in columns. Activation is indicated in black. Colored arrowheads correspond to trials shown in **C**. Right: Graph of connections between followers. The color indicates strength. Purple: trigger. Any followers detected but not connected to other followers are placed below. **B.** Left: Same as **A** after clustering and sorting trials and followers. Right: same as in A but removing all connections except the top 5%. Entries and followers are colored by follower clustering. Note how clusters share strong connections. **C** Sequence from trials corresponding to arrowheads in **A** and **B**. Left: spikes are colored by follower clustering. Graphs show in gray those followers not spiking in the trial. From [1].

well with binary data (follower active or follower absent). Manual examination indicated the followers were well grouped into groups of 5-10 neurons for this and other simulations.

K-modes analysis uncovered clusters that were made up of followers that had the tendency to activate together in the same trials (Fig. 5.2C left). It is important to note that activity, and not connectivity, was used in both the detection and clustering of followers.

After translating the follower clusters onto the graph of synaptic connections, I found that the connections within each cluster were often strong. They corresponded to the tail of the distribution of connection strengths (Fig. 5.2B right and Fig. 5.2C right). In light of this, I refer to each of these clusters of followers, defined by their activity, as a "sub-network", and I refer to the parts of the spiking sequence that correspond to their activation as a "sub-sequence". Identical breakdowns into sub-sequences were found in all of the other simulations, despite having varying numbers of followers that covered two orders of magnitude (Fig. A.1).

Sub-network activations

Despite the fact that the connections between neurons were chosen at random, it was still possible to identify and locate these separate sub-networks inside the full model. Remember that I define a "strong" connection as one in the top 0.3 percent of the distribution. I found strong connections more often inside sub-networks than between pairs of sub-networks. This segregation provides a mechanistic basis for the relative independence of activation of sub-networks (Fig. 5.3A). Notably, the very few connections between two sub-networks that are comparably strong do not necessarily ensure the reliable transmission of activity between them. This is because of the unexpected network effects that emerge from recurrent interactions (Fig. 5.3B).

In order to quantify the degree of interdependence of activation between sub-networks, I chose the two sub-networks with the biggest number of followers in each simulation (**a** and **b**). I then categorized each trial as one of four possible scenarios with respect to these sub-networks: **a** and **b** activate together, **a** alone, **b** alone, or full failure to propagate (Fig. 5.3C schematic). A sub-network is considered to be active according if at least forty percent of its followers emit at least one action potential throughout the trial. This classification allowed me to evaluate the probability of each scenario, which I quantified as the entropy. Given that spontaneous firing rates affect the reliability of followers, I evaluated how this value affects the entropy of these sub-networks. In Fig. 5.3C, I show the entropy of the two biggest sub-networks as a function of the mean firing rate of the corresponding simulation. I found that, at low firing rates, the entropy is around one bit, which indicates that two scenarios are equally probable (partial or full failure). The maximum expected entropy is two bits, which corresponds to the situation where the probability of each of the four events is equal. When there is a greater amount of spontaneous activity, the entropy increases toward this maximum. This tendency towards maximum entropy is compatible with recurrent connectivity and spontaneous activity being modeled as random. This prevents biases of activation of one sub-network over the other and moves the trend closer to maximal entropy. In a nutshell, in spite of the extensive recurrent connectivity, the state of the network has the ability to activate or deactivate some parts of a sequence while leaving other parts unaffected.

In conclusion, I found that sequences could be broken down based on the coactivation of followers, which showed sub-networks of followers that were strongly connected to one another. Surprisingly, strong connections are not a guarantee of propagation, and although some pairs of sub-networks shared strong connections, these were sparse and could frequently fail to propagate. As a consequence, subsequences could appear independently of one another. I hypothesized that this apparent failure really provides flexibility to the system, making it feasible for the sub-networks to function as separate channels of propagation.

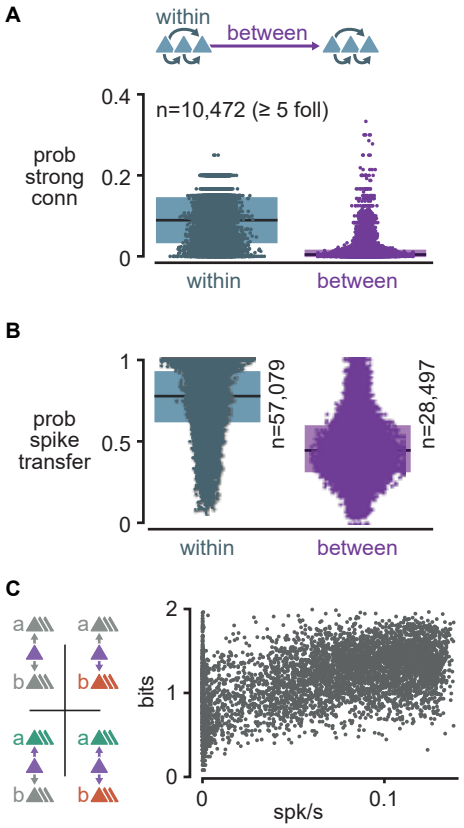

Figure 5.3: Connectivity behind the activation of sub-networks. **A.** Relationship between follower clusters and strong connections. Probability of a strong connection (top 0.3%) between excitatory followers of different (between) or the same (within) clusters. Connections pooled across all 6,000 simulations resulting in 10,472 clusters. **B.** Probability of a spike transfer: postsynaptic spike given a presynaptic spike in the same trial. Pooled for excitatory followers across all 6,000 simulations. Boxes indicate median and [25th, 75th] percentiles. **C.** Relationship between mean firing rate and entropy of cluster activation. Left: schematic of the four possible patterns of activations when considering only two sub-networks. Right: The plot displays the entropy across trials for each simulation consisting of at least two sub-networks (n=5,371), plotted as a function of the average firing rate of the network. From [1].

5.3 Sparse external input can halt or facilitate the propagation of activity

How does the segregation into sub-networks affect sequence propagation? I propose that the sparse and strong connections between pairs of sub-networks could be of vital significance for the process of sequence propagation since they allow activity to traverse from one sub-network to another. Furthermore, the few recipients of those connections in a sub-network are entry points into the sub-network, and their excitability and conductance state may be critical to ensure the activation of the rest of the sub-network. As a result, I refer to those neurons that serve as access points as the "gates" of their respective sub-networks. Using these gates, I discovered that I could regulate the activation of sub-sequences and, as a result, the course that the sequence took as it propagated through the network by sending a single external spike to each of these gates.

I started by locating the neurons inside each sub-network that had the smallest median spike delay from their respective triggers. This allowed me to choose which neurons would serve as the gates for each and every simulation. I changed the proportion of excitatory and inhibitory conductances that gate neurons got so that I could examine how their excitability would influence the propagation of single spikes across the remainder of the sub-network. I re-run a random selection of 2,000 of the initial simulations, but this time I added a spike coming from an ideal external source to a gate that was chosen at random (Fig. 5.4A). This ideal external spike arrived at the same time as the trigger neuron was forced to generate a spike. The ideal external spike had the greatest value from the excitatory or inhibitory conductance distributions (Fig. 4.3A). The intensity of the spontaneous activity was not modified. Subsequently, I calculated the variation in activation resulting from the modified conductance state of the gate in comparison to the initial simulation ($\Delta a/a_0$, %) (Fig. 5.4B). The degree of activation of each sub-network (a) was established as the proportion of trials (out of 100) in which the sub-network is considered active.

I discovered that external inputs to gate neurons are extremely effective in altering the propagation through their sub-network. Applying a single external spike with an excitatory conductance was able to double the likelihood that the sub-network activated in 55% of the simulations (Fig. 5.4BC). On the other hand, inhibition was even more effective. A single external spike with a maximal inhibitory conductance could reduce the likelihood that the sub-network activated by 50% in 74% of the simulations, and it was able to completely and reliably halt the sub-network activation in 26% of the simulations.

Timing and strength

Although maximal external inputs at the time of the trigger spike were indeed effective in controlling sub-network activation, it is possible that the relative strength and timing might affect this effectiveness. To examine the effect of these parameters,

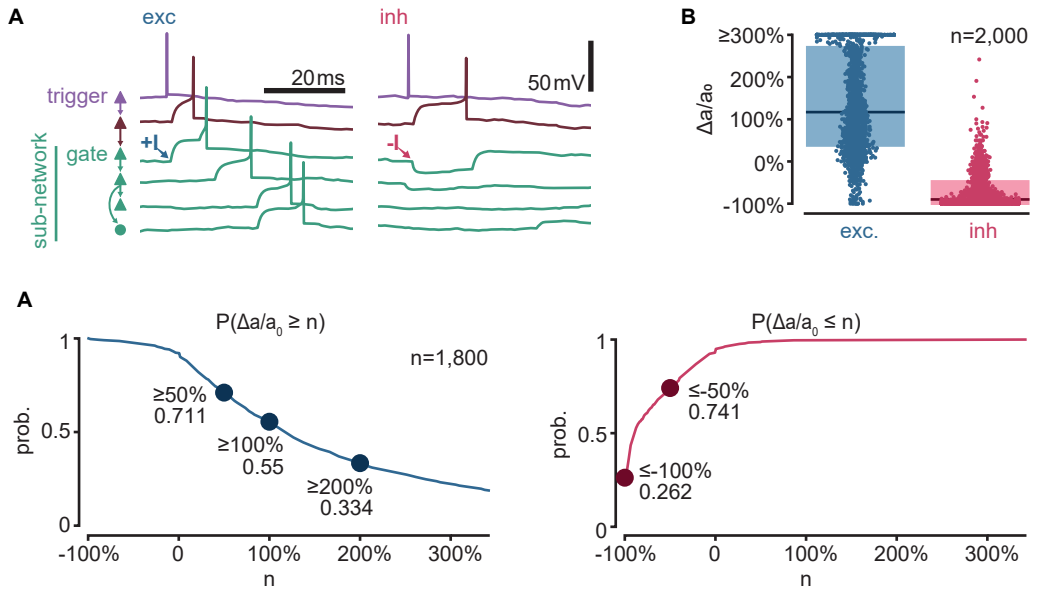

Figure 5.4: Gating sub-network activations. **A.** The voltage profiles of six neurons arranged in a sequence (with the same network and color scheme as in Fig. 5.2B) are presented for two simulation trials. In these trials, the first neuron in the green sub-network received either a solitary excitatory input (+I) or an inhibitory input (-I) from an external source, as indicated by the arrows. **B.** The percentage increase in sub-network activation above the baseline level ($\Delta a/a_0$, fold change, %) is presented for sub-networks chosen randomly from the pool of 6,000 original simulations. In each trial, the model gate neuron of the sub-network received an extra input from an external source at the onset of the simulation. **C.** Left: the likelihood of an excitatory input onto a gate neuron resulting in a minimum n-fold increase in the probability of activation within its corresponding sub-network. Right: the same is shown for inhibitory inputs. From [1].

I ran a series of simulations where I varied them systematically. Because of the computational cost of these simulations, I focused on a few network instantiations. Here I present the results for the network shown in Fig. 5.2B in detail and the averages for 16 additional networks. The details of all tested networks can be seen in Fig. A.2.

I experimented with changing both the sign (hyperpolarizing or depolarizing) and amplitude (strength) of the conductance of the external input that was sent to the gate. I also swept its relative time with respect to trigger activation. Only conductances that correspond to single connections were investigated in accordance with the estimations of connection strengths in the turtle cortex (Fig. 4.3). I found that the activation of each separate sub-network was subject to three distinct phenomena ($\Delta a/a_0$, Fig. 5.5AB).

First, an inhibitory input of any intensity within a temporal window of about 100 milliseconds before the trigger may acutely limit sub-network activity, often completely preventing propagation. This can occur regardless of the strength of the inhibitory input.

Second, an excitatory input with a strength that is somewhere in the tail of the distribution but not necessarily at the tip can facilitate the sequence propagation. The expected spike timing of the gate neuron serves as the core point of reference for this temporal facilitation window.

Lastly, an excitatory input with a very strong amplitude might hinder propagation if it takes place a significant amount of time before the activation of the trigger neuron. This surprising behavior of the sub-network is likely the result of the single-cell adaptive properties as reported experimentally [37]. In the case of the adaptation, the external input, by possibly triggering the gate and some of the neurons in its sub-network too early, activates adaptive currents, resulting in an increased threshold for spiking, which hinders the probability of propagation at the time of the trial.

I repeated these analyses on simulated networks with heterogeneous connectivity, spanning orders of magnitude in their number of followers and baseline firing rates. I observed the same three phenomena (Fig. A.2 and Fig. A.3). When evaluated relative to the trigger spike, the mean of these activation maps across all of the tested sub-networks suggested that halting is most effective in a time frame of 70 ms, which is shorter than that for facilitation of 100 ms (Fig. 5.5C). However, when the expected spike time of the gate is taken into account, the temporal windows become much larger (Fig. 5.5D). Interestingly, despite the fact that the temporal window for halting seems to have a hard temporal limit that is specified by the expected time of spike of the gate neuron, the temporal window for facilitation expands as the input amplitude increases.

Lastly, I investigated the effect that modifying the gate of one sub-network had on the probability of another sub-network activating (Fig. 5.5EF). It is interesting to note that the effect between the different sub-networks is not always symmetrical. In the case of this example network, the activation of the orange sub-network is

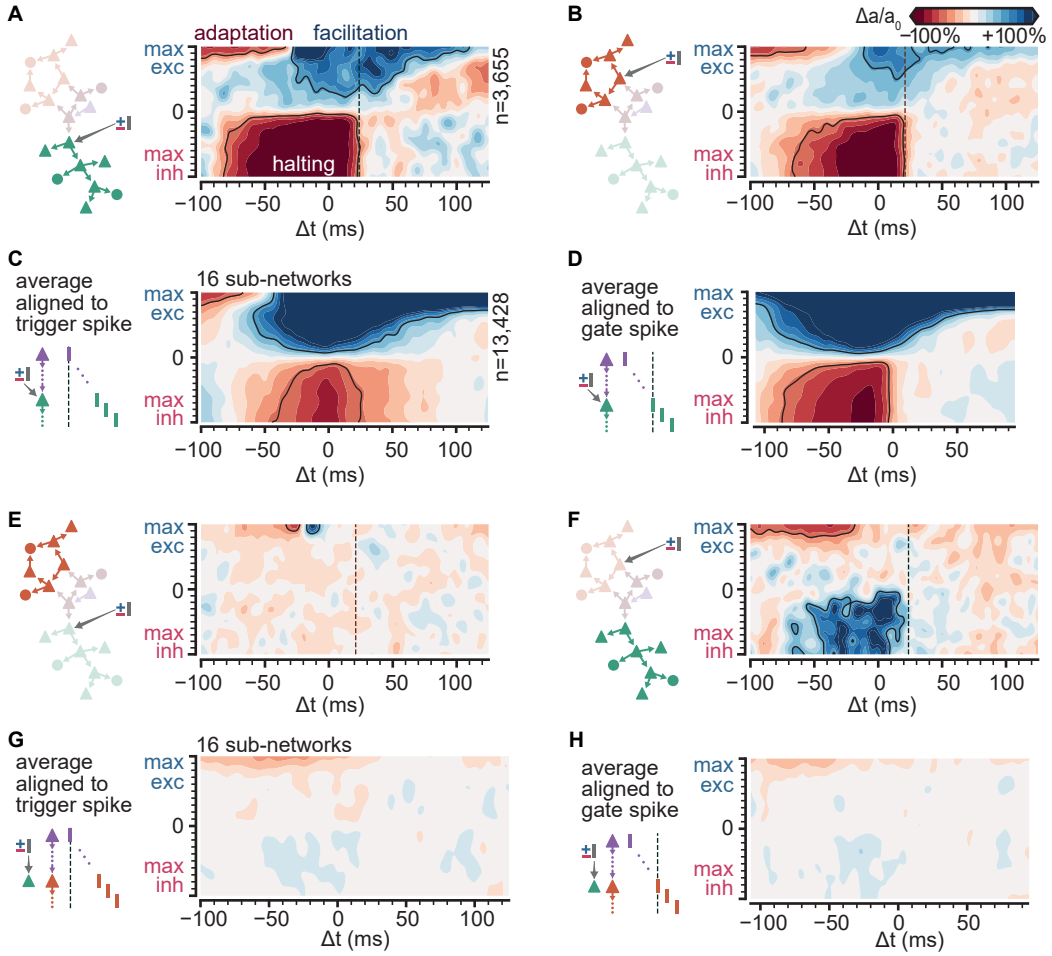

Figure 5.5: Effect of timing and strength on sub-network activations. Color bar as in B for all panels. **A.** Diagram illustrating the experimental procedure and the resulting fold change in activation ($\Delta a/a_0$, fold change) for a sub-network when manipulating its gate (same network as Fig. 5.2B). The arrow indicates the stimulated gate. Δt denotes the delay between the trigger spike and external input. Solid outlines highlight combinations of strength and timing that result in facilitation (+50% $\Delta a/a_0$), halting (-50% $\Delta a/a_0$), or adaptation (-50% $\Delta a/a_0$). The median gate spike time is indicated by a dashed line. **B.** Same as **A** for a different sub-network. **C.** The mean map of the change in activation ($\Delta a/a_0$, fold change, %) was determined for 16 sub-networks derived from 8 distinct networks encompassing between 5 and 462 followers and a baseline firing rate between 0.01 and 0.1 spk/s. The maps were aligned relative to the trigger spike before averaging ($\Delta t = 0$ in **A**). The solid outlines denote the ±50 % $\Delta a/a_0$ range. **D.** Same as **C**, but maps were aligned to each gate median spike time before averaging (dashed line in **A**). **E.** Same as **A** for the orange sub-network when manipulating the green gate. **F.** Same as **B** for the green sub-network when manipulating the orange gate. **G.** Average map of the change in activation ($\Delta a/a_0$, %) from 16 sub-networks when manipulating a gate from a different sub-network in the same network. Same networks and alignment as **C**. **H.** Same as G but maps were aligned to each gate median spike time before averaging (as in **D**). From [1].

almost completely independent of the facilitation or halting of the green sub-network (Fig. 5.5E). Surprisingly, the gating of the orange sub-network makes it easier for the green sub-network to become activated (Fig. 5.5F bottom right). This enhanced activation through halting strongly suggests that the orange sub-network generally has an inhibiting effect on the green one under normal conditions. Furthermore, robust excitatory inputs that have the potential to trigger the orange sub-network too early also have the effect of inhibiting the activation of the green sub-network (Fig. 5.5F top left). Across all pairs of sub-networks that I explored, I observed a variety of intricate interactions between them (Fig. A.2 and Fig. A.3). Nevertheless, the interactions between the sub-networks do not seem to follow any particular pattern (Fig. 5.5GH), which suggests that they are dependent on the specific implementation of the random connections in each network.

In conclusion, I found gate neurons in various sub-networks. These are neurons whose activity is essential for the reliable activation of the sub-network they are a part of. External single spikes have the ability to affect the conductance state of the membrane of gate neurons, and as a result, they may either halt or promote the activation of particular sub-networks. The impact of these spikes is determined by the timing of those spikes in relation to the trigger spike of the sequence. Furthermore, the activation of one sub-network may have an effect on other sub-networks via lateral connections, which might result in complicated interactions that are not necessarily reciprocal in nature and depend on the specific network instantiation.

5.4 Sequences from multiple triggers reliably activate combinations of followers

My exploration with gate neurons suggested that recurrent interactions across parallel running sub-sequences might have surprising side effects in the sequential propagation of spikes. Next, I set to investigate if and how different parallel running sequences might interact with one another. I ran a fresh batch of simulations with several trigger neurons being activated at once. The results of analyzing the responding neurons showed that sequences produced by coactivated trigger neurons interact frequently and predictably, resulting in a combinatorial definition of the neurons that are considered to be followers.

For the new set of simulations that I describe below, I employed the same instantiated network (same as Fig. 5.2B), but I verified the results that I describe here are general in multiple other network instantiations summarized in Fig. A.4.

I started by simulating 100 trials in each simulation where a single spike was forced onto a single one of 2,000 neurons that were chosen at random to serve as triggers. Next, I generated 5,000 random pairs from those trigger neurons and found that the likelihood of two excitatory neurons sharing followers is extremely small (Fig. 5.6AB), with just 3.12% of all tested pairings of triggers sharing one or more followers, and no pair ever showing more than 10 followers.

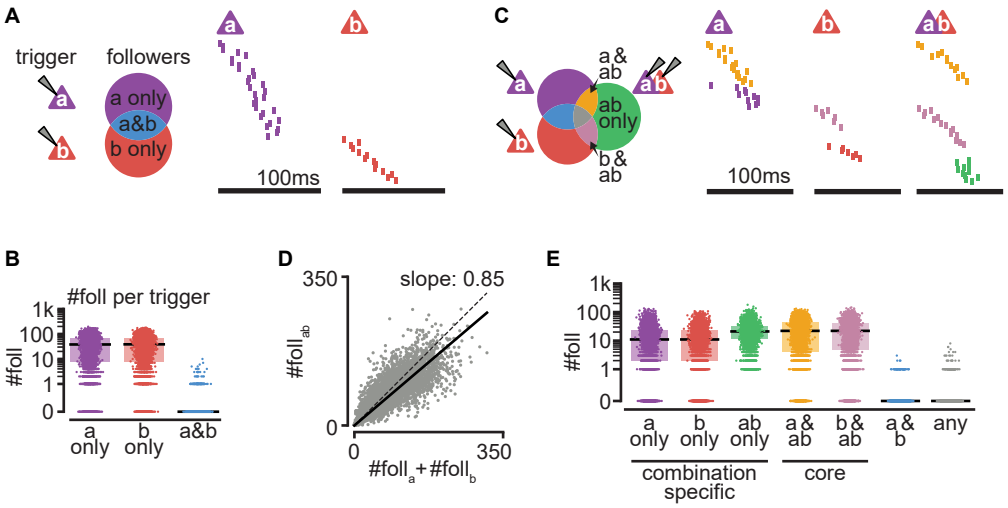

Figure 5.6: Coactivation of multiple triggers. **A.** Left: diagram illustrating the follower categories as a function of the trigger neuron: purple is followers to **a**, red are followers to **b**, and blue are followers to both **a** and **b**. Right: examples of sequences generated by the activation of **a** or **b**. The followers are arranged in order of their trigger class (left) and median spike time. The triggers do not share any followers, as indicated by the absence of blue spikes. **B.** Count of followers for each trigger class in **A** across all 5,000 randomly chosen pairs of simulations. Boxes indicate median and [25th, 75th] percentiles. **C.** Left: diagram illustrating the follower categories as a function of the trigger: a trigger neuron activating in isolation (**a** or **b**) or the coactivation of two triggers together (**ab**). The colors blue, purple, and red correspond to those in **A**. On the right-hand side, the same trials as in **A** and a trial under coactivation are presented. The vertical sorting of followers is identical in all trials. **D.** The number of followers that are present when two triggers are coactivated together in relation to the total sum of followers for each of the triggers when activated separately. The solid line shows a linear fit where the intercept is set to zero. The dashed line indicates the diagonal. **E.** Count of detected followers for each trigger category in **C** pulled from all tested trigger pairs (n=5,000). From [1].

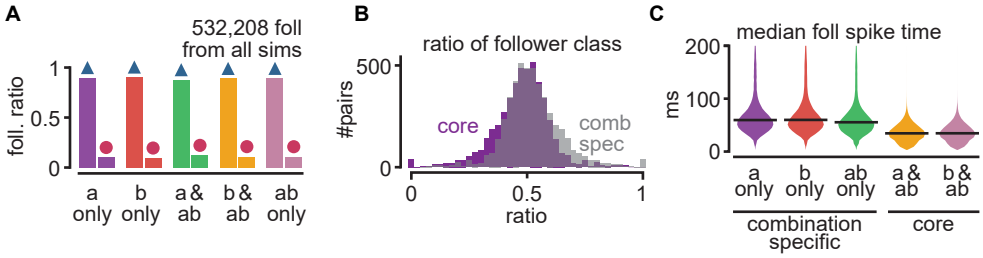

Figure 5.7: Properties of combination-specific and core followers. **A.** Proportion exc/inh followers in each follower category (see Fig. 5.6C). Pooled from 5,000 random pairs of triggers. **B.** The ratio of core and combination-specific followers for all 5,000 random pairs of triggers. **C.** The median time of spiking, relative to trigger, for followers in each category after pooling from all simulations (5,000 simulations; 532,208 followers). Lines indicate the median. From [1].

Trigger coactivations

Next, I generated a new simulation for each randomly selected pair of trigger neurons, where they were both forced to emit a spike at the same time over the usual 100 trials. When I compared the followers that are produced by these double-trigger coactivations to those produced by single-trigger activations, I arrived at distinct groups of followers (Fig. 5.6C): those that are activated in response to one of the single triggers (**a** and **b**) and those that are activated in response to the coactivation of both triggers (**a&b**).

The number of followers that are produced by the coactivating **a** or **b** is less than the total number of followers obtained by either **a** or **b** on their own (Fig. 5.6D and Fig. A.4A). I had already found that **a** and **b** seldom share followers. Consequently, the resultant sublinear activity implies frequent lateral inhibition, as I had already described in the previous section for sub-networks (Fig. 5.5F).

Consistent with this lateral inhibition and sub-linear summation, neurons that typically follow a particular single trigger can become unresponsive or unreliable if that trigger neuron is coactivated with a new trigger neuron (**a only** and **b only**, Fig. 5.6E, Fig. A.4B). More surprisingly, some neurons may not be present following a spike of either trigger neuron on its own, but they may become reliable followers under the activation of the two trigger neurons (**ab only**). In both cases, since the activity of these followers is dependent on whether each of the trigger neurons is quiet or active, I refer to them as "combination-specific followers". In summary, there are three types of combination-specific followers: **a only**, **b only**, and **ab only**. In contrast to combination-specific followers, there are neurons that act as followers of a trigger regardless of whether or not the second trigger neuron is activated. I refer to these as "core followers": **a&ab** and **b&ab**.

Excitatory and inhibitory followers were equally likely to be core or combination-specific (Fig. 5.7A). Under trigger coactivation, combination-specific followers account for nearly half of all followers combined (Fig. 5.7B). In a typical sequence, core

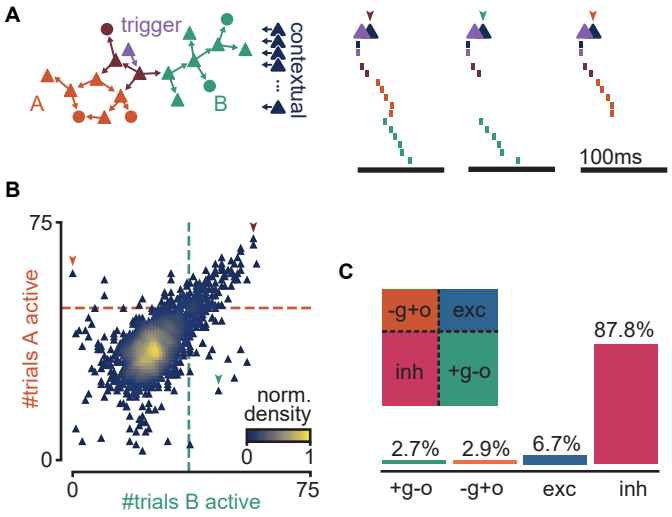

Figure 5.8: Contextual activation of sub-networks. **A.** Left: the trigger neuron in purple, its followers, the strong connections between them (indicated by arrows), and the other neurons in the network (colored in dark blue). The follower colors correspond to those used in Fig. 5.2B. Right: sequences triggered by the simultaneous coactivation of the trigger neuron (purple) and one additional contextual neuron that facilitates propagation. The contextual neuron facilitates propagation in both sub-networks on the left, only in the green sub-network in the middle, or only in the orange sub-network on the right. **B.** The impact of coactivating the trigger neuron with 2,000 randomly chosen contextual neurons on the activation of the green and orange sub-networks is presented in the left panel of Figure A. The baseline activation of each sub-network is indicated by the dashed lines. Neurons are color-coded according to their local density, and colored arrows correspond to the sequences shown in **A**. **C.** The ratio of random triggers for each one of the quadrants in **B** (inset schematic). From [1].

followers will activate early on, while combination-specific followers will become active with a longer delay from the trigger event (Fig. 5.7C). This relationship between time and reliability is in agreement with the observation, both in experiments and in the model, that, under single-trigger activation, follower identity was more reliable in the first few ranks of the sequence [37], (Fig. 4.7A). Indeed, under single-trigger activations, different neurons in the network may also spontaneously spike from trial to trial, resulting in the activation, in each case, of different combination-specific followers but always the same core followers. Conversely, the fact that combination-specific followers tend to activate later in the sequence indicates that the regulation of sub-sequences through lateral cortical interactions is most effective on followers that are further from the trigger neuron in terms of the number of synaptic jumps.

Contextual control of sub-sequences

In section 5.3, I described how sub-sequences could be reliably modulated through single external inputs. The new observation that lateral intracortical

interactions could also be reliable prompted me to study their effect at the level of sub-sequences. With this goal, I returned to the trigger in the network in Fig. 5.2B as an example of a sequence with a known sub-sequence decomposition. I generated a new set of 2,000 simulations where this known trigger was coactivated with a randomly picked excitatory neuron over 100 trials. I term these additional triggers "contextual" since they impose a fixed context from trial to trial that is different from the ever-changing random spontaneous activity. Next, I leveraged the already classified green and orange sub-networks to examine the impact that coactivation had on the sequence (Fig. 5.8A).

When I evaluated the number of these contextualized trials in which each of the sub-sequences became active, I discovered that the majority of contextual neurons either stimulate or inhibit both sub-networks in an equal manner (diagonal trend, Fig. 5.8B). The most common effect, by far, was lateral inhibition, with about 87.8 percent of the contextual neurons inducing a reduction of the likelihood that both sub-networks become active, as can be seen in the bottom left quadrant of figure Fig. 5.8B (see also Fig. 5.8C). This pervasive inhibition was already observed in the sub-linear summation of follower counts under the coactivation of multiple triggers (Fig. 5.6D).

On the other hand, lateral interactions could also result, more rarely, in excitation. About 6.7 percent of the interactions resulted in equal excitation of both branches, as can be observed in the upper right quadrant of Fig. 5.8B (see also Fig. 5.8C).

The tendency towards symmetrical effects on both branches is most likely the result of the random connectivity in the model together with random spontaneous activity, which, under sufficient sampling, will prevent a bias towards the activation of one branch over the other. This balancing effect also reflects the fact that, although the activation of one sub-network may strongly affect another, the average effect across multiple instantiations of the network results in a net zero effect (Fig. 5.5G right, Fig. 5.5H). I constrained the main parameters of my model by biological measurements, but I lacked detailed information on the fine structure of its connectivity. It is possible that if certain connectivity patterns are more likely than others in the turtle cortex, they may introduce asymmetries in these results.

Interestingly, the randomness of network instantiation creates heterogeneities in the connectivity that result in some variance away from the diagonal. Those off-the-diagonal contextual neurons thus had asymmetric effects on the two studied sub-networks. They could inhibit the activation of one of the two sub-networks while, at the same time, promoting propagation along the other one. These contextual neurons correspond to the bottom right and top left quadrants of Fig. 5.8B and account for a smaller percentage of the sampled contextual neurons (about 3 percent each, Fig. 5.8C).

The fact that just a small percent of contextual neurons are arbiters of propagation along one and only one of the two sub-networks might hint at a highly specialized kind of routing. This fine control of propagation is in line with my earlier observation that excitatory neurons, on the whole, are particularly tolerant to ongoing spontaneous

network activity thanks to their huge number in the circuit and relatively lower connection probabilities when compared to inhibitory neurons (Fig. 4.7). Remember that the model network that I present here aims to match estimates of neural density in the biological system, corresponding to a 2x2 mm sheet of turtle cortex, and it contains 100k model neurons (Fig. 4.2A). Based on the size of this network, and the percentages of contextual neurons that can switch propagation along the two sub-sequences, I estimate that, in any given trial, the path of activation through the network might be determined by the activation or silence of approximately 5,000 contextual neurons. Thus, although proportionally small, there exists a large reservoir of combinations that may determine the routing of the single initial trigger spike.

In summary, my exploration of multi-trigger interactions suggests that although it is uncommon for sequences that were begun by distinct triggers to overlap, interactions between multiple parallel running sequences are very frequent. Most importantly, lateral interactions between multiple sequences are reliable across trials, resulting in the reliable activation or silencing of different groups of followers. These interactions take the form of halting or promoting propagation along specific sub-sequences. Lateral inhibition is the most common form of interaction, hinting at a balancing effect in the network. Nonetheless, lateral excitation and asymmetric effects on multiple sub-networks are also present. As a consequence of this, the state of a relatively small number of contextual neurons in the background network can selectively gate propagation and fix the pathway that the initial spike will take.

Taken together, my findings point to a series of mechanisms of recurrent interactions that can lead to the reliable and highly selective routing of spikes within the sparse network of strong connections. This reliable routing may form a basis upon which cortical computation may be implemented.

5.5 Discussion

In Chapter 4, I introduced a model of the turtle dorsal cortex that was constrained by experimental measurements, recapitulated several properties of spiking cortical sequences, and predicted others. In this chapter, I have studied in detail how repeatable sequences in this model can be reliably modified.

First, by exploring how followers might fail to activate in some trials, I showed that sequences might be broken down into several sub-sequences. These sub-sequences reflect a sub-network of followers sharing strong connections. Under the random connectivity rules of the model, sub-networks form on their own without any manual adjustment of the network architecture. Next, I showed that the state of a few gate neurons could determine the activation of individual sub-sequences. These gate neurons, and by extension, sequence propagation, could be controlled through external inputs and ongoing recurrent activity. When I studied recurrent interactions between the sequences of multiple trigger neurons, I found that trigger neurons only rarely share followers with one another but that their sequences often interact.

Importantly, these interactions are reliable, leading to the presence of combination-specific followers. Since followers are selective in the specific responses of multiple triggers, the system has access to a large combinatorial repertoire of activations while maintaining its reliability and specificity. Lastly, I discovered that the precise route that signal propagation takes through the network is shaped by the activation of a very small proportion of neurons in the network. In summary, I described the mechanisms through which reliable paths of propagation can be established and controlled within a cortical network.

Routing on a network of strong connections

Based on the findings described in this chapter, it seems that the interruption of propagating spiking sequences may be a defining characteristic of the flexibility of the system. This flexibility appears to the external eye as failures to propagate activity, but they might, in fact, be the result of routing the activity. For example, variations in the context of the activity from one trial to the next may often channel propagation away from specific neurons, producing what an external observer would define as followers with a low modulation of firing rate. The model that I presented here provides a plausible circuit mechanism for the generation of spiking sequences with flexible trajectories. Indeed, flexible sequences have been observed in the rat hippocampus and have been related to the flexible planning of spatial routes during memory and goal-directed tasks, which result in multiple distinct replay sequences of spikes from place cells in the same environment [191, 192].

Based on my observations on propagation, I hypothesize that the dynamic routing of cortical activity may be implemented on two key factors: the presence, for each pathway, of a few sites that may act as points of failure; and the circuit mechanisms that can activate those gates reliably. In the model that I studied, the strength of some connections in the network is what made reliable propagation possible, while the sparsity of those same connections creates groups of neurons (sub-networks) that can receive activity through only a limited number of entry points. These entry points, or "gates", are perfect points of failure where propagation may be halted or promoted. Gates may be opened or closed selectively depending on the timing and strength of external input. In addition, ongoing parallel sequences may also have an effect on the route that information takes through the recurrent network by generating lateral excitation (cooperation) or inhibition (competition) across the various sub-networks in the network.

The key characteristic for the form of routing that I propose here might is that it can be reliably controlled from trial to trial. This reliability makes it a great candidate for implementing robust mappings from inputs to outputs, that is, implementing cortical computations. Indeed, prior theoretical research on propagation within recurrent networks characterized by irregular background activity showed that the interplay of gating and parallel pathways might be leveraged to construct logic statements, switches, and even memory units (Fig. 1.4), illustrating how complex computations can be built on the basis of reliable gating [12].

Interestingly, the most recent advancements in machine learning are the consequence of making use of gating in order to explicitly restrict and redirect the flow of information. For example, gating can be used to achieve context-dependent learning that avoids catastrophic forgetting [193], to implement various forms of working memory within recurrent artificial neural networks [194], and to produce mechanisms of attention that have proven to be extremely powerful [195].

Other frameworks for routing

The routing of a single spike propagating in the model as a reliable spiking sequence is connected to previous theoretical frameworks of mechanisms of cortical communication, namely: synfire chains, neuronal avalanches, stochastic resonance, and the temporary breaking of excitatory/inhibitory balance. See section 3.3 for detailed descriptions. Next, I briefly describe the core similarities and differences between these frameworks and a routing mechanism based on sparse strong connections.

A traditional model of reliable signal propagation is that of synfire chains (see section 3.3.1). Under this model, pools of several hundred neurons are connected in a chain with divergent-convergent connections between each pair of consecutive groups. Each group then activates synchronously and, in doing so, ensures postsynaptic spikes in the next group. Theoretical analysis shows that this mechanism can tolerate the partial failure of activation and of synchrony within any one group [14, 15, 138]. Moreover, research conducted on modeling studies has incorporated synfire chains in recurrent networks, providing valuable insights into the specific network structures required for effective propagation and gating of activity through this mechanism [15, 142, 143, 159].

A closely related model, polychronous chains, replaced synchrony with precisely timed patterns of activations in order to compensate for variance in the distribution of synaptic delays [7] (see section 3.3.2. In the case of polychronous groups, applying spike-timing-dependent plasticity results in the reinforcement of those time-locked patterns, which leads to very strong synapses. However, synfire and polychronous chains are a largely theoretical model, with only limited or indirect biological evidence for their existence [129, 135, 136, 140, 147, 173].

In the model that I described in this chapter, activity propagates over synaptic chains that are sparse ($\leq$ 2 connections, Fig. 4.11BC, Fig. 4.12D) yet powerful enough to cause a postsynaptic neuron to fire without the need for the synchronous firing of a large group of presynaptic neurons. This is in contrast to classical synfire chains, in which connectivity is highly structured. The unstructured sparsity of strong connections facilitates the selective halting of spike propagation by even modest inputs, such as a single external spike (Fig. 5.4A) or sequence-to-sequence interactions (Fig. 5.6C), and it leads late followers to be more prone to modulation (Fig. 4.7, Fig. 5.8B). Contrast this to synfire chains, where propagation is so robust that requires a highly distorted input pattern to halt it [15]. Interestingly, the unstructured random connectivity in my model manages to capture very accurately

that 80% of trigger neurons in the experiments were able to start sequences even when activated in isolation ([37] Fig. 4.7E), which suggests that, at least in turtle cortex, sparse strong connections are highly distributed creating "thin" synaptic chains. Synfire chains, on the contrary, are defined to activate only under the presence of synchronous volleys. Taken together, my model proposes that successful propagation could entail a great deal fewer neurons and spikes than was previously expected by synfire models. This changes the predictions for the activity that should be observable in the cortex and has substantial ramifications in terms of the kind of analysis that we leverage when studying cortical synaptic propagation.

Both in the model and in the turtle experiments, a single spike from a random pyramidal neuron could trigger several ones in the surrounding network, which can be seen as a signature of neuronal avalanches [196]. Note that this multiplicative effect never leads to uncontrolled excitation. Avalanches are supposed to appear when a network is in a "critical" state, that is, a dynamical regime between explosive amplification and decay. There are theoretical studies that suggest that this regime can maximize information transmission in recurrent networks [197]. Previous studies of the turtle cortex that used visual stimulation with an eye-attached ex vivo preparation observed scale-free statistics of population activity, a known hallmark of criticality [62]. These studies did not address the reliability of these avalanches in terms of which neurons activated and in which order across trials. Indeed cortical criticality does not imply reliability, but it also does not rule it out. Earlier studies of LFP in slice cultures showed that LFP avalanches could be repeatable and temporally precise, which may suggest the presence of underlying sequences of spikes [196].

The conclusions of my modeling that suggest sequences of spikes are the manifestation of routing via strong connections may thus be considered an alternative interpretation to avalanches that is compatible with the theory of cortical criticality. Note that I propose the halting of a sequence of spikes to be a form of contextual control of the path of propagation, i.e., routing. My simulations with a fixed external input (Fig. 5.4) and contextual activity (Fig. 5.8) showed that single sub-sequences could be reliably controlled. This aspect differs from that of avalanches, where they are treated as a statistical property of a stochastic branching process where the halting of a sequence is purely probabilistic. My model suggests, instead, that once part of the context of the network is fixed, halting also becomes fixed. Note that the number of neurons that influence a given sub-sequence may be relatively small (in the order of 6% of the pyramidal population, see Fig. 5.8). My model thus predicts that, as more of the contextual network state becomes fixed from trial to trial, sequence statistics become less variable and move away from the traditional power-law distributions often predicted from a stochastic branching process. This prediction is important from an experimental point of view, where we might use pharmacological agents to control the network state or increase the number of trigger neurons that are coactivated (Fig. 5.8).

Another statistical framework that has been used to study neuronal signal propagation is stochastic resonance. This phenomenon describes the way in which certain noise properties may increase instead of decrease the response of a system to certain stimuli [198]. In terms of networks of neurons, modeling studies have suggested that long-tailed distributions of strengths could generate the chance for "aperiodic stochastic resonance" where weak input support optimal spike transition through strong connections [199, 200].

The stochastic resonance between heterogeneous synaptic strengths is consistent with the results of my simulations with and without weak connections. By truncating the network connectivity, I found that there exist different input ranges in which weak connections have the potential to either boost or reduce reliability in terms of the number of reliably responding neurons (Fig. 4.12E). However, my research goes beyond this global form of control and zooms into the finer mechanisms of selectivity of particular sub-networks. These are influenced by a small number of neurons within the recurrent circuit (Fig. 5.8B) and by sparse but temporally constrained external inputs (Fig. 5.5). In the same vein as in the comparison with criticality and synfire chains, the conclusions of my model point to cortical routing mechanisms that are fine-grained and reliable instead of simply stochastic and coarse.

Finally, the routing that I present here is related to the balance of inhibition and excitation often assumed for cortical circuits [161]. In spite of a single spike reliably triggering a large number of subsequent spikes in the network, there was never a case of explosive amplification in either the model or the turtle experiments. On the other hand, the frequent lateral inhibition onto sub-sequences (Fig. 5.8B) seems to cause a net non-linear summing of responses to multiple simultaneous excitatory spikes (Fig. 5.6D). Furthermore, changes in excitatory and inhibitory conductances in certain gate neurons within temporally-constrained windows resulted in the reliable gating of activity (Fig. 5.4). Taken together, these results support the hypothesis that sequences function under a regime in which excitatory inputs are frequently canceled by local inhibition.

The model presented here was not intentionally built such that it would exhibit a balance of inhibition and excitation. Yet, this kind of frequent lateral inhibition and routing via the manipulation of excitatory and inhibitory conductances is related to the phenomenon known as "precise balance," which is a type of excitatory/inhibitory balancing that occurs within small temporal windows [201]. Earlier theoretical work proposed that the gating of signals could be arbitrated through the local alteration of excitatory/inhibitory balance [13]. However, these studies relied on firing rate to convey information, accounting for temporal variance and large neuronal populations. Thus, my work extends these ideas of balance disruption to work with just a few neurons in shorter timescales.

In conclusion, by studying the properties of sequences of spikes in a model of the turtle dorsal cortex, I have arrived at a mechanism of cortical communication based on sparse strong connections that allow for reliable and flexible routing. The sparsity of the connections allows for the selective halting or promotion of signal propagation

through mechanisms of external and recurrent control. The form of routing that I propose is fine-grained, acting at the level of a few neurons and sparse paths of propagation, which distinguishes it from previous theoretical frameworks relying on statistical and population effects, such as synfire chains, neuronal avalanches, stochastic resonance, and the temporary breaking of excitatory/inhibitory balance. The conclusions from my computational research have implications for the activity that should be observable in the cortex, with substantial ramifications for the analysis of cortical synaptic propagation.

6 Inter-hemispheric competition during REM sleep in the lizard brain

Remark: Some of the methods, results, and figures in this chapter are part of an article entitled *Inter-hemispheric competition during REM sleep*, which was written together with my collaborators Dr. Lorenz Fenk and Prof. Dr. Gilles Laurent [2]. The article has been reviewed and published in the journal *Nature*.

Dr. Lorenz Fenk carried out all the experiments and collected all the data. I developed and applied analysis methods and produced visualizations and quantifications of the data. Dr. Lorenz Fenk, Prof. Dr. Gilles Laurent, and I discussed and interpreted the results together. All methods and figures from this article that are part of this chapter were my contribution to the article unless specifically mentioned otherwise.

6.1 Overview

During sleep, the cortical electroencephalograms of mammals can be broken down into two distinct stages: rapid-eye-movement (REM) sleep and non-REM (NREM) sleep. See section 3.2 for a detailed description. REM sleep is distinguished by high-frequency electrical signals and fast eye movements [202–204]. NREM is also known as "slow-wave sleep" due to the presence of low-frequency activity. The possible functions of REM sleep are still largely unknown, although there is some evidence that suggests a possible link with emotional memories [122–124]. On the other hand, there is evidence that non-rapid eye movement (NREM) sleep is involved in the reactivation and consolidation of specific types of memories in both rats and humans [205–208].

The existence of biphasic sleep also in birds, non-avian reptiles, and, more recently uncovered, fish, suggests the possibility of common roots of these brain states that could go back to the first vertebrates 500 million years ago [24, 30, 127]. If this common ancestry of biphasic sleep is correct, comparative studies in systems that reflect a variety of animal lineages may be able to assist us in better comprehending not just the evolution of sleep but also the roles it serves and the mechanistic foundations upon which it is based.

Recordings of local field potential (LFP) of *Pogona vitticeps*, a diurnal agamid lizard, revealed these two stages of regular electrophysiological sleep in these lizards [32]. The NREM periods last about a minute and contain multiple sharp-wave ripples

that occur irregularly, with an inter-event interval of about 0.5–2s. The REM periods contain faster awake-like activity that is accompanied by rapid-eye movements. Both REM and NREM last for about 80 seconds each, resulting in a full cycle of just under three minutes. These two periods alternate on a consistent basis during the course of the night [32]. These recordings were made in the dorsal-ventricular ridge (DVR), a sub-cortical region, in the in vivo lizard. NREM activity in DVR is predominantly quiet except for the presence of sharp-wave ripples, which originate in a subarea recently found to be the homolog to mammalian claustrum [30]. These claustral SWR then propagate through the adjacent anterior dorsal ventricular ridge (aDVR) (Fig. 6.1A).

6.1.1 Experimental observations and computational challenges

Having established claustrum as the origin of sharp-wave-ripples during NREM, my collaborators now turned to examine the claustral LFP signals during REM sleep. They collected new data from the in vivo *Pogona vitticeps*, and I developed computational methods to analyze it.

When examining claustral LFP activity during REM sleep, my collaborator, Dr. Lorenz Fenk, observed that it was populated by many single events that occurred at very high rates, consistent in their general shape but variable in amplitude (Fig. 6.1B bottom left). These LFP events presented a sharp downward (negative) deflection, often followed by a slower recovery (positive) phase. Due to their shape, we named them Sharp Negatives (SNs).

My collaborators obtained bilateral recordings containing simultaneous LFP signals from the left and right claustra and found that the regular alternation between REM and NREM activity was closely coordinated between the two hemispheres (Fig. 6.1B bottom). Note that hemispheres in the reptilian brain are connected through several forebrain commissures [209], in contrast to the *corpus callosum* present in placental mammals [210–212]. However, zooming into the finer detail of the signals revealed different coordination of electrophysiological signatures within periods of REM and NREM. While the Sharp Wave Ripples that populate the NREM periods did not appear synchronized bilaterally, periods of REM displayed close coordination of the timing of SNs.

With these initial observations, my collaborators and I set out to investigate the coordination of SNs during REM in more detail. In this chapter, I describe the computational methods and analysis that I developed to quantify this coordination. These methods enabled us to quantify the competitive dynamics of signal propagation across multiple sub-cortical areas during REM, despite several large-data challenges.

Computational challenges

Electrophysiological recordings of local field potentials (LFP) can provide valuable insights into the functioning of the brain by describing the electrical activity of neuronal populations at high temporal resolutions, allowing us to study brain

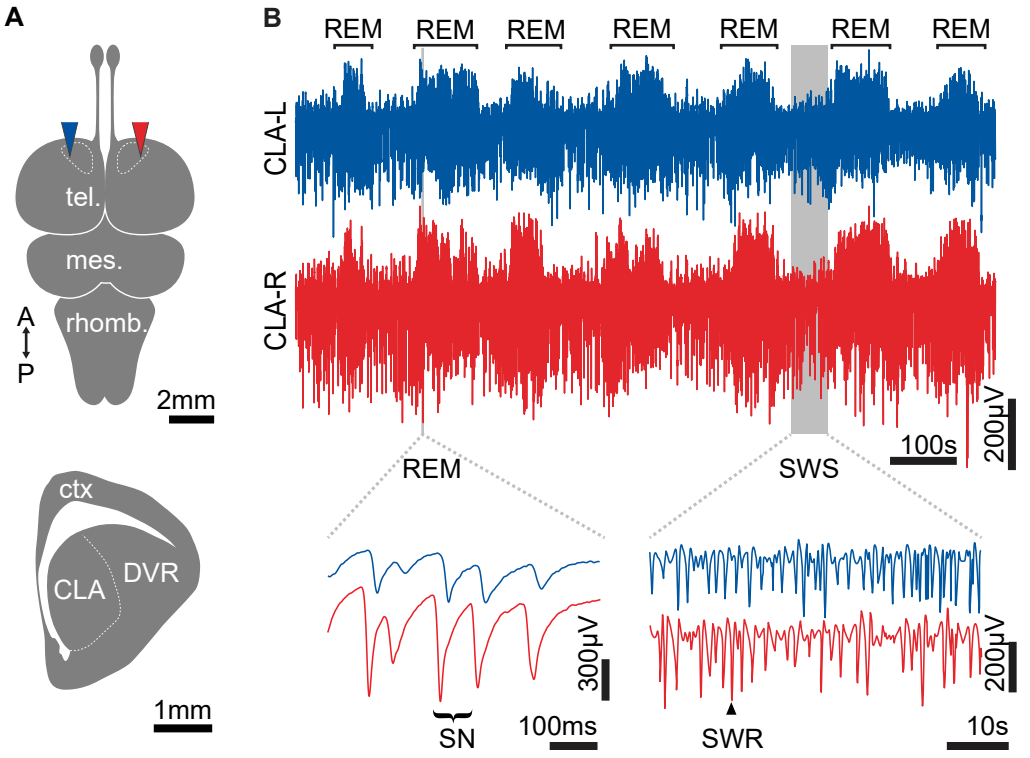

Figure 6.1: Sharp negative field-potential events populate Claustrum REM activity. **A.** Top: Schematic of dorsal view of the *Pogona vitticeps* brain (Telencephalon; Mesencephalon; Rhombencephalon). A: anterior; P: posterior. The dashed line indicates the claustrum in both hemispheres, and arrowheads indicate recording sites. Bottom: Schematic of a coronal section of the telencephalon in the right hemisphere showing Claustrum (CLA), the dorsal-ventricular ridge (DVR), and cortex (ctx). **B.** Top: Eight minutes of sleep local field potential (LFP) recording from sites in A. Note bilateral coordination of REM and NREM alternation. Bottom Left: Zoom-in of LFP during REM sleep with a single Sharp Negative (SN) highlighted. Note the close coordination between left and right claustral SNs. Bottom Right: Zoom-in of LFP during NREM with a single Sharp Wave Ripple (SWR) highlighted. Note SWRs are not coordinated across the two hemispheres. From [2].

processes in detail [119]. However, the large data sizes produced by long recordings, coupled with the need for processing at fine temporal scales, present significant challenges for data analysis.

One major challenge with LFP data is the sampling rate, which, in the recordings that I describe in this chapter, was 30–32 kHz. At the same time, recordings were of over 10 hours in duration since *Pogona vitticeps* is a diurnal animal displaying almost constant sleep throughout the night. Consequently, these recordings provided us with a high number of samples for statistical tests but also generated an enormous amount of data (in the order of 10^9 samples per recording per channel), often running into more than 3 TB of data for every single recording. Manual processing of such an amount of data is simply not feasible. To overcome these challenges, we needed to turn to efficient computational methods that could analyze the data in sections using out-of-core processing techniques.

Another challenge is developing algorithms to detect events and match them in LFP recordings. Considering the size of single SNs, detecting and identifying events at a high temporal resolution was critical to studying coordination in detail. However, most algorithms that detect features of LFP are tailored for the recording probe, area, animal, and protocol of the experiments, making it non-trivial to scale for new experimental questions [32, 119]. Newer methods relying on deep learning approaches to detect, for instance, sharp-wave-ripples, are more general but require that expert humans manually label a large corpus for training the algorithm [213]. For the case of SNs, we needed to identify and extract these features from the LFP recordings and analyze them in relation to our research questions on the bilateral coordination of REM sleep. Consequently, I developed and implemented new methods tailored to our specific datasets while ensuring scalability with data size.

Finally, in order to show that the quantifications and conclusions that we extracted from the data were general enough, it was key that my methods performed well in multiple recordings. This included recordings from different animals with different degrees of noise and recording quality, as multiple factors can affect the quality of the LFP recordings. It was critical that this variability did not impact the accuracy of the analyses. Furthermore, in order for the methods to scale well and be as objective as possible, I had to find ways to normalize the data in order to minimize the amount of per-experiment tweaking of parameters. Our methods had to be robust and adaptable, ensuring consistent and accurate results across multiple datasets.

In summary, the study of LFP recordings presents many challenges for data analysis, including large data sizes, high sampling rates, and the need for analysis at fine temporal scales. In this chapter, I present methods for efficient and accurate analysis of LFP data, including algorithms for SN detection and matching. These methods showed good scalability with data size and robustness to noise and recording quality and could be tailored to our particular scientific questions. Relying on these computational methods, we were able to quantify the tight temporal coordination

between the claustra in the left and right hemispheres during REM periods which in turn revealed signatures of a competitive process with slow dynamics on the scale of minutes involving nuclei in the midbrain.

6.2 Sharp Negative deflections in claustral activity during REM

In order to study the properties of SNs and their relationship to REM sleep in the lizard, we required a method to detect single SNs. Importantly, the initial observations suggested that SNs had a timescale in the order of tens of milliseconds, while the overnight recordings consisted of several hours, which made detection by hand unsuitable. Consequently, I developed and implemented an algorithm to detect individual events that leveraged the sharp phase of the SN.

I first low-pass filtered the raw LFP signal to 40 Hz and extracted its first and second derivatives (Fig. 6.2A top). The procedure involved extracting peaks from the time series and defining a specific pattern consisting of three peaks: a negative first derivative peak surrounded by negative and positive second derivative peaks (Fig. 6.2A bottom). The two peaks in the second derivative correspond to the start and end of the sharp downward deflection of an SN.

Small negative deflections of variable amplitude and duration are common in LFP signals, often the result of electrical noise or other LFP events. To remove false positives, I first extracted two properties of each potential SN: the amplitude and the duration of the negative deflection (Fig. 6.2B). I generated an estimate of the distribution of amplitude and duration for noisy deflections and only considered the SNs with a low probability (p ¡ 0.025) on the cumulative distribution function of the noise. In order to estimate the distribution of small LFP deflections that were mistakenly identified as SNs, I inverted the signal and applied the same triplet peak detection method to detect sharp positive deflections. Since such deflections do not exist in the signal, any positive events detected in this manner were a result of LFP noise. These events were then used to determine the minimum thresholds for amplitude and duration on the SNs that had been originally detected.

The resulting method of SN detection allowed me to isolate single events (Fig. 6.2B) and estimate, for a full 9h recording, the distribution of SN properties (Fig. 6.2C). SNs displayed a skewed long-tailed distribution of amplitudes but a surprisingly uniform distribution of durations, with most SNs taking about 12 ms. The inter-event interval was also skewed and displayed a mode at about 40 ms (median: 60.2 ms; [25th, 75th] percentiles: [39.8, 110.5] ms). Importantly, since my method leveraged the signal to estimate the distribution of noise (rather than fixed thresholds), it adapted well to multiple recordings across multiple animals (Fig. 6.2C) and revealed consistent distributions of SN properties.

Previous results had shown that REM (in DVR and claustrum) is dominated by spectral power in the beta band (12–30 Hz) [30, 32] (Fig. 6.3A). Since the median of SN inter-event-interval fell right within the beta range (40 ms, 25 Hz), we

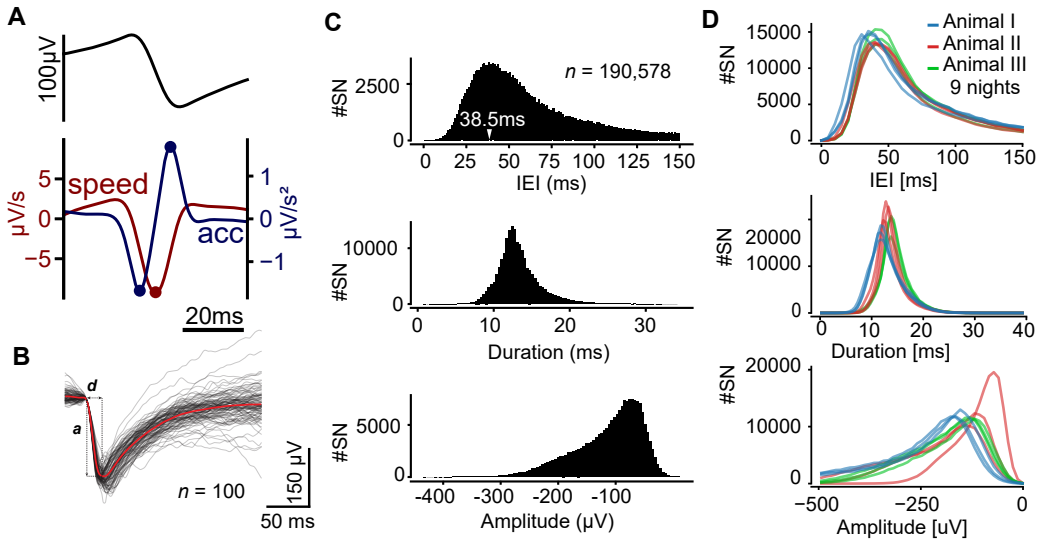

Figure 6.2: Sharp negative detection and properties. **A.** Top: Schematic of zoom-in into a single SN. Bottom: First and second derivative (speed and acc) of the LFP. Dots indicate peaks used to detect SNs. **B.** 100 superimposed SNs with their average in red. Arrows indicate the amplitude (a) and duration (d) of the falling phase. **C.** Statistics for all SNs detected in 9 hours of sleep of a single animal. Top to bottom: Distribution of inter-event intervals (IEI), duration, and amplitude (as in B). IEI distribution truncated at 150ms. **D.** Same as C for 9 different nights of three animals. From [2].

estimated that SNs could be the main contributor to beta spectral power. Using the detected SNs, I could estimate the instantaneous SN rate and compare it with the instantaneous spectral power of the LFP in the beta band (Fig. 6.3B). Note that I normalized beta power so 1 corresponded to its peak value in the recording. The two metrics of LFP were indeed highly correlated across all of our recordings (Fig. 6.3C), suggesting that electrical activity in the claustrum during REM is dominated by SNs and their inter-event interval mode.

LFP reflects mainly synaptic activity in the surrounding population of neurons [119]. Could these sharp and short events reflect strong synaptic input? If so, then spiking activity in the claustrum might be linked to single SNs during REM. Dr. Lorenz Fenk detected spikes from single units using Kilosort2, the ecephys pipeline, and Phy (Fig. 6.4A) [1]. Having access to the entire population of SNs, I could then use the spike times and SN times to extract per-SN-time-histograms (Fig. 6.4B). These histograms showed that single units often fired zero or one action potential per SN and that these spikes were strongly aligned to the SN's descending phase. Given the short duration of SNs (12 ms), the spikes from multiple units occurred within very few ms of one another. Thus, SN events seem to be the changes in extracellular potential brought on by a net-depolarizing current in claustrum neurons, likely reflecting synaptic input. This input then causes the coordinated firing of claustrum units.

6.3 Coordination of activity between the left and right claustra during NREM and REM

I next analyzed the bilateral coordination of sleep-related processes using simultaneous recordings from the left and right claustra Fig. 6.1B.

Re-synchronizing recorded signals

In order to quantify the coordination of LFP, I first processed the data to establish a common temporal frame of reference. For different recordings, my collaborators used either a combination of multiple 32-channel NeuroNexus probes or multiple 384-channel Neuropixel probes [214]. While recordings using multiple NeuroNexus probes shared a single recording setup with a single clock, recordings using multiple Neuropixel probes had a different sampling clock per probe. Although Neuropixel clocks were calibrated together (to 30,000 Hz), small differences in the calibration process (in the order of 0.1 Hz) caused a drift on the sampling process across probes that resulted in a few seconds for the 11-9hour recordings that I analyzed. These small sampling differences became a limiting factor for quantitatively comparing short-duration bilateral phenomena (like REM SNs).

[1] see `https://github.com/jenniferColonell/ecephys_spike_sorting` and `https://github.com/cortex-lab/phy`

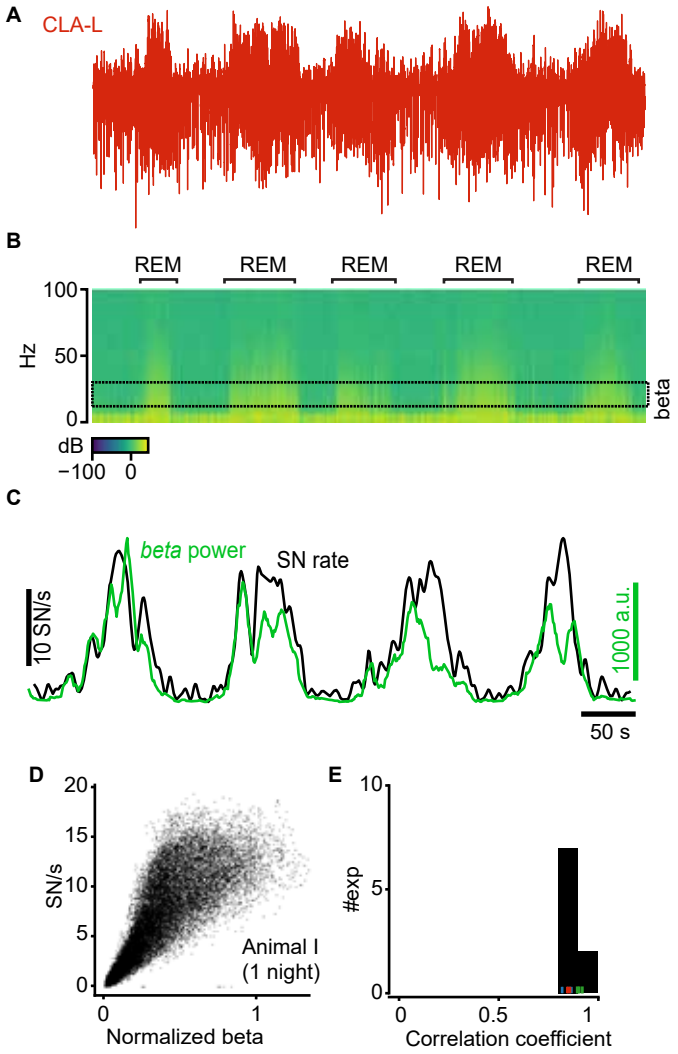

Figure 6.3: Sharp negatives make up beta spectral power during REM sleep. **A.** Top: LFP recorded in the claustrum. Bottom: Corresponding band spectrogram (0.1–100 Hz). Beta band (13–30 Hz) defines the electrophysiological signature of claustrum REM. **B.** Instantaneous rate of SNs and spectral power of the beta band. **C.** Left: Instantaneous rate of SNs and corresponding beta power of the LFP. Each data point represents a sample obtained every second, taken from a sliding window of 10 seconds, during a 9-hour period of sleep in a single animal. Right: Values of Pearson's correlation coefficient corresponding to linear fits between instantaneous SN rate and beta power for 9 different recordings of three animals. From [2].

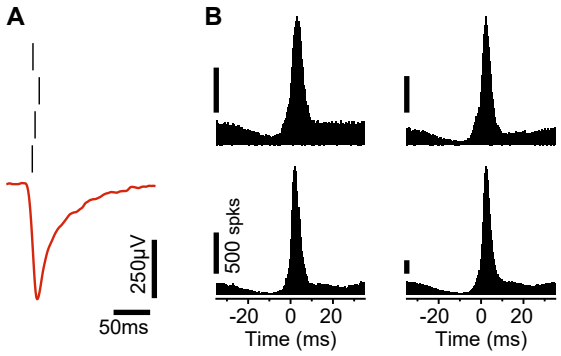

Figure 6.4: Spiking aligned to the downward phase of sharp negatives. **A.** Single example SN (blue) and action potentials from 4 sorted single units (vertical bars). **B.** Histograms of action potential times for the same units in B. Spikes are relative to the time of peak first derivative of the SNs (n = 100,632 events). These units spiked within an SN at least once with probabilities 14-43%. However, the probability of seeing more than one spike was only 0.3-3%. From [2].

In order to compare multi-probe data, I developed a tool pipeline that could load the raw LFP recording, including an artificial 1 Hz square pulse signal recorded in one of the channels. The pipeline then used the onset and offset of the square pulse to linearly interpolate each probe signal to a common 30 kHz sampling frame of reference without drift. Since NeuroNexus probes did not suffer from this limitation, I used those as ground truth to verify the interpolation process. This interpolation process enabled us to take advantage of the higher-yield and longer Neuropixel probes throughout the rest of this study.

Lack of SWR coordination during NREM

To verify the uncoordinated presence of SWR during NREM, I implemented an algorithm to automatically detect single events in a single channel. While previous attempts at SWR automatic detection had relied on template matching or multichannel supervised learning methods [28, 213], we found that SWR signatures in the claustrum of *Pogona vitticeps* were sufficiently clean to apply a simpler method. Using a single-channel recording, I downsampled the signal, low-passed filtered it to 30 Hz, and used an established peak detection that finds local maxima by comparison of neighboring values. This simplified approach was important for the speed of processing, considering the huge number of single SWRs in a single night (around 3,000).

Having access to single SWR times during the night, I extracted average bilateral LFP at each one and constructed average shapes across multiple recordings Fig. 6.5A. Extracting the average LFP of one hemisphere triggered on the SWR times detected on the other hemisphere resulted in a flat shape. After detecting SWR times separately in each hemisphere, I constructed auto-correlograms, revealing a short

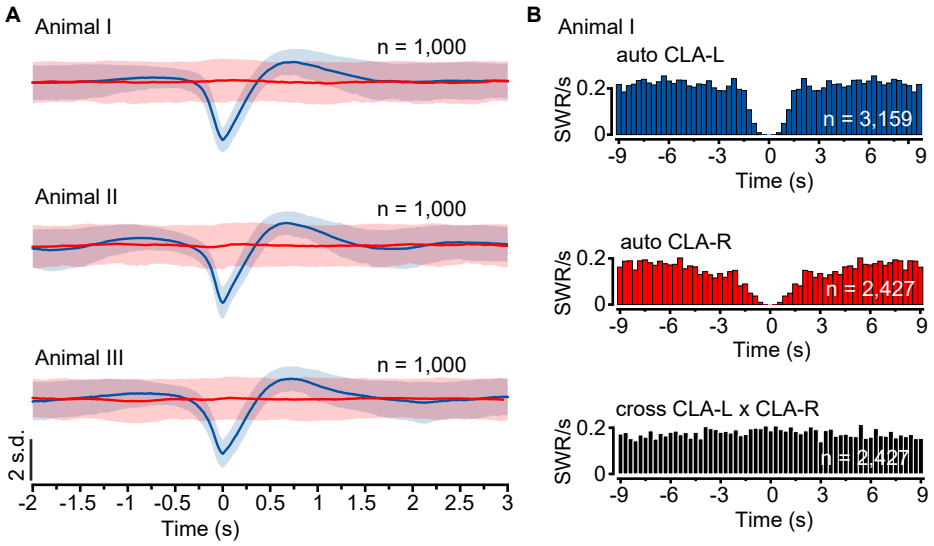

Figure 6.5: Absence of bilateral coordination between claustral SWR during NREM sleep. **A.** LFP average and standard deviation in 5 seconds around the time of 1000 SWRs randomly picked in a 9hrs of sleep. SWRs were detected on the left (blue) hemisphere. Top to bottom: three different animals showing no bilateral coordination of SWRs. **B.** Top and middle: Auto-correlograms of SWR times detected in a single hemisphere. Bottom: Cross-correlogram of SWR times between left and right hemispheres. From [2].

refractory period of around 1.5 seconds, with a lack of strong regularity beyond that (Fig. 6.5B). On the other hand, cross-correlograms displayed a complete flat profile indicating that SWR times are not correlated inter-hemispherically.

These observations were consistent with the independent production of SWRs in each claustrum as observed in slices [30] and the apparent absence of contralateral projections between the two claustra in both *Pogona vitticeps* and mammals [30,215].

Cross-correlations during REM

Having the LFP claustral signals in the same time frame, I turned to cross-correlations to quantify the short-term bilateral coordination during REM. Sliding cross-correlations computed with lags of less than 50 ms yielded clear periods of near-zero correlations corresponding to NREM and of high correlations corresponding to REM (Fig. 6.6A top). However, the precise lag corresponding to the high-correlation periods was difficult to establish. I found that we could leverage the defining feature of SNs, their sharp negative deflections, to obtain higher temporal resolution in our estimates. Because claustral LFP tends to be flat except for the sharp SN deflection and slow recovery phase, the first derivative of the signal produced sharp peaks and otherwise low values (Fig. 6.2). I developed a pipeline that downsampled the signal to 1 kHz, low-pass filtered it to 40 Hz, z-scored it, and then extracted its first derivative. I then computed lagged cross-correlations between the first derivative of both signals. Indeed, I found that this cross-correlation on the first derivative

yielded much more precise estimates of the lags corresponding to high correlations (Fig. 6.6A bottom).

Formally, cross-correlations were computed according to:

$$c(t,l) = \sum_{tau=-w/2}^{+w/2} c_L(t+\tau) \cdot c_R(t+\tau+l)$$

where l is the lag, t is the time of the signal, w is the size of the window (10 seconds), and c_L and c_R are the first derivatives of the z-scored and filtered claustral signals.

To account for rare LFP artifacts (for instance, from motion), I normalized cross-correlations to their value in the top 99.9% percentile. From this lagged cross-correlation, we observed that REM periods were precisely coordinated but clearly not simultaneous. The signals were offset by a lag of nearly 20 ms (median: 19.3 ms; [25th, 75th] percentiles: [15.4, 24.5] ms). Most strikingly, this lag was fixed throughout the night but changed in sign, with alternations between 20 ms (indicating the signal in the left hemisphere preceded the right one) and −20 ms (right preceded left). These switches of signs did not happen with apparent regularity but rather could happen between REM episodes or multiple times within a single REM episode.

To obtain which lags best capture the relationship between two signals for multiple recordings, I extracted the distribution of peak-correlation lags (Fig. 6.6B). I extracted, for every time sample in the lagged cross-correlation of the first derivatives, the lag that yielded the highest correlation. This produced a time series of best lags. Because I extracted cross-correlations for the full experiment, which included periods of low activity (such as NREM), I discarded lags with correlations in the bottom 75% percentile. I also discarded lags at the extremes of the range that could be over-represented due to the finite exploration of lags with caused boundary artifacts. Finally, I extracted the histogram of these lags for several recordings, which revealed surprisingly conserved lag distributions across multiple nights of multiple animals, including recordings with different probe types.

Algorithm to bilaterally match SNs

Interestingly, the z-scored beta spectral power of each claustral signal seemed to correlate with the changes in lag sign, where the beta power of the leading side was higher than that of the following side (Fig. 6.6C). We had already established that SNs were the main contributors to beta power unilaterally (Fig. 6.3C). To clarify the relationship between bilateral SNs and bilateral beta power, we need a method to automatically match bilateral pairs of SNs in order to evaluate differences in their properties that could explain the differences in beta power.

I next describe the algorithm of bilateral SN matching as schematized in Fig. 6.7A. I first used the same method described in the previous section (Fig. 6.2A) to detect

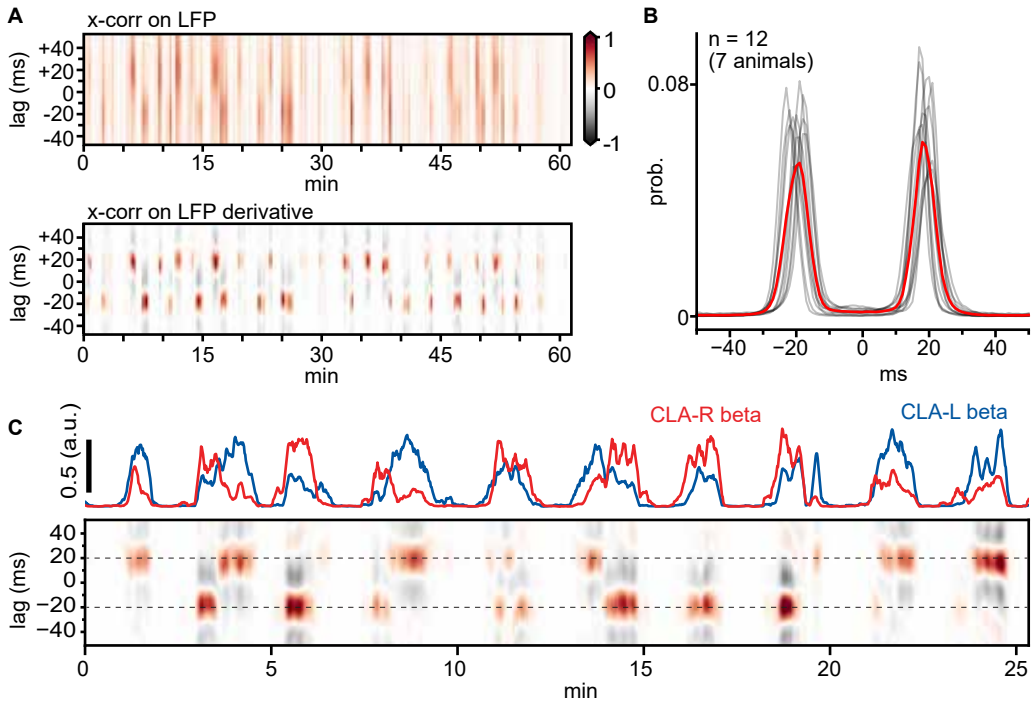

Figure 6.6: Cross-correlation analysis of short-timescale bilateral coordination during REM sleep. **A.** Top: Lagged cross-correlation of bilateral claustral LFP. The correlation value was normalized to the maximum in the recording. Bottom: Same but using the first derivative of each LFP time series. **B.** Distributions of peak-correlation lag between left and right claustral LFP. Note symmetric peaks at ±20ms. Red: average. Extracted from 12 recordings (7 animals) of 9h of sleep each. **C.** Top: Spectral power in the beta band for left and right claustral LFP. Bottom: Lagged cross-correlation of the first derivative for the same period. From [2].

single SNs in both signals independently. Note that, possibly due to LFP noise or limits of my method, not all SNs are detected, meaning that the total number of SNs in each hemisphere is very similar but not exactly equal. To reduce the number of falsely unmatched SNs, I relaxed the limits on the comparison to the LFP noise distribution so that a bilateral pair of SNs was considered valid if, after matching (as described below), at least one of the two SNs was above noise levels (see SN detection in section 6.2).

When considering all potential SN matches, not all are possible: because time only flows in one direction, once left and right SN at a given time point have been matched, none of the subsequent matches can include an SN that happened before and one that happened after that time point. The result is that the combination of possible matches creates a directed acyclic graph that moves forward in time with its nodes represented as open circles in Fig. 6.7A (middle). The problem of SN matching is then finding the best possible path in this graph. Because every node can be mapped to a time and a lag, I could assign them a value from the corresponding cross-correlation of the first derivative. To increase temporal resolution in the lag dimension, I extracted the lagged cross-correlation with a sliding window of 100ms and multiplied it by an exponential kernel with a time constant of 10ms. That is, nodes may have positive values (high local cross-correlation) or negative values (anti-correlations). Based on our previous observation that ±20ms captures the best lags (Fig. 6.6) with small variance, I could further limit the evaluation of potential SN pairs to those with a lag up to 50 ms.

By propagating the accumulated value (sum) backward from the latest possible nodes, I construct a simple greedy algorithm that detects the path of maximum accumulated value (magenta in Fig. 6.7A bottom). Although the complexity of comparing all possible SN pairs would make the algorithm quadratic ($\mathcal{O}(n_L \cdot n_R)$, where n_L and n_R are the number of SNs detected on each side), the limits on lags and time result in an approximately linear complexity (approx. $\mathcal{O}(\frac{n_L+n_R}{2})$). With these optimizations, I could detect and match SNs in consecutive sections of 1 hour.

Matched SN pairs

From my SN matching algorithm, we estimated that out of the roughly 200,000 clearly detectable SNs in 9 hours of sleep, about 85 to 90% of the events could be matched (Fig. 6.2B).

Note that beta spectral power captures the amplitude of oscillations in the beta band, which can be affected by both the rate and the amplitude of SNs. I next classified SN pairs by their relative lag (negative: right leads, positive: left leads), and calculated their pairwise differences in amplitude (Fig. 6.7C). Note that I used SN amplitudes after z-scoring each LFP time series independently since LFP absolute amplitude can be strongly affected by variability in the recording probe and site. Indeed, we observed that SNs that were on the leading side showed a consistently higher amplitude than their pairs on the following side. Note that SN amplitude distributions are skewed (Fig. 6.2C), which results in also skewed

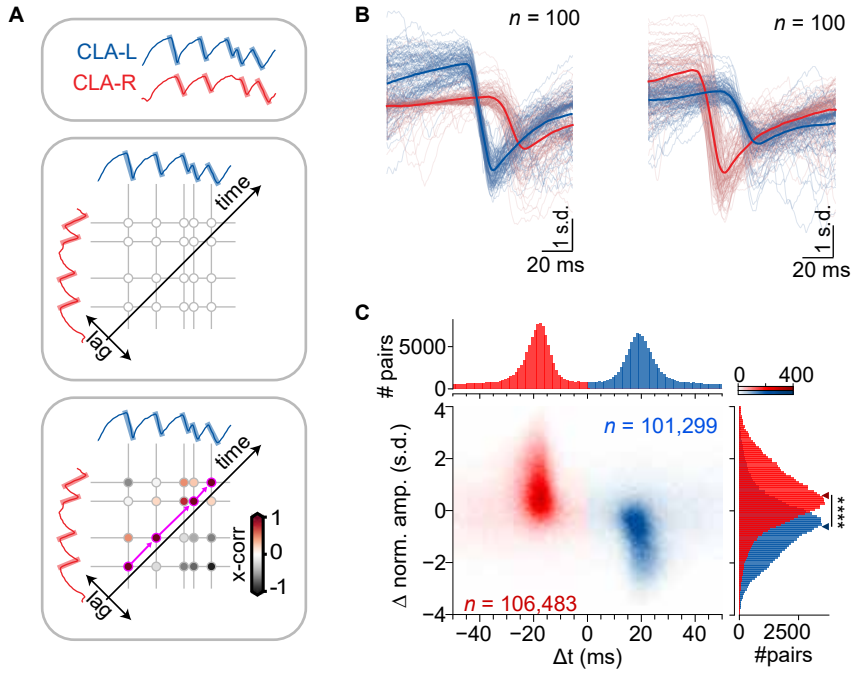

Figure 6.7: Bilateral SN matching. **A.** Schematic of the algorithm to bilaterally match SN. Single SNs were detected independently in each channel. The value of every potential pair of SNs was taken from the cross-correlation on the first derivative of the signal. The final matching traced a path that could only move forward in time and maximized the total sum of SN pair values. **B.** Example of 100 matched bilateral SNs for left-leading periods and for right-leading periods. The thick line is the average. **C.** 2d distribution of differences between amplitudes and time lag for over 200,000 bilateral pairs of SNs. Note SN amplitudes are first z-scored using single-side data. Histograms indicate marginal distributions. The color indicates side leadership for each pair. Note SNs on the leading side have higher amplitudes than their contralateral partner. Mann-Whitney U test, U = 2312573585.5, ****P = 0.0. Arrowheads indicate means. From [2].

distributions of amplitude difference, with tails in opposite directions (Fig. 6.7C right marginal distributions). We concluded that it was a change in amplitude, not in rate, of SNs that caused the enhanced beta power of the leading claustrum.

If claustral SNs are the product of synaptic input, then higher amplitude may correspond to a stronger input. To evaluate the effect of this potential high synaptic input, I extracted spike-time histograms for sorted single units in the left claustrum. Spike-time histograms relative to SNs detected in the same (left) claustrum showed tight alignment to the downward phase of SNs again (Fig. 6.8A, compare to Fig. 6.4; different units). Histograms relative to SNs in the opposite (right) claustrum, however, showed much-reduced firing modulation and two small peaks (Fig. 6.8B) corresponding to ±20 ms. Importantly, the earliest of both peaks (-20ms) was

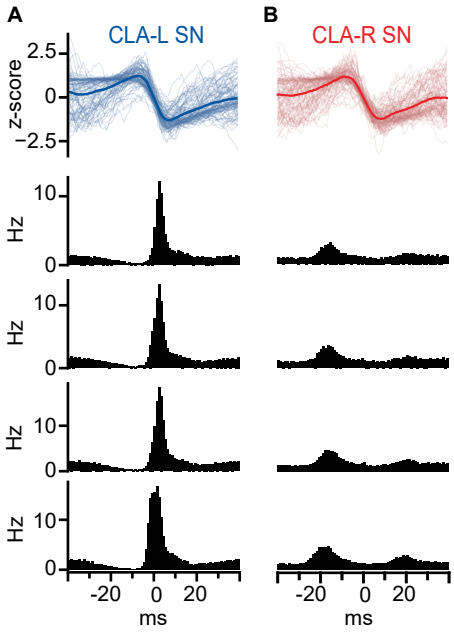

Figure 6.8: Claustral spiking is most common during ipsilateral SN leadership. **A.** Top: example 100 SNs detected in the left claustrum. The thick line indicates average. Rows: Histograms of spike times for 4 different units in the left claustrum. Spike times are relative to the time of peak first derivative of the temporally closest SN in the left claustrum. **B.** Top: example 100 SNs detected in the right claustrum. The thick line indicates average. Rows: Histograms of spike times for the same 4 units in the left claustrum shown in A. Spike times are relative to the closest SN in the right claustrum. Note the increase of firing rate is locked to the time of the corresponding left SN (−20 ms and 20 ms) and is more likely to spike when the left claustrum is leading (−20 ms). From [2].

consistently the highest one, suggesting that SNs drive spiking more strongly when the corresponding claustrum is leading and less strongly when it is following.

In summary, we observed matching bilateral pairs of claustral SNs, showing fixed lag with alternating positive or negative signs and a stronger signal on the lead side. The lags showing the same absolute value suggested that the circuitry projecting onto the claustrum (and driving claustral spiking) might also be hemispherically symmetric. The difference in leadership and their relationship to amplitude and spiking suggested that these input signals could be the result of a competitive process upstream of the bilateral claustra. Under this competition, the stronger side might impose its output on the weaker side momentarily before losing its dominant position.

6.4 Dynamics of left and right claustral leadership.

We next looked at the dynamics of bilateral dominance in order to better describe and understand this hypothetical competitive process. I quantified the switching of dominance at increasing time scales: that of single REM periods, between consecutive REM periods, and for the entire night of sleep.

Quantifying leadership

First, we required a method to strictly classify epochs of dominance. Indeed, cross-correlations suggested that dominance had a time scale bigger than the duration of single SNs, but that it could also change rapidly within a single REM period. While dominance was often stable over seconds, sometimes it could not be clearly defined by eye, with SN leadership seemingly jumping back and forth between left and right (Fig. 6.9A yellow period). These periods where leadership is not settled often corresponded to spectral beta power balanced between the two claustra.

To provide a simple score of the level of leadership between the two claustra at any single moment, I leveraged the cross-correlation of the first derivative of their LFP at the two most common lags (-20 ms and $+20$ ms). Note that because SNs show a long-tailed distribution of amplitudes but an almost fixed distribution of durations, the distribution of values of the first derivative was also long-tailed, as was their cross-correlation. A common approach to long-tailed distributions is to work on log scales. However, because the two signals could move in opposite directions, the cross-correlation could contain negative values, which prevented a log transformation. Consequently, before calculating the cross-correlation, I clipped the first derivative of the signal to be ≤ 0, which ensures that the resulting cross-correlation must be ≥ 0. In practice, small noisy deflections of the LFP resulted in the cross-correlation being > 0, which allowed me to take the base-10 logarithm of the resulting values.

Formally, I defined the leadership score as:

$$s_{\pm 20}(t) = \log_{10}\left(\sum_{tau=-w/2}^{+w/2} [c_L(t+\tau)]^- \cdot [c_R(t+\tau \pm 20)]^-\right)$$

where t is the time of the (left) signal, w is the size of the window (10 seconds), and c_L and c_R are the first derivatives of the z-scored, and low-pass filtered claustral signals.

The two resulting time series of leadership score (s_{+20} and s_{-20}) defined a 2-D state space for the full experiment (Fig. 6.9B). This state space always displayed three areas of high density: low score for both lags, corresponding to NREM, or high score for one lag, and medium score for the other, corresponding to left or right dominance during REM. Periods of undefined dominance were captured as transitions between these two REM states (see the cyan trace in Fig. 6.9B). Simple linear thresholds could segregate these areas to establish a strict classification of dominance. Because the signal could cross these thresholds very transiently, which

resulted in an over-fragmentation of the state classification, I implemented the thresholds so that single detours that left and re-entered the same state for a short duration (< 3 seconds) would be ignored.

By applying this state space visualization and state definition, I could then quantify the properties of periods of dominance for multiple recordings. Periods of undefined dominance (as in Fig. 6.9A) were indeed rare: an average REM episode contained approximately 48 seconds of left or right dominance and only 4 seconds of undefined dominance. The short duration of undefined periods suggests the hypothetical competitive process can quickly settle its winner-take-all dynamics within very few seconds and then sustain that state for a long time period. The majority of REM periods displayed one or two epochs of different dominance, that is, zero or one switch, with rarer REM periods showing up to 8 switches (Fig. 6.9C). Those REM periods with zero dominance switches lasted about 60 seconds (mean: 57 s, std: 16.6 s, n = 59), while those with a single switch took 90 seconds (mean: 88.5 s, std: 19.6 s, n=103). Note the diagonal trend in Fig. 6.9D, indicating that many REM periods had a similar length independent of the number of switches, which suggests some type of adjustment where the duration of multiple dominance periods within the same REM adds up sub-linearly.

Dominance across multiple REM/NREM cycles

Next, I investigated whether dominance switching showed dynamics with a time scale longer than single REM periods. Using my state classification method, I could extract dominance for all 210 REM periods in 9 hours of sleep (Fig. 6.10A). The type of dominance at the end of one REM episode was often different from the type of dominance in the immediately subsequent episode. The opposite seemed more common: because most REM periods showed a single switch, I repeatedly observed contiguous REM periods with the same first and second type of dominance. This relationship was common but not stationary throughout the recording, as dominance durations waxed and waned and flipped ordering.

To quantify the relationship between dominance order and dominance duration across consecutive REM periods, I calculated, for each REM period, when it showed left or right dominance relative to the REM center. Note that these values did not always exist since many REM periods showed no switching. I then computed, for each dominance type separately, a running average using a sliding window of 12 minutes and ignoring missing values (Fig. 6.10B). The two resulting time series displayed regular fluctuations in opposing directions, indicating slow changes in dominance duration, and often crossed, indicating a flip of dominance order. I then extracted auto-correlations of these time series across multiple recordings and observed a moderate positive correlation across all animals around 30 minutes.

All together, dominance switching seemed to display slow dynamics in the order of several minutes and multiple consecutive REM periods, suggesting that the hypothetical competitive process works at a timescale longer than that of the REM/NREM cycle.

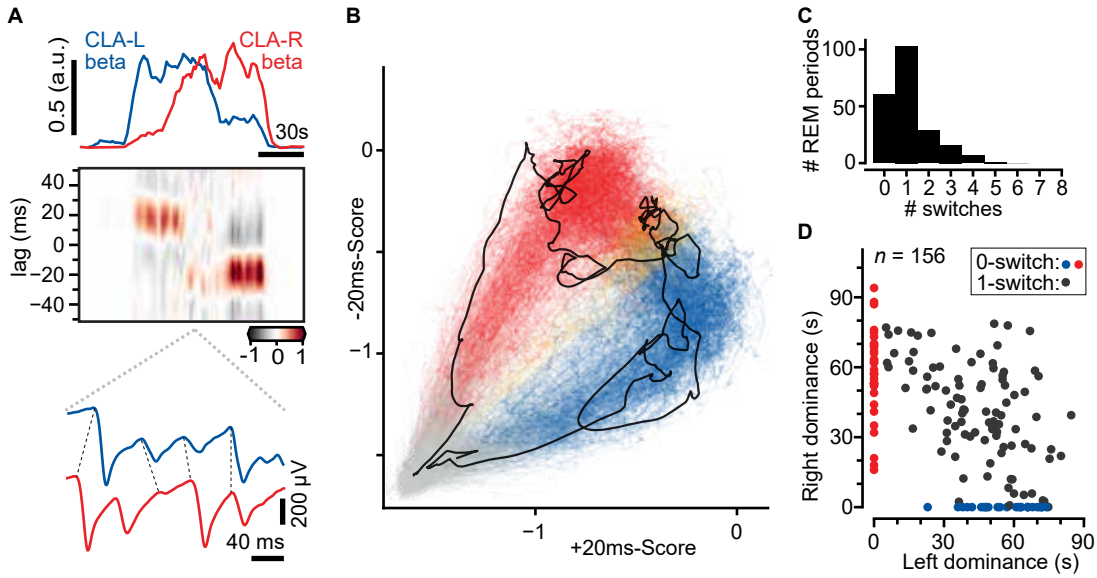

Figure 6.9: Dominance switching within REM periods. **A.** Top: Spectral power in the beta band for left and right claustra over one REM period. Middle: Lagged cross-correlation of the first derivative for the same period. Colored underline as in B. Bottom: Zoom-in of bilateral SN relationship showing a period where dominance is not clearly established. Note unestablished leadership corresponds to a moment of balanced beta power (Top). **B.** State space of cross-correlations. Each data point corresponds to a single sample in time taken every 100 ms. Red and blue coloring indicates right and left leadership. Yellow indicates undefined dominance. Gray indicates NREM periods. Highlighted in black is the REM period shown in A. **C.** Histogram of REM periods by the total number of switches happening for a single experiment. **D.** Duration of leadership periods for REM periods with 0 or 1 switch. Black indicates those with a single switch. Red and Blue indicate those with a single type of leadership. Each dot is one period. The sum of X and Y indicates the total duration of a period. Note that REM periods with one switch have an average duration of 90 seconds on average, while episodes with no switch tend to last 60 seconds. From [2].

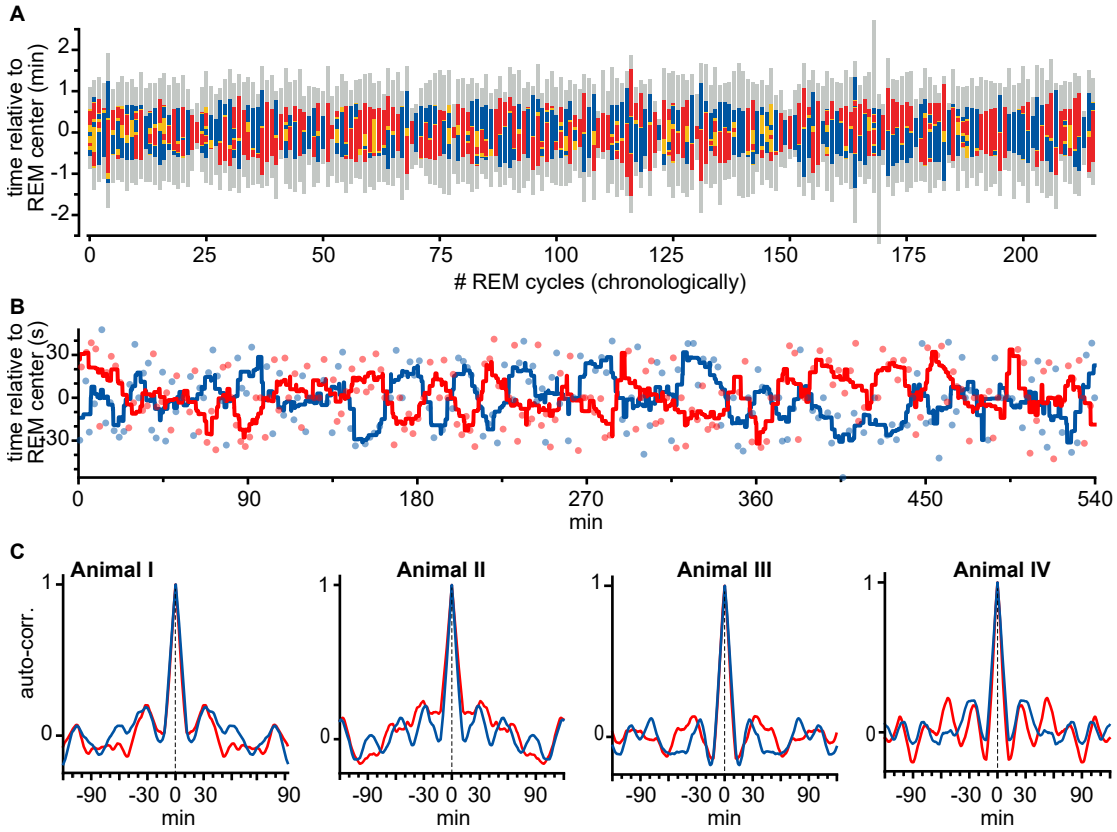

Figure 6.10: Slow dynamics of leadership across REM periods. **A.** Time spent in each type of leadership for 220 consecutive sleep cycles in 9 hours of sleep. The color indicates the leadership state (blue: left; red: right; yellow: not established; gray: NREM). **B.** Slow alternation in the propensity of each type of dominance over 9 hours of sleep. Each REM period is represented by two dots with the same X value, corresponding to left (blue) or right (red) dominance. Each dot indicates the center time of a dominance period relative to its REM center time). Each line shows a running average in a 12-minute window for each dominance period. Note regular slow fluctuations longer than a single REM cycle where dominance change in its order within the period and its duration. **C.** Auto-correlations of running averages of dominance centers (same as B). Note that all four animals display positive auto-correlations near the 30-minute mark. Animal A corresponds to the data shown in B. From [2].

Dominance throughout the night

Finally, I explored whether the slow dominance dynamics were stationary throughout the night. Cross-correlations early in the recording revealed that REM-to-REM changes of dominance were common, as well as within-REM switches (Fig. 6.11A top). On the other hand, cross-correlations late in the recording often showed extended epochs where contiguous REM periods shared the same type of dominance, with little or no switching (Fig. 6.11A bottom).

To quantify this effect, I turned to the four hours immediately before lights were turned on during the experiment and extracted, in a sliding window, the percent of the time that the animal spent in each type of dominance (Fig. 6.11B). The result showed extended epochs of mostly single dominance.

This result suggests that not only the dynamics of dominance switching are slow, but they are governed by an even slower process, in the scale of hours, that makes the winner-take-all dynamics volatile early during sleep and more robust as the night progresses.

In summary, by quantifying dominance periods at the scale of single REM periods, multiple REM periods, and full night, we observed signatures of a multi-scale competitive process. While we observe dominance in the form of single SNs, which take a few milliseconds, leadership often takes several seconds, suggesting the mechanism driving the competition might be different than those producing the inputs that result in claustral SNs. The patterning of leadership switching involves several minutes, affecting multiple consecutive REM periods. Finally, this patterning of leadership evolves slowly through the night, at the scale of hours, with early periods showing more frequent switches than late periods.

6.5 Coordination between REM activity in the claustrum and an ipsilateral midbrain nucleus

The properties and dynamics of bilateral claustral coordination during REM suggested input from an upstream region, likely in the form of a mirror-symmetric circuit implementing bilateral competition. Note that the two claustra do not share direct bilateral connections. Using the quantifications developed in the previous sections, we could quickly evaluate the results from lesion studies. We discarded several brain areas with known ipsi- and contralateral connections as lesioning them did not show an effect on the bilateral leadership switching, including the cortex, amygdala, and parts of the striatum (see Extended Data Fig. 5 in [2]). My collaborators then turned to retrograde tracing and electrophysiological recordings to explore the activity of nuclei in the junction between the midbrain and hindbrain, which are known to be involved in REM control in mammals [124, 216, 217]. Traversing the Optic Tectum with long Neuropixel probes, they found activity in the mesencephalon that closely resembled the SNs that we studied in the ipsilateral claustrum (Fig. 6.12). The recording site corresponded to a nucleus that they

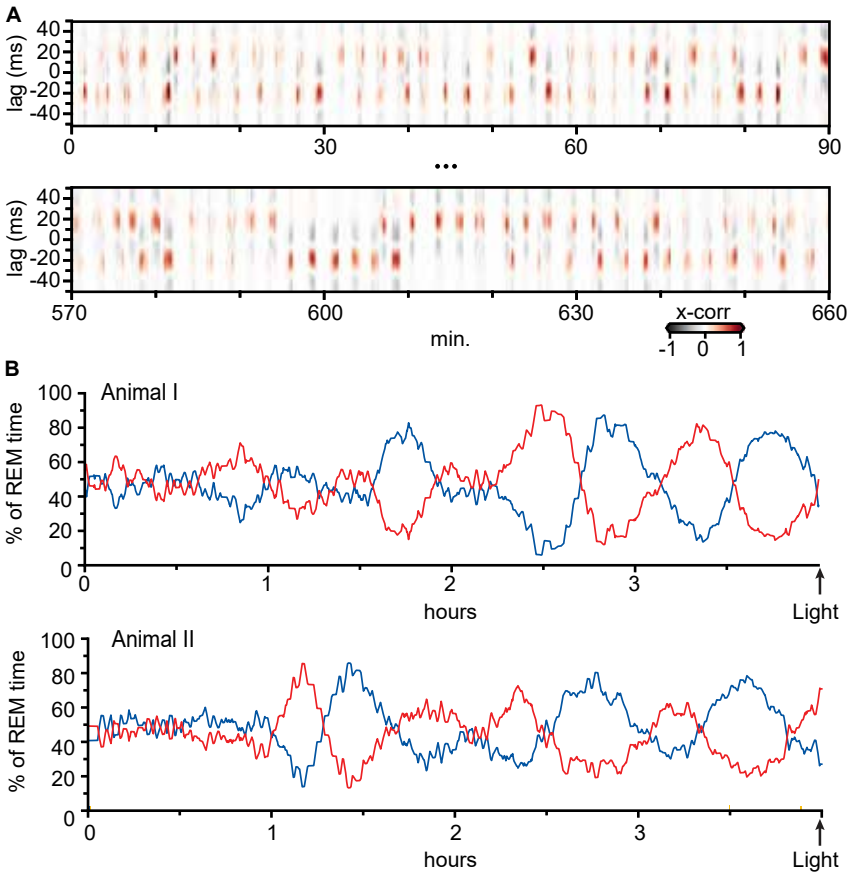

Figure 6.11: Changing dynamics of leadership throughout the night. **A.** Cross-correlation of LFP first derivative during the first and the last 90 minutes of the same recording. Note leadership switching from one REM period to the next is more common early in the night, while leadership remains over multiple REM cycles late in the night. **B.** Percent of REM time spent in each type of leadership (red: right, blue: left) during the 4 final hours of sleep immediately before lights are turned on. Note that both animals display slow fluctuations that result from a tendency towards one-sided dominance over multiple consecutive REM periods. From [2].

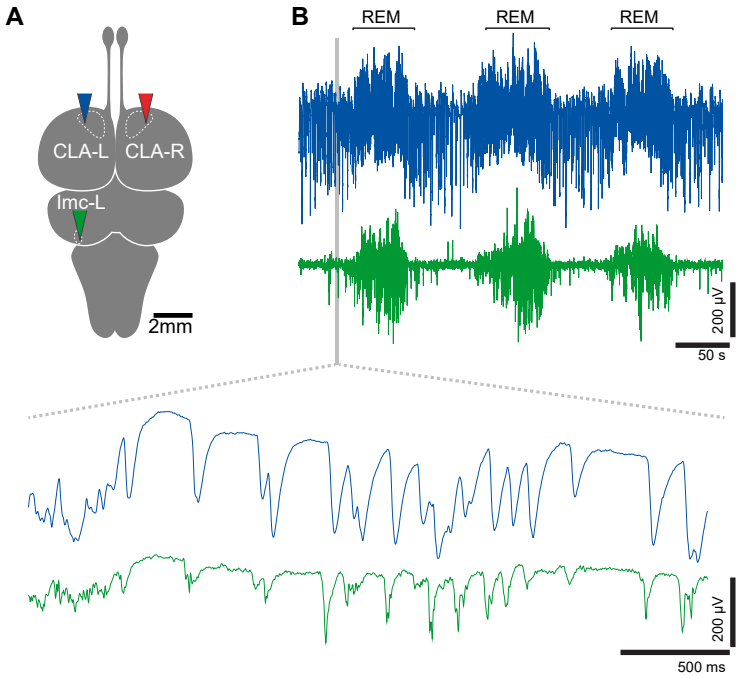

Figure 6.12: Recordings in Imc reveal close coordination with ipsilateral claustral SNs. A. Schematic of dorsal view of the *Pogona vitticeps* brain. The dashed line indicates the left and right claustra and left Imc. Arrowheads indicate recording sites. **B.** Example LFP recorded in left claustrum (blue) and in left Imc (green). LFP was low-pass filtered at 100 Hz. Zoom-in shows closely coordinated LFP activity between both. From [2].

identified as the reptilian homolog of the avian nucleus isthmi pars magnocellularis (Imc). This nucleus, composed of large and sparse GABAergic neurons, is part of a complex of isthmic nuclei, including the cholinergic isthmi pars parvocellularis, that, together with Optic Tectum, have been intensively studied in avians for their role in the selection of attention to competing visual stimuli [218–222].

Quantifying Imc-Claustrum coordination

In order to evaluate pairwise relationships between Imc and its ipsi- and contralateral claustrum, I generalized the pipeline of data quantification that I had developed for comparing bilateral claustrum recordings. This included the interpolation method to correct for Neuropixel diverging sampling frequencies and the cross-correlation of the first derivative (both described above in section 6.3). The resulting cross-correlation indicated that Imc activity preceded that in its ipsilateral claustrum by 30 ms (median: −30 ms; [25th, 75th] percentiles: [-33.0, -27.0] ms) and its contralateral claustrum by 50 ms (median: −45 ms; [25th, 75th] percentiles: [-51, -26] ms) (Fig. 6.13A). These numbers are self-consistent with our previous observation that the two claustra shared a 20 ms delay between them (Fig. 6.6).

Importantly, when I extracted distributions of peak-correlations, we observed that Imc displayed a single modal peak with respect to its ipsilateral claustrum (Fig. 6.13B), indicating a fixed time delay irrespective of leadership. More importantly, the distribution of peak-correlations between Imc and contralateral claustrum also displayed a single prominent peak (at −50 ms). If Imc was active independent of leadership, with a fixed delay to its ipsilateral claustrum, we should have observed, in its relationship to the contralateral claustrum, two peaks: one at 50 ms (30+20, which we observed) and one at 10 ms (30-20, which we did not). This single delay contrasts with the balanced double peak of delays between claustra and indicates that ipsilateral Imc shows sharp deflections of LFP only when its ipsilateral claustrum is leading (Fig. 6.13C).

In order to verify that Imc shows SNs only when its ipsilateral claustrum leads, I extracted single claustral SN times (method described in section 6.2) and extracted the LFP in all three channels aligned to these times (Fig. 6.13C). We observed strong LFP deflections in Imc consistently and only when its ipsilateral claustrum was leading its contralateral claustrum. Note that these deflections are very precisely timed, which I could confirm by extracting the average LFP shape in Imc triggered on leading SNs of ipsilateral claustrum (Fig. 6.13D). Altogether, these temporal relationships suggested that the apparent rivalry between the claustra may really be the manifestation of a bilateral competitive process in the midbrain involving the GABAergic Imc and excitatory inputs into claustra.

Quantifying bilateral Imc-Imc relationship

Intrigued by the relationship between bilateral Imc, my collaborators performed a new set of recordings which included both Imc and one claustrum (Fig. 6.14A). Due to spatial and surgical limitations, only three Neuropixel probes could be inserted into the brain simultaneously. I then processed the data and extracted cross-correlations between each Imc and the left claustrum (Fig. 6.14B) and observed the same delays, with activity in Imc ipsilateral to claustrum preceding by about 30 ms, and contralateral Imc by 50 ms.

Cross-correlations between both Imc did not show high values or consistent delays. Indeed, when we inspected the cross-correlations with claustrum more carefully, we observed that only one of the two Imc would be highly correlated with claustrum at any given time, but never both (see zoom-in in Fig. 6.14B). I further quantified this by extracting the cross-correlations at the relevant delays (−30 ms and −50 ms) and calculating the correlation between them (Fig. 6.14C). The result was a mild anti-correlation (Pearson's r = -0.149), with most samples over time being axis-aligned. That is, correlated LFP activity with claustrum was mutually exclusive between both Imc nuclei.

Spiking in Imc

Next, we inspected firing activity within the bilateral Imc nuclei from single-unit spikes (Fig. 6.15A). I binned spike counts and smoothed them with a Gaussian filter

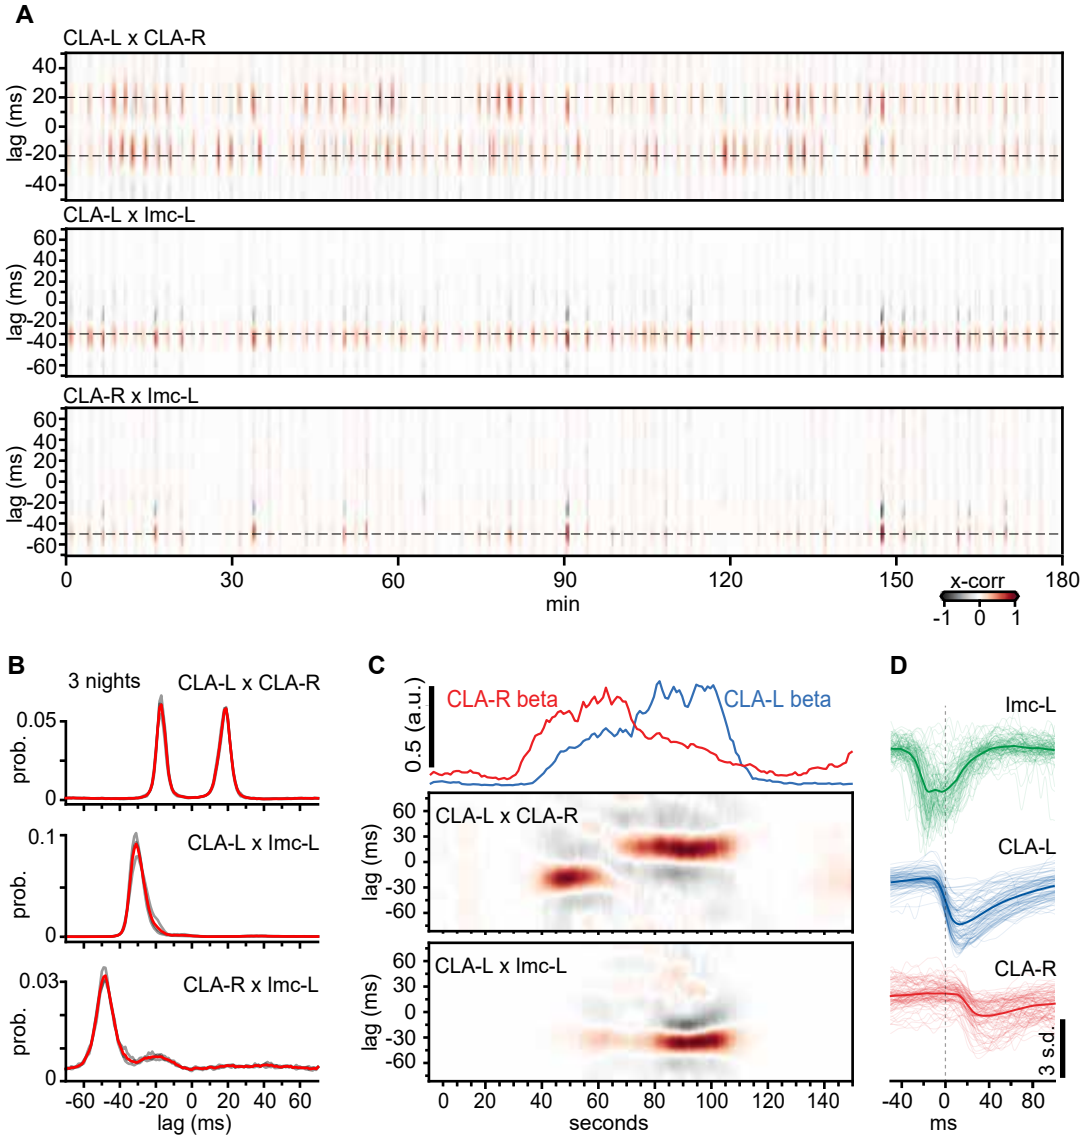

Figure 6.13: Coordination of bilateral claustra and Imc. **A.** Cross-correlation of LFP first derivative between all three recorded areas (bilateral claustra and left Imc). **B.** Distributions of peak-correlation lag between all three recorded areas. Note left and right claustra display symmetric peaks at ±20ms, but correlations with Imc yield a single peak with a larger, negative value, indicating Imc precedes both claustra with a fixed delay. Imc precedes ipsilateral claustrum by 30ms and contralateral claustrum by 50ms (30 + 20). **C.** Top: Spectral power in the beta band for left and right claustra over a single REM cycle. Middle: Cross-correlation of LFP first derivative recorded in both claustra. Bottom: Cross-correlation of LFP first derivative recorded in the left claustrum and left Imc. **D.** 100 superimposed SNs detected in the left claustrum (blue) and the corresponding LFP in ipsilateral Imc (green) and contralateral claustrum (red). Thick lines indicate averages. From [2].

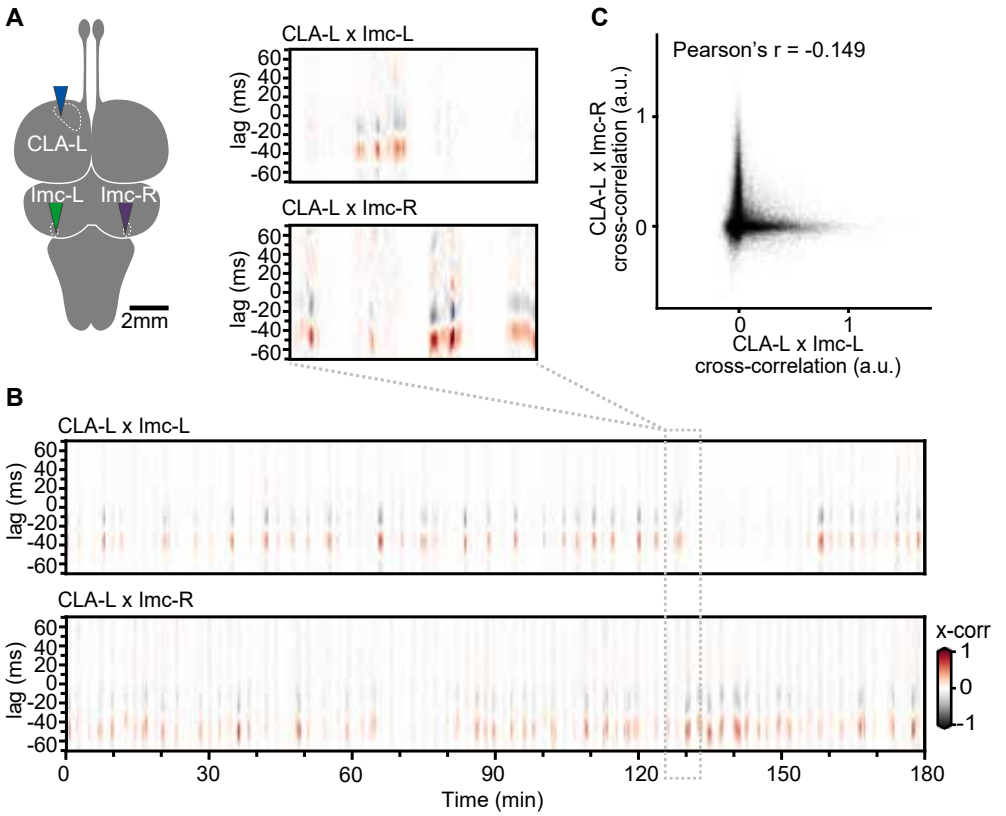

Figure 6.14: Mutual exclusion between bilateral Imc nuclei in their correlation to a single claustrum. **A.** Schematic of dorsal view of the *Pogona vitticeps* brain. The dashed line indicates the left claustra and left and right Imc. Arrowheads indicate recording sites. **B.** Cross-correlation of LFP first derivative between left claustrum and ipsi- and contralateral Imc during the same 3 hours of sleep. Zoom-in: claustrum correlates strongly with only one of the two Imc nuclei at any time. **C.** Cross-correlation value of the same left claustrum to both Imc nuclei. Each dot is a sample taken every 100 ms during 9 hours of sleep (n = 130911). Note that samples are axis-aligned and do not populate the top right quadrant, indicative of mutual exclusion. Pearson's r = -0.149, P = 0.0. From [2].

to produce an estimate of instantaneous firing rates that could be averaged for an entire Imc population.

We observed that Imc units spiked almost exclusively during REM periods and remained silent during NREM. Importantly, firing in one of the Imc nuclei was suppressed whenever the other one showed high firing rates (Fig. 6.15B), consistent with my previous analysis on the cross-correlation of the LFP. Note that the main contributor to LFP is expected to be synaptic activity [119]. This analysis confirmed that not only putative synaptic activity (LFP) in bilateral Imc is mutually exclusive, but also its spiking output. Importantly the timing of switches between activations of Imc seemed to match power changes in a single claustrum (Fig. 6.15A top), although we could not directly evaluate the relationship to claustral leadership due to recording limitations.

Additionally, we found that firing rate domination between Imc tended to become longer towards the later stages of the experiment. These longer periods of domination were consistent with my quantifications of the slow dynamics of claustrum leadership (see Fig. 6.11), where a single claustrum tended to lead for many consecutive REM periods late in the night.

Altogether, our results strongly suggested that the patterns of claustral domination reflect the output of a competitive process involving bilateral Imc.

Imc lesions

To test the role of Imc nuclei on bilateral claustral leadership, my collaborators used ibotenic acid to unilaterally lesion Imc and recorded from both claustra (Fig. 6.16A). I used my data processing and analysis pipeline to extract cross-correlations of the first derivatives of the recorded LFP. We again found strong correlations during all REM periods, but now only one of the two claustra, that in the non-lesioned side, led the other one. This was visible in a uni-modal distribution of peak-correlation lags Fig. 6.16B (8 nights from 3 animals).

Remember that beta power in a given claustrum LFP takes an intermediate value when that claustrum is following and a high value when it is following (see Fig. 6.3). Thus, we expected that lesioned and non-lesion claustra show only one of the two levels of beta. However, due to the variability of recording quality across probes, we had to normalize beta to its maximum value (for all previous analyses, see section 6.2), which, in this case, makes the differentiation of intermediate and high values of beta impossible. Consequently, I normalized beta to its lowest 10th percentile, which corresponded to NREM periods. The resulting normalized beta powers confirmed that the claustrum in the non-lesioned side consistently remained at a relatively higher beta than its partner.

If SNs compose most of claustral REM activity and they reflect synaptic input, it is possible that sustained dominance of one of the two hemispheres may lead to alterations of REM duration. In order to quantify the effects on the REM/NREM cycle, I algorithmically extracted REM durations from multiple recordings. I

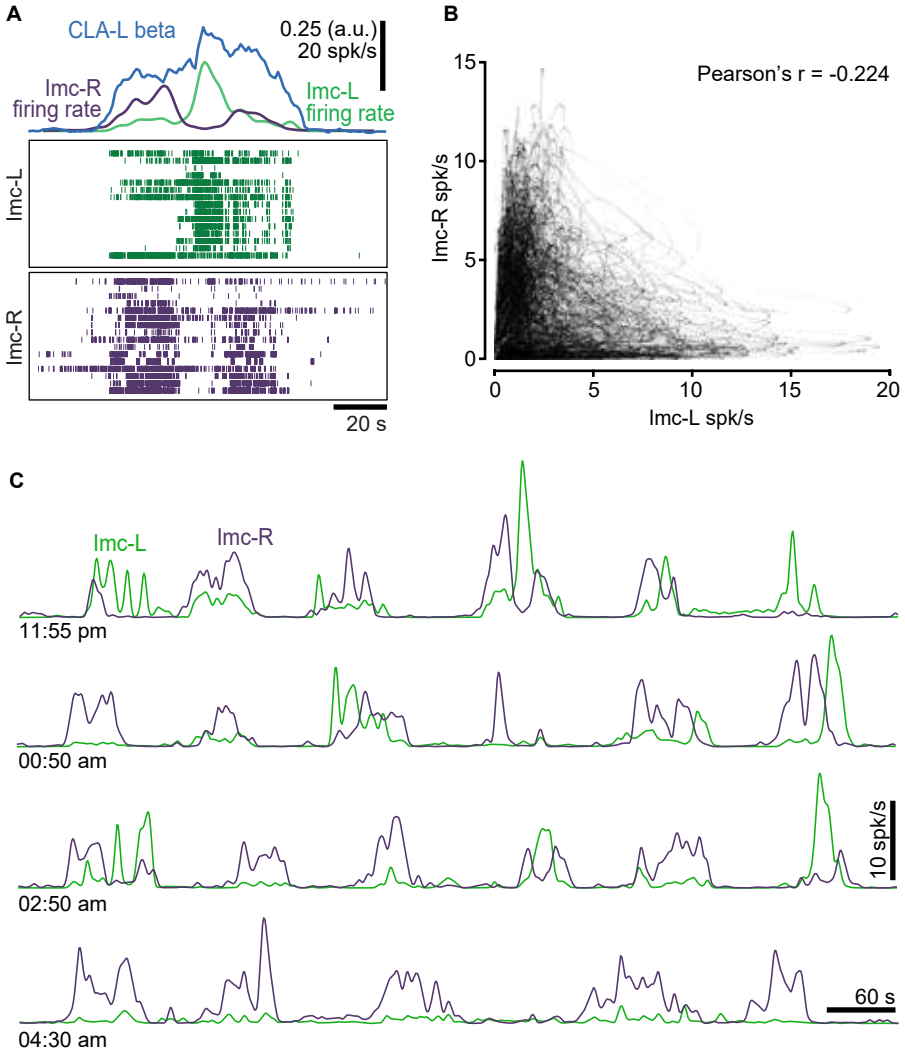

Figure 6.15: Antagonistic spiking of single Imc units during REM. **A.** Top: Spectral power in the beta band for left claustrum over a single REM cycle (blue) and corresponding firing rates for both left (green) and right (purple) Imc nuclei. Middle and center: Spike raster of sorted single units in left (green) and right (purple) Imc nuclei. **B.** Firing rate of left and right Imc, with each dot representing a sample taken every 100 ms. Only periods of REM sleep over 9 hours of continuous sleep were used (n = 130,911). Pearson's $r = -0.224$, $P = 0.0$. **C.** Mean firing rates of sorted left (green) and right (purple) Imc units over four different periods during the night. Note both nuclei show periodicity of spiking increases corresponding to REM cycles and are mostly silent during NREM. Also note the sustained firing of right Imc in the last section of the experiment, consistent with longer periods of dominance (Fig. 6.11). From [2].

extracted beta spectral power (12–30 Hz) using a 10 s sliding window at 1 s steps and defined REM as periods where power was above its 15th percentile. Note that simple thresholding easily leads to over-fragmentation that results from the noisy crossing of the threshold, especially in transition periods between NREM and REM and in transition periods between leadership modes (periods of undefined dominance, see Fig. 6.9), which can strongly bias our estimations of REM duration. Thus, I implemented an adaptive thresholding strategy that ignored any intervals that abandoned a state (either REM or NREM) and then re-entered it for less than 15 seconds.

Leveraging this algorithmic definition of REM, we found that animals with a lesion to one of their two Imc nuclei displayed consistently shorter REM periods than animals without any lesion. The mean REM duration in control animals was 65 s (n = 12), while it resulted in 33 s for lesioned animals (n = 8) (Welch's two-sided t-test, P = 0.000028, t = -6.992). Interestingly, these shorter REM periods did not lead to a faster REM/NREM cycle but rather were compensated by an excess of NREM, which resulted in shorter total REM time throughout the night: control animals had an average REM of 4 hours and 40 minutes out of 9 hours of sleep, while lesioned animals showed, on average, only 1 hour and 37 minutes of REM (P = 0.000062, t = -7.204). This could be interpreted as Imc activity not being directly involved in the transition of NREM and REM but that this may instead be mediated by other midbrain nuclei or upstream areas.

6.6 Discussion

In this chapter, I have described the computational methods of analysis that I developed to process, quantify and visualize claustral LFP activity during REM, recorded by my collaborators. These analyses revealed the existence of competing dynamics between the two hemispheres of the lizard brain, a phenomenon absent during NREM.

We observed that very short synchronized bursts of firing occurred during REM sleep in the claustral neuronal population, which coincided with strong negative LFP events (Fig. 6.2). Most of the REM sleep in the claustrum is comprised of these Sharp Negative events (SNs), which underlie the characteristic beta component of the REM LFP spectral power (Fig. 6.3). I algorithmically detected SNs in each hemisphere and then bilaterally matched them (Fig. 6.7). In contrast to sharp-wave ripples (SWR) that occur during NREM, we found that SNs were tightly correlated across the two claustra. Extracting cross-correlations (Fig. 6.6) and visualizing the state space of bilateral activity (Fig. 6.9), we observed precise time delays and changes in leadership. Switches of dominance were fast, with very short periods where left- and right-leading SNs were mixed. The periods of dominance between switches displayed timescales longer than those of normal REM/NREM cycle (Fig. 6.10) and correlated with reversals of amplitude dominance (Fig. 6.7C), which suggested that the two claustra are engaged in a type of competition of winner-take-all type.

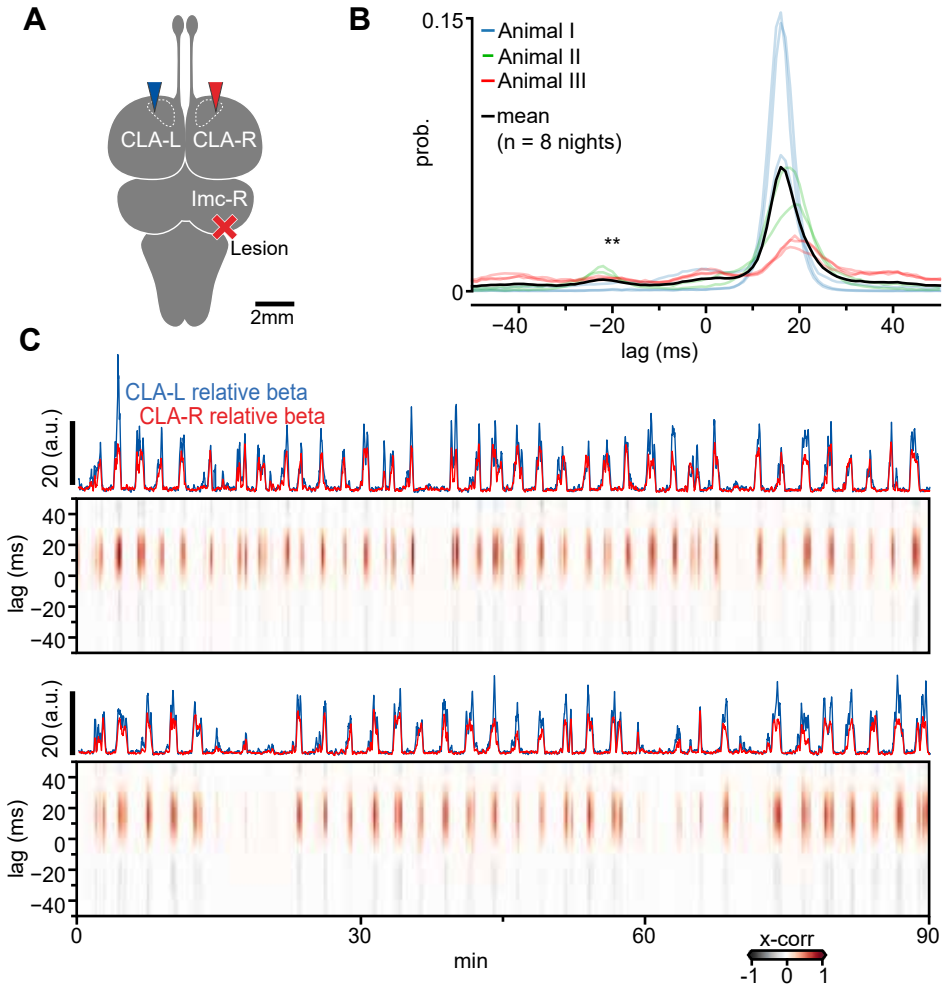

Figure 6.16: Effects of Imc unilateral lesion. **A.** Schematic of dorsal view of the *Pogona vitticeps* brain. The dashed line indicates the left and right claustra. Arrowheads indicate recording sites. The red cross indicates a lesion on the right Imc. **B.** Estimated distributions for lags showing peak correlation between right and left claustral LFP. Note the absence of peaks at -20ms indicating exclusive leadership of the left claustrum. Wilcoxon one-sided signed-rank test, W = 36, **P = 0.00390625. The average is shown in black (8 nights from 3 animals with unilateral Imc lesions). **C.** Cross-correlation of LFP first derivative between left and right claustra during 3 hours of sleep. Red and blue traces indicate beta spectral power in both claustra normalized to their minimum value. From [2].

However, this competition did not originate in the telencephalon; rather, correlations of claustral activity with recordings in the mesencephalon (Fig. 6.13) suggested the involvement of the ipsilateral GABAergic nucleus isthmi pars parvocellularis (Imc). Under unilateral lesions of Imc, cross-correlations revealed that the contralateral claustrum consistently dominated its partner (Fig. 6.16).

What appeared as a coordination of two claustra during REM sleep resulted from a competition occurring in the midbrain. Lagged correlations between Imc and claustrum suggest that the results of this competition are then propagated to the claustra with a seemingly fixed temporal delay. This propagation happened bilaterally and, possibly, via the thalamic nucleus rotundus [223, 224]. As a consequence of this delayed propagation, the left and right LFP of both claustra co-varied, reflecting a common source. The reduced amplitude of SNs on the lagging side, together with the reduced spiking, suggest a weaker gain on the path that crosses the mid-line.

The complex of isthmic nuclei is thought to be involved in bottom-up attention and gaze control in birds and is strongly linked to the optic tectum [218, 220, 222, 225]. Classic divisions of these nuclei often differentiate between the GABAergic Imc and the cholinergic–glutamatergic Ipc (isthmi parvocelullaris), and, depending on the species and labeling method, includes the cholinergic SLu (isthmi semilunaris). There are likely homologs of these avian (reptilian) isthmic nuclei found in fish [226–228], amphibians [229], and non-avian reptiles [230, 231], as well as mammals (parabigeminal nucleus).

Ipc forms a strong positive feedback loop with the ipsilateral optic tectum. The Imc is responsible for implementing broad lateral inhibition on excitatory neurons in the ipsilateral Ipc and optic tectum. Ipc, in turn, forms a strong positive feedback loop with the ipsilateral optic tectum. This circuitry has been studied in the context of two stimuli falling onto one retina at the same time. In that setup, the circuit goes through a type of competitive interaction known as winner-take-all, which means that when two stimuli are present, attention and gaze are not directed towards a weighted mean of the two stimuli, but instead, the most powerful one is selected. This competition may be traced back to the Imc because of its inhibitory nature and connectivity pattern with both the optic tectum and the Ipc ([219, 220]). Recent experimental evidence from the barn owl suggests that this competition also takes place between contralateral visual stimuli and, intriguingly, also between auditory stimuli [232]. Interestingly, the claustrum has been hypothesized to be involved in attention-related tasks, too [233, 234].

Claustral connectivity has become clearer in recent years, but our knowledge of its function remains largely incomplete. The claustrum is broadly connected, with strong connections bilaterally onto cortical regions, but there is no evidence for direct projection between the two claustra [215, 234, 234, 235]. Recent evidence has pointed towards an active state of the claustrum in NREM in mammals and reptiles [30, 31].

With the analysis presented in this chapter, we have shown that the bilateral claustra receive a common input signal from the midbrain with a fixed delay. Given

that this input triggers claustral spiking that is almost synchronous in the descending phase of SNs (Fig. 6.4 and Fig. 6.8), it is possible that the two claustra are forwarding the signal toward other brain areas. In that case, diverse brain areas, especially cortical ones, may receive a duplicated copy of the same signal with a fixed delay. Furthermore, this delay may repeatedly switch signs as the competition within the midbrain evolves. Furthermore, during NREM, claustral neurons fire during SWR that are temporally uncorrelated across hemispheres.

The practical implications of bilateral claustrum spiking for a downstream area are intriguing but unclear. On the one hand, through my modeling of spike propagation in the turtle cortex, I showed that the timing of an incoming external spike relative to a cortical spike might produce reliable gating (halting or promotion) of spiking sequences (Fig. 5.5). It is thus possible that the cortico-claustral interaction, combined with this temporal switching of a signal, might implement a form of reliable gating to coordinate cortical activity. On the other hand, the delays that my analysis uncovered fall within the timescale often associated with spike-timing-dependent plasticity rules, which might suggest a role of claustrum in propagating and consolidating memory during REM [122–124] or in triggering synaptic homeostasis within the cortex.

In summary, through the development of algorithms and computational methods of analysis, I provided quantifications of the propagation of neuronal activity across multiple sub-cortical areas in the sleeping lizard. These analyses allowed me and my collaborators to describe in unprecedented detail the electrophysiological nature of REM sleep and to provide evidence for a dynamic competitive process deep in the midbrain with effects in the telencephalon. Overall, my work provided mechanistic insights into the behavior of these areas, with potential implications for their role within reptilian sleep. Further research on the interaction between all of these areas across a variety of species will be needed to obtain a full mechanistic picture of biphasic sleep and its evolution in the vertebrate lineage.

7 Discussion and outlook

In my research, I investigated how signals propagate within the reptilian brain using computational techniques.

First, I built and explored a biologically-constrained network model to study how sequences of spikes might be triggered and sustained within the turtle visual cortex (Chapter 4). Then, I used that model to investigate how spiking activity might be reliably and flexibly routed in a recurrent cortical network (Chapter 5). Finally, I applied computational analysis techniques to quantify area-to-area interactions of sub-cortical regions of the sleeping lizard (Chapter 6).

In this chapter, I first summarize the main results related to the modeling of the turtle cortex and the routing of activity, discussing them together and in relation to existing literature, and then sketch some of the future directions of research (section 7.1). I then do the same for my work analyzing the brain activity of the sleeping lizard (section 7.2). Finally, I broaden the scope of the discussion to offer some reflections on the general framework and main lessons from this thesis (section 7.3).

7.1 Modeling propagation in the turtle cortex

In Chapters 4 and 5, I studied the impact of individual spikes in cortical networks, specifically how single spikes can trigger repeatable firing sequences and how these could be reliably controlled.

First, I developed a model based on experimental data from the turtle cortex that recapitulated experimentally observed spiking sequences [37]. This model made quantitative predictions about the properties of sequences and predicted their presence even under in vivo levels of spontaneous activity. In addition, the results of these simulations suggested that the strongest connections, even if rare, are the main backbone supporting propagation of spiking sequences. On the other hand, the abundant weak connections regulate the reliability of propagation.

By analyzing how activity propagates in the model, I found that spiking sequences might be broken down into several sub-sequences, with each sub-sequence representing a sub-network composed of followers linked by strong connections. The sparsity of these strong connections enabled the selective activation or silencing of different sub-networks through small external inputs. By studying interactions of multiple sequences, I found that although spiking sequences rarely overlap, they often interact in the form of lateral inhibition or excitation. As a result, the final path of propagation of a single spike in the cortex is influenced by the activation of

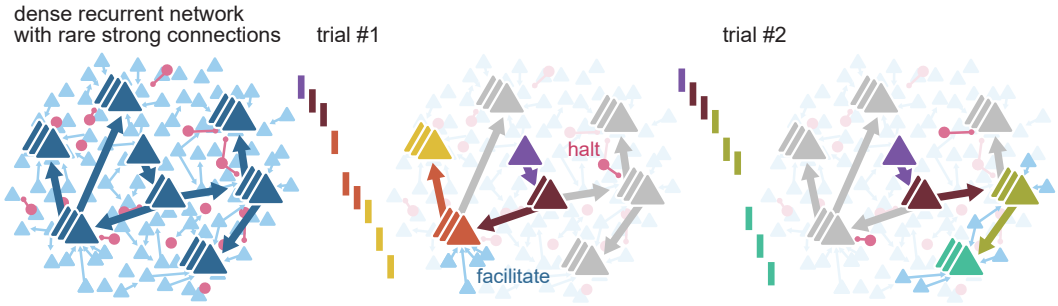

Figure 7.1: Routing using sparse strong connectivity. The routing mechanism in the cortex might rely on sparse strong connectivity. As activity spreads, spikes tend to travel through the strongest connections, which provide reliability. Because these connections are rare, they create sub-networks with gate neurons that can be leveraged to flexibly facilitate or halt propagation. The halting of propagation may manifest as partial failures in the expected sequence of spikes. The context formed by the ongoing activity in the recurrent network as well as sparse external inputs defines the final route for each spike in the network. From [1].

a small percentage of other neurons in the network, suggesting a highly specific but combinatorially powerful mechanism. In summary, I found that few external inputs and recurrent interactions could reliably control the path of propagation of a single spike in the cortex, a process that I termed "routing" (Fig. 7.1).

7.1.1 Routing and cortical function.

Sequences of action potentials have been observed in the cortices of mammals and correlated with memory or spatial tasks [126, 149–152]. They have also been observed within the stimuli-evoked and spontaneous activity in sensory cortices of rodents [148, 153–156]. The model that I introduced in Chapter 4 predicts that sequences might be triggered by single spikes in the in vivo turtle cortex. However, the functional role of spiking cortical sequences is still largely unknown.

Sequences have been hypothesized to be part of a packet-based scheme for coding information and communication within the cortex [156]. Within this framework, sequences would form the first part of each packet, presumably using precise spike times to encode additional information to the main body of the packet. A different theory has proposed that sequences are spontaneously generated by the brain and that each element in the sequence can be plastically linked to different items [149]. In this way, sequences may act as a kind of "index" that connects multiple items that are distributed over several cortical areas in order to compose a full multi-faceted concept.

Through my modeling research, I came to the conclusion that sequences in the turtle cortex are likely the expression of strong synaptic connections (Chapter 4), which enable the reliable routing of activity (Chapter 5). This reliable and flexible propagation of activity essentially maps one pattern of neuronal activity onto another,

enabling the implementation of computations. For instance, theoretical studies using firing rates and precise timing have shown that flexible gating can be used to implement logic gates, switches, and even memory units [12, 13, 15]. Thus, we might ask what computations may be at play in the visual cortex of turtles and how these are related to routing.

I introduced the concepts of internal and external frames to address neuroscientific questions in section 1.2. My work so far stems from an internal frame that studies brain activity with respect to its internal mechanisms and dynamics. In this section, I describe how the lessons from that internal frame can be linked to an external one that aims to understand how brain activity relates to its environment. In particular, I will propose the hypothesis that the reliable routing of sequential cortical activity using single spikes might implement the reliable tuning of cortical neurons to specific stimulus properties.

Routing and single-cell tuning

The turtle dorsal cortex is known to respond to visual stimulation with wave-like patterns of activity [114, 176]. Contrary to the mammalian cortex, the turtle visual cortex lacks certain organizational features, such as retinotopy, and most neurons display receptive fields that cover the entire visual field [26]. However, turtle cortical neurons still display orientation selectivity with a wide range of selectivity indices [26]. Since my model suggests that excitatory neurons are primarily driven via single strong connections (Fig. 4.11), and these neurons possibly display orientation tuning, it follows that like-tuned neurons might share strong connections.

Interestingly, evidence from the primary visual cortex of mice suggests that very strong connections might be key in defining orientation tuning [182]. In a series of experiments in vivo, visual stimuli were found to cause a sustained depolarization of L2/3 pyramidal neurons that happened independently of the orientation of the stimulus. Only when the stimuli matched a specific orientation did those depolarized neurons fire action potentials. Using a combination of in vivo calcium imaging and slice multi whole-cell patch-clamp, the authors found that the broadly-tuned depolarization was caused via a majority of weak connections, while the orientation-selective activations were triggered by sparse strong connections. These findings are consistent with predictions from my model where weak connectivity can amplify the external input, increasing the resting potential of neurons (Fig. 4.4CD) and, in doing so, alter the effectiveness of strong connections and modulate follower reliability (Fig. 4.12). The predictions from my model regarding the roles of weak and strong connectivity in routing may thus have a counterpart in the voltage dynamics of neurons and their relationship to external stimuli.

Stimulus "tuning" is a concept that takes an external view at describing neuronal activity (see section 1.2). As a result, it only addresses specific neuronal activations. If we say that a neuron is tuned to a particular orientation, what can we tell about the presynaptic sub-network that leads it to spike or about the postsynaptic sub-network that will receive its outputs? Tuning, in other words, addresses only a portion of the

recurrent pipeline of transformations that link stimuli and behavior. By describing the routing of spikes, I provide a mechanistic, internal view that is complementary to that of tuning and that can address these causal questions. Understanding routing will provide us with a way to examine how tuning is implemented, bridging both the internal and external frame to visual processing in neuronal systems.

Computation through lateral interactions

My modeling results go beyond the focus of how single cells become activated and include the interaction between multiple cells. Indeed, parallel running sequences have been observed experimentally in the ex vivo mouse auditory cortex [154], and the in vivo rat somatosensory cortex [181]. Experiments in the ex vivo turtle showed that the coactivation of multiple trigger neurons in the turtle cortex ex vivo does not simply add up to the total number of followers [37]. My model recapitulates this observation (Fig. 5.6D) and provides a mechanistic explanation in terms of lateral excitatory and inhibitory interactions between sub-networks of strong connections (Fig. 5.8). Furthermore, by exploring the effect of timing on the gating of single sub-networks in the model, I found that the temporal window for the effective halting of activity (approx. 100 ms, Fig. 5.5A) might be about two times longer than the standard deviation of follower spike times (approx. 50 ms, Fig. 4.8), meaning that sequences display a relative temporal precision higher than that required for controlled routing. In other words, my model suggests that parallel running sequences not only frequently interact but that these interactions are reliable.

The reliable interaction between multiple sequences results, in the model, in the detection of late-activating combination-specific followers (Fig. 5.7C). These followers activate reliably and exclusively if a particular pair of trigger neurons are in a particular state; for example, one trigger must activate, and the other must remain silent. If we consider that trigger and follower neurons may display some form of tuning to stimulus properties, then the tuning of combination-specific followers may be expressed as the result of a logical operation on the tuning of their corresponding triggers. From the external perspective of tuning in sensory cortices, lateral excitation or inhibition between multiple sub-networks are thus forms of computations on the properties of stimuli.

In my model, lateral inhibition between sub-networks is the most common form of lateral interactions (Fig. 5.5F, Fig. 5.8B). Indeed, lateral inhibition in mouse olfactory cortex has been shown to be triggered by early responding pyramidal cells, implementing concentration-invariant coding [236]. Similarly, single L2/3 pyramidal neurons in the mouse primary visual cortex frequently trigger lateral inhibition onto other neurons with overlapping tuning. This tuning-specific lateral inhibition has been termed a form of "feature competition" between stimuli [93].

On the other hand, my model also suggests that multiple sub-networks may introduce highly-specific, sparse, but powerful lateral excitation (Fig. 5.8B). These interactions might correspond to the rarer case of "feature amplification" [93]. Alternatively, one might interpret these reliable excitatory interactions as a mechanism

to positively link together multiple stimulus properties. The binding of multiple stimuli to compose a single perceptual item is known as the "binding problem" and has been the object of theories of cortical operation, including some based on synfire chains [16, 137] (see section 3.3.1). The activation of a combination-specific follower in my model may correspond to an instance of binding of multiple stimuli. Compared to the large populations and synchronous firing involved in synfire chains, the routing on sparse strong connections that I propose may require the reliable activation of only a few neurons to solve the binding problem. Note that combination-specific followers might be defined by the activation of very few trigger cells, while the experimental estimates of neuronal density that constrained my model described a rather dense large population of excitatory neurons in the turtle dorsal cortex (Fig. 4.2A). Consequently, the circuitry of the turtle dorsal cortex may display a very high combinatorial capacity with respect to the space of stimulus properties.

In conclusion, by studying the properties of spiking sequences in a computational model, I found that sequences are likely the expression of reliable routing of activity. This reliable activation of neurons of a sensory area of the cortex likely relates to the properties of the stimulus presented, i.e., the tuning of these neurons. The recurrent interactions that lead to the reliable activation or silencing of specific neurons thus might implement tuning computations such as feature competition, amplification, or stimulus binding. Thus, understanding the mechanisms by which spikes are routed provides us with an alternative perspective to describe cortical computations and to understand how they are implemented. By starting within the internal frame of local interactions of cortical activity, we may reach insights into the phenomena of the external frame of visual processing.

7.1.2 Outlook

Looking forward, there are several promising directions for further research on the mechanisms supporting reliable propagation of activity in cortical networks. Here, I briefly describe two such directions, both stemming from limitations of the current model: extension of these results to other animal species and switching from a static network to a plastic one.

7.1.2.1 Beyond turtles

As I described in section 1.3, a comparative approach to neuroscience requires investigating diverse animal species. Consequently, it becomes important to explore how these results extend to other species beyond the turtle. Indeed, the model that I presented here was tightly constrained by experimental measurements from the turtle cortex (Fig. 4.2 and Fig. 4.3), which might diverge from the visual cortex of other species in multiple aspects. Two aspects that may play a particularly pivotal role in my conclusions are features of connectivity and operating regimes.

Several of the fundamental elements of the connectivity in the model introduced in this thesis for the turtle cortex are present in other animals and brain areas. For instance, long-tailed distributions of synaptic strengths have been reported for the cortices of many mammalian species [68, 74, 101, 182]. There are also spatial features of connectivity shared between rodent and turtle cortices. For instance, the rodent auditory cortex is known to display Gaussian profiles of distance-dependent connectivity such as those used in this thesis (Fig. 4.2F) [237], and sensory afferent fibers lay a gradient of en-passant synapses in the three-layered rodent piriform cortex, as it does in the three-layered turtle dorsal cortex (not modeled here) [50]. However, it is difficult to make detailed comparisons between different cortical areas and species because the available data is limited, especially for reptiles.

Much more is known for the detailed connectivity in rodents [68, 74], and it is possible that some features that I did not explore in my model might play a key role in the reliability of sequence propagation in other species. For instance, the particular thickness and length of the tail of the distribution of strong connections, i.e., how strong and common are the strongest connections, is of particular interest since my model showed that truncating this tail prevented follower activations (Fig. 4.12). However, the effect of an EPSP caused by a synaptic activation depends not only on the synaptic strength but needs to be considered in conjunction with the passive properties and the expected conductance state of the postsynaptic neurons. Thus, species-specific electrophysiological properties of excitatory neurons need to be taken into account.

Another key element that I explored in this model was the robustness of sequences to increased baseline firing rates, where I showed that even under average turtle in vivo levels of activity, sequences are present and detectable. Average firing rates in the in vivo turtle are estimated in the range of [0.02, 0.09] spk/s. Although firing rates in mammals are typically considered much higher, these are often reported under some form of stimulus presentation.

When looking at spontaneous activity, mean firing rates below 1 spk/s have been observed in multiple mammalian neocortical areas [238, 239]. Indeed, this quiescent default state has inspired many theories of the sparse-firing operation of the cortex [172]. Using multi-area neuropixel data, a recent study estimates that up to 40% of neurons within a column of in vivo mouse sensory cortex might display mean firing rates below 1 spk/s [240]. Nonetheless, the turtle cortex displays very strongly adaptive firing [26], which may not extend to other animals and areas, so firing-rate operating regimes might still be highly relevant.

As part of my modeling work, I explored a wider range of firing rate regimes (Fig. 4.9). I concluded that sequences might be present under mean firing rates up to two orders of magnitude higher than the turtle cortex but that these sequences may not be detectable. Indeed, spiking sequences have been reported in multiple mammalian cortical structures, but their discovery typically requires some form of external correlates, such as decoding the location of the animal, presenting a visual stimulus, or triggering a memory task [126, 149, 150, 152, 153, 156]. Repeat-

able sequences have been found within the spontaneous activity in the cortices of anesthetized and awake rats right after the transition from down-states (when firing rates are low) to up-states (when firing becomes sustained), but they could not be detected within the up-state [155]. Indeed, the unsupervised discovery of sequences in spontaneous activity is still a largely unresolved problem, but one that could radically change our understanding of cortical communication [11, 154, 155, 241, 242].

In summary, to explore the extent to which these differences affect sequential propagation over a sparse network of strong synaptic connections, one would need to develop a different species-specific model that includes known features of connectivity and consider the operating regime of the network. Having constrained my model through data and literature on the turtle cortex, my work lays the ground for other modeling attempts, which, when compared, may highlight potential commonalities and differences in the forms of communication across cortices as they diverged evolutionarily.

7.1.2.2 Plasticity

Another area of interest is the effect of plasticity mechanisms on sequences of spikes (Fig. 7.2). Specifically, mechanisms of activity-dependent synaptic plasticity, which refers to the changes in synaptic strengths as a function of the patterns of pre and postsynaptic activity, may interact in complex ways with the reliable routing of spikes that I studied in this thesis. Indeed, the brains of animals are subject to changes through experience, which are thought to affect the structure of their connectivity [121]. Since the cortex that I studied is a sensory one, it is likely to be exposed to structured activity throughout the life of the animal. Furthermore, activity-dependent plasticity mechanisms may act even at the stage of development, resulting in a network configuration that is shaped and prepared by spontaneous activity [8, 243]. Therefore, investigating how plasticity rules might shape the emergence and unfolding of sequential activity in recurrent network models is highly relevant to understanding cortical development and function. Such a line of research would likely involve investigating the effects of different plasticity rules on the formation and propagation of sequences from single-spike triggers.

We might distinguish between two different spatial scales at which plasticity could affect sequential propagation: coarse scales, affecting the statistical distribution of different connectivity parameters, and fine scales, affecting the detailed connectivity motifs. Compounded on the question of spatial scale, we might want also to consider the differences between spontaneous (uncorrelated) activity or stimulus-driven (structured) inputs.

On the one hand, plastic mechanisms may induce the general statistics of synaptic connectivity that my model has shown are sufficient to trigger sequences (Fig. 7.2, left). Most importantly, long-tailed distributions of synaptic strengths were key in enabling reliable propagation (Fig. 4.12). In addition, the rare and strong synapses were distributed within the network and not concentrated on a few hub neurons so

that multiple spike transfers could happen sequentially (Fig. 4.10). At this scale, one might investigate how developmental processes lead to a network configuration matching these statistical properties. In addition, homeostatic processes may need to act to maintain that configuration stable throughout the life of the animal, even under the presence of spontaneous noisy activity.

On the other hand, plasticity may alter the fine network motifs without changing the statistics (Fig. 7.2, right). Since strong connections are rare, adding or removing a single one can drastically affect the stability of follower neurons. Repetitive routing may therefore lead to plastic changes that turn dynamic gating (halting or promoting propagation, Fig. 7.1) into a fixed feature of the network. Furthermore, plasticity might tweak weak excitatory and inhibitory connections, which my model has shown play a key role in establishing the final path of propagation through lateral interactions between sequences (Fig. 5.8). At this scale, one might investigate how robust individual sequences are to plastic changes when we consider the spontaneous, uncorrelated activity. Another interesting question is how these sequences might be shaped by structured input, which would reflect the capacity of the system to implement new computations (mappings from inputs to outputs) under a learning paradigm.

Theoretical studies have shown that combining spike-timing-dependent plasticity with heterosynaptic competition can lead to a binarization of the distribution of synaptic strength, creating a single unary sequence of strong connections within the network [140]. Interestingly, this approach relied on random inputs and may hint at the mechanisms that establish coarse-scale statistics. Other modeling work has shown that small discrete-time networks of binary neurons can also produce similar connectivity [133, 141]. However, in the turtle cortex, about 80% of trigger neurons showed multiple followers [37], which suggests that strong connections are widely distributed and makes it unlikely that the network forms a single unbroken unary sequence. In addition, these models typically focus on capturing the generation of neural sequences in the HVC area of the songbird system [165], which is known to crystallize and remain fixed with respect to the main courtship song throughout the adult life of the animal, and thus the models rarely address questions of changing connectivity through learning.

It would also be interesting to avoid the full binarization of connections, given that my model suggests that weak connections play a role in establishing sequence-to-sequence interactions, and instead study the effects of a continuous distribution where changes in weak connectivity may affect sequence propagation. Furthermore, while useful to illustrate simple mechanisms, binary neuronal models and discrete-time networks might incur an oversimplification of the network dynamics and prevent quantitative experimental predictions. Indeed, some of the key insights in my modeling work come from grounding my model on the available biological data following a bottom-up approach to modeling (section 1.4).

Other theoretical work has shown that plastic rules can ingrain sequential activity into more realistic network models, but they almost exclusively rely on the use of

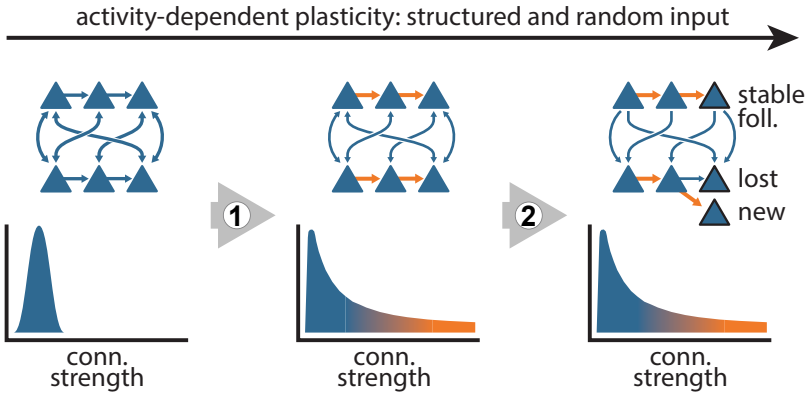

Figure 7.2: Plastic effects on sequential propagation. Effects of plasticity might be highly relevant for sequential propagation at coarse and fine spatial scales. **(1).** Plasticity may shape the long-tail distributions of synaptic strengths measured in the adult turtle cortex. My model showed that these distributions and random connectivity are enough to produce sequential activity. We might ask what plastic and homeostatic mechanisms give rise to these distributions in the first place and whether they change throughout the life of the animal. **(2).** Plasticity may shape the paths of strong connections without changing the general network statistics. Changing the strength of one or a few strong connections may alter the follower composition of a sequence and reshape the path that activity takes. In addition, changes to the weak and dense excitatory connectivity and inhibitory connections may alter the interactions between multiple sequences favoring competition over cooperation or vice versa.

sequential inputs in order to create these structures [174,244,245]. Although thalamic inputs to the turtle cortex display a directional gradient on connectivity [50] and axonal projections of principal neurons are also directionally biased [28], sequences of followers do not necessarily align with this axis [37]. It seems unlikely that external input into the dorsal cortex is consistently sequential in nature through the development and life of the animal. Thus, one might want to combine some of these approaches with those of self-organization and homeostatic mechanisms that produce the right coarse-scale statistics even without structured input.

In conclusion, understanding the role of activity-dependent plasticity in generating and modifying sequences of spikes might be key to understanding cortical development and function. Plasticity may affect sequential propagation in different spatial scales, from whole distributions to fine motifs of connectivity. We may need to consider the robustness with which plastic sequences tolerate noisy activity and the flexibility with which they can adapt to changing structured, stimulus-driven inputs. Existing plasticity models of self-organization, homeostasis, and stimulus-driven learning might be combined with details from the turtle cortex to understand how sparse networks of strong connections arise and change with experience. These structural changes will, in turn, modify the capacity of the system to implement computations through reliable and flexible routing of activity.

7.2 Analyzing propagation in the lizard brain

In Chapter 6, I introduced computational methods that helped my experimental collaborators and me investigate multi-area coordination of neuronal activity during the rapid-eye-movement (REM) phase of sleep in the *Pogona vitticeps* lizard.

By analyzing large night-long recordings of local field potentials (LFP) from the two hemispheres, we observed that bilateral claustra, a sub-cortical region much more accessible in lizards than in rodents, displayed tight temporal coordination during REM sleep. Using lagged cross-correlations on the derivative of the signal as well as diverse algorithms to detect and match bilateral signals, we observed that most of the LFP signal during REM was composed of single Sharp Negative events (SNs) that were coordinated in time across the two sides. Importantly, this coordination was lagged by 20 ms, and the leading side changed throughout the recording, with dynamics slower than those of the REM/NREM cycle. Additionally, we showed that bilateral sharp-wave-ripples (SWR) happening during NREM were not coordinated [30]. Thus, the two claustra regularly switch between acting independently and entering a mode of shared activity. During that mode of coordination, the activity in the two claustra reflects a competitive process that switches which SNs lead and display a stronger amplitude.

This competition did not occur directly between the two claustra or the telencephalic hemispheres but rather in the midbrain. I then extended these methods to new recordings of a GABAergic nucleus in the midbrain, the Imc in reptiles and birds, a homolog to the parabigeminal nucleus in mammals [226–231, 246]. By extracting correlations and detecting SNs in healthy animals and in animals with Imc lesions, we established that this nucleus shows highly correlated activity preceding its ipsilateral claustral SNs and that it is necessary for establishing the lag leadership with the contralateral claustrum.

Our findings contribute to a comparative approach to understanding the evolution of sleep. They also emphasize the importance of exploring common behaviors in diverse and experimentally advantageous animal models to comprehend brain function and the diverse mechanisms that implement it. The fact that biphasic sleep has been observed in a diverse range of animal lineages, including birds, reptiles, and fish, suggests that it may have evolved early in the vertebrate lineage [24, 30, 32, 127], and thus may have important implications for understanding the origins and functions of sleep. It will be interesting to see how future research on diverse animal lineages may help to further clarify the evolution of sleep [17, 24].

Overall, the computational methods that I developed provided important insights into the nature of sleep in sub-cortical areas of the lizard brain. These methods helped us describe mechanisms of propagation of neuronal activity across hemispheres during sleep and quantify winner-take-all dynamics that involve tight temporal coordination of a few milliseconds, developing over periods of multiple hours. Further development in computational methods to dissect the functions and mechanisms

underlying biphasic sleep may provide important new insights into the functions and evolution of sleep more broadly.

7.2.1 Outlook

These experiments and the analysis of the data they produced revealed striking coordination across multiple bilateral, sub-cortical areas, with quantifiable dynamics that describe a competitive process. However, these also opened several questions where computational approaches might shine a light or guide future experiments.

As I introduced in section 1.4, one of the goals of mechanistic computational modeling is to integrate empirical observations and explore the set of biological solutions that are consistent with these observations in order to guide future experiments. Importantly, computational models allow us to explore plausible biological mechanisms that are of difficult experimental access or for which we have mostly indirect measurements. I illustrated the value of this approach in my modeling work on the turtle dorsal cortex (section 4.1.1), and I propose that a similar approach might be useful in exploring the circuitry underlying bilateral coordination during sleep. In particular, there are two open questions that might be addressed through modeling (schematized in Fig. 7.3A): the circuitry underlying the hemispheric competitive process during REM and its interactions with REM/NREM alternations. These are closely related questions since our analysis shows that unilateral lesions of Imc not only ablate dominance from the lesioned hemisphere but also affect the duration of REM and NREM periods. How is Imc connected, and what interactions with other nuclei affect dominance and the REM/NREM alternation? We might gain insights into these questions by leveraging methods of dynamical systems analysis constrained by evidence from anatomical tracing.

Previous and new tracing by my collaborator Dr. Lorenz Fenk on *Pogona vitticeps* [30], as well as other studies in mice [247], indicate that the claustrum does not receive direct input from the Imc (homolog to the parabigeminal nucleus in mammals), an area whose integrity is however required for establishing REM dominance. It also suggests that Imc likely does not project to its contralateral partner. Anatomical and functional evidence from birds and other reptiles indicate that the Imc is part of an interconnected network with the cholinergic nucleus isthmi pars parvocellularis (Ipc) and the optic tectum (homolog to the superior colliculus in mammals), which is a complex, layered structure containing excitatory and inhibitory neurons [225, 232, 248, 249]. This sub-network has been shown to mediate winner-take-all dynamics when multiple visual stimuli are presented to the same hemisphere in birds [220]. Additionally, this sub-network may present contralateral projections, at least in rodents [250].

We might leverage the tools from dynamical systems analysis in order to study how the interaction between the inhibitory and excitatory populations across all of these nuclei may result in the dynamics of dominance that I have described (section 6.4). For example, this approach has been applied to studying the relationship between

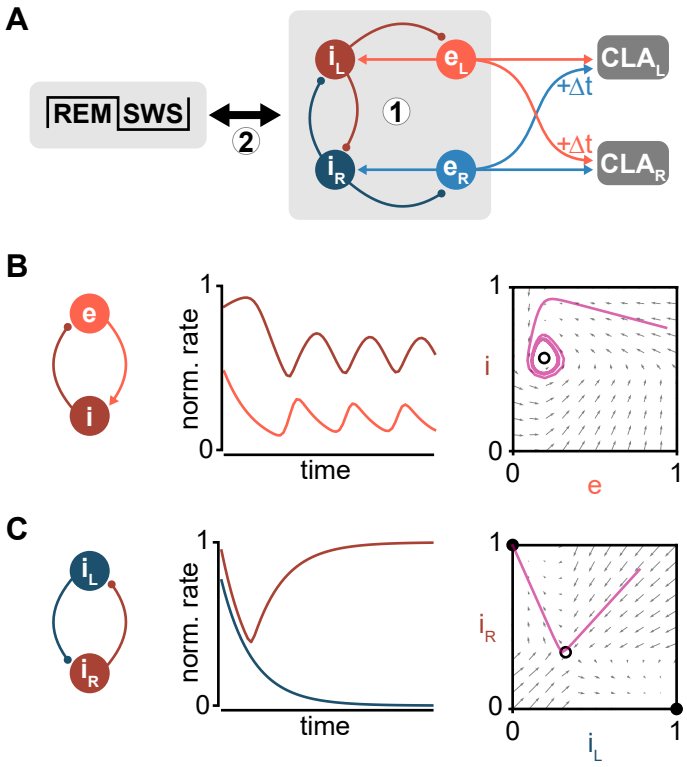

Figure 7.3: Modeling inter-hemispheric competition. **A.** Empirical questions that a data-constrained model could inform: **(1)** What connectivity between the multiple midbrain and hindbrain nuclei related to Imc might lead to the patterns of dominance observed in claustrum? **(2)** How does dominance interact with the REM and NREM alternation? The populations of neurons (circles) might be distributed among multiple nuclei (including Imc, Ipc, and Optic Tectum) for which connectivity is only partially known. e/i: excitation/inhibition. L/R: Left/Right. CLA: Claustrum. **B.** Connectivity motif implementing oscillations. Left: motif schematic of an excitatory population that triggers feedback inhibition. Middle: simulated firing rate of both populations. Right: phase-plane plot of the motif with an example trajectory (same as Middle, magenta). Arrows indicate the gradient of population activity. An unstable fixed point (open circle) attracts the trajectory and creates a cycle that results in the regular oscillation of firing rates. **C.** Connectivity motif implementing competition. Left: motif schematic of two mutually-inhibiting populations. Middle: simulated firing rate of both populations. Right: phase-plane plot of the motif with an example trajectory (same as Middle, magenta). An unstable saddle point (open circle) attracts the trajectory and then redirects it towards one of two opposing stable fixed points (closed circles).

rodent hippocampus and neocortex during NREM, revealing signatures of two different modes of excitable dynamics [120]. That work leveraged a Wilson–Cowan-like model with adaptive properties to study the effect of sharp-wave ripples and used distributions of network states extracted from biological data to constrain the model parameters that best explained the dynamics. In the case of claustral dominance, two dynamical motifs seem relevant (Fig. 7.3B-C): oscillators through feedback inhibition and competition through mutual inhibition.

Fast oscillatory dynamics might underlie the regular SNs at a 40-millisecond period while being nested or driven by a slower oscillator that alternates REM and NREM activity at an 80-second period. Competition through mutual inhibition, on the other hand, can create stable dominance of one hemisphere over the other, with adaptation or input noise enabling the leading side to switch by lowering the energy barrier [251].

The particular time scales of these motifs and how they interact with one another in a full circuit could be constrained by measuring the statistics from the data, such as the distribution of dominance durations (Fig. 6.9D). The result of such constraining might shed light on the patterns of connectivity and physiological properties that most likely underlie the phenomena that we found.

In conclusion, the use of computational modeling might be a valuable next step in exploring the complex circuitry and biological mechanisms that underlie REM activity in *Pogona vitticeps* and other species. By leveraging dynamical systems analysis, statistical measurements from data, and anatomical evidence across related species, we might identify plausible mechanisms and dynamics that underlie the bilateral coordination of sub-cortical areas during sleep. By integrating empirical observations and exploring the set of biological solutions that are consistent with these observations in a formal framework, computational modeling can guide future experiments and help us better understand the complex dynamics of the sleeping brain.

7.3 Final remarks

In conclusion, this thesis has investigated some of the fundamental processes of propagation of neuronal activity in the reptilian brain using computational modeling and analysis of electrophysiological data. The findings presented here provide insights into the mechanisms that underlie information processing and communication within the reptilian brain.

I found that the propagation of single spikes in the turtle cortex is influenced by the biophysical properties of the network, with a key role in the rare but powerful connections present in multiple animal systems. This sparse network of strong connections enables the reliable routing of signals as small as a single spike. I have also shown how the spiking state of very few other neurons in the network, together with external inputs, determines the eventual path for each spike, offering insights into the context-dependent nature of the propagation of neuronal activity.

Together with my experimental collaborators, I also found that brain activity in the lizard brain during sleep involves inter-area coordination. This coordination happened between neuronal populations in the two hemispheres, across multiple sub-cortical regions, claustrum, and the isthmic nuclei at least. The cross-area propagation of activity involved tight temporal delays, in the scale of milliseconds, nested within the minute-long oscillatory REM/NREM rhythm and influenced by even longer winner-take-all dynamics.

One of the key contributions of this thesis is the identification of the factors that influence the propagation of neuronal activity within the cortex and across sub-cortical areas. These results provide insights into how neural circuits are organized and how they interact together even in the absence of external stimuli, thus directly addressing fundamental questions about how neurons communicate. This work highlights the importance of taking an internal frame in neuroscience research, focusing on the intrinsic and spontaneous patterns of neural activity rather than solely in relation to the external environment. By examining how neurons and neural networks communicate with each other, we can gain a better understanding of the core causal mechanisms through which the brain processes information, providing an alternative, complementary perspective of the relationship between brain activity, stimuli, and behavior.

Reptiles are remarkable models for neuroscientific research, with their diverse nervous systems, unique physiological characteristics, and interesting behaviors providing valuable study subjects for understanding the operation of the brain. By studying reptiles in greater depth, we can gain insights into the evolution and function of the nervous system and how it has been adapted and modified throughout evolution. While this study has specifically focused on turtles and lizards, the methods and insights gained from this research might be applied to, or at least guide further research in, other animal models. The computational techniques of bottom-up modeling and data analysis described in this thesis are highly adaptable

and might be used to investigate a range of questions regarding the propagation of neuronal activity.

In conclusion, my research has contributed theoretical understanding, predictions, and quantitative evidence on how neuronal activity propagates in the reptilian brain, providing new insights into the workings of this fascinating machinery. It is my hope that my findings and methods will inspire further research in this area and contribute to a deeper understanding of the complex mechanisms that underlie neural activity in reptiles and other species.

A Appendix: Multiple network instantiations

In sections 5.2, 5.3, and 5.4 of Chapter 5, I focus on a single instantiation of a model network to explore the routing of activity in the turtle cortex. This was for simplicity of illustration and analysis. Nonetheless, to verify the results remained consistent and general, I repeated these simulations and analyses using many instantiations of the random connectivity model. Note that, depending on the explored question, I adjusted the number of different networks due to the high computational cost of each one of these simulations (approximately 30 minutes per simulation) and of the analyses that followed them. This appendix includes detailed plots of multiple instantiations of the network and indicates how they were summarized in the corresponding section of my main results.

Fig. A.1 contains examples of clustering followers by their activity, revealing "sub-networks" of followers that are most likely to be connected through strong connections (see Fig. 5.2C). These are just a few examples of the original 6000 simulations that were generated for studying sequence generation in Chapter 4. Note the number of followers spans two orders of magnitude, yet clustering consistently reveals strongly connected sub-networks. These, as well as the rest of the original 6000 simulations, were analyzed to quantify their sub-network connectivities and entropy of activations (Fig. 5.3), as described in the relevant section 5.2.

Fig. A.2 and Fig. A.3 contain the results of simulations from multiple instantiations for the study of the control of sequences through an external input (see Fig. 5.5AB). The average of these maps of fold change in activation is included in Fig. 5.5CDGH and described in the relevant section 5.3.

Fig. A.4 contains the results of simulations from multiple instantiations for the study of interactions of multiple trigger neurons (see Fig. 5.6DE). Due to the computationally heavy simulation time, I reduced the number of pairs of neurons tested per network to 800 (from the original 2000), but I instantiated a total of 50 different networks. For the results, I pooled followers across all of the resulting 4000 pairs. The conclusions of these simulations are a perfect match to those described in the results section 5.4.

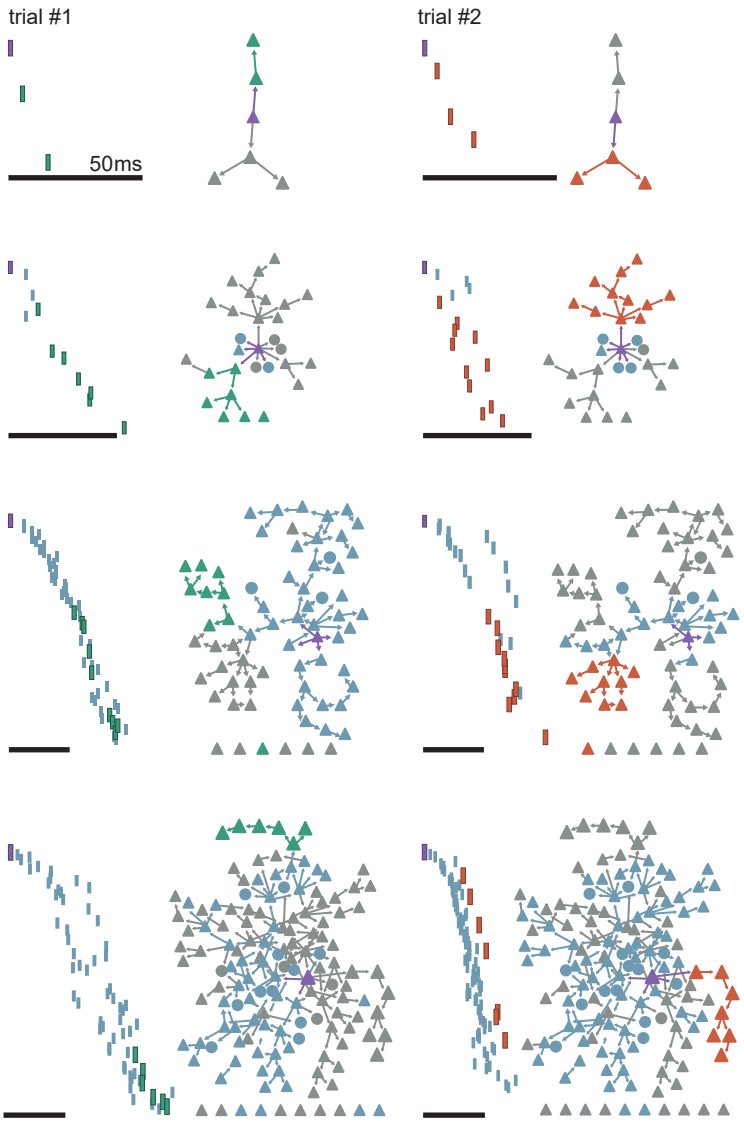

Figure A.1: Clustering of followers in multiple networks. Columns show two trials of different example simulations of different networks where the trigger neuron activated different sequence sizes (5, 28, 68, and 145 followers, top to bottom). Only the top 5% connections are shown. Followers were clustered by activity. Green and orange: example active sub-networks. Blue: other active sub-networks. Gray: inactive followers. From [1].

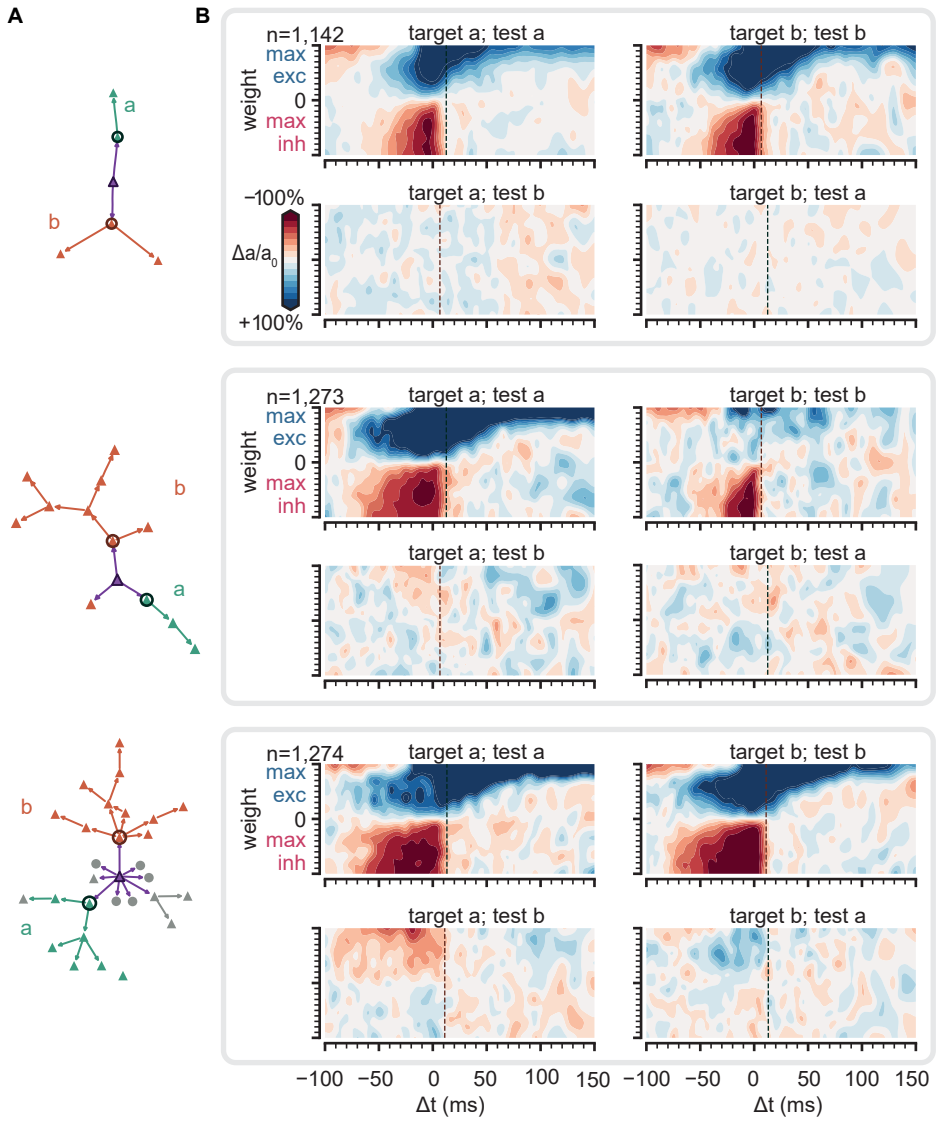

Figure A.2: Gating through external input. **A.** Graphs of strong connections between followers for networks where the trigger neuron had an increasing number of followers. Due to space, this figure continued in Fig. A.3. For each instantiation, two example sub-networks are colored and labeled a and b. Circles indicate sub-network gates. Followers in other sub-networks are in gray. **B.** Map of the fold change in activation ($\Delta a/a_0$, %) on sub-network a or b (test) as a function of external input strength and timing onto the gate of sub-network a or b (target). From [1].

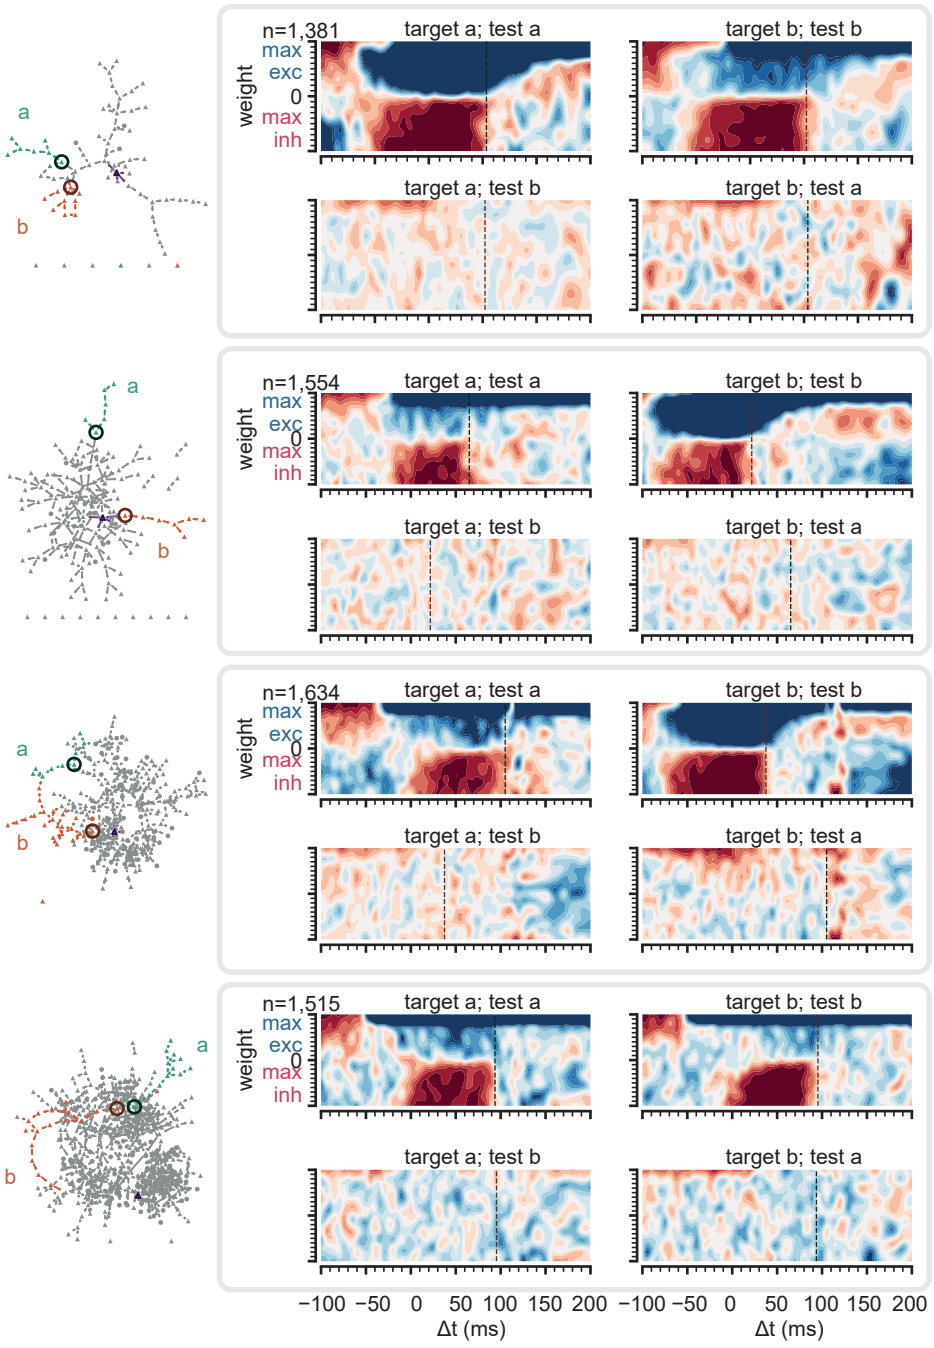

Figure A.3: Gating through external input (cont.). Continuation of Fig. A.2.

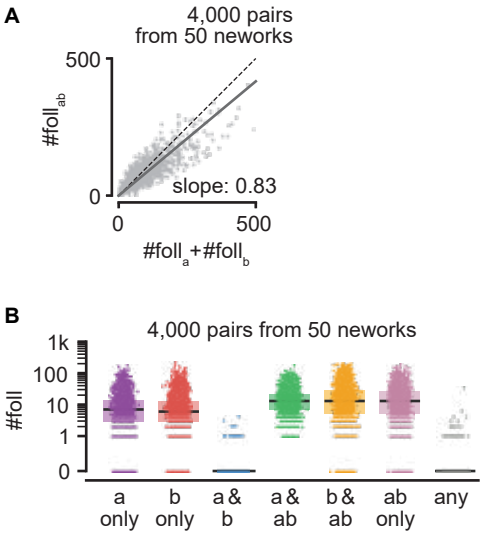

Figure A.4: Interactions of multiple trigger neurons. **A.** The number of model followers detected under the coactivation of two triggers in relation to the sum of followers detected for each trigger activated in isolation. Same as Fig. 5.6D but with follower pairs pooled across 50 randomly picked networks (800 pairs per network). **B.** Count of detected followers for each trigger category after pooling all simulations. Same as Fig. 5.6E but with follower pairs pooled as in A. From [1].

Bibliography

[1] J. L. Riquelme, M. Hemberger, G. Laurent, and J. Gjorgjieva, "Single spikes drive sequential propagation and routing of activity in a cortical network," *eLife*, vol. 12, p. e79928, Feb. 2023, publisher: eLife Sciences Publications, Ltd.

[2] L. A. Fenk, J. L. Riquelme, and G. Laurent, "Interhemispheric competition during sleep," *Nature*, Feb. 2023.

[3] R. Brette, "Philosophy of the Spike: Rate-Based vs. Spike-Based Theories of the Brain," *Frontiers in Systems Neuroscience*, vol. 9, Nov. 2015.

[4] D. H. Hubel and T. N. Wiesel, "Receptive fields of single neurones in the cat's striate cortex," *The Journal of Physiology*, vol. 148, no. 3, pp. 574–591, Oct. 1959.

[5] J. D. Roitman and M. N. Shadlen, "Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination Reaction Time Task," *Journal of Neuroscience*, vol. 22, no. 21, pp. 9475–9489, Nov. 2002, publisher: Society for Neuroscience Section: ARTICLE.

[6] R. Brette, "Is coding a relevant metaphor for the brain?" *Behavioral and Brain Sciences*, vol. 42, p. e215, 2019, publisher: Cambridge University Press.

[7] E. M. Izhikevich, "Polychronization: computation with spikes," *Neural Computation*, vol. 18, no. 2, pp. 245–282, Feb. 2006.

[8] J. H. Kirchner and J. Gjorgjieva, "Emergence of local and global synaptic organization on cortical dendrites," *Nature Communications*, vol. 12, no. 1, p. 4005, Dec. 2021.

[9] G. Buzsáki, *The brain from inside out*. New York, NY: Oxford University Press, 2019.

[10] D. Kobak, W. Brendel, C. Constantinidis, C. E. Feierstein, A. Kepecs, Z. F. Mainen *et al.*, "Demixed principal component analysis of neural population data," *eLife*, vol. 5, p. e10989, Apr. 2016, publisher: eLife Sciences Publications, Ltd.

[11] A. H. Williams, B. Poole, N. Maheswaranathan, A. K. Dhawale, T. Fisher, C. D. Wilson *et al.*, "Discovering precise temporal patterns in large-scale

neural recordings through robust and interpretable time warping," *Neuron*, vol. 105, no. 2, pp. 246–259, 2020, publisher: Elsevier.

[12] T. P. Vogels and L. F. Abbott, "Signal propagation and logic gating in networks of integrate-and-fire neurons," *The Journal of Neuroscience*, vol. 25, no. 46, p. 10786, Nov. 2005.

[13] ——, "Gating multiple signals through detailed balance of excitation and inhibition in spiking networks," *Nature Neuroscience*, vol. 12, no. 4, pp. 483–491, Apr. 2009, number: 4 Publisher: Nature Publishing Group.

[14] M. Abeles, *Corticonics: Neural Circuits of the Cerebral Cortex*, 1st ed. Cambridge University Press, Feb. 1991.

[15] A. Kumar, S. Rotter, and A. Aertsen, "Spiking activity propagation in neuronal networks: reconciling different perspectives on neural coding," *Nature Reviews Neuroscience*, vol. 11, no. 9, pp. 615–627, Sep. 2010, number: 9 Publisher: Nature Publishing Group.

[16] E. Bienenstock, "A model of neocortex," *Network: Computation in Neural Systems*, vol. 6, no. 2, pp. 179–224, Jan. 1995.

[17] G. Laurent, "On the value of model diversity in neuroscience," *Nature Reviews Neuroscience*, vol. 21, no. 8, pp. 395–396, Aug. 2020.

[18] P. Cisek, "Evolution of behavioural control from chordates to primates," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 377, no. 1844, p. 20200522, Feb. 2022.

[19] ——, "Resynthesizing behavior through phylogenetic refinement," *Attention, Perception, & Psychophysics*, vol. 81, no. 7, pp. 2265–2287, Oct. 2019.

[20] D. Hain, T. Gallego-Flores, M. Klinkmann, A. Macias, E. Ciirdaeva, A. Arends *et al.*, "Molecular diversity and evolution of neuron types in the amniote brain," *Science*, vol. 377, no. 6610, p. eabp8202, Sep. 2022.

[21] G. F. Striedter and R. G. Northcutt, *Brains Through Time: A Natural History of Vertebrates*, 1st ed. Oxford University Press, Dec. 2019.

[22] B. Hommel, C. S. Chapman, P. Cisek, H. F. Neyedli, J.-H. Song, and T. N. Welsh, "No one knows what attention is," *Attention, Perception, & Psychophysics*, vol. 81, no. 7, pp. 2288–2303, Oct. 2019.

[23] L. Baciadonna, F. M. Cornero, N. J. Emery, and N. S. Clayton, "Convergent evolution of complex cognition: Insights from the field of avian cognition into the study of self-awareness," *Learning & Behavior*, vol. 49, no. 1, pp. 9–22, Mar. 2021.

[24] P.-A. Libourel, B. Barrillot, S. Arthaud, B. Massot, A.-L. Morel, O. Beuf *et al.*, "Partial homologies between sleep states in lizards, mammals, and birds suggest a complex evolution of sleep states in amniotes," *PLOS Biology*, vol. 16, no. 10, p. e2005982, Oct. 2018.

[25] R. K. Naumann, J. M. Ondracek, S. Reiter, M. Shein-Idelson, M. A. Tosches, T. M. Yamawaki, and G. Laurent, "The reptilian brain," *Current Biology*, vol. 25, no. 8, pp. R317–R321, Apr. 2015.

[26] J. Fournier, C. M. Müller, I. Schneider, and G. Laurent, "Spatial Information in a Non-retinotopic Visual Cortex," *Neuron*, vol. 97, no. 1, pp. 164–180.e7, Jan. 2018.

[27] S. Guirado and J. Carlos Dávila, "Evolution of the Optic Tectum in Amniotes," in *Encyclopedia of Neuroscience*, M. D. Binder, N. Hirokawa, and U. Windhorst, Eds. Berlin, Heidelberg: Springer, 2009, pp. 1375–1380.

[28] M. Shein-Idelson, L. Pammer, M. Hemberger, and G. Laurent, "Large-scale mapping of cortical synaptic projections with extracellular electrode arrays," *Nature Methods*, vol. 14, no. 9, pp. 882–890, Sep. 2017.

[29] W. L. Shew, H. Yang, S. Yu, R. Roy, and D. Plenz, "Information capacity and transmission are maximized in balanced cortical networks with neuronal avalanches," *Journal of Neuroscience*, vol. 31, no. 1, pp. 55–63, Jan. 2011.

[30] H. Norimoto, L. A. Fenk, H.-H. Li, M. A. Tosches, T. Gallego-Flores, D. Hain *et al.*, "A claustrum in reptiles and its role in slow-wave sleep," *Nature*, vol. 578, no. 7795, pp. 413–418, Feb. 2020.

[31] K. Narikiyo, R. Mizuguchi, A. Ajima, M. Shiozaki, H. Hamanaka, J. P. Johansen *et al.*, "The claustrum coordinates cortical slow-wave activity," *Nature Neuroscience*, vol. 23, no. 6, pp. 741–753, Jun. 2020.

[32] M. Shein-Idelson, J. M. Ondracek, H.-P. Liaw, S. Reiter, and G. Laurent, "Slow waves, sharp waves, ripples, and REM in sleeping dragons," *Science*, vol. 352, no. 6285, pp. 590–595, Apr. 2016.

[33] A. R. Adamantidis, C. Gutierrez Herrera, and T. C. Gent, "Oscillating circuitries in the sleeping brain," *Nature Reviews Neuroscience*, vol. 20, no. 12, pp. 746–762, Dec. 2019, number: 12 Publisher: Nature Publishing Group.

[34] D. Levenstein, V. A. Alvarez, A. Amarasingham, H. Azab, Z. S. Chen, R. C. Gerkin *et al.*, "On the Role of Theory and Modeling in Neuroscience," *Journal of Neuroscience*, vol. 43, no. 7, pp. 1074–1088, Feb. 2023, publisher: Society for Neuroscience Section: Viewpoints.

[35] R. Brette and W. Gerstner, "Adaptive exponential integrate-and-fire model as an effective description of neuronal activity," *Journal of Neurophysiology*, vol. 94, no. 5, pp. 3637–3642, Nov. 2005.

[36] C. Miehl, S. Onasch, D. Festa, and J. Gjorgjieva, "Formation and computational implications of assemblies in neural circuits," *The Journal of Physiology*, vol. n/a, no. n/a, Sep. 2022, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1113/JP282750.

[37] M. Hemberger, M. Shein-Idelson, L. Pammer, and G. Laurent, "Reliable Sequential Activation of Neural Assemblies by Single Pyramidal Cells in a Three-Layered Cortex," *Neuron*, vol. 104, no. 2, pp. 353–369.e5, Oct. 2019.

[38] J. L. Riquelme and J. Gjorgjieva, "Towards readable code in neuroscience," *Nature Reviews Neuroscience*, vol. 22, no. 5, pp. 257–258, May 2021, number: 5 Publisher: Nature Publishing Group.

[39] C. Linssen, M. E. Lepperød, J. Mitchell, J. Pronold, J. M. Eppler, C. Keup *et al.*, "NEST 2.16.0," Aug. 2018.

[40] L. Puelles, J. Sandoval, A. Ayad, R. Corral, A. Alonso, J. Ferran, and M. Martínez-de-la Torre, "The Pallium in Reptiles and Birds in the Light of the Updated Tetrapartite Pallium Model," in *Evolution of Nervous Systems: Second Edition*, Jan. 2017, pp. 519–555, journal Abbreviation: Evolution of Nervous Systems: Second Edition.

[41] M. A. Tosches, T. M. Yamawaki, R. K. Naumann, A. A. Jacobi, G. Tushev, and G. Laurent, "Evolution of pallium, hippocampus, and cortical cell types revealed by single-cell transcriptomics in reptiles," *Science*, vol. 360, no. 6391, pp. 881–888, May 2018.

[42] R. G. Northcutt, "Variation in Reptilian Brains and Cognition," *Brain, Behavior and Evolution*, vol. 82, no. 1, pp. 45–54, 2013, publisher: Karger Publishers.

[43] L. Bruce, "Evolution of the Nervous System in Reptiles," in *Evolution of Nervous Systems*. Elsevier, 2007, pp. 125–156.

[44] G. Laurent, J. Fournier, M. Hemberger, C. Müller, R. Naumann, J. M. Ondracek *et al.*, "Cortical Evolution: Introduction to the Reptilian Cortex," in *Micro-, Meso- and Macro-Dynamics of the Brain*, ser. Research and Perspectives in Neurosciences, G. Buzsáki and Y. Christen, Eds. Cham: Springer International Publishing, 2016, pp. 23–33.

[45] L. Medina and A. Abellán, "Development and evolution of the pallium," *Seminars in Cell & Developmental Biology*, vol. 20, no. 6, pp. 698–711, Aug. 2009.

[46] N. Moreno, R. Morona, J. M. López, and A. González, "Subdivisions of the turtle Pseudemys scripta subpallium based on the expression of regulatory genes and neuronal markers," *The Journal of Comparative Neurology*, vol. 518, no. 24, pp. 4877–4902, Dec. 2010.

[47] G. F. Striedter, "The telencephalon of tetrapods in evolution," *Brain, Behavior and Evolution*, vol. 49, no. 4, pp. 179–213, 1997.

[48] E. D. Jarvis, O. Güntürkün, L. Bruce, A. Csillag, H. Karten, W. Kuenzel *et al.*, "Avian brains and a new understanding of vertebrate brain evolution," *Nature Reviews Neuroscience*, vol. 6, no. 2, pp. 151–159, Feb. 2005, number: 2 Publisher: Nature Publishing Group.

[49] J. F. Montiel, N. A. Vasistha, F. Garcia-Moreno, and Z. Molnár, "From sauropsids to mammals and back: New approaches to comparative cortical development," *Journal of Comparative Neurology*, vol. 524, no. 3, pp. 630–645, 2016, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.23871.

[50] J. Fournier, C. M. Müller, and G. Laurent, "Looking for the roots of cortical sensory computation in three-layered cortices," *Current Opinion in Neurobiology*, vol. 31, pp. 119–126, Apr. 2015.

[51] F. Aboitiz, J. Montiel, D. Morales, and M. Concha, "Evolutionary divergence of the reptilian and the mammalian brains: considerations on connectivity and development," *Brain Research. Brain Research Reviews*, vol. 39, no. 2-3, pp. 141–153, Sep. 2002.

[52] A. Reiner, "Neurotransmitter organization and connections of turtle cortex: implications for the evolution of mammalian isocortex," *Comparative Biochemistry and Physiology Part A: Physiology*, vol. 104, no. 4, pp. 735–748, Apr. 1993.

[53] T. Crockett, N. Wright, S. Thornquist, M. Ariel, and R. Wessel, "Turtle Dorsal Cortex Pyramidal Neurons Comprise Two Distinct Cell Types with Indistinguishable Visual Responses," *PLOS ONE*, vol. 10, no. 12, p. e0144012, Dec. 2015.

[54] P. H. Desan, "The Organization of the Cerebral Cortex of the Pond Turtle, Pseudemys Scripta Elegans," Ph.D., Harvard University, United States – Massachusetts, 1984, iSBN: 9798641918532.

[55] J. B. Colombe, J. Sylvester, J. Block, and P. S. Ulinski, "Subpial and stellate cells: Two populations of interneurons in turtle visual cortex," *Journal of Comparative Neurology*, vol. 471, no. 3, pp. 333–351, 2004, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.20037.

[56] M. E. Larkum, S. Watanabe, N. Lasser-Ross, P. Rhodes, and W. N. Ross, "Dendritic properties of turtle pyramidal neurons," *Journal of Neurophysiology*, vol. 99, no. 2, pp. 683–694, Feb. 2008.

[57] A. R. Kriegstein and B. W. Connors, "Cellular physiology of the turtle visual cortex: synaptic properties and intrinsic circuitry," *The Journal of*

Neuroscience: The Official Journal of the Society for Neuroscience, vol. 6, no. 1, pp. 178–191, Jan. 1986.

[58] B. W. Connors and A. R. Kriegstein, "Cellular physiology of the turtle visual cortex: distinctive properties of pyramidal and stellate neurons," *Journal of Neuroscience*, vol. 6, no. 1, pp. 164–177, Jan. 1986, publisher: Society for Neuroscience Section: Articles.

[59] A. R. Kriegstein, "Synaptic responses of cortical pyramidal neurons to light stimulation in the isolated turtle visual system," *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, vol. 7, no. 8, pp. 2488–2492, Aug. 1987.

[60] B. Haider, M. Häusser, and M. Carandini, "Inhibition dominates sensory responses in the awake cortex," *Nature*, vol. 493, no. 7430, pp. 97–100, Jan. 2013, number: 7430 Publisher: Nature Publishing Group.

[61] D. A. Belkin, "Anoxia: Tolerance in Reptiles," *Science*, vol. 139, no. 3554, pp. 492–493, Feb. 1963, publisher: American Association for the Advancement of Science.

[62] W. L. Shew, W. P. Clawson, J. Pobst, Y. Karimipanah, N. C. Wright, and R. Wessel, "Adaptation to sensory input tunes visual cortex to criticality," *Nature Physics*, vol. 11, no. 8, pp. 659–663, Aug. 2015.

[63] G. D. Field, J. L. Gauthier, A. Sher, M. Greschner, T. A. Machado, L. H. Jepson *et al.*, "Functional connectivity in the retina at the resolution of photoreceptors," *Nature*, vol. 467, no. 7316, pp. 673–677, Oct. 2010, number: 7316 Publisher: Nature Publishing Group.

[64] G. Turrigiano, "Homeostatic synaptic plasticity: local and global mechanisms for stabilizing neuronal function," *Cold Spring Harbor Perspectives in Biology*, vol. 4, no. 1, p. a005736, Jan. 2012.

[65] M. Hemberger, L. Pammer, and G. Laurent, "Comparative approaches to cortical microcircuits," *Current Opinion in Neurobiology*, vol. 41, pp. 24–30, Dec. 2016.

[66] Y. K. Hong, C. O. Lacefield, C. C. Rodgers, and R. M. Bruno, "Sensation, movement and learning in the absence of barrel cortex," *Nature*, vol. 561, no. 7724, pp. 542–546, Sep. 2018, number: 7724 Publisher: Nature Publishing Group.

[67] C. P. J. de Kock, R. M. Bruno, H. Spors, and B. Sakmann, "Layer- and cell-type-specific suprathreshold stimulus representation in rat primary somatosensory cortex," *The Journal of Physiology*, vol. 581, no. Pt 1, pp. 139–154, May 2007.

[68] S. Song, P. J. Sjöström, M. Reigl, S. Nelson, and D. B. Chklovskii, "Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits," *PLOS Biology*, vol. 3, no. 3, p. e68, Mar. 2005, publisher: Public Library of Science.

[69] B. Sakmann, "Patch pipettes are more useful than initially thought: simultaneous pre- and postsynaptic recording from mammalian CNS synapses in vitro and in vivo," *Pflugers Archiv: European Journal of Physiology*, vol. 453, no. 3, pp. 249–259, Dec. 2006.

[70] H. Markram, J. Lübke, M. Frotscher, A. Roth, and B. Sakmann, "Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex." *The Journal of Physiology*, vol. 500, no. 2, pp. 409–440, 1997, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1113/jphysiol.1997.sp022031.

[71] A. M. Thomson and J. Deuchars, "Synaptic interactions in neocortical local circuits: dual intracellular recordings in vitro," *Cerebral Cortex (New York, N.Y.: 1991)*, vol. 7, no. 6, pp. 510–522, Sep. 1997.

[72] D. Feldmeyer, V. Egger, J. Lübke, and B. Sakmann, "Reliable synaptic connections between pairs of excitatory layer 4 neurones within a single 'barrel' of developing rat somatosensory cortex," *The Journal of Physiology*, vol. 521, no. Pt 1, pp. 169–190, Nov. 1999.

[73] R. Perin, T. K. Berger, and H. Markram, "A synaptic organizing principle for cortical neuronal groups," *Proceedings of the National Academy of Sciences*, vol. 108, no. 13, pp. 5419–5424, Mar. 2011, publisher: Proceedings of the National Academy of Sciences.

[74] X. Jiang, S. Shen, C. R. Cadwell, P. Berens, F. Sinz, A. S. Ecker *et al.*, "Principles of connectivity among morphologically defined cell types in adult neocortex," *Science*, vol. 350, no. 6264, p. aac9462, Nov. 2015, publisher: American Association for the Advancement of Science.

[75] V. Egger, D. Feldmeyer, and B. Sakmann, "Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex," *Nature Neuroscience*, vol. 2, no. 12, pp. 1098–1105, Dec. 1999, number: 12 Publisher: Nature Publishing Group.

[76] M. Helmstaedter, C. P. J. de Kock, D. Feldmeyer, R. M. Bruno, and B. Sakmann, "Reconstruction of an average cortical column in silico," *Brain Research Reviews*, vol. 55, no. 2, pp. 193–203, Oct. 2007.

[77] H. Markram, E. Muller, S. Ramaswamy, M. W. Reimann, M. Abdellah, C. A. Sanchez *et al.*, "Reconstruction and Simulation of Neocortical Microcircuitry," *Cell*, vol. 163, no. 2, pp. 456–492, Oct. 2015, publisher: Elsevier.

[78] M. W. Reimann, S. Bolaños-Puchet, J.-D. Courcol, D. E. Santander, A. Arnaudon, B. Coste *et al.*, "Modeling and Simulation of Rat Non-Barrel Somatosensory Cortex. Part I: Modeling Anatomy," Aug. 2022, pages: 2022.08.11.503144 Section: New Results.

[79] C. Stringer, M. Pachitariu, N. Steinmetz, C. B. Reddy, M. Carandini, and K. D. Harris, "Spontaneous behaviors drive multidimensional, brainwide activity," *Science*, vol. 364, no. 6437, p. eaav7893, 2019, publisher: American Association for the Advancement of Science.

[80] S. B. Hofer, H. Ko, B. Pichler, J. Vogelstein, H. Ros, H. Zeng *et al.*, "Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex," *Nature Neuroscience*, vol. 14, no. 8, pp. 1045–1052, Aug. 2011, number: 8 Publisher: Nature Publishing Group.

[81] G. Tamás, A. L. Simon, Anna, and J. Szabadics, "Identified Sources and Targets of Slow Inhibition in the Neocortex," *Science*, vol. 299, no. 5614, pp. 1902–1905, Mar. 2003, publisher: American Association for the Advancement of Science.

[82] Y. Han, J. M. Kebschull, R. A. A. Campbell, D. Cowan, F. Imhof, A. M. Zador, and T. D. Mrsic-Flogel, "The logic of single-cell projections from visual cortex," *Nature*, vol. 556, no. 7699, pp. 51–56, Apr. 2018, number: 7699 Publisher: Nature Publishing Group.

[83] R. Tremblay, S. Lee, and B. Rudy, "GABAergic Interneurons in the Neocortex: From Cellular Properties to Circuits," *Neuron*, vol. 91, no. 2, pp. 260–292, Jul. 2016.

[84] B. Rudy, G. Fishell, S. Lee, and J. Hjerling-Leffler, "Three groups of interneurons account for nearly 100% of neocortical GABAergic neurons," *Developmental Neurobiology*, vol. 71, no. 1, pp. 45–61, 2011, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/dneu.20853.

[85] A. Reyes, R. Lujan, A. Rozov, N. Burnashev, P. Somogyi, and B. Sakmann, "Target-cell-specific facilitation and depression in neocortical circuits," *Nature Neuroscience*, vol. 1, no. 4, pp. 279–285, Aug. 1998, number: 4 Publisher: Nature Publishing Group.

[86] M. Avermann, C. Tomm, C. Mateo, W. Gerstner, and C. C. H. Petersen, "Microcircuits of excitatory and inhibitory neurons in layer 2/3 of mouse barrel cortex," *Journal of Neurophysiology*, vol. 107, no. 11, pp. 3116–3134, Jun. 2012, publisher: American Physiological Society.

[87] G. Silberberg and H. Markram, "Disynaptic Inhibition between Neocortical Pyramidal Cells Mediated by Martinotti Cells," *Neuron*, vol. 53, no. 5, pp. 735–746, Mar. 2007.

[88] D. F. English, S. McKenzie, T. Evans, K. Kim, E. Yoon, and G. Buzsáki, "Pyramidal Cell-Interneuron Circuit Architecture and Dynamics in Hippocampal Networks," *Neuron*, vol. 96, no. 2, pp. 505–520.e7, Oct. 2017.

[89] F. Pouille and M. Scanziani, "Routing of spike series by dynamic circuits in the hippocampus," *Nature*, vol. 429, no. 6993, pp. 717–723, Jun. 2004, number: 6993 Publisher: Nature Publishing Group.

[90] T. Klausberger and P. Somogyi, "Neuronal Diversity and Temporal Dynamics: The Unity of Hippocampal Circuit Operations," *Science*, vol. 321, no. 5885, pp. 53–57, Jul. 2008, publisher: American Association for the Advancement of Science.

[91] J. J. Couey, A. Witoelar, S.-J. Zhang, K. Zheng, J. Ye, B. Dunn *et al.*, "Recurrent inhibitory circuitry as a mechanism for grid formation," *Nature Neuroscience*, vol. 16, no. 3, pp. 318–324, Mar. 2013, number: 3 Publisher: Nature Publishing Group.

[92] G. Laurent, M. Stopfer, R. W. Friedrich, M. I. Rabinovich, A. Volkovskii, and H. D. Abarbanel, "Odor Encoding as an Active, Dynamical Process: Experiments, Computation, and Theory," *Annual Review of Neuroscience*, vol. 24, no. 1, pp. 263–297, 2001, _eprint: https://doi.org/10.1146/annurev.neuro.24.1.263.

[93] S. N. Chettih and C. D. Harvey, "Single-neuron perturbations reveal feature-specific competition in V1," *Nature*, vol. 567, no. 7748, pp. 334–340, Mar. 2019.

[94] D. Feldmeyer, J. Lübke, and B. Sakmann, "Efficacy and connectivity of intracolumnar pairs of layer 2/3 pyramidal cells in the barrel cortex of juvenile rats," *The Journal of Physiology*, vol. 575, no. Pt 2, pp. 583–602, Sep. 2006.

[95] G. Buzsáki and K. Mizuseki, "The log-dynamic brain: how skewed distributions affect network operations," *Nature Reviews Neuroscience*, vol. 15, no. 4, pp. 264–278, Apr. 2014, number: 4 Publisher: Nature Publishing Group.

[96] S. Dorkenwald, N. L. Turner, T. Macrina, K. Lee, R. Lu, J. Wu *et al.*, "Binary and analog variation of synapses between cortical pyramidal neurons," Neuroscience, preprint, Dec. 2019.

[97] G. Molnár, S. Oláh, G. Komlósi, M. Füle, J. Szabadics, C. Varga *et al.*, "Complex Events Initiated by Individual Spikes in the Human Cerebral Cortex," *PLOS Biology*, vol. 6, no. 9, p. e222, Sep. 2008, publisher: Public Library of Science.

[98] D. A. Henze, L. Wittner, and G. Buzsáki, "Single granule cells reliably discharge targets in the hippocampal CA3 network in vivo," *Nature Neuroscience*,

vol. 5, no. 8, pp. 790–795, Aug. 2002, number: 8 Publisher: Nature Publishing Group.

[99] N. P. Vyleta, C. Borges-Merjane, and P. Jonas, "Plasticity-dependent, full detonation at hippocampal mossy fiber–CA3 pyramidal neuron synapses," *eLife*, vol. 5, p. e17977, Oct. 2016, publisher: eLife Sciences Publications, Ltd.

[100] M. Frotscher, "Mossy fibres form synapses with identified pyramidal basket cells in the CA3 region of the guinea-pig hippocampus: a combined Golgi-electron microscope study," *Journal of Neurocytology*, vol. 14, no. 2, pp. 245–259, Apr. 1985.

[101] S. Lefort, C. Tomm, J.-C. Floyd Sarria, and C. C. H. Petersen, "The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex," *Neuron*, vol. 61, no. 2, pp. 301–316, Jan. 2009.

[102] S. J. Guzman, A. Schlögl, M. Frotscher, and P. Jonas, "Synaptic mechanisms of pattern completion in the hippocampal CA3 network," *Science (New York, N.Y.)*, vol. 353, no. 6304, pp. 1117–1123, Sep. 2016.

[103] A. Granda and C. Dvorak, "Vision in Turtles," in *The Visual System in Vertebrates*, ser. Handbook of Sensory Physiology, F. Crescitelli, C. Dvorak, D. Eder, A. Granda, D. Hamasaki, K. Holmberg *et al.*, Eds. Berlin, Heidelberg: Springer, 1977, pp. 451–495.

[104] N. Azuma, "Molecular Cell Biology on Morphogenesis of the Fovea and Evolution of the Central Vision," *Japanese Journal of Ophthalmology*, vol. 45, no. 3, p. 325, May 2001.

[105] J. Ammermüller and H. Kolb, "Functional architecture of the turtle retina," *Progress in Retinal and Eye Research*, vol. 15, no. 2, pp. 393–433, Jan. 1996.

[106] A. H. Bass and R. G. Northcutt, "Retinal recipient nuclei in the painted turtle, Chrysemys picta: an autoradiographic and HRP study," *The Journal of Comparative Neurology*, vol. 199, no. 1, pp. 97–112, Jun. 1981.

[107] P. S. Ulinski and J. Nautiyal, "Organization of retinogeniculate projections in turtles of the genera Pseudemys and Chrysemys," *Journal of Comparative Neurology*, vol. 276, no. 1, pp. 92–112, 1988, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.902760107.

[108] K. A. Mulligan and P. S. Ulinski, "Organization of geniculocortical projections in turtles: Isoazimuth lamellae in the visual cortex," *Journal of Comparative Neurology*, vol. 296, no. 4, pp. 531–547, 1990, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.902960403.

[109] L. M. Smith, F. F. Ebner, and M. Colonnier, "The thalamocortical projection in Pseudemys turtles: A quantitative electron microscopic study,"

Journal of Comparative Neurology, vol. 190, no. 3, pp. 445–461, 1980, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cne.901900304.

[110] P. Z. Mazurskaya, "Organization of receptive fields in the forebrain ofEmys orbicularis," *Neuroscience and Behavioral Physiology*, vol. 6, no. 4, pp. 311–318, Oct. 1973.

[111] ——, "Retinal projection in the forebrain of Emys orbicularis," *Neuroscience and Behavioral Physiology*, vol. 6, no. 1, pp. 75–82, Jan. 1973.

[112] G. B. Ermentrout and D. Kleinfeld, "Traveling Electrical Waves in Cortex: Insights from Phase Dynamics and Speculation on a Computational Role," *Neuron*, vol. 29, no. 1, pp. 33–44, Jan. 2001.

[113] L. Muller, F. Chavane, J. Reynolds, and T. J. Sejnowski, "Cortical travelling waves: mechanisms and computational principles," *Nature Reviews Neuroscience*, vol. 19, no. 5, pp. 255–268, May 2018, number: 5 Publisher: Nature Publishing Group.

[114] J. C. Prechtl, L. B. Cohen, B. Pesaran, P. P. Mitra, and D. Kleinfeld, "Visual stimuli induce waves of electrical activity in turtle cortex," *Proceedings of the National Academy of Sciences*, vol. 94, no. 14, pp. 7621–7626, Jul. 1997, publisher: Proceedings of the National Academy of Sciences.

[115] D. M. Senseman and K. A. Robbins, "High-speed VSD imaging of visually evoked cortical waves: decomposition into intra- and intercortical wave motions," *Journal of Neurophysiology*, vol. 87, no. 3, pp. 1499–1514, Mar. 2002.

[116] U. Mitzdorf, "Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena," *Physiological Reviews*, vol. 65, no. 1, pp. 37–100, Jan. 1985.

[117] I. Schneider, "Visually-evoked dynamics in turtle dorsal cortex examined using large-scale, chronic electrocorticogram recordings," Sep. 2015.

[118] G. Buzsáki, "Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning," *Hippocampus*, vol. 25, no. 10, pp. 1073–1188, 2015, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/hipo.22488.

[119] G. Buzsáki, C. A. Anastassiou, and C. Koch, "The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes," *Nature Reviews Neuroscience*, vol. 13, no. 6, pp. 407–420, Jun. 2012.

[120] D. Levenstein, G. Buzsáki, and J. Rinzel, "NREM sleep in the rodent neocortex and hippocampus reflects excitable dynamics," *Nature Communications*, vol. 10, no. 1, p. 2478, Jun. 2019, number: 1 Publisher: Nature Publishing Group.

[121] L. F. Abbott and S. B. Nelson, "Synaptic plasticity: taming the beast," *Nature Neuroscience*, vol. 3, no. S11, pp. 1178–1183, Nov. 2000.

[122] U. Wagner, S. Gais, and J. Born, "Emotional Memory Formation Is Enhanced across Sleep Intervals with High Amounts of Rapid Eye Movement Sleep," *Learning & Memory*, vol. 8, no. 2, pp. 112–119, Mar. 2001.

[123] J. G. Klinzing, N. Niethard, and J. Born, "Mechanisms of systems memory consolidation during sleep," *Nature Neuroscience*, vol. 22, no. 10, pp. 1598–1610, Oct. 2019.

[124] J. Peever and P. M. Fuller, "The Biology of REM Sleep," *Current Biology*, vol. 27, no. 22, pp. R1237–R1248, Nov. 2017.

[125] S. P. Jadhav, C. Kemere, P. W. German, and L. M. Frank, "Awake Hippocampal Sharp-Wave Ripples Support Spatial Memory," *Science*, vol. 336, no. 6087, pp. 1454–1458, Jun. 2012.

[126] K. Diba and G. Buzsáki, "Forward and reverse hippocampal place-cell sequences during ripples," *Nature Neuroscience*, vol. 10, no. 10, pp. 1241–1242, Oct. 2007.

[127] L. C. Leung, G. X. Wang, R. Madelaine, G. Skariah, K. Kawakami, K. Deisseroth *et al.*, "Neural signatures of sleep in zebrafish," *Nature*, vol. 571, no. 7764, pp. 198–204, Jul. 2019.

[128] Y. Ikegaya, T. Sasaki, D. Ishikawa, N. Honma, K. Tao, N. Takahashi *et al.*, "Interpyramid spike transmission stabilizes the sparseness of recurrent network activity," *Cerebral Cortex*, vol. 23, no. 2, pp. 293–304, Feb. 2013.

[129] R. Egger, Y. Tupikov, M. Elmaleh, K. A. Katlowitz, S. E. Benezra, M. A. Picardo *et al.*, "Local Axonal Conduction Shapes the Spatiotemporal Properties of Neural Sequences," *Cell*, vol. 183, no. 2, pp. 537–548.e12, Oct. 2020.

[130] R. J. Douglas, C. Koch, M. Mahowald, K. A. Martin, and H. H. Suarez, "Recurrent excitation in neocortical circuits," *Science (New York, N.Y.)*, vol. 269, no. 5226, pp. 981–985, Aug. 1995.

[131] J. S. Griffith, "On the stability of brain-like structures," *Biophysical Journal*, vol. 3, no. 4, pp. 299–308, Jul. 1963.

[132] M. Abeles, *Local Cortical Circuits*, ser. Studies of Brain Function, V. Braitenberg, H. B. Barlow, H. Bullock, E. Florey, O.-J. Grüsser, and A. Peters, Eds. Berlin, Heidelberg: Springer, 1982, vol. 6.

[133] P. Zheng and J. Triesch, "Robust development of synfire chains from multiple plasticity mechanisms," *Frontiers in Computational Neuroscience*, vol. 8, 2014.

[134] P. König, A. K. Engel, and W. Singer, "Integrator or coincidence detector? The role of the cortical neuron revisited," *Trends in Neurosciences*, vol. 19, no. 4, pp. 130–137, Apr. 1996.

[135] Y. Prut, E. Vaadia, H. Bergman, I. Haalman, H. Slovin, and M. Abeles, "Spatiotemporal structure of cortical activity: properties and behavioral relevance," *Journal of Neurophysiology*, vol. 79, no. 6, pp. 2857–2874, Jun. 1998.

[136] M. Abeles, H. Bergman, E. Margalit, and E. Vaadia, "Spatiotemporal firing patterns in the frontal cortex of behaving monkeys," *Journal of Neurophysiology*, vol. 70, no. 4, pp. 1629–1638, Oct. 1993.

[137] M. Abeles, G. Hayon, and D. Lehmann, "Modeling compositionality by dynamic binding of synfire chains," *Journal of Computational Neuroscience*, vol. 17, no. 2, pp. 179–201, Sep. 2004.

[138] M. Diesmann, M.-O. Gewaltig, and A. Aertsen, "Stable propagation of synchronous spiking in cortical neural networks," *Nature*, vol. 402, no. 6761, pp. 529–533, Dec. 1999, number: 6761 Publisher: Nature Publishing Group.

[139] L. Montangie, C. Miehl, and J. Gjorgjieva, "Autonomous emergence of connectivity assemblies via spike triplet interactions," *PLOS Computational Biology*, vol. 16, no. 5, p. e1007835, May 2020, publisher: Public Library of Science.

[140] I. R. Fiete, W. Senn, C. Z. Wang, and R. H. Hahnloser, "Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity," *Neuron*, vol. 65, no. 4, pp. 563–576, Feb. 2010.

[141] A. Lazar, G. Pipa, and J. Triesch, "SORN: a self-organizing recurrent neural network," *Frontiers in Computational Neuroscience*, vol. 3, 2009.

[142] Y. Aviel, C. Mehring, M. Abeles, and D. Horn, "On Embedding Synfire Chains in a Balanced Network," *Neural Computation*, vol. 15, no. 6, pp. 1321–1340, Jun. 2003.

[143] C. Mehring, U. Hehl, M. Kubo, M. Diesmann, and A. Aertsen, "Activity dynamics and propagation of synchronous spiking in locally connected random networks," *Biological Cybernetics*, vol. 88, no. 5, pp. 395–408, May 2003.

[144] A. Kumar, S. Rotter, and A. Aertsen, "Conditions for Propagating Synchronous Spiking and Asynchronous Firing Rates in a Cortical Network Model," *The Journal of Neuroscience*, vol. 28, no. 20, pp. 5268–5280, May 2008.

[145] A. Riehle, S. Grün, M. Diesmann, and A. Aertsen, "Spike Synchronization and Rate Modulation Differentially Involved in Motor Cortical Function," *Science*, vol. 278, no. 5345, pp. 1950–1953, Dec. 1997, publisher: American Association for the Advancement of Science.

[146] Y. Ikegaya, G. Aaron, R. Cossart, D. Aronov, I. Lampl, D. Ferster, and R. Yuste, "Synfire chains and cortical songs: temporal modules of cortical activity," *Science*, vol. 304, no. 5670, pp. 559–564, Apr. 2004.

[147] M. W. Oram, M. C. Wiener, R. Lestienne, and B. J. Richmond, "Stochastic nature of precisely timed spike patterns in visual system neuronal responses," *Journal of Neurophysiology*, vol. 81, no. 6, pp. 3021–3033, Jun. 1999.

[148] J.-M. Fellous, P. H. Tiesinga, P. J. Thomas, and T. J. Sejnowski, "Discovering spike patterns in neuronal responses," *Journal of Neuroscience*, vol. 24, no. 12, pp. 2989–3001, 2004, publisher: Soc Neuroscience.

[149] G. Buzsáki and D. Tingley, "Space and Time: The Hippocampus as a Sequence Generator," *Trends in Cognitive Sciences*, vol. 22, no. 10, pp. 853–869, Oct. 2018.

[150] G. Dragoi and S. Tonegawa, "Preplay of future place cell sequences by hippocampal cellular assemblies," *Nature*, vol. 469, no. 7330, pp. 397–401, Jan. 2011.

[151] M. N. Modi, A. K. Dhawale, and U. S. Bhalla, "CA1 cell activity sequences emerge after reorganization of network correlation structure during associative learning," *eLife*, vol. 3, p. e01982, Mar. 2014.

[152] A. P. Vaz, J. H. Wittig, S. K. Inati, and K. A. Zaghloul, "Replay of cortical spiking sequences during human memory retrieval," *Science*, vol. 367, no. 6482, pp. 1131–1134, Mar. 2020.

[153] L. Carrillo-Reid, J.-e. K. Miller, J. P. Hamm, J. Jackson, and R. Yuste, "Endogenous sequential cortical activity evoked by visual stimuli," *Journal of Neuroscience*, vol. 35, no. 23, pp. 8813–8828, Jun. 2015.

[154] J. B. Dechery and J. N. MacLean, "Emergent cortical circuit dynamics contain dense, interwoven ensembles of spike sequences," *Journal of Neurophysiology*, vol. 118, no. 3, pp. 1914–1925, Sep. 2017.

[155] A. Luczak and J. N. MacLean, "Default activity patterns at the neocortical microcircuit level," *Frontiers in Integrative Neuroscience*, vol. 6, 2012.

[156] A. Luczak, B. L. McNaughton, and K. D. Harris, "Packet-based communication in the cortex," *Nature Reviews Neuroscience*, vol. 16, no. 12, pp. 745–755, Dec. 2015.

[157] W. Gerstner, *Neuronal Dynamics: From Single Neurons To Networks And Models Of Cognition*. Cambridge: Cambridge University Press, Sep. 2014.

[158] M. London, A. Roth, L. Beeren, M. Häusser, and P. E. Latham, "Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex," *Nature*, vol. 466, no. 7302, pp. 123–127, Jul. 2010.

[159] J. Kremkow, A. Aertsen, and A. Kumar, "Gating of signal propagation in spiking neural networks by balanced and correlated excitation and inhibition," *Journal of Neuroscience*, vol. 30, no. 47, pp. 15 760–15 768, Nov. 2010.

[160] Z. F. Mainen and T. J. Sejnowski, "Reliability of spike timing in neocortical neurons," *Science*, vol. 268, no. 5216, pp. 1503–1506, Jun. 1995.

[161] A. Renart, J. de la Rocha, P. Bartho, L. Hollender, N. Parga, A. Reyes, and K. D. Harris, "The asynchronous state in cortical circuits," *Science*, vol. 327, no. 5965, pp. 587–590, Jan. 2010.

[162] M. N. Shadlen and W. T. Newsome, "Noise, neural codes and cortical organization," *Current Opinion in Neurobiology*, vol. 4, no. 4, pp. 569–579, Aug. 1994.

[163] W. Bair and C. Koch, "Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey," *Neural Computation*, vol. 8, no. 6, pp. 1185–1202, Aug. 1996.

[164] K. A. Bolding and K. M. Franks, "Complementary codes for odor identity and intensity in olfactory cortex," *eLife*, vol. 6, p. e22630, Apr. 2017.

[165] R. H. R. Hahnloser, A. A. Kozhevnikov, and M. S. Fee, "An ultra-sparse code underlies the generation of neural sequences in a songbird," *Nature*, vol. 419, no. 6902, pp. 65–70, Sep. 2002.

[166] M. Wehr and G. Laurent, "Odour encoding by temporal sequences of firing in oscillating neural assemblies," *Nature*, vol. 384, no. 6605, pp. 162–166, Nov. 1996.

[167] M. Wehr and A. M. Zador, "Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex," *Nature*, vol. 426, no. 6965, pp. 442–446, Nov. 2003.

[168] M. Brecht, M. Schneider, B. Sakmann, and T. W. Margrie, "Whisker movements evoked by stimulation of single pyramidal cells in rat motor cortex," *Nature*, vol. 427, no. 6976, pp. 704–710, Feb. 2004.

[169] G. Doron, M. von Heimendahl, P. Schlattmann, A. Houweling, and M. Brecht, "Spiking irregularity and frequency modulate the behavioral report of single-neuron stimulation," *Neuron*, vol. 81, no. 3, pp. 653–663, Feb. 2014.

[170] A. R. Houweling and M. Brecht, "Behavioural report of single neuron stimulation in somatosensory cortex," *Nature*, vol. 451, no. 7174, pp. 65–68, Jan. 2008.

[171] A. Kwan and Y. Dan, "Dissection of cortical microcircuits by single-neuron stimulation in vivo," *Current Biology*, vol. 22, no. 16, pp. 1459–1467, Aug. 2012.

[172] J. Wolfe, A. R. Houweling, and M. Brecht, "Sparse and powerful cortical spikes," *Current Opinion in Neurobiology*, vol. 20, no. 3, pp. 306–312, Jun. 2010.

[173] M. A. Long, D. Z. Jin, and M. S. Fee, "Support for a synaptic chain model of neuronal sequence generation," *Nature*, vol. 468, no. 7322, pp. 394–399, Nov. 2010.

[174] A. Maes, M. Barahona, and C. Clopath, "Learning spatiotemporal signals using a recurrent spiking network that discretizes time," *PLOS Computational Biology*, vol. 16, no. 1, p. e1007606, Jan. 2020.

[175] K. Rajan, C. Harvey, and D. Tank, "Recurrent Network Models of Sequence Generation and Memory," *Neuron*, vol. 90, no. 1, pp. 128–142, Apr. 2016.

[176] Z. Nenadic, B. K. Ghosh, and P. Ulinski, "Propagating waves in visual cortex: a large-scale model of turtle visual cortex," *Journal of computational neuroscience*, vol. 14, no. 2, pp. 161–184, Apr. 2003, place: United States.

[177] D. M. Iascone, Y. Li, U. Sümbül, M. Doron, H. Chen, V. Andreu *et al.*, "Whole-Neuron Synaptic Mapping Reveals Spatially Precise Excitatory/Inhibitory Balance Limiting Dendritic and Somatic Spiking," *Neuron*, vol. 106, no. 4, pp. 566–578.e8, May 2020.

[178] M. Kajiwara, R. Nomura, F. Goetze, T. Akutsu, and M. Shimono, "Inhibitory neurons are a Central Controlling regulator in the effective cortical microconnectome," Neuroscience, preprint, Feb. 2020.

[179] A. Rubinski and N. E. Ziv, "Remodeling and Tenacity of Inhibitory Synapses: Relationships with Network Activity and Neighboring Excitatory Synapses," *PLOS Computational Biology*, vol. 11, no. 11, p. e1004632, Nov. 2015.

[180] N. C. Wright and R. Wessel, "Network activity influences the subthreshold and spiking visual responses of pyramidal neurons in the three-layer turtle cortex," *Journal of Neurophysiology*, vol. 118, no. 4, pp. 2142–2155, Oct. 2017.

[181] A. Luczak, P. Barthó, S. L. Marguet, G. Buzsáki, and K. D. Harris, "Sequential structure of neocortical spontaneous activity *in vivo*," *Proceedings of the National Academy of Sciences*, vol. 104, no. 1, pp. 347–352, Jan. 2007.

[182] L. Cossell, M. F. Iacaruso, D. R. Muir, R. Houlton, E. N. Sader, H. Ko *et al.*, "Functional organization of excitatory synaptic strength in primary visual cortex," *Nature*, vol. 518, no. 7539, pp. 399–403, Feb. 2015.

[183] R. Sayer, M. Friedlander, and S. Redman, "The time course and amplitude of EPSPs evoked at synapses between pairs of CA3/CA1 neurons in the hippocampal slice," *The Journal of Neuroscience*, vol. 10, no. 3, pp. 826–836, Mar. 1990.

[184] A. Shapson-Coe, M. Januszewski, D. R. Berger, A. Pope, Y. Wu, T. Blakely *et al.*, "A connectomic study of a petascale fragment of human cerebral cortex," Neuroscience, preprint, May 2021.

[185] K. M. Franks and J. S. Isaacson, "Strong Single-Fiber Sensory Inputs to Olfactory Cortex: Implications for Olfactory Coding," *Neuron*, vol. 49, no. 3, pp. 357–363, Feb. 2006.

[186] H. Setareh, M. Deger, C. C. H. Petersen, and W. Gerstner, "Cortical Dynamics in Presence of Assemblies of Densely Connected Weight-Hub Neurons," *Frontiers in Computational Neuroscience*, vol. 11, p. 52, Jun. 2017.

[187] R. V. Rikhye, M. Yildirim, M. Hu, V. Breton-Provencher, and M. Sur, "Reliable Sensory Processing in Mouse Visual Cortex through Cooperative Interactions between Somatostatin and Parvalbumin Interneurons," *The Journal of Neuroscience*, vol. 41, no. 42, pp. 8761–8778, Oct. 2021.

[188] B. Scholl, C. I. Thomas, M. A. Ryan, N. Kamasawa, and D. Fitzpatrick, "Cortical response selectivity derives from strength in numbers of synapses," *Nature*, vol. 590, no. 7844, pp. 111–114, Feb. 2021.

[189] T. Ishikawa and Y. Ikegaya, "Locally sequential synaptic reactivation during hippocampal ripples," *Science Advances*, vol. 6, no. 7, p. eaay1492, Feb. 2020.

[190] A. Kumar, S. Schrader, A. Aertsen, and S. Rotter, "The high-conductance state of cortical networks," *Neural Computation*, vol. 20, no. 1, pp. 1–43, Jan. 2008.

[191] B. E. Pfeiffer and D. J. Foster, "Hippocampal place-cell sequences depict future paths to remembered goals," *Nature*, vol. 497, no. 7447, pp. 74–79, May 2013.

[192] J. Widloski and D. J. Foster, "Flexible rerouting of hippocampal replay sequences around changing barriers in the absence of global place field remapping," *Neuron*, p. S089662732200109X, Feb. 2022.

[193] N. Y. Masse, G. D. Grant, and D. J. Freedman, "Alleviating catastrophic forgetting using context-dependent gating and synaptic stabilization," *Proceedings of the National Academy of Sciences*, vol. 115, no. 44, Oct. 2018.

[194] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling," *arXiv*, 2014, publisher: arXiv Version Number: 1.

[195] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez *et al.*, "Attention Is All You Need," *arXiv*, 2017, publisher: arXiv Version Number: 5.

[196] J. M. Beggs and D. Plenz, "Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures," *The Journal of Neuroscience*, vol. 24, no. 22, p. 5216, Jun. 2004.

[197] J. M. Beggs, "The criticality hypothesis: how local cortical networks might optimize information processing," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 366, no. 1864, pp. 329–343, Feb. 2008.

[198] J. J. Collins, C. C. Chow, and T. T. Imhoff, "Stochastic resonance without tuning," *Nature*, vol. 376, no. 6537, pp. 236–238, Jul. 1995.

[199] ——, "Aperiodic stochastic resonance in excitable systems," *Physical Review E*, vol. 52, no. 4, pp. R3321–R3324, Oct. 1995, publisher: American Physical Society.

[200] J.-n. Teramae, Y. Tsubo, and T. Fukai, "Optimal spike-based communication in excitable networks with strong-sparse and weak-dense links," *Scientific Reports*, vol. 2, no. 1, p. 485, Dec. 2012.

[201] A. Bhatia, S. Moza, and U. S. Bhalla, "Precise excitation-inhibition balance controls gain and timing in the hippocampus," *eLife*, vol. 8, p. e43415, Apr. 2019.

[202] E. Aserinsky and N. Kleitman, "Regularly Occurring Periods of Eye Motility, and Concomitant Phenomena, During Sleep," *Science*, vol. 118, no. 3062, pp. 273–274, Sep. 1953.

[203] W. Dement, "The occurence of low voltage, fast, electroencephalogram patterns during behavioral sleep in the cat," *Electroencephalography and Clinical Neurophysiology*, vol. 10, no. 2, pp. 291–296, May 1958.

[204] M. Jouvet, F. Michel, and J. Courjon, "[On a stage of rapid cerebral electrical activity in the course of physiological sleep]," *Comptes Rendus Des Seances De La Societe De Biologie Et De Ses Filiales*, vol. 153, pp. 1024–1028, 1959.

[205] M. A. Wilson and B. L. McNaughton, "Reactivation of Hippocampal Ensemble Memories During Sleep," *Science*, vol. 265, no. 5172, pp. 676–679, Jul. 1994.

[206] B. Rasch, C. Büchel, S. Gais, and J. Born, "Odor Cues During Slow-Wave Sleep Prompt Declarative Memory Consolidation," *Science*, vol. 315, no. 5817, pp. 1426–1429, Mar. 2007.

[207] G. Girardeau, K. Benchenane, S. I. Wiener, G. Buzsáki, and M. B. Zugaro, "Selective suppression of hippocampal ripples impairs spatial memory," *Nature Neuroscience*, vol. 12, no. 10, pp. 1222–1223, Oct. 2009.

[208] V. Ego-Stengel and M. A. Wilson, "Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat," *Hippocampus*, pp. NA–NA, 2009.

[209] R. SuÃ¡rez, I. Gobius, and L. J. Richards, "Evolution and development of interhemispheric connections in the vertebrate forebrain," *Frontiers in Human Neuroscience*, vol. 8, Jul. 2014.

[210] R. Owen, "On the structure of the brain in marsupial animals," *Philosophical Transactions of the Royal Society of London*, vol. 127, pp. 87–96, Dec. 1837.

[211] W. H. Flower, "XIII. On the commissures of the cerebral hemispheres of the marsupialia and monotremata as compared with those of the placental mammals," *Philosophical Transactions of the Royal Society of London*, vol. 155, pp. 633–651, Dec. 1865.

[212] F. Aboitiz and J. Montiel, "One hundred million years of interhemispheric communication: the history of the corpus callosum," *Brazilian Journal of Medical and Biological Research*, vol. 36, no. 4, pp. 409–420, Apr. 2003.

[213] A. Navas-Olive, R. Amaducci, M.-T. Jurado-Parras, E. R. Sebastian, and L. M. de la Prida, "Deep learning-based feature extraction for prediction and interpretation of sharp-wave ripples in the rodent hippocampus," *eLife*, vol. 11, p. e77772, Sep. 2022.

[214] J. J. Jun, N. A. Steinmetz, J. H. Siegle, D. J. Denman, M. Bauza, B. Barbarits *et al.*, "Fully integrated silicon probes for high-density recording of neural activity," *Nature*, vol. 551, no. 7679, pp. 232–236, Nov. 2017.

[215] Q. Wang, L. Ng, J. A. Harris, D. Feng, Y. Li, J. J. Royall *et al.*, "Organization of the connections between claustrum and cortex in the mouse: Connections between mouse claustrum and cortex," *Journal of Comparative Neurology*, vol. 525, no. 6, pp. 1317–1346, Apr. 2017.

[216] F. Weber and Y. Dan, "Circuit-based interrogation of sleep control," *Nature*, vol. 538, no. 7623, pp. 51–59, Oct. 2016.

[217] C. B. Saper and P. M. Fuller, "Wake–sleep circuitry: an overview," *Current Opinion in Neurobiology*, vol. 44, pp. 186–192, Jun. 2017.

[218] Y. Wang, D. E. Major, and H. J. Karten, "Morphology and connections of nucleus isthmi pars magnocellularis in chicks (Gallus gallus)," *The Journal of Comparative Neurology*, vol. 469, no. 2, pp. 275–297, Feb. 2004.

[219] G. Marin, C. Salas, E. Sentis, X. Rojas, J. C. Letelier, and J. Mpodozis, "A Cholinergic Gating Mechanism Controlled by Competitive Interactions in the Optic Tectum of the Pigeon," *Journal of Neuroscience*, vol. 27, no. 30, pp. 8112–8121, Jul. 2007.

[220] S. P. Mysore and E. I. Knudsen, "A shared inhibitory circuit for both exogenous and endogenous control of stimulus selection," *Nature Neuroscience*, vol. 16, no. 4, pp. 473–478, Apr. 2013.

[221] F. Garrido-Charad, T. Vega-Zuniga, C. Gutiérrez-Ibáñez, P. Fernandez, L. López-Jury, C. González-Cabrera *et al.*, ""Shepherd's crook" neurons drive and synchronize the enhancing and suppressive mechanisms of the midbrain stimulus selection network," *Proceedings of the National Academy of Sciences*, vol. 115, no. 32, Aug. 2018.

[222] E. I. Knudsen, "Neural Circuits That Mediate Selective Attention: A Comparative Perspective," *Trends in Neurosciences*, vol. 41, no. 11, pp. 789–805, Nov. 2018.

[223] S. Guirado, M. A. Real, and J. C. Dávila, "The ascending tectofugal visual system in amniotes: New insights," *Brain Research Bulletin*, vol. 66, no. 4-6, pp. 290–296, Sep. 2005.

[224] E. Gruberg, E. Dudkin, Y. Wang, G. Marin, C. Salas, E. Sentis *et al.*, "Influencing and Interpreting Visual Input: The Role of a Visual Feedback System," *Journal of Neuroscience*, vol. 26, no. 41, pp. 10 368–10 371, Oct. 2006.

[225] G. Marin, "Oscillatory Bursts in the Optic Tectum of Birds Represent Re-Entrant Signals from the Nucleus Isthmi Pars Parvocellularis," *Journal of Neuroscience*, vol. 25, no. 30, pp. 7081–7089, Jul. 2005.

[226] S. J. Zottoli, K. J. Rhodes, J. G. Corrodi, and E. J. Mufson, "Putative cholinergic projections from the nucleus isthmi and the nucleus reticularis mesencephali to the optic tectum in the goldfish (Carassius auratus)," *The Journal of Comparative Neurology*, vol. 273, no. 3, pp. 385–398, Jul. 1988.

[227] P. M. Henriques, N. Rahman, S. E. Jackson, and I. H. Bianco, "Nucleus Isthmi Is Required to Sustain Target Pursuit during Visually Guided Prey-Catching," *Current Biology*, vol. 29, no. 11, pp. 1771–1786.e5, Jun. 2019.

[228] A. M. Fernandes, D. S. Mearns, J. C. Donovan, J. Larsch, T. O. Helmbrecht, Y. Kölsch *et al.*, "Neural circuitry for stimulus selection in the zebrafish visual system," *Neuron*, vol. 109, no. 5, pp. 805–822.e6, Mar. 2021.

[229] E. R. Gruberg and S. B. Udin, "Topographic projections between the nucleus isthmi and the tectum of the frog rana pipiens," *The Journal of Comparative Neurology*, vol. 179, no. 3, pp. 487–500, Jun. 1978.

[230] M. I. Sereno and P. S. Ulinski, "Caudal topographic nucleus isthmi and the rostral nontopographic nucleus isthmi in the turtle,pseudemys scripta," *The Journal of Comparative Neurology*, vol. 261, no. 3, pp. 319–346, Jul. 1987.

[231] M. G. Belekhova and N. B. Kenigfest, "Turtle isthmic complex of visual nuclei: Immunohistochemistry of gamma-aminobutyric acid, choline acetyltransferase, calcium-binding proteins and histochemistry of cytochrome oxidase activity," *Journal of Evolutionary Biochemistry and Physiology*, vol. 50, no. 5, pp. 435–447, Sep. 2014.

[232] H. Schryver and S. Mysore, "Spatial Dependence of Stimulus Competition in the Avian Nucleus Isthmi Pars Magnocellularis," *Brain, Behavior and Evolution*, vol. 93, no. 2-3, pp. 137–151, 2019.

[233] Y. Goll, G. Atlan, and A. Citri, "Attention: the claustrum," *Trends in Neurosciences*, vol. 38, no. 8, pp. 486–495, Aug. 2015.

[234] J. Jackson, J. B. Smith, and A. K. Lee, "The Anatomy and Physiology of Claustrum-Cortex Interactions," *Annual Review of Neuroscience*, vol. 43, no. 1, pp. 231–247, Jul. 2020.

[235] F. C. Crick and C. Koch, "What is the function of the claustrum?" *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 360, no. 1458, pp. 1271–1279, Jun. 2005.

[236] K. A. Bolding and K. M. Franks, "Recurrent cortical circuits implement concentration-invariant odor coding," *Science*, vol. 361, no. 6407, p. eaat6904, Sep. 2018.

[237] R. B. Levy and A. D. Reyes, "Spatial Profile of Excitatory and Inhibitory Synaptic Connectivity in Mouse Primary Auditory Cortex," *Journal of Neuroscience*, vol. 32, no. 16, pp. 5609–5619, Apr. 2012.

[238] A. L. Barth and J. F. Poulet, "Experimental evidence for sparse firing in the neocortex," *Trends in Neurosciences*, vol. 35, no. 6, pp. 345–355, Jun. 2012.

[239] S. Shoham, D. H. O'Connor, and R. Segev, "How silent is the brain: is there a "dark matter" problem in neuroscience?" *Journal of Comparative Physiology A*, vol. 192, no. 8, pp. 777–784, Aug. 2006.

[240] D. Levenstein and M. Okun, "Logarithmically scaled, gamma distributed neuronal spiking," *The Journal of Physiology*, 2022, publisher: Wiley Online Library.

[241] E. L. Mackevicius, A. H. Bahle, A. H. Williams, S. Gu, N. I. Denisenko, M. S. Goldman, and M. S. Fee, "Unsupervised discovery of temporal sequences in high-dimensional datasets, with applications to neuroscience," *Elife*, vol. 8, p. e38471, 2019, publisher: eLife Sciences Publications, Ltd.

[242] A. H. Williams, T. H. Kim, F. Wang, S. Vyas, S. I. Ryu, K. V. Shenoy *et al.*, "Unsupervised discovery of demixed, low-dimensional neural dynamics across multiple timescales through tensor component analysis," *Neuron*, vol. 98, no. 6, pp. 1099–1115, 2018, publisher: Elsevier.

[243] L. M. Richter and J. Gjorgjieva, "Understanding neural circuit development through theory and models," *Current Opinion in Neurobiology*, vol. 46, pp. 39–47, Oct. 2017.

[244] C. Clopath, L. Büsing, E. Vasilaki, and W. Gerstner, "Connectivity reflects coding: a model of voltage-based STDP with homeostasis," *Nature Neuroscience*, vol. 13, no. 3, pp. 344–352, Mar. 2010, number: 3 Publisher: Nature Publishing Group.

[245] M. Loidolt, L. Rudelt, and V. Priesemann, "Sequence memory in recurrent neuronal network can develop without structured input," Sep. 2020, pages: 2020.09.15.297580 Section: New Results.

[246] A. M. Graybiel, "A satellite system of the superior colliculus: the parabigeminal nucleus and its projections to the superficial collicular layers," *Brain Research*, vol. 145, no. 2, pp. 365–374, Apr. 1978.

[247] Q. Wang, Y. Wang, H.-C. Kuo, P. Xie, X. Kuang, K. E. Hirokawa *et al.*, "Regional and cell-type-specific afferent and efferent projections of the mouse claustrum," *Cell Reports*, vol. 42, no. 2, Feb. 2023, publisher: Elsevier.

[248] B. Hellmann, M. Manns, and O. Güntürkün, "Nucleus isthmi, pars semilunaris as a key component of the tectofugal visual system in pigeons: Nucleus Semilunaris in Pigeons," *Journal of Comparative Neurology*, vol. 436, no. 2, pp. 153–166, Jul. 2001.

[249] N. R. Mahajan and S. P. Mysore, "Combinatorial Neural Inhibition for Stimulus Selection across Space," *Cell Reports*, vol. 25, no. 5, pp. 1158–1170.e9, Oct. 2018.

[250] A. Deichler, D. Carrasco, L. Lopez-Jury, T. Vega-Zuniga, N. Márquez, J. Mpodozis, and G. J. Marín, "A specialized reciprocal connectivity suggests a link between the mechanisms by which the superior colliculus and parabigeminal nucleus produce defensive behaviors in rodents," *Scientific Reports*, vol. 10, no. 1, p. 16220, Dec. 2020.

[251] R. Moreno-Bote, J. Rinzel, and N. Rubin, "Noise-Induced Alternations in an Attractor Network Model of Perceptual Bistability," *Journal of Neurophysiology*, vol. 98, no. 3, pp. 1125–1139, Sep. 2007, publisher: American Physiological Society.