

Bursting dynamics and topological structure of in vitro neuronal networks

Dissertation
zur Erlangung des mathematisch-naturwissenschaftlichen Doktorgrades
"Doctor rerum naturalium"
an der Georg-August-Universität Göttingen

vorgelegt von
Frank Olav Stetter
aus
München

Göttingen, 2012

Mitglied des Betreuungsausschusses (Referent):

Prof. Dr. Theo Geisel, Nichtlineare Dynamik, Max-Planck-Institut für Dynamik und Selbstorganisation

Mitglied des Betreuungsausschusses (Koreferent):

Prof. Elisha Moses, Physics of Complex Systems, Weizmann Institute of Science, Rehovot, Israel

Mitglied des Betreuungsausschusses:

Prof. Dr. Fred Wolf, Nichtlineare Dynamik, Max-Planck-Institut für Dynamik und Selbstorganisation

Tag der mündlichen Prüfung:

Errata

- Introduction, page 1: The direction of signaling between axons and dendrites was stated ambiguously.
The sentence should have been: “Also since Ramón y Cajal, we know that the communication between two cells via a synapse is a one-way process, where an electrical signal from the *axon* of a pre-synaptic cell is communicated to the *dendrites* of a post-synaptic cell.”
- Higher order degree, page 20: The stated formula was incorrect.
The higher order degree is correctly given by $k_i^{(n)} = \sum_j (A^n)_{ji}$, and in particular the second order degree is given by $k_i^{(2)} = \sum_j (A^2)_{ji}$.

Contents

1	Introduction	1
I	Burst onset spiking dynamics	9
2	Discrete time dynamics	11
2.1	Model characteristics	11
2.2	Two types of outcomes	12
2.3	Different topologies	14
2.4	The expected value of the first firing time	16
2.4.1	Prediction of the first firing time	18
2.4.2	Higher order degree	19
2.4.3	Degree-degree correlations	20
2.5	Comparison of numerical and analytical results	21
2.5.1	Random graphs	21
2.5.2	Locally clustered graphs	23
2.5.3	Higher order correlations	23
3	Quorum percolation	27
3.1	Model characteristics	27
3.2	Deriving the degree distribution	29
3.3	Effect of degree-degree correlations	30
4	Leaky integrate-and-fire dynamics	33
4.1	Model characteristics	33
4.2	Adaptive burst rates	34
4.3	Random and local topologies	36
4.4	Analytical prediction of the first firing time	37
4.5	Comparison to simulated data	39
II	Reconstruction of network connectivity	43
5	Neuronal activity as shaped by network structure	45
5.1	Experimental recordings of calcium fluorescence	46
5.2	Simulation of calcium fluorescence	46
5.3	Burst characteristics for different network topologies	49

6	Measuring causality	51
6.1	Linear measures	51
6.2	Non-linear measures	52
6.3	Measuring reconstruction performance	53
6.4	Extending Transfer Entropy	54
6.4.1	Same-bin interactions	54
6.4.2	Identification of the dynamical state	55
6.4.3	Generalized Transfer Entropy	57
6.4.4	Causality measures in comparison to GTE	57
7	GTE reconstruction of simulated synaptic connectivity	59
7.1	GTE performance in clustered networks	59
7.1.1	Non-locally clustered networks	59
7.1.2	Locally clustered networks	61
7.2	Dependence of reconstructed quantities on threshold	63
7.3	Dependence on causality measure and conditioning	64
7.3.1	Non-locally clustered networks	65
7.3.2	Locally clustered networks	65
7.4	Comparison to Bayesian approach	67
7.5	Dependence on extension to Transfer Entropy	67
7.6	Dependence on sample number	69
7.7	Dependence on system size	69
8	Experimental recordings of in vitro networks	73
8.1	Excitatory connectivity of in vitro networks	73
8.1.1	Cultured network at 12 DIV	73
8.1.2	Cultured network at 9 DIV	75
8.2	State-dependency in experimental recordings	75
9	Extending GTE to inhibitory connections	79
9.1	Reconstruction of both inhibitory and excitatory links	79
9.2	Identification of inhibitory links	80
10	Summary	83
III	Predicting dynamics based on reconstructed topology	85
11	The interplay between network structure and dynamics	87
11.1	Determining the onset time from data	87
11.2	The distribution of onset times	88
12	Network reconstruction can predict onset times	91
12.1	Onset time prediction in simulated networks	91
12.2	Onset time prediction in experimental recordings	93
IV	Discussion	95
13	Structure and dynamics	97
13.1	First firing times in silico determined by structure	97

13.2	Connectivity can be reconstructed based on dynamics	99
13.3	First firing times in vitro can be predicted based on GTE reconstruction	104
13.4	Summary	104
14	Outlook	105
14.1	Technological advances	105
14.2	Future work	105
V	Appendix	109
A	Experimental preparation	111
B	Estimators of probability densities	113
B.1	Binless estimators	113
B.2	Application to the estimation of GTE	114
	List of figures	117
	Bibliography	119

Chapter 1

Introduction

Hippocrates (460–379 B.C.) was among the first to describe the brain as the seat of sensation and intelligence [1], but it is only since the 19th century and the work of Ramón y Cajal, that we have been able to isolate the *neuron* as the principal unit of computation in the brain and the *synapse* as the principal mode of interaction among neurons [2]. These two components together give rise to the *neuronal networks* of the human brain, linking together approximately 86 billion neurons, which are supported by about the same number of non-neuronal cells [3–5]. Also since Ramón y Cajal, we know that the communication between two cells via a synapse is a one-way process, with the *dendrites* of the pre-synaptic cell providing the input in the form of an electrical signal, and the *axon* of a post-synaptic cell acting as the output. This electrical pulse in neurons is known as the *action potential* (often simply called a *spike* or neuronal *firing*). It is mediated by *ion channels* in the neuron's membrane, and was studied quantitatively by Hodgkin and Huxley in the squid axon [6,7].

The temporal patterns of action potentials across the population of neurons constitute the *dynamics* of the network. Numerous studies have revealed the *sparsity* of the neural code, i. e. the relatively low average firing frequency of the neuronal population [8]. For instance in rat somatosensory cortex, neurons were found to be active only once for every 20s on average [9,10]. The neuronal network in the brain is also sparse in terms of the connectivity, i. e. the overall probability of connection between two neurons [11]. The reason for both kinds of sparsity is likely the high metabolic cost of the brain itself and its neuronal activity, which is reflected in the fact that even though the brain contributes only about 2% of body mass in adult humans, the energy consumption of the metabolism of the brain at rest is about 20% [12,13].

This sparsity is one of the reasons why there is an increasing interest in neuroscience to uncover the *structure* of neuronal networks, the wiring diagram of neuronal circuitry. This is especially pursued on the large scale of structures in the brain, building on historic findings which have correlated specific neuronal functions with certain regions of the brain by studying patients with lesions in these areas [4,14,15]. A wealth of studies using functional magnetic resonance imaging (fMRI) have shown in recent years that metabolic activity in different regions of the brain correlates with behavioral patterns such as decision-making, learning and even states of self-reflection [16–19]. However, while such studies are valuable to determine in which regions of the brain a given activity is prevalently carried out, they do not allow us to study the microscopic dynamics. This makes it difficult, based on the described methods, to answer at some point what may be called the ultimate question of modern neuroscience: Not *where* does the brain carry out its function, but *how*.

On a finer length scale, it is known that neuronal processes in the brain find a path across centimeters of length guided by chemical gradients and other mechanisms of axon guidance, yielding an enormous degree of *specificity* of the connection structure [20–23]. Recently, in networks




Figure 1.1: Bright-field image of a fully grown *in vitro* network of hippocampal rat neurons at 21 days *in vitro*. Image taken using a difference interference contrast microscope at the Max Planck Institute for Biochemistry, Martinsried.

of dissociated neurons *in vitro*, Feinerman et al. succeeded in building basic logic devices by making use of patterned substrates to achieve a guided network topology [24]. Thus, is the key to understanding the brain in the neuronal circuitry, i. e. in the network structure of synaptic connections between neurons?

A tight inter-relation between structure and dynamics

There is some evidence that the network structure is at least not the only important key for neuronal function. The latter is represented by the dynamics of the neuronal system, but while the network structure underlies the dynamics, it does not determine it directly. It has been shown that even topologically static networks of neurons are able to display a wide range of activity patterns, with dynamics that are typically complex and highly irregular [25–28]. As a consequence of this richness in dynamics, even in networks of highly simplified model neurons, very different dynamical patterns can exist on the same structural network substrate [29,30]. The structure of a network can thus be said to restrict, but not determine dynamics.

In particular, it was recently shown that the same structural network can display a range of dynamical patterns which can be switched by way of brief perturbations, without altering the neuronal connectivity [31,32]. The function that a network can perform is therefore much richer than its structural foundation. We conclude that the synaptic connectivity of a network is not the full story — to learn about the way that neuronal function is embedded in neuronal networks, it is necessary to study structure and dynamics together, and to understand the complex inter-relation between the two, rather than each of them in isolation.

In this thesis, we aim to improve our knowledge about this relation between structure and dynamics *in vitro* by analyzing and modeling signs of non-randomness in neuronal activity patterns. We will combine analytical and numerical analysis to answer the following questions: To what degree do aspects of the structural topology determine the observed patterns of activity in neuronal networks? What can we learn about the structure of the neuronal network based on observing its activity?

We hope that the answers we give for simplified systems can in the future contribute to the knowledge about the development and evolution of synaptic structure in the behaving organism.

In vitro networks of cultured neurons

In order to make these questions tractable, in the present work we focus on networks of dissociated neuronal networks in vitro, which in recent years have become an important model system for the study of connectivity and spiking activity [33–43]. In these *cultured networks*, neurons are taken out of embryonic tissue by means of a mechanic and enzymatic protocol and are placed on a substrate, initially in isolation from each other. Within hours, however, neurons start to form axons and dendrites. In the absence of external guidance cues or external stimuli as present in the brain, the neurons connect and give rise to the formation of a self-organized network of neurons [44–46] (see Fig. 1.1 for an example of a fully developed cultured network).

Coinciding with the maturation of the synapses, spontaneous electrical activity (action potentials) starts to arise in the cultured cells and the whole network begins to be active [47,48]. As such, in vitro neuronal networks have the potential to better reveal the relationship between structure and dynamics that might be hidden in the intact brain by layers of additional complexity which is absent in vitro. Thus, the activity which is observed in vitro can be seen as a window into the more complex dynamics in vivo.

In particular, while the resulting wiring diagram in vitro is artificial due to the different chemical environment during development and is thus different from the intact tissue [49], some features of the connectivity might be preserved. In this case, it would be highly interesting to understand the cause of this similarity. Are there intrinsic mechanisms in the neurons which guide the connections? These are questions to which the answer is currently unknown, but which we are beginning to understand more by way of the techniques described in this thesis.

More generally, in the absence of stimuli, how much does the network structure determine the network dynamics? And on the other hand, can we study non-random activity patterns in vitro and learn about the self-organized structure of the network?

Bursting dynamics in vitro

Probably the most prominent feature of neuronal activity in vitro are *network bursts*: The firing activity in the population of neurons is increasingly synchronized with higher age of the culture, until at about 4–6 days in vitro, when the network finally displays spiking activity that recruits a large fraction of the neurons in the network for up to 200ms of high-frequency, synchronous firing (see Fig. 1.2A) [34,52]. So-called *full bursts* ignite the whole network, with individual neurons often spiking multiple times.

Bursts have been found in a range of neuronal systems in vitro, mostly studied in cultures of dissociated hippocampal neurons [38,42,53–55] and dissociated cortical neurons [37,43,46,52,56,57], but have also been found in slices of hippocampal tissue [58,59] and leech ganglia [60].

The onset of a burst shows a shape which can be approximated very well by an exponential rise (see Fig. 1.2B), indicating that just before a burst, the neurons in the network are highly excitable and activity spreads in an avalanche-like manner through the population. This activity quickly reaches saturating levels and then decays back to baseline firing rates when fatigue effects — potentially combined with the delayed onset of firing of the inhibitory sub-population — set in.

In bursts, despite the homogeneous limit of global activation, the onset is actually determined by diverse and precise patterns [37,61]. Eytan and Marom found that in cultured cortical networks the temporal structure of activity preceding a network burst was surprisingly hierarchical rather than random. A small subset of neurons, called the *leader neurons*, were firing consistently tens of milliseconds before others [51]. Early firing is therefore not based purely on chance — there is order in the spike times at the burst onset. The activity of these neurons could indeed be used to predict an upcoming network burst event [50,51]. This early firing property of a given neuron was

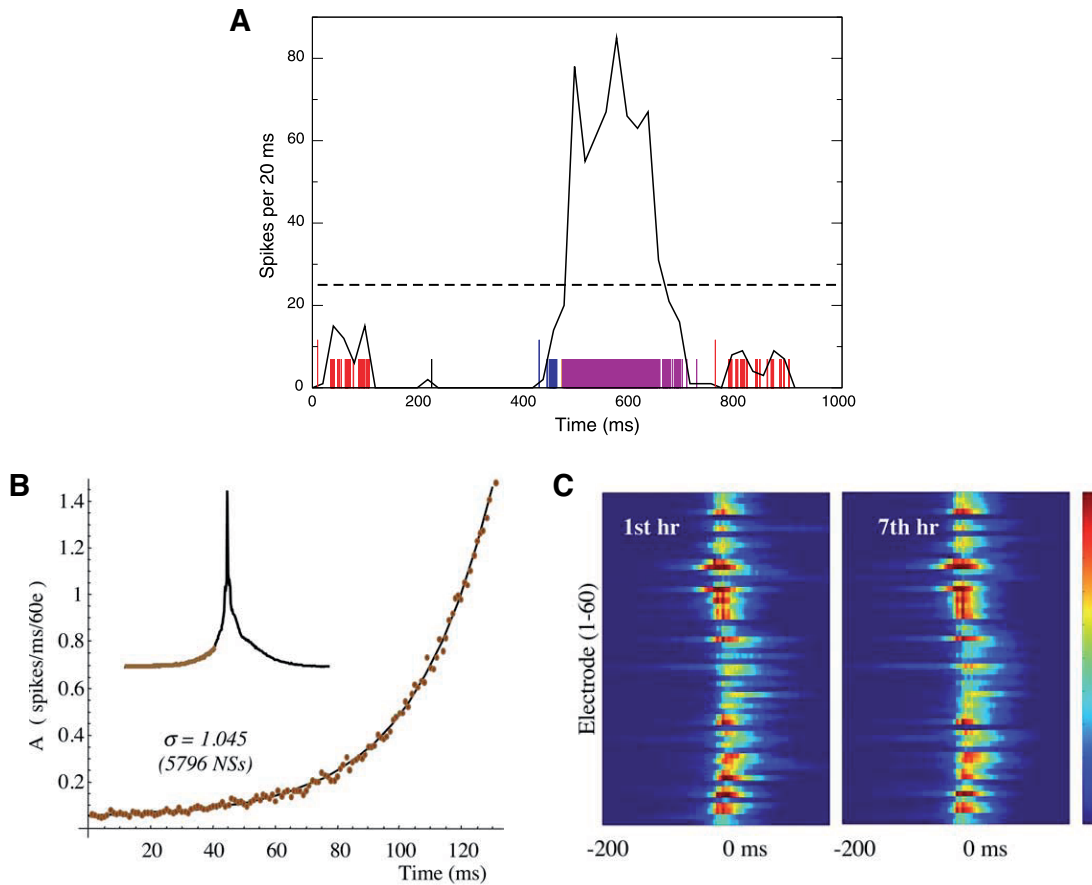


Figure 1.2: Previous studies of burst onset dynamics in vitro. (A) The temporal development of population firing rate in terms of number of spikes across all recording electrodes within a window of 20ms (black line). Exact spike times are indicated by vertical lines on the bottom of the panel. Aborted bursts (red spikes) remain under the threshold (dashed line) while full bursts recruit the whole network and are characterized by an onset phase (blue vertical lines) and a bursting phase (violet vertical lines). The first recorded spike (*leader*) of each population event is shown as a larger vertical line. Shown is a recording from a multi-electrode array in vitro. Panel reproduced with permission from [50]. (B) The average firing during burst onset of 5796 network spikes together with an exponential fit (black line). Inset shows the time-dependent averaged firing frequency of the whole burst. (C) Stability of firing rate heterogeneity is shown as the averaged probability of firing at the onset of a burst for bursts recorded during one hour for intervals separated by 6 hours. Panels B and C reproduced with permission from [51].

found to be stable across many hours [51,62], as shown in Fig. 1.2C. But what are the characteristics that cause a neuron to be active earlier than others?

Eckmann et al. hypothesize that there must exist a sub-network of leaders spread homogeneously across the culture, which are active first, and then serve to ignite the rest of the network [50]. Yet, the nature and origin of this sub-network was unclear. For instance, is the non-randomness in dynamics necessarily reflected in non-randomness in the network structure? It might be that the non-stationary condition inherent in the burst onset serves to amplify finite-size variations in the network, without systematic network structure.

There is some indication a priori that, if such a subset of neurons exists, the topological structure in the network is an important factor in determining the firing dynamics in this regard. For instance, Soriano et al. showed in a study on externally stimulated cultures that the probability of activation of a neuron was determined by the number of synaptic inputs [63]. In detailed simulations of neuronal systems, a positive correlation between leadership and the number of inputs was found empirically by training a linear classifier on the parameters of the model neurons [64]. A number of studies have examined the ability of an element of a network to synchronize its activity with the global population, all indicating that the network topology has an influence [65–71]. Furthermore, Schwab et al. showed in a theoretical study that the existence of communities of neurons of high in-degree determines whether the network is able to display rhythmic bursting activity [72]. However, is this equally true at the onset of spontaneous, aperiodic burst events in neuronal networks? Indeed, whether or not it is possible to predict the dynamics of neuronal firing based on the network, i. e. the degree to which the network topology constrains the onset dynamics, is unclear.

In Part I of this thesis, we address the question of what the determining factors in mathematical models of neuronal networks are, that cause a neuron to be among the subset of early firing cells. Studying simple mathematical neuron models, we obtain closed expressions for the expected time of first firing. To first order in terms of the network topology surrounding a given neuron, the first firing time is determined by a logarithmic dependence on the number of input connections a neuron has, confirming indeed a strong influence of the network topology. We also obtain higher order approximations of the surrounding topology of a given neuron in analogy to a Bethe-Peierls approach [73,74], which allows for an accurate prediction of the expected first firing time of each neuron. While the first firing time for a neuron in the context of a particular burst is of course a stochastic process, the expected first firing time is found to be a deterministic function of the structural topology.

This structure, however, is in the case of experimentally recorded networks unknown a priori. As the structure constrains the firing dynamics of the network, at least to some degree, an intriguing question is whether the dynamics constrains in any way the set of possible underlying network structures.

Identification of the structural topology

At present, reliable statistical measures to reconstruct the connectivity in large-scale networks are unavailable. This is despite the considerable efforts made in neuroanatomy, electrophysiology and imaging [75–81] since 1986, when White et al. presented the first major systematic result by extracting an exhaustive map of the neuronal connectivity in *Caenorhabditis elegans* based on serial electron microscopy [82]. However, such a reconstruction has remained elusive for any other organism to the present day [83].

It is known that the same structural topology can give rise to different effective connectivities represented by patterns of activity [32], and the challenge is to find, if possible, based on observations of the dynamics of the network, a statistical description that identifies the unique structural topology as closely as possible. While it is clear that network structure can constrain the possible

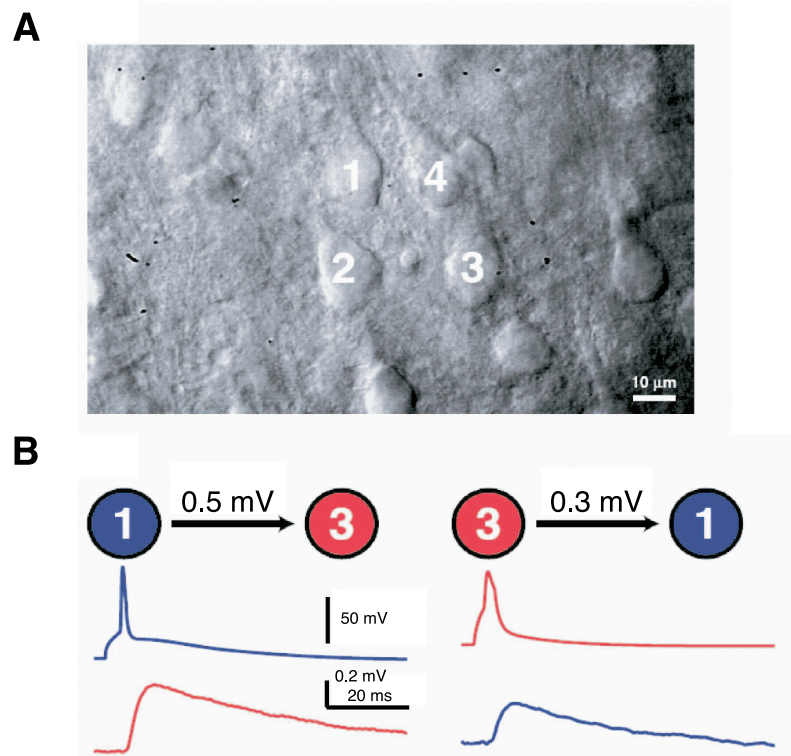


Figure 1.3: Patch-clamp recordings by Song et al. to extract neuronal connectivity. (A) Four example pyramidal cells in cultured tissue from layer V of rat visual cortex. (B) The left panel shows a current-clamp recording of the activity in cell 3 (shown in red) just after a spike was triggered in cell 1 (shown in blue). The sharp peak in cell 1 corresponds to the action potential, after which the membrane voltage decays back to the resting level, and the more broad, less pronounced peak in cell 3 below corresponds to the post-synaptic potential due synaptic currents from cell 1. The right panel shows the reversed direction, showing the activity in cell 1 after a spike occurred in cell 3. Both panels reproduced with permission from [84].

patterns of network activity, to identify the effect of this restriction in a non-trivially small number of signals (neurons) is far from obvious due to the interactions that make up the network, enabling dynamics that are much richer than a static network might at first glance be able to sustain. In a sense, this is the inverse of the question stated above: We now ask whether it is possible to predict the network topology based on observing non-randomness in the activity, given the fact that the neuronal dynamics is non-linear and non-stationary.

A related question is whether it is at all possible, based on observing the network dynamics, to distinguish connectivity motifs which lead to very similar dynamical motifs. Consider the case of *indirect causation*, where in one scenario, neurons A and B are connected by a direct link, $A \rightarrow B$, and in another scenario the connection is indirect via a third neuron C, as in $A \rightarrow C \rightarrow B$. Another important case is *co-causation*, where one aims to distinguish the cases $A \rightarrow B$ and $A \leftarrow C \rightarrow B$. A reconstruction algorithm must be able to distinguish the two cases reliably, even in the common case of calcium imaging recordings where the temporal resolution of the recording is such that it cannot resolve the precise temporal order of activity.

To this day, the single most reliable technique to extract synaptic connectivity is based on multiple simultaneous patch-clamp recordings, developed in the 1980s by Sakmann and Neher [85, 86]. In this approach, fluctuations in the membrane voltage of multiple neurons are recorded simultaneously, and the response of neurons to a current-induced spike in one cell is used to determine the existence and strength of mono-synaptic connectivity between them. Fig. 1.3A shows four example cells in a cultured slice of rat visual cortex. They are simultaneously recorded and the height of the post-synaptic potential is taken as evidence for the strength of the synaptic connection (see Fig. 1.3B), if the time delay between pre- and post-synaptic events is short enough such that mono-synaptic transmission is indeed ensured. Even though impressive numbers of pairs of neurons have been studied across cultures [49, 84] the method is nevertheless extremely difficult to scale to a larger number of neurons in the same culture.

On the other end of the length scale, a wealth of studies endeavors to reconstruct connectivity between regions of the brain, the majority of which is based on electroencephalography (EEG) or fMRI signals [87]. Most of them are based on cross-correlation algorithms to find connected voxels or using generalizations of this approach [88–92]. The macro-scale reconstruction results are proposed to aid in the diagnosis of Alzheimer’s disease [93], chronic pain [94, 95], multiple sclerosis [96] and epilepsy [97]. Functional MRI has also been used to study changes in large scale brain connectivity in patients with severe physical brain injuries [98, 99].

Significant progress has also been made to reconstruct the connectivity based on three dimensional images of tissue obtained by way of electron-microscopy, particularly in the tracing procedure employed to label axons and dendrites [77, 79, 100, 101]. While determining if neuronal branches are near is increasingly accurate, it is still difficult to determine the presence and weight of a putative synaptic connection. Significant work was also done by explicitly assuming a specific mathematical model of the underlying dynamics of the system [102, 103]. One popular model is the class of Generalized Linear Models [104, 105] which have indeed been combined with Bayesian inference to reconstruct first spike times from calcium fluorescence time series and then to infer the connectivity [106]. While achieving very good performance for simulated data sets, the performance for experimental recordings is difficult to estimate due to the constraints the model imposes, which may or may not be accurate for some or all of the cells under observation.

To obtain an actual circuit diagram with single-cell resolution in larger neuronal assemblies of hundreds or thousands of neurons it is generally regarded that statistical tools have to be employed on the basis of recordings of the activity of a population of neurons [107–110]. In this context, it is necessary to define what constitutes a *causal influence* of one neuron on another, a question with no single answer [111]. In the 1950s and 60s Granger and Wiener developed a measure of causality based on predictability: a time series A was said to be causally influencing another time series B when the fact of observing A was helpful in predicting the future states of B [112, 113]. This concept of causality is still used today, and it coincided with Shannon’s development of information theory [114]. Even though Shannon himself was criticizing the trend [115], his ideas were soon used in a broad range of different applications, including the forecasting of time series. Importantly for the present work, in 2000 Schreiber developed a model-free generalization of Granger’s work called *Transfer Entropy* [116].

As described above, a reliable description of the network topology of neuronal networks in vitro is necessary. To this end, in Part II of this thesis we aim at reconstructing, based on the observed dynamics of a neuronal network, the structural connectivity of the network. As a first step, we analyze how much information can be retrieved about the network topology based on the fact of it displaying bursting spiking dynamics, given the significant influence of the topology on the burst onset dynamics described above. It turns out, however, that very different networks can display bursting dynamics of the same inter-burst interval distributions, and thus a classification of

the network structure based on bursting alone is impossible.

To identify mono-synaptic interactions based on the complete observed dynamics, we propose a measure of causality based on Transfer Entropy. Crucially for reliable reconstruction, this formalism is modified to capture the interactions between neurons that are faster than what the recording is able to resolve.

Analyzing the effective connectivity in the different dynamical states displayed by the network, which in the present application are identified indirectly via the population averaged signal, we find dramatically different connectivities for each state. The degree of overlap between the effective connectivity and the structural connectivity is also different. Crucially for our question, we find that there exists a regime where the mapping between dynamics and structural topology is more direct than for other dynamical states, namely the regime of firing rates lower than in the case of synchronous bursting, and higher than the resting state. Conditioning out analysis to this regime results in a very good reconstruction performance. We call this novel measure of causal interaction, incorporating instantaneous interactions as well as state-dependency, *Generalized Transfer Entropy* (GTE).

Its performance is systematically studied using simulated signals, revealing a reliable reconstruction quality in terms of the rate of correctly identified links. In the special case of simple networks with a highly local connectivity, even linear causality measures can yield a reasonable reconstruction quality, but only GTE is performing well across network topologies. Moreover, we show that in contrast to other methods, GTE is unbiased in the sense that reconstructed topological quantities show a good linear correlation with the ground truth values. Thus, we have found a reliable way of identifying the structural topology from the observed dynamics.

Closing the loop

An interesting question is whether, making use of the inferred topology, it is possible to predict the onset time of neuronal activation preceding a network burst. Indeed, in Part III of this work we show that it is possible to predict the burst onset dynamics this way. A very precise prediction is achieved for simulated networks, and we demonstrate that early firing is indeed at least in part a result of the network topology also in biological networks.

The analysis verifies that the work presented in this thesis is self-consistent. Additionally, we show that even though the reconstruction of connectivity and the detection of burst onset times for individual neurons are subject to noise, the errors in prediction are low enough to allow for the presented analysis.

In short, in this thesis we study the tight inter-relation between structure and dynamics in networks of neurons *in vitro*. Despite the highly complex nature of this relation, we are going to show that much can be learned from the structure about the dynamics and vice versa, if the complex nature of the interaction is taken into account. We hope that the present work will, aside from the practical results described above and in the following chapters, lead to a deeper understanding of the co-evolution of structure and dynamics in neuronal networks.

Part I

Burst onset spiking dynamics

Chapter 2

Discrete time dynamics

As a first step towards the analysis of the temporal order of activation preceding a network burst in cultured networks, we analyze activity patterns in discrete-time mathematical models of neurons. While the correspondence of the parameters of these models is less direct than in the continuous-time model which will be considered later (see Chapter 4) they provide a simplified mathematical framework, making the study of firing times more tractable.

Discrete-time integrate-and-fire models have repeatedly been used to understand the dynamics of burst and avalanche dynamics and generation [117–122]. Of particular interest was the characteristics of the model that would lead to a distribution of avalanche sizes that follows a power-law distribution which had been found in certain biological systems, most notably in neuronal cultures of neocortical neurons [123] or cultured slices of coronal tissue [61]. More recently, such avalanches have also been found in vivo in anesthetized cats [124].

While the model system we are studying is taken from the same literature [119], the question we are asking is different. Rather than studying avalanche size distributions, we are here interested in the onset dynamics of such avalanches or burst, to address the question of which neurons in the network model would fire first during a time period immediately preceding a burst. The mathematical simplicity of the discrete-time integrate-and-fire model will allow us to determine analytically what the common characteristics of early firing neurons are. In particular, we will determine whether the dynamics of the system in this regard can be described in terms of the network structure.

2.1 Model characteristics

In discrete-time integrate-and-fire models, each node i is characterized by a variable $\vartheta_i(t)$, corresponding to the membrane potential of real neurons. The value of this potential is initially drawn from a uniform distribution on the interval $[0, 1[$.

The N nodes are connected on a directed graph, characterized by a number of C connections between them. In this thesis, we always consider the more biologically realistic case where the network is sparse rather than fully connected as in [120], i. e. where $C \ll N(N - 1)$ [13]. This graph is represented formally by a binary adjacency matrix A , with element $A_{ij} = 1$ if there exists a link $i \rightarrow j$ and $A_{ij} = 0$ otherwise. The in-degree k_i of each node is then given by $k_i = \sum_j A_{ji}$, and the out-degree correspondingly by $k_i^{out} = \sum_j A_{ij}$. We denote the average in-degree across the population by $\mu = 1/N \sum_i k_i$.

Initially, in each time step a node is chosen at random, and it receives an external input current of homogeneous amplitude, such that $I_i^{ext}(t) \in \{0, I_0^{ext}\}$. In the present work, we used $I_0^{ext} = 0.025$

as in [120].

$$\vartheta_i^{init}(t+1) = \vartheta_i(t) + I_i^{ext}(t) \quad (2.1)$$

This is done until there is at least one node in the network with $\vartheta_i^{init}(t) > 1$, corresponding to the first spike or action potential in the network. Each preceding time step is computed the following way:¹

$$\vartheta_i(t+1) = \vartheta_i(t) - \Theta(\vartheta_i(t) - 1) + \alpha \sum_{j=1}^N A_{ij} \zeta_j(t-1) \quad (2.2)$$

The Heaviside function $\Theta(\vartheta_i(t) - 1)$ denotes the mechanism of spike generation and the reset of membrane potential.² If such a threshold crossing occurs at time t in node i , it releases a current to its post-synaptic neighbors, and the binary indicator function $\zeta_i(t)$ is set to 1. Otherwise, $\zeta_i(t) = 0$. The amplitude that is added to the membrane potential of the post-synaptic neighbors is the connection strength, α . As the onset of a burst is governed by excitatory activity, we will only consider connections of a homogeneous, positive connection strength $\alpha = 0.1$. Note that the dynamical system described here is unit-less in the sense that the sub-threshold voltage of the membrane potential is bound in the interval $[0, 1[$, and therefore the connection strength α is also unit-less.

The simulation ends when there is a time step without any active node, as the external stimulation is considered to be slow and weak in comparison with the internal synaptic dynamics [119]. This results in a clear temporal separation of simulated events into a noise-driven regime, according to $\vartheta_i^{init}(t)$, and a purely deterministic regime of higher neuronal activity described by $\vartheta_i(t)$.

2.2 Two types of outcomes

The general dynamical nature of the resulting activity depends crucially on what would be the *branching ratio* in a branching process, i. e. the expected number of post-synaptic neurons that become active upon activation of one pre-synaptic neuron [125]. It is known that in a purely random or fully connected graph, a branching ratio of exactly 1 results in a *critical* state of the network, where the activity of a single node can have the full range of effects from none to the excitation of the whole network [119, 120, 126]. Indeed, the regime of a branching ratio close to 1 has been shown to be optimal for information transmission [127].

A system is *super-critical* if the branching ratio is greater, and *sub-critical* if it is less than one, corresponding to the average tendency of the number of active neurons to grow or to shrink over time.

We will now develop the corresponding branching ratio in the present model. If the distribution of membrane voltage across the unit interval is uniform, then the probability that the resulting voltage is super-threshold, after adding one amount α to it, is given by α itself, as long as $\alpha \leq 1$. It follows that in such a case the branching ratio, which we will denote by γ , is given by:

$$\gamma = \alpha \frac{C}{N} = \alpha \mu. \quad (2.3)$$

¹Importantly in these models, $\vartheta_i(t)$ may not depend on the membrane potential of any other node at the same time step $\vartheta_j(t)$, $j \neq i$ because then the dynamics would depend of the order in which the nodes are evaluated.

²As the membrane potential ϑ is upon spiking reduced by 1 in the next time step. This means that ϑ is not constrained to the unit interval, if a node receives a synaptic input potential greater than 1, and will continue firing for some time even if the input current were to be switched off.

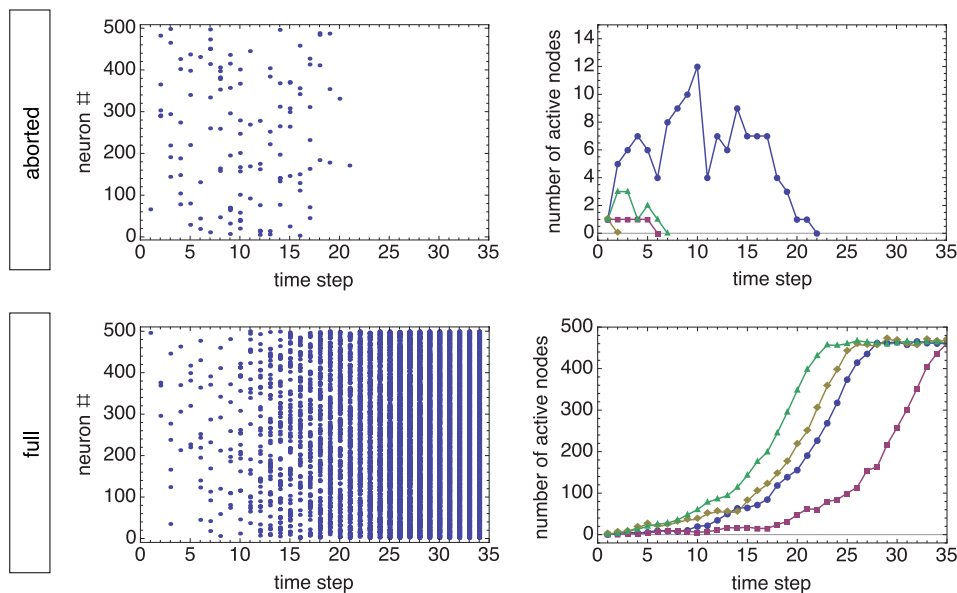


Figure 2.1: Full and aborted bursts in a discrete-time integrate-and-fire model. The top row shows examples for aborted bursts (see text), and the bottom row shows full bursts. For each case, the left column displays the activity of one example event as a raster plot, and the right column shows four example events in terms of the number of activated nodes per time step. The activity on the left hand side corresponds to the blue curve on the right hand side panel. Note that the vertical axis on the two plots on the right hand side is different. Note also that the activity shown in the top left panel is an exceptionally long aborted burst — over 90% of aborted bursts last shorter than 10 time steps in this simulation. Shown are simulations based on a random graph with $N = 500$ nodes and a branching ratio of $\gamma = 1.2$.

To get the experimentally observed behavior [50] of the co-existence of both full bursts and aborted bursts (activity that excites a portion of the network, but that does not propagate through the whole network) we therefore constrain the model to the parameter regime of $\gamma > 1$.

Note that we are here only modeling the onset behavior of a burst. We therefore identify the simulated firing activity which settles in a steady, high firing rate regime with the experimentally observed full burst, and simulated activity that falls back to zero firing rate with the aborted bursts. In a real network, factors not included in our model, such as synaptic depression due to depletion of neurotransmitter vesicles upon repeated firing would bring the activity back to the baseline firing rate [128, 129].

Examples of realizations of such activity is shown in Fig. 2.1, showing the two distinct types of activity that the models displays. On the one hand, the population firing rate as a function of time (which by construction always starts at one node, see above) either decays to zero in the immediately preceding time step (indicative of the failure of the first node to elicit activity in any other) or grows shortly but then decays back to zero a small number of time steps afterwards.

Indeed, we find that for 3000 simulations, the activity always decayed back to zero if at no point in time there were more than 10% of the network activated (not shown). On the other hand, activity that had reached the high firing rate steady state never decayed back to zero for the 500 time steps that were simulated (not shown), indicating a clear separation between the two activity patterns.

The reason for this separation of dynamical behaviors is found in the probabilistic nature of activation. Again assuming a uniform distribution of membrane potential and denoting the probability of connection $p = C/(N(N-1))$, then the average number of nodes $M(t+1)$ that will be activated by one pre-synaptic node ($M(t) = 1$) is given by a binomial expression:

$$E(M(t+1)) = \sum_{k=0}^{N-1} k P(M(t+1) = k) = \sum_{k=0}^{N-1} k \binom{N-1}{k} (p\alpha)^k (1-p\alpha)^{N-k-1} \quad (2.4)$$

One consequence of this is that there is a sizable probability that the burst will abort in just the second time step:

$$P(M(t+1) = 0) = \binom{N-1}{0} (1-p\alpha)^{N-1} = (1-p\alpha)^{N-1}. \quad (2.5)$$

For the network shown in Fig. 2.1, $\gamma = 1.2$ and $p = 0.028$ and therefore $P(M(t+1) = 0) = 0.30$. The origin of the separation of activity in the two regimes (*full* and *aborted*) is now clear: For low numbers of activated nodes $M \ll N$, the activity is dominated by finite size effects, resulting in a significant probability to be absorbed by the boundary at $M = 0$. In the limit $M \rightarrow N$, the activity is dominated by the supercritical branching ratio $\gamma > 1$ and its probability to abort will be zero.

2.3 Different topologies

So far we have been considering Erdős-Renyí random graphs, i. e. graphs that are completely determined by the number of nodes N and the probability of connection $p = C/(N(N-1))$. Additionally, we now consider networks that have a topologically non-random structure by making the probability of connection depend on the spatial distance between two neurons. Such a dependence on distance has been observed in slices of cortical tissue [84].

To generate networks with distance-dependent topology, we first assign each node a two dimensional position on a square of $d = 1\text{mm}$ length. The assignment is random with the only restriction being a minimum distance of $d_{min} = 10\mu\text{m}$ to previously placed neurons, corresponding to the typical diameter of the cell bodies of pyramidal neurons [33, 130, 131]. The connections are then drawn from a connection probability shaped by a Gaussian kernel of the following form, given a characteristic length scale of the system $\lambda = d/\sqrt{N}$ and the distance r between two neurons,

$$p(r) = \exp\left(-\frac{r}{\lambda}\right)^2. \quad (2.6)$$

To generate the graph, we randomly select two nodes i and j with $A_{ij} = 0$, draw a random number w from a uniform distribution on the unit interval and generate the link $i \rightarrow j$ if $w < p(r)$. This process is repeated until the target number of connections C is reached.

The effect of introducing, by way of the process described above, non-random topological structure into the network dynamics given by Eq. 2.2 is shown in Fig. 2.2. The top panel shows how the introduction of the distance dependent connectivity leads to nodes that have a higher in-degree if they are clustered spatially — in the random connectivity, there is by construction no correlation between position and in-degree. A side effect of the local connectivity kernel is the introduction of degree-degree correlations: The in-degree of node connected to a high in-degree node tends to be high as well. With the parameters of d , λ and C presented above, we reached typically values of degree-degree correlation of $C_1 \approx 0.4$.

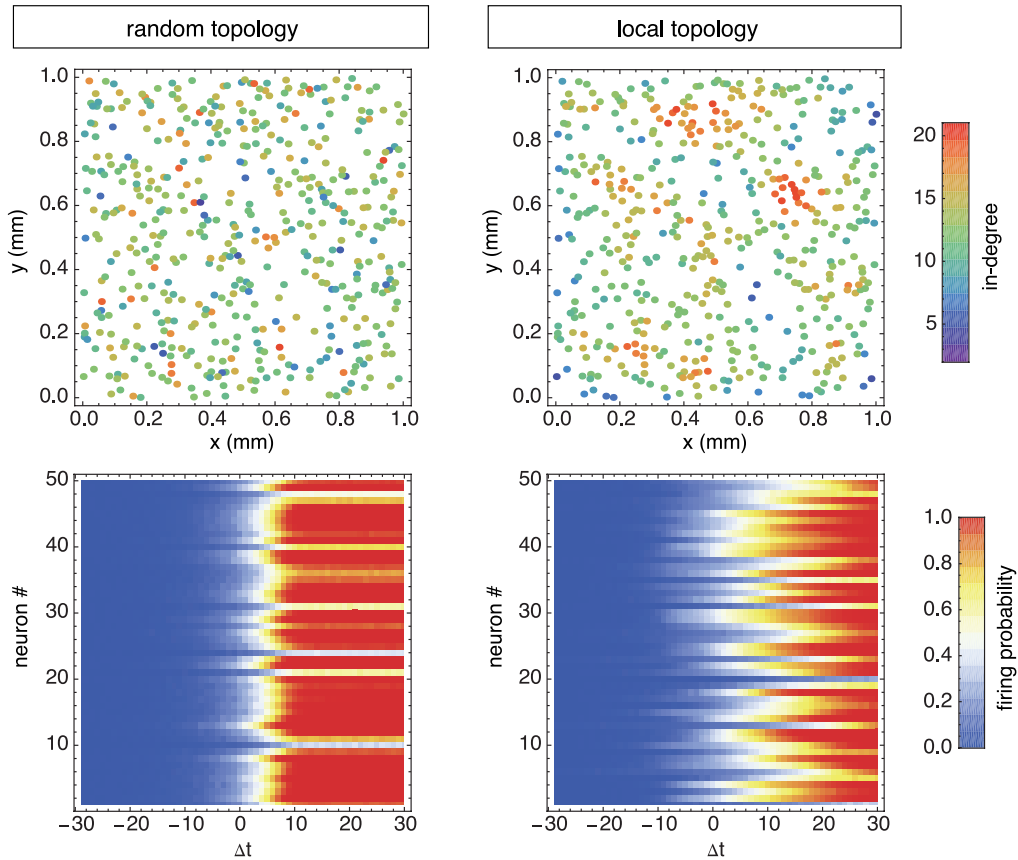


Figure 2.2: The effect of topological structure on onset firing probability. Top panel shows simulated neurons on planar substrate, left on a random graph, right with local connection probability. Color coded is their in-degree k as generated by their respective connection probabilities (see text). The bottom panel shows the firing probability of 50 example neurons at time Δt relative to the onset of a full burst. Shown are simulations with $\gamma = 1.2$.

Local connectivity kernels also bring about *clustering* in the adjacency matrix. Generalizing the clustering coefficient to directed networks, the full clustering index was introduced in [132] as:

$$C_i^D = \frac{(A + A^T)_{ii}^3}{2(d_i^{tot}(d_i^{tot} - 1) - 2d_i^{bidir})} \quad (2.7)$$

Here we have used the total degree $d_i^{tot} = d_i^{in} + d_i^{out}$ and the number of bidirectional connections of a node $d_i^{bidir} = \sum_i (A^2)_{ii}$. In the following, we will denote by CC the population average of C_i^D .

To study the differences in neuronal dynamics before the onset of a burst, we have computed the average firing frequency by averaging over simulated full bursts, defining $\Delta t = 0$ as the point where for the first time in a given simulation instance, more than 20% of the neurons in the network were active. We have plotted the resulting firing probabilities in the bottom panel of Fig. 2.2. Due to the steep rise in the average firing frequency across the network during the burst onset, all of the neurons show a monotonically rising firing probability as time increases. The firing frequency for the majority of nodes quickly approaches one at about $\Delta t = 8$. There are, however, a number of nodes which do not reach the maximum firing frequency of 1 spike per time step and by construction this set must contain all those nodes with an in-degree k_i of less than $1/\alpha$. For the random graph shown in the left panel of Fig. 2.2, the number of these nodes is given by:

$$N_{f < 1} = N \sum_{k=0}^{\lceil 1/\alpha - 1 \rceil} \binom{N-1}{k} p^k (1-p)^{N-k-1} \approx 0.239 N. \quad (2.8)$$

Interestingly, regardless of the fact that the peak firing probability is different for some nodes, in the case of the random graph the transition between zero and maximum firing probability occurs between $\Delta t = 0$ and $\Delta t = 10$ for all neurons (see Fig. 2.2, bottom left panel). This situation changes dramatically when analyzing in the same way the dynamical behavior of the locally clustered topology (see Fig. 2.2, bottom right panel). Now there is a shift in time from node to node which can be of the order of 20 time steps. As the width in real time of one discrete time step can be identified approximately with the duration of synaptic transmission (2–4 ms [128, 133]), this matches the experimental observation of Eytan and Marom that early firing neurons are “consistently recruited tens of milliseconds before others” [51].

2.4 The expected value of the first firing time

We now aim to quantify the above observations in terms of the expected value of the first firing time of each neuron in the simulated networks. First we calculate from the preceding analysis, resulting in the time-dependent firing probability $P_i^f(t)$, the probability of each neuron i to fire for the first time relative to the burst onset, $P_i^{f1}(t)$.

$$P_i^{f1}(t) = P_i^f(t) \prod_{t'=-\infty}^{t-1} P_i^f(t') \quad (2.9)$$

The result is plotted, again on the left hand side for the random graph and on the right hand side for the local graph, in the top panel of Fig. 2.3. The plots are generated from the same simulations as the ones used for Fig. 2.2.

The expected first firing time \hat{t}_i — the property on which we will focus our attention for the rest of this chapter — can of course be directly calculated from Eq. 2.9 as the first moment of $P_i^{f1}(t)$. In the bottom panel of Fig. 2.3, \hat{t}_i is plotted against in-degree for all the nodes in the simulations. Neurons with a higher degree have in general a lower first firing time, indicative of the fact that they tend to experience higher than average input currents.

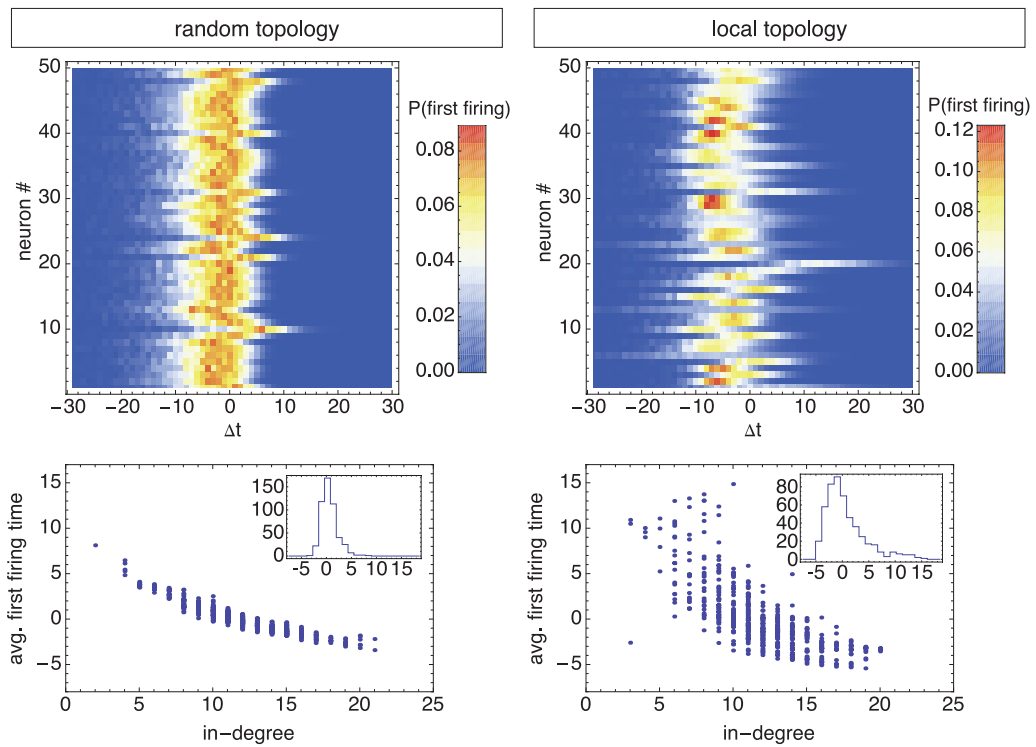


Figure 2.3: Topological structure enhances variability in onset first firing probability. Top panel shows the probability of 50 example neurons to fire for the first time during the onset of a burst extracted from simulations. The bottom panel shows, for each of the $N = 500$ neurons in the simulated networks, the expected value of the first firing (time relative to the mean of the population) against the in-degree of the respective node. The inset shows a histogram of expected first firing time across the population, with the number of occurrence on the vertical axis. Shown are simulations with $\gamma = 1.2$, identical to those of Fig. 2.2.

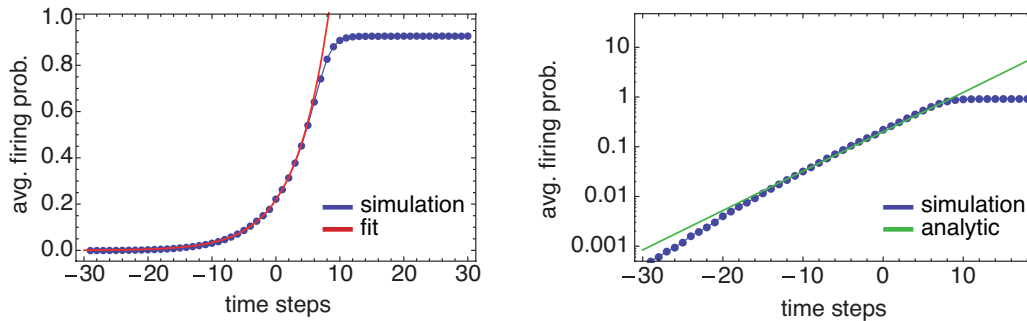


Figure 2.4: Rise before burst is exponential and can be predicted analytically. Shown are population-averaged firing probabilities of bursts simulated on a random graph with $\gamma = 1.2$, compared to an exponential fit (left panel, red curve) and the analytical prediction developed in the main text (right panel, green curve).

In a random graph, the spread in expected first firing times is moderate (standard deviation of 1.59), as shown in the histogram in the inset of the bottom left panel of Fig. 2.3. The width of this distribution is significantly increased compared with a standard deviation of 3.81 as indicated by the bottom right panel of Fig. 2.3.

2.4.1 Prediction of the first firing time

To better understand the characteristics that cause neurons to fire earlier than others, we now develop an analytical expression for the expected value of their first firing time. From Eq. 2.2 we know that the time that a neuron fires for the first time \hat{t} is completely determined by its synaptic input, the threshold (here set to 1) and the initial value of the membrane voltage $\vartheta_i(t_0)$.

$$\vartheta_i(t_0) + \sum_{t=t_0}^{\hat{t}} I(t) = 1 \quad (2.10)$$

We now approximate the discrete-time input current $I(t)$ by a continuous mean-field firing frequency $F(t)$, i.e.

$$\vartheta_i(t_0) + \alpha k_i \int_{t_0}^{\hat{t}} F(t) dt = 1. \quad (2.11)$$

From the experimental evidence we know that the initial rise in the population average of the firing frequency is exponential (see Fig. 1.2B). Indeed we find that we can such a curve even to the numerical simulations discussed above. An example of this is shown in the left panel of Fig. 2.4. It is clear that the onset of a burst in our simulations can be described accurately by an exponential.

We can actually predict the form of this exponential rise from the model parameters. To estimate the parameters a and b in $F(t) = a b^t$, we need only to assume the rise to be comparable to a branching process and to remember how Δt was defined earlier. The former means ignoring the existence of loops and bidirectional connections, which is a good approximation as long as the average firing frequency is low, such that the resulting effect of loops on the activity pattern is negligible. Note that the approximation will also be more precise in random graphs, where the probability of loops and bidirectional connections is lower overall. Given this assumption, $b = \gamma$ by definition (see above and Eq. 2.3). The second parameter a is then simply given by the way we

defined the burst onset: $\Delta t = 0$ is the point in time where $P_0 = 20\%$ of nodes in the network were active.

$$F(t) = P_0 \gamma^t \quad (2.12)$$

The result for an example network is plotted in the right panel of Fig. 2.4. Despite the somewhat crude approximations made, there is a very good agreement between simulation and prediction over an interval of about 20 time steps. This covers with good accuracy the relevant region between small firing rates on the left hand side and high firing rates after the first firing time on the right side of this interval. We note however that the firing rate on the left side is lower than expected. This is due to the bias of only examining activity that developed into a full burst. Therefore, the average activity in the initial region, where the activity is decisive for determining whether it will develop into a full burst or abort (due to finite size effects, see discussion above) is overestimated and the slope higher than expected from the exponential rise.

We can now plug the above approximation of Eq. 2.12 for $F(t)$ into Eq. 2.11 (omitting the index i for readability):

$$\gamma^{\hat{t}} - \gamma^{t_0} = \frac{(1 - \vartheta_0) \log \gamma}{\alpha P_0 k} \quad (2.13)$$

This we can solve, substituting $\Gamma = 1/\log(\gamma)$:

$$\hat{t}(\vartheta_0, k) = \Gamma \log \left\{ \frac{1 - \vartheta_0}{\alpha P_0 \Gamma k} + \exp^{t_0/\Gamma} \right\}. \quad (2.14)$$

We can average out the dependence of \hat{t} on ϑ_0 knowing that the probability distribution of ϑ_0 is drawn from a uniform distribution on the unit interval at the start of each event.

$$\begin{aligned} \hat{t}(k) &= \int_0^1 d\vartheta_0 \hat{t}(\vartheta_0, k) \\ &= -\Gamma \{ D k \log(D k) - (1 + D k) \log(1 + D k) + 1 + \log(\alpha P_0 \Gamma k) \} \end{aligned} \quad (2.15)$$

Note that we have substituted $D = \alpha P_0 \Gamma \exp(t_0/\Gamma)$. The first two terms correspond to the start of the integration from t_0 and are typically small compared to the other terms, because D is small. Neglecting those terms, we end up with the following equation:

$$\hat{t}(k) = \frac{1}{\log \gamma} \left\{ \log \frac{\log \gamma}{\alpha P_0 k} - 1 \right\}. \quad (2.16)$$

This equation corresponds to the solution if we had taken the limit $t_0 \rightarrow -\infty$. The precise value for t_0 might be determined from the right panel of Fig. 2.4 where the average firing rate deviates from the exponential rise. We will not attempt to determine its value analytically here. As a choice of t_0 corresponds to an offset in time common to all neurons in the network (as long as $t_0 \ll \hat{t}(k)$ from Eq. 2.16 for any k in the network), we will fix this degree of freedom by setting the mean of the population to zero.

2.4.2 Higher order degree

So far we have only considered the simplest form of a mean-field approximation, namely a neuron receiving the input current directly from the mean field. In the actual network, the input a node receives comes of course from its neighbors (see Fig. 2.5 for an illustration). We therefore want to consider also the next neighbors in our calculation, analogously to a Bethe-Peierls approximation [73,74]. This can be done very simply by realizing that under static conditions, the




Figure 2.5: Illustration of first and second order mean-field approximation. On the left, only the neuron in question (shown in dark blue) is modeled, receiving input directly from the mean field, simply multiplied by its input degree. On the right, we include the nearest neighbors (shown in light blue) explicitly, who are now the ones receiving the mean field input. This is indicated by the dashed circle, with neurons within the circle are treated explicitly, while neurons outside the circle are treated as being part of the mean field.

output firing rate of the model neurons is just the inverse of the number of input spikes that a neuron needs to once cross the unit interval between zero and threshold:

$$f_i^{\text{out}} = \frac{\alpha k_i f_i^{\text{in}}}{1} \quad (2.17)$$

An obvious but instructive consequence of this is that the ratio between stationary output and input rate of a representative neuron with an average number of inputs $k_i = \mu$ is just given by γ .

The second-order equivalent to Eq. 2.11 can then be written using the *second order degree* $k_i^{(2)} = (A^2)_{ii}$. Note that this ignores all links that are not *feed forward*, i. e. that deviate from the directionality illustrated in Fig. 2.5, from the mean field to the node i in question.

$$\vartheta_i(t_0) + \alpha^2 k_i^{(2)} \int_{t_0}^{\hat{t}} F(t) dt = 1 \quad (2.18)$$

As the difference to Eq. 2.11 is constant with respect to the integration over both time and ϑ_0 , we can immediately write down the solution by comparing with Eq. 2.16:

$$\hat{t}^{(2)}(k) = \frac{1}{\log \gamma} \left\{ \log \frac{\log \gamma}{\alpha^2 P_0 k^{(2)}} - 1 \right\}. \quad (2.19)$$

2.4.3 Degree-degree correlations

We note that while the first order solution of Eq. 2.11 depends on the degree k_i only, Eq. 2.19 requires the knowledge of the second order degree $k_i^{(2)}$, and thus the whole adjacency matrix A . We can find an approximate solution between these extremes by making use of Pearson's correlation coefficient between the degree of connected nodes, which we will denote by the *degree correlation* C_1 . This coefficient C_1 is zero in random networks, but positive in networks such as the locally clustered networks discussed above.

Let k_A and k_B be the degree of two nodes. We assume that they are drawn from Gaussian distributions with means μ_A and μ_B and variances σ_A^2 and σ_B^2 . We will assume that they are correlated with a given Pearson's coefficient C_1 . We know then that their joint probability distribution is the following [134]:

$$P(k_A, k_B) = \frac{1}{2\pi\sigma_A\sigma_B\sqrt{1-C_1^2}} \exp\left\{-\frac{z}{2(1-C_1^2)}\right\}$$

$$\text{with } z = \frac{(k_A - \mu_A)^2}{\sigma_A^2} - \frac{2C_1(k_A - \mu_A)(k_B - \mu_B)}{\sigma_A\sigma_B} + \frac{(k_B - \mu_B)^2}{\sigma_B^2}. \quad (2.20)$$

Since k_A and k_B come from the same underlying process generating degrees, we will assume their mean and standard deviation to be the same, so that $\mu_A = \mu_B = \mu$ and $\sigma_A = \sigma_B = \sigma$. It is then straightforward to derive $P(k_B|k_A)$ since we know $P(k_A)$:

$$P(k_B|k_A) = \frac{1}{\sqrt{2\pi(1-C_1^2)}\sigma} \exp\left\{-\frac{(C_1(k_A - \mu) - (k_B - \mu))^2}{2\sigma^2(1-C_1^2)}\right\} \quad (2.21)$$

The maximum of this distribution is therefore at the point where the denominator inside of the exponential term is zero, and the degree that corresponds to this peak is given by:

$$k_B^{\max}(k_A, C_1) = C_1(k_A - \mu) + \mu \quad (2.22)$$

As expected, k_B^{\max} is equal to μ in random networks ($C_1 = 0$) and between k_A and μ for positive degree-degree correlated networks such as the locally structured networks discussed above.

Eq. 2.22 leads to the following approximation, simplifying the second-order one of Eq. 2.19: Neuron i with degree k_i receives input from its neighbors, which we assume to all have equal degree $k^{(\text{corr})} = C_1(k_i - \mu) + \mu$ as above. As a result, we have a second-order equation that requires much less knowledge of the whole adjacency matrix, but only of an average statistic of the topology of the network, namely the degree-degree correlation coefficient C_1 .

$$\hat{t}^{(\text{corr})}(k) = \frac{1}{\log \gamma} \left\{ \log \frac{\log \gamma}{\alpha^2 P_0 k^{(\text{corr})}} - 1 \right\} \quad (2.23)$$

Comparing this with the first-order solution of Eq. 2.16 we find that this yields an effective degree, k_i^{eff} , different from the structural degree k_i :

$$k_i^{\text{eff}} = \alpha k_i [C_1(k_i - \mu) + \mu] \quad (2.24)$$

Note that the effective degree has a quadratic dependence on k_i , corresponding e.g. for $C_1 > 0$ to the fact that high-degree nodes receive input from nodes whose degree is also above the mean on average.

2.5 Comparison of numerical and analytical results

The analytical predictions of the expected first firing time developed in the last section can be directly compared to the numerical results (such as the data of the bottom panel of Fig. 2.3).

2.5.1 Random graphs

In a random graph, the expected onset time can be calculated very precisely using the analytical formalism developed above (see Fig. 2.6). For both values of γ shown, the first order approximation according to Eq. 2.16 yields a good prediction of the numerically evaluated onset time \hat{t} across the range of onset times, with an uncertainty of approximately half a time step. This prediction performance is even increased when going to the second order approximation according to Eq. 2.19,

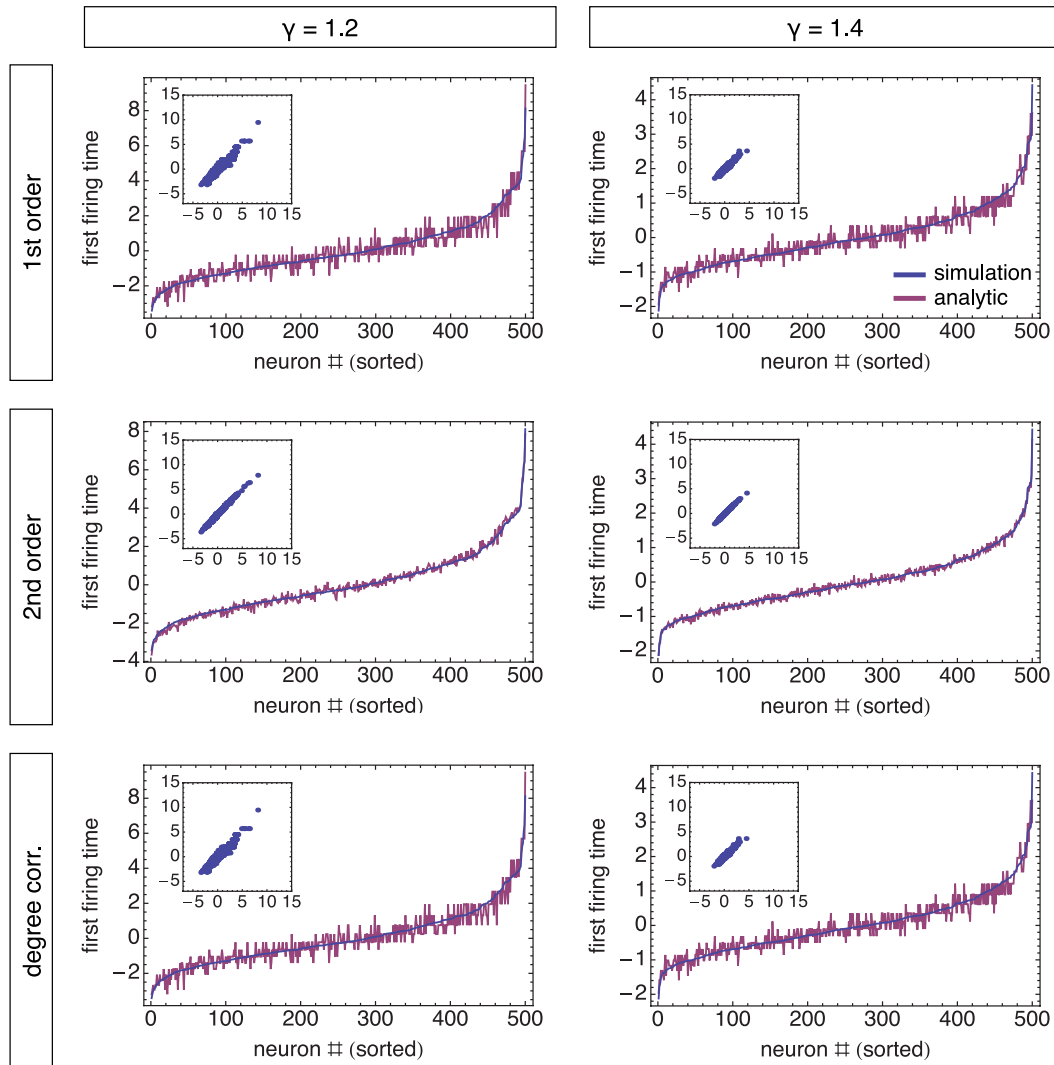


Figure 2.6: The expected first firing time in a random graph can be predicted analytically. Shown are, for the two cases $\gamma = 1.2$ (left column) and $\gamma = 1.4$ (right column) numerical results and analytical prediction with different methodologies. In each plot, the expected first firing time was evaluated from simulations (see text) and the neurons were sorted according to this time on the x-axis, yielding the blue curves. In each row, a different analytical prediction method is used and compared to the numerical times. The first order approximation (red curves, top panel), the second order approximation (middle panel) and the degree-correlation approximation (bottom panel). Insets are scatter plots of the same data points, showing numerical (horizontal axis) versus analytically predicted onset time (vertical axis) in units of time steps.

reducing the uncertainty to approximately 0.3 time steps (middle panel). Again, this accuracy is conserved across the population, including the outliers of onset time of very early and very late firing neurons. The bottom graph of Fig. 2.6, showing the prediction using the degree-correlation approximation of Eq. 2.23, is by construction identical to the top graph for the first order solution, because $k_i^{\text{eff}} = k_i$ for negligible degree-degree correlation (see Eq. 2.24).

In all cases, the Pearson's correlation coefficient of numerical against analytical value (see insets) is $p > 0.96$, in the case of the second order approximation in fact $p > 0.99$. This highlights the very good quality of the analytical predictions in the case of random graphs.

2.5.2 Locally clustered graphs

The situation is more complicated for locally clustered networks, as introduced in Chapter 2.3. The left column of Fig. 2.7 illustrates this, plotting the prediction performance analogously to the case of random graphs (see Fig. 2.6). The first order approximation (Fig. 2.7, top left panel) displays a reduced slope of the predicted onset time, indicative of a reduced predicted variance of the distribution of onset times across the network. Especially for the neurons with the latest firing times, the prediction is off by over 10 time steps. There is, however, still a good linear correlation between measured and calculated onset time, yielding a Pearson's correlation coefficient of $p = 0.76$. The same observations are true for the case of $\gamma = 1.4$ (Fig. 2.7, top right panel), even if the reduced overall spread of onset times due to the faster rise of the overall activity results in a slightly better correlation of $p = 0.83$.

The second order approximation (Fig. 2.7, middle row) yields a much better approximation for the majority of neurons: The predicted onset time follows the numerical result much better due to an increased slope for both cases of γ plotted. We also observe a correction at the end of late onset times, even if the amount of this correction is still insufficient.

The degree-correlated approximation of Eq. 2.23 yields an interesting middle ground between first and second order. Taking into account the degree-correlations in the network does yield the increase in slope observed for the second order approximation and therefore a large improvement in prediction quality for the large majority of neurons. It however does not show any improvements at the right tail of late onset times, indicating that these nodes are exceptional in that the degree correlation of these nodes is different (in fact, lower) than the population average used for estimating the coefficient C_1 .

2.5.3 Higher order correlations

The last observation leads to the question if we need to include higher order terms in our calculations, rather than just second order. It is easy to see that we can generalize Eq. 2.16 and Eq. 2.19 to a solution of arbitrary order n (with $n \in \mathbb{N} \setminus \{0\}$) in the following way:

$$\hat{t}^{(n)}(k) = \frac{1}{\log \gamma} \left\{ \log \frac{\log \gamma}{\alpha^n P_0 k^{(n)}} - 1 \right\} \quad \text{with } k^{(n)} = (A^n)_{ii}. \quad (2.25)$$

Clearly, taking into account the assumptions that led to Eq. 2.19, there is a trade-off: A low order will ignore higher-order correlations in the topology, while the higher-order solutions ignore loops and feed-backward links.

The top panel of Fig. 2.8 shows the result of Eq. 2.25 for $n = 5$ for the local networks of Fig. 2.7. Now, despite the fact that the prediction accuracy in the range of medium first firing times is virtually unchanged, we see a dramatic improvement in predicting the behavior of the very late firing time neurons. The correlation coefficient is significantly higher than in the case of $n = 2$, resulting in $p = 0.859$ for $\gamma = 1.2$ and $p = 0.926$ for the case of $\gamma = 1.4$.

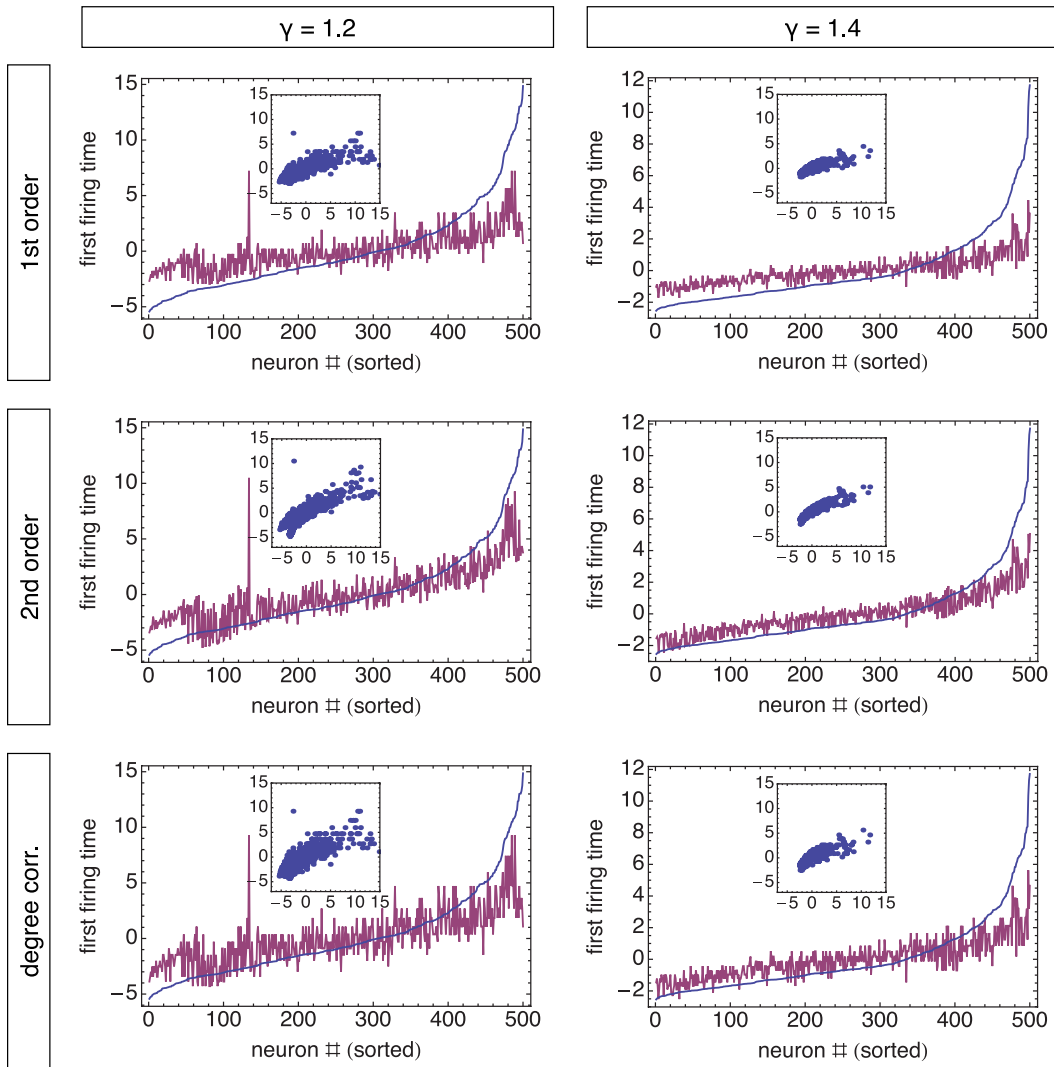


Figure 2.7: First firing time in a locally clustered graph can be approximated analytically. Layout of plots and prediction methods used are identical to the case of Fig. 2.6. The best accuracy is achieved in the regions of moderate onset time.

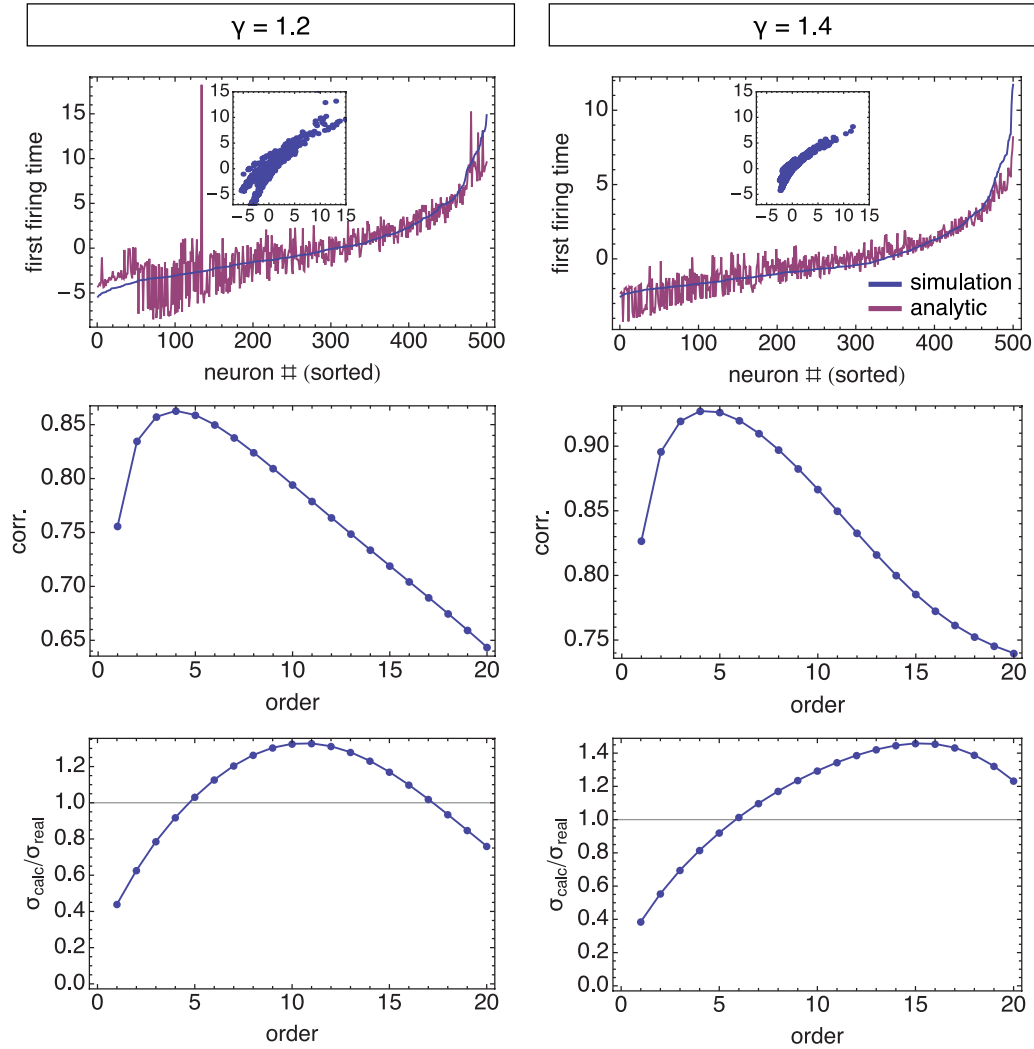


Figure 2.8: Dependence of the predicted first firing time on the order of topological expansion. The top panel shows comparisons of the numerically evaluated expected first firing time (data identical to Fig. 2.7) with the calculated value using the 5th order approximation. Insets are scatter plots of the same data points showing numerical (horizontal axis) against calculated (vertical axis) expected first firing time in units of time steps. The middle row shows, for both values of γ , the correlation between numerical and predicted first firing time against the order n of the calculation. The bottom row shows the precision of the calculated standard deviation of the distribution of first firing times in terms of the fraction $\sigma(\text{analytical})/\sigma(\text{numerical})$. A fraction of 1 indicates an exact prediction of the standard deviation, while e.g. a fraction greater than 1 indicates that the predicted first firing times have a broader distribution than the numerically evaluated ones. Shown are simulations on a locally clustered graph.

Note that going to higher order beyond $n = 3$ has virtually no effect on the random graphs of Fig. 2.6 beyond increasing the accuracy in prediction. For $n = 5$, the resulting correlation coefficient is $p = 0.998$ for $\gamma = 1.2$ and $p = 0.997$ for the case of $\gamma = 1.4$.

Going back to the local topology, we are now in a position to examine the trade-off described above. The correlation between numerically measured and analytically predicted first firing times has a peak around $n = 4$ (see Fig. 2.8, middle row). Beyond the correlation itself, the width of the resulting distribution of expected first firing time across the population has its maximum in the same region, crossing the optimal ratio of 1 between $n = 5$ and $n = 6$ (see Fig. 2.8, bottom row). There is clearly an optimal order around $n = 5$, which was also shown above, even if this optimal n is not known a priori. This indicates the high quality of analytical prediction of the expected onset time in discrete-time integrate-and-fire models, given that the structural neighborhood (given by the order n) is taken into account.

Chapter 3

Quorum percolation

There has recently been an interest in the neurosciences in *percolation* models, i. e. studying an initially unconnected graph which shows, as an increasing number of vertices are added to it, a phase transition in terms of the largest connected group of nodes. At this transition point, structures appear that span the entire system, where before there were only isolated clusters of nodes.

Percolation models have long been studied for understanding phase transitions in random media [135–137]. More recently and closer to the topic of this thesis, percolation theory has been applied to study the activity dynamics of neuronal cultures [40, 63, 138].

In our study [139] we have applied this knowledge to understand the burst onset dynamics in a *quorum percolation* model of neuronal activity. As will be discussed in detail below, the quorum percolation model relies on a fraction of nodes that is excited externally to start activity. Consequently, it is suited to the analysis of activity following global stimulation, as indeed it was developed in the context of in vitro cultures and their response to global bath stimulation [63, 138, 140]. It is interesting to analyze the onset dynamics under these circumstances, in comparison with the results of Chapter 2, given the finding that leaders were identical in spontaneous and evoked activity [51].

3.1 Model characteristics

Just like the integrate-and-fire dynamics in Chapter 2, quorum percolation (QP) is a discrete-time model on a directed graph. Each node on the graph is characterized by a binary state variable $s_i(t)$.

$$s_i(t+1) = s_i(t) + (1 - s_i(t)) \Theta \left\{ \sum_j A_{ij} s_j(t) - m \right\} \quad (3.1)$$

Consequently, a node is switched to 1 (the *on* state) when at least a *quorum* of m active input nodes is achieved. Afterwards, it remains in this state for all following time steps. There is no noise in the system, so the latter effect is the only way for a node to switch to the on state.

The graph employed, represented by the adjacency matrix A , is random apart from a prescribed, Gaussian degree distribution with mean μ_k and standard deviation σ_k which is cut at both ends, such that the degree distribution is zero below k_{min} and above k_{max} . We will always assume $m > k_{min}$ for otherwise there would be nodes which can never be activated by their presynaptic neighbors. For each node, a degree k_i is drawn from this distribution, and a number of links are randomly connected from i to the rest of the network.

A simulation consists of generating a network of nodes (all in the off state), switching a random fraction f of nodes to the on state and evaluating Eq. 3.1 until the number of active nodes is

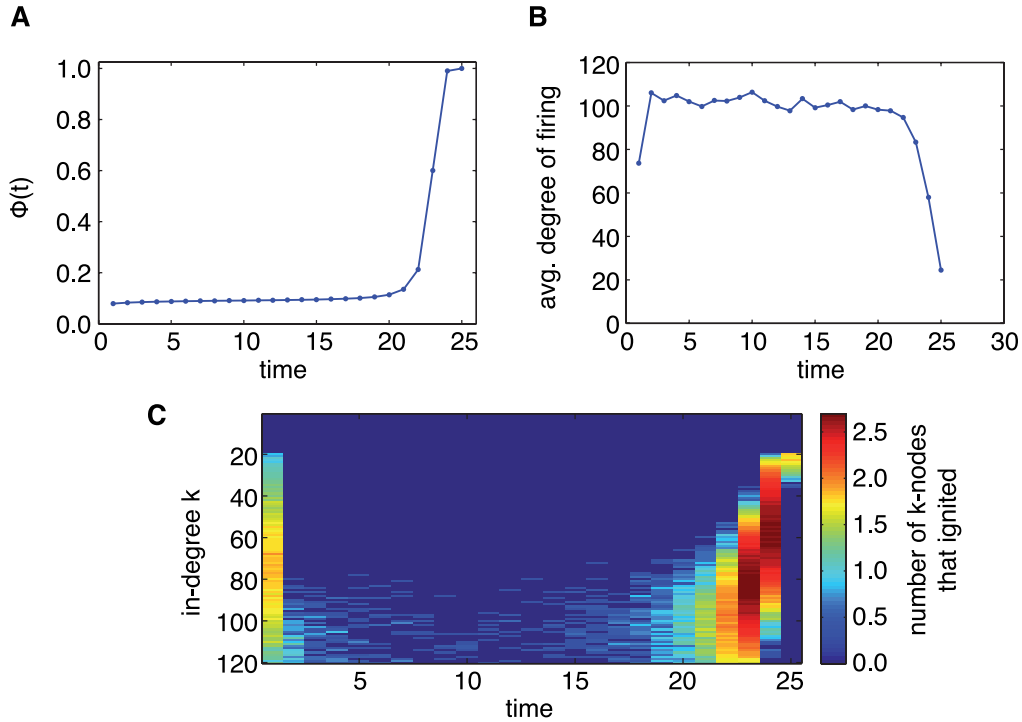


Figure 3.1: A typical realization of the quorum percolation model. (A) Plotted is the fraction of neurons in the on state $\Phi(t)$ as a function of time step. (B) For the same simulation, the average in-degree k of the neurons which were activated in the respective time step. (C) For the same simulation, the color-coded histogram of the degree of the activated nodes in each time step, which makes this plot a detail view of panel B. The parameters used for this simulation were $N = 50,000$, $m = 15$, $f = 0.079$, $k_{min} = 20$, $\mu_k = 75$, $\sigma_k = 25$, $k_{max} = 120$.

constant for one time step.

A typical realization of this process in terms of the fraction of active nodes in the network $\Phi(t)$ is shown in Fig. 3.1A. The initial value at $t = 1$ is given by the parameter f , corresponding to a random fraction of nodes which are set to on. By virtue of the QP model, $\Phi(t)$ must rise strictly monotonically or abort, and there is indeed a slightly upward tilted slope for the following time steps. At about $t = 21$ there is a sharp rise in the number of activated neurons, reaching $\Phi(t) = 1$ at time step 25. Analyzing the dynamics in terms of the average degree of activated nodes in a given time step (see Fig. 3.1B) we find that the high-degree nodes tend to be activated earlier. The average degree then decreases until reaching its lowest value at the last step. The last value is indeed less than half of the average degree in the network, indicating that the nodes activated in the last time step had an in-degree just barely over k_{min} . This is confirmed when plotting the actual degree distributions of activated nodes in Fig. 3.1C. The first time step is by construction a realization of the bounded Gaussian degree distribution described above, which is followed by a long period of sparse, high-degree activation and then a short period with high numbers of activated neurons of rapidly decreasing average degree.

It turns out that this model can have two distinct outcomes, as in the case of the discrete-time integrate-and-fire model discussed in the preceding chapter. Finite-size effects can lead to an aborted burst when not enough nodes are recruited initially to display the sharp rise seen in

Fig. 3.1A.

3.2 Deriving the degree distribution

(The work of the manuscript [139] was a collaborative effort, with the main contribution of the author of this thesis being the comparison of QP to integrate-and-fire models and the extension to structured topologies presented in the next section. The following analytical calculation is primarily the work of Tsvi Tlusty.)

We can understand the behavior of this system better if we split the nodes into k -classes of all nodes having a certain number of input connections k . We then denote by $\Phi_k(t)$ the firing probability of this k -class. It follows that the total firing probability is given by:

$$\Phi(t) = \sum_k p_k \Phi_k(t). \quad (3.2)$$

The mean-field solution for these classes can then be calculated using Ψ_k , the probability of a neuron from a k -class to have more than m active inputs.

$$\Phi_k(t+1) = f + (1-f) \Psi_k(\Phi(t)); \quad \text{with } \Psi_k = \sum_{l=m}^k \binom{k}{l} \Phi^l (1-\Phi)^{k-l} \quad (3.3)$$

In the limit of large μ_k and large m , we can replace Ψ_k by a step function $\Theta(k\Phi(t) - m)$, corresponding to the point where $\Phi(t)$ is just large enough (super-threshold) such that $k\Phi(t) = m$.

$$\begin{aligned} \Phi_k(t+1) &= f + (1-f) \sum_k p_k \Theta(k - m/\Phi(t)) \\ &= f + (1-f) P\left(\frac{m}{\Phi_t}\right) \end{aligned} \quad (3.4)$$

Here we have made use of the cumulative degree distribution $P(k) = \sum_k^\infty p_k$. The time continuous variety of this equation is given by:

$$\dot{\Phi}(t) = f - \Phi(t) + (1-f) P(m/\Phi(t)) \quad (3.5)$$

The beauty of this treatment is that we can now derive analytically the required degree distribution which is needed to get the (initially) exponential shape found in [51]. If the onset is of the shape $\Phi(t) \propto \exp(\alpha t)$ for some growth constant α , then plugging this into Eq. 3.5 yields the following:

$$P\left(\frac{m}{\Phi(t)}\right) = \frac{(1+\alpha)\Phi(t) - f}{1-f} \quad (3.6)$$

We can substitute $q = m/\Phi(t)$ to find a power-law in the degree distribution.

$$p_k = -\frac{d}{dq} P(q)|_{q=k} = \frac{(1+\alpha)m}{1-f} \frac{1}{k^2} \quad (3.7)$$

Note, however, that this calculation is only valid for the early onset of a burst, and we still expect the majority of the nodes to follow a Gaussian degree distribution [63, 138]. The solution is therefore a Gaussian distribution with a right tail that has a power-law shape with $p_k = B k^{-2}$. The two distributions are joined together at a certain degree k_{tail} by the expectation that 10% of the neurons are part of the power-law distribution, as the burst grows (after the *pre-burst*) by a factor of 10 [50, 139]. This distribution is shown in Fig. 3.2A.

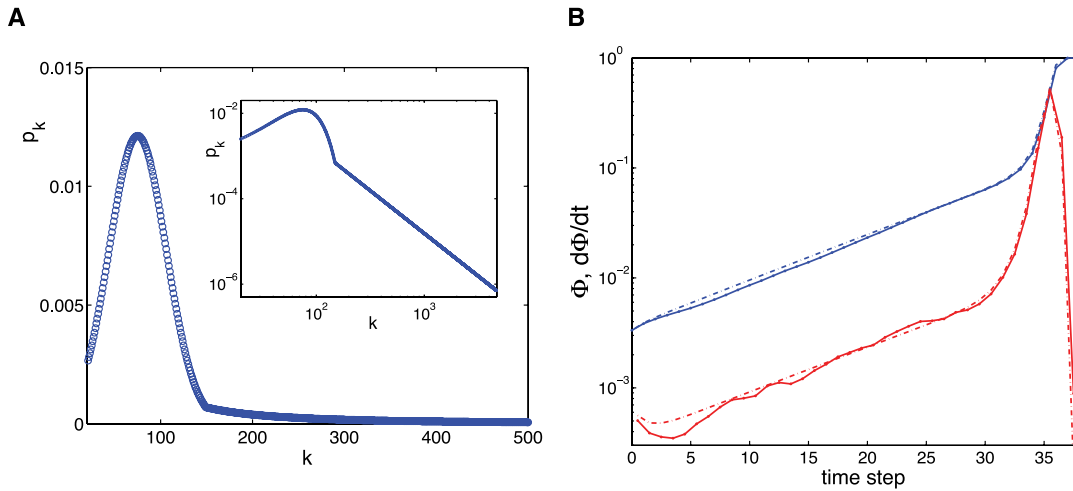


Figure 3.2: Gaussian degree distribution combined with a power-law right tail. (A) The degree distribution for a network of $N = 500,000$ neurons in linear and log-log scale (inset). The Gaussian distribution on the left (with parameters $\mu_k = 75$, $\sigma_k = 31$ and $k_{min} = 20$) and the power-law tail on the right hand side (parameters $B = 15.65$ and $k_{max} = 4,680$) are joined at $k_{tail} = 150$. (B) A typical simulation with the degree distribution shown in panel A. The fraction of active nodes $\Phi(t)$ (solid blue line) is plotted together with its temporal derivative (solid red line). The dashed lines are based on iterating a mean-field approximation based on Eq. 3.3 (derivation not shown, see [139]). Both panels reproduced with permission from [139].

The effect of this degree distribution on the percolation dynamics is indeed as expected, with an exponential growth phase until about $t = 30$, followed by a super-exponential rise phase until $\Phi(t) = 1$ (see Fig. 3.2B).

3.3 Effect of degree-degree correlations

To compare these results with those of the discrete-time integrate-and-fire model, where we saw an exponential growth even in random graphs, it is necessary to examine if deviations from the random graph — apart from a skewed degree distribution as discussed above — can lead to similar effects. In the integrate-and-fire case, positive degree-degree correlations caused an increase in variance in the onset times on individual neurons, and might therefore lead to a temporal development of the average activity $\Phi(t)$ which is better described by an exponential.

For the networks to display a given value of the degree-degree correlation coefficient C_1 , a random pair of links is selected and crossed, such that $A \rightarrow B$ and $C \rightarrow D$ becomes $A \rightarrow D$ and $C \rightarrow B$. This change is accepted if the resulting C_1 of the network is closer to the target value, and reversed otherwise. This procedure is repeated until the target value is reached with an error of less than 10^{-3} . A large number of bursts are then simulated in the resulting network. Each burst was time shifted such that it crossed the fraction of $\Phi = 0.2$ at time zero, and extended on both ends by constant extrapolation to allow for averaging over all simulated bursts.

It turns out that such a change towards higher C_1 does not increase the exponential phase of the burst onset, but only marginally increases the slope of the onset of activity, shortening the time until the network reaches the fully saturated state at $\Phi(t) = 1$ (see Fig. 3.3). The reason for this can be found in the work of Tlusty and Eckmann in [141]. They showed that for small networks

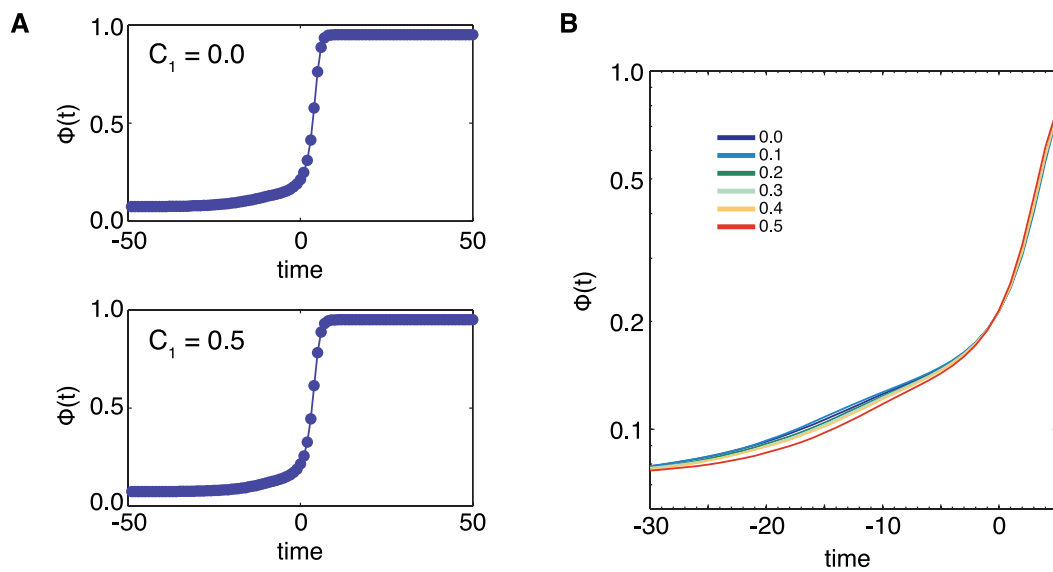


Figure 3.3: Dependence of $\Phi(t)$ on degree correlations in the underlying network. (A) Plotted are traces of $\Phi(t)$, averaged over 1000 bursts. Simulation were performed with networks of $N = 5,000$ which were based on a random graph (top panel) or engineered to display an degree-degree correlation coefficient C_1 of 0.5 (bottom panel). (B) Magnified plot of the onset of the bursts for networks with degree correlation ranging from 0 (blue line) to 0.5 (red line), again averaged over 1000 bursts each.

such as those in the simulations and those studies experimentally, where the number of nodes N is much smaller than a critical number N_* , the network structure is therefore effectively free of loops. This results in a quasi-random percolation dynamics, despite the topological structure imposed here.

We conclude that in quorum percolation models of the relevant system sizes, early firing time is determined by the in-degree k — or equivalently by voltage threshold (see [50, 139]) here set to 1 — and not by the higher-order topological structure in the network. This is in contrast to the results obtained in the discrete-time integrate-and-fire models. Their activity patterns are more similar to spontaneous neuronal activity where the larger topological neighborhood determines the expected firing time.

Chapter 4

Leaky integrate-and-fire dynamics

Building on the two previous chapters, we now study the onset dynamics in a biologically more realistic model system, to test the previous findings in a setting which allows for direct comparison of parameters and time scales with the biologically measurable quantities. The gold standard of such mathematical neuron models is the leaky integrate-and-fire model [142, 143] originally developed by Lapicque in 1907 [144]. It has since been generalized to include dynamic synaptic conductances and other detailed parts of the system, and its dynamical behavior has been studied in great detail [66, 145–147].

Importantly for the purpose of the present study, Tsodyks et al. have developed a model of a *depressive synapse*, in which the synaptic strength is reduced after repeated firing of the pre-synaptic neuron, caused by a depletion of neurotransmitter vesicles at the synaptic cleft [128, 129, 148, 149]. The pool of available neurotransmitters is refilled by diffusion of vesicles to the synapse, and thus the effective synaptic strength recovers over time. This class of synapses is known to generate bursts as seen in cultured networks, i. e. activity that synchronizes across the population over a short time window of less than 200ms [129, 150].

Such a system allows for the analysis of the spontaneous activity of a large number of biologically more realistic model neurons, to study what the determinants of early onset firing are, and to what extent the results are comparable to the discrete-time case.

4.1 Model characteristics

In contrast to the previous chapters, the leaky integrate-and-fire (LIF) model¹ is a continuous time model. Its sub-threshold dynamics are defined by the following differential equation for the membrane voltage $V_i(t)$ of a given neuron i [128, 129]:

$$\tau_m \frac{dV_i(t)}{dt} = -V_i + \frac{I_{syn}(t)}{g_l}. \quad (4.1)$$

Accordingly, the membrane voltage will decay (leak) with the membrane constant $\tau_m = 20\text{ms}$ to a resting voltage (here set to zero without loss of generality). The only sub-threshold deviations from this exponential decay can come from synaptic currents $I_{syn}(t)$ that the cell receives, resulting in a change in membrane voltage according to the leak conductance $g_l = 50\text{pS}$. If the membrane voltage of this neuron reaches a given threshold $\Theta = 20\text{mV}$, the voltage is reset to zero for the refractory period $\tau_0 = 20\text{ms}$ and post-synaptic neurons will receive an input synaptic current.

¹Simulations of the leaky integrate-and-fire model are performed using the Neural Simulation Technology (NEST) [151–153].

The synaptic current is itself given by a differential equation, determined by an exponential decay according to the synaptic time constant $\tau_{syn} = 2.0\text{ms}$ and inputs from pre-synaptic neighbors or external inputs, both resulting in exponentially decaying pulses of current:

$$\frac{dI_{syn}(t)}{dt} = -\frac{I_{syn}}{\tau_s} + \alpha_{int} \sum_{j=1}^N \sum_k A_{ji} E_{ji}(t) \delta(t - t_j^k - t_d) + \alpha_{ext} \sum_l \delta(t - t_{ext,i}^l - t_d). \quad (4.2)$$

The spike times of a neuron j form the set $\{t_j^k\}$, arriving after a delay time $t_d = 1.5\text{ms}$ in the form of input current at neuron i if the two cells are connected as determined by the adjacency matrix A , i. e. if $A_{ji} = 1$, and if there is a non-vanishing fraction of available neurotransmitters at this synapse, i. e. if $E_{ji}(t) > 0$. In such a case, the amplitude of the current pulse is further determined by the (homogeneous and constant) internal synaptic weight α_{int} . Note that the parameters are standard values for leaky integrate-and-fire models [66, 154], and that the determination of the exact value of α_{int} will be discussed below.

On the other hand, current pulses may arrive at neuron i from the outside, mimicking background activity and thermal noise in the ion channels of each neuron [48]. The time points of these spikes form the set $\{t_{ext,i}^l\}$. They are drawn independently according to an average rate $\nu_{ext} = 1.6\text{Hz}$ and with constant amplitude $\alpha_{ext} = 4.0\text{pA}$.

The fraction of available neurotransmitters at a given synapse is modeled according to the following formula, differentiating between resources in the *effective* state $E_{ji}(t)$, which is responsible for mediating the actual synaptic current, the *recovered* state $R_{ji}(t)$, and the implied *inactive* state $I_{ji}(t) = 1 - R_{ji} - E_{ji}$ [128, 129].

$$\begin{aligned} \frac{dE_{ji}(t)}{dt} &= -\frac{E_{ji}}{\tau_{inact}} + U R_{ji} \sum_k \delta(t - t_j^k) \\ \frac{dR_{ji}(t)}{dt} &= -\frac{1}{\tau_{rec}} (1 - R_{ji} - E_{ji}) \end{aligned} \quad (4.3)$$

After a pre-synaptic spike, a fraction $U = 0.3$ of the synaptic resources of a given synapse in the recovered state is entering the effective state, and are then converted to the inactive state with the rate $1/\tau_{inact}$, where $\tau_{inact} = 3.0\text{ms}$. From the inactive state, synaptic resources are transferred back to the recovered state with slow time constant $\tau_{rec} = 500\text{ms}$ [129].

4.2 Adaptive burst rates

The rate of bursting in the network ν_{burst} is typically 0.1Hz for in vitro cultures, and in particular of the experiments studied in Part II of this thesis. To obtain such a rate in our simulations, we employ an iterative scheme to find an internal synaptic strength consistent with the targeted burst rate.

Starting from an initial value of $\alpha_{int} = 5.0\text{pA}$, the synaptic strength is iteratively increased until the measured burst rate in the simulation changes for the first time. In the following steps, we use linear extrapolation of the last two burst rates and synaptic weights to predict the weight consistent with the targeted burst rate. This is repeated until the value of α_{int} is such that the network bursting rate is $0.10 \pm 0.01\text{Hz}$.

The distribution of the fraction of active nodes across time is bimodal, with the most probable fraction being less than 10%, and with more than 90% within a fully developed burst. Consequently, we measure the burst rate by defining a burst to occur when more than 40% of the population is active in a window of 50ms (the time scale of burst onset, see Fig. 1.2). Given that less than 0.5% of

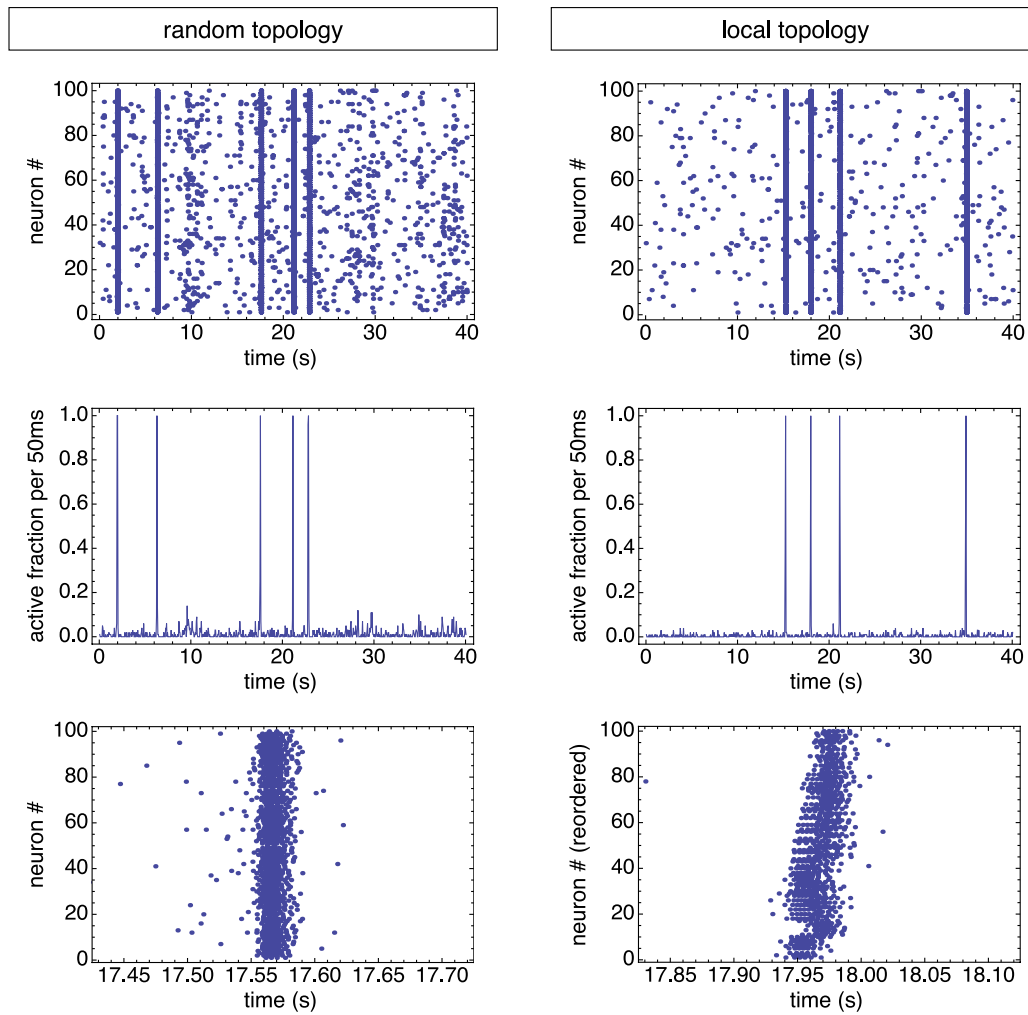


Figure 4.1: Spiking activity and bursting in networks with random and local topology. The raster plots of spiking activity across the population (top row) show a switching behavior between low-frequency asynchronous spiking and synchronized burst-like dynamics. The middle row shows the same data, but plotted as a function of the fraction of nodes that were active within a window of 50ms. Simulation were performed on a random graph (left column) and a local topology generated with $\lambda = 0.25\text{mm}$ (right column). Burst rates were $0.10 \pm 0.01\text{Hz}$ (averaged over 60 minutes of simulated time) in both cases. The bottom panels are raster plots, magnified to show a detailed view of an example burst along a temporal window of 300ms. In the case of the local topology (bottom right panel) neurons have a 2D-spatial position which correlates with the topology, as neurons that are closer to each other in Euclidean space are more likely to be connected. Reordering them according to their y-coordinate in Euclidian space reveals that there is a spatial component also to burst generation. This particular burst is ignited in the lower part of the culture and then elicits behavior in the upper part. Note that such a spatial component is not present in the random graph, and therefore a reordering would have been a random permutation in the case of the random graph.

data points show an active fraction between 0.2 and 0.8, this measure of bursting rate is robust against selection of the threshold.

The result of such a simulation on a random graph is shown in the top left panel of Fig. 4.1. As expected from the results of [128, 129], the raster plot reveals a two-state dynamics, switching from non-bursting to bursting regime. This behavior is again plotted below in terms of the fraction of active nodes in the network over time (see the left middle panel of Fig. 4.1). The active fraction fluctuates around zero, but then shows peaks during bursts, where consistently over 95% of nodes are firing at least once in a 50ms time window, confirming that the bursts in the simulations are indeed recruiting the whole network as seen in vitro [51, 61].

4.3 Random and local topologies

So far we have only considered leaky integrate-and-fire neurons on a random graph. In real biological networks in vitro, there often are correlations between the spatial position of a node, and the topology of the network graph that surrounds it, resulting in a connection probability that depends on the spatial distance between two connected nodes [37, 49, 84].

To account for possible locality of connections, we will consider — in addition to the random graphs discussed above — networks where the probability density of connection between two neurons i and j at positions \vec{x}_i and \vec{x}_j has a Gaussian shape, characterized by the length scale λ . We therefore assign each neuron in the network a random position on a two-dimensional square of side length 1mm, drawing each coordinate independently from a uniform distribution. Coordinates are rejected only if they would place a neuron closer than a minimum distance of $d_{min} = 10\mu\text{m}$ to previously placed cells.

$$P(r) = P_0(N, \lambda) \exp \left\{ -\frac{(\vec{x}_i - \vec{x}_j)^2}{\lambda^2} \right\} \quad (4.4)$$

To have networks that are local, but still comparable to the random case, the average probability of connection is set to be equal in both cases. This ensures that the random network is indeed the limit of the local networks for $\lambda \rightarrow \infty$. Therefore, we employ a two-step algorithm: In the first step, we set $P_0 = 1$, generate the connections and the resulting number of connections C_0 is recorded. Then, the links are discarded and regenerated with a new $P_0 = C/C_0$, where C is the average target number of links given by the average probability of connection p via $C = pN(N-1)$. Throughout this work, we only consider networks and values of λ where $C_0 > C$ and thus $P_0 < 1$.

The net result of this procedure is an adjacency matrix A where the topology is correlated with the spatial position of the neurons. The simulation of spiking activity is then performed as for the random graph. Note that the adaptive algorithm described above resulted in significantly weaker connection strength α_{int} for the local connectivities, yielding for instance values of 5.2pA for $\lambda = 0.25\text{mm}$ and 6.5pA for $\lambda = 0.75\text{mm}$, while yielding 6.6pA in the case of the random graph. For details see Table 5.1.

The resulting spiking activity is plotted in the right column of panels in Fig. 4.1. Analogously to the random graph, the activity is clustered between two distinct types of activity: low-frequency spontaneous activity and high-frequency synchronized bursts (see top right panel of Fig. 4.1). The same is visible in the fraction of active nodes, with firing probabilities again reaching above 95% as in the random case (see Fig. 4.1, middle right panel). Interestingly, the spatial aspect of the network topology results in bursting dynamics that now have a spatial component: When reordering the neurons according to their position in Euclidean space, the geometrical development of firing activity within a network burst becomes apparent. For instance, the example burst shown in the

bottom right panel of Fig. 4.1 is ignited in the bottom half of the network and proceeds through to elicit neurons in the top half of the simulated culture.

4.4 Analytical prediction of the first firing time

It turns out that the expected first firing time of a neuron in the simulated networks can be predicted analytically to a good degree of accuracy. Proceeding in a similar manner as in Chapter 2, we now develop a solution of this problem. In general, such a solution is only possible in the context of constant or periodic Gaussian stochastic inputs [155–157] which is inherently not applicable at the burst onset. Thus, the following calculations are based on a noise-free mean-field ansatz.

In a first order approximation, we assume a monotonically rising, exponential firing frequency of the mean field resulting in a synaptic current from one pre-synaptic neuron $I_{syn,mf}(t) = \alpha a \exp(bt)$. The parameter b is given from the experiment as $b = 45.0\text{s}^{-1}$ [51], and without loss of generality we set $a = 1$, which will be discussed later in the text.

We can simplify the question of solving Eq. 4.1 for the membrane voltage to the question of at which point in time the input current $I_{syn}(t)$ will be strong enough such that the voltage reaches the threshold Θ . We will denote this threshold current $\hat{I} = \Theta/R$ and its position in time by $\hat{t}^{(1)}$ such that $I_{syn}(\hat{t}^{(1)}) = \hat{I}$. We can immediately solve for the first firing time $\hat{t}^{(1)}$:

$$\hat{I} = \alpha a k \exp(b \hat{t}^{(1)}) \Rightarrow \hat{t}^{(1)} = \frac{1}{b} \log \frac{\hat{I}}{\alpha a k}. \quad (4.5)$$

To take the nearest neighbors into account in a second-order approximation in analogy with Chapter 2, we make use of the fact that the input-output firing rate under stationary conditions (constant input current I_{in}) is known [143].

$$f(I_{in}) = H(R I_{in} - \Theta) \left[\tau_0 + \tau_m \log \frac{R I_{in}}{R I_{in} - \Theta} \right]^{-1} \quad (4.6)$$

Here we have denoted by H the Heaviside step function, by R the membrane resistance and by τ_0 the refractory period. Again assuming that the change in current is slow compared to the membrane time constant, we can write down the input current into a given neuron due to one input neuron, which in turn has $k^{(2)}$ inputs:

$$I_{syn}^{(1)}(k^{(2)}, t) = \alpha H(\alpha a R k^{(2)} \exp(bt) - \Theta) \left[\tau_0 + \tau_m \log \frac{\alpha a R k^{(2)} \exp(bt)}{\alpha a R k^{(2)} \exp(bt) - \Theta} \right]^{-1} \quad (4.7)$$

This relationship is plotted in Fig. 4.2A for different values of the second-order input degree $k^{(2)}$. For the same exponentially rising firing frequency of the mean field, a higher number of the second-order degree serves to amplify the current from the mean field and thus leads to an earlier firing onset of the second-order neuron. After an approximately linear period in each case, the first-order input current saturates as the second-order firing frequency approaches its theoretical limit of $1/\tau_0$.

For the networks used in the simulations shown in Fig. 4.1, the shown degrees correspond approximately to the range of degrees present in the network (mean degree 12.1, standard deviation 3.4). Thus, the expected first firing time of a neuron with only one input connection would be in a range of about 30ms, depending on its second-order degree.

If we further assume that the average degree $\mu_k \gg 1$, we can approximate the second-order degree distribution by a delta distribution. Thus, we can finally write the synaptic current due to

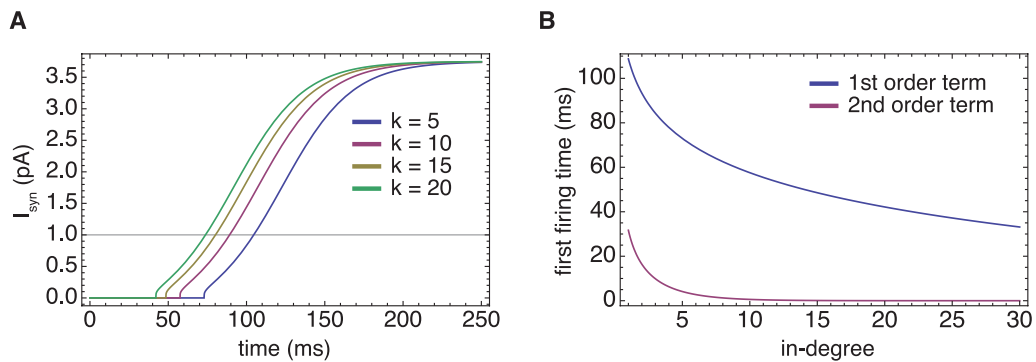


Figure 4.2: Analytical prediction of the first spike time. (A) For different values of the second-order input degree, the synaptic current into a first-order neuron from one pre-synaptic second-order neuron receiving exponentially rising input currents from the mean field is plotted. For a higher number of the second-order degree the second-order neuron starts to fire earlier in time, showing an approximately linear rise of the firing frequency, and thus the output current for about 50ms, after which the output current saturates. The horizontal line indicates the threshold current \hat{I} , i. e. the current which would be sufficient to drive the first-order neuron to spike from the plotted current alone. (B) The two terms of the second order solution for the expected first firing time. The first term is shown in blue, the second term is shown in red. Note that we refer to the first term as the first-order term because it is the first-order solution, even if it depends on the second-order degree when it is part of the second-order solution. The sum of the two terms would give the solution itself, for the respective degrees.

the second-order input neurons as simply a multiplication of Eq. 4.7 by the first-order input degree.

$$I_{syn}^{(1)}(k^{(1)}, k^{(2)}, t) = \alpha k^{(1)} H(\alpha a R k^{(2)} \exp(bt) - \Theta) \left[\tau_0 + \tau_m \log \frac{\alpha a R k^{(2)} \exp(bt)}{\alpha a R k^{(2)} \exp(bt) - \Theta} \right]^{-1} \quad (4.8)$$

For the regime in time where the Heaviside function is positive, i. e. $\alpha a R k^{(2)} \exp(bt) > \Theta$, which can be rewritten using Eq. 4.5 as $t > \hat{t}^{(1)}(k^{(2)})$, the solution with regard to the onset time can be calculated by way of standard algebraic manipulations, yielding:

$$\hat{t}^{(2)} = \frac{1}{b} \log \left\{ \frac{\Theta e^D}{\alpha a R k^{(2)} (e^D - 1)} \right\} \quad \text{with} \quad D = \frac{\alpha k^{(1)}}{\hat{t} \tau_m} - \frac{\tau_0}{\tau_m}. \quad (4.9)$$

We can rewrite this equation in a more intuitive form.

$$\hat{t}^{(2)} = \hat{t}^{(1)}(k^{(2)}) + \frac{1}{b} \log \left\{ \frac{e^D}{e^D - 1} \right\} \quad (4.10)$$

The two terms of this sum are plotted in Fig. 4.2B. Both terms depend on the respective degree, but the second-order term (the second term shown above) shows a strong dependence on the degree only for low degrees $k < \mu_k$, and approaches a constant value for higher degree.

4.5 Comparison to simulated data

To allow for the comparison of the analytical solutions Eq. 4.5 and Eq. 4.10 with numerical simulations, the first firing times need to be extracted from the simulations. This is done in an analogous way to the case of the discrete-time integrate-and-fire-case, i. e. by extracting from the observed firing rates relative to a burst the probability to fire for the first time, and from this in turn the expected first firing times.

For simulations lasting 60 minutes of simulated real time, we detect burst events as events when the fraction of active neurons crossed from below a threshold of 20% of neurons with a time window of 5ms [50]. We then pruned events such that those remaining had a minimum distance of 200ms from each other, in order to avoid the effect of transient firing rate fluctuations shortly before or after adjacent events in our analysis. In a next step, the threshold crossing time was calculated with sub-sampled resolution using a linear interpolation of the averaged signal, and the individual firing rate signal of neurons was aligned such that the threshold crossing of the population average occurred at time $t = 0$. Thus, this averaging resulted in an average burst across events for each neuron. As described above (see also Eq. 2.9) this enables the extraction of the probability to fire for the first time relative to the burst onset.

As the choice of $t = 0$ at the burst threshold is arbitrary, we proceed again (cf. Section 2.4.1) by fixing the mean of the distributions of both numerical and analytically determined onset times to zero. The initially arbitrary choice of a is thus fixed, since any value of a corresponds to a constant time offset due to $a \exp(bt) = \exp(b(t + t_{\text{off}}))$ with $t_{\text{off}} = a/b$.

The result of this analysis for the random graph, i. e. for activity patterns as shown in the left column of Fig. 4.1, is shown in Fig. 4.3. For the first-order solution, the prediction of onset times follows the numerical data closely, with an error shown as the standard deviation of the difference between simulation and prediction of 3.42ms. Furthermore, there is also a very clear linear correlation between the numerical and analytically evaluated onset times (see inset of the left panel of Fig. 4.3), yielding a Pearson's correlation coefficient of 0.835.

The second-order solution further improves on the first-order prediction, resulting in a correlation coefficient of 0.877 as shown in the inset of the right panel of Fig. 4.3. On the other hand,




Figure 4.3: Prediction of the expected first firing time at the onset of a burst in a random graph of leaky integrate-and-fire neurons. For neurons sorted according to their expected first firing time (blue line) as extracted from simulations, the analytical prediction is plotted according to the first-order (left panel) or second-order (right panel) solution. The insets are scatter plots of numerical (horizontal axis) against analytically predicted first firing time (vertical axis) in milliseconds relative to the burst onset.

the standard deviation of the prediction error is 3.90ms, so there is a slight increase in the average difference between simulation and prediction due to an increased slope of the prediction across the sorted population.

When comparing the prediction errors with Chapter 2, it is worth pointing out that the prediction is made more difficult by an increased heterogeneity in the model parameters: Contrary to the discrete time dynamics, the simulation is not reset after each burst, so that temporal long-range correlations of the activity of consecutive burst events cannot exist in the discrete-time models studied before. In the leaky integrate-and-fire case, such correlations do exist while being neglected in the analytical solution presented above. This results for instance in a positive correlation between the number of action potentials in one burst and the time it takes until the next burst is elicited, as neurotransmitter resources are more depleted and take a longer time to be regenerated. The analytical calculations also neglect the small but non-vanishing firing probability of neurons outside of a burst, due to their external inputs or non-bursting internal activity, which can affect the numerical evaluation of the first spiking time due to bursting (see above).

In the case of local topologies generated according to Section 4.3 for two cases of λ , Fig. 4.4 shows the numerical result next to the analytical predictions. For the first case of $\lambda = 0.25\text{mm}$, the first-order solution shows a good agreement, with a Pearson's correlation coefficient between simulation and prediction of 0.827 and a standard deviation of the prediction error of 4.59ms (see top left panel of Fig. 4.4). Note that the deviation is partly dominated by a single outlier as visible by the peak at very late onset times — the standard deviation of the prediction error when removing this single neuron was only 3.67ms. The second order solution is slightly improved in the sense of an improved linear correlation with a Pearson's coefficient of 0.888. The prediction error is at a standard deviation of 7.24ms, or 4.53ms without the outlier neuron as above, indicating a local topology with correlations that go beyond first order and the outlined approximation to second order (using $k^{(2)}$ as the average degree of the neighbors of a given node).

The case of $\lambda = 0.75\text{mm}$ is a network in between the very local topology described above, where connectivity in the 1mm^2 plane is dominated by the spatial local neighborhood, and the random topology discussed before ($\lambda \rightarrow \infty$). The prediction displays a high linear correlation with

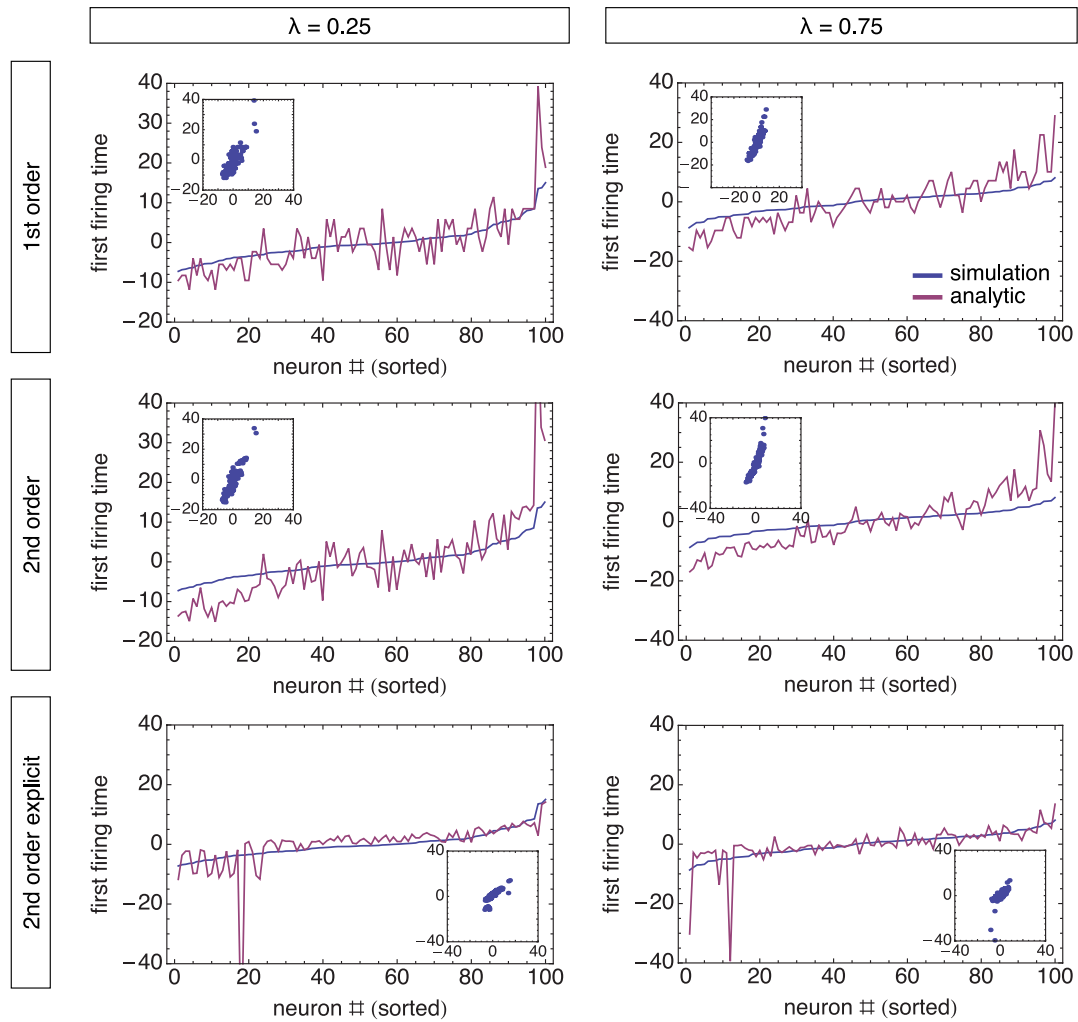


Figure 4.4: Prediction of the expected first firing time at the onset of a burst in a locally structured graph of leaky integrate-and-fire neurons for different values of the length scale of typical connections λ (in millimeters). For neurons sorted according to their expected first firing time (blue line) as extracted from simulations, the analytical prediction is plotted according to the first-order (top row), second-order (middle row) or “explicit” 2nd order solution (bottom row). The insets are scatter plots of numerically (horizontal axis) against analytically predicted first firing time (vertical axis) in milliseconds relative to the burst onset.

coefficients of 0.888 for the first order solution and 0.909 for the second order solution. As before however, the prediction error increases from first to second order, here from 5.39ms to 6.68ms (24%).

Overall the local topologies display a picture that is similar to the random topology: A good prediction for the first order solution, showing a very good linear correlation (Pearson's correlation of $p > 0.8$ in all cases) and is even improved when going to second order. However, the second order prediction slightly overestimated the width of the distribution of onset times.

The source of the remaining error in predicting the expected first firing time can be better understood — and the prediction quality further improved — in the following way: In the outliers pointed to above, the proposed error is caused by assuming that the individual degrees of the neighbors of a neuron can be substituted by the average degree, denoted above as $k^{(2)}$. However, finite-size effects in the network simulated here ($N = 100$) can potentially lead to large errors. However one can calculate the numerical solution to the analog of Eq. 4.10 without this averaging step.

$$\hat{I} = \alpha \sum_{j \in C} H(R k_j I_{mf}(\hat{t}^{(2e)}) - \Theta) \left[\tau_0 + \tau_m \log \frac{R k_j I_{mf}(\hat{t}^{(2e)})}{R k_j I_{mf}(\hat{t}^{(2e)}) - \Theta} \right]^{-1} \quad (4.11)$$

Here we have denoted with C the set of indices of neighbors of the neuron of interest, and by $I_{mf}(t)$ the time-dependent current from the mean field to the neighbors due to a single second-order connection. In effect, k_j are the number of input connections of the neighbors. Note that the first order degree $k^{(1)}$ is, in this equation, present only implicitly as the length of the set C .

The bottom row of Fig. 4.4 shows the results of numerically solving Eq. 4.11 for the onset time $\hat{t}^{(2e)}$, denoted as such as the “second-order explicit” solution. For both $\lambda = 0.25\text{mm}$ and $\lambda = 0.75\text{mm}$, the reconstructed values follow the simulation results much better, and the prediction error is dramatically reduced, with 69% and 81% of neurons showing a prediction error of less than 3ms. On the other hand, there are some neurons (less than 3%) whose predicted onset time is more than 20ms earlier than the the simulated one, showing higher order correlations in the topology, leading to a small number of outlier nodes whose connectivity in terms of the input degree of their neighbors is low, but who are linked to tightly connected, early firing groups.

Part II

Reconstruction of network connectivity

Chapter 5

Neuronal activity as shaped by network structure

The overarching finding of the previous study of mathematical models of neuronal connectivity and activity was that the dynamical nature of the burst onset is dominated by the connectivity matrix of the underlying network, highlighting the effect of network structure on spiking dynamics. Given this strong interdependence, we seek here to extract the structural connectivity from recordings. Such a reconstruction, ultimately aiming for a map of the circuitry of the brain, has remained a hard, and largely unsolved problem to the present day, as described in Chapter 1.

The complexity in this question arises from the fact that any measure of causal influence is a priori not necessarily a measure of synaptic connectivity. In general, there are three forms of interaction in the study of neuronal ensembles, called *structural*, *functional* and *effective* connectivity [88, 158, 159]. Whereas functional connectivity refers to correlation between signals, effective connectivity is the result of a causality measure such as Transfer Entropy. Structural connectivity, on the other hand, refers to the underlying physical network structure determined by synaptic connections. As will be shown below, effective connectivity as a representation of the ongoing network dynamics can be measured in a number of ways, and knowledge of the system at hand is necessary to determine if and when this effective connectivity might overlap with the structural connectivity.

The object of any measure of (effective) connectivity is a measure of the activity or the dynamics of the individual parts of a dynamical system. Here, we will use calcium imaging recordings as an indirect, but experimentally reliable indicator of time-dependent neuronal activity. Calcium fluorescence imaging has in recent years become a standard tool to quantify activity of a population of neurons [138, 160–162]. In this technique, the fluorescent dye is inserted into the culture medium where it diffuses through the membranes of neurons into the cell bodies. Depending on the calcium concentration inside of a given cell, Ca^{2+} ions can bind to the dye. In the event of an action potential, calcium channels open, the calcium concentration inside the cell transiently increases and more calcium binds to the dye, and unbinds some time later. The rates of binding and unbinding are such that during an action potential, the concentration of calcium bound to the dye rises very steeply on a time scale of 1–3ms and decays on a much slower time scale of about 1–2s depending on the dye used [163–166]. The net effect of this is that the amplitude of observed fluorescence traces show an approximately linear relationship to the density of action potentials in time [59], making calcium imaging an indirect but reliable indicator of neuronal activity.

5.1 Experimental recordings of calcium fluorescence

In this work, we make use of the fluorescent dye Fluo-4 [167], to record the calcium activity from cultured in vitro networks from hippocampal neurons from rat embryonal tissue. These neurons have been dissociated from the brain tissue, have grown in a nutritious medium onto a glass substrate and have spontaneously formed networks. After 3–4 days in vitro (DIV) the neurons begin to be spontaneously active, showing asynchronous activity patterns that synchronize to bursts after about 7 days in vitro [168]. As the onset of a burst is dominated by excitatory connections, we will focus on networks where no inhibition is present (with the exception of Chapter 9). In the culture, this is done by adding a concentration of $40\mu\text{M}$ bicuculline, a GABA_A antagonist. See Appendix A for details on the experimental preparation.

An example of such a culture is shown in Fig. 5.1A. Neurons, having a typical diameter of $10\mu\text{m}$, are clearly visible as white round dots in the bright field image. A region of interest (ROI) is selected manually for each cell that is distanced enough from other neurons such that the ROI selections do not overlap.

The fluorescence activity can thus be recorded over time for hundreds or thousands of neurons simultaneously, at a temporal resolution of 50Hz. Five example traces from a set of $N = 1720$ neurons are shown in the left panel of Fig. 5.1B. Baseline activity is separated by bursts, after which the fluorescence slowly decays back to the resting level. Beside the bursts, small fluctuations can be seen, corresponding to low-frequency asynchronous activity.

This bursting activity is even more clearly visible in the population average of the fluorescence (see left panel of Fig. 5.1C). The switching between asynchronous and synchronous phases of activity can then be conveniently plotted as the histogram of the mean fluorescence over the full duration of the recording (see left panel of Fig. 5.1D). The majority of time points exhibit low-level activity, as indicated by the Gaussian-like peak on the left side of the histogram. This Gaussian corresponds to regions in time where the population firing rate is low. The other end of the histogram, the high-fluorescence tail, corresponds to network bursts and thus a peak in the mean fluorescence of the culture. As can be seen from the histogram and the traces of mean fluorescence, this event is less frequent by about two orders of magnitude, but will nevertheless dominate network activity and fluorescence while it is active.

5.2 Simulation of calcium fluorescence

To complement the experimental recordings and to be able to benchmark the performance of causality measures, we again refer to simulated time series. The basis of these simulations is the standard current-based leaky integrate-and-fire model as discussed in Chapter 4, and as defined by the equations 4.1, 4.2 and 4.3.

Simulations are performed on a network of $N = 100$ neurons on one of two families of networks: *Locally clustered networks* (see Section 4.3) which are generated according to a Gaussian kernel dictating the decay of connection probability with distance, and *non-locally clustered networks*. The latter correspond to networks which are engineered to display a given value of the average clustering coefficient, which is defined on the graph topology alone, without a dependence of the connectivity on the spatial position. Local clustering has been observed in neuronal systems [49, 84], as well as non-local clustering [22, 170], making both important model systems for real biological networks.

As the networks in the models discussed in this thesis are all directed networks, a generalization of the clustering index for undirected networks must be chosen. As before, we will make use the full clustering coefficient (see Eq. 2.7).

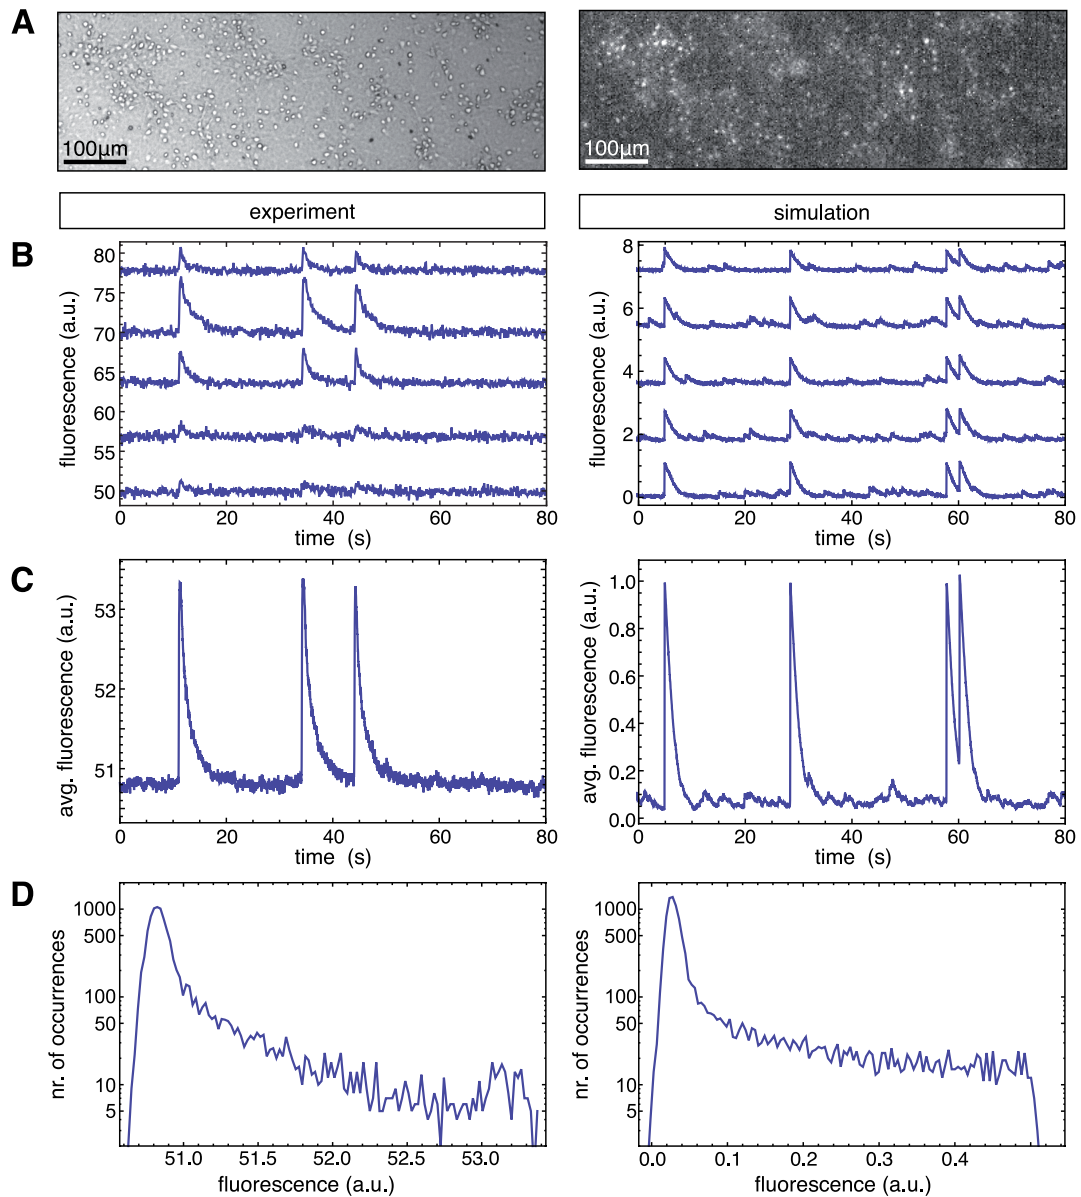


Figure 5.1: Calcium imaging recordings from in vitro cultures and simulations. (A) Bright field (left panel) and dark field fluorescence image (right panel) of a small subset of a cultured neuronal network at DIV 12. Bright spots in both panels correspond to neuronal cell bodies. (B) Five example traces of manually selected neurons of the culture shown in panel A (left panel) and of a simulated network (right panel). (C) The population average of the fluorescence, highlighting the presence of network burst events, of the networks shown in panel B. (D) A histogram of the average fluorescence as shown in panel C over 60 minutes of recorded (left panel) or simulated time (right panel). Figure reproduced with permission from [169].

To generate non-locally clustered networks, first a random graph is generated according to given probability of connection p . Then two links, $A \rightarrow B$ and $C \rightarrow D$ with A, B, C and D all being different neurons, are chosen at random and crossed, resulting in the connections $A \rightarrow D$ and $C \rightarrow B$. If the full clustering index CC of this new topology is closer to the target value, the change is kept and otherwise discarded. This procedure is repeated until the target clustering coefficient is reached with an accuracy greater than 0.1%.

Note that both the non-locally clustered networks and the locally clustered networks contain the random graph as a limiting case. As discussed before, the local networks display this limit as the length scale of the system $\lambda \rightarrow \infty$, whereas the non-locally clustered networks display it at $CC = p$, i. e. the clustering coefficient of a random graph is by definition given by the probability of connection (see Eq. 2.7).

The simulation of calcium fluorescence time series, for a sample point acquired at time t , is based on a standard model from the literature [171] and thus on the spike times from the leaky integrate-and-fire model. The calcium concentration bound to the simulated fluorescent dye is determined by an exponential decay with slow time constant $\tau_{Ca} = 1.0s$ and instantaneous positive offsets of amplitude $A_{Ca} = 50\mu M$ for each of the n_t action potentials that the respective neuron elicits. We have omitted here for readability the index i for each neuron.

$$[Ca^{2+}]_t - [Ca^{2+}]_{t-1} = -\frac{\Delta t}{\tau_{Ca}}[Ca^{2+}]_{t-1} + A_{Ca} n_t \quad (5.1)$$

The calcium concentration is convolved with a saturating Hill function with saturation concentration $K_d = 300\mu M$, and Gaussian white noise η_t with mean zero and a standard deviation of 0.03 is added, corresponding to spontaneous fluctuations of the fluorescence and camera (photodetector) noise.

$$F_t = \frac{[Ca^{2+}]_t}{[Ca^{2+}]_t + K_d} + \eta_t \quad (5.2)$$

To further increase the realism in the simulations, we have (unless otherwise noted) included an artifact for light scattered between ROIs which are close to each other [172, 173]. The light scattering cannot be fully eliminated using deconvolution techniques due to the limited field of view, inaccuracies in inferring the value of the scattering radius λ_{sc} and anisotropies in the recording medium. The parameters of the simulation are $\lambda_{sc} = 0.15mm$ and $A_{sc} = 0.15$, with the latter denoting the (residual) amplitude of this artifact.

$$F_{i,t}^{sc} = F_{i,t} + A_{sc} \sum_{j=1, j \neq i}^N F_{j,t} \exp \left\{ - (d_{ij}/\lambda_{sc})^2 \right\} \quad (5.3)$$

Here d_{ij} denotes the spatial distance between neuron i and neuron j . The length scale used for simulations in this work is shown compared to an example network in Fig. 5.2.

Five example traces of a resulting simulation are plotted in the right panel of Fig. 5.1B. Similarly to the case of real recordings, periods of relative silence are broken by network bursts, visible as peaks in the simulated time series. The bursts are again visible more clearly in the population average of the fluorescence (see the right panel of Fig. 5.1C). The histogram in the right panel of Fig. 5.1D highlights the similar structure of the average fluorescence over time when compared to the real recordings.

Note that for an improved sampling of probability distributions, our analysis will later be based on the difference signal $\Delta F_t = F_t - F_{t-1}$ rather than the fluorescence time series itself.




Figure 5.2: Radius of the light scattering artifact used in simulations of calcium fluorescence. Plotted are cell positions of a network of neurons (black dots) on a square of length 1mm, and the scattering radius λ_{sc} superimposed over three example neurons (shown in light green). Figure reproduced with permission from [169].

non-locally clustered			locally clustered		
CC	$\langle \alpha_{int} \rangle$ (pA)	s.d. of α_{int}	λ (mm)	$\langle \alpha_{int} \rangle$ (pA)	s.d. of α_{int}
0.1	6.604	0.146	0.25	5.207	0.171
0.2	6.156	0.124	0.50	6.241	0.166
0.3	5.719	0.054	0.75	6.481	0.150
0.4	5.361	0.113	1.00	6.556	0.230
0.5	5.274	0.067	1.25	6.505	0.158
0.6	5.214	0.209	1.50	6.519	0.113

Table 5.1: Synaptic weights in networks of different topological indices. For networks generated according to a given topological index, CC or λ , the average synaptic weight α_{int} and the corresponding standard deviation, across 6 networks of the given type, is shown. For details on the adaptive simulation see Section 4.2. Table reproduced with permission from [169].

5.3 Burst characteristics for different network topologies

An interesting first question is how much information can be retrieved about the statistical properties of the underlying network topology from the observation of bursts alone. If, for instance, only locally clustered networks were showing bursting behavior, we might derive the degree of clustering from the histograms of fluorescence as shown in Fig. 5.1D.

This question can be studied by systematically by varying the topological parameters in the network ensembles described above and observing the resulting activity (see Fig. 5.3). Three networks, each prescribed with very different values of the full clustering coefficient, are simulated according to the adaptive procedure (see Section 4.2) aiming for a target bursting rate of 0.1 Hz. Despite the fact that the underlying network as well as the synaptic weights in the simulation are different (see Table 5.1 for the actual values) there is no clear distinction to be made from the activity patterns shown. Furthermore, even the distributions of inter-burst intervals (IBIs) are strikingly similar in shape for all networks. The same is true when comparing simulations for the locally clustered ensemble, varying the length scale λ (not shown).

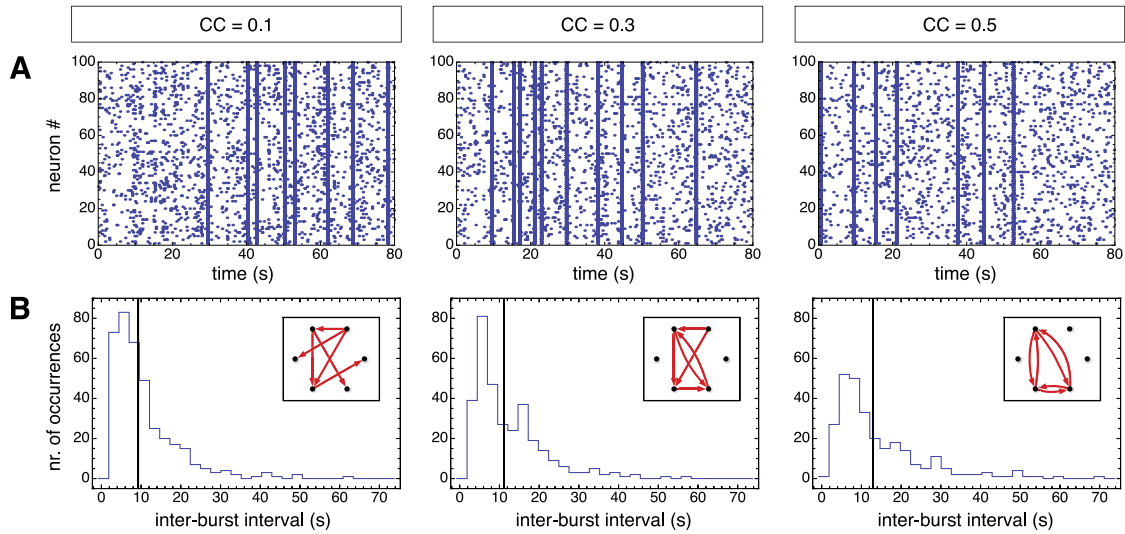


Figure 5.3: Spiking activity showing bursting for a range of clustering coefficients. (A) For three values of CC (0.1 corresponding to a random graph) we plot the resulting spiking activity, for a target bursting rate of 0.1Hz. (B) Inter-burst interval (IBI) distributions for the three simulations shown in panel A, with the vertical line indicating the mean of the distribution. Insets show cartoon networks having the same clustering coefficient as those of the simulated networks. Figure reproduced with permission from [169].

We conclude that the observation of bursting patterns alone does not allow for the determination of the network topology, and that more precise measures, relating neuronal activity with the underlying topology, are necessary.

Chapter 6

Measuring causality

The definition of any measure of effective connectivity, relating dynamics to interactions, is built on a specific definition of causal interactions.

In this chapter, we will present standard measures of causality, and then develop the novel approach of this thesis. Note that because of the improvements we will develop later, the formulas given in the following two sections are given as in the literature — the formulas which will be used in Chapter 8 and thereafter, to compare the reconstruction quality of different approaches, are stated in Section 6.4.4.

6.1 Linear measures

As outlined in Chapter 1, the first attempts to measure causality were based on the assumption of the linearity of the causal effect. The most well-known approach, still widely used today, is based on Pearson's cross correlation coefficient. This coefficient is used to compute the cross-correlogram, i. e. the correlation between the target time series x (consisting of S data points) and the putative, time-shifted source time series y for different lags. The peak in the cross-correlation is identified to adjust for a causal shift in the time series.

$$XC_{Y \rightarrow X} = \max_{\Delta t=1 \dots t_{max}} \left\{ \text{corr} \left(x_S^{(S-\Delta t)}, y_{S-\Delta t}^{(S-\Delta t)} \right) \right\} \quad (6.1)$$

Here we denote by $x_t^{(l)}$ the vector in time of length l , namely $\{x_t, x_{t-1}, \dots, x_{t-l+1}\}$. In practice, the cross-correlation for a given shift decays rapidly for larger lag times. In this work we always use a value of $t_{max} = 60\text{ms}$.

The cross-correlation method (XC) has been used to measure causal influences in particular based on macro-scale recordings [89, 91, 92], but also in studies with single-cell precision based on fluorescence imaging data or multi-electrode recordings [38, 42, 103, 158, 174].

Later, Granger developed what became known as *Granger causality* (GC) based on the Wiener-Granger concept of causality as improved prediction [113]. GC has in recent years been applied neurosciences infer connectivity structures [175–177].

GC is defined as follows: Assuming that the present value of a time series x_t is determined by its last k time steps, this present value can be predicted up to an error η_t^0 .

$$x_t = \sum_{l=1}^k a_l^0 x_{t-l} + \eta_t^0. \quad (6.2)$$

This prediction can be compared to another linear forecasting model, namely one that takes into account the putative source time series y_t , resulting in a prediction error η_t^1 .

$$x_t = \sum_{l=1}^k a_l^1 x_{t-l} + \sum_{m=1}^k b_m^1 y_{t-m} + \eta_t^1 \quad (6.3)$$

Granger causality (GC) is defined as the log ratio between the variances of these prediction errors, as expressed by the covariance matrix Γ^0 for the univariate regression and Γ^1 for the bivariate case:

$$\text{GC}_{Y \rightarrow X} = \log \frac{(\Gamma^0)_{0,0} + (\Gamma^0)_{1,1}}{(\Gamma^1)_{0,0} + (\Gamma^1)_{1,1}}. \quad (6.4)$$

6.2 Non-linear measures

Since a linearity assumption limits the applicability of a measure, generalizations to the non-linear case were developed by making use of information theory. Generalizing the cross-correlation measure XC, the *Mutual Information* between two stochastic variables A and B is defined as the Kullback-Leibler distance between the joint probability distribution $P(A, B)$ and the product distribution $P(A)P(B)$ [178]. The two distributions are identical, and thus the distance vanishes, only if the two variables are independent of each other. Thus, Mutual Information (MI) measures the degree to which they are dependent on each other. For two time series x_t and y_t , the MI score is defined, analogously to XC, as the maximum found for a range of time lags.

$$\text{MI}_{Y \rightarrow X} = \max_{\Delta t=1 \dots t_{max}} \left\{ \sum P(x_n, y_{n-\Delta t}) \log_2 \frac{P(x_n, y_{n-\Delta t})}{P(x_n)P(y_{n-\Delta t})} \right\} \quad (6.5)$$

The sum is defined over all possible values of x_t and y_t , with t_{max} defined as above for XC. The number of all possible values is necessarily finite as the plug-in estimation of probabilities is used, counting frequencies in a number of $b = 20$ bins in each dimension. Note that the choice of this particular bin number is motivated by the low dimensionality of the joint probability distribution $P(x_n, y_{n-\Delta t})$, which allows for a relatively high bin number. We have however checked for the robustness of the MI score regarding this number, finding no significant change for the range 20 ± 5 . Details on the estimation of probability densities and alternative estimation methods are presented in Appendix B. The logarithm with the base of 2 ensures that the unit of MI is bits.

Just as the linear measures described above, MI was applied in numerous studies to infer connectivity in neuronal systems [159, 179, 180], but also in the context of genetic networks [181].

In 2000, Schreiber published a new measure of causality as a generalization of Granger causality called *Transfer Entropy* (TE) [116, 182]. TE has been used to infer connections in gene regulatory networks [183] as well as in neuronal systems based on spike times [184–186] and on multi-unit signals of neuronal activity [187–189].

If x_t and y_t are stationary Markov processes of some finite length k , such that $P(x_{t+1}|x_t^{(k+1)}) = P(x_{t+1}|x_t^{(k)})$, then we can calculate the Kullback-Leibler distance between two models of the process X , namely one based on its past alone $P(x_{t+1}|x_t^{(k)})$ and one including the past of the process Y in the model, $P(x_{t+1}|x_t^{(k)}, y_t^{(k)})$.

$$\text{TE}_{Y \rightarrow X} = \sum P(x_{t+1}, x_t^{(k)}, y_t^{(k)}) \log_2 \frac{P(x_{t+1}|x_t^{(k)}, y_t^{(k)})}{P(x_{t+1}|x_t^{(k)})} \quad (6.6)$$

The sum is defined over all possible values of x_{t+1} , $x_t^{(k)}$ and $y_t^{(k)}$. Binning is again used to estimate the transition probabilities from time series, but as the dimensionality is higher especially for Markov orders $k > 1$, we use $b = 3$ bins. An estimation with low bin number can actually be

seen as a symbolic representation of the fluorescence time series, which has been shown to be an effective strategy to estimate TE [190]. For a discussion of alternative strategies to estimate conditional probability densities without binning, see Appendix B.

It is easy to show that TE is equal to the difference of two Mutual Information terms by expanding the conditional probabilities. Here, we have denoted with $I(A, B) = H(A) - H(A|B)$ the Mutual Information between time series A and B , without varying the lag as done above in Eq. 6.5.

$$\begin{aligned} \text{TE}_{Y \rightarrow X} &= \sum P(x_{t+1}, x_t^{(k)}, y_t^{(k)}) \log_2 \frac{P(x_{n+1}, x_t^{(k)}, y_t^{(k)}) P(x_t^{(k)}) P(x_{t+1})}{P(x_t^{(k)}, y_t^{(k)}) P(x_{n+1}, x_t^{(k)}) P(x_{t+1})} \\ &= I(x_{t+1}; x_t^{(k)}, y_t^{(k)}) - I(x_{t+1}; x_t^{(k)}) \end{aligned} \quad (6.7)$$

Interestingly, Barnett et al. showed recently that TE and Granger causality are indeed equivalent for Gaussian random variables [191].

In summary, TE is advantageous to other measures of causality as it relies on the established notion of Wiener-Granger causality, but extends to a non-linear regime. It requires — apart from the very weak assumption of a Markov process of some order k , an upper bound of which can easily be estimated from the expected leak time constants — no assumptions about the dynamics of calcium fluorescence or of neuronal firing.

6.3 Measuring reconstruction performance

Before elaborating further on solutions to the network reconstruction problem, it is important to define a measure of the quality of a reconstruction, i. e. to which extent the value a given causality measure — such as those described above — overlaps with the structural, mono-synaptic connectivity that we aim to identify.

The central question is whether a reconstructed link is indeed present in the ground truth topology (a *true positive*) or is spuriously introduced by way of an imperfect reconstruction procedure (a *false positive*). Directly related is the question of the correct or incorrect non-existence of links which are absent from the reconstruction (*true negatives* and *false negatives*, respectively). This of course assumes that the ground truth is known, limiting the analysis to the extent of simulated data.

Such an overlap between true and reconstructed variables is traditionally and conveniently quantified with a *Receiver operating characteristic* (ROC) [178, 192, 193]. Let Θ_{rec} be an arbitrary threshold, then we define the reconstructed adjacency matrix as the thresholded matrix of the measure of causality for each pair of links. Formally, for two neurons i and j , denoting by Z_{ij} the value of the causality measure in question (for the link $i \rightarrow j$), the corresponding entry of the binary, reconstructed adjacency matrix is $(A_{\text{rec}})_{ij} = H(Z_{ij} - \Theta_{\text{rec}})$. Thus, for each value of the threshold Θ_{rec} we obtain a binary, reconstructed adjacency matrix A_{rec} , which we compare to the ground-truth adjacency matrix A to obtain, for instance, the fraction of true positives TP and the fraction of false positives FP.

$$\text{TP} = P(t = 1 | r = 1) \quad (6.8)$$

$$\text{FP} = P(t = 0 | r = 1) \quad (6.9)$$

Here we have denoted with t the true value of a given link, and with r its reconstructed counterpart. Analogously, we define the fraction of true negatives TN and of false negatives FN.

$$\text{TN} = P(t = 0|r = 0) \quad (6.10)$$

$$\text{FN} = P(t = 1|r = 0) \quad (6.11)$$

An ROC curve is then given by plotting TP on the vertical axis against FP on the horizontal axis, for the entire range of possible values of the threshold Θ_{rec} . The result is a curve from the point (0,0) for thresholds above the maximum of Z to the point (1,1) for thresholds below the minimum of Z . Values on the diagonal correspond to the characteristic of a random choice, with ROCs above the diagonal corresponding to reconstructions which are better than random.

To be able to quantify the reconstruction quality across networks in a single number, we define the fraction of true positives at a value of 10% of false positives as the *reconstruction performance*.

6.4 Extending Transfer Entropy

It turns out that applying TE directly to the analysis of calcium fluorescence time series yields unsatisfactory results, as will be shown later. There are two general properties of the dynamical system under study which need to be incorporated. This will be the purpose of the two following sections.

6.4.1 Same-bin interactions

As described above, calcium fluorescence imaging allows for the recording of hundreds or thousands of neurons simultaneously. A downside of this technique, however, is that in practice the rate of image acquisition cannot be increased arbitrarily. While the current generation of cameras is able to record in excess of 100 frames per second (fps), their efficiency of recording is limited such that an ever greater light intensity is needed to get images of reasonable signal-to-noise ratio. Yet as the light needed for fluorescent dyes is from the UV spectrum, increased light intensity leads to bleaching, i. e. neurons are damaged up to the point where the maximum recording duration in high frame rate two-photon recording can be below one minute [131, 194–196].

In the experiments shown in Fig. 5.1, a recording rate of 50Hz was used, corresponding to the typical rate of recording in such experiments [24, 106, 162, 171, 197]. This allows for continuous recording of 1 hour without any signs of bleaching or altered firing dynamics (see Fig. A.1).

Therefore, each image frame in the recording has a temporal width of 20ms, much larger than the synaptic delay time, which can be faster than 2ms [198–200]. Recognizing this fact, we define a first step towards a new causality measure by allowing for instantaneous (*same-bin*) interactions between the Markov processes by modifying Eq. 6.6.

$$\text{TE}_{Y \rightarrow X}^{\text{same-bin}} = \sum P(\Delta x_{t+1}, \Delta x_t^{(k)}, \Delta y_{t+1}^{(k)}) \log_2 \frac{P(\Delta x_{n+1} | \Delta x_t^{(k)}, \Delta y_{t+1}^{(k)})}{P(\Delta x_{n+1} | \Delta x_t^{(k)})} \quad (6.12)$$

Note that the first index of the process Y is now $t + 1$. This allows to capture interactions on a time scale shorter than the sampling width, thereby increasing dramatically the information gained for a given number of samples (examples of reconstruction performance, comparing TE and $\text{TE}_{\text{same-bin}}$, will be shown in Section 7.5). As described above (see Section 5.2), TE is computed based on the difference signal Δx_t , as its distribution is more Gaussian-like, which improves sampling for a given number of bins.

6.4.2 Identification of the dynamical state

One of the assumptions behind Transfer Entropy is the stationarity of the time series in question. Due to the switching behavior seen in neuronal networks *in vitro*, this assumption is valid only to some degree — the time series is at least approximately stationary in piece-wise manner, between and within bursts. Across these different phases, the dynamical nature of activity will be different — varying for instance from synchronous to asynchronous firing times — and this dynamical nature will be reflected in different networks of effective connectivity.

We can identify this dynamical nature in the histogram of the population mean of the fluorescence (see Fig. 6.1A), showing a peak on the left corresponding to low average fluorescence caused by low-frequency, asynchronous firing, and a long tail on the right corresponding to ongoing network bursts and high values of the average fluorescence. We denote this average fluorescence by g_t , based on the fluorescence time series of individual neurons $x_{i,t}$:

$$g_t = \frac{1}{N} \sum_{i=1}^N x_{i,t}. \quad (6.13)$$

To calculate the effective connectivity for each of the different parts of these dynamical states separately, we define as $\text{TE}^{\text{state-dep.}}$ the effective connectivity based on Eq. 6.12 for all data points where the mean fluorescence g_t is within a given range (indicated in Fig. 6.1 by roman numerals) $R_i = [g^{\text{low}}, g^{\text{high}}]$, $i = \{\text{I, II, ..., VII}\}$. Effectively, this means that we calculate the effective connectivity $\text{TE}^{\text{state-dep.}}$ based on the sets of points in time $T = \{t : g_i^{\text{low}} \leq g_t \leq g_i^{\text{high}}\}$.

$$\text{TE}_{Y \rightarrow X}^{\text{state-dep.}}(R) = \sum P(\Delta x_{t+1}, \Delta x_t^{(k)}, \Delta y_{t+1}^{(k)} | g_t \in R) \log_2 \frac{P(\Delta x_{n+1} | \Delta x_t^{(k)}, \Delta y_{t+1}^{(k)}, g_t \in R)}{P(\Delta x_{n+1} | \Delta x_t^{(k)}, g_t \in R)} \quad (6.14)$$

Interestingly, the reconstruction performance as defined in Section 6.3 depends heavily on the range which is used to determine the effective connectivities (see Fig. 6.1B, blue curve). The reconstruction performance shows a very high value of 70% of true positives at 10% of true positives for the lowest region I. For the ranges II to IV, the performance is acceptable at about 40% TP, while decreasing back to chance level (10% TP) for the highest ranges.

However these estimates are based on sample numbers which differ by large amounts — by almost two orders of magnitude when comparing range VII to range I — which makes the comparison of reconstruction performance biased since the transition probabilities underlying TE can be better estimated based on more data points. When correcting for this bias by restricting the calculation in each region to be based on an identical number of data points, a different picture emerges (see Fig. 6.1B, red curve). The reconstruction performance of range I is now poor at about 20% TP, while the range with the highest reconstruction performance is now range II, shifting the maximum towards higher values of the mean fluorescence. The higher regions have generally slightly lower performance as in the case of full numbers. Note that the performance generally has to be lower, as a lower number of samples was used in all regions that was lower (or equal in the case of region VII) than before.

The actual reconstructed networks (see Fig. 6.1C, plotted with the restriction to equal sample numbers) confirm the intuition that the effective connectivity depends on the dynamical state of the network as revealed by the average fluorescence. The lowest range I corresponds to a noise-dominated part of the dynamics, where the stochastic firing is not sufficient to drive spiking activity that is informative about the topology in the fluorescence time series. On the other end of the spectrum, the high ranges VI and VII are dominated by synchronous groups of neurons firing together within a network burst (which is due to a higher cross-correlation of the fluorescence signals within the highly inter-connected groups visible in the network reconstructions as dense

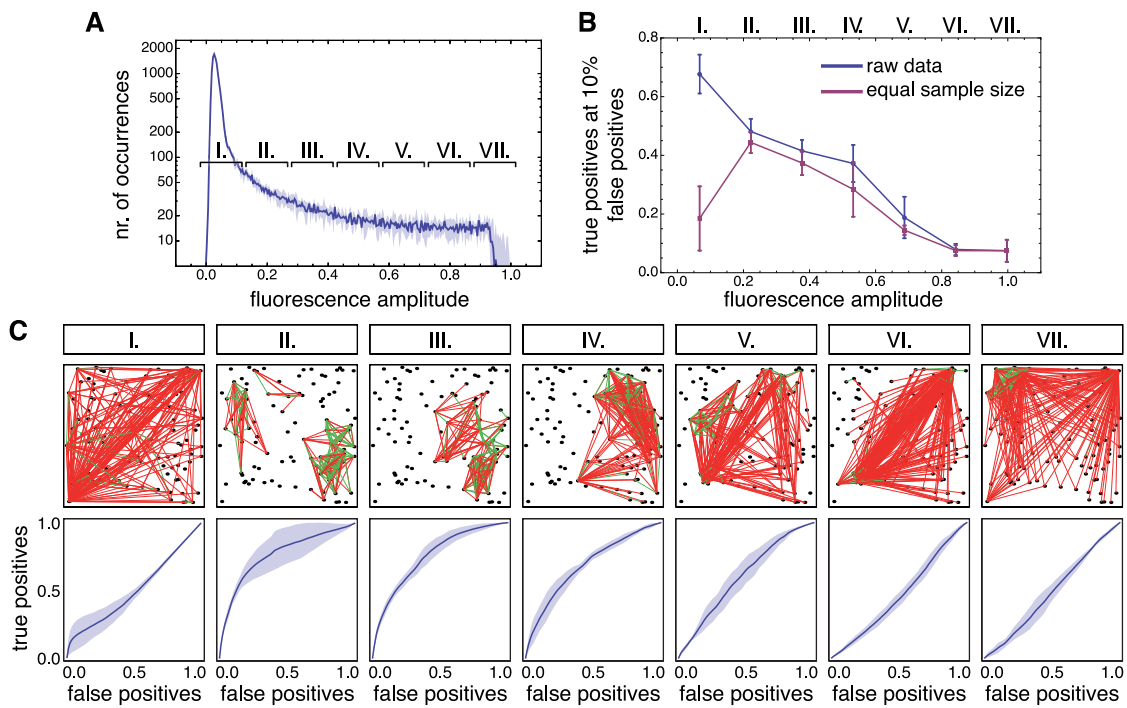


Figure 6.1: Effective connectivity is heavily dependent on the dynamical state of the network. (A) The histogram of the population average of the (simulated) fluorescence signal g_t across 60 minutes of simulated real time, representing the switching behavior from low-frequency firing activity (region I) to bursting activity (region VII). (B) For data points corresponding to each of the seven regions indicated on top in roman numerals, the reconstruction performance in terms of the fraction of true positives at 10% of false positives is shown in blue. The red line shows the same performance, but restricting the number of data points in each region to the number of data points in region VII, thereby using an equal number of data points for each region. (C) For each of the latter reconstructions based on an equal number of data points, the links corresponding to the top 10% of TE values of an example network is shown, plotting false positives in red and true positives in green (top row). Below are plotted, again for each region, the ROC curves across six networks. In all panels, shaded areas and error bars correspond to 95% confidence intervals across six networks. Figure reproduced with permission from [169].

communities, not shown). Only in the middle ranges, and especially range II, is the spiking activity in the system of a frequency high enough to be dominated by mono-synaptic interactions (rather than isolated firing events driven by noise), and it is in these ranges that indeed a reconstruction with a good overlap between effective and structural connectivity is achieved.

6.4.3 Generalized Transfer Entropy

Transfer Entropy is a measure of causality which has been used with much success for the analysis of network connectivity. For the present application, we have described an extension to same-bin interactions in Eq. 6.12 and have discussed the state-dependency according to Eq. 6.14 above. As a result, we combine the individual advantages of the lower ranges of the average fluorescence — corresponding to dynamical ranges of asynchronous activity of low to medium firing rates — by considering in our analysis all data points when the average fluorescence is below a *conditioning level* \tilde{g} .

We therefore defined in [169] the *Generalized Transfer Entropy* (GTE) of Markov order k :

$$\text{GTE}_{Y \rightarrow X} = \sum P(\Delta x_{t+1}, \Delta x_t^{(k)}, \Delta y_{t+1}^{(k)} | g_t < \tilde{g}) \log_2 \frac{P(\Delta x_{n+1} | \Delta x_t^{(k)}, \Delta y_{t+1}^{(k)}, g_t < \tilde{g})}{P(\Delta x_{n+1} | \Delta x_t^{(k)}, g_t < \tilde{g})}. \quad (6.15)$$

The conditioning level is consequently chosen, in the following sections, to be on the right of the Gaussian-like peak in the fluorescence histogram by one standard deviation, where the histogram deviates from a Gaussian fit. The reconstruction is however robust against the precise choice of this conditioning level, as will be shown later in Section 7.3.

6.4.4 Causality measures in comparison to GTE

One of the advantages of the improvements to TE we developed above is that they are general enough to be applied to other measures of causality as well. Including these improvements in all cases in the following chapters allows for the unbiased comparison of different measures.

In the case of linear measures, the conditioning is applied by considering only the subset of data points where the average fluorescence is below the conditioning level \tilde{g} , i. e. all time points $T = \{t : g_t < \tilde{g}\}$ such that the signal is given by $\tilde{x} = \{x_t : t \in T\}$. The time series $\Delta \tilde{x}$ then refers, as before, to the difference signal of \tilde{x} .

$$\text{XC}_{Y \rightarrow X} = \max_{\Delta t=0 \dots t_{max}} \left\{ \text{corr} \left(\Delta \tilde{x}_S^{(S-\Delta t)}, \Delta \tilde{x}_{S-\Delta t}^{(S-\Delta t)} \right) \right\} \quad (6.16)$$

Granger Causality is defined similarly to the literature notation defined in [113], note only the use of the conditioned time series \tilde{x}_t and \tilde{y}_t and that the index of the first term in the sum over \tilde{y}_t is zero. Apart from these modifications, GC is defined as described in detail in Section 6.1.

$$\begin{aligned} \Delta \tilde{x}_t &= \sum_{l=1}^k a_l^0 \Delta \tilde{x}_{t-l} + \eta_t^0 \\ \Delta \tilde{x}_t &= \sum_{l=1}^k a_l^1 \Delta \tilde{x}_{t-l} + \sum_{m=0}^{k-1} b_m^1 \Delta \tilde{y}_{t-m} + \eta_t^1 \\ \Rightarrow \text{GC}_{Y \rightarrow X} &= \log \frac{(\Gamma^0)_{0,0} + (\Gamma^0)_{1,1}}{(\Gamma^1)_{0,0} + (\Gamma^1)_{1,1}} \end{aligned} \quad (6.17)$$

Mutual Information is defined in the language of probability theory, therefore we can apply the extension to the difference signal and conditioning in the same notation as it was originally developed for GTE.

$$\text{MI}_{Y \rightarrow X} = \max_{\Delta t=0 \dots t_{max}} \left\{ \sum P(\Delta x_n, \Delta y_{n-\Delta t} | g_t < \tilde{g}) \log_2 \frac{P(\Delta x_n, \Delta y_{n-\Delta t} | g_t < \tilde{g})}{P(\Delta x_n | g_t < \tilde{g}) P(\Delta y_{n-\Delta t} | g_t < \tilde{g})} \right\} \quad (6.18)$$

The thus modified causality measures allow for a fair comparison between GTE and other methods when studying their respective reconstruction performance.

Chapter 7

GTE reconstruction of simulated synaptic connectivity

Having developed an appropriate measure of causality which is applicable to the system of calcium fluorescence time series of neuronal networks, we will now analyze and quantify its performance under various experimentally relevant conditions, focussing as a first step on the reconstruction of excitatory connections only. Next to the goal of reconstructing a large fraction of true positives at a small fraction of false positives, a question of particular importance is whether it is possible — with GTE as described in the previous chapter, or with other measures of effective connectivity — to reconstruct the statistics of the topological properties in the network.

7.1 GTE performance in clustered networks

7.1.1 Non-locally clustered networks

In non-locally clustered networks as described in Section 5.2, GTE achieves a very good reconstruction performance, displaying a sharp rise in the ROC curve for high thresholds (see Fig. 7.1A). The performance level, as given by the fraction of TP at a fraction of 10% of FP, is 75%, yielding true positives with excellent accuracy. This performance is indeed even better than other studies based on TE using precise spike times [185]. Only for lower thresholds and therefore higher numbers of reconstructed links does the fraction of false positives start to increase as well, leading to a drop in ROC slope. This performance is conserved across networks, as demonstrated by the small confidence intervals around the mean.

Indeed in these networks, not only is the reconstruction performance very good, but the reconstructed network matches the ground truth one very well in terms of the distributions of topological indices) — as will be shown later, especially in the case of linear measures, the former is not a sufficient condition for the latter. The two distributions for the clustering coefficient are very similar in terms of mean and variance, even if a small underestimation of CC is visible. The same is true for the distribution of in-degrees. Note the noise in the histograms due to the fact that there is only one data point per neuron in these cases. There is also a slight under-estimation of the average spatial distance of the connections in the network (see Fig. 7.1B, bottom panel). Indeed there is mostly an over-estimation of the frequency of low-distance links within a distance interval indicating that this is an effect caused by light-scattering artifacts ($\lambda_{sc} = 0.15\text{mm}$, see Section 5.2).

Importantly, GTE is able to reconstruct the clustering index of networks engineered to display a wide range of CC values with good accuracy (see Fig. 7.1C), as quantified by a Pearson correlation

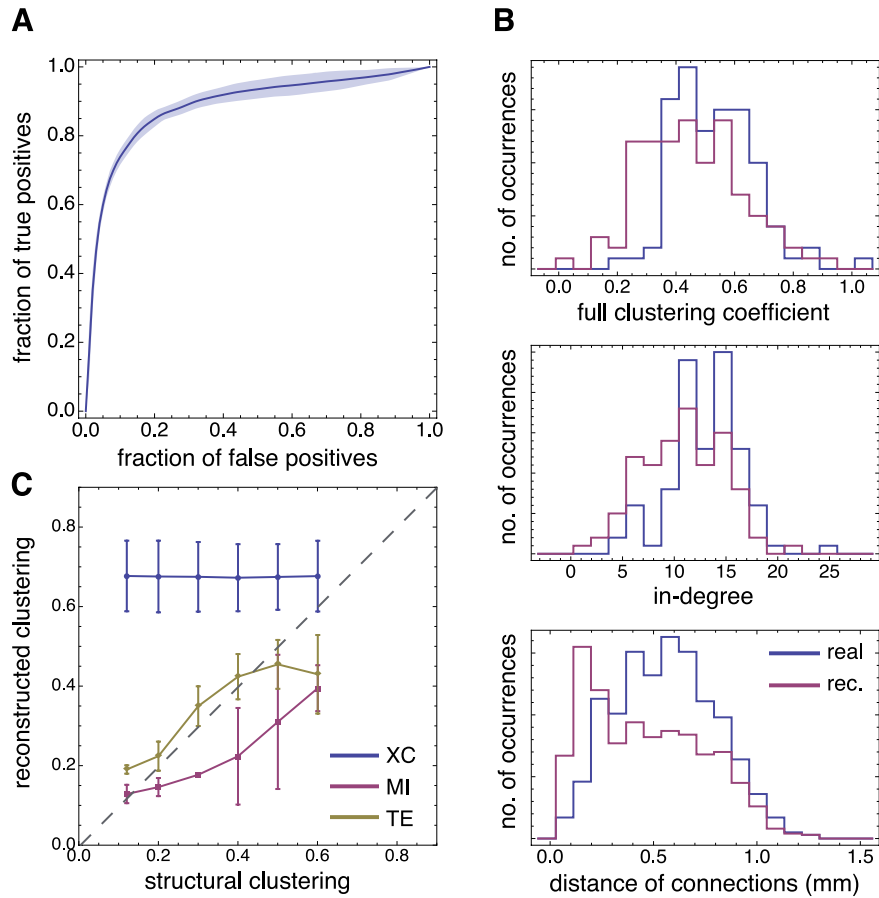


Figure 7.1: Reconstruction performance in non-locally clustered topologies. (A) Receiver-operating characteristic (ROC) of the GTE reconstruction based on simulated fluorescent time series from non-locally clustered networks. Shaded areas correspond to 95% confidence intervals across six networks. (B) In one example network, the real (blue curve) and reconstructed (red) distributions of topological properties of the network are shown, namely the full clustering coefficient CC (top panel), the in-degree (middle) and the spatial distance between connected neurons (bottom panel). Reconstructed network based on the top 10% of GTE values. (C) Correlation between real and reconstructed clustering coefficients, for a range of networks engineered to display a given clustering coefficient between 0.1 and 0.6. Lines correspond to different measures of causality, namely XC (blue curve), MI (red curve) and GTE (yellow curve). Error bars correspond to 95% confidence intervals across six networks, and the dashed line marks the identity (perfect result). As before, reconstructed networks are based on a network of $N = 100$ nodes with a clustering coefficient of $CC = 0.5$, with the exception of panel C. The conditioning level used was $\tilde{g} = 0.112$. Figure reproduced with permission from [169].

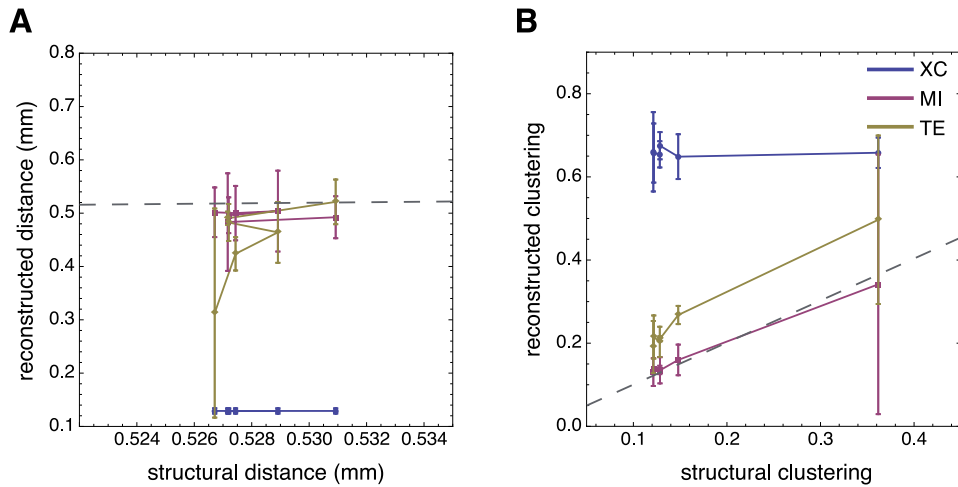


Figure 7.2: Reconstruction quality of the complementary topological properties, i. e. (A) the clustering index of networks generated according to a given length scale and (B) the average connection distance in networks engineered to display a given value of CC. Plots are based on the top 10% of links, with conditioning levels of $\tilde{g} = 0.084$ for the non-locally clustered networks of panel A, and of $\tilde{g} = 0.112$ for the local networks of panel B. Lines correspond to different measures of causality, namely XC (blue curve), MI (red curve) and GTE (yellow curve). Error bars correspond to 95% confidence intervals across six networks, and the dashed line marks the identity (perfect result). Note that as the value of the horizontal axis is first averaged across networks and then sorted according to the ground truth value, the linear interpolation indicated by the lines need not be single-valued. Figure reproduced with permission from [169].

coefficient of $r = 0.92$. The linear measure reconstructs clustered networks with $CC \approx 0.7$ throughout, and therefore cannot be used to determine the clustering level from the fluorescence signal. The non-linear measure MI does show a clear correlation of reconstructed to ground truth CC, but underestimates the true value for higher CC values.

To complete this analysis, in Fig. 7.2A the correlation of the reconstructed average connection length is plotted against the structural value. For non-locally clustered topologies there is no connection between connection probability and spatial distance, and is therefore given purely by the finite-size fluctuations of different realizations of the random graph. Nevertheless, it is important to note that the non-linear measures MI and GTE do provide an estimate of the average connection distance with reasonable accuracy. Note also that the reconstructed length scale is dramatically underestimated using the XC measure as before.

7.1.2 Locally clustered networks

For the case of locally clustered networks, i. e. networks created with a connection probability which depends on the spatial distance as described in Section 2.3, reconstruction quality displays a broader range of values — due to the finite size of the network relative to the length scale λ — while still displaying a high average performance of 60% TP at 10% FP (see Fig. 7.3A). In local networks, this reconstruction performance can even be assessed visually: Fig. 7.3B shows an example of a reconstructed network in relation to the structural ground-truth network. Many links are correctly identified by GTE, with a small fraction of spurious long-range links and a bias towards more

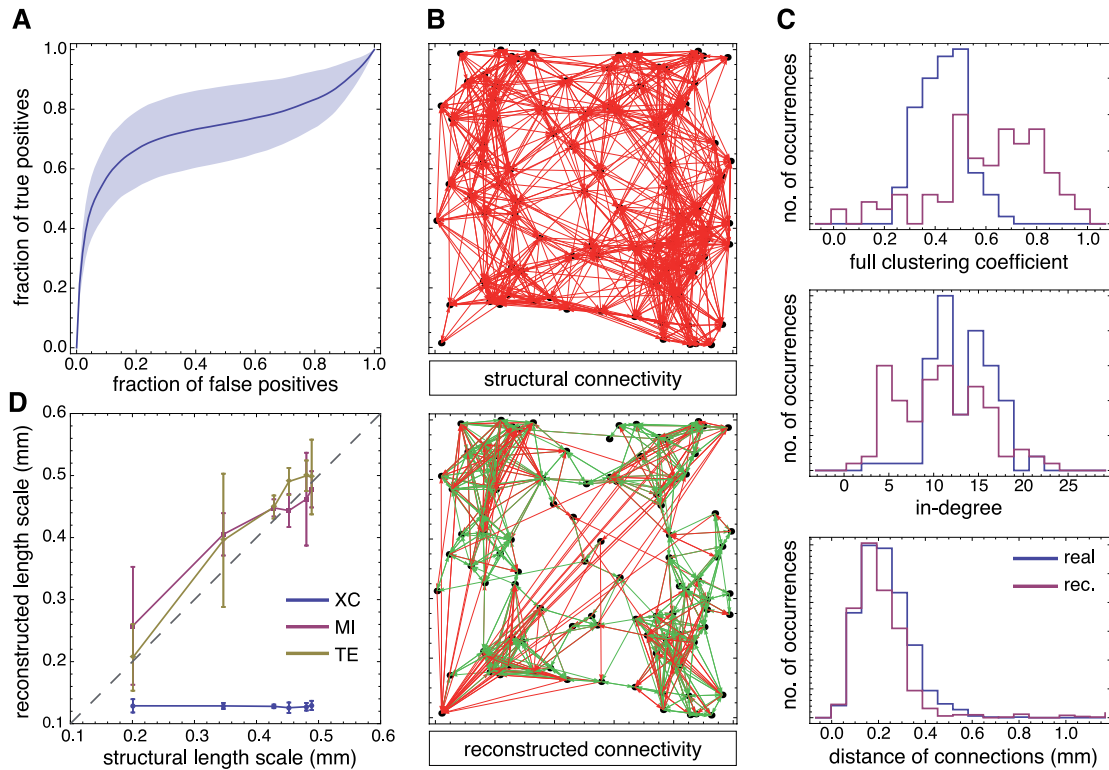


Figure 7.3: Reconstruction performance in locally clustered topologies. (A) ROC curve of the GTE reconstruction based on simulated fluorescent time series from locally clustered networks. Shaded areas correspond to 95% confidence intervals across six networks. (B) Network graph of the ground truth network with links shown in red (top panel) and the reconstructed network (based on the top 10% of GTE values) with true positives shown in green and false positives shown in red (bottom panel). (C) In one example network, the real (blue curve) and reconstructed (red) distributions of topological properties of the network are shown, namely the full clustering coefficient CC (top panel), the in-degree (middle) and the spatial distance between connected neurons (bottom panel). Reconstructed network again based on the top 10% of GTE values. (D) Correlation between real and reconstructed average connection lengths, in terms of spatial distance, for a range of networks engineered according to a Gaussian profile of the probability of connection with length scale λ . Lines correspond to different measures of causality, namely XC (blue curve), MI (red curve) and GTE (yellow curve). Error bars correspond to 95% confidence intervals across six networks, and the dashed line marks the identity (perfect result). As before, reconstructed networks are based on the top 10% of links. Shown are results based on a network of $N = 100$ nodes with a length scale of $\lambda = 0.25\text{mm}$, with the exception of panel C. The conditioning level used was $\tilde{g} = 0.112$. Figure reproduced with permission from [169].

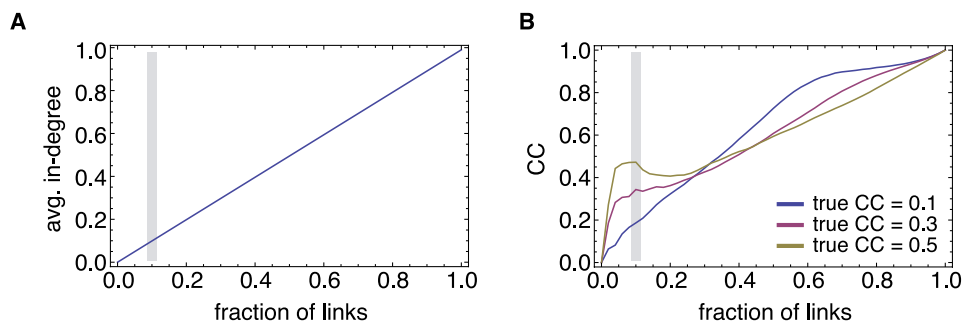


Figure 7.4: Dependence of reconstructed average degree and clustering coefficient on the threshold Θ_{rec} . (A) Plotted is the average degree of a reconstruction based on a threshold Θ_{rec} such that a given fraction of links reconstructed (horizontal axis). (B) For three example networks where the ground truth topology had a given value of CC, the reconstructed value of CC is plotted. In both panels, the light grey bar corresponds to the fraction of links present in the ground truth topology, and typically chosen for reconstructions in the simulated networks (see text). Figure reproduced with permission from [169].

tightly clustered groups of links.

The reconstruction performance of estimating topological quantities is plotted in Fig. 7.3C. The clustering bias visible in the network graph is indeed present, yielding an average CC of 0.59 instead of 0.41 for the structural network. The degree distribution is broadened, and the distribution shows a small bias towards lower connection distances. This bias is, however, not an effect consistently present across realizations of networks (see Fig. 7.3D). Instead, on average there is a very good, linear correlation (Pearson’s correlation of $r = 0.97$) between real and reconstructed length scales for both MI and GTE. Only XC shows a reconstructed length scale which is constant and irrespective of the average connection distance of the ground truth network, in analogy to the case of locally clustered networks in Fig. 7.1C.

A lower value of the connection length scale λ in a network naturally yields a higher value of the clustering index CC. It is therefore important that GTE shows a good linear correlation between real and reconstructed CC with a correlation coefficient of $r = 0.98$ (see Fig. 7.2B). Interestingly, in this situation MI actually performs better than GTE, despite having a lower reconstruction performance overall. Note also that the linear measure again shows a reconstructed value that does not depend on the ground truth topology.

7.2 Dependence of reconstructed quantities on threshold

It is important to note that the reconstructed quantities do of course depend on the chosen threshold Θ_{rec} . As the threshold directly determines the number of links in the network, the relation between average degree of the reconstruction and fraction of links reconstructed is by construction linear (see Fig. 7.4A).

The dependence of the reconstructed clustering coefficient on Θ_{rec} is less trivial. For a random graph, the reconstructed clustering coefficient is by definition (see Eq. 2.7) given by the fraction of links itself, because in a random network the latter is identical to the probability of closing a given loop, and thus on the probability of connection. The deviations from the diagonal therefore correspond to deviations of the reconstructed network from the random graph. The reconstructed

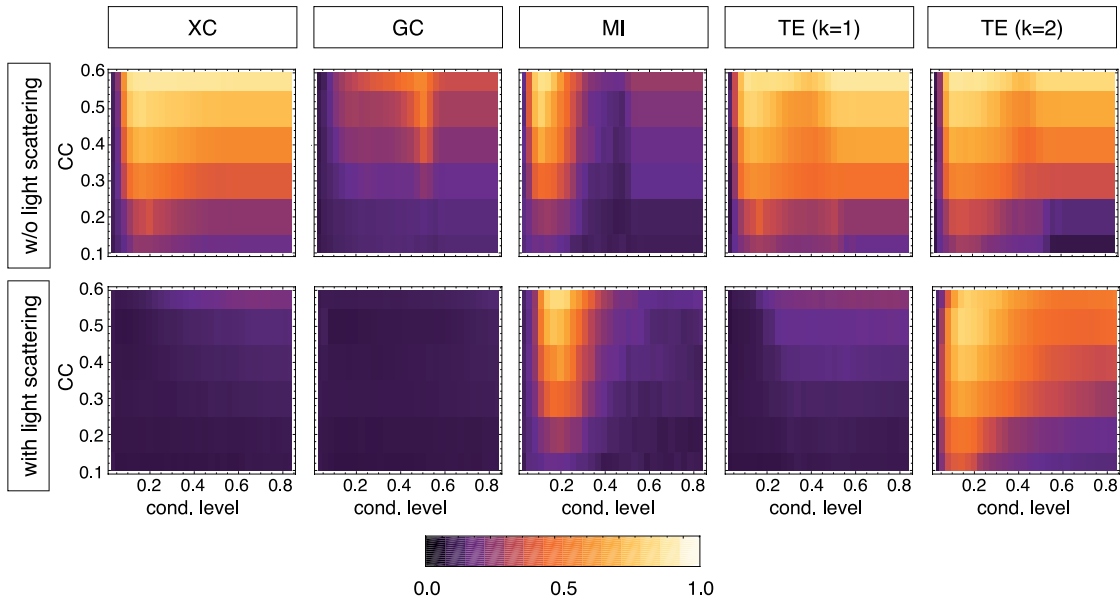


Figure 7.5: Reconstruction performance across causality measures and conditioning levels in the non-locally clustered ensemble. Shown are reconstruction based on simulation without light scattering (top row) and with light scattering (bottom row). Each column corresponds to a different causality measure, namely (from left to right) XC, GC, MI and GTE of Markov order 1 and of Markov order 2. In each panel, the average reconstruction performance across three networks is plotted as color code, for a given conditioning level \tilde{g} (horizontal axis) and a certain clustering index of the underlying networks (vertical axis). Figure reproduced with permission from [169].

value of CC, for higher values of the true CC, shows either a peak or plateau. Interestingly and usefully for the application discussed here, there is a good agreement between real and reconstructed CC between a fraction of positives between 5% and 10% (perhaps not too surprisingly, given the good reconstruction performance described above). This corresponds well to the fraction of links actually chosen in panels B and C of Fig. 7.1, motivated by the expected number of ground-truth links.

7.3 Dependence on causality measure and conditioning

GTE is able to quantitatively extract important features of the network topology across a wide range of networks, as demonstrated by Figures 7.1C and 7.3D and by Fig. 7.3 and as discussed above. In this section the detailed dependence on the value of the conditioning level, the causality measure used and the existence of light scattering artifacts will be discussed. The latter is conceptually important, because it introduces an experimentally unavoidable artifact which is small in amplitude, but linear in nature. Such artifacts are common in a range of recordings and causality analysis, because some form of cross-talk between time series (recording sites) is often present [189,201,202] and may, due to the inherent additivity in terms of signal amplitude, prove hazardous for causality measures relying on the assumed linearity of the interaction.

7.3.1 Non-locally clustered networks

In networks from the non-locally clustered network ensemble and without the light scattering artifact, a very good reconstruction performance is achieved across the higher values of CC and for a broad range of conditioning levels (see top row of Fig. 7.5). For lower values of CC, the range of conditioning levels with reasonable reconstruction gets smaller, showing an unsatisfactory reconstruction at $CC = 0.1$, corresponding approximately to the random graph. Interestingly, there is no correlation between reconstruction performance and reconstructed topological indices (compare Fig. 7.1C and Fig. 7.2A), indicating that the high-clustering topology is indeed an artifact of XC itself.

The GC reconstruction shows poor performance overall, with only a small range of conditioning levels with good reconstruction for higher CCs, and a poor performance for networks of low CCs, indicating the inadequacy of regression-based measures for the inherently non-linear neuronal interaction dynamics. It is then not surprising that MI is performing much better, with a range of conditioning levels for high values of CC that only diminishes for the smallest values of the clustering coefficient. GTE of Markov order $k = 1$, as well as for $k = 2$, show a very good reconstruction performance across conditioning levels. The reconstruction performance for the lowest value $CC = 0.1$ is comparable to the XC reconstruction.

If the light-scattering artifact — which is always present in real recordings — is included in the simulation of the fluorescence time series, a dramatically different picture emerges (see bottom row of Fig. 7.5). The reconstructions based on XC and GC now shows virtually no resemblance to the structural topology, reflected in a random performance across all values of CC and for all conditioning levels. This confirms the point made above in that the introduction of a weak but linear interaction term is enough to completely derail causality measures which assume linearity.

Consistently, then, is that the reconstruction performance based on MI is not dramatically changed, as it is a measure from information theory which is not based on a linearity assumption. Interestingly though, the reconstruction performance of GTE of Markov order $k = 1$ is dramatically reduced as well. Finally, the only measure for which a range of conditioning levels exists which yields consistent reconstruction performance across all values of CC, is GTE of Markov order $k = 2$, as indicated by the yellow stripe in the bottom right panel of Fig. 7.5 at $\tilde{g} \approx 0.1$.

It can be concluded that, in order to achieve a reconstruction that is robust across all values of the clustering index CC and including the experimentally realistic light scattering artifact, two ingredients (in addition to the improvements as described in Section 6.4) are needed: First, the measure of causality needs to be non-linear, and second, the measure must be able to take into account interactions on more than one time scale.

7.3.2 Locally clustered networks

For locally clustered networks, a similar picture emerges as in the case of non-locally clustered networks, even if the overall reconstruction quality is less optimal across all causality measures (see Fig. 7.6). This is due to the near-randomness of networks with characteristic length scales λ which are comparable or even larger than the system size $d = 1\text{mm}$.

Across causality measures, without the light-scattering artifact there exists a range of conditioning levels for which a good quality of reconstruction can be achieved for the lowest, most local value of $\lambda = 0.25$. For higher values of λ and therefore less locally structured networks, the reconstruction quality decays. This must indeed be that case as the networks are indeed identical random graphs for $CC = p$ and $\lambda \rightarrow \infty$.

With light scattering, the picture changes insofar as all measures — with the exception of GTE of higher Markov order — show a reconstruction quality which is very good for the most local

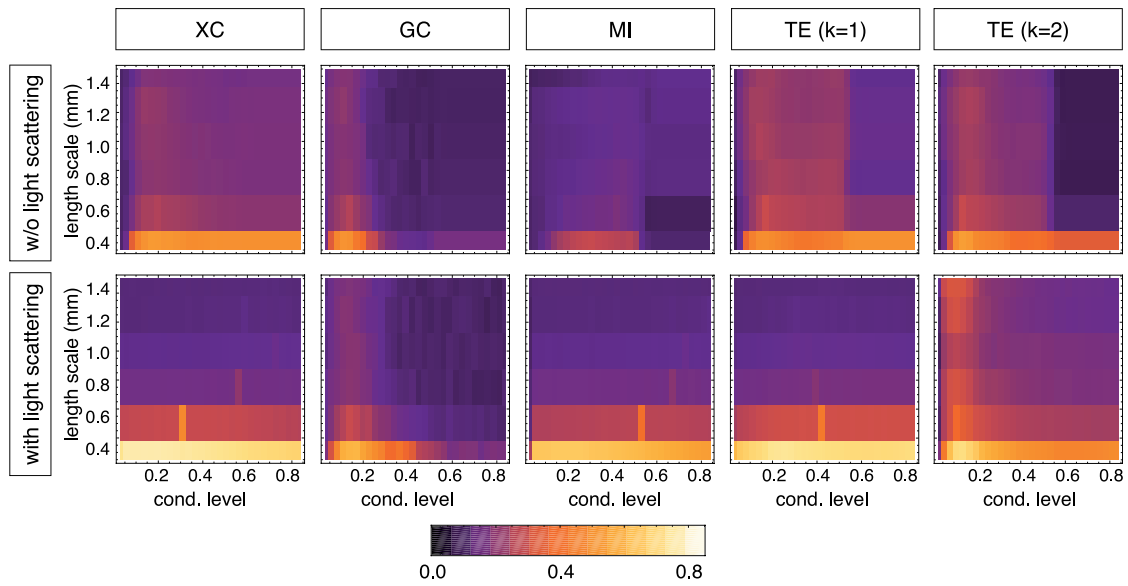


Figure 7.6: Reconstruction performance across causality measures and conditioning levels in the locally clustered ensemble. Shown are reconstruction based on simulation without light scattering (top row) and with light scattering (bottom row). Each column corresponds to a different causality measure, namely (from left to right) XC, GC, MI and GTE of Markov order $k = 1$ and for $k = 2$. In each panel, the average reconstruction performance across three networks is plotted as color code, for a given conditioning level \tilde{g} (horizontal axis) and a characteristic length scale λ of the underlying networks (vertical axis). Figure reproduced with permission from [169].

networks, but is not robust to an increase in λ . This good performance for $\lambda = 0.25$ is explained by an overlap of the structural connectivity and the erroneously inferred effective connectivity caused by light scattering artifacts ($\lambda_{sc} = 0.15$). The light scattering leads to an increase in spurious, local inferred links which are consistent with the structural connectivity only in this special case — for higher values of λ this overlap is no longer present and the reconstruction performance vanishes accordingly. Only GTE of Markov order $k = 2$ shows good reconstruction performance of about 40% of TP not only for the most local networks, but shows a range at around $\tilde{g} \approx 0.1$ where a reasonable reconstruction is possible across all networks.

Note that for both the non-locally clustered ensemble discussed above as well as for the locally clustered ensemble, the range of optimal conditioning levels is consistent with the logic developed in Section 6.4.2 as it is a value on the right side of the Gaussian-like side of the histogram, where the distribution deviates from a Gaussian form — see Fig. 6.1A for a histogram of fluorescence levels from the local clustering ensemble.

7.4 Comparison to Bayesian approach

Previously, the probably most established approach to reconstruct the connectivity of a neuronal assembly based on calcium imaging was developed by Mishchencko et al. [106]. It makes use of a sophisticated Bayesian inference algorithm to first detect precise spike times, and in a second stage to infer the parameters of an assumed generalized linear model (GLM). While for the case of simulations, based on a GLM and the same model of calcium fluorescence as used in the present work, the reconstruction performance has indeed been shown to be very good, even distinguishing excitatory and inhibitory connection with a good degree of accuracy [106].

The problem of Bayesian methods is founded in the reliance of the reconstruction algorithm on a specific mathematical model of neuronal dynamics and calcium fluorescence. The inference of the parameters of the (correct) model can indeed be highly accurate, but there is always the danger that the assumed dynamics are too restrictive for a data set from experimentally recorded networks. The problem is amplified when using an inference using multiple steps, as done in [106], where errors and biases can propagate to disrupt the final reconstruction result in a way that is extremely hard to disentangle a posteriori.

In an example case, we have applied the algorithm of [106] to a network from the locally clustered network ensemble; see Fig. 7.7A for a raster plot of spike times and Fig. 7.7B for a plot of the resulting fluorescence signals from three example neurons in the network. We find that the inferred spike times, using the algorithm of Mishchencko et al. as published in detail in [171], display a different pattern than the ground-truth spike times (see Fig. 7.7C). While the network bursts are correctly identified, indicated by many inferred action potentials around the time of bursting, the inferred rate of firing in between bursts is virtually zero. For one particular neuron, on the other hand, there is an homogeneously high firing rate overall, with a firing rate that exceeds the true rate by far.

These biases result in a effective connectivity, represented by the connection weights of the inferred GLM, that is dominated by these biases of spike detection. As a result, the overlap between the inferred weights and the structural topology is small, and the ROC curve is very close to the diagonal representing a random choice (see Fig. 7.7D).

7.5 Dependence on extension to Transfer Entropy

The previous sections have demonstrated the improvements of GTE as compared to traditional causality measures such as a cross-correlation based approach, Mutual Information or Granger

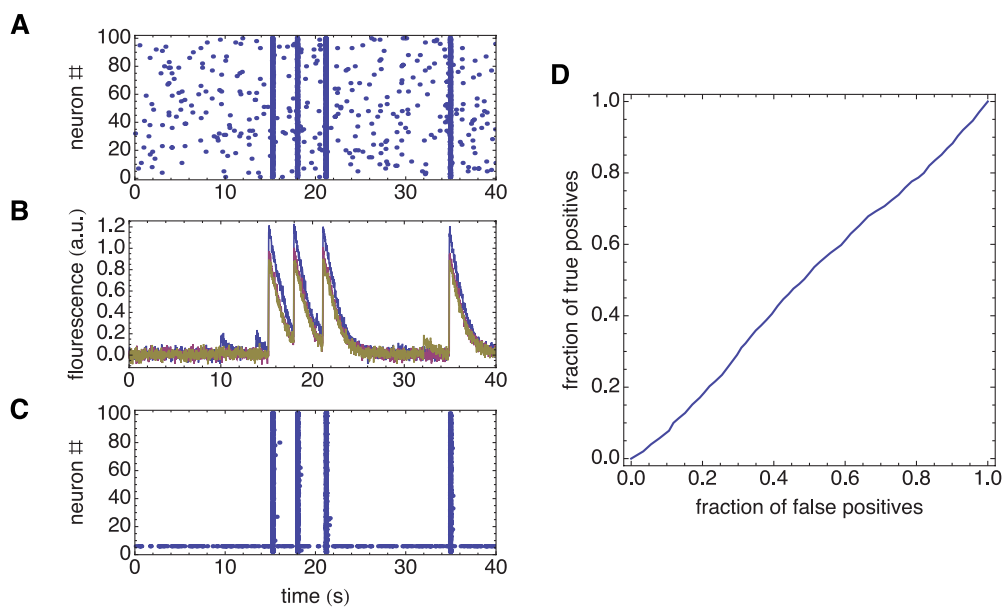


Figure 7.7: Reconstruction performance of the Bayesian approach by Mishchenko et al. in [106]. (A) Ground-truth spike times for a simulated example network from the locally clustered ensemble shown in Fig. 7.3. (B) For the same network, three example fluorescence signals showing bursts and sparse, non-bursting activity. (C) Reconstructed spike times according to [106]. (D) ROC curve of the reconstruction based on the reconstructed synaptic weights, showing a reconstruction quality that is close to random.

Causality. Here, the performance gain by each of these individual improvements to GTE will be quantified.

The reconstruction quality of raw TE (based on the difference signal) is, for both network ensembles, below the diagonal in the ROC plot, indicating a reconstruction quality which is worse than random choice (see Fig. 7.8A). This lack in performance is due to misinterpreted, spurious interactions because TE ignores the time delay which is most informative and which in the simulated recordings belongs to a delay of zero image frames. Consequently, after introducing instantaneous interactions, the performance level dramatically increases by over 50% of true positives (at 10% of false positives). The effect of only conditioning TE, again neglecting instantaneous interactions, results in an increase of about 40% TP in the non-locally clustered ensemble, and of about 20% TP in the locally clustered ensemble. Only when combining the two improvements to GTE as defined in Eq. 6.15 can a performance level of over 50% TP be achieved in both ensembles.

An alternative way of quantifying the quality of a reconstruction is by plotting *Positive prediction curves* (PPCs) [185]. For these plots, denoting with #TP the (absolute) number of true positives and by #FP the number of false positives in a given reconstruction, Garofalo et al. define the *true-false ratio* TFR as the fractional difference of the two.

$$\text{TFR} = \frac{\#TP - \#FP}{\#TP + \#FP} \quad (7.1)$$

TFR therefore is a measure of how much more likely it is for a randomly chosen link from the reconstruction to be a true positive, rather than a false positive. Note that the denominator $\#TP + \#FP$ is simply the total number of reconstructed links.

Fig. 7.9 displays the PPCs for the reconstructions shown in Fig. 7.8A. Apart from the raw TE reconstruction, the TFR is positive for a range of the total number of reconstructed links. Note that this already constitutes a marked improvement over the results of [185]. In particular, the reconstructions based on $\text{TE}^{\text{same-bin}}$ shows good performance. GTE itself shows a peak at a very good TFR of about 0.5, and stays positive until about 1,500 links (15%), highlighting the fact that true positives outnumber false positives for a broad range of the selected threshold Θ_{rec} . For high values of the number of reconstructed links, the TFR by construction converges for any measure to the rate $\text{TFR}_{\infty} = 2p - 1$, which for the networks studied here is equal to -0.67.

7.6 Dependence on sample number

The reconstruction performance also depends on the number of samples available. The reconstructions shown in this chapter were performed with 60 minutes of simulated time at an image rate of 50Hz, resulting in a sample number for each neuron of $S = 180,000$ samples. As fewer samples deteriorate the estimates of the transition probabilities, the reconstruction performance decreases with lower numbers of samples (see Fig. 7.8B). Plotted is the performance against sample number divisor s , i. e. by using only a reduced number of samples $S' = S/s$ for the GTE reconstruction. After a reduction in sample size by a factor of 5 (36,000 samples), the performance level of GTE is about 40%, and still shows a performance significantly better than random choice for a sample number divisor of $s = 15$. A longer recording time should lead to an even better reconstruction (see insets in Fig. 7.8B) yet such a recording is presently unattainable (see discussion in Section 6.4.1).

7.7 Dependence on system size

An interesting question is what the limit of performance is as the size of the system increases, especially in light of the fact that the real recording that will be analyzed in Chapter 8 contain on

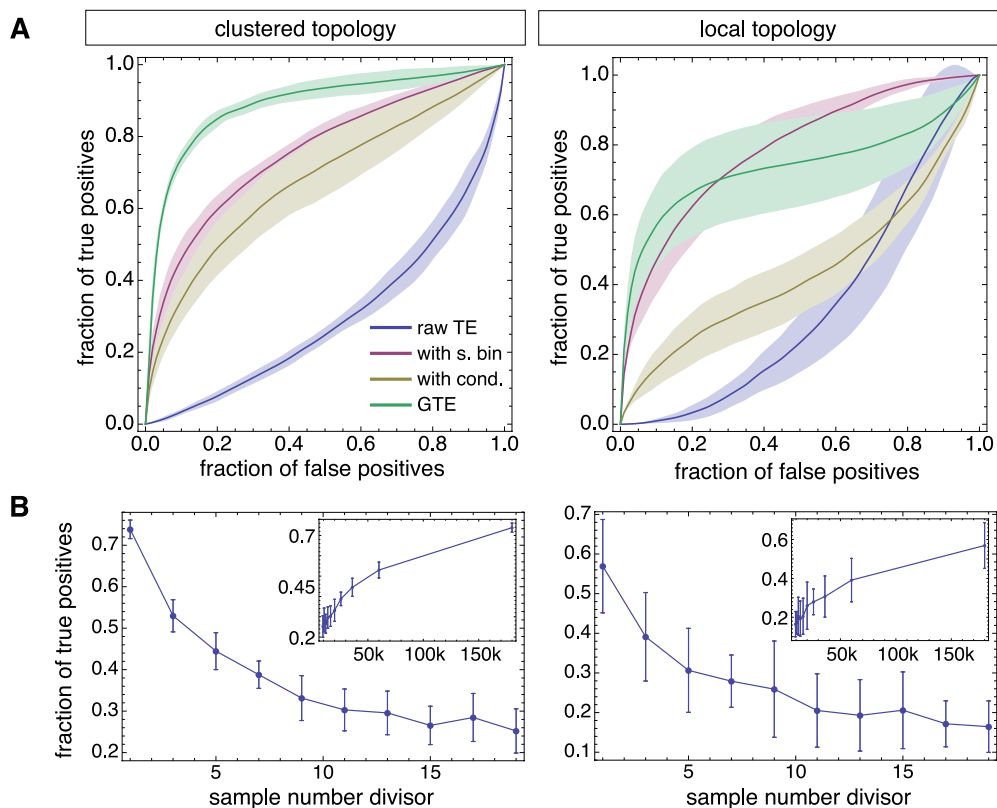


Figure 7.8: Reconstruction performance of individual extensions of Transfer Entropy, and due to sample number. (A) ROC curves for different extensions of TE, for the networks from the non-locally clustered ensemble as shown in Fig. 7.1 (left panel) and from the locally clustered ensemble as in Fig. 7.1 (right panel). Curves are based on the original formulation of TE based on the difference signal of the fluorescence (blue curve), including same-bin interactions (red curve), alternatively including conditioning (yellow) and including both, yielding GTE itself (green curve). (B) The dependence of the reconstruction performance of GTE across six networks, from both ensembles as in panel A, against sample number divisor (see main text). Insets show that same data, but plotted with the sample number itself as the horizontal axis. Shaded areas and error bars correspond to 95% confidence intervals across six networks. Figure reproduced with permission from [169].

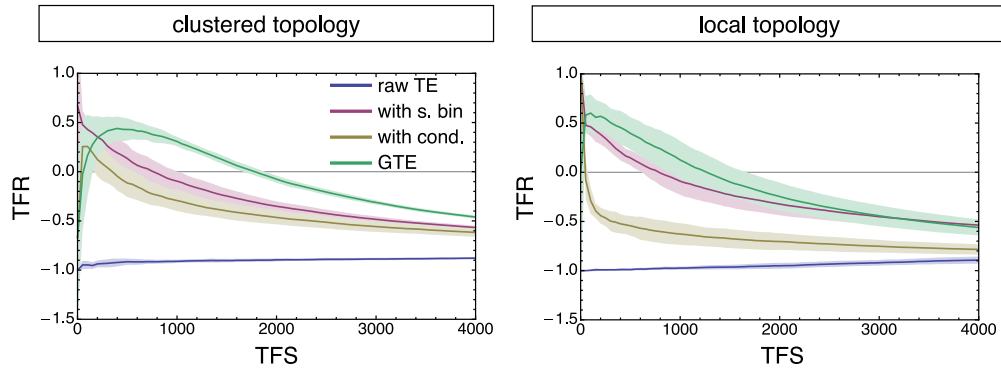


Figure 7.9: Positive-prediction curves individual extensions based on Transfer Entropy, showing the result from six networks from the non-locally clustered network ensemble (left panel) and from the locally clustered ensemble (right panel). Underlying reconstructions and color code identical to Fig. 7.8A. Figure reproduced with permission from [169].




Figure 7.10: Dependence of GTE reconstruction quality on conditioning level and system size. Shown is the reconstruction quality in terms of fraction of true positives at 10% of false positives. Note the logarithmic vertical scale, encompassing network sizes ranging from $N_{\min} = 30$ (bottom of panel) to $N_{\max} = 4,000$ (top of panel).

the order of a thousand neurons (regions of interest).

The dependence of reconstruction performance of a local topology when changing system size N from tens to thousands is obtained by rescaling the associated length scales λ and λ_{sc} appropriately with a factor of $1/\sqrt{N}$. This simulates a recording obtained simply by increasing the recorded area, “zooming out” of the culture such that more neurons fit inside the field of view. Note that this results in a total number of connections that scales linearly with the number of neurons, rather than quadratically if the overall probability of connection was constant.

It turns out that the reconstruction performance actually increases with system size (see Fig. 7.10). At the same time, the range of conditioning levels is broadened for which an excellent reconstruction performance is achieved.

Chapter 8

Experimental recordings of in vitro networks

Plots of experimental recordings of in vitro networks were already, presented in Chapter 5, highlighting the similarity of the statistics of mean fluorescence between simulations and recordings. The developed causality measure GTE was benchmarked in Chapter 7, allowing now for the application of GTE to these experimental recordings. We note that as stated previously, these are early mature cultures between 9 and 12 days in vitro (DIV), where inhibition was blocked such that the network activity was driven exclusively by excitatory connections.

8.1 Excitatory connectivity of in vitro networks

8.1.1 Cultured network at 12 DIV

The resulting network activity for a culture of 12 DIV was displayed in Fig. 5.1. Applying GTE, to the fluorescence recordings obtained by selecting manually a number of $N = 1720$ clearly distinguishable neurons (regions of interest), yielded, after thresholding for the top 5% of GTE values reconstructed networks (see Fig. 8.1A). The threshold was chosen such that the average degree is consistent with that found in the literature [55,63].

The GTE reconstruction reveals a markedly non-trivial network topology (see top panel of Fig. 8.1B). To analyze the reconstruction result — given that, for experimental recordings, the ground truth topology is of course not known — the distributions of topological indices are compared against the distributions of the same indices after randomizing the adjacency matrix in one of two ways: A *partial randomization* is done by randomizing the target of all links, thereby keeping the number of outgoing links for all nodes constant; and a *full randomization*, which randomizes the source and target of all links, keeping only the total number of links in the network preserved. This procedure allows for the differentiation of different biases in the reconstruction.

In the case of the connection lengths, the GTE reconstruction displays a broadened degree distribution with 6.3% of neurons having degree $k > 200$, much more than would be expected from a Poisson degree distribution as found in a random graph (see top left panel of Fig. 8.1B). On the other hand, we find no sign of a scale-free degree distribution (i. e. a power-law distribution [203]) reported elsewhere [51,174,204] as there are no neurons in the network with degree $k > 350$.

The distribution of connection lengths is markedly non-local, with only a minor, statistically insignificant tendency towards lower average values. The distribution of CC values across the population is significantly higher than in a random graph with a mean of 0.09 (compared to a mean

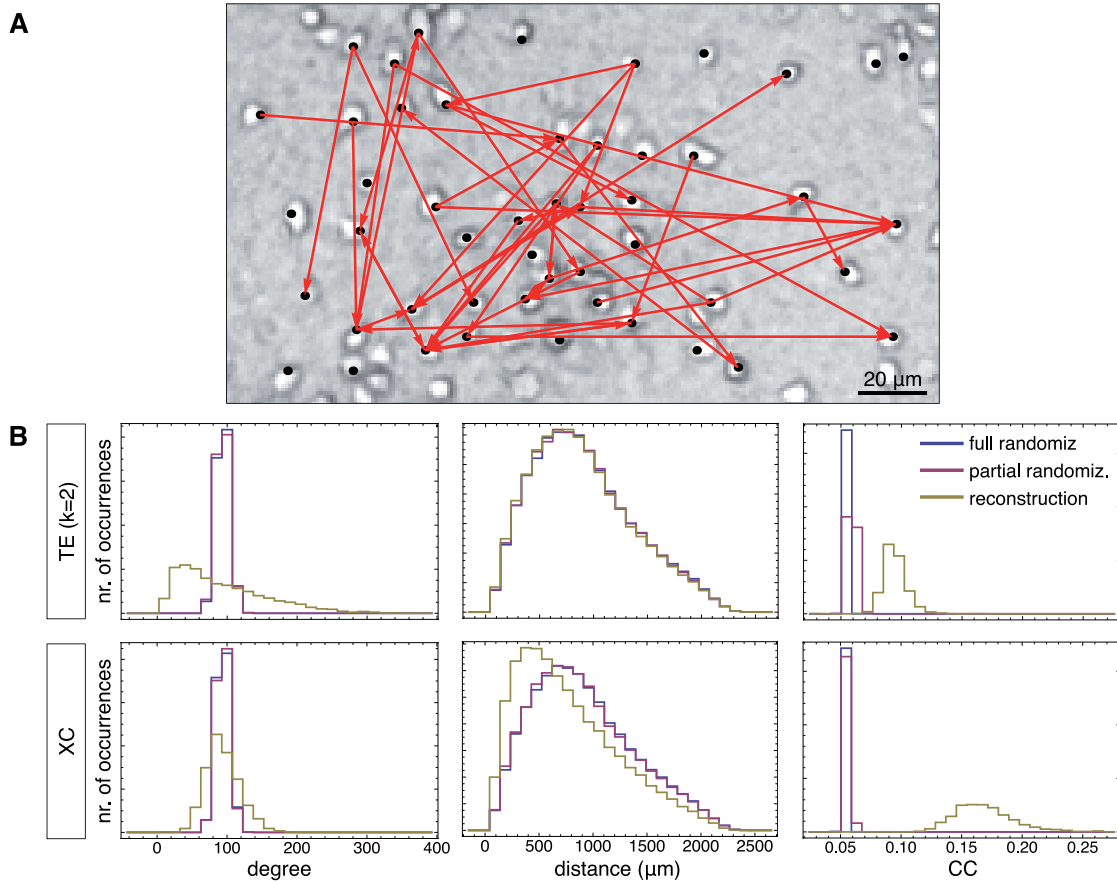


Figure 8.1: GTE reconstruction of the excitatory connectivity of an in vitro network at 12 DIV. (A) For an example subset of selected neurons (indicated by black dots) in the culture (49 out of a total number of $N = 1720$) the top 5% of reconstructed links are plotted (red arrows). (B) Histograms of reconstructed topological quantities, namely the in-degree (left column), connection length (middle column) and full clustering index are plotted for a reconstruction using GTE (top row) and XC (bottom row). Each panel compares the actual reconstruction (yellow) with two randomizations; a partial randomization which randomizes only the target of all links in the reconstructed network (red) and a full randomization maintaining only the total number of links in the network with randomizes source and targets of all links (shown in blue). Figure reproduced with permission from [169].

CC of 0.05). The mean of the partial randomization is not significantly different from the random level, indicating that there is genuine clustering present in the network, beyond just the broadened degree distribution.

The reconstruction based on XC shows different properties (see bottom panel of Fig. 8.1B). The degree distribution here is much less broadened compared to the GTE reconstruction, and the distribution of connection lengths is visibly shifted towards smaller connection lengths and a more local topology. This effect is complemented by a much larger average clustering coefficient of $CC = 0.17$. Interestingly, these effects are analogous to the error symptoms discovered for synthetic data in Section 7.1, where XC — independently from the ground truth topology — inferred an artificially high clustering coefficient and a low average connection distance.

8.1.2 Cultured network at 9 DIV

Calcium fluorescence recordings obtained from a younger network of hippocampal neurons at 9 days in vitro is shown in Fig. 8.2A. The signals demonstrate that bursting activity has already developed. Note that signs of sub-bursting firing activity is sometimes visible in the time before the bursts.

A subset of the GTE reconstruction of the same recording is plotted in Fig. 8.2B. Interestingly, the reconstructed properties share many similarities with the ones from the culture at 12 DIV (see top row of Fig. 8.2C). The degree distribution is again broadened, and the average CC again slightly elevated as compared to the random graph with the same number of connections; in fact the average clustering coefficient is now 0.11, higher than in the case of 12 DIV. Interestingly, the probability for connections to display a small connection length is now elevated, likely due to the younger age of the culture.

The XC reconstruction shows the same symptoms as for the case of the culture of 12 DIV, displaying a sharper degree distribution than GTE, and more importantly a much more local, more highly clustered connectivity (see bottom row of Fig. 8.2C).

Note that while the non-local topology is in contrast with connectivities found in slices of cortical tissue [49] it is entirely consistent with connection lengths obtained from GFP staining of dissociated cultures such as the ones used here (see Fig. A.1).

8.2 State-dependency in experimental recordings

The choice of correct conditioning level for the reconstructions studied above was motivated by the arguments developed for the state-dependency seen in simulated data (see Fig. 6.1). This logic was transferred to the analysis of experimental recordings due to the general nature of the logic developed, and based on the similar structure in histograms of average fluorescence (see Fig. 5.1D) indicating a general similarity in firing and fluorescence dynamics in simulated and recorded data sets. It is nevertheless interesting to compare state-dependency between simulations and recordings.

We first note that, comparing the two experimental recordings shown above to a simulated fluorescence time series, the similarity in general behavior of the fluorescence dynamics (see top row of Fig. 8.3). After splitting the time series, in analogous manner to the procedure of Fig. 6.1, into four regions I to IV, the resulting networks can be analyzed regarding different topological quantities (see middle and bottom row of Fig. 8.3).

For each of the four regions and each of the three data sets, the average clustering coefficient CC and the average connection length are plotted. Note that the thresholds used for experimental recordings and simulations are different (5% and 10%, respectively) due to the different system

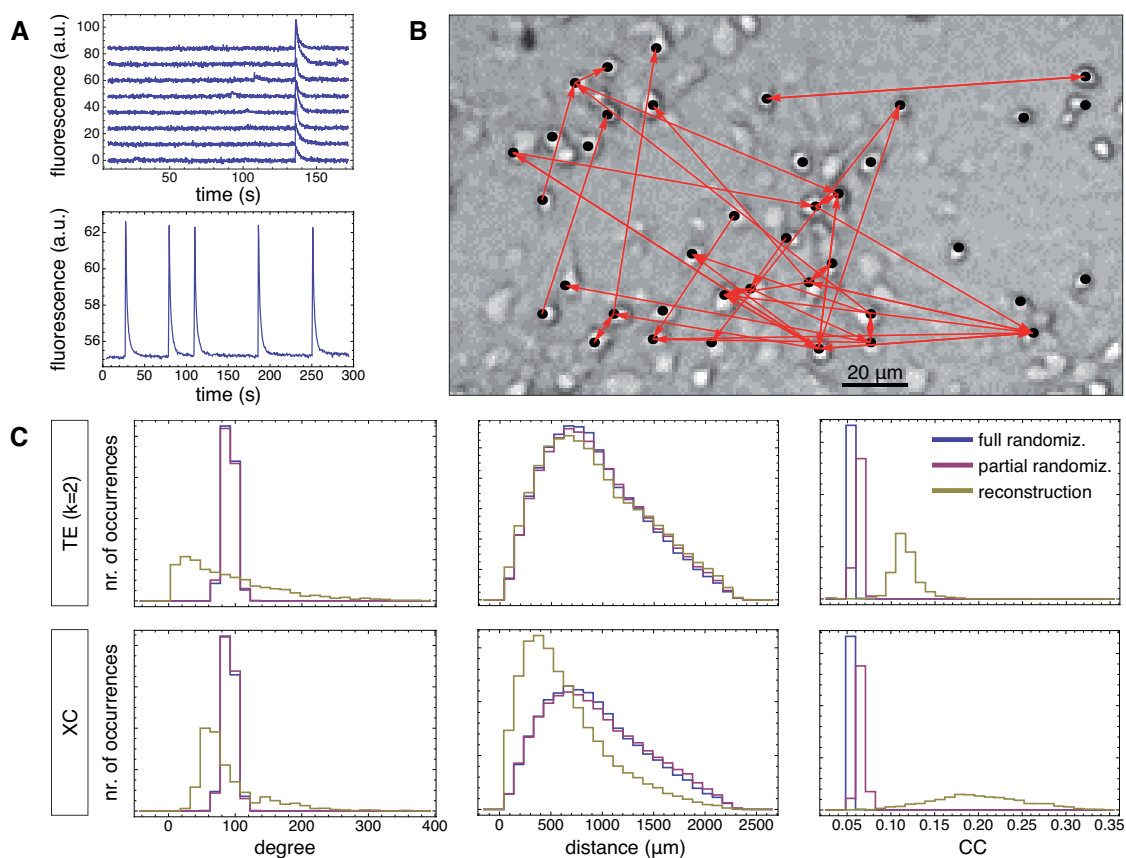


Figure 8.2: GTE reconstruction of the excitatory connectivity of an in vitro network at 9 DIV. (A) Example traces of the fluorescence signal of five example neurons (top panel) is plotted together with the population average (bottom panel). (B) For an example subset of selected neurons (indicated by black dots) in the culture (39 out of a total number of $N = 1668$) the top 5% of reconstructed links are plotted (red arrows). (C) Histograms of reconstructed topological quantities, namely the in-degree (left column), connection length (middle column) and full clustering index are plotted for a reconstruction using GTE (top row) and XC (bottom row). Each panel compares the actual reconstruction (yellow) with two randomizations; a partial randomization which randomizes only the target of all links in the reconstructed network (red) and a full randomization maintaining only the total number of links in the network with randomizes source and targets of all links (shown in blue). Figure reproduced with permission from [169].

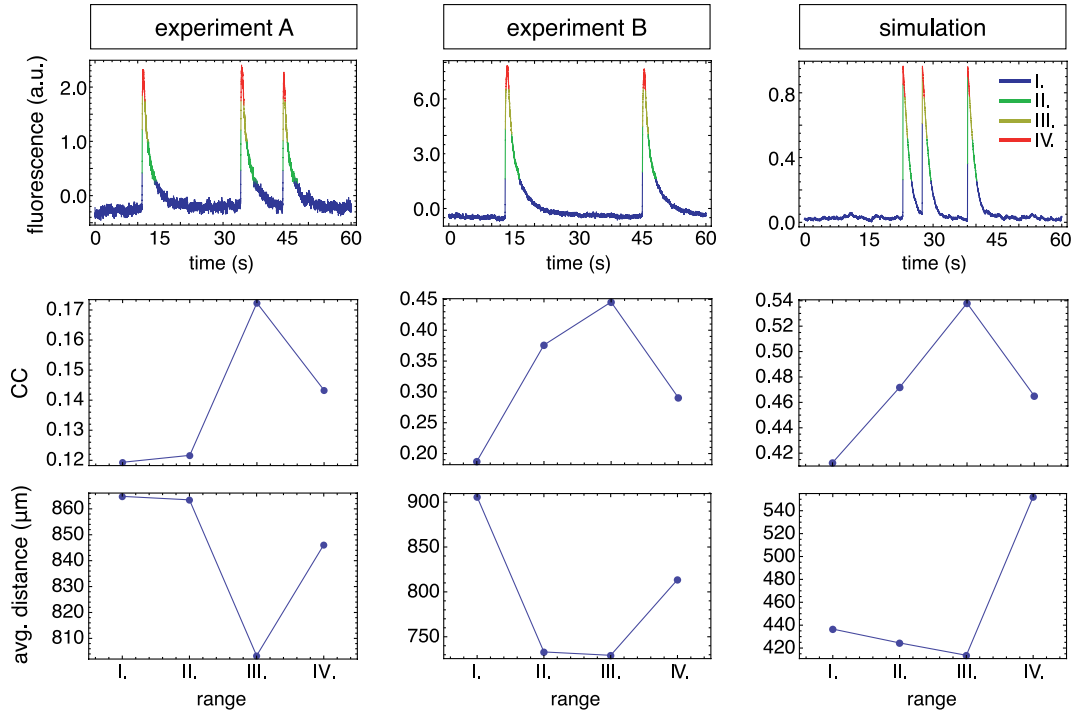


Figure 8.3: Dependency of reconstructed topological quantities on dynamical state in experimental recordings and simulations. For three data sets, both from experimental recordings and simulations, the average histograms are plotted (top row), color-coded (from blue to red) to indicate four ranges of average fluorescence. For each of these ranges, the average clustering coefficient CC (middle row) and the average connection length (bottom row) are shown, based on reconstructions according to $TE^{\text{state-dep}}$. (see Eq. 6.14). The data sets used are “experiment A”, corresponding to the 12 DIV recording shown in Fig. 8.1, “experiment B”, denoting the 9 DIV recording shown in Fig. 8.2, and “simulation” (right column) which is the locally clustered topology shown in Fig. 7.3. Figure reproduced with permission from [169].

size, as discussed above. Interestingly, it turns out that while individual CC values are different, the general structure — with the lowest CC value being calculated in region I, followed by a rise in region II and a peak in region III — is conserved across data sets. Similarly, the lowest average connection length is observed in range III, again across all datasets. These results further support the argument developed that the conditioning can indeed be transferred to fluorescence signals from experimental recordings.

Chapter 9

Extending GTE to inhibitory connections

So far, only excitatory connections have been considered. The simulations were done using only excitatory links, and the experimental recordings were performed with cultured treated with bicuculline to block inhibitory synapses. In real networks, however, about 20% of neurons are inhibitory neurons, [47, 63] so a natural question is whether or not inhibitory interactions can be revealed with the methodology presented above.

The work in this chapter has been done in cooperation with Javier Orlandi of the University of Barcelona and is to be considered preliminary.

9.1 Reconstruction of both inhibitory and excitatory links

In principle, there is no reason why GTE should not be able to reliably reconstruct inhibitory connections. Because GTE measures causal interactions and precisely because of its independence of a specific model, there are in principle no modifications necessary to reconstruct inhibitory connections. All that is necessary is an influence of the respective interaction on the transition probabilities of the fluorescence time series, and such an effect is clearly present both for excitation and inhibition due to the change in post-synaptic firing probability upon pre-synaptic firing. GTE should, in effect, be related more to the *absolute value* of the strength of a synaptic interaction. Note that this effect, if true, is very positive for the extension to inhibitory connections as it actually negates the need for a modifying GTE, but it poses the new problem, discussed later in the next section, of differentiating excitatory and inhibitory links.

The stated hypothesis can be tested in a minimal, non-bursting network model, consisting of $N = 100$ neurons, of which 25 neurons have excitatory projections, and 25 neurons have inhibitory projections (see left panel of Fig. 9.1). Each neuron from the sub-populations projects to other 5 neurons, which in turn do not have any outgoing connections, making the topology free of loops. Each neuron in the network receives independent Poissonian input spike trains of rate $\nu_{ext} = 30\text{Hz}$ and weight $\alpha_{ext} = 1\text{pA}$, while the internal synaptic weight is set to $\alpha_{ext} = 2.8\text{pA}$. The synaptic time constant of inhibitory connections is $\tau_s = 6\text{ms}$, twice the value of excitatory connections. All other parameters are as stated above for the purely excitatory networks. Note that because there are no bursts in this minimal model, no conditioning is used when evaluating GTE.

The resulting GTE values are plotted in the right panel of Fig. 9.1. The links in the ground-truth topology are reconstructed with excellent accuracy, separated by a factor of 2-5 in terms of GTE value from the background. Due to the low temporal resolution of the recording, thus making the

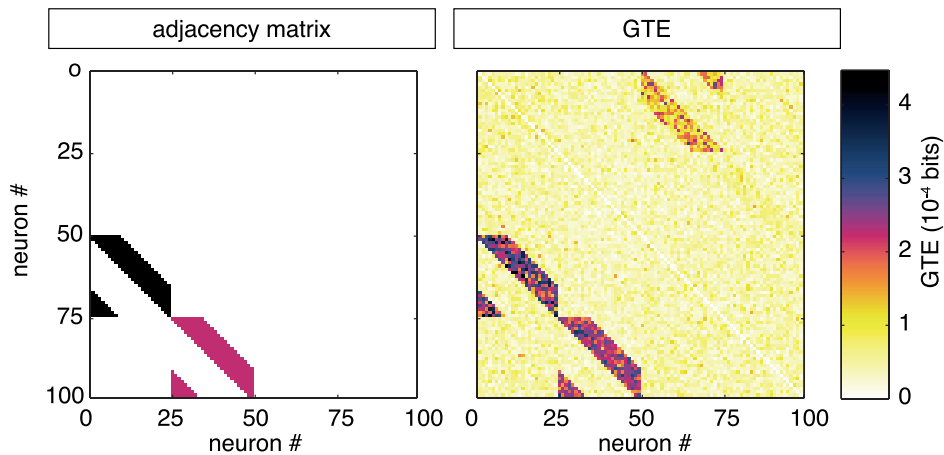


Figure 9.1: Minimal network model with both excitatory and inhibitory links present, showing the adjacency matrix (left) and the reconstructed GTE values (right panel). In the adjacency matrix, excitatory links are shown in black, and inhibitory links are shown in red. Note that links are defined here as column-to-row, i. e. a link $i \rightarrow j$ corresponds in the matrix plot to row j and column i .

identification of directionality of links more difficult, the “mirrored” connections of especially the excitatory ground-truth links also yield elevated GTE values. This is however not the case for the inhibitory sub-population, which is most likely due to the longer time constants of post-synaptic inhibitory potentials.

Interestingly, and confirming the hypothesis stated above, there is no significant difference in GTE amplitude between the true positive links for excitation and inhibition.

9.2 Identification of inhibitory links

Being able to reconstruct both excitatory and inhibitory links is only part of the task of extending GTE to excitatory connections. As GTE yields the absolute value of a causal interaction in terms of the Kullback-Leibler distance between two conditional probability distributions, it is a priori not clear whether the distinction between excitatory and inhibitory links can be made.

Note that the nature of the influence of the type of the connection on the transition probabilities is by definition different, at least on average or disregarding noise: Excitation leads to an increase in post-synaptic firing probability, and thus potentially to an increase in post-synaptic fluorescence amplitude. Inhibition, on the other hand, leads to a decrease in post-synaptic firing probability, thus potentially avoiding a post-synaptic increase in post-synaptic fluorescence amplitude, and thereby leading actually to a decrease in post-synaptic fluorescence amplitude due to its exponential decay.

Thus, it should be possible to differentiate the two types of connections not based on their GTE score, but on the individual terms of the sum that makes up GTE (cfr. Eq. 6.15). This hypothesis is tested, again for the case of the minimal model network, in Fig. 9.2. The logic developed above is clearly visible, with excitatory and inhibitory links showing opposite average values for terms corresponding to simultaneous firing (label “PNNPN” in Fig. 9.2) or firing only of the source node (label “NNNPN”). Naturally and as stated before, the connection between fluorescence and firing can only be made on average due to the effect of noise.

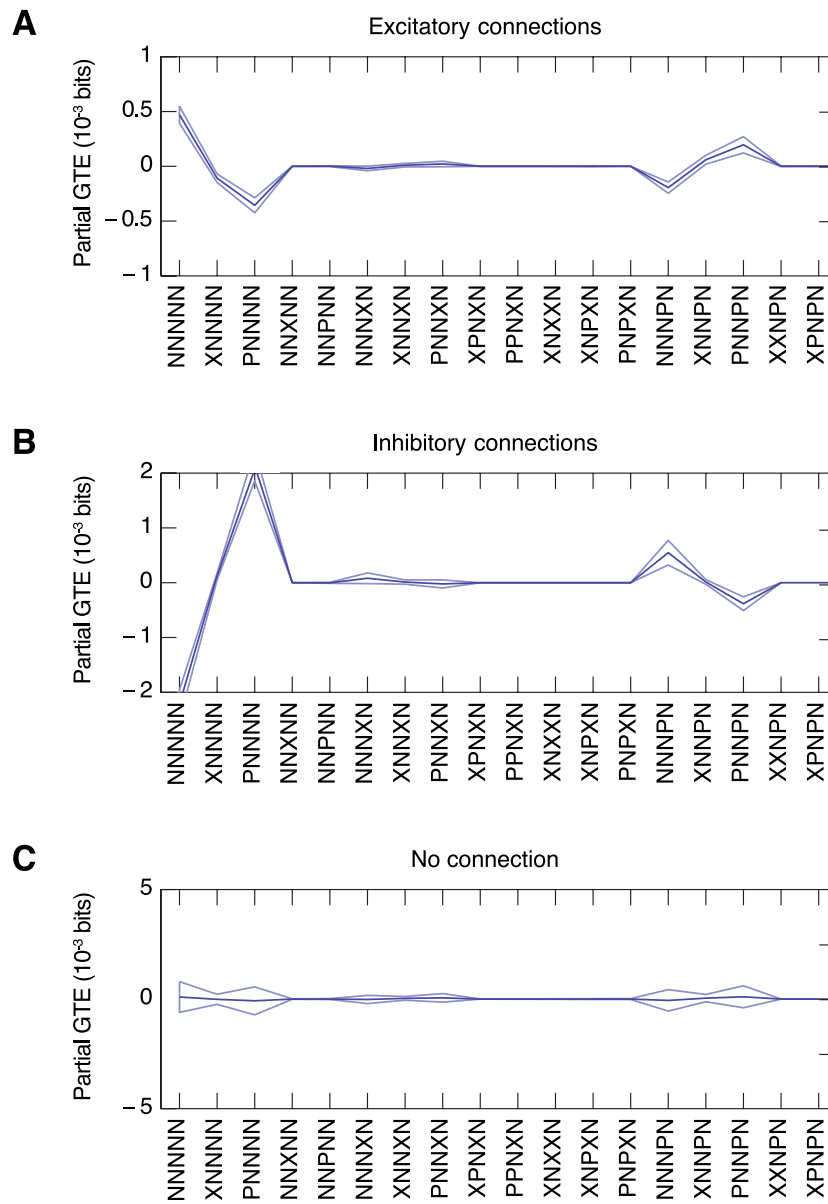


Figure 9.2: Separation of the GTE scores from the analysis of Fig. 9.1 into different terms, for all the excitatory links (panel A), the the inhibitory links (panel B) and the false links (panel C). GTE is calculated with Markov order $k = 2$ and $b = 3$ bins, the latter indicated by “P” (positive), “X” (neutral) and “N” (negative) according to the slope of the fluorescence signal. The five-letter labels then correspond in turn, for source Y and target X , to $x_t, x_{t-1}, x_{t-2}, y_t$ and y_{t-1} . A label of “PPPNX” thus stands for the GTE term corresponding to a continuous rise of the target signal, and a sudden decline of the source signal. Note that only a selection of the $3^5 = 243$ terms is shown. The mean across the respective terms of all relevant links is indicated by the dark blue line, with the light blue lines indicating the width of twice of the standard deviation.

Based on this analysis, at least in the minimal network model the distinction between excitatory and inhibitory synapses can clearly be made, with very high accuracy. Note that we have not even exploited the possibility of combining different terms of GTE to improve the accuracy, neither have we made use of the fact that synapses of neurons are prevailingly all excitatory or all inhibitory (often called, slightly incorrectly, Dale's law [200,205]). Both aspects could be used to improve the identification of excitatory versus inhibitory links, especially for experimental recordings.

Chapter 10

Summary

Existing measures of causality are inadequate to measure interactions in the complex data sets of fluorescence imaging recordings due to non-linear dynamics, non-stationarity and insufficient temporal precision of typical recordings. Analyzing these problems, we have developed a simple but crucial extension to Transfer Entropy, a causality measure from information theory, which we call Generalized Transfer Entropy.

We have systematically tested and benchmarked GTE to reconstruct excitatory connections based on simulated fluorescence signals, finding very good performance across networks and network ensembles. GTE even allowed for the reconstruction of statistical properties of the network topology such as the clustering index or the typical connection length. Furthermore, the robustness of GET was extensively demonstrated.

Applying these results and the developed method to the analysis of experimental recordings, we found non-trivial topological features including a broadened degree distribution and an elevated clustering coefficient.

We have also presented preliminary results of extending GTE for use with networks which have both excitatory and inhibitory links, showing that, at least in model networks, the new measure is suitable for that purpose as well.

Apart from the direct results presented here regarding calcium fluorescence imaging, we point out that due to its general form, GTE could directly be applied to spike data or entirely different time series. It is therefore our hope that GTE will find applications, beyond neuroscience, in the field of statistical network reconstruction in general.

Part III

Predicting dynamics based on reconstructed topology

Chapter 11

The interplay between network structure and dynamics

A natural question arising from the previous parts of this thesis — having developed analytical predictions of the dynamics in model neuronal networks in Part I, and having developed an approximate method to reconstruct the topology based on the neuronal dynamics in Part II — is whether it is possible to predict the dynamics of the network based on the *reconstructed* network topology. This is especially interesting in light of the finding of Part I, where in all model systems studied, the common determining factor for the expected first firing time was the local network topology surrounding the neuron in question.

11.1 Determining the onset time from data

In a first step towards answering this question, the average first firing time at the onset of a burst for each neuron has to be estimated from the data. Given the nature of the experimental recordings, spike times are not available directly, but only affecting the amplitude of the calcium fluorescence recordings indirectly (see Chapter 5). Consequently, our analysis must go beyond what was done in Part I, where the probability to fire for the first time at the onset of a burst $P_i^{f1}(t)$ was available directly (see for instance Eq. 2.9). Here, what is available is only the shape of the fluorescence signal during a burst, which increases, starting from the baseline level, at a time we call the *onset time* t_0 . This onset time might, relative to the population mean, be different for each neuron in the network.

In the following, the onset time will be extracted from the fluorescence time series. To compare its value with the analytical predictions developed earlier, we will assume that the value of the onset time corresponds at least approximately to the expected first firing time. This is motivated by the fact that the change in firing rate at the onset of a burst is much faster than the decay of the fluorescence signal. This allows for the fluorescence amplitude — given the experimental observation that the latter is approximately directly proportional to the number of action potentials involved [59, 81] — to be taken as an approximate measure for the firing rate of each neuron. Thus, a shift in time of the fluorescence amplitude before a burst compared to the population average can be seen as a shift in firing rate exactly corresponding to the distribution of values of \hat{t} as developed in Part I.

The reliable extraction of the onset time requires one to first determine the average burst shape for each neuron, relative to the burst shape of the population average. To this end, the population averaged signal was thresholded at a level Θ_F to determine burst onset times. The individual burst




Figure 11.1: Fitting alpha functions to average bursts to determine onset times. Shown are, for a simulated network (left) and an experimental recording (right panel), the averaged bursts for four example neurons (indicated as colored dots, with each color corresponding to a different neuron) together with the fitted curves (shown as solid lines in the same color code). Insets are histograms of the resulting onset times t_0 across the respective population. The time $t = 0$ (vertical line) indicates the point where the globally averaged burst (averaged not only across burst events but also across the population) crosses the respective thresholds, which were $\Theta_F = 0.6$ for the simulation and $\Theta_F = 40\%$ of the range for experimental recordings (traces here are shown after subtraction of the mean value).

events for each neuron were then aligned in time by determining the time of threshold crossing to sub-millisecond precision by way of linear interpolation of the fluorescence signal. By averaging across events, this results in an average burst shape for each neuron. To represent the rise of the burst from a baseline level F_0 , an alpha function $\chi(t)$, characterized by the onset time t_0 , the amplitude A and time constants τ_R for the rise and τ_D for the decay, is then fitted to these burst shapes.

$$\chi(t) = F_0 + A \frac{\exp(-(t - t_0)/\tau_D)}{\exp(-(t - t_0)/\tau_R) + 1} \quad (11.1)$$

The result of this procedure for four example neurons of a simulation based on a locally clustered topology (as shown in the right panel of Fig. 5.1B) is plotted in the left panel of Fig. 11.1. The fit indicates a good overlap between the data and the assumed shape $\chi(t)$, resulting in a small mean-squared error across the population (not shown).

In the case of experimental recordings, corresponding to the cultured network analyzed in Fig. 8.2, the onset is slower (see right panel of Fig. 11.1), resulting in a rise time constant of $\tau_R = 0.043 \pm 0.004\text{s}$ (mean \pm standard deviation) compared to the simulations with a mean value of $\tau_R = 0.010\text{s}$. The rise time constant was fixed in both cases to $\tau_R = 2.0\text{s}$.

11.2 The distribution of onset times

The distribution of onset times for simulated networks was quite narrow, with a range of values with $t_0 = -2.0 \pm 3.3\text{ms}$ (mean \pm standard deviation), as shown in the inset of the left panel of Fig. 11.1.

For experimental recordings, the distribution was more Gaussian-like and much more broad (see inset of the right panel of Fig. 11.1), yielding onset times of $t_0 = -3.6 \pm 15.8\text{ms}$. Note that the number of burst events from which the average shape was determined (31 events) was much lower

than in the case of the simulation (250 events), due to a low burst rate for this culture. This fact may contribute to this broad distribution, in addition to the heterogeneity expected in biological neuronal networks.

Chapter 12

Network reconstruction can predict onset times

Having extracted the burst onset time for each neuron, a quantitative comparison to the predicted onset times — based on the reconstructed topologies using GTE from Part II — can be made.

12.1 Onset time prediction in simulated networks

In a simulation, the onset time of each neuron can be assessed with good accuracy from the shape of the average burst event. The predicted value for the onset time can then be calculated based on the first order approximation to the leaky integrate-and-fire dynamics of Eq. 4.5, and on the GTE-reconstructed network topology yielding, for a given threshold Θ_{TE} , the in-degrees k_i for all neurons.

For a fraction of 20% of links, the resulting scatter plot relating measured and predicted onset time is shown in the top left panel of Fig. 12.1. The width of the distribution of onset times is over-estimated by a factor of about two. A correlation between measured (“real”) and predicted values of a Pearson’s coefficient of 0.29 is achieved. This correlation depends on the fraction of links taken from the GTE reconstruction (see top right panel of Fig. 12.1).

The bottom panel of Fig. 12.1 shows the same network and simulations, but predicting the onset time with the second order approximation of Eq. 4.10, thereby necessarily taking a larger topological neighborhood of each node into account. Taking for the degree distributions found for the top 20% of links, this results in an excellent correlation coefficient of 0.76, indicating a high precision of the predicted onset times and by extension the reconstructed topology. Note also the range of real and reconstructed onset times is now virtually identical.

It turns out that the quality of this prediction depends crucially on the fraction of links in the reconstructed network (see bottom right panel of Fig. 12.1). Starting from a low fractions of links, the resulting correlation coefficient rises up to a peak at 0.76 for a fraction of links of 0.2, and decreases for higher fractions. The reason for this behavior is most likely in the strong (negative logarithmic) dependence of the prediction of the expected first firing time \hat{t} on a low numbers of links, and therefore the sensitivity of the onset time prediction to errors in the reconstructed degree distribution for that low fraction of links and therefore a low average degree of each node. For higher numbers of links, the true positives of the reconstruction are washed out by a large fraction of false positives (see e.g. Fig. 7.4).

Nevertheless, given the number of assumptions involved in deriving both the measured and the predicted onset times, this high correlation constitutes a remarkable agreement.

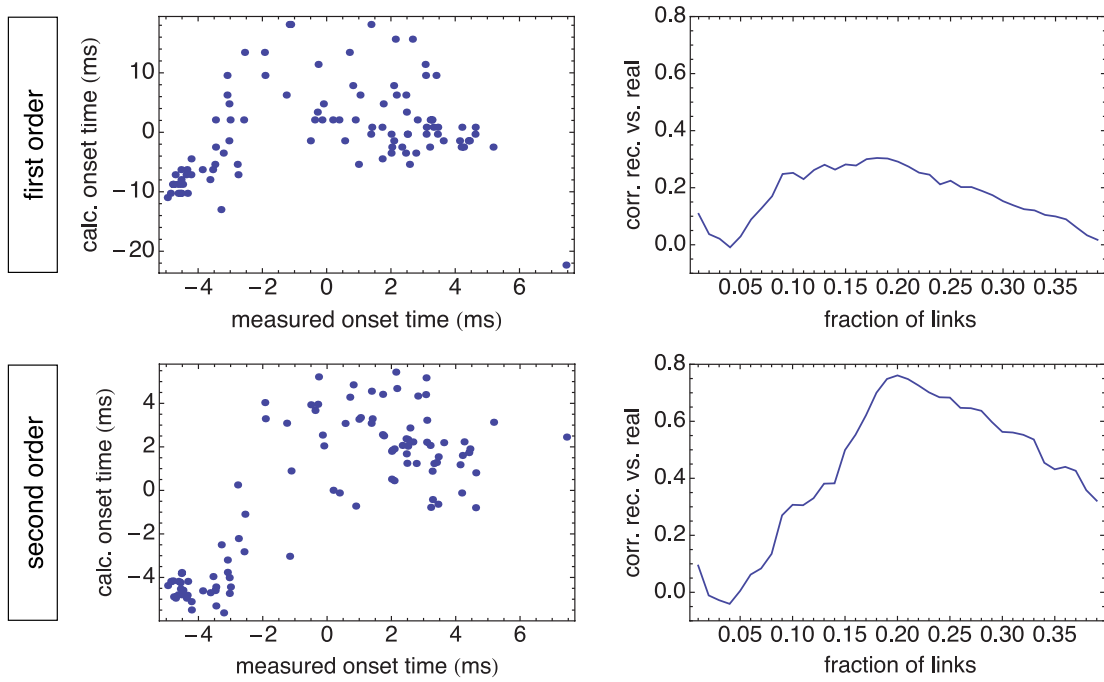


Figure 12.1: Comparison between measured and predicted onset time in a non-locally clustered, simulated network. For the network obtained from the top 20% of links, the predicted onset times are plotted in the left column against measured onset time, according to the first order equation $\hat{t}^{(1)}$ (top panel) and the second order equation $\hat{t}^{(2)}$ (bottom panel). In the right column, the Pearson's correlation coefficient between measured and predicted onset time (of the corresponding order, see left column) is plotted for different numbers of included links, i. e. thresholds Θ_{TE} .




Figure 12.2: Comparison between measured and predicted onset time in a cultured network at 9 DIV of dissociated neurons. For the network obtained from the top 5% of links, the predicted onset times are plotted against measured onset time, according to the first order equation $\hat{t}^{(1)}$ (left panel). The Pearson's correlation coefficient between measured and predicted onset time is shown for different numbers of included links, i. e. thresholds Θ_{TE} (right panel).

12.2 Onset time prediction in experimental recordings

For an experimental recording, the first order correlation between measured and predicted onset time is weaker than for simulations, yielding a Pearson's coefficient of 0.24 (see left panel of Fig. 12.2). Also, the width of the distribution of onset times is again overestimated, resulting in a standard deviation of the prediction of 27.3ms against value of 10.1ms for the measured distribution.

It is interesting to see that the dependence of the correlation on the number of links has a similar shape than for the case of simulations in the sense that there is a drop at low numbers of links, followed by a peak at higher numbers (see right panel of Fig. 12.2). The fact that the second order prediction results in a reduced correlation coefficient of 0.11 (not shown) seems to confirm this hypothesis, as it depends on a larger topological neighborhood of the neuron in question and is therefore more susceptible to errors in the reconstructed topology.

Despite the fact that value of the Pearson's coefficient is lower, we find that it is statistically highly significant, with an estimated error of ± 0.0027 obtained via standard bootstrapping.

Part IV

Discussion

Chapter 13

Structure and dynamics

In networks of neurons, there is a complex relation between the structure of the network, given by the synaptic connections, and the dynamics of the network, given by patterns of action potentials in the population. Structural patterns do not determine, but rather restrict the possible dynamics, and vice versa.

In the present work, this complex relation was studied *in silico* and *in vitro* based on signatures of non-randomness in the activity, such as bursting events and early firing neurons. Both the effects of structure on the dynamics, as well as what can be learned about the structure by observing the dynamics, were studied to enable a better analysis of neuronal networks *in vitro* and, in the future, *in vivo*.

13.1 First firing times in silico determined by structure

Spontaneous activity in discrete time

The onset of network bursts shows surprising signs of non-randomness in the order of activation. A small subset of neurons consistently fires ahead of the population burst [50,51]. The question is whether this early firing can be explained based on the topological structure of the underlying network.

We studied various mathematical model networks, highlighting different aspects of the neuronal dynamics. In the simple but illustrative model of discrete-time integrate-and-fire neurons, the origin of the non-random, hierarchical mode of activity at the onset of a burst was investigated. We could show that the exponential shape of the population burst onset can be predicted from the corresponding branching process dynamics and thus, directly from the known model parameters. The expected value of the first firing time can then be calculated as a function of the input current that a given neuron receives. We demonstrated that this first firing time, being indicative of a neuron's ability to fire before the bulk of the network and thus potentially act as a leader of activation, can be calculated directly from the basic model parameters. To first order in terms of the topological neighborhood, we find analytically a negative logarithmic dependence of the first firing time on the in-degree k . The numerical analysis based on simulations confirms this prediction and thus the strong influence of the network structure on firing activity in the population.

We showed how the introduction of non-random topological structure increases the range of expected first firing times in the network to tens of milliseconds, as reported in experimental studies [51]. Taking into account a larger topological neighborhood of each cell, the quality of prediction was further improved. In particular in non-random topologies, we found that early firing times are determined with very good accuracy by the network topology, if a higher order of

the included neighborhood is taken into account in the calculation.

Evoked activity in a quorum percolation model

In quorum percolation (QP) systems, the firing dynamics were studied after the excitation of a significant fraction of nodes. The expected firing times were again found to be primarily dependent on the local topology, with neurons with higher in-degree firing earlier than others.

We demonstrated analytically that in such a system it is possible to determine the degree distributions that are necessary to have an initially exponential rise in the population firing rate at the onset of a burst. The result is a Gaussian degree distribution with a power-law tail, such that a small group of neurons with very large degrees does exist in the network. As described in detail later, we do not find such a power-law tail in the structural topology of biological networks, and we will discuss the reasons behind this discrepancy.

Apart from the skewed degree distributions, the influence of higher order correlation on the dynamics was demonstrated to be small. In fact, it was shown by Tlusty and Eckmann that percolation graphs of the studied size behave essentially as if they were free of loops [141]. It is interesting that the higher-order correlations in the topology induced by a positive degree-degree correlation, which go beyond the presence of loops, did not have a stronger effect on the dynamics. This indicates that the graph also behaves as random in this regard, contrary to the effect seen in the integrate-and-fire model. We will discuss the difference in findings between QP and integrate-and-fire below, after presenting findings obtained using a model in continuous time.

Spontaneous activity in leaky integrate-and-fire models

The leaky integrate-and-fire (LIF) model is a biologically more realistic model system, as it allows for quantitative comparisons with all experimentally observed parameters and includes non-trivial neuronal dynamics in the response to stimuli. In this thesis, we studied the effect of variability of the expected first firing times preceding a burst in this system, finding analytical solutions that to first order parallel the results from the discrete time models. We find, however, that the higher order solutions differ from their discrete time counterparts due to the non-trivial input-output transfer function of the firing rate in the LIF model. This results in a pronounced effect of higher order correlations in the topology for low numbers of the input degree k , indicating a high first order input degree as a necessary condition for a neuron to fire earlier than others.

We showed that the expected first firing time can be predicted with high accuracy, despite the inherent non-stationarity of the burst onset, by taking into account the first and second order topology around each neuron in the network. The dominating influences on firing time are the input degree of a node and the topological surroundings. If a node is highly connected, and its neighbors are highly connected as well, there is a high likelihood that the neuron will belong to the group of early firing neurons. These findings are consistent with earlier studies pointing out the importance of the network topology for synchronization behavior [67].

Interestingly, in the integrate-and-fire models the calculated onset time had a logarithmic shape depending on the in-degree (see for instance Eq. 2.16 for the discrete-time integrate-and-fire and Eq. 4.5 for the leaky integrate-and-fire first-order solutions). Even with a Gaussian degree distribution, this would lead to a skewed first firing time distribution that displays a long tail of late firing times, in agreement with experimental observations [62].

We point out an important difference between the integrate-and-fire models and the quorum percolation (QP) case: The former simulates spontaneous activity, whereas the latter requires the ignition of a significant fraction of the population before activity can propagate. As such, QP is a model for externally evoked, stimulated activity. In this sense, it is interesting that only in the

integrate-and-fire models are higher-order correlations an important factor. This indicates that integrate-and-fire models are not only more biologically realistic, but are, crucially for the present application, better suited for the study of the onset dynamics because the onset is not skipped by externally igniting a fraction of the network. It was found in recordings using multi-electrode array that the neurons firing earlier at the onset of a burst event in the culture are identical to those firing earlier than others in response to stimuli [51,206], so there is necessarily some overlap between the two regimes, as indicated in the present analytical work on the dependence on the first-order degree in both cases. However, the dynamics of the onset of spontaneous bursting activity can only be fully understood by including the higher order topological neighborhood and thus higher order topological correlations.

This answers a crucial question stated by Eckmann et al. in [50] about whether the early firing neurons are part of a special sub-network, distributed throughout in the culture. We find that the early firing neurons are part of a highly-connected subset of the population, and thus directly by virtue of their higher degree more likely to be connected to each other. We will discuss more features of this sub-network later.

A question that remains is whether or not the effect of early firing neurons is possible in random networks in the experimentally observed magnitude. An indication that this is not the case can be found in the necessity to include higher order terms in the analytical calculations for both discrete and continuous time models, because thus, the higher order degree distributions do not average out as can be expected in a purely random graph. We found that the topological neighborhood also determines early firing in a topologically non-random but otherwise homogenous network. A possibility is, however, that in vitro parameters like the voltage threshold might be correlated with the structure in such a way that neurons with a higher number of connections display a higher threshold. Such a mechanism does indeed exist, and is called *homeostasis*, i. e. the fact that neurons which receive a high average input current down-regulate their synaptic efficacy over time [207–209]. This principle is nevertheless compatible with the expected case that neurons with a higher local connectivity have a higher firing rate, because the range of adaption in terms of the synaptic weight is relatively small. While for instance the input degree varies by a factor of over 20 in the in vitro networks (see Fig. 8.1B), the output firing rate of a neuron exposed to an extended period of complete elimination of the input current due to tetrodotoxin-blocked synaptic transmission shows, after wash-out, an increase in firing rate of only a factor of about 2 [35,210]. Thus, homeostasis is too weak to account for the observed variability in expected first firing times.

Nevertheless, we can not rule out such a possibility in general, without knowledge about the structure that underlies the observed dynamical features in vitro.

13.2 Connectivity can be reconstructed based on dynamics

Thus, the findings discussed above motivated the study of the synaptic connectivity of in vitro networks based on large-scale observations of the network dynamics. We focus here on recordings of calcium fluorescence imaging, which allows for the simultaneous recording of a large population of neurons for up to 60 minutes without damaging the cell culture. The fact that the network topology poses restrictions on the system dynamics, as discussed above, leads to the question of whether the observed dynamics might in turn restrict the compatible network structure, which would make a reconstruction possible.

The knowledge about the topology in cultured networks is limited at present, as most studies in these systems are either lacking in single-cell resolution, restricted to low numbers of cells such as patch-clamp techniques [49,84], or are based on a specific model of the underlying dynamics, such as assuming linearity as in [174] or inferring the parameters of a generalized linear model in a

Bayesian sense as in [106,211]. Specific models are problematic in a conceptual sense, when striving towards the solution with the least number of assumptions necessary, but also for more practical reasons, because an explicit model necessarily restricts the compatible dynamics of each cell to a given model domain. Due to the high variability inherent in biological networks, deviations from the assumed dynamics will occur and will introduce biases which can be difficult to isolate due to the complex architecture of the algorithm.

GTE and the state of the art

By way of improving on the methods described above, we have proposed a novel model-free measure of causality rooted in information theory as an extension of Transfer Entropy, which we have called Generalized Transfer Entropy (GTE). This measure is inherently non-linear and directly applicable to the signal of calcium fluorescence, despite the low temporal resolution of typical calcium recordings, and without the need to first infer the set of spike times in the network (using for example a method as in [171]). Even without precise spike times, GTE achieves a reconstruction performance in excess of those found in other studies based on spike times thanks to the inclusion of instantaneous interactions in the model-free framework and the possibility to take into account the inherent non-stationarity of the network activity due to bursting. The latter is done by conditioning the analysis on temporal ranges of relatively low-frequency spiking activity, thereby focussing on the causal effect of asynchronous mono-synaptic interactions. The effect of structure on the dynamics is highly complex, and thus the effective connectivity does not in general reflect mono-synaptic aspects of the network activity but higher-order correlations, but it is nevertheless possible to find ranges in the dynamical regime of the population where, given an appropriate measure of effective connectivity, there is an overlap between structural and effective connectivity. This overlap can indeed be used, by GTE, to reconstruct values for true (directed) links that are significantly higher than for false links (see Chapter 7 of this thesis and detailed discussion below).

For comparison, we have applied the Bayesian method of [171] — this being the first step in the reconstruction method presented in [106] — to infer the spike times from the simulated fluorescence signal. However, the algorithm failed to identify inter-burst activity for all but one neuron in the network, yielding a reconstruction quality very close to the random result (see Fig. 7.7). The systematic bias in detecting spike times was due to an erroneous inference of the change in calcium concentration caused by a single action potential — for signals with a high maximum fluorescence amplitude, network bursts were identified as single-neuron events. Given that the presented case is just one possible problem in relation to one of a large number of parameters in the model, it seems unclear how such biases are to be avoided when analyzing experimental data sets. We point out that while the simulation that we used to generate surrogate data sets consist of a number of parameters as well, what is needed is a measure of causality that is model-free, as is GTE.

Another alternative to GTE, the method of reconstructing functional connectivity by Bonifazi et al. [174], is based on cross-correlations and is indeed similar to the XC score as defined in Eq. 6.16. To allow for comparisons of the reconstruction performance, we point out the key differences: The algorithm of Bonifazi does not include conditioning, and potential links corresponding to a peak in the cross-correlogram at a time lag of zero are pruned from the reconstructed network. As discussed in Section 6.4.1, in our case the latter leads to a dramatic decrease in reconstruction quality, regardless of whether the cross-correlogram is computed based on the raw fluorescence (as in [174]) or on the conditioned signal (as in the present work). In the latter case, for instance, the performance level of a reconstruction from the non-locally clustered ensemble as in Fig. 7.1A is reduced from 80% to 35% (not shown). As such, while the method proposed in [174] may help in identifying hubs of functional connectivity, it is not able to identify the structural connectivity of a

neuronal population.

We note that due to the fact that GTE is model-free, it is immediately applicable not only to recordings using different fluorescent dyes or different microscopy set-ups, but also to the analysis of spike times or other time series, even to time series unrelated to neuroscience. It is also directly applicable to in vivo recordings, and can benefit from future improvements in imaging technology [131,212,213] or by way of opto-genetical tools [77,214], as will be discussed in Chapter 14.

GTE yields reliable reconstruction

To be able to benchmark GTE, we have generated surrogate fluorescence time series using an integrate-and-fire model with short-term synaptic depression. We also included a light scattering artifact, that makes the simulated fluorescence signal significantly more biologically and experimentally realistic than competitor approaches. We have then systematically studied the performance of GTE in Part II, indicating the robustness of the GTE measure to parametric variations, for instance in conditioning level or system size. It was indeed found that GTE is robust to the light scattering artifacts where linear measures of causality are not, which is a finding that very likely parallels the study of [189], where TE was found to be robust to cross-talk between EEG electrodes.

A common argument against pairwise causality measures is their supposed inability to distinguish between different connection motifs. It is indeed interesting to examine two cases, namely the case of *indirect causation*, differentiating the two scenarios $A \rightarrow C$ and $A \rightarrow B \rightarrow C$, and *co-causation*, in which the challenge is to differentiate $B \leftarrow A \rightarrow C$ and $B \leftrightarrow C$. It might be asked whether in particular the same-bin interactions included in GTE might not render such a differentiation impossible. While it is true that higher-order correlations in the signal are, with the exception of conditioning which depends on the whole population, not included in the GTE measure, it is precisely the purpose of conditioning to be able to exclusively analyze a subset of time points where the activity in the network is dominated by mono-synaptic, i. e. pairwise interactions. To exemplify this, we have identified, for an example network from the locally-clustered topologies (see Fig. 7.3B), each case where an co-causation motif occurred, in both the true and the reconstructed topology, and where there were no links $B \rightarrow C$ or $C \rightarrow B$. We then analyzed, in the reconstructed network, how often a link $B \rightarrow C$ was spuriously inferred, finding the correct result in 79% of the cases (11314 out of 14325). Similarly we studied the case of indirect causation, finding that in only 2954 out of 14206 cases (21%) was the incorrect link identified as positive by the reconstruction.

We further studied whether this reconstruction quality extends to chains of length 4 as well, i. e. $A \rightarrow B \rightarrow C \rightarrow D$. For the same example network we have examined all instances where such a chain existed in the structural topology, and where there were no shortcuts, i. e. where the links $A \rightarrow C$, $A \rightarrow D$ or $B \rightarrow D$ were not present. From the remaining set, in analogy to the cases presented above we examined the cases where the chain was also present in the reconstructed topology, and finally evaluated how often a spurious connection $A \rightarrow D$ was inferred. We found again a high fraction of correctly inferred non-existent links of 87% (62 out of 71 cases), confirming the high reliability of GTE in inferring connectivity motifs even in higher order.

We have also shown, based on simulated calcium fluorescence, that GTE can correctly estimate the level of clustering in a network (see Fig. 7.1C). Note that what was plotted was the full clustering coefficient, according to the definition by Garofalo [132]. There are four other generalization of the clustering index to the case of directed networks, emphasizing for instance the existence of loops or indirect links (so-called Middleman clustering). We have checked for the linear relation between true and reconstructed clustering to hold for all five of these definitions, finding a virtually identical performance in each case. This further indicates the ability of GTE to reconstruct connectivity motifs without bias. The only exception to this that we found, is the systematic

over-representation of bidirectional connections in the reconstructed networks, this being the trade-off for the reconstruction performance gain due to the inclusion in GTE of instantaneous interactions.

The only requirement for our method is that in the conditioned regions, there is still significant spiking activity; otherwise, there is little useful information about the interactions in the signal between bursts. However, we note that on the one hand, the conditioned regime does contain the early burst onset, which should prove informative while not being synchronous, and, on the other hand, activity patterns without activity between bursts is rare; see e.g. [53], where even in the 8% of cultures with epileptic-like activity patterns a number of inter-burst spikes are visible.

We have also demonstrated the increasing performance level of GTE with larger system size, contrary to the behavior found in pairwise maximum entropy models [215]. Thus, we find GTE to be applicable without restriction to experimental recordings.

Reconstruction of in vitro networks

We have applied GTE to experimental recordings of the calcium fluorescence activity of in vitro cultures with purely excitatory connections, as a first step. In cultures ranging from 9 to 12 days in vitro, we find a highly non-random but also non-local connectivity, characterized by a broadened degree distribution and an elevated clustering coefficient. Thus despite the self-organized nature of development in vitro in the absence of sensory stimuli, we find a rich topological structure.

The reconstructed topology is consistent with available results from other studies based on studying the connectivity of a small number of neurons across various networks of cultured tissue. These studies have revealed, however, a systematic over-representation (compared to the random network) of a number of clustering motifs [49, 84]. The origin of this elevated clustering cannot be reliably identified at this point — a plausible candidate mechanism of axonal guidance would be *fasciculation*, which provides a weak binding force of new axons to already existing ones [216–218], and may enhance this clustering property not only in vivo, but also in vitro.

The topology we found is also in agreement with the argument detailed above: Early firing neurons, in the experimentally observed variability of firing times, exist not in random networks but in structured networks, as now confirmed by the reconstruction.

Interestingly, the corresponding reconstruction based on the cross-correlogram between two fluorescence signals shows symptoms in parallel to the simulations, showing a local structure of connections and a high clustering coefficient (see Fig. 7.3D). This is especially interesting in the light of [174], where strongly clustered hub cells were found using a cross-correlation based algorithm.

The exact numbers for topological properties of the network depend on the fraction of links accepted as positives (see Fig. 7.4) which can a priori only be estimated from the expected numbers of links for each system from the literature. We note, however, that even if an estimate of the probability of connection was unavailable for some system, a relative comparison between networks, retaining the same number of reconstructed links would still be possible. Also we note that the reconstructed topological properties are relatively robust against change in threshold — when doubling the fraction of included links from 5% to 10% in the experimental recording of Fig. 8.1, the average clustering coefficient changed from 0.090 ± 0.010 (mean \pm standard deviation) to 0.157 ± 0.011 , whereas in a random graph, an increase by 100% would be expected. The average connection length changed only marginally from $852 \pm 463\mu\text{m}$ to $855 \pm 464\mu\text{m}$.

We have checked for consistency of the reconstructed, non-local topology by way of GFP staining (see Fig. A.1C), finding good agreement. Such a finding is also confirmed by the available literature on this subject [41, 197, 219].

Finally, we have found no evidence for the studied biological networks to be structurally

scale-free, i. e. showing a structural degree distribution following a power-law [203]. The question of whether cultured networks from dissociated neurons or from slices show a scale-free degree distribution is currently highly debated [38,51,174,204]. The varied technical and mathematical sophistication of the applied methods suggests that scale-free distribution can arise in functional connectivity of networks which are not inherently scale-free, but this remains to be tested. What we do find in the cultured networks studied in this thesis is the *small world* property, where the co-existence of long and short-range connections lead to a small average path length between any two neurons in the network [220,221].

In the QP model, only a scale free degree distribution was consistent with the exponential rise of the average firing frequency at the onset of a burst. We can now conclude that we found no such degree distribution in the structural connectivity, but hypothesize that, in agreement with the work by Bonifazi et al., the functional degree distribution is likely scale free as a result of the globally synchronous onset stimulus.

Based on the GTE reconstruction, we find hints that another factor may contribute to the appearance of the sub-network of early firing neurons and their homogenous distribution in the culture [50]: We observed a small but statistically significant correlation of $p = 0.16$ for the culture of Fig. 8.1 between in-degree of a neuron and the average length of its incoming connections. For the younger culture of Fig. 8.2 on the other hand, we found no correlation, yielding $p = -0.02$. This opens up the intriguing possibility that after a certain maturation, during which there is longer axonal and dendritic outgrowth, this subnetwork spans the entire culture due to its longer processes. The origin of these findings would be interesting to examine more in detail in the future, and indeed with GTE we now have the tool to study it.

Reconstructing inhibitory connections

In a minimal network model, we have demonstrated that GTE is in principle capable of detecting excitatory as well as inhibitory links without modification; this can be seen as a demonstration of the advantage of the model-free design. We showed that, similar to other studies using Transfer Entropy [185,186], the non-linear measure of causality used results in a measure related to the absolute value of the synaptic strength. We also showed how in principle, the two types of connection could be distinguished after the reconstruction based on the individual terms of the GTE formula.

For the case that, for whichever reason, a reconstruction of both inhibition and excitation is not possible, or if it turns out that the distinction between the two cannot be made based on GTE, we propose an alternative two-step strategy: First, all the excitatory links are reconstructed while blocking inhibition as done above, and then, after wash-out, the inhibitory links are reconstructed in an isolated manner by blocking excitatory links. In the latter case, spontaneous activity would have to be restored chemically as in [222].

We have indeed obtained, in cooperation with Javier Orlandi from the University of Barcelona, preliminary results based on simulations showing that such a two-step reconstruction can yield very good reconstruction results. In the simulated networks, which were from the non-locally clustered ensemble, both the excitatory and inhibitory links could be identified based on their respective recordings with a performance level of more than 70% each. A caveat is, however, that while the elicitation of spontaneous activity has been demonstrated in cultured slices [222], this has to our knowledge not been attempted in dissociated networks.

13.3 First firing times in vitro can be predicted based on GTE reconstruction

We have combined the results described above to test the possibility of predicting the average first firing time, in simulations as well as in cultured networks, based on the reconstructed topology. Since we found that structure restricts the compatible dynamics in vitro and vice versa, we asked whether we can indeed use the inferred topology to predict, at least to some degree, the observed dynamics.

We found a remarkably good agreement for simulations ($p = 0.76$) for an elevated number of links, which was chosen due to the sensitivity of the analytical solution to low degree numbers. For the experimental recordings, a linear correlation with $p = 0.244$ was achieved. In light of the approximations involved in both the predicted onset time, as well as in extracting the onset time from experimental recordings with low signal-to-noise ratio, this can be seen as a surprisingly high value. Statistical significance was tested and confirmed using a standard bootstrapping method.

The fact that the resulting analytics overestimate the range of expected onset times can likely be attributed at least in part to homeostatic effects, i. e. the fact that neurons which receive a high average input current down-regulate their synaptic efficacy over time [35, 209], as discussed in detail above.

The fact that we were able to close the structure–dynamics loop this way opens up the intriguing question: Is it possible to predict the response of the system to arbitrary stimuli? For instance, Shahaf et al. showed that in vitro, the temporal order of neuronal activation can be used to reliably predict the spatial location of a stimulating electrode [223]. Based on the work presented above, by making use of the analytical work and the reconstructed topology, this temporal order of activation might be predicted, thus gaining further insight into the onset dynamics in neuronal networks.

13.4 Summary

In the present work, the interplay between dynamics and network topology in in vitro networks of neurons has been studied from two different, reciprocal perspectives.

In model systems of neuronal dynamics, the expected first firing time of a given neuron relative to the population was studied and calculated analytically, in order to understand the driving forces behind the hierarchical development of bursting activity (Part I). As the firing dynamics were shown to be dependent on the network topology, we developed statistical methods to quantify causal interactions and to infer the synaptic topology from recordings of the activity. We then applied these measures to experimental recordings of the calcium fluorescence signal of cultured neurons (Part II). Further, it was shown in Part III that the two perspectives can be joined to predict aspects of the dynamics in simulated as well as in experimental recordings based on the inferred network topology.

We hope that the findings presented in this thesis will help to further illuminate the development of structure in neuronal networks in vitro as well as in vivo, by aiding in the determination of the individual underlying mechanisms that shape the circuitry in the brain, and ultimately neuronal function in the brain.

Chapter 14

Outlook

In the future, we hope that the results of this thesis will be useful in the study of neuronal dynamics and the underlying network structure. We here discuss the potential advantages of other technologies to further develop the current work.

14.1 Technological advances

Progress in a number of fields promises to open up the possibility to probe neuronal systems on an unprecedented temporal and spatial resolution, facilitating the analysis of the structure as well as the dynamical properties of neuronal networks.

For instance, advances in two-photon imaging technology may allow for the recording of calcium imaging signals at a frame rate comparable, or even higher, than the rates in this work [160, 196]. Two-photon microscopy is not constrained to the surface, but is able to record from cells inside of neuronal tissue [194, 195]. Further, two-photon microscopy is a versatile tool that can be used to stimulate as well as record from neuronal populations [202] or can be used to guide patch-clamp recordings to calculate functional connectivities [81].

Another promising field for the analysis of neuronal populations is opto-genetics, where genetically encoded fluorescent markers are attached to living cells [214, 224, 225], eliminating the problem of bleaching and resulting cell damage as discussed in the present work. Opto-genetical tools have indeed been used to reveal connectivity in *Drosophila* [226], or to identify the location of synapses [227], which has the potential to be used in conjunction with statistical inference methods and anatomical data to reconstruct synaptic connectivities [228].

The combination of the two fields can result in obvious synergies, which are beginning to be exploited to extract information about neuronal circuitry [77, 229]. Already today for instance, it is possible to record from a large part of brain of a genetically modified, behaving zebrafish [230].

14.2 Future work

Presently, the key challenges which remain regarding the reconstruction of synaptic connectivity based on GTE are the extension to inhibition and the improved experimental benchmarking of the formalism. For the former, inhibition has already been studied in small model networks in this thesis (see Section 9.2) and preliminary analysis of a two-step approach has been discussed above. In the future, a deeper analysis of the differentiation and the analysis of experimental recordings without blocking of inhibitory synapses would be of great value for applications, as it would allow to reconstruct the complete connectivity map of a given neuronal network. It would

also allow for a new analysis of the firing time dynamics, because while the onset of a burst is dominated by excitatory synaptic currents, inhibition has been shown to also play a role in the temporal organization of activity [43, 122].

In a next step, an improved assessment of the prediction accuracy for experimental recording would be desirable. A number of strategies might be employed. The most precise ground-truth test would be attained by recording the fluorescence activity of a culture of neurons, and then to patch, in as many iterations as possible for a given network, two neurons to determine whether they are connected mono-synaptically. To maximize the statistical information gained from such test, the selection of neurons should be unbiased and random. However a biased course of action, aiming to confirm the strongest GTE connections, is also feasible. Such a strategy might also make use of parallel-patch techniques which are starting to be available [231].

In the short term, a simple experiment might be to mechanically divide a culture of neurons, and to assess the GTE connectivity before and after the cut. For a given number of reconstructed links C' , resulting in an average probability of a reconstructed connection $p' = C'/N(N-1)$, an estimate of the fraction of false positives FP (see Eq. 6.9) is given simply by the number of spuriously reconstructed links joining the two parts of the culture. Together with an estimate of the true probability of connection p , an ROC curve can be calculated:

$$\text{TP} = \frac{p' - \text{FP}(1-p)}{p}. \quad (14.1)$$

Since the likelihood of a spurious link across populations is for all practical applications smaller than a spurious link within a population, the resulting ROC curve is a lower bound to the true reconstruction performance.

In the long term, due to the generality of the measure we have proposed, we hope that it will find applications in reconstructing functional connectivity in a range of systems. For the application of synaptic connectives among neuronal cell populations, the combination of imaging technologies with opto-genetic markers is a promising route, for which the developed measure should be ideal. This is especially the case since, even though we have restricted ourselves in this work to the simpler case of planar *in vitro* networks, GTE can be immediately applied to the analysis of *in vivo* recordings. For *in vivo* systems which do not show bursting dynamics, we expect the reconstruction performance to be actually higher than shown for the *in vitro* systems. The strength of the conditioning approach is then still in exposing the state dependency of the effective connectivity; *in vivo*, this conditioning might be applied not only regarding the fluorescence level but also regarding the type of stimulus provided.

We hope the formalisms developed in this thesis will be helpful to study the topological structure in neuronal networks in more detail, starting with the *in vitro* situation. A particularly interesting aspect would be the analysis of pathological connectivity profiles, which are currently studied only in the context of functional interactions [78, 232], and the role of hub neurons characterized by their high in-degree — similar to the early firing neurons studied in this thesis — which have been proposed to play a major role in epilepsy [233].

Another key challenge of neuroscience where we hope to be able to contribute in the future is the study of the development of neuronal circuits. In this thesis, we have considered short-term synaptic plasticity, since it is a prominent feature of the firing dynamics of neuronal networks [128, 149]. In the brain, the network connectivity is evolving over time, shaped not only by transient effects but by a co-evolution of structure and dynamics that has been called symbiotic [234, 235]. For instance, after the initial circuitry has been set up by processes of axon guidance exogenous to the neurons, the development of pattern formation in the early visual system depends critically on waves of spontaneous activity to act as guiding system to align and correlate the synaptic connections of neighboring cells [236]. In the adult brain, the synaptic connectivity in the brain is

continued to be shaped by dynamical processes dependent on learning and experiences [234,237, 238]. For instance by studying the statistical properties of a neuronal network during growth or after changes due to learning, we hope that the techniques developed in this thesis will aid in the study of the dynamical nature of the structural connectivity.

Part V

Appendix

Appendix A

Experimental preparation

In this chapter an overview is given of the experimental procedure to obtain, from in vitro cell cultures, the calcium imaging recordings studied earlier. The experimental preparation follows in general the process outlined in [63], and a detailed description and list of chemicals is given in [169].

The experimental recordings analyzed in this thesis were from dissociated cortical neurons of Sprague-Dawley embryonic rats at 19 days of development. Neurons were dissociated and plated on glass, which had been coated with Poly-L-Lysine to enhance neuronal adhesion on the glass surface. Thus, after 5 days of incubation, a typical cell density between $500\text{--}700\text{mm}^2$ is achieved.

It is known that spontaneous activity in the networks begins to occur at 3–4 days in vitro (DIV) and that the *GABA switch*, the point where the GABAergic synapses change from being excitatory to inhibitory, occurs at 6–7 DIV. [54, 63, 239] The networks studied for this thesis are recorded at 9–12 DIV. At this point, neurons are well connected, with connection lengths in the range of millimeters (see Fig. A.1C).

For recording calcium fluorescence, cultures were incubated for 60 minutes in the presence of 0.4% of the calcium dye Fluo-4-AM (Invitrogen), after which the dye was washed off and the medium replaced. Next, the culture was placed in a Zeiss inverted microscope equipped with a Hamamatsu Orca Flash 2.8 high-speed CMOS camera, allowing for the acquisition of a field of view of approximately $2.2 \times 1.1\text{mm}^2$ at a resolution of 648×312 with 8 bit greyscale colors. Images were initially recorded at 100Hz and then resampled to 50Hz which — besides the benefits in signal-to-noise ratio — is preferential as it allows for a better comparison with other work using time widths of 20ms [24, 106, 162, 171, 197].

This setting allowed for the stable recording of activity for up to 60 minutes without bleaching or drift in the fluorescence signal (see Fig. A.1A). Note that in this recording, as Part II of the present work, activity which is based on excitatory synapses only. Thus, before recording, inhibitory synaptic transmission had to be blocked by applying a concentration of $40\mu\text{M}$ bicuculline, a GABA_A antagonist. This blocking lead typically increased the fluorescence amplitude of burst events, while lowering the burst rate, as shown in the example traces in Fig. A.1B. This result is consistent with [54] — non-withstanding a large variability across cultures, for instance, in about 29% of cultures (2 out of 7) studied in the context of the present work, the burst rate was actually higher with inhibition blocked than before.

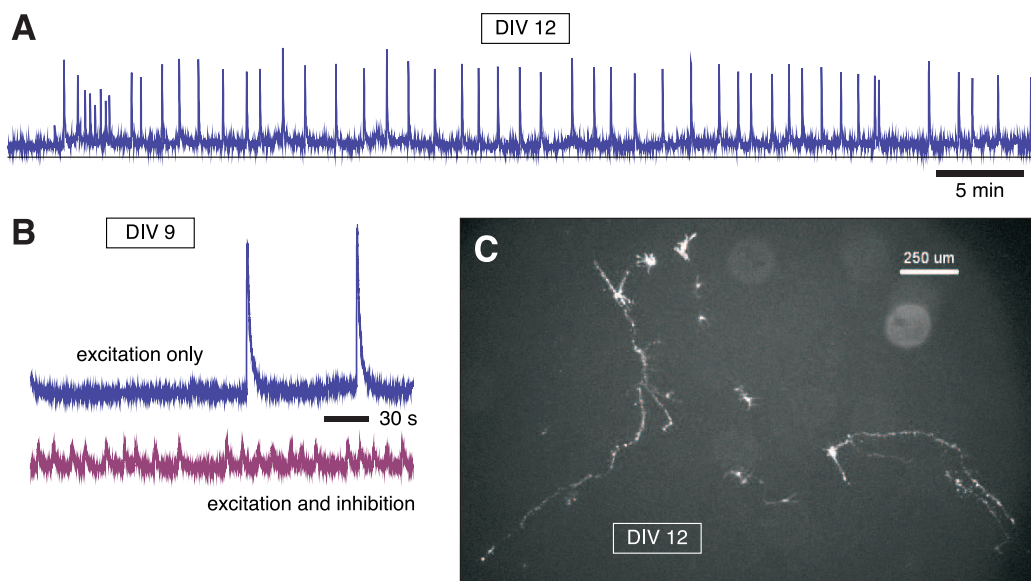


Figure A.1: (A) Long-term stability over 40 minutes of an example experimental recording at 12 DIV (inhibition blocked) at 100 frames/second. (B) Switching dynamics with (top panel) and before adding bicuculline to block inhibitory synapses (bottom panel), of an example culture recorded at 9 DIV. (C) GFP-labelled example cells demonstrating typical axon lengths at 12 DIV. Figure reproduced with permission from [169].

Appendix B

Estimators of probability densities

The causality measures discussed in Part II depend by definition on probability distributions, which are a priori unknown and must be inferred from samples data points. The problem of estimating probability densities is a fundamental one, and has received much attention, especially in the context of estimating Mutual Information and Transfer Entropy [116, 182, 240–242].

For the common case of a signal which is defined in discrete time points (determined by a constant sampling rate) and a continuous amplitude such as membrane voltage or fluorescence signal, there are two types of strategies applicable: Either the continuous amplitude is discretized into a range of intervals, typically of equal width, and a correction (*de-biasing*) is applied afterwards to estimate the true value of the entropy [243–245] or, secondly, indirect methods to estimate entropy based on statistical properties can be employed (*binless estimators*, see discussion below).

De-biasing is, for the application presented in this thesis, of lesser interest as for the reconstruction the absolute values of Transfer Entropy are irrelevant — what is relevant is the rank ordering of the values corresponding to each putative link in the network. We have checked however, that a straight-forward extrapolation of TE estimates to infinite sample number according to [246] has no significant effect on reconstruction performance.

B.1 Binless estimators

We will here consider a class of binless estimators of entropy which are based on nearest-neighbor statistics,¹ originally developed by Kozachenko and Leonenko in 1987 [248]. For a number of S samples, the set $\{\lambda_i\}$, where λ_i is the distance of each sample point x_i to its nearest neighbor, yields an asymptotically unbiased expression of the entropy [248, 249].

$$H \approx \frac{1}{S} \sum_{i=1}^S \log_2 \lambda_i + \log_2 [2(S-1)] + \frac{\gamma}{\log 2} \quad (\text{B.1})$$

Here, γ denotes the Euler-Mascheroni constant.

Estimating Mutual Information requires the estimation of the difference of two entropies, which is sensitive to small errors in estimation. To address this problem, Kraskov et al. developed an estimator based on k nearest neighbors [250]. It is based on the maximum norm in the joint space $Z = (X, Y)$, defined as $\|z - z'\| = \max\{\|x - x'\|, \|y - y'\|\}$. Let then be $\epsilon(i)/2$ the distance from sample point z_i to its k th neighbor, and $\epsilon_x(i)/2$ and $\epsilon_y(i)/2$ the distances, projected on the respective subspaces. We then denote by $n_x(i)$ the number of samples whose distance from x_i

¹The nearest neighbor search was implemented using the exact limit of the Fast Library for Approximate Nearest Neighbors (FLANN) of Muja et al. [247].




Figure B.1: Comparison of GTE reconstruction performance using different estimators of probability densities. For a simulation from the non-locally clustered network ensemble ($CC = 0.5$) as in Fig. 7.1A, the reconstruction performance is plotted as a function of sample number, in a linear scale (left panel) and a log-linear scale (right panel).

is less than $\epsilon(i)/2$, and analogously for the subspace Y . Kraskov et al. show that an estimate of the Mutual Information can be derived using the digamma function $\psi(x)$, with angled brackets indicating the average across sample points.

$$I(X, Y) = \psi(k) - \langle \psi(n_x + 1) + \psi(n_y + 1) \rangle + \psi(S) \quad (\text{B.2})$$

Intuitively, this estimator is preferable to estimate Mutual Information as, due to the projection of nearest-neighbor distances on the individual subspaces, the remaining bias is not any more different as it would be when applying a Leonenko-type estimator to samples of different dimensionality.

A further refinement in the context of nearest-neighbor strategies for estimating Transfer Entropy was proposed in 2007 by Frenzel and Pompe [251]. Their estimate for Partial Mutual Information $I(A; B|C)$ can be directly applied to Transfer Entropy, as can be easily shown:

$$I(y_{t-1}^{(k)}; x_t | x_{t-1}^{(k)}) = I(x_t; x_{t-1}^{(k)}, y_{t-1}^{(k)}) - I(x_t; x_{t-1}^{(k)}) = \text{TE}_{Y \rightarrow X}. \quad (\text{B.3})$$

Partial Mutual Information (PMI) is estimated by the following equation, making use of the harmonic numbers $h(n) = -\sum_{i=1}^n i^{-1}$, where N_{\dots} are again the respective projections on the sub-spaces [251].

$$I(X; Y|Z) = \langle h(N_{xz}(t)) + h(N_{yz}(t)) - h(N_z(t)) \rangle - h(k-1) \quad (\text{B.4})$$

B.2 Application to the estimation of GTE

To estimate GTE of Markov order k from fluorescence time series, the joint probability densities must be sampled in a sample space of dimensionality $2k + 1$. As the most robust performance in the present case is achieved for Markov order $k = 2$ (see Part II), the sample space is 5-dimensional.

In principle, it is known that the resulting probability distributions are dependent on the estimation technique used [182, 241]. For the purpose of network reconstructions, only the resulting rank ordering of estimated values (rather than the absolute value) is relevant; nonetheless here we analyze the influence of different estimation techniques on the resulting reconstruction quality.

We have systematically studied the achieved performance level for each of the estimators described above (see Fig. B.1). At a fraction of 10% of false positives, 10% of true positives

correspond to the performance of a random choice, and both the Leonenko (Eq. B.1) and Kraskov (Eq. B.2) are below this level in performance. Of the binless estimators, only the one according to Frenzel et al. (Eq. B.4) is achieving performance levels higher than random choice. None of them however is as data-efficient as the plug-in estimate using binning, indicating the difficulty of estimating probability densities of high dimensionality and limited sample number.

As the estimation of Transfer Entropy relies on the difference of individual estimates, it might be interesting to compare the estimates with estimates based on appropriately shuffled data (bootstrapping), generating new data sets while maintaining the auto-correlation function [252,253]. Such a procedure has the advantage of getting rid of the problem of estimating densities with different dimensionality entirely. It turns out however that such a procedure has no significant benefit in terms of the performance level (not shown).

It is also worth pointing out, that the computational time needed for the Leonenko estimator was higher than binning by an order of magnitude, and higher by three orders of magnitude for the Kraskov and Frenzel estimators.

In summary, due to the high dimensionality and the finite number of experimentally available sample points, we have with the exception of this chapter relied on the plug-in estimation method, which derives its high performance from the transformation of the fluorescence time series in a quasi-symbolic representation (due to low bin number). Such a transformation has been shown to be highly efficient in terms of the needed number of samples to achieve a certain accuracy in the TE estimate [190].

List of figures

1.1	Bright-field image of a fully grown in vitro network of hippocampal rat neurons	2
1.2	Previous studies of burst onset dynamics in vitro	4
1.3	Patch-clamp recordings by Song et al. to extract neuronal connectivity	6
2.1	Full and aborted bursts in a discrete-time integrate-and-fire model	13
2.2	The effect of topological structure on onset firing probability	15
2.3	Topological structure enhances variability in onset first firing probability	17
2.4	Rise before burst is exponential and can be predicted analytically, in a discrete-time I&F model	18
2.5	Illustration of first and second order mean-field approximation	20
2.6	The expected first firing time in a random graph can be predicted analytically	22
2.7	First firing time in a locally clustered graph can be approximated analytically	24
2.8	Dependence of the predicted first firing time on the order of topological expansion	25
3.1	A typical realization of the quorum percolation model	28
3.2	Gaussian degree distribution combined with a power-law right tail	30
3.3	Dependence of the burst onset in the QP model on degree correlations	31
4.1	Spiking activity and bursting in networks with random and local topology	35
4.2	Analytical prediction of the first spike time in a leaky I&F model	38
4.3	Prediction of the first firing time at the onset of a burst in a random graph of leaky integrate-and-fire neurons	40
4.4	Prediction of the first firing time at the onset of a burst in a locally structured graph of leaky integrate-and-fire neurons	41
5.1	Calcium imaging recordings from in vitro cultures and simulations	47
5.2	Radius of the light scattering artifact used in simulations	49
5.3	Spiking activity showing bursting for a range of clustering coefficients	50
6.1	Effective connectivity depends on the dynamical state of the network	56
7.1	Reconstruction performance in non-locally clustered topologies	60
7.2	Reconstruction quality of the complementary topological properties	61
7.3	Reconstruction performance in locally clustered topologies	62
7.4	Dependence of reconstructed average degree and clustering coefficient on the threshold	63
7.5	Reconstruction performance across causality measures and conditioning levels in the non-locally clustered ensemble	64
7.6	Reconstruction performance across causality measures and conditioning levels in the locally clustered ensemble	66

7.7	Reconstruction performance of the Bayesian approach by Mishchenko et al.	68
7.8	Reconstruction performance of individual extensions of Transfer Entropy, and due to sample number	70
7.9	Positive-prediction curves individual extensions based on Transfer Entropy	71
7.10	Dependence of GTE reconstruction quality on conditioning level and system size	71
8.1	GTE reconstruction of the excitatory connectivity of an in vitro network at 12 DIV	74
8.2	GTE reconstruction of the excitatory connectivity of an in vitro network at 9 DIV	76
8.3	Dependency of reconstructed topological quantities on dynamical state in experimental recordings and simulations	77
9.1	Reconstruction for a minimal network model with both excitatory and inhibitory links	80
9.2	Separation of the GTE scores from the analysis of Fig. 9.1 into different terms	81
11.1	Fitting alpha functions to average bursts to determine onset times	88
12.1	Comparison between measured and predicted onset time in a non-locally clustered, simulated network	92
12.2	Comparison between measured and predicted onset time in a cultured network at 9 DIV of dissociated neurons	93
A.1	Influence of bleaching and bicuculline on recorded activity, and axon outgrowth at typical ages of cultured networks	112
B.1	Comparison of GTE reconstruction performance using different estimators of probability densities	114

Bibliography

- [1] Finger S (1994) *Origins of neuroscience: A history of explorations into brain function*. New York: Oxford University Press.
- [2] Ramón y Cajal S (1967) The structure and connexions of neurons. In: *Nobel Lectures: Physiology or medicine 1901-1921*, Amsterdam: Elsevier Publishing Company.
- [3] Lent R, Azevedo FAC, Andrade-Moraes CH, Pinto AVO (2011) How many neurons do you have? Some dogmas of quantitative neuroscience under revision. *European Journal of Neuroscience* 35: 1–9.
- [4] Mai JK, Paxinos G (2011) *The human nervous system*. London: Academic Press, Inc., 3rd edition.
- [5] Herculano-Houzel S (2012) The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 109: 10661–10668.
- [6] Hodgkin AL, Huxley AF (1952) Currents carried by sodium and potassium ions through the membrane of the giant axon of *Loligo*. *Journal of Physiology* 116: 449–472.
- [7] Hodgkin AL, Huxley AF (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology* 117: 500–544.
- [8] Shoham S, O'Connor DH, Segev R (2006) How silent is the brain: Is there a “dark matter” problem in neuroscience? *Journal of Comparative Physiology A* 192: 777–784.
- [9] Brecht M, Sakmann B (2002) Dynamic representation of whisker deflection by synaptic potentials in spiny stellate and pyramidal cells in the barrels and septa of layer 4 rat somatosensory cortex. *The Journal of physiology* 543: 49–70.
- [10] Brecht M, Roth A, Sakmann B (2003) Dynamic receptive fields of reconstructed pyramidal cells in layers 3 and 2 of rat somatosensory barrel cortex. *The Journal of physiology* 553: 243–265.
- [11] Braitenberg V, Schüz A (1998) *Cortex: Statistics and geometry of neuronal connectivity*. New York: Springer, 2nd edition.
- [12] Attwell D, Laughlin SB (2001) An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow and Metabolism* 21: 1133–1145.
- [13] Bullmore E, Sporns O (2012) The economy of brain network organization. *Nature Reviews Neuroscience* 13: 336–349.
- [14] James W (1890) *The principles of psychology*. New York/London: Holt and Macmillan.

- [15] Gazzaniga MS, Ivry RB, Mangun GR (2004) *Cognitive neuroscience: The biology of the mind*. Cambridge: MIT Press, 3rd edition.
- [16] Shallice T (1988) *From neuropsychology to mental structure*. Cambridge University Press.
- [17] Pessoa L, Gutierrez E, Bandettini PA, Ungerleider LG (2002) Neural correlates of visual working memory: fMRI amplitude predicts task performance. *Neuron* 35: 975–987.
- [18] Johnson SC, Baxter LC, Wilder LS, Pipe JG, Heiserman JE, et al. (2002) Neural correlates of self-reflection. *Brain* 125: 1808–1814.
- [19] Treadway MT, Buckholz JW, Cowan RL, Woodward ND, Li R, et al. (2012) Dopaminergic mechanisms of individual differences in human effort-based decision-making. *Journal of Neuroscience* 32: 6170–6176.
- [20] Shepherd GM (1988) *The synaptic organization of the brain*. Oxford University Press.
- [21] Gupta A, Wang Y, Markram H (2000) Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. *Science* 287: 273–278.
- [22] Perin R, Berger TK, Markram H (2011) A synaptic organizing principle for cortical neuronal groups. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 108: 5419–5424.
- [23] Ko H, Hofer SB, Pichler B, Buchanan KA, Sjöström PJ, et al. (2011) Functional specificity of local synaptic connections in neocortical networks. *Nature* 473: 87–91.
- [24] Feinerman O, Rotem A, Moses E (2008) Reliable neuronal logic devices from patterned hippocampal cultures. *Nature Physics* 4: 967–973.
- [25] Battaglia D, Brunel N, Hansel D (2007) Temporal decorrelation of collective oscillations in neural networks with local inhibition and long-range excitation. *Physical Review Letters* 99: 238106.
- [26] Jahnke S, Memmesheimer RM, Timme M (2009) How chaotic is the balanced state? *Frontiers in Computational Neuroscience* 3: 1–15.
- [27] Monteforte M, Wolf F (2010) Dynamical entropy production in spiking neuron networks in the balanced state. *Physical Review Letters* 105: 268104.
- [28] Battaglia D, Hansel D (2011) Synchronous chaos and broad band gamma rhythm in a minimal multi-layer model of primary visual cortex. *PLoS Computational Biology* 7: e1002176.
- [29] Golubitsky M, Stewart I (2003) *The symmetry perspective: From equilibrium to chaos in phase space and physical space*. Basel: Birkhäuser Verlag.
- [30] Neves FS, Timme M (2012) Computation by switching in complex networks of states. *Physical Review Letters* 109: 018701.
- [31] Timme M, Wolf F, Geisel T (2002) Coexistence of regular and irregular dynamics in complex networks of pulse-coupled oscillators. *Physical Review Letters* 89: 258701.
- [32] Battaglia D, Witt A, Wolf F, Geisel T (2012) Dynamic effective connectivity of inter-areal brain circuits. *PLoS Computational Biology* 8: e1002438.
- [33] Dichter MA (1978) Rat cortical neurons in cell culture: Culture methods, cell morphology, electrophysiology, and synapse formation. *Brain Research* 149: 279–293.

- [34] Kamioka H, Maeda E, Jimbo Y, Robinson HP, Kawana A (1996) Spontaneous periodic synchronized bursting during formation of mature patterns of connections in cortical cultures. *Neuroscience Letters* 206: 109–112.
- [35] Desai NS, Rutherford LC, Turrigiano GG (1999) Plasticity in the intrinsic excitability of cortical pyramidal neurons. *Nature Neuroscience* 2: 515–520.
- [36] Lau P, Bi G (2005) Synaptic mechanisms of persistent reverberatory activity in neuronal networks. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 102: 10333–10338.
- [37] Wagenaar DA, Pine J, Potter SM (2006) An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neuroscience* 7: 1–18.
- [38] Srinivas KV, Jain R, Saurav S, Sikdar SK (2007) Small-world network topology of hippocampal neuronal network is lost, in an in vitro glutamate injury model of epilepsy. *The European Journal of Neuroscience* 25: 3276–3286.
- [39] Feinerman O, Segal M, Moses E (2007) Identification and dynamics of spontaneous burst initiation zones in unidimensional neuronal cultures. *Journal of Neurophysiology* 97: 2937–2948.
- [40] Eckmann JP, Feinerman O, Gruendlinger L, Moses E, Soriano J, et al. (2007) The physics of living neural networks. *Physics Reports* 449: 54–76.
- [41] Erickson J, Tooker A, Tai YC, Pine J (2008) Caged neuron MEA: A system for long-term investigation of cultured neural network connectivity. *Journal of Neuroscience Methods* 175: 1–16.
- [42] Gandolfo M, Maccione A, Tedesco M, Martinoia S, Berdondini L (2010) Tracking burst patterns in hippocampal cultures with high-density CMOS-MEAs. *Journal of Neural Engineering* 7: 056001.
- [43] Baltz T, de Lima AD, Voigt T (2010) Contribution of GABAergic interneurons to the development of spontaneous activity patterns in cultured neocortical networks. *Frontiers in Cellular Neuroscience* 4: 1–17.
- [44] Huettner JE, Baughman RW (1986) Primary culture of identified neurons from the visual cortex of postnatal rats. *Journal of Neuroscience* 6: 3044–3060.
- [45] Shefi O, Golding I, Segev R, Ben-Jacob E, Ayali A (2002) Morphological characterization of in vitro neuronal networks. *Physical Review E* 66: 021905.
- [46] van Pelt J, Corner MA, Wolters PS, Rutten WLC, Ramakers GJA (2004) Longterm stability and developmental changes in spontaneous network burst firing patterns in dissociated rat cerebral cortex cell cultures on multielectrode arrays. *Neuroscience Letters* 361: 86–89.
- [47] Marom S, Shahaf G (2002) Development, learning and memory in large random networks of cortical neurons: Lessons beyond anatomy. *Quarterly Reviews of Biophysics* 35: 63–87.
- [48] Faisal AA (2012) Noise in neurons and other constraints. In: Le Novère N, editor, *Computational Systems Neurobiology*, Dordrecht: Springer Netherlands. pp. 227–257.
- [49] Holmgren C, Harkany T, Svennenfors B, Zilberter Y (2003) Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *Journal of Physiology* 551: 139–153.

- [50] Eckmann JP, Jacobi S, Marom S, Moses E, Zbinden C (2008) Leader neurons in population bursts of 2D living neural networks. *New Journal of Physics* 10: 1–19.
- [51] Eytan D, Marom S (2006) Dynamics and effective topology underlying synchronization in networks of cortical neurons. *Journal of Neuroscience* 26: 8465–8476.
- [52] Maeda E, Robinson H, Kawana A (1995) The mechanisms of generation and propagation of synchronized bursting in developing networks of cortical neurons. *Journal of Neuroscience* 15: 6834–6845.
- [53] Arnold F, Hofmann F, Bengtson C, Wittmann M, Vanhoutte P, et al. (2005) Microelectrode array recordings of cultured hippocampal networks reveal a simple model for transcription and protein synthesis-dependent plasticity. *Journal of Physiology* 564: 3–19.
- [54] Cohen E, Ivenshitz M, Amor-Baroukh V, Greenberger V, Segal M (2008) Determinants of spontaneous activity in networks of cultured hippocampus. *Brain Research* 1235: 21–30.
- [55] Jacobi S, Soriano J, Segal M, Moses E (2009) BDNF and NT-3 increase excitatory input connectivity in rat hippocampal cultures. *The European Journal of Neuroscience* 30: 998–1010.
- [56] Pasquale V, Martinoia S, Chiappalone M (2010) A self-adapting approach for the detection of bursts and network bursts in neuronal cultures. *Journal of Computational Neuroscience* 29: 213–229.
- [57] Kanagasabapathi TT, Massobrio P, Barone RA, Tedesco M, Martinoia S, et al. (2012) Functional connectivity and dynamics of cortical–thalamic networks co-cultured in a dual compartment device. *Journal of Neural Engineering* 9: 036010.
- [58] Kleindienst T, Winnubst J, Roth-Alpermann C (2011) Activity-dependent clustering of functional synaptic inputs on developing hippocampal dendrites. *Neuron* 72: 1012–1024.
- [59] Takano H, McCartney M, Ortinski PI, Yue C, Putt ME, et al. (2012) Deterministic and stochastic neuronal contributions to distinct synchronous CA3 network bursts. *Journal of Neuroscience* 32: 4743–4754.
- [60] Mazzone A, Broccard FD, Garcia-Perez E, Bonifazi P, Ruaro ME, et al. (2007) On the dynamics of the spontaneous activity in neuronal networks. *PLoS ONE* 2: e439.
- [61] Beggs JM, Plenz D (2004) Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures. *Journal of Neuroscience* 24: 5216–5229.
- [62] Ham MI, Bettencourt LM, McDaniel FD, Gross GW (2008) Spontaneous coordinated activity in cultured networks: Analysis of multiple ignition sites, primary circuits, and burst phase delay distributions. *Journal of Computational Neuroscience* 24: 346–357.
- [63] Soriano J, Martinez MR, Tlustý T, Moses E (2008) Development of input connections in neural cultures. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 105: 13758–13763.
- [64] Zbinden C (2011) Leader neurons in leaky integrate and fire neural network simulations. *Journal of Computational Neuroscience* 31: 285–304.
- [65] Golomb D, Hansel D (2000) The number of synaptic inputs and the synchrony of large, sparse neuronal networks. *Neural Computation* 12: 1095–1139.

- [66] Brunel N (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience* 8: 183–208.
- [67] Belykh I, De Lange E, Hasler M (2005) Synchronization of bursting neurons: What matters in the network topology. *Physical Review Letters* 94: 188101.
- [68] Chavez M, Besserve M, Le Van Quyen M (2011) Dynamics of excitable neural networks with heterogeneous connectivity. *Progress in Biophysics and Molecular Biology* 105: 29–33.
- [69] Zhao L, Bryce BI, Netoff T, Nykamp DQ (2011) Synchronization from second order network connectivity statistics. *Frontiers in Computational Neuroscience* 5: 1–16.
- [70] Maheswaranathan N, Ferrari S, VanDongen AM, Henriquez CS (2012) Emergent bursting and synchrony in computer simulations of neuronal cultures. *Frontiers in Computational Neuroscience* 6: 1–11.
- [71] Gosak M, Markovič R, Marhl M (2012) The role of neural architecture and the speed of signal propagation in the process of synchronization of bursting neurons. *Physica A: Statistical Mechanics and its Applications* 391: 2764–2770.
- [72] Schwab DJ, Bruinsma RF, Feldman JL, Levine AJ (2010) Rhythmogenic neuronal networks, emergent leaders, and k-cores. *Physical Review E* 82: 051911.
- [73] Bethe H (1935) Statistical theory of superlattices. *Proceedings of the Royal Society of London Series A, Mathematical and Physical Sciences* 150: 552–575.
- [74] Weiss P (1948) The application of the Bethe-Peierls method to ferromagnetism. *Physical Review* 74: 1493–1504.
- [75] Sporns O, Kötter R (2004) Motifs in brain networks. *PLoS Biology* 2: e369.
- [76] Sporns O, Tononi G, Kötter R (2005) The human connectome: A structural description of the human brain. *PLoS Computational Biology* 1: e42.
- [77] Arenkiel BR, Ehlers MD (2009) Molecular genetics and imaging technologies for circuit-based neuroanatomy. *Nature* 461: 900–907.
- [78] Gong G, Rosa-Neto P, Carbonell F, Chen ZJ, He Y, et al. (2009) Age- and gender-related differences in the cortical anatomical network. *Journal of Neuroscience* 29: 15684–15693.
- [79] Lichtman JW, Denk W (2011) The big and the small: Challenges of imaging the brain's circuits. *Science* 334: 618–623.
- [80] Wedeen VJ, Rosene DL, Wang R, Dai G, Mortazavi F, et al. (2012) The geometric structure of the brain fiber pathways. *Science* 335: 1628–1634.
- [81] Kwan AC, Dan Y (2012) Dissection of cortical microcircuits by single-neuron stimulation in vivo. *Current Biology* : In press.
- [82] White JG, Southgate E, Thomson JN, Brenner S (1986) The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Philosophical transactions of the Royal Society of London Series B, Biological sciences* 314: 1–340.
- [83] Lakadamyali M (2012) High resolution imaging of neuronal connectivity. *Journal of Microscopy* : In press.

- [84] Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB (2005) Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biology* 3: e68.
- [85] Sigworth FJ, Neher E (1980) Single Na⁺ channel currents observed in cultured rat muscle cells. *Nature* 277: 447–449.
- [86] Neher E, Sakmann B (1984) Patch clamp techniques for studying ionic channels in excitable membranes. *Annual Review in Physiology* 72: 455–472.
- [87] Bandettini PA (2012) Twenty years of functional MRI: The science and the stories. *NeuroImage* 62: 575–588.
- [88] Friston KJ (1994) Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping* 2: 56–78.
- [89] Ferguson MA, Anderson JS (2012) Dynamical stability of intrinsic connectivity networks. *NeuroImage* 59: 4022–4031.
- [90] Gallos LK, Sigman M, Makse HA (2012) The conundrum of functional brain networks: Small-world efficiency or fractal modularity. *Frontiers in Physiology* 3: 1–9.
- [91] Narayana S, Laird AR, Tandon N, Franklin C, Lancaster JL, et al. (2012) Electrophysiological and functional connectivity of the human supplementary motor area. *NeuroImage* 62: 250–265.
- [92] Lowell SY, Reynolds RC, Chen G, Horwitz B, Ludlow CL (2012) Functional connectivity and laterality of the motor and sensory components in the volitional swallowing network. *Experimental Brain Research* 219: 85–96.
- [93] Matthews PM, Filippini N, Dounaud G (2012) Brain structural and functional connectivity and the progression of neuropathology in Alzheimer's disease. *Journal of Alzheimer's Disease* 30: 1–10.
- [94] Baliki MN, Petre B, Torbey S, Herrmann KM, Huang L, et al. (2012) Corticostriatal functional connectivity predicts transition to chronic back pain. *Nature Neuroscience* 15: 1117–1119.
- [95] Saab CY (2012) Pain-related changes in the brain: Diagnostic and therapeutic potentials. *Current Biology* : In press.
- [96] Hawellek DJ, Hipp JF, Lewis CM, Corbetta M, Engel AK (2011) Increased functional connectivity indicates the severity of cognitive impairment in multiple sclerosis. *Proceedings Of The National Academy Of Sciences Of The United States Of America* : 1–6.
- [97] Emiliano S, Giampaolo V, Daniela M, Nicola P, Alfonso C, et al. (2012) Cerebro-cerebellar functional connectivity profile of an epilepsy patient with periventricular nodular heterotopia. *Epilepsy Research* 4751: 1–4.
- [98] Boly M, Massimini M, Garrido MI, Gosseries O, Noirhomme Q, et al. (2012) Brain connectivity in disorders of consciousness. *Brain Connectivity* 2: 1–10.
- [99] Ovadia-Caro S, Nir Y, Soddu A, Ramot M, Hesselmann G, et al. (2012) Reduction in inter-hemispheric connectivity in disorders of consciousness. *PLoS ONE* 7: e37238.
- [100] Helmstaedter M, Briggman KL, Denk W (2011) High-accuracy neurite reconstruction for high-throughput neuroanatomy. *Nature Neuroscience* 14: 1081–1088.

- [101] Bock DD, Lee WCA, Kerlin AM, Andermann ML, Hood G, et al. (2011) Network anatomy and in vivo physiology of visual cortical neurons. *Nature* 471: 177–184.
- [102] Shandilya SG, Timme M (2011) Inferring network topology from complex dynamics. *New Journal of Physics* 13: 013004–013016.
- [103] Biffi E, Menegon A, Regalia G, Maida S, Ferrigno G, et al. (2011) A new cross-correlation algorithm for the analysis of in vitro neuronal network activity aimed at pharmacological studies. *Journal of Neuroscience Methods* 199: 321–327.
- [104] Truccolo W, Eden U, Fellows MR, Donoghue JP, Brown EN (2005) A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology* 93: 1074–1089.
- [105] Pillow J, Shlens J, Paninski L, Sher A, Litke A (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454: 995–1000.
- [106] Mishchencko Y, Vogelstein JT, Paninski L (2011) A Bayesian approach for inferring neuronal connectivity from calcium fluorescent imaging data. *Annals of Applied Statistics* 5: 1229–1261.
- [107] Quiroga RQ, Panzeri S (2009) Extracting information from neuronal populations: Information theory and decoding approaches. *Nature Reviews Neuroscience* 10: 173–185.
- [108] Ostwald D, Bagshaw AP (2011) Information theoretic approaches to functional neuroimaging. *Magnetic Resonance Imaging* 29: 1417–1428.
- [109] van Ooyen A (2011) Using theoretical models to analyse neural development. *Nature Reviews Neuroscience* 12: 311–326.
- [110] Fairhall A, Shea-Brown E, Barreiro A (2012) Information theoretic approaches to understanding circuit function. *Current Opinion in Neurobiology* 22: 1–7.
- [111] Russell B (1912) On the notion of cause. *Proceedings of the Aristotelian Society* 1: 1–26.
- [112] Wiener N (1956) The theory of prediction. In: Beckenbach EF, editor, *Modern Mathematics for Engineers*, McGraw-Hill, New York. pp. 165–190.
- [113] Granger C (1969) Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society* 37: 424–438.
- [114] Shannon CE (1948) A mathematical theory of communication. *The Bell System Technical Journal* 27: 379–423,623–656.
- [115] Shannon CE (1956) The bandwagon. *IRE Transactions on Information Theory* 2: 3.
- [116] Schreiber T (2000) Measuring information transfer. *Physical Review Letters* 85: 461–464.
- [117] Bak P, Tang C, Wiesenfeld K (1988) Self-organized criticality. *Physical Review A* 38: 364–374.
- [118] Zapperi S, Lauritsen K, Stanley H (1995) Self-organized branching processes: Mean-field theory for avalanches. *Physical Review Letters* 75: 4071–4074.
- [119] Eurich CW, Herrmann JM, Ernst UA (2002) Finite-size effects of avalanche dynamics. *Physical Review E* 66: 066137.

- [120] Levina A, Herrmann JM, Geisel T (2007) Dynamical synapses causing self-organized criticality in neural networks. *Nature Physics* 3: 857–860.
- [121] De Arcangelis L, Herrmann HJ (2010) Learning as a phenomenon occurring in a critical state. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 107: 3977–3981.
- [122] Lombardi F, Herrmann HJ, Perrone-Campano C, Plenz D, De Arcangelis L (2012) Balance between excitation and inhibition controls the temporal organization of neuronal avalanches. *Physical Review Letters* 108: 228703.
- [123] Beggs JM, Plenz D (2003) Neuronal avalanches in neocortical circuits. *Journal of Neuroscience* 23: 11167.
- [124] Hahn G, Petermann T, Havenith MN, Yu S, Singer W, et al. (2010) Neuronal avalanches in spontaneous activity in vivo. *Journal of Neurophysiology* 104: 3312–3322.
- [125] Athreya KB, Ney PE (2004) *Branching processes*. Mineola, New York: Dover Publications.
- [126] Plenz D (2010) A leak-proof model. *Nature Physics* 6: 717.
- [127] Kinouchi O, Copelli M (2006) Optimal dynamical range of excitable networks at criticality. *Nature Physics* 2: 348–351.
- [128] Tsodyks M, Markram H (1997) The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 94: 719–723.
- [129] Tsodyks M, Uziel A, Markram H (2000) Synchrony generation in recurrent networks with frequency-dependent synapses. *Journal of Neuroscience* 20: 1–5.
- [130] Stosiek C, Garaschuk O, Holthoff K, Konnerth A (2003) In vivo two-photon calcium imaging of neuronal networks. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 100: 7319–7324.
- [131] Grewe BF, Langer D, Kasper H, Kampa BM, Helmchen F (2010) High-speed in vivo calcium imaging reveals neuronal network activity with near-millisecond precision. *Nature Methods* 7: 399–405.
- [132] Fagiolo G (2007) Clustering in complex directed networks. *Physical Review E* 76: 026107.
- [133] Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402: 529–533.
- [134] Whittaker E, G R (1967) *The calculus of observations: An introduction to numerical analysis*. New York: Dover Publications, 4th edition.
- [135] Shante V, Kirkpatrick S (1971) An introduction to percolation theory. *Advances in Physics* 20: 325–357.
- [136] Grimmett G (1999) *Percolation*, volume 321 of *Grundlehren der mathematischen Wissenschaften*. New York: Springer, 2nd edition.
- [137] Kesten H (2006) What is percolation? *Notices of the American Mathematical Society* 53: 1–2.
- [138] Breskin I, Soriano J, Moses E, Tlusty T (2006) Percolation in living neural networks. *Physical Review Letters* 97: 188102.

- [139] Eckmann JP, Moses E, Stetter O, Tlusty T, Zbinden C (2010) Leaders of neuronal cultures in a quorum percolation model. *Frontiers in Computational Neuroscience* 4: 1–12.
- [140] Cohen O, Keselman A, Moses E, Martinez MR, Soriano J, et al. (2010) Quorum percolation in living neural networks. *Europhysics Letters* 89: 1–6.
- [141] Tlusty T, Eckmann JP (2009) Remarks on bootstrap percolation in metric networks. *Journal of Physics A: Mathematical and Theoretical* 42: 205004.
- [142] Abbott L (1999) Lapicque’s introduction of the integrate-and-fire model neuron (1907). *Brain Research Bulletin* 50: 303–304.
- [143] Gerstner W, Kistler W (2002) *Spiking neuron models. Single neurons, populations, plasticity.* Cambridge University Press.
- [144] Lapicque L (1907) Recherches quantitatives sur l’excitation électrique des nerfs traitée comme une polarisation. *Journal de physiologie et de pathologie générale* 9: 620–635.
- [145] Amit DJ, Tsodyks M (1991) Quantitative study of attractor neural network retrieving at low spike rates: I. Substrate. *Network: Computation in Neural Systems* 2: 259–273.
- [146] Amit DJ, Tsodyks M (1991) Quantitative study of attractor neural networks retrieving at low spike rates: II. Low-rate retrieval in symmetric networks. *Network: Computation in Neural Systems* 2: 275–294.
- [147] Brunel N, Hakim V (1999) Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Computation* 11: 1621–1671.
- [148] Markram H, Tsodyks M (1996) Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature* 382: 807–810.
- [149] Abbott L, Varela J, Sen K, Nelson S (1997) Synaptic depression and cortical gain control. *Science* 275: 220–224.
- [150] Loebel A, Tsodyks M (2002) Computation by ensemble synchronization in recurrent networks with synaptic depression. *Journal of Computational Neuroscience* 13: 111–124.
- [151] Gewaltig MO, Diesmann M (2007). NEST (NEural Simulation Tool). URL [http://www.scholarpedia.org/article/NEST_\(NEural_Simulation_Tool\)](http://www.scholarpedia.org/article/NEST_(NEural_Simulation_Tool)).
- [152] Eppler JM, Helias M, Muller E, Diesmann M, Gewaltig MO (2009) PyNEST: A convenient interface to the NEST simulator. *Frontiers in Neuroinformatics* 2: 4–12.
- [153] Gewaltig MO, Morrison A, Plesser HE (2012) NEST by example: An introduction to the neural simulation tool NEST. In: Le Novère N, editor, *Computational Systems Neurobiology*, Dordrecht: Springer Netherlands. pp. 533–558.
- [154] Rotter S, Diesmann M (1999) Exact digital simulation of time-invariant linear systems with applications to neuronal modeling. *Biological Cybernetics* 81: 381–402.
- [155] Plesser HE, Gerstner W (2000) Noise in integrate-and-fire neurons: From stochastic input to escape rates. *Neural Computation* 12: 367–384.
- [156] Amemori K, Ishii S (2001) Gaussian process approach to spiking neurons for inhomogeneous Poisson inputs. *Neural Computation* 13: 2763–2797.

- [157] Lindner B, García-Ojalvo J, Neiman A, Schimansky-Geier L (2004) Effects of noise in excitable systems. *Physics Reports* 392: 321–424.
- [158] Aertsen A, Gerstein G, Habib MK, Palm G (1989) Dynamics of neuronal firing correlation: Modulation of effective connectivity. *Journal of Neuroscience* 61: 900–917.
- [159] Sporns O, Tononi G, Edelman GM (2000) Theoretical neuroanatomy: Relating anatomical and functional connectivity in graphs and cortical connection matrices. *Cerebral Cortex* 10: 127–141.
- [160] Kerr JN, Greenberg DS, Helmchen F (2005) Imaging input and output of neocortical networks in vivo. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 102: 14063–14068.
- [161] Feinerman O, Moses E (2006) Transport of information along inidimensional layered networks of dissociated hippocampal neurons and implications for rate coding. *Journal of Neuroscience* 26: 4526–4534.
- [162] Sasaki T, Matsuki N, Ikegaya Y (2007) Metastability of active CA3 networks. *Journal of Neuroscience* 27: 517–528.
- [163] Kreitzler AC, Gee KR, Archer EA, Regehr WG (2000) Monitoring presynaptic calcium dynamics in projection fibers by in vivo loading of a novel calcium indicator. *Neuron* 27: 25–32.
- [164] Yasuda R, Nimchinsky E, Scheuss V, Pologruto T, Oertner T, et al. (2004) Imaging calcium concentration dynamics in small neuronal compartments. *Science Signaling* 219: p15.
- [165] Zahradníková AJ, Poláková E, Zahradník I, Zahradníková A (2007) Kinetics of calcium spikes in rat cardiac myocytes. *Journal of Physiology* 578: 677–691.
- [166] Bruton JD, Cheng AJ, Westerblad H (2012) Methods to detect Ca²⁺ in living cells. In: Islam MS, editor, *Calcium signaling*, Dordrecht: Springer Netherlands. pp. 27–43.
- [167] Gee KR, Brown KA, Chen WNU, Bishop-Stewart J, Gray D, et al. (2000) Chemical and physiological characterization of fluo-4 Ca²⁺-indicator dyes. *Cell Calcium* 27: 97–106.
- [168] Opitz T, de Lima AD, Voigt T (2002) Spontaneous development of synchronous oscillatory activity during maturation of cortical networks in vitro. *Journal of Neurophysiology* 88: 2196–2206.
- [169] Stetter O, Battaglia D, Soriano J, Geisel T (2012) Model-free reconstruction of excitatory neuronal connectivity from calcium imaging signals. *PLoS Computational Biology* 8: e1002653.
- [170] Kalisman N, Silberberg G, Markram H (2005) The neocortical microcircuit as a tabula rasa. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 102: 880–885.
- [171] Vogelstein JT, Watson BO, Packer AM, Yuste R, Jedynak B, et al. (2009) Spike inference from calcium imaging using sequential Monte Carlo methods. *Biophysical Journal* 97: 636–655.
- [172] Ishimaru A (1999) *Wave propagation and scattering in random media*. Wiley-IEEE Press.
- [173] Pyhtila JW, Ma H, Simnick AJ, Chilkoti A, Wax A (2006) Analysis of long range correlations due to coherent light scattering from in-vitro cell arrays using angle-resolved low coherence interferometry. *Journal of Biomedical Optics* 11: 034022.

- [174] Bonifazi P, Goldin M, Picardo MA, Jorquera I, Cattani A, et al. (2009) GABAergic hub neurons orchestrate synchrony in developing hippocampal networks. *Science* 326: 1419–1424.
- [175] Cadotte AJ, DeMarse TB, He P, Ding M (2008) Causal measures of structure and plasticity in simulated and living neural networks. *PLoS ONE* 3: e3355.
- [176] Ladroue C, Guo S, Kendrick K, Feng J (2009) Beyond element-wise interactions: Identifying complex interactions in biological processes. *PLoS ONE* 4: e6899.
- [177] Kispersky T, Gutierrez GJ, Marder E (2011) Functional connectivity in a rhythmic inhibitory circuit using Granger causality. *Neural Systems & Circuits* 1: 1–15.
- [178] MacKay DJC (2003) *Information theory, inference, and learning algorithms*. Cambridge University Press.
- [179] Quinn CJ, Coleman TP, Kiyavash N, Hatsopoulos NG (2010) Estimating the directed information to infer causal relationships in ensemble neural spike train recordings. *Journal of Computational Neuroscience* 30: 17–44.
- [180] Singh A, Lesica NA (2010) Incremental mutual information: A new method for characterizing the strength and dynamics of connections in Neuronal circuits. *PLoS Computational Biology* 6: e1001035.
- [181] Margolin AA, Nemenman I, Basso K, Wiggins C, Stolovitzky G, et al. (2006) ARACNE: An algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context. *BMC Bioinformatics* 7: S7.
- [182] Kaiser A, Schreiber T (2002) Information transfer in continuous processes. *Physica D: Nonlinear Phenomena* 166: 43–62.
- [183] Tung TQ, Ryu T, Lee KH, Lee D (2007) Inferring gene regulatory networks from microarray time series data using transfer entropy. *Twentieth IEEE International Symposium on Computer-Based Medical Systems, CBMS '07* 22: 383–388.
- [184] Gourevitch B, Eggermont J (2007) Evaluating information transfer between auditory cortical neurons. *Journal of Neurophysiology* 97: 2533–2543.
- [185] Garofalo M, Nieuwenhuis T, Massobrio P, Martinoia S (2009) Evaluation of the performance of information theory-based methods and cross-correlation to estimate the functional connectivity in cortical networks. *PLoS ONE* 4: e6482.
- [186] Ito S, Hansen ME, Heiland R, Lumsdaine A, Litke AM, et al. (2011) Extending transfer entropy improves identification of effective connectivity in a spiking cortical network model. *PLoS ONE* 6: e27431.
- [187] Besserve M, Schölkopf B, Logothetis NK, Panzeri S (2010) Causal relationships between frequency bands of extracellular signals in visual cortex revealed by an information theoretic analysis. *Journal of Computational Neuroscience* 29: 547–566.
- [188] Wibral M, Rahm B, Rieder M, Lindner M, Vicente R, et al. (2011) Transfer entropy in magnetoencephalographic data: Quantifying information flow in cortical and cerebellar networks. *Progress in Biophysics and Molecular Biology* 105: 80–97.
- [189] Vicente R, Wibral M, Lindner M, Pipa G (2011) Transfer entropy: A model-free measure of effective connectivity for the neurosciences. *Journal of Computational Neuroscience* 30: 45–67.

- [190] Staniek M, Lehnertz K (2008) Symbolic transfer entropy. *Physical Review Letters* 100: 158101.
- [191] Barnett L (2009) Granger causality and transfer entropy are equivalent for Gaussian variables. *Physical Review Letters* 103: 238701.
- [192] Egan JP (1975) *Signal detection theory and ROC analysis*. New York: Academic Press, Inc.
- [193] Fawcett T (2006) An introduction to ROC analysis. *Pattern recognition letters* 27: 861–874.
- [194] Helmchen F, Denk W (2005) Deep tissue two-photon microscopy. *Nature Methods* 2: 932–940.
- [195] Holekamp T, Turaga D, Holy T (2008) Fast three-dimensional fluorescence imaging of activity in neural populations by objective-coupled planar illumination microscopy. *Neuron* 57: 661–672.
- [196] Scanziani M, Häusser M (2009) Electrophysiology in the age of light. *Nature* 461: 930–939.
- [197] Feinerman O, Segal M, Moses E (2005) Signal propagation along unidimensional neuronal networks. *Journal of Neurophysiology* 94: 3406–3416.
- [198] Zucker RS (1996) Exocytosis: A molecular review and physiological perspective. *Neuron* 17: 1049–1055.
- [199] Kuhn A, Aertsen A, Rotter S (2004) Neuronal integration of synaptic input in the fluctuation-driven regime. *Journal of Neuroscience* 24: 2345–2356.
- [200] Bear MF, Connors BW, Paradiso MA (2007) *Neuroscience. Exploring the brain*. Baltimore: Lippincott Williams & Wilkins, 3rd edition.
- [201] Sullivan MR, Nimmerjahn A, Sarkisov DV, Helmchen F, Wang SSH (2005) In vivo calcium imaging of circuit activity in cerebellar cortex. *Journal of Neurophysiology* 94: 1636–1644.
- [202] Zhu P, Fajardo O, Shum J, Zhang Schärer YP, Friedrich RW (2012) High-resolution optical control of spatiotemporal neuronal activity patterns in zebrafish using a digital micromirror device. *Nature Protocols* 7: 1410–1425.
- [203] Barabási AL, Albert R (1999) Emergence of scaling in random networks. *Science* 286: 509–512.
- [204] Pajevic S, Plenz D (2009) Efficient network reconstruction from dynamical cascades identifies small-world topology of neuronal avalanches. *PLoS Computational Biology* 5: e1000271.
- [205] Strata P, Harvey R (1999) Dale’s principle. *Brain Research Bulletin* 50: 349–350.
- [206] Luczak A, MacLean JN (2012) Default activity patterns at the neocortical microcircuit level. *Frontiers in Integrative Neuroscience* 6: 1–6.
- [207] Laughlin SB (1989) The role of sensory adaptation in the retina. *Journal of Experimental Biology* 146: 39–62.
- [208] Burrone J, Murthy VN (2003) Synaptic gain control and homeostasis. *Current Opinion in Neurobiology* 13: 560–567.
- [209] Remme MWH, Wadman WJ (2012) Homeostatic scaling of excitability in recurrent neural networks. *PLoS Computational Biology* 8: e1002494.
- [210] Turrigiano GG, Leslie KR, Desai NS, Rutherford LC, Nelson SB (1998) Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391: 892–896.

- [211] Cocco S, Leibler S, Monasson R (2009) Neuronal couplings between retinal ganglion cells inferred by efficient inverse statistical physics method. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 106: 14058–14062.
- [212] Greenberg DS, Houweling AR, Kerr JND (2008) Population imaging of ongoing neuronal activity in the visual cortex of awake rats. *Nature Neuroscience* 11: 749–751.
- [213] Chemla S, Chavane F (2010) Voltage-sensitive dye imaging: Technique review and models. *Journal of Physiology Paris* 104: 40–50.
- [214] Bernstein JG, Boyden ES (2011) Optogenetic tools for analyzing the neural circuits of behavior. *Trends in Cognitive Sciences* 15: 592–600.
- [215] Roudi Y, Nirenberg S, Latham PE (2009) Pairwise maximum entropy models for studying large biological systems: When they can work and when they can't. *PLoS Computational Biology* 5: e1000380.
- [216] Stallcup WB, Beasley L (1985) Involvement of the nerve growth factor-inducible large external glycoprotein (NILE) in neurite fasciculation in primary cultures of rat brain. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 82: 1276–1280.
- [217] Tessier-Lavigne M, Goodman C (1996) The molecular biology of axon guidance. *Science* 274: 1123–1133.
- [218] Petrinovic MM, Duncan CS, Bourikas D, Weinman O, Montani L, et al. (2010) Neuronal Nogo-A regulates neurite fasciculation, branching and extension in the developing nervous system. *Science Signaling* 137: 2539–2550.
- [219] Kriegstein AR, Dichter MA (1983) Morphological classification of rat cortical neurons in cell culture. *Journal of Neuroscience* 3: 1634–1647.
- [220] Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. *Nature* 393: 440–442.
- [221] Downes JH, Hammond MW, Xydas D, Spencer MC, Baccera VM, et al. (2012) Emergence of a small-world functional network in cultured neurons. *PLoS Computational Biology* 8: e1002522.
- [222] Bartos M, Vida I, Jonas P (2007) Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nature Reviews Neuroscience* 8: 45–56.
- [223] Shahaf G, Eytan D, Gal A, Kermany E, Lyakhov V, et al. (2008) Order-based representation in random networks of cortical neurons. *PLoS Computational Biology* 4: e1000228.
- [224] Yizhar O, Fenno LE, Davidson TJ, Mogri M, Deisseroth K (2011) Optogenetics in neural systems. *Neuron* 71: 9–34.
- [225] Tiesinga PHE (2012) Motifs in health and disease: The promise of circuit interrogation by optogenetics. *European Journal of Neuroscience* 36: 2260–2272.
- [226] Kohsaka H, Okusawa S, Itakura Y, Fushiki A, Nose A (2012) Development of larval motor circuits in *Drosophila*. *Development, Growth & Differentiation* 54: 408–419.
- [227] Zhao T, Petralia RS, Yu Y, Peng H, Myers E, et al. (2011) mGRASP enables mapping mammalian synaptic connectivity with light microscopy. *Nature Methods* 9: 96–U139.

- [228] Mishchenko Y, Paninski L (2012) A Bayesian compressed-sensing approach for reconstructing neural connectivity from subsampled anatomical data. *Journal of Computational Neuroscience* : In press.
- [229] Strutz A, Völler T, Riemensperger T, Fiala A, Sachse S (2012) Calcium imaging of neural activity in the olfactory system of *Drosophila*. In: Martin JR, editor, *Genetically encoded functional indicators*, Totowa, NJ: Humana Press. pp. 43–70.
- [230] Ahrens MB, Li JM, Orger MB, Robson DN, Schier AF, et al. (2012) Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature* 485: 471–477.
- [231] Stoelzle S, Obergrussberger A, Brüggemann A, Haarmann C, George M, et al. (2011) State-of-the-art automated patch clamp devices: Heat activation, action potentials, and high throughput in ion channel screening. *Frontiers in Pharmacology* 2: 76.
- [232] Neul JL (2011) Unfolding neurodevelopmental disorders: The mystery of developing connections. *Nature Medicine* 17: 1353–1355.
- [233] Morgan R, Soltesz I (2008) Nonrandom connectivity of the epileptic dentate gyrus predicts a major role for neuronal hubs in seizures. *Proceedings Of The National Academy Of Sciences Of The United States Of America* 105: 6179–6184.
- [234] Song S, Miller KD, Abbott L (2000) Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience* 3: 919–926.
- [235] Rubinov M, Sporns O, van Leeuwen C, Breakspear M (2009) Symbiotic relationship between brain structure and dynamics. *BMC Neuroscience* 10: 1–18.
- [236] Thivierge JP (2009) How does non-random spontaneous activity contribute to brain development? *Neural Networks* 22: 901–912.
- [237] Kempter R, Gerstner W, van Hemmen JL (1999) Hebbian learning and spiking neurons. *Physical Review E* 59: 4498–4514.
- [238] Caroni P, Donato F, Muller D (2012) Structural plasticity upon learning: Regulation and functions. *Nature Reviews Neuroscience* 13: 478–490.
- [239] Ganguly K, Schinder AF, Wong ST, Poo Mm (2001) GABA itself promotes the developmental switch of neuronal GABAergic responses from excitation to inhibition. *Cell* 105: 521–532.
- [240] Paninski L (2003) Estimation of entropy and mutual information. *Neural Computation* 15: 1191–1253.
- [241] de Matos Simoes R, Emmert-Streib F (2011) Influence of statistical estimators of mutual information and data heterogeneity on the inference of gene regulatory networks. *PLoS ONE* 6: e29279.
- [242] Lindner M, Vicente R, Priesemann V (2011) TRENTOOL: A Matlab open source toolbox to analyse information flow in time series data with transfer entropy. *BMC Neuroscience* 12: 1–22.
- [243] Strong S, Koberle R, de Ruyter van Steveninck RR, Bialek W (1998) Entropy and information in neural spike trains. *Physical Review Letters* 80: 197–200.
- [244] Nemenman I, Bialek W, de Ruyter van Steveninck RR (2004) Entropy and information in neural spike trains: Progress on the sampling problem. *Physical Review E* 69: 056111.

- [245] Panzeri S, Senatore R, Montemurro MA, Petersen RS (2007) Correcting for the sampling bias problem in spike train information measures. *Journal of Neurophysiology* 98: 1064–1072.
- [246] Treves A, Panzeri S (1995) The upward bias in measures of information derived from limited data samples. *Neural Computation* 7: 399–407.
- [247] Muja M, Lowe DG (2009) Fast approximate nearest neighbors with automatic algorithm configuration. *International Conference on Computer Vision Theory and Applications VISSAPP'09* : 331–340.
- [248] Kozachenko L, Leonenko N (1987) Sample estimate of the entropy of a random vector. *Problems of Information Transmission* 23: 95–101.
- [249] Victor JD (2002) Binless strategies for estimation of information from neural data. *Physical Review E* 66: 051903.
- [250] Kraskov A, Stögbauer H, Grassberger P (2004) Estimating mutual information. *Physical Review E* 69: 066138.
- [251] Frenzel S, Pompe B (2007) Partial mutual information for coupling analysis of multivariate time series. *Physical Review Letters* 99: 204101.
- [252] Schreiber T, Schmitz A (1996) Improved surrogate data for nonlinearity tests. *Physical Review Letters* 77: 635–638.
- [253] Schreiber T, Schmitz A (2000) Surrogate time series. *Physica D: Nonlinear Phenomena* 142: 346–382.

Acknowledgments

I would like to thank many people, especially Theo Geisel for being my advisor and bringing together a wonderful group of people in his department, stimulating scientific exchange and open collaboration. I would like to thank Demian Battaglia for teaching me a great many things about how to work scientifically, and Anna Levina for advising me on the initial part of my thesis. Making these years especially interesting and enjoyable, in science as well as in real life, and last but not least to some of them for reading and correcting this document, I would like to thank my colleagues, especially Wolfgang Keil, Jakob Metzger, Christoph Kirst, Andreas Sorge, Max Puelma Touzel and everyone else.

I very much thank Elisha Moses and Elad Schneidman at the Weizmann Institute of Science and Christian Leibold of the Bernstein Center for Computational Neuroscience Munich for having me as a guest and giving me the opportunity to take part in the scientific work of their respective groups. I would like to thank Jordi Soriano and Javier Orlandi for the stimulating collaboration on neuronal networks in vitro, and Jean-Pierre Eckmann and Cyrille Zbinden for valuable and memorable discussions on the same subject.

I would like to thank Fred Wolf for being part of my thesis committee and his advice throughout the years of my thesis, the GGNB team for organizing our graduate school and the Minerva Foundation for generously supporting my stay in Israel.

Finally, and most importantly, I would like to thank my parents for their support and Nadine for everything.

I hereby declare that this thesis has been written independently and with no other sources and aids than quoted.

Frank Olav Stetter
August 29, 2012