UiO : University of Oslo

Mikkel Elle Lepperød

Dissecting neuronal circuits for navigation in experiments and models

Thesis submitted for the degree of Philosophiae Doctor

Institute of Basic Medical Sciences
Faculty of Medicine

Centre for Integrative Neuroplasticity (CINPLA)

2020

Series of dissertations submitted to the
Faculty of Medicine, University of Oslo

ISBN 978-82-8377-722-2

Cover: Hanne Baadsgaard Utigard.
Print production: Reprosentralen, University of Oslo.

Dedicated to my daughters Alicia and Leah

Preface

> Use the talents you possess, for the woods would be very silent if no birds sang except the best
>
> *Henry Van Dyke*

This thesis is submitted in partial fulfillment of the requirements for the degree of *Philosophiae Doctor* at the University of Oslo. The research presented here was conducted at the University of Oslo under the supervision of associate professor Torkel Hafting, Professor Marianne Fyhn, Professor Gaute Einevoll and associate professor Trygve Solstad. This work was supported by the Norwegian Research Council through grant 231248 and 217920.

The thesis is a collection of three papers, presented in chronological order of writing. The papers are preceded by an introductory chapter that relates them together and provides background information and motivation for the work. Two of the papers are joint work with Ane Charlotte Christensen and colleagues. I am the sole author of the remaining paper apart from senior authors.

Acknowledgements

I first want to thank my supervisors and mentors: Marianne Fyhn for always being there, listening, bringing support and enthusiasm, for being a role model in science and life. For teaching me experimental neuroscience and about the brain in general. Torkel Hafting for always being positive and understanding, for allowing me to pursue my own goals, for teaching me about neuroscience and the importance of never giving up on getting details "just right". For letting me evolve into an independent scientist and for showing that it is possible to do excellent science, be super fit and being a great dad at the same time. Gaute Einevoll for introducing me to computational neuroscience and believing that I was fit for a PhD, and for telling old stories of Norwegian climbing history and generally spreading good scientific standards combined with good mood and lots of jokes. Very good. Trygve Solstad for enthusiastic discussions on every aspect of neuroscience, and giving critical reviews of ideas and providing insights in computational neuroscience at large. Stefan Leutgeb for allowing me to stay in your lab and being a role model in how to conduct science and leading a focused group, and showing that it is possible to be on top on everything and at the same time having children and hobbies. Konrad Kording for teaching me about causal inference, how to conduct science and how to write scientific papers. You have been an immense inspiration, by being outspoken about the importance of doing open and clear research and of the ethics one often forget when having

"scientific tunnel vision". Remember, twitter is not reality. Marc de Kamps for letting me visit your lab, all our conversations and inviting me to collaborate on MIIND. I really enjoyed my stay and our collaboration and have learned allot from you.

Of my colleagues I especially want to thank: Kristian Lensjø for teaching me about neuroscience and physiology in general, for always taking time and for being part of our exclusive "wining club", it sure helps to co-wine in desperate times. Ane Charlotte Christensen for collaborating on two of the papers in this thesis, for having high scientific standards, always ready to discuss methods and results, or the latest articles and being a living library of entorhinal and hippocampus literature. Thank you for all the fun times at the office. Alessio Buccino for collaborating on all our coding projects, for showing me that it is possible to both get much done and do it great at the same time. For a great time in San Diego, and for all our time surfing together. Malin Bennum Røe for always helping out without expecting anything in return, and for always being fun to hang out with. Thank you for your self-distance. Svenn Arne Dragly for teaching me about the world of coding, forcing high coding standards and always ready to discuss or explore new ideas. Milad Mobarhan for teaching me high coding standards, and showing high fashion standards, for being inclusive and fun to hang out with, both during and off working hours. Simen Tennø for showing high moral during our drink.py CINPLA indoctrination, and for being fun to collaborate with on coding projects. Andreas Solbrå for being the strongest guy in CINPLA both in math and in lifting metal objects from the floor. Ivan Myhre Vinje for making UiO a fun and enjoyable place to work in, and spending time after work hours and for introducing me to the fantasy book genre. Alexander Stasik for all our discussions, for always being interested in crazy ideas. I always learn something new when we talk, and I look forward to working with you more closely in the coming years. Tristan Stöber for always being positive and enthusiastic about science, for being fun to hang out with and including everyone around you. In general the entire lab for making CINPLA a place where I look forward to coming to work every day.

I also want to thank my friends and colleagues at UCSD, Marta Sabariego, Anna-Lena Schlenner, Vincent Luczak and Sílvia Viana da Silva for making my family feel at home and including me in the lunch club and for all the fun time spent after work.

My father Einar Lepperød and mother Torgunn Elle for always being open minded and raising me in an environment where I have felt free to take my own decisions, my elder brother Espen Lepperød for being a role model in life, my little brother Mattis Lepperød for inspiring me and keeping with your principles.

My friend Lennart Nordenflycht for always being there with understanding and patience. And of course, my fellow adventurists and friends Adam Palmer, Jarle Haugan, Anders Rummelhoff. I hope we soon are traveling for our next adventure, a finished PhD shall entail more time climbing, going to concerts, or generally having more fun! My friend an colleague Eivind Hennestad for making studies, work and free time more interesting, I really enjoy all our discussions and I hope they will continue. Rest in peace Peter Holum Normann, wish you

were still here.

My soon to be wife Maria Hultman, for being the foundation in my life. I could not have done this without you, for being patient and understanding when times are rough, and constantly reminding me that life is more than science and achievements. For being the love of my life and being the mother of our beloved daughters. My daughters Alicia Hultman Lepperød and Leah Hultman Lepperød for reminding me what really matters in life.

Summary

Shaped by evolution, many brain areas and basic elements present in the human brain, are also found in many species such as rodents. This similarity gives an opportunity to study less complex brains as models for the human brain. Still, how neural networks ultimately govern perception, cognition, and action is one of the most elusive questions in contemporary neuroscience. Fortunately, some neural networks are recently shown to be governed by simple dynamical properties relative to their underlying complexity. This reduced dimensionality allows a combination of animal models and relatively simple mathematical models to understand information processing in the brain. The grid cell network, underlying elements of navigation, express low dimensional activity, however, the underlying mechanisms are still unknown. Whether the activity is governed by oscillations, projected spatial input or connectivity, and how activity remains stable across time and space remains to be revealed.

Using chronically implanted electrodes in the brain area of medial entorhinal cortex (MEC) in rats, we study grid cells that form hexagonal patterns of activity when the animal is searching for food rewards in an enclosed environment. In MEC, strong oscillatory activity, known as theta oscillations (4 - 12 Hz), is observed and is driven by pace making cells in the medial septal area (MSA). These oscillations are believed to play a crucial role in how the hexagonal pattern of grid cells emerge. To assess this proposed spatio-temporal relationship, we endow inhibitory neurons in MSA with light sensitive ion channels allowing precise control of oscillatory activity by pulsing laser light through an implanted optic fibre in MSA, a method known as optogenetics. In Paper I, we find that pacing oscillations at different frequencies, vanishes endogenous oscillatory activity, while the spatial pattern of grid cells remains. The results strongly indicate that spatial and temporal activity of grid cells can be dissociated, thus, oscillatory activity of grid cells do not underlie their hexagonal pattern.

Grid cells show exceptional stability of activity at the population level. This is proposedly due to stable connectivity underlying population activity which attracts and confines single neuron activity, essentially reducing the dimensionality of the population. We proposed that a specialized form of extracellular matrix called perineuronal nets (PNNs) contribute to stability in connectivity among grid cells. PNNs regulate the extent to which neurons can change their connectivity with other neurons, known as synaptic plasticity. In Paper II, we inject a bacterial enzyme, known as Chondroitinase ABC (chABC),

into MEC breaking down PNNs. We observe that the pairwise spatial correlations of grid cells and their pairwise temporal correlations were significantly reduced in animals lacking PNNs. Results strengthen our hypothesis that PNNs restrict plasticity in MEC and support the stable, low dimensional activity of grid cells.

Much of the literature concerning pairwise activity of grid cells together with the results presented in this thesis, suggest rigid connectivity to underlie their activity. Grid cells are identified based on their function (*in vivo*) and can thus not be directly studied in extracted brain tissue (*in vitro*). Conclusive evidence whether grid cells emerge due to connectivity or are e.g. subject to projected place specificity from other brain areas remains unresolved. Using pairwise recorded spiking activity to infer connectivity poses difficulty as correlations between neurons do not imply causal interactions. Naive inference of connectivity may thus reflect spurious correlation when background activity confounds the system under study. To overcome these challenges we attempted in Paper III to derive a method based on the instrumental variable technique, commonly used in econometrics. By combining pairwise recording of spikes with optogenetics the method allows inference of causal transmission probability between neurons. We find that the method shows promising results in an example of three neurons where it manages to remain causally valid while in naive methods validity is hampered by confounders.

In conclusion, the results in the present thesis contribute to the understanding of mechanisms that underlie brain activity and how this may explain cognitive functions such as navigation and memory processing. As with any scientific study, the research raises more questions that deserve further investigation.

Sammendrag

Formet av evolusjon, mange hjerneområder og grunnleggende elementer som er til stede i den menneskelige hjernen, finnes også i mange arter, som for eksempel gnagere. Denne likheten gir en mulighet til å studere mindre sammensatte hjerner som modeller for den menneskelige hjernen. Fortsatt, hvordan nevrale nettverk til slutt styrer persepsjon, erkjennelse og handling er et av de mest unnvikende spørsmålene i samtidens nevrovitenskap. Heldigvis er det nylig vist at noen nevrale nettverk styres av enkle dynamiske egenskaper i forhold til deres underliggende kompleksitet. Denne reduserte dimensjonalitet tillater en kombinasjon av dyremodeller og relativt enkle matematiske modeller for å forstå informasjonsbehandling i hjernen. Et spesielt nettverk, under fokus i denne avhandlingen som underligger navigasjonselementer, uttrykker lav dimensjonal aktivitet, men de underliggende mekanismene er fremdeles ukjente. Hvorvidt aktiviteten er styrt av svingninger, projisert stedsspesifisitet eller lokal tilkobling, og hvordan aktiviteten forblir stabil over tid og rom gjenstår å avsløre.

Ved å bruke kronisk implanterte elektroder i hjerneområdet medial entorhinal cortex (MEC) hos rotter, studerer vi gitterceller som danner sekskantede aktivitetsmønstre når dyret søker matbelønning i et lukket miljø. I MEC observeres sterk oscillerende aktivitet, kjent som teta-svingninger (4 - 12 Hz),

og drives av tempoet fra celler i det mediale septale området (MSA). Disse svingningene antas å spille en avgjørende rolle i hvordan det sekskantede mønsteret til gitterceller dukker opp. For å vurdere dette foreslåtte romtidsmessige forholdet, gir vi hemmende nevroner i MSA lysfølsomme ionekanaler som tillater presis kontroll av oscillerende aktivitet, ved å pulse laserlys gjennom en implantert optisk fiber i MSA, en metode kjent som optogenetikk. I artikkel 1 finner vi at stimuleringspulser ved forskjellige frekvenser fjærner endogen svingningsaktivitet, mens det romlige mønsteret til gitterceller forblir. Resultatene indikerer sterkt at romlig og tidsmessig aktivitet i gitterceller kan dissosieres, og derfor gir ikke oscillerende aktivitet til gitterceller grunnlag for deres sekskantede mønster.

Gitterceller viser eksepsjonell stabilitet av aktivitet på populasjonsnivå. Dette er antagelig på grunn av stabil underliggende tilkobling på tvers av celler som skaper tiltrekkning i populasjonsaktivitet og begrenser nevronaktivitet, noe som i vesentlig grad reduserer dimensjonaliteten til populasjonen. Vi foreslo at en spesialisert form for ekstracellulær matrise kalt perinevrale nett (PNN-er) bidrar til den stabile tilkoblingen blant gittercellene. PNN-er regulerer i hvilken grad nevroner kan endre sin forbindelse med andre nevroner, kjent som synaptisk plastisitet. I artikkel 2 injiserer vi et bakterieenzym, kjent som Chondroitinase ABC (chABC), i MEC som bryter ned PNN-er. Vi observerer at de parvise romlige korrelasjonene av gitterceller og deres parvise temporale korrelasjoner var betydelig redusert hos dyr som manglet PNN. Resultat styrker hypotesen vår om at PNN-er begrenser plastisiteten i MEC og støtter den stabile, lavdimensjonelle aktiviteten til gitterceller.

Mye av litteraturen om parvis aktivitet av gitterceller sammen med resultatene presentert i denne oppgaven, tyder på at bevarte tilkoblinger over tid ligger til grunn for deres aktivitet. Gitterceller identifiseres basert på deres funksjon (*in vivo*) og kan dermed ikke studeres direkte i ekstrahert hjernevev (*in vitro*). Bevis på om gitterceller oppstår på grunn av tilkobling eller er f.eks. projisert stedsspesifisitet fra andre hjerneområder forblir uavklart. Å bruke parvis registrert fyringsaktivitet for å utlede tilkoblingen gir vanskeligheter ettersom korrelasjoner mellom nevroner ikke innebærer årsaksinteraksjoner. Naiv inferens av tilkobling kan således reflektere falsk korrelasjon når bakgrunnsaktivitet konfunderer systemet som studeres. For å overvinne disse utfordringene forsøkte vi i artikkel 3 å utlede en metode basert på instrumentell variabel teknikk, vanligvis brukt i økonometri. Ved å kombinere parvis opptak av fyringsaktivitet med optogenetikk, tillater metoden inferens av kausal transmisjonssannsynlighet mellom nevroner. Vi finner at metoden viser lovende resultater i et eksempel på tre nevroner der den forblir årsaksmessig gyldig, mens i naive metoder hindres gyldigheten av konfundering.

Avslutningsvis bidrar resultatene i denne avhandlingen til forståelsen av mekanismer som ligger til grunn for hjerneaktivitet og hvordan dette kan forklare kognitive funksjoner som navigasjon og hukommelsesprosessering. Som med enhver vitenskapelig studie reiser forskningen flere spørsmål som fortjener videre undersøkelse.

List of Papers

Paper I

Lepperød, M.E., Christensen, A.C., Lensjø, K.K., Buccino, A.P., Yu, J., Fyhn, M., Hafting, T. 'Optogenetic pacing of medial septum PV cells disrupts temporal but not spatial firing in grid cells'. *Submitted for publication, in review.*

Paper II

Christensen, A.C., Lensjø, K.K., Lepperød, M.E., Dragly, S-A., Sutterud, H., Blackstad, J.S., Fyhn, M., Hafting, T. 'Perineuronal nets stablilize the grid cell network'. *Submitted for publication, invited for resubmission, minor revision.*

Paper III

Lepperød, M.E., Hafting, T., Fyhn, M., Kording, K.P. 'Inferring causal transmission probability from pairwise recordings and optogenetics'. *In progress.*

Contents

Preface iii

List of Papers ix

Contents xi

1 Background 1
- 1.1 Introduction . . . 1
- 1.2 Aims . . . 2
- 1.3 Neurons . . . 3
- 1.4 Neural networks . . . 10
- 1.5 Causality in neuroscience . . . 23
- 1.6 Estimating connectivity . . . 28

2 Objectives of Papers 31

3 Summary of Papers 33

4 Discussion 39

5 Future perspectives 49

6 Conclusion 53

Bibliography 55

Papers 76

I Optogenetic pacing of medial septum PV cells disrupts temporal but not spatial firing in grid cells 77

II Perineuronal nets stablilize the grid cell network 113

III Inferring causal transmission probability from pairwise recordings and optogenetics 159

Chapter 1

Background

> You', your joys and your sorrows, your memories and your ambitions, your sense of personal identity and free will, are in fact no more than the behavior of a vast assembly of nerve cells
>
> Francis Crick

1.1 Introduction

You are reading this by means of a system that has been shaped through millions of years of evolution - the human brain. Through eons of time, it has formed into the most complex entity in the known universe - or at least, so thinks the brain about the brain. Being isometrically scaled up from primates[1], the human brain is at the end of a long range of species struggling to survive with the means given by its ancestors. With no little amount of dramatic events, it has thus come to both fit inside your skull and be able to comprehend these words. As successful species survived there were little to no revision on parts that were already working sufficiently well[2]. Many of the brain regions that developed in early mammalian history are therefore highly preserved across species. For example, the brain area known as the hippocampus is found in nearly all vertebrate brains[3]. This gives us an opportunity to study simpler brains as relevant models of the human brain.

The brain is bounded in size, to support flexible navigation in complex environments and be manageable for mothers during birth, but at the same time be large enough to hold hardware (or wetware) that performs complex calculations. Compared with many other species, humans are born early in terms of development, completely reliable on caretakers years after birth. Being of comparable size to other organs in your body, the brain requires the most energy, consisting of about 86 billion neurons it draws about 20% of your life support. Each neuron has on average about 10 000 synaptic connections, making the brain a vast interconnected network of neurons.

Neurons communicate by means of electrical impulses, known as action potentials, governed by ionic currents that flow through ion channels - which is preserved throughout evolution in the vertebrate brain. From the vast complexity of molecules to the intricate interactions that enable humans to scale the steepest

[1] Azevedo et al., 2009 [2] Northcutt, 2002 [3] Witter et al., 2017

mountainsides of earth, brains operate on multiple scales. Spatially, all the way from picometer scale of molecular dynamics of ions, nanometer scaled ion channels, micrometer scaled neurons to the millimeter scale of neural networks, to meter scale brains. Temporally, ranging from nanosecond scale of ions moving through ion channels along an electrical or chemical gradient, millisecond scale of action potentials to networks storing memories spanning years. Due to the complexity of dealing with such a variety of spatio-temporal scales, the study of the brain spans multiple scientific fields, like biology, chemistry, philosophy, psychology, engineering, physics and mathematics.

Experiments in neuroscience are typically very complex, usually combining behavioral monitoring and various measurements of brain activity, recently in genetically modified animals. This involves both small-scale physics on particle interactions and electric fields, meso-scale complex systems-interactions of neurons and macro scaled interactions between brain regions. Zooming out even further is theory of learning and behavior, ranging from cognitive to social sciences. On each scale or level, complex computational models are used, of chemical interactions inside the cell, or single cell models using mathematical tools such as dynamical systems and network theory, or normative computational models of learning and behavior. This is why neuroscience is said to be one of the most difficult scientific fields currently undertaken by human kind - the endeavor of many brains to understand the brain.

The brain has been of interest for several millennia, with the earliest recorded reference to the brain about 30-25th century BC[4]. With the advent of technology such as the microscope, research has increased in intensity over the last century. Many brain areas have been cataloged, much of the nerves are traced, and the electrical activity of neurons is to a large part understood. But still, little is known about how the brain works as a whole, how the components in the brain co-orchestrate to give a stable and consistent integration of sensory input, produce coherent thoughts and plan and perform actions - ultimately providing the basis for consciousness. How then, are we to proceed scientifically in order to further understand the brain?

1.2 Aims

In this thesis I will outline the background and history to give context to three manuscripts which describes the work I have done during my PhD period in more detail. I introduce neurons and how their activity can be measured in living animals. I aim to give a brief overview of electrical activity in neurons, with focus on how they can be modeled mathematically. Moreover to show that seemingly complex systems of millions of interacting elements, the ion channels, give rise to a fairly simple low dimensional dynamical system, producing action potentials.

Then, I present neural networks and aim to show that at least in some particular brain systems, there is evidence of a similar low dimensional

[4]Kamp et al., 2011

representation as in the single neuron, but now, instead of ion channels, the acting elements are neurons communicating through action potentials, finally producing population-level representations of sensory input. With these ideas at hand I present one particular neural network, the grid cell network. I aim to provide theory together with experimental findings that indicates that this network indeed form a low dimensional representation of space and that it underlies an important element in self-localization, namely, path integration. Within this framework I present two papers, 1 and 2.

In Paper I we aim to perturb brain wide temporal activity while simultaneously monitoring grid cells to investigate whether temporal and spatial aspects in this cell type can be dissociated.

In Paper II we aim to assess the effect on stability of grid cells when removing an extracellular matrix surrounding cell bodies.

Questions in neuroscience are inherently causal, and it is easy to land on wrong conclusions when analysing observational data. Therefore, in the final part of the thesis I aim to introduce methods from causal inference that are commonly used in other fields such as economy and machine learning, but less so in neuroscience. With these methods at hand I aim to show that it is possible to form causal hypothesis about brain function.

In the final Paper III, we aim to use a particular method from causal inference, the instrumental variable technique, to infer causal interactions between neurons under large scale perturbations.

Finally I aim to discuss the work and conclude that the future of neuroscience can benefit from a population view in combination with causal inference in order to lure out the inner workings of the brain.

1.3 Neurons

How information was transmitted throughout the vertebrate brain was one of the primary questions that drove early neuroscience. The widely held view at the time was that the brain was made up of a single network of nerve fibers that were all directly connected to one another. In order to investigate this view, Camillo Golgi stained brain tissue with silver nitrate, labelling a random subset of neurons. With this sparse labeling, it was possible to see not only nerve fibres, but also nerve cells or neurons, however how they communicated remained unclear. It was still possible that neurons formed a continuum where each neuron's activity was continually altered by the activity of others. However, this seemed not to be the case as neurons was later shown to be disconnected by a synaptic cleft, indicated by Fridjof Nansen[5], later confirmed by Ramon Y. Cajal[6]. This finding had vast implications in neuroscience, and still has. It hinted to the single neuron as an independent computational unit, and gave rise to the neuron doctrine[7], which led to intensified invested research in single neuron dynamics.

[5]Nansen, 1887 [6]López-Muñoz et al., 2006 [7]Yuste, 2015

The brain is typically separated into cortical and subcortical areas, where cortical areas are subdivided into layers containing different neuronal types. Throughout the brain, neurons and non-neuronal glial cells make up the gray matter and axons make up the white matter which together form the main part of the brain tissue. Neurons and glial cells represent about the same amount of gray matter[8], where the former is believed to be the main component in information processing and the latter to be a support system. Neurons consist of three parts, a soma, an axon and dendrites (Fig. 1.1A). The dendrites represent the most complex structure of the neuron, and receives most of its input. Inputs are received mainly through chemical synapses, transmitting signal molecules called neurotransmitters between a presynaptic neuron and its postsynaptic recipient or with electrical synapses called gap junctions. In chemical synapses, synaptic transmission is activated by large deflections in the membrane potential of neurons called action potentials (or spikes), generated in the soma of a presynaptic neuron (the axon hillock to be precise), travelling down axons, and ultimately reaching synapses activating transmission of neurotransmitters (Fig. 1.1A). Neurons typically have either inhibitory or excitatory synapses, in which neurotransmitters activate ion channels in the membrane that decrease or increase the postsynaptic membrane potential (PSP) respectively (Fig. 1.1C excitatory PSP). The most common neurotransmitters are γ-aminobutyric acid (GABA) and glutamate for inhibitory and excitatory synapses, respectively. Most of the input are integrated in dendrites, but synapses can also be found directly on the soma, depending on the type of presynaptic neuron.

1.3.1 Measuring neural activity

Neurons communicate through action potentials which are recognizable as spikes in the extracellular potential surpassing the amplitude of the (baseline) chatter generated by background neural activity. In the work presented in this thesis we have measured brain activity by extracellular electrophysiology with tetrodes, and in what follows I introduce the basic concepts underlying this measurement method.

After synaptic activation, currents travel through the membrane. These currents are made up of charged ions that enter the dendritic tree or soma due to opening of ion channels. When ions enter the top of a dendritic tree and travel down towards the soma, an electric sink is generated at the input and a source can be found in the soma. This activity, generates an electrical field which can be picked up by electrodes placed in the extracellular space. If neurons are aligned, such that the dendrites stretches in the same general direction, and receives synchronous input, the fields are superposed and an even larger deflection in potential can be measured. When recording from the brain extracellularly, synchronous population activity gives rise the local field potential (LFP), which can be observed with low-pass filtering of the signals as their frequency content is below 300 Hz, (Fig. 1.2B).

[8] Bartheld et al., 2016

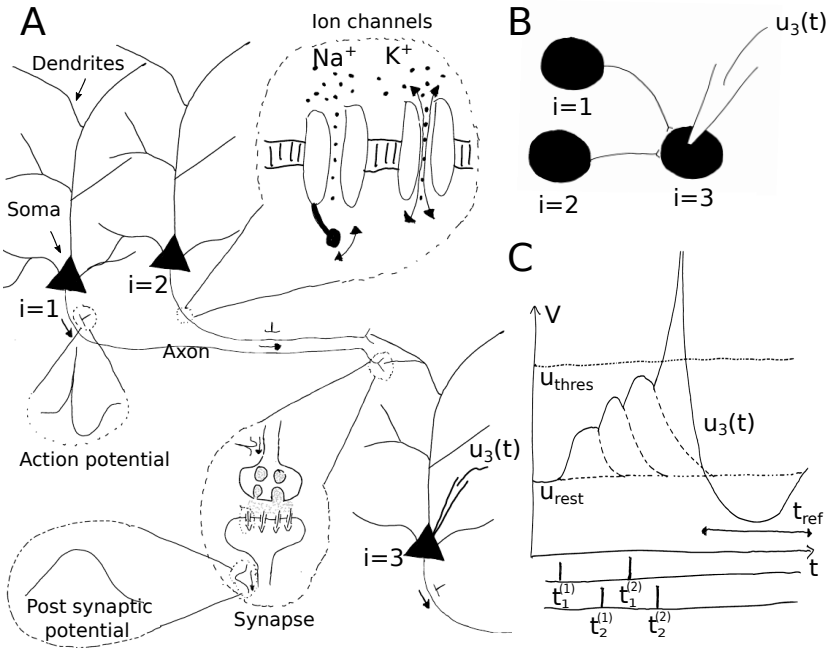

Figure 1.1: **Basic intracellular neurophysiology.** **A** Three pyramidal neurons in a small network (i=1-3). Action potentials travel down the axon and induce release of neurotransmitters in the synapse. Postsynaptic potential (PSP) then rise as a result of neurotransmitters opening ion channels in the synapse. The two main ion channels involved in the action potential of neurons are voltage gated sodium (Na^+) channels and potassium (K^+) channels. **B** Simplified circuit of the three neurons to the left as conceptualized with point neuron models, with an electrode measuring the membrane potential of neuron i=3. **C** Membrane potential ($u_3(t)$) measured at the soma of neuron $i = 3$ upon spikes t_i^k received from neurons $i = 1, 2$. Resting membrane potential (u_{rest}), threshold for action potential (u_{thres}) and refractory time t_{ref}

When the signal is high pass filtered (> 300 Hz), the activity obtained is referred to as multi unit activity (MUA), a signal containing the spiking activity of many neurons (Fig. 1.2B). This signal consists of spiking activity (and high frequency noise) from all neurons in a close vicinity[9]. In order to recognize and quantify single neuron activity relative to behavior we have to distinguish which spikes are being measured from which neuron. To this end, we preferably have as many electrode channels as possible in close vicinity, to get a detailed "electrical view" of the nervous tissue. With multiple channels measuring one neuron, it is possible to better identify individual neurons (Fig. 1.2B). Multiple adjacent recording sites are then used to distinguish which spikes comes from which neurons in a process called spike sorting[i]. As the amplitude of the electrical potential falls quickly with distance, it is possible to distinguish neurons by their amplitude across electrodes (Fig. 1.2B). Moreover, it is possible to separate inhibitory and excitatory neurons from their shape of the extracellularly recorded action potential[10]. However, to ever increase the number and density of electrode channels is technologically challenging, both in terms of fitting them into a tiny brain, but also amplification and storage of the numerous channels. In addition, an extra complication is implanting electrodes that can stay fixed in one brain area and drift minimally. Therefore, bundles of four fine wires called tetrodes are still often preferred when recording from freely behaving rats over long periods of time. Being on the lower range of channel count they can still isolate single neurons, however, to a less extent than with more recently developed high density electrodes e.g. Neuropixels[11].

1.3.2 Modelling the neuron

The proteins within the neuronal membrane known as ion channels allow ions to flow across the membrane (Fig. 1.1A). Transport molecules use energy to pump ions against their electrochmical gradient and ion channels allow bulk flow of ions down their gradient. Some of the ion channels are passive (open all the time), some are active (open and close as a function of membrane potential, or chemical species such as calcium or neurotransmitters), however it is not the single ion channel that gives rise to neuron dynamics - it is the collective activity of thousands of ion channels. Therefore, it is more reasonable to speak about populations of ion channels and which currents they give rise to. The existence of ion channels was not known in early history of neuroscience as they were simply too small to study. Hodgkin and Huxley[13], thus studied the total currents flowing in a closed system using a voltage clamp. This method allowed for clamping the membrane potential to a steady value by combining one electrode to inject currents and another to record the corresponding membrane potential. Using long wires along the axon, they ensured that no axial current

[9]Buzsáki, 2004 [10]Barthó et al., 2004 [11]Jun et al., 2017 [13]Hodgkin and Huxley, 1952

[i]For details on how spike sorting was performed in the works presented in this thesis see the methods sections of the respective manuscripts.

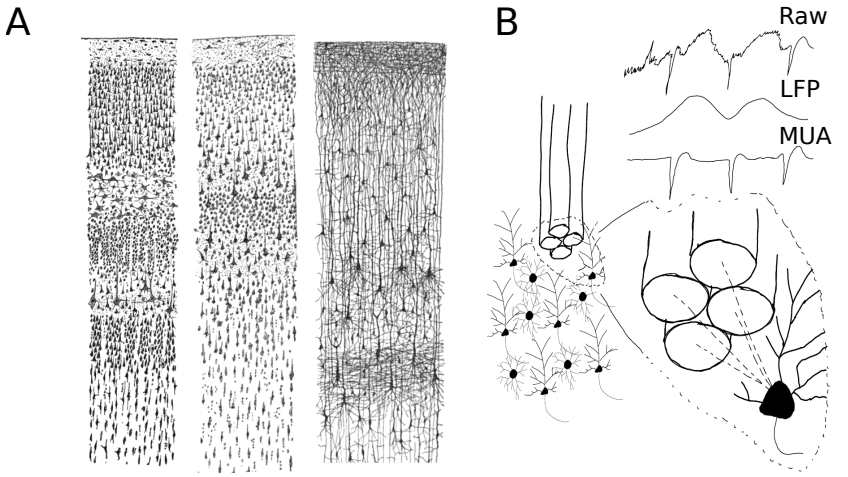

Figure 1.2: **Basic extracellular neurophysiology. A** Neurons in different cortical layers, left: visual cortex of a human adult, middle: motor cortex of a human adult, right: cortex of a 1 1/2 month old infant, adapted from Cajal[12]. **B** Tetrode, consisting of 4 electrodes, recording neuronal activity. The raw signal recorded from tetrodes is low pass filtered to obtain the local field potential (LFP) and high pass filtered to obtain multi unit activity (MUA). To separate the signal from different neurons from the MUA, the 4 separate channels on the tetrode is finally triangulated by amplitude.

occurred, known as space clamp configuration. They varied the voltage across the membrane and measured the current flow of different ions to establish a mathematical model explaining the mechanisms underlying the action potential. In other words, by using system-wide perturbations they were able to understand activity of populations of ion channels that make up the dynamics of the system at large. Finally, they reached the conclusion that the action potential was governed by the following equations:

$$C\dot{V} = I - I_{Na} - I_K - I_L, \tag{1.1}$$

where

$$\begin{aligned} I_{Na} &= \bar{g}_{Na} m^3 h \, (V - E_{Na}) \\ I_K &= \bar{g}_K n^4 \, (V - E_K) \\ I_L &= g_L \, (V - E_L) \end{aligned}$$

and

$$\begin{aligned} \dot{n} &= \alpha_n(V)(1-n) - \beta_n(V)n \\ \dot{m} &= \alpha_m(V)(1-m) - \beta_m(V)m \\ \dot{h} &= \alpha_h(V)(1-h) - \beta_h(V)h \end{aligned}$$

Here, $\dot{V}$ is the time derivative of the membrane potential, C is the membrane capacitance, I is the total ionic current, I_x is the specific current of ionic species $x \in \{Na, K\}$ or of leak L. The functions α, β describe the transition rates between open and closed states of the channels which were fitted to experimental results. The main point for showing the equations here is the fact that the systems state is uniquely determined by the membrane potential V and the gating variables n, m and h. This dynamical system is described by its state and the law given in Eq. (1.1) determining how the state variables evolves over time.

The millions of elements that make up the axon were thus described with high accuracy by only a small set of equations. This is known as dimensionality reduction. Typically, in systems with many interacting elements, the activity from a vantage point is contained in only a few dimensions. If each membrane channel (dimension) is taken into account it can be very hard to understand activity at a larger scale. Taken together, their activity, if acting in concert, might hold a much simpler story. For instance, the Hodgkin and Huxley model is given by 4 dimensions V, n, m and h.

Intense effort was invested by Hodgkin and Huxley to unveil the dynamics of the axon and how it conducts action potentials. Paradoxically, in modern computational neuroscience the axon is typically not included in network models for just this reason - we know enough about the axon to confidently model it as a simple transmission delay. However, the same dynamics can also be used to model the neuronal membrane and laid the foundation for using point neuron models in models of neural networks (Fig. 1.1B). Here, the dendritic structure and axons are ignored while only the synapse and neuronal membrane potential is modeled[ii]. By studying how neurons enter the spiking event, known as bifurcations, it was later understood that spiking behavior can be governed by just a few types of bifurcations. More recent models have shown that neuron's spiking dynamics can also be reduced to just a two coupled equations[15,16].

1.3.3 Postsynaptic potentials and integrate-and-fire neurons

By using an intracellular electrode it is possible to measure the difference in electrical potential between the interior and the exterior of the membrane which we will denote $u(t)$ at time t (Fig. 1.1). When neuron i recieves no input, it has a resting potential u_{rest} and upon an input, say from neuron j, the PSP is given by $\epsilon_{ij}(t) = u_i(t) - u_{rest}$. Following Gerstner et al.[16], when receiving multiple spikes k from multiple neurons arriving at time t_j^k the membrane potential is given by the sum

$$u_i(t) = \sum_j \sum_k \epsilon_{ij}(t - t_j^k) + u_{rest} \tag{1.2}$$

[15] Izhikevich, 2007 [16] Gerstner et al., 2014

[ii] To ignore the dendritic structure is not trivial as with the axon, and recent work has shown that in some neurons, dendrites can perform non linear integration which might have a significant impact on network computations[14].

The action potential has a remarkably stable shape across neurons, indicating that information is not carried by the shape of the action potential. Rather, information is carried by whether or not a spike is present, the inter arrival time and number of spikes received per unit time (rate). Therefore, the widely used, integrate-and-fire models of neurons does not model the shape of the action potential as in the Hodgkin and Huxley model, but just the time in which it is received and generated. When the membrane potential reaches a threshold u_{thres}, the neuron i is said to fire a spike, the time t_i^k is stored and the membrane potential is reset to rest $u(t) \rightarrow u_{rest}$. Action potentials are governed by active ion channels which requires some time to open and close, after an action potential is generated, there is a time period where the neuron is unable to generate a new spike, known as the refrectory period t_{ref}. In the integrate-and-fire models the refractory period is typically modeled "brute-force style" by rejecting all input to the respective neurons just after a spike event.

The most simple neuron model is the leaky integrate-and-fire model, which sees the neuronal membrane as a simple RC circuit and is given by

$$\tau_m \frac{\mathrm{d}u}{\mathrm{d}t} = -(u - u_{rest}) + RI \tag{1.3}$$

Here $\tau_m = RC$ is the membrane time constant, R is the resistance, and C capacitance. I is an input current which can represent a synaptic input or any other external currents (e.g. stimulation currents through an intracellular electrode).

In Eq. (1.3) the right hand side is linear. Therefore, when combined with the firing threshold u_{thres}, the potential u returns to rest from any value $u < u_{thres}$ when no stimulus is applied ($I = 0$). In real neurons however, the upswing towards threshold is nonlinear[17] and is much better matched by the exponential integrate-and-fire neuron[16,17].

Being a variant of the leaky integrate-and-fire neuron, the exponential integrate-and-fire model is given by

$$\tau_m \frac{\mathrm{d}u}{\mathrm{d}t} = -(u - u_{rest}) + \Delta_T \exp\left(\frac{u - u_{rh}}{\Delta_T}\right) + RI. \tag{1.4}$$

Here, Δ_T is a sharpness parameter which determines how the neuron enters the spiking event where the sodium activation variable are approximated by an exponential function[18]. The "rehobase threshold" u_{rh} marks the voltage level in which the exponential term becomes positive and thus the action potential is initiated, in addition, a hard "cutoff threshold" u_{thres} is also used as the level in which the potential is reset to rest $u(t) \rightarrow u_{rest}$.

Even though it is less physiologically accurate, the the leaky integrate-and-fire neuron is very simple and fast to simulate. Therefore it is often used in studies that investigate general aspects of interactions in neural networks. For these reasons we use the leaky integrate-and-fire neuron in Paper III. When

[17]Badel et al., 2008 [16]Gerstner et al., 2014 [17]Badel et al., 2008 [18]Fourcaud-Trocmé et al., 2003

simulations are supposed to match experiments more closely, models that include more of the details observed in experiments are often used. In Paper II, we use the exponential integrate-and-fire for those reasons, which is shown to match physiological properties of the stellate cell in medial entorhinal cortex[19].

1.3.4 Synapse models

In synaptic transmission, a presynaptic spike depolarizes the synaptic terminal, leading to an influx of calcium through presynaptic calcium channels which cause neurotransmitters to be released into the synaptic cleft. The neurotransmitters bind to postsynaptic channels which open and allow ionic currents to flow across the membrane. Modeling the entire process mathematically would be rather tedious and is typically circumvented by describing transmitter-activated ion channels as time-dependent conductivity $g_{syn}(t)$ similarly as in the Hodgkin and Huxley model. When a family of ion-channels open, a current pass through the membrane which depends on the difference between its reversal potential E_{syn} and the membrane potential u given by

$$I_{syn} = g_{syn}(t)(u(t) - E_{syn}) \tag{1.5}$$

For inhibitory synapses E_{syn} is typically negative and for excitatory synapses $E_{syn} = 0$. Now, the function $g_{syn}(t)$ can be used to model the response dynamics of synapses and ultimately provide bases for the PSP.

In Paper II we use these conductance-based synapses in a model of a neural network that govern navigation - a grid cell model. In Paper III we use a simplified version of Eq. (1.5) which is not dependent on the membrane potential and g is considered a current rather than a conductance, known as current-based synapses.

1.4 Neural networks

Investigating the workings of single neurons might give a hint to how we may proceed to understand some neural networks. By using systems wide perturbations, a single neuron's means of communication was understood by a quite simplistic model, given the vast complexity of its underlying elements - the Hodgkin Huxley model.

Given a brain with restricted resources, utilizing individual neurons to represent external variables can give high information yield at low cost[iii] by representing N dimensions with N neurons - each neuronal state represent a variable in each dimension. However, in highly unpredictable environments, such systems can be vulnerable to noise, such as abrupt sensory input e.g. induced by a sudden presence of danger, as there is no restoring dynamics. Therefore, it is not

[19] Pastoll et al., 2013

[iii] Where cost is defined in terms of energy expenditure, e.g. using one neuron to represent one item cost less energy than using two neurons to represent one item.

unlikely that through evolution, some brain systems, such as systems underlying integration or memory, have settled on a trade-off between representational dimensionality and stability, by for instance utilizing attractive states that arise in dynamical systems[20]. Moreover, given the fast time constant of single neurons it can be difficult to account for the slow timescale of behavior. A single neuron integrates complex input and affects other neurons in its network. It contributes to a population activity which in turn affects single neuron activity. Much is known about single neurons, which has been studied for about a century, however, how the interplay in large populations of neurons is governed, is still under heavy investigation in contemporary neuroscience.

Below I will introduce a neural network which is the main focus in this thesis, a network that presumably governs an important aspect of navigation in rodents - the grid cell network.

1.4.1 Navigation

Being predator or prey, self-localization is of primary importance for an animal, and it is not unlikely that stable systems excelling at this task have been favored through natural selection. When entering a novel environment one cannot rely on previously learned relations to environmental cues. In order to have spatially stable neuronal representations of the environment, the animal must somehow simultaneously deduce its position relative to cues and keep track of location relative to its outset.

There are at least two ways in which to self-localize. One is to carefully observe the surroundings and generate a map. Then in order to self localize one can triangulate relative to environmental cues, similar to celestial navigation. Another way, is to keep track of direction and distance traveled. Imagine walking with your eyes shut, if you know the initial position, direction and speed, it is possible to sum up distance traveled in all directions and then infer your final destination, known as dead-reckoning or path integration.

It is reasonable to believe that both these methods are utilized in animals, and there is evidence of both. In early work from Tolman, rats were subjected to a maze with only one northern entry point leading to a goal in form of a treat at a north-eastern point relative to the starting location, the animal was familiarized with the maze and its surroundings (Fig. 1.3A left). Then, after 4 days of training, the maze was exchanged with a sunburst maze having multiple entries, but with the northern entry closed (Fig. 1.3A right). Which direction did the rat choose? If it associated the entry point with the goal, it would choose entries adjacent to the northern entry (denoted R in Fig. 1.3A right), if however, it associated the general direction towards the goal it would choose the entry that led in that direction (denoted R in Fig. 1.3A right). With probability above chance levels, the rat chose the direction in which laid on a straight line towards the goal. With this experiment, Tolman[21] concluded that the animal must keep

[20] Yoon et al., 2013 [21] Tolman et al., 1946

some abstract representation of the environment which he termed a cognitive map.

In another experiment[22], the strong maternal instinct in gerbils was exploited. By moving a pup out of its nest, the mother went searching for its lost pup, being in complete darkness it had to rely on path integration in order to find its way home (Fig. 1.3B). To make sure the mother did not find its way home by odor, the nest was removed as soon as the mother ventured out, and it did indeed return to where the old nest was located to begin with.

Together, these experiments indicated that rodents can use both methods of localization, path integration and mapping distal cues. Moreover, it is likely that they use them in combination. In order to use an internal map, it has to first be generated. When only using path integration, errors accumulate during navigation as small errors are integrated and accumulates over time. This problem is known in the robotics literature as simultaneous localization and mapping (SLAM) and is an active area of research.

1.4.2 Place cells and grid cells

How are these navigational issues solved in the rat brain? The area known as the hippocampus was implied to play a crucial role in spatial memory after the seminal work by Milner[23] on the patient Henry Molaison (HM). Patient HM had his hippocampi bilatterally lesioned to treat severe epilepsy. After surgery, he was unable to form new episodic memories of the type what, where and when. Later, O'Keefe and Dostrovsky[24] began an electrophysiological survey of the hippocampus in freely moving rats. With chronically implanted tetrodes in the hippocampus CA1 in rats (Fig. 1.3C-D) it was possible to monitor neuron activity when the animal was foraging for food in an enclosed environment. Indeed, neurons with positional preferences were identified, termed place cells, as their spikes were clustered in a few particular areas in two dimensional space, termed place fields (Fig. 1.3G).

Since different place cells had place fields at different locations together with the previously established relation between hippocampus and memory, O'Keefe and Nadel[25] hypothesized that place cells formed the cognitive map postulated by Tolman - as abstract representations of space forming the context which episodic memories are embedded in. Subsequent work showed that place cells were under the control of distal visual landmarks[26], but remained stable when these cues were removed[27] or the animal was placed in the dark[28]. Place fields are thus not just simple correlates of spatial cues or specific external landmarks. Place cells are, to a large degree, unrelated to specific sensory inputs and thus seems to represent more abstract conceptions of space[29].

Determinants of stability in place fields were hypothesized to be due to path integration[25]. In order to update the map, information about spatial movement

[22]Mittelstaedt and Mittelstaedt, 1980 [23]Scoville and Milner, 1957 [24]O'Keefe and Dostrovsky, 1971 [25]O'keefe and Nadel, 1978 [26]O'Keefe and Conway, 1978 [27]O'Keefe and Speakman, 1987 [28]Quirk et al., 1990 [29]Knierim and Hamilton, 2011 [25]O'keefe and Nadel, 1978

and objects in the surrounding seems necessary. The entorhinal cortical (EC) area was thus investigated as it has strong projections to the hippocampus (Fig. 1.3E). Indeed, neurons that had representations in space was later found in the Medial and Lateral areas of EC termed MEC and LEC respectively. While LEC show neural correlates about objects in space, MEC was found to hold information about movement in space.

Investigating the upstream cortical area of hippocampus, cells with multiple firing fields were identified in the MEC[30–32]. These neurons were termed grid cells as their triangular patterns tessellate the explored environment forming hexagonal grids with firing fields at the vertices (Fig. 1.3F). Grid cells have from the outset been suspected to perform path integration[31], and does indeed provide input to the hippocampus[33].

Large changes in either environmental cues[26,34] or sensory inputs[35] can alter locations of place fields, termed remapping[36,37] - which is thought to be the cellular substrate for discriminating similar memories as well as context-specific memories[29]. Once the place fields settle, place cells show stability in terms of distance traveled, independent of sensory input[35,38]. However, the spatial locations of grid fields seems relatively insensitive to particulars in the environment, such as stability of their fields in cue-poor environments and darkness[31], while spatial responses in LEC and the hippocampus are more prone to change upon environmental manipulations[29,39]. Grid cells can rotate together with salient external cues[31], resize their scale in response to rescaling of a familiar environment[40] and show firing-rate modulations[41]. However, the relative insensitivity of grid cells to external cues suggest that self-motion is the primary determinant of grid cell firing. Therefore, it is likely that the grid cells compute or is subject to a path integrated estimate of the animal's position. For example, self-motion cues, such as the vestibular senses, optic flow, and proprioception, provide information about velocity that is integrated over time to yield an estimate of position. However, direct evidence of the role of grid cells in path integration is still lacking.

1.4.3 Continuous attractor model of grid cells

When placed in different environments, neighboring populations of grid cells show a coherent shift of grid fields[42] that retains their relative spatio-temporal relationships[20]. In contrast, place cells act more independently and remap individually and relatively unpredictable. Moreover, grid cells are different from place cells in that they have tightly organized multiple firing fields and seems to form a population code with distinct modularity[40,43], spread over the dorso-ventral (DV) axis of MEC (Fig. 1.4B). Within one module, grid cells show

[30]Fyhn et al., 2004 [31]Hafting et al., 2005 [32]Sargolini et al., 2006 [31]Hafting et al., 2005 [33]Rowland et al., 2018 [26]O'Keefe and Conway, 1978 [34]O'Keefe and Burgess, 1996 [35]Ravassard et al., 2013 [36]Kubie, 1983 [37]Muller and Kubie, 1987 [29]Knierim and Hamilton, 2011 [35]Ravassard et al., 2013 [38]Gothard et al., 1996 [31]Hafting et al., 2005 [29]Knierim and Hamilton, 2011 [39]Derdikman and Knierim, 2014 [31]Hafting et al., 2005 [40]Barry et al., 2007 [41]Savelli et al., 2008 [42]Fyhn et al., 2007 [20]Yoon et al., 2013 [40]Barry et al., 2007 [43]Stensola et al., 2012

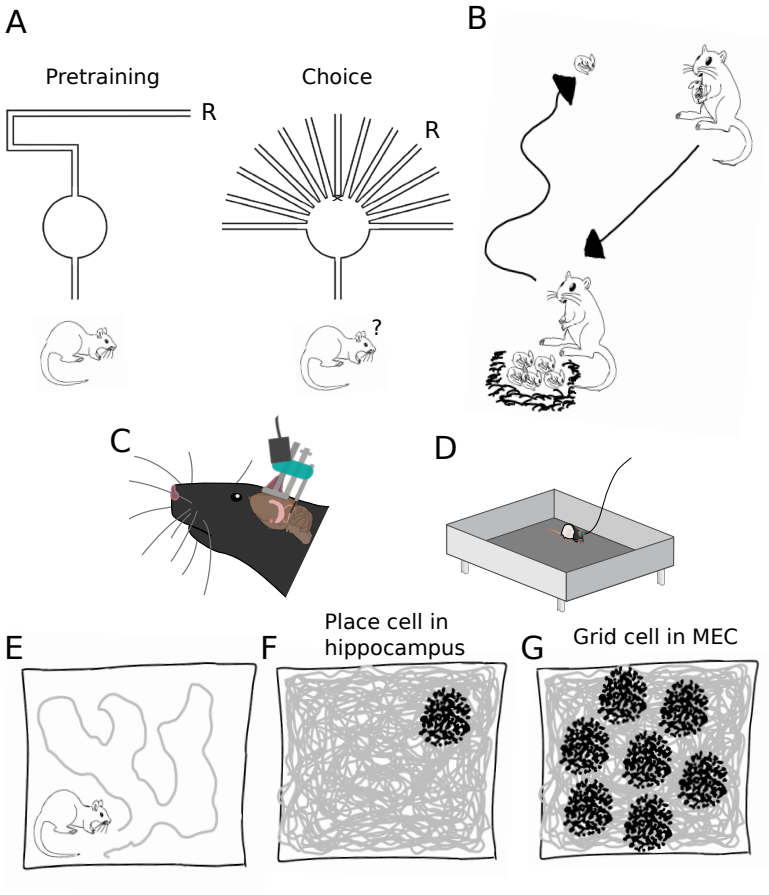

Figure 1.3: **Self localization in mammals**. **A** Experiment performed by Tolman[21] showing that rats can perform map-based navigation. In the pretraining maze (left) the animal learn the location of a reward (R). After four days the maze is changed to the choice version (right), where there is multiple exits, while the exit that was available during training is closed. If rats associate the general direction of the reward they choose a rightwards directing exit. **B** Experiment performed by Mittelstaedt[22] showing that gerbils can navigate by means of path integration. A pup was removed from the nest, after the mother locates its pup she managed to find the way home without relying on landmarks, as the test was performed in darkness, or by odor, as the nest was relocated after the mother ventured out. **C** By placing tetrodes, held with a microdrive, in the hippocampus or MEC of rats it is possible to record the extracellular activity of neurons. **D** The microdrive is conneced to an acquisition system while the animal explores its environment, in this case a square box. **E** When recording the postion simultaneously with neuron activity it is possible to find **F** place cells in the hippocampus and **G** grid cells in MEC.

similar spacing, orientation and field size, differing only in their relative phase (Fig. 1.4B). This indicates that grid cells are part of a larger population dynamic, each module forming a distinct population, in line with continuous attractor models of grid cells.

One way to impose stability in a dynamical system of neurons is to organize the coupling such that the system settles in a stable state. These states are called attractive states when the system can withstand perturbations by settling back to the same state. Many types of attractors can be found, determined by the dynamics of the system under study. In some systems, stable states of patterns form through instabilities or bifurcations.

Much like all the ion channels that gives rise to a collective dynamic in one neuron generating the action potential, here many neurons give a collective dynamic of the network state. Although, the comparison does not hold for long as ion channels differ from neurons in many ways, in particular, neurons are connected and communicate in complex ways, such as with different types of synapses (inhibitory, excitatory, slow and fast etc.). How the connections make up the network can determine the ways neurons are "allowed" to behave or self organize. Early work with neural networks assumed the neurons made up a continuous neural sheet - termed neural fields, which is still an active area of research. This excitable medium is then given communication rules by means of a connectivity kernel that can give rise to a plethora of patterns[44]. The type of pattern-forming network typically used in order to model the population dynamics of grid cells, known as continuous attractor networks (CANs), stems back to early work by Wilson and Cowan[45,46] and Amari[47] (see Deco et al.[48] for a review). Following Ocko et al.[49], a CAN model is typically of the form

$$\frac{\mathrm{d}s(u)}{\mathrm{d}t} = -\frac{s(u)}{\tau_m} + G\left(\int_\Omega W(u,u')s(u)du'\right). \tag{1.6}$$

Here $s(u)$ is the activity of a neuron at position u which lives in the continuous space $\Omega \in \mathrm{R}$, while τ_m is the neuronal membrane time constant. Moreover, G is a nonlinear activity function, typically a rectifying function $G(x) = max(0, x)$, and W gives the connection strength between u and u' which is typically translation invariant i.e. $W(u, u') = J(u - u')$.

When J has long range inhibition and nearby excitation e.g. a Mexican hat function, the system can sustain stable attractor "bumps" (Fig. 1.4C). Early work introduced these models with Ω as a periodic ring to describe the head direction system[50]. With Ω as a 2D sheet, the system Eq. (1.6) can settle into an hexagonal pattern. Much like the Turing pattern formation[51], in which a system of two chemical species self organize into an hexagonal pattern where the inhibitory diffusion rate is larger than an excitatory diffusion rate. After the discovery of grid cells, these models were adapted to model grid cells[52–55]. From the perspective of a neuron at position u it allows nearby neurons to be

[44]Ermentrout, 1998 [45]Wilson and Cowan, 1972 [46]Wilson and Cowan, 1973 [47]Amari, 1977 [48]Deco et al., 2008 [49]Ocko et al., 2018 [50]Zhang, 1996 [51]Turing, 1952 [52]Samsonovich and McNaughton, 1997 [53]Fuhs, 2006 [54]McNaughton et al., 2006 [55]Burak and Fiete, 2009

active, but inhibits those in less proximity, until the end of its reach. There, other neurons may be active and a self-organization occurs, each position tries to dominate a local ring, and the optimal[56] distribution of these rings is a hexagonal pattern which the system settles into (Fig. 1.4C).

There are several ways to impose a hexagonal pattern in a CAN model, the most common is using multiple bumps in a periodic or toroidal sheet, and a-periodic sheets[55], or a single bump in a twisted torus topology[57] (see Zilli[58] for a review).

In Paper II we use the exponential integrate-and-fire neuron in a multi bump CAN model to investigate the effects of changes in excitatory to inhibitory synaptic strength and changes in capacitance on inhibitory neurons.

1.4.4 Using grid cells to estimate position

When recording dorsal to more ventral locations of MEC, grid cells increase in scale[61] (Fig. 1.4A left). The populations along the dorso-ventral axis, when combined, can theoretically provide a position code and a distance between two locations[62]. Moreover, decoding location can be done optimally when the grid scales are independent modules[59,63], indeed, the grid scales are discretized[43]. With a position code at hand it is further possible to estimate a distance vector given by the difference between two position vectors[59]. This position vector is hypothesized to give the basis for vector-based navigation[64] which is shown in a normative model of the grid cell network to aid learning in navigational problems[65].

Whether grid cells stem from place cells or the other way around is a much debated topic. Grid cells are indicated to be primary determinants of place cell firing[60,66,67] (Fig. 1.4A). Place cells form before grid cells in the developing rat[68,69], and when inactivating the hippocampus grid cells disappears[70]. Moreover, given the projections from hippocampus to MEC[3], it is not unlikely that grid cells can be dependent on place cell input. Following this line of thought, other models of grid cells derive the pattern from place cell input, where path integration is typically thought to be done in place cells. Kropff and Treves[71] considered grid cells to strengthen their synaptic input from place cells when co-active. A sparse representation was then forced through firing rate adaptation in grid cells leading to optimization of grid field packing in 2D and thus a hexagonal pattern. This model was later updated to conform to the population activity of grid cells forming a continuous attractor[72]. Dordek et al.[73] considered a single layer feed-forward neural network with Hebbian learning (strengthening co active neurons) where grid cells formed when a non-negative input was forced and it was noted that the architecture resembles principal component analysis

[56]Chang and Wang, 2010 [55]Burak and Fiete, 2009 [57]Guanella et al., 2007 [58]Zilli, 2011 [61]Brun et al., 2008 [62]Fiete et al., 2008 [59]Stemmler et al., 2015 [63]Mathis et al., 2012 [43]Stensola et al., 2012 [59]Stemmler et al., 2015 [64]Bush et al., 2015 [65]Banino et al., 2018 [60]Solstad et al., 2006 [66]Gil et al., 2017 [67]Hales et al., 2014 [68]Langston et al., 2010 [69]Wills et al., 2010 [70]Bonnevie et al., 2013 [3]Witter et al., 2017 [71]Kropff and Treves, 2008 [72]Si and Treves, 2013 [73]Dordek et al., 2016

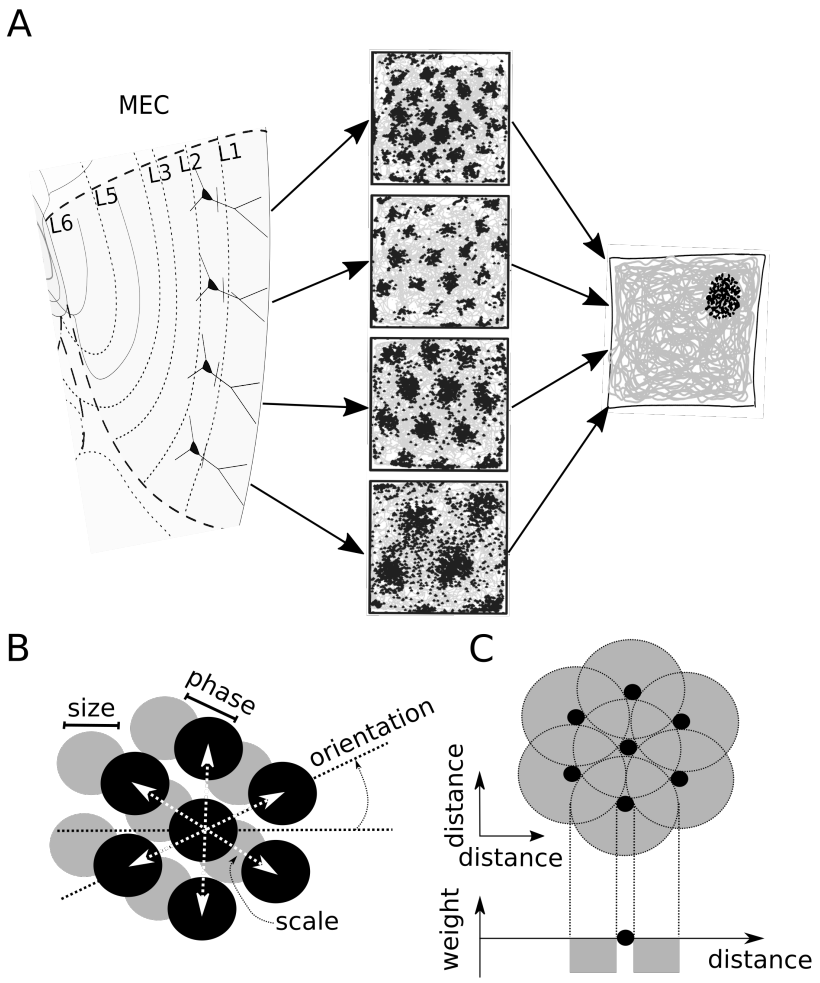

Figure 1.4: **Position estimates from grid cells and the principle of the CAN model A** Sagittal view of MEC with four neurons in Layer 2 highlighted and their activity from a recording when the rat was running in a square box. To the right a hypothetical place cell based on input from these four neurons. Grid cells in the same location (module) have similar spacing, but the scale of grid cells increase along the dorso-ventral axis.[43], By combining these modules it is possible to accurately decode position[59], as indicated with arrows ending in a positional code, which can be interpreted as place cells[60]. **B** Within a module, grid cells share scale, size and orientation but is shifted in phase. **C** Simple illustration on how to achieve grid cells with CAN models. Within a neural sheet (distance x distance in neural space) inhibitory connectivity is distance dependent making up a doughnut around neurons (distant inhibition). With this local inhibitory connectivity, neurons self organize in activity and form an hexagonal pattern (black dots).

(PCA). Stachenfeld et al.[74] used reinforcement learning where the hippocampus was considered a predictive map and grid cells then formed in eigenvectors of the place cell activity also with a non-negativity constraint. Banino et al.[65] and Cueva and Wei[75] used velocity input in a recurrent neural network and obtained both place like activity and grid like activity. Moreover, in an attempt to unify these different views Sorscher et al.[76] proposed that grid cells form through pattern formation dynamics that arise when error in estimated position from place-field-like structures is minimized and a non-negativity constraint is forced. All in all, it is not unlikely that information travels back and fourth between hippocampus and MEC and that the entire parahippocampal structure must rather be seen as one single system given the reciprocal connectivity between hippocampus and MEC and that animals are likely to simultaneously update some type of map and use path integration.

Adding to this discussion, in Paper II, we observe detriments in place cells and grid cells after a bilateral local perturbation in MEC, indicating that there is at least some dependency in place cells to changes in activity in MEC.

1.4.5 Path integration in CANs

Since position is the integral of velocity, a neuron that can maintain spatially stable spiking can update its encoded position by perfectly integrating its velocity input. Thus, when a stable hexagonal pattern is formed in a CAN, if shifted according to velocity, it may be used to integrate self motion cues. In continuous attractors, states of the pattern at different positions in the neural sheet are identical, thus, by feeding the network with directional input modulated by speed, it is possible to map the neural pattern to a spatial pattern - effectively integrating the input (path integration).

Many models of the grid cell network now have some form of attractor dynamics to match the experimental findings of grid cell modularity. However, how the pattern is shifted in neural space relative to self-movement in order to perform path integration is an active area of debate[58,77].

There are several ways to shift the pattern, here I will introduce the two most common ones. The first uses conjunctive direction-velocity[49,52,55] neurons with firing rates given by

$$s_x(u) = v_x s(u). \tag{1.7}$$

Here, v_x is the speed in direction $x \in \{east, west, north, south\}$ and $s(u)$ is given by Eq. (1.6). Then these conjunctive rates are added in Eq. (1.6) as multiple layers of direction-selective speed-modulated neurons that feed into the network which maintain the hexagonal pattern.

$$\frac{\mathrm{d}s(u)}{\mathrm{d}t} = -\frac{s(u)}{\tau_m} + G\left(\int_\Omega W(u,u')s(u)du'\right) + \beta\sum_x s_x(u+\Delta u_x). \tag{1.8}$$

[74]Stachenfeld et al., 2017 [65]Banino et al., 2018 [75]Cueva and Wei, 2018 [76]Sorscher et al., 2019 [58]Zilli, 2011 [77]Giocomo et al., 2011 [49]Ocko et al., 2018 [52]Samsonovich and McNaughton, 1997 [55]Burak and Fiete, 2009

Here, Δu_x bias the connectivity from direction selective neurons in the preferred direction x and as long as the strength parameter β is small enough such that the attractive state is not diminished, its firing rate modulates which direction the pattern moves.

1.4.6 Theta oscillations and MEC

Rhythmic neural activity in the form of brain oscillations are proposed to be essential to synchronize activity for efficient population coding. Oscillations in the theta frequency range (4-12 Hz) are particularly prominent in the MEC of rats exploring their environments. Theta activity is proposed to govern sequences of neural activity in the hippocampus and entorhinal cortex, possibly facilitating plasticity and strengthening connections within neural ensembles[78].

Theta oscillations in MEC are associated with the medial septum and the diagonal band of Broca, together forming the medial septal area (MSA)[79,80]. The MSA contains multiple cell types that together drive the pace-making activity[81,82], which drive theta oscillations in MEC by means of glutamatergic, cholinergic and GABAergic projecting neurons[83]. The GABAergic projections predominantly terminate on inhibitory interneurons, believed to be the main driver of rhythmic activity in MEC through its coordination of the local inhibitory circuitry[84,85] (Fig. 1.5). The glutamatergic projecting neurons seems to encode speed in their firing rate picked up downstream in the hippocampus[86] and MEC[87]. This speed code is indicated to carry speed information to grid cells[53,55]. The cholinergic projection to MEC seems less relevant for grid cells, but rather modulates behaviours associated with novelty or anxiety[88]. Theta nested gamma oscillations has been proposed to underlie compression of sequences which is multiplexed and sent to hippocampus[89]. Pharmacological inactivation of MSA disrupts theta oscillations in MEC and impairs grid cell activity[90,91]. And lastly, passive transport of animals remove the relation between speed and theta frequency and power together with vanished spatial representations in grid cells[92].

Early on after the grid cells were first observed, the oscillatory activity of these cells have been an active research topic[90–94]. In addition to the possibility of positional decoding from multiple grid fields, information about location within single fields is also observed. First discovered in place cells[95], spiking within a place field is closely related to where the spike falls relative to the phase of theta - spikes fall progressively earlier in the theta phase with distance traveled within a field, a phenomenon known as phase precession. Later observed

[78] Buzsáki and Moser, 2013 [79] Petsche et al., 1962 [80] Mitchell et al., 1982 [81] Müller and Remy, 2017 [82] Leão et al., 2014 [83] Manns et al., 2001 [84] Gonzalez-Sulser et al., 2014 [85] Unal et al., 2015 [86] Fuhrmann et al., 2015 [87] Justus et al., 2016 [53] Fuhs, 2006 [55] Burak and Fiete, 2009 [88] Carpenter et al., 2017 [89] Colgin et al., 2009 [90] Koenig et al., 2011 [91] Brandon et al., 2011 [92] Winter et al., 2015 [90] Koenig et al., 2011 [91] Brandon et al., 2011 [92] Winter et al., 2015 [93] Burgess et al., 2007 [94] Schmidt-Hieber and Häusser, 2013 [95] O'Keefe and Recce, 1993

in grid cells[96], this phase code can be used to decode a positional reference of location within fields[97].

Taken together, the studies mentioned above indicate that theta oscillations are important or even necessary in order to perform path integration or otherwise support a stable representation of space in grid cells. However, even if it is possible to decode position from phase precession, this is not necessarily relevant in a specific animal. Moreover, large-scale unspecific perturbations, such as with pharmacological silencing[90,91], can be miss-informative as they affect off target neurons or brain areas (discussed further in Section 1.5).

In Paper I we investigate whether the temporal activity of grid cells, i.e. theta oscillations and phase precession can be dissociated with spatial specific activity.

1.4.7 Path integration by oscillatory interference

Another model of grid cells focus on the fact that grid cells phase precess and generates grid fields by means of oscillatory interference (OI). By combining velocity controlled oscillators (VCO) with an oscillatory background, e.g. driven by MSA input to MEC, the difference in frequency in single neurons would give rise to a constructive interference pattern that both gives phase precession and hexagonal grid activity, known as OI models. Here, the directional velocity (VCO) $v_\phi(t)$ with direction-selectivity ϕ is defined as

$$v_\phi(t) = s(t)\cos(d(t) - \phi) \tag{1.9}$$

where $s(t)$ is the animal's speed and $d(t)$ the direction at time t. Consider two oscillators, one with a constant frequency ω_2, e.g. from MSA, the other with frequency $\omega_1 = \omega_2 + \beta v_\phi$. Since an oscillator's frequency ω is the time derivative of its phase $\frac{\mathrm{d}\Phi}{\mathrm{d}t} = \omega$ the derivative of their phase difference is the difference in frequencies $\frac{\mathrm{d}(\Phi_1 - \Phi_2)}{\mathrm{d}t} = \omega_1 - \omega_2$. Thus, the phase difference encodes position as $\Phi_1 - \Phi_2 = \beta \int v_\phi$ and is proportional to distance traveled in direction ϕ. In order to cover multiple directions, Eq. (1.9) can be used to make multiple VCOs, which again can be used to model the membrane potential in grid cells[93].

However, as the OI model is a single neuron model it does not account for the the population code within a single grid module. Recent work performed patching of cells in MEC *in vivo* in order to look more closely at how the membrane potential in grid cells behave during field crossings. They found that the potential ramped up during field entry, as expected according to CAN models, however, they also exhibited clear theta modulated spikes which also phase precessed[94,98]. To account for both phase precession, ramping and the population dynamics of grid cells a new model was proposed which combined OI with CAN in so called hybrid OI-CAN models[94,99]. Briefly, a single bump CAN

[96] Hafting et al., 2008 [97] Reifenstein et al., 2012 [90] Koenig et al., 2011 [91] Brandon et al., 2011 [93] Burgess et al., 2007 [94] Schmidt-Hieber and Häusser, 2013 [98] Domnisoru et al., 2013 [94] Schmidt-Hieber and Häusser, 2013 [99] Bush and Burgess, 2014

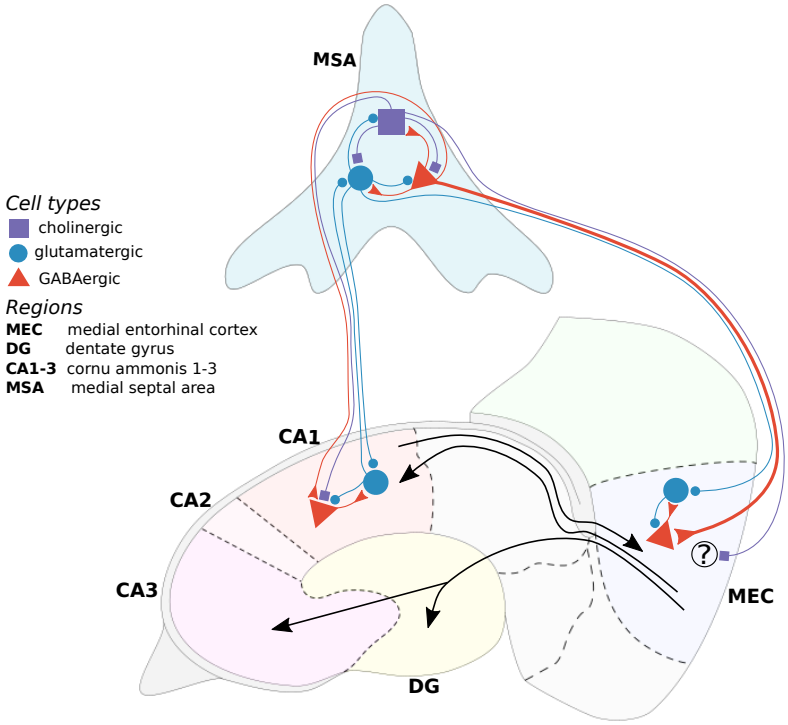

Figure 1.5: **Outline of connections between medial septum area, hippocampus and entorhinal cortex** Medial septal area contains cholinergic, glutamatergic and GABAergic cells which are interconnected and projects to hippocampus and entorhinal cortex. The strongest MSA to MEC projection is the GABAergic, mainly PV^+ interneurons which we stimulate in Paper I, which terminate on GABAergic cells in MEC and CA1. The glutamatergic neurons of MSA project to glutamatergic neurons in CA1 and MEC. While the cholinergic neurons of MSA projects to GABAergic neurons in CA1, it is not fully determined their targets in MEC.

is used to maintain a static pattern, then VCOs are used to focus the position of the bump in relation to movement.

With the OI-CAN models a causal relationship is proposed between phase precession and path integration. That is, phase precession and the appearance of grid fields is proposed to be associated to such an extent that if neurons cease to phase precess, they also loose their distinct spatial firing patterns. However, it is through the difference in base frequency and VCO frequency that phase precession occurs and in order to assess the spatio-temporal relation proposed in OI-CAN models it is not enough to only change the base frequency. In Paper I we thus investigate whether the hypothesis from hybrid OI-CAN models is true by pacing PV^+ neurons in MSA, both driving the brain-wide theta oscillations and grid cell oscillations, by locking grid cells temporally to the pacing, phase precession was abolished, allowing us to test directly the hypothesis from the OI-CAN model.

1.4.8 Stability in the grid cell network

The patterns of grid fields are remarkably stable. When exposed to novel environments, the pairwise activity of grid cells closely resembles that would be expected of continuous attractors[20]. Similarly, pairwise correlations are also stable in different stages of sleep[100,101]. Moreover, the location of the fields remains unaltered during local perturbations of different cell types in MEC besides changes to grid field firing rates or increased number of out-of-field spikes[102–104]. These findings together point towards the grid cells forming an attractive network state as postulated by CAN models. However, whether or not this observed activity of grid cells stem from connectivity is still not known.

1.4.8.1 Perineuronal nets

If the grid cell network of MEC is determined by its connectivity, is it there from birth or is it developed? During development, the brain goes through periods of varying levels of plasticity. In early development, many sensory areas goes through a critical period with high levels of plasticity where adequate sensory input determines functional aspects of the circuits[105,106]. Place cells and grid cells form at different time periods during development, where head direction and place fields form as early as 16 days after birth (P16)[68,69]. Grid fields are less distinct at this stage and does not show clear hexagonal periodic patterns before P20-P22[107].

Alongside development of grid cells, a specialized form of extracellular matrix aggregates known as perineuronal nets (PNNs) surrounds the soma and proximal dendrites of different sets of neurons, mainly the parvalbumin positive (PV^+) interneurons in MEC[108]. The PV^+ neurons are believed to give recurrent

[20]Yoon et al., 2013 [100]Gardner et al., 2019 [101]Trettel et al., 2019 [102]Miao et al., 2017 [103]Kanter et al., 2017 [104]Zutshi et al., 2018 [105]Hensch, 2005 [106]Espinosa and Stryker, 2012 [68]Langston et al., 2010 [69]Wills et al., 2010 [107]Moser et al., 2015 [108]Lensjø et al., 2017

inhibition in the grid cell network[109,110]. In the neocortex, the PV^+ cells mature in parallel with the assembly of the PNNs late in postnatal development as the critical period comes to an end[111]. The PV^+ cells have been proposed as key regulators of plasticity, both during development and in adulthood[105,112].

PNNs are thought to restrict plasticity in the adult brain, both by acting as a structural barrier preventing synapse formation and by effecting PV cell physiology[113].

This role of PNNs with regards to plasticity is supported by the fact that degradation of PNNs in adult animals by the enzyme Chondroitinase ABC (chABC) increases plasticity levels in several brain areas, including the visual and auditory cortices[114–116], perirhinal cortex[117] and amygdala[118]. Removing PNNs disrupts the recall of remote memories, and at a cellular level reduces excitability and spiking activity of putative PV^+ neurons[115,119].

Given that PNNs emerge during development at the same time as grid cells, and PNNs are thought to modulate the interactions of principal and GABA-ergic interneurons, there are good reasons to expect that PNNs may play a role in the grid cell network. There is a tight, developmentally regulated association of PNNs with PV^+ interneurons in MEC[108], and there is relatively low degree of remapping in mature grid cells compared to hippocampal spatial neurons. How this stability of spatial representations in MEC arises developmentally has remained unanswered. Moreover, whether there is a PNN-regulated confinement of plasticity is not known.

In Paper II we aim to study the effects of PNNs on the function of grid cells in the entorhinal cortex.

1.5 Causality in neuroscience

In neuroscience, we are currently undergoing a revolution in amount of methods to measure and perturb neural systems.

The brain is a distributed system of neural networks and when measuring brain activity in combination with behavior it is difficult to know that the activity in the measured brain area actually causes the observed behavior. For instance, the area under consideration might be downstream of the actual area that causes behavior and might thus be spuriously correlated. Moreover, one system might be a partly interacting element in a larger, system-wide computational effort to cause a certain behavior. The interactions can be complex and non-linear and, most importantly, we do not know *a priori* how a brain area supports any given behavior and may thus not account for interactions. Therefore, by simply investigating correlations between behavior and brain activity might give a completely false mechanistic, or causal, idea of how the brain works.

[109]Couey et al., 2013 [110]Buetfering et al., 2014 [111]Hockfield et al., 1990 [105]Hensch, 2005 [112]Donato et al., 2013 [113]Fawcett et al., 2019 [114]Pizzorusso, 2002 [115]Lensjø et al., 2016 [116]Happel et al., 2014 [117]Romberg et al., 2013 [118]Gogolla et al., 2009 [115]Lensjø et al., 2016 [119]Thompson et al., 2017 [108]Lensjø et al., 2017

Questions of causation started as a philosophical debate[120], and causal inference as a statistical tool is a fairly new development, as formal methods to establish causal relations has been unavailable. In the latter half of the 20th century this changed, thanks to the work of pioneering methodologists such as Donald Rubin[121], Judea Pearl[122] and others[123].

Historically, scientific inquiry has been hampered by a disability to speak rigorously about causality[124], much attributed to Pearson, inventor of the Pearson correlation coefficient. In his opinion, the closest one could get at causal inference was through high correlation. However, when two things A, B, for example have a common cause C, their correlation can be spurious. There may be no causal interaction between A and B, they are correlated only because of C - known as a confounder. In complex systems with multiple interactions, confounders can be abundant, and if not controlled for, they will impede attempts at understanding the underlying mechanics or causal relations. One of the tools that sits at the heart of causal inference is the graphical model known as a directed acyclic graph (DAG). With this graph, a model can be stated rigorously not only with mathematics but through a graphical illustration that omits any assumptions about specific functional relations. For this reason, it is in my view a tool which may become invaluable in neuroscience. It can be used to clearly communicate hypothesis about causation, and as a starting point when designing experiments in combination with mathematical models. The first use of a DAG was actually in biology[125], used to estimate the amount of heritage versus environmental factors in how the fur was patterned in ginea pigs. The novel framework was unfortunately forgotten[iv] only to be put in use mostly in fields outside biology, such as in epidemiology, econometrics and machine learning (e.g. causal Bayesian networks).

A DAG is a formalization of causal connections in a system, that is, the arrows denotes statistical association and a parameter β is typically associated with the strength of causation which we wish to estimate. An example of a DAG is the following graph

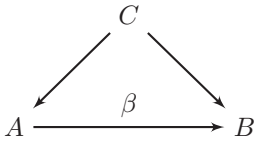

When drawing a DAG, one always starts with an exposure (A) and an outcome (B), and then the other factors in the system, such as the confounder (C). Even if A is associated with B i.e. the conditional probability of B given A is larger than zero, or $\beta = P(B|A) > 0$, we cannot know if this represents a causal relationship

[120]Philosophy Archive, 2001 [121]Rubin, 1974 [122]Pearl, 2000 [123]Spirtes et al., 1993 [124]Pearl and Mackenzie, 2018 [125]Wright, 1920

[iv]Not only forgotten, but cast aside by followers of Pearson's ideological scientific view according to Pearl[124]

because of the confounding factor C which occludes the estimation of β - if truly zero, there is no causal relation between A and B. Let us imagine that A, B, C represents a system in the brain, for example, A represents grid cells, B represents path integration, and C represents some other brain system that can affect both, and we do not necessarily know about it - to keep with the subjects in this thesis. Furthermore, we want to establish that grid cells cause path integration. To this end we need some type of intervention on A, which can be formalized with the do-operator[122] as $P(B|do(A))$ meaning the conditional probability of B given experimental control over A. When confounders are present, typically $P(B|do(A)) \neq P(B|A)$, since experimental control of A removes confounding effects, graphically presented as removing the arrows from C to A, B.

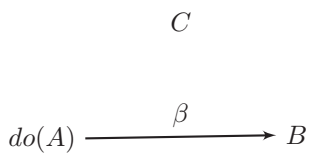

1.5.1 Interventions in neuroscience

One of the most important tools in neuroscience is the ability to intervene in systems activity. Historically, the only tools available were coarse, such as lesions and pharamcological interventions that affected entire brain areas. These interventions have both a coarse spatial and temporal extent. To get a more finely tuned temporal extent, electrical stimulation has been used, however, this has no neuron type specificity. This lack of experimental control was recognized as a problem already in the 70's by Francis Crick[126], and now we can give an example with a DAG on why this can be problematic

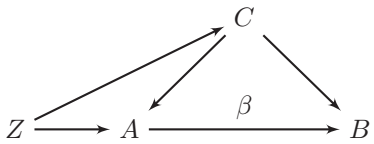

Here, an intervention is denoted by Z, but with lack of experimental control one cannot rule out the possibility that the intervention also interacts with a confounder and thus further occludes the attempt at estimating the causal effect between A and B.

When there are no interventions available, or only imperfect interventions, the data is called observational. If it is not possible to isolate the causal effect under investigation, it is called unidentifiable. However, it can sometimes be possible to circumvent introducing an intervention, or otherwise improve issues of a certain intervention. By drawing an ever more detailed DAG of the system, e.g. identifying some sub-systems affecting others, it is sometimes possible to estimate $P(B|do(A))$ using do-calculus[122]. Or, with quasi-experimental methods, it is

[122]Pearl, 2000 [126]Crick, 1979 [122]Pearl, 2000

possible to emulate randomized controlled trials (RCT) by utilizing randomness innate in the system under consideration.

In order to get causally valid understanding of the brain we should improve both experimental techniques and estimation techniques, preferably such that the two goes hand in hand when starting analysis of experimental data. One experimental refinement would be the ability to affect subsets of neuronal types[v]. Light would be the optimal activation medium as it does not have any natural modes of interaction in the normal brain, but how to achieve this has been elusive up until the recent discovery of light sensitive proteins.

1.5.2 Stimulating neurons with optogenetics

Microbial opsins of an ancient gene family adapted from organisms such as algae and archaebacteria was found to be light sensitive. With each gene encoding a distinct protein that directly elicits electrical current across cellular membranes in response to light, make up the fundamental ingredients in the tool-set known as optogenetics[129]. Here light-driven ion pumps and channels known as opsins provide a rapid, specific transport of ions across the membrane[130,131], enabling neural activity to be driven or silenced by lightFig. 1.6. These are typically introduced by viral vectors which carries the genes for opsin-clones to the cell which are further mass-produced by the cells' own machinery. In order to activate opsins, the regular method is to implant an optic fiber attached to a laser or micro LEDs into the brain-area of interest producing light at the opsins preferred wavelength.

In this thesis we use optogenetics in an experimental setting and a modeling setting in Paper I and Paper III respectively. Furthermore, we use methods from causal inference specifically in Paper III.

1.5.3 Delivery of opsins to neuron specific targets

In order to induce expression of opsins in nerve cells, several methods can be utilized. They can be genetically "knocked in" from birth in loci that confer cell-type specific expression, such as under the parvalbumin promoter to facilitate parvalbumin specific expression. However, having these exogenous proteins expressed in the membrane from birth might lead to cytotoxicity and thereby dysfunction or even cell death[129]. To lessen the risk for cytotoxicity and achieve spatial specificity of expression it is common to deliver the genes that code for a specific opsin by means of viral vectors. Viruses typically used in neuroscience are modified to be replication deficient and unable to multiply and thus not infectious. By gentically modifying viruses it is possible to delete their genes for

[129]Deisseroth, 2015 [130]Zemelman et al., 2002 [131]Boyden et al., 2005 [129]Deisseroth, 2015

[v]Neuronal genotype does not necessarily identify the functional aspects of neurons, for example, both cells expressing calbindin[127] and reelin[128] are found to be grid cells, if these distinct cell types represent two distinct functional types remains to be determined.

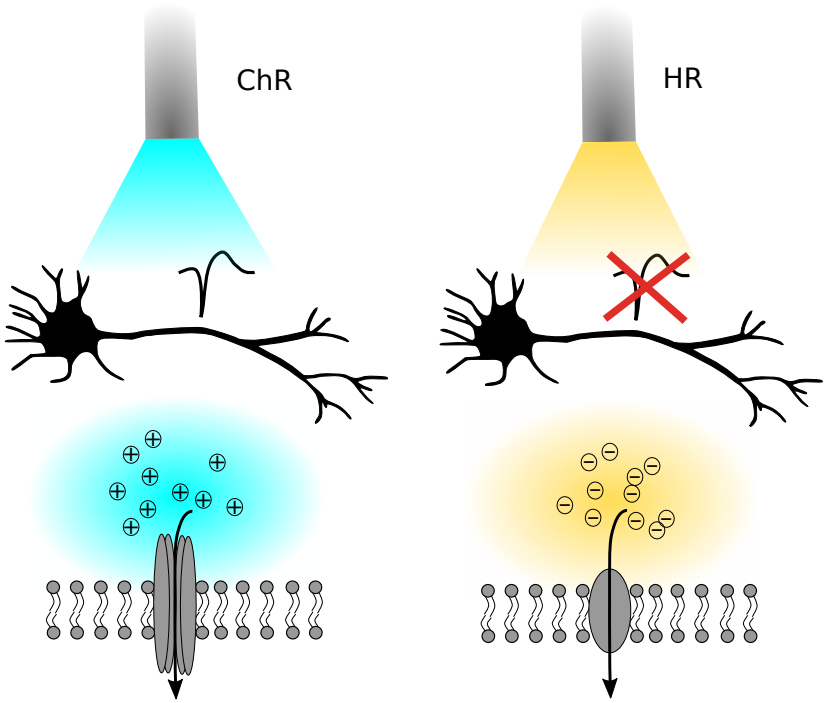

Figure 1.6: **Optogenetics: activation of light sensitive ion channels and pumps** The two main proteins used in optogenetics is the ion-channel channelrhodopsin (ChR) and the ion-pump halorhodopsin (HR). ChR is activated by blue light and upon activation typically opens for nonspecific cations depolarizing the cell. HR on the other hand is activated by yellow/green light and pumps chloride ions through the membrane and thus hyperpolarize the cell.

self-replication and insert genes that inscribe how to produce certain proteins - such as opsins. The most common is the AAV virus, due to its ease of production and relative harmlessness. In order to express genes in neuron specific targets, one can either use promotors that have specific selectivity engraved in the DNA of the viral vector or one may use a Cre-lox system. With a Cre-lox system, a kind of "gene key hole" is knocked in from birth, and the viral vectors only has to encode the "key" in order to be neuron specific. The Cre-lox system is more versatile, as for example it can require two types of "keys" to be present at the same time, but if only single type selectivity is aimed for, a promotor might often be the best alternative - if the desired promotor is available that is.

As there were no PV^+ promotors available for rats at the outset of Paper I, we used a Cre-lox system in a novel rat line[132] that has Cre expressed in PV^+ neurons, allowing us to selectively express opsins in PV^+ neurons.

[132] Yu et al., 2018

1.6 Estimating connectivity

Activity of a network can be heavily influenced or directed by connectivity. In MEC, the leading hypothesis on how grid cells are formed through pattern formation is that of connectivity - forming continuous attractors. One of the difficulties in obtaining connectivity experimentally from grid cells, is that they are functionally identified. For example, when investigating slice preparations from MEC it is possible to identify connectivity by tracing the axonal targets. Once the slice is prepared, it is difficult to know if the neuron under study actually was a grid cell - which is only possible to observe *in vivo*. The functional strength of connections need measurements of electrical activity such as with patch clamp techniques[109]. There have been several attempts at inferring connectivity from statistical analysis of spike trains of simultaneously recorded neurons, in which the membrane potential is considered a latent variable. By using cross correlations between neurons where the effect of overlapping firing fields are removed, grid cells with nearby spatial phases show positive correlations, with no or negative correlation further away[133,134], reminiscent of a Mexican hat type connectivity. These noise correlations, however, cannot separate out other known sources of correlations, such as theta oscillations and head direction tuning. To circumvent this issue Dunn et al.[133] found significant residual functional coupling using a kinetic Ising model, that was positive for grid cells with nearby phases and negative for larger phase differences. However, systematic errors in connectivity inference is evident in networks with strong recurrent connectivity - when the true network dynamics and the generative model assumed for inference mismatch[135].

Recently, connectivity inference was attempted by means of cross correlations in combination with optogenetic stimuli[136]. However, large spread of activation upon light stimulation can confound connectivity inference and it is thus illusive if correlated activity is due to co-stimulated neurons, not necessarily measured, or the activity itself. Understanding how neural activity reflect connectivity is ultimately a causal question and it is not sufficient to look at correlates[137,138]. Confounders are crucial to take into account when estimating the amount of influence between neurons. Consider a simple system of three neurons A,B,C where A and B have a common cause U and A is physically connected to C but B is not. Due to U, A and B are correlated ($\mathrm{cov}\,(A,B) \neq 0$), which will again induce a spurious correlation between B and C ($\mathrm{cov}\,(B,C) \neq 0$). With each neuron in the brain having an order of 10^4 number of connections, confounders can be abundant. Synaptic strength is typically distributed log-normally[139], where most connections are weak, but with a heavy tailed distribution, some are considerably large. It usually takes many spikes to induce a postsynaptic action potential, and indeed cross correlations between spike trains are typically small[140]. With a heavy tail, some connections can however, be very influential.

[109]Couey et al., 2013 [133]Dunn et al., 2015 [134]Tocker et al., 2015 [133]Dunn et al., 2015 [135]Das and Fiete, 2019 [136]English et al., 2017 [137]Marinescu et al., 2018 [138]Mehler and Kording, 2018 [139]Buzsáki and Mizuseki, 2014 [140]Cohen and Kohn, 2011

With many long ranging connections across brain regions one may never be certain that all confounders are taken into account even if a large population of local neurons are recorded. When a population supports network states that are attractive, such as in a CAN network, the network state itself can act as a confounder[135].

When data originates from randomized perturbations, correlations reflect causal interactions. Randomization is typically induced with randomized controlled trials (RCTs) which can be understood as a randomized perturbation. (Z) which removes confounding, illustrated by the lack of arrows going from C to A, B

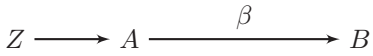

In an RCT, we can see the randomization as a random treatment assignment indicator, in the binary case 1: treatment, 0: placebo. To emulate RCT the assignement indicator can be approximated, known as in instrumental variable (IV), and have a similar effect on a DAG as an RCT - when employed correctly, incoming arrows from a confounder is abolished. This technique was utilized in Paper III.

1.6.1 Perturbing neural systems

Randomizing activity of a putative presynaptic neuron is key when estimating connectivity between neurons. Perturbations are thus used to separate imposed activity from the confounding background. If perturbations only affect one neuron by e.g. patch clamp technique, it is possible to infer causal connectivity. Using a controlled graded stimulus S with strength s driving the potential v_X, of neuron X such that $v_X(t) = f(s,t)$ for some function f, and time t allows controlled spike times $t_k^X, k = 1, 2, 3, ..., N_X$ with N_X spikes. The causal effect β of neuron X on the postsynaptic potential v of neuron Y (a proxy for the synaptic strength as measured by e.g. the peak EPSP) can be estimated by the treatment effect

$$\beta = v(t_k^X + \Delta t) - v'(t_k^X + \Delta t). \tag{1.10}$$

Here a treatment is referred to as a presynaptic spike t_k^X which is induced by stimulations, and the response in v_Y is measured after Δt. The term v' is the counterfactual value of v i.e. the value v would have had, at the exact time of measurement, if there were no stimulation $s = 0$ i.e. no presynaptic spike. The counterfactual, although, is not available, but may be estimated by a baseline

[135] Das and Fiete, 2019

value v_0 just before the treatment, estimating β with the average treatment effect

$$\hat{\beta} = \frac{1}{N_X} \sum_{i}^{N_X} \left(v(t_k^X + \Delta t) - v_0 \right). \tag{1.11}$$

In cases where the membrane potential is not available, but only the spikes t_k^X, t_k^Y it is natural to define a causal effect as the interventional probability $\beta = P(y|do(x))$[122]. Which reads: the probability of a spike y, when a spike x is forced by the experimenter as indicated by the $do(\cdot)$ operator. In Paper III we attempt to estimate the interventional probability when using optogenetics and derive a method that, based on instrumental variables can give causal estimates of transmission probability between neurons.

[122] Pearl, 2000

Chapter 2

Objectives of Papers

In this thesis the following aims are sought:

Paper I aims to investigate if the temporal and spatial activity in grid cells can be dissociated, and whether it is possible to perturb the temporal oscillation in grid cells while the spatial representation remains stable.

Paper II aims to investigate how removal of perineuronal nets affects the grid cell network, consistent with the idea that the appearance of PNNs stabilizes the network during grid cell development.

Paper III aims to establish a method that can utilize optogenetics and causal inference to abolish the confounding factors present when estimating transmission probability from pairwise spiking activity of recorded neurons.

Chapter 3

Summary of Papers

Paper I

Oscillations in the theta frequency band (4-12Hz), are particularly prominent during animal exploration and are proposed to be essential for spatial navigation and memory processing[78]. When silencing the medial septal area (MSA), theta oscillations are disrupted alongside degradation in hexagonal firing patterns of grid cells in the medial entorhinal cortex (MEC)[90,91]. Both speed-rate correlations and grid cell firing ceased when disrupting the speed-frequency and speed-amplitude correlations in theta oscillations when animals were passively exploring an environment[92]. Combined, these studies propose that theta oscillations have a causal role for the emergence of grid cells and thereby essential for animal navigation.

A phenomenon known as phase precession where grid cells fire progressively earlier in each theta cycle traversing through the grid field, has been taken as evidence for a crucial link between theta activity and grid cell firing patterns[93,96]. Recently, direct measurement of grid cell membrane potentials has been shown to support a model that combine phase precession from oscillatory interference (OI) together with a continuous attractor network (CAN)[94,98].

There is also evidence that grid cells are reliant on input from the hippocampus[70]. There are strong projections from MSA to the hippocampus, thus, silencing MSA[90,91] may indirectly affect other potential sources of input to grid cells. Further, studies with measurements of membrane potentials[94,98] lack interventions, thus, direct experimental evidence for the hypothesis that theta oscillations are crucial for grid cell firing is still lacking.

In this work we set out to investigate the relation between temporal oscillations in grid cells and their spatial representation. Are these two aspects intertwined or can they be dissociated?

MSA may have a critical role in controlling grid cell firing as grid cells in superficial layers of MEC are interconnected through local interneuron circuits[109,141], targeted by inhibitory projections from MSA, particularly the projections from parvalbumine positive interneurons (PV^+). Selective activation of these projections through inhibitory neurons of MSA opens the possibility to directly assess the relation to grid cells spatial representations. Using a PVCre rat line allowed cell-type specific optogenetic activation of PV^+ cells in MSA. We robustly drove oscillations in the local field potentials (LFP) in MEC at different

[78]Buzsáki and Moser, 2013 [90]Koenig et al., 2011 [91]Brandon et al., 2011 [92]Winter et al., 2015 [93]Burgess et al., 2007 [96]Hafting et al., 2008 [94]Schmidt-Hieber and Häusser, 2013 [98]Domnisoru et al., 2013 [70]Bonnevie et al., 2013 [90]Koenig et al., 2011 [91]Brandon et al., 2011 [94]Schmidt-Hieber and Häusser, 2013 [98]Domnisoru et al., 2013 [109]Couey et al., 2013 [141]Fuchs et al., 2016

frequencies, completely abolishing the endogenous theta activity. We found that optogenetic stimulation of MSA PV^+ cell projections reliably activated grid cells through disinhibition from local interneurons. Grid cells fixed their firing to the stimuli and were thus not able to phase precess. This gave an opportunity to test directly if the animal may sustain spatial stable patterns without performing phase precession - testing directly the hypothesis given by OI type models.

During our intervention, grid fields remained remarkably spatially stable, with some increased out-of-field firing during the time window of elevated response. Since phase precession was abolished we concluded that phase precession is not important for the spatial firing of grid cells.

While speed correlations in the local field potential was absent during stimulation, the grid cells remained speed modulated. Therefore, speed information in grid cells are unlikely to be obtained from theta oscillations. The correlation between speed information in theta oscillations and grid cell firing found by Winter et al.[92] is indicated to be spurious. Taken together, we show that theta oscillations and grid cell spatial activity patterns can be dissociated, suggesting that the spatial coding of grid cells are independent of the temporal code, contradictory to established theory of the emergence of grid cells and path integration.

Paper II

The ability to encode novel environments without compromising old memories are vital for survival and cognitive functions. Spatially tuned neurons in the hippocampus and entorhinal cortex are key units for navigation and spatial memory. Neurons in the medial entorhinal cortex (MEC) represent information about self-location[30] where grid cells have multiple, spatially specific firing fields forming a characteristic hexagonal pattern spanning the entire surface of the area visited by the animal[31].

The network that controls grid cell spiking appears to rely on recurrent inhibition for the specific activity of grid cells. Stellate cells, one of the two principal cell types that display grid cell firing[33], are connected via parvalbumin expressing (PV^+) inhibitory interneurons[109,141]. Whether PV^+ inhibitory neurons in MEC play a role for development of the grid cell network remains elusive.

In rodents, grid cell firing patterns emerge around postnatal day 16-18 (P16-P18) and transition over time from being unstable and non-periodic to highly regular, reaching adult level grid scores around P28-P34[68,142]. A hallmark of maturing PV^+ cells is that aggregates of specialized extracellular matrix, called perineuronal nets (PNNs), condense on the cell soma and proximal dendrites, leaving openings only for synaptic connections[143]. PNNs are believed to help stabilize the activity of PV^+ cells by supporting synaptic integrity and limiting

[92]Winter et al., 2015 [30]Fyhn et al., 2004 [31]Hafting et al., 2005 [33]Rowland et al., 2018 [109]Couey et al., 2013 [141]Fuchs et al., 2016 [68]Langston et al., 2010 [142]Bjerknes et al., 2014 [143]Dityatev and Schachner, 2003

synaptogenesis, in addition to supporting PV^+ cell physiology[144]. By the time PNNs in sensory cortex are fully mature, plasticity in the local network is strongly reduced. However, juvenile levels of plasticity can be reinstated by experimentally removing PNNs in adult animals, which both increases structural plasticity and reduces inhibitory spiking[114–116]. Interestingly, the timeline for maturation of PNNs in MEC coincides with the timeline for the development of grid cell firing[108,142]. This co-occurrence suggests that grid cell activity could be shaped during the developmental period with high levels of plasticity, and that the later presence of PNNs ensures stability of established synaptic connections, and thus maintains the integrity of the network and the spatio-temporal relationship between grid cells.

Once established, the periodic spiking pattern of grid cells is remarkably stable when animals revisit the same environment. Local perturbations of different cell types in MEC cause changes to grid field spike rates or an increased number of out-of-field spikes, but the location of the fields remains unaltered[102–104]. Furthermore, when placed in a different environment, neighboring populations of grid cells show a coherent shift of grid fields[42] that retains their relative spatio-temporal relationships[20]. This suggests that the mature grid cell network is relatively hardwired. However, the mechanism by which information processing is altered at the neuronal and network level by PNNs remains rather elusive. We thus set out to examine the functional consequences of PNN removal on neuronal representation of space in adult rat MEC.

To test if PNNs support the stability of the grid cell network, we experimentally disrupted PNNs in MEC of adult rats and recorded from single units while animals explored a familiar arena or a novel environment. We observed reduced inhibitory spiking activity when PNNs were removed, and grid cells displayed reduced spatial specificity and spatial information in the familiar environment. Theta power was altered after PNN removal which is also observed during learning, indicating a shift in network state. Moreover, we showed with simulations that altered synaptic weights can give rise to reduced inhibition in a continuous attractor network. This is in line with the role for inhibitory neurons in shaping grid cell activity. When the MEC network was challenged with new information by letting animals explore a novel environment, both the pairwise spatial correlations of grid cells and their pairwise temporal correlations were significantly reduced in animals lacking PNNs. Results indicate that the novel place code remained unstable and the spatio-temporal relationship between grid cells was impaired. The exposure to a novel environment also destabilized the subsequent representation of the familiar arena, suggesting that PNNs are important for maintaining consistent grid cell representations when a previous environment is revisited.

Grid cells provide input to the hippocampus[33] and are assumed to be the

[144]Wang and Fawcett, 2012 [114]Pizzorusso, 2002 [115]Lensjø et al., 2016 [116]Happel et al., 2014 [108]Lensjø et al., 2017 [142]Bjerknes et al., 2014 [102]Miao et al., 2017 [103]Kanter et al., 2017 [104]Zutshi et al., 2018 [42]Fyhn et al., 2007 [20]Yoon et al., 2013 [33]Rowland et al., 2018

primary determinant of hippocampal place cell firing[60,66,67]. However, this notion has been challenged by the fact that place cells appear before grid cells during development[68,69]. Finally, we recorded place cells from hippocampal area CA1 when PNNs were removed in MEC. Our data shows that the local changes we observed in MEC were also reflected in place cell coding, supporting the idea that the stability and high spatial specificity of grid cell representations are necessary to provide accurate spatial information to place cells.

Because the PV^+ cell type has been shown to play a critical role in spatial representations in grid cells in MEC, and because narrow spike waveforms are associated with fast-spiking interneurons, we conclude that alteration of PNNs directly affects PV^+ cells, which in turn alters spatial representations. Together, our data shows that the presence of PNNs ensures precise spatial and temporal coding needed to maintain the network configuration of grid cell and place cell circuits.

Paper III

A central goal of neuroscience, arguably, is to understand the mechanisms or causal chains that give rise to activity in the brain, to perception, cognition, and action. Typical studies using e.g. electrophysiology or calcium imaging only record from a small subset of all neurons with countless connections. Obtaining estimates of connectivity that reflects the underlying mechanisms in such complex systems is hard because of the numerous ways the contributing elements may interact internally[145]. Even if we could record all neurons at the same time, estimating causality and producing a mechanistic understanding would be extremely challenging[135].

Data obtained from pairwise activity of neurons is observational, which means that it does not result from randomized perturbations. With such data at hand, to understand the effect of one neuron on the other it is not sufficient to know the correlations between them. In such cases, we can never know to which level the observed activity was caused by other observed activity, or by unobserved confounding activity. If mechanisms are estimated from observational data in the presence of confounding, the consequence may be large errors and incorrect conclusions[146]. Unobserved neural activity confounds estimates of causal interactions and makes it difficult to estimate underlying mechanisms.

Confounding is the big threat to causal validity[122] irrespective of the use of simple regression techniques or advanced functional connectivity techniques[147–150]. To estimate connectivity, it is first and foremost important that the used signals reflect cause and effect, therefore we use the term causal transmission probability. Naïve regressions in partially observed systems will generally not reveal causality.

[60]Solstad et al., 2006 [66]Gil et al., 2017 [67]Hales et al., 2014 [68]Langston et al., 2010 [69]Wills et al., 2010 [145]Jonas and Kording, 2017 [135]Das and Fiete, 2019 [146]Angrist and Pischke, 2008 [122]Pearl, 2000 [147]Stevenson et al., 2008 [148]Honey et al., 2009 [149]Aitchison and Lengyel, 2017 [150]Pfau et al., 2013

To estimate causal relationships between neurons, stimulating the presynaptic neuron is the gold standard. In fact, a common definition of causality is in terms of the effect of changing one variable in the system, independently of all the other variables – an intervention[122]. If we stimulate single neurons, the ability to estimate causal relationships by regression is within reach. However, this is experimentally challenging and yields low cell count because it requires intracellular, juxtacellular or two-photon stimulation[151–154]. Because gold-standard perturbations are challenging, it is necessary and highly desirable if causality could be obtained from optogenetic stimulation in combination with neural recordings of large populations of neurons[130,131].

Interpreting the results from optogenetic stimulation in terms of causal interactions is difficult. In most experimental settings, optogenetic stimulation will affect many neurons simultaneously. Hence, the stimulus will produce a distributed pattern of activity. This distributed pattern of stimulation produces activity which then percolates through the network of neurons. Thus any postsynaptic activity induced by stimulation could in principle come from any of the stimulated neurons, introducing problematic confounding.

For insights into how we may resolve the confounding problem induced by optogenetic stimulation, we look to other fields that have addressed this problem of endogenous regressors. The inference of causality from observational data is addressed in the fields of statistics[122], machine learning[155] and econometrics[146]. These fields have extensively worked on methods to estimate causality in the face of potential confounding and may offer us clues on how to solve our problems.

In the final work we describe a method for inferring causal interactions between neurons from spike train data. We demonstrate problems with naive estimators of causal effects under influence of confounders, and propose a method based on instrumental variables (IVs) and neural refractoriness.

We present a novel application of IV causal analysis to neural data. Specifically, we use IV analysis to estimate connection strengths between neurons which are being recorded via extracellular electrodes while a subset of those neurons are stimulated with relatively local, one-photon optogenetic excitation. We demonstrate the method on a toy simulation of three LIF neurons. The proposed method shows clear advantages over previously used naive methods based on correlations.

This manuscript is in preparation and the method remains to be tested in a larger simulated network.

[122]Pearl, 2000 [151]Pinault, 1996 [152]Lerman et al., 2017 [153]Nikolenko et al., 2007 [154]Emiliani et al., 2015 [130]Zemelman et al., 2002 [131]Boyden et al., 2005 [122]Pearl, 2000 [155]Peters et al., 2017 [146]Angrist and Pischke, 2008

Chapter 4

Discussion

In this thesis I have presented a brief introduction to experiments and theoretical models in neuroscience. I introduced how the electrical activity of single neurons was found to behave in a simple manner relative to the complexity of its structure. By utilizing populations of ion channels that gives rise to currents, the neuron produces action potentials. There has been much focus on the computational aspects of single neurons, much due to technological limitations of measuring only a handful of neurons[7]. Neurons have been found as specific as only responding to images of Halle Berry[156], termed grandmother cells, from the idea that one neuron is responsible for holding information regarding your grandmother. Population theory, interpret this finding as neuron activity manifested by population activity and is part of an emergent signal of the population representing the person. However, the same neuron can also take part in another population coding for another person or event.

Many networks in the brain seem to form individual computational units, such as the grid cell network that form distinct modules of populations that share orientation, scale etc. This indicates that the brain is a network of networks and by identifying functionality of these networks we might be able to simplify models such that they do not have take into account the details of single neurons - much like the single neuron models does not take into account the complexity of ion channels. One example of such a model is the CAN model of grid cells (Eq. (1.6) in Section 1.4.3).

Paper I

In Paper I we investigated the hypothesis of a causal relationship between theta oscillations and the spatial representation of grid cells, as postulated by experimental findings[90–92], grid cell models by oscillatory interference (OI)[93] and hybrid OI-CAN models[94,99]. While experimentally induced alterations of theta oscillations in MEC impaired phase precession and speed-modulation of grid cells, their spatial representations were intact. Our findings thus show that the spatial and temporal firing pattern of grid cells can be dissociated.

By pacing MSA PV^+ cells we were able to shift oscillations in MEC to higher frequencies and abolish the endogenous theta. While grid cells were strongly modulated by the new oscillatory frequency, this did not disrupt their overall spatial firing patterns, although they showed slightly reduced peak activity and spatial information rate. Optogenetic stimulation of PV^+ cells activated grid cells

[7]Yuste, 2015 [156]Quiroga et al., 2005 [90]Koenig et al., 2011 [91]Brandon et al., 2011 [92]Winter et al., 2015 [93]Burgess et al., 2007 [94]Schmidt-Hieber and Häusser, 2013 [99]Bush and Burgess, 2014

through disinhibition, where grid cells showed a minor increase in out-of-field spike probability, indicating that the local inhibitory network determines the position of firing fields. We add support to the notion that inhibitory interneurons in superficial layers of MEC recieve monosynaptic input from PV^+ cells in MSA[84,141,157]. Moreover, the rapid return of activity to the initial state after stimulation indicate that the grid cell system is supported by attractor dynamics where the inhibitory circuit of MEC are tightly connected to grid cells[19,109]. Both grid cells and narrow spiking cells showed substantial phase locking to the stimulation frequency with the effect of impaired grid cell phase precession. However, cessation of phase precession did not affect the spatial position of grid fields, strongly suggesting that grid cells do not rely on phase precession to maintain stable spatial representations, thus falsifying the assumptions made in OI type models[93,94,99].

The stability of grid cells have been shown in multiple studies[20,100–102,104] and here we also observe that the grid cell system is stable under disinhibition. Although optogenetic stimulation initiated a brief increase in grid cell activity causing increased out-of-field spiking, the fields remained remarkably stable when looking at time-averaged rate maps. In many cases we observed a reduction in the peak of grid cell rate maps during stimulation, despite robust excitatory responses 5-10 ms after each stimulation pulse. This was likely caused by an inhibitory response to disinhibition of excitatory cells which again activated the inhibitory network due to widespread recurrent connections. Ultimately this may have caused the grid cells to be more strongly inhibited just after the initial response to laser pulses, compared to normal activity.

The results presented in Paper I together with literature of grid cell stability may provide insight to the computational models of grid cells on multiple aspects, both in terms of stable spatial representations under neuronal response but also under phase locking oscillations. Typically, continuous attractor network models (CANs) are associated with path integration arising from velocity input where direction and speed are provided into the network. Even though these models can sustain oscillatory input[19], they do not normally account for temporal activity in grid cells such as phase precession, with the exception of Navratilova et al.[158]. Resting on attractor dynamics, CAN models would likely be able to sustain a stable spatial pattern during interventions similar to those presented here.

Recent theoretical work[159,160] links the temporal relation between MSA pacemaking theta and the resonance and rebound spiking found in stellate cells, to underlie spatial representation, phase precession and theta cycle skipping of grid cells. Taken together, these models strongly indicate that theta oscillations are closely related to the grid cells spatial representations. This is in contradiction to the experimental findings reported here and it will therefore be interesting

[84]Gonzalez-Sulser et al., 2014 [141]Fuchs et al., 2016 [157]Freund and Antal, 1988 [19]Pastoll et al., 2013 [109]Couey et al., 2013 [93]Burgess et al., 2007 [94]Schmidt-Hieber and Häusser, 2013 [99]Bush and Burgess, 2014 [20]Yoon et al., 2013 [100]Gardner et al., 2019 [101]Trettel et al., 2019 [102]Miao et al., 2017 [104]Zutshi et al., 2018 [19]Pastoll et al., 2013 [158]Navratilova et al., 2011 [159]Hasselmo, 2014 [160]Shay et al., 2016

to see if such models can accommodate the dissociation between temporal and spatial grid cell properties.

Recent normative models[65,73–75] does not rely on biologically detailed mechanistic assumptions, but may thus still be used to test the stability of the grid code. If these models are unable to sustain stable patterns under strong perturbations, additional detail may be added into the models, such as inhibitory connections or architecture that gives rise to attractor dynamics. This type of stability testing can create an avenue to additionally inform modellers to create models that better conform to experimental data.

It is well established that frequency and amplitude of the theta rhythm in the field potential increase with running speed[161–164], suggesting that theta oscillations play a role in representation of speed. Passively transporting an animal in an arena does not reduce overall theta rhythms, but eliminates the linear velocity modulation of theta[92]. Our recordings show that the strong and positive correlation between running speed and theta frequency is disrupted when we pace MSA PV^+ neurons with optogenetic stimulation. Interestingly, the neurons' firing rates are still modulated by speed during optogenetic stimulation. The speed signal necessary to update a path integrator may still be provided. This signal could be provided by glutamatergic projections from MSA[87], but it is likely that the glutamatergic neurons are also heavily affected by the stimulation of PV^+ cells in MSA as these are connected[81]. Importantly, the response effect we see in grid cells are highly likely to be due to disinhibition as response dynamics were similar for local stimulations of projecting synapses from MSA in MEC.

Previous studies show that the animal's position on a linear track can be decoded using spike phase relations[97], and that this phase relation is maintained in 2D environments[165]. During Baseline II we found a change in the distribution of precession/recession, indicating that phase relations can vary with time. As Baseline I and II differ in many of the recorded parameters, the change between baseline sessions is likely due to the stimulation session in between. Spike-relation to phase may thus be modulated by MSA inputs, possibly providing control of how a neuron's activity relates to the phase of theta. This modulation might indicate that spikes from grid cells are more robustly related to space than theta phase. Using phase relations as a code for spatial position might thus be problematic. If, for example, there is some context dependence to how a neuron's activity relates to the phase of theta, the information of position might be difficult to decode from phase alone.

Whether oscillations in MEC may contribute to other behavioral aspects such as spatial memory or decision making remains to be shown. At the extreme of scepticism, phase precession and recession in MEC might just be an epiphenomenon. If, however, theta oscillations underlie aspects of behavior, our results may indicate that theta coordination is more important for

[65]Banino et al., 2018 [73]Dordek et al., 2016 [74]Stachenfeld et al., 2017 [75]Cueva and Wei, 2018 [161]Jeewajee et al., 2008 [162]Maurer et al., 2005 [163]Rivas et al., 1996 [164]Whishaw and Vanderwolf, 1973 [92]Winter et al., 2015 [87]Justus et al., 2016 [81]Müller and Remy, 2017 [97]Reifenstein et al., 2012 [165]Reifenstein et al., 2014

memory than navigation, as memory depends primarily on internally generated neuronal sequences[78]. Colgin et al.[89] propose that theta oscillations are carrying multiplexed information in theta-gamma coupling. This could be tested experimentally by finding behavior that is proposed to depend on theta-gamma coupling and combine with stimulation of MSA which should eradicate this coupling. To further determine the nature of spike-phase relations with respect to position, future studies could investigate more closely the stability of phase relations during exploration of novel environments[20], or spike-phase relations during context-dependent tasks[166].

Paper II

A network of reciprocally connected interneurons and principal excitatory cells is thought to be the basis of the grid cell continuous attractor network. In Paper II we assessed the effect on stability in MEC when removing PNNs. We show that degradation of PNNs changes grid cell network dynamics by altering the temporal relationship between grid cells and impairing representations of novel environments. Results from our simulations of the grid cell network in a continuous attractor model indicate that reducing the level of inhibition is sufficient to produce impairments in grid cell spatial specificity, similar to what we observe in experimental data after removal of PNNs, and that this can arise from reduced excitatory to inhibitory connection strength. These results support the notion that PNNs stabilize the network through its effects on inhibitory, putative PV^+ neurons, and suggest a prominent role for PNNs in maintaining the structural and functional organization that shapes the activity of microcircuits in MEC.

Many of the components in PNNs are found in the general extracellular matrix as well. It is worth mentioning that the enzyme ChABC also digests the sugar chains outside the PNNs. ChABC treatment may have a general effect on network function, and cannot be used as a definite proof that a particular neural function is supported by PNNs. Other methods of disrupting PNNs exist, where genes coding for proteins specific for PNNs are knocked out[167]. However, given the stability and long life time of PNNs it takes several weeks before the PNNs vanish in a knock-out model, in contrast to the immediate effect of enzymatic treatments. It can thus be challenging to control when the PNN components are properly disrupted, in addition to control for possible compensatory mechanisms that can be activated during this long time window. If side effects occur in specific matrix component knock-outs, possibly due to various compensatory effects, this needs to be further studied in order to assess the specificity of the PNN disruption. ChABC is widely used to study PNNs both functionally and anatomically and the use of chABC has led to valuable insight into the functional role of PNNs in the brain[108,119]. Thus, we chose to use chABC in this study.

[78]Buzsáki and Moser, 2013 [89]Colgin et al., 2009 [20]Yoon et al., 2013 [166]Aronov et al., 2017 [167]Rowlands et al., 2018 [108]Lensjø et al., 2017 [119]Thompson et al., 2017

If we assume that grid cell activity is governed by continuous attractor dynamics which is maintained by connectivity, we might expect that the effects from removing PNNs in MEC is graded. In a recent study by Chaudhuri et al.[168] a low dimensional representation of the head direction system was obtained through a novel dimensionality reduction technique on the head direction system. They showed that this system had activity confined within a one dimensional ring, which have properties of an attractor typically found in dynamical systems. The ring was stable during sleep, indicating that the low dimensional representation is a fundamental property of this network. They also performed simulations that indicates that the grid cell activity is confined within a torus. Together with findings that the grid cell network is remarkably stable[20,100,101], this indicates that the grid cell system has attractive properties. In such attractor systems the activity of one neuron is stable due to the restorative force governed by the population. This might indicate that the system does not require a stabilizing element over a relative short timescale (days), but might need an extra element of stabilization over the timescale of weeks, or during significant perturbations of systems activity. With this taken into account, we might not expect to see abrupt changes in grid cell representations after PNN removal as found in Paper II. It might be that in order to see a stronger effect in grid cell spatial representations, e.g. completely abolishing grid cell hexagonal firing, we would need both an opening of plasticity and a second large-scale perturbation that force the network to reshape. This could be accomplished through larger change in sensory input such as solving a behavioral task or by using optogenetics. As shown in Lensjø et al.[115] the change in plasticity after PNN removal is not observed until there is a major change in the input to the system/cortical area under investigation (in that case eye-closeure).

Acute non-specific disruption of PV^+ interneuron function by pharmacogenetic manipulations have been shown to impair grid cell firing patterns[102]. Similarly we see a reduction in firing activity of putative PV^+ neurons and one might think that the effects we see are explained by reduced activity in PV^+ cells alone. In the familiar environment, we found reduced spatial specificity, maximum firing rate (both spatial and temporal), and a reduction in activity inside grid fields. This is different from the results by Miao et al.[102]. When acutely reducing PV^+ cell activity they observed an increase in unspecific firing outside fields, and activity inside fields were unchanged. If stable attractor states would rely only on PV^+ cell activity we should see reduced spatial stability also in the familiar environment prior to novel environment exposure. Indeed, Miao and colleagues find reduced spatial stability in grid cells as long as PV^+ cells are inhibited, but this is a much larger reduction in PV^+ cell activity compared to the one we see in our data. We do not see this instability in familiar environments, and the change in stability we observe happens only during and right after the network encodes new information. Is it possible to separate changes in activity from PV^+ cells and changes in plasticity? Changes in one may lead to

[168]Chaudhuri et al., 2019 [20]Yoon et al., 2013 [100]Gardner et al., 2019 [101]Trettel et al., 2019 [115]Lensjø et al., 2016 [102]Miao et al., 2017 [102]Miao et al., 2017

changes in the other e.g. Vogels et al.[169] shows that inhibitory plasticity regulates population activity. Moreover, changes in activity due to pharmacogenetic manipulations are shown to cause changes in PNN and plasticity[170,171]. We thus find it unlikely that activity levels of PV^+ cells is sufficient to explain our results from novel environment experiments, and that removal of PNNs causes alterations in network properties that can not be explained solely by an acute reduction in PV^+ cell activity.

When counting the number of excitatory and inhibitory synapses we found an effect mainly on inhibitory to inhibitory connections. However, the observed changes in synaptic connections is not necessarily representative of the results obtained electrophysiologically as the animals were only placed in their home cage after surgery for a shorter time period compared to the animals undergoing the full experiment with electrophysiology. Thus, we might expect larger changes in synaptic reorganization if the animals were subjected to larger environmental changes or behavioral opportunities.

Computational models have shown that inhibitory connections are sufficient to generate the hexagonal structure of the grid pattern[19,55,109]. Thus, the dispersed grid cell patterns we observed after PNN removal is likely a consequence of reduced inhibition, especially since the increased out-of-field spiking occurs where grid cells presumably rely on inhibitory domination in normal circumstances. In line with experimental data, we observed that lowering excitatory synaptic strength onto inhibitory neurons caused reduced inhibitory spiking in the continuous attractor model. It also led to a decrease in mean spike rate of excitatory neurons, similar to what we observed in the experimental data. This result is, at first glance, counterintuitive as reduced inhibition would be expected to increase excitatory activity. However, the dynamical interactions in the model and the spatial outreach of connections have a large influence on activity levels of both inhibitory and excitatory neuron populations.

When modelling a system, one often use the most simplistic alternative. To impose as little assumptions as possible on how the grid fields in neuronal space is mapped to movement space, we did not include path integration. If we assume that the hexagonal pattern is perfectly mapped to movement space we would expect the fields observed in the model to be that observed experimentally. If path integration would be included we would likely use either speed and directional input as in Burak and Fiete[55] or place cell input as in Dunn et al.[172]. The quality of path integration would depend on several factors that are different in the two simulated scenarios (with and without PNNs) that would affect path integration differently. Spike irregularity has an effect on quality of path integration[55] and this is likely to be different with different levels of inhibition. Strength of speed and directional input might affect the attractor state differently as the attracting strength of the network is lower with lower inhibition, at a certain point of reduced inhibition the attractor vanish. This

[169]Vogels et al., 2011 [170]Cisneros-Franco and Villers-Sidani, 2019 [171]Devienne et al., 2019 [19]Pastoll et al., 2013 [55]Burak and Fiete, 2009 [109]Couey et al., 2013 [55]Burak and Fiete, 2009 [172]Dunn et al., 2017 [55]Burak and Fiete, 2009

would also be an issue if input was in the form of place cell input[172].

When choosing the neuronal model type we had several options, spiking neurons or rate-based neurons. We wanted to explore how changes in capacitance affected grid fields and we therefore chose spiking neurons. Since the exponential integrate-and-fire neuron has been used previously[19], we also chose this particular neuron model as it has been fit to experimental values of stellate cells, although we chose values for resistance and capacitance based on Tewari et al.[173]. However, we recognize that using a leaky integrate-and-fire neuron might have given similar results, but the edges of firing fields might have been different. In the exponential integrate-and-fire model the threshold is not modeled as a set value in which spike response occurs, but rather a rheobase value in which the exponential part of the model kicks in. Due to this fact, the edges of the firing fields might be expected to extend more graded in the exponential model if compared with a leaky model. We recognize that this could have been quantified, but it was not prioritized in the current study.

Paper III

In Paper III we present a novel application of Instrumental Variable (IV) causal analysis to neural data. We use IV analysis to estimate connection strengths between neurons which are being recorded via extracellular electrodes while a subset of those neurons are stimulated with relatively local, one-photon optogenetic excitation. We demonstrate problems with naive estimators of causal effects under influence of confounders, and show clear advantages when using the IV method over previously used naive methods based on correlations. We have found that this approach performs considerably better than the naïve method. Moreover, we have found that neither a naïve linear regression model nor a naïve cross correlation method produce reliable estimates of connectivity between neuron pairs. The IV approach effectively reverse engineers causality by looking at the response that is missing because of refractoriness which effectively allows better estimates of causal effects.

The IV method shows promise in being able to give insight in population connectivity, which is at best difficult with single neuron methods, such as with patch clamp or two-photon optogenetic stimulation. For example, if estimating the connectivity of the grid cell population one would require a large number of pairs in which connectivity is estimated.

At the moment, we have no ground-truth dataset at hand to test our technique and compare with other approaches. Ideally, we would have known causal effects from single-cell stimulation (e.g. from two-photon optogenetics) to establish causal effects. Such data should contain many randomly distributed, short and intensive stimulation trials combined with traditional optogenetics, designed in a way where refractoriness matter. Most optogenetic protocols use set stimulation frequency, which is not ideal for the IV method. To the best of our knowledge, a

[172]Dunn et al., 2017 [19]Pastoll et al., 2013 [173]Tewari et al., 2018

dataset with randomly distributed optogenetic stimulation, with short stimulation pulses on the timescale of the refractory period, is currently not available and prevents us from testing how good our estimator would work on experimental data. Future experiments are needed to obtain reliable insights into the validity and robustness of the IV method.

One aspect that might intervene with the predictive power of the IV method presented in Paper III is mainly effects due to optogenetic stimulation. When stimulating many neurons at the same time their co-activity or co-activity of downstream targets might change their strength of connections due to the stimulation. Especially considering the requirement of large number of stimulations to obtain sufficient predictive power. Moreover, if neurons that are co-stimulated synchronize over time, due to the highly synchronized input to many neurons at the same time, the Markov assumption (that each trial is independent from the previous) might be violated.

Also, stimulating multiple neurons may give overload of synaptic response such that a single spike will never make a difference. In such scenarios, the downstream neuron would respond on every trial, and this can be tested. Importantly, the stimulation introduce a different network setting than what would be expected from baseline activity.

With these considerations, the method might require additional considerations.

To maintain a simplistic description of neural dynamics we used current-based synapses with alpha-based shape and short synaptic time constants. This was chosen in order to reduce variance in transmission strength, as this synapse type is not dependent on the membrane potential. If conductance based synapses would have been chosen, the input current would vary with the membrane potential and the IV estimate would likely have larger variance, but this remains to be tested. The effect of conductance based synapses on transmission probability estimates should be quantified at a later stage of IV method development.

For the refractory period to be a good instrument, it is necessary that it is not overly affected by the network activity. This will clearly be problematic in many cases. After all, network activity affects neuron activity and hence refractoriness. However, there are multiple scenarios where refractoriness will be a good instrument. For example, if we have balanced excitation and inhibition, we may expect largely independent refractory states of individual neurons. If a neuron biophysically implements something like conditional Poisson spiking, its refractory states will be random conditioned on the network state. Importantly, we may expect the phase of a neuron to be far more random than the activity of the network as a whole.

In the proposed IV approach it is necessary to manually set the response windows of the stimulation, both in the presynaptic and postsynaptic neuron. If this could be done automatically, that would be an advantage.

One of the strengths of the IV estimator presented here is that it only requires one pair to be recorded because we can utilize the randomness of the refractory periods along with random stimulations. Under those assumptions, the IV estimator can produce actual causal estimates. One popular way of estimating

causal effects is fitting generalized linear models (GLMs) to simultaneously recorded neuron activities[174,175]. GLMs are multiple nonlinear regressions and require multiple neurons to perform well. Even if activity from all neurons were recorded, GLMs might fail to estimate causal connections[135]. Complete recordings are not possible in the mammalian brain, especially not in primates, where recordings include only a very small subset of the neurons involved in the actual computation. When using GLMs one may accurately estimate latency distributions and sequences of spikes from individual neurons. These ideas should, arguably, be merged with IV approaches.

[174]Pillow et al., 2008 [175]Roudi et al., 2009 [135]Das and Fiete, 2019

Chapter 5

Future perspectives

Multimodal models

Brains are robust and can easily cope with changes, e.g. novel environmental or contextual challenges. Therefore, neurons do not have only one mode of action. If we measure brain activity from the same brain area in different environments or contexts, we would expect neural activity also to be different[176]. Because of this, saying that the brain encodes stimuli might be an inaccurate statement[177] - as decoding would depend on environmental or contextual aspects. To simplify matters, it is common to isolate behavior to simplistic proxies of natural behavior and study one type of stimulus at a time.

One of the strengths of the work of Hodgkin and Huxley was that they initially modelled the action potential, but then, using this theory they were able to accurately predict how fast an action potential travels down an axon[13]. That is, their theory was multi modal. Given population theory, it is not enough that the proposed model solves a task in only one context, it should be able to also perform in other contextual settings. For example grid cells are active during decision making in auditory tasks[166], this has yet to be supported in grid cell models, but a recent framework where grid cells emerge from minimizing location-prediction in an arbitrary space might be a good candidate[76].

Novelty also has little representation in grid cell models, however, some models can already be interesting to test against experimental data. When rodents are introduced to a novel environment, grid cells realign by a uniform shift in phase and orientation[42], temporarily expand in scale alongside reduced spatial stability[178], and are shaped by environmental geometry[179]. The CAN model can show stretching of the single neuron response by an amplitude modulation of head direction inputs tuned to the relevant head direction[55]. In a recent study it is shown how landmark cells, such as border cells or place cells can be used to stabilize the grid pattern in a CAN model and perform error correction[49]. In this model, stable grid representations are obtained over time as input from landmark cells are learned, thus in familiar environments a stable representation is already achieved but in novel environments this has to be relearned - possibly explaining how realignments and the initial weak stability occur in novel environments. Notably, if grid cells utilize plasticity to stabilize representations in novel environments as would be in Ocko et al.[49], these are unlikely to be related to PNNs in the first place as that would reduce plasticity levels.

[176]Fusi et al., 2016 [177]Brette, 2018 [13]Hodgkin and Huxley, 1952 [166]Aronov et al., 2017 [76]Sorscher et al., 2019 [42]Fyhn et al., 2007 [178]Barry et al., 2012 [179]Krupic et al., 2015 [55]Burak and Fiete, 2009 [49]Ocko et al., 2018 [49]Ocko et al., 2018

Another recent model of single grid cells uses excitatory and inhibitory plasticity[180] to show that grid representations arise through self organization. They reproduce findings by Carpenter et al.[88] where grid cells merge local representations to form a global representation of localized environments when a localization border is removed. The model also achieves unstable patterns when learning new environments, but disregards the population representation evident in experiments with pairwise comparison[20,42].

Exploring the underlying principles which give rise to observed changes in grid representations, normative models might be good candidates where reinforcement learning is utilized such as Stachenfeld et al.[74], which does account for reshaping geometries, or maybe Banino et al.[65] where novelty signals may possibly be implemented to serve exploration in more realistic environments.

Causally understanding brain activity

Questions of causation are important if we are to understand the mechanisms underlying brain activity. Recent considerations suggest that even if we can causally identify relations between neurons we still might not be able to understand it's complexities. According to Ramaswamy[181] there is an algorithmic barrier to understanding the causal link between neural circuits and behavior, indicating that we need exponentially many experiments in neurons - briefly said, it would take more than a lifetime of experiments to understand just one brain. Even with relatively simple models of learning, such as convolutional neural networks that can learn to recognize cats and dogs from images, it can be very hard to go backwards - understanding the rule of learning by only looking at the network layout and output is hard[182]. Even if we cannot directly understand how the brain generates behavior through solely using experiments, we might be able to build models that narrow down the scope to emergent population activity of neural circuits that underlie computations. If we can identify networks that can be reduced to low dimensional models, we might be able to circumvent the problem introduced by Ramaswamy[181]. Instead of considering all the neurons in the brain we might only consider populations, by using dimensionality reduction methods like those proposed by Chaudhuri et al.[168]. With population models, it might be possible to understand how different populations causally interact and give rise to behavior since $N_{populations} << N_{neurons}$. When mapping out neural connections, e.g. with the IV method, we can focus on which neural circuits communicate on a large scale, and the circuit motifs. These are less likely to change with learning as opposed to the detailed neuron to neuron connections which are likely to be learned[182]. The framework in which to include these populations could for example be causal inference with frameworks such as DAGs, where each population is represented by a node and their activity is represented in such a way that the causal interactions between these populations

[180]Weber and Sprekeler, 2018 [88]Carpenter et al., 2017 [20]Yoon et al., 2013 [42]Fyhn et al., 2007 [74]Stachenfeld et al., 2017 [65]Banino et al., 2018 [181]Ramaswamy, 2019 [182]Lillicrap and Kording, 2019 [181]Ramaswamy, 2019 [168]Chaudhuri et al., 2019 [182]Lillicrap and Kording, 2019

can be estimated. However, networks are context dependent, where neurons show mixed selectivity[176], and whether models that respect this multi modality remains simple enough to be tractable to support a brain wide causal understanding, as suggested above, deserves attention.

Adaptation in CAN models

There are two major classes of principal neurons in MEC layer II, pyramidal cells and stellate cells (SCs), and reportedly two sub classes in each major class[141]. A large part of the grid cells in MEC are stellate cells[33]. Stellate cells can be identified by morphology and electrophysiological features, among them, adaptation which cause burst firing measured in inter spike interval (ISI) due to depolarization[109].

Adaptation has been investigated with neural field theory where a common aim is to determine stability in neural fields[183,184]. With adaptation a large variety of pattern formation occurs, in some cases it destabilizes stationary bumps, as each excitatory cell will as soon as excited, be diminished due to adaptation. In a grid cell model governed by CAN, the "dominant" grid cell will lose its competitive drive and a new cell will start being excited and thus take the lead. When this is happening throughout the network, a continuous attractor state may not emerge.

This raises the questions whether stellate cells can support continuous attractor dynamics due to adaptation. Moreover, stellate cells exhibit a wide variety of response properties which are typically neglected in modelling studies such as depolarizing after potential[141,185], sub-threshold membrane potential oscillations (MPOs)[185,186] and post inhibitory rebound spikes[160,187]. To assess whether it is possible to obtain a stable grid cell representation with adaptation, and other properties of the stellate cell, future work might use the adaptive exponential integrate-and-fire neuron[188] to model the stellate cell. Of course, this would not be a proof that adapting neurons can not show grids through CAN, but might show that this is a topic that deserves more attention.

Issues with optogenetics

Confounding is a major issue when regular optogenetics is employed to infer connectivity between neurons as in e.g.[136]. There will be a population response even though the light intensity falls quickly ($\approx 1/r^2$) with distance r from the end of the optic fiber. Since the illuminated area of brain tissue will expand by $\approx r^2$ the number of illuminated neurons per area is constant as a function of distance. Neuronal spiking is believed to be induced by small membrane potential fluctuations as neural networks *in vivo* are observed to be in a high conductance

[176]Fusi et al., 2016 [141]Fuchs et al., 2016 [33]Rowland et al., 2018 [109]Couey et al., 2013 [183]Coombes and Owen, 2005 [184]Ermentrout et al., 2014 [141]Fuchs et al., 2016 [185]Pastoll et al., 2012 [185]Pastoll et al., 2012 [186]Erchova et al., 2004 [160]Shay et al., 2016 [187]Ferrante et al., 2016 [188]Brette and Gerstner, 2005 [136]English et al., 2017

state where the membrane potential is generally close to threshold[189–191]. The population response from optogenetic stimulation will thus increase with distance since the photo induced current as a function of photo intensity falls proportional to a logarithmic function[192]. Moreover, indications that light may travel through the dendrites themselves[193], suggests a large spread of activation upon light stimulation.

[189]Destexhe et al., 2003 [190]Léger et al., 2005 [191]Kumar et al., 2008 [192]Wang et al., 2007 [193]Thunemann et al., 2018

Chapter 6

Conclusion

In conclusion, the results in the present thesis contribute to the understanding of mechanisms that underlie brain activity and how this may explain cognitive functions such as navigation and memory processing. Paper I contributes to the study of temporal activity in the brain, strongly suggesting oscillations not to underlie the hexagonal pattern formed by grid cells. Paper II contributes to the field of plasticity mechanisms and indicates that PNNs stabilize low dimensional activity of grid cells by restricting changes in neuronal connectivity. Paper III contributes to the field of connectivity inference which is crucial to assess structure versus function in neural networks.

As with any scientific study, the research raises more questions that deserve further investigation. From Paper I, whether oscillations in MEC may contribute to other behavioral aspects such as spatial memory or decision making remains to be shown. From Paper II, specifically which connectivity motifs are degraded, how much and the mode of action remains unclear. Moreover, if it is possible to completely eradicate the hexagonal pattern of grid cells returning the MEC network to it's juvenile state remains to be shown. From Paper III, how the method performs in real biological networks, and if it can be utilized to confirm the hypothesis that connectivity underlies grid cell low dimensional activity is still not clear.

Bibliography

1. Frederico AC Azevedo, Ludmila RB Carvalho, Lea T Grinberg, José Marcelo Farfel, Renata EL Ferretti, Renata EP Leite, Roberto Lent, Suzana Herculano-Houzel, et al. "Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain". In: *Journal of Comparative Neurology* 513.5 (2009), pp. 532–541.
2. R. G. Northcutt. "Understanding Vertebrate Brain Evolution". In: *Integrative and Comparative Biology* 42.4 (Aug. 2002), pp. 743–756. DOI: 10.1093/icb/42.4.743. URL: https://doi.org/10.1093/icb/42.4.743.
3. Menno P. Witter, Thanh P. Doan, Bente Jacobsen, Eirik S. Nilssen, and Shinya Ohara. "Architecture of the Entorhinal Cortex A Review of Entorhinal Anatomy in Rodents with Some Comparative Notes". In: *Frontiers in Systems Neuroscience* 11 (June 2017). DOI: 10.3389/fnsys.2017.00046. URL: https://doi.org/10.3389/fnsys.2017.00046.
4. M. Kamp, Y. Tahsim-Oglou, H.-J. Steiger, and D. Hänggi. "Traumatic Brain Injuries in the Ancient Egypt: Insights from the Edwin Smith Papyrus". In: *Central European Neurosurgery* 72.S 01 (Dec. 2011), pp. 001–001. DOI: 10.1055/s-0031-1275746. URL: https://doi.org/10.1055/s-0031-1275746.
5. Fridtjof Nansen. *The structure and combination of the histological elements of the central nervous system.* J. Grieg, 1887.
6. Francisco López-Muñoz, Jesús Boya, and Cecilio Alamo. "Neuron theory, the cornerstone of neuroscience, on the centenary of the Nobel Prize award to Santiago Ramón y Cajal". In: *Brain Research Bulletin* 70.4-6 (Oct. 2006), pp. 391–405. DOI: 10.1016/j.brainresbull.2006.07.010. URL: https://doi.org/10.1016/j.brainresbull.2006.07.010.
7. Rafael Yuste. "From the neuron doctrine to neural networks". In: *Nature Reviews Neuroscience* 16.8 (July 2015), pp. 487–497. DOI: 10.1038/nrn3962. URL: https://doi.org/10.1038/nrn3962.
8. Christopher S von Bartheld, Jami Bahney, and Suzana Herculano-Houzel. "The search for true numbers of neurons and glial cells in the human brain: a review of 150 years of cell counting". In: *Journal of Comparative Neurology* 524.18 (2016), pp. 3865–3895.
9. György Buzsáki. "Large-scale recording of neuronal ensembles". In: *Nature Neuroscience* 7.5 (Apr. 2004), pp. 446–451. DOI: 10.1038/nn1233. URL: https://doi.org/10.1038/nn1233.

10. Peter Barthó, Hajime Hirase, Lenaïc Monconduit, Michael Zugaro, Kenneth D. Harris, and György Buzsáki. "Characterization of Neocortical Principal Cells and Interneurons by Network Interactions and Extracellular Features". In: *Journal of Neurophysiology* 92.1 (July 2004), pp. 600–608. DOI: 10.1152/jn.01170.2003. URL: https://doi.org/10.1152/jn.01170.2003.

11. James J. Jun et al. "Fully integrated silicon probes for high-density recording of neural activity". In: *Nature* 551.7679 (Nov. 2017), pp. 232–236. DOI: 10.1038/nature24636. URL: https://doi.org/10.1038/nature24636.

12. Santiago Ramón y Cajal. *Comparative study of the sensory areas of the human cortex*. Clark University, 1899.

13. A. L. Hodgkin and A. F. Huxley. "A quantitative description of membrane current and its application to conduction and excitation in nerve". In: *The Journal of Physiology* 117.4 (Aug. 1952), pp. 500–544. DOI: 10.1113/jphysiol.1952.sp004764. URL: https://doi.org/10.1113/jphysiol.1952.sp004764.

14. Albert Gidon, Timothy Adam Zolnik, Pawel Fidzinski, Felix Bolduan, Athanasia Papoutsi, Panayiota Poirazi, Martin Holtkamp, Imre Vida, and Matthew Evan Larkum. "Dendritic action potentials and computation in human layer 2/3 cortical neurons". In: *Science* 367.6473 (Jan. 2020), pp. 83–87. DOI: 10.1126/science.aax6239. URL: https://doi.org/10.1126/science.aax6239.

15. Eugene M Izhikevich. *Dynamical systems in neuroscience*. MIT press, 2007.

16. Wulfram Gerstner, Werner M Kistler, Richard Naud, and Liam Paninski. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.

17. Laurent Badel, Sandrine Lefort, Thomas K. Berger, Carl C. H. Petersen, Wulfram Gerstner, and Magnus J. E. Richardson. "Extracting non-linear integrate-and-fire models from experimental data using dynamic I–V curves". In: *Biological Cybernetics* 99.4-5 (Nov. 2008), pp. 361–370. DOI: 10.1007/s00422-008-0259-4. URL: https://doi.org/10.1007/s00422-008-0259-4.

18. Nicolas Fourcaud-Trocmé, David Hansel, Carl van Vreeswijk, and Nicolas Brunel. "How Spike Generation Mechanisms Determine the Neuronal Response to Fluctuating Inputs". In: *The Journal of Neuroscience* 23.37 (Dec. 2003), pp. 11628–11640. DOI: 10.1523/jneurosci.23-37-11628.2003. URL: https://doi.org/10.1523/jneurosci.23-37-11628.2003.

19. Hugh Pastoll, Lukas Solanka, Mark C.W. van Rossum, and Matthew F. Nolan. "Feedback Inhibition Enables Theta-Nested Gamma Oscillations and Grid Firing Fields". In: *Neuron* 77.1 (Jan. 2013), pp. 141–154. DOI: 10.1016/j.neuron.2012.11.032. URL: https://doi.org/10.1016/j.neuron.2012.11.032.

20. KiJung Yoon, Michael A Buice, Caswell Barry, Robin Hayman, Neil Burgess, and Ila R Fiete. “Specific evidence of low-dimensional continuous attractor dynamics in grid cells”. In: *Nature Neuroscience* 16.8 (July 2013), pp. 1077–1084. DOI: 10.1038/nn.3450. URL: https://doi.org/10.1038/nn.3450.

21. E. C. Tolman, B. F. Ritchie, and D. Kalish. “Studies in spatial learning. I. Orientation and the short-cut.” In: *Journal of Experimental Psychology* 36.1 (1946), pp. 13–24. DOI: 10.1037/h0053944. URL: https://doi.org/10.1037/h0053944.

22. M. -L. Mittelstaedt and H. Mittelstaedt. “Homing by path integration in a mammal”. In: *Naturwissenschaften* 67.11 (Nov. 1980), pp. 566–567. DOI: 10.1007/bf00450672. URL: https://doi.org/10.1007/bf00450672.

23. W. B. Scoville and B. Milner. “Loss of recent memory after bilateral hippocampal lesions”. In: *Journal of Neurology, Neurosurgery & Psychiatry* 20.1 (Feb. 1957), pp. 11–21. DOI: 10.1136/jnnp.20.1.11. URL: https://doi.org/10.1136/jnnp.20.1.11.

24. J. O’Keefe and J. Dostrovsky. “The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat”. In: *Brain Research* 34.1 (Nov. 1971), pp. 171–175. DOI: 10.1016/0006-8993(71)90358-1. URL: https://doi.org/10.1016/0006-8993(71)90358-1.

25. John O’keefe and Lynn Nadel. *The hippocampus as a cognitive map*. Oxford: Clarendon Press, 1978.

26. J. O’Keefe and D.H. Conway. “Hippocampal place units in the freely moving rat: Why they fire where they fire”. In: *Experimental Brain Research* 31.4 (Apr. 1978). DOI: 10.1007/bf00239813. URL: https://doi.org/10.1007/bf00239813.

27. J. O’Keefe and A. Speakman. “Single unit activity in the rat hippocampus during a spatial memory task”. In: *Experimental Brain Research* 68.1 (Sept. 1987). DOI: 10.1007/bf00255230. URL: https://doi.org/10.1007/bf00255230.

28. GJ Quirk, RU Muller, and JL Kubie. “The firing of hippocampal place cells in the dark depends on the rat’s recent experience”. In: *The Journal of Neuroscience* 10.6 (June 1990), pp. 2008–2017. DOI: 10.1523/jneurosci.10-06-02008.1990. URL: https://doi.org/10.1523/jneurosci.10-06-02008.1990.

29. James J. Knierim and Derek A. Hamilton. “Framing Spatial Cognition: Neural Representations of Proximal and Distal Frames of Reference and Their Roles in Navigation”. In: *Physiological Reviews* 91.4 (Oct. 2011), pp. 1245–1279. DOI: 10.1152/physrev.00021.2010. URL: https://doi.org/10.1152/physrev.00021.2010.

30. Marianne Fyhn, Sturla Molden, Menno P Witter, Edvard I Moser, and May-Britt Moser. “Spatial Representation in the Entorhinal Cortex”. In: *Science* 305.5688 (Aug. 2004), pp. 1258–1264. DOI: 10.1126/science.1099901. URL: https://doi.org/10.1126/science.1099901.

31. Torkel Hafting, Marianne Fyhn, Sturla Molden, May-Britt Moser, and Edvard I. Moser. “Microstructure of a spatial map in the entorhinal cortex”. In: *Nature* 436.7052 (June 2005), pp. 801–806. DOI: 10.1038/nature03721. URL: https://doi.org/10.1038/nature03721.

32. Francesca Sargolini, Marianne Fyhn, Torkel Hafting, Bruce L McNaughton, Menno P Witter, May-Britt Moser, and Edvard I Moser. “Conjunctive Representation of Position, Direction, and Velocity in Entorhinal Cortex”. In: *Science* 312.5774 (May 2006), pp. 758–762. DOI: 10.1126/science.1125572. URL: https://doi.org/10.1126/science.1125572.

33. David C Rowland, Horst A Obenhaus, Emilie R Skytøen, Qiangwei Zhang, Cliff G Kentros, Edvard I Moser, and May-Britt Moser. “Functional properties of stellate cells in medial entorhinal cortex layer II”. In: *eLife* 7 (Sept. 2018). DOI: 10.7554/elife.36664. URL: https://doi.org/10.7554/elife.36664.

34. John O’Keefe and Neil Burgess. “Geometric determinants of the place fields of hippocampal neurons”. In: *Nature* 381.6581 (May 1996), pp. 425–428. DOI: 10.1038/381425a0. URL: https://doi.org/10.1038/381425a0.

35. P. Ravassard, A. Kees, B. Willers, D. Ho, D. Aharoni, J. Cushman, Z. M. Aghajan, and M. R. Mehta. “Multisensory Control of Hippocampal Spatiotemporal Selectivity”. In: *Science* 340.6138 (May 2013), pp. 1342–1346. DOI: 10.1126/science.1232655. URL: https://doi.org/10.1126/science.1232655.

36. JL Kubie. “Sensory-behavioral correlates in individual hippocampus neurons in three situations: Space and context”. In: *Neurobiology of the Hippocampus* (1983).

37. RU Muller and JL Kubie. “The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells”. In: *The Journal of Neuroscience* 7.7 (July 1987), pp. 1951–1968. DOI: 10.1523/jneurosci.07-07-01951.1987. URL: https://doi.org/10.1523/jneurosci.07-07-01951.1987.

38. KM Gothard, WE Skaggs, KM Moore, and BL McNaughton. “Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task”. In: *The Journal of Neuroscience* 16.2 (Jan. 1996), pp. 823–835. DOI: 10.1523/jneurosci.16-02-00823.1996. URL: https://doi.org/10.1523/jneurosci.16-02-00823.1996.

39. Dori Derdikman and James J Knierim. *Space, time and memory in the hippocampal formation*. Springer, 2014.

40. Caswell Barry, Robin Hayman, Neil Burgess, and Kathryn J Jeffery. "Experience-dependent rescaling of entorhinal grids". In: *Nature Neuroscience* 10.6 (May 2007), pp. 682–684. DOI: 10.1038/nn1905. URL: https://doi.org/10.1038/nn1905.

41. Francesco Savelli, D. Yoganarasimha, and James J. Knierim. "Influence of boundary removal on the spatial representations of the medial entorhinal cortex". In: *Hippocampus* 18.12 (Dec. 2008), pp. 1270–1282. DOI: 10.1002/hipo.20511. URL: https://doi.org/10.1002/hipo.20511.

42. Marianne Fyhn, Torkel Hafting, Alessandro Treves, May-Britt Moser, and Edvard I. Moser. "Hippocampal remapping and grid realignment in entorhinal cortex". In: *Nature* 446.7132 (Feb. 2007), pp. 190–194. DOI: 10.1038/nature05601. URL: https://doi.org/10.1038/nature05601.

43. Hanne Stensola, Tor Stensola, Trygve Solstad, Kristian Frøland, May-Britt Moser, and Edvard I. Moser. "The entorhinal grid map is discretized". In: *Nature* 492.7427 (Dec. 2012), pp. 72–78. DOI: 10.1038/nature11649. URL: https://doi.org/10.1038/nature11649.

44. Bard Ermentrout. "Neural networks as spatio-temporal pattern-forming systems". In: *Reports on Progress in Physics* 61.4 (Apr. 1998), pp. 353–430. DOI: 10.1088/0034-4885/61/4/002. URL: https://doi.org/10.1088/0034-4885/61/4/002.

45. Hugh R. Wilson and Jack D. Cowan. "Excitatory and Inhibitory Interactions in Localized Populations of Model Neurons". In: *Biophysical Journal* 12.1 (Jan. 1972), pp. 1–24. DOI: 10.1016/s0006-3495(72)86068-5. URL: https://doi.org/10.1016/s0006-3495(72)86068-5.

46. H. R. Wilson and J. D. Cowan. "A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue". In: *Kybernetik* 13.2 (Sept. 1973), pp. 55–80. DOI: 10.1007/bf00288786. URL: https://doi.org/10.1007/bf00288786.

47. Shun-ichi Amari. "Dynamics of pattern formation in lateral-inhibition type neural fields". In: *Biological Cybernetics* 27.2 (1977), pp. 77–87. DOI: 10.1007/bf00337259. URL: https://doi.org/10.1007/bf00337259.

48. Gustavo Deco, Viktor K. Jirsa, Peter A. Robinson, Michael Breakspear, and Karl Friston. "The Dynamic Brain: From Spiking Neurons to Neural Masses and Cortical Fields". In: *PLoS Computational Biology* 4.8 (Aug. 2008). Ed. by Olaf Sporns, e1000092. DOI: 10.1371/journal.pcbi.1000092. URL: https://doi.org/10.1371/journal.pcbi.1000092.

49. Samuel A. Ocko, Kiah Hardcastle, Lisa M. Giocomo, and Surya Ganguli. "Emergent elasticity in the neural code for space". In: *Proceedings of the National Academy of Sciences* 115.50 (Nov. 2018), E11798–E11806. DOI: 10.1073/pnas.1805959115. URL: https://doi.org/10.1073/pnas.1805959115.

50. K Zhang. “Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory”. In: *The Journal of Neuroscience* 16.6 (Mar. 1996), pp. 2112–2126. DOI: 10.1523/jneurosci.16-06-02112.1996. URL: https://doi.org/10.1523/jneurosci.16-06-02112.1996.

51. Alan Mathison Turing. “The chemical basis of morphogenesis”. In: *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 237.641 (Aug. 1952), pp. 37–72. DOI: 10.1098/rstb.1952.0012. URL: https://doi.org/10.1098/rstb.1952.0012.

52. Alexei Samsonovich and Bruce L. McNaughton. “Path Integration and Cognitive Mapping in a Continuous Attractor Neural Network Model”. In: *The Journal of Neuroscience* 17.15 (Aug. 1997), pp. 5900–5920. DOI: 10.1523/jneurosci.17-15-05900.1997. URL: https://doi.org/10.1523/jneurosci.17-15-05900.1997.

53. M. C. Fuhs. “A Spin Glass Model of Path Integration in Rat Medial Entorhinal Cortex”. In: *Journal of Neuroscience* 26.16 (Apr. 2006), pp. 4266–4276. DOI: 10.1523/jneurosci.4353-05.2006. URL: https://doi.org/10.1523/jneurosci.4353-05.2006.

54. Bruce L. McNaughton, Francesco P. Battaglia, Ole Jensen, Edvard I Moser, and May-Britt Moser. “Path integration and the neural basis of the ’cognitive map’”. In: *Nature Reviews Neuroscience* 7.8 (Aug. 2006), pp. 663–678. DOI: 10.1038/nrn1932. URL: https://doi.org/10.1038/nrn1932.

55. Yoram Burak and Ila R. Fiete. “Accurate Path Integration in Continuous Attractor Network Models of Grid Cells”. In: *PLoS Computational Biology* 5.2 (Feb. 2009). Ed. by Olaf Sporns, e1000291. DOI: 10.1371/journal.pcbi.1000291. URL: https://doi.org/10.1371/journal.pcbi.1000291.

56. Hai-Chau Chang and Lih-Chung Wang. “A simple proof of Thue’s Theorem on circle packing”. In: *arXiv preprint arXiv:1009.4322* (2010).

57. Alexis Guanella, Daniel Kiper, and Paul Verschure. “A model of grid cells based on a twisted torus topology”. In: *International Journal of Neural Systems* 17.04 (Aug. 2007), pp. 231–240. DOI: 10.1142/s0129065707001093. URL: https://doi.org/10.1142/s0129065707001093.

58. Eric A. Zilli. “Models of Grid Cell Spatial Firing Published 2005–2011”. In: *Frontiers in Neural Circuits* 6 (2011). DOI: 10.3389/fncir.2012.00016. URL: https://doi.org/10.3389/fncir.2012.00016.

59. Martin Stemmler, Alexander Mathis, and Andreas V. M. Herz. “Connecting multiple spatial scales to decode the population activity of grid cells”. In: *Science Advances* 1.11 (Dec. 2015), e1500816. DOI: 10.1126/science.1500816. URL: https://doi.org/10.1126/science.1500816.

60. Trygve Solstad, Edvard I. Moser, and Gaute T. Einevoll. “From grid cells to place cells: A mathematical model”. In: *Hippocampus* 16.12 (Dec. 2006), pp. 1026–1031. DOI: 10.1002/hipo.20244. URL: https://doi.org/10.1002/hipo.20244.

61. Vegard Heimly Brun, Trygve Solstad, Kirsten Brun Kjelstrup, Marianne Fyhn, Menno P. Witter, Edvard I. Moser, and May-Britt Moser. "Progressive increase in grid scale from dorsal to ventral medial entorhinal cortex". In: *Hippocampus* 18.12 (Dec. 2008), pp. 1200–1212. DOI: 10.1002/hipo.20504. URL: https://doi.org/10.1002/hipo.20504.

62. I. R. Fiete, Y. Burak, and T. Brookings. "What Grid Cells Convey about Rat Location". In: *Journal of Neuroscience* 28.27 (July 2008), pp. 6858–6871. DOI: 10.1523/jneurosci.5684-07.2008. URL: https://doi.org/10.1523/jneurosci.5684-07.2008.

63. Alexander Mathis, Andreas V. M. Herz, and Martin Stemmler. "Optimal Population Codes for Space: Grid Cells Outperform Place Cells". In: *Neural Computation* 24.9 (Sept. 2012), pp. 2280–2317. DOI: 10.1162/neco_a_00319. URL: https://doi.org/10.1162/neco_a_00319.

64. Daniel Bush, Caswell Barry, Daniel Manson, and Neil Burgess. "Using Grid Cells for Navigation". In: *Neuron* 87.3 (Aug. 2015), pp. 507–520. DOI: 10.1016/j.neuron.2015.07.006. URL: https://doi.org/10.1016/j.neuron.2015.07.006.

65. Andrea Banino et al. "Vector-based navigation using grid-like representations in artificial agents". In: *Nature* 557.7705 (May 2018), pp. 429–433. DOI: 10.1038/s41586-018-0102-6. URL: https://doi.org/10.1038/s41586-018-0102-6.

66. Mariana Gil, Mihai Ancau, Magdalene I. Schlesiger, Angela Neitz, Kevin Allen, Rodrigo J. De Marco, and Hannah Monyer. "Impaired path integration in mice with disrupted grid cell firing". In: *Nature Neuroscience* 21.1 (Dec. 2017), pp. 81–91. DOI: 10.1038/s41593-017-0039-3. URL: https://doi.org/10.1038/s41593-017-0039-3.

67. Jena B. Hales, Magdalene I. Schlesiger, Jill K. Leutgeb, Larry R. Squire, Stefan Leutgeb, and Robert E. Clark. "Medial Entorhinal Cortex Lesions Only Partially Disrupt Hippocampal Place Cells and Hippocampus-Dependent Place Memory". In: *Cell Reports* 9.3 (Nov. 2014), pp. 893–901. DOI: 10.1016/j.celrep.2014.10.009. URL: https://doi.org/10.1016/j.celrep.2014.10.009.

68. R. F. Langston, J. A. Ainge, J. J. Couey, C. B. Canto, T. L. Bjerknes, M. P. Witter, E. I. Moser, and M.-B. Moser. "Development of the Spatial Representation System in the Rat". In: *Science* 328.5985 (June 2010), pp. 1576–1580. DOI: 10.1126/science.1188210. URL: https://doi.org/10.1126/science.1188210.

69. T. J. Wills, F. Cacucci, N. Burgess, and J. O'Keefe. "Development of the Hippocampal Cognitive Map in Preweanling Rats". In: *Science* 328.5985 (June 2010), pp. 1573–1576. DOI: 10.1126/science.1188224. URL: https://doi.org/10.1126/science.1188224.

70. Tora Bonnevie, Benjamin Dunn, Marianne Fyhn, Torkel Hafting, Dori Derdikman, John L Kubie, Yasser Roudi, Edvard I Moser, and May-Britt Moser. "Grid cells require excitatory drive from the hippocampus". In: *Nature Neuroscience* 16.3 (Jan. 2013), pp. 309–317. DOI: 10.1038/nn.3311. URL: https://doi.org/10.1038/nn.3311.

71. Emilio Kropff and Alessandro Treves. "The emergence of grid cells: Intelligent design or just adaptation?" In: *Hippocampus* 18.12 (Dec. 2008), pp. 1256–1269. DOI: 10.1002/hipo.20520. URL: https://doi.org/10.1002/hipo.20520.

72. Bailu Si and Alessandro Treves. "A model for the differentiation between grid and conjunctive units in medial entorhinal cortex". In: *Hippocampus* 23.12 (Sept. 2013), pp. 1410–1424. DOI: 10.1002/hipo.22194. URL: https://doi.org/10.1002/hipo.22194.

73. Yedidyah Dordek, Daniel Soudry, Ron Meir, and Dori Derdikman. "Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis". In: *eLife* 5 (Mar. 2016). DOI: 10.7554/elife.10094. URL: https://doi.org/10.7554/elife.10094.

74. Kimberly L Stachenfeld, Matthew M Botvinick, and Samuel J Gershman. "The hippocampus as a predictive map". In: *Nature Neuroscience* 20.11 (Oct. 2017), pp. 1643–1653. DOI: 10.1038/nn.4650. URL: https://doi.org/10.1038/nn.4650.

75. Christopher J Cueva and Xue-Xin Wei. "Emergence of grid-like representations by training recurrent neural networks to perform spatial localization". In: *arXiv preprint arXiv:1803.07770* (2018).

76. Ben Sorscher, Gabriel Mel, Surya Ganguli, and Samuel Ocko. "A unified theory for the origin of grid cells through the lens of pattern formation". In: *Advances in Neural Information Processing Systems 32*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett. Curran Associates, Inc., 2019, pp. 10003–10013. URL: http://papers.nips.cc/paper/9191-a-unified-theory-for-the-origin-of-grid-cells-through-the-lens-of-pattern-formation.pdf.

77. Lisa M. Giocomo, May-Britt Moser, and Edvard I. Moser. "Computational Models of Grid Cells". In: *Neuron* 71.4 (Aug. 2011), pp. 589–603. DOI: 10.1016/j.neuron.2011.07.023. URL: https://doi.org/10.1016/j.neuron.2011.07.023.

78. György Buzsáki and Edvard I Moser. "Memory, navigation and theta rhythm in the hippocampal-entorhinal system". In: *Nature Neuroscience* 16.2 (Jan. 2013), pp. 130–138. DOI: 10.1038/nn.3304. URL: https://doi.org/10.1038/nn.3304.

79. H. Petsche, Ch. Stumpf, and G. Gogolak. "The significance of the rabbit's septum as a relay station between the midbrain and the hippocampus I. The control of hippocampus arousal activity by the septum cells". In: *Electroencephalography and Clinical Neurophysiology* 14.2 (Apr. 1962), pp. 202–211. DOI: 10.1016/0013-4694(62)90030-5. URL: https://doi.org/10.1016/0013-4694(62)90030-5.

80. SJ Mitchell, JN Rawlins, O Steward, and DS Olton. "Medial septal area lesions disrupt theta rhythm and cholinergic staining in medial entorhinal cortex and produce impaired radial arm maze behavior in rats". In: *The Journal of Neuroscience* 2.3 (Mar. 1982), pp. 292–302. DOI: 10.1523/jneurosci.02-03-00292.1982. URL: https://doi.org/10.1523/jneurosci.02-03-00292.1982.

81. Christina Müller and Stefan Remy. "Septo–hippocampal interaction". In: *Cell and Tissue Research* 373.3 (Dec. 2017), pp. 565–575. DOI: 10.1007/s00441-017-2745-2. URL: https://doi.org/10.1007/s00441-017-2745-2.

82. Richardson N Leão, Zé H Targino, Luis V Colom, and André Fisahn. "Interconnection and synchronization of neuronal populations in the mouse medial septum/diagonal band of Broca". In: *Journal of neurophysiology* 113.3 (2014), pp. 971–980. DOI: 10.1152/jn.00367.2014. URL: https://doi.org/10.1152/jn.00367.2014.

83. I.D Manns, L Mainville, and B.E Jones. "Evidence for glutamate, in addition to acetylcholine and GABA, neurotransmitter synthesis in basal forebrain neurons projecting to the entorhinal cortex". In: *Neuroscience* 107.2 (Nov. 2001), pp. 249–263. DOI: 10.1016/s0306-4522(01)00302-5. URL: https://doi.org/10.1016/s0306-4522(01)00302-5.

84. Alfredo Gonzalez-Sulser, Daniel Parthier, Antonio Candela, Christina McClure, Hugh Pastoll, Derek Garden, Gülşen Sürmeli, and Matthew F Nolan. "GABAergic projections from the medial septum selectively inhibit interneurons in the medial entorhinal cortex". In: *Journal of Neuroscience* 34.50 (2014), pp. 16739–16743.

85. Gunes Unal, Abhilasha Joshi, Tim J Viney, Viktor Kis, and Peter Somogyi. "Synaptic targets of medial septal projections in the hippocampus and extrahippocampal cortices of the mouse". In: *Journal of Neuroscience* 35.48 (2015), pp. 15812–15826.

86. Falko Fuhrmann et al. "Locomotion, Theta Oscillations, and the Speed-Correlated Firing of Hippocampal Neurons Are Controlled by a Medial Septal Glutamatergic Circuit". In: *Neuron* 86.5 (June 2015), pp. 1253–1264. DOI: 10.1016/j.neuron.2015.05.001. URL: https://doi.org/10.1016/j.neuron.2015.05.001.

87. Daniel Justus et al. "Glutamatergic synaptic integration of locomotion speed via septoentorhinal projections". In: *Nature Neuroscience* 20.1 (Nov. 2016), pp. 16–19. DOI: 10.1038/nn.4447. URL: https://doi.org/10.1038/nn.4447.

88. Francis Carpenter, Neil Burgess, and Caswell Barry. “Modulating medial septal cholinergic activity reduces medial entorhinal theta frequency without affecting speed or grid coding”. In: *Scientific Reports* 7.1 (Nov. 2017). DOI: 10.1038/s41598-017-15100-6. URL: https://doi.org/10.1038/s41598-017-15100-6.

89. Laura Lee Colgin, Tobias Denninger, Marianne Fyhn, Torkel Hafting, Tora Bonnevie, Ole Jensen, May-Britt Moser, and Edvard I. Moser. “Frequency of gamma oscillations routes flow of information in the hippocampus”. In: *Nature* 462.7271 (Nov. 2009), pp. 353–357. DOI: 10.1038/nature08573. URL: https://doi.org/10.1038/nature08573.

90. J. Koenig, A. N. Linder, J. K. Leutgeb, and S. Leutgeb. “The Spatial Periodicity of Grid Cells Is Not Sustained During Reduced Theta Oscillations”. In: *Science* 332.6029 (Apr. 2011), pp. 592–595. DOI: 10.1126/science.1201685. URL: https://doi.org/10.1126/science.1201685.

91. M. P. Brandon, A. R. Bogaard, C. P. Libby, M. A. Connerney, K. Gupta, and M. E. Hasselmo. “Reduction of Theta Rhythm Dissociates Grid Cell Spatial Periodicity from Directional Tuning”. In: *Science* 332.6029 (Apr. 2011), pp. 595–599. DOI: 10.1126/science.1201652. URL: https://doi.org/10.1126/science.1201652.

92. Shawn S. Winter, Max L. Mehlman, Benjamin J. Clark, and Jeffrey S. Taube. “Passive Transport Disrupts Grid Signals in the Parahippocampal Cortex”. In: *Current Biology* 25.19 (Oct. 2015), pp. 2493–2502. DOI: 10.1016/j.cub.2015.08.034. URL: https://doi.org/10.1016/j.cub.2015.08.034.

93. Neil Burgess, Caswell Barry, and John O’Keefe. “An oscillatory interference model of grid cell firing”. In: *Hippocampus* 17.9 (2007), pp. 801–812. DOI: 10.1002/hipo.20327. URL: https://doi.org/10.1002/hipo.20327.

94. Christoph Schmidt-Hieber and Michael Häusser. “Cellular mechanisms of spatial navigation in the medial entorhinal cortex”. In: *Nature Neuroscience* 16.3 (Feb. 2013), pp. 325–331. DOI: 10.1038/nn.3340. URL: https://doi.org/10.1038/nn.3340.

95. John O’Keefe and Michael L. Recce. “Phase relationship between hippocampal place units and the EEG theta rhythm”. In: *Hippocampus* 3.3 (July 1993), pp. 317–330. DOI: 10.1002/hipo.450030307. URL: https://doi.org/10.1002/hipo.450030307.

96. Torkel Hafting, Marianne Fyhn, Tora Bonnevie, May-Britt Moser, and Edvard I. Moser. “Hippocampus-independent phase precession in entorhinal grid cells”. In: *Nature* 453.7199 (May 2008), pp. 1248–1252. DOI: 10.1038/nature06957. URL: https://doi.org/10.1038/nature06957.

97. E. T. Reifenstein, R. Kempter, S. Schreiber, M. B. Stemmler, and A. V. M. Herz. “Grid cells in rat entorhinal cortex encode physical space with independent firing fields and phase precession at the single-trial level”. In: *Proceedings of the National Academy of Sciences* 109.16 (Apr. 2012),

pp. 6301–6306. DOI: 10.1073/pnas.1109599109. URL: https://doi.org/10.1073/pnas.1109599109.

98. Cristina Domnisoru, Amina A. Kinkhabwala, and David W. Tank. "Membrane potential dynamics of grid cells". In: *Nature* 495.7440 (Feb. 2013), pp. 199–204. DOI: 10.1038/nature11973. URL: https://doi.org/10.1038/nature11973.

99. Daniel Bush and Neil Burgess. "A Hybrid Oscillatory Interference/Continuous Attractor Network Model of Grid Cell Firing". In: *The Journal of Neuroscience* 34.14 (Apr. 2014), pp. 5065–5079. DOI: 10.1523/jneurosci.4017-13.2014. URL: https://doi.org/10.1523/jneurosci.4017-13.2014.

100. Richard J. Gardner, Li Lu, Tanja Wernle, May-Britt Moser, and Edvard I. Moser. "Correlation structure of grid cells is preserved during sleep". In: *Nature Neuroscience* 22.4 (Mar. 2019), pp. 598–608. DOI: 10.1038/s41593-019-0360-0. URL: https://doi.org/10.1038/s41593-019-0360-0.

101. Sean G. Trettel, John B. Trimper, Ernie Hwaun, Ila R. Fiete, and Laura Lee Colgin. "Grid cell co-activity patterns during sleep reflect spatial overlap of grid fields during active behaviors". In: *Nature Neuroscience* 22.4 (Mar. 2019), pp. 609–617. DOI: 10.1038/s41593-019-0359-6. URL: https://doi.org/10.1038/s41593-019-0359-6.

102. Chenglin Miao, Qichen Cao, May-Britt Moser, and Edvard I. Moser. "Parvalbumin and Somatostatin Interneurons Control Different Space-Coding Networks in the Medial Entorhinal Cortex". In: *Cell* 171.3 (Oct. 2017), 507–521.e17. DOI: 10.1016/j.cell.2017.08.050. URL: https://doi.org/10.1016/j.cell.2017.08.050.

103. Benjamin R. Kanter, Christine M. Lykken, Daniel Avesar, Aldis Weible, Jasmine Dickinson, Benjamin Dunn, Nils Z. Borgesius, Yasser Roudi, and Clifford G. Kentros. "A Novel Mechanism for the Grid-to-Place Cell Transformation Revealed by Transgenic Depolarization of Medial Entorhinal Cortex Layer II". In: *Neuron* 93.6 (Mar. 2017), 1480–1492.e6. DOI: 10.1016/j.neuron.2017.03.001. URL: https://doi.org/10.1016/j.neuron.2017.03.001.

104. Ipshita Zutshi, Maylin L. Fu, Varoth Lilascharoen, Jill K. Leutgeb, Byung Kook Lim, and Stefan Leutgeb. "Recurrent circuits within medial entorhinal cortex superficial layers support grid cell firing". In: *Nature Communications* 9.1 (Sept. 2018). DOI: 10.1038/s41467-018-06104-5. URL: https://doi.org/10.1038/s41467-018-06104-5.

105. Takao K. Hensch. "Critical period plasticity in local cortical circuits". In: *Nature Reviews Neuroscience* 6.11 (Nov. 2005), pp. 877–888. DOI: 10.1038/nrn1787. URL: https://doi.org/10.1038/nrn1787.

106. J. Sebastian Espinosa and Michael P. Stryker. "Development and Plasticity of the Primary Visual Cortex". In: *Neuron* 75.2 (July 2012), pp. 230–249. DOI: 10.1016/j.neuron.2012.06.009. URL: https://doi.org/10.1016/j.neuron.2012.06.009.

107. May-Britt Moser, David C. Rowland, and Edvard I. Moser. “Place Cells, Grid Cells, and Memory”. In: *Cold Spring Harbor Perspectives in Biology* 7.2 (Feb. 2015), a021808. DOI: 10.1101/cshperspect.a021808. URL: https://doi.org/10.1101/cshperspect.a021808.

108. Kristian Kinden Lensjø, Ane Charlotte Christensen, Simen Tennøe, Marianne Fyhn, and Torkel Hafting. “Differential Expression and Cell-Type Specificity of Perineuronal Nets in Hippocampus, Medial Entorhinal Cortex, and Visual Cortex Examined in the Rat and Mouse”. In: *eneuro* 4.3 (2017), ENEURO.0379–16.2017. DOI: 10.1523/eneuro.0379-16.2017. URL: https://doi.org/10.1523/eneuro.0379-16.2017.

109. Jonathan J Couey et al. “Recurrent inhibitory circuitry as a mechanism for grid formation”. In: *Nature Neuroscience* 16.3 (Jan. 2013), pp. 318–324. DOI: 10.1038/nn.3310. URL: https://doi.org/10.1038/nn.3310.

110. Christina Buetfering, Kevin Allen, and Hannah Monyer. “Parvalbumin interneurons provide grid cell–driven recurrent inhibition in the medial entorhinal cortex”. In: *Nature Neuroscience* 17.5 (Apr. 2014), pp. 710–718. DOI: 10.1038/nn.3696. URL: https://doi.org/10.1038/nn.3696.

111. S. Hockfield, R.G. Kalb, S. Zaremba, and H. Fryer. “Expression of Neural Proteoglycans Correlates with the Acquisition of Mature Neuronal Properties in the Mammalian Brain”. In: *Cold Spring Harbor Symposia on Quantitative Biology* 55 (Jan. 1990), pp. 505–514. DOI: 10.1101/sqb.1990.055.01.049. URL: https://doi.org/10.1101/sqb.1990.055.01.049.

112. Flavio Donato, Santiago Belluco Rompani, and Pico Caroni. “Parvalbumin-expressing basket-cell network plasticity induced by experience regulates adult learning”. In: *Nature* 504.7479 (Dec. 2013), pp. 272–276. DOI: 10.1038/nature12866. URL: https://doi.org/10.1038/nature12866.

113. James W. Fawcett, Toshitaka Oohashi, and Tommaso Pizzorusso. “The roles of perineuronal nets and the perinodal extracellular matrix in neuronal function”. In: *Nature Reviews Neuroscience* 20.8 (July 2019), pp. 451–465. DOI: 10.1038/s41583-019-0196-3. URL: https://doi.org/10.1038/s41583-019-0196-3.

114. T. Pizzorusso. “Reactivation of Ocular Dominance Plasticity in the Adult Visual Cortex”. In: *Science* 298.5596 (Nov. 2002), pp. 1248–1251. DOI: 10.1126/science.1072699. URL: https://doi.org/10.1126/science.1072699.

115. Kristian Kinden Lensjø, Mikkel Elle Lepperød, Gunnar Dick, Torkel Hafting, and Marianne Fyhn. “Removal of Perineuronal Nets Unlocks Juvenile Plasticity Through Network Mechanisms of Decreased Inhibition and Increased Gamma Activity”. In: *The Journal of Neuroscience* 37.5 (Dec. 2016), pp. 1269–1283. DOI: 10.1523/jneurosci.2504-16.2016. URL: https://doi.org/10.1523/jneurosci.2504-16.2016.

116. M. F. K. Happel, H. Niekisch, L. L. Castiblanco Rivera, F. W. Ohl, M. Deliano, and R. Frischknecht. “Enhanced cognitive flexibility in reversal learning induced by removal of the extracellular matrix in auditory cortex”. In: *Proceedings of the National Academy of Sciences* 111.7 (Feb. 2014), pp. 2800–2805. DOI: 10.1073/pnas.1310272111. URL: https://doi.org/10.1073/pnas.1310272111.

117. C. Romberg et al. “Depletion of Perineuronal Nets Enhances Recognition Memory and Long-Term Depression in the Perirhinal Cortex”. In: *Journal of Neuroscience* 33.16 (Apr. 2013), pp. 7057–7065. DOI: 10.1523/jneurosci.6267-11.2013. URL: https://doi.org/10.1523/jneurosci.6267-11.2013.

118. N. Gogolla, P. Caroni, A. Luthi, and C. Herry. “Perineuronal Nets Protect Fear Memories from Erasure”. In: *Science* 325.5945 (Sept. 2009), pp. 1258–1261. DOI: 10.1126/science.1174146. URL: https://doi.org/10.1126/science.1174146.

119. Elise Holter Thompson, Kristian Kinden Lensjø, Mattis Brænne Wigestrand, Anders Malthe-Sørenssen, Torkel Hafting, and Marianne Fyhn. “Removal of perineuronal nets disrupts recall of a remote fear memory”. In: *Proceedings of the National Academy of Sciences* 115.3 (Dec. 2017), pp. 607–612. DOI: 10.1073/pnas.1713530115. URL: https://doi.org/10.1073/pnas.1713530115.

120. Stanford Encyclopedia of Philosophy Archive. *David Hume (accessed January 24, 2020)*. 2001. URL: https://plato.stanford.edu/archives/spr2017/entries/hume/#Cau.

121. Donald B. Rubin. “Estimating causal effects of treatments in randomized and nonrandomized studies.” In: *Journal of Educational Psychology* 66.5 (1974), pp. 688–701. DOI: 10.1037/h0037350. URL: https://doi.org/10.1037/h0037350.

122. Judea Pearl. “Causality”. In: *Causality, by Judea Pearl, pp. 400. ISBN 0521773628. Cambridge, UK: Cambridge University Press, March 2000.* (2000), p. 400.

123. Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. Springer New York, 1993. DOI: 10.1007/978-1-4612-2748-9. URL: https://doi.org/10.1007/978-1-4612-2748-9.

124. Judea Pearl and Dana Mackenzie. *The Book of Why: The New Science of Cause and Effect*. Basic Books, 2018.

125. S. Wright. “The Relative Importance of Heredity and Environment in Determining the Piebald Pattern of Guinea-Pigs”. In: *Proceedings of the National Academy of Sciences* 6.6 (June 1920), pp. 320–332. DOI: 10.1073/pnas.6.6.320. URL: https://doi.org/10.1073/pnas.6.6.320.

126. F. H. C. Crick. “Thinking about the Brain”. In: *Scientific American* 241.3 (1979), pp. 219–233. ISSN: 00368733, 19467087. URL: http://www.jstor.org/stable/24965297.

127. T. Kitamura, M. Pignatelli, J. Suh, K. Kohara, A. Yoshiki, K. Abe, and S. Tonegawa. “Island Cells Control Temporal Association Memory”. In: *Science* 343.6173 (Jan. 2014), pp. 896–901. DOI: 10.1126/science.1244634. URL: https://doi.org/10.1126/science.1244634.

128. C. Pesold, F. Impagnatiello, M. G. Pisu, D. P. Uzunov, E. Costa, A. Guidotti, and H. J. Caruncho. “Reelin is preferentially expressed in neurons synthesizing -aminobutyric acid in cortex and hippocampus of adult rats”. In: *Proceedings of the National Academy of Sciences* 95.6 (Mar. 1998), pp. 3221–3226. DOI: 10.1073/pnas.95.6.3221. URL: https://doi.org/10.1073/pnas.95.6.3221.

129. Karl Deisseroth. “Optogenetics: 10 years of microbial opsins in neuroscience”. In: *Nature Neuroscience* 18.9 (Aug. 2015), pp. 1213–1225. DOI: 10.1038/nn.4091. URL: https://doi.org/10.1038/nn.4091.

130. Boris V. Zemelman, Georgia A. Lee, Minna Ng, and Gero Miesenböck. “Selective Photostimulation of Genetically ChARGed Neurons”. In: *Neuron* 33.1 (Jan. 2002), pp. 15–22. DOI: 10.1016/s0896-6273(01)00574-8. URL: https://doi.org/10.1016/s0896-6273(01)00574-8.

131. Edward S Boyden, Feng Zhang, Ernst Bamberg, Georg Nagel, and Karl Deisseroth. “Millisecond-timescale, genetically targeted optical control of neural activity”. In: *Nature Neuroscience* 8.9 (Aug. 2005), pp. 1263–1268. DOI: 10.1038/nn1525. URL: https://doi.org/10.1038/nn1525.

132. Jai Y. Yu et al. “Knock-in rats expressing Cre and Flp recombinases at the Parvalbumin locus”. In: (Aug. 2018). DOI: 10.1101/386474. URL: https://doi.org/10.1101/386474.

133. Benjamin Dunn, Maria Mørreaunet, and Yasser Roudi. “Correlations and Functional Connections in a Population of Grid Cells”. In: *PLOS Computational Biology* 11.2 (Feb. 2015). Ed. by Olaf Sporns, e1004052. DOI: 10.1371/journal.pcbi.1004052. URL: https://doi.org/10.1371/journal.pcbi.1004052.

134. Gilad Tocker, Omri Barak, and Dori Derdikman. “Grid cells correlation structure suggests organized feedforward projections into superficial layers of the medial entorhinal cortex”. In: *Hippocampus* 25.12 (July 2015), pp. 1599–1613. DOI: 10.1002/hipo.22481. URL: https://doi.org/10.1002/hipo.22481.

135. Abhranil Das and Ila R Fiete. “Systematic errors in connectivity inferred from activity in strongly coupled recurrent circuits”. In: (Jan. 2019). DOI: 10.1101/512053. URL: https://doi.org/10.1101/512053.

136. Daniel Fine English, Sam McKenzie, Talfan Evans, Kanghwan Kim, Euisik Yoon, and György Buzsáki. “Pyramidal Cell-Interneuron Circuit Architecture and Dynamics in Hippocampal Networks”. In: *Neuron* 96.2 (Oct. 2017), 505–520.e7. DOI: 10.1016/j.neuron.2017.09.033. URL: https://doi.org/10.1016/j.neuron.2017.09.033.

137. Ioana E. Marinescu, Patrick N. Lawlor, and Konrad P. Kording. "Quasi-experimental causality in neuroscience and behavioural research". In: *Nature Human Behaviour* 2.12 (Nov. 2018), pp. 891–898. DOI: 10.1038/s41562-018-0466-5. URL: https://doi.org/10.1038/s41562-018-0466-5.

138. David Marc Anton Mehler and Konrad Paul Kording. "The lure of causal statements: Rampant mis-inference of causality in estimated connectivity". In: *arXiv preprint arXiv:1812.03363* (2018).

139. György Buzsáki and Kenji Mizuseki. "The log-dynamic brain: how skewed distributions affect network operations". In: *Nature Reviews Neuroscience* 15.4 (Feb. 2014), pp. 264–278. DOI: 10.1038/nrn3687. URL: https://doi.org/10.1038/nrn3687.

140. Marlene R Cohen and Adam Kohn. "Measuring and interpreting neuronal correlations". In: *Nature Neuroscience* 14.7 (June 2011), pp. 811–819. DOI: 10.1038/nn.2842. URL: https://doi.org/10.1038/nn.2842.

141. Elke C. Fuchs, Angela Neitz, Roberta Pinna, Sarah Melzer, Antonio Caputi, and Hannah Monyer. "Local and Distant Input Controlling Excitation in Layer II of the Medial Entorhinal Cortex". In: *Neuron* 89.1 (Jan. 2016), pp. 194–208. DOI: 10.1016/j.neuron.2015.11.029. URL: https://doi.org/10.1016/j.neuron.2015.11.029.

142. Tale L. Bjerknes, Edvard I. Moser, and May-Britt Moser. "Representation of Geometric Borders in the Developing Rat". In: *Neuron* 82.1 (Apr. 2014), pp. 71–78. DOI: 10.1016/j.neuron.2014.02.014. URL: https://doi.org/10.1016/j.neuron.2014.02.014.

143. Alexander Dityatev and Melitta Schachner. "Extracellular matrix molecules and synaptic plasticity". In: *Nature Reviews Neuroscience* 4.6 (June 2003), pp. 456–468. DOI: 10.1038/nrn1115. URL: https://doi.org/10.1038/nrn1115.

144. Difei Wang and James Fawcett. "The perineuronal net and the control of CNS plasticity". In: *Cell and Tissue Research* 349.1 (Mar. 2012), pp. 147–160. DOI: 10.1007/s00441-012-1375-y. URL: https://doi.org/10.1007/s00441-012-1375-y.

145. Eric Jonas and Konrad Paul Kording. "Could a Neuroscientist Understand a Microprocessor?" In: *PLoS Comput. Biol.* 13.1 (2017), pp. 1–24. ISSN: 15537358. DOI: 10.1371/journal.pcbi.1005268.

146. Joshua D Angrist and Jörn-Steffen Pischke. *Mostly harmless econometrics: An empiricist's companion*. Princeton university press, 2008.

147. Ian H Stevenson, James M Rebesco, Lee E Miller, and Konrad P Körding. "Inferring functional connections between neurons". In: *Current opinion in neurobiology* 18.6 (2008), pp. 582–588.

148. CJ Honey, O Sporns, Leila Cammoun, Xavier Gigandet, Jean-Philippe Thiran, Reto Meuli, and Patric Hagmann. "Predicting human resting-state functional connectivity from structural connectivity". In: *Proceedings of the National Academy of Sciences* 106.6 (2009), pp. 2035–2040.

149. Laurence Aitchison and Máté Lengyel. “With or without you: predictive coding and Bayesian inference in the brain”. In: *Current opinion in neurobiology* 46 (2017), pp. 219–227.

150. David Pfau, Eftychios A Pnevmatikakis, and Liam Paninski. “Robust learning of low-dimensional dynamics from large neural ensembles”. In: *Advances in neural information processing systems*. 2013, pp. 2391–2399.

151. Didier Pinault. “A novel single-cell staining procedure performed in vivo under electrophysiological control: morpho-functional features of juxtacellularly labeled thalamic cells and other central neurons with biocytin or Neurobiotin”. In: *Journal of neuroscience methods* 65.2 (1996), pp. 113–136.

152. Gilad M Lerman, Jonathan V Gill, Dmitry Rinberg, and Shy Shoham. “Two Photon Holographic Stimulation System for Cellular-Resolution Interrogation of Olfactory Coding”. In: *Optics and the Brain*. Optical Society of America. 2017, BrM3B–5.

153. Volodymyr Nikolenko, Kira E Poskanzer, and Rafael Yuste. “Two-photon photostimulation and imaging of neural circuits”. In: *Nature methods* 4.11 (2007), p. 943.

154. Valentina Emiliani, Adam E Cohen, Karl Deisseroth, and Michael Häusser. “All-optical interrogation of neural circuits”. In: *Journal of Neuroscience* 35.41 (2015), pp. 13917–13926.

155. Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. MIT Press, 2017.

156. R. Quian Quiroga, L. Reddy, G. Kreiman, C. Koch, and I. Fried. “Invariant visual representation by single neurons in the human brain”. In: *Nature* 435.7045 (June 2005), pp. 1102–1107. DOI: 10.1038/nature03687. URL: https://doi.org/10.1038/nature03687.

157. Tamás F. Freund and Miklós Antal. “GABA-containing neurons in the septum control inhibitory interneurons in the hippocampus”. In: *Nature* 336.6195 (Nov. 1988), pp. 170–173. DOI: 10.1038/336170a0. URL: https://doi.org/10.1038/336170a0.

158. Zaneta Navratilova, Lisa M. Giocomo, Jean-Marc Fellous, Michael E Hasselmo, and Bruce L. McNaughton. “Phase precession and variable spatial scaling in a periodic attractor map model of medial entorhinal grid cells with realistic after-spike dynamics”. In: *Hippocampus* 22.4 (Apr. 2011), pp. 772–789. DOI: 10.1002/hipo.20939. URL: https://doi.org/10.1002/hipo.20939.

159. Michael E. Hasselmo. “Neuronal rebound spiking, resonance frequency and theta cycle skipping may contribute to grid cell firing in medial entorhinal cortex”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 369.1635 (Feb. 2014), p. 20120523. DOI: 10.1098/rstb.2012.0523. URL: https://doi.org/10.1098/rstb.2012.0523.

160. Christopher F. Shay, Michele Ferrante, G. William Chapman, and Michael E. Hasselmo. “Rebound spiking in layer II medial entorhinal cortex stellate cells: Possible mechanism of grid cell function”. In: *Neurobiology of Learning and Memory* 129 (Mar. 2016), pp. 83–98. DOI: 10.1016/j.nlm.2015.09.004. URL: https://doi.org/10.1016/j.nlm.2015.09.004.

161. A. Jeewajee, C. Lever, S. Burton, J. O’Keefe, and N. Burgess. “Environmental novelty is signaled by reduction of the hippocampal theta frequency”. In: *Hippocampus* 18.4 (2008), pp. 340–348. DOI: 10.1002/hipo.20394. URL: https://doi.org/10.1002/hipo.20394.

162. Andrew P. Maurer, Shea R. VanRhoads, Gary R. Sutherland, Peter Lipa, and Bruce L. McNaughton. “Self-motion and the origin of differential spatial scaling along the septo-temporal axis of the hippocampus”. In: *Hippocampus* 15.7 (2005), pp. 841–852. DOI: 10.1002/hipo.20114. URL: https://doi.org/10.1002/hipo.20114.

163. J. Rivas, J.M. Gaztelu, and E. Garcia-Austt. “Changes in hippocampal cell discharge patterns and theta rhythm spectral properties as a function of walking velocity in the guinea pig”. In: *Experimental Brain Research* 108.1 (Feb. 1996). DOI: 10.1007/bf00242908. URL: https://doi.org/10.1007/bf00242908.

164. I.Q. Whishaw and C.H. Vanderwolf. “Hippocampal EEG and behavior: Change in amplitude and frequency of RSA (Theta rhythm) associated with spontaneous and learned movement patterns in rats and cats”. In: *Behavioral Biology* 8.4 (Apr. 1973), pp. 461–484. DOI: 10.1016/s0091-6773(73)80041-0. URL: https://doi.org/10.1016/s0091-6773(73)80041-0.

165. Eric Reifenstein, Martin Stemmler, Andreas V. M. Herz, Richard Kempter, and Susanne Schreiber. “Movement Dependence and Layer Specificity of Entorhinal Phase Precession in Two-Dimensional Environments”. In: *PLoS ONE* 9.6 (June 2014). Ed. by Maurice J. Chacron, e100638. DOI: 10.1371/journal.pone.0100638. URL: https://doi.org/10.1371/journal.pone.0100638.

166. Dmitriy Aronov, Rhino Nevers, and David W. Tank. “Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit”. In: *Nature* 543.7647 (Mar. 2017), pp. 719–722. DOI: 10.1038/nature21692. URL: https://doi.org/10.1038/nature21692.

167. Daire Rowlands, Kristian K. Lensjø, Tovy Dinh, Sujeong Yang, Melissa R. Andrews, Torkel Hafting, Marianne Fyhn, James W. Fawcett, and Gunnar Dick. “Aggrecan Directs Extracellular Matrix-Mediated Neuronal Plasticity”. In: *The Journal of Neuroscience* 38.47 (Oct. 2018), pp. 10102–10113. DOI: 10.1523/jneurosci.1122-18.2018. URL: https://doi.org/10.1523/jneurosci.1122-18.2018.

168. Rishidev Chaudhuri, Berk Gerçek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. "The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep". In: *Nature Neuroscience* 22.9 (Aug. 2019), pp. 1512–1520. DOI: 10.1038/s41593-019-0460-x. URL: https://doi.org/10.1038/s41593-019-0460-x.

169. T. P. Vogels, H. Sprekeler, F. Zenke, C. Clopath, and W. Gerstner. "Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks". In: *Science* 334.6062 (Nov. 2011), pp. 1569–1573. DOI: 10.1126/science.1211095. URL: https://doi.org/10.1126/science.1211095.

170. J. Miguel Cisneros-Franco and Étienne de Villers-Sidani. "Reactivation of critical period plasticity in adult auditory cortex through chemogenetic silencing of parvalbumin-positive interneurons". In: *Proceedings of the National Academy of Sciences* 116.52 (Dec. 2019), pp. 26329–26331. DOI: 10.1073/pnas.1913227117. URL: https://doi.org/10.1073/pnas.1913227117.

171. Gabrielle Devienne et al. "Regulation of perineuronal nets in the adult cortex by the electrical activity of parvalbumin interneurons". In: (June 2019). DOI: 10.1101/671719. URL: https://doi.org/10.1101/671719.

172. Benjamin Dunn, Daniel Wennberg, Ziwei Huang, and Yasser Roudi. "Grid cells show field-to-field variability and this explains the aperiodic response of inhibitory interneurons". In: *arXiv preprint arXiv:1701.04893* (2017).

173. Bhanu P. Tewari, Lata Chaunsali, Susan L. Campbell, Dipan C. Patel, Adam E. Goode, and Harald Sontheimer. "Perineuronal nets decrease membrane capacitance of peritumoral fast spiking interneurons in a model of epilepsy". In: *Nature Communications* 9.1 (Nov. 2018). DOI: 10.1038/s41467-018-07113-0. URL: https://doi.org/10.1038/s41467-018-07113-0.

174. Jonathan W Pillow, Jonathon Shlens, Liam Paninski, Alexander Sher, Alan M Litke, EJ Chichilnisky, and Eero P Simoncelli. "Spatio-temporal correlations and visual signalling in a complete neuronal population". In: *Nature* 454.7207 (2008), p. 995.

175. Yasser Roudi, Joanna Tyrcha, and John Hertz. "Ising model for neural data: model quality and approximate methods for extracting functional connectivity". In: *Physical Review E* 79.5 (2009), p. 051915.

176. Stefano Fusi, Earl K Miller, and Mattia Rigotti. "Why neurons mix: high dimensionality for higher cognition". In: *Current Opinion in Neurobiology* 37 (Apr. 2016), pp. 66–74. DOI: 10.1016/j.conb.2016.01.010. URL: https://doi.org/10.1016/j.conb.2016.01.010.

177. Romain Brette. "Is coding a relevant metaphor for the brain?" In: *Behavioral and Brain Sciences* 42 (July 2018). DOI: 10.1017/s0140525x19000049. URL: https://doi.org/10.1017/s0140525x19000049.

178. C. Barry, L. L. Ginzberg, J. O'Keefe, and N. Burgess. "Grid cell firing patterns signal environmental novelty by expansion". In: *Proceedings of the National Academy of Sciences* 109.43 (Oct. 2012), pp. 17687–17692. DOI: 10.1073/pnas.1209918109. URL: https://doi.org/10.1073/pnas.1209918109.

179. Julija Krupic, Marius Bauza, Stephen Burton, Caswell Barry, and John O'Keefe. "Grid cell symmetry is shaped by environmental geometry". In: *Nature* 518.7538 (Feb. 2015), pp. 232–235. DOI: 10.1038/nature14153. URL: https://doi.org/10.1038/nature14153.

180. Simon Nikolaus Weber and Henning Sprekeler. "Learning place cells, grid cells and invariances with excitatory and inhibitory plasticity". In: *eLife* 7 (Feb. 2018). DOI: 10.7554/elife.34560. URL: https://doi.org/10.7554/elife.34560.

181. Venkatakrishnan Ramaswamy. "An Algorithmic Barrier to Neural Circuit Understanding". In: (May 2019). DOI: 10.1101/639724. URL: https://doi.org/10.1101/639724.

182. Timothy P. Lillicrap and Konrad P. Kording. *What does it mean to understand a neural network?* 2019. eprint: arXiv:1907.06374.

183. S. Coombes and M. R. Owen. "Bumps, Breathers, and Waves in a Neural Network with Spike Frequency Adaptation". In: *Physical Review Letters* 94.14 (Apr. 2005). DOI: 10.1103/physrevlett.94.148102. URL: https://doi.org/10.1103/physrevlett.94.148102.

184. G Bard Ermentrout, Stefanos E Folias, and Zachary P Kilpatrick. "Spatiotemporal pattern formation in neural fields with linear adaptation". In: *Neural Fields*. Springer, 2014, pp. 119–151.

185. Hugh Pastoll, Helen L. Ramsden, and Matthew F. Nolan. "Intrinsic electrophysiological properties of entorhinal cortex stellate cells and their contribution to grid cell firing fields". In: *Frontiers in Neural Circuits* 6 (2012). DOI: 10.3389/fncir.2012.00017. URL: https://doi.org/10.3389/fncir.2012.00017.

186. I. Erchova, G. Kreck, U. Heinemann, and A. V. M. Herz. "Dynamics of rat entorhinal cortex layer II and III cells: characteristics of membrane potential resonance at rest predict oscillation properties near threshold". In: *The Journal of Physiology* 560.1 (Oct. 2004), pp. 89–110. DOI: 10.1113/jphysiol.2004.069930. URL: https://doi.org/10.1113/jphysiol.2004.069930.

187. Michele Ferrante, Christopher F. Shay, Yusuke Tsuno, G. William Chapman, and Michael E. Hasselmo. "Post-Inhibitory Rebound Spikes in Rat Medial Entorhinal Layer II/III Principal Cells: In Vivo, In Vitro, and Computational Modeling Characterization". In: *Cerebral Cortex* (Mar. 2016), bhw058. DOI: 10.1093/cercor/bhw058. URL: https://doi.org/10.1093/cercor/bhw058.

188. Romain Brette and Wulfram Gerstner. “Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity”. In: *Journal of Neurophysiology* 94.5 (Nov. 2005), pp. 3637–3642. DOI: 10.1152/jn.00686.2005. URL: https://doi.org/10.1152/jn.00686.2005.

189. Alain Destexhe, Michael Rudolph, and Denis Paré. “The high-conductance state of neocortical neurons in vivo”. In: *Nature Reviews Neuroscience* 4.9 (Sept. 2003), pp. 739–751. DOI: 10.1038/nrn1198. URL: https://doi.org/10.1038/nrn1198.

190. Jean-François Léger, Edward A. Stern, Ad Aertsen, and Detlef Heck. “Synaptic Integration in Rat Frontal Cortex Shaped by Network Activity”. In: *Journal of Neurophysiology* 93.1 (Jan. 2005), pp. 281–293. DOI: 10.1152/jn.00067.2003. URL: https://doi.org/10.1152/jn.00067.2003.

191. Arvind Kumar, Sven Schrader, Ad Aertsen, and Stefan Rotter. “The High-Conductance State of Cortical Networks”. In: *Neural Computation* 20.1 (Jan. 2008), pp. 1–43. DOI: 10.1162/neco.2008.20.1.1. URL: https://doi.org/10.1162/neco.2008.20.1.1.

192. H. Wang et al. “High-speed mapping of synaptic connectivity using photostimulation in Channelrhodopsin-2 transgenic mice”. In: *Proceedings of the National Academy of Sciences* 104.19 (May 2007), pp. 8143–8148. DOI: 10.1073/pnas.0700384104. URL: https://doi.org/10.1073/pnas.0700384104.

193. Martin Thunemann et al. “Does Light Propagate Better Along Pyramidal Apical Dendrites in Cerebral Cortex?” In: *Optics and the Brain*. Optical Society of America. 2018, JW3A–56.