THE UNIVERSITY
of EDINBURGH

Functional relevance of inhibitory and disinhibitory circuits in signal propagation in recurrent neuronal networks

Marzena Maria Bihun

Doctor of Philosophy
Institute for Adaptive and Neural Computation
School of Informatics
University of Edinburgh
2017

Abstract

Cell assemblies are considered to be physiological as well as functional units in the brain. A repetitive and stereotypical sequential activation of many neurons was observed, but the mechanisms underlying it are not well understood. Feedforward networks, such as synfire chains, with the pools of excitatory neurons unidirectionally connected and facilitating signal transmission in a cascade-like fashion were proposed to model such sequential activity. When embedded in a recurrent network, these were shown to destabilise the whole network's activity, challenging the suitability of the model. Here, we investigate a feedforward chain of excitatory pools enriched by inhibitory pools that provide disynaptic feedforward inhibition. We show that when embedded in a recurrent network of spiking neurons, such an augmented chain is capable of robust signal propagation. We then investigate the influence of overlapping two chains on the signal transmission as well as the stability of the host network. While shared excitatory pools turn out to be detrimental to global stability, inhibitory overlap implicitly realises the motif of lateral inhibition, which, if moderate, maintains the stability but if substantial, it silences the whole network activity including the signal. Addition of a disinhibitory pathway along the chain proves to rescue the signal transmission by transforming a strong inhibitory wave into a disinhibitory one, which specifically guards the excitatory pools from receiving excessive inhibition and thereby allowing them to remain responsive to the forthcoming activation. Disinhibitory circuits not only improve the signal transmission, but can also control it via a gating mechanism. We demonstrate that by manipulating a firing threshold of the disinhibitory neurons, the signal transmission can be enabled or completely blocked. This mechanism corresponds to cholinergic modulation, which was shown to be signalled by volume as well as phasic transmission and variably target classes of neurons. Furthermore, we show that modulation of the feedforward inhibition circuit can promote generating spontaneous replay at the absence of external inputs. This mechanism, however, tends to also cause global instabilities.

Overall, these results underscore the importance of inhibitory neuron populations in controlling signal propagation in cell assemblies as well as global stability. Specific inhibitory circuits, when controlled by neuromodulatory systems, can robustly guide or block the signals and invoke replay. This mounts to evidence that the population of interneurons is diverse and can be best categorised by neurons' specific circuit functions as well as their responsiveness to neuromodulators.

Acknowledgements

First and foremost, I wish to wholeheartedly thank my primary supervisor, Dr. Matthias Hennig. Completing this thesis would have not been possible without his unending patience, tremendous academic support and the most valuable guidance. I also want to thank the second supervisor, Dr. Emma Wood as well as Dr. Mark van Rossum for their immensely useful advice throughout the thesis.

Many thanks to the Doctoral Training Centre in Neuroinformatics (DTC) for providing an inspiring and unique environment.

Finally, I wish to thank my family which, although far away, would always encourage and support me in so many ways.

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.

(*Marzena Maria Bihun*)

Table of Contents

1 Introduction **1**

1.1 Structure of this thesis 3

2 Background **4**

2.1 Cell assemblies in the brain 4

2.1.1 Theoretical models of cell assemblies 5

2.1.2 Experimental evidence for sequential activity 6

2.1.3 Difficulties with detecting sequences 7

2.2 Computational models of feedforward networks 8

2.2.1 Spiking balanced networks 8

2.2.2 A generic model of feedforward networks 10

2.2.3 Two modes of propagation 11

2.2.4 Chain–rest of the network relationship 13

2.2.5 Current vs conductance-based synapses 15

2.2.6 Recruitment of inhibition 16

2.2.7 Gating mechanisms 17

2.2.8 Summary of the models and proposed augmentations 18

2.3 Neural inhibition 21

2.3.1 Beyond balancing excitation 21

2.3.2 A brief history of inhibition 22

2.3.3 Classification of interneurons 23

2.3.4 Inhibitory circuits 25

2.3.5 Inhibition as a better modulator than excitation 27

2.4 Disinhibition 28

2.4.1 Disinhibitory circuit 28

2.4.2 Disinhibitory circuit as a modulator 29

2.4.3 Disinhibitory pathways found in brain 29

2.5 Neuromodulatory cholinergic system 34
2.5.1 Cholinergic signalling 35
2.5.2 Action of Acetylcholine 36
2.5.3 Acetylcholine and disinhibition 37
2.5.4 Computational models of cholinergic modulation 38
2.5.5 Summary 39

3 Methods **40**
3.1 Introduction 40
3.2 Neuron model 41
3.2.1 Synapses 41
3.3 Network characteristics 43
3.4 Feedforward chain characteristics 44
3.4.1 Embedding the chain in the network 45
3.4.2 The nature of overlaps 45
3.5 Input injection 46
3.6 Generating data via simulations 47
3.7 Simulation tools 47

4 Signal propagation in overlapping chains with disinhibitory pathways **48**
4.1 One embedded chain 48
4.2 Definitions 50
4.2.1 Excitatory and inhibitory haloes 52
4.3 Chains with disinhibitory pathways 53
4.4 Embedding two non-overlapping chains 57
4.5 Embedding two overlapping chains 58
4.5.1 Activity of the rest of the network 61
4.6 Evolution of a pulse packet along the chain 63
4.6.1 State space analysis for basic and guarded chain conditions . . 65
4.6.2 The shapes of trajectories for all combinations of overlaps . . 66
4.6.3 Successful trajectories 66
4.6.4 Halo failure trajectories 69
4.6.5 Explosion failure trajectories 70
4.7 Can one predict pulse packet's fate? 72
4.8 Injecting input to two chains simultaneously 75
4.9 Splitting the inhibitory overlap 76

4.9.1 Specific pool allocation 76
4.9.2 Specific pool allocation with separate overlap levels 78
4.10 Conclusions 80

5 Signal gating by cholinergic modulation of disinhibitory pathways 82
5.1 Introduction 82
5.1.1 Control of disinhibition as a gating mechanism 83
5.1.2 Default state of a gate 84
5.2 The actual contribution of disinhibitory pathways 85
5.3 Methods 86
5.3.1 Modification of disinhibitory pathways 87
5.3.2 Choosing the levels of overlaps 88
5.4 Cholinergic signalling via volume transmission 91
5.4.1 Location-specific volume transmission 92
5.5 Cholinergic signalling via phasic transmission 94
5.6 Volume and phasic transmission comparison 96
5.7 Discussion 97

6 Cholinergic modulation of feedforward inhibition to invoke replay 98
6.1 Introduction 98
6.2 Methods 100
6.3 Cholinergic modulation of feedforward inhibition 101
6.4 Synfire chain explosions 102
6.5 Modes of modulation to avoid the instabilities 104
6.5.1 Modulation targeting a subset of layers 104
6.5.2 Modulation of a fraction of neurons in pools 105
6.6 Overlapping chains 106
6.7 Discussion 107
6.7.1 Two types of inhibition 107

7 General discussion 109
7.1 Summary of the thesis 109
7.2 Relevance of the presented models 110
7.3 Inhibitory overlaps between cell assemblies 111
7.4 Should cortical activity be dominated by inhibition? 113
7.5 The significance of disinhibition 114

7.5.1 Can disinhibition be equivalent to increased excitation? 115
7.6 Mapping neurons to functions? 116
7.6.1 Sampling problem 116
7.6.2 Neuron vs network doctrine 118
7.6.3 Building blocks of networks 119
7.7 Biological realism of interneuron classes in spiking networks 120
7.8 Neuromodulation 121
7.9 Back to cell assemblies – how can they be created? 122

Bibliography **124**

List of Tables

2.1 Comparison between models extending the synfire chain architecture . 20
2.2 Selected disinhibitory circuits found in the Central Nervous System . 33

3.1 Values used in the network model. 42

4.1 Definitions of the terms introduced in Chapter 4. 51

7.1 An overview of sample sizes in selected interneuron studies 118

Chapter 1

Introduction

Sequential activation of neurons has emerged as a ubiquitous feature of network activity. Repetitive sequences have been detected in many brain regions including hippocampus and cortex and during various behaviours such as working memory or decision making (Pastalkova et al., 2008; Harvey et al., 2012). Elucidating the theoretical underpinnings of these is an important issue in contemporary neuroscience and will be investigated in computational models in this thesis.

Groups of interconnected and coactivated neurons are commonly conceptualised as *cell assemblies* which, unlike the individual neurons, are considered to be the physiological and functional units in the brain (Yuste, 2015). Over the decades, a great wealth of theoretical models of cell assemblies emerged. The ones with the feedforward architectures aim to model the sequential activation and are the main topic of this thesis. A classical model of feedforward networks – a synfire chain – is composed of exclusively excitatory neurons arranged in unidirectionally connected layers (Abeles, 1991). To date, numerous extensions of synfire chains have been thoroughly studied in spiking networks and several limitations of these were identified (Mehring et al., 2003). Such models are commonly criticised for their stereotypical activity and the ability to rapidly synchronise and destabilise the whole network in which they are embedded if their parameters are not carefully controlled. It is also often claimed that such highly ordered architectures lack biological realism (Rajan et al., 2016).

Although the great majority of neurons in the brain are excitatory and these are the actual carriers of the signals, the importance of inhibition should by no means be deprecated. A solid and constantly growing body of evidence reveals that the role of inhibition in regulating network activity is immense and often critical (Tremblay et al., 2016). Depending on the connectivity patterns between the cells, inhibitory

neurons can control network's oscillations, gate the signals in excitatory populations, control the temporal precision of spikes, or normalise the activity of groups of neurons (Roux and Buzsáki, 2015). Many contemporary studies explore the non-overlapping, genetically-defined subpopulations of inhibitory neurons which appear to be highly specialised for different, well-defined tasks in the microcircuits (Rudy et al., 2011).

We argue that there is a considerable discrepancy between the theoretical modelling of feedforward networks and the recent experimental findings. Theoretical models appear to lag behind the current trend underscoring the importance of neural inhibition. Synfire chains composed of exclusively excitatory neurons are a prime example of it.

In this thesis, we seek to bring more of biological realism into the theoretical models of feedforward networks in the hope of mitigating the known limitations. Firstly, we embed synfire chains in recurrent networks and secondly, we aim to incorporate stereotypical inhibitory circuits into the chain's architecture and systematically study their influence on the signal propagation and network stability. We focus on lateral and disynaptic feedforward inhibition as well as disinhibition. Furthermore, we aim to identify possible mechanisms of gating the signals transmitted by the excitatory populations, so that the signal propagation can be flexibly controlled by a 'switch'. In search of such switch we explore different modes of neuromodulation targeting exclusively inhibitory circuits.

Cholinergic modulation will be investigated as an important source of neuromodulation. Acetylcholine was shown to be crucially implicated in attention, memory, coordination of behavioural state and arousal (Hasselmo, 2006). Numerous studies revealed that the inhibitory neurons richly express cholinergic receptors and acetylcholine was found to directly control the disinhibitory mechanisms in the cortex (Letzkus et al., 2011). Moreover, cholinergic transmission was found to operate on many time scales – from hours to milliseconds – offering a whole range of potential modulatory modes (Sarter et al., 2009).

Cholinergic modulation is also involved in invoking spontaneous replay – a sequential activation that is internally generated at the absence of input (Gais and Born, 2004). Replay is observed during animal's sleep and quiet wakefulness but never during the fully awake state (Carr et al., 2011). Noteworthy, it was shown that the awake and sleep states differ considerably in the levels of acetylcholine available in the extracellular space (Lee and Dan, 2012). Thus, it is hypothesised that the level of cholinergic modulation can realise a switch between the awake and sleep state and influence the circuits in such a way that the spontaneous replay emerges only in the sleep state.

1.1 Structure of this thesis

This thesis is structured in the following way. Chapter 2 contains the literature review, which firstly discusses cell assemblies, their feedforward subcategory and the experimental evidence for such assemblies. Then, the theoretical modelling of the neural networks is discussed, with emphasis on the features of the already existing computational models of the feedforward networks followed by a critical analysis and a list of suggested improvements inspired by the physiology of inhibitory neurons. Then, the selected aspects of neural inhibition, disinhibition and cholinergic modulation are discussed, as they are crucial to justifying the choices made in the modelling work as well as in evaluating the overall significance and novelty of the results. Chapter 3 outlines the methods that were used in all the variants of the model presented in this thesis. Chapter 4 presents the data and analysis obtained from the simulations of the feedforward networks embedded in a spiking balanced network with the addition of the disinhibitory pathways and sharing the pools of neurons to various degrees. We show that the level and type of overlaps between the feedforward chains significantly influence the signal propagation as well as the stability of the whole host network. We also show the functional role of the disinhibitory pathways along the chains. Chapter 5 investigates the volume and phasic modes of cholinergic transmission and their role in modulating disinhibitory pathways to control the signal propagation along the feedforward chains. Modulation of pools of inhibitory neurons controlling feedforward inhibition is described in Chapter 6. We show that disabling the feedforward inhibition enables weak and essentially random signals to invoke spontaneous replay along the chain. Lastly, Chapter 7 discusses the results and their relevance to the current state of knowledge and points at the future avenues that are worth pursuing in order to advance the knowledge about the role of inhibitory pathways in cell assemblies and in cortical networks.

Chapter 2

Background

2.1 Cell assemblies in the brain

One of the main quests in neurobiology involves understanding of how the memory is being represented and managed in the brain. To date, experimentalists and theoreticians have unravelled a great wealth of mechanisms governing memory formation and retrieval from subcellular to network level (Poo et al., 2016). Noteworthy, many of the contemporary discoveries still give credit to two old theories that surfaced back in times when even the notion of neural inhibition was still unknown[1].

Engram theory – as formulated by Richard Semon (1921) – postulated that any memory is physically stored in brain and represented by a coactivation of a group of neurons termed an *engram*. Semon aptly predicted that learning induces persistent physical/chemical changes in these neurons and subsequent partial reactivation of these may lead to a full memory retrieval – a phenomenon that nowadays is known as *pattern completion* (Tonegawa et al., 2015).

In similar vein, Donald Hebb in his seminal work introduced a term *cell assembly* to denote a group of strongly interconnected neurons that represent a distinct cognitive entity and proposed a mechanism based on synaptic plasticity as a substrate of memory formation (Hebb, 1949). Most, if not all contemporary plasticity rules, stem from the Hebbian learning rule, which, overly simplified, asserts that *cells that fire together, wire together*. Besides neuroscience, Hebbian theory also immensely contributed to the fields of artificial neural networks and machine learning.

Nowadays, the term *cell assembly* (as well as cell/neural ensemble) is widely used to describe spatiotemporal orchestration of activity of interconnected neurons and pro-

[1]The discovery of inhibition is discussed in Section 2.3.2

vides a useful conceptual framework to study integration of distributed neuronal activity (Varela et al., 2001). Moreover, the cell assemblies, not the individual neurons, are now believed to be the physiological units of the brain which generate and sustain the functional properties as well as the dynamical states of the entire system (Yuste, 2015; Buzsáki, 2010).

2.1.1 Theoretical models of cell assemblies

Over the decades, numerous models of cell assemblies have been proposed. Based on their architecture, these can be grouped into two general types: *feedforward* networks, which are predominantly governed by the one-way connections, and *recurrent* networks, also called attractor-based models, which essentially rely on a positive feedback to sustain their activity (Yuste, 2015).

The attractor-based architectures are well-studied and frequently employed to solve problems involving working memory, path integration, and the head direction cells system, to list a few (Durstewitz et al., 2000; Samsonovich and McNaughton, 1997; Redish et al., 1996). On the other hand, certain limitations of such models were identified, especially the need for the fine-tuning of the parameters as well as low variability in neuronal response properties (Shafi et al., 2007; Lim and Goldman, 2012). The neuronal variability is commonly measured by the coefficient of variation C_V of the neuron's interspike intervals distribution and defined by the ratio of the standard deviation to the mean of this distribution. Cortical neurons were shown to have high C_V values implying that their firing patterns are very irregular, a property the attractor-based models fail to reproduce (Softky and Koch, 1993). In the context of the working memory, however, one solution to alleviate the variability problem involved the inclusion of short-term plasticity to the model (Hansel and Mato, 2013).

Finally, it should be clearly stressed, that these general types are actually the two extremes along a continuum. A fully recurrent network, such as a Hopfield network, can be initiated in many ways, but the resultant dynamics will always push the network to enter a stable state and remain there infinitely. In contrast, fully feedforward networks lack such fixed points and their activity keeps transitioning from one state to another. In between, however, there exist mixed models which enable sequential transitions from one fixed point to another (Kleinfeld, 1986; Sompolinsky and Kanter, 1986), or fully recurrent models which functionally in fact behave akin to the purely feedforward networks (Goldman, 2009).

No attempts will be made to directly compare the two architectures and evaluate their performance. Throughout the whole thesis only the former – feedforward networks models – will be under thorough investigation, whereas the purely recurrent ones will be considered as out of scope.

2.1.2 Experimental evidence for sequential activity

The feedforward networks are by no means merely the abstract, theoretical creations as there is ample evidence for the transient sequential activity found in many brain regions across various species. Neural sequences are now considered to be a common feature of the network activity during a whole range of behaviours. Below, the prominent examples of such sequences are briefly reviewed, followed by a commentary concerning the general difficulties with detecting such activity in neuronal networks.

Birdsong

Birdsong, commonly studied on the zebra finch, is one of the most well-understood sequential activity found in brain and is characterised by the extremely precise spiking patterns (Hahnloser et al., 2002). The songs are internally generated in the high vocal centre (HVC) – an analogue of the mammalian premotor cortex – where the principal neurons produce sparse, time-locked bursts of activity that are stereotyped from trial to trial and each neuron typically produces only one such burst during a song. Subsequent studies revealed that inhibitory interneurons' activity is stereotyped as well (Markowitz et al., 2015).

Hippocampal and cortical sequences

Hippocampal place cells, which encode particular regions in space, produce sequences while the animal explores the environment, so that based on the neurons' firing patterns it is possible to reproduce the animal's trajectory (O'Keefe and Dostrovsky, 1971). Such sequences are believed to rely on the external inputs such as environmental or self-motion cues as they are activated during the behaviour. In contrast, after the experience and typically during sleep or rest periods, such sequences reappear in a rapid, compressed way (Buzsáki, 2010). Such reactivation of the previously experienced trajectories was termed as a spontaneous or offline replay (Carr et al., 2011) and it was shown to appear in either forward or reverse fashion (Diba and Buzsáki, 2007). On top, a preplay – an activation of a future trajectory – was also reported in animals that were

resting before the exploration of a novel environment (Dragoi and Tonegawa, 2011). Neurons termed time cells that encode the time course, rather than the animal's location during a task were also shown to fire in a sequential fashion (MacDonald et al., 2011). Lastly, yet another group of neurons produced internally-generated sequential activations that were related to the future choices of a behaving animal (Pastalkova et al., 2008). All these examples demonstrate that the sequential activation is ubiquitous and comes in two forms – either it is stimulus-driven or internally-generated.

Sequential activation is by no means limited to the hippocampal system. In rodent posterior parietal cortex and visual cortices various choice or stimulus-specific sequential patterns were also reported (Harvey et al., 2012; Luongo et al., 2016; Sadovsky and MacLean, 2014; Carrillo-Reid et al., 2015). In the rat medial prefrontal cortex, on the other hand, an offline replay of sequential motifs related to the previously experienced rule learning sessions was observed during animal's sleep (Peyrache et al., 2009).

2.1.3 Difficulties with detecting sequences

For the sake of completeness, it should be remarked that finding solid evidence for the sequential cell assembly activation is fraught with difficulties. Although the cell assemblies theories were postulated many decades ago, for years the lack of appropriate technology to record many neurons simultaneously in behaving animals hindered the discovery. Nowadays not only hundreds or even thousands of neurons can be recorded at a time (Stevenson and Kording, 2011), but also individual neurons or groups of such can be manipulated by virtue of the optogenetic techniques (Buzsáki, 2004; Deisseroth, 2011). The experimental techniques, however, is not the only limitation in finding the true sequential activation. Detecting robust sequences in non-categorised datasets of spike trains is not a trivial task and it heavily relies on complex statistical methods and hypotheses testing.

A publication reporting the detection of cortical songs (a counterpart of the birdsong) attracted a lot of attention as the authors boldly claimed to have found precise repetitions of sequential activation in the visual cortex (Ikegaya et al., 2004). Based on the motifs found in vitro and composed of individual cells' repetitive firings, they developed an indirect method to search for similar motifs in vivo. These were indeed found in the visual cortex of an anaesthetised cat and it was concluded that sequential activation can be internally generated there. Later, the flaws in the detection method were identified (Roxin et al., 2008; Mokeichev et al., 2007), and it was argued that the

detected motifs in vivo can in fact occur by chance. Thus, this report shall not be used as a piece of evidence for the actual sequential activation in the cortex.

The discovery of the hippocampal preplay has also been greeted with a deal of scepticism (Dragoi and Tonegawa, 2011). A subsequent study attempted to reproduce the results and its conclusions rejected the idea of preplay in favour of the sequence forming exclusively during the actual experience (Silva et al., 2015). Does it mean that the methodology in the original study was again faulty and the very idea of preplay downright wrong? Not necessarily. One of the clear differences between the studies is that in the first one, the 'novel' environment was actually an extension of a familiar track, which could possibly influence the 'reuse' of the already existing representations (Eichenbaum, 2015). The broader question here is whether and how the pre-existing structures influence the new encoding. Another study indeed identified two types of neurons – the rigid and plastic ones – that are differentially prone to undergoing modifications during the novel experience (Grosmark and Buzsáki, 2016), suggesting that the observed preplay might have depended on such rigid cells. The problem of cell assembly creation and modification will be revisited later.

In conclusion, although some reports remain controversial, the majority of accounts of sequential activation are still based on sound methodologies and imply that the sequential activation indeed robustly occurs in various brain regions and can be regarded as a common feature of the neuronal networks. The following sections will review a selection of the feedforward network models which are believed to provide the theoretical underpinnings of the sequential activation.

2.2 Computational models of feedforward networks

2.2.1 Spiking balanced networks

Although there is a wealth of models generating seuqential activation (Rajan et al., 2016), only the ones based on the integrate-and-fire (IAF) neuron model are being investigated here. Integrate-and-fire model is widely used in theoretical neuroscience and was first proposed by Lapicque (1907), long before the action potential mechanism was understood (see Figure 2.1). It provides a useful simplification of the underlying biophysical mechanisms by treating the model neuron as an electric circuit consisting of a resistor R and capacitor C in parallel (Abbott, 1999).

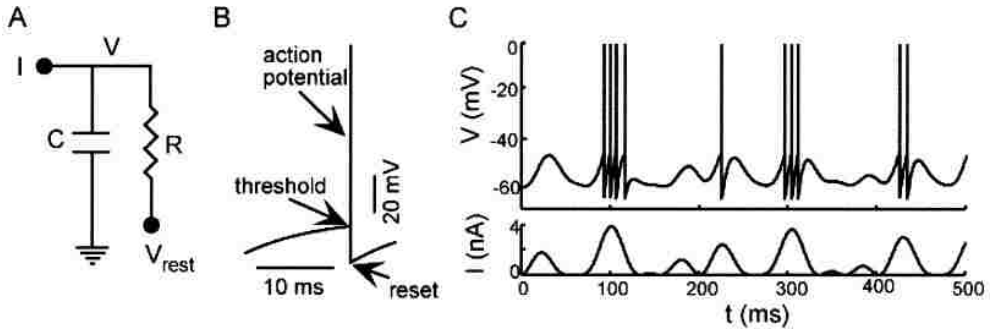

Figure 2.1: The IAF model. A: An equivalent RC circuit. V is the membrane potential, V_{rest} is the resting membrane potential, and I is an injected current. B: The voltage trajectory of the model. When V reaches a threshold value, an action potential is generated and V is reset. C: The upper trace is the membrane potential and the bottom trace is the input current. Diagram from (Abbott, 1999).

Networks composed of such IAF spiking neurons have emerged as a standard theoretical model and are ubiquitously adapted in both analytical and numerical studies. These networks comprise two populations of excitatory and inhibitory neurons which are sparsely (typically 1-10%) and randomly connected via the weighted synapses. Based on the cortical measurements, 80% of neurons are excitatory and the inhibitory weights are stronger so that the two forces are balanced. Depending on the amount of external noise injected into the system and the ratio between the excitatory and inhibitory weights, different regimes of activity were observed. The network's activity can be either asynchronous or synchronous from a population viewpoint and either regular or irregular from a neuron viewpoint (Brunel, 2000). The asynchronous irregular (AI) regime is considered to model a default state of the cortical networks.

The notion of the balance between excitation and inhibition (E/I balance) was introduced to generate the irregular responses of the individual neurons and it assumes that the positive and negative inputs cancel each other. If excitation dominates, the neuron produces a fairly regular firing patterns in spite of high levels of noise. If the excitation and inhibition are balanced on a slower scale, yet remain uncorrelated on the faster one – then the net input current is dominated by those faster fluctuations and the membrane potential follows a random walk towards the firing threshold, resulting in an output spike train with Poisson statistics (Denève and Machens, 2016).

2.2.2 A generic model of feedforward networks

A basic model of a feedforward network (FFN) was proposed by Abeles (1991) and termed a *synfire chain* to highlight the synchronous firing patterns of neurons along the chain. Essentially, it is a divergent-convergent network composed of groups (also called layers) of excitatory neurons connected in an exclusively feedforward manner. A synfire chain, as shown in Figure 2.2, can be described by its width w – a number of neurons in every group; and length l – a number of such groups. A given neuron in a group i receives connections from all the neurons in the preceding group $i-1$; there are neither lateral connections within the groups, nor the feedback ones between the groups.

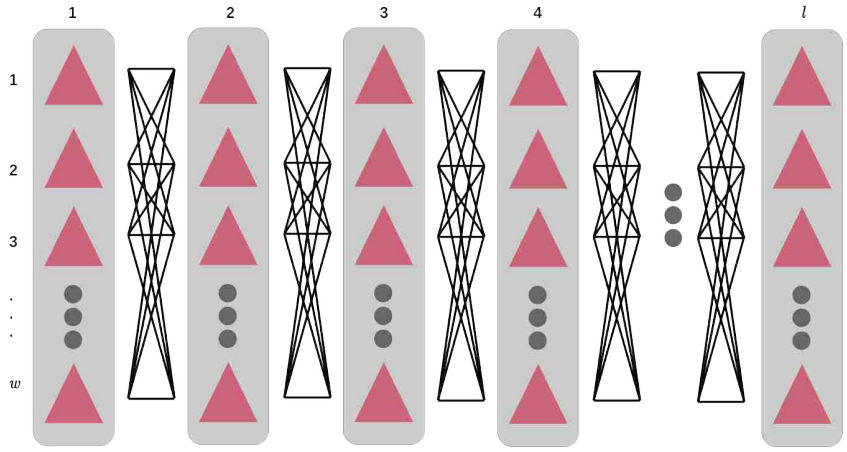

Figure 2.2: A diagram depicting a synfire chain with a w width and l length. All neurons are excitatory pyramidal cells and the consecutive layers are connected in an all-to-all fashion.

Although the synfire chain might appear to have a rather simple architecture and resulting dynamics, numerous implementations and extensions involving isolated as well as embedded variants focused on distinct aspects, stumbled upon various problems and operated in different parameter spaces. A systematic overview of these models can be found in reviews (Kumar et al., 2010; Vogels et al., 2005; Destexhe and Contreras, 2006), in the following sections, however, the goal is to identify a set of key features that can be used to describe and assess these models and are relevant to the extensions presented in this thesis. As these features are tightly coupled to one another, some models, principles or problems will be mentioned several times.

2.2.3 Two modes of propagation

Firstly, the signal transmission along the chain can be realised in two distinct modes: synchronous spiking or asynchronous firing rate propagation; which correspond to two widely discussed paradigms of neural coding (Shadlen and Newsome, 1994; Perkel and Bullock, 1968).

The *rate coding* paradigm postulates that all the information about a stimulus is contained in the firing rate of the population of neurons. Single neuron responses were shown to be variable from trial to trial and only averaging a given population may carry information (Vogels et al., 2005). An opposing view – the *temporal coding* paradigm – assumes that precise spike timing is a significant element in neural coding (Softky and Koch, 1993), as a number of experimental studies have found (Butts et al., 2007). Although deciphering and understanding the neural code is one of the central problems in neuroscience, there are no definite answers yet which paradigm is closer to truth or whether the two can actually coexist in different systems (Luczak et al., 2015; Brette, 2015).

2.2.3.1 Synchronous spiking

The majority of models reviewed here relate to the synchronous spiking mode. This mode is initiated by an arrival of a strong transient input at the first layer of the synfire chain, which makes the neurons fire in unison to form a so-called *pulse packet* (Diesmann et al., 1999), characterised by its activity a – number of spikes and σ – temporal dispersion measured by the standard deviation of the underlying pulse density (more details on the pulse packet characteristics are given in Section 4.6). The pulse packet's activity in the first layer propagates to the second layer after one synaptic delay and then further down the chain. The signal transmission in this mode is thus characterised by a rapid and transient pulse traversing the layers and engaging usually most of the neurons belonging to the chain.

2.2.3.2 Asynchronous firing rate propagation

The second mode of signal transmission involves propagating firing rates, rather than transient pulses, along the layers in a feedforward chain. Because the neurons in the deep layers of synfire chains by definition share a lot (if not all) of their input connections, preventing synchronisation is the main concern in this mode.

This problem was tackled by van Rossum et al. (2002), who implemented an isolated feedforward chain of excitatory neurons. At the absence of external noise, only a strong input into the first layer would cause the chain to activate, resembling a synfire pulse packet activation. Addition of a noisy background current with a positive mean changed the mode of propagation, so that also weak inputs would reach the distant layers and the rates would be faithfully transmitted along the chain. However, the background current needed two components adjusted to be rendered useful. Firstly, the mean current, which mostly brought the average neuron membrane potential closer to the threshold, so that the rate propagation would proceed faster. And secondly, the noise component, which ensured that the neurons are in different states when the signal arrives preventing them from acting in synchrony. With only the mean current present, the synchrony would still arise, whereas only noisy component would cause the strong thresholding of inputs so that weak firing rates would fail to propagate.

Another model attempted to incorporate the notion of E/I balance to examine its role in the firing rate propagation along feedforward networks (Litvak et al., 2003). In this implementation, an isolated synfire chain consisted of layers of 3000 excitatory and 3000 inhibitory neurons and each neuron maintained 10% connectivity with the preceding layer. Input was modelled as long, uncorrelated Poisson spike trains (half excitatory, half inhibitory) to keep a precise balance, where the net synaptic current was zero, and the response was driven entirely by the variance of the membrane potential. It was shown that it is difficult to transmit the firing rates beyond the third layer as deeper layers would fire independently of the input firing rate. The neurons belonging to the same, deep layers also developed synchrony. Removing a precise balance in input spike trains caused an emergence of strong synchronisation resembling a pulse packet travelling along the synfire chain.

A model of firing rate propagation in a feedforward chain within a recurrent network has also been proposed (Vogels and Abbott, 2005). In this implementation, a synfire chain was not artificially embedded into an otherwise random structure – instead, the already existing connections were chosen to form a chain. Since the network was sparsely connected (2%), not many neurons shared large pools of inputs and outputs, and a width of a chain was set to only 33 excitatory neurons and in deep layers each neuron would receive at least 3 input connections from the preceding layer, rendering it a rather 'diluted' version of a synfire chain. To compensate for the small number of connections, synaptic weights along the chain were strengthened up to 13 times compared to the rest of the synapses in the whole network. With these adjust-

ments, it was demonstrated that firing rates can be propagated in a feedforward manner in sparse networks and they are capable of performing simple logical operations.

2.2.4 Chain–rest of the network relationship

Another key feature of a model is the way how the relationship between the chain and its surrounding network is being modelled. In the actual brain, every local circuitry is in fact only a small part of an enormous recurrent network and apart from local connections, it also receives lots of inputs from distant locations and sends many outputs there. Theoretical models have no capacity to precisely and faithfully recreate the whole system, and the activity of *the rest of the brain* is usually simplified or removed entirely. Concerning the feedforward chains, these can be either isolated entities or embedded in a bigger, recurrent network and the examples of the two classes were already given above.

It should be highlighted, that the *chain – rest of the network* relationship is essentially *mutual* and both entities influence one another:

– The network provides external as well as feedback inputs in a form of the balancing excitation and inhibition.

– The chain has a potential of destabilising the whole network through its non-random connectivity and transient synchronous events.

Below the two points are discussed in greater detail.

2.2.4.1 Network's influence on chain

By isolating a chain, that is by creating it as a stand-alone entity, the relationship between the chain and the rest of the network is completely abolished. In consequence, some compensatory adjustments need to be added to the model, as already seen in the above-mentioned examples, where the fine-tuned noise or balancing inhibition had to be present in order to keep the neurons' activity asynchronous (van Rossum et al., 2002; Litvak et al., 2003). Such setup, however, is too simplified and biologically unrealistic as the networks are recurrent and their noisy inputs are essentially *internal*, not *external* and in other words, the noise cannot be adjusted *to* the network, as it is actually fixed *by* the network (Vogels and Abbott, 2005).

From a certain point of view, it may look like a paradox that a finite network itself can provide the noise ensuring the asynchronous E/I balance. Neighbouring neurons inevitably share some fractions of their presynaptic inputs, yet they can produce un-

correlated spike trains (Ecker et al., 2010). Should not it be the external noise that adds extra randomness to the system that could help to scramble the otherwise correlated signals? Surprisingly, it was shown that the spike-train correlations in the simulated finite-size recurrent networks are much smaller compared to the case with the spike trains generated by large but finite-size non-interacting Poissonian processes (Tetzlaff et al., 2004). A follow-up study identified the negative feedback, not the E/I balance, to be a sufficient cause of efficient suppression of correlations between the neighbouring neurons (Tetzlaff et al., 2012). Negative feedback is a product of the network's own activity, a fact which proves that the recurrent networks indeed can decorrelate the spike trains by their own means. Thus, to ensure the asynchronous state during the signal propagation along the feedforward chain, a large enough host network should provide the negative feedback to its sub-elements rather than relying on the externally generated random inputs.

2.2.4.2 Chain's influence on network

The overall problem here is how can a sparse, randomly connected network in an asynchronous irregular (AI) state accommodate a synchronous transient within a subnetwork without simultaneous switching to a global synchronous state. It was shown that pulse packet's (PP) activity excites background neurons after one synaptic delay – a phenomenon termed a 'halo' of the travelling PP (Kumar et al., 2008a). Aviel et al. (2003) embedded a synfire chain of width 250 neurons into a large balanced network to find out that such setup leads to large transient global oscillations. Another study embedded a synfire chain of width 300 into a modified version of balanced networks – a locally connected random network (LCRN) (Mehring et al., 2003), where neurons are defined by their location on a 2D plane and connectivity is established with Gaussian kernels so that the probability of connection decreases with distance. Without embedding a synfire chain, just a mere simultaneous forced firing of 100–500 neighbouring neurons led to serious disruptions of the surrounding region. Excitatory and inhibitory neurons would fire in unison, causing a so-called *synfire explosion* followed by a silent period caused by a wave of inhibition quenching all the network's activity. Only after recovering from that period, the whole network would return to its default, AI state.

In the previous sections, another model was discussed which studied firing rate propagation in a recurrent network where no synfire explosions were reported (Vogels and Abbott, 2005). The reason for that is that the chain used was of a small size (33 neurons per layer), the layers had 11% connectivity and the synapses along the chain

were significantly strengthened. Under such conditions, an injection of a pulse packet activated only 33 neurons which fed to another 33 neurons and so forth via strong synapses, but the background was fed only via weak synapses. In setups with synfire explosions a *chain neuron* would feed to the next layer and to the background (non-chain neurons) via the synapses of equal strengths and, essentially, each layer would be much wider than 33.

Synfire explosion is definitely an undesirable phenomenon in a robust model and various solutions were proposed to tackle this. In the next section we will discuss two distinctive types of synapses used in models that might play an important role in synfire explosion generation. Other proposed ways of avoiding such instabilities involve adding inhibitory control along the chain (discussed in the forthcoming section), fine-tuning of network parameters, introducing heterogeneities or spreading the layers across the network to minimise the interference of the haloes from individual layers.

Instabilities caused by a strong halo is an extreme case of the chain's influence on the host network. In more general terms, the addition of an embedded feedforward connectivity into a recurrent network can be seen as introducing non-random elements into an otherwise random structure. As the AI state of a random network is believed to faithfully recreate a typical cortical activity, embedding a synfire chain automatically disturbs the randomness that might in turn compromise network's variability, general E/I balance and asynchronous activity. On the other hand, it should be stressed that the cortex itself is not perfectly random. Various studies revealed that cortical micro-circuits display many stereotyped motifs and neurons seem to be clustered to a higher degree than it would be expected in a purely random network (Song et al., 2005). Thus, a random network with the embedded synfire chain might be seen as a plausible example of a cortical network with some degree of specific circuitry, much more than an idealised random and sparse balanced network.

2.2.5 Current vs conductance-based synapses

As outlined above, embedding synfire chains in a recurrent network uncovered a problem with the explosions of activity. As indicated by Kumar et al. (2008a), the cause of the troubles might, at least partially, lie in the very choice of the computational model used in the simulations, namely current-based synapses.

There are two ways of modelling synaptic inputs – they can be either voltage-independent (as in the current-based model) or voltage-dependent (as in the conduc-

tance-based model). The current-based model (CUBA) is popular due to its relative simplicity and linear dynamics that can facilitate derivation of analytical solutions (Cavallari et al., 2014). The conductance-based model (COBA) is more biologically grounded, as, for example, it can reproduce a so-called high-conductance state, a state when a neuron under bombardment of inputs decreases its membrane input resistance 3 to 5–fold (Destexhe et al., 2003). The two models were exhaustively compared in a single-cell (Kuhn et al., 2004) and in a network setting (Cavallari et al., 2014; Kumar et al., 2008b) and indeed it was demonstrated that the two models differ considerably, especially in the second order statistics of neural population interactions.

And how does it affect the behaviour of synfire chains? Kumar and colleagues (Kumar et al., 2008a) explained how the dynamics of current-based synapses can lead to the synfire explosions. The arrival of a pulse packet invokes large compound EPSPs in the chain neurons whose decay is governed by a membrane time constant, which in the studies reporting the explosions was set to 10ms (Aviel et al., 2003; Mehring et al., 2003). As a result, the chain neurons emit multiple spikes, which also activate the background neurons leading to a global explosion. A large time constant, however, is not realistic and in a biological cell there is only a short conductance transient and the long-lasting compound PSPs are shortened at the high conductance states (Kuhn et al., 2004). Thus, in the case with conductance-based synapses, large compound PSPs are abolished as a strong transient input affects the integrative properties, the effective time constant is shortened and the neurons respond with one spike only to a strong input, as demonstrated by Kumar et al. (2008a).

2.2.6 Recruitment of inhibition

The original framework of synfire chains involves only excitatory neurons, as shown in most of the above-mentioned models. In a random balanced network, where all 4 types of connections between the excitatory and inhibitory populations exist ($E \to E$, $E \to I$, $I \to E$, $I \to I$) it may look somewhat bizarre to modify only one type: the $E \to E$ connections. It should come as no surprise that the balance is disturbed, sometimes even catastrophically, when only the excitatory signals get rapidly amplified without the counteracting inhibitory ones. Studies showed that inhibition tightly follows excitation (Wehr and Zador, 2003) as well as distant inputs recruit both, excitation and inhibition, leaving the E/I balance approximately unperturbed (Isaacson and Scanziani, 2011; Xue et al., 2014). Thus, even the very injection of the excitatory pulse packet into a pool of

local excitatory neurons might be considered as a breach of a general rule of balancing inputs in the cortex. Further divergent-convergent dense $E \rightarrow E$ connectivity only makes matters worse, upsetting the balance even more, as clearly demonstrated in the example networks with the occurrence of synfire chain explosions (Aviel et al., 2003; Mehring et al., 2003). Yet, only a few models explicitly employed inhibitory neurons to improve signal propagation and render the model more biologically plausible.

The circuitry of feedforward inhibition (details of this circuit are given in Section 2.3) added to an isolated synfire chain were studied by Kremkow et al. (2010b). They demonstrated that such disynaptic inhibition ensures that only strong and synchronous signals get propagated which eliminates a problem of spontaneous signals triggering the chain creating false-positive responses. This study also embedded such chain into a recurrent network, but only 3 layers were involved, which could not elicit a synfire explosion.

Another model that incorporated inhibition along a synfire chain aimed to reproduce the birdsong (Cannon et al., 2015). The model itself is an isolated case where excitatory neurons are grouped into pools which are then grouped into zones. Feedforward connectivity is realised between the groups belonging to successive zones, which are arranged into a circle. A signal effectively travels along a spiral, visiting the zones several times via different groups. The crucial element of the setup are the local inhibitory pools for each excitatory zone, which do not receive the actual signal, but are just activated by the zone's activity. It was demonstrated that the presence of such dedicated inhibition improved spiking synchrony and consistency across the trials.

2.2.7 Gating mechanisms

All the models reviewed so far are focused on a general signal propagation, where the goal is to find a set of conditions that allows for a high fidelity signal transmission without upsetting the host network (if one exists). Once such setup is found, it is implicitly assumed that the signalling pathways are always in a ready state, or in other words, a given pathway, or a gate to it, is open. The actual cortical networks, however, are very flexible and take part in various tasks under different conditions and states. Thus, the signals not only need to be faithfully transmitted at the presence of an input, but they also might need to be dynamically changed, rerouted or even completely stopped (Chatham and Badre, 2015). In essence, identical inputs, depending on a current 'con-

text', might lead to different responses and in order to obtain such flexibility, control or gating mechanisms are required.

In principle, there are two contrasting ways of implementing the gating mechanism: either the gate is open by default (as in all the models discussed so far) and some extra mechanism needs to be added to close it, or the gate is persistently closed and in order to open it, certain modifications need to be applied. (More details on gating principle will be discussed in Chapter 5).

The principle of a gate closed by default was implemented in the context of firing rates propagation (Vogels and Abbott, 2009) within a 2D balanced network, where a subset of inhibitory neurons had distance-dependent connectivity, and the rest of the neurons were connected randomly. The signal path was composed of only 2 layers – sender and receiver. The sender contained exclusively excitatory neurons which projected onto either excitatory or inhibitory receiver neurons, so that they maintained a so-called detailed balance. In a default, balanced mode, the rate propagation would fail, as the receiver neurons would receive the signal monosynaptically which was then balanced out by inhibition delivered disynaptically via local connections. Only in the unbalanced state achieved by the decreased inhibition, the gate would open and the rate propagation would take place. This framework, however, failed to block the transient signals – whenever the sender would rapidly increase its firing rate, the disynaptic inhibition would be too slow to balance this upsurge out and the receiver would still respond to the signal.

The gating of pulse packets was tackled by Kremkow et al. (2010a). In their implementation, the signal path was composed of three layers: sender, gate and receiver embedded in an LCRN and the gating was controlled by the delayed feedforward inhibition (Kremkow et al., 2010b). They presented the idea of exploiting the arrival time of the inhibition to open and close the gates – if the inhibition arrived at the gate before the actual signal, the signal would be successfully blocked, if afterwards, the signal would freely pass. The arrival time of inhibition was controlled externally by injecting extra current, so it remained an open question what biologically plausible mechanism could implement this sort of precise control of the timing.

2.2.8 Summary of the models and proposed augmentations

How to compare the models and, most importantly, what can one learn from them? Table 2.1 summarises the chosen models with regard to the key features discussed in

the previous sections. What can be striking is a high variability of parameters and assumptions used. Some networks were small, some feedforward chains were embedded in huge recurrent networks, some required fine-tuning, some destabilised the whole host networks. Several points are concluded from the overview presented here.

1. An isolated case of a feedforward chain proved to be too idealised, although very useful for formulating the basic properties of the phenomenon. Disengaging the chain from its host network annihilates the mutual influence between the two. Without this, important features are missed out and the conclusions based on the studies on the isolated models do not necessarily find agreement with the actual networks found in brain.

2. Biological plausibility of a model is an important factor and the conductance-based synapse model (COBA) is preferred despite its non-linearity and higher complexity compared to the CUBA. It can be argued that the model does not need to reflect all the biological features, but the COBA synapse adds a lot more realistic behaviour, especially during the transient high-conductance states, essential for studying synchronous pulse packet propagation within a recurrent network.

3. It is strongly stressed that the use of inhibition should be regarded as another key factor in signal propagation modelling. When a chain is embedded in a balanced network, it cannot exclusively engage the excitatory population. A rapid amplification of excitatory signal should always be accompanied by the inhibitory activity to keep the network balanced and at the same time let the pulse packet traverse the feedforward structure. Feedforward inhibition is a good candidate to be integrated into the feedforward networks.

4. Overall, propagating signals across recurrent networks proves to be difficult, and it will be claimed, that this fact can be turned into a very desirable feature. In a network that embeds multiple chains or assemblies, by defaut they should be kept quiet. If the signal propagation was easy and possible in a wide range of parameters, spontaneous signal amplification would blur the signal transmission fidelity. It is argued that the default state of cortical networks is to prevent the assemblies from firing and only dedicated gating/control signal should have the power to turn the signal on.

. Model	Isolated?	Layer size	Inhibition in chain	Synapse model	Propagation mode	Comment
van Rossum et al. (2002)	Isolated	20	None	CUBA	Firing rate	Fine-tuning of noise
Litvak et al. (2003)	Isolated	6000	50%	CUBA	Firing rate	Fine-tuning of E/I balance
Vogels and Abbott (2005)	Embedded in BN	33	None	CUBA COBA	Both	Strengthened synapses
Aviel et al. (2003)	Embedded in BN	250	None	CUBA	Pulse packet	Global synchronisation
Mehring et al. (2003)	Embedded in LCRN	250	None	CUBA	Pulse packet	Synfire chain explosions
Kumar et al. (2008a)	Embedded in LCRN	300	None	COBA	Both	heterogeneities & COBA eliminates explosions
Cannon et al. (2015)	Isolated	100	Local shared pools	CUBA	Pulse packet	Birdsong sequence, spiralling chain
Kremkow et al. (2010a)	Embedded in LCRN	125	20%	COBA	Gating of pulse packet	Delayed inhibition, 3 layers
Vogels and Abbott (2009)	Embedded in LCRN	728 and 536	14% (in 2nd layer)	COBA	Gating of firing rates	2 layers, detailed balance
Kremkow et al. (2010b)	Isolated	125	20%	COBA	Pulse packet	Feedforward inhibition

Table 2.1: Comparison between various models extending the synfire chain architecture. Abbreviations: BN: balanced network, LCRN: locally connected random network, CUBA: current-based model, COBA: conductance-based model.

2.3 Neural inhibition

Most neurons in the CNS are excitatory and it is their activity that triggers other activities or is transformed into behaviour. However, a network composed of only excitatory units has a very limited range of behaviours and even a simplistic simulation inevitably leads to an observation that the excitation tends to spread and amplify nearly infinitely if left uncontrolled.

One can compare the activity of excitatory neurons to a herd of galloping horses: it will always push forward if given an opportunity, powerful and energetic but with no specificity or direction. It is the presence of inhibition that keeps tight reins on this darting activity. It creates obstacles, counterweights and gates in order to control the magnitude, timing, and direction of the flow of excitation. To fully understand the functionality of cortical activity, an appreciation of the less numerous players – inhibitory neurons – is indispensable.

2.3.1 Beyond balancing excitation

Neuronal excitation and inhibition are inseparable and a healthy, well-functioning network is believed to maintain a balance between the two. It was demonstrated that inhibition tightly follows excitation – whenever there is an increase or decrease in excitatory activity, the inhibitory activation faithfully follows suit (Wehr and Zador, 2003). Additionally, it was revealed that during sensory processing in the neocortex, presentation of a sensory stimulus invariably recruits inhibition in addition to excitation, leaving the E/I balance approximately unperturbed (Isaacson and Scanziani, 2011; Xue et al., 2014). Conversely, several reports suggested that the elevated E/I ratio and impaired inhibitory circuits may be responsible for the clinical features found in autism, schizophrenia, intellectual disabilities and epilepsy (Yizhar et al., 2011; Marín, 2012; Cossart et al., 2001; Engel, 1996).

The role of neural inhibition, however, goes far beyond a mere balancing of its rampant companion. Inhibitory neurons (also called interneurons) play a key role in various forms of network oscillations and synchronising the spiking of principal neurons which serves to coordinate communication between different brain areas. Depending on the connectivity patterns between the neurons, inhibition can gate the signals, control the temporal precision of spikes or normalise the activity of groups of neurons (Roux and Buzsáki, 2015; Kepecs and Fishell, 2014). Interneurons are also believed to be a dominant factor in mediating the selectivity of projection neurons and adjusting

their input/output relationship (Letzkus et al., 2015).

The microcircuits and several functions of inhibitory neurons will be discussed in more detail in the subsequent sections. Before that, let us briefly focus on the history of the discovery of inhibition which might shed some light on why the role of inhibition had been ignored for decades and only recently it started attracting the well-deserved attention which, beyond any doubt, advanced the understanding of neural circuits immensely (for a more comprehensive history of inhibition refer to (Fishell and Rudy, 2011; DeFelipe, 2002)).

2.3.2 A brief history of inhibition

When back in 19th century Camillo Golgi developed his ground-breaking staining method (De Carlos and Borrell, 2007), he was the first one to suggest the two main types of neurons – motor neurons (type I) with long axons and sensory neurons (type II) with short axons (DeFelipe, 2002). Santiago Ramón y Cajal rejected such motor/sensory differentiation and proposed his own terms instead, which were based solely on the morphology of neurons: "cells with a long axon" and "cells with a short axon" (y Cajal, 1995). For many decades afterwards, the terms *short-axon cells* and *interneurons* would be used synonymously.

Noteworthy, at Cajal's times the very idea of inhibition was still non-existing. Cajal reasoned that the short-axon cells might work as *"condensers, or accumulators, of nervous energy"* that would support the signal transmission along the principal cells by *boosting* their energy (DeFelipe, 2002). The possibility of decreasing or counteracting such energy was not taken into account at all. As a matter of fact, it was not until the late 60s and early 70s when immunocytochemical studies confirmed that most interneurons release a neurotransmitter called GABA – γ-aminobutyric acid – which *reduces* the neuronal excitability (Bowery and Smart, 2006). Only then the notion of neural inhibition, as an opposite to excitation and signalled by the GABA neurotransmitter and to a much smaller extent, glycine, was established. Nowadays it is estimated that inhibitory neurons make up roughly 15-30% of all neurons found in the CNS, with the actual percentage varying between the cortical layers, areas and species (Markram et al., 2004). They are still commonly called *interneurons* as they usually wire around local cells and their axons terminate in the same region where their dendritic arborisation is (with some clear exceptions). It should also be mentioned that some classes of interneurons are of an excitatory type, but here we aim to focus only on the inhibitory

ones, which means that any occurrence of the term 'interneuron' will imply the inhibitory interneuron, unless explicitly stated. Excitatory glutamatergic neurons, on the other hand, are commonly called principal cells and since their most numerous subgroup comprises the pyramidal neurons, all these terms are often used interchangeably.

2.3.3 Classification of interneurons

The Francis Crick's famous statement *"If you want to understand function, study structure"* (Crick, 1988) proved to be not the most useful advice for the community researching the interneurons as the whole population of interneurons is tremendously heterogeneous on the level of morphological, molecular and physiological features (Markram et al., 2004). A plethora of data collected over the decades and via independent research programmes rendered the problem of classification remarkably difficult as multiple approaches were used to systematise and name the individual subtypes. Although in 2008 a group of experts was convened to establish a common terminology and classification system for the whole interneuron population, a consensus is not reached yet (Ascoli et al., 2008). A study employing a web-based interactive system asked 42 experts in the field to classify 320 cortical interneurons to uncover a level of agreement regarding the assignment of the morphological features between the leading neuroscientists (DeFelipe et al., 2013). While some neuron subtypes were correctly classified with a high level of agreement (such as Chandelier or Martinotti cells), other subtypes – ironically including 'common type' and 'common basket' cells – turned out to be highly confusing even for the experts.

In contrast to this intricate, multi-dimensional systematisation, there exists a classification that creates non-overlapping groups of nearly all interneurons and is based on only four biochemical markers: the calcium-binding protein parvalbumin (PV), the neuropeptide somatostatin (SOM), the ionotropic serotonin receptor 5HT3a (5HT3aR) and the vasoactive intestinal polypeptide (VIP) (Rudy et al., 2011). While the PV and the SOM markers create one group each, the 5HT3aR marker demarcates two further groups – gathering the neurons that co-express VIP or not.

The second classification is widely exploited in contemporary studies harnessing optogenetic (Deisseroth, 2011) or two-photon calcium imaging techniques together with the recent advances in mouse genetics where the SOM, VIP or PV-expressing neurons can be selectively targeted and manipulated both *in vitro* and *in vivo* by virtue of genetically engineered cre-driver mice lines (Taniguchi et al., 2011).

PV-positive interneurons

Interneurons expressing calcium-binding protein parvalbumin are the largest group (40% of all interneurons) and they all exhibit fast-spiking firing pattern. There are two most common types of the PV-positive (from now on PV+) interneurons: basket and chandelier cells, whose connectivity patterns and electrophysiological properties vary (Rudy et al., 2011). Overall, the PV+ interneurons are believed to be specialised for speed, efficiency and temporal precision (Hu et al., 2014).

PV+ neurons receive inhibitory inputs primarily from other PV+ neurons (via both synapses and gap junctions) as well as from the SOM and VIP neurons. On top, some autapses were found on the basket cells in the cortical layer 5 that are believed to provide the fastest and most reliable form of feedback inhibition (Bacci et al., 2005). The output of PV+ cells targets the perisomatic region of pyramidal neurons including the cell body, the axon initial segment and the proximal apical and basal dendrites (Freund and Katona, 2007).

In the rodent hippocampus it was found that the strength and time course of PV-mediated inhibition decays with distance, a feature which presumably leads to higher precision in principal cell spike times (Strüber et al., 2015).

SOM interneurons

Somatostatin-expressing interneurons (SOM and occasionally also called SST) account for 30% of interneurons. These neurons, often represented by the dendritic-targeting, regular-spiking Martinotti cells are believed to provide more graded inhibition and control the inputs of the pyramidal cells (as opposed to PV+ cells that control the output owing to the perisomatic inhibition) (Wang et al., 2004). SOM interneurons do not inhibit each other, but reach all the other inhibitory populations (Pfeffer et al., 2013). Their connectivity with local pyramidal neurons is very dense, as they were shown to be synaptically connected to virtually all neighbouring cells, forming a so-called blanket of inhibition (Karnani et al., 2014).

In the cortex, SOM interneurons mediate surround suppression of visual responses (Zhang et al., 2014), provide disynaptic inhibition between neighbouring pyramidal cells (Silberberg and Markram, 2007), whereas in the hippocampus they were shown to control burst firing (Kepecs and Fishell, 2014). A subtype of SOM interneurons residing in layer 4 was found to specialise in disinhibition of local principal cells via the inhibition of fast-spiking PV+ cells (Xu et al., 2013). Another study showed that

the SOM and PV+ neurons are activated during different phases of a behavioural task (approaching a reward zone vs leaving it), suggesting that interneuron subtypes can specialise in temporal regulation of the flow of information during behavioural events (Kvitsiani et al., 2013).

5HT3aR-expressing interneurons

The last group, the 5HT3aR-expressing interneurons, associating 30% of all interneurons, comprises VIP (40%) and non-VIP neurons (60%).

VIP neurons are often associated with disinhibition as they tend to inhibit other groups of inhibitory neurons: the SOM and to a smaller degree PV+ interneurons, whereas they rarely directly target pyramidal cells (Hangya et al., 2014; Jackson et al., 2016).

Non-VIP neurons mostly comprise neurogliaform cells and a prominent example of those is a class of interneurons found in the cortical layer 1. Many studies refer to those neurons as L1 as indeed, in the whole layer apart from the 5HT3aR-expressing interneurons there is only a small group of SOM interneurons. Interestingly, L1 neurons were also found to be implicated in mediating disinhibition in the context of associative fear learning (Letzkus et al., 2011).

Summary of the subgroups

The big, overly simplified, picture that emerges from the studies on the interneuron subgroups is as follows: PV+ interneurons are a fast-spiking population that controls the output of pyramidal cells via perisomatic inhibition. SOM interneurons are dendrite-targeting, regular-spiking neurons that control the inputs of the pyramidal cells. 5HT3aR-expressing interneurons, on the other hand, are believed to predominantly mediate disinhibition.

2.3.4 Inhibitory circuits

Inhibitory neurons can form various stereotypical circuits with pyramidal neurons that are widely found in the brain. Depending on the connectivity patterns between the neurons, one can define the following motifs: feedforward inhibition, feedback inhibition, lateral inhibition, direct inhibition and disinhibition. Direct inhibition is self-explanatory, whereas the forthcoming section will thoroughly review the disinhibitory circuit. The remaining patterns are briefly discussed below.

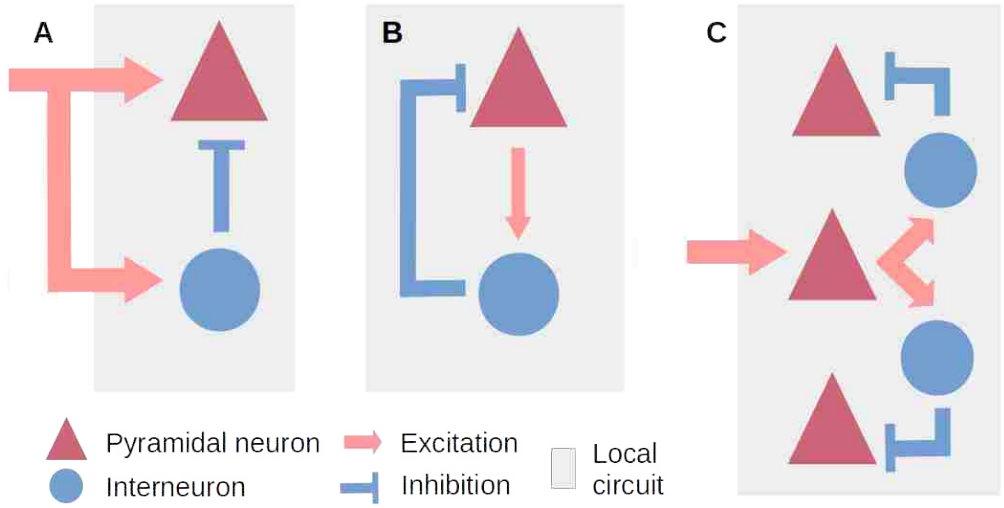

Figure 2.3: Main circuits involving inhibition. A. Feedforward inhibition. B. Feedback inhibition. C. Lateral inhibition.

Feedforward inhibition

Feedforward inhibition was already mentioned in the context of the synfire chain models (Kremkow et al., 2010b; Vogels and Abbott, 2009). In this scenario (Figure 2.3 A), the principal neurons receive an external signal and inhibition activated by the same source. The duration of the time window between the arrival of both determines how the signals will be integrated. Although inhibition is delivered disynaptically, owing to the lower firing threshold and more efficient synapses it may potentially arrive simultaneously or even before the signal, preventing the neuron from firing. When it arrives with a short delay, it can improve precision of the evoked spiking and filter out weak signals (Roux and Buzsáki, 2015). It was shown that such inhibition improves temporal precision in the auditory cortex (Wehr and Zador, 2003).

Feedback inhibition

In the feedback inhibition circuit (Figure 2.3 B), the principal neurons that trigger the inhibition are the ones who receive it via a feedback loop. Such circuitry provides a form of a regulatory mechanism akin to a thermostat, because the elevated activity of principal cells will inevitably cause an elevated inhibition fed back into the same neurons. Such reciprocal excitatory–inhibitory connections are ubiquitously found in various brain regions (Markram et al., 2004; Roux and Buzsáki, 2015).

Lateral inhibition

An extension of feedback inhibition is lateral inhibition (Figure 2.3 C). This occurs when the activation of principal cells recruits interneurons, which in turn suppress the activity of the neighbouring principal cells of a similar function. Such circuitry can realise neuronal competition and a winner-takes-all computational motif (Maass, 2000) to assure that only one assembly is active at a time.

2.3.5 Inhibition as a better modulator than excitation

Lastly, yet another important feature of the nature of inhibitory activity as a whole, is the fact that it appears to be more suitable than the excitation to provide modulation.

One study asked whether excitatory or inhibitory inputs can better gate the signal propagation along the synfire chains (Shinozaki et al., 2007). The chain was an isolated entity composed of Hodgkin-Huxley (HH) neurons which are more biologically grounded than the integrate-and-fire neurons (Izhikevich, 2004). A given layer, besides the pulse packet from the preceding layer, received a modulatory injection of either inhibitory or excitatory inputs at various times relative to the arrival of the pulse packet. It was demonstrated that when the inhibitory modulation arrived a few milliseconds before the packet, the signal transmission improved, whereas when it coincided with the packet's arrival, the transmission was suppressed. The excitatory modulation, on the other hand, failed to suppress the signal at all times, and managed to enhance it efficiently only when their arrival times coincided. When the excitatory modulation was arriving slightly before the packet, the modulatory signal would behave as a pulse packet itself and thus override and corrupt the timing information of the actual signal.

These results can be summarised as follows: inhibition, if timed appropriately, is capable of providing a bidirectional modulation: either by suppressing or boosting the signal without overriding its temporal information. Excitation, on the other hand, can only work unidirectionally: it merely helps to strengthen the signal. Because the excitatory modulation is composed of actual spikes, it can be disguised as a signal itself, causing a confusion between what is being transmitted and what is merely a modulator. Thus, it should be concluded that excitatory neurons do not provide the flexibility needed to modulate the signal propagation, and it is the inhibitory population that appears to be better equipped to control the transmission bidirectionally.

2.4 Disinhibition

The previous section presented a number of stereotyped circuits involving interneurons that are repeated across the brain and used to perform specific computations. Although all were shown to be powerful in many ways, here only the motif of disinhibition will be discussed in a greater detail.

2.4.1 Disinhibitory circuit

As outlined in earlier sections, persistent disharmony in the E/I balance is typically associated with a network dysfunction or disease. However, a *transient* disruption of this balance mediated by disinhibition is considered to be a potent mechanism for processing information in the brain networks. In principle, the balance can be disrupted by either increasing the excitation, or reducing the inhibition and it can be mediated by a whole range of mechanisms (Froemke, 2015). Here, however, the focus is placed solely on the disinhibition defined as *"a transient and selective break in the excitatory/inhibitory balance caused by reduced firing of different interneuron groups"* (Letzkus et al., 2015).

Disinhibition involves the removal of an already existing inhibitory control from a target cell and the elementary form of a disinhibitory circuit comprises 3 neurons (or groups of these): a principal cell that is under inhibitory control, an interneuron providing this control and another interneuron that removes it (Figure 2.4 A). Thus, a net effect of activating such pathway is an increased activity of the excitatory cell.

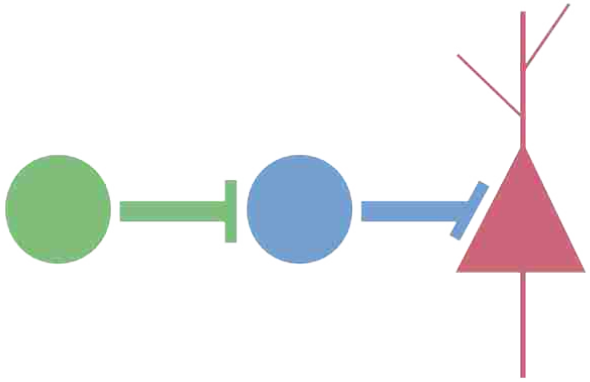

Figure 2.4: A basic diagram of a disinhibitory microcircuit. A pyramidal neuron is disinhibited by an interneuron (green) which inhibits another interneuron (blue) which holds the pyramidal neuron under the inhibitory control.

2.4.2 Disinhibitory circuit as a modulator

It has been already indicated that the inhibitory activity is a suitable candidate to provide modulation and the disinhibition is one form of realising this. As mentioned above, a net effect of disinhibition is an increased activity of the excitatory cells and this is the key element that renders disinhibition an attractive mechanism for modulating pyramidal cells' responses. It does not *cause* firing in itself, but rather it is *permissive* for strong activation of excitatory neurons (Poorthuis et al., 2014). By silencing the inhibitory inputs, disinhibition allows other systems/excitatory inputs to take over the control of the firing of the target cell.

This passive role in inducing increased spiking activity is exploited in signal gating, which is a common computational task performed by the disinhibitory circuits. Signal gating by disinhibition was recognised as the central influence that the striatum (a part of the basal ganglia) exerts on the motor system. Inhibitory striatal inputs relieve the target neurons from the ongoing inhibition and thus enable the premotor circuits to work (Chevalier and Deniau, 1990). Another example of gating by disinhibition was uncovered in the hippocampus, where the long-range entorhinal inhibitory inputs were shown to target the local CA1 interneurons and thereby causing transient disinhibition of the CA1 pyramidal neurons during memory tasks (Basu et al., 2016). When such disinhibition was finely timed with the incoming inputs from the CA3 fibres, more specific memories could be formed. At the absence of such synchrony, the mice exhibited inappropriate context-dependent responses as well as they did not distinguish novel from familiar objects very well. Thus, it was concluded that the disinhibitory circuit increases the precision of the hippocampus-based memory associations.

2.4.3 Disinhibitory pathways found in brain

A concise view on the role of disinhibitory mechanisms in neural circuits has only started emerging and to date, disinhibition has been shown to be implicated in associative learning, attention, social behaviour and spatial navigation (Letzkus et al., 2015). Many contemporary studies exploit the SOM, VIP and PV+ interneuron differentiation and scrutinise the disinhibitory circuits between these subgroups. Below a short review of the selected disinhibitory circuits found in the brain is laid out. It is by no means exhaustive, as the goal is solely to point at the mounting evidence showing the importance and ubiquity of the disinhibitory motif (for a more systematic review refer to Letzkus et al. (2015)).

The disinhibitory pathways $VIP \dashv SOM$ as well as $VIP \dashv PV+$ were uncovered in the medial prefrontal cortex and the auditory cortex during sensory processing (Pi et al., 2013). VIP stimulation was shown to modulate the gain of the auditory cortical responses. On the behavioural level, the recruitment of VIPs was strongly correlated with the reinforcement signals – punishment events – when the mice made a mistake in a discrimination task. The $VIP \dashv SOM$ pathway in the V1 is also linked to the top-down attention (Zhang et al., 2014) of visual processing. The pathway was shown to be activated by the inputs from the cingulate part of the mouse frontal cortex and the VIP and SOM responses were crucial for the correct visual discrimination. In similar vein, it was suggested that the $VIP \dashv SOM$ disinhibition might in fact create a so-called *spotlight of attention* by selectively overriding the lateral disynaptic inhibition between the pyramidal cells and thereby enhancing their responses (Karnani et al., 2016). This mechanism was conceptualised as a process of *opening the holes in the blanket of inhibition* and it directly refers to the already discussed view on the innervation pattern of the SOM interneurons which were shown to target nearly all the local excitatory cells.

Disinhibitory circuits often occur in the context of the cortical cross-modality. Although the brain contains specialised, anatomically isolated areas for processing different sensory inputs – visual, auditory, somatosensory cortices – it was found that the cross-modal interactions take place much earlier than previously thought (Kayser et al., 2005). Numerous studies showed that primary cortices – the visual V1 and the auditory A1 – are mutually innervated and on the behavioural level, the presence of sound improved the visual responses (Iurilli et al., 2012). On the microcircuit level, it was revealed that the targets of axons originating from the A1 are the L1 non-VIP neurons in V1, whose activation sharpened the orientation selectivity of pyramidal neurons via inhibitory and disinhibitory effects (Ibrahim et al., 2016).

Similarly to the visual-auditory cross-modality, several studies revealed that locomotion also improves the visual responses. Again, disinhibitory circuits were dissected, suggesting that locomotion activates VIP neurons in the V1, which in turn disinhibit local pyramidal cells via SOM inhibition (Fu et al., 2014). Subsequent optogenetic activation of the VIP neurons was shown to recreate the effects at the absence of the actual locomotion suggesting that it is necessary and sufficient to obtain improved visual responses (Fu et al., 2015).

Although the emerging picture suggests that via disinhibitory effects mediated by the VIPs, the SOM interneurons should decrease their activity, some studies contradict such conclusions. In the case of locomotion, one study showed the opposite effect (Polack et al., 2013), while another one reported two types of the responses in SOM neurons (Reimer et al., 2014). To reconcile these contradicting views, Dipoppa et al. (2016) suggested that the recruitment of disinhibition critically depends on the type of the visual stimuli. Since VIP and SOM in fact inhibit each other, the net effect of disinhibition can widely vary depending on many factors. Overall, it should also be noted that many of the aforementioned studies used very small samples, for example 28 VIPs pooled from 4 mice, 44 VIPs pooled from 7 mice and 11 SOM neurons with undisclosed number of mice (Fu et al., 2014). The misclassification of the SOM interneurons in cre-mice was also reported (Hu et al., 2013), a fact which could possibly explain the two contrasting responses of the SOM neurons (Reimer et al., 2014).

The last example of cross-modality comes from the interactions between the primary motor M1 and sensory S1 cortices (Lee et al., 2013). It was revealed that the pyramidal cells in M1 target the VIP interneurons in S1 which then activate the local disinhibitory pathway via the SOM interneurons. The interaction was shown to enhance the sensory processing in the cortex during whisking.

Finally, numerous disinhibitory pathways were found to be activated in the context of the fear learning, as already shown in the aforementioned study on hippocampal gating (Basu et al., 2016). In the auditory cortex, it was shown that the acquisition associative fear-memories critically depends on the recruitment of a disinhibitory circuit (Letzkus et al., 2011). The signalling pathway was meticulously dissected and it was shown that the circuit is activated by the cholinergic inputs originated from the basal forebrain and targeting the L1 interneurons via the nicotinic receptors. These L1 neurons then inhibit the PV+ interneurons in the L2/3 which in turn disinhibit the pyramidal neurons. In amygdala, a similar pattern was observed during the foot-shocks (Wolff et al., 2014) where both groups, SOM and PV+ interneurons, were shown to be strongly inhibited and thus enhanced the responses of the pyramidal neurons. The origin of the disinhibition, however, was not uncovered. Lastly, a mutli-area pathway was dissected during the defensive behaviour (freezing) which again, pointed at possible disinhibitory mechanisms (Tovote et al., 2016). The pathway consisted of GABAergic

inputs from the central nucleus of amygdala to the ventrolateral midbrain PAG region, which resulted in an increased activity of the local excitatory neurons, enabling their outputs reach the pre-motor targets in the magnocellular nucleus of the medulla.

2.4.3.1 Summary of discussed disinhibitory circuits identified in brain

We have presented a whole range of disinhibitory pathways found in numerous brain regions. The recurring motif across all of them is the consistent recruitment of the same, specific interneuron classes. It is either SOM and PV+ interneurons that provide the inhibitory control to the pyramidal cells, and the VIP and L1 interneurons that remove this control via the disinhibition (Figure 2.5). Importantly, the signal activating the local disinhibitory pathways typically comes from the external source – either from another cortex (such as the input from the M1 which activates the disinhibitory pathway in the S1) or the cholinergic basal forebrain fibres that directly innervate the VIP and L1 interneurons. Table 2.2 provides a more detailed and systematic review of the discussed disinhibitory pathways.

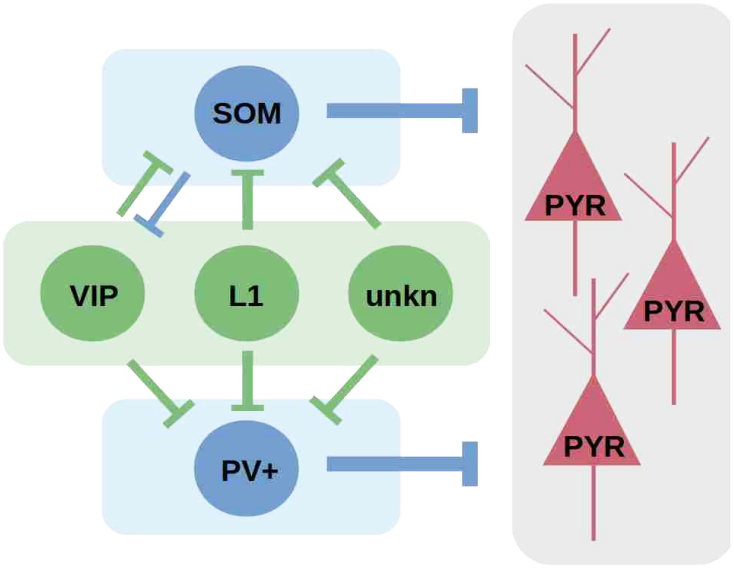

Figure 2.5: Simplified view of the selected disinhibitory microcircuits found in the brain. VIP, L1 and the unclassified interneurons belong to the disinhibitory group, whereas the SOM and PV+ hold the inhibitory control over the pyramidal cells. The $SOM \dashv VIP$ connection refers to the finding about the reciprocal interaction between the two groups (Dipoppa et al., 2016). Anatomically, there is also a $SOM \dashv PV$ connection which is omitted in the diagram since any of the discussed studies mentioned it.

Study	Brain region	Circuit found	Mode	Comment
Letzkus et al. (2011)	Acx	$BF \rightarrow nAChR \rightarrow L1 \dashv PV \dashv PYR$	awake	Fear conditioning. BF cholinergic signalling
Pi et al. (2013)	Acx & mPFC	$VIP \dashv SOM \dashv PYR$ $VIP \dashv PV \dashv PYR$	awake & in vitro	VIPs controlled by the reinforcement signals
Lee et al. (2013)	M1 & S1	$PYR(M1) \rightarrow VIP(S1) \dashv SOM(S1) \dashv PYR(S1)$	in vivo	Motor cortex input controls processing in sensory cortex
Wolff et al. (2014)	Basolateral amygdala	$unkn \dashv SOM \dashv PYR$ $unkn \dashv PV \dashv PYR$	in vivo	Fear conditioning
Fu et al. (2015)	V1	$VIP \dashv SOM \dashv PYR$	awake	Increased V1 responses independent of locomotion
Basu et al. (2016)	EC & CA1	IN (EC) $\dashv$ IN (CA1) $\dashv$ PYR (CA1)	in vivo	Disinhibition increases specificity of memory
Karnani et al. (2016)	L2/3 of V1 & S1	$VIP \dashv SOM \dashv PYR$	in vivo & in vitro	Disinhibition opens holes in the blanket of inh.
Ibrahim et al. (2016)	V1 & A1	$PYR(A1) \rightarrow L1(V1) \dashv SOM/PV \dashv PYR$	in vivo	Cross-modulation
Dipoppa et al. (2016)	V1	$VIP \dashv SOM \dashv PYR$ $SOM \dashv VIP$	in vivo	Extended model of disinhibitory control

Table 2.2: Selected disinhibitory circuits found in the Central Nervous System. Abbreviations: Str: Striatum; Acx: Auditory Cortex; BF: basal forebrain; ACh: Acetylcholine; nAChR: nicotinic cholinergic receptors; L1: interneurons in cortical Layer 1; mPFC: medial prefrontal cortex; EC: Entorhinal Cortex; CA1: region in hippocampus, IN: unspecified interneuron, unkn: unknown source of inhibition.

2.5 Neuromodulatory cholinergic system

Neuromodulation is a term used to describe various phenomena that modify electrical properties of a neuron, be it changing its excitability, altering presynaptic neurotransmitter release or inducing synaptic plasticity (Kaczmarek and Levitan, 1987). The difference between neurotransmission and neuromodulation is that the effect of the latter is not *directly* excitatory or inhibitory (mediated through ionotropic receptors), but instead it involves modification of a cell's response to a subsequent stimulation (Picciotto et al., 2012). Neuromodulation is critically involved in adjusting neural networks' activity to the current behavioural requirements and providing flexibility so much needed in the ever-changing environments. Although there are many neuromodulators found in the brain (Lee and Dan, 2012), here, the focus is placed solely on acetylcholine (ACh), which is considered to be crucially implicated in various cognitive functions including attention, memory, coordination of behavioural state and arousal (Hasselmo, 2006; Picciotto et al., 2012).

Peculiarly, acetylcholine was the first substance to be identified as a neurotransmitter. Back in 1921, German physiologist Otto Loewi (allegedly following an idea that occurred to him in a dream) performed a ground-breaking experiment where he demonstrated that electrical stimulation of the vagus nerve slows down the heartbeat by releasing a chemical agent – originally referred to as *Vagusstoff (vagus substance* in English) – which later turned out to be acetylcholine (Loewi, 1921; McCoy and Tan, 2014). Indeed, acetylcholine is a fast-acting neurotransmitter at the skeletal neuromuscular junction and in the autonomic ganglia, but in the central nervous system it acts predominantly as a neuromodulator. Only this role will be examined here.

Neuromodulators, unlike neurotransmitters, are confined to well-defined systems across the CNS. Acetylcholine is principally released by neurons originating from the two: basal forebrain (BF) and brain stem cholinergic systems (Newman et al., 2012). Neurons in the latter system primarily reach the thalamus, basal ganglia, and to a smaller extent they also innervate the basal forebrain and neocortex (Mesulam et al., 1983). The basal forebrain cholinergic system, on the other hand, contains projection neurons that innervate neocortex, hippocampus, entorhinal cortex, amygdala and olfactory bulb, among others (Mesulam et al., 1983). In the neocortex, BF axons project to all 6 layers and it should be highlighted that the BF input does not only comprise cholinergic neurons – in some parts of medial and lateral prefrontal cortex, BF fibres were meticulously measured and it was revealed that cholinergic terminals

represented merely $\sim 19\%$ of all terminals, whereas the glutamatergic and GABAergic ones represented $\sim 15\%$ and $\sim 52\%$, respectively (Henny and Jones, 2008). For the sake of completeness, it should be remarked that various brain areas including basal ganglia and prefrontal cortex also host local cholinergic neurons, some of which also co-release GABA (Thiele, 2013). However, unless stated otherwise, cholinergic modulation typically refers to the one stemming from the basal forebrain or the brain stem systems.

2.5.1 Cholinergic signalling

Classically, the cholinergic system is described as a *diffuse* or *reticular* cortical projection system (Sarter et al., 2016), highlighting the fact that the distribution of acetylcholine is non-specific and uniformly broadcast across the cortical regions. Many studies employing microdialysis techniques supported this idea by demonstrating that the *in vivo* measurements of cortical ACh do not considerably differ across a range of various behaviours (Zaborszky et al., 2015b). Another traditional view assumes that the cholinergic signalling is primarily carried out via diffuse extrasynaptic modulation coined *volume transmission*, as opposed to a *wired transmission* provided by the regular synapses or gap junctions (Agnati et al., 1995). Indeed, cholinergic fibres reaching the cortex have many axonal varicosities not associated with postsynaptic densities and they do not form synaptic contacts with the target neurons. As a result, the signalling is thought to be relatively slow and lacking spatial and temporal precision, a fact which attracted some scepticism about acetylcholine's role in attentional and memory tasks, which are thought to demand a lot more accuracy (Thiele, 2013).

Nowadays, however, these traditional views are being challenged. Recent studies employing 3D reconstruction and retrograde tracing of BF cholinergic (and non-cholinergic) neurons to various cortical areas concluded that these projections are not strictly diffusive and uniform. Instead, they are topologically organised into the segregated and overlapping pools of neurons, potentially broadcasting different signals from specific locations in the BF (Zaborszky et al., 2015a).

The view concerning the tonic ACh release via the volume transmission as the only mode of the cholinergic signalling also underwent adjustments. Several studies reported regionally–specific phasic signalling in the cortex coined *cholinergic transients*, proving that the ACh release can also be fast (on a few milliseconds scale) and precise (Sarter et al., 2009; Hangya et al., 2015). In the context of cue detection, phasic

signalling was found in the medial prefrontal cortex (mPFC) (Parikh et al., 2007). The phasic transients were present only in the trials with the correctly reported cues and the follow-up study employed optogenetic methods to uncover the causal role of such transients (Gritton et al., 2016). It was demonstrated that the suppression of cholinergic transmission was directly related to the reduced hit rates, whereas the photostimulation of either cholinergic soma in the BF or cholinergic terminals in the mPFC increased the number of false alarms in the non-cued trials. These studies, however, did not investigate the specific cellular effects of the cholinergic transients.

Nevertheless, cholinergic signalling is now considered to work on multiple time scales and provide a wide range of specificity and precision. Slow and fast ACh release is thought to be mediated by distinct synaptic mechanisms which involve two separate families of cholinergic receptors: metabotropic muscarinic (mAChRs) and ionotropic nicotinic (nAChRs) receptors.

Metabotropic muscarinic receptors are second messenger, seven-transmembrane, G protein-coupled receptors (GPCRs). Apart from the endogenous acetylcholine, they are also activated by muscarine and inhibited by atropine and scopolamine (Eglen, 2005). There are 5 subtypes of muscarinic receptors which are coupled either to G_q proteins (M1, M3 and M5 subtypes, referred to as M1-type) or to $G_{i/o}$ proteins (M2 and M4 subtypes, or M2-type) (Thiele, 2013). Muscarinic receptors are located pre and postsynaptically and affect mainly potassium and calcium channels. On the level of cellular effects, they were shown to cause depolarisation, reduction of spike frequency adaptation, increased excitability and spontaneous activity (Thiele, 2013).

Ionotropic nicotinic receptor channels are activated by nicotine and inhibited by mecamylamine. Among a large variety of their subtypes (Albuquerque et al., 2009), two main groups can be distinguished: the low affinity homomeric $\alpha 7$ receptors and the high affinity heteromeric $\alpha 4\beta 2$ receptors (Arroyo et al., 2014) which were shown to mediate fast and slow nicotinic responses respectively (Arroyo et al., 2012).

2.5.2 Action of Acetylcholine

As already stated, BF cholinergic fibres reach all 6 layers of the cortical mantle, but the projections are not uniformly spaced among them. Moreover, the distinct classes of neurons show very diverse expressions of numerous cholinergic receptors, so that the ACh can impact the local circuits in multiple ways and degrees. On a large network scale, acetylcholine was shown to enhance cortical sensory processing by improving a

signal-to-noise ratio (Sato et al., 1987) and causing cortical desynchronisation (Harris and Thiele, 2011). Reported effects of acetylcholine on the pyramidal neurons include direct depolarisation and reduction of cortico-cortical input by muscarinic receptor activation, enhancement of thalamo-cortical input via presynaptic nicotinic receptors or lowering the firing threshold of the target neurons (Poorthuis et al., 2014; Martinello et al., 2015). Here, the emphasis is put on the cholinergic modulation of the cortical inhibition. Although the knowledge of the effects of neuromodulators on specific interneuron subtypes is still incomplete (Kruglikov and Rudy, 2008), differential recruitment of distinct interneuron types is emerging as an important neuromodulatory mechanism. Indeed, local inhibitory neurons are thought to be the main targets of neuromodulation (Bacci et al., 2005), a fact clearly demonstrated in the context of cholinergic attentional modulation in macaque V4 (Mitchell et al., 2007).

In vivo studies revealed that the PV+ interneurons in layer 2/3 of anaesthetised mouse V1 do not directly respond to ACh released after the BF stimulation, but increase their activity at low ACh concentrations due to the muscarinic effects on the pyramidal cells. At higher ACh concentrations, however, PV+ cells attenuate their responses due to the activation of VIP and L1 interneurons via nicotinic receptors (Alitto and Dan, 2012). Noteworthy, the study did not detect any effects on the SOM interneurons, presumably due to the fact that the SOM responses were shown to be suppressed under anaesthesia (Urban-Ciecko and Barth, 2016).

2.5.3 Acetylcholine and disinhibition

The findings presented above lead to the conclusion that acetylcholine is very likely to invoke local disinhibitory pathways. Both, VIP and L1, populations were shown to be directly activated by the BF cholinergic inputs via the nicotinic receptors. These interneuron groups, in turn, preferentially target other interneurons and are implicated in numerous disinhibitory pathways as laid out in the earlier sections.

The study which dissected the signalling pathway during the auditory fear conditioning clearly demonstrated that the $L1 \dashv PV \dashv PYR$ pathway is in fact directly activated by the basal forebrain cholinergic inputs via the nicotinic signalling (Letzkus et al., 2011). In the context of locomotion, increased responses of VIPs were shown to be caused by exactly the same, $BF \rightarrow nAChRs$ signalling (Fu et al., 2014).

Finally, another already-discussed study which uncovered the $VIP \rightarrow SOM/PV$ pathways, showed that the recruitment of VIPs was strongly correlated with the rein-

forcement signals – punishment events – when the mice made a mistake in a discrimination task (Pi et al., 2013). These reinforcement signals were shown in a similar study to also rapidly activate the BF cholinergic fibres (Hangya et al., 2015). It is well plausible, that these in turn activate the disinhibitory pathways via VIP activation to enable learning.

2.5.4 Computational models of cholinergic modulation

Computational models of hippocampal circuits numerously employed cholinergic modulation to explain the theta rhythms formation and their role in supporting learning via the synaptic plasticity (Newman et al., 2012). Most of such models, however, do not differentiate between muscarinic and nicotinic effects and focus only on the volume transmission. Therefore, the implementation of cholinergic effects usually involves a uniform modulation of cellular parameters of all the neurons in a local circuit.

For instance, a simple model composed of a single layer of rate coded point neurons was used to demonstrate how the acetylcholine might tune the networks for a more effective processing by improving the signal-to-noise ratio (Hasselmo et al., 1992). The modulated variable in the model was the ratio between the external (thalamocortical) input and recurrent feedback reaching the excitatory neurons. High levels of ACh were shown to increase the efficacy of the feedforward signals, whereas the absence of modulation caused a significant interference between the feedback and feedforward signals impairing the memory encoding. Such model is admittedly over-simplistic, as the network itself is an isolated entity of the rate units. Also, the modulation evenly targets exclusively excitatory neurons and only these neurons receive the thalamic input.

One model explicitly harnessed the disinhibitory pathway targeted by cholinergic modulation in the context of visual selective attention (Sridharan and Knudsen, 2015). The investigated network, however, was again not composed of the spiking neurons, but was approximated with the mean field approach using two equations. Selective disinhibition was shown to be able to affect differential processing of information that was prioritised prospectively or retrospectively by selective attention.

In summary, to date, cholinergic modulation has been modelled as a global modification of neurons' properties and targetting primarily excitatory population. No model harnessed the recent findings concerning the differential recruitment of interneurons by acetylcholine in spiking networks, nor compared the volume with phasic transmission as possible modes of cholinergic signalling.

2.5.5 Summary

Cholinergic system can be summarised in the following points:

1. Although classically of diffusive nature, cholinergic projections originating from the basal forebrain are clustered and reach various, well-defined subareas of the cortex.

2. Cholinergic signalling is facilitated via slow volume transmission as well as fast, phasic ACh release via classical synaptic transmission.

3. Principal neurons as well as various classes of interneurons express a wide range of sets of muscarinic and nicotinic receptors found both pre- and postsynaptically, which vary across the layers as well as across cortical regions.

4. There is an immense repertoire of possible mechanisms that allow acetylcholine to selectively and precisely alter the neural circuits.

5. A recurring motif of cholinergic action is selectively targeting interneurons facilitating local disinhibitory effects - L1 and VIPs. Thus, disinhibition emerges as a potent mechanism activated by the cholinergic signalling via the nicotinic receptors.

6. The existing computational models typically do not differentiate between the receptor types and the cholinergic modulation is modelled by a uniform, global change affecting all the neurons or is targeted only at the excitatory population. No computational model investigated the role of acetylcholine in controlling the disinhibitory pathways in the spiking networks.

Chapter 3

Methods

3.1 Introduction

In this chapter, the main aspects of methodology and the reasoning behind the choices made that are common to all the variants of the model presented in chapters 4, 5, and 6 are explained. All the specific details of the further extensions are laid out in individual chapters for better clarity.

The model presented in this thesis is heavily based on the models developed by Kremkow and colleagues (Kremkow et al., 2010b,a). Their models, firstly, incorporate the notion of the disynaptic feedforward inhibition (FFI) into the synfire chain architecture and secondly, exploit the delay of such FFI to realise the gating mechanism. We aim to harness such architecture and the gating principle in the long chains of arbitrary length and crucially embedded in a random, recurrent network, as opposed to the isolated chains (Kremkow et al., 2010b) or short, 3-layer chains, embedded in the locally connected random networks (LCRNs) (Kremkow et al., 2010a) already scrutinised. We also aim to embed two chains into a network and study the functionality of the overlapping pools of neurons between them.

The majority of the basic properties of the network is modelled similarly as in the previous studies (Kremkow et al., 2010a,b) and all the values of the key model parameters are collected in Table 3.1.

3.2 Neuron model

The neurons were modelled as leaky, conductance-based integrate-and-fire (IAF) neurons with the subthreshold dynamics of the membrane potential described by the following equation:

$$C\frac{d}{dt}V^i(t) + G_{rest}[V^i(t) - V_{rest}] = I^i_{syn},$$

where the I^i_{syn} is the total synaptic input current into neuron i, and the capacitance C and the leak conductance G_{rest} are the passive electrical properties of its membrane at rest (V_{rest}). As soon as the membrane potential reached the threshold ($V_{th} = -57mV$), a spike was emitted, the membrane potential reset to its resting value ($V_{rest} = -70mV$), and the synaptic integration was suspended for the refractory period of 2ms.

Since the neural inhibition is of particular interest in this thesis, a certain limitation of the IAF model shall be mentioned. The dynamics of the model implies that a neuron can emit a spike only at the presence of a sufficient number of inputs. While this is generally true for excitatory neurons, many inhibitory neurons were found to be also active spontaneously, at the absence of any inputs (Häusser and Clark, 1997; Frank and Mendelowitz, 2012). Although such intrinsic firing is expected to have a considerable impact on the activity of the entire networks, this cannot be captured by the IAF model employed here. This limitation will be revisited later.

3.2.1 Synapses

Synaptic inputs were modelled as transient conductance changes, using exponential functions with $\tau_{exc} = 1.5ms$ and $\tau_{inh} = 10ms$. Various values of weights of inhibitory (J_{IE} and J_{II}) as well as excitatory (J_{EE} and J_{EI}) synapses were used, depending on the size of the network as well as the ratio between the strength of external noise and recurrent connections. The choice of the conductance-based synapses was based on the findings discussed in the previous chapter. Most importantly, this model is more biologically grounded than the current-based equivalent, and appears to be more suited to model high conductance states and prevent from the formation of synfire chain explosions.

The total synaptic current into the neuron i was given by:

$$I^i_{syn}(t) = -G^i_{exc}(t)[V^i(t) - V_{exc}] - G^i_{inh}(t)[V^i(t) - V_{inh}],$$

where the reversal potentials of the excitatory and inhibitory synaptic currents were set to $V_{exc} = 0mV$ and $V_{inh} = -80mV$ respectively. The total excitatory conductance in neuron i denoted by $G^i_{exc}(t)$ was given by the following:

$$G^i_{exc}(t) = \sum_{j=1}^{k_{exc}+K_{ext}} \sum_k g_{exc} \exp((t - t^j_k - D_{exc})/\tau_{exc}),$$

where the outer sum runs over all excitatory synapses, including the ones providing the external noise onto the neuron i and the inner sum runs over the spikes arriving at a given synapse. Transmission delay for all excitatory synapses was set to $D_{exc} = 2ms$. Similarly, the inhibitory conductance $G^i_{inh}(t)$ in the neuron i was expressed as follows:

$$G^i_{inh}(t) = \sum_{j=1}^{k_{inh}} \sum_k g_{inh} \exp((t - t^j_k - D_{inh})/\tau_{inh}),$$

with the transmission delay for all inhibitory synapses set to $D_{inh} = 3ms$. There were no external sources of inhibition.

The main reason for introducing different delays for excitatory and inhibitory synapses was to reliably induce a delay in inhibitory responses along the feedforward structures. Moreover, such a discrepancy was also prompted by certain electrophysiological properties of neurons such as conduction velocity and gating of receptors, suggesting that excitatory signalling tends to be faster.

$C = 290pF$	Capacitance
$G_{rest} = 29nS$	Leak conductance
$\tau_{exc} = 1.5ms$	Excitatory synaptic time constant
$\tau_{inh} = 10ms$	Inhibitory synaptic time constant
$D_{exc} = 2ms$	Excitatory synaptic delay
$D_{inh} = 3ms$	Inhibitory synaptic delay
$V_{exc} = 0mV$	Excitatory reversal potential
$V_{inh} = -80mV$	Inhibitory reversal potential
$V_{th} = -57mV$	Firing threshold
$V_{rest} = -70mV$	Resting potential
$N_{exc} = 5000$	Number of excitatory neurons
$N_{inh} = 1250$	Number of inhibitory neurons

Table 3.1: Values of the parameters used in the network model.

3.3 Network characteristics

The network was composed of excitatory (80%) and inhibitory (20%) neurons which were sparsely connected with the 5% probability of connection (fixed in-degree). The excitatory noise was injected to all the neurons so that the network maintained the low firing rate activity (1–2Hz) at the absence of a strong input.

The connectivity in the host network was set to be purely random, in contrast to Kremkow et al. (2010a) and Mehring et al. (2003), whose models employed the locally connected random networks (LCRNs). In this setting, every neuron was defined by its location on a 2D plane and the connectivity profiles were defined with the Gaussian filters, so that neurons were more likely to connect to the nearby neighbours and inhibitory neurons had considerably narrower Gaussian kernels to model a localised mode of inhibition. Here, in order to make as little assumptions as possible, the location-based connections were not employed.

Kremkow's 3-layer chain was embedded in such a way that the three layers were pooled from separate, distant locations. Then, the activation of the first layer influenced the activity of the second one mostly via the within-chain connections. The probability of other connections (both excitatory and inhibitory) via the rest-of-the-chain was very small. Clearly, such isolation prevented the network from forming synfire explosions and unwanted interference of the signal. Such setup, however, is believed to minimise the recurrent influence of the chains' activity back to itself. A pulse packet traversing the chain also influences the background neurons which, inevitably, also influence the chain in a feedback loop. In a fully recurrent network, a strong synchronised transient cannot traverse the network without changing it – and then remain unaffected by this change. When the individual layers are located far away from each other on a 2D plane, the probability of recurrent connectivity between them is low.

Localisation can also make the embedding of longer, multiple, and overlapping chains unnecessarily complicated – while 3 non-overlapping layers can be placed far away from each other without the risk of interference, 10–15 layers would need some denser allocation. Then one could define various patterns of allocating consecutive layers – they could be placed locally next to each other to form a spiral or rows, or they could be allocated randomly, with some fixed distance between the consecutive layers. This aspect was already briefly discussed in (Mehring et al., 2003). Embedding multiple chains would open doors for another set of patterns one could implement. It is conceivable that depending on different patterns of allocation or overlap, the chain–

background influence would vary so that a mode of computing overlaps would need to become yet another network parameter. Here it is assumed that delving into such details is unnecessary and definitely out of scope.

3.4 Feedforward chain characteristics

The main feature of the models studied here is the embedded chain of groups of neurons that are expected to propagate the sequential activity in a cascade-like fashion. The previous studies investigated the case only with 3 groups, but here the concept of temporal gating is adapted in the extended scenario. We will use an arbitrary number of groups, which from now on will be called layers. These layers were connected to form a chain using the aforementioned principle of the disynaptic feedforward inhibition in order to introduce the delay between the excitation and inhibition arriving at the layer (Figure 3.1). Each layer consisted of 100 excitatory and 25 inhibitory neurons. Every neuron in any layer (excluding the first one) received input from 60 (out of 100) excitatory inputs from the preceding layer. Every excitatory neuron in any layer received 15 (out of 25) inhibitory inputs from the same layer.

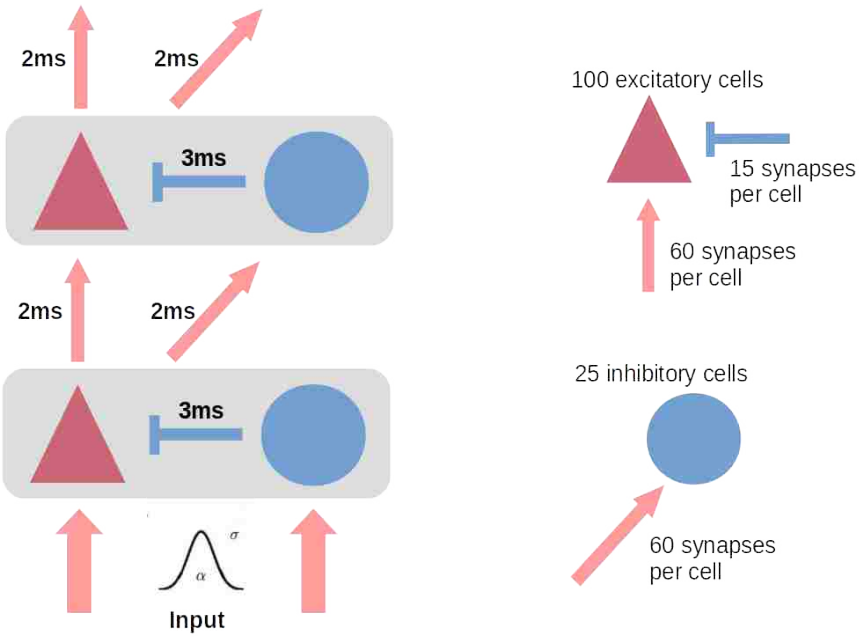

Figure 3.1: Connectivity along the chain. Each layer consists of 100 excitatory and 25 inhibitory neurons. The connectivity between the consecutive layers is on the level of 60%. Input reaches both, excitatory and inhibitory pools.

3.4.1 Embedding the chain in the network

Since the host network was random, the selection of neurons to build a chain was also random. To ensure that all neurons maintain the 5% global connectivity (fixed in-degree), the in-chain neurons received fewer connections from the background. This rule was consistently applied in all the further modifications of the model, even if not mentioned explicitly.

No restrictions were imposed on pooling the neurons to set up the random connections onto the in-chain neurons, except for not allowing the self-connections (autapses). This freedom entailed that the pool of presynaptic neurons could contain both non-chain as well as in-chain neurons, meaning that an individual neuron in, for example, 6^{th} layer, could receive direct inputs from neurons belonging to the 1^{st}, 3^{rd} or 8^{th} layers, apart from the signal pathway containing synapses from neurons in the preceding 5^{th} layer. Prohibiting such short-circuiting within the chain could potentially cause problems in cases with embedding long, multiple and non-overlapping chains, as the number of non-chain neurons there would be considerably smaller. Any restriction would need to be parameterised and could possibly cause extra non-random pathways between the non and in-chain populations.

3.4.2 The nature of overlaps

Two embedded chains with overlapping pools were generated the following way. Excitatory and inhibitory overlaps between the two chains were defined and processed separately. The levels of excitatory overlaps were selected from the following set: [0%, 10%, 20%, 30%, 40%, 50%]. The maximum overlap level was chosen to be 50% as any higher level would create too similar chains. Inhibitory levels of overlap, however, were allowed to go up to 100%. Even though technically the inhibitory pools participate in the chain architecture, the propagating signal is carried solely by the excitatory neurons that feed the activity forward. Inhibitory neurons provide control within a layer and as such do not transmit the signal. Due to their controlling rather than relaying role, it was assumed that two chains in principle could share the same inhibitory population. It could be argued that for the sake of completeness, excitatory overlaps should also reach the natural limit of 100%. This point will be revisited in the forthcoming chapter.

Most importantly, the overlaps were defined on the level of the whole chain, not the individual layers. That means that the 10% overlap does not entail that, for example,

5^{th} layer in the first chain shares 10% of its neurons with the 5^{th} layer in the second chain. It was assumed that such setting is highly ordered and thus less probable to be found in real networks. A random assignment was used instead which involved selecting a pool of neurons to be shared between the chains and each chain assigned individual neurons to layers independently and randomly. As a result, a given neuron could for example belong to the 3^{rd} layer in the first chain and to the 12^{th} layer in the second chain. Similarly, 100% inhibitory overlap does not entail that the connectivity is exactly the same in both chains. Rather, it means that every layer in one chain is a mixture of neurons belonging to potentially all the layers in the second chain and every such neuron is involved in the two connectivity patterns along both chains.

3.5 Input injection

Previous studies on signal propagation, in order to activate the chain, either all the neurons in the first layer were forced to spike simultaneously, regardless of their membrane potential (Mehring et al., 2003), or they received an input in a form of a volley of spikes (Kremkow et al., 2010a). Here, the latter approach was used, which was extended to include the inhibitory population in the first layer to also receive the volley.

In simulations it was shown that simultaneous activation of a large group of exclusively excitatory neurons tends to be detrimental to the whole network's stability (Mehring et al., 2003). Similarly, experimental results suggest that in the neocortex the projection axons invariably recruit inhibition in addition to excitation (Isaacson and Scanziani, 2011; Xue et al., 2014), whereas thalamic input onto the cortex was shown to recruit inhibition more strongly that excitation (Cruikshank et al., 2007). Thus, the very injection of excitatory pulse packet into a pool of solely excitatory neurons might be considered as a breach of a general rule of balancing inputs in the cortex.

The volley is characterised by two parameters: α – number of spikes; and σ – temporal dispersion. In most simulations, the values were chosen to ensure that the chain is reliably activated on most trials. For each trial a new volley of spikes for each projection neuron was randomly generated instead of fixing one instance of a volley to avoid the variability as it was done by Kremkow et al. (2010b). Tests were run to examine network's behaviour when a fixed volley was repeatedly used. Indeed, some volleys tended to lead to successful signal propagations more frequently than others, but qualitatively the results were similar. It is implausible that an external input is fixed – when distant axons and multiple synapses are involved, there is an intrinsic

variability due to the variable synaptic delays, release probability and others. Since the goal was to capture a general behaviour of the system with all its randomness, the volleys were generated separately for each trial.

3.6 Generating data via simulations

The networks studied here are highly non-deterministic and contain many sources of variability, such as random noise, recurrent random connectivity or injected inputs. Since the identical setups might lead to drastically different behaviours, it was vital to generate enough data to capture the whole range of possible outcomes.

Most of the simulations involved creating an instance of a network with a predefined set of parameters and recording its activity during signal propagation along the embedded chains. The following protocol was obeyed to obtain the data. For every set of parameters, 100 instances of the network were created with random connectivity and random pools of neurons assigned to the embedded chains. Then, each such instance was recorded for the time window of 21 seconds, filled with 20 injections of randomly generated volleys of spikes into the first layer of the chain. As a result, for each setup 2000 trials of signal injection coming from 100 individual networks were available and the setup's characteristics were expressed by the probabilities.

3.7 Simulation tools

All network simulations were written in Python and run using the NEST 2.10 simulator (Gewaltig and Diesmann, 2007). Simulation management was provided by the Python package NeuroTools (http://neuralensemble.org/NeuroTools/). Data analysis was carried out in Python using libraries SciPy, NumPy and visualised with the help of the library Matplotlib.

Chapter 4

Signal propagation in overlapping chains with disinhibitory pathways

Chapter summary

In this chapter, we exploit the principle of disynaptic feedforward inhibition and build a long chain with layers of both excitatory and inhibitory pools embedded into a purely random network. We then add disinhibitory pathways to demonstrate that these can be a powerful mean of control of signal propagation across random networks. We also study two chains sharing the pools of neurons and demonstrate stark differences in functional relevance of excitatory and inhibitory overlaps between the two embedded chains. Specifically, we demonstrate that the inhibitory overlap between the chains stabilises the network and protects it from synfire chain explosions by realising the motif of lateral inhibition. We also show that inhibitory and disinhibitory overlaps should be considered as separate entities.

4.1 One embedded chain

Firstly, a single chain of 15 layers of both excitatory and inhibitory pools (as described in Section 3.4) was embedded into a random recurrent network and a strong volley of spikes was repeatedly injected into its first layer to elicit the mode of synchronous spiking propagation. The results obtained from 2000 signal injections are shown in Figure 4.1.

As expected, a transient activity lasting $\sim 50ms$ was observed along the chain, where a strong pulse packet traversed the subsequent layers in a cascade-like fashion.

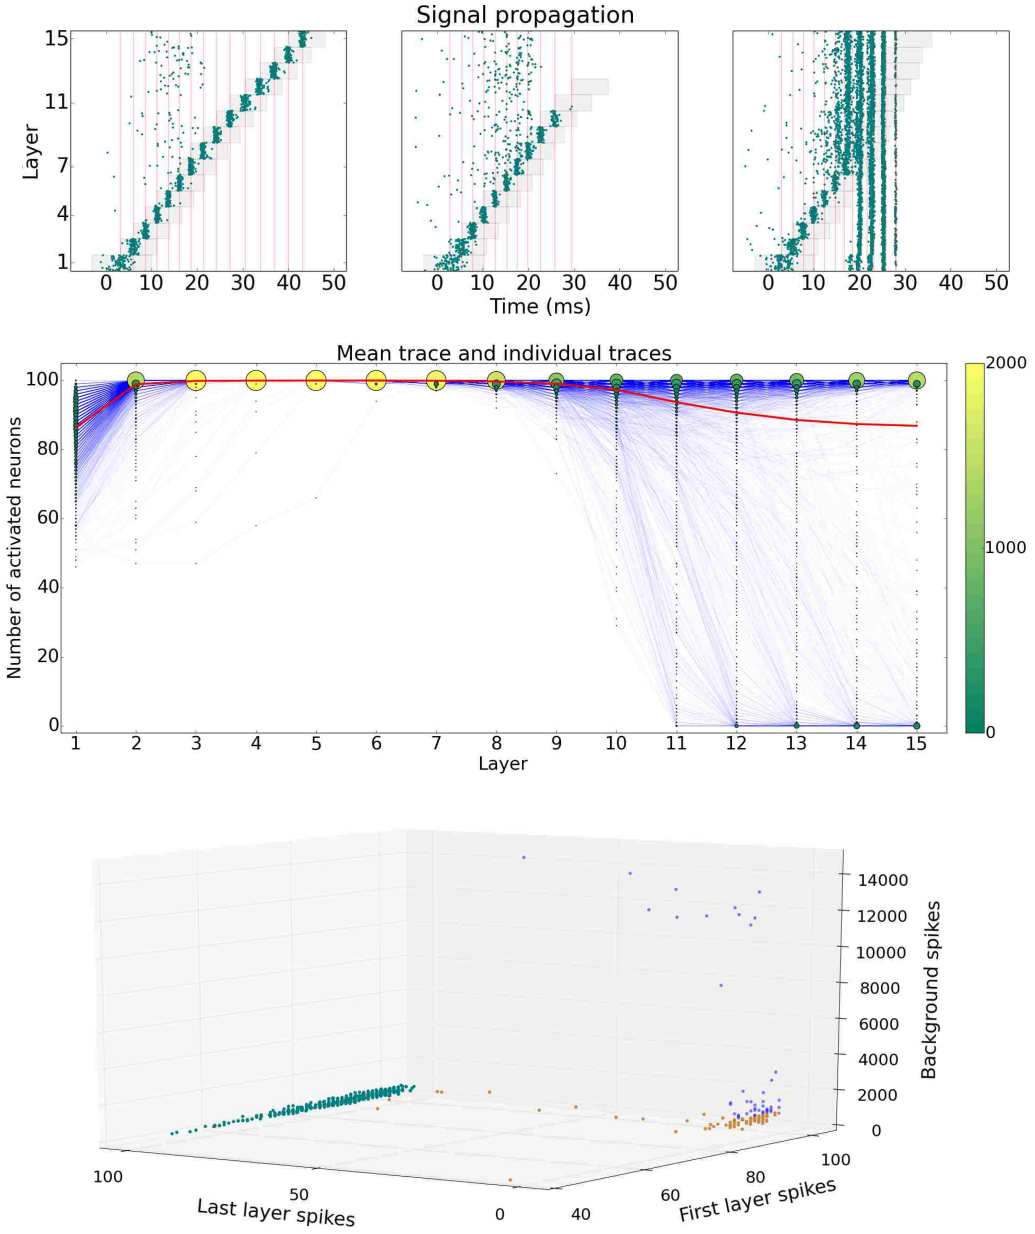

Figure 4.1: Signal propagation along a chain embedded in a random network. Top: Three possible outcomes of signal propagation, starting from left: successful propagation, failed propagation due to a strong inhibitory halo, and failed propagation due to a synfire chain explosion. Middle: Individual traces (blue) and the average trace of the signal propagation along 15 layers (red) for 2000 trials. Colour and size of the circles correspond to the number of trials. Bottom: Relationship between the activity in first layer, last layer and the background. Successful trials (teal) are clearly separated, the borderline between the halo (orange) and explosion (blue) failures is arbitrary and set to 700 spikes.

Overall, the signal was successfully propagated up to the final layer in 86% of trials, 10% led to synfire explosions, and in the remaining 4% of trials the signal declined along the chain before reaching the final layer. These results demonstrate that it is possible to transmit a signal along the chain within a random network, extending the previous reports which studied the locally connected random networks (LCRNs) with and without inhibitory pools within the chain (Kremkow et al., 2010a; Kumar et al., 2008a; Mehring et al., 2003).

4.2 Definitions

In order to describe and evaluate the results presented in Figure 4.1, it is useful to define certain activity characteristics that will also be relevant in the forthcoming extensions of the basic model.

First of all, one can characterise three fates of the signal – it can be either successfully propagated, it can decline before reaching the chain's final layer, or it can destabilise the whole network by causing a synfire chain explosion, as shown in the top panel in Figure 4.1. For the sake of completeness, one can define yet another fate – an *outset failure* – which denotes the case when the signal dies out immediately after the injection. This fate, however, is irrelevant here, as it refers to the signal's initiation, not its propagation across the network and the interactions with it. As a matter of fact, in all the experiments described here the input was chosen to be strong enough to minimise the chances of the outset failure occurrence, but due to the system's intrinsic randomness, a small number of trials (up to 3 per 2000) still failed to initiate the signal. These trials were removed from the analysis unless explicitly mentioned.

A successful signal propagation is the main feature of interest and it can be expressed in many ways. The middle panel in Figure 4.1 displays a mean trace (red) which is composed of arithmetic means of the numbers of neurons that fired in each layer in all 2000 trials (blue traces). The mean trace captures the average behaviour of the pulse packet – after initiation in the first layer, the number of active neurons increases to reach 100% up to the 6^{th} layer and afterwards the pulse packet undergoes disturbances which is reflected in the decreased number of active neurons in the subsequent layers. In the last, 15^{th} layer, the average number of active neurons equals 87% and this can be used as a measure of the chain's success rate. It can be argued though, that such measure is misleading as the underlying distribution is rather bimodal. Successful propagations typically activate 95-100%, whereas many failed ones exactly

0% of neurons and the arithmetic average between them removes this information. Instead, a ratio of trials that led to successful propagations can be given, which is still very much comparable to the above success rate and here equals 86%. This measure is similar to the chain's *survival probability* which was introduced by Gewaltig et al. (2001). From now on, the arithmetic success rate will be reported by default and the survival probability will be given only if strongly divergent from the first measure.

Mean trace	Arithmetic mean of the activity (number of neurons that emitted a spike) along the chain.
Success rate	Average number of active neurons in the last layer (last value in the mean trace).
Successful propagation	A signal propagation with at least 90% neurons active in the last layer.
Survival rate	A ratio of successful propagations.
Background spikes	A number of spikes that occurred in the non-chain excitatory neurons within a 50*ms* time window after the input injection.
Explosion failure	A failed trial due to the synfire chain explosion (less than 90% neurons active in the last layer and background spikes at least 700).
Halo failure	A failed trial due to the inhibitory halo (less than 90% neurons active in the last layer and background spikes less than 700).
Outset failure	A failed trial due to a faulty initiation (less than 40 neurons active in layer 5).

Table 4.1: Definitions of the terms introduced in this chapter.

Another important aspect of signal propagation is chain's influence on the host network. As already discussed in Section 2.2.4, a pulse packet travelling along the chain excites background neurons after one synaptic delay, creating an excitatory halo of its activity, which is then followed by an inhibitory halo. As a matter of fact, these haloes are the main cause of the signal's failure – too strong excitatory halo causes synfire chain explosions (as shown in Figure 4.1, top panel, right graph) , whereas too strong inhibitory halo hyperpolarises neurons rendering the chain, as well as the whole network, unresponsive to the following stimuli (as shown in Figure 4.1, top panel,

middle graph). To quantify the activity of the host network, a number of spikes that occurred in the background, non-chain excitatory neurons was obtained.

All the definitions of the aforementioned terms are included in Table 4.1. It can be argued that the thresholds chosen are arbitrary, but as shown in Figure 4.1C, the three groups are well separated: successful trials reside in the bottom left corner, failed by halo trials on the bottom right corner and failed by explosion trials on the top region of the plot. It should be noted that indeed, some explosions reside in the bottom right corner too (700 on z-axis upwards) and are somewhat artificially cut from the halo failures.

4.2.1 Excitatory and inhibitory haloes

When examining the top panel in Figure 4.1, one can notice an increased number of spikes across many layers within the $10-20ms$ time window, followed by a sharp cut around the $20ms$ time point. This behaviour is in fact a hallmark of the haloes of the travelling pulse packet. Firstly, the network is dominated by the additional excitation which in turn recruits additional inhibition that can subsequently dominate the network.

Figure 4.2 displays the traces of conductances and membrane potentials of three arbitrary excitatory neurons from the 10^{th} layer during the activation of the chain. Initially, these neurons began to receive an increased amount of excitatory inputs via the excitatory halo, but also the amount of inhibition followed suit. At around $21ms$ time, one neuron emitted a spike, akin to many other neurons across the network, after which the inhibitory inputs became stronger for all the three neurons. This was the time when the inhibitory halo dominated the whole network. At around $32ms$ time, the actual pulse packet arrived at the layer and only one neuron correctly responded to the activation while the two remaining ones were too hyperpolarised to reach the firing threshold and emit a spike. On top, the pulse packet arriving from the preceding, 9^{th} layer was already slightly weaker compared to the packet visiting the earlier layers (actual data not shown, but the middle panel in Figure 4.1 displays the overall decay of the number of activated neurons around the 9^{th} layer). This rendered it even harder to excite the neurons influenced by the inhibitory halo. As expected, the pulse packet failed to propagate any further and this example was labelled as the halo failure.

In the case of explosion halo (data not shown), the initial additional excitation recruits too many neurons at the same time and the subsequent inhibitory halo fails to

smoothly balance this, and instead a rapid sequence of *explosion – silence – explosion* appears, as shown in the top right panel of Figure 4.1.

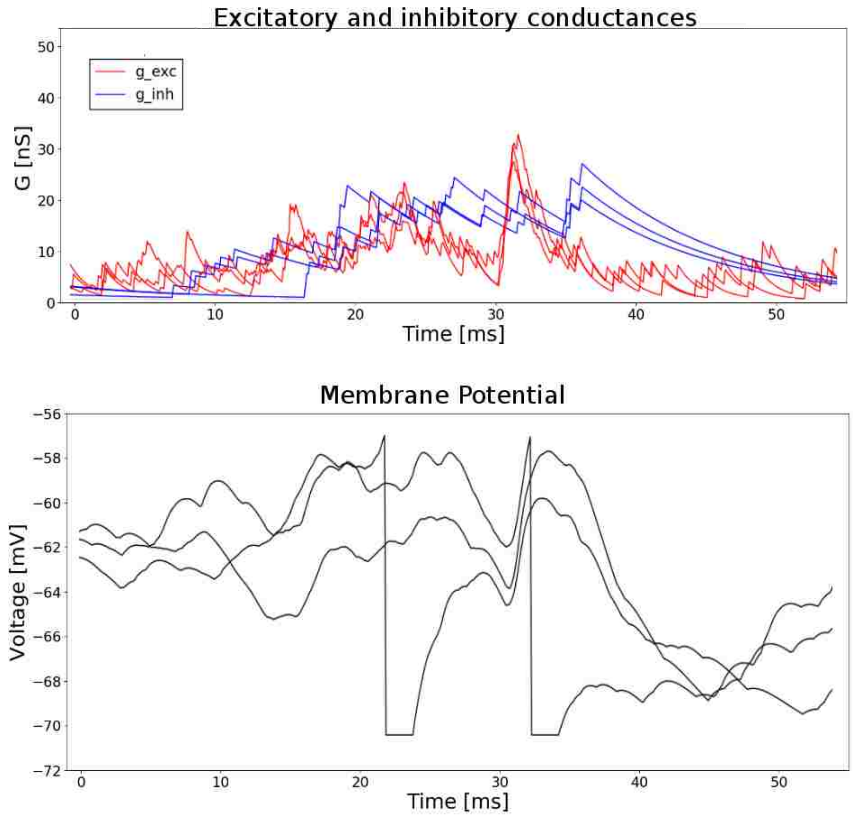

Figure 4.2: Excitatory and inhibitory conductances (top) and membrane potential (bottom) traces of three excitatory neurons from the 10^{th} layer of a chain during the signal propagation. The input was injected into the first layer at $0ms$ time point.

4.3 Chains with disinhibitory pathways

Although the explosions are the most severe and least desired outcomes of the synfire chain activation, the other cause of the failure – inhibitory halo – was tackled first. It was hypothesised that the principle of disinhibition might protect the chain from the inhibitory halo that spreads across the network and hyperpolarises the neurons rendering them unresponsive to the subsequent activation. Disinhibitory connectivity was expected to take advantage of this global inhibitory wave to specifically disinhibit neurons along the chain instead of directing excessive inhibition at them. The diagram of the chain with extra disinhibitory pathways is shown in Figure 4.3.

Since a fully operational disinhibitory pathway requires three synaptic contacts:

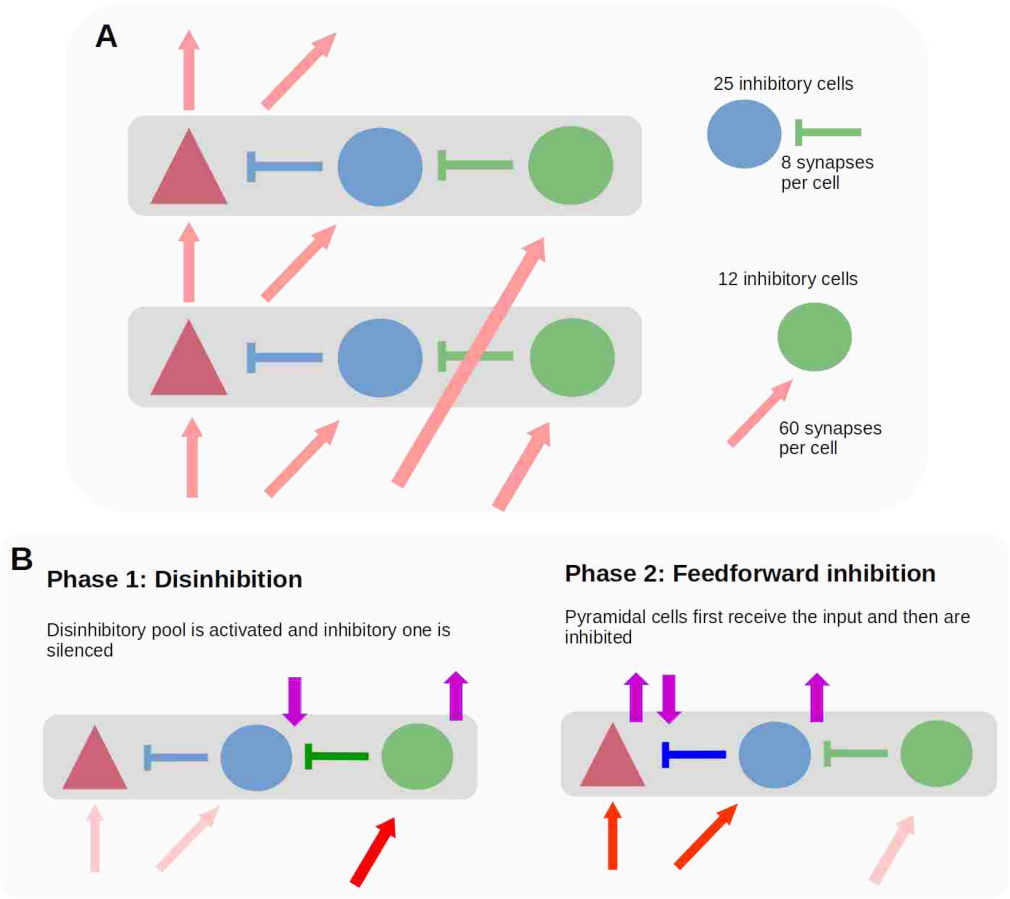

Figure 4.3: Disinhibitory pathway added to the chain. A: a schematic of the circuit. B: The same schematic, but split into 2 temporal phases. Firstly, disinhibitory pool deactivates the inhibitory one, and thus pyramidal cells are 'guarded' against the excessive inhibition. In the second phase, the feedforward inhibition motif is activated, similar to the circuitry without the disinhibition. Purple arrows denote the activation (up) and deactivation (down) of the pools.

$input \rightarrow IN \dashv IN \dashv PYR$, and the inhibitory synaptic delay is longer than the excitatory one. In the model described here, the excitatory pool in the 1^{st} layer connects to the disinhibitory pool in the 4^{th} layer (three layers ahead), the pool in the 2^{nd} one reaches the 5^{th} layer and so forth. Disinhibitory pools consist of 12 inhibitory neurons which connect to the chain's inhibitory pool such that each neuron receives inputs from 8 out of 12 disinhibitory neurons. As a consequence, inhibitory pools within the chain now play a dual role: firstly, they are inhibited due to the disinhibitory pathway in order to serve disinhibition to the excitatory pools, and soon after they receive the actual signal so they can provide disynaptic inhibition to the same excitatory pools. As

a result, the excitatory neurons within a chain are expected to be protected from the excessive hyperpolarisation before the arrival of the signal so that they can respond to it as expected.

Another possibility to include the disinhibitory pathway is to build it independently of the already existing connectivity in the chain as shown in Figure 4.4. In that way, the inhibitory pool would be released from its dual role in providing both disinhibition and feedforward inhibition, as a dedicated pathway comprising two inhibitory pools would realise disinhibition exclusively. Both dedicated and shared pathways were tested and qualitatively, they gave similar results, with the dedicated variant being slightly more effective. It is argued, however, that this variant recruits unnecessarily too many inhibitory neurons per layer (one pool for disynaptic inhibition and two pools for disinhibition) and thus was not further examined.

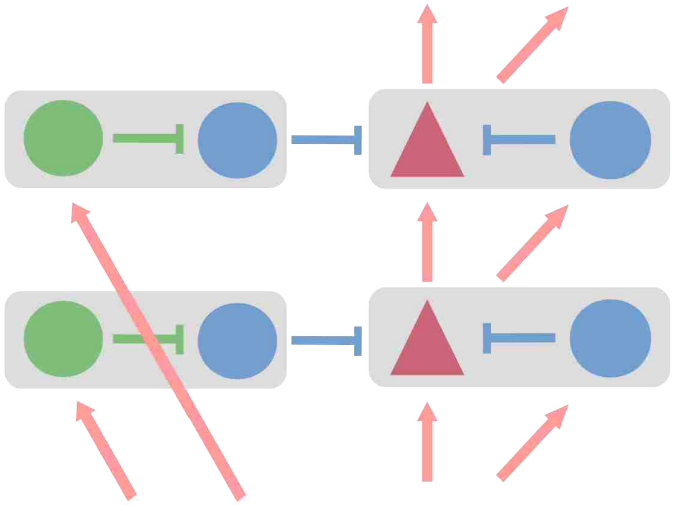

Figure 4.4: Dedicated disinhibitory pathway added to the chain. Separate inhibitory pools implement disynaptic inhibition and disinhibition. This variant was not examined due to a large number of inhibitory neurons recruited.

The results from the simulations involving a chain with a shared disinhibitory pathway are shown in Figure 4.5. Surprisingly, not only the halo failures were removed, but also this setup did not lead to a single synfire chain explosion. It should be noted that the disinhibitory pools not only specifically connect to the inhibitory pools within the chain, but they also maintain a global 5% connectivity within the whole network. Thus, when activated, they also inhibit the background neurons which would otherwise fire and could potentially cause synfire explosions.

As shown in the top subplot in Figure 4.5, the individual traces and their mean trace

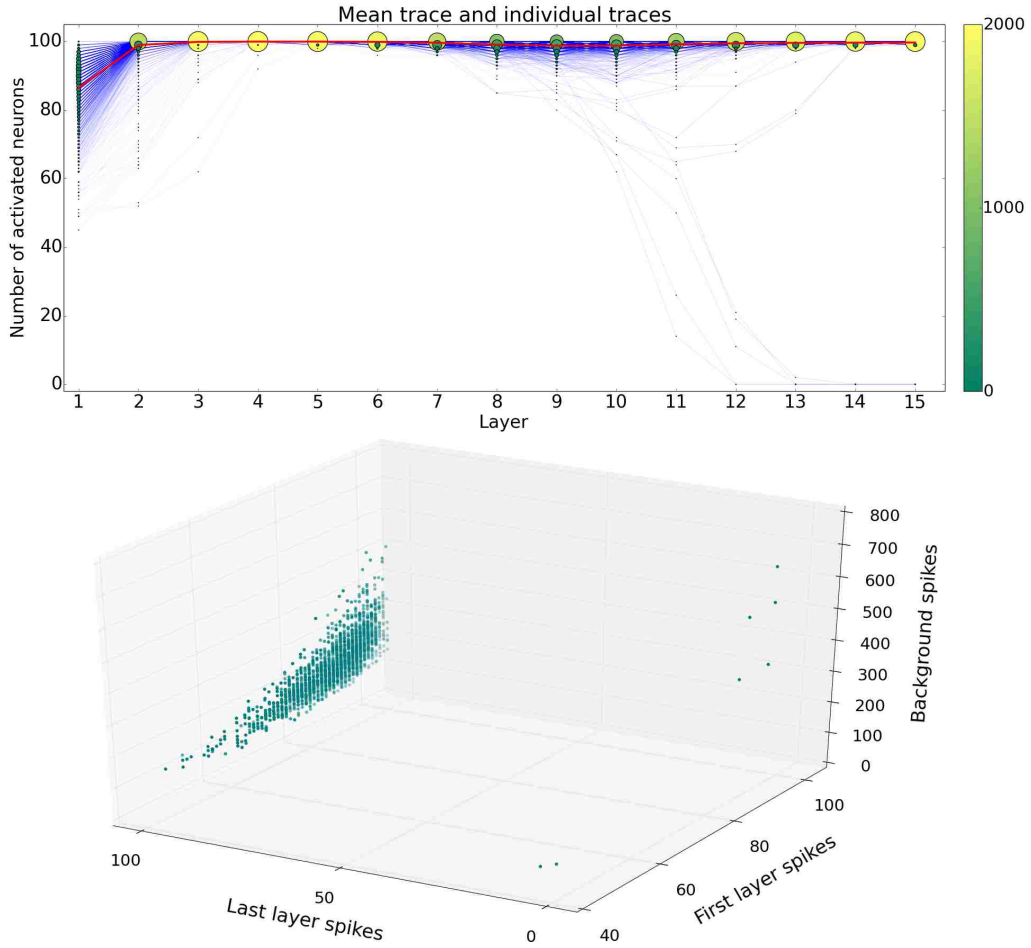

Figure 4.5: Signal transmission after adding disinhibitory pathway to the chain. Top: Individual traces (blue) and the average trace of the signal propagation along 15 layers (red) for 2000 trials. Colour and size of the circles correspond to the number of trials. Bottom: Relationship between the first layer, last layer and the background activity.

differ from the ones in the case without the disinhibitory pathways. Most importantly, much fewer trials failed to propagate past the 7^{th} layer. Around that area, many traces nevertheless decreased the number of their active neurons, but towards the end of the chain, the signal again got stronger to fire with 100% of neurons. This behaviour demonstrates the core effect of disinhibition – the neurons indeed got hyperpolarised due to the inhibitory halo, some of them even failed to fire, but overall, enough neurons managed to withstand this wave and the signal did not diminish. Disinhibition takes advantage of the increased global inhibition and transforms it into a power that locally cancels out the inhibitory control held over the excitatory neurons.

The bottom subplot in Figure 4.5 displays the relationship between the chain's success and the background spikes. All the successful trials invoked less than 700

spikes in the background, the value chosen to demarcate the borderline between the halo and explosion failures.

From this point forward, a chain without disinhibition will be referred to as a *basic chain*, whereas the chain with the shared disinhibitory pathway will be referred to as a *guarded chain*, to highlight the protective role of the disinhibitory pathway against the global instabilities.

4.4 Embedding two non-overlapping chains

Having established that firstly, signal propagation in a random network is possible and secondly, that specific disinhibition protects the chain and the whole network from synfire chain explosions and thus improves the overall success rate, the focus was placed on the possible interactions between multiple chains embedded in the same network.

To study this case, two chains were embedded in the network and the experimental procedure remained the same, that is a volley of spikes was injected only to one chain and this chain's activity was analysed. When analysing the background activity, both non-chain neurons as well as the neurons belonging to the second, inactivated chain, were included.

Initially, the two chains were non-overlapping. Both, basic and guarded chains were considered and the comparison between the conditions with one and two chains embedded is shown in Figure 4.6.

Surprisingly, the addition of an extra basic chain improved the signal propagation – 95% of trials were successful, 5% led to explosions and 0.5% trials failed due to the inhibitory halo (as compared to 86%, 10% and 4% respectively in the one chain condition). The background activity during the signal transmission also decreased at the presence of the second chain. For the guarded chain conditions, the difference was minimal as in both setups the success rate was excellent and the average number of background spikes relatively low.

The reason why the addition of the second basic chain improved the signal transmission by eliminating both, explosion and halo failures is the fact that adding non-random connectivity implicitly removes the random connectivity. Globally, the connectivity is maintained on the 5% level and the neurons belonging to the chain have a large proportion of their incoming and outgoing connections 'locked' along the chain. Thus, when only one chain is embedded, the signal along the chain activates the back-

ground which immediately interacts back with the chain. When the two chains are embedded, however, the feedback loop is longer, as the signals that reach the second chain remain within this chain and only a smaller part of the signal feeds back to the first chain.

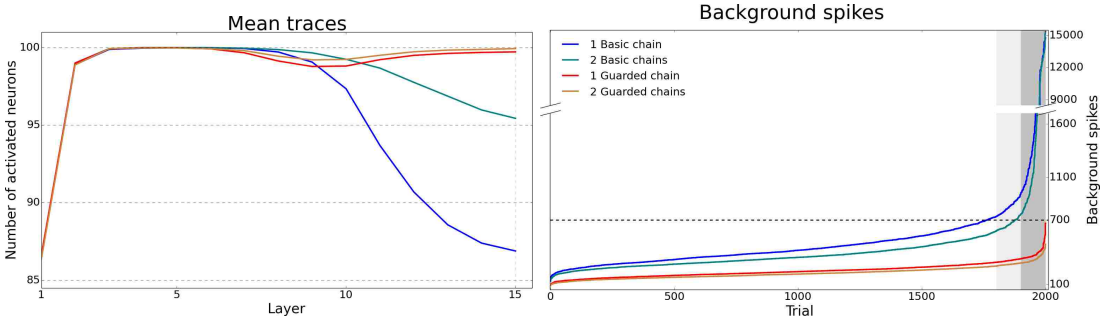

Figure 4.6: One vs two embedded chains for basic and guarded conditions. Left: mean traces. Right: Background spike counts for all 2000 trials per condition, sorted in ascending order. Y-axis is not continuous to capture the average distribution as well as the extreme counts. Bright and dark grey shading denote the top 10% and 5% trials (200 and 100). Dashed line along the 700 spikes denotes the borderline between the halo and explosion fails.

4.5 Embedding two overlapping chains

It is desirable that neurons in networks are flexible and take part in encoding or computing of more than just one entity or task. If neurons were able to encode only one entity, networks containing large cell assemblies would have a very limited capacity which would render them rather inefficient computational units. The problem of networks' capacity to embed multiple synfire chains was previously studied (Trengove et al., 2013), where the symmetry between the excitatory and inhibitory overlaps was assumed. Here, the network's capacity is out of scope and the focus is placed solely on the excitatory and inhibitory overlaps between the two embedded chains. Since the excitation and inhibition have distinctive functions in the network, it was hypothesised that excitatory and inhibitory overlaps will have contrasting influence on the signal transmission and network's stability and thus should be studied in separation.

Both basic and guarded chain conditions were considered. The grid plots in Figure 4.7 display the signal transmission success rates depending on the levels of excitatory (up to 50%) and inhibitory (up to 100%) overlaps. Each square displays the success rate for a given setup obtained from 2000 trials.

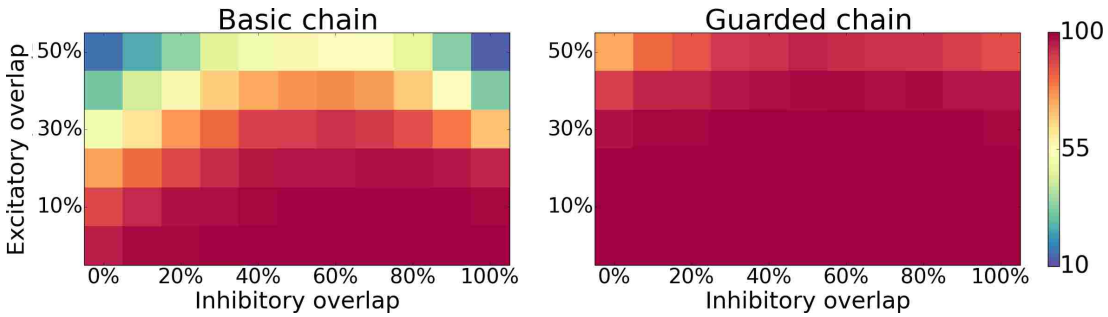

Figure 4.7: Success rates for basic chain and guarded chain conditions for the whole range of excitatory and inhibitory overlaps.

The first observation is a clear difference between the basic and guarded chain conditions. Although qualitatively similar, quantitatively, the guarded chains proved to be more resistant against the disturbances caused by the overlapping non-random pathways. The grid plot in Figure 4.8 depicts the difference in the success rates after the addition of the disinhibitory pathways. In the extreme case, disinhibition improved the success rate by 75%, and on average – by 19.5%.

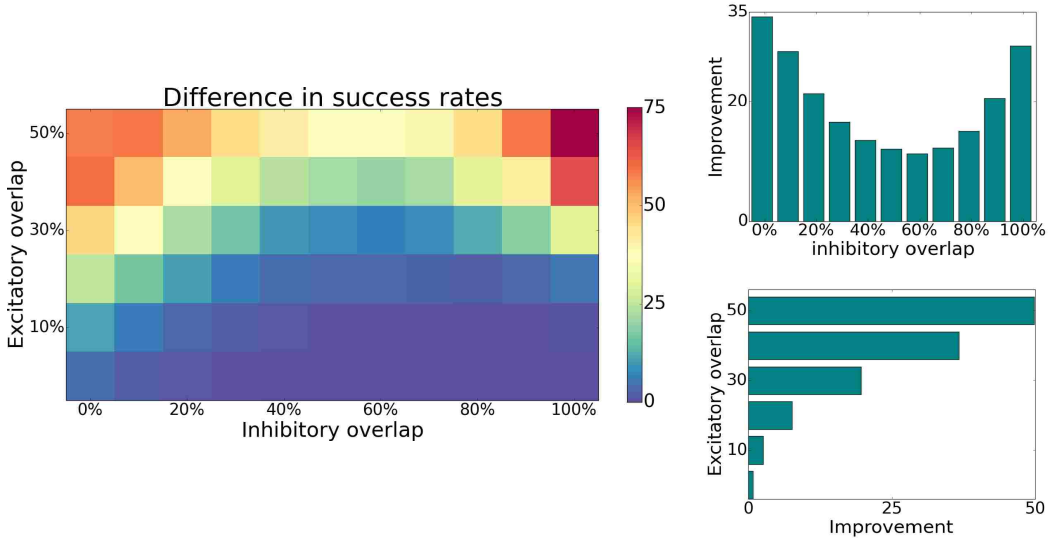

Figure 4.8: Difference in the success rates after adding disinhibitory pathways. Left: a grid plot of the differences (guarded minus basic chain condition). Right: the average improvement for each level of inhibitory (top) and excitatory (bottom) overlap.

Regardless of the presence of the disinhibitory pathways, an asymmetry between the excitatory and inhibitory overlaps is evident. Increased excitatory overlap (EO) impairs the signal transmission at all times, whereas the effect of inhibitory overlap (IO) is more complex as it depends on the level of EO, but in most cases inhibitory

overlap is beneficial for the signal transmission.

Excitatory overlap has very severe effects when applied on its own (first column in the grid plots). For the basic chain condition, two non-overlapping chains setup yields a high success rate of 95%, 10% overlap decreases the rate down to 89%, 30% – 51% and 50% EO causes the success rate to go as low as 16% and the main cause of the failure are the synfire chain explosions. When the inhibitory overlap is applied on its own (bottom row in the grid plots), the success rate is very close to 100% for all the levels of IO. When both overlaps are at their maximum, the success rate is extremely low and its main cause are the halo failures. The forthcoming sections will explore in more detail the relation between the combination of overlaps and the explosion and halo failures occurrence.

The above observations firmly confirm the hypothesis that the excitatory and inhibitory overlaps should be treated as two individual factors. They play different roles in the circuits and the combination of two can affect the network immensely. Were the level of overlaps treated as a single feature, then by looking at the diagonal of the grid plots in Figure 4.7, one would conclude that in the basic chain condition, the setups with more than 20% overlap are unstable. This could lead to the conclusion that network's capacity is rather low as it cannot remain stable with highly overlapping assemblies. Releasing the symmetry assumption opens up the possibility that even without disinhibitory pathways, the chains with the substantial excitatory overlap can still robustly facilitate signal transmission, provided that the inhibitory overlap is considerably higher.

Remarkably, inhibitory overlap implicitly realised the motif of lateral inhibition – when one chain is active, the pool of shared inhibitory neurons prevents the other chain from firing. As discussed in Section 3.4.2, the overlaps were set on the level of the whole chain, not the individual layers. Thus, in the case of two chains with 100% inhibitory overlap, activating the first layer of the first chain did not activate the first layer of the second chain via the overlap. Instead, due to the random allocation, these activated neurons could potentially reside in any layer in the second chain, so that every layer received a part of the inhibitory signal. Then, when the signal propagated further along the first chain, all the layers in the second chain would evenly receive the parts of the inhibitory waves created by the activated chain. In that way, the entire chain was under the inhibitory control throughout the whole activation time. However, such excessive inhibitory control via the 100% overlap was the main cause of the low success rate in the setup with 50% excitatory overlap. In this case, the neurons not

only received the global inhibitory halo, but also the inhibitory signals produced by the implicit activation of the second, overlapping chain. As a result, the chain neurons were too hyperpolarised to correctly respond the to arriving signal.

4.5.1 Activity of the rest of the network

A closer look at the activity of the rest of the network reveals that the low success rates at the top left corner (high EO, low IO) have different causes than the ones at the top right corner (both EO and IO high). The plots in Figure 4.9 present the background spike counts for all the levels of excitatory overlaps and the 4 levels of inhibitory overlaps (one per plot). Clearly, high excitatory overlap activates the whole network, whereas the inhibitory overlap silences it.

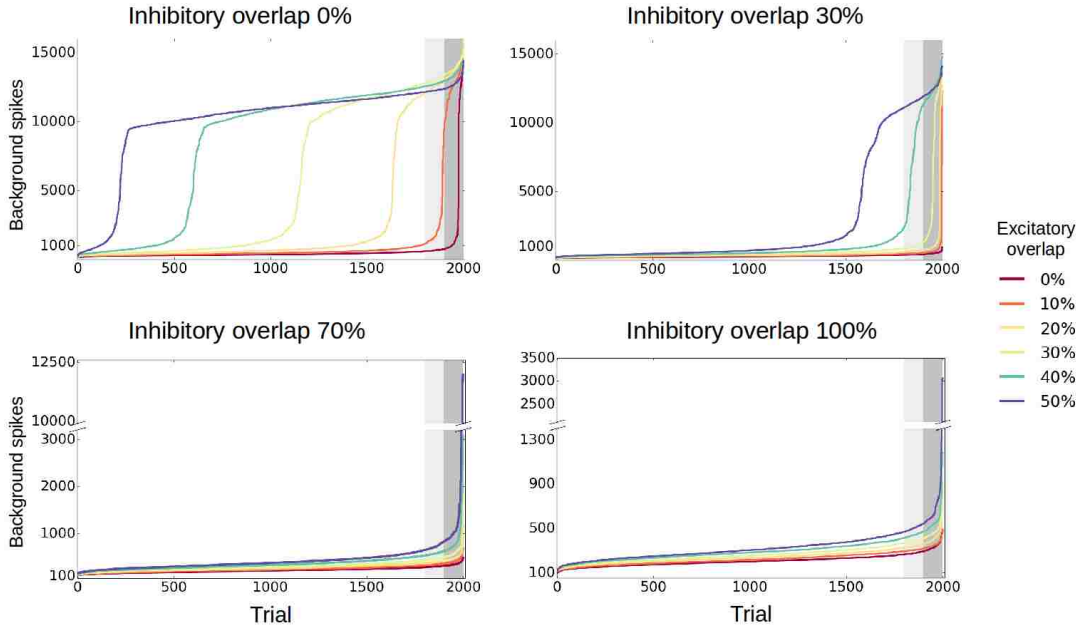

Figure 4.9: Sorted background spike counts in the basic chains condition for 4 levels of inhibitory and all levels of excitatory overlaps. Note the different scales and maximal values on the y-axis. Excitatory overlap causes a strong activation of the whole network, while the inhibitory one silences it. Even for the 100% IO there are still explosion failures, but with much smaller number of spikes, compared to lower IO levels.

Excessive activation and silencing of the network are the signatures of the explosion and halo failures. The grid plots in Figure 4.10 display the probability of occurrence of them depending on the condition and the overlap levels. As expected, excitatory overlap that is not counteracted by the IO disbalances the network and its

activity explodes. When the high EO is accompanied by a much higher IO, the whole network and the chain get silenced, promoting the halo failures.

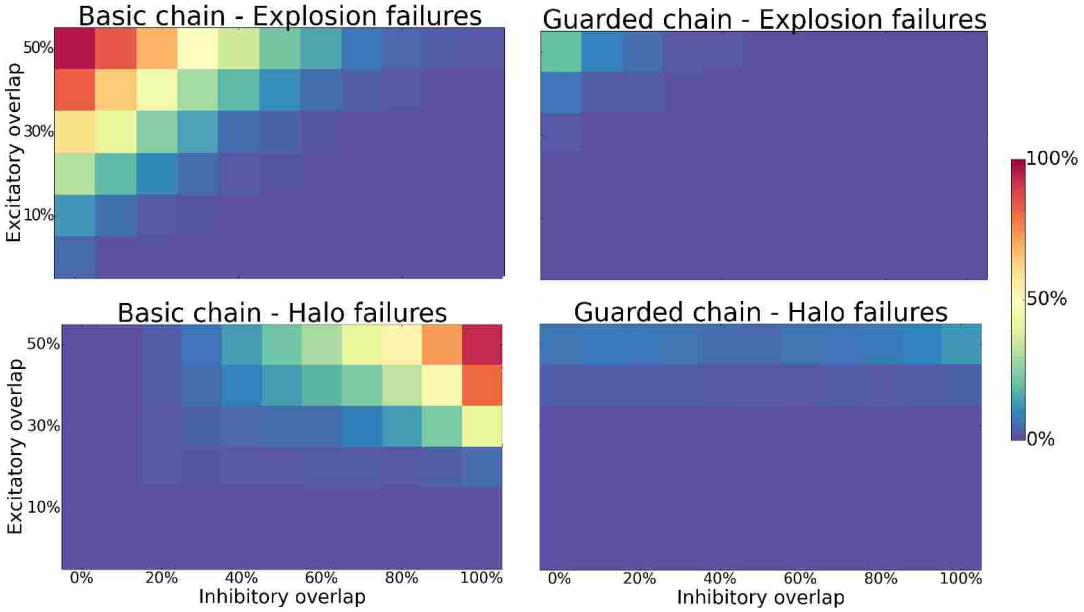

Figure 4.10: Halo vs explosion failures. Top row: explosion failures, bottom row: halo failures for the basic chain (left column) and the guarded chain (right column) conditions.

These grid plots show the average activation of the non-chain neurons, defined from the perspective of the active chain. In fact, those non-chain neurons are either the background neurons or the ones belonging to the second, not activated, chain. In principle, the excessive activation of both neuron classes is detrimental to the signal transmission and network's stability. Since the second chain is a cell assembly encoding some information and composed of highly non-random circuits, it is useful to analyse its behaviour separately. With the non–zero excitatory overlaps, the second chain naturally becomes partially activated, but the question is what happens to the remaining, non-overlapping part of the chain. If it also fires, then the total amount of activation might reach the levels where this chain should be labelled as active, not silent (although the individual gates would not necessarily fire in a cascades-like fashion, so the temporal patterns would be scrambled). Nevertheless, this is not a desirable setting, as an activation of one chain is not supposed to trigger yet another assembly as this corrupts the overall signal and recruits too many local resources for a single signal transmission. An ideal setup would cause the signal-carrying neurons to activate, while keeping the rest on hold. The plots in Figure 4.11 explore the behaviour of the second chain during the 50ms time window after the signal injection into the first,

active chain. It was assumed that due to the intrinsic randomness, even a stable setup can at times switch to an unstable state via the synfire chain explosions. In order to account for that, for each setup the 5% of the trials (100 out of 2000 trials) with the highest values reached were removed to omit the clear outliers that would distort the final average number.

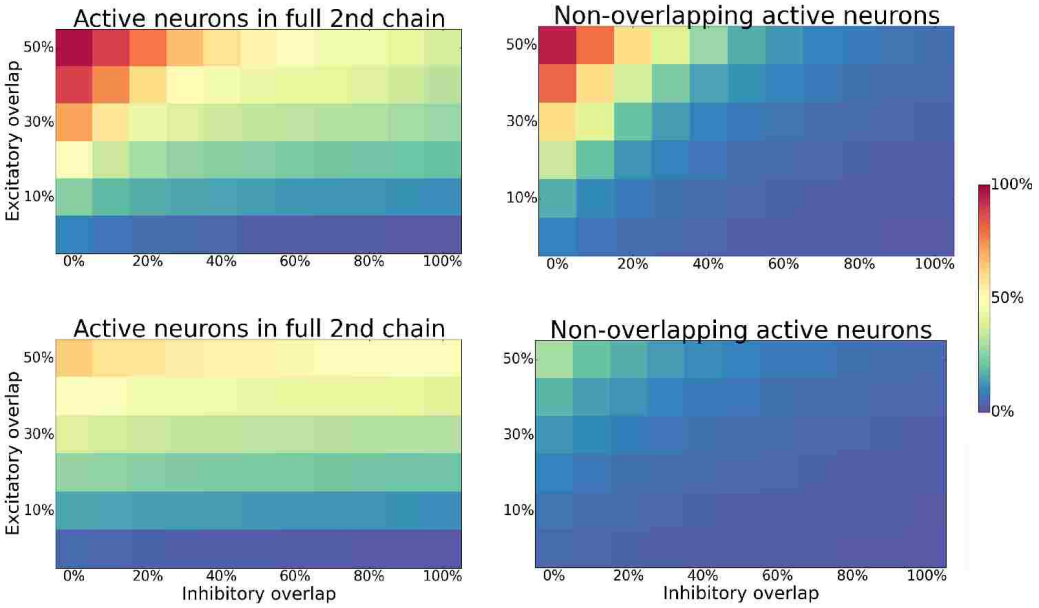

Figure 4.11: Second chain activation during the signal transmission. Top row: basic chain, bottom row: guarded chain condition. First column: the percentage of active neurons in the whole second chain including the overlap. Second column: the percentage of active the non-overlapping neurons.

Top plots in the left column in Figure 4.11 show the levels of activation of whole second chain (top: basic, bottom: guarded chain condition). The horizontal stripes reflect the fact that the overlapping neurons got activated via a signal traversing the first chain. On the left, only the non-overlapping neurons are shown. Overall, only in the setups with the prevalence of synfire chain explosions the activation of the second chain turned out to be very pronounced. High inhibitory overlap effectively silenced the non-overlapping neurons in the second chain.

4.6 Evolution of a pulse packet along the chain

To quantify the degree of synchrony of a signal propagated along a chain, a term *pulse packet* was introduced (Diesmann et al., 1999). A pulse packet, similarly to the volley

of spikes used as an input, is characterised by its activity a – a number of spikes; and σ – temporal dispersion, measured by the standard deviation of the underlying pulse density. These two variables form a 2-dimensional state space, commonly used to study the pulse packet dynamics, where the evolution of synchronous activity along the layers of a chain is plotted as a trajectory. In studies of isolated synfire chains, it was shown that there are two fixed points in such space: an attractor and a saddle point (Figure 4.12 left). A separatrix running through this saddle separates the state space into two regimes. In the basin of attraction, all trajectories converge into the attractor, which denotes the successful signal transmissions. Packets starting their evolution outside of this basin, thin and die out already after a few layers of the chain.

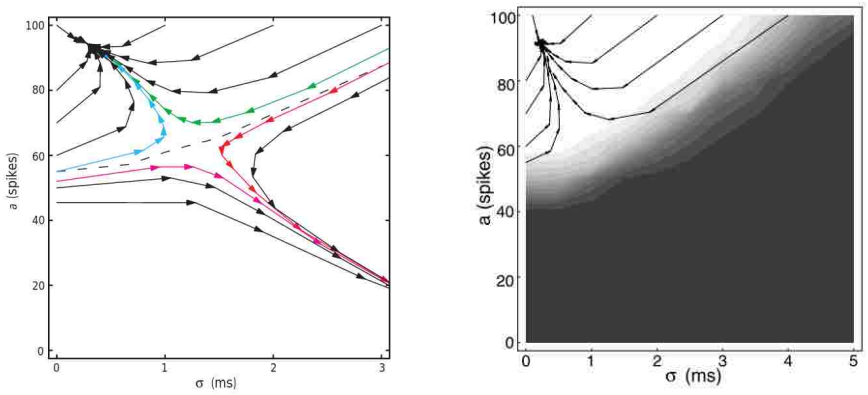

Figure 4.12: State space portrait of synfire chain activity. Left: state space for an isolated chain. Blue and green trajectories are successful propagations, red and pink trajectories are failures. Between the two classes there is a separatrix. Figure from (Diesmann et al., 1999). Right: state space for an isolated chain with a strong background input. The space is covered by a gradient from white (100%) to black (0%) denoting the packet's survival probability. Figure from (Gewaltig et al., 2001).

The main observation here is that the fate of the pulse packet is binary and determined by its initial conditions defined by only two variables. The very shape and position of the separatrix and the basin of attraction are determined by the network's and neurons' properties, but once they are set, the pulse packet's trajectory will depend solely on the a and σ. This behaviour, first shown by Diesmann et al. (1999), was later formulated by a simplified mathematical approach applying the Fokker-Planck equations (Câteau and Fukai, 2001). Subsequent studies investigated an isolated chain with a strong background input mimicking the host network (Gewaltig et al., 2001) to reveal that the system is no longer deterministic. Along the separatrix the fate of pulse

packets is best characterised by a gradient of survival probabilities (Fig. 4.12 right).

4.6.1 State space analysis for basic and guarded chain conditions

Here, the goal was to apply the state space analysis to the behaviour of a pulse packet travelling along a chain that is no longer isolated, but embedded in a random network.

It should be remarked that the variable a that measures the number of spikes per layer, does not explicitly specify whether the spikes come from different neurons or whether some neurons do emit more than one spike. This issue was briefly investigated and it was concluded that the event of multiple spikes emitted by one neuron is extremely rare and equating the variable a with the number of active neurons should be deemed acceptable.

Visualising the state space spanned between the a and σ helps to gain more insight into the pulse packet's evolution in embedded chains as demonstrated in Figure 4.13.

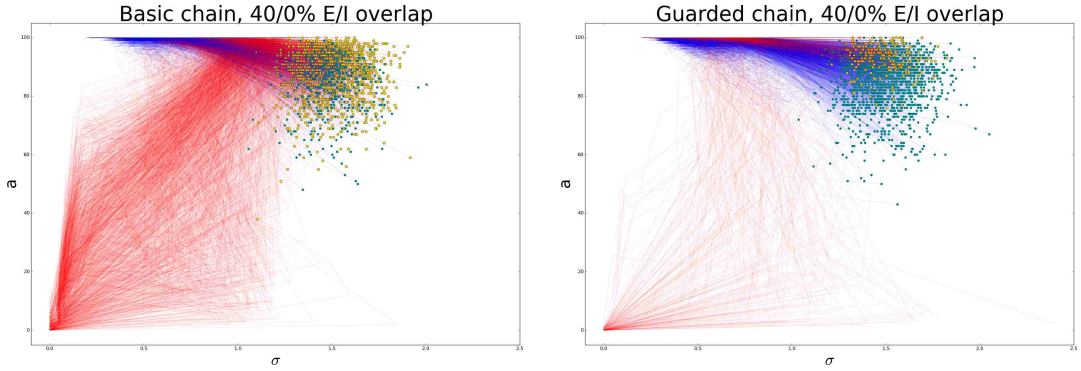

Figure 4.13: Examples of state spaces for basic and guarded chain conditions. Failed (red) and successful (blue) trajectories are marked with starting points (yellow and teal circles respectively). There is no separatrix that could separate the two classes.

First of all, the colour-coding of failed and successful trajectories reveals that it is impossible to find an attractor and a separatrix between the two classes of trajectories. They all originate in a similar area, travel to the same place (which corresponds to the attractor in the isolated chain condition – high a and narrow σ), but then the failed trajectories are pushed away from that area to pave their way to the 'failure' region where they terminate. As expected, the system is stochastic and the variables a and σ cannot alone predict the fate of a pulse packet.

4.6.2 The shapes of trajectories for all combinations of overlaps

As already demonstrated, success rates as well as the probability of halo and explosion failures occurrence depend on the levels of overlaps. It was hypothesised that these dependencies might be reflected in the patterns of the trajectories across the $(a - \sigma)$ spaces. To investigate this, successful and failed trajectories were separately plotted for all the overlap combinations for both, basic and guarded chain conditions and arranged akin to the grid plots for easy comparison. Figure 4.14 shows the successful trajectories for both conditions, whereas Figure 4.15 shows all the failed ones.

Trajectories of successful propagations, on the other hand, look similar to the trajectories inside the basin of attraction (as in Figure 4.12) – they quickly reach the attractor region and remain there. The shapes of failed trajectories vary considerably across the overlap combinations and they clearly reflect the occurrences of the halo and explosion failures. Both types of failures have distinct underlying causes and thus the pulse packet's trajectory is affected in a different manner. The plots also display the outset failures to demonstrate that their trajectories are akin to the ones that originate outside of the basin of attraction in the deterministic setup (as in Figure 4.12).

4.6.3 Successful trajectories

In all the successful trajectories, after just a few layers the activity of a pulse packet increases to 100% and the packet becomes narrower, regardless of the number of spikes in the first layer. Then, the pulse packet either remains that way until the end, or it begins to lose a few spikes and become slightly wider, which is illustrated by trajectories forming various shapes of loops or zigzags, as shown in Figure 4.16. Different levels of overlaps were selected in order to display the distinct trajectory shapes.

These shapes are in fact a signature of an influence of the inhibitory halo: near the middle of the chain, the neurons get hyperpolarised due to the global inhibitory halo and some of them fail to respond to the pulse packet activation. Because the basic chains lack the protection against the halo, they usually formed irregular trajectories with their endpoints scattered. Guarded chains, on the other hand, formed mostly the trajectories with loops and the endpoints concentrated near the maximal pulse packet's activity.

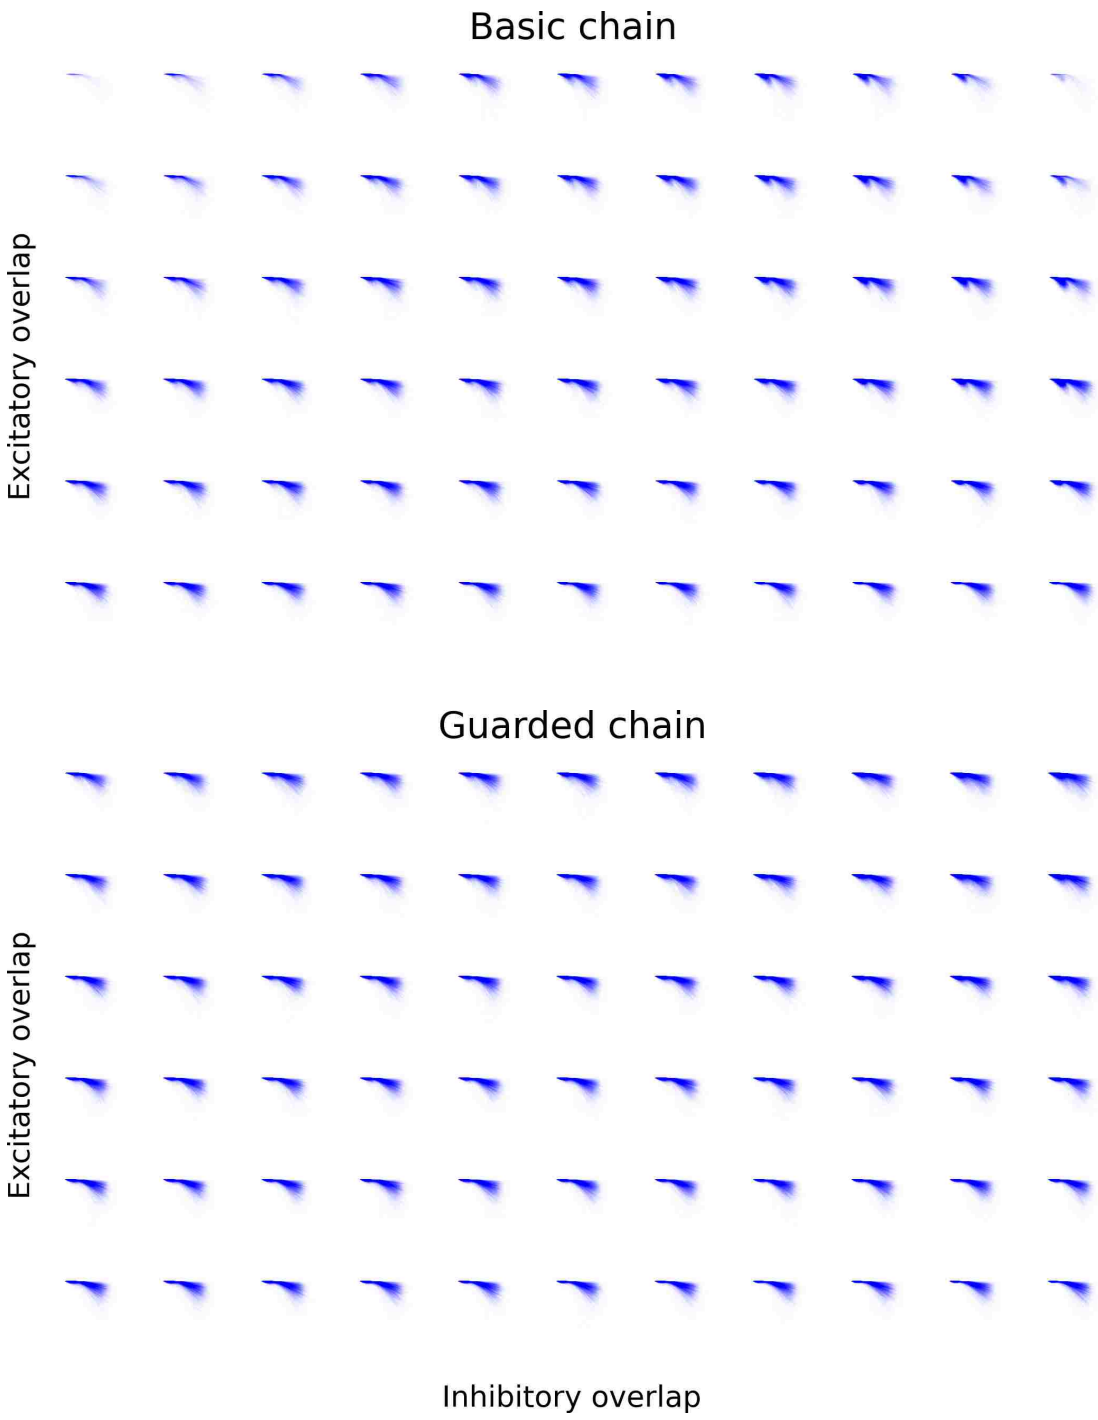

Figure 4.14: State space diagrams for the basic (top) and guarded (bottom) chain conditions: trajectories of successful propagations for all the overlap combinations.

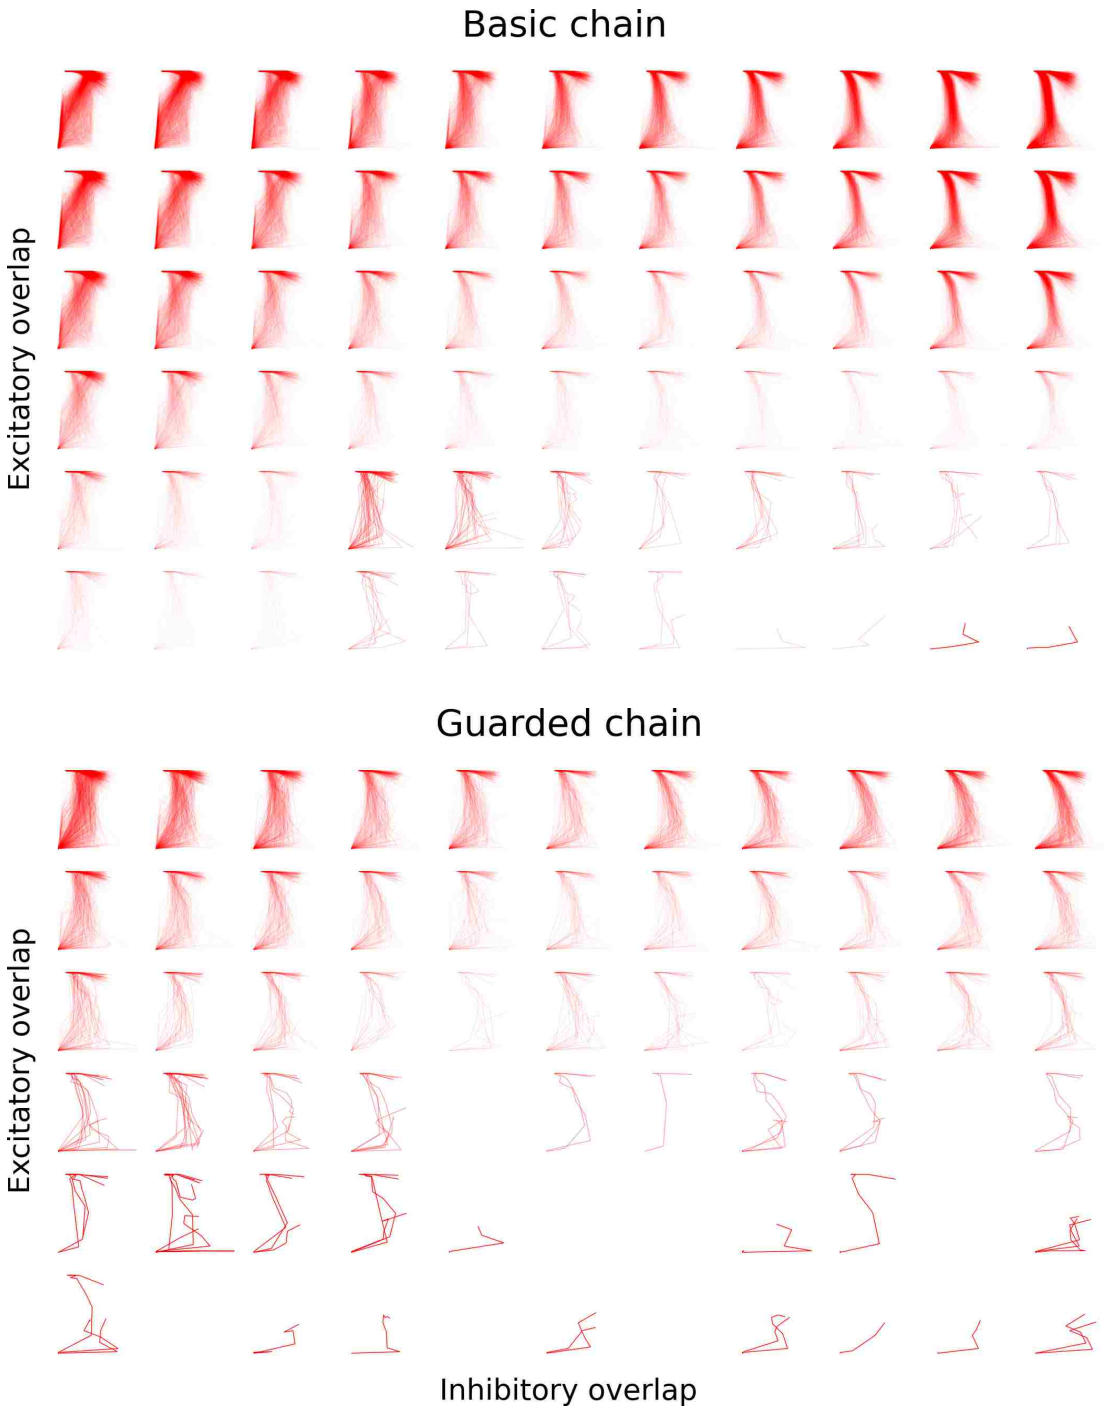

Figure 4.15: State space diagrams for the basic (top) and guarded (bottom) chain conditions: trajectories of failed propagations for all the overlap combinations. In some subplots individual trajectories were highlighted for better visibility. Note the presence of the outset failures (bottom right corner for the guarded chain condition).

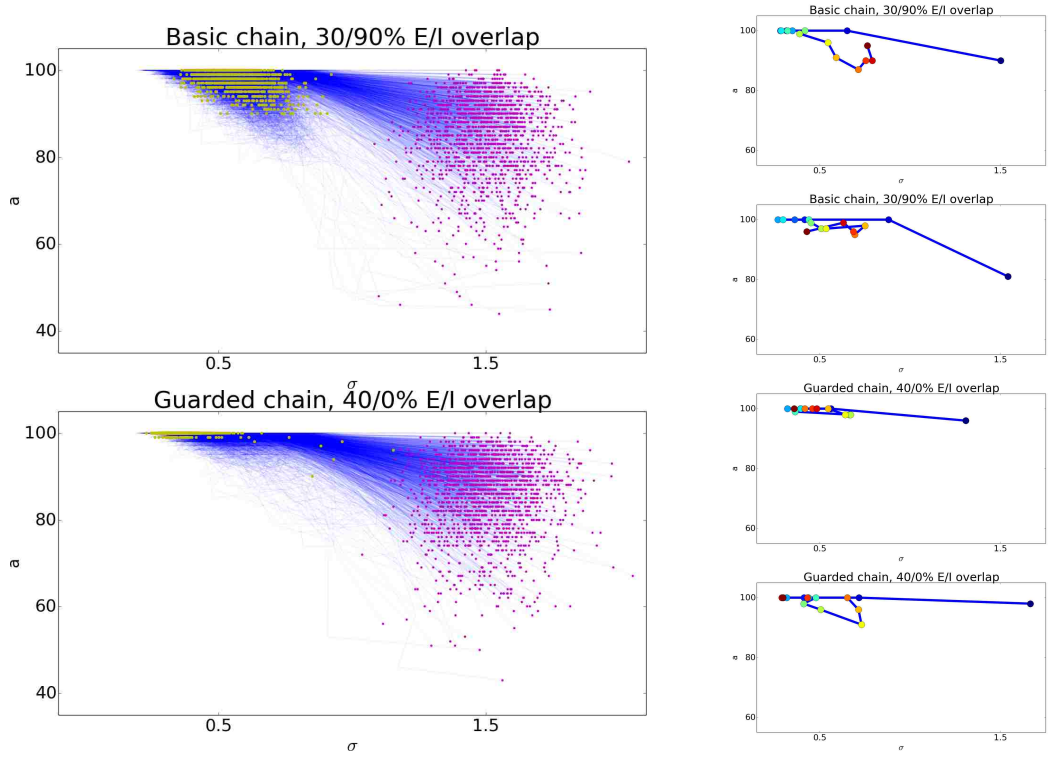

Figure 4.16: Successful trajectories. Left: Two examples of setups with different trajectory silhouettes. Purple and yellow dots denote the start and end points respectively. Right: Individual trajectories pooled from the setups on the left. Subsequent chain layers are marked by colour-coded circles: first layer is marked by a dark blue circle and the end one by a brown circle via the rainbow gradient.

4.6.4 Halo failure trajectories

The trajectory of a pulse packet that failed to propagate due to the inhibitory halo is characterised by the direction it follows when being pushed away from the attractor area (top left corner). Both parameters are initially affected simultaneously – a descends and σ grows and the trajectory follows the diagonal until it reaches the bottom area of the plot. The resulting silhouette forms an inverted Z–shape. As shown in Figure 4.17 left, some trajectories terminated all the way along the diagonal, but the vast majority of trajectories ended near the bottom left corner marked by $a = 0$ and $\sigma = 0$. This fact is captured in the plot by many red lines directed at the several dots at that corner.

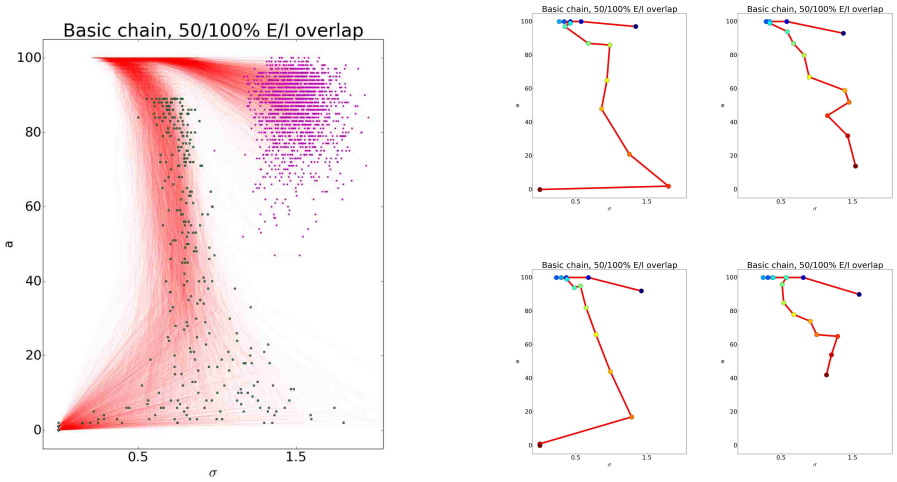

Figure 4.17: Trajectories of halo failures. Left: Halo failure trajectories for the Basic chain with 30/0% E/I overlap. Right: individual trajectories. Subsequent chain layers are denoted by colour-coded circles: first layer is marked by a dark blue circle and the final one by a brown circle via the rainbow gradient.

4.6.5 Explosion failure trajectories

Similarly to the halo failures, trajectories of explosion failure propagations can also be identified by their behaviour when being pushed away from the attractor area. As expected, at the outset of the explosion, the number of spikes per gate remain 100%. Only the σ becomes wider as illustrated by the first half of the trajectory – it is first reaching the attractor corner and then heading rightwards in a straight line. Then, the number of spikes sharply decreases and the σ becomes narrow again, causing the trajectory to move via diagonal towards the bottom left side where it terminates. The overall silhouette forms a 7–shape and the distribution of the end points reveals some signs of regularity. One can identify two sharp spires in the silhouette – a horizontal one at the top, which is a signature of the outset of the explosion, and the second, vertical one - at the bottom left corner, along which many trajectories terminated. One might also notice a hyperbolic curve starting from the tip of the lower spire towards the bottom, also dotted with many trajectory end points.

These regularities created by the trajectories endpoints are in fact an artefact of a pulse packet detection algorithm. The rightmost plots in Figure 4.18 depict the elements used in the algorithm. Starting from the first layer at the bottom, the shaded area marks the time window where the spikes were counted and the red line denotes the pulse packet's mean spike time. This mean is then used as a starting point for the

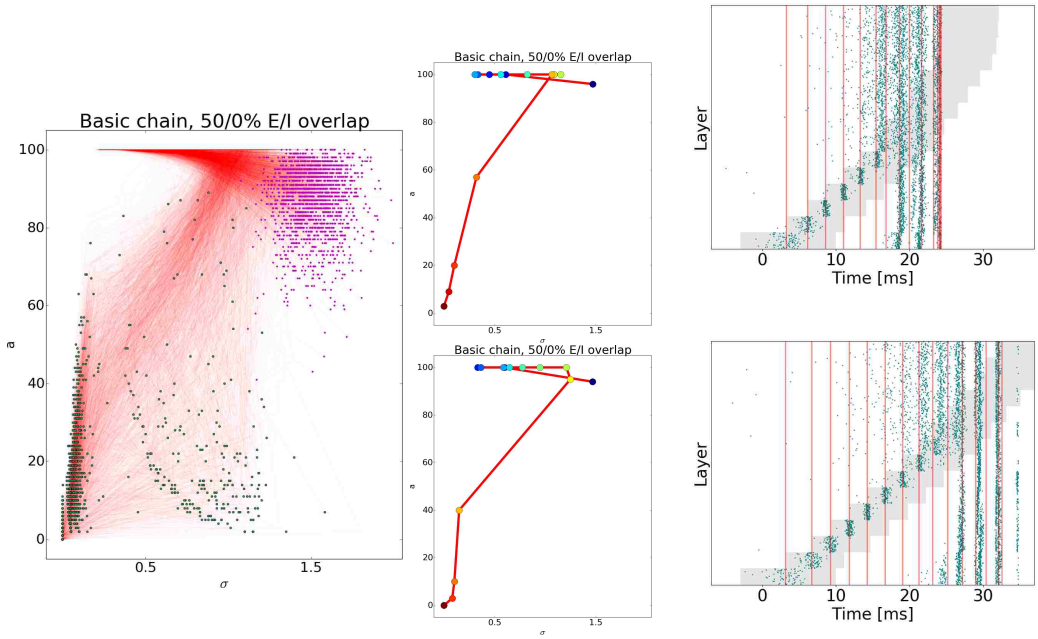

Figure 4.18: Trajectories of explosion failures. Left: explosion failure trajectories for the Basic chain with 50/0% E/I overlap. Middle: individual trajectories. Subsequent chain layers are denoted by colour-coded circles: first layer is marked by a dark blue circle and the final one by a brown circle via the rainbow gradient. Right: depiction of a pulse packet detection algorithm. Starting from the bottom (first layer) a shaded rectangle marks the area where the spikes were counted to form a pulse packet. Red lines denote the centre of mass for each layer.

time window for the subsequent layer. The algorithm was shown to work well with the not-exploding cases and different sizes of the time window showed to have only a minimal effect on the packet's final parameters: a and σ. When the explosion appears, however, the spikes in different layers appear simultaneously, and the mean spike times (the red lines) are very close to each other and the time windows of consecutive layers are nearly the same. As a result, the last layer, which defines the trajectory endpoint, can capture one or two waves of explosion composed of multiple spikes appearing nearly simultaneously. If one wave is captured (Figure 4.18 top right), the trajectory endpoint would lie along the vertical spire, as there are many spikes, but the σ is close to zero. If two waves are captured (Figure 4.18 bottom right) - the σ is larger and the endpoint lies along the hyperbolic curve.

If the algorithm was changed to differently treat the overlapping means (for example stopping as soon as the explosion is detected and zeroing all the following values)

or the borders of the time window redefined, the final results would differ quantitatively. Since the very appearance of synfire explosion is highly undesirable, there is no need to elaborate on how to count spikes in pulse packets when the network enters the explosion phase. More important is that with the current algorithm this state has a clear signature and the further details of it are negligible.

4.7 Can one predict pulse packet's fate?

Although it is apparent that a clear separation between the failed and successful trials is impossible on an $a - \sigma$ state space diagram, one could still reason whether there is some other measure that sets the two apart. Since the system explored here is not deterministic, we can ask whether there is some set of features that can be used to predict the probability of packet's success or failure.

By visual inspection one can conclude that the failed trajectories usually start with a high a. In Figure 4.19, every trajectory and its endpoint was assigned a colour based on the number of activated neurons a in the first layer. The plot on the left displays a condition where the majority of trials are failures, whereas on the right – successes. In both cases the failures are dominated by the blue trajectories denoting the trials with large initial a.

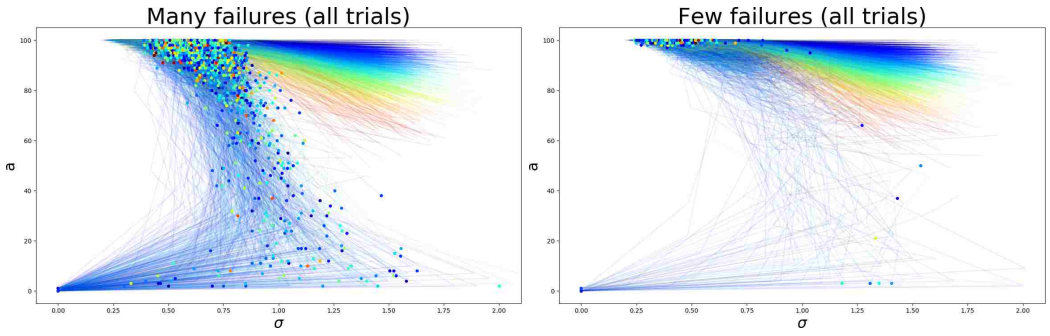

Figure 4.19: Colour-coded trajectories by the number of spikes elicited in the first gate. Coloured dots denote the endpoints of the trajectories.

This intuition was checked for both, basic and guarded chain conditions across all the overlap combinations. In principle, in all the trials, the injected volley of spikes was drawn from the same distribution. Due to the randomness of the volley and the current state of the network, a volley would activate a variable number of neurons to create a pulse packet of a wide range of temporal dispersion. Across all the setups, pulse packet's parameters a and σ in the first layer were similar (data not shown) with the

average a ranging between 85–87 neurons. When the trials were divided into groups of failed and successful propagations, it was revealed that the failed trials contain on average 6 neurons more (Figure 4.20). Every square in the grid plots was computed with the amount of data that was available for a given condition. In some conditions more than 80% of trials were failures, for some only a handful or none failures were available. All the extreme values in the plots stem from such cases.

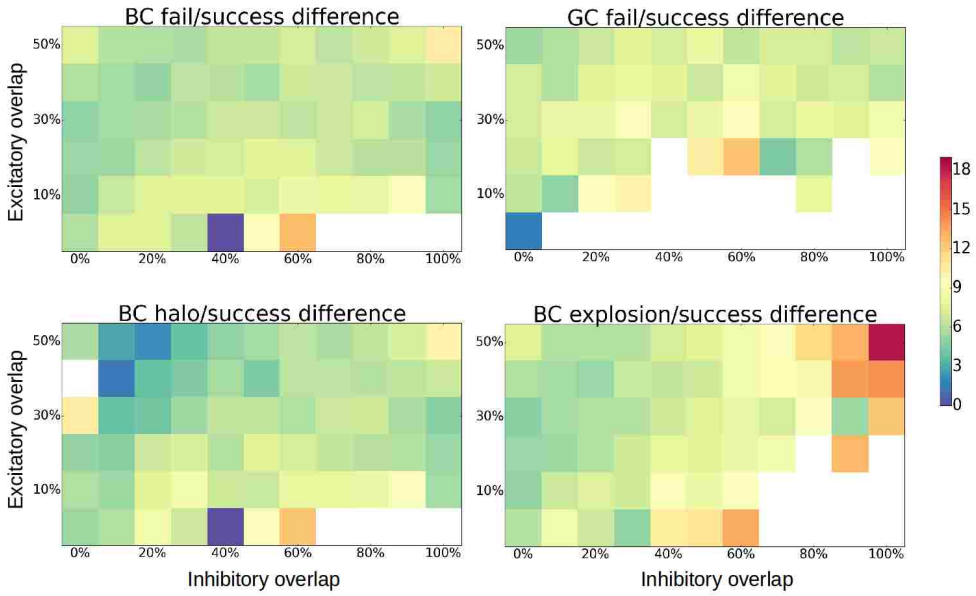

Figure 4.20: Top: Difference in the mean number of active neurons in the first layer between the failed and successful trials for the basic chain (BC, left) and guarded chain (GC, right) conditions. Bottom: Basic chain condition only, the same difference as above, but including only the halo failures (left) and explosion failures (right).

In the halo/success difference plot, there is a darker region on the top left corner, meaning that in these setups the halo failures, on average, were triggered by weaker pulse packet activations. It should be borne in mind that this region was dominated by the synfire chain explosions and not many samples of halo failures were available. On top, this also reflects the instability of these setups – most of the stronger activations led to the explosion and only much weaker, yet marginally stronger than the baseline led to the halo failures. Similarly, in the explosion/success plot, the top right corner features the examples of considerably stronger activations leading to the explosions. Again, this region was dominated by the halo failures and due to the strong inhibitory control, only the activation of particularly many neurons in the first layer could cause the synfire chain explosion.

Although the above reasoning suggests that the failed propagations are the ones which, on average, elicit more spikes in the first layer, the inverse is not always true. Many successful trials also have a high number of spikes in the first layer. The plots in Figure 4.21 display the ratio between failed and successful trajectories depending on the activity in the first layer. As already mentioned, some setups contained a vast majority of failures or successes and overall, the sample sizes were often too small to draw decisive conclusions.

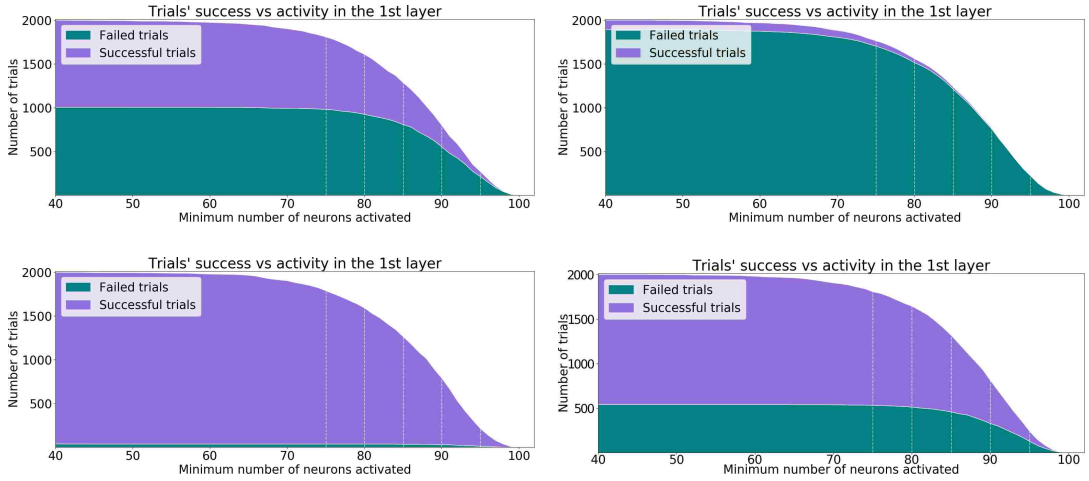

Figure 4.21: Ratio of successful vs failed trials depending on the number of neurons active in first gate for the basic (top) and guarded (bottom) chain condition and four selected combinations of overlaps.

The top left subplot displays a setup with a nearly perfect ratio: 50% of trials were successful, and the other 50% were failed. As the number of active neurons in the first layer is increased, the ratio changes for failures' favour. The trials with 95 or more active neurons are nearly all failures.

Overall, it is demonstrated that when an injected input elicits a strong pulse packet already in the first layer which is rapidly amplified, it is more prone to destabilise the network either by inducing explosion or a global wave of inhibition abolishing the signal. One suggestion to avoid instabilities is to use the inputs which on average do not elicit strong responses in the first layer. But if a weaker input is chosen, the probability of the outset failures increases, so instead of instabilities, the signal fails to fully enter the chain.

4.8 Injecting input to two chains simultaneously

By now, all the setups involved solely one chain receiving the input at a time. Here, we briefly examined a condition where two chains received the signal of identical strength simultaneously. Since injecting two strong inputs into a network is very likely to induce synfire chain explosions, only the regimes that normally are resilient to explosion formation were examined. A small batch of simulations was run, with 5 instances of the network per condition.

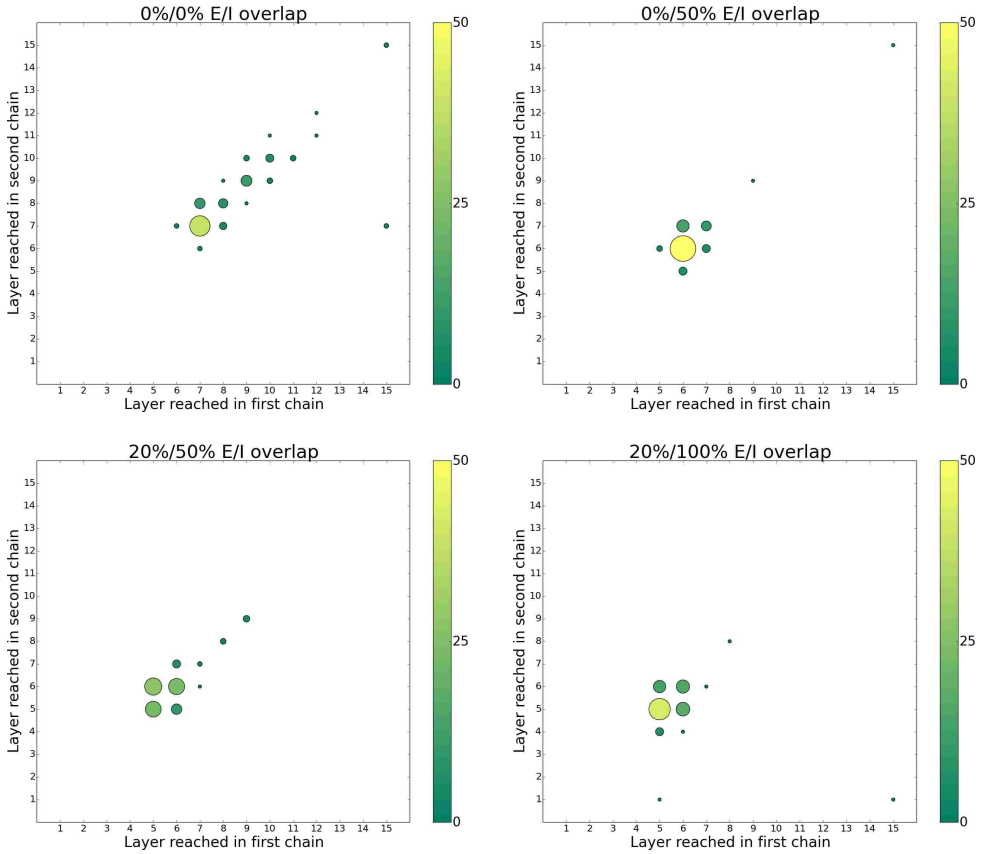

Figure 4.22: Layers reached by the pulse packets along the two chains for four combinations of overlaps. Size and colour of the circles correspond to the number of input injections (out of 100).

As demonstrated in Figure 4.22, most of the pulse packets failed to reach the final layer in either chain. In fact, the signal turned out to die out in both chain at the same stage, which is indicated by the data points lined up along the diagonal. Inhibitory overlap was shown to impair the signal propagation, due to the lateral inhibition reaching both chains.

It is hypothesised that biasing one signal should help to select one chain to permit

the full propagation while suppressing the other one. In order to test it, however, bigger networks are believed to be more suitable since the injection of two pulse packets invokes a large number of synchronised spikes. Thus, the issue of multiple simultaneous signals will not be investigated here any further.

4.9 Splitting the inhibitory overlap

All the data generated and analysed so far relied on the default algorithm to implement the overlaps between the chains, which involved creating a single random pool of neurons to be shared by the two chains which were then assigned to layers and pools independently for each chain. In the case of the guarded chains this simplification introduced a possible problem, since the inhibitory population is essentially composed of two distinct pools: the inhibitory pool of neurons realising feedforward inhibition (FFI) and the disinhibitory pool of neurons along the disinhibitory pathway, as shown in the left panel in Figure 4.23. The default algorithm implicitly allowed for a random pool allocation and as a result, individual neurons would simultaneously participate in the two distinctive circuits realising FFI in one chain and disinhibition in another (right panel in Figure 4.23). It can be well criticised that it is not biologically realistic that a given inhibitory neuron has two different functional roles within multiple assemblies it is a part of. Experimental evidence suggests that various classes of interneurons have well-defined microcircuits and connectivity patterns between each other and with the principal neurons (as discussed in 2.3.3). Furthermore, it is not desirable in the microcircuits to have too much such cross-talks. The chains discussed here have specific connectivity and synaptic delays to serve a tight control over the timing and order of the signals traversing the circuit. Neurons that in one assembly are wired to fire in a given order and in another one in reverse, can cause interference and corrupt the signal transmission.

4.9.1 Specific pool allocation

It was hypothesised that the signal corruption and interference affected the success rates in the guarded chain condition in the setups with both overlaps high (top right corner in the grid plot of Figure 4.7). Admittedly, for these setups the overall success rates were already very high, but still consistently lower than the ones of the setups residing in the middle and the bottom regions of the grid plot. The sole cause of the

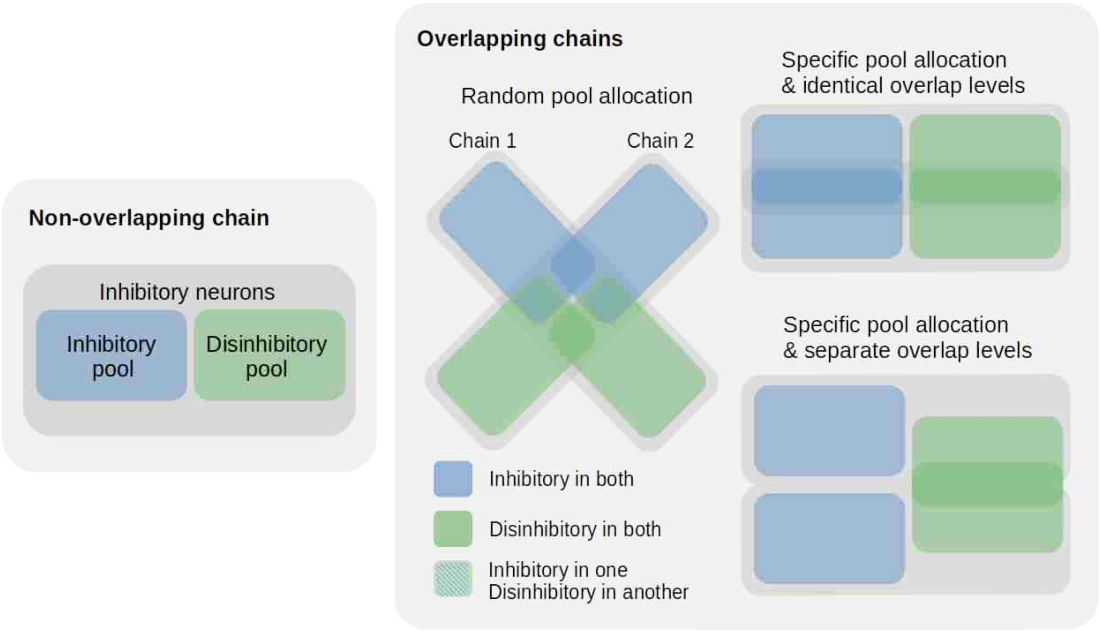

Figure 4.23: Different ways of generating overlaps of inhibitory populations. Left panel: A diagram of an inhibitory population of a single, non-overlapping chain. It is composed of two separate pools: neurons implementing disynaptic feedforward inhibition (FFI) and neurons belonging to the disinhibitory pathway. Right panel: 3 ways of generating inhibitory overlaps between two chains. Left: random pool allocation causes individual neurons to end up in two different pools and are forced to implement both, FFI and disinhibition. Top right: Specific pool allocation makes sure that the shared neurons implement either FFI or disinhibition and never both. The levels of overlaps for both pools are equal. Bottom right: overlap levels for inhibitory and disinhibitory pools can be set separately.

failures were inhibitory haloes, suggesting that the present disinhibitory pathways did not manage to sufficiently protect the chain from the excessive wave of inhibition.

To check this possibility, the algorithm for generating the inhibitory overlap was modified to ensure the specific pool allocation, so that two separate pools of neurons to be shared by the chains were created. First pool contained the neurons to be assigned exclusively to the inhibitory pools realising the FFI, and the second pool contained the neurons to be used to exclusively populate the disinhibitory pools (as illustrated in the right panel in Figure 4.23). Another batch of simulations was ran, but this time it did not span the whole overlap space. Specifically, only the setups with at least 30% excitatory and 10% inhibitory overlap were under examination, because the setups with less than 30% excitatory overlap yielded very high success rates already for both,

the basic and guarded chain condition and thus no effect was expected to be observed.

As shown in Figure 4.24, implementing the specific pool allocation caused only a marginal improvement in the region where the signal transmission was already high (30-40% EO). However, the key observation is that the separation indeed improves the signal transmission in setups with 50% EO and high IO. It proves that when both overlaps are large, separating the inhibitory neurons by their function in a microcircuit is beneficial to the signal propagation. It also demonstrates that the circuits can function more efficiently when interneurons implement only one type of computation at a time, even if they belong to more than one cell assembly.

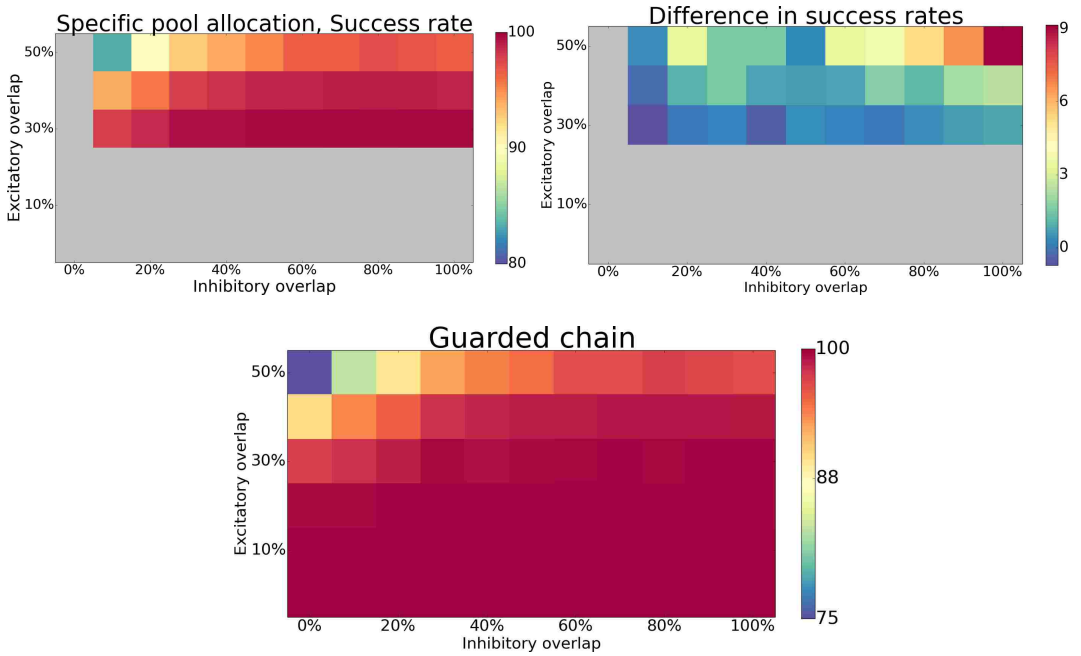

Figure 4.24: Success rates for the guarded chain condition with a modified algorithm creating inhibitory overlaps. Top left: Success rates. First column was omitted as it involved IO of 0%. Top right: Difference between the default and modified algorithm. Bottom: two grid plots combined. Three top rows come from the simulations with the modified algorithm, the three bottom rows come from the default algorithm. Notice different ranges on colour bars chosen for better visualisation.

4.9.2 Specific pool allocation with separate overlap levels

The results above demonstrated that the signal propagation works better if the interneurons are allocated to only one computation – disynaptic feedforward inhibition or disinhibition – when they belong to two cell assemblies. This observation opens up an-

other possibility, not discussed so far – in principle, two chains can maintain three, not two, types of overlaps between each other: the excitatory, the inhibitory and the disinhibitory one. By now, when the 100% inhibitory overlap was reported, it implicitly assumed the 100% overlap in both pools of interneurons comprising the chain. In general, high inhibitory overlap was shown to protect the chain from the synfire explosions and the presence of the disinhibitory pathways helped to overcome the upsurge of the global inhibition. But the effects of treating the two inhibitory overlaps separately, as illustrated in the right panel in Figure 4.23, were not explored.

Two boundary conditions from the top corners in the grid plots were examined: first one involved the setup with 50% excitatory overlap and 0% inhibitory overlap. In the default cases it was shown that the success rate is low and that the main cause of the failure is the synfire chain explosion. The second condition involved again 50% excitatory overlap, but full, 100% inhibitory one which was shown to protect the chain from the explosions and the haloes were the causes of failures. In both cases the EO and IO were fixed and the disinhibitory overlap (DO) varied from 0% up to 100%.

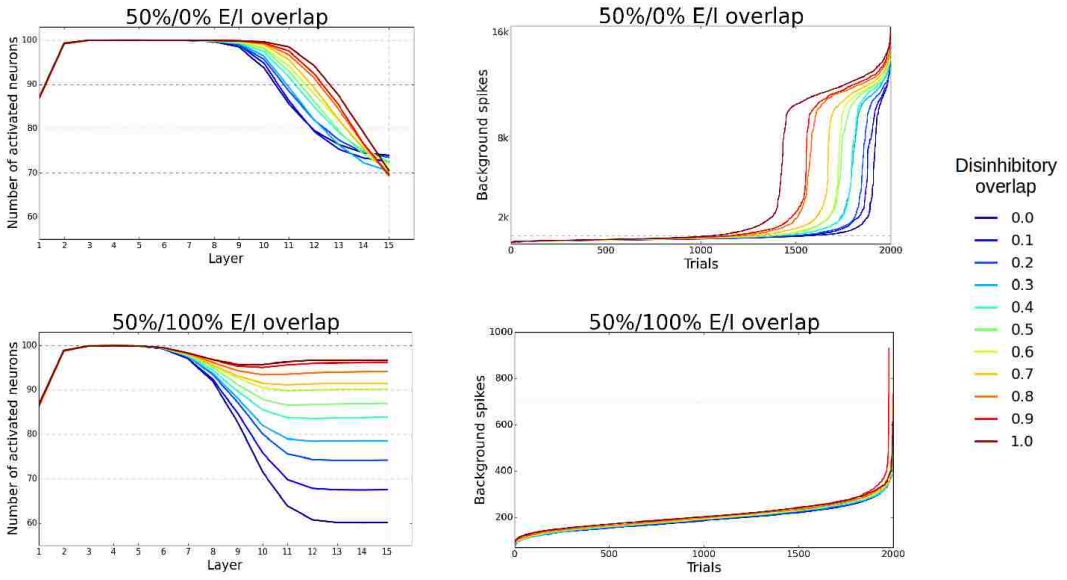

Figure 4.25: Signal propagation (left) and background spikes (right) for different levels of disinhibitory overlaps for the 50% EO and 0% IO (top row) and 50% EO and 100% IO (bottom row).

As expected, higher disinhibitory overlap for the 0% IO case did not help the overall signal transmission (Figure 4.25 top left). This setup was dominated by the instabilities caused by the synfire chain explosions and transforming global inhibition

into disinhibition to release even more excitation not only failed to improve the signal propagation, but also caused more occurrences of the explosions (Figure 4.25 top right). The setup with 100% IO benefited from the higher disinhibitory overlap (Figure 4.25 bottom left) by effectively turning the halo failures into successful trials. Releasing the excitation in the high DO cases, however, led to a few occurrences of mild explosions (Figure 4.25 bottom right).

These results clearly demonstrate the fact, that the disinhibition can destabilise the whole network if too much excitation is being released. In order to safely and robustly benefit from the advantages that the disinhibition offers, the network should be in regimes dominated by inhibition so that releasing extra excitation does not compromise the overall stability.

4.10 Conclusions

The main findings from this chapter can be summarised as follows:

1. It is possible to create long chains composed of layers of excitatory and inhibitory populations realising a principle of disynaptic feedforward inhibition in a fully recurrent network. When the network remains in a low firing regime, synfire explosions appear only rarely.

2. Multiple chains can be embedded into a random network and they can be overlapping or non-overlapping. Excitatory overlap is detrimental to the signal propagation and network's stability (synfire explosions), whereas inhibitory overlaps can be as high as 100% and typically improve the signal success rate via the lateral inhibition motif.

3. Addition of disinhibitory pathways into a chain stabilises the network during the signal transmission and rescues the signal from dying out when global inhibition is strong. Conversely, disinhibitory pathways can destabilise the whole network in regimes with weak inhibition by releasing too much excitatory activity.

4. External stimulus should target both excitatory and inhibitory local neurons not to disturb the E/I balance.

5. When interneurons participate in multiple assemblies, they should realise only one type of computation in each assembly to avoid interference. Thus, it appears

that the interneuron groups are the most efficient when they are specialised in terms of the roles they play in the network.

6. Although the excitatory neurons are the carriers of the signals, manipulations of inhibitory connectivity patterns play a key role in controlling the signal transmission and stabilising the whole network. These observations underscore the importance of neuronal inhibition which shall not be considered as a mere provider of a force to balance excitation.

Chapter 5

Signal gating by cholinergic modulation of disinhibitory pathways

Chapter summary

In the previous chapter we have explored overlapping feedforward chains with the inhibitory as well as disinhibitory pools. The disinhibitory pathway was shown to be central in improving signal transmission success rate in the setups dominated by the inhibitory haloes. Here, the goal is to introduce the principle of gating to control the effectiveness of the disinhibitory pools activity to flexibly propagate or block the signal. We first show that cholinergic modulation targeting the disinhibitory pools can robustly act as a switch. Then we test whether location-specific modulation can modulate individual assemblies while ignoring the others. Both, volume and phasic transmission modes were shown to implement gating, with the phasic transmission providing higher, but not perfect, levels of precision.

5.1 Introduction

Earlier chapters briefly outlined the importance of incorporating gating mechanisms to obtain flexibility of neural networks' responses when exposed to the ever-changing flow of information. As shown in the context of working memory in the prefrontal cortex, the gates can be located at different stages of information processing (Chatham and Badre, 2015). The activity of basal ganglia (BG) was shown to act as a gate on the input to, as well as on the output out of the working memory system. In terms of the input gating, it is commonly assumed that a so-called early selection performed

by a top-down attentional system filters out all the irrelevant sensory signals at their very arrival (Jacoby et al., 1999). Studies on primates performing a context-dependent selection task, however, found no evidence to support this hypothesis and showed that even irrelevant cues get integrated into the circuitry processing which in the later stages leads towards a choice (Mante et al., 2013). Thus, early gating off the sensory inputs, at least in the working memory context, might not be biologically realistic.

Nevertheless, the fact that the BG input can play a role of a gate on the incoming or local signals illustrates another principle – that the information about the signal importance (be it cue's relevance or saliency) essentially comes from another brain region. To analyse the activity of a local network, the metaphor of the black box can be employed. A given black box receives some inputs and is expected to process them to produce an output. The question is which piece of information or element of processing belongs to the 'input' part and what is intrinsic, that belongs to the internal toolkit of this black box. It is well established that the neocortex behaves differentially during various behavioural states suggesting that a variable determining the 'current state' should belong to the input ingredients (Doiron et al., 2016).

5.1.1 Control of disinhibition as a gating mechanism

The black box metaphor can readily be applied to the disinhibitory pathways along the synfire chain, which were thoroughly studied in the previous chapter. In the guarded chain model, these were an integral part of the circuitry itself, so that the chain was unconditionally helping itself to transmit the signal and the elements were tightly timed. No external input was determining the signal's fate, only the internal toolbox which, admittedly, was better equipped than the one of the basic chain.

To implement the gating mechanism, however, the control and decision whether the signal shall pass or halt should come from the outside and be independent of the chain circuitry. In that way, depending on some external input, the disinhibitory mechanism should be on or off, determining whether the actual signal should be propagated along the chain or not. In principle, three modes of implementing the control of the disinhibitory mechanism can be distinguished, as shown in Figure 5.1:

Intrinsic mechanism Disinhibitory pathways are integrated in the chain circuitry and provide a scaffolding that guards the signal from the globally increased inhibition. This was implemented in the previous chapter in the guarded chains.

Controlled by direct external input Disinhibition is activated by external projections.

Controlled by neuromodulation Instead of the direct synaptic projections, disinhibition is activated by neuromodulatory mechanisms, both via volume transmission and phasic inputs.

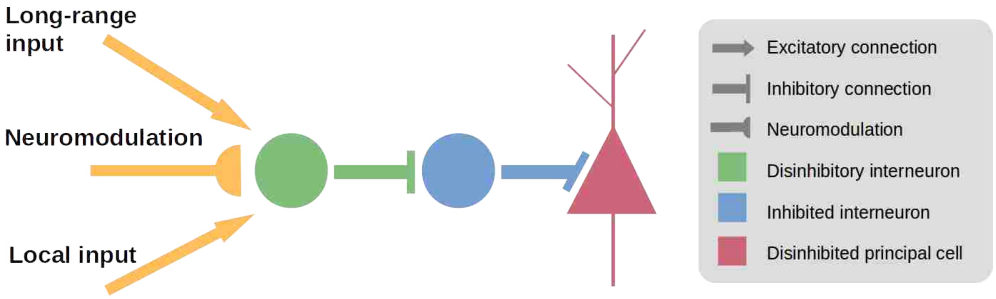

Figure 5.1: Disinhibitory circuit with three possible sources of controlling signals.

It is argued that in order to realise the gating mechanism, the control of the disinhibitory pathways should be transferred from the chain to the external entity.

5.1.2 Default state of a gate

Two possible ways of realising the gating mechanism – with the gate open or closed by default – were already discussed and the examples of models with both variants were given. But what are the functional differences between the two frameworks?

Let's assume that within a given network as many as 10 assemblies are embedded and they are all open in the default state. When there are many inputs arriving, possibly even simultaneously, the gating mechanism would need to make sure that only one gate is open, which translates to extra work on the remaining gates the amount of which scales up with the number of incoming signals. Were the 10 gates closed by default, the extra work would have to be done solely on the gate that is expected to get opened, regardless of the total number of inputs.

Although the closed-gates variant implies that by default a given network is not tuned for optimal signal processing, Newman et al. (2012) argue that this can actually be very advantageous, for example while resting or consolidating memories, when tuning out the external stimuli is desirable and best if achieved effortlessly. Ultimately, it can be argued that depending on the task, brain region or species, either framework can be more relevant. The applicability of either framework can be assessed by estimating the cost of having a false-positive or false-negative response. If a missed signal is to be avoided at all costs, persistently open systems are more suitable, whereas if precision and low error rate are sought-after, closed systems should be favoured.

5.2 The actual contribution of disinhibitory pathways

Before embarking on studying the modulation, we tested the actual contribution of the disinhibitory pathways, which are composed of 2 types of synaptic connections: $E \rightarrow I$ and then $I \dashv I$. We individually removed either type from the chain circuitry to see what their partial contributions are in the signal transmission. After removing the $E \rightarrow I$ connections (Figure 5.2 top left), the chain still contained disinhibitory pools, called unactivated pools, as these were not activated by the chain's own activity. This setup aimed to check whether the background activity is sufficient to trigger the disinhibitory mechanism. The second condition involved a so-called non-specific activation (Figure 5.2 top right). The chain still sent the signal to inhibitory pools, but these were random and not explicitly connected to the inhibitory pools within the chain. The goal was to verify the role of specificity of the disinhibitory connections.

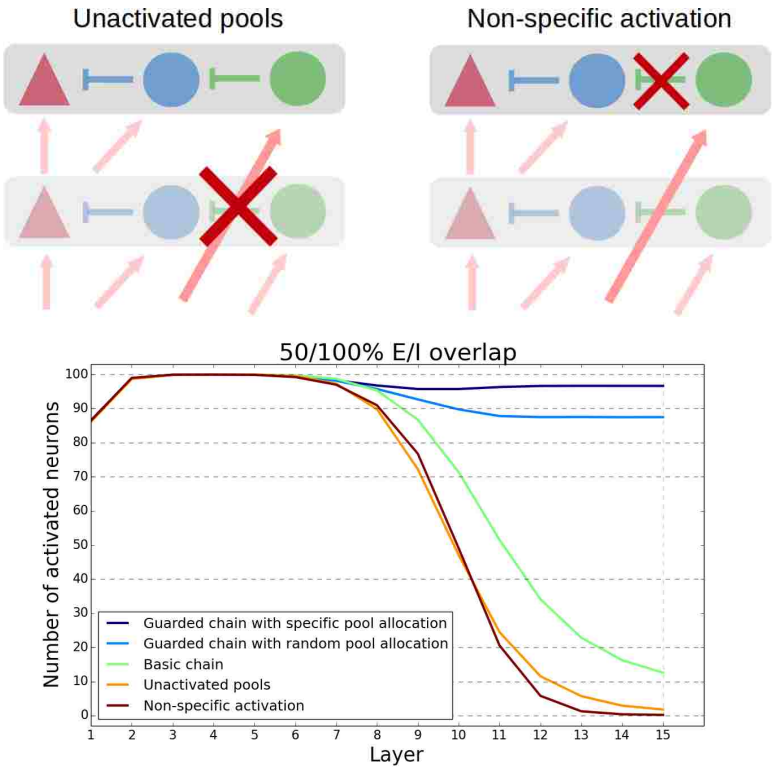

Figure 5.2: Partitioning the disinhibitory pathway. Top left: a chain with unactivated disinhibitory pools. Top right: a chain with non-specific activation of inhibitory pools. Bottom: a comparison between basic, guarded and chains with partitioned disinhibitory pathways in 50/100% E/I overlap condition.

Both setups were tested with high E/I overlaps: 50% excitatory and 100% inhibitory, to ensure a clear difference between the success rates in basic and guarded

chains. As demonstrated in Figure 5.2, only the full disinhibitory pathways provided substantial improvement in the transmission success rates. Partitioned pathways, not only failed to improve the rates, but actually impaired the already poor signal propagation even further. The condition with non-specific activation involved activating an extra number of inhibitory neurons without providing the protecting disinhibitory scaffolding to the chain. As a result, the chain had to withstand an even stronger wave of inhibition, which inevitably led to a complete failure in signal transmission. The unactivated pools, although providing the scaffolding, also worsened the transmission. This was due to the fact that the disinhibitory pathways could be activated only by the random background signals, which naturally were independent of the timing of the signals traversing the chain. The chain, however, to fully benefit from the disinhibitory protection needs precise, time-locked activations strictly occurring before the arrival of the actual signal.

Thus, we have shown that the disinhibitory mechanism in order to be effective needs to fulfil two conditions. Firstly, it should steer the already existing inhibitory wave into inhibitory pools to transform it into disinhibition and secondly, these created disinhibitory signals should be well-timed to effectively provide the protection to the target circuitry.

5.3 Methods

Here, the closed framework will be under investigation and the modulation of disinhibitory pathways will provide a mechanism to open a gate to control propagation of a signal already traversing the chain. The default state will correspond to the basic chains discussed in the previous chapter, whereas the opening of the gate will in essence activate the guarding disinhibitory pathway to facilitate the signal transmission. We will only focus on the modulatory control of disinhibition since there is ample evidence for cholinergic activation of disinhibitory pathways across many brain parts (as outlined in sections 2.4.3 and 2.5.2).

For the sake of completeness, it should be clearly stated that the modulation targeting only one subpopulation of interneurons is a considerable simplification. As already outlined in Chapter 2, neuromodulation is an umbrella term that encompasses various phenomena that alter the electrical properties of neurons. Indeed, interactions associated with acetylcholine are complex and multi-faceted as there is a whole repertoire of cholinergic receptors residing in pre– and postsynaptic sites of various neuron types

and in diverse concentrations. Although a clear-cut neuromodulation of one neuron subtype is biologically improbable, such a simplification allows the model to remain simple and tractable.

5.3.1 Modification of disinhibitory pathways

Having established that only the full disinhibitory pathway improves the signal transmission thanks to its specificity and transforming the inhibition into disinhibition, we modified the pathway so that it could be a target of cholinergic modulation.

The core modification involves decreasing the number of $E \rightarrow I$ synapses along the disinhibitory pathway. In the guarded chain each disinhibitory neuron received 60 synapses from the excitatory chain neurons to ensure that enough neurons emit a spike to disable the inhibitory pools. In the modulation condition, the number of synapses is decreased down to 20–30, as shown in Figure 5.3. Disinhibitory neurons still receive the activation from the chain, but it is too weak to evoke robust responses, and the success rates are similar to those of the basic chains. It is the role of modulation to activate the full disinhibitory pathway.

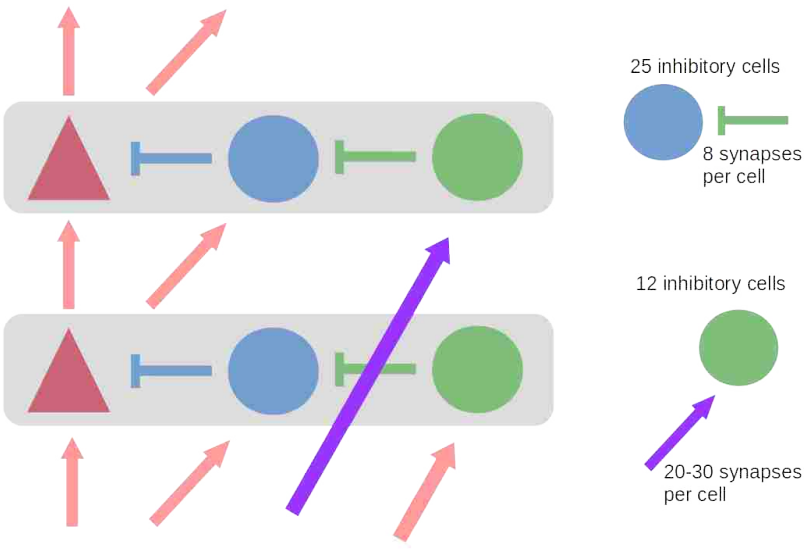

Figure 5.3: Diagram of the guarded chain with modified disinhibitory pathways. Purple arrows denote the altered connections. In the modulation conditions, each disinhibitory neuron receives a decreased number of synapses ranging from 20 to 30 (instead of 60) from the chain.

The reason for keeping the full disinhibitory pathways and only attenuating them is the fact that modulation itself is not expected to activate neurons, it can merely

influence neurons' responses to the incoming activation. If there is no input, even the most powerful modulation should have no effect, as it can only come into play when accompanied by the actual spiking activity. Furthermore, such setup ensures that the disinhibitory pathways will be activated at the expected points of time, strictly depending on the signals already traversing the chain.

5.3.2 Choosing the levels of overlaps

In the previous chapter we showed that when two chains are embedded, the levels of overlaps influence the signal transmission and overall, guarded chains had a much wider parameter space that led to high success rates. In order to model a clear transition from a fully closed to a fully opened gate, a setup with such levels of overlaps had to be found that, firstly, the success rate is very low in the basic chain and very high in the guarded chain; and secondly, the network itself is stable in both variants.

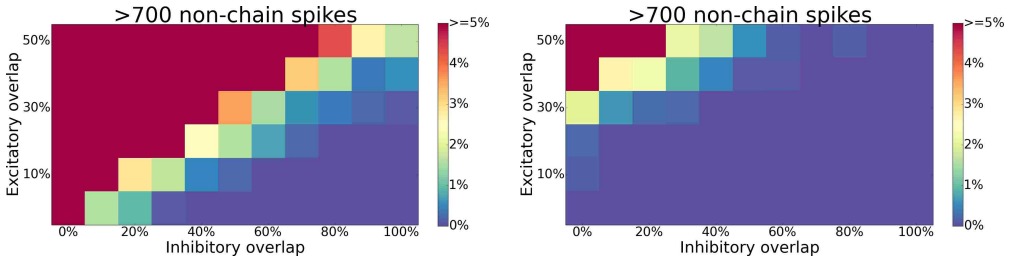

Figure 5.4: Chances of the synfire chain explosion for basic and guarded chains. Note the maximum value set at 5%.

High excitatory and low inhibitory overlap (top left corner in the grid plots in Figure 5.4) is a region where the chances of synfire explosion are very high. It is assumed that a robust system does not allow this to occur and therefore this region shall be labelled as unstable and not likely to be biologically relevant, although the higher the excitatory overlap, the higher the network capacity, as the neurons can potentially participate in more than one chain, increasing the maximum number of embedded chains. Bottom left corner represents the regime where multiple chains share very little or no neurons, compromising the network's capacity. On top of that, the chances of synfire chains are non-zero, rendering this regime unstable. The regimes in the bottom right corner still compromise the capacity, as the excitatory overlaps are low; the whole network, however, is stable and synfire chains do not appear. But it should be noted that the signal transmission succes rate is already near 100% in the basic chain. As the excitatory overlap increases (top right corner), the network is still stable but the dense

inhibition hinders the signal transmission in the basic chain. Disinhibitory pathway improves the signal transmission while keeping the whole network stable. And the relatively high excitatory overlap, potentially improves the network's capacity.

Therefore, the chosen regime to be explored is the one with firstly, a relatively high excitatory overlap (30%–50%) and secondly, a high inhibitory overlap ($>= 50\%$). This regime does not allow for synfire explosions to appear, the addition of the disinhibitory pathways proved to improve the signal transmission considerably and network's capacity is potentially high.

5.3.2.1 Setups for uniform volume transmission

Initially, for the guarded chain conditions, the setups with specific pool allocation and identical overlap levels for inhibitory and disinhibitory populations were chosen, so that the two embedded networks maintained 50% excitatory, 100% inhibitory and 100% disinhibitory overlap. This setup, in the basic chain condition, yielded the lowest success rate, rendering it an optimal default state representing a closed gate.

5.3.2.2 Setups for location-specific volume transmission

The setup above – with the full disinhibitory overlap, might be problematic when two chains are expected to be differentially modulated. If disinhibitory pools are to be utilised as *switches* in the gating mechanism, is it acceptable that they are fully overlapping? If such switches are shared through 100% overlap, then modulating one chain implicitly entails modulating the other chain too. It is conceivable that the modulation can target only one chain and as such, the switches should be private, that is having no overlaps with other ones.

It can also be argued that the 50% excitatory overlap is rather extreme and it might be desirable to explore smaller levels of overlap to find a general behaviour of the networks rather than exploiting a specific point in the parameter space.

Meeting the above requirements turned out to be problematic, because such setups yield a decent success rate in the basic condition, so they could not be used as a default, closed system. The reason for that is that the global inhibition is not as strong during the chain activation so that many signals manage to get propagated even without the disinhibitory scaffolding. One way to introduce more global inhibition is to increase the global inhibitory connectivity. All the networks studied here maintained 5% global connectivity and we explored the networks with denser connectivity onto inhibitory

population to find the setups with a very low success rate in the default state to then apply modulation in order to open the gate. As a result, 4 different setups were chosen for further investigation:

– One setup with 6.0% global inhibition and 100% inhibitory overlap

– Three setups with 6.5% global inhibition and 85%, 90% and 100% inhibitory overlap respectively.

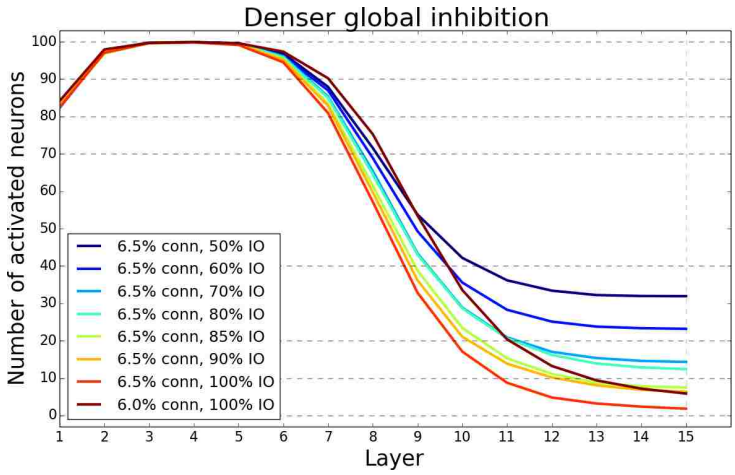

Figure 5.5: Default state in the setups with 6% and 6.5% global inhibition with various inhibitory overlaps. The smaller the overlap, the better signal success rate. Only the setups with the success rate below 10% were selected to study the cholinergic modulation.

Figure 5.5 shows the mean traces in the default state for the chosen setups as well as the rejected setups with 6.5% global inhibition but lower inhibitory overlaps. Setups with denser, 7% connectivity were also explored, where the lowest success rate in the default state was shown to be around 10%. However, when the full disinhibitory pathway was added (60 synapses per neuron, as in the guarded chains), the success rate did not improve up to $> 90\%$, but merely to 50% (data now shown).

It might also be desirable to explore the effects of various levels of inhibitory overlaps on the specificity of cholinergic modulation. The chosen setups above have very high inhibitory overlap because all the setups with the IO of 80% or less yielded too high success rate in the default state as shown in Figure 5.5. Thus, here we will not explore the cases with low or no inhibitory overlaps. This issue will be revisited later.

5.4 Cholinergic signalling via volume transmission

First of all, we modelled cholinergic modulation signalled by non-specific and uniform volume transmission. It was shown that acetylcholine lowers the firing threshold of the target neurons, which makes them fire more effectively.

Only the disinhibitory pools – switches – were under modulation in the activated chain (as depicted in the diagram in Figure 5.7 C). Neurons in these switches had lowered (or increased) firing threshold throughout the whole simulation time, not only during the input injections. Setups with various numbers of synapses along the disinhibitory path were investigated. The simulations followed the same protocol as in the previous chapter, and as already stated, the excitatory overlap between the chains was set to 50% and the inhibitory one – to 100%.

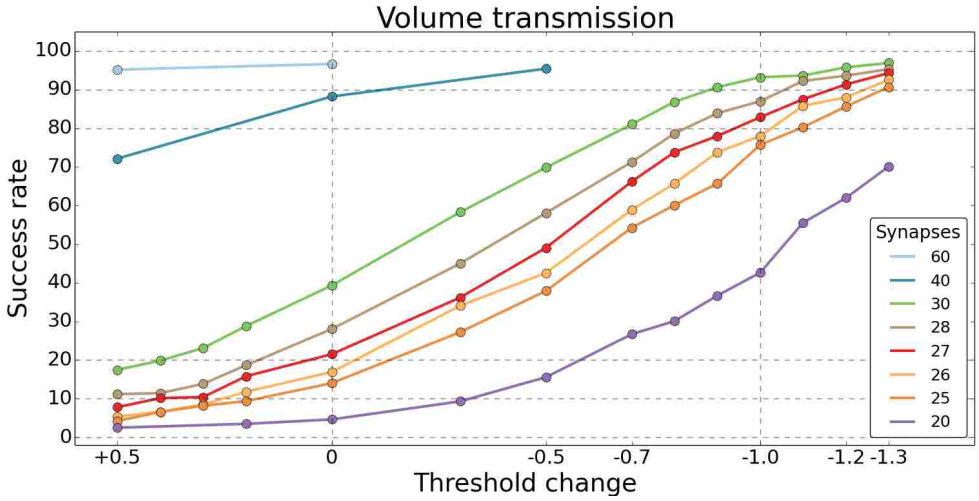

Figure 5.6: Success rate as a function of the firing threshold change in the disinhibitory neurons for setups with variable number of $E \rightarrow I$ synapses in the disinhibitory pathway. Circles denote the actual data obtained from simulations, the lines are interpolated.

As shown in Figure 5.6, by manipulating the firing threshold of the disinhibitory pools, it is possible to obtain a robust switch to perform a graded gating of the signals along the synfire chains augmented with disynaptic inhibition. What is striking is how small the range of the threshold change needs to be in this setup to obtain a near perfect switch from a completely blocked gate to an opened gate with almost 100% success rate.

5.4.1 Location-specific volume transmission

Having demonstrated that the cholinergic modulation modelled as the manipulation of the firing threshold of disinhibitory neurons can work as a robust gating mechanism, we asked how specific such modulation can be. The example above involved the case when both, inhibitory and disinhibitory pools of the two chains were fully overlapping. Here, the goal was to employ the private switches (non-overlapping disinhibitory pools) and modulate them separately. This allowed us to examine the old view about the diffusive nature of the basal forebrain projection into the cortex. Although classically it is assumed that acetylcholine is released uniformly across the cortices, recent studies showed that the projection axons reach specific sites and thus there is a possibility to send different signals of variable magnitude to different regions (Zaborszky et al., 2015a). As a result, the modulation does not need to reach the areas equally.

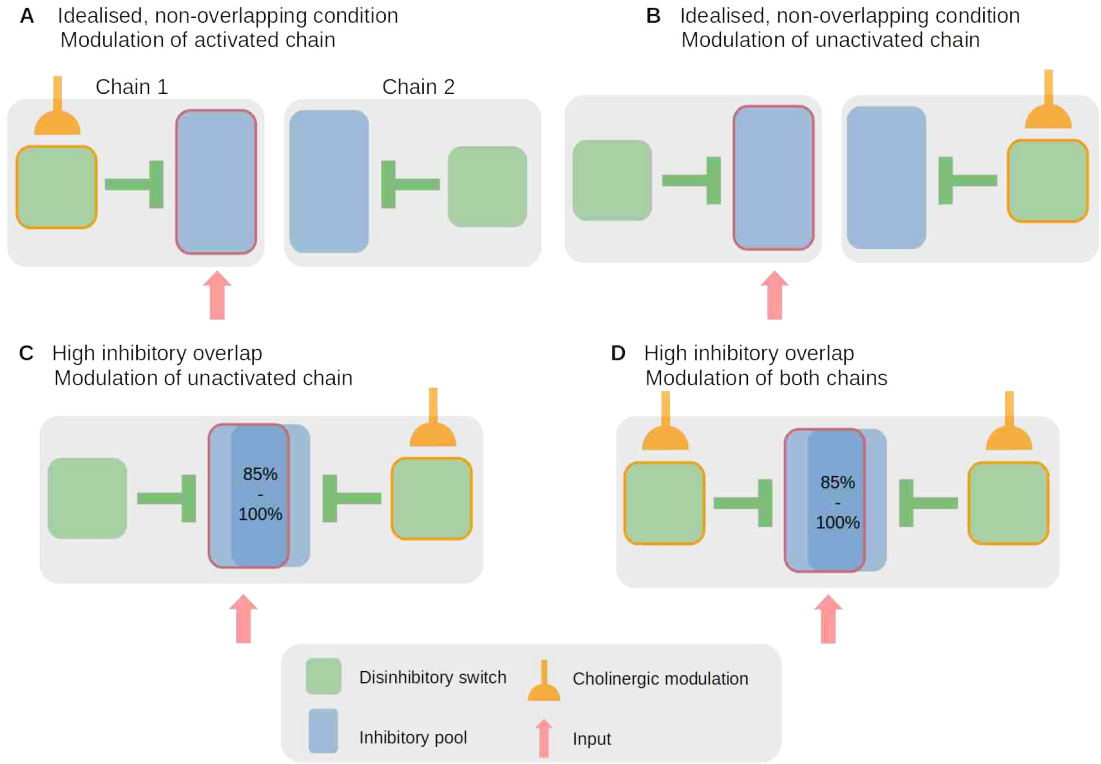

Figure 5.7: Inhibitory and disinhibitory pools within two chains under cholinergic modulation. A. Idealised setup where one chain receives the input and the second one is modulated. B. The actual setups under investigation: disinhibitory switches are private, but the inhibitory pools maintain 85%-100% overlap. When one chain receives the input and the other one is modulated, the targets of the two are mixed. C. Modulating two switches during the activation of one chain.

In principle, the goal was to explore the influence of the modulation of one chain on the activity of the other one. As shown in Figure 5.7 A, there can be two chains with private switches, and only the first chain receives the activation, whereas the second one receives the modulation. The question is to what extent the modulation of one chain will affect the second chain if they are embedded in the same local network. If the modulation of one chain does not affect the signal propagation along the other one, we can label it as a perfect separation. The setups under investigation, however, are the ones depicted in Figure 5.7 B. Although the switches are private, their targets – inhibitory pools – are heavily overlapping. Then the modulation of one chain inevitably can influence the signal propagation along the second one. To find out the extent of this influence, we fixed the number of synapses reaching the disinhibitory neurons to 25 per neuron and we modulated the switch either of the activated or the unactivated chain.

It was revealed that modulating the unactivated chain can immensely influence the signal transmission. As demonstrated in Figure 5.8, when a given strength of modulation (expressed by the threshold change) fully opens the gate in the activated chain (success rate above 90%), the same strength applied on the unactivated chain can still cause the gate to be opened in 30%– 60% of trials. Overall, no setup was found with a near perfect separation.

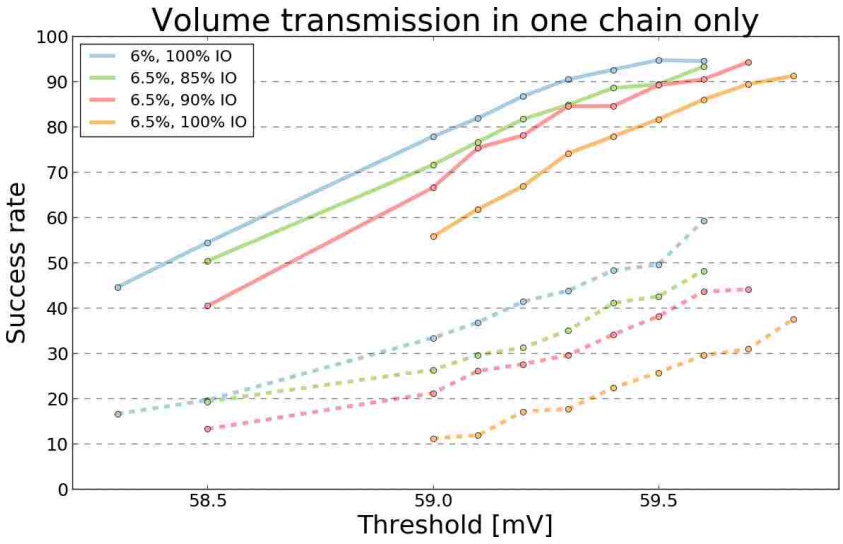

Figure 5.8: Modulation of one chain only. Solid lines denote the cases with modulation and input targeting the same chain, dashed lines denote the cases when modulation targeted the unactivated chain.

5.4.1.1 Modulation of two switches simultaneously

The above example tackled the cases when only one switch was under modulation – either of a chain receiving the input or the unactivated one. The separation of responses were not very well, since, although the switches are private, they still target the inhibitory pools, which share 85-100% of their neurons. We asked whether there is a further improvement in signal propagation if both switches are modulated simultaneously (as shown in Figure 5.7 C). We extended the protocol above to also modulate the two switches simultaneously. Various levels of thresholds were chosen in order to explore the parameter space.

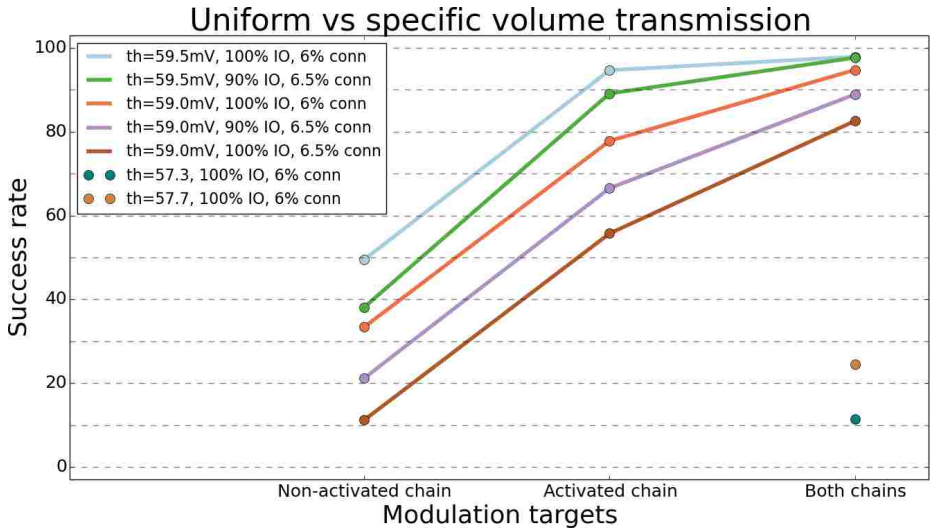

Figure 5.9: Modulation of one or two switches simultaneously. The orange and teal dots in the bottom right corner show that if the modulation is weak, even targeting two sets of switches will not be sufficient to open the gate.

Not surprisingly, modulation of both switches improves the signal propagation, when the strength of modulation is weak. When modulating the switches along the activated chain already robustly opens the gate, modulating another set of switches improves the signal transmission only modestly.

5.5 Cholinergic signalling via phasic transmission

Finally, we sought to explore another mode of cholinergic transmission. It was found that the ACh release can also be fast and precise (Gritton et al., 2016; Sarter et al., 2009), contradicting the notion of the diffusive nature of cholinergic projections. Thus, we asked whether in the case of disinhibitory control on the synfire chains such fast

signalling can indeed provide more precise responses when targeting only one switch but not the other. Again, an assumption that the cholinergic receptors reside exclusively on the presynaptic terminals of the disihibitory neurons was made. In reality, however, the $\alpha 7$ receptors, which are believed to mediate fast responses, were found on both, pre– and postsynaptic sites (Picciotto et al., 2012).

Phasic input was modelled as a volley of spikes arriving at a certain time at the disinhibitory pools. The volley was a Gaussian, similar to the input volley, and described by α: number of spikes reaching each neuron, and σ: standard deviation of the volley. The values of both, α and σ, were varied to explore their influence on the activating of the switches. The volleys were targeting only a subset of layers, since the first 3 layers are lacking the disinhibitory pools and it was shown that the global inhibitory wave affects mostly the middle layers along the chain. Many setups were tested and the results presented here come from the setup where 6 middle layers were receiving the phasic inputs and the timing of these was set to arrive well ahead of the actual signal along the chain, with the 2ms increase in the subsequent layer.

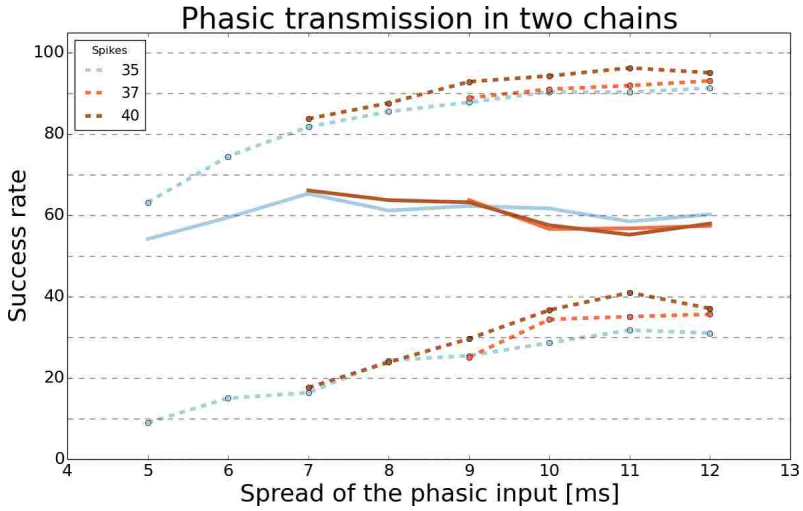

Figure 5.10: Phasic transmission improvement. Top dashed lines denote the cases with modulation and input targeting the same chain, bottom dashed lines denote the cases when modulation targeted the unactivated chain. Solid lines in the middle denote the difference between the two cases.

We show that more spread injections work better, presumably due to better coverage of the time window when the signal might occur. Too specific pulses might miss the signal which in fact can appear within a few millisecond jitter. The volleys essentially come from the external sources and their timing is fixed and not dependent on the signals traversing the chain. Too strong and too dispersed modulatory signals cause

the other, not modulated chain to also improve its success rate, so there should be some trade-off between the strength and dispersion of the phasic input.

5.6 Volume and phasic transmission comparison

Finally, we compared the two modes of transmission, which are controlled by different sets of parameters – volume transmission involves modifying the firing threshold throughout the entire simulation time, whereas in the phasic transmission mode the volleys of spikes are sent to the chosen pools to elicit spikes. The goals of both modes, however, are the same: firstly, to open the gate along the activated and modulated chain and, secondly, provide some level of specificity when the modulation targets the unactivated chain.

Two plots in Figure 5.11 collate the success rates obtained in the volume and phasic transmission setups. In the left panel, success rates in the modulated chain are plotted against the rates in the unmodulated one. The right-hand side plot displays the success rates in the modulated chain versus the difference. This plot highlights the fact that the relation between the success rate and the difference is not perfectly linear – the best precision can be obtained in regimes that open the gate with the 80% efficiency, but if the above 90% efficiency is sought-after, the precision decreases.

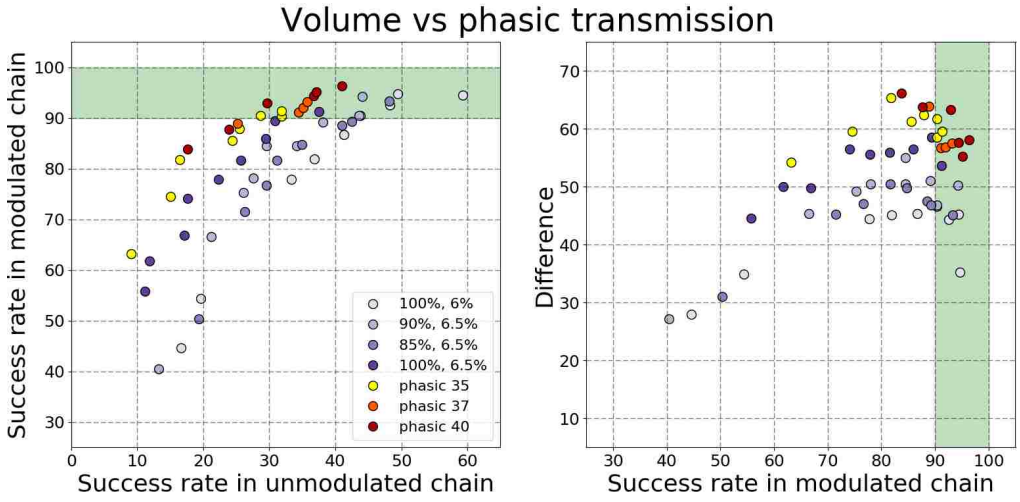

Figure 5.11: Volume vs phasic transmission. The same data, two ways of showing the relation. Blue dots denote volume transmission, warm reds – phasic setups. Left: success rate in the modulated vs unmodulated chain. Right: success rate in the modulated chain vs the difference in rates (modulated – unmodulated). Green shading denotes the regions where the gate in the modulated chain is open (>90% success rate).

The plots reveal that all the setups with phasic transmission are moderately but consistently better than the volume ones. Does it mean that the phasic mode should be regarded as the more efficient one? Not necessarily. It should be remarked that the volume transmission is controlled by only one parameter affecting all the neurons in the disinhibitory pools, whereas the phasic one is more complex, and, most importantly, targets only a subset of disinhibitory pools. It can be argued that introducing more complexity and optimisation only to obtain a rather modest improvement is unnecessary and somewhat troublesome.

Finally, although a perfect separation was not found, it should be clearly highlighted that the setups that were under investigation involved as high as 85-100% level of inhibitory overlap. Although the disinhibitory switches were private, their target populations were highly overlapping. Then, even with such a high level of overlap, as high as 65% separation was achieved to boost one cell assembly while ignoring the other.

5.7 Discussion

We have shown that in principle, a uniform global modulation targeting the disinhibitory pathways is capable of controlling the signal propagation along the feedforward chains. In more detailed scenarios, where the goal was to modulate one assembly while ignoring the other, such a global mode turned out to be insufficient. The phasic cholinergic transients improved the response separation, but only to some extent.

The core of the issue lies in the fact that although the modulated disinhibitory signals are specific, they target the same, globally shared inhibitory blanket. Such setup obviously lacks the sought-for specificity. It has been suggested that the solution to this problem might actually be found on the subcellular level, namely in the dendritic branches of pyramidal cells where the SOM interneurons' axons terminate. Recently, it was proposed that although the $SOM \dashv PYR$ connectivity appears to form the indiscriminate blanket of inhibition, this connectivity can actually be sparse when examined at the individual branches (Yang et al., 2016). This observation was exploited in a computational model which demonstrated that if the specific input pathways cluster on separate dendritic branches of the pyramidal neurons, these pathways can be robustly gated by the disinhibitory mechanisms. Importantly, such branch-specific disinhibition was shown to be possible also with a dense, blanket-like connectivity.

Chapter 6

Cholinergic modulation of feedforward inhibition to invoke replay

Chapter summary

In this chapter we explore the notion of spontaneous replay and possible mechanisms to generate one in the feedforward chain model presented in this thesis. We show that the modulation of a subpopulation of interneurons providing disynaptic feedforward inhibition (FFI) proves to be the key element in invoking replay. Since this circuit also protects the network from the synfire chain explosions by providing balancing inhibition, weakening the FFI promotes not only replay but also explosions. In order to find stable setups we explore various patterns of selecting only subpopulations of the interneurons. We conclude that both tasks, providing balancing FFI and ensuring global stability should be realised by separate pools of interneurons and hypothesise that the SOM and PV+ interneuron classes might be suitable candidates.

6.1 Introduction

The mode of the synchronous spiking propagation triggered by a strong input – volley of spikes – was extensively explored in the previous chapters and the modulation of the disinhibitory pathways was shown to realise the gating of signals along the chains of excitatory and inhibitory pools. In this chapter we focus on the question, whether such chains can also be modulated to invoke an internally generated replay.

Internally generated or spontanous replay is defined here as a type of an activation that is generated at the absence of an external, specific input. It is implicitly assumed

that the signal pathways are already embedded into a random network, but the baseline activity is too weak to turn them on. Such a replay refers to an activity that is typically observed during animal's sleep or quiet wakefulness, when an animal indeed tunes out the external stimuli (Carr et al., 2011).

Technically, this definition of replay calls for a requirement that the state that enables the formation of replay has to be crucially different that the awake state characterised by network's vigilance and responsiveness to the incoming stimuli. It is not desirable to have an awake state facilitating both: input-driven and internally-generated activation, because that would cause confusion whether an activation was caused by an actual input and is relevant or whether it can be disregarded due to its spontaneous formation. Thus, certain manipulations need to be applied to turn the awake network into a sleep–like state, when the replay is allowed to take place. Since it is assumed that there is no external input reaching the network in that state, it is irrelevant whether the robust input-driven activation is possible or not.

Neuromodulation is a straightforward candidate for controlling network states as these were shown to be characterised by different levels of neuromodulators available in the extracellular space (Lee and Dan, 2012). Microdialysis measurements of acetylcholine in the hippocampus of freely moving rats (Kametani and Kawamura, 1990) and cats (Marrosu et al., 1995) showed high levels of acetylcholine release during this active waking behaviour. During sleep, on the other hand, the levels of cortical acetylcholine are considerably lower. Thus, it was hypothesised that the sleep-like state can be achieved by modifying the level of available acetylcholine in the local circuits.

The question is what exactly such modulation is supposed to achieve and which network elements should be its targets. In principle, the network receiving only excitatory background noise and maintaining low firing rates is expected to integrate and amplify a random signal along the feedforward circuitry. By default, the balanced connectivity ensures that any upsurge of excitatory activity is immediately followed by the inhibition, impeding the integration of stronger signals out of random fluctuations. Thus, breaking a balance might be an option to overcome this obstacle. Again, this can be achieved by either modifying excitation or inhibition. Since the expectation is to amplify a weak signal, the excitation should become stronger or inhibition weaker. Cholinergic modulation of excitatory neurons was numerously reported, but it would always improve neurons' excitability in the states with high ACh levels. Here, the goal is to modulate a network to turn it from the default, awake state into the sleep-like state, thus, the biologically realistic modulation of excitatory population would

in fact have to involve decreasing the activity of the excitatory cells. Increasing the excitatory background noise is also not realistic since in the sleep state the networks are less active. In fact, injection of extra excitation was already examined and it was demonstrated that when the network increases its firing rates from 5Hz up to 12Hz, spontaneous replay can robustly occur (Chenkov et al., 2016). Here, it is argued that such mode does not account for the biological characteristics of the sleep-like state.

Modulation of inhibition remains as a potential candidate. In the previous chapter we discussed the disinhibition as a powerful mean of passively increasing the activity of excitatory neurons. Here, however, this mechanism is inappropriate, since in order to apply disinhibition to gate the signal, the signal itself should already be existing. Generating replay involves creating a signal out of random noise and disinhibition as such cannot help. Also, first three layers of the chain do not have disinhibitory pools so it is impossible to modulate the initial part of the chain via disinhibition.

Notably, the chains modelled here contain yet another type of inhibitory pools – the interneurons that realise the disynaptic feedforward inhibition. As remarked by Kremkow et al. (2010b), the addition of the FFI into the synfire chain results in powerful filtering out the weak signals and ensures that random background fluctuations do not trigger the signal transmission along the chain causing false-positive responses (Tetzlaff et al., 2002). Such spontaneous firing is naturally unacceptable during the awake state, but in the sleep-like state this is exactly what is needed to invoke replay. Thus, it was hypothesised that during the sleep-like state the FFI mechanisms are relaxed so that the random network fluctuations are sufficient to trigger signal amplification along the chain.

6.2 Methods

In the previous chapter, the default state of the network involved a closed gate and the application of cholinergic modulation opened it to enable signal transmission. In the case of replay, the setup is inverse. The default state examined so far corresponds to the awake state, which is implicitly already under the influence of the cholinergic modulation thanks to the high levels of extracellular acetylcholine. The modulation examined here essentially involves *removing* the acetylcholine from the local circuits which presumably will decrease inhibitory neurons' efficacy and is thought to work exclusively via the volume transmission. Phasic transmission was shown to involve precise transient pulses correlated with the behaviour and as such, the awake state can

be described as dominated by the periods without the cholinergic transients in local circuits. Thus, modelling the sleep-like state without the transients would in fact look too similar.

Only the inhibitory neurons in the FFI pools were under modulation. Their firing threshold was raised so that they fired less reliably. The expectation was that at the event of random upsurges of excitatory activity, weak signals would have a chance to integrate inside an excitatory pool within a given layer and then activate the next layer's excitatory pool and FFI interneurons pool. Since these interneurons would be under modulation, they would need more inputs to reach the threshold and thus their firing would be delayed or abolished. As a result, the random signal would gain extra time to integrate and again reach the next layer until it gets amplified to become a strong volley.

To detect the replay events, networks with the two embedded guarded chains were simulated for 60 seconds. In contrast to the previous protocols, only the background excitatory noise was delivered to all the neurons so that the network maintained the low firing rate activity (1–2Hz). For each investigated condition, 10 network instances were created.

6.3 Cholinergic modulation of feedforward inhibition

Firstly, an idealised setting was examined. Although two chains were embedded in the network, these were non-overlapping and only one chain underwent modulation, that is all the interneurons in the FFI pools had their firing threshold increased.

The spontaneous replay indeed occurred along the modulated chain, as shown in Figure 6.1, whereas no replay event was recorded along the second, non-modulated chain.

This result underscores the importance of the FFI circuit in controlling the width of the time window for integrating the signals. In the awake state, this filter works powerfully to ensure that only the strong volleys are capable of activating the chain, while the random fluctuations remain with no opportunity to invoke false positive responses. Modulation of the FFI supports the formation of spontaneous replay out of much weaker signals which is not desirable in the awake state, but biologically realistic during sleep and quiet wakefulness.

The right hand side plot in Figure 6.1 displays a rather abnormal instance of a replay event. The signal amplification occurred in multiple sites simultaneously creating

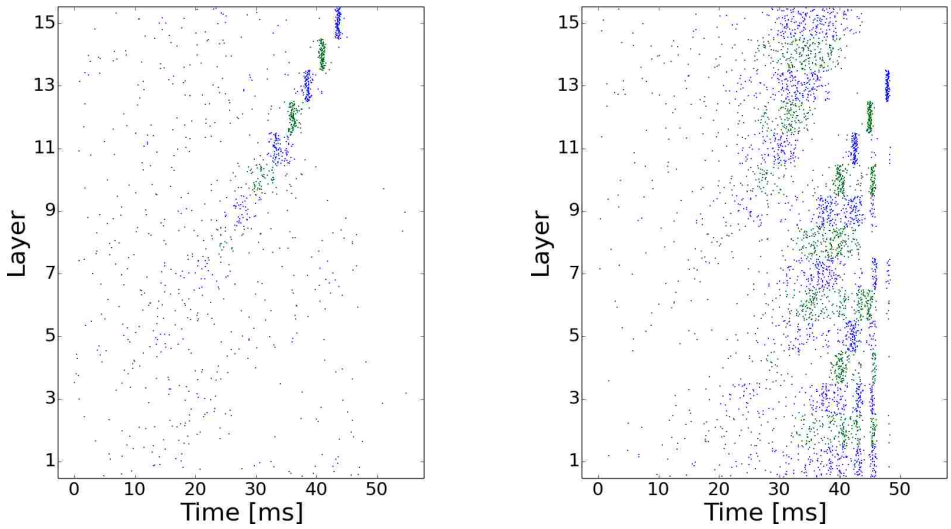

Figure 6.1: Spontaneous replay during the modulation of FFI. Left: an ideal replay event initiated in the middle of the chain out of random fluctuations. Right: signal amplification at multiple sites simultaneously.

two strains of cascades and a mild explosion in the initial layers of the chain, since the neurons there fired in unison. Such events, although undesirable, are the straightforward results of the uniform modulation of the whole chain – the random signals can independently enter the chain at any layer, especially during strong global fluctuations which spread across the whole network.

6.4 Synfire chain explosions

The abnormal replay events are not the only unwanted result of the FFI modulation. The main and most severe side effect of such intervention is the reduced amount of global inhibition which results in frequent occurrences of spontaneous synfire chain explosions, as shown in Figure 6.2.

Since the whole chain is evenly modulated, a signal amplification can be initialised in any layer. As a result, a strong global fluctuation can trigger such amplification in two or more sites simultaneously. Already weakened global inhibition has no chance to balance such multiple upsurges and instead of replay, an explosion spreads across the whole network.

Four levels of spontaneous activation can be singled out. The replay can be reliably invoked without causing global instabilities, as shown in the left hand side plot in Figure 6.1. Right hand side plot displays the second level of activation – a replay with

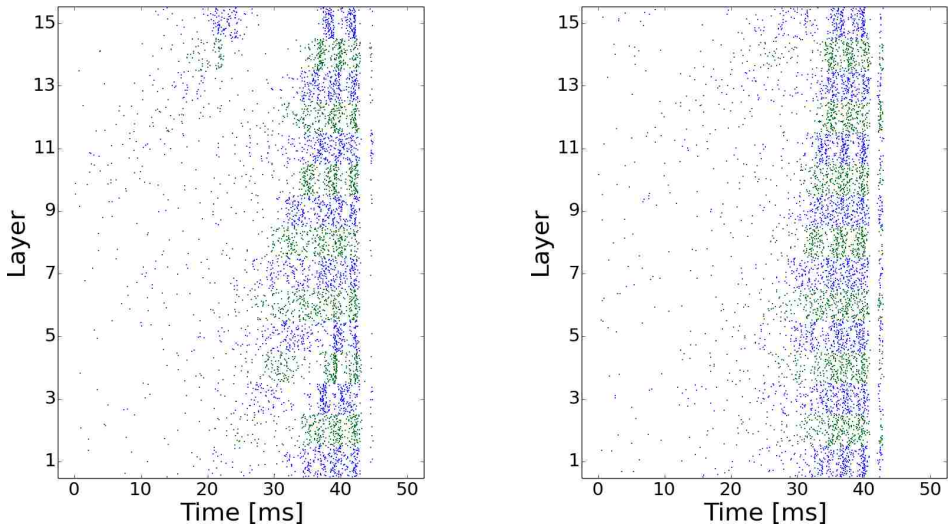

Figure 6.2: Spontaneous synfire chain explosion. Left: an explosion triggered by a cascade-like activity. Right: an explosion without the preceding cascades.

multiple cascades. Then, the emerging volley can cause global explosion, as shown in Figure 6.2 left, and finally, the explosion can occur at the absence of any preceding cascade-like behaviour (Figure 6.2 right). Since any occurrence of explosions is adverse, only two types of events will be of interest here: either robust replays or any events that involve explosions.

To search for the replay events, various levels of the firing threshold were tested. For each level, the number of occurrences of both, replay and explosions were counted and it was revealed that explosions emerge considerably more frequently than the replay events, as shown in Figure 6.3.

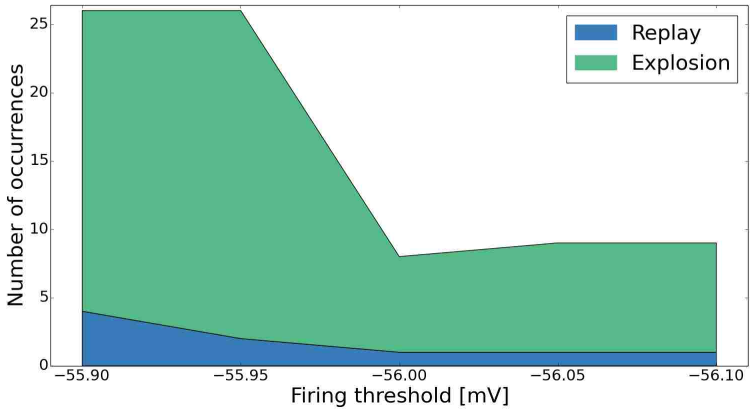

Figure 6.3: Spontaneous synfire chain explosion vs replay for different levels of the firing threshold. The default value of the firing threshold was -57mV.

The setups which, in principle, allow for the occurrences of replay, but are dominated by the random explosion events are evidently unacceptable. The ideal setup should facilitate random replay events and remain stable at all times.

6.5 Modes of modulation to avoid the instabilities

The next goal was to find a way of modulating a sufficient amount of FFI interneurons to still promote the replay formation but at the same time to minimise the chances of the synfire chain explosion events. A straightforward way of avoiding instabilities is to modulate fewer FFI interneurons. Below two ways of achieving this are scrutinised.

6.5.1 Modulation targeting a subset of layers

Firstly, we selected only the initial segment of the chain to be the target of the modulation. In principle, we need to find a setup to let weak and dispersed pulse packets accumulate along the chain. Wide pulse packets were shown to always get narrower and more precise when they traverse the chain (see the state space diagram for synfire chains in Figure 4.12). In essence, only at the entrance of the chain the FFI is required to be relaxed. Down the chain, there is no longer a need for the wider window of opportunity as the signal is narrow enough to withstand the powerful filtering. On top, reliable FFI is necessary to balance the sudden amplification of the excitatory activity.

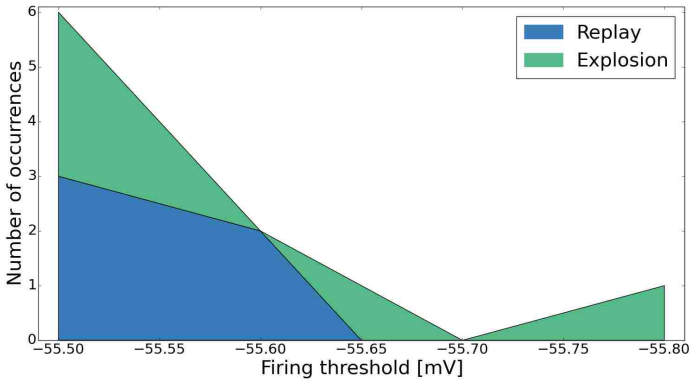

Figure 6.4: Spontaneous synfire chain explosion vs replay for different levels of neurons' firing threshold. Modulation targeting only the initial 10 layers of the chain.

Figure 6.4 displays the number of occurrences of spontaneous events when only the first 10 layers are modulated for various levels of firing threshold. Although the

ratio between the replay and explosion events occurrences improved, the results are still not fully satisfying. As soon as the setup allows for the replay events, these are inevitably accompanied by the explosions.

Such setup can be conceptualised as an analogue of opening holes in the blanket of inhibition (Karnani et al., 2016). One only needs to open a hole in the 'blanket of FFI' to let a wide signal in, which will subsequently get narrower once it is inside the chain. Although it can be criticised that it is unrealistic and somewhat rigid to only modulate the selected segment of the chain, it may still have a biological relevance. The washout of modulators does not happen uniformly so the holes can randomly arise along a chain. Also, the cortex is composed of 6 layers with different stereotyped connectivity patterns. It is plausible that the FFI gets weakened only in some layers where the initial segment of the chain resides – such as L4 which was shown to receive inputs from the thalamus (Cruikshank et al., 2007).

One can ask whether there is the minimum number of layers that need to be modulated to let wide signals in, or parameterise the width and strength of the such wide signals. The above analysis involved 10 layers under modulation, the setups with 9, 8 and 7 layers were also investigated (data not shown) and in all cases the setups were found where both, replay and explosion events emerged. Importantly, no setup was found that only the replay events were robustly present. Weak modulation caused no events at all, whereas as soon as the level of modulation allowed the replay events to emerge, these were always accompanied by the explosions. Because no satisfactory setup was identified, no thorough analysis was carried out to establish a minimum number of layers and the width of the signals needed for the replay formation.

6.5.2 Modulation of a fraction of neurons in pools

Another way of selecting a 'sufficient number' of the FFI interneurons to be the target of modulation is to pick only a fixed percentage of these in each layer under the assumption to model the division of labour. While some neurons undergo cholinergic modulation, some might not, so that they can remain responsive during the amplification to protect the whole network from synfire chain explosions.

Simulations were run with the 50%, 60%, 65% and 70% of the FFI interneurons modulated. While for the first two setups no events were detected, the 65% and 70% ones exposed a number of synfire chain explosions. No replay events were recorded at all. It was hypothesised that it is more probable to form an explosion as this can arise

when the amplification takes place in several places simultaneously. The replay event, however, necessitates only one site where the signal becomes narrower and stronger.

This mode was rejected as a way of invoking replay while avoiding explosions.

6.6 Overlapping chains

So far only the idealised setup was examined, where the embedded two chains were non-overlapping and the modulation affected only one of them. In the previous chapters, however, we examined two chains with substantial overlaps and showed that these have a lot of influence on the signal propagation and overall stability. Thus, it was expected that these might also play an important role in invoking replay as well as synfire chain explosion events is the sleep-like state. Because the excitatory overlap in general promotes explosions, here only the inhibitory overlap was under consideration.

Two setups were tested. The first one involved modulation targeting the full chain and the second setup involved modulating 10 initial layers. In both setups the inhibitory overlap was set at 50%, whereas the excitatory as well as disinhibitory overlap remained at 0%. The levels of firing threshold were set at the same levels as in the cases with the 0% inhibitory overlap discussed in the previous sections to enable a direct comparison.

Figure 6.5 compares the counts of replay and explosion occurrences for the setups with 0% inhibitory overlap (data taken from Figures 6.3 and 6.4) and the newly generated setups, with 50% overlap. The inhibitory overlap turned out to be detrimental to the network's stability. Interestingly, in both setups, the overlap had a different effect. When the modulation was applied to the whole chain, the number of both, replay and explosion events raised. When the modulation targeted only the initial segment, the presence of overlap only promoted more explosions, but the number of occurrences of replay events remained the same.

These results are not surprising. Overlapping pools of FFI provide lateral inhibition between the chains and thereby protect the network from the synfire chain explosions. While the FFI is relaxed, this protection no longer works and the instabilities dominate the network.

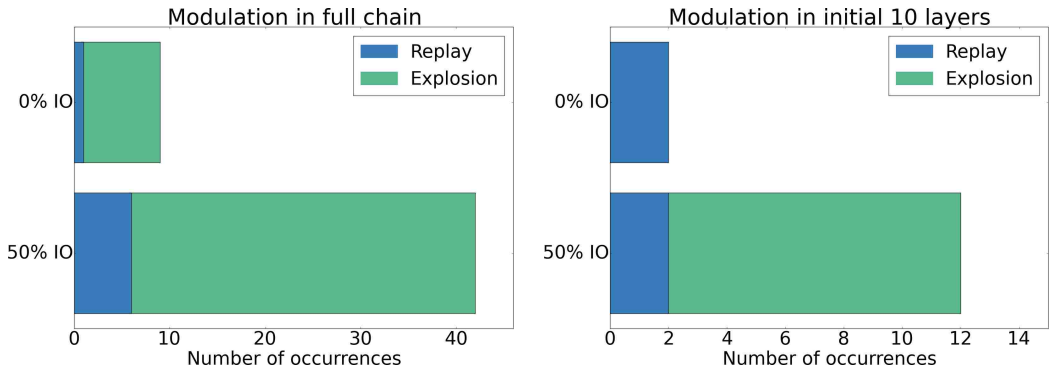

Figure 6.5: Spontaneous synfire chain explosion and replay event occurrences for two levels of inhibitory overlap (IO). Left: Modulation targeting the full chain. Right: Modulation targeting the initial 10 layers of the chain.

6.7 Discussion

We have shown that in principle, the modulation of feedforward inhibition can invoke spontaneous replay. However, relaxing FFI implies decreasing the amount of global and lateral inhibition, which was shown to result in instabilities in form of synfire chain explosions. Without accompanying reliable inhibition, rapidly amplified excitation spreads uncontrollably.

6.7.1 Two types of inhibition

How to safely relax the FFI without destabilising the whole network? One solution would be to introduce another source of reliable inhibition apart from the modulated FFI. Since the interneuron population is highly heterogenous, such an extension of the model could be readily justifiable and backed by the experimental data. Specifically, the PV+ interneurons, unlike the SOM ones, typically do not directly respond to modulation and their physiological properties are specialized for rapid signaling (Hu et al., 2014). In the presented model, all the inhibitory synapses had their delay set to 3ms, so that all the inhibitory signaling was carried out at the same speed, which was also slower than the excitatory one (2ms synaptic delay). It is proposed that another group of interneurons loosely mimicking the PV+ interneurons as expressed by shortening their synaptic delay would potentially provide a faster component of inhibition that would in turn improve the network's stability. The future extension of the model should verify these claims.

Curiously, in the theoretical studies the division of slow and fast feedback inhi-

bition was proposed to be responsible for different aspects of neuronal activity. Fast inhibition was suggested to be *for coding*, whereas slow inhibition – *for computation* (Denève and Machens, 2016). The fast component provides stability, tightly tracks excitation and rapidly balances the activity, whereas the slow one can realise more specific computations on longer time scales, for instance the FFI. This reasoning was backed up by the models involving the rate neurons, so it remains to be discovered how this idea can be practically realised in the spiking networks.

Chapter 7

General discussion

7.1 Summary of the thesis

In this thesis we have explored the propagation of the synfire chain-like activity augmented by the feedforward inhibition (FFI) embedded in the random spiking networks. We have shown that the signal propagation is possible in the chains embedded in the recurrent networks and the addition of disinhibition improves the transmission. By overlapping the chains, a lateral connectivity between the two assemblies is implicitly created. While the excitatory overlap is detrimental, the inhibitory overlap ensures that while one chain is active, the remaining chains stay suppressed. Such lateral inhibition was shown to robustly protect the network from instabilities.

We then studied the cholinergic neuromodulation to gate the signal on and off where the disinhibitory pools were its sole targets. Modulation was modelled either as a manipulation of the firing threshold via the volume transmission or as a transient direct depolarisation mimicking the phasic transmission. We show that both modes can powerfully realise the gating mechanism along the feedforward networks. Modulating one assembly while ignoring the other turned out to be problematic due to the fact that the modulated disinhibitory switches targeted the same, shared inhibition.

Modulation of the feedforward inhibition proved to be suitable to invoke spontaneous replay, but at the same time it also increased the chances of synfire chain explosions. We have explored various ways of modulating the FFI while keeping the network stable, but it appeared that the replay would always be accompanied by the explosion risk. We hypothesised that the solution to this problem might lie in introducing fast and reliable inhibition that would not be subject to modulation.

Altogether, we have shown that by incorporating various inhibitory mechanisms

inspired by the experimental results, one can obtain not only stability, but also robust control over the signal transmission along the multiple synfire chains in random spiking networks. Inhibitory and disinhibitory circuits are excellent targets for modulation because they only passively promote spiking activity. Lastly, when neurons participate in multiple cells assemblies, they should invariably realise only one type of computation in all the assemblies.

7.2 Relevance of the presented models

'The key test of the value of a theory is not necessarily whether it predicts something new, but whether it makes postdictions that generalise to other systems and provide valuable new ways of thinking.' (Marder, 2015)

How to evaluate the work presented in this thesis in the light of the above quotation? We believe to have demonstrated the importance of studying neural inhibition even in the scenarios where the primary goal is to shape the excitatory activity. Exploration of inihibitory circuits should be considered an essential part of studying neural networks. Nevertheless, certain prediction can be formed based on the presented results.

The first prediction relates to the disinhibitory circuit. As thoroughly reviewed in Section 2.4.3, many disinhibitory motifs were dissected in various parts of brain and typically involved VIP and SOM interneurons that increased and decreased their activity, respectively. Some studies, however, showed an inverted effect, suggesting that a model with a one-way disinhibitory path is too simplistic and might require an extension (Dipoppa et al., 2016). How do our data relate to this? It should be borne in mind that these experimental studies employed the data recorded over long periods of time, when an animal was, for instance, running on a ball or watching some visual stimuli (Fu et al., 2014). The disinhibitory circuit we developed, on the other hand, operates on a milisecond timescale. Activated disinhibitory neurons reliably silence inhibitory pools, which just a few miliseconds later become equally active, as they are also involved in a yet another circuit realising FFI. Thus, it is predicted that certain disinhibitory effects can operate only on a small scale, and what is captured on a big scale, might not necessarily reflect the actual functionalities. It is well plausible that an increased activity of SOM interneurons is in fact a consequence of a well-timed disinhibition, that controlled the excitability of the SOMs, as shown in our examples with the guarded chains, where disinhibitory pools (putative VIPs) reliably protected

the circuit so that it could successfully transmit the signal and thus activate the (putative SOM) interneurons. Further physiological experiments should examine the SOM – VIP circuits on a milisecond scale to verify the results obtained on a bigger scale.

Regarding overlapping cell assemblies, we have shown that two assemblies can share a considerable proportion of excitatory neurons and virtually an entire inhibitory neurons population. Thus, our model appears to confirm the notion of a blanket of inhibition (Karnani et al., 2014) and the idea of harnessing the disinhibition to locally override this blanket (Karnani et al., 2016). However, as revealed in Chapter 5, densely shared inhibition poses a problem in modulation of individual assemblies via inhibitiory circuits and a solution involving dendritic computation was proposed (Yang et al., 2016) In detail, it is expected that interneurons from the global blanket are not only locally connected to specific neurons to form specialised circuits, but also connect to these neurons on specific dendritic branches. Thanks to such organisation, 10 inputs entering the same branch would have a stronger effect on a cell than 10 inputs entering a cell in random, distant branches. Future computational studies should examine to what extent such a solution improves the specificity of inputs filtering and where exactly in circuits such differentiation should be placed. Physiological studies, on the other hand, could uncover the actual distribution of synaptic contacts on the dendritic branches between various interneuron groups and pyramidal neurons.

Finally, we have shown that, in principle, a spontaneous replay can be invoked by silencing a group of interneurons realising the circuit of feedforward inhibition (FFI). Silencing a random pool of inhibitiory or disinhibitory neurons showed no effect. This result could be readily tested experimentally. Given a network that was shown to exhibit some sequential activity, it is expected that by inactivating (for example by harnessing optogenetics) certain groups of interneurons, one can obtain variable levels of sequential spontaneous activations. Since the SOM interneurons were shown to realise the FFI, it is expected that their decreased activity should be correlated with the increased appearance of sequential activity. The VIP and PV+ groups are expected to show no such relation.

7.3 Inhibitory overlaps between cell assemblies

One of the main results in this thesis is that the levels of inhibitory, unlike excitatory, overlaps between the embedded chains were critical in providing stability of the networks and flexibility in signal transmission.

Firstly, in Chapter 4 we have shown that the high inhibitory overlap protects the whole network from the synfire chain explosions. Then, in Chapter 5, disinhibitory pathways were chosen to be the targets of modulation and the separation of inhibitory and disinhibitory overlaps turned out to be essential. In the scenarios with modulating only one chain, high inhibitory overlaps proved to be problematic. Finally, Chapter 6 demonstrated that when the feedforward inhibition is modulated in one chain, the risk of explosions increases and the inhibitory overlap only escalates the problem.

It appears that the high inhibitory overlap is beneficial in some setups, while it becomes problematic in others. Lateral inhibition is essential in providing global stability, but if too much inhibition is shared, it becomes troublesome to harness the inhibition-targeting modulation to affect only one cell assembly while ignoring the rest. A shared, global inhibition seems to agree with the recent studies demonstrating that the SOM interneurons are locally densely connected, targeting virtually all the neighbouring principal cells, and thereby creating a 'blanket of inhibition' (Karnani et al., 2014). Similarly, other studies revealed that the highly interconnected *operational hubs* (following the small-world wording), which orchestrate the whole network's activity, are usually the GABAergic interneurons, not the principal cells (Cossart, 2014; Bonifazi et al., 2009).

These findings support only the results, where the lateral inhibition proved to be beneficial. But how to obtain the specificity within a dense, indiscriminate inhibitory connectivity? As already mentioned in Chapter 5, the computations on the dendritic branches might reconcile the contrasting results (Yang et al., 2016). In brief, it was demonstrated in a computational model that when specific pathways terminate on separate branches, the signals can be successfully gated by the disinhibitory mechanisms. Future studies should resolve, whether the compartmental models are indeed necessary to obtain the specificity between the cell assemblies or whether these effects can still be reproduced in the less complex, spiking networks.

Another extension of the current model should definitely involve studying much bigger networks with the capacity to embed several cell assemblies simultaneously. This would allow to construct overlaps spanning more than 2 assemblies, which can potentially shed more light on the functionality and limitations of the globally shared blanket of inhibition.

Nevertheless, studying two overlapping cell assemblies proved to be immensely valuable. This setup enabled us to uncover many problems and mechanisms that would have never appeared if only one feedforward chain was under investigation. This ob-

servation extends the view about the cell assemblies studied in isolation. These should not only be embedded in a recurrent network, but should also interact with other assemblies so that the full range of behaviours is exposed.

7.4 Should cortical activity be dominated by inhibition?

The networks that were under investigation can be described as random balanced networks with the embedded non-random elements in a form of feedforward chains. We have shown that the presence of such non-random elements has a potential to destabilise the whole network due to a strong signal amplification along the divergent-convergent connections along the chain.

One can single out two types of regimes depending on the amount of global inhibition that arises during the amplification. If a network lacks inhibitory countermeasures, a sudden upsurge of excitation can cause synfire chain explosions and such regimes can be labelled as dominated by excitation. Conversely, the regimes with strong lateral inhibition causing inhibitory haloes can be labelled as dominated by inhibition. Both regimes are not ideal: explosions are the events that should be avoided at all costs, whereas silencing of the whole network is nothing else but throwing a baby out with the bathwater. The addition of disinhibitory pathway rescues the signal, but it should also be used with care. Disinhibition involves removing the inhibition, and naturally, if too much inhibition is being removed, the regime can no longer be dominated by inhibition and can become susceptible to explosions.

This issue is figuratively depicted in Figure 7.1. The regimes dominated by the effective excitation are coloured red, and the ones dominated by the effective inhibition – blue. In the middle, where the two forces are balanced, stable propagation can take place. Green arrows denote the action of disinhibition – it involves removing a part of inhibition, so it shifts the regime horizontally towards the redder region. Two leftmost arrows represent the negative effect of disinhibition – due to the removal of inhibition, excitation becomes dominant and thus the risk of explosions raises. Two rightmost arrow represent safe application of disinhibition – after the removal of some amount of inhibition, it is still dominant so that two effects are achieved: firstly, the risk of explosion is negligible and secondly, disinhibition can selectively gate the actual signal.

Thus, it is suggested that the regimes strongly dominated by inhibition provide better conditions for efficient signal transmission than the balanced ones. In order to

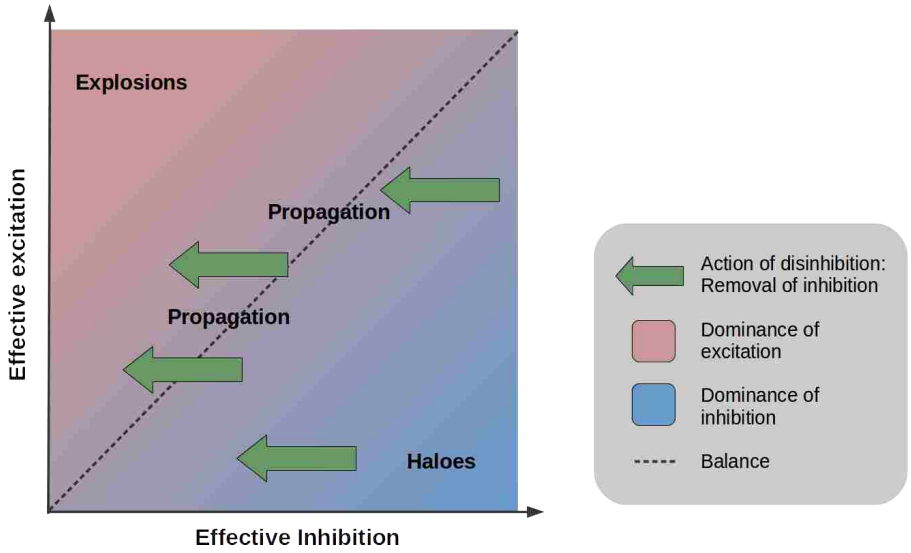

Figure 7.1: A simplified diagram of regimes with various E/I ratio during the signal amplification. Green arrows symbolise the action of disinhibition: it removes inhibition causing the regime to shift horizontally. In the balanced regimes, disinhibition can cause instabilities, whereas the safest region to apply disinhibition is the one dominated by inhibitory haloes.

modulate the signal and briefly break the E/I balance via the disinhibitory mechanisms, the system should have sufficient amount of inhibition to forgo without compromising the global stability.

7.5 The significance of disinhibition

The role of disinhibition was already discussed and we have shown that the disinhibitory pathways can powerfully gate the signals when controlled by neuromodulation. Disinhibition serves as a mean to briefly break the E/I balance in order to provide a window of opportunity for excitation to pass through. In principle, however, the balance can potentially be broken either by decreasing inhibition or increasing excitation. Before discussing this issue further, the very notion of *balance* should be clarified.

The term E/I balance may potentially cause confusion as it might entail that it is created out of two symmetric forces. The E/I balance was in fact compared to the yin yang balance taken from the Chinese philosophy by Northoff (2013). While inhibition indeed cannot live without excitation, the converse is not entirely correct and it was actually phrased that the *sparsened excitation wouldn't do without preceding inhibition* (Northoff, 2013). Throughout this thesis, it was numerously reiterated that the role of

inhibition goes far beyond a mere balancing excitation. In fact, there is a fundamental asymmetry between excitation and inhibition: while inhibition is capable of tracking excitation when a neuron is stimulated, excitation does not compensate for inhibition when the neuron is suppressed (Denève and Machens, 2016). Neurons communicate via spikes, which are 'positive' events and only the excitation can bring the membrane potential closer to the firing threshold. Thus, it has a causal power to induce spikes, whereas inhibition can directly only cause a neuron NOT to spike.

7.5.1 Can disinhibition be equivalent to increased excitation?

In the light of the above E/I asymmetry, a short answer to the question whether disinhibition can be equivalent to increased excitation is: no.

Excitation and inhibition have different tasks and roles to play and as such, breaking the balance between the two is not only about the arithmetics, that is achieving some disbalanced E/I fraction like 1.2:1 which, naturally, can be obtained in two ways: either by increasing the excitatory contributions or decreasing the inhibitory ones. Because the contributions are asymmetric, adding a pool of excitatory contributions which have a causal power is profoundly different from removing contributions that lack such power but can suppress and modulate the signals instead. Disinhibition assumes the removal of the brake holding the already existing drive. Increasing excitation assumes that more drive is being added which has a causal power to evoke spikes and then inhibition.

The reasoning above is rather speculative and only the actual analysis can prove the claims. In the context of the spiking networks we can ask what effect on the spike trains the disinhibition has, compared to the increased excitation. For example, as mentioned earlier, the negative feedback was found to be responsible for the decorrelation of the neighbouring neurons in the network (Tetzlaff et al., 2012). If disinhibition implies removing inhibition, we could describe the decorrelation as a function of the amount of the removed inhibitory inputs via disinhibition and then compare it to the neurons with extra excitatory inputs. If disinhibition indeed is different from the increased excitation, it should be reflected in neurons variability and correlations.

As a matter of fact, selective disinhibition was already compared to other mechanisms in the context of predictive and post hoc attentional selection using the mean field approach (Sridharan and Knudsen, 2015). Specifically, two mechanisms based on excitation were tested: multiplicative input gain and adding an excitatory bias cur-

rent. The input gain mechanism proved to be unable to enhance the responses in the post hoc scenarios, whereas adding the extra current lifted the baseline activity prior to the stimulus presentation increasing the risk of saturation. In conclusion, disinhibitory mechanism was indicated to be the most suited solution to the problem.

7.6 Mapping neurons to functions?

In this thesis, a great deal of experimental evidence was used to justify the extensions of the standard synfire chain model. Although we were inspired by the findings about the interneuron classes, in the model we chiefly employed the connectivity patterns between these, not the detailed characteristics of neurons' physiology.

We show for instance, that the modulation of the disinhibitory pathway is a powerful mean of control of the signal transmission and the circuits involving the VIP or L1 interneurons were a clear inspiration for this model. Although many contemporary studies aim at mapping interneuron classes to function, we deliberately refrain from this. Below, several issues supporting this stance are discussed.

7.6.1 Sampling problem

One problem involves the statistical considerations, as it was communicated that the average statistical power of studies in neuroscience is very low, which severely undermines their credibility (Button et al., 2013). A low statistical power, usually due to small sample sizes or small effects, not only reduces the chances of detecting a true effect, but it also decreases the probability that the statistically significant findings reflect a true effect. How does it relate to the studies on interneurons? Can the findings concerning various subgroups and their functions be considered reliable? Here, only the sample sizes used in the reviewed studies will be scrutinised.

First of all, the interneuron population is inherently small and on average, only 20% of all cortical neurons are inhibitory. Thus, a blind pooling of 1000 neurons would contain a modest sample of 200 interneurons. Then, if we follow a rough estimate, $40\% = 80$ of these should be PV+, $30\% = 60$ should be SOM and $15\% = 30$ should be VIPs (Rudy et al., 2011). Naturally, the actual numbers will differ across layers, regions and species. For instance, it was estimated that in the L2/3 of mice S1 (primary somatosensory cortex) 50% of interneurons are 5HT3aR-expressing (both VIP and non-VIP), and PV+ and SOM account for roughly 25% each (Rudy et al., 2011). Thus,

in a pool of 1000 neurons one would expect 100 5HT3aR-expressing cells, 50 PV+ and 50 SOM cells. Clearly, such estimates should not be taken as definite, especially in relation to studies that used cre-based strategies. Nevertheless, one can reason that in order to image/photostimulate 80 PV+ cells, one needs to cover an area/volume resided by 1000 neurons (again, under the assumption that all the neurons are evenly spaced). Nevertheless, it should be appreciated that the chances of obtaining a representative sample of a given interneuron group are rather low and may widely vary depending on the interneuron type, area or cortical depth.

Another fact undermining the value of the samples is that the present-day cre-mice lines cannot label two or more genetically-defined cell types simultaneously, unless specifically crossed. Hence, in most of the reviewed studies which compared the activity of SOM with PV+ or VIP classes, different and multiple organisms were used to pool individual classes. For instance, one study reported that 6 PV-Cre driver mice were needed to obtain 23 cells and another 6 SOM-Cre driver mice for further 33 cells (Kvitsiani et al., 2013). It remains controversial whether such samples are indeed representative and whether low variability between the animals can be safely assumed. Clearly, individual neurons pooled from different animals are definitely not neighbours, they do not share inputs or outputs, by no means are connected to each other and were not necessarily recorded during the same brain state.

Table 7.1 presents the sample sizes selected from the studies on interneuron functions that used cre-driver mice to label selected subgroups. What is striking is that indeed the sample sizes are alarmingly low. The last study pooled considerably more cells than the rest (Dipoppa et al., 2016), and interestingly, the authors reported significant differences in their results compared to the studies that used much smaller samples. Altogether, it is recommended to wait for studies with higher statistical power reproducing the results before accepting the conclusions as final and definitive.

Study	Technique	Samples description
Kvitsiani et al. (2013)	Optogenetics and tetrode recordings	Optogenetic stimulation: All cells=1339, PV=23 from 6 mice, SOM=13+22 from 6 mice. Behavioural task: All cells = 1034, PV=14 from 4 mice, SOM=31 from 6 mice
Polack et al. (2013)	Two-photon guided whole cell recordings	PV=9, SOM=10 Number of animals not disclosed.
Zhang et al. (2014)	Optogenetic inactivation	PV=9, SOM=10, VIP=9 "Data were from more than three mice in each group"
Fu et al. (2014)	Optogenetics & two-photon Ca^{2+} imaging	First set: VIP=28, non-VIP=77 from 4 mice Second set: VIP=44, non-VIP=76 from 7 mice Locomotion effects: VIP=21, SOM=11, PV=40. Number of mice not disclosed.
Karnani et al. (2016)	Optogenetics & two-photon Ca^{2+} imaging	Morphology analysis: VIP=19, SOM=8, PV=7 Behavioural task: 11 experiments, total VIP=68 VIP activation: 27 experiments, 638 cells, 5% disinhibited, $<1\%$ inhibited. Number of mice not disclosed.
Dipoppa et al. (2016)	Two-photon Ca^{2+} imaging	Pyr=5556, PV=192, VIP=633, SOM=525 from different cortical depths. Number of mice not disclosed.

Table 7.1: Sample size in selected interneuron studies. Selected, chronologically ordered, studies that used cre-driver mice to study functions of interneuron groups.

7.6.2 Neuron vs network doctrine

The history of neuroscience is actually a history of techniques as most of the major breakthroughs were due to a new technique that expanded the possibilities of investigation (Yuste, 2015). As a matter of fact, one cannot discover anything that is beyond the capacity of the method/equipment used. Small-scale single cell recordings over the decades contributed to the ideas which can be labelled as the *neuron doctrine* which asserts that single neurons are the structural and functional units of the nervous system. With the advent of more advanced techniques and theoretical models, cell assemblies

gained more recognition and, as already stated, these are nowadays considered to be the functional units of the nervous system, reinforcing a so-called *network doctrine* (Yuste, 2015).

Presently, the biochemical markers – SOM, PV, VIP – dominate the narrative of the exploration of the neuronal inhibition. Do studies on these markers capture the real picture of the interneuron population? Or is it merely yet another step in the history of methods, which will become obsolete as soon as a more advanced technique is developed?

It is argued here that dissecting the functions of the separate, genetically-defined interneuron subgroups, which are pooled from multiple animals, can be compared to the methodology of the single neuron doctrine – a some sort of a *single–subgroup doctrine*. Similarly, Kumar et al. (2013) rejected the view that in complex and recurrent networks like the brain one can find a *single neuron-type/single-function relationship*. Only in the case of a purely feedforward network one can fully learn the interactions between the elements solely by modulating one element at a time. In a recurrent network of n elements, however, this simplistic approach is insufficient, as in order to fully uncover the actual interactions between all the elements, one would need to simultaneously modulate all combinations of $1, 2, 3, ..., (n-1)$ elements. The total number of such subsets is given by Bell's number, which grows faster than exponentially with n (Kumar et al., 2013).

7.6.3 Building blocks of networks

Instead of finding a mapping, we claim to have studied so-called building blocks of neuronal networks, which can be used to perform canonical computations (Miller, 2016). Circuit motifs like lateral and feedforward inhibition or disinhibition can be useful in controlling the cortical processing and some neuron classes indeed repeatedly take part in these. However, when these building blocks are assembled together and are under the influence of multiple neuromodulators, some novel emergent properties might appear that may change the responses of individual neurons/classes considerably. Thus, a rigid class-to-function mapping should not be regarded as a final goal in deciphering such a complex, multi-dimensional network like the brain.

Ironically, the presented results inadvertently seem to validate the intuition about three separate tasks for inhibition. Firstly, we have shown a massive benefit of introducing a dedicated group of neurons that serve disinhibition and is subject to neuro-

modulation. Then, we demonstrated that relaxing feedforward inhibition can invoke replay, but at the same time it destabilises the whole network. Thus, it was hypothesised that yet another type of inhibition – fast and not subject to modulation – might be necessary to maintain the global stability. Disinhibition is typically mapped to the VIP/L1 classes, FFI and lateral inhibition usually involves the SOM interneurons, and the fast inhibition is usually associated with the PV+ interneurons. The keyword here is *usually*, as there is plenty of examples where such clear separation is invalid. For instance, VIP and SOM classes inhibit each other and under some conditions it might appear that the SOM interneurons serve disinhibition (Dipoppa et al., 2016).

Thus, it is asserted that the complex behaviour of the networks is the effect of the interplay of various building blocks. The total functionality of the system cannot be reduced to some arithmetical sum of the functions of individual elements. Instead, it is the effect of the emergent properties of the recurrently connected combination of elements.

7.7 Biological realism of interneuron classes in spiking networks

The current model can be readily extended to explicitly model the physiology of individual interneuron classes to reproduce some experimental results. In Chapter 6, the differentiation of fast and slow inhibition was already discussed and it was suggested that introducing these two modes via modelling the distinct physiology of the SOM and PV+ interneuron classes might be a necessary next step to realise.

Another issue involves topology. In the model, although we imposed distinct connectivity patterns, all neurons maintained random, 5% global connectivity with the fixed number of inputs (in-degree). The experiments, however, suggest that there is a much higher specificity in the patterns. For instance, the VIP interneurons target only other interneuron classes and avoid the excitatory cells, whereas the SOM interneurons seem to connect to nearly all neighbouring pyramidal cells (Pfeffer et al., 2013; Karnani et al., 2014).

It should be borne in mind that the spiking balanced networks were shown to maintain the AI (asynchronous irregular) state under the assumption that the connectivity is sparse and the neurons' in-degree is homogenous. Recently it was shown that by relaxing this rule, that is by allowing the number of inputs be drawn from a wide dis-

tribution, the E/I balance is disrupted (Landau et al., 2016). Homeostatic plasticity and spike-frequency adaptation were demonstrated to mitigate the problem.

Another consideration refers to the fact that the interneurons do fire intrinsically, at the absence of any inputs. As already mentioned, this feature is not integrated into the IAF neuron model, but might have serious consequences in reality. We have dealt with various examples where the haloes were able to disrupt the whole network's activity. Once the interneurons received strong inhibition, they would get hyperpolarised and become unresponsive to the further activation. Also, inactivating the FFI reduced the total amount of inhibition and promoted the formation of explosions in the scenarios with the replay. It is plausible that allowing interneurons to emit spikes intrinsically would increase the amount of global inhibition and presumably help in preventing the abnormal behaviours. This prediction, however, is impossible to verify with the current IAF model.

7.8 Neuromodulation

We have investigated two modes of cholinergic signalling: volume and phasic transmission. Although the phasic transmission proved to be consistently more efficient in modulating specific cell assemblies, the difference was rather minuscule. The core of the problem turned out to be related to the shared inhibitory pools and neither mode had capacity to overcome it. Thus, our results cannot be used to compare the two modes. In principle, the volume transmission mode proved to be sufficient to control the specific disinhibitory pathways and as such, there is no reason to employ more specific, well-timed phasic transmission instead.

Although we assumed that it is the cholinergic modulation that is being investigated, the studied mechanisms can as well apply to other neuromodulators. Specifically, the disinhibitory switch that was the main target of the cholinergic modulation in the model, was hypothesised to be composed of the VIP or L1 interneuron classes. Noteworthy, both classes were shown to also robustly express the ionotropic serotonin receptor 5HT3a (Lee et al., 2010). Thus, it is plausible that the serotonergic modulatory system may also exert control on the disinhibitory signalling in the cortex. Serotonin is critically involved in various cortical functions and was found to be implicated in cognition, mood and impulse control (Celada et al., 2013). The combination of two neuromodulators regulating a circuit is a common motif and the cholinergic together with norepinephrine modulation was already proposed to be responsible for controlling

two types of uncertainty (Angela and Dayan, 2005).

7.9 Back to cell assemblies – how can they be created?

Although classically it is assumed that the assembly formation takes place during the actual experience, it somewhat implicitly assumes that at the time of learning the network encoding the memory is perfectly random and uniform, which naturally cannot be true. Neuronal networks are in a constant flux, the synaptic weights continually change, and neurons and synapses are being reused in the ever-changing environments and contexts (McKenzie and Eichenbaum, 2011). In such setting, a local network will always be tainted with the already existing non-random structures which were indeed shown to influence the recruitment of neurons to form new assemblies during the encoding novel memories (Grosmark and Buzsáki, 2016; Holtmaat and Caroni, 2016).

However, the very fact that the non-random elements are plausible to be found does not immediately entail that these elements can be large and highly ordered. Synfire chains are very specific structures which require hundreds of neurons and thousands of synapses to become a fully operational feedforward network. Recent theoretical work on the spike timing dependent plasticity (STDP) demonstrated that the synfire chains are in fact likely to be formed in random recurrent networks (Tannenbaum and Burak, 2016). Remarkably, it was suggested that such structures can also emerge autonomously, without the need to provide structured inputs to the network during learning phases. This study, however, investigated only the formation of the standard synfire chains, composed exclusively of excitatory neurons. Excitatory synapses were also the sole targets of the plasticity dynamics.

Inhibitory synapses, on the other hand, were also found to be plastic, and it is believed that it is a crucial feature in maintaining the E/I balance (Vogels et al., 2011; D'amour and Froemke, 2015). Also, various inhibitory and disinhibitory mechanisms were indicated as the key factors in controlling network activity and thus in shaping cell assembly formation (Holtmaat and Caroni, 2016). Furthermore, neuromodulators were numerously shown to be implicated in plasticity including STDP and long term potentiation and depression (LTP and LTD) (Hasselmo, 2006; Seol et al., 2007; Pawlak et al., 2010).

In conclusion, standard synfire chains appear to be likely to be formed in random networks. Since the inhibitory mechanisms and neuromodulation targeting the interneuron pools were shown to be the key factors in plasticity and assemblies for-

mation, it is speculated that the augmented synfire chains should also be likely to appear. Overall, the presence of inhibitory and disinhibitory pathways that are controlled by neuromodulators might not only support the signal transmission, but actually contribute to the learning mechanisms, such as strengthening the connections along the chosen chain out of many overlapping assemblies. Future studies should extend the present model by adding synaptic plasticity to elucidate potential effects of inhibitory mechanisms on the chain formation.

Bibliography

Abbott, L. F. (1999). Lapicque's introduction of the integrate-and-fire model neuron (1907). *Brain research bulletin*, 50(5):303–304.

Abbott, L. F. (2008). Theoretical neuroscience rising. *Neuron*, 60(3):489–495.

Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex*. Cambridge University Press.

Agnati, L., Zoli, M., Strömberg, I., and Fuxe, K. (1995). Intercellular communication in the brain: wiring versus volume transmission. *Neuroscience*, 69(3):711–726.

Albuquerque, E. X., Pereira, E. F., Alkondon, M., and Rogers, S. W. (2009). Mammalian nicotinic acetylcholine receptors: from structure to function. *Physiological reviews*, 89(1):73–120.

Alitto, H. J. and Dan, Y. (2012). Cell-type-specific modulation of neocortical activity by basal forebrain input. *Frontiers in systems neuroscience*, 6.

Angela, J. Y. and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, 46(4):681–692.

Arroyo, S., Bennett, C., Aziz, D., Brown, S. P., and Hestrin, S. (2012). Prolonged disynaptic inhibition in the cortex mediated by slow, non-α7 nicotinic excitation of a specific subset of cortical interneurons. *The Journal of Neuroscience*, 32(11):3859–3864.

Arroyo, S., Bennett, C., and Hestrin, S. (2014). Nicotinic modulation of cortical circuits. *Frontiers in neural circuits*, 8.

Ascoli, G. A., Alonso-Nanclares, L., Anderson, S. A., Barrionuevo, G., Benavides-Piccione, R., Burkhalter, A., Buzsáki, G., Cauli, B., DeFelipe, J., Fairén, A., et al.

(2008). Petilla terminology: nomenclature of features of gabaergic interneurons of the cerebral cortex. *Nature Reviews Neuroscience*, 9(7):557–568.

Aviel, Y., Mehring, C., Abeles, M., and Horn, D. (2003). On embedding synfire chains in a balanced network. *Neural computation*, 15(6):1321–1340.

Bacci, A., Huguenard, J. R., and Prince, D. A. (2005). Modulation of neocortical interneurons: extrinsic influences and exercises in self-control. *Trends in neurosciences*, 28(11):602–610.

Basu, J., Zaremba, J. D., Cheung, S. K., Hitti, F. L., Zemelman, B. V., Losonczy, A., and Siegelbaum, S. A. (2016). Gating of hippocampal activity, plasticity, and memory by entorhinal cortex long-range inhibition. *Science*, 351(6269):aaa5694.

Bonifazi, P., Goldin, M., Picardo, M., Jorquera, I., Cattani, A., Bianconi, G., Represa, A., Ben-Ari, Y., and Cossart, R. (2009). Gabaergic hub neurons orchestrate synchrony in developing hippocampal networks. *Science*, 326(5958):1419–1424.

Bowery, N. and Smart, T. (2006). Gaba and glycine as neurotransmitters: a brief history. *British journal of pharmacology*, 147(S1):S109–S119.

Brette, R. (2015). Philosophy of the spike: Rate-based vs. spike-based theories of the brain. *Frontiers in systems neuroscience*, 9.

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

Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S., and Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5):365–376.

Butts, D. A., Weng, C., Jin, J., Yeh, C.-I., Lesica, N. A., Alonso, J.-M., and Stanley, G. B. (2007). Temporal precision in the neural code and the timescales of natural vision. *Nature*, 449(7158):92–95.

Buzsáki, G. (2004). Large-scale recording of neuronal ensembles. *Nature neuroscience*, 7(5):446–451.

Buzsáki, G. (2010). Neural syntax: cell assemblies, synapsembles, and readers. *Neuron*, 68(3):362–385.

Cannon, J., Kopell, N., Gardner, T., and Markowitz, J. (2015). Neural sequence generation using spatiotemporal patterns of inhibition. *PLoS Comput Biol*, 11(11):e1004581.

Carr, M. F., Jadhav, S. P., and Frank, L. M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature neuroscience*, 14(2):147–153.

Carrillo-Reid, L., Miller, J.-e. K., Hamm, J. P., Jackson, J., and Yuste, R. (2015). Endogenous sequential cortical activity evoked by visual stimuli. *The Journal of Neuroscience*, 35(23):8813–8828.

Câteau, H. and Fukai, T. (2001). Fokker–planck approach to the pulse packet propagation in synfire chain. *Neural networks*, 14(6):675–685.

Cavallari, S., Panzeri, S., and Mazzoni, A. (2014). Comparison of the dynamics of neural interactions between current-based and conductance-based integrate-and-fire recurrent networks. *Front Neural Circuits*, 8:12.

Celada, P., Puig, M. V., and Artigas, F. (2013). Serotonin modulation of cortical neurons and networks. *Frontiers in Integrative Neuroscience*, 7.

Chatham, C. H. and Badre, D. (2015). Multiple gates on working memory. *Current opinion in behavioral sciences*, 1:23–31.

Chenkov, N., Sprekeler, H., and Kempter, R. (2016). Memory replay in balanced recurrent networks. *bioRxiv*, page 069641.

Chevalier, G. and Deniau, J. (1990). Disinhibition as a basic process in the expression of striatal functions. *Trends in neurosciences*, 13(7):277–280.

Cossart, R. (2014). Operational hub cells: a morpho-physiologically diverse class of gabaergic neurons united by a common function. *Current opinion in neurobiology*, 26:51–56.

Cossart, R., Dinocourt, C., Hirsch, J., Merchan-Perez, A., De Felipe, J., Ben-Ari, Y., Esclapez, M., and Bernard, C. (2001). Dendritic but not somatic gabaergic inhibition is decreased in experimental epilepsy. *Nature neuroscience*, 4(1):52–62.

Crick, F. (1988). *What mad pursuit: A Personal View of Scientific Discovery*. Basic Books.

Cruikshank, S. J., Lewis, T. J., and Connors, B. W. (2007). Synaptic basis for intense thalamocortical activation of feedforward inhibitory cells in neocortex. *Nature neuroscience*, 10(4):462–468.

De Carlos, J. A. and Borrell, J. (2007). A historical reflection of the contributions of cajal and golgi to the foundations of neuroscience. *Brain research reviews*, 55(1):8–16.

DeFelipe, J. (2002). Cortical interneurons: from cajal to 2001. *Progress in brain research*, 136:215–238.

DeFelipe, J., López-Cruz, P. L., Benavides-Piccione, R., Bielza, C., Larrañaga, P., Anderson, S., Burkhalter, A., Cauli, B., Fairén, A., Feldmeyer, D., et al. (2013). New insights into the classification and nomenclature of cortical gabaergic interneurons. *Nature Reviews Neuroscience*, 14(3):202–216.

Deisseroth, K. (2011). Optogenetics. *Nature methods*, 8(1):26–29.

Denève, S. and Machens, C. K. (2016). Efficient codes and balanced networks. *Nature neuroscience*, 19(3):375–382.

Destexhe, A. and Contreras, D. (2006). Neuronal computations with stochastic network states. *Science*, 314(5796):85–90.

Destexhe, A., Rudolph, M., and Paré, D. (2003). The high-conductance state of neocortical neurons in vivo. *Nature reviews neuroscience*, 4(9):739–751.

Diba, K. and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature neuroscience*, 10(10):1241–1242.

Diesmann, M., Gewaltig, M.-O., and Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402(6761):529–533.

Dipoppa, M., Ranson, A., Krumin, M., Pachitariu, M., Carandini, M., and Harris, K. D. (2016). Vision and locomotion shape the interactions between neuron types in mouse visual cortex. *bioRxiv*.

Doiron, B., Litwin-Kumar, A., Rosenbaum, R., Ocker, G. K., and Josić, K. (2016). The mechanics of state-dependent neural correlations. *Nature neuroscience*, 19(3):383–393.

Dragoi, G. and Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469(7330):397–401.

Durstewitz, D., Seamans, J. K., and Sejnowski, T. J. (2000). Neurocomputational models of working memory. *Nature neuroscience*, 3:1184–1191.

D'amour, J. A. and Froemke, R. C. (2015). Inhibitory and excitatory spike-timing-dependent plasticity in the auditory cortex. *Neuron*, 86(2):514–528.

Ecker, A. S., Berens, P., Keliris, G. A., Bethge, M., Logothetis, N. K., and Tolias, A. S. (2010). Decorrelated neuronal firing in cortical microcircuits. *Science*, 327(5965):584–587.

Eglen, R. M. (2005). Muscarinic receptor subtype pharmacology and physiology. *Progress in medicinal chemistry*, 43:105–136.

Eichenbaum, H. (2015). Does the hippocampus preplay memories? *Nature neuroscience*, 18(12):1701–1702.

Engel, J. (1996). Excitation and inhibition in epilepsy. *Canadian Journal of Neurological Sciences/Journal Canadien des Sciences Neurologiques*, 23(03):167–174.

Fishell, G. and Rudy, B. (2011). Mechanisms of inhibition within the telencephalon:"where the wild things are". *Annual review of neuroscience*, 34:535.

Frank, J. G. and Mendelowitz, D. (2012). Synaptic and intrinsic activation of gabaergic neurons in the cardiorespiratory brainstem network. *PloS one*, 7(5):e36459.

Freund, T. F. and Katona, I. (2007). Perisomatic inhibition. *Neuron*, 56(1):33–42.

Froemke, R. C. (2015). Plasticity of cortical excitatory-inhibitory balance. *Annual review of neuroscience*, 38:195.

Fu, Y., Kaneko, M., Tang, Y., Alvarez-Buylla, A., and Stryker, M. P. (2015). A cortical disinhibitory circuit for enhancing adult plasticity. *Elife*, 4:e05558.

Fu, Y., Tucciarone, J. M., Espinosa, J. S., Sheng, N., Darcy, D. P., Nicoll, R. A., Huang, Z. J., and Stryker, M. P. (2014). A cortical circuit for gain control by behavioral state. *Cell*, 156(6):1139–1152.

Gais, S. and Born, J. (2004). Low acetylcholine during slow-wave sleep is critical for declarative memory consolidation. *Proceedings of the National Academy of Sciences*, 101(7):2140–2144.

Gewaltig, M.-O. and Diesmann, M. (2007). Nest (neural simulation tool). *Scholarpedia*, 2(4):1430.

Gewaltig, M.-O., Diesmann, M., and Aertsen, A. (2001). Propagation of cortical synfire activity: survival probability in single trials and stability in the mean. *Neural networks*, 14(6):657–673.

Goldman, M. S. (2009). Memory without feedback in a neural network. *Neuron*, 61(4):621–634.

Gritton, H. J., Howe, W. M., Mallory, C. S., Hetrick, V. L., Berke, J. D., and Sarter, M. (2016). Cortical cholinergic signaling controls the detection of cues. *Proceedings of the National Academy of Sciences*, page 201516134.

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

Hahnloser, R. H., Kozhevnikov, A. A., and Fee, M. S. (2002). An ultra-sparse code underliesthe generation of neural sequences in a songbird. *Nature*, 419(6902):65–70.

Hangya, B., Pi, H.-J., Kvitsiani, D., Ranade, S. P., and Kepecs, A. (2014). From circuit motifs to computations: mapping the behavioral repertoire of cortical interneurons. *Current opinion in neurobiology*, 26:117–124.

Hangya, B., Ranade, S. P., Lorenc, M., and Kepecs, A. (2015). Central cholinergic neurons are rapidly recruited by reinforcement feedback. *Cell*, 162(5):1155–1168.

Hansel, D. and Mato, G. (2013). Short-term plasticity explains irregular persistent activity in working memory tasks. *The Journal of Neuroscience*, 33(1):133–149.

Harris, K. D. and Thiele, A. (2011). Cortical state and attention. *Nature reviews neuroscience*, 12(9):509–523.

Harvey, C. D., Coen, P., and Tank, D. W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*, 484(7392):62–68.

Hasselmo, M. E. (2006). The role of acetylcholine in learning and memory. *Current Opinion in Neurobiology*, 16:710–715.

Hasselmo, M. E., Anderson, B. P., and Bower, J. M. (1992). Cholinergic modulation of cortical associative memory function. *Journal of Neurophysiology*, 67(5):1230–1246.

Häusser, M. and Clark, B. A. (1997). Tonic synaptic inhibition modulates neuronal output pattern and spatiotemporal synaptic integration. *Neuron*, 19(3):665–678.

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. Wiley, New York.

Henny, P. and Jones, B. E. (2008). Projections from basal forebrain to prefrontal cortex comprise cholinergic, gabaergic and glutamatergic inputs to pyramidal cells or interneurons. *European Journal of Neuroscience*, 27(3):654–670.

Holtmaat, A. and Caroni, P. (2016). Functional and structural underpinnings of neuronal assembly formation in learning. *Nature Neuroscience*.

Hu, H., Cavendish, J. Z., and Agmon, A. (2013). Not all that glitters is gold: off-target recombination in the somatostatin–ires-cre mouse line labels a subset of fast-spiking interneurons. *Frontiers in neural circuits*, 7:195.

Hu, H., Gan, J., and Jonas, P. (2014). Fast-spiking, parvalbumin+ gabaergic interneurons: From cellular design to microcircuit function. *Science*, 345(6196):1255263.

Ibrahim, L. A., Mesik, L., Ji, X.-y., Fang, Q., Li, H.-f., Li, Y.-t., Zingg, B., Zhang, L. I., and Tao, H. W. (2016). Cross-modality sharpening of visual cortical processing through layer-1-mediated inhibition and disinhibition. *Neuron*, 89(5):1031–1045.

Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Ferster, D., and Yuste, R. (2004). Synfire chains and cortical songs: temporal modules of cortical activity. *Science*, 304(5670):559–564.

Isaacson, J. S. and Scanziani, M. (2011). How inhibition shapes cortical activity. *Neuron*, 72(2):231–243.

Iurilli, G., Ghezzi, D., Olcese, U., Lassi, G., Nazzaro, C., Tonini, R., Tucci, V., Benfenati, F., and Medini, P. (2012). Sound-driven synaptic inhibition in primary visual cortex. *Neuron*, 73(4):814–828.

Izhikevich, E. M. (2004). Which model to use for cortical spiking neurons? *IEEE transactions on neural networks*, 15(5):1063–1070.

Jackson, J., Ayzenshtat, I., Karnani, M. M., and Yuste, R. (2016). Vip+ interneurons control neocortical activity across brain states. *Journal of neurophysiology*, 115(6):3008–3017.

Jacoby, L. L., Kelley, C. M., and McElree, B. D. (1999). The role of cognitive control: Early selection versus late correction.

Kaczmarek, L. K. and Levitan, I. B. (1987). *Neuromodulation: the biochemical control of neuronal excitability*. Oxford University Press, USA.

Kametani, H. and Kawamura, H. (1990). Alterations in acetylcholine release in the rat hippocampus during sleep-wakefulness detected by intracerebral dialysis. *Life sciences*, 47(5):421–426.

Karnani, M. M., Agetsuma, M., and Yuste, R. (2014). A blanket of inhibition: functional inferences from dense inhibitory connectivity. *Current opinion in neurobiology*, 26:96–102.

Karnani, M. M., Jackson, J., Ayzenshtat, I., Sichani, A. H., Manoocheri, K., Kim, S., and Yuste, R. (2016). Opening holes in the blanket of inhibition: Localized lateral disinhibition by vip interneurons. *The Journal of Neuroscience*, 36(12):3471–3480.

Kayser, C., Petkov, C. I., Augath, M., and Logothetis, N. K. (2005). Integration of touch and sound in auditory cortex. *Neuron*, 48(2):373–384.

Kepecs, A. and Fishell, G. (2014). Interneuron cell types are fit to function. *Nature*, 505(7483):318–326.

Kleinfeld, D. (1986). Sequential state generation by model neural networks. *Proceedings of the National Academy of Sciences*, 83(24):9469–9473.

Kremkow, J., Aertsen, A., and Kumar, A. (2010a). Gating of signal propagation in spiking neural networks by balanced and correlated excitation and inhibition. *The Journal of neuroscience*, 30(47):15760–15768.

Kremkow, J., Perrinet, L. U., Masson, G. S., and Aertsen, A. (2010b). Functional consequences of correlated excitatory and inhibitory conductances in cortical networks. *Journal of computational neuroscience*, 28(3):579–594.

Kruglikov, I. and Rudy, B. (2008). Perisomatic gaba release and thalamocortical integration onto neocortical excitatory cells are regulated by neuromodulators. *Neuron*, 58(6):911–924.

Kuhn, A., Aertsen, A., and Rotter, S. (2004). Neuronal integration of synaptic input in the fluctuation-driven regime. *The Journal of neuroscience*, 24(10):2345–2356.

Kumar, A., Rotter, S., and Aertsen, A. (2008a). Conditions for propagating synchronous spiking and asynchronous firing rates in a cortical network model. *The Journal of neuroscience*, 28(20):5268–5280.

Kumar, A., Rotter, S., and Aertsen, A. (2010). Spiking activity propagation in neuronal networks: reconciling different perspectives on neural coding. *Nature reviews neuroscience*, 11(9):615–627.

Kumar, A., Schrader, S., Aertsen, A., and Rotter, S. (2008b). The high-conductance state of cortical networks. *Neural computation*, 20(1):1–43.

Kumar, A., Vlachos, I., Aertsen, A., and Boucsein, C. (2013). Challenges of understanding brain function by selective modulation of neuronal subpopulations. *Trends in neurosciences*, 36(10):579–586.

Kvitsiani, D., Ranade, S., Hangya, B., Taniguchi, H., Huang, J., and Kepecs, A. (2013). Distinct behavioural and network correlates of two interneuron types in prefrontal cortex. *Nature*, 498(7454):363–366.

Landau, I. D., Egger, R., Dercksen, V. J., Oberlaender, M., and Sompolinsky, H. (2016). The impact of structural heterogeneity on excitation-inhibition balance in cortical networks. *Neuron*, 92.

Lapicque, L. (1907). Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation. *J. Physiol. Pathol. Gen*, 9(1):620–635.

Lee, S., Hjerling-Leffler, J., Zagha, E., Fishell, G., and Rudy, B. (2010). The largest group of superficial neocortical gabaergic interneurons expresses ionotropic serotonin receptors. *The Journal of Neuroscience*, 30(50):16796–16808.

Lee, S., Kruglikov, I., Huang, Z. J., Fishell, G., and Rudy, B. (2013). A disinhibitory circuit mediates motor integration in the somatosensory cortex. *Nature neuroscience*, 16(11):1662–1670.

Lee, S.-H. and Dan, Y. (2012). Neuromodulation of brain states. *Neuron*, 76(1):209–222.

Letzkus, J. J., Wolff, S. B., and Lüthi, A. (2015). Disinhibition, a circuit mechanism for associative learning and memory. *Neuron*, 88(2):264–276.

Letzkus, J. J., Wolff, S. B., Meyer, E. M., Tovote, P., Courtin, J., Herry, C., and Lüthi, A. (2011). A disinhibitory microcircuit for associative fear learning in the auditory cortex. *Nature*, 480(7377):331–335.

Lim, S. and Goldman, M. S. (2012). Noise tolerance of attractor and feedforward memory models. *Neural computation*, 24(2):332–390.

Litvak, V., Sompolinsky, H., Segev, I., and Abeles, M. (2003). On the transmission of rate code in long feedforward networks with excitatory–inhibitory balance. *The Journal of neuroscience*, 23(7):3006–3015.

Loewi, O. (1921). Über humorale übertragbarkeit der herznervenwirkung. *Pflügers Archiv European Journal of Physiology*, 189(1):239–242.

Luczak, A., McNaughton, B. L., and Harris, K. D. (2015). Packet-based communication in the cortex. *Nature Reviews Neuroscience*.

Luongo, F. J., Zimmerman, C. A., Horn, M. E., and Sohal, V. S. (2016). Correlations between prefrontal neurons form a small-world network that optimizes the generation of multineuron sequences of activity. *Journal of neurophysiology*, 115(5):2359–2375.

Maass, W. (2000). On the computational power of winner-take-all. *Neural computation*, 12(11):2519–2535.

MacDonald, C. J., Lepage, K. Q., Eden, U. T., and Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron*, 71(4):737–749.

Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84.

Marder, E. (2015). Understanding brains: details, intuition, and big data. *PLoS Biol*, 13(5):e1002147.

Marín, O. (2012). Interneuron dysfunction in psychiatric disorders. *Nature Reviews Neuroscience*, 13(2):107–120.

Markowitz, J. E., Liberti III, W. A., Guitchounts, G., Velho, T., Lois, C., and Gardner, T. J. (2015). Mesoscopic patterns of neural activity support songbird cortical sequences. *PLoS Biol*, 13(6):e1002158.

Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nature Reviews Neuroscience*, 5(10):793–807.

Marrosu, F., Portas, C., Mascia, M. S., Casu, M. A., Fà, M., Giagheddu, M., Imperato, A., and Gessa, G. L. (1995). Microdialysis measurement of cortical and hippocampal acetylcholine release during sleep-wake cycle in freely moving cats. *Brain research*, 671(2):329–332.

Martinello, K., Huang, Z., Lujan, R., Tran, B., Watanabe, M., Cooper, E. C., Brown, D. A., and Shah, M. M. (2015). Cholinergic afferent stimulation induces axonal function plasticity in adult hippocampal granule cells. *Neuron*, 85(2):346–363.

McCoy, A. N. and Tan, Y. S. (2014). Otto loewi (1873–1961): Dreamer and nobel laureate. *Singapore medical journal*, 55(1):3.

McKenzie, S. and Eichenbaum, H. (2011). Consolidation and reconsolidation: two lives of memories? *Neuron*, 71(2):224–233.

Mehring, C., Hehl, U., Kubo, M., Diesmann, M., and Aertsen, A. (2003). Activity dynamics and propagation of synchronous spiking in locally connected random networks. *Biological cybernetics*, 88(5):395–408.

Mesulam, M., Mufson, E., Wainer, B., and Levey, A. (1983). Central cholinergic pathways in the rat: an overview based on an alternative nomenclature (ch1–ch6). *Neuroscience*, 10(4):1185–1201.

Miller, K. D. (2016). Canonical computations of cerebral cortex. *Current opinion in neurobiology*, 37:75–84.

Mitchell, J. F., Sundberg, K. A., and Reynolds, J. H. (2007). Differential attention-dependent response modulation across cell classes in macaque visual area v4. *Neuron*, 55(1):131–141.

Mokeichev, A., Okun, M., Barak, O., Katz, Y., Ben-Shahar, O., and Lampl, I. (2007). Stochastic emergence of repeating cortical motifs in spontaneous membrane potential fluctuations in vivo. *Neuron*, 53(3):413–425.

Newman, E. L., Gupta, K., Climer, J. R., Monaghan, C. K., and Hasselmo, M. E. (2012). Cholinergic modulation of cognitive processing: insights drawn from computational models. *Frontiers in behavioral neuroscience*, 6.

Northoff, G. (2013). *Unlocking the brain: Volume 1: Coding*. Oxford University Press.

O'Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain research*, 34(1):171–175.

Parikh, V., Kozak, R., Martinez, V., and Sarter, M. (2007). Prefrontal acetylcholine release controls cue detection on multiple timescales. *Neuron*, 56(1):141–154.

Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, 321(5894):1322–1327.

Pawlak, V., Wickens, J. R., Kirkwood, A., and Kerr, J. N. (2010). Timing is not everything: neuromodulation opens the stdp gate. *Frontiers in synaptic neuroscience*, 2.

Perkel, D. H. and Bullock, T. H. (1968). Neural coding. *Neurosciences Research Program Bulletin*.

Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S. I., and Battaglia, F. P. (2009). Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature neuroscience*, 12(7):919–926.

Pfeffer, C. K., Xue, M., He, M., Huang, Z. J., and Scanziani, M. (2013). Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nature neuroscience*, 16(8):1068–1076.

Pi, H.-J., Hangya, B., Kvitsiani, D., Sanders, J. I., Huang, Z. J., and Kepecs, A. (2013). Cortical interneurons that specialize in disinhibitory control. *Nature*, 503(7477):521–524.

Picciotto, M. R., Higley, M. J., and Mineur, Y. S. (2012). Acetylcholine as a neuromodulator: cholinergic signaling shapes nervous system function and behavior. *Neuron*, 76(1):116–129.

Polack, P.-O., Friedman, J., and Golshani, P. (2013). Cellular mechanisms of brain state-dependent gain modulation in visual cortex. *Nature neuroscience*, 16(9):1331–1339.

Poo, M.-m., Pignatelli, M., Ryan, T. J., Tonegawa, S., Bonhoeffer, T., Martin, K. C., Rudenko, A., Tsai, L.-H., Tsien, R. W., Fishell, G., et al. (2016). What is memory? the present state of the engram. *BMC biology*, 14(1):1.

Poorthuis, R. B., Enke, L., and Letzkus, J. J. (2014). Cholinergic circuit modulation through differential recruitment of neocortical interneuron types during behaviour. *The Journal of physiology*, 592(19):4155–4164.

Rajan, K., Harvey, C. D., and Tank, D. W. (2016). Recurrent network models of sequence generation and memory. *Neuron*, 90(1):128–142.

Redish, A. D., Elga, A. N., and Touretzky, D. S. (1996). A coupled attractor model of the rodent head direction system. *Network: Computation in Neural Systems*, 7(4):671–685.

Reimer, J., Froudarakis, E., Cadwell, C. R., Yatsenko, D., Denfield, G. H., and Tolias, A. S. (2014). Pupil fluctuations track fast switching of cortical states during quiet wakefulness. *Neuron*, 84(2):355–362.

Roux, L. and Buzsáki, G. (2015). Tasks for inhibitory interneurons in intact brain circuits. *Neuropharmacology*, 88:10–23.

Roxin, A., Hakim, V., and Brunel, N. (2008). The statistics of repeating patterns of cortical activity can be reproduced by a model network of stochastic binary neurons. *The Journal of neuroscience*, 28(42):10734–10745.

Rudy, B., Fishell, G., Lee, S., and Hjerling-Leffler, J. (2011). Three groups of interneurons account for nearly 100% of neocortical gabaergic neurons. *Developmental neurobiology*, 71(1):45–61.

Sadovsky, A. J. and MacLean, J. N. (2014). Mouse visual neocortex supports multiple stereotyped patterns of microcircuit activity. *The Journal of Neuroscience*, 34(23):7769–7777.

Samsonovich, A. and McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *The Journal of neuroscience*, 17(15):5900–5920.

Sarter, M., Lustig, C., Berry, A. S., Gritton, H., Howe, W. M., and Parikh, V. (2016). What do phasic cholinergic signals do? *Neurobiology of learning and memory*, 130:135–141.

Sarter, M., Parikh, V., and Howe, W. M. (2009). Phasic acetylcholine release and the volume transmission hypothesis: time to move on. *Nat Rev Neurosci*, 10(5):383–390.

Sato, H., Hata, Y., Masui, H., and Tsumoto, T. (1987). A functional role of cholinergic innervation to neurons in the cat visual cortex. *Journal of Neurophysiology*, 58(4):765–780.

Semon, R. W. (1921). *The mneme*. G. Allen & Unwin Limited.

Seol, G. H., Ziburkus, J., Huang, S., Song, L., Kim, I. T., Takamiya, K., Huganir, R. L., Lee, H.-K., and Kirkwood, A. (2007). Neuromodulators control the polarity of spike-timing-dependent synaptic plasticity. *Neuron*, 55(6):919–929.

Shadlen, M. N. and Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Current opinion in neurobiology*, 4(4):569–579.

Shafi, M., Zhou, Y., Quintana, J., Chow, C., Fuster, J., and Bodner, M. (2007). Variability in neuronal activity in primate cortex during working memory tasks. *Neuroscience*, 146(3):1082–1108.

Shinozaki, T., Câteau, H., Urakubo, H., and Okada, M. (2007). Controlling synfire chain by inhibitory synaptic input. *Journal of the Physical Society of Japan*, 76(4):044806.

Silberberg, G. and Markram, H. (2007). Disynaptic inhibition between neocortical pyramidal cells mediated by martinotti cells. *Neuron*, 53(5):735–746.

Silva, D., Feng, T., and Foster, D. J. (2015). Trajectory events across hippocampal place cells require previous experience. *Nature neuroscience*, 18(12):1772–1779.

Softky, W. R. and Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random epsps. *The Journal of Neuroscience*, 13(1):334–350.

Sompolinsky, H. and Kanter, I. (1986). Temporal association in asymmetric neural networks. *Physical Review Letters*, 57(22):2861.

Song, S., Sjöström, P. J., Reigl, M., Nelson, S., and Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol*, 3(3):e68.

Sridharan, D. and Knudsen, E. I. (2015). Selective disinhibition: A unified neural mechanism for predictive and post hoc attentional selection. *Vision research*, 116:194–209.

Stevenson, I. H. and Kording, K. P. (2011). How advances in neural recording affect data analysis. *Nature neuroscience*, 14(2):139–142.

Strüber, M., Jonas, P., and Bartos, M. (2015). Strength and duration of perisomatic gabaergic inhibition depend on distance between synaptically connected cells. *Proceedings of the National Academy of Sciences*, 112(4):1220–1225.

Taniguchi, H., He, M., Wu, P., Kim, S., Paik, R., Sugino, K., Kvitsani, D., Fu, Y., Lu, J., Lin, Y., et al. (2011). A resource of cre driver lines for genetic targeting of gabaergic neurons in cerebral cortex. *Neuron*, 71(6):995–1013.

Tannenbaum, N. R. and Burak, Y. (2016). Shaping neural circuits by high order synaptic interactions. *arXiv preprint arXiv:1605.03005*.

Tetzlaff, T., Geisel, T., and Diesmann, M. (2002). The ground state of cortical feedforward networks. *Neurocomputing*, 44:673–678.

Tetzlaff, T., Helias, M., Einevoll, G. T., and Diesmann, M. (2012). Decorrelation of neural-network activity by inhibitory feedback. *PLoS Comput Biol*, 8(8):e1002596.

Tetzlaff, T., Morrison, A., Geisel, T., and Diesmann, M. (2004). Consequences of realistic network size on the stability of embedded synfire chains. *Neurocomputing*, 58:117–121.

Thiele, A. (2013). Muscarinic signaling in the brain. *Annual review of neuroscience*, 36:271–294.

Tonegawa, S., Liu, X., Ramirez, S., and Redondo, R. (2015). Memory engram cells have come of age. *Neuron*, 87(5):918–931.

Tovote, P., Esposito, M. S., Botta, P., Chaudun, F., Fadok, J. P., Markovic, M., Wolff, S. B., Ramakrishnan, C., Fenno, L., Deisseroth, K., et al. (2016). Midbrain circuits for defensive behaviour. *Nature*.

Tremblay, R., Lee, S., and Rudy, B. (2016). {GABAergic} interneurons in the neocortex: From cellular properties to circuits. *Neuron*, 91(2):260 – 292.

Trengove, C., van Leeuwen, C., and Diesmann, M. (2013). High-capacity embedding of synfire chains in a cortical network model. *Journal of computational neuroscience*, 34(2):185–209.

Urban-Ciecko, J. and Barth, A. L. (2016). Somatostatin-expressing neurons in cortical networks. *Nature Reviews Neuroscience*.

van Rossum, M. C., Turrigiano, G. G., and Nelson, S. B. (2002). Fast propagation of firing rates through layered networks of noisy neurons. *The Journal of neuroscience*, 22(5):1956–1966.

Varela, F., Lachaux, J.-P., Rodriguez, E., and Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nature reviews neuroscience*, 2(4):229–239.

Vogels, T., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science*, 334(6062):1569–1573.

Vogels, T. P. and Abbott, L. (2009). Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nature neuroscience*, 12(4):483–491.

Vogels, T. P. and Abbott, L. F. (2005). Signal propagation and logic gating in networks of integrate-and-fire neurons. *The Journal of neuroscience*, 25(46):10786–10795.

Vogels, T. P., Rajan, K., and Abbott, L. (2005). Neural network dynamics. *Annu. Rev. Neurosci.*, 28:357–376.

Wang, Y., Toledo-Rodriguez, M., Gupta, A., Wu, C., Silberberg, G., Luo, J., and Markram, H. (2004). Anatomical, physiological and molecular properties of mar-

tinotti cells in the somatosensory cortex of the juvenile rat. *The Journal of physiology*, 561(1):65–90.

Wehr, M. and Zador, A. M. (2003). Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature*, 426(6965):442–446.

Wolff, S. B., Gründemann, J., Tovote, P., Krabbe, S., Jacobson, G. A., Müller, C., Herry, C., Ehrlich, I., Friedrich, R. W., Letzkus, J. J., et al. (2014). Amygdala interneuron subtypes control fear learning through disinhibition. *Nature*, 509(7501):453–458.

Xu, H., Jeong, H.-Y., Tremblay, R., and Rudy, B. (2013). Neocortical somatostatin-expressing gabaergic interneurons disinhibit the thalamorecipient layer 4. *Neuron*, 77(1):155–167.

Xue, M., Atallah, B. V., and Scanziani, M. (2014). Equalizing excitation-inhibition ratios across visual cortical neurons. *Nature*, 511(7511):596–600.

y Cajal, S. R. (1995). *Histology of the nervous system of man and vertebrates*, volume 1. Oxford University Press, USA.

Yang, G. R., Murray, J. D., and Wang, X.-J. (2016). A dendritic disinhibitory circuit mechanism for pathway-specific gating. *bioRxiv*, page 041673.

Yizhar, O., Fenno, L. E., Prigge, M., Schneider, F., Davidson, T. J., O'Shea, D. J., Sohal, V. S., Goshen, I., Finkelstein, J., Paz, J. T., et al. (2011). Neocortical excitation/inhibition balance in information processing and social dysfunction. *Nature*, 477(7363):171–178.

Yuste, R. (2015). From the neuron doctrine to neural networks. *Nature Reviews Neuroscience*, 16(8):487–497.

Zaborszky, L., Csordas, A., Mosca, K., Kim, J., Gielow, M. R., Vadasz, C., and Nadasdy, Z. (2015a). Neurons in the basal forebrain project to the cortex in a complex topographic organization that reflects corticocortical connectivity patterns: an experimental study based on retrograde tracing and 3d reconstruction. *Cerebral Cortex*, 25(1):118–137.

Zaborszky, L., Duque, A., Gielow, M., Gombkoto, P., Nadasdy, Z., and Somogyi, J. (2015b). Organization of the basal forebrain cholinergic projection system: specific or diffuse. *The rat nervous system*, pages 491–507.

Zhang, S., Xu, M., Kamigaki, T., Do, J. P. H., Chang, W.-C., Jenvay, S., Miyamichi, K., Luo, L., and Dan, Y. (2014). Long-range and local circuits for top-down modulation of visual cortex processing. *Science*, 345(6197):660–665.