The Influence of Synaptic Remodelling on Neuronal Network Activity and Function

Pia Siegele

Master of Science

Cognitive Science

School of Informatics

University of Edinburgh

2020

Abstract

Synapses, the connections between neurons, regulate the flow of information within neuronal networks in the brain. For a long time, it has thus been thought that these synaptic connections only change when new information is incorporated into a network. Surprising therefore are recent observations that synapses within neuronal networks are constantly remodelled at high rates, although the activity and function of the networks stay remarkably stable. This opens up the question of how this stability is maintained despite the volatility of the synaptic connections.

A possible answer to this question comes from recent observations that neuronal networks in the brain exhibit sloppiness, which suggests that not all network components have an equal influence on a network's overall activity and function. This led to the idea that the overall behaviour of a network is strongly influenced by a small subset of its synapses, which might stabilise the overall network, whereas most other synapses are free to change. The identity of these stabilising synapses has however not been elucidated yet. In this dissertation, we investigate the influence of different types of synapses on the activity of a computational spiking neuronal network (SNN) model. Combining and extending recent findings, we for the first time show that a very small subgroup of synapses indeed has a disproportionately large influence on the activity of such a neuronal network. This subgroup consists of highly active inhibitory synapses with large synaptic weights. We then test the effect of rewiring these synapses on the representation of information within an SNN and find that, although it affects the activity and information content of individual neurons, this does not alter the performance of a decoder designed to read out information from the network. Overall, this suggests that the small subgroup of inhibitory synapses might have a stabilising influence on the activity of a neuronal network, but their role with regard to a network's function remains to be further elucidated.

In general, computational studies such as ours provide insights into the dynamics of neuronal networks and the principles of brain function beyond what is currently directly experimentally accessible in the brain.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(*Pia Siegele*)

Acknowledgements

First and foremost, I would like to thank my supervisor Dr Matthias Hennig for his guidance, support and valuable feedback throughout this project. Thank you for the many interesting discussions and for answering all my questions. I would also like to thank Dr Gianluigi Mongillo and Professor Yonatan Loewenstein for sending me the code of their random network model in C, which helped me reproduce their model in Python. Finally, I would like to thank the *Studienstiftung des Deutschen Volkes* for providing the financial support that enabled me to do the Cognitive Science MSc programme.

Table of Contents

Abbreviations

E	excitatory
EE	excitatory-to-excitatory
EI	excitatory-to-inhibitory
LGN	lateral geniculate nucleus
I	inhibitory
IE	inhibitory-to-excitatory
II	inhibitory-to-inhibitory
LIF neuron	leaky integrate-and-fire neuron
MAE	mean absolute error
OSI	orientation selectivity index
PO	preferred orientation
SNN	spiking neuronal network

Chapter 1

Introduction

Synapses, the structures at which electrical signals are transmitted from one neuron to the other, are crucial components of the neuronal networks in our brain. They regulate information flow not only by relaying signals between neurons, but also by determining whether a signal promotes a neuron's activity (excitatory signal) or attenuates it (inhibitory signal), and how much influence a signal has. Consequently, it has traditionally been thought that the specific patterns and properties of synaptic connections determine the way information is processed and stored within these networks, and any changes in synaptic connections have been thought to reflect the incorporation of new information into the network [1–3]. Surprising, therefore, are recent observations that synapses in the brain constantly change at a high rate, without necessarily altering network function or adding/removing stored information [1, 4, 5]. Understanding how neuronal network stability is maintained despite the volatility of synaptic connections promises to give important insights into the dynamics of neuronal networks and the principles of brain function.

This dissertation mainly builds on two recent findings that offer clues about how the contradiction of a stable network with unstable synapses might be resolved. Firstly, Panas et al found that neuronal networks exhibit sloppiness [6]. Sloppiness is a concept from complex systems theory which describes systems whose overall behaviour depends only on a small subset of components, called 'stiff' parameters [7–9]. Most other components, called 'sloppy' parameters, are free to change without significantly affecting the overall system. In the context of neuronal networks, it seems that the stiff parameters correspond to a small number of highly active neurons, which have been proposed to support and maintain overall network stability [6]. Secondly, Mongillo et al observed that rewiring all inhibitory synapses within a computational neuronal

network model strongly affected overall network activity and lead to the loss of stored information, whereas rewiring excitatory synapses, which made up a much larger proportion of the network, had little effect overall [10]. The two findings above suggest the following: Firstly, not all neurons and their synapses seem to have an equal influence on the overall activity and function of a network. Secondly, a small subset of highly active neurons, and therefore also their synapses, seem to have a large influence [6]. And thirdly, the precise locations of inhibitory synaptic connections within a network seem to matter more than the precise locations of excitatory connections [10]. In this dissertation, we aimed to combine and explicitly investigate these three points, by testing the hypothesis that neuronal network stability is supported by small number of inhibitory synapses which transmit signals of highly active neurons.

To investigate our hypothesis, we systematically assess the effect of rewiring different groups of synapses on the overall activity of a spiking neuronal network (SNN) model, which is a simplified computational model of a neuronal network in the brain. We show that rewiring a small subgroup of inhibitory synapses has a disproportionately large influence on the activity of the overall network, whereas rewiring most other synapses has little effect. In line with our hypothesis, this subgroup of inhibitory synapses transmits signals of highly active neurons, but is also characterised by high synaptic weights. Due to their strong influence on the overall network, we argue that overall network stability can be maintained by keeping these synapses stable. Finally, we assess the effect of synaptic rewiring on an SNN that models how visual information is represented in the brain. We find that, although rewiring the subgroup of inhibitory synapses affects the network's activity, this has surprisingly little effect on its function, as it does not significantly affect the performance of a decoder designed to recover information from the network. Nevertheless, the small subgroup of inhibitory synapses might have other roles in the brain, beyond what we have modelled here, which might include supporting stable storage or parallel processing of information.

This dissertation is structured as follows. After providing some background on the fundamental principles of neuronal communication, Chapter 2 will review evidence of high synaptic volatility in the brain, and discuss how the concept of sloppiness might explain how neuronal networks maintain global stability despite the volatility of their synapses. Chapter 3 will describe the computational models and methods used to evaluate the effect of synaptic remodelling. Chapter 4 contains experiments and results. Finally, Chapter 5 will summarise and discuss our results, consider limitations, point out future work, and present our conclusions.

Chapter 2

Background

In a typical network in the brain, every neuron receives thousands of synaptic inputs from other neurons [11, 12]. Each incoming signal triggers a transient increase or decrease in the neuron's membrane potential, depending on whether it stems from an excitatory or inhibitory synapse respectively. If multiple incoming signals manage to increase a neuron's membrane potential above a certain threshold, one or multiple action potentials will be triggered. Action potentials are sudden spikes in electrical potential, and the stronger a signal, the higher the frequency at which spikes will be generated. Action potentials are initiated at a neuron's cell body and can travel long distances along its axon. A neuron's synapses are located at the end of the axon, and every time an action potential arrives, the synapses transmit the signal to all connected neurons. Synapses are unidirectional and weighted, which means that the signal will always flow from the presynaptic to the postsynaptic neuron, and the amount of influence the signal has on a postsynaptic neuron depends on the weight of each individual synapse. The type of synapse is determined by the identity of the presynaptic neuron; excitatory neurons send signals via excitatory synapses and inhibitory neurons via inhibitory synapses [11, 13]. Overall, a neuron thus performs computations by integrating many different input signals and distributing its outputs to a selection of other neurons. Which inputs a neuron receives, where it sends its outputs, and how each signal is weighted is determined by the synapses within the network.

Understanding the way synapses connect neurons to form the many different types of neuronal networks in the brain seems to be one of the keys to understanding how the brain carries out the vast array of computations that underlie its function [14–18]. Therefore, a lot of research in past decades has gone towards studying the connectivity structure of different networks in the brain and relating it to their activity and

functions [18–20]. Until recently, the general consensus in most studies, which has also served as inspiration for the construction of artificial neural networks, has been that once a network's function has been established, it remains relatively static unless learning or a new experience takes place [2, 3]. In this case, synaptic connections will change to enable the network to adapt and incorporate any new information, a process called synaptic plasticity [2,21,22]. However, recent research has shown that neuronal networks are much more dynamic than that, with synaptic changes occurring at a high rate even when no learning takes place [1,4,5]. These observations indicate that, given the remarkable stability of neuronal networks in the brain, not every synaptic change is likely to alter network function or add/remove stored information [4, 10, 23–25]. This highlights that our understanding of neuronal network dynamics in the brain is still incomplete.

The following section of this chapter will review evidence of high synaptic volatility in the brain. The second section will then present evidence that suggests how the brain might maintain stability despite the dynamic nature of its components.

2.1 Synaptic Volatility in the Brain

The fact that synaptic changes enable networks to adapt and learn is a well-accepted phenomenon and has been extensively studied [2, 21, 22, 26–28]. Studies investigating the effect of sensory deprivation, such as the long-term closing of one eye, have shown that cortical networks can significantly reorganise their synaptic connections to adapt to lasting changes in network inputs [28–34]. Furthermore, learning seems to induce changes to synaptic weights [35–37] and cause the formation of new synapses [38–41] in relevant brain regions. Moreover, experimental weakening or deletion of selected synapses can erase specific memories or learned skills [36, 42, 43], and inhibition of synaptic changes has been shown to prevent learning [37, 44].

Many of the studies investigating the link between synaptic changes and learning conducted measurements only at limited time points and often averaged over multiple laboratory animals, which likely contributed to the perception that synaptic changes predominantly occur due to, and contribute to, network adaptation and learning [3]. However, enabled by the development of modern imaging techniques, more and more studies have started to track synaptic changes in living animals over longer time periods. These studies revealed that, even under so-called 'baseline' conditions in which animals do not experience or learn anything new, synapses in the brain seem to undergo

significant remodelling. Estimates of synaptic volatility vary depending on brain area, age of animal (all studies below were done in adult mice), and imaging technique used. Regarding synapse formation and disappearance, older studies seem to report lower turnover rates ranging from 4% turnover within one month in the visual cortex [45] over 3-5% within two weeks in the barrel cortex [46] to 27% within one month in the somatosensory cortex [47]. Other, mainly newer, studies have reported much higher rates of 20% turnover within one day in the barrel cortex [34], 31% turnover within four days in the auditory cortex [5], and 100% turnover within 3-4 weeks in the hippocampus [4]. These studies imaged synaptic changes at much smaller intervals (1-3 days) than the older studies with lower estimates (2 weeks). The difference in estimates may thus partly be attributable to the differences in imaging intervals, which is supported by one study's observation that increasing imaging intervals decreased their estimate of synaptic volatility [46]. This could suggest that a proportion of synapses in the brain appear and disappear at high rates, which would likely not be detectable at large intervals, while other synapses might be more stable. Even synapses that persist for many weeks have however been reported to fluctuate in weight under baseline conditions [48]. Interestingly, despite the volatility of individual synapses, the overall synapse numbers, as well as the distribution of synaptic weights seem to stay constant within neuronal networks [4, 34, 48, 49], a principle which we will adhere to in the computational models used in this dissertation.

One possible explanation of high synaptic volatility in neuronal networks is that, even under baseline conditions, the brain is extremely sensitive to small fluctuations in environment that are hard to control for in experiments [3]. This would mean that all observed synaptic changes are essentially due to plasticity mechanisms. This is unlikely however, as plasticity mechanisms are activity-dependent, and studies which blocked neuronal activity in slice cultures [49–51] and *in vivo* [52] still observed substantial synaptic volatility, albeit with altered dynamics. Whether the spontaneous fluctuations of synaptic connections within neural networks serve a specific function is still unclear. One factor might simply be biological constraints, as synapses themselves are complex structures composed of many dynamic elements, such as receptor molecules and adhesion proteins, that are constantly turned over [53–55].

One should note that a major limitation of the studies above is that they almost exclusively focus on excitatory connections. Excitatory synapses are much easier to monitor because most are situated on dendritic spines, which are small protrusions on the surface of neurons [56,57]. Inhibitory synapses are much harder to detect and track

as they can be located on any part of the postsynaptic neuron [58, 59]. It is therefore unclear whether inhibitory synapses display similar dynamics to excitatory synapses, although one study reported that inhibitory synapses frequently seem to disappear and reappear again at the same sites on postsynaptic neurons [59].

Overall, the evidence above indicates that neuronal networks in the brain have to constantly deal with both spontaneous and plasticity-induced changes to their synaptic connections, while keeping their overall network stable. Even learning-induced changes that add information to the network should not destabilise overall network function, disturb other ongoing processes, or lead to the loss of previously stored information. How the brain manages to maintain overall network stability in the face of synaptic volatility is still poorly understood. One candidate explanation, which we will further explore in this dissertation, is sloppiness. The next section will explain what sloppiness is and review evidence of sloppiness in neuronal networks.

2.2 Sloppiness in Neuronal Networks

Where does the concept of sloppiness stem from? Mathematical and computational models are invaluable tools that help researchers characterise, study and understand complex systems [60–62]. These models often contain large amounts of parameters that need to be estimated from experimental data. This is often difficult because the models tend to be loosely constrained, which means that many different parameter combinations can fit the data well [62, 63]. The key observation that led to the idea of sloppiness is that, in many cases, each of these possible parameter combinations seem to yield the same accurate and useful predictions [7, 61, 62, 64–67]. This indicates that the precise value of these 'sloppy' parameters is not important for the overall system's behaviour [7]. However, although most parameters can vary greatly, the predictions of these models seem to depend on a small subset of tightly constrained, 'stiff' parameter combinations, and changes in these will strongly affect a model's behaviour (Figure 2.1). This phenomenon, called sloppiness, has been characterised in many models of complex systems, such as models of biochemical signalling networks [62, 64], interatomic potentials [65], radioactive decay [66], or insect flight [67]. The ubiquity of sloppiness in complex systems models indicates that the underlying systems also consist of components with unequal influence on the overall system, with a small subset of stiff components tightly controlling system behaviour, while the other, sloppy, components are free to vary [66].

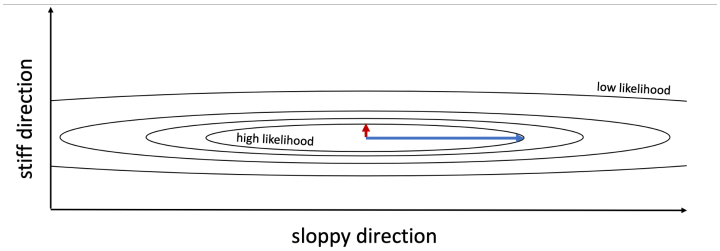

Figure 2.1: Schematic drawing of the surface contours of the likelihood function of a sloppy model. Along sloppy directions, the likelihood is essentially flat and parameter changes along these directions (blue arrow) do not significantly affect the model. Parameter changes along stiff directions (red arrow), where the likelihood is strongly curved, lead to large changes in model behaviour. In many models, the stiff/sloppy directions consist of combinations of multiple parameters. In models of neuronal networks, however, they interestingly correspond to individual network parameters [6]. Figure adapted from [68]

Sloppiness in neuronal networks was first observed by Panas et al in 2015 [6]. The authors identified stiff and sloppy components in pairwise maximum entropy models fitted to the recordings of the activity of neuronal networks in culture and in the visual cortex of macaques' brains. Interestingly, the stiff dimensions of the models were consistently associated with a small number of highly active neurons within the networks, which exhibited stable activity over time. The neurons associated with the sloppy directions in contrast strongly fluctuated in activity. Importantly, the overall activity of the networks remained stable over time, indicating that network stability is maintained by the small number of highly active stiff neurons while the other, sloppy neurons are free to change [6]. This is supported by earlier observations that neuronal networks in the hippocampus of rats contain a large number of neurons with highly variable activity, whereas a small proportion of highly active neurons, in line with the overall population activity, seem to remain stable over time [69]. Similarly, Okun et al [70] report that the visual cortex of mice and monkeys contains "chorister" neurons whose activity is highly correlated with the activity of the overall neuronal population, whereas the activity of other "soloist" neurons is only weakly correlated. This suggests that the "chorister" neurons correspond to stiff neurons that strongly influence overall population activity, whereas sloppy "soloist" neurons are free to change.

Neuronal networks do not only seem to be sloppy in their electrical activity but also in the way they represent information. In the hippocampus, neurons called place cells represent information about the location a person/animal is currently in [71]. When

monitoring the activity of place cells in the hippocampus of mice while these were running along a linear track, Ziv et al [72] noticed that 15-25% of place cells consistently represented the animals' locations over several weeks, whereas other cells only represented spatial information within one or few recording sessions. Nevertheless, a decoder trained on data recorded on one day could still reliably predict an animal's location from data recorded on a different day [72]. This indicates that a small number of functionally stable neurons seem to suffice for reliable encoding and representation of information within networks over time.

As explained in section 2.1, neuronal networks have to constantly deal with spontaneous and plasticity-induced changes to their components, which should not destabilise overall network function, disturb other ongoing processes or lead to the loss of previously stored information. Sloppiness seems to provide a candidate explanation for how this is possible. During sleep, place cells in the hippocampus replay spatial information learned during the day, which is thought to support the consolidation of spatial memories [73, 74]. When recording the activity of place cells during sleep, before and after rats learned to navigate a new environment, Grosmark and Buzsaki [75] noticed that the activity of some slow-firing neurons was strongly altered by the new experience, which suggests that plasticity took place in their synaptic connections. In contrast, the activity of other, fast-firing, neurons remained unaltered by the new experience [75]. This suggests that these highly active neurons form a stable subnetwork that maintains overall network stability while plasticity takes place in other neurons, thus allowing the network to process and learn new information without compromising its overall function [75].

A similar principle has also been observed with regards to the processing of sensory information in the brain. Ponce-Alvarez et al [76] recorded the activity of neurons in the auditory cortex of mice, both under baseline conditions and in response to sound. Similar to Panas et al [6] above, the authors then identified stiff and sloppy directions in pairwise maximum entropy models fitted to the recordings [76]. In line with previous observations, the stiff directions seemed to correspond to a minority of highly active neurons with low fluctuations in activity, whereas the activity of sloppy neurons varied greatly. Importantly, the variability in the activity of sloppy neurons not only consisted of spontaneous fluctuations but also represented the brain network's response to sound [76]. Furthermore, the activity of stiff neurons was not simply static, but rather exhibited small periodic changes which correlated with changes in cortical state of the network. Cortical states are different modes of overall network activity in the

brain, with different frequencies of electrical activity and different levels of synchronisation across neurons [77]. Interestingly, a network's cortical state seems to influence how information is processed within the network [78]. Taken together, this indicates that stiff neurons control overall network activity, which allows sensory information to be processed by sloppy neurons without perturbing the whole network. Furthermore, by controlling changes in cortical state, stiff neurons might provide more than just generic stability, but might rather supply a framework that controls how information is processed within the network [76].

Overall, the evidence above suggests that within a neuronal network, a small subset of neurons, and thus also their synapses, have a disproportionately high influence on the overall network activity. The stability of these neurons and their synaptic connections seems to ensure overall network stability and allow other components to change spontaneously and to process information. Not much is currently known about the identity of these stabilising stiff neurons and the properties of their synapses, except that stiff neurons tend to be highly active, i.e. fire action potentials at high rates [6, 69, 75, 76]. This is because the types of experimental recordings used in the studies above, which allow researchers to monitor the activity of a large number of neurons in a network at the same time, do not permit any detailed characterisation of individual neurons [79, 80].

To gain further insights into neuronal network dynamics and the properties and roles of individual neurons beyond what is currently directly measurable in the brain, computational models of neuronal networks can be immensely useful tools. In the case of the identity of the stabilising stiff neurons and synapses, a computational study by Mongillo et al [10] provides clues that stiff neurons might be inhibitory neurons. The authors built a spiking neuronal network model with parameters determined by experimental measurements in the mouse brain [10, 81], and found that complete rewiring of inhibitory synapses in the model strongly altered overall network activity and led to the loss of stored memories. In contrast, completely rewiring all excitatory synapses had little effect overall, although these constituted the majority of synapses within the network [10]. As in the living brain, the inhibitory neurons within the network model tended to have higher firing rates than the excitatory neurons [10]. Mongillo et al's observation thus agrees with the fact that stiff neurons seem to be highly active. In fact, Mongillo et al argue that the higher firing rates of inhibitory neurons are the determining feature that provides inhibitory synapses with a high influence on the overall network, but they do not explicitly test this in their model. Furthermore,

although inhibitory synapses are less frequent than excitatory ones in the model, they still constitute a fairly large proportion (25%) of synapses. This opens up the question of whether maybe not all, but only a fraction, of these inhibitory synapses are stiff, perhaps the ones that transmit signals of especially active inhibitory neurons.

In this dissertation, we expand upon Mongillo et al's [10] research and explicitly test whether overall network activity is determined by a small number of inhibitory synapses that transmit signals of highly active neurons. To do so, we reproduce Mongillo et al's [10] spiking neuronal network model and systematically assess the effect of rewiring selected synapses on the network's overall activity. We find that, although synapses of inhibitory neurons with high firing rates seem to have a large influence on the overall network, firing rate is not the only determining factor. The synapses with the highest influence seem to be inhibitory synapses characterised by a high presynaptic firing rate and high synaptic weight. Finally, we test what effects the impact of synaptic rewiring has on the way information is represented in a neuronal network. The following chapter will provide details of our computational neuronal network models, and will explain the methods we used to evaluate the effect of synaptic remodelling on the networks' activity and function.

Chapter 3

Methods

3.1 The Spiking Neuronal Network Models

In order to investigate the effect of synaptic remodelling on the activity and function of neuronal networks, we used two computational spiking neuronal network (SNN) models, which are networks composed of simulated spiking (action-potential-producing) neurons. The first model (random network model), a randomly connected SNN previously used by Mongillo et al [10], allowed us to test the effect of synaptic remodelling on the activity of a neuronal network. The second model (ring model), a modified version of the randomly connected network, models the phenomenon of orientation selectivity in the visual cortex (explained in detail in section 3.1.2 below), and allowed us to test the effect of synaptic remodelling on the representation of information within a network.

Both SNNs are composed of N_E excitatory (E) and N_I inhibitory (I) leaky integrate-and-fire (LIF) neurons. LIF neurons are simplified models of biological neurons, which model the state of a neuron in terms of its membrane potential at a given time t [82] (Figure 3.1A). In our model, the membrane potential $v_a^i(t)$ of a neuron i in population a (= E;I) evolves according to

$$\dot{v}_a^i(t) = -\frac{v_a^i(t)}{\tau_m} + h_a^i(t) + \frac{H_a^{ext,i}}{\tau_m} \tag{3.1}$$

where $i = 1, ..., N_a$ [10]. Equation 3.1 shows that the membrane potential of a neuron is influenced both by recurrent synaptic inputs $h_a^i(t)$ from other neurons within the network, as well as by a constant external input $H_a^{ext,i}$ that represents input from neurons in other brain regions [10] (Figure 3.1A). The time constant τ_m ensures that, without any new inputs, the membrane potential decays towards the neuron's resting potential

ν_{rest} (= 0mV in this case).

Any time a neuron's membrane potential reaches a fixed threshold θ, an action potential (also called spike) is generated and the neuron's membrane potential is reset to its reset potential ν_{reset}:

$$\nu_a^i(t^+) = \nu_{reset} \quad \text{if} \quad \nu_a^i(t) \geq \theta \tag{3.2}$$

The LIF model does not explicitly model the membrane potential dynamics underlying each action potential, but rather just records the times at which action potentials are triggered. This is because, in biological neurons, the membrane voltage-changes during an action potential are highly stereotyped, and information in the brain is thought to be transmitted by the timing of action potentials and the rate at which they are generated [82]. Therefore, it is considered sufficient for many purposes to only model sub-threshold membrane dynamics, as this saves computational resources and enables much more efficient simulations of large networks [82].

Every time an action potential has been generated, the LIF neuron remains at its reset potential ν_{reset} for the duration of a refractory period τ_{arp} during which no new spikes can be generated. After this, equation 3.1 resumes from ν_{reset} [10]. At the beginning of each simulation, the starting value $\nu_a^i(0)$ of each neuron was initialised randomly.

The recurrent synaptic input $h_a^i(t)$ that a given neuron i in population a receives from other neurons in the network is given by

$$h_a^i(t) = \sum_{j=1}^{N_E} c_{Ea}^{ji} w_{Ea}^{ji} \sum_k \delta(t - t_{E,k}^j) - \sum_{j=1}^{N_I} c_{Ia}^{ji} w_{Ia}^{ji} \sum_k \delta(t - t_{I,k}^j) \tag{3.3}$$

where $c_{ba}^{ji} = 1$ if a synaptic connection exists from neuron j in population b to neuron i in population a, otherwise $c_{ba}^{ji} = 0$ [10]. The model does not allow connections of a neuron with itself. w_{ba}^{ji} denotes the weight of each synapse. The first part of the equation corresponds to all excitatory inputs received by neuron i, the second part corresponds to the inhibitory inputs. The sums over j are sums over all presynaptic neurons in the corresponding population, and the sums over k are over all action potential emission times $t_{E,k}^j$ of neuron j [10]. The δ function is the Kronecker delta, which means that $\delta(t - t_{b,k}^j) = 1$ if neuron j generated an action potential at time t, and $\delta(t - t_{b,k}^j) = 0$ otherwise (Figure 3.1A). The synaptic inputs are transmitted from a presynaptic to a postsynaptic neuron with a small delay τ_{delay}.

Our two models, the random network model and the ring model, are both composed of LIF neurons as defined above, but differ slightly in their external inputs and

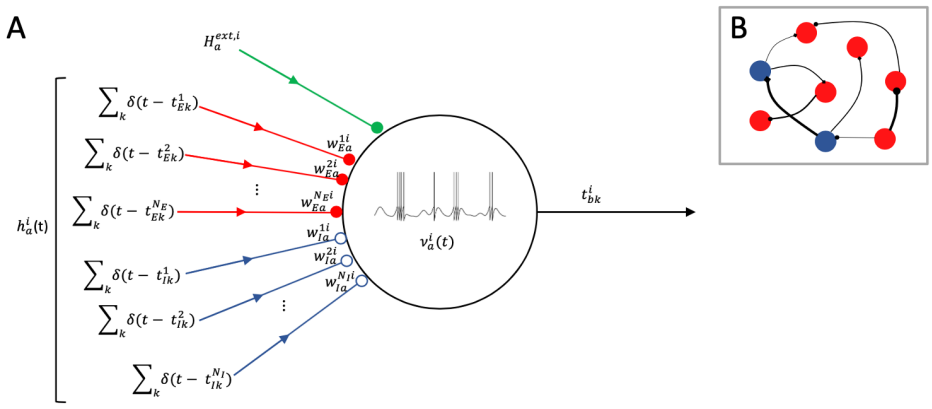

Figure 3.1: A: Schematic diagram of a leaky integrate-and-fire (LIF) neuron. LIF neurons model the state of a neuron i in terms of its membrane potential $v_a^i(t)$. The LIF neuron receives external inputs (green) as well as recurrent excitatory (red) and inhibitory (blue) synaptic inputs from other neurons in the network, and outputs a train of action potentials. Variable names are defined in the main text. Figure adapted from [84]. B: Schematic drawing of the random network model composed of excitatory (red) and inhibitory (blue) LIF neurons. Adapted from [10].

synaptic connections, which will be explained below. All parameter values for both networks are reported in Tables A.1-3 in the appendix. Both models were implemented in Python 3.7 using the Brian2 simulator package (version 2.3) [83].

3.1.1 The Random Network Model

The random network model consists of N_E = 32,000 excitatory and N_I = 8,000 inhibitory neurons [10] (Figure 3.1B), which reflects the fact that networks in the brain contain more excitatory than inhibitory neurons [85, 86]. Each neuron i within a population a (= E;I) receives the same constant external input, $H_a^{ext,i} = H_a^{ext}$. The synaptic connections between neurons were defined as follows [10]. The c_{ba}^{ji} values were randomly set to 1 with probability C_{ba}. The C_{ba} values for the different synaptic populations range from 0.2 to 0.4 (Table A.1), which means that the network is relatively sparsely connected. The synaptic weights w_{ba}^{ji} were independently drawn from a lognormal distribution, which corresponds to the way synaptic weights seem to be distributed in the brain [87]. The distribution is parametrised by

$$p(x) = \frac{1}{\sigma_{ba} x \sqrt{2\pi}} \exp\left(\frac{-(ln(x) - \mu_{ba})^2}{2\sigma_{ba}^2}\right) \tag{3.4}$$

The values for μ_{ba} and σ_{ba} were chosen such that the mean and variance of the resulting weight distribution align with the distribution of synaptic weights measured in the auditory and barrel cortices of mice [10, 22, 81]. Because the mean synaptic input that a given neuron receives depends on the number of neurons N_a within the network, as well as the connection probabilities C_{ba} [88], all synaptic weights were scaled by

$$\tilde{w}_{ba}^{ji} = \frac{w_{ba}^{ji}}{\sqrt{C_{ba}N_a}} \tag{3.5}$$

3.1.2 The Ring Model

Area V1 of the visual cortex, a brain area that processes visual information from our eyes, contains excitatory neurons that selectively respond if we view lines or bars of light oriented at specific angles [90]. Each of these neurons has a preferred orientation (PO), and will respond most strongly if a line at this orientation is shown, whereas lines perpendicular to the neuron's PO will suppress the neuron's activity (Figure 3.2A) [90]. Together, a population of these orientation selective cells (also called 'simple cells') will consist of neurons with different POs such that they cover all possible angles.

Since its discovery, the phenomenon of orientation selectivity has been extensively studied as a canonical example of how neuronal circuits carry out computations in the brain [91–94], and many computational models of orientation selectivity have been

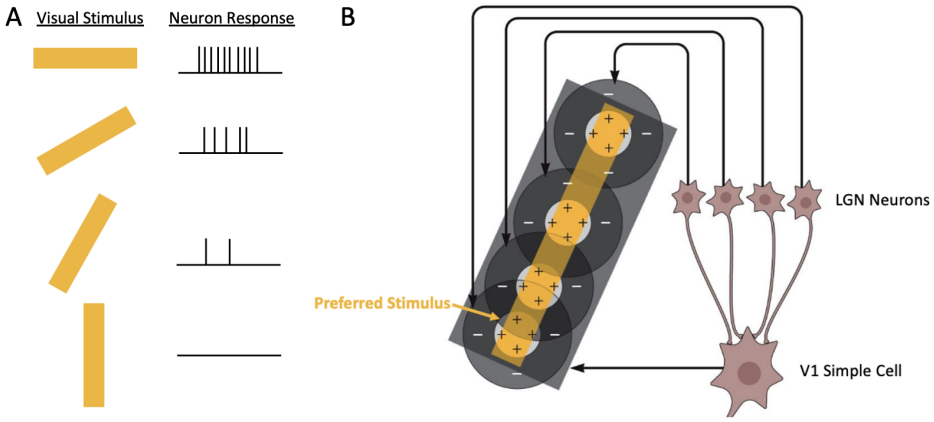

Figure 3.2: Orientation selective neurons (simple cells) in the visual cortex. A: Bars of light (yellow) at different orientations elicit very different responses in a given simple cell. This neuron preferentially responds to horizontal bars. A stronger response means that the neuron fires more action potentials (vertical black bars). B: The response properties of a simple cell are thought to arise at least partly from the way its inputs from neurons in the LGN are organised. Adapted from [11, 89].

proposed [94–99]. It is generally accepted that the phenomenon arises at least partly from the way the neurons' excitatory inputs from the lateral geniculate nucleus (LGN; a brain area) are organised [94,95] (Figure 3.2B). Neurons in the LGN each respond to points of light at a specific location in our visual field. In the visual cortex, a given simple cell will receive inputs from multiple excitatory LGN cells that respond to points located next to each other such that they form a line (Figure 3.2B). The input to a simple cell will now be strongest if a bar of light falls exactly on this line such that all LGN cells that send inputs to the simple cell are excited [95]. The further the bar will be off the line, the weaker is the input which the simple cell will receive from the LGN.

However, orientation selectivity is not only thought to arise due to feedforward inputs from the LGN. Lateral connections between neurons within the V1 (excitatory simple cells & inhibitory neurons) are thought to ensure contrast-invariance, which means that simple cells respond in the same way to stimuli with low and high contrast, as well as amplification and sharpening of the neurons' responses [94, 97, 99]. It is still debated to which extent these lateral connections within the networks of different species (mice, cats, monkeys, humans, ...) are random or feature-specific. In the latter case, the connection probability between two neurons would depend on the difference in their POs [97,99].

In our study, we wanted to test the effect of synaptic rewiring on the orientation selectivity of neurons in the visual cortex. For this purpose, we decided to use a network model that is as closely related to our random model as possible, to maximise the comparability of our results. Our model is therefore composed of the same LIF neurons with the same connection probabilities and biologically plausible distributions of synaptic weights. Because the population of orientation selective neurons responding to a given area of the visual field is likely a lot smaller than 40,000 neurons, the size of our random model, we reduced the number of neurons in our ring model by a factor of 10, thus N_E = 3,200 and N_I = 800. The model again receives constant external excitatory inputs H_a^{ext} that are the same for every neuron i within a population a (= E;I). In addition, the model receives excitatory inputs from the LGN. Inhibitory neurons each receive a combined input of r_I^{LGN} = 80Hz from the LGN. Excitatory neurons receive an input of on average $\bar{r}_E^{LGN}$ = 100Hz, but the input to each individual excitatory neuron i depends on the neuron's PO and the orientation of the stimulus α [99], according to

$$r_E^{LGN,i} = \bar{r}_E^{LGN}\left[1 + \eta_{LGN}\cos(2(\alpha - PO_i))\right] \tag{3.6}$$

where $i = 1, ..., N_E$. η_{LGN} controls the amount of input tuning, i.e. how strongly the in-

put changes for different $(s - PO_i)$. The spike times of the LGN neurons are modelled as Poisson point processes, where $r_E^{LGN,i}$ is the driving rate and Poisson spike trains are generated with the corresponding rate at each point in time. This represents a good approximation of stochastic neuronal firing in the brain [82]. Excitatory and inhibitory neurons all receive their LGN inputs via excitatory synapses of strength w_{LGN}. The POs of the excitatory neurons were defined such that the excitatory population covered the range of angles between 0 and 180 degrees at evenly spaced intervals.

We tested two versions of the ring model, one with the same random lateral connectivity as in the random network model (Section 3.1.1 & Figure 3.1B), and one with feature-specific excitatory-to-excitatory (EE) connections. We implemented the feature-specific connections by varying the probability C_{EE}^{ij} of a given EE connection depending on the difference in PO between two excitatory neurons i and j. This was defined such that the average EE connection probability C_{EE} remained the same as in the randomly connected networks:

$$C_{EE}^{ij} = C_{EE}[1 + \eta_{FS} \cos(2(PO_i - PO_j))] \tag{3.7}$$

where $i, j = 1, ..., N_E$; $i \neq j$. η_{FS} represents the amount of tuning of the feature-specific connections.

For each of the two versions of the ring model, the parameters η_{LGN} and η_{FS} were defined such that the response of the networks to a given stimulus was approximately equal. The first version of the model received strongly tuned input from the LGN ($\eta_{LGN} = 1$), and the EE connections were completely random ($\eta_{FS} = 0$) (Figure 3.3A). To match the response properties of the first version, the second version received only a weakly tuned signal from the LGN ($\eta_{LGN} = 0.2$), which was then sharpened by feature-sepcific EE connections within the network ($\eta_{FS} = 0.5$) (Figure 3.3B).

3.2 Synaptic Rewiring

In order to assess the influence of different types of synapses on the activity/function of our networks, we systematically rewired different groups of synapses and evaluated the effect this had on the networks. This section will explain how we rewired the different groups of synapses, while the next section will describe the methods we used to measure the effect of synaptic rewiring.

When rewiring synapses, we either completely rewired a whole synaptic population or a subpart of a synaptic population. The four synaptic populations within our model

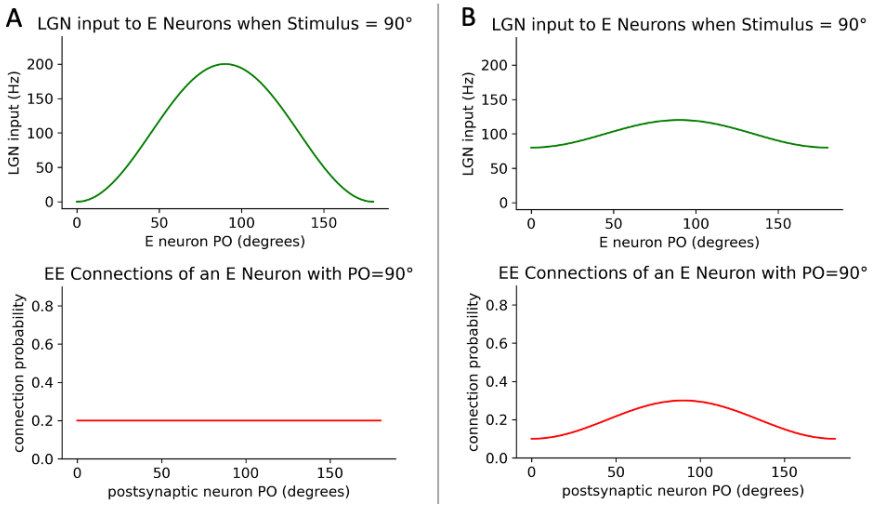

Figure 3.3: Tuning of LGN inputs and EE connectivity in the two versions of the ring model. A: The first model receives strongly tuned inputs from the LGN, and EE connections are uniformly random. B: The second model receives weakly tuned LGN inputs and has feature-specific EE connections.

are excitatory-to-excitatory (EE), excitatory-to-inhibitory (EI), inhibitory-to-excitatory (IE) and inhibitory-to-inhibitory (II) synapses. When completely rewiring one of these populations, we simply removed all existing synapses of that population within the network, and generated them anew by the random process described in Section 3.1.1. In some experiments, we only rewired a subpopulation of synapses, selected by some criterion (presynaptic firing rate, synaptic weight, synaptic impacts). In this case, the presynaptic neuron of each of these synapses was kept fixed, while we randomly chose a new postsynaptic neuron for each synapse.

3.3 Evaluation

3.3.1 Evaluating Activity Changes

To measure the effect of synaptic rewiring on the activity of a neuronal network, we compared the firing rate vectors of our models before and after rewiring. The firing rate vectors are N_a-dimensional, with each entry corresponding to the firing rate of a given neuron within the population a (=E;I).

To obtain the firing rate vectors for the random network model, the simulation was briefly run until $t = 1$s to ensure that the network reached a stable firing regime. Then, the network's state was stored, and the simulation was continued until $t = 11$s. The

firing rate vectors for the excitatory and inhibitory neuron populations of the network under default conditions were computed by calculating the firing rate (spikes per second) of each neuron within the time interval $(1, 11)$s. The network was then reset to the previously stored state at t = 1s, synapses were rewired, the simulation was run again until t = 11s, and the firing rate vectors for the network under rewired conditions were computed as before. The simulation length was longer than in Mongillo et al's [10] experiments (3s), but was chosen because it allowed a more reliable estimation of individual neurons' firing rates and thus produced more robust results.

To obtain the firing rate vectors for the ring model, the network was initialised, the network state stored, and a simulation was run. The simulation consisted of showing the network a sequence of 16 stimuli, representing orientations between 0 and 180 degrees, for 3.5s each. Firing rate vectors were computed separately for each stimulus, for the interval (0.5,3.5)s after stimulus onset. The first 0.5s were discarded to allow the network's response to transition from one stimulus to the next. The firing rate vectors after rewiring were again obtained by resetting the network, rewiring the synapses and re-running the simulation. Because the LGN inputs to the ring model are generated by a Poisson process, they present a source of randomness in the simulations. We thus ensured that exactly the same LGN input was fed into the network for default and rewired simulations, which allowed us to rule out any firing rate differences caused by random differences in LGN inputs.

To quantify the difference in firing rates due to rewiring, we computed the cosine similarity between the corresponding firing rate vectors. We opted to use cosine similarity, rather than the Pearson's correlation coefficient used by Mongillo et al [10], because the measure considers firing rate changes of neurons relative to the mean of the population, rather than absolute changes [100], and is thus more meaningful when comparing the effect of synaptic rewiring on populations with different means. (This is valid in our case because rewiring neither changes the mean nor the distribution of firing rates.) For each of our experiments, we verified that the two measures revealed the same trends. A more detailed justification, as well as an example of the results of experiment 4.1.1 with the correlation coefficient can be found in Appendix B.1.

3.3.2 Evaluating Ring Model Function

To evaluate the effect of synaptic rewiring on the function of the ring model, we firstly investigated how rewiring affected the representation of different stimuli within the

network. This was done by comparing the orientation selectivity indices (OSIs) of the network's excitatory neurons before and after rewiring. We then tested whether any shifts in the representation of the stimuli affected the performance of a decoder designed to recover stimulus information from the network's activity.

3.3.2.1 Comparing Orientation Selectivity Indices

A neuron's OSI measures how selectively the neuron responds to stimuli oriented at angles close to its PO. It takes values between 0 (not selective) and 1 (very selective). We computed each neuron's OSI as

$$OSI = \frac{\sqrt{(\sum_k R(\alpha_k)\sin(2\alpha_k))^2 + (\sum_k R(\alpha_k)\cos(2\alpha_k))^2}}{\sum_k R(\alpha_k)} \tag{3.8}$$

where $R(\alpha_k)$ signifies the neuron's firing rate in response to stimulus k [101]. $R(\alpha_k)$ was extracted from the network's firing rate vectors, computed as explained above. To quantify any shift in OSIs, we calculated Pearson's correlation coefficient of the network's OSIs before and after rewiring.

3.3.2.2 Decoder

To recover stimulus information from the network's activity, we designed a simple linear decoder. This type of decoder was chosen because it can be biologically interpreted as a readout neuron that integrates inputs from a subset of the population by computing a weighted sum [102]. The decoder has the form

$$y(\alpha) = M^T r(\alpha) \tag{3.9}$$

where $r(\alpha)$ is the normalised excitatory firing rate vector of the network's response to stimulus α. M is a $N_E \times 2$-dimensional weight matrix, where each row corresponds to a two-dimensional vector oriented at the corresponding neuron's PO. $y(\alpha)$ is thus also a two-dimensional vector, and the orientation of this vector gives the decoder's prediction of the stimulus orientation. The decoder may also read out information from only a subset of the excitatory population, in this case only the corresponding subsets of the firing rate vector and the weight matrix M are used.

Chapter 4

Experiments and Results

4.1 The Random Network Model

To test the effect of synaptic rewiring on the overall activity of a neuronal network, we successfully reproduced the random network model from Mongillo et al [10]. As in Mongillo et al [10], our network has lognormal firing rate distributions (Figure 4.1A) that are comparable to the firing rate distributions measured in the brain [10, 81, 87]. Like in biological neuronal networks [81], the mean and median firing rate of the

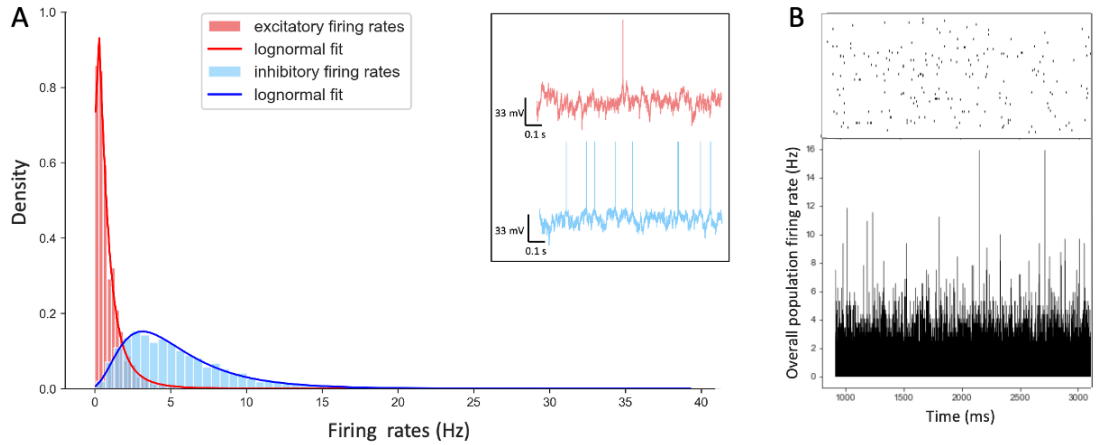

Figure 4.1: Neuronal activity of the random network model. A: Firing rate distributions of the excitatory (red) and inhibitory (blue) neuronal populations. The inset shows membrane potential traces of a randomly chosen excitatory and inhibitory neuron. B: Overall excitatory population activity within the network. The top panel shows the spike times for 100 randomly selected excitatory neurons. The dynamics of the inhibitory population are similar, just with generally higher firing rates.

network's inhibitory neurons (mean: 5.45 Hz; median: 4.6 Hz) is higher than that of the excitatory neurons (0.90 Hz; 0.60 Hz). The network operates in an asynchronous irregular regime (Figure 4.1B), in which excitation and inhibition are balanced and the average input to the neurons is slightly below threshold [10, 103]. The neurons' activity is thus driven by random fluctuations in their input [103]. In the subsequent experiments, we systematically rewired groups of synapses within the network, and evaluated the effect this had on the network's overall activity.

4.1.1 Rewiring Inhibitory Synapses Affects Network Activity

As a starting point, we repeated one of Mongillo et al's [10] experiments and evaluated the effect of rewiring entire synaptic populations (EE, EI, IE, or II) by comparing the network's firing rate vectors before and after rewiring. In line with Mongillo et al's [10] observations, we found that rewiring excitatory synapses (EE and EI) had little effect on the overall activity of the network, although these constitute the majority of the network's synapses (Figure 4.2). In contrast, rewiring inhibitory synapses, more specifically IE synapses, had a large effect on the activity of the excitatory population of the network (Figure 4.2). Rewiring II synapses had a modest effect on both the inhibitory and excitatory firing rate vectors. The means and distributions of firing rates (Kolmogorov-Smirnov; $p > 0.2$) were not affected by any rewiring, which is in line with experimental observations in biological networks [6].

In general, any changes to the overall activity of a network's excitatory neurons is thought to have a larger effect on the population readout than changes to the activity of

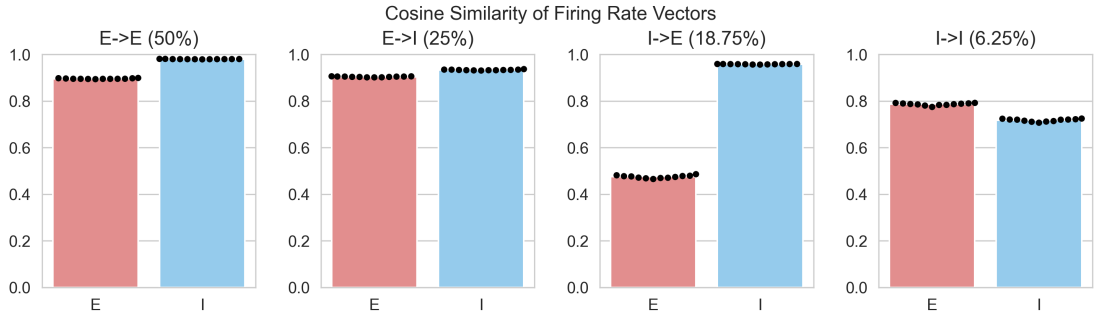

Figure 4.2: Cosine similarities of the excitatory (red) and inhibitory (blue) firing rate vectors before and after rewiring different synaptic populations. The numbers in brackets indicate the proportion of synapses of the given population within the network. Data from n=12 independent network simulations.

inhibitory neurons. This is because excitatory neurons form long-range connections to other brain regions, whereas inhibitory neurons are thought to only connect locally to neurons within a network [104]. Thus, other brain regions are thought to read out the information of a network from the activity of its excitatory neurons. In the experiments below, we therefore mainly focus on the strong influence of IE synapses on the overall activity of the network's excitatory neurons, but also examine the more moderate effect of II synapses on the overall network.

4.1.2 Inhibitory Synapses of Highly Active Presynaptic Neurons Have a Strong Influence

Because the stiff parameters in neuronal networks seem to be associated with highly active neurons [6, 69, 75, 76], we wanted to test whether the synapses that transmit signals of highly active inhibitory neurons have a disproportionately large influence on the activity of the network. We thus tested the effect of rewiring different proportions of inhibitory synapses, chosen such that these proportions always constituted the synapses of the inhibitory neurons with the highest firing rates (Figure 4.3). The results indicate that inhibitory synapses of highly active neurons indeed seem to have a larger influence on the excitatory population activity than the synapses of inhibitory neurons with lower rates (Figure 4.3).

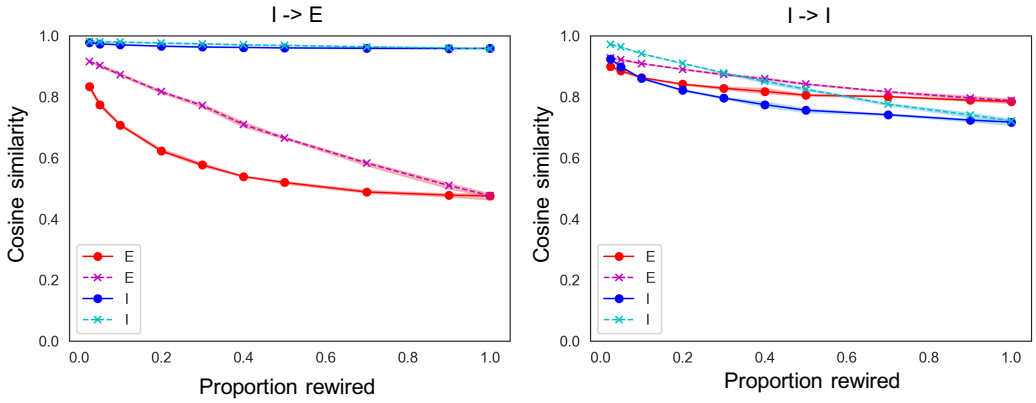

Figure 4.3: Effect of rewiring inhibitory synapses by presynaptic firing rate (solid lines) compared to rewiring the same proportion of randomly selected synapses (dashed lines). Effect on the excitatory population activity is indicated in red, effect on inhibitory activity is in blue. Due to the low variability of results between independent simulations in the previous experiment (Figure 4.2), the sample size was reduced to n=3 in this experiment. Markers: mean; shading: range.

4.1.3 Firing Rate is Not the Whole Story

In their paper, Mongillo et al [10] propose that the high influence of inhibitory synapses on the overall network activity arises due to the generally higher firing rate of inhibitory neurons in the model [10]. Our observations above seem to support this claim. However, it is not clear whether the high influence of inhibitory synapses stems from their high activity alone. To investigate this, we changed the ratio of excitatory to inhibitory neuron numbers in our model, which changed the difference between the model's mean excitatory and inhibitory firing rates (Table 4.1, Figure 4.4A). We then re-evaluated the effect of completely rewiring one of the synaptic populations (EE, EI, IE, II) on the model's firing rate vectors (as in Section 4.1.1).

If the difference between the mean population firing rates is the only factor that determines the effect of synaptic rewiring, we would expect to see a complete reversal of the effect when swapping the mean firing rates of inhibitory and excitatory neurons. This however does not seem to be the case (Figure 4.4B). The effect of EI rewiring on the inhibitory population, after swapping the difference in mean firing rates, is not as strong as the effect of IE rewiring on the excitatory population in the original model. Similarly, the effect of EE rewiring does not mirror the modest effect of II rewiring on both populations in the old model. This indicates that firing rate is not the whole story, but rather that other properties of the network's inhibitory synapses also play a part in the effect of inhibitory synaptic rewiring.

model	1	2	3	4
N_E	32,000	20,000	8,000	4,000
N_I	8,000	20,000	32,000	36,000
mean E firing (Hz)	0.90	1.20	2.53	4.50
mean I firing (Hz)	5.45	2.02	1.68	0.98

Table 4.1: Neuron numbers and mean excitatory (E) and inhibitory (I) firing rates in four different versions of our random network model. Model 1 is the original version.

4.1.4 Synaptic Weights Matter

The parameters for the lognormal synaptic weight distributions for each synaptic population were taken from experimental measurements in the mouse auditory and barrel cortices (EPSP/IPSP measurements) [10, 22, 81]. Although the distributions have sim-

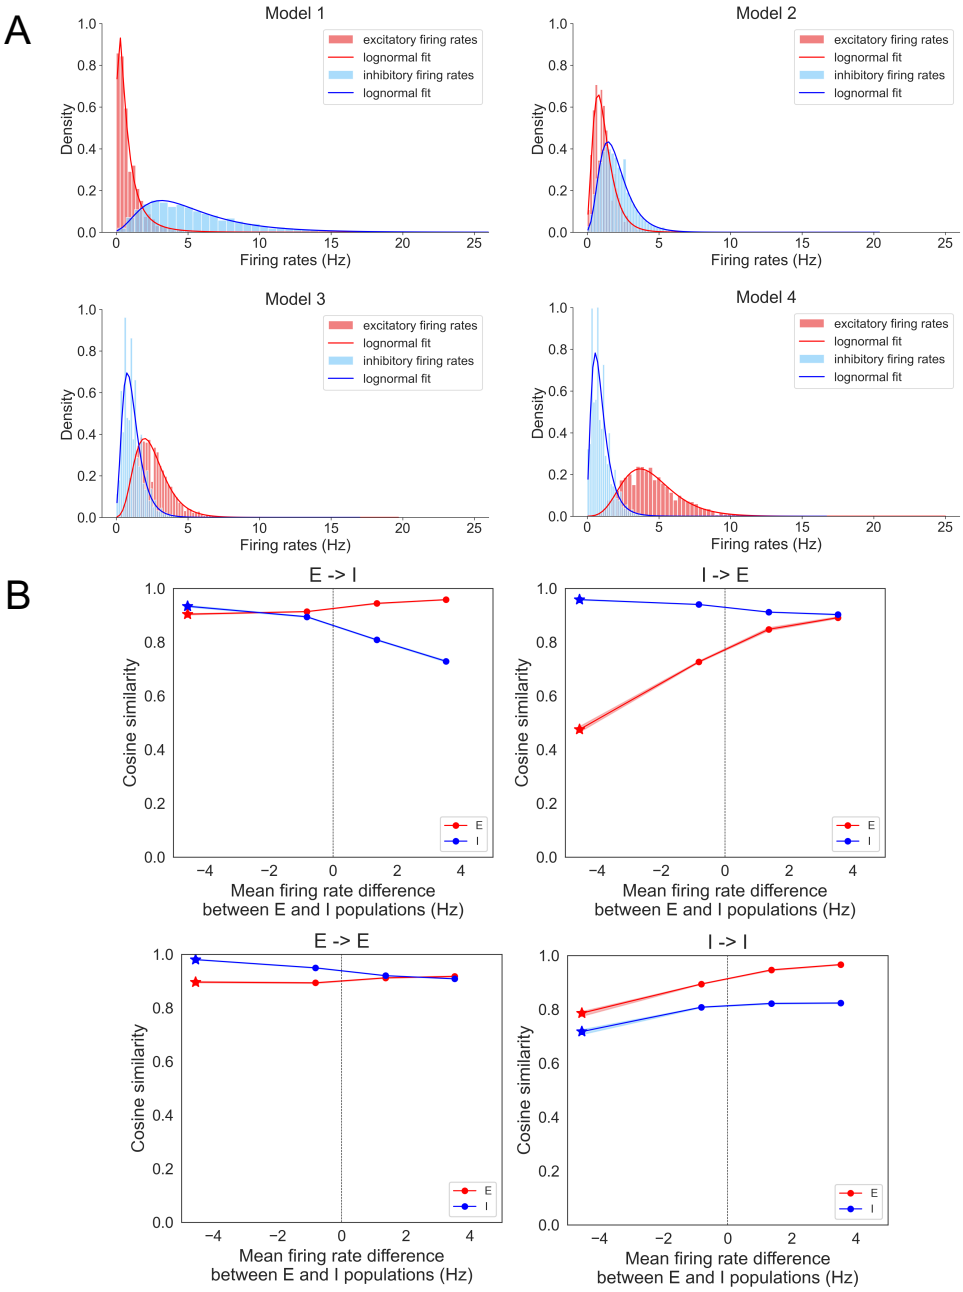

Figure 4.4: Impact of firing rate differences on the effect of synaptic rewiring. A: Distribution of inhibitory (blue) and excitatory (red) firing rates in four different versions of our random network model. B: Effect of rewiring entire synaptic populations on the firing rate vectors of the four versions of the model, as a function of the difference in mean excitatory (E) and inhibitory (I) firing rates. If the difference in mean firing rates was the only factor to determine the effect of rewiring, the EI and IE plots, and the EE and II plots should be mirror images of each other. N=3. Markers: mean; shading: range. Star-shaped marker: values for original model (Model 1).

ilar medians (EE: 0.27mV, EI: 0.36mV , IE: 0.42mV , II: 0.39mV), the distribution of IE weights has a heavier right tail (Figure 4.5A), also indicated by a higher kurtosis (EE: 30, EI: 12.13, IE: 82.90, II: 29.27) and skewness (EE: 3.61, EI: 2.44, IE: 5.57, II: 3.54). This means that the population of IE synapses contains more synapses with large weights. We thus wondered whether synaptic weight played a role in the influence of inhibitory synapses on the activity of our network. To test this, we repeated the experiment described in Section 4.1.2, but this time rewired synapses with the highest weights first. This resulted in a very similar effect to the one observed when rewiring by firing rate (Figure 4.3), both when rewiring IE synapses (Figure 4.5B) and II synapses (Supplementary Figure B.2A). Inhibitory synapses with a high weight had a larger influence on excitatory population activity than synapses with lower weights. This indicates that not only the firing rate of a synapse's presynaptic neuron, but also the weight of the synapse, determines its influence on the activity of a network.

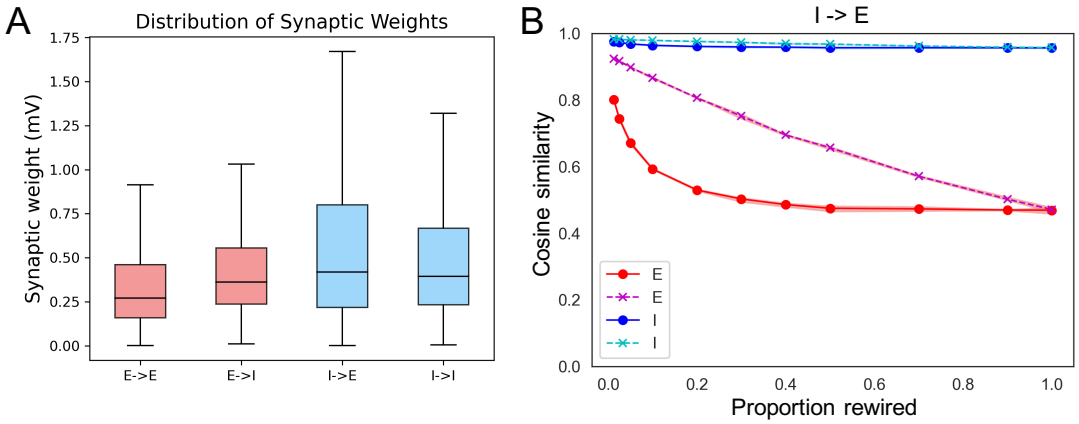

Figure 4.5: Influence of IE synaptic weight. A: Synaptic weight distribution of each of the different synaptic populations. Box extends from Q1 to Q3, line indicates median. Whisker range: (Q1 - 1.5IQR; Q3 + 1.5IQR), where IQR = Q3-Q1. B: Effect of rewiring IE synapses by synaptic weight (solid lines) compared to rewiring the same proportion of randomly selected synapses (dashed lines). Effect on the excitatory population activity is indicated in red, effect on inhibitory activity is in blue. N=3. Markers: mean; shading: range.

4.1.5 Firing Rate and Synaptic Weight Interact

The firing rate of a given excitatory neuron within our network negatively correlates with both the summed firing rate of all its presynaptic inhibitory neurons (Spearman's

correlation: -0.57), as well as the summed weight of all its incoming IE synapses (-0.69) (Figure 4.6A). We wanted to investigate whether these two factors interact. To combine the two, we defined the 'impact' of a synapse as: synapse weight $\times$ presynaptic firing rate. The correlation of the excitatory firing rates with the total IE synaptic impact received (-0.84) is higher than with one of the two factors alone (Figure 4.6A). This suggests that they might indeed interact in determining the influence of a given IE synapse on the overall activity of the network's excitatory neurons.

If presynaptic firing rate and weight interact, rewiring IE synapses with high synaptic impacts should cause a stronger effect than rewiring the same proportion of synapses with high presynaptic firing rates or weights alone. This indeed is the case (Figure 4.6B). In fact, rewiring 30% of the network's highest impact IE synapses (Figure 4.6B, solid arrow) results in an effect that is as strong as when rewiring 100% of IE synapses. To achieve the same effect, about 50% of IE synapses with high weights or 90% of IE synapses with high firing rates would have to be rewired. Moreover, rewiring 1.25% of highest impact IE synapses (Figure 4.6B, dashed arrow), which only constitute 0.23% of all synapses, already achieves an effect that is stronger than when rewiring 100%

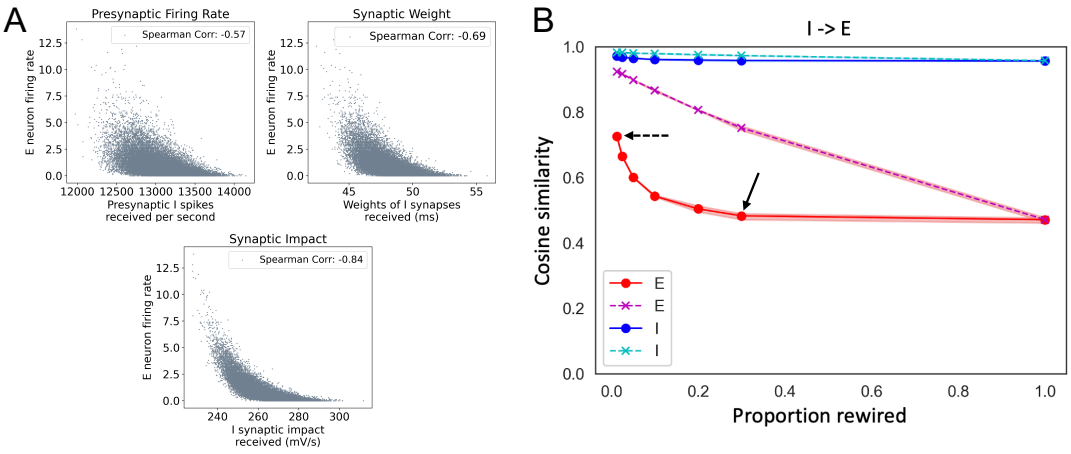

Figure 4.6: Influence of IE synaptic impact. A: Correlation of excitatory (E) neuron firing rate with the total firing rate of the presynaptic inhibitory (I) neurons, the total weight of I synapses received and the total synaptic input received. B: Effect of rewiring IE synapses by synaptic impact (solid lines) compared to rewiring the same proportion of randomly selected synapses (dashed lines). Effect on the excitatory population activity is indicated in red, effect on inhibitory activity is in blue. The arrows point at the effect of rewiring 1.25% (dashed) and 30% (solid line) of IE synapses. N=3. Markers: mean; shading: range.

of any of the other synaptic populations (EE, EI, II) (Figure 4.2). Again, rewiring II synapses by synaptic impact resulted in a similar, but more moderate effect on the excitatory and inhibitory populations of the model (Supplementary Figure B.2B). In line with our hypothesis, this indicates that there indeed is a small subset of inhibitory synapses that strongly influence the activity of the overall network. This small population seems to be characterised by a high synaptic impact.

4.2 The Ring Model

Our experiments on the random network model indicate that a small proportion of inhibitory synapses with high synaptic impacts has a disproportionately large influence on the activity of the overall network. To test whether the same principle applies to a network's function, we built a ring model of orientation selectivity with similar parameters as the random network model. The only major differences are that the model is smaller by a factor of 10, and that it receives inputs from the LGN, which convey information about the stimulus orientation (details in Methods 3.1.2). Furthermore, while one version (version 1) of the model contains completely random synaptic connections, the other version (version 2) has feature-specific EE connections that vary depending on the difference between two neuron's preferred orientations (POs). Like the random network model, the ring model has lognormal firing rate distributions (Figure 4.7A).

The ring network models the response of excitatory simple cells in the visual cortex when we view lines oriented at different angles. Each simple cell has a PO at which it will respond most strongly, whereas it will only respond weakly to orientations orthogonal to its PO. How selective a given neuron is for orientations close to its PO is characterised by the neuron's tuning curve (Figure 4.7B). When a given angle is shown to the network, the activity of the network's overall excitatory population forms a bump, with neurons with a PO close to the stimulus orientation being most active (Figure 4.7C). Both versions of our ring model display contrast invariance, which means that a range of different LGN input strengths ($\bar{r}_E^{LGN} = [60Hz, 110Hz]$; $r_I^{LGN} = 0.8 \times \bar{r}_E^{LGN}$) elicits a network response with the same magnitude and shape (as in Figure 4.7C).

To investigate the influence of different types of synapses on the ring model, we first assessed the effect of synaptic rewiring on the model's activity, and the orientation selectivity of its excitatory neurons. We then assessed the effect of synaptic rewiring on the performance of a decoder designed to read out stimulus information from the network's activity.

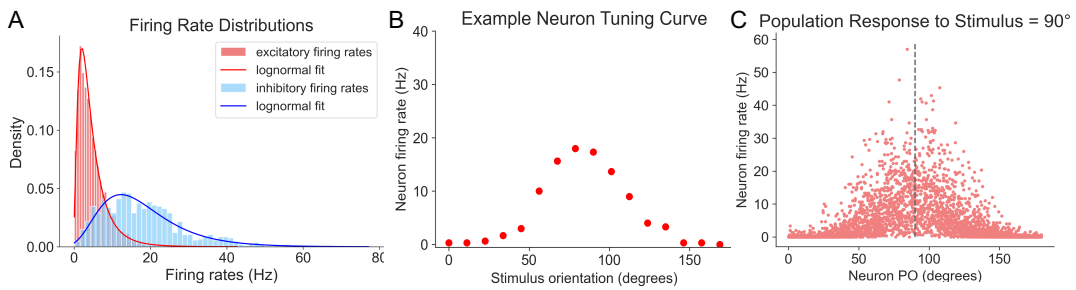

Figure 4.7: Characteristics of the Ring Model. A: Firing rate distributions of the network's excitatory (red) and inhibitory (blue) neuronal populations. The distribution shows each neuron's average response to 16 different stimuli. B: Example of a response of a single excitatory neuron to stimuli at 16 different orientations. The curve that is formed by the neuron's responses is called the neuron's tuning curve. The preferred orientation (PO) of the displayed neuron is close to 90 degrees. C: Response of the network's excitatory population to a stimulus oriented at 90 degrees, as a function of each neuron's PO. The figures shown are from the ring model version 1 with no feature-specific connectivity, but are the same for the other model version.

4.2.1 Synaptic Rewiring Affects Network Activity and Orientation Selectivity

As a starting point, we assessed the effect of rewiring entire synaptic populations (EE, EI, IE, II) on the ring model. In addition to measuring the effect on the network's activity as before, we measured changes in the orientation selectivity of the model's excitatory neurons by comparing their orientation selectivity indices (OSIs) [99] before and after rewiring. The effect of synaptic rewiring on the model's firing rate vectors followed the same trend as in the random network model (Figure 4.8A). Rewiring IE synapses elicited the largest effect on the activity of the excitatory population and rewiring II synapses had a modest effect on both the excitatory and inhibitory populations, whereas rewiring excitatory synapses (EE, EI) barely had any effect on the model's activity. As before, the means and distributions of firing rates remained unaltered (Kolmogorov-Smirnov; $p > 0.4$). The effect of rewiring was similar in both versions of our model, which indicates that the presence of feature-specific excitatory connections does not significantly alter the influence of given synapses on the network.

Synaptic rewiring also caused a shift in the OSIs of the network's excitatory neurons. As with the network activity, the largest shift was induced by rewiring IE synapses (Figure 4.8B), and the effect was similar for both model versions. The distribution of

OSIs did not change. The OSI quantifies how broad or narrow a neuron's tuning curve is. Neurons with OSIs close to 1 have a very narrow tuning curve and therefore selectively respond to stimuli at their PO, whereas neurons with lower OSIs have wider tuning curves and are less selective, and their responses thus contain less information about the current stimulus orientation. In the case of IE rewiring, neurons' OSIs often shifted by values of 0.2 or more (Figure 4.8B). This suggests that IE rewiring might change the information content of individual neurons, with some neurons containing more precise information about the stimulus angles after rewiring, and others less.

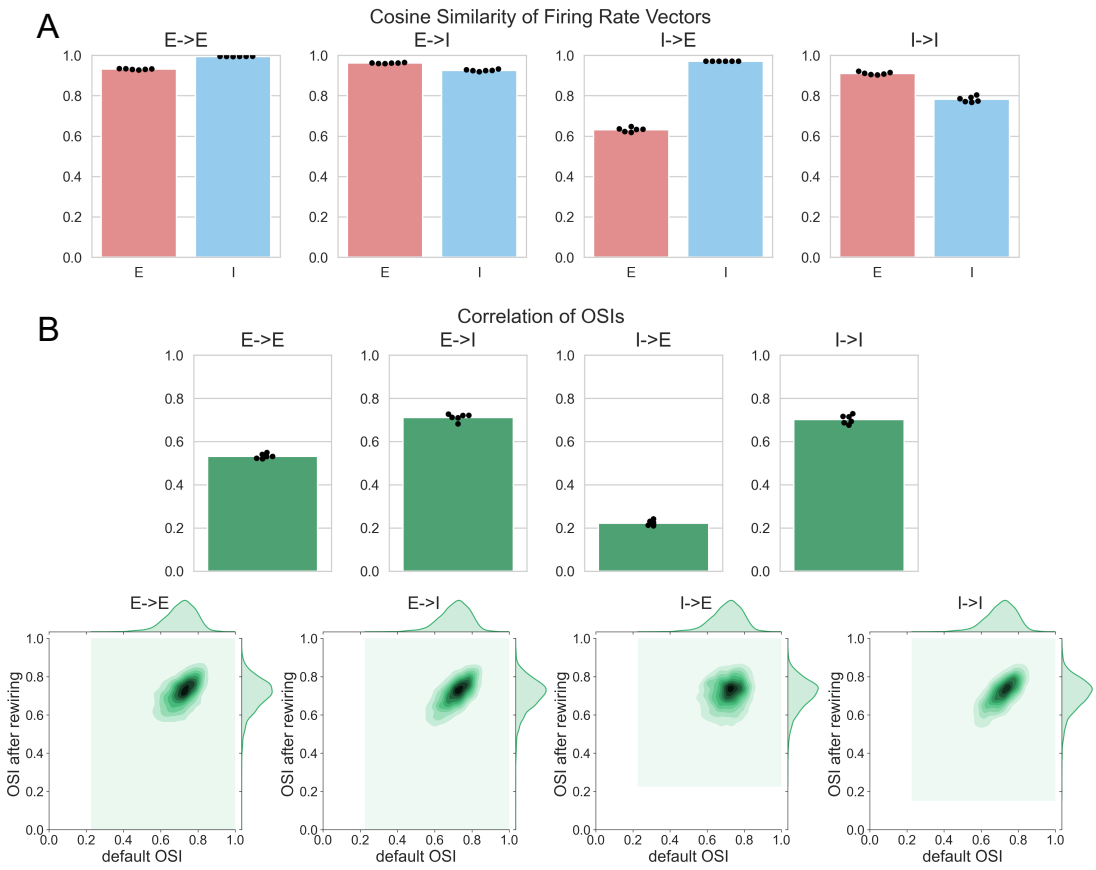

Figure 4.8: Effect of synaptic rewiring on the ring model version 1. A: Cosine similarities of the excitatory (red) and inhibitory (blue) average firing rate vectors before and after rewiring different synaptic populations. The average was taken over the network's response to 16 different stimuli. N = 6 independent simulations. B: Correlation of excitatory neuron OSIs before and after rewiring. N = 6. The bottom row provides more detail on how the OSIs shifted for each type of rewiring. Density plots on top and right axes show OSI distributions before and after rewiring. Synaptic rewiring elicited similar effects on the model version 2 (Supplementary Figure B.3).

4.2.2 Shift in Activity, but not in OSIs, Affects Decoder Performance

To test whether the shifts in neuronal activity and orientation selectivity affect the representation of stimulus information in the ring model, we constructed a linear decoder designed to read out the stimulus orientation from the network's activity. Because excitatory neurons form long-range connections to other brain regions, whereas inhibitory neurons mainly connect to neurons locally within a network [104], our decoder only received inputs from excitatory neurons. The decoder can thus be biologically interpreted as a readout neuron that is situated in a downstream brain region and reads out the information represented by the network of orientation-selective simple cells. Because the network contains excitatory neurons with a wide range of mean firing rates (Figure 4.7A), we normalised the firing rate of each neuron with the neuron's mean firing rate before feeding it into the decoder. This can biologically be interpreted as the decoder neuron adapting its synaptic weights to the mean firing rates of the synapses' presynaptic neurons.

The decoder achieves a mean absolute error (MAE) and mean decoding variance ranging from 8.0 ± 2.46 degrees (MAE $\pm$ variance) and 0.53 ± 0.054 degrees (mean decoding variance $\pm$ variance), when reading out from a population subset of 0.5% (16 neurons), to 0.23 ± 0.0034 degrees and $5.2 \times 10^{-4} \pm 6.5 \times 10^{-8}$ degrees, when reading out from 100% (3,200 neurons) of the population (evaluated for 16 orientations between 0 and 180 degrees; N = 6 independent simulations). Initial exploration of the decoder performance on different population subsets indicated that reading out from a subset of neurons with high OSIs, instead of a random subset, did not significantly improve the decoder's performance. This suggests that a shift in OSIs due to rewiring might not affect the decoder much. To test this, we evaluated the performance of the decoder on different random population subsets of the network, before and after rewiring different proportions of IE synapses. As in the experiment in Section 4.1.5, we always rewired the subset of IE synapses with the highest synaptic impacts. Because synaptic rewiring not only leads to a shift in OSIs but also in neuron firing rates, we tested two different cases: In the one case (constant normalisation), we normalised the neuron firing rates after rewiring with their original mean firing rates before rewiring. In the other case (adapted normalisation), we adapted the normalisation and used the neuron's new means after rewiring to normalise their firing rates after rewiring. The latter case corresponds to a decoder neuron that would be able to adapt its synaptic weights to changes in the mean rates of its individual synaptic inputs, a mechanism

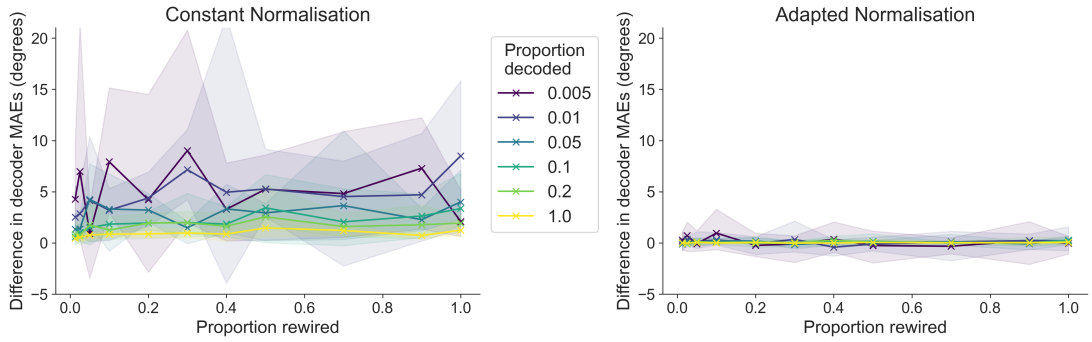

Figure 4.9: Difference in decoder mean absolute error (MAE) before and after rewiring different proportions of IE synapses in the ring model version 1. Positive differences indicate worse decoder performance after rewiring. On the left, the neuron's mean firing rates before rewiring were used to normalise both the firing rate vectors before and after rewiring. On the right, the normalisation was adapted and the neuron's firing rates after rewiring were normalised with their means after rewiring. N = 6. Markers: mean; shading: range.

which possibly exists in the brain [105–107].

IE rewiring significantly worsened the MAE of the decoder with constant normalisation, but not the performance of the decoder with adapted normalisation in our ring model version 1 (Figure 4.9). The same trend was observed for the mean decoding variance (Supplementary Figure B.4). This indicates that the shift in neuron firing rates induced by IE rewiring has the potential to affect the ability of a downstream decoder that does not compensate for shifts in mean firing rates. As no rewiring effects on the decoder could be observed when the normalisation of the firing rates was adapted, the shift in OSIs induced by IE rewiring does not seem to affect decoder performance. One should note however that even in the case of the decoder with constant normalisation, rewiring only has a significant effect if a small subset ($\leq 20\%$) of the population is used as input to the decoder (Figure 4.9). When decoding from larger subsets, the difference became negligible. Overall, the results indicate that the shift in orientation selectivities, within the range it occurs in our network, does not affect the representation of information within the ring model enough to alter the population readout. Only the shift in firing rates seems to affect the readout, but this can be compensated by simply scaling the decoder neuron's synaptic weights according to the mean input rate to its synapses, or by reading out from a large enough subset of the population.

Chapter 5

Discussion and Conclusions

In this dissertation, we investigated the influence of different types of synapses on the activity and function of a neuronal network. This was motivated by recent observations that neuronal networks are sloppy, which means that their overall behaviour seems to depend on a small number of 'stiff' components, whereas the other 'sloppy' components are free to change without significantly affecting the overall network [6, 72, 75, 76]. In the brain, neuronal networks constantly face both spontaneous and plasticity-induced changes to their synaptic connections. Sloppiness provides a candidate explanation for how neuronal networks allow for these changes, while not letting them destabilise overall network function, disturb other ongoing processes or cause the loss of stored information. The idea is that a small number of stable, stiff synapses might support overall network stability and allow other synapses to change.

Experimental studies that analysed recordings of neuronal activity in the brain suggest that stiff synapses might be associated with highly active neurons [6, 69, 70, 76]. However, due to constraints of current experimental recording techniques, it is difficult to further characterise the properties and role of stiff network components directly in the brain, without having specific clues about what to look for. This is where computational models of neuronal networks can be immensely useful tools. In this dissertation, we built on a previous computational study by Mongillo et al [10], in which the authors observed that rewiring all inhibitory synapses within a spiking neuronal network (SNN) model strongly affected the activity of the overall network, whereas rewiring excitatory synapses had little effect overall. In our experiments, we investigated the influence of different types of synapses on the activity and function of two SNN models in more detail. Combining the experimental observations [6, 69, 70, 76] and Mongillo et al's [10] results, our hypothesis was that a small subgroup of inhibitory synapses

that transmits signals of highly active neurons has a disproportionately high influence on the activity and function of neuronal networks, and might thus constitute the stiff components that control overall network stability.

To investigate this hypothesis, we first measured the effect of rewiring different groups of synapses on the overall activity of a randomly connected network model. This revealed that a small subgroup of inhibitory synapses indeed had a disproportionately large effect on the activity of the network. This small subgroup was composed of IE synapses with high synaptic impacts, which we defined as presynaptic firing rate $\times$ synaptic weight. The influence of these IE synapses was concentrated on the activity of the network's excitatory neurons, which constitute the majority of the network and are thought to convey the network's information content to other brain regions. The activity of the network's inhibitory neurons was most strongly influenced by a subgroup of II synapses, also characterised by high synaptic impacts. Our findings therefore support our hypothesis, but highlight that not only the firing rate of a synapse's presynaptic neuron, but also the synaptic weight seem to determine the influence of a synapse on the activity of the overall network. This is new knowledge, as Mongillo et al [10] had originally hypothesised that only the firing rate of the presynaptic neuron determines the influence of a synapse. We moreover show for the first time that, in line with the idea of sloppiness, a very small proportion of synapses indeed has a disproportionately large influence on the activity of a spiking neuronal network model.

Following these findings, we then tested the effect of synaptic rewiring on the function of a ring model, which models the phenomenon of orientation selectivity in the visual cortex. As in the random network model, rewiring IE synapses caused a strong shift in the network's excitatory activity, and moreover altered the strength of orientation selectivity of individual excitatory neurons (measured by the OSI). To test whether these shifts affected the representation of stimulus information in the ring model, we constructed a biologically interpretable linear decoder designed to read out the stimulus orientation from the network's activity. Against our initial intuition, the shift in OSIs induced by IE rewiring did not affect the representation of information within the ring model enough to alter the population readout. In contrast, the shift in firing rates induced by IE rewiring affected the readout, but this could be compensated by scaling the decoder's synaptic weights according to the mean input rates to its synapses, or by reading out from a large enough subset of the network's excitatory population.

Overall, our experiments support the idea that neuronal networks are sloppy, in that their activity seems to be strongly influenced by a small number of stiff inhibitory

synapses with high synaptic impacts, whereas most other sloppy synapses can change without significantly affecting the overall network. Therefore, if a network's stiff synapses remain stable, the activity of the overall network is expected to remain stable too. It is however less clear what implications this sloppiness has for the function of a neuronal network, as we could not detect a significant influence of stiff synapses on the representation of information within our ring model.

In light of our findings, this section will discuss a number of open questions that remain: How do the results of our simulations fit in with experimental observations in the brain? Is the high influence of a stiff synapse on a neuronal network determined only by its high synaptic impact, or does the synapses' inhibitory, rather than excitatory, nature play a role? And finally, what is the role of stiff synapses in neuronal networks? Before discussing these questions, it will be useful to consider the limitations of our computational models that relate to the biophysical complexity of biological neuronal networks. These will be explained below, before the open questions will be discussed.

5.1 Model Limitations

Our SNN models are constructed to align with experimental data from the brain in several ways: The models' firing rate distributions are lognormal, with higher inhibitory than excitatory rates, and their spiking patterns reproduce the asynchronous spiking behaviour of neuronal networks in the cortex [10, 87, 96]. Moreover, the synaptic weight distributions and connection probabilities match experimental measurements from the mouse barrel and auditory cortex [10, 22, 81]. For investigating the dynamics of large neuronal networks, the leaky integrate-and-fire (LIF) neurons, which are the building blocks of our SNN models, are generally considered good approximations of biological neurons, despite the fact that they make a number of simplifications (discussed below) [82, 84, 96]. These simplifications can be beneficial in several ways, as they improve computational efficiency, and make it easier to understand the dynamics of the network and relate them to individual features of the model [84]. Furthermore, more complicated models, such as the Hodgkin-Huxley model [108], contain large numbers of parameters that are often hard to estimate [84]. Nevertheless, it is important to keep in mind that neuronal networks in the brain are incredibly complex, and that our models do not account for a number of aspects that contribute to this complexity. The ones most relevant for our study will be pointed out below.

The first simplification that LIF neurons make is that they model neurons as a single

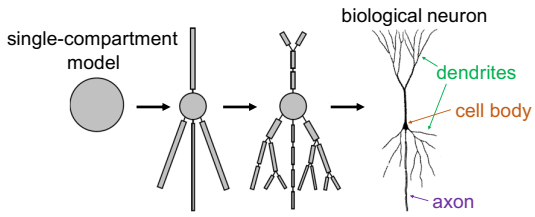

Figure 5.1: Single-compartment models, such as our LIF neurons (left), can be extended to include multiple compartments and approximate the physical structure of a neuron (right) to varying degrees of complexity. Figure adapted from [82].

compartment. This means that the state of a neuron at a given time point is described by a single variable, the membrane potential $v(t)$. However, biological neurons have a complex physical structure that consists of a cell body and long, branching protrusions that form the neuron's dendrites and axon (Figure 5.1, right). The membrane potential of a neuron varies along these protrusions, and incoming synapses can have different influences on the activity of a neuron depending on where they are located; a synapse on the far end of a dendrite for example is likely to have less of an influence than a synapse located close to the neuron's cell body [82, 109, 110]. In our SNN models, IE synapses with high synaptic impacts had the largest influence on the activity of the overall network. It is likely that, for biological networks, the definition of the synaptic impact would have to be extended to take the location of a given synapse into account. To test this, one could use multi-compartment models, which approximate the physical structure of individual neurons to varying degrees of complexity, by modelling the membrane potential of a neuron in terms of multiple, coupled compartments [82] (Figure 5.1). The challenge here however is to model these compartments in a way that still permits simulations of large enough neuronal networks.

Secondly, our models assume that the intrinsic membrane potential dynamics of each neuron within a network are the same. Neuronal networks in the brain however contain a wide range of different neuron types, with different membrane potential dynamics, responsiveness and spiking behaviours [111–113]. For example, some types of inhibitory neurons tend to have higher firing rates than others solely due to their intrinsic properties [81, 113]. Since firing rate is a factor that determines the influence of a synapse in our model, it would be relevant to extend the model to include these different types of inhibitory neurons. This could be achieved by using the set of LIF neuron types recently defined by Teeter et al [113], which model the different neuron types characterised in the brain. Perhaps, this would allow one to match the high

influence IE synapses in our model to a specific inhibitory neuron type. A biological study could then specifically investigate the dynamics of the synapses of this inhibitory neuron type further.

Thirdly, our models use relatively simple approximations of the dynamics of individual synapses. Our synapses are current-based, which means that every time a presynaptic neuron fires, the effect on the membrane potential of a postsynaptic neuron is the same [114]. In biological networks however, the effect of a synapse at time t depends on the membrane potential of the postsynaptic neuron at that time t [82]. The dynamics of neuronal network models with conductance-based synapses, which take this property into account, are slightly different to models with current-based synapses [114]. It would therefore be interesting to test whether our predictions hold for models with conductance-based synapses. Furthermore, conductance-based synapses allow refining the model to include different subtypes of excitatory and inhibitory synapses, that are different depending on the types of neurotransmitter receptors present at the synapse [82].

Overall, these limitations highlight that neuronal networks are complex, and that there are multiple factors that could affect the influence of a given synapse on the overall network, beyond what we have included in our models. However, our models constitute an important starting point that provides a fundamental understanding of the dynamics of neuronal networks in relation to their synapses. As a next step, our models could be gradually refined, perhaps starting by adding conductance-based synapses and then different neuronal cell types. At each added level of complexity, one could test whether our findings still hold, which would help to gradually build up an understanding of the influence of different synapses in more complex networks.

5.2 Open Questions

5.2.1 How do our Results Fit in with Experimental Observations?

Our experiments indicate that overall network activity might be kept stable by a small number of stiff, high impact IE synapses. This implies that these, too, would thus have to be stable themselves. Despite generally high synaptic volatility in the brain, some synapses seem to indeed be more stable than others [5, 47, 51]. These stable synapses tend to be larger, which is indicative of a high synaptic weight [5, 47, 51]. As synaptic weight is one factor that contributes to a high synaptic impact in our model, this thus

supports the idea that high impact synapses are more stable. It has to be kept in mind however that the studies cited above exclusively focused on excitatory synapses as these are easier to monitor. It is not clear yet whether the same principle applies to inhibitory synapses, and whether inhibitory synapses are generally more stable than excitatory synapses. Despite the technical difficulties, a systematic comparison of the stability of inhibitory with excitatory synapses in the brain would thus be beneficial.

Despite the activity changes induced by IE rewiring, our decoder was able to reliably read out stimulus information from our ring model. This could be achieved either by reading out from a large enough subset of the population or by adapting the decoder's synaptic weights to changes in the mean rates of its individual inputs. In the brain, the latter would require neurons to monitor the average activity of each of their synapses, and individually adapt their synaptic weight if lasting changes to the activity occur. Alongside other homeostatic mechanisms, biological neurons indeed seem to scale their synaptic weights to adapt to changes in their inputs [106, 115–118], although it has been debated whether this occurs at the level of individual synapses [107] or whether neurons only collectively scale all their synapses when changes in their overall input occur [115, 119]. If the former was the case, this might present a possible mechanism that supports robust network function in the face of synaptic volatility.

5.2.2 Does Inhibition Play a Role in Determining the Influence of a Synapse?

Inhibitory synapses within our network models generally have higher synaptic impacts than excitatory synapses. This is due to the higher firing rates of the inhibitory neurons in our models, as well as the heavier tail of the IE synaptic weight distribution. The question therefore is: does the effect of inhibitory synaptic rewiring on the activity of our networks solely stem from their higher synaptic impacts, or does the synapses' inhibitory nature play a role? The former would likely be the case if inhibitory and excitatory synaptic signals had symmetric effects on the dynamics of a network. However, due to the way neurons integrate synaptic signals, the effect of inhibition and excitation is not strictly symmetrical. Multiple excitatory inputs arriving at the same time, or in quick succession, can only increase a neuron's membrane potential until it reaches the neuron's threshold, where an action potential will be triggered and the neuron's membrane potential will be reset. In contrast, multiple inhibitory inputs can, in theory, decrease the membrane potential of a neuron in our model indefinitely. The

neuron threshold thus adds asymmetry to the system. One should note however that the effect of inhibitory inputs is not strictly the same in our model as in biological networks, as the influence of a synapse depends on the current membrane potential of a biological neuron, and inhibitory inputs can only decrease a neuron's membrane potential up to a certain point (another argument for adding conductance-based synapses to our models, see Section 5.1). Nevertheless, the argument with the threshold still holds.

To quantify the asymmetric influence of excitation and inhibition, it would be useful to modify our random network model such that both the firing rate and synaptic weight distributions are either swapped or equal, and then re-evaluate the effect of synaptic rewiring. I briefly attempted to simulate this, by swapping all parameters for excitatory and inhibitory neurons or by making them equal, but could not obtain a stable network model in either case. With considerably more parameter fine-tuning, it is likely feasible to obtain a network with equal firing rate and synaptic weight distributions, which would make it possible to investigate how much of the effect of synaptic rewiring stems from the synapses' inhibitory nature rather than their high synaptic impacts.

5.2.3 What is the Role of Stiff Synapses in Neuronal Networks?

Although rewiring IE synapses affected the firing rates and the tuning curves (measured by the OSI) of individual neurons in our ring model, the latter did not have a significant effect on the population readout. This might indicate that information can be robustly represented in the face of synaptic volatility, and does not require stabilising stiff synapses, as long as downstream neurons read out information from a large enough subset of the population or dynamically adapt their synaptic weights to changes in the mean rate of their inputs. One should however keep in mind that the task that we set our ring model and decoder was relatively simple. Firstly, the network in our ring model is relatively strongly driven by its inputs from the LGN, which stayed the same before and after rewiring. Secondly, networks in the brain tend to perform multiple computations within the same network. In addition to information about the stimulus orientation, area V1 of the visual cortex for example also processes information about colour or the direction of stimulus movement [11].

Stiff neurons and synapses might play a role in ensuring that networks can perform multiple computations in parallel, without them interfering with each other. Evidence that supports this idea comes from a recent study by Ponce-Alvarez et al [76]

in which the authors analysed a pairwise maximum entropy model fitted to recordings of neuronal activity in area A1 of the mouse auditory cortex (Area A1 is the auditory counterpart of the visual area V1). Interestingly, neurons that represented information about sound were associated with sloppy directions of the model [76]. In contrast, neurons associated with stiff directions exhibited stable activity, were less responsive to sound, and seemed to control the networks' cortical state, which is the overall mode of network activity [76,77]. In light of these observations, the authors propose that stiff neurons control and stabilise the activity of the overall network, such that sloppy neurons can locally respond to and process information about sound, without influencing the whole network [76]. This might permit multiple processes to take place in parallel, without interfering with each other. The authors also propose that the synapses of sloppy neurons might be subject to more synaptic plasticity, which would allow the network to locally adapt its computations or store new information, without interfering with other network functions [76]. This fits with our observation that, as long as a small subgroup of stiff synapses is stable, other synapses can undergo changes without affecting the overall network.

Stiff synapses might also play a role in the storage of information within a network. In addition to their random network model, Mongillo et al [10] also analysed a spiking attractor network, which is an SNN that can store multiple patterns of activity as memories. A given stored activity pattern can be activated by briefly increasing the external input to a subset of neurons that, in the given activity pattern, should be highly active [10]. The network will then settle into the activity pattern (also called attractor state), even if the input is restored to baseline levels, until a new memory pattern is activated by selectively increasing the input to a different subset of neurons again [10]. The authors tested the effect of rewiring all EE synapses within this network, which barely had any effect. In contrast, rewiring all II synapses led to loss of the stored memory patterns [10]. This suggests that the stability of inhibitory synapses might be more important for the storage of information within a neuronal network than the stability of excitatory synapses. Mongillo et al [10] however did not separately test the effect of EI and IE synaptic rewiring in the network, the latter of which would have been especially interesting in light of our findings.

Overall, although stiff synapses did not seem to have a significant influence on the representation of information within our ring model, other studies indicate that they might play a role in the stable storage [10] or the parallel processing [76] of information within networks.

5.3 Conclusion

All in all, in this dissertation, we systematically assess the effect of rewiring different groups of synapses on the activity and function of two spiking neuronal network models. We show for the first time that a very small subgroup of inhibitory synapses has a disproportionately large influence on the overall activity of such a neuronal network. This subgroup consists of IE synapses with high synaptic impacts, which we define as presynaptic firing rate $\times$ synaptic weight. This extends previous findings by Mongillo et al [10] and demonstrates that not only firing rate but also synaptic weight determine the influence of a synapse. Building on the idea of sloppiness in neuronal networks, the small subgroup of high impact IE synapses might be the stiff synapses that keep a network stable, while allowing other synapses to change. This might explain how neuronal networks maintain stability despite high rates of synaptic volatility.

In a ring model of orientation selectivity, we show that the subgroup of high impact IE synapses has a high influence on the mean activity of the network's neurons, but does not significantly affect the way they represent information. As long as downstream neurons that read out this information can dynamically adapt to fluctuations in the network's mean firing rates, a network such as the ring model might not need to rely on stabilising stiff synapses. However, stiff synapses might have other functions within neuronal networks in the brain, beyond what we have modelled here. These might include supporting stable storage or parallel processing of information within networks.

Overall, our computational study enhances our understanding of the dynamics of neuronal networks. Our experiments provide a starting point upon which future research can build, for example by gradually adding more complexity to our models and thus refining our characterisation of the properties and role of stiff synapses. This might then lead to the design of targeted biological experiments, that can characterise the role of stiff synapses directly in the brain. In the long run, this promises to contribute to our understanding of the principles of brain function.

Bibliography

[1] Gianluigi Mongillo, Simon Rumpel, and Yonatan Loewenstein. Intrinsic volatility of synaptic connections — a challenge to the synaptic trace theory of memory. *Current Opinion in Neurobiology*, 46:7–13, 2017.

[2] Pico Caroni, Flavio Donato, and Dominique Muller. Structural plasticity upon learning: regulation and functions. *Nature Reviews Neuroscience*, 13(7):478–490, 2012.

[3] Anna R. Chambers and Simon Rumpel. A stable brain from unstable components: Emerging concepts and implications for neural computation. *Neuroscience*, 357:172–184, 2017.

[4] Alessio Attardo, James E Fitzgerald, and Mark J Schnitzer. Impermanence of dendritic spines in live adult ca1 hippocampus. *Nature*, 523(7562):592–596, 2015.

[5] Yonatan Loewenstein, Uri Yanover, and Simon Rumpel. Predicting the dynamics of network connectivity in the neocortex. *Journal of Neuroscience*, 35(36):12535–12544, 2015.

[6] Dagmara Panas, Hayder Amin, Alessandro Maccione, Oliver Muthmann, Mark van Rossum, Luca Berdondini, and Matthias H. Hennig. Sloppiness in spontaneously active neuronal networks. *Journal of Neuroscience*, 35(22):8480–8492, 2015.

[7] Bryan C. Daniels, Yan Jiun Chen, James P. Sethna, Ryan N. Gutenkunst, and Christopher R. Myers. Sloppiness, robustness, and evolvability in systems biology. *Current Opinion in Biotechnology*, 19(4):389–395, 2008.

[8] Ryan N. Gutenkunst, Joshua J. Waterfall, Fergal P. Casey, Kevin S. Brown, Christopher R. Myers, and James P. Sethna. Universally sloppy parameter sen-

sitivities in systems biology models. *PLoS Computational Biology*, 3(10):1871–1878, 2007.

[9] Mark K. Transtrum, Benjamin B. Machta, Kevin S. Brown, Bryan C. Daniels, Christopher R. Myers, and James P. Sethna. Perspective: Sloppiness and emergent theories in physics, biology, and beyond. *Journal of Chemical Physics*, 143(1), 2015.

[10] Gianluigi Mongillo, Simon Rumpel, and Yonatan Loewenstein. Inhibitory connectivity defines the realm of excitatory plasticity. *Nature Neuroscience*, 21(10):1463–1470, 2018.

[11] Mark F. Bear, Barry W. Connors, and Michael R. Paradiso. *Neuroscience - Exploring the Brain*. Wolters Kluwer, 4th edition, 2016.

[12] Jeff Hawkins and Subutai Ahmad. Why neurons have thousands of synapses, a theory of sequence memory in neocortex. *Frontiers in neural circuits*, 10:23, 2016.

[13] John Carew Eccles. From electrical to chemical transmission in the central nervous system: the closing address of the sir henry dale centennial symposium cambridge, 19 september 1975. *Notes and records of the Royal Society of London*, 30(2):219–230, 1976.

[14] D. Marr and T. Poggio. From understanding computation to understanding neural circuitry. Technical report, Massachusetts Institute of Technology, USA, 1976.

[15] Rafael Yuste. From the neuron doctrine to neural networks. *Nature reviews neuroscience*, 16(8):487–497, 2015.

[16] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.

[17] Danielle S Bassett and Michael S Gazzaniga. Understanding complexity in the human brain. *Trends in cognitive sciences*, 15(5):200–209, 2011.

[18] Leonidas MA Richter and Julijana Gjorgjieva. Understanding neural circuit development through theory and models. *Current Opinion in Neurobiology*, 46:39–47, 2017.

[19] Sadra Sadeh, Claudia Clopath, and Stefan Rotter. Emergence of functional specificity in balanced networks with synaptic plasticity. *PLoS computational biology*, 11(6):e1004307, 2015.

[20] Sarah EV Richards and Stephen D Van Hooser. Neural architecture: from cells to circuits. *Journal of neurophysiology*, 120(2):854–866, 2018.

[21] Ami Citri and Robert C Malenka. Synaptic plasticity: multiple forms, functions, and mechanisms. *Neuropsychopharmacology*, 33(1):18–41, 2008.

[22] Anthony Holtmaat and Karel Svoboda. Experience-dependent structural synaptic plasticity in the mammalian brain. *Nature Reviews Neuroscience*, 10(9):647–658, 2009.

[23] Gina Turrigiano. Too Many Cooks? Intrinsic and Synaptic Homeostatic Mechanisms in Cortical Circuit Refinement. *Annual Review of Neuroscience*, 34(1):89–103, 2011.

[24] David J Margolis, Henry Lütcke, Kristina Schulz, Florent Haiss, Bruno Weber, Sebastian Kügler, Mazahir T Hasan, and Fritjof Helmchen. Reorganization of cortical population activity imaged throughout long-term sensory deprivation. *Nature neuroscience*, 15(11):1539, 2012.

[25] Daniel Huber, DA Gutnisky, Simon Peron, Daniel H O'Connor, J Simon Wiegert, Lin Tian, TG Oertner, LL Looger, and K Svoboda. Multiple dynamic representations in the motor cortex during sensorimotor learning. *Nature*, 484(7395):473–478, 2012.

[26] Wickliffe C Abraham, Owen D Jones, and David L Glanzman. Is plasticity of synapses the mechanism of long-term memory storage? *NPJ science of learning*, 4(1):1–10, 2019.

[27] SJ Martin and RGM Morris. New life in an old idea: the synaptic plasticity and memory hypothesis revisited. *Hippocampus*, 12(5):609–636, 2002.

[28] Kalen P Berry and Elly Nedivi. Experience-dependent structural plasticity in the visual system. *Annual review of vision science*, 2:17–35, 2016.

[29] Torsten N Wiesel and David H Hubel. Single-cell responses in striate cortex of kittens deprived of vision in one eye. *Journal of neurophysiology*, 26(6):1003–1017, 1963.

[30] Daniëlle van Versendaal, Rajeev Rajendran, M Hadi Saiepour, Jan Klooster, Laura Smit-Rigter, Jean-Pierre Sommeijer, Chris I De Zeeuw, Sonja B Hofer, J Alexander Heimel, and Christiaan N Levelt. Elimination of inhibitory synapses is a major component of adult ocular dominance plasticity. *Neuron*, 74(2):374–383, 2012.

[31] Yanmei Zhou, Baoling Lai, and Wen-Biao Gan. Monocular deprivation induces dendritic spine elimination in the developing mouse visual cortex. *Scientific Reports*, 7(1):1–11, 2017.

[32] Arianna Maffei, Sacha B Nelson, and Gina G Turrigiano. Selective reconfiguration of layer 4 visual cortical circuitry by visual deprivation. *Nature neuroscience*, 7(12):1353–1359, 2004.

[33] Gerald T Finnerty, Langdon SE Roberts, and Barry W Connors. Sensory experience modifies the short-term dynamics of neocortical synapses. *Nature*, 400(6742):367–371, 1999.

[34] Joshua T Trachtenberg, Brian E Chen, Graham W Knott, Guoping Feng, Joshua R Sanes, Egbert Welker, and Karel Svoboda. Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. *Nature*, 420(6917):788–794, 2002.

[35] Mengia-S Rioult-Pedotti, Daniel Friedman, Grzegorz Hess, and John P Donoghue. Strengthening of horizontal cortical connections following skill learning. *Nature neuroscience*, 1(3):230–234, 1998.

[36] Woong Bin Kim and Jun-Hyeong Cho. Encoding of discriminative fear memory by input-specific ltp in the amygdala. *Neuron*, 95(5):1129–1146, 2017.

[37] Akiko Hayashi-Takagi, Sho Yagishita, Mayumi Nakamura, Fukutoshi Shirai, Yi I Wu, Amanda L Loshbaugh, Brian Kuhlman, Klaus M Hahn, and Haruo Kasai. Labelling and optical erasure of synaptic memory traces in the motor cortex. *Nature*, 525(7569):333–338, 2015.

[38] Cora Sau Wan Lai, Thomas F Franke, and Wen-Biao Gan. Opposite effects of fear conditioning and extinction on dendritic spine remodelling. *Nature*, 483(7387):87–91, 2012.

[39] Min Fu, Xinzhu Yu, Ju Lu, and Yi Zuo. Repetitive motor learning induces coordinated formation of clustered dendritic spines in vivo. *Nature*, 483(7387):92–95, 2012.

[40] Tonghui Xu, Xinzhu Yu, Andrew J Perlik, Willie F Tobin, Jonathan A Zweig, Kelly Tennant, Theresa Jones, and Yi Zuo. Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature*, 462(7275):915–919, 2009.

[41] Kaja Ewa Moczulska, Juliane Tinter-Thiede, Manuel Peter, Lyubov Ushakova, Tanja Wernle, Brice Bathellier, and Simon Rumpel. Dynamics of dendritic spines in the mouse auditory cortex during memory formation and memory recall. *Proceedings of the National Academy of Sciences*, 110(45):18315–18320, 2013.

[42] Sadegh Nabavi, Rocky Fox, Christophe D Proulx, John Y Lin, Roger Y Tsien, and Roberto Malinow. Engineering a memory with ltd and ltp. *Nature*, 511(7509):348–352, 2014.

[43] Kareem Abdou, Mohammad Shehata, Kiriko Choko, Hirofumi Nishizono, Mina Matsuo, Shin-ichi Muramatsu, and Kaoru Inokuchi. Synapse-specific representation of the identity of overlapping memory engrams. *Science*, 360(6394):1227–1231, 2018.

[44] Mauro Costa-Mattioli, Delphine Gobert, Elad Stern, Karine Gamache, Rodney Colina, Claudio Cuello, Wayne Sossin, Randal Kaufman, Jerry Pelletier, Kobi Rosenblum, et al. eif2α phosphorylation bidirectionally regulates the switch from short-to long-term synaptic plasticity and memory. *Cell*, 129(1):195–206, 2007.

[45] Jaime Grutzendler, Narayanan Kasthuri, and Wen-Biao Gan. Long-term dendritic spine stability in the adult cortex. *Nature*, 420(6917):812–816, 2002.

[46] Yi Zuo, Aerie Lin, Paul Chang, and Wen-Biao Gan. Development of long-term dendritic spine stability in diverse regions of cerebral cortex. *Neuron*, 46(2):181–189, 2005.

[47] Anthony JGD Holtmaat, Joshua T Trachtenberg, Linda Wilbrecht, Gordon M Shepherd, Xiaoqun Zhang, Graham W Knott, and Karel Svoboda. Transient

and persistent dendritic spines in the neocortex in vivo. *Neuron*, 45(2):279–291, 2005.

[48] Yonatan Loewenstein, Annerose Kuras, and Simon Rumpel. Multiplicative dynamics underlie the emergence of the log-normal distribution of spine sizes in the neocortex in vivo. *Journal of Neuroscience*, 31(26):9481–9488, 2011.

[49] Amir Minerbi, Roni Kahana, Larissa Goldfeld, Maya Kaufman, Shimon Marom, and Noam E Ziv. Long-term relationships between synaptic tenacity, synaptic remodeling, and network activity. *PLoS biology*, 7(6), 2009.

[50] Nobuaki Yasumatsu, Masanori Matsuzaki, Takashi Miyazaki, Jun Noguchi, and Haruo Kasai. Principles of long-term dynamics of dendritic spines. *Journal of Neuroscience*, 28(50):13592–13608, 2008.

[51] Dylan P Quinn, Annette Kolar, Sydney A Harris, Michael Wigerius, James P Fawcett, and Stefan R Krueger. The stability of glutamatergic synapses is independent of activity level, but predicted by synapse size. *Frontiers in Cellular Neuroscience*, 13:291, 2019.

[52] Akira Nagaoka, Hiroaki Takehara, Akiko Hayashi-Takagi, Jun Noguchi, Kazuhiko Ishii, Fukutoshi Shirai, Sho Yagishita, Takanori Akagi, Takanori Ichiki, and Haruo Kasai. Abnormal intrinsic dynamics of dendritic spines in a fragile x syndrome mouse model in vivo. *Scientific reports*, 6(1):1–11, 2016.

[53] Akihiro Inoue and Shigeo Okabe. The dynamic organization of postsynaptic proteins: translocating molecules regulate synaptic function. *Current opinion in neurobiology*, 13(3):332–340, 2003.

[54] Daniel Choquet and Antoine Triller. The role of receptor diffusion in the organization of the postsynaptic membrane. *Nature Reviews Neuroscience*, 4(4):251–265, 2003.

[55] Seth GN Grant. Synapse molecular complexity and the plasticity behaviour problem. *Brain and Neuroscience Advances*, 2:2398212818810685, 2018.

[56] K Harris and SB Kater. Dendritic spines: cellular specializations imparting both stability and flexibility to synaptic function. *Annual review of neuroscience*, 17(1):341–371, 1994.

[57] Alan Peters. Examining neocortical circuits: some background and facts. *Journal of neurocytology*, 31(3-5):183–193, 2002.

[58] Henry Markram, Maria Toledo-Rodriguez, Yun Wang, Anirudh Gupta, Gilad Silberberg, and Caizhi Wu. Interneurons of the neocortical inhibitory system. *Nature reviews neuroscience*, 5(10):793–807, 2004.

[59] Katherine L Villa, Kalen P Berry, Jaichandar Subramanian, Jae Won Cha, Won Chan Oh, Hyung-Bae Kwon, Yoshiyuki Kubota, Peter TC So, and Elly Nedivi. Inhibitory synapses are repeatedly assembled and removed at persistent sites in vivo. *Neuron*, 89(4):756–769, 2016.

[60] G Wayne Brodland. How computational models can help unlock biological systems. In *Seminars in cell & developmental biology*, volume 47, pages 62–73. Elsevier, 2015.

[61] Benjamin B Machta, Ricky Chachra, Mark K Transtrum, and James P Sethna. Parameter space compression underlies emergent theories and predictive models. *Science*, 342(6158):604–607, 2013.

[62] Kevin S Brown and James P Sethna. Statistical mechanical approaches to models with many poorly known parameters. *Physical review E*, 68(2):021904, 2003.

[63] Brian K Mannakee, Aaron P Ragsdale, Mark K Transtrum, and Ryan N Gutenkunst. Sloppiness and the geometry of parameter space. In *Uncertainty in Biology*, pages 271–299. Springer, 2016.

[64] Kevin S Brown, Colin C Hill, Guillermo A Calero, Christopher R Myers, Kelvin H Lee, James P Sethna, and Richard A Cerione. The statistical mechanics of complex signaling networks: nerve growth factor signaling. *Physical biology*, 1(3):184, 2004.

[65] Søren L Frederiksen, Karsten W Jacobsen, Kevin S Brown, and James P Sethna. Bayesian ensemble approach to error estimation of interatomic potentials. *Physical review letters*, 93(16):165501, 2004.

[66] Joshua J Waterfall, Fergal P Casey, Ryan N Gutenkunst, Kevin S Brown, Christopher R Myers, Piet W Brouwer, Veit Elser, and James P Sethna. Sloppy-

model universality class and the vandermonde matrix. *Physical review letters*, 97(15):150601, 2006.

[67] Gordon J Berman and Z Jane Wang. Energy-minimizing kinematics in hovering insect flight. *Journal of Fluid Mechanics*, 582:153–168, 2007.

[68] James P Sethna. What are sloppy models? 2008. URL: http://www.lassp.cornell.edu/sethna/ Sloppy/WhatAreSloppyModels.html. Accessed 2020-04-01.

[69] Kenji Mizuseki and György Buzsáki. Preconfigured, skewed distribution of firing rates in the hippocampus and entorhinal cortex. *Cell Reports*, 4(5):1010–1021, 2013.

[70] Michael Okun, Nicholas A. Steinmetz, Lee Cossell, M. Florencia Iacaruso, Ho Ko, Péter Barthó, Tirin Moore, Sonja B. Hofer, Thomas D. Mrsic-Flogel, Matteo Carandini, and Kenneth D. Harris. Diverse coupling of neurons to populations in sensory cortex. *Nature*, 521(7553):511–515, 2015.

[71] John O'keefe and Lynn Nadel. *The hippocampus as a cognitive map*. Oxford: Clarendon Press, 1978.

[72] Yaniv Ziv, Laurie D Burns, Eric D Cocker, Elizabeth O Hamel, Kunal K Ghosh, Lacey J Kitch, Abbas El Gamal, and Mark J Schnitzer. Long-term dynamics of ca1 hippocampal place codes. *Nature neuroscience*, 16(3):264, 2013.

[73] Matthew A Wilson and Bruce L McNaughton. Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172):676–679, 1994.

[74] H Freyja Ólafsdóttir, Daniel Bush, and Caswell Barry. The role of hippocampal replay in memory and planning. *Current Biology*, 28(1):R37–R50, 2018.

[75] Andres D. Grosmark and György Buzsáki. Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science*, 351(6280):1440–1443, 2016.

[76] Adrián Ponce-Alvarez, Gabriela Mochol, Ainhoa Hermoso-Mendizabal, Jaime De La Rocha, and Gustavo Deco. Cortical state transitions and stimulus response evolve along stiff and sloppy parameter dimensions, respectively. *eLife*, 9:e53268, 2020.

[77] James FA Poulet and Sylvain Crochet. The cortical states of wakefulness. *Frontiers in systems neuroscience*, 12:64, 2019.

[78] Marius Pachitariu, Dmitry R Lyamzin, Maneesh Sahani, and Nicholas A Lesica. State-dependent population coding in primary auditory cortex. *Journal of Neuroscience*, 35(5):2058–2073, 2015.

[79] Marie Engelene J Obien, Kosmas Deligkaris, Torsten Bullmann, Douglas J Bakkum, and Urs Frey. Revealing neuronal function through microelectrode array recordings. *Frontiers in neuroscience*, 8:423, 2015.

[80] Christine Grienberger and Arthur Konnerth. Imaging calcium in neurons. *Neuron*, 73(5):862–885, 2012.

[81] Michael Avermann, Christian Tomm, Celine Mateo, Wulfram Gerstner, and Carl CH Petersen. Microcircuits of excitatory and inhibitory neurons in layer 2/3 of mouse barrel cortex. *Journal of neurophysiology*, 107(11):3116–3134, 2012.

[82] Peter Dayan and Laurence F Abbott. *Theoretical neuroscience: computational and mathematical modeling of neural systems*. MIT press, 2001.

[83] Marcel Stimberg, Romain Brette, and Dan FM Goodman. Brian 2, an intuitive and efficient neural simulator. *Elife*, 8, 2019.

[84] Anthony N Burkitt. A review of the integrate-and-fire neuron model: I. homogeneous synaptic input. *Biological cybernetics*, 95(1):1–19, 2006.

[85] Attila I Gulyás, Manuel Megıas, Zsuzsa Emri, and Tamás F Freund. Total number and ratio of excitatory and inhibitory synapses converging onto single interneurons of different types in the ca1 area of the rat hippocampus. *Journal of Neuroscience*, 19(22):10082–10097, 1999.

[86] Hanno S Meyer, Daniel Schwarz, Verena C Wimmer, Arno C Schmitt, Jason ND Kerr, Bert Sakmann, and Moritz Helmstaedter. Inhibitory interneurons in a cortical column form hot zones of inhibition in layers 2 and 5a. *Proceedings of the National Academy of Sciences*, 108(40):16807–16812, 2011.

[87] György Buzsáki and Kenji Mizuseki. The log-dynamic brain: How skewed distributions affect network operations. *Nature Reviews Neuroscience*, 15(4):264–278, 2014.

[88] 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.

[89] E.B. Goldstein. *Sensation and Perception*. Cengage Learning, 2009.

[90] David H Hubel and Torsten N Wiesel. Receptive fields of single neurones in the cat's striate cortex. *The Journal of physiology*, 148(3):574, 1959.

[91] David C Somers, Sacha B Nelson, and Mriganka Sur. An emergent model of orientation selectivity in cat visual cortical simple cells. *Journal of Neuroscience*, 15(8):5448–5465, 1995.

[92] Rodney J Douglas, Kevan AC Martin, and David Whitteridge. A canonical microcircuit for neocortex. *Neural computation*, 1(4):480–488, 1989.

[93] David Ferster and Kenneth D Miller. Neural mechanisms of orientation selectivity in the visual cortex. *Annual review of neuroscience*, 23(1):441–471, 2000.

[94] Nicholas J Priebe. Mechanisms of orientation selectivity in the primary visual cortex. *Annual review of vision science*, 2:85–107, 2016.

[95] David H Hubel and Torsten N Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, 160(1):106, 1962.

[96] Paul C. Bressloff, Neil W Bressloff, and Jack D. Cowan. Dynamical mechanism for sharp orientation tuning in an integrate-and-fire model of a cortical hypercolumn. *Neural computation*, 12(11):2473–2511, 2000.

[97] Peggy Seriès, Peter E Latham, and Alexandre Pouget. Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. *Nature neuroscience*, 7(10):1129–1135, 2004.

[98] David Hansel and Carl van Vreeswijk. The mechanism of orientation selectivity in primary visual cortex without a functional map. *Journal of Neuroscience*, 32(12):4049–4064, 2012.

[99] Sadra Sadeh, Claudia Clopath, and Stefan Rotter. Processing of feature selectivity in cortical networks with specific connectivity. *PloS one*, 10(6):e0127547, 2015.

[100] William P Jones and George W Furnas. Pictures of relevance: A geometric analysis of similarity measures. *Journal of the American society for information science*, 38(6):420–442, 1987.

[101] Benjamin Scholl, Andrew YY Tan, Joseph Corey, and Nicholas J Priebe. Emergence of orientation selectivity in the mammalian visual pathway. *Journal of Neuroscience*, 33(26):10616–10624, 2013.

[102] Michael E. Rule, Adrianna R. Loback, Dhruva V. Raman, Laura N. Driscoll, Christopher D. Harvey, and Timothy O'Leary. Stable task information from an unstable neural population. *Elife*, page 9:e51121, 2020.

[103] Nicolas Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of computational neuroscience*, 8(3):183–208, 2000.

[104] Luiz Pessoa. Understanding brain networks and brain organization. *Physics of life reviews*, 11(3):400–435, 2014.

[105] Jonathan Cannon and Paul Miller. Stable control of firing rate mean and variance by dual homeostatic mechanisms. *The Journal of Mathematical Neuroscience*, 7(1):1, 2017.

[106] Gina Turrigiano. Homeostatic synaptic plasticity: Local and global mechanisms for stabilizing neuronal function. *Cold Spring Harbor Perspectives in Biology*, 4(1):1–18, 2012.

[107] Qingming Hou, Dawei Zhang, Larissa Jarzylo, Richard L Huganir, and Heng-Ye Man. Homeostatic regulation of ampa receptor expression at single hippocampal synapses. *Proceedings of the National Academy of Sciences*, 105(2):775–780, 2008.

[108] Alan L Hodgkin and Andrew F Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, 117(4):500, 1952.

[109] Wilfrid Rall. *The theoretical foundation of dendritic function: selected papers of Wilfrid Rall with commentaries*. MIT press, 1995.

[110] Stephen R Williams and Greg J Stuart. Dependence of epsp efficacy on synapse location in neocortical pyramidal neurons. *Science*, 295(5561):1907–1910, 2002.

[111] Richard H Masland. Neuronal cell types. *Current Biology*, 14(13):R497–R500, 2004.

[112] Allen Institute for Brain Science. Allen cell types database. 2016. http://celltypes.brain-map.org/.

[113] Corinne Teeter, Ramakrishnan Iyer, Vilas Menon, Nathan Gouwens, David Feng, Jim Berg, Aaron Szafer, Nicholas Cain, Hongkui Zeng, Michael Hawrylycz, et al. Generalized leaky integrate-and-fire models classify multiple neuron types. *Nature communications*, 9(1):1–15, 2018.

[114] Stefano Cavallari, Stefano Panzeri, and Alberto Mazzoni. Comparison of the dynamics of neural interactions between current-based and conductance-based integrate-and-fire recurrent networks. *Frontiers in neural circuits*, 8:12, 2014.

[115] Gina G Turrigiano. The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell*, 135(3):422–435, 2008.

[116] Arianna Maffei and Alfredo Fontanini. Network homeostasis: a matter of coordination. *Current Opinion in Neurobiology*, 19(2):168–173, 2009.

[117] Hey-Kyoung Lee and Alfredo Kirkwood. Mechanisms of homeostatic synaptic plasticity in vivo. *Frontiers in Cellular Neuroscience*, 13:520, 2019.

[118] Nathalia Vitureira, Mathieu Letellier, and Yukiko Goda. Homeostatic synaptic plasticity: from single synapses to neural circuits. *Current opinion in neurobiology*, 22(3):516–521, 2012.

[119] Keiji Ibata, Qian Sun, and Gina G Turrigiano. Rapid synaptic scaling induced by changes in postsynaptic firing. *Neuron*, 57(6):819–826, 2008.

[120] Leo Egghe and Loet Leydesdorff. The relation between pearson's correlation coefficient r and salton's cosine measure. *Journal of the American Society for information Science and Technology*, 60(5):1027–1036, 2009.

Appendix A

Model Parameters

A.1 Single-Cell Parameters

single-cell parameters (E and I)		
v_{rest}	0 mV	resting potential
θ	33 mV	spike emission threshold
v_{reset}	25.75 mV	reset potential
τ_m	10 ms	membrane time constant
τ_{arp}	1 ms	absolute refractory period
τ_{delay}	0.01ms	synaptic delay

Table A.1: Single-cell parameters for the LIF neurons in both the random network model and the ring model. Parameters are the same as in Mongillo et al [10]. The high value of the reset potential is somewhat unusual, but was kept to increase the comparability of our results with those of Mongillo et al [10]. The high reset parameter had no effect on the nature of our results. To confirm this, I ran some of our experiments with an altered version of our model with v_{reset} = 0 mV, which produced the same results.

A.2 Network Parameters

network parameters (random network)		
N_E	32,000	number of E neurons
N_I	8,000	number of I neurons
H_E^{ext}	72.6 mV	external excitatory input to E neurons
H_I^{ext}	57.8 mV	external excitatory input to I neurons
C_{EE}	0.2	probability of E $\rightarrow$ E connection
C_{EI}	0.4	probability of E $\rightarrow$ I connection
C_{IE}	0.3	probability of I $\rightarrow$ E connection
C_{II}	0.4	probability of I $\rightarrow$ I connection
network parameters (ring model)		
N_E	3,200	number of E neurons
N_I	800	number of I neurons
H_E^{ext}	33 mV	external excitatory input to E neurons
H_I^{ext}	26.4 mV	external excitatory input to I neurons
$\bar{r}_E^{LGN}$	100Hz	rate of average LGN inputs to E neurons
r_I^{LGN}	80Hz	rate of LGN inputs to I neurons
w_{LGN}	0.1 mV	synaptic efficacy of input synapses from the LGN
η_{LGN} (vers.1)	1	tuning of LGN inputs for model version 1
η_{LGN} (vers.2)	0.2	tuning of LGN inputs for model version 2
η_{FS} (vers.1)	0	tuning of feature-specific EE connections for model version 1
η_{FS} (vers.2)	0.5	tuning of feature-specific EE connections for model version 1
C_{EE}	0.2	average probability of E $\rightarrow$ E connection
C_{EI}	0.4	probability of E $\rightarrow$ I connection
C_{IE}	0.3	probability of I $\rightarrow$ E connection
C_{II}	0.4	probability of I $\rightarrow$ I connection

Table A.2: Network parameters for the random network and ring models. Parameters for the random network model are the same as in Mongillo et al [10], except that the input H_E^{ext} to the excitatory neurons is slightly lower. This was done to make the simulation more stable over longer run times, as the network with the original input values tended to collapse after about 3s. Parameters for the ring model were adapted to accommodate the smaller network size and the additional LGN inputs.

A.3 Synaptic Weights

synaptic efficacies (both models)		
W_{EE}	0.37 mV	mean E $\rightarrow$ E synaptic efficacy
W_{EI}	0.66 mV	mean E $\rightarrow$ I synaptic efficacy
W_{IE}	0.44 mV	mean I $\rightarrow$ E synaptic efficacy
W_{II}	0.54 mV	mean I $\rightarrow$ I synaptic efficacy
μ_{EE}	-13.9 mV	mu E $\rightarrow$ E synaptic efficacy
μ_{EI}	-15.69 mV	mu E $\rightarrow$ I synaptic efficacy
μ_{IE}	7.16 mV	mu I $\rightarrow$ E synaptic efficacy
μ_{II}	-12.9 mV	mu I $\rightarrow$ I synaptic efficacy
σ_{EE}	26.07 mV	sigma E $\rightarrow$ E synaptic efficacy
σ_{EI}	20.79 mV	sigma E $\rightarrow$ I synaptic efficacy
σ_{IE}	31.68 mV	sigma I $\rightarrow$ E synaptic efficacy
σ_{II}	25.74 mV	sigma I $\rightarrow$ I synaptic efficacy

Table A.3: Parameters for the lognormal distribution of synaptic weights in the random network model and ring models. Parameters are the same as in Mongillo et al [10].

Appendix B

Supplementary Figures

B.1 Cosine Similarity vs Correlation Coefficient

The purpose of the Supplementary Figure B.1 below is to illustrate why, in contrast to Mongillo et al [10], we decided to use cosine similarity instead of Pearson's correlation coefficient to quantify the effect of synaptic rewiring on the population activity.

Because the mean and the distribution of the firing rate vectors do not change due to rewiring, computing the correlation coefficient between the two firing rate vectors is equivalent to computing the cosine similarity of the vectors after their arithmetic mean has been subtracted [100,120]. This means that the correlation coefficient places larger emphasis on the absolute changes of firing rates of the individual neurons, while the cosine similarity considers relative changes. This is preferable in our case, because the cosine similarity allows a more meaningful comparison of the effect of synaptic rewiring on populations with different means.

The two measures overall showed the same trends in all our experiments. The only difference is that the correlation coefficient indicates a larger effect of II rewiring on the firing rates of the inhibitory population (Figure B.1A). This is mainly because the inhibitory neurons in our model have a larger mean firing rate than excitatory neurons, which permits larger fluctuations in neuronal activity per se. When considering the neuron's firing rates relative to the population mean (Figure B.1B), it becomes apparent that II rewiring in fact does not seem to affect the inhibitory population much more than the excitatory population.

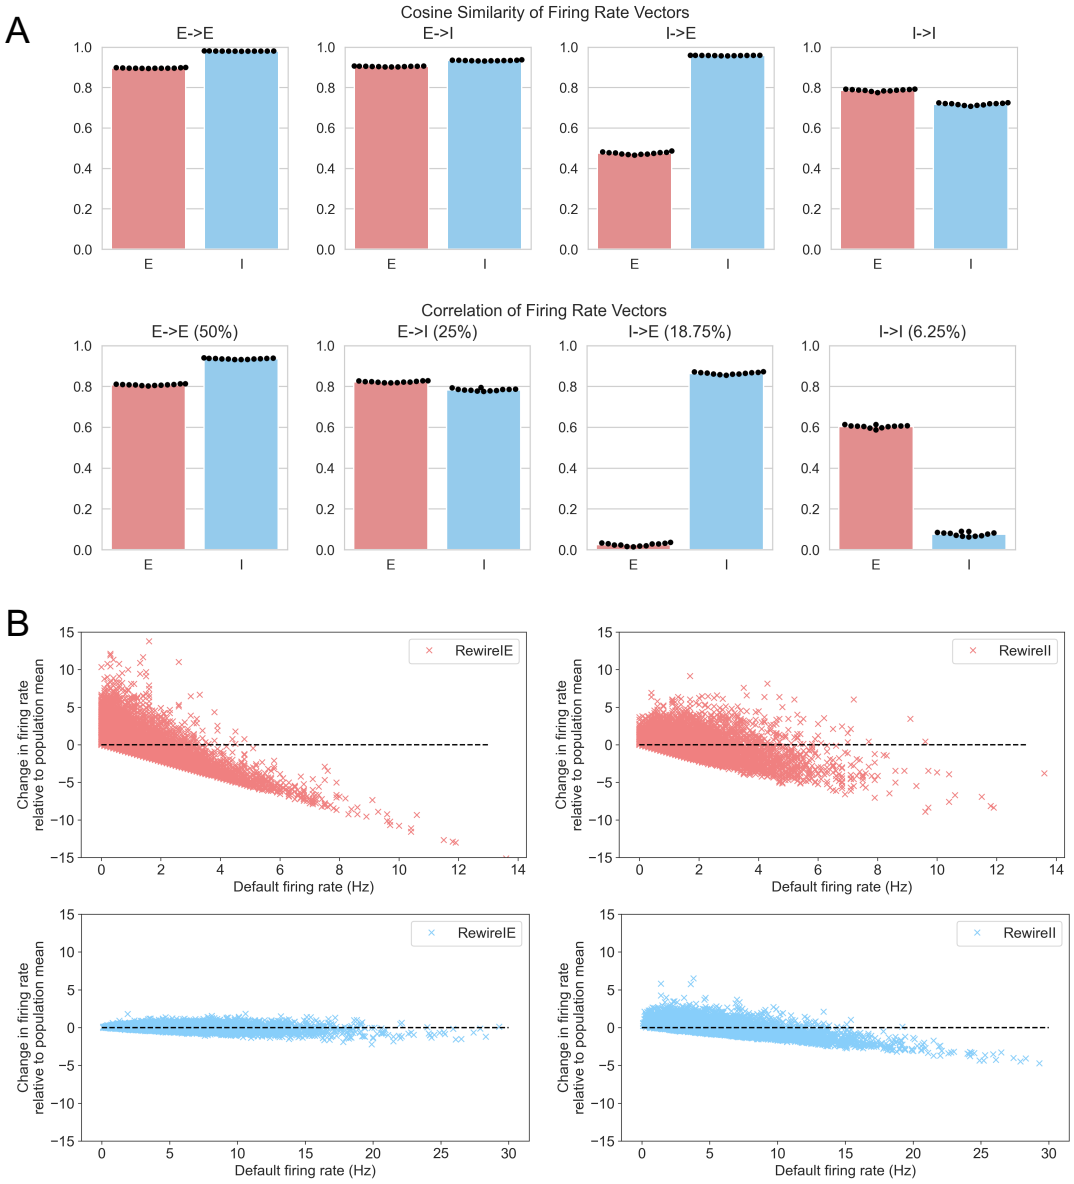

Figure B.1: A: Firing rate changes due to rewiring measured by the cosine similarity (top) and Pearson's correlation coefficient (bottom). Figure corresponds to the experiment described in Section 4.1.1. B: Relative firing rate changes of individual excitatory (red) and inhibitory (blue) neurons as a function of their firing rates under default conditions. Although the correlation coefficient (A, bottom) suggests a strong absolute effect of II rewiring on the inhibitory population, the effect is much weaker when considered in relation to the mean firing rate of the inhibitory population (B). We therefore decided to use cosine similarity instead of the correlation coefficient as a measure of population firing rate changes.

B.2 Effects of II Synaptic Rewiring on the Random Network Model

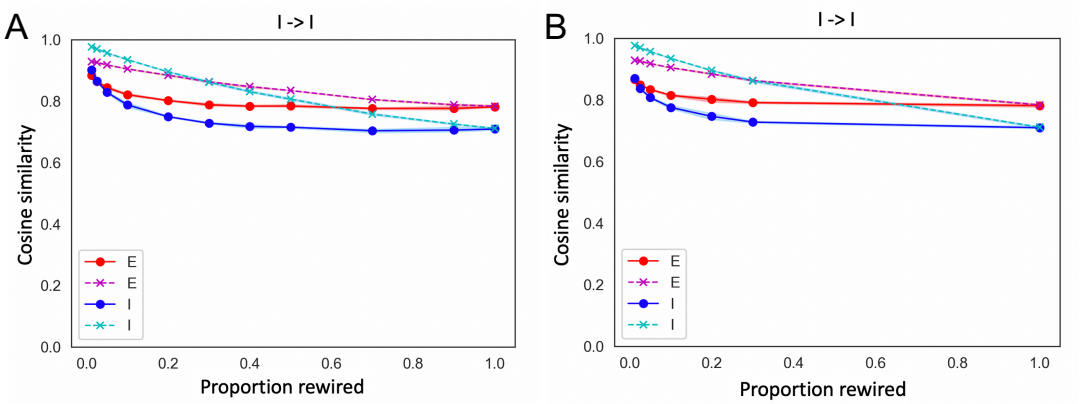

Figure B.2: Effects of rewiring II synapses by weight (A, solid lines) and synaptic impact (B, solid lines) compared to rewiring the same proportion of randomly selected II synapses. Effect on the excitatory population is in red, effect on inhibitory activity is in blue. Figures correspond to Figures 4.5B and 4.6 respectively.

B.3 Effects of Synaptic Rewiring on the Ring Model Version 2

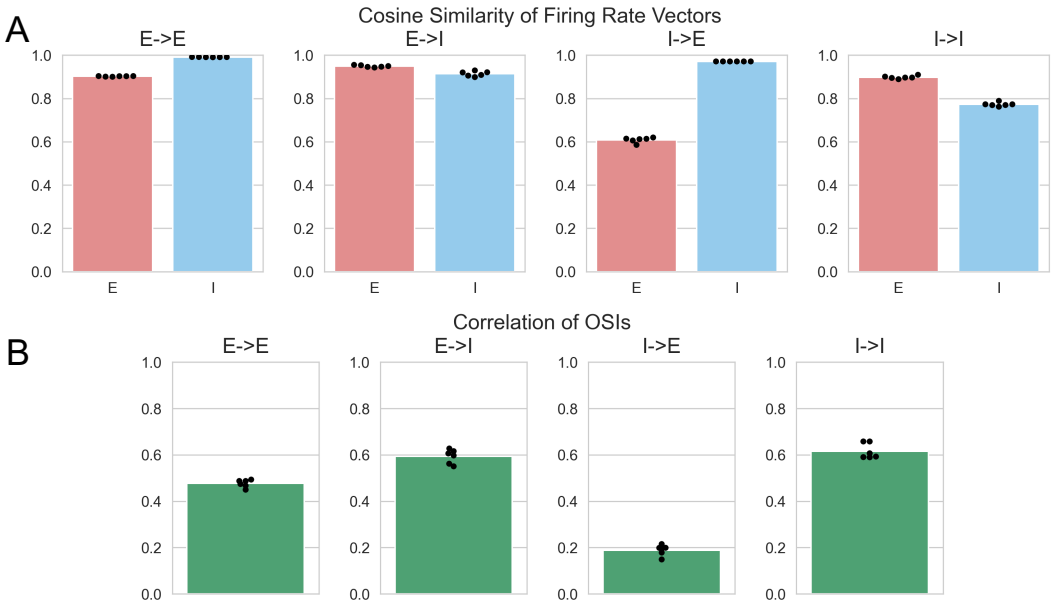

Figure B.3: Effect of synaptic rewiring the ring model version 2. A: Cosine similarities of the excitatory (red) and inhibitory (blue) average firing rate vectors before and after rewiring different synaptic populations. The average was taken over the network's response to 16 different stimuli. N = 6. B: Correlation of excitatory neuron OSIs before and after rewiring. N = 6.

B.4 Effect of IE Rewiring on Decoder Variance

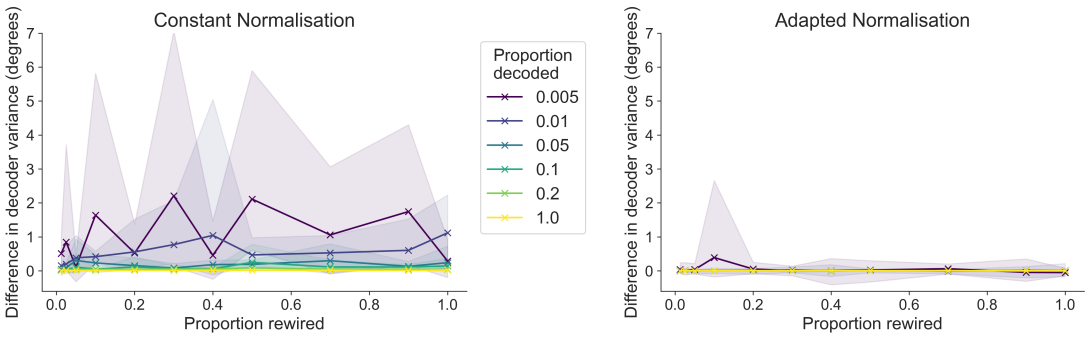

Figure B.4: Difference in decoder variance before and after rewiring different proportions of IE synapses. On the left, the neuron's mean firing rates before rewiring were used to normalise both the firing rate vector before and after rewiring. On the right, the normalisation was adapted and the neuron's firing rates after rewiring were normalised with their means after rewiring. N = 6. Markers: mean; shading: range.