MASTER THESIS

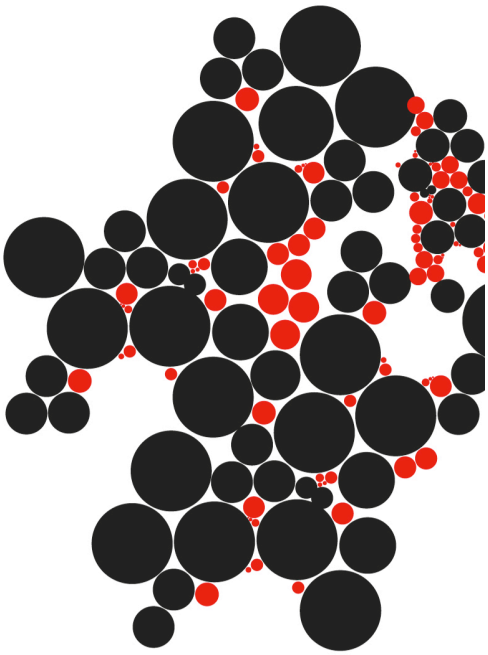

ADAPTATION TO LOW-FREQUENT STIMULI IN COMPUTATIONAL AND CULTURED NEURONAL NETWORKS

Jelle Dijkstra

FACULTY OF ELECTRICAL ENGINEERING, MATHEMATICS AND COMPUTER SCIENCE
BIOMEDICAL SIGNALS AND SYSTEMS

EXAMINATION COMMITTEE
Prof.dr. Richard J.A. van Wezel
Dr.ir. Joost le Feber
Dr.ir. T. Heida
Dr. A. Zagaris

DOCUMENT NUMBER
BSS – 14-05

UNIVERSITY OF TWENTE.

28-03-2014

Samenvatting

Men schat het totale aantal neuronen in het menselijke brein op ongeveer 10^{11}. Neuronen zijn gespecialiseerde cellen die elektrisch en chemisch met elkaar communiceren via verbindingen genaamd synapsen. Deze verbindingen veranderen een verzameling cellen in een neuraal netwerk dat complex genoeg is om over zijn eigen complexiteit na te denken. Het begrijpen van zulke netwerken is een veelbelovend doel in neurowetenschappen. Enerzijds kunnen hierdoor mogelijkerwijs neurale aandoeningen zoals Parkinson, epilepsie en dementie beter doorgrond en daarmee behandeld worden. Andererzijds kan de natuur als bron van inspiratie dienen voor uitdagingen op het gebied van kunstmatige intelligentie. De immense complexiteit van biologische neurale netwerken maakt het echter erg lastig om ze volledig te doorgronden. In dit onderzoek beschouwen we daarom slechts één van hun vele eigenschappen, namelijk het aanpassen aan stimulatie, een specifieke vorm van geheugen. Het is algemeen aanvaard dat geheugen in neurale netwerken gegeven wordt door synaptische plasticiteit, d.w.z. dat de synapsen kunnen veranderen in sterkte. Plasticiteit kan optreden bij intrinsieke activiteit of externe input op het netwerk.

Er zijn verschillende manieren om biologische netwerken toegankelijker te maken voor onderzoek. Er kunnen computermodellen van netwerken worden gemaakt zodat hun gedrag gesimuleerd kan worden. Een dergelijk computermodel kan perfect geobserveerd en gecontroleerd worden maar is altijd een sterke abstractie van de werkelijkheid. Een andere manier is een platform, genaamd “multi-electrode array” (MEA). Dit is een raster van zeer kleine elektrodes waarop geïsoleerde, levende neurale cellen worden geplateerd. Deze cellen vormen opnieuw verbindingen en tonen spontane activiteit, genaamd “network bursts”. De elektrodes kunnen deze activiteit meten en daarnaast het netwerk ook elektrisch stimuleren. Neuronen in de buurt van de elektrodes kunnen geactiveerd worden wat tot netwerk-wijde activiteit oftewel “stimulus responses” kan leiden.

We vermoeden dat de spontane activiteit en synaptische sterktes in een gekweekt netwerk een balans vormen. Deze balans kan mogelijkerwijs worden verstoord door een netwerk herhaaldelijk te stimuleren, waardoor “stimulus responses” opgewekt worden die niet passen bij de heersende spontane activiteit in het netwerk. Na het stimuleren zouden dan sporen van de stimulus responsie zichtbaar moeten zijn in de spontane activiteit van het netwerk. Deze hypothese werd getest door middel van een maat voor de gelijkenis van “network bursts” met “stimulus responses”. Aan het begin van het onderzoek stond een kleinschalig computermodel, gemaakt uit bestaande neuronen-, synaps- en plasticiteitsmodellen. Dit model ontwikkelde een stabiel evenwicht tussen intrinsieke activiteit en synaptische sterktes. Na het bereiken van dit evenwicht hebben we de gesimuleerde netwerken 900 keer gestimuleerd. De gemiddelde gelijkenis van de responsies op deze stimuli met de spontane activiteit ná het stimuleren was significant hoger dan de gemiddelde gelijkenis tussen de stimulus responsies en de spontane activiteit vóór het stimuleren. Een soortgelijk protocol hebben we toen toegepast in negen verschillende experimenten, waarvoor vier individuele, gekweekte MEA netwerken werden gebruikt. Aansluitend werd dezelfde analyse toegepast als op de gesimuleerde data. De experimentele data volgde dezelfde tendens als de door simulaties verkregen data maar was te inconsistent om statistisch significant te zijn.

Summary

It is estimated that the human brain consists of approximately 10^{11} individual neurons, cells with remarkable features that form electrical and chemical connections among each other known as synapses. These connections turn a cluster of nerve cells into a neuronal network that is complex enough to think about its own complexity. Comprehending the mechanisms of the neuronal structures in living beings is a rewarding task in neuroscience. It may aid in understanding and therefore treating neurological disorders, such as Parkinson's disease, epilepsy and dementia. In addition, artificial neuronal networks may be fruitful for problems such as image processing, pattern recognition and artificial intelligence in general. The great complexity of biological neuronal networks however makes understanding them very challenging. In this research, we therefore focus on one of the many aspects of neuronal networks, namely their adaptation to stimulation, which is a special form of memory. It is generally accepted that memory in neuronal networks is enabled by the ability of the synapses to change in strength, called synaptic plasticity. Plasticity can be induced by intrinsic activity or external input to the network.

There are several ways to make biological neuronal networks more accessible for research. One of them are computational models that simulate their behavior. Computer models are perfectly observable and controllable but they are always a strong abstraction of reality, since even if we could make realistic models of all the processes in a living network, the model would become too complex to be useful in this study. Another option closer to biological reality is a platform called multi-electrode array (MEA). This is basically a grid of electrodes on which isolated, living nerve cells are plated. These cells start to reconnect, forming a 2D network and show spontaneous synchronized activity called network bursts. The electrodes can measure this spontaneous activity and can also be used for electrical stimulation. Neurons near the electrodes can be activated and cause network wide responses known as stimulus responses.

We hypothesized, that the spontaneous bursts and synaptic connections in a cultured network form a balance. Repeatedly stimulating a network and thereby inducing stimulus responses that do not agree with the prevailing spontaneous activity may disturb this balance. After stimulation, traces of the stimulus response should then be visible in the spontaneous activity of the network after stimulation. This hypothesis was examined by utilizing a measure of similarity between network bursts and stimulus responses. We started our investigations with a small-scale computational network made of existing neuron, synapse and long-term plasticity models. This model developed a stable equilibrium between intrinsic activity and synaptic strengths. We then applied one phase of low-frequent stimulation with 900 stimuli to the model. The average similarity of the stimulus responses with network bursts after stimulation was significantly higher than the average similarity of the stimulus responses with the network bursts before stimulation. A similar protocol was then applied to four distinct MEA cultures in a total of nine experiments and the same analysis was performed on the measurement data. This data followed the same trend as the data gained by the computational model, but was too inconsistent to be statistically significant.

Table of Contents

1. Introduction

Single neurons have been studied thoroughly in the past decades and their functionality is quite well understood, both phenomenologically as well as mechanistically. Any kind of behavior seen in living organisms however requires a network of many mutually connected neurons that operate together. Such neuronal networks show properties that far surpass the capabilities of single neurons and are of great interest in neuroscience. Understanding the dynamics and mechanics of neuronal networks is relevant in two ways. First, knowing how a system should work helps in identifying the problem in cases it does not function properly. Second, neuronal networks show information processing abilities that exceed those of technical systems in specialized tasks. For these engineering applications, nature may be a source of inspiration.

Despite extensive research, the level of understanding of neuronal networks found in living organisms is still basic. Possible explanations for this are their very small scale, which cause difficulties in interfacing, as well as their shear complexity. The complexity can be reduced by focusing on single features observed in neuronal networks. One of the many topics of interest is memory function. On a functional level, memory can be described as the concept of storing and recalling information. In the broadest sense however, memory means that the current behavior of a network depends on prior input and therefore requires (semi-)permanent changes in the network. This is essential for many higher-order features, e.g. any kind of adaptive system.

1.1. Synaptic plasticity

In order to show memory, external input to a network must lead to long-lasting changes. Input could cause structural changes in a network, e.g. new axons and synapses could grow or synapses could disappear. Another possibility is a change in the synaptic strengths of existing connections, also known as synaptic plasticity. The fact that the timescale of learning identified in biological networks is much smaller than the timescale of possible structural changes, implies that a mere change in synaptic strengths within an otherwise static network is sufficient for memory function. This is in agreement with a wide variety of theoretical work, most famously the work of D. Hebb, who came to the conclusion that the connections between neurons, which are simultaneously active, increase in strength [25]. His learning rule, known as Hebb's rule, has been adapted to artificial neural networks such as Hopfield networks and is capable of achieving desired input-output relations by locally adjusting the synaptic strengths according to the desired relation.

A mechanism that causes long-lasting strength changes in biological synapses is called Spike-timing-dependent plasticity (STDP). It states that the strength of synapses change according to the precise timing between the arrival of presynaptic action potentials and the action potentials fired by the postsynaptic neuron. If a neuron B spikes very shortly after receiving synaptic input from neuron A, the connection from A to B strengthens. If the timing is reversed, the connection weakens. In the past decades, STDP has been researched extensively on a cellular level [2, 30] and the gained results are valid in small biological networks as well [3]. Spike-timing-dependent plasticity is intuitive in that synapses at

which the postsynaptic neuron fires before the arrival of a presynaptic spike seem to be obsolete while synapses at which the presynaptic spike often induces a postsynaptic spike seem to be vital. Still, the challenge remains to show whether and, if so, how this low-level mechanism can explain the memory effects seen in neuronal networks.

1.2. Cultured neuronal networks

A substantial difficulty when studying naturally grown neuronal networks is the accessibility of individual neurons. The neuronal density in the cortex of primates can exceed 100 million neurons/g [6] which makes it difficult to simultaneously access a set of single neurons *in vivo*. Visual and electrical accessibility is improved by studying natural neuronal networks *in vitro* using two dimensional cultured networks grown on multi-electrode arrays (MEAs) [24]. Cultured networks consist of hundreds of thousands of isolated neurons, usually taken from the cortex of newborn rats. After the cells are plated on the MEA, they start to grow new axons and synapses, thereby forming new connections with each other. Every neuron has axons that, on average, connect to the dendrites of approximately one thousand other neurons [13]. It is relatively easy to culture, maintain and interface neuronal networks on MEAs, which have therefore become a major platform for studies on this topic.

Beneath the neurons is a grid of electrodes. The diameter of one electrode lies in the range 10–50 µm, which is approximately the size of neuronal cell bodies in mammalian brains and have a spacing of 100–500 µm. Regular MEAs are about 8×8 electrodes in size while high-density MEAs can have grids of up to 64×64 electrodes. Ideally, every electrode is covered with exactly one neuron. The electrodes measure changes in the neuron's extracellular potential caused by ions flowing through the membrane and can thereby detect action potentials fired by the according neuron. In addition, the electrodes can be used for local stimulation of a network by inducing potential gradients in the medium surrounding the cells. Several experiments aimed at showing network plasticity have been conducted on cultured networks. The authors of [31,34,42] showed that applying external input to cultured neuronal networks can cause network-wide changes and thereby confirm that MEAs are suitable platforms for studying memory.

Bursting

A significant characteristic of cultured networks is the spontaneous bursting behavior. During a burst, a large part of the cultured network synchronizes and becomes active simultaneously, as exemplary shown in figure 1.1a. The frequency and shape of network bursts show a substantial development over time and appear to stabilize after three to four weeks in vitro [19,20]. The exact origin of the bursts is unknown but research implies that the bursting starts locally and then spreads out to the rest of the network [8,22]. The amount of activity within a burst therefore depends on the connectivity of the network. When there are only few, weak connections between the neurons, the activity will die out quickly whereas many, strong connections can cause the whole network to become active. At the same time, the spontaneous activity alters this connectivity since neurons spiking in a small temporal window can induce long-lasting changes in the synapses. A theory to explain memory effects within a cultured network states that the activity and connectivity in a network form a balance or equilibrium, such that the occurring activity patterns maintain the connectivity that supports this specific activity [21,42]. The authors suggest that the current equilibrium may be altered by stimulation, if the stimulation

induces stimulus responses that do not agree with the spontaneously occurring patterns. As discussed in the articles, stimulation may drive the network to a different equilibrium in case a network incorporates multiple equilibria. But also in case of a single equilibrium, stimulation may temporary disturb the equilibrium. If the time-scale of the network dynamics is relatively large compared to observation period, this disturbance may become visible after stimulation.

In one of the experiments described in [42], three phases of stimulation with four stimulation periods each were applied to cultured neuronal networks. The same stimulation electrode was used for the first and third stimulation phase, and a different electrode was used for the second stimulation phase. The change in connectivity due to the stimulation by means of conditional firing probabilities (CFP) [20] is shown in figure 1.2. The first and second stimulation phases have a strong effect on the connectivity in the network. Stimulating the first electrode again however in the third phase has no significant effect on the connectivity anymore. The authors note that the change in connectivity at the end of the first stimulation phase is less than the change in connectivity induced by the first two stimulation periods in this phase and interpret this as follows: "These [the first

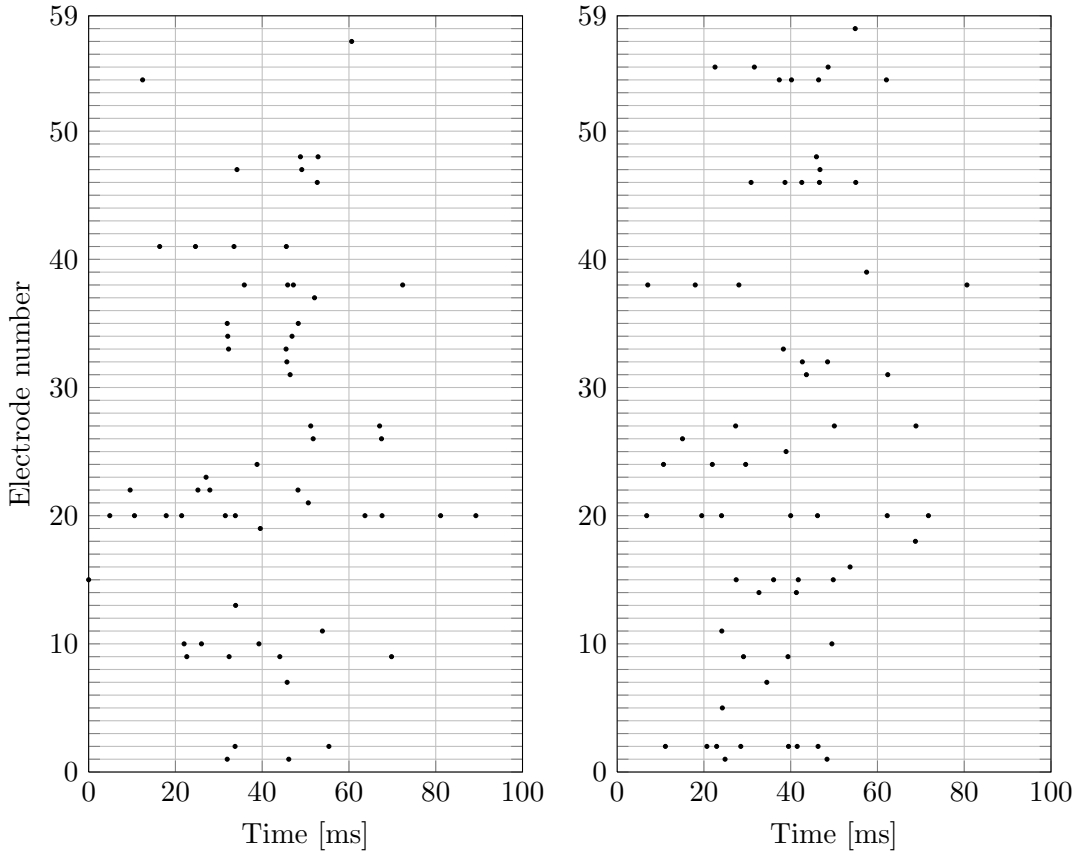

(a) Network burst (b) Stimulus response

Figure 1.1.: (a) shows a network burst that occurred during spontaneous activity of a cultured neuronal network ($t = 0$ is the time instance of the first spike in the burst), (b) shows a stimulus response in the same cultured network ($t = 0$ is the moment of stimulation). Every dot depicts an action potential detected at one of the 60 electrodes of the underlying multi-electrode array at a given time.

two] stimulation periods pushed the network towards a new balance, which supported the activity patterns that are evoked during stimulation. Because this new balance supported the induced activity patterns, the subsequent stimulation did not cause new connectivity changes. In other words it memorized the applied input." [42, p. 56] On the following page they say about the third stimulation phase: "When we reapplied stimulation on stimulation electrode A we did not observe this potentiation [the strengthening of the connections seen in the second period of the first two stimulation phases]. [...] This could indicate that the memory developed for electrode A is not forgotten and maintained over a period of 5 hours."

Practical limitations

Working with MEAs yields some practical issues for plasticity studies. Despite the large number and small size of the electrodes, the mapping between electrodes and neurons is far from one-to-one. Just a very small fraction of the action potentials in the network is recorded and a single electrode may occasionally sense the action potentials of multiple neurons. When stimulating the network through the electrodes, several neurons and axons are stimulated simultaneously. In addition, networks are always changing and it is impossible to freeze the current state or restore a previous state. This makes it impossible

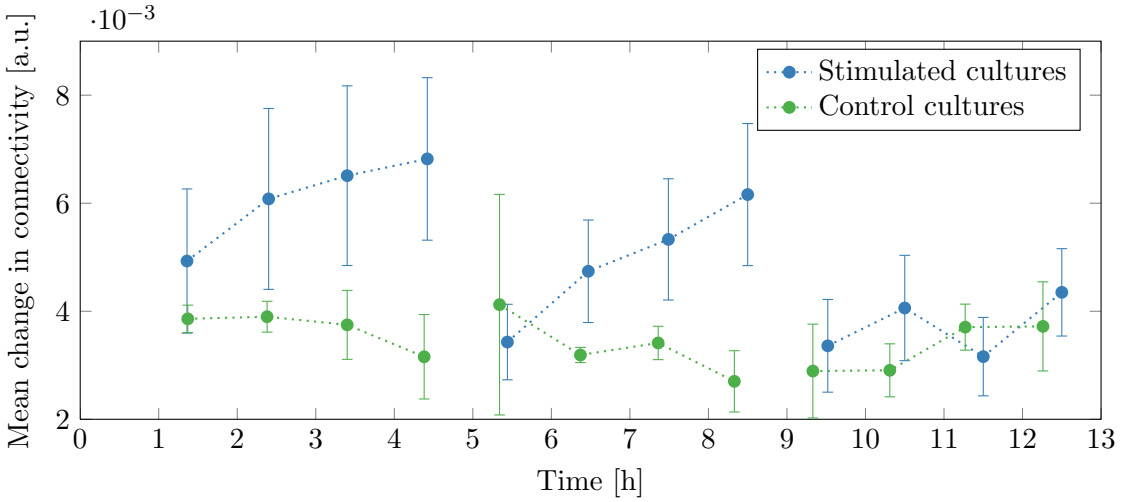

Figure 1.2.: Mean changes in connectivity in cultured neuronal networks due to low-frequent stimulation at two individual electrodes ($n = 7$), compared to the mean changes in connectivity in unstimulated control cultures ($n = 3$). Three phases of stimulation were applied, indicated by the three groups visible in the plot. Every phase consisted of four stimulation periods. The same stimulation electrode was used for all the stimulation periods in the first and third stimulation phase, a different electrode in the second stimulation phase. The connectivities in the cultures were quantified using the method of conditional firing probabilities (CFP), leading to weight matrices [20]. The change in connectivity during a stimulation/control phase is plotted as the Euclidean distance between the weight matrix after a stimulation period and the weight matrix at the beginning of the according stimulation phase. The error bars show the standard error of the mean. Recreated from the experimental data of [42].

to apply different stimuli to the very same network. The biggest problem related to memory however is that there is currently no feasible method to simultaneously measure the strength of several synaptic weights in a MEA culture. Only the functional connections between neurons recorded by the electrodes can be described in statistical terms [20]. A computer simulation overcomes these issues but requires a high degree of abstraction and simplification.

1.3. Computational model

The modeling study [39] showed that the properties of cultured neuronal networks important for this research, i.e. bursting and adaption to stimuli, can be reproduced by a small-scale computational model. A simulation yields a number of advantages compared to an in vitro experiment. First of all a simulation can be faster. The exact speed of a simulation strongly depends on the level of realism, the size of the network and the machine the simulation runs on. A simulation of the full brain for instance is still far from real-time, even on a Beowulf cluster with 60 computers [17]. The simulation developed by [39] can run approximately real-time on a mediocre desktop machine. Yet, the effective speed compared with real experiments is much higher, since there is no preparation time, multiple machines can be used simultaneously, etc. Second there is controllability. In a real network, every experiment is different. The network changes over time in an uncontrolled and unknown manner and it is impossible to store and recall a previous state. In a simulation, we can rerun experiments with just very subtle changes and thereby observe their specific effects. With respect to stimulation, only areas around the electrodes can be activated in a MEA culture. The effective stimulation range is determined by the gradient of the electric field that is applied. Not only all the neurons but also the axons within that range are stimulated. In a simulation on the other hand, neurons can be observed and stimulated individually. Last, there is the issue of accessibility. The focus of this thesis is on memory and therefore the development of synaptic weights. Currently, there is no method to simultaneously measure the synaptic weights in vitro, only statistical approaches that describe the cross-correlation between activity measured at the electrodes are available. In a simulation we can not only distinctly record the synaptic weights, we also know the axonal delays between neurons and all the other parameters that may be important for their development.

The problem with a computational model however is that not all the physiological effects that play a role in a biological neuronal network can be included, simply because they are not all known. On top of that, many effects such as STDP are studied in very narrow contexts, usually on cell-level and their exact implementation relies on parameters which are unknown in general. But even if we had the knowledge to make a realistic model of a cultured network, it would be too complex to be practically useful in this study. The computational network that is used is therefore not supposed to accurately mimic a cultured network. Instead, it is a suggestive model for the mechanisms that may take place in a cultured network as well as the data analysis techniques suitable to quantify them.

1.4. Research questions

A cultured neuronal network usually shows network responses to stimuli at some of the electrodes. An exemplary stimulus response is visualized in figure 1.1b. These responses are similar to the network bursts seen in spontaneous activity with respect to the amount

of action potentials and the duration. If the connectivity and activity in a network indeed form a balance that can be disturbed by stimulation, this gives rise to the idea that, after repeated stimulation, the new connectivity should incorporate the induced activity and that the new the spontaneous activity should show traces of the stimulus responses. This leads to the following research question:

> **Do the stimulus responses in a cultured neuronal network, induced by repeated, low-frequent stimuli at a single electrode, show a higher similarity with the spontaneous activity after stimulation than with the spontaneous activity before stimulation?**

Due to practical limitations of a MEA experiment, we start our investigations with a computer simulated network. This leads to the following sub-questions:

1. Can a small-scale computational network form a stable equilibrium between activity and connectivity, such that the synaptic weights do not show significant changes on a timescale similar to the length of MEA experiments (hours to days)?
2. How does the stimulus response evolve during repeated stimulation of the computational network?
3. How can the similarity between network bursts and stimulus responses be measured and is this similarity significantly altered by the stimulation?
4. Can the results gained by simulations be applied to cultured neuronal networks?

2. Methods

2.1. Computational model

The art of modeling is to choose the minimum level of detail that is able to produce the functionality relevant for answering our research questions. Non relevant phenomena not only slow down the simulation, they may also obscure the origin of observed behavior. Our simulation is based on the simulation developed by [39], which has its origins in [10]. The model consists of several submodels found in literature that are merged together. In the following sections, we describe our modeling and implementation choices.

2.1.1. Structure

In agreement with [39], we modeled a very small network of 100 neurons. On the one hand, this allowed for fast simulations, on the other hand, all neurons and synapses could be observed simultaneously on an individual level. We did not see a reason to alter the ratio of 4-to-1 between excitatory and inhibitory neurons (see section 2.1.3), which is in agreement with biological data [28].

A cultured network consists of up to 150 000 neurons [24]. Every neuron is, on average, connected to ca. 1000 other neurons [13]. This corresponds to a density of less than 1 %. To compensate for the small amount of neurons, the synaptic density in [39] is increased to 50 %, such that every neuron is on average connected to 50 other neurons instead of only one. A density of 100 % was avoided since this leads to solely bilateral connections. In addition to the increased density of the connections, the synaptic efficacies, i.e. the amplitude of post synaptic potentials induced by the synapses, are significantly higher than those found for real synapses, see section 2.1.3. In accordance with the findings in [39], we modeled the axons as delays that define the time between the spiking of a neuron and the arrival of the corresponding action potential at the synapse. The delays are drawn from the continuous uniform distribution $[1\,\text{ms}, 10\,\text{ms}]$. Delays are a computational burden, but we assumed that they are essential since exact timings play a crucial role in STDP. We did allow neurons to form a synaptic connection with themselves, called autapses, but with our plasticity mechanism, their synaptic strength quickly approaches zero as described in section 2.1.3.

In [39], neuron positions are not incorporated into the model. Therefore, both the chance of two neurons being connected as well as the axonal delay of the connection are independent of the neurons. From the results, we did not conclude that the positions are relevant and since we wanted to keep the model as simple as possible, we did not include them either.

2.1.2. Neurons

Probably the most famous neuron model is the Hodgin-Huxley model, published in 1952 [12]. The benefit of being very close to biological reality is counteracted by its simulation complexity. Since then, several models have been published that gain computational efficiency by abandoning physiological realism. A very promising and therefore widely

used model is the Izhikevich neuron [14, 15]. It is based on the bifurcation theory applied to dynamical systems and is able to reproduce all the dynamics shown by neurons found in the mammalian brain, yet it is computationally effective. The status of a neuron is described by two variables, its membrane potential v and a recovery value u. v is a one-to-one representation of its natural counterpart and denoted in mV. The behavior is then defined by a system of two differential equations:

$$\frac{\mathrm{d}v}{\mathrm{d}t} = 0.04v^2 + 5v + 140 - u + I$$
$$\frac{\mathrm{d}u}{\mathrm{d}t} = a \cdot (b \cdot v - u)$$

I is the input current and the four parameters a, b, c, d define the neuron type. A neuron fires an action potentials when $v \geq 30\,\mathrm{mV}$. This is not a threshold as used by integrate-and-fire models but simply the upper bound of the voltage after which the two variables are reset:

$$\text{if} \quad v \geq 30\,\mathrm{mV}, \quad \text{then} \begin{cases} v \leftarrow c \\ u \leftarrow u + d \end{cases}$$

In [14], the author shows an exemplary usage of his model in a simulation with 1000 neurons. To create heterogeneity, every neuron has its own parameters and therefore dynamics. The parameters of the excitatory neurons (see section 2.1.3) are a linear interpolation between the parameters of *regular spiking* (RS) and *chattering* (CH) neurons, biased towards RS neurons. For every neuron, a variable r is chosen from the uniform distribution $U(0, 1)$ and then squared to bias it towards zero. The four parameters a, b, c, d of the neuron are then calculated as $a = (1 - r) \cdot a_{\mathrm{RS}} + r \cdot a_{\mathrm{CH}}$, $b = (1 - r) \cdot b_{\mathrm{RS}} + r \cdot b_{\mathrm{CH}}$ etc. A value of $r = 0$ therefore corresponds to a RS neuron, a value of $r = 1$ to a CH neuron. Likewise, the parameters of the inhibitory neurons are chosen randomly between the parameters of *fast spiking* (FS) and *low threshold spiking* (LTS) neurons, without a bias towards either type.

The forward Euler method is used to solve the differential equations in the discrete time domain because it is simple to add noise, see section 2.1.4. Since the voltage v of a neuron is changing relatively fast compared to its recovery variable u, v is updated twice as often with half the time step used for u. In the example, the author of [14] uses an update sequence that, looks like

for all neurons n **do**
$v_n \leftarrow v_n + 0.5 \cdot dt \cdot (0.04v_n^2 + 5v_n + 140.0 - u_n + I_n)$
$v_n \leftarrow v_n + 0.5 \cdot dt \cdot (0.04v_n^2 + 5v_n + 140.0 - u_n + I_n)$
$u_n \leftarrow u_n + dt \cdot a_n \cdot (b_n \cdot v_n - u_n)$
end for

The order of this updating has been changed to

for all neurons n **do**
$v_n \leftarrow v_n + 0.5 \cdot dt \cdot (0.04v_n^2 + 5v_n + 140.0 - u_n + I_n)$
$u_n \leftarrow u_n + dt \cdot a_n \cdot (b_n \cdot v_n - u_n)$
$v_n \leftarrow v_n + 0.5 \cdot dt \cdot (0.04v_n^2 + 5v_n + 140.0 - u_n + I_n)$
end for

since updating u in between v showed better results, see appendix A.1.

Stimulation

When it comes to stimulation, the computer simulation fully shows its potential since individual neurons can be stimulated with arbitrary stimulation patterns. Technically, a neuron is stimulated by setting its potential to 30 mV, such that an action potential is detected at the corresponding neuron the next time it is updated.

2.1.3. Synapses

Synapses connect the outgoing axons of one neuron with the dendrites of a target neuron and can be either electrical or chemical. We modeled exclusively chemical synapses since they are predominant in natural networks. According to biological observations, the axons of one neuron either connect solely to inhibitory or solely to excitatory synapses. Inhibitory synapses cause hyperpolarization of the postsynaptic neuron's membrane potential while excitatory synapses cause depolarization and thereby facilitate postsynaptic action potentials. Neurons, whose axons connect to excitatory synapses are called excitatory neurons and likewise neurons with axons connecting to solely inhibitory synapses are called inhibitory neurons.

The synaptic efficacy can change on a short or long time-scale. Long-term plasticity effects are necessary in this research since long-lasting changes are required for network adaptation, but also the short-term plasticity effects may be indispensable. Long-term plasticity effects occur mainly during bursts and the shape of bursts is strongly affected by short-term plasticity. For instance, it is not purely the activation of inhibitory synapses that ends a burst, but also the short-term depression of excitatory synapses caused by a depletion of the available neurotransmitters. Deactivation of inhibitory synapses has an effect on the burst duration and intensity but does not lead to endless bursting [18]. Synapses can also temporary increase in strength when activated in short succession, a mechanism called short-term facilitation. It is plausible that short-term facilitation has a significant effect on the shape of a burst and thereby long-term plasticity as well.

Accordingly, a complex synapse model that incorporates both forms of short-term plasticity [23] is used in [39]. Two variables, y and R model the state of a synapse[1]. A synapse has a pool of available neurotransmitter and when triggered by an incoming action potential, part of this neurotransmitter pool is released. The variable y_n is the fraction of the available neurotransmitter pool that is released by the nth spike arriving at the synapse. In compliance with short-term facilitation, it increases when presynaptic spikes arrive in quick succession and returns back to its initial value U over time. R_n describes the relative amount of neurotransmitter available at the synapse after the nth spike. R decreases every time a spike arrives at the synapse, causing the release of neurotransmitters, and, over time, increases back to one.

$$y_{n+1} = y_n \cdot \exp\left(\frac{-\Delta t}{\tau_{\text{facil}}}\right) + U \cdot \left(1 - y_n \cdot \exp\left(\frac{-\Delta t}{\tau_{\text{facil}}}\right)\right)$$

$$R_{n+1} = R_n \cdot (1 - y_{n+1}) \cdot \exp\left(\frac{-\Delta t}{\tau_{\text{rec}}}\right) + 1 - \exp\left(\frac{-\Delta t}{\tau_{\text{rec}}}\right)$$

τ_{facil} and τ_{rec} are the time constants for facilitation and recovery respectively and Δt denotes the time that has elapsed since the previous spike. The variable R_n describes the relative amount of neurotransmitter available *after* the nth spike. In order to calculate

[1] The original paper calls this variable u but we want to avoid confusion with the recovery variable u from the neuron model.

the amount of neurotransmitter released by a spike, we need to know the amount of neurotransmitter available *at* the arrival of the spike. The only change would be

$$R^*_{n+1} = R^*_n \cdot (1 - y_n) \cdot \exp\left(\frac{-\Delta t}{\tau_{\text{rec}}}\right) + 1 - \exp\left(\frac{-\Delta t}{\tau_{\text{rec}}}\right)$$

which is minor since the difference between consecutive y's is small. The relative amount of neurotransmitters released by a spike is estimated by $y_n \cdot R_n$.

Neurotransmitters activate selective ion channels at the postsynaptic terminal, which change the conductivity to that ion. This conductivity change follows a certain profile, which may be modeled as the difference of two exponentials that describe the increase and decrease of the conductivity [4]. In common with [39], we use a simpler function and state that the arrival of an action potential at the synapse leads to an instantaneous increase in conductivity proportional to the amount of neurotransmitter released. The conductivity shows an exponential decay with a time constant τ_{syn}. This time depends on the neurotransmitter that is utilized by the synapse. Similar to [39], we assume that all the synapses are of the type AMPA and GABA_A which have time constants of 5.26 ms and 5.6 ms respectively [7]. For simplicity, all synapses use a τ_{syn} of 5 ms.

Dendrites

A change in conductivity causes a current flow into or out of the dendritic tree. In reality, the magnitude of this current depends on the potential difference across the cell membrane. We ignore this since the fluctuations in the cell membrane potential are rather small as long as the cell does not fire an action potential, and say that the current through the postsynaptic terminal is given by

$$\frac{\mathrm{d}I_{\text{syn}}(t)}{\mathrm{d}t} = -\frac{I_{\text{syn}}(t)}{\tau_{\text{syn}}} + A \cdot S(t) \cdot y_n \cdot R_n \cdot \delta(t - t_n)$$

The first term on the right side describes the exponential decay of the current, the second term any instantaneous increases of the current due to an action potential arriving at $t = t_n$. The current is proportional to the amount of released neurotransmitters ($y_n \cdot R_n$), a constant scaling factor which depends on the synapse type (A) and $S(t)$. S is the relative synaptic strength or weight in the interval $[0, 1]$ for excitatory synapses, or [-1, 0] for inhibitory synapses. This synaptic strength is a function of time since it is affected by the long-term plasticity mechanism described in the following section. Since $S(t)$, y_n and R_n are relative variables bounded at ± 1, the scaling factor A equals the maximum current that can be "generated" by a synapse.

Dendrites cause the summation and filtering of all the currents arriving at a neuron [4]. As in [39], we did not explicitly include dendrites in our model. In order to determine the input to a neuron at a given time, we simply add all the momentary currents of the synapses onto this neuron. The distributive property of the exponential decay makes this a computational relief: instead of calculating the decay at all the synapses individually, we can sum them first and then apply the exponential decrease to the total input current of a neuron.

Long-term plasticity

As a long-term plasticity mechanism, we use spike-timing-dependent plasticity, which can cause both an in- and decrease in synaptic strength. It is unknown to which extend STDP

can account for the long-term synaptic plasticity seen in biological neuronal networks, but it is shown in [16] that the famous BCM learning rule (named for E. Bienenstock, L. Cooper and P. Munro), which relates synaptic changes to the firing frequency of pre- and postsynaptic neurons, is not necessarily a different mechanism and may be compatible with STDP.

Many different STDP models exist, for a review see [27]. Most of them have in common that they relate the change in synaptic strength of excitatory synapses to the relative timing between the spikes of the pre- and postsynaptic neurons. If the presynaptic neuron fires before the postsynaptic neuron, the strength is increased and vice versa, as depicted in figure 2.1. Models are constructed by fitting a function to the measured data. In the simplest case, the function is a combination of two exponential functions that are symmetric with respect to the origin. More complex functions introduce asymmetry between the change in strength for potentiation and depression. In addition, a continuous function may be used in order to avoid the discontinuity at $t = 0$. Similar to [39], we use one exponential function for either side. Both exponentials have the same time constant, but we do allow asymmetry in the magnitude of change for potentiation and depression of the synapse. Our relative synaptic efficacies are denoted by S. In accordance with equation (10) in [27, p. 465] this leads to

$$\Delta S = \begin{cases} F_-(S) \cdot \exp\left(-\dfrac{\Delta t_{\mathrm{STDP}}}{\tau}\right) & \text{if } t_{\mathrm{pre}} > t_{\mathrm{post}} \\ F_+(S) \cdot \exp\left(-\dfrac{\Delta t_{\mathrm{STDP}}}{\tau}\right) & \text{if } t_{\mathrm{pre}} < t_{\mathrm{post}} \end{cases} \tag{2.1}$$

where Δt_{STDP} is the absolute time between a pre- and postsynaptic spike. F_- and F_+ are factors that determine the magnitude of change.

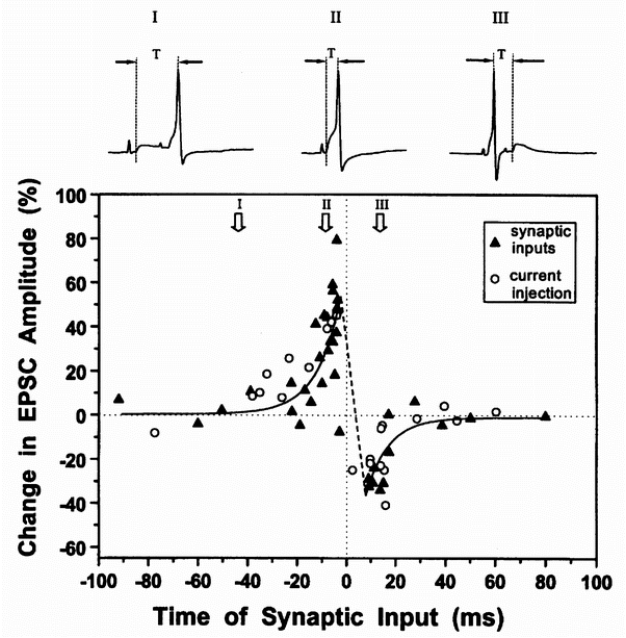

Figure 2.1.: Measurement data for synaptic potentiation/depression as a function of relative spiking between the pre- and postsynaptic neuron. Negative and positive times are fitted separately with an exponential function. Reprinted with permission from [44].

In general, F is a function of the current synaptic strength and this dependence is used to distinguish between additive and multiplicative models. In an additive model, the weight change does not depend on the current weight, i.e. F is independent of S. In multiplicative models on the other hand, F is proportional to the current weight, i.e. $F_-(S) \propto S$ and $F_+(S) \propto (1-S)$. This can be generalized to the form $F_-(S) \propto S^\mu$ and $F_+(S) \propto (1-S)^\mu$. In the additive model, $\mu = 0$ and in the multiplicative model, $\mu = 1$. Any value in between describes an intermediate form [27]. Additive models tend to lead to bimodal synaptic distributions, since strong synapses are likely to become even stronger, while a decrease in synaptic strength precedes further depression [32]. Multiplicative models are self-regulating since the stronger a synapse, the slower it will increase and the weaker a synapse, the slower it will decrease in strength. This can lead to a unimodal weight distribution in the network, which seems to be more plausible from a biological standpoint and is in agreement with experimental data obtained presented in [35]. The unimodal weight distribution found in natural networks suggests that multiplicative models are a better representation of reality. Experimental data of synaptic changes induced by paired spikes imply that synaptic depression is indeed multiplicative [2, 26]. The data for potentiation is less clear and we therefore used the same updating rule as the used for the model in [39], i.e. a combination of the multiplicative and additive model. Equation 2.1 becomes

$$\Delta S = \begin{cases} A_- \cdot S(t) \cdot \exp\left(-\dfrac{\Delta t_{\mathrm{STDP}}}{\tau}\right) & \text{if } t_{\mathrm{pre}} > t_{\mathrm{post}} \\ A_+ \cdot \exp\left(-\dfrac{\Delta t_{\mathrm{STDP}}}{\tau}\right) & \text{if } t_{\mathrm{pre}} < t_{\mathrm{post}} \end{cases}$$

with S bounded at one. Additive increase and multiplicative decrease is sufficient to be self-regulating and as is demonstrated by its use in TCP (Transmission Control Protocol) congestion avoidance. The results presented in [39] support the fact that this model can lead to unimodal synaptic weight distributions. In their model, the authors did not limit the synaptic weights however, leading to a small number of ever-growing synapses, especially after stimulation.

Another model choice for STDP is the spike pairing scheme that specifies which spike

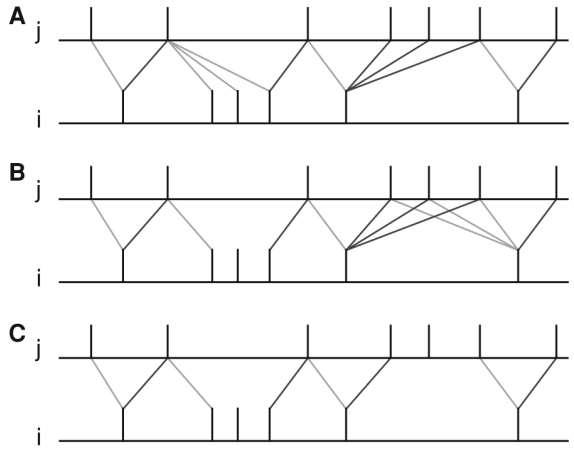

Figure 2.2.: Three different nearest neighbor schemes for a presynaptic neuron j and postsynaptic neuron i. We chose the reduced symmetric scheme C for our model. Used with permission from [27].

pairs cause a change in synaptic strength. For instance in an all-to-all scheme, all the spike pairs with a time difference within the STDP-window contribute to a weight change. In a nearest neighbor scheme, only the closest pre- and postsynaptic spikes form pairs. Still, this leaves room for different interpretations as shown in figure 2.2. Both experimental measurements [9] and theoretical considerations [16] lead to the conclusion that a reduced symmetric scheme resembles reality best. We therefore decided to use the reduced symmetric scheme (C), which is different from the model in [39], for which the symmetric scheme (A) was used.

Finding STDP rules for inhibitory synapses is difficult if we were to follow the same logic as for excitatory synapses. Since the "purpose" of an inhibitory synapse is to hinder the target neuron from firing an action potential, the magnitude of its strength should increase whenever the target neuron does not fire after the inhibitory neuron spiked. But it is impossible to give an exact timing for when a neuron is not-firing. Therefore, the relative strengths of all inhibitory synapses are fixed to -0.5.

As mentioned in section 2.1.1, we allow neurons to connect with themselves, forming autapses. Whenever an autapse fires an action potential, this action potential will arrive at the synapse a couple of milliseconds later, according to the axonal delay. From the synapse's point of view, there was a very recent postsynaptic spike. This leads to a decrease in synaptic strength in accordance with the STDP rule. The synaptic weights of autapses therefore quickly approach zero in our model.

2.1.4. Intrinsic activity

During moments of no stimulation, there is no external input to the network. Yet, we want the network to show spontaneous activity, similar to the spontaneous activity seen in cultured networks, so we have to add intrinsic activity in some form. There are several ways to initiate spontaneous activity and bursting in a computational network [10]. Methods found in literature include the use of pacemaker neurons, which are neurons that fire intrinsically on a regular basis and initiate network responses, as well as noise injection through the synapses of the neurons. In the latter case, noise can be added either as white noise or as Poisson spike trains. The level of noise added in the second case must be high enough to cause spontaneous spiking of cells, making them intrinsically active. Functionally, they are therefore similar to pacemaker neurons. The difference lies in the fact that in the case of noise injection, usually all neurons receive the same amount of noise, whereas pacemaker neurons make up only a small part of the neurons in a network and fire with a regular frequency. In [39] satisfying results are obtained with Gaussian noise added to the input currents of all neurons and we did not see a reason to deviate from this method. The membrane potential of a neuron is then updated utilizing the Euler-Maruyama method according to:

for all neurons n **do**

$$v_n \leftarrow v_n + 0.5 \cdot dt \cdot (0.04v_n^2 + 5v_n + 140.0 - u_n + I_n) + \sqrt{0.5 \cdot dt} \cdot N(0, \sigma)$$

$$u_n \leftarrow u_n + dt \cdot a_n \cdot (b_n \cdot v_n - u_n)$$

$$v_n \leftarrow v_n + 0.5 \cdot dt \cdot (0.04v_n^2 + 5v_n + 140.0 - u_n + I_n) + \sqrt{0.5 \cdot dt} \cdot N(0, \sigma)$$

end for

The model in [39] uses the same amount of noise for inhibitory and excitatory neurons. Since inhibitory neurons are triggered more easily due to different parameter settings, the spiking frequency of inhibitory neurons was much higher. We therefore chose individual noise levels for inhibitory and excitatory neurons.

2.1.5. Parameter choices and initial values

The parameters values of the different models as well as their origin are summarized in table 2.3. The initial values of the variables can be found in table 2.2. In case multiple values are listed, different values were used for different experiments.

2.1.6. Simulations

In this section, we list all the different simulations that have been run. The first subquestion of this research deals with the development of an equilibrium between the activity and connectivity in the computational network. We therefore start with networks that are unstimulated and observe their "natural" development over time.

Unstimulated networks

We first simulated networks that use the parameter settings defined as "standard parameter settings" in [39]. Based on intermediate results, the following parameters/initial values were then varied

- Noise on inhibitory neurons
- Initial weights of excitatory synapses
- Axonal delays
- Neuron types

We use six different network settings as shown in table 2.1 with three different network realizations each, therefore simulating 18 different networks for 10 days of "network time" each. Since the axonal delays, neuron parameters and connections between neurons are chosen at random, the different realizations were obtained by reseeding the random number generator before generating the network. After the network has been generated, individual parts of the network can be reinitialized without affecting the rest of the network.

Using the standard settings, some of the inhibitory neurons spike with a very high frequency compared to other neurons (see figure 3.1). These neurons are low-threshold spiking (LTS) neurons, which, as their name indicates, need very little input to fire. We simulated the same networks again, but this time with all inhibitory neurons set to

Table 2.1.: The six different network settings used for simulations without stimulation.

Description	Noise σ exc./inh.[1]	Initial excitatory weigths	Axonal delays	Neuron types
Standard settings	2.2/2.2	0.5	$\sim U(1,10)$	all
No LTS[2]	2.2/2.2	0.5	$\sim U(1,10)$	no LTS[2]
Less inhibitory noise	2.2/0.8	0.5	$\sim U(1,10)$	all
Different noise realization	2.2/0.8	0.5	$\sim U(1,10)$	all
Random initial weights	2.2/0.8	$\sim U(0,1)$	$\sim U(1,10)$	all
Larger axonal delays	2.2/0.8	0.5	$\sim U(1,20)$	all

[1] On excitatory/inhbitory neurons. In $\mathrm{mV}/\sqrt{\mathrm{ms}}$.
[2] LTS = low threshold spiking

fast-spiking neurons. Instead of removing the LTS neurons, the amount of spikes fired by inhibitory neurons can also be decreased by reducing the amount of noise that is applied to them. In the third setting the standard deviation of the noise that is added to the potentials of the inhibitory neurons was therefore reduced to 40 % of its original value. To examine the effect of the noise applied to the network, the previous networks were simulated again but the random number generator was reseeded after network initialization. This led to a different realization of the noise that is applied to membrane potentials. In addition to using a different noise realization, we also simulated three networks with different initial weights. Instead of initializing the weights of all excitatory synapses to 0.5, they were chosen randomly from the uniform interval $[0, 1]$. The bursts in the previous networks are relatively short compared to the bursts in cultured networks. With the last network settings, we examined the effect of longer axonal delays by running three simulations with axonal delays that, on average, were twice as large as the axonal delays in the previous simulations.

Stimulated networks

The protocol used to determine the effect of stimulation on the simulated networks was very similar to the measurement protocol described in section 2.2.3, i.e. it consisted of three phases: measuring the spontaneous activity in the network, stimulate the network with low-frequent stimuli and than measuring the spontaneous activity again. It differed in two ways. First, three hours of spontaneous activity were measured before and after stimulation instead of only one because the computational networks showed a lower bursting frequency than their biological counterparts. Second, we simultaneously stimulated three excitatory neurons instead of one electrode. Three neurons were the least number of neurons leading to a consistent stimulus response.

As the initial networks, we used the three networks with less noise applied to the inhibitory neurons (third row in table 2.1) after the weights were allowed to settle for 10 days. Four simulations with different stimulation neurons were run for each network, leading to a total of 12 simulations. Since all the neurons are statistically the same, we chose excitatory neuron 1 to 3 as the stimulation neurons for the first simulation of a network, 4 to 6 for the second, etc.

2.1.7. Implementation

Due to personal experience and speed considerations, the whole simulation was written in C++. A simple text-based user interface and graphical output allow to control and observe the simulation while it is in progress. Commands can also be executed from a file. This way, experiment protocols can be written in an easily interpretable format. The proper working of the simulation was verified as described in appendix A.

The computer simulation developed for this research gives us a very flexible platform to run specifically designed experiments. Any internal state of the simulation that is modeled, such as membrane voltages, action potentials and synaptic weights can be observed and recorded to a file for post processing. The CPU of the computer mostly determines the speed of the simulation. On a mediocre personal computer with a 2.66 GHz CPU, 133 seconds of network time can be simulated in one second of real-time, i.e. the simulation runs more than 100 times faster than real-time.

Table 2.2.: Initial values of all the variables used in the individual models.

	Parameter	Sym.	Init. value(s)		Origin
Neuron					
	Membrane pot. [mV]	v	-65		[14]
	Recovery variable	u	$-65 \cdot b$		[14]
Synapses			*Excitatory*	*Inhibitory*	
	Synaptic strength	S	0.5, $\sim U(0,1)$	-0.5	[39]
	Neurotrans. utilization	y	U	U	[23]
	Neurotrans. pool	R	1	1	[23]

Table 2.3.: Values of all the parameters used in the simulation.

	Parameter	Sym.	Value(s)				Origin
Structure							
	Number of neurons	N	100				[39]
	Num. of exc. neurons	N_e	80				[39]
	Num. of inh. neurons	N_i	20				[39]
	Connection density	˜	50 %				[39]
	Axonal delays	˜	$\sim U(1,10)$, $\sim U(1,20)$				[39]
Neuron			*Excitatory*		*Inhibitory*		
	Auxiliary variable	r	$\sim U(0,1)$		$\sim U(0,1)$		
	Recovery speed of u	a	0.02		$0.1 - 0.08 \cdot r$		[14]
	Sensitivity of u to v	b	0.2		$0.25 + 0.05 \cdot r$		[14]
	Reset value of v [mV]	c	$-65 + 15 \cdot r^2$		-65		[14]
	Reset value of d	d	$8 - 6 \cdot r^2$		2		[14]
Synapses			*EE*	*EI*	*IE*	*II*	
	Resting value of y	U	0.59	0.049	0.16	0.25	[10, 11, 23]
	Max. efficacy [$\mathrm{V\,s^{-1}}$]	A	10.8	32.4	43.2	43.2	[38, 39][1]
	Facil. const. [ms]	τ_{facil}	0	1797	376	21	[10, 11, 23]
	Recovery const. [ms]	τ_{rec}	813	399	45	706	[10, 11, 23]
	Synaptic const. [ms]	τ_{syn}	5	5	5	5	[7]
STDP							
	Potentiation rate	A_+	0.005				[36]
	Depression rate[2]	A_-	$A_+ \cdot 1.05 \cdot 2 = 0.0105$				[36, 39]
	Time constant [ms]	τ_{STDP}	20				[36]
Noise							
	On exc. neurons[3]	σ	2.2				[39]
	On inh. neurons[3]	σ	0.88, 2.2				[14, 39][4]

[1] The ratios between the different synapses are from [38]. The values have been scaled in [39] to compensate for the small network.

[2] According to the model in [36], the depression rate is scaled by a factor 1.05 compared to the potentiation rate. The factor 2 is introduced in [39] since the change in weight is proportional to the current weight, which is expected to be 0.5 on average.

[3] The unit of this parameter is $\mathrm{mV}/\sqrt{\mathrm{ms}}$.

[4] A value of 2.2 is used in [39]. In a minimal working example provided by the author of the neuron model [14], the inhibitory neurons receive 40 % of the noise of the excitatory neurons. $2.2 \cdot 0.4 = 0.88$.

2.2. Experimental setup

Based on the results gained by the simulations (section 3.1), we designed an experiment to see whether the findings on simulated neuronal networks are applicable to in vitro neuronal networks as well. This section describes the complete experiment, from the preparation of the cultures to the execution of the measurement protocol.

2.2.1. Cultures

We used commercially available MEAs (Multi Channel Systems, Reutlingen, Germany) with 60 electrodes in a 8×8 grid configuration as shown in figure 2.3. The electrodes are 30 µm in diameter and have a pitch of 200 µm. A glass ring with an inner diameter of 19 mm is glued on top of the MEA. For better adhesion of the cells, the MEAs were coated with polyethylene-imine (PEI) before plating.

Cortical cells were obtained from Wistar rats one day after birth. After physically isolating the cortex, it was first dissociated chemically using Trypsin and then mechanically by trituration. This way a cell solution with a concentration of 3×10^6 cells/ml was obtained. MEAs were plated with approximately 50 µl of this solution, i.e. approximately 100 000 cells. After allowing the cells to adhere for 2 h in an incubator, loose cells were removed by rinsing with 500–700 µl of Dulbecco's Modified Eagle Medium (DMEM). Finally, 600 µl of R12 [29] was added.

Cultures were stored in an incubator preserving a temperature of 36 °C, 100 % humidity and a CO_2 level of 5 %. The growth medium was refreshed three times a week by removing 300 µL of the old medium and then adding 400 µL of fresh R12 medium.

2.2.2. Hardware and data acquisition

During the measurement, the conditions of the incubator were mimicked. The MEAs were temperature regulated at 36 °C using a TC01 temperature controller (Multi Channel

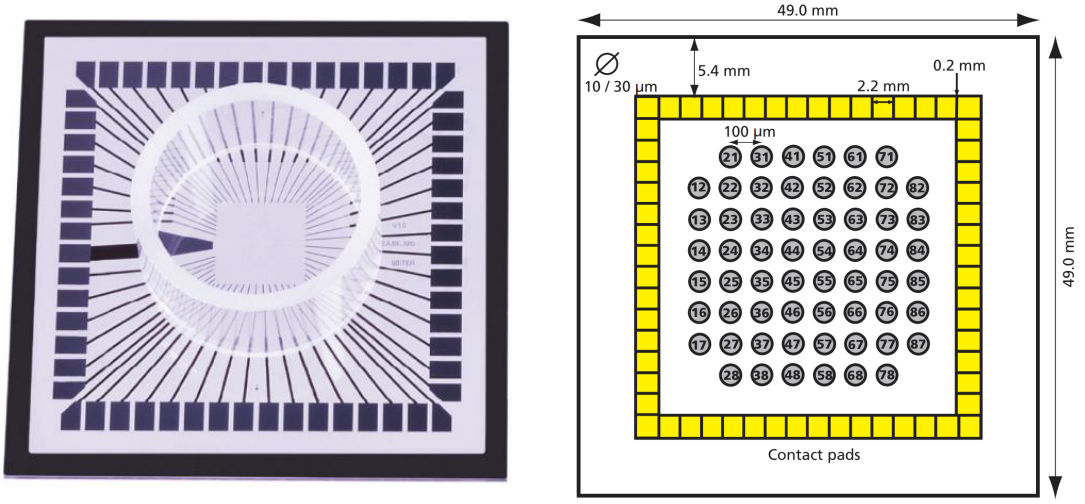

(a) (b)

Figure 2.3.: (a) shows the multi-electrode array used in our experiments. (b) illustrates its electrode configuration and numbering. From the Multi Channel Systems website/MEA manual, with permission.

Systems) and covered by a case made of acrylic glass. Fresh moisturized air with 5 % CO_2 was continuously fed to the interior of the case.

A combination of a 1060BC preamplifier and FA60s filter amplifier (both from Multi Channel Systems) filtered the potentials measured at all the electrodes with a 100 Hz to 6 kHz bandpass filter and amplified them by a factor of 1000. After that, the signals were sampled at a frequency of 16 kHz using a PCI-6024E DAQ card (National Instruments, Austin, TX). Stimuli were applied with a STG1002 stimulus generator (Multi Channel Systems).

Data acquisition and stimulation protocols were performed with a custom made Labview program that ran on a desktop PC. Spikes were detected by continuously estimating the individual noise levels at all electrodes. Whenever the amplitude of the signal at an electrode exceeded 5.5 times the standard deviation of the noise at that electrode, an event was detected. For every event, the time-stamp, electrode number, noise level and a small section of the waveform were stored to the hard disk for subsequent analysis. In most cases, an event resembles an action potential, yet off-line spike-validation was done later on as described in section 2.3.

2.2.3. Measurement protocol

A fairly simple measurement protocol was used in order to effectuate a comparison between the simulation and the actual experiments. One experiment contained only one stimulation period. Before and after stimulation, the spontaneous activity of the network was recorded. We used single-electrode, low-frequency stimulation with an inter stimulus interval (ISI) of 8 s. The ISI was chosen as a compromise between two conflicting issues. On the one hand we wanted to give the network enough time to recover from its previous stimulus before stimulating again, which takes approximately 10 s in the simulation. On the other hand we wanted to avoid network bursts in between the stimulus responses, such that stimulus responses make up a substantial part of the network activity during the stimulation phase. Two types of stimulation are common in MEA studies, low-frequent and tetanic stimulation. It is not clear whether one type of stimulation has a larger effect on connectivity changes. Tetanic stimulation seems to predominate in literature but this may be because tetanic stimulation was used in early long-term-potentiation and long-term-depression research on cell-level. Due to its results there, it may have simply been adopted to cultured networks. In low-frequent stimulation, part of the network is activated once and from there, the existing pathways in the network determine the further propagation of the burst. In tetanic stimulation, where multiple stimuli are applied in quick succession, the first stimulus in such a train may already be sufficient to start the burst. Parts of

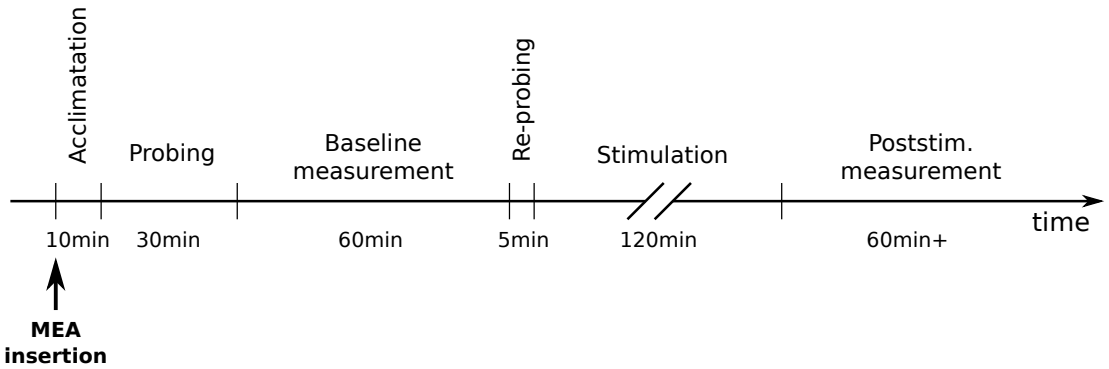

Figure 2.4.: Overview of the measurement protocol used in the MEA experiments.

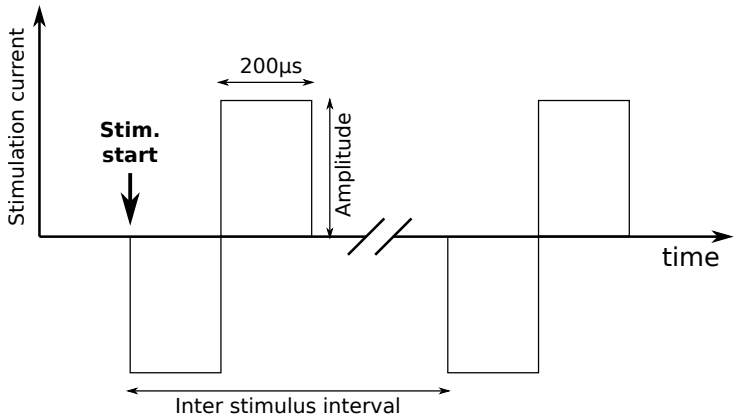

Figure 2.5.: Biphasic waveform pattern used for stimulation in the MEA experiments.

the network are then stimulated again while the "natural" stimulus response is already in progress. On the one hand, this may lead to stronger changes in the connectivity because pathways may be activated that do not exist yet. On the other hand, spatio temporal activity patterns may show up in the stimulus response that originate from restimulating the network and not from actual pathways in the network. We therefore decided to use low-frequent stimulation.

The protocol is visualized in figure 2.4. After installing the MEA in the measurement setup, we waited for at least 10 minutes in order for the culture to adjust to its new environment. We then started a probing sequence. During probing, all the electrodes were stimulated two times with three different amplitudes, resulting in a total of six stimulations. Since the contact impedances of the electrodes vary, we stimulated with a fixed current. A biphasic stimulation waveform as shown in figure 2.5 avoided the undesired buildup of charge at the electrodes. The three stimulation amplitudes were 12 µA, 16 µA and 24 µA and the interval between two stimulations was 4 s, which is half of the ISI used for later stimulations. The reason for the initial probing was twofold. First, it showed whether the network is healthy and responds to stimuli. Second, it allowed a first selection of electrodes showing a consistent and strong response to stimuli.

In case of a suitable culture, approximately one hour of spontaneous activity was measured, serving as a baseline. Now the candidates for stimulation that were found in the first probing sequence were re-probed. This time, five stimuli were applied with an ISI of 8 s. A single amplitude was chosen for all stimuli on all electrodes. The reason was that the amplitude of the stimulation current needed for a reliable network activation was usually consistent for all electrodes within one culture. We selected the electrode showing the best response in this second probing session as the stimulation electrode, based on visual inspection of the post-stimulus time histograms (PSTHs). It was stimulated for 2 h with a frequency of 0.125 Hz, resulting in 900 stimuli. After stimulation, at least 1 h of spontaneous activity (in practice usually more) was measured. Table 2.4 summarizes the stimulation settings for the probing, re-probing and stimulation periods.

2.3. Data analysis

The experimental setup described above generates a list of events, where every event is likely to be an action potential measured at one of the electrodes. Every possible action potential is first validated as described in the section below, leading to a revised list of spikes. Every spike is a combination of the electrode number that recorded it and the

time instance it was recorded. With the computational model, we can also generate a list of action potentials (that do not need any validation). Every action potential in the simulation consists of the time instance at which it occurred and the neuron instead of electrode number. For the rest of this chapter, the term neuron may refer to an actual neuron in the simulation or to an electrode in the measurement setup, which most likely records action potentials from a single neuron.

This raw spike data is typically viewed in a raster plot where every spike is depicted with a dot or bar. Because of the immense amount of spikes and the fact that these raster plots can expose only very coarse properties of the network, further processing needs to be done in order to produce meaningful, comparable and interpretable results. This chapter gives an overview of the analytical methods that were used in our research. Most important is the comparison of spontaneous network activity (network bursts) and induced network activity (stimulus responses). In accordance with [33], we will both denote as a synchronous bursting events (SBEs).

2.3.1. Spike validation

The events recorded by our measurement setup as described in section 2.2.2 usually correspond with action potentials fired by neurons located on top of the electrodes. Sometimes however events are just artifacts, for instance induced by stimulation, and we therefore check every event for validity. Spikes are validated based on (a) the amount of simultaneous spikes at a certain time instance as suggested in [37] and (b) the waveform of the spike as suggested in [40]. If more than four spikes were recorded at the same sampling interval ($T = 1/16\,\mathrm{s}$), all spikes were removed since they are most likely a stimulus artifact. In addition, according to figure 1 of [40], the waveform of the action potential had to satisfy the following conditions:

- In a window of $\pm 1\,\mathrm{ms}$ around the center of the detection, there may not be an additional peak with an absolute value that is greater or equal to the amplitude of the detected spike.
- In the same window, but outside the interval $\pm 0.5\,\mathrm{ms}$ around the center of the detection, the measured potential may not exceed 50 % of its peak value, retaining the sign of the peak.
- In the same window, but outside the interval $\pm 0.15\,\mathrm{ms}$ around the center of the detection, the measured potential may not exceed 90 % of its peak value, retaining the sign of the peak.

The last two criteria imply that the action potential shows a sharp peak. This validation step led to the removal of approximately 0.20 % of all spikes.

Table 2.4.: Stimulation settings during the probing, re-probing and stimulation periods of the experiment.

Period	Electrodes	Amplitude [µA]	Stimuli per electrode	ISI [s]
Probing	All	12, 16 and 24	2×3	4
Re-probing	Selection	One of the above	5	8
Stimulation	One	One of the above	900	8

2.3.2. Stimulus validation

In vitro, one can distinguish between direct and late responses. The direct response consists of action potentials that are fired because of direct stimulation by the electrode(s) and therefore do not depend on synaptic transmission. Action potentials that do require synaptic transmission form the late response [41]. Even if a stimulus causes a late response most of the time, it can sometimes fail to do so. This is especially true for in vitro experiments but also in the simulations, a stimulus may not lead to network activation, e.g. if the network bursted very recently and the excitatory synapses are still depleted. Since we believe that activity should involve synaptic action in order to cause changes in synaptic weights, we implemented the following algorithm to sort out stimuli that did not lead to a sufficient response.

For every stimulus in an experiment, the amount of spikes in the response is calculated. Graphs of these vectors for the in vitro measurements are shown in appendix C. This vector is then smoothed using a median filter with a window size of 50. The reason for using a median filter is that it is affected less by outliers in the data than for instance a moving average filter. Now for every stimulus response it is checked that

- The amount of spikes in the stimulus response is at least half the smoothed value for this response.
- The amount of spikes in the stimulus response is at least 15.

The first criterion sorts out stimulus responses that strongly differ from the preceding and succeeding stimulus responses. The second criterion was added in case the stimulus response fails to induce network responses over a long time in which case the first criterion does not apply.

2.3.3. Burst detection

An important step in the analysis is the detection of bursts during moments of spontaneous activity. Bursts are vaguely defined as moments where a large part of the network synchronizes, i.e. many electrodes detect activity simultaneously. Many different burst detection algorithms can be found in literature. We use the algorithm in [8], that first splits the action potentials into subsets and then checks for every subset whether a bursting condition occurs. For details on the implementation, see the next section.

2.3.4. Burst leaders

The idea of burst leaders is that most bursts in a network are preceded by a less-intense preburst and that the first electrode or neuron that is active in such a preburst may convey information about the origin and realization of a burst [8]. The algorithm as described in [8] uses three parameters, n_{burst}, δt_{burst} and $\delta t_{isolated}$. The authors chose these parameters individually for different culture types used in the experiments and the origin of the parameters is not clear. The algorithm has therefore been adapted such that it uses dynamic thresholds.

Spikes are first subdivided into groups, such that the time between the last spike of one group and the first spike of the next group is $> \delta t_{isolated}$. The rationale is that if the interval between two spikes is larger than $\delta t_{isolated}$, the spikes cannot belong to a single burst because there cannot be any "memory" between the spikes [8]. This parameter is therefore basically determined by the sum of the axonal and synaptic delay between neurons. For the simulation we know that the maximum axonal delay is 10 ms (20 ms in

one of the simulations). After a spike arrives at a synapse, the postsynaptic potential will build up and then decrease again. The timescale of this process is a couple of milliseconds, so we choose $\delta t_{isolated} = 15\,\text{ms}$ ($\delta t_{isolated} = 25\,\text{ms}$). For the culture, a similar argumentation holds. Since we cannot observe all the activity in the network, and there may be unrecorded spikes in between two recorded spikes, a value twice as large was chosen: $\delta t_{isolated} = 30\,\text{ms}$.

A subgroup is said to contain a burst, if within the group at least n_{burst} spikes occur within an interval of length δt_{burst}. In this case, the first interval that fulfills this criterion in the group is determined and the first electrode/neuron in this interval denotes the start of the burst. All the spikes before this spike make up the preburst. If the preburst exists, i.e. it contains at least one spike, the electrode/neuron firing the first spike in the preburst defines the leader of that burst [8]. Instead of choosing fixed parameters, we chose both parameters dynamically for every subgroup. n_{burst} is set to the number of unique neurons participating in the subgroup, but at least half of the total number of active neurons in the network. If for example activity was recorded at 54 electrodes during an experiment, n_{burst} was at least 27. For the analysis of the experimental data, the time interval δt_{burst} was set to one fifth of the total subgroup length but at least to 10 ms. For the computational data it was set to one tenth of the total subgroup length.

Similarity

In [8] the leader distribution $\vec{f}$ is defined as a vector with length equal to the number of electrodes/neurons in the network. Every component of the vector denotes the number of bursts led by the corresponding electrode/neuron within a certain period of spontaneous activity. The similarity between two leader distributions is then defined as their normalized scalar product

$$S = \frac{\vec{f_1} \cdot \vec{f_2}}{|\vec{f_1}| \cdot |\vec{f_2}|}$$

The measure is interpreted as follows: "The similarity is bounded between 0 and 1, where $S = 1$ means identical trigger distributions, while a similarity of 0 means that all the triggers have been completely replace by new trigger electrodes." [8, p. 13]

2.3.5. Similarity between synchronous bursting events

Our research focuses on the development of stimulus responses and the similarity between stimulus responses and the bursts in spontaneous activity. We therefore need a significant and easily interpretable measure that compares firing patterns that are realized by neuronal networks. The method of *inter-SBE correlation matrices* [33] is very promising, since it combines spatial and temporal information of activity patterns. Summarized, two SBEs are compared as follows: For every neuron in the network, a vector is created that describes the activity of the neuron in time relative to the start of the SBE. These activity vectors are constructed by first making a histogram of the spike timings of a neuron of that SBE using a time bin of 1 ms. The vectors are essentially binary because the time bin is small compared to the time interval between two consecutive spikes of a single neuron. The vectors are then smoothed using a Gaussian filter. To measure the similarity between two SBEs, for every neuron the cross-correlations between the activity vectors of either SBE are calculated and summed, leading to a function that denotes the correlation of the two SBEs as a function of time-lag. The similarity is then the absolute maximum of this

function, ignoring the time-lag at which this maximum occurs. This way, the measure becomes insensitive to global shifts of the individual SBEs.

Normalization

In [33], the cross-correlation for every electrode is normalized, such that it has a maximum of one in case the firing patterns of the electrodes are identical except for a shift in time. Every electrode can therefore contribute one to the total similarity measure, leading to a maximal similarity that equals the amount of active neurons in a SBE.

In order to allow a proper comparison among the individual simulations/experiments, we normalize the measure so that it ranges between zero and one. The similarity is one if and only if the firing patterns are exactly the same, except for a global time-lag. Normalization could be achieved by dividing the similarity by the number of active neurons in the SBE, which lays equal weight on all the neurons, no matter their firing rate. Instead, we chose to weight the cross-correlation of the individual neurons according to their relative amount of activity within the burst. The amount of activity per neuron is defined as the mean number of spikes they fire in the two SBEs. The total activity is then the sum of all these individual activities.

Smoothing of activity vectors

In the original paper, the width of the Gaussian filter is not specified. Instead, the authors state that the width is chosen according to the firing rate. On the one hand this makes sense: In case the smoothing is independent of the firing rate, two SBEs will be more likely to show a high similarity if their spikes are dense. On the other hand, we observed that some stimulus responses showed similar late responses, but the delay between the stimulus and this late response varied. If the width of the Gaussian filter varied based on the firing rate, both late responses would be smoothed differently, leading to a decrease in similarity. Likewise, It is difficult to define the exact edges of a burst, but they do have an effect on the average firing rate of the neurons during the burst. The algorithm to extract the bursts would then affect the similarity between the bursts. We therefore used a fixed width of the Gaussian filter, defined by its standard deviation $\sigma = 3\,\text{ms}$.

Algorithm

The adapted algorithm for two very small SBEs is visualized in figure 2.6. The top row shows the spike raster plots of the individual SBEs. The average activities of the electrodes are 1.5 and 2.5, resulting in relative activities of $\frac{1.5}{4} = 0.375$ and $\frac{2.5}{4} = 0.625$. Next are the four activity vectors, two per SBE. The following two graphs show the normalized cross-correlations of the paired activity vectors. The cross-correlations are scaled by the relative spiking frequency of the respective neuron, which is not done in the original paper. The two normalized and scaled cross-correlations are summed and shown in the last figure. The similarity is then defined as the maximum of this function, in this case ≈ 0.833. This maximum occurs at a time lag of $-5\,\text{ms}$ which means that the second SBE has a time-lag $5\,\text{ms}$ compared to the first SBE. This time-lag is not used in our analysis.

A MATLAB implementation of the algorithm that creates a similarity matrix from a list of SBEs is shown in appendix B. The entries of the similarity matrix are defined as the mutual similarities between every SBE combination. Therefore, the matrix is symmetric and its diagonal consists of ones. Computationally, the presented algorithm can be improved by removing any trailing zeros of the activity vectors, which significantly

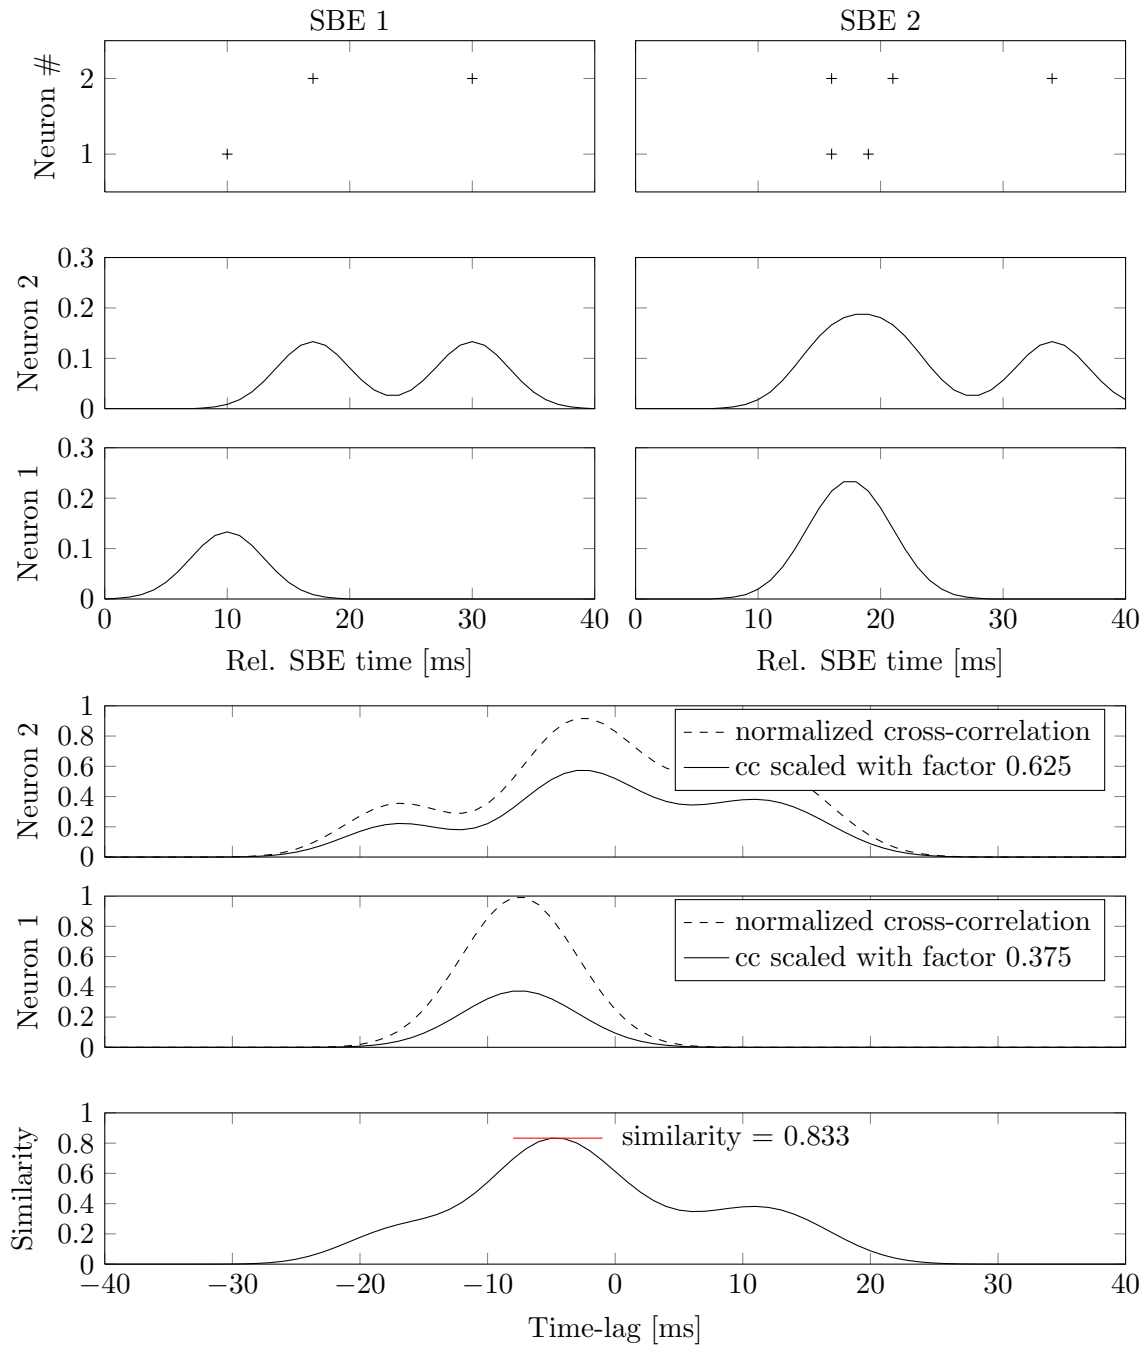

Figure 2.6.: Calculation of the similarity between two synchronous bursting events (SBEs). The top row shows the spike raster plots of both SBEs. Per SBE, an activity vector is calculated for every neuron by smoothing a histogram of the timings with a Gaussian filter, as shown in rows 2 and 3. Per neuron, the mutual normalized cross-correlations are calculated and scaled according to the relative activity of the neuron as plotted in the two rows below. Neuron 1 spikes 1.5 times on average while neuron 2 spikes 2.5 times. Their relative activities are therefore $\frac{1.5}{4} = 0.375$ and $\frac{2.5}{4} = 0.625$. The cross-correlations are summed and the similarity between SBE 1 and SBE 2 is defined as the maximum of this summed cross-correlation as depicted in the last row.

increases the speed of calculating the cross-correlations. The cross-correlations then vary in size and need to be centered before summing them. In addition, the normalization of the cross-correlations has to be done manually.

2.3.6. Euclidean distance between weight matrices

The Euclidean distance is a metric to compare two vectors with the same dimension. The Euclidean distance between two vectors is defined as

$$\mathrm{ED}(\vec{p}, \vec{q}) = \sqrt{\sum_i (p_i - q_i)^2}$$

The idea can be extended to compare equally sized matrices

$$\mathrm{ED}(\boldsymbol{A}, \boldsymbol{B}) = \sqrt{\sum_i \sum_j (\boldsymbol{A}_{i,j} - \boldsymbol{B}_{i,j})^2}$$

The significance of the measure is limited but we use it to get an overview of the temporal development of weight matrices by measuring their distance as a function of time. In accordance with [39], $\mathrm{ED}_0(t')$ denotes the Euclidean distance between the synaptic weight matrices at $t = 0$ and $t = t'$.

3. Results

3.1. Simulation results

This section describes the results gained by the simulations described in section 2.1.6. The first section deals with the development of unstimulated networks, the second part describes the effect of stimulation.

3.1.1. Development of an unstimulated network

Standard settings

Figure 3.1 shows the spike raster plot of 10 minutes of activity for one of the network realizations with standard settings and a close up of the first out of four bursts. To get a first impression of the development of the synaptic weights, we use the Euclidean distance measure ED_0. This measure is plotted for all three realizations in figure 3.2. It can be seen that the Euclidean distance changes very quickly at the beginning. The reason for this is that the average synaptic weight quickly decreased from the initial value of 0.5 to approximately 0.35 as will be illustrated in greater detail shortly. After approximately two

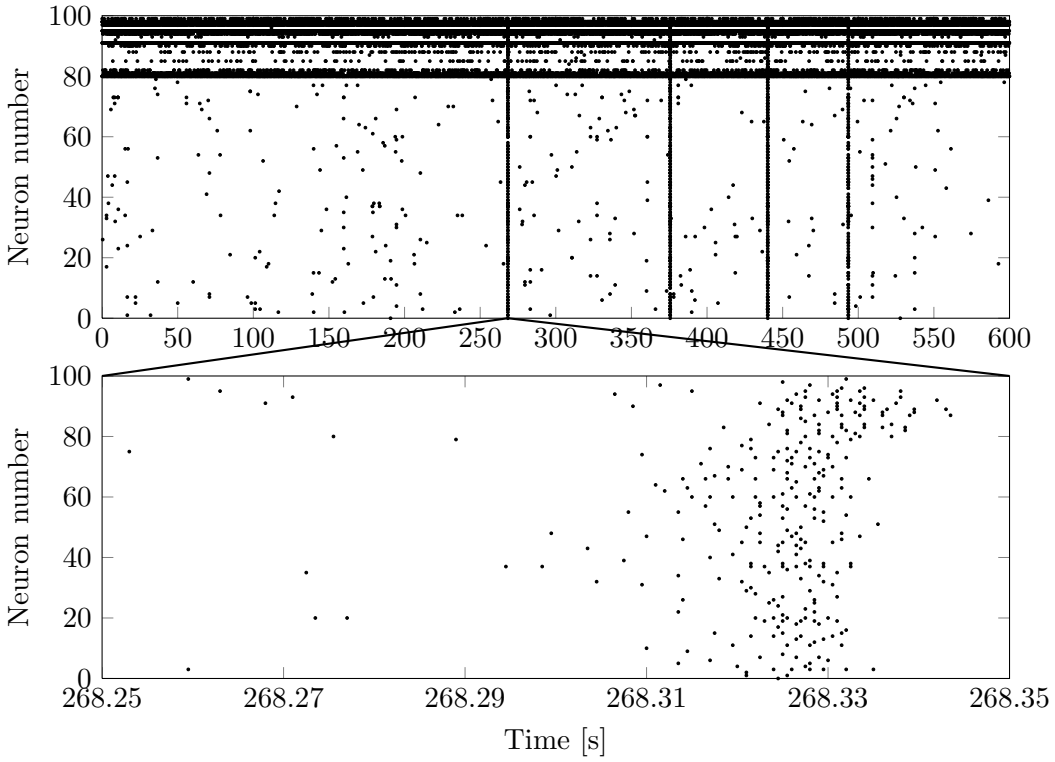

Figure 3.1.: Spike raster plot example for 10 minutes of activity for one of the network realizations with standard settings after one day of activity. Four bursts can be clearly identified.

days, the Euclidean distance appears to settle. The three networks then start to behave differently. While for the first network the Euclidean distance remains stable over the full ten days, it further increases for the third network after four days of spontaneous activity. For the second network realization, an increase occurs after approximately nine days.

We continue with the development of the bursting rate, i.e. the amount of network bursts that occur per time. The bursting rates for the same three network realizations are plotted in figure 3.3. After a relatively high bursting rate of approximately 1.35 bursts per minute at the beginning, the frequency drops to approximately 0.2 bursts per minute within half a day for all three networks. The third network, which showed a sudden increase in ED_0 after five days, also shows an increase in the bursting rate from that moment on and reaches a rate of 0.4 bursts per minute after ten days. The bursting rate of the second realization increases more gradually and becomes approximately 0.4 bursts per minute also. The first network, which showed the least variation in the Euclidean distance of the weights, also has the most consistent bursting rate of the three networks. It stays at approximately 0.25 bursts per minute during the full simulation period.

We can examine the network bursts more closely by using the concept of burst leaders, which are neurons that fire first in a burst (section 2.3.4). For every network realization, the five neurons that initiate the most bursts in the total simulation period are determined. The leader frequency is then calculated for each of those leaders by using only the bursts initiated by that neuron. Figure 3.4 shows the temporal development of those leader

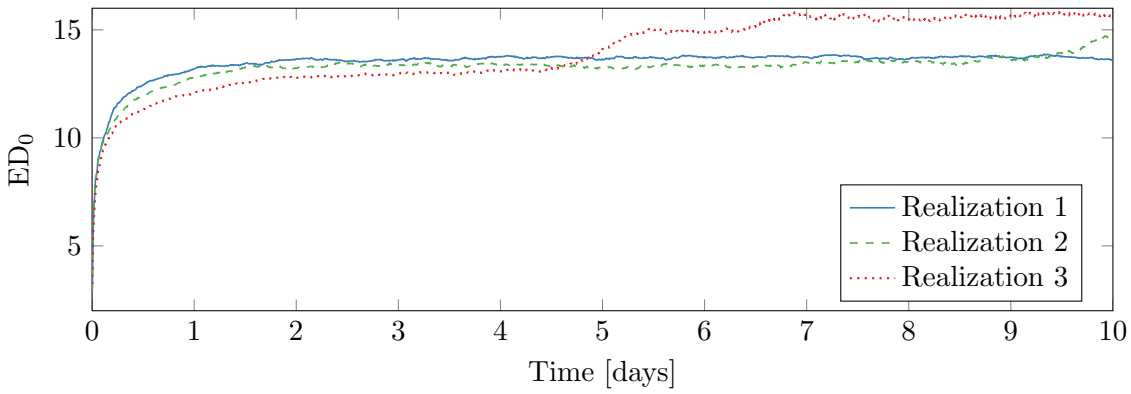

Figure 3.2.: Development of the Euclidean distance ED_0 over ten days for the three different network realizations with standard settings.

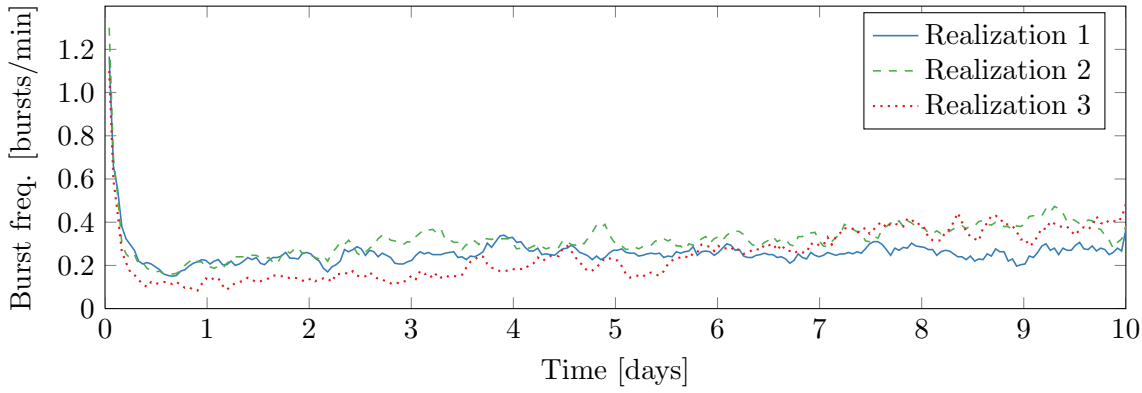

Figure 3.3.: Development of the burst frequency over ten days for the three network realizations with standard settings. The plots were created by making a histogram of the burst timings with a bin size of 1 h and then applying a moving average filter with a width of 5 h.

frequencies for all three realizations. The sum of the five most occurring leaders is plotted as a solid line, and the total bursting frequency as plotted in figure 3.3 is indicated with a dashed line. For the 1st network, there is no clear temporal development in the bursting rates, except for a limited increase for one of the leaders between days six and eight. The top five leaders account for approximately 50 % of all the bursts. The 2nd network realization shows a similar behavior. One of the neurons becomes a strong leader for approximately two days and one neuron starts to lead more bursts after nine days. The 3rd realization has the lowest bursting frequency of all networks in the beginning. Two neurons however start to initiate a significant amount of bursts after three and six days respectively. In the last three days, the bursts initiated by these two neurons account for almost all of the bursts in the network.

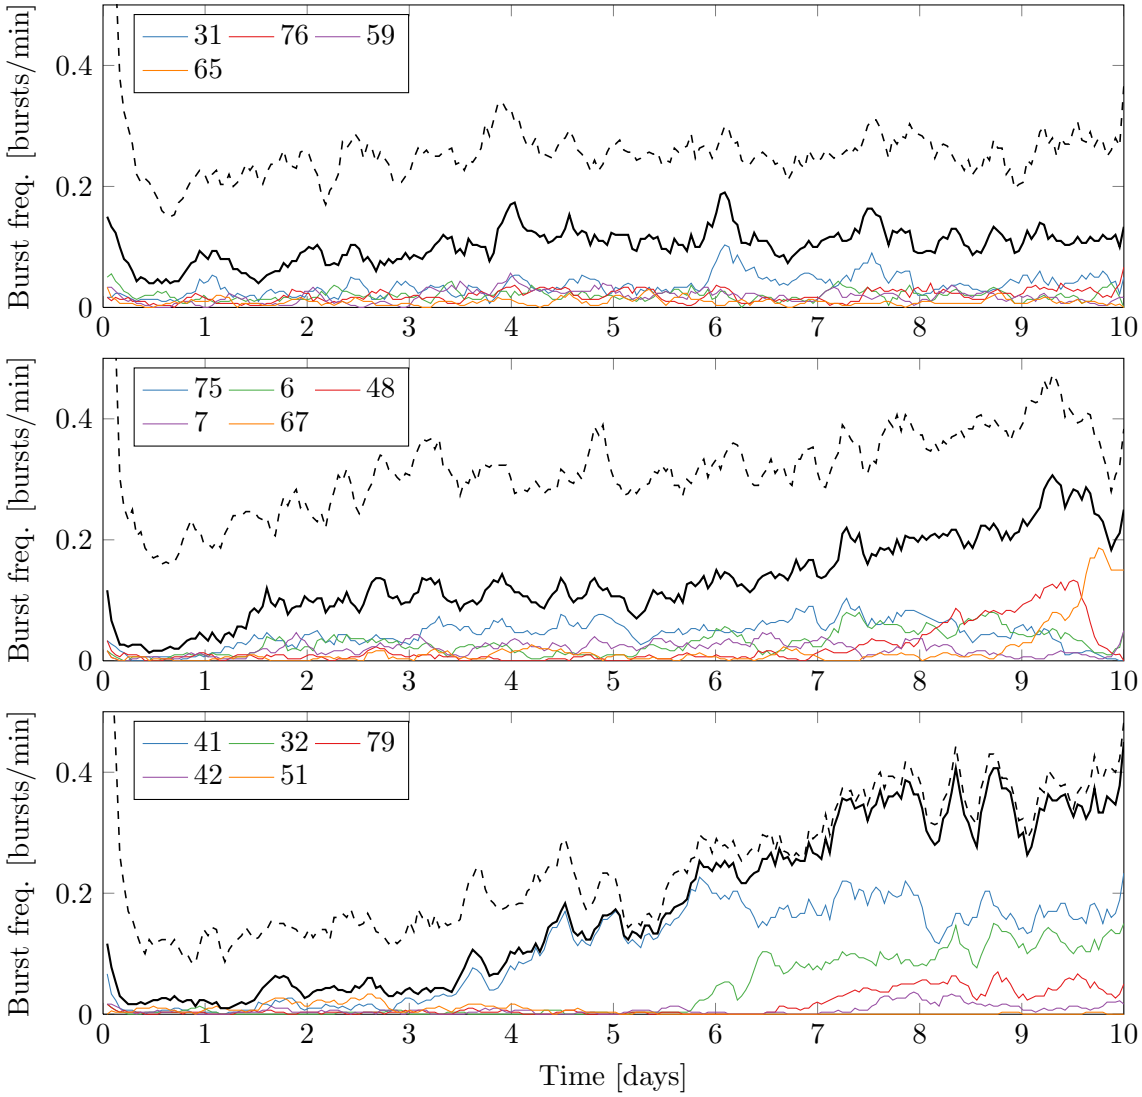

Figure 3.4.: The three plots show the development of the leader frequencies for the five most present leaders per network realization. The sum of the five leader frequencies is plotted with a thick line whereas the dashed line indicates the total bursting frequency in the network. The bursts were binned using a binsize of 1 h and then smoothed with a moving average filter of width 5 h. The legend shows the neuron number of the corresponding leader.

Table 3.1.: Statistical parameters for the weight distribution of the three network realization with standard settings after ten days of activity.

Network	Syn. types[1]	Mean	Median	Variance	Skewness
Realization 1	EE & EI	0.343	0.327	0.021	0.771
	EE	0.303	0.298	0.013	0.372
	EI	0.511	0.503	0.021	0.631
Realization 2	EE & EI	0.351	0.339	0.031	0.929
	EE	0.310	0.299	0.025	1.269
	EI	0.519	0.504	0.022	0.867
Realization 3	EE & EI	0.352	0.321	0.040	1.231
	EE	0.310	0.288	0.033	1.739
	EI	0.528	0.516	0.032	0.553

[1] EE = synapses from excitatory neurons to excitatory neurons, EI = synapses from excitatory neurons to inhibitory neurons.

Weight distribution in the networks with standard settings

The small size of the network allowed us to look at individual synapses simultaneously using a color plot. Figure 3.5a depicts the relative synaptic weights for the three networks with standard settings after ten days of spontaneous activity. Only the excitatory synapses are shown since the weights of the inhibitory synapses are constant at -0.5. It can be seen that the synapses that connect to inhibitory postsynaptic neurons (EI synapses) are on average stronger than the synapses that connect to excitatory postsynaptic neurons (EE synapses). In network realizations 2 and 3, there are some neurons that have very strong outgoing and weak ingoing synapses. In network 2 this is neuron 67. For the 3rd network, neurons 32 and 41 can be clearly identified as such while neurons 55 and 79 show a similar but less distinct trend. We saw before (figure 3.4) that after ten days, there is one neuron in the 2nd network that leads almost half of all bursts (orange line). This graph corresponds to neuron 67. The most present leaders in the third network are 41 (blue line), 32 (green line) and 79 (red line). Based on this connection and to simplify further notation, we define "leader neurons" as neurons of which at least 80 % of the outgoing synapses have a strength ≥ 0.9. By this definition, neurons 41 and 32 of the 3rd network are leader neurons, while 55 and 79 are not.

Histograms of the three weight distributions are shown in figure 3.5b. The total histogram (black) has been subdivided into EE synapses (gray) and EI synapses (white) and their statistical parameters are shown in table 3.1. The visual impression that EI synapses are on average stronger than EE synapses is confirmed. The extreme synaptic weights present in networks 2 and 3 show up as peaks in the histograms. Realization 2 has a peak only at the maximum value, while in the 3rd realization, there is also a peak for synaptic weights with a value close to zero. In network realization 1, there are no peaks at the boundaries and its synaptic weights form a unimodal distribution.

No LTS neurons

The characteristics for all the unstimulated networks that have been simulated are shown in table 3.2. If we compare the three networks without LTS neurons to the three networks with standard settings, we see that the first two realizations show a strong increase in the

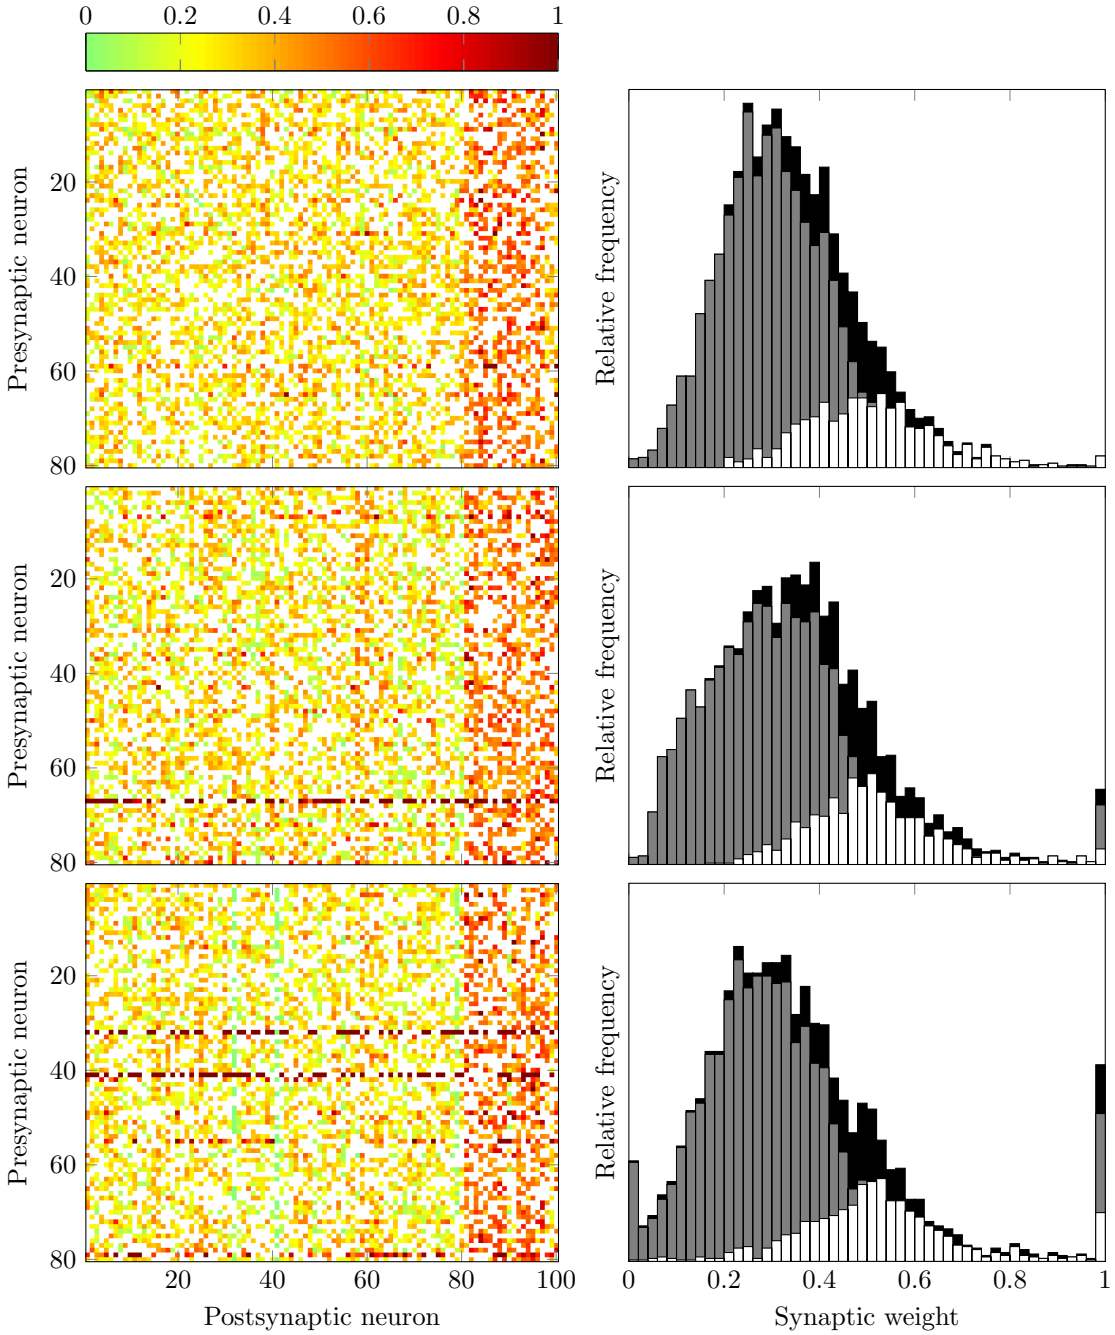

(a) White cells indicate that there is no connection between the neurons.

(b) Gray = EE synapses only, white = EI synapses only, black = both.

Figure 3.5.: Relative weights of the excitatory synapses for all three network realizations after ten days of spontaneous activity. The histogram differentiates between synapses from an excitatory neuron to an excitatory neuron (EE) and synapses from an excitatory neuron to an inhibitory neuron (EI).

number of bursts, which on average contain more spikes fired by excitatory neurons. Less bursts occur in the 3rd network but their average length increases. No leader neurons exist any more after ten days of activity and there is less variance in the synaptic weights while their mean does not change significantly.

Reduced noise on inhibitory neurons

As can be seen from table 3.2, the effect of reducing the noise applied to the inhibitory neurons is similar to removing the LTS neurons, i.e. the first two networks show an increased amount of bursts with on average more spikes, while the number of bursts in the 3rd network decreases and their average length increases. Again, no leader neurons form in 10 days of spontaneous activity. Figure 3.6 shows 20 randomly selected synapses in the first of these realizations. It appears that the synapses stabilize within one day and then only fluctuate around their final value without changing significantly. This weight distribution has been analyzed further by plotting ED_0 but resetting $t = 0$ every two days. This plot is shown for all three network realizations with less inhibitory noise in figure 3.7. Apparently, the random fluctuations around the final value of a synapse are such that the euclidean distance between two weight matrices approaches approximately 1.

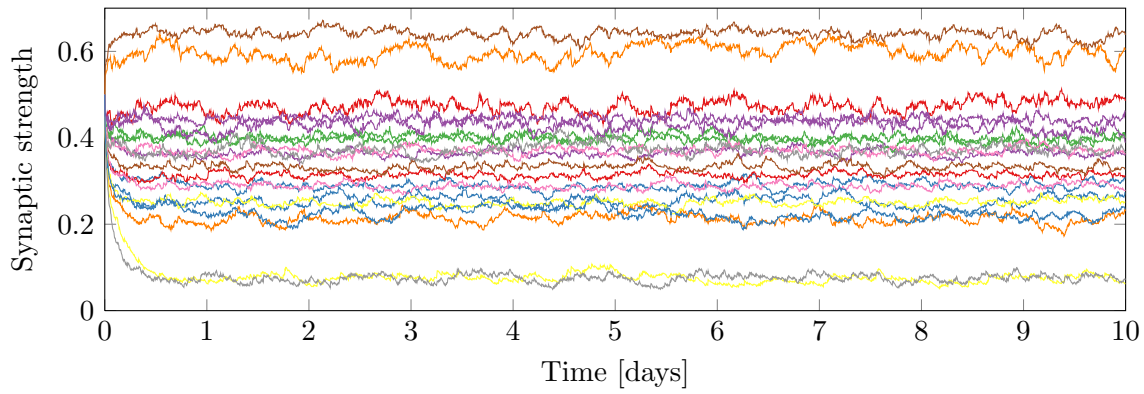

Figure 3.6.: Development of 20 random synapses in the first network realization of the networks with less noise on inhibitory neurons.

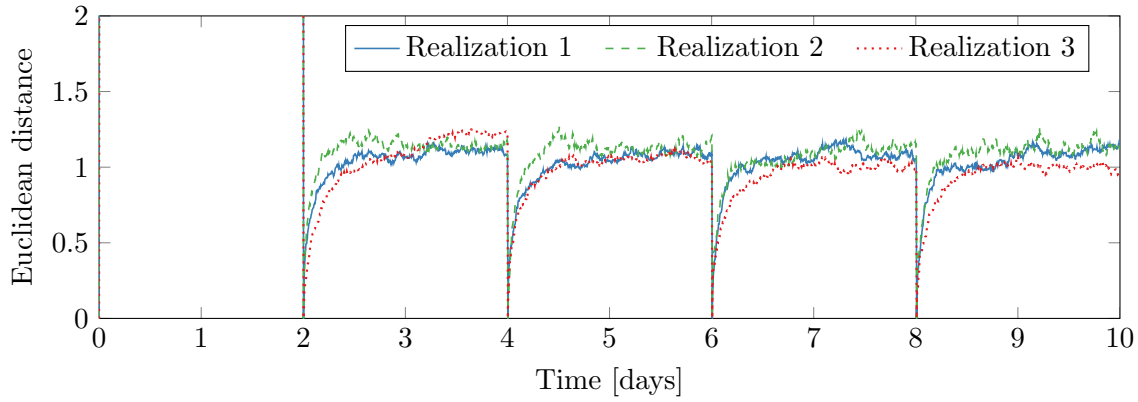

Figure 3.7.: Euclidean distance between the current weight matrix and the weight matrix at the last even day for the three network realizations with less noise on the inhibitory neurons.

Table 3.2.: Network characteristics for all 18 network realizations simulated for 10 days.

Network (sim. #)	Bursts	Avg. burst length	Spikes per burst[1]	Avg. weight EE synapses[2,3]	Leader neurons[3]	EE/EI corr. coef.[2,4]
Standard settings						
Realization 1 (1)	3738	65 ms	219	0.303 ±0.115	0	-0.75/-0.73
Realization 2 (2)	4620	63 ms	221	0.310 ±0.157	1	-0.57/-0.66
Realization 3 (3)	3585	47 ms	235	0.310 ±0.180	2	-0.49/-0.57
No LTS						
Realization 1 (4)	5176	64 ms	328	0.304 ±0.101	0	-0.82/-0.75
Realization 2 (5)	7929	64 ms	361	0.311 ±0.101	0	-0.81/-0.82
Realization 3 (6)	3275	66 ms	336	0.309 ±0.095	0	-0.83/-0.75
Less inh. noise						
Realization 1 (7)	4325	65 ms	298	0.304 ±0.103	0	-0.80/-0.69
Realization 2 (8)	5478	67 ms	318	0.313 ±0.099	0	-0.80/-0.81
Realization 3 (9)	2859	68 ms	310	0.312 ±0.092	0	-0.83/-0.80
Diff. noise real.						
Realization 1 (10)	4235	65 ms	298	0.303 ±0.103	0	-0.80/-0.63
Realization 2 (11)	5226	67 ms	317	0.311 ±0.100	0	-0.79/-0.82
Realization 3 (12)	2712	68 ms	310	0.313 ±0.093	0	-0.83/-0.82
Rnd. init. weights						
Realization 1 (13)	4275	65 ms	298	0.304 ±0.103	0	-0.80/-0.69
Realization 2 (14)	5483	67 ms	316	0.311 ±0.099	0	-0.80/-0.81
Realization 3 (15)	2904	68 ms	310	0.313 ±0.093	0	-0.83/-0.80
Larger axon. del.						
Realization 1 (16)	4795	74 ms	223	0.256 ±0.192	1	-0.61/-0.67
Realization 2 (17)	5618	85 ms	238	0.273 ±0.129	0	-0.80/-0.80
Realization 3 (18)	2960	90 ms	224	0.276 ±0.122	0	-0.82/-0.85

[1] Only the spikes of excitatory neurons are counted.

[2] EE synapses are synapses where both the pre- and postsynaptic neuron are excitatory. For EI synapses, the presynaptic neuron is excitatory and the postsynaptic neuron inhibitory.

[3] After 10 days of activity.

[4] This value is the correlation coefficient between the axonal delays in a network and the corresponding synaptic strengths after 10 days of activity for EE and EI synapses.

Different noise realization

The three previous networks were simulated again but with different noise realizations. As shown in table 3.2, the activity characteristics per realization barely change. We will now look at the effect of different noise on the weight development. We calculated the Euclidean distances between the weight matrices of these networks and the previous networks per realization as a function of time. Figure 3.8 shows these differences for all three network realizations. In the 2nd realization, the weight difference is approximately one during the whole simulation. Before, we showed that even after the synapses have reached a stable distribution, they still show random fluctuations, such that the euclidean distance between two weight matrices more than half a day apart is approximately 1. A Euclidean distance of approximately 1 therefore indicates that the means of the synaptic weights are the same. In the 3rd realization, the Euclidean distance between the weights is 1 in the beginning and end. In between, it increases to 1.5 for about five days. The 1st realization shows

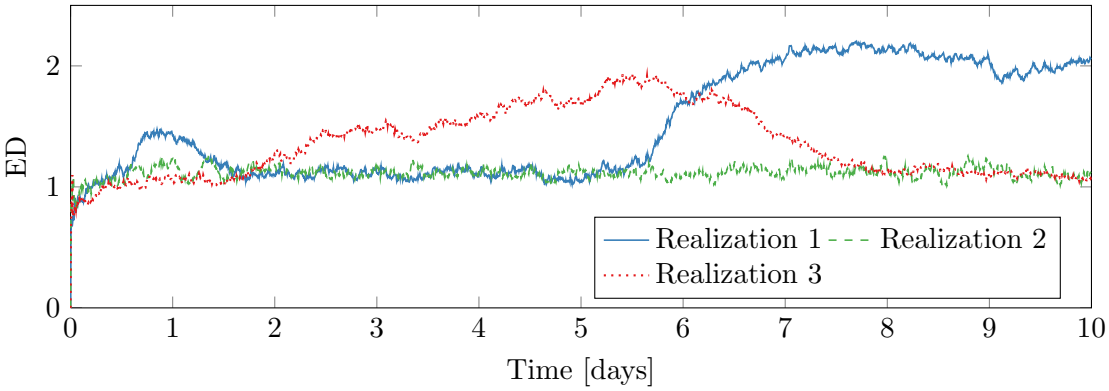

Figure 3.8.: Development of the Euclidean distances between the networks 7-9 and 10-12 from table 3.2, paired per realization. The networks are equal except for different noise applied to the neuron potentials.

an increase of the Euclidean distance to 2 after six days, which lasts until the end of the simulation. We visualize this change by calculating the average weight difference for all synapses of this realization between day 6 and 10 as shown in figure 3.9. The difference between most of the weights has an average close to zero. Only for the ingoing synapses of one of the inhibitory neurons there is a significant difference. In one of the networks, the synaptic weights are close to zero, while in the other network they are around 0.4.

Different initial weights

Three networks were simulated with the initial synaptic weights for the excitatory synapses chosen randomly from the uniform distribution $U(0, 1)$. The characteristics of the networks with these settings do not show a significant difference from the characteristics of the networks with the previous two simulation settings. We will again use the Euclidean distance measure to see if there is a change in weight development as plotted in figure 3.10. In the beginning, the difference in synaptic weights is large due to the different initial weights. Within half a day, the Euclidean distance between the weight matrices has

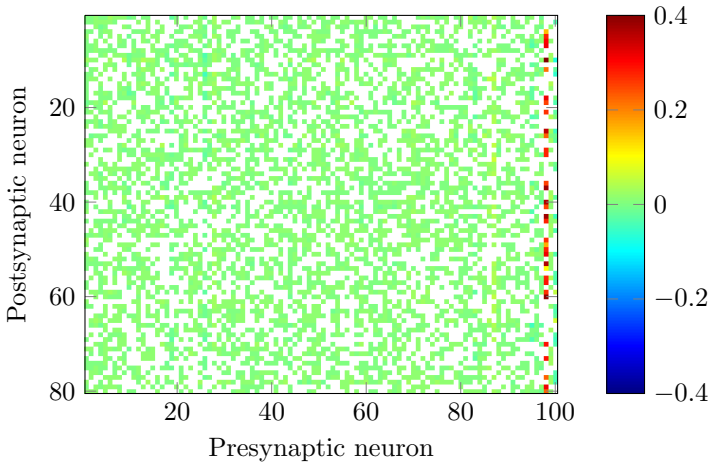

Figure 3.9.: Average weight differences between networks 7 and 10 from table 3.2 after 6 h of simulation. The networks are equal except for different noise applied to the neuron potentials.

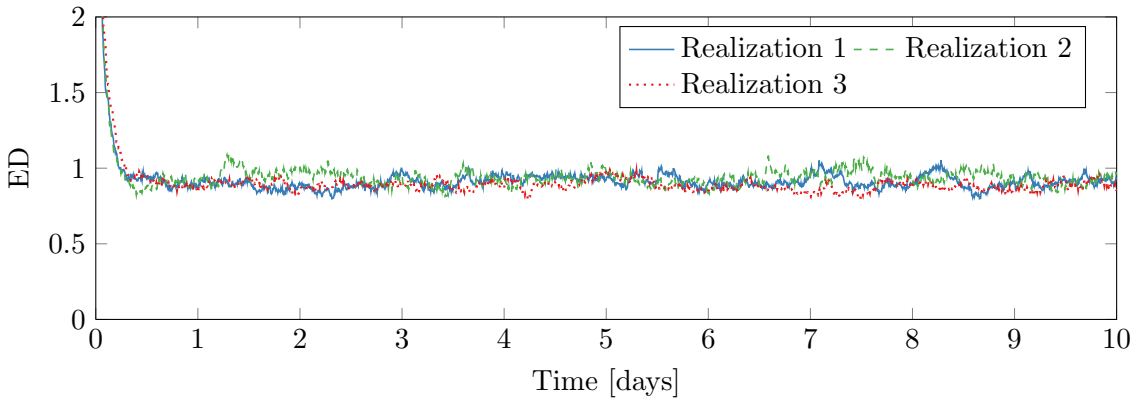

Figure 3.10.: Development of the Euclidean distances between the networks 7-9 and 13-15 from table 3.2, paired per realization. In one of the networks, the initial weights of the excitatory synapses are set to 0.5, in the other network they are initialized randomly from the uniform unit interval. In addition, different noise is applied to the neuron potentials. At $t = 0$, ED is approximately 20 for all plots.

decreased to approximately 1 for all three realizations, indicating that the same weight distribution has been reached.

Increased axonal delays

Three network realizations were simulated for 10 days with the axonal delays uniformly distributed on the interval $[1, 20]$ instead of $[1, 10]$. Compared to the previous three simulation settings, the number of bursts increases. Bursts are wider on average while the amount of excitatory spikes per burst decreases. The weights of the excitatory synapses are lower than in the previous simulations and network realization 1 develops a leader neuron.

Correlation between synaptic weights and axonal delays

The last column in table 3.2 lists the cross correlation of the weights in the network after 10 days of simulation with the corresponding axonal delays, subdivided for EI and EE synapses. Only those neuron pairs that have a connection are used for the calculation. This correlation is strongly negative, indicating that two neurons with a long axonal path between them on average have a weak synapse and vive versa. This is visualized in the scatter plot 3.11, which shows all the axonal delay / synaptic strength pairs in one of the network realizations.

3.1.2. The effect of stimulation

The following three sections describe the effect of stimulation on the synaptic weights, the spontaneous activity in the network, the burst leaders and the development of the stimulus response for the 12 simulations that involve stimulation as described in section 2.1.6.

Synaptic weights

The change in synaptic weights is first summarized by means of the Euclidean distance ED_0 where $t = 0$ denotes the start of the simulation protocol. Two exemplary plots

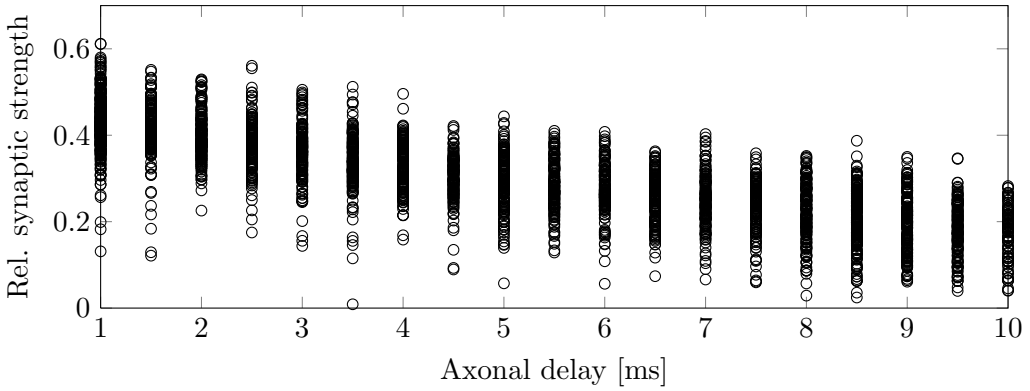

Figure 3.11.: Scatter plot depicting the strength / axonal delay pair for every synapse with an excitatory pre- and postsynaptic neuron in one of the networks with less inhibitory noise after 10 days of spontaneous activity.

are shown in figure 3.12. Figure 3.13a shows the mean ED_0 of all 12 simulations. The standard deviation for the first 2 h is very small since only three unique networks were used in the 12 simulations and their activity before the stimulation is fully deterministic. Since the maximum amount of change in the Euclidean distance varies strongly among the simulations, we also calculated a normalized Euclidean distance, which mean is plotted in figure 3.13b.

In the first two hours without stimulation, the Euclidean distance increases slightly to about 1 due to the individual fluctuations of the synaptic weights. When the stimulation starts, ED_0 quickly increases. The rate of change decreases during stimulation and it appears that the Euclidean distance reaches its maximum after approximately 2 h of stimulation. After the stimulation has stopped, ED_0 decreases again, but after 3 h of spontaneous activity, its value is on average still approximately 70 % of its maximum. The relative rate at which ED_0 decreases after the stimulation has stopped shows a large variation among the simulations.

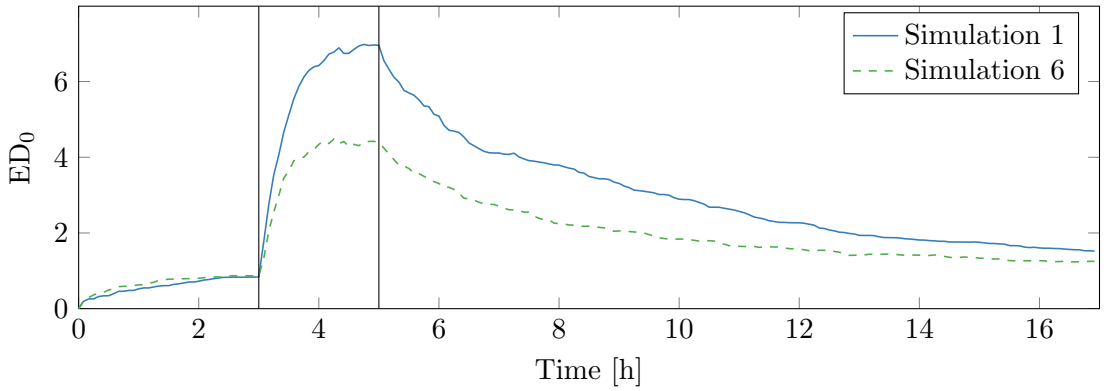

Figure 3.12.: ED_0 plots for two of the 12 simulations. The stimulation of three neurons of the network started at $t = 3$ h and lasted for 2 h as indicated by the vertical lines. The stimulation frequency was 0.125 Hz.

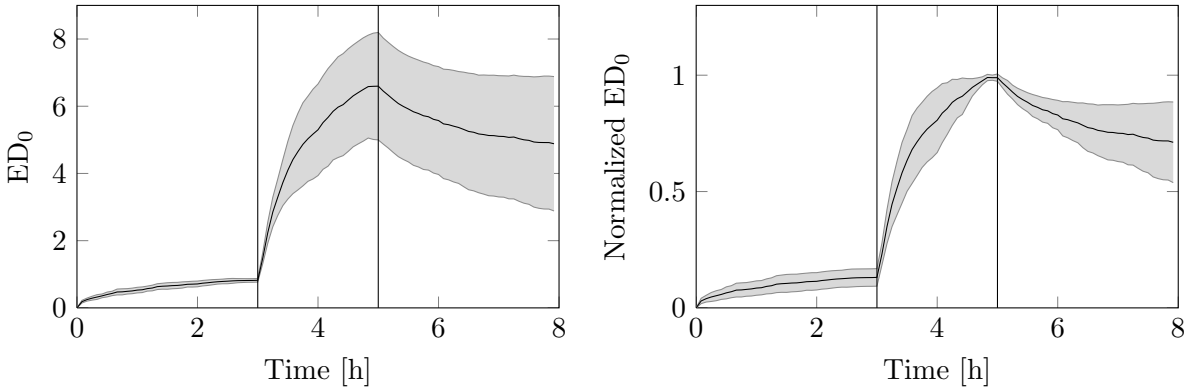

(a) Mean of the Euclidean distance for all 12 simulations.

(b) Mean of the normalized Euclidean distance. Normalization is done per simulation by dividing ED_0 by its maximum in that simulation.

Figure 3.13.: Effect of stimulation on the Euclidean distance and its normalization, averaged over 12 simulations. The shaded areas indicate the unbiased estimation of the standard deviation. The vertical lines depict the stimulation period.

Stimulus responses and network bursts

We denote by synchronous bursting events (SBEs) both the network bursts during spontaneous activity and the network responses due to stimuli. For every simulation, the SBEs can be categorized by three types:

- Bursts that occurred before the stimulation phase (prestimulus bursts)
- Stimulus responses
- Bursts that occurred after the stimulation phase (poststimulus bursts)

The similarity between two SBEs is measured utilizing the algorithm described in section 2.3.5. Figure 3.14 shows a matrix with the similarities between all the SBEs for one of the simulations. Matrices are calculated for every simulation and then used for further analysis.

The first measure is the average similarity between a stimulus response and all the other stimulus responses of that simulation. Figure 3.15 shows this average similarity for all stimulus responses of simulation 7. The horizontal line depicts the median of the last 300 responses. The average similarity starts at approximately 0.94 and asymptotically approaches its maximum value of 0.975, which is reached after about 250 responses. The mean of this measure for all twelve simulations is shown in figure 3.16. The shape is very similar and for approximately the first 30 stimulus responses, the average similarity is more than one standard deviation below the asymptotic value.

We continue with the average similarity between a stimulus response and all the pre- or poststimulus bursts that occur during one simulation, as exemplary shown for simulation 1 in figure 3.17. The average similarities of the stimulus responses with both the prestimulus and poststimulus bursts increase over time. But while the similarities are comparable at first, the similarity with the poststimulus bursts shows a stronger increase and after circa 200 stimulus responses, the average similarity with the poststimulus bursts is clearly higher. There appears to be a high correlation between pre- and poststimulus bursts, i.e. if a stimulus response has a low similarity with the prestimulus bursts, the similarity

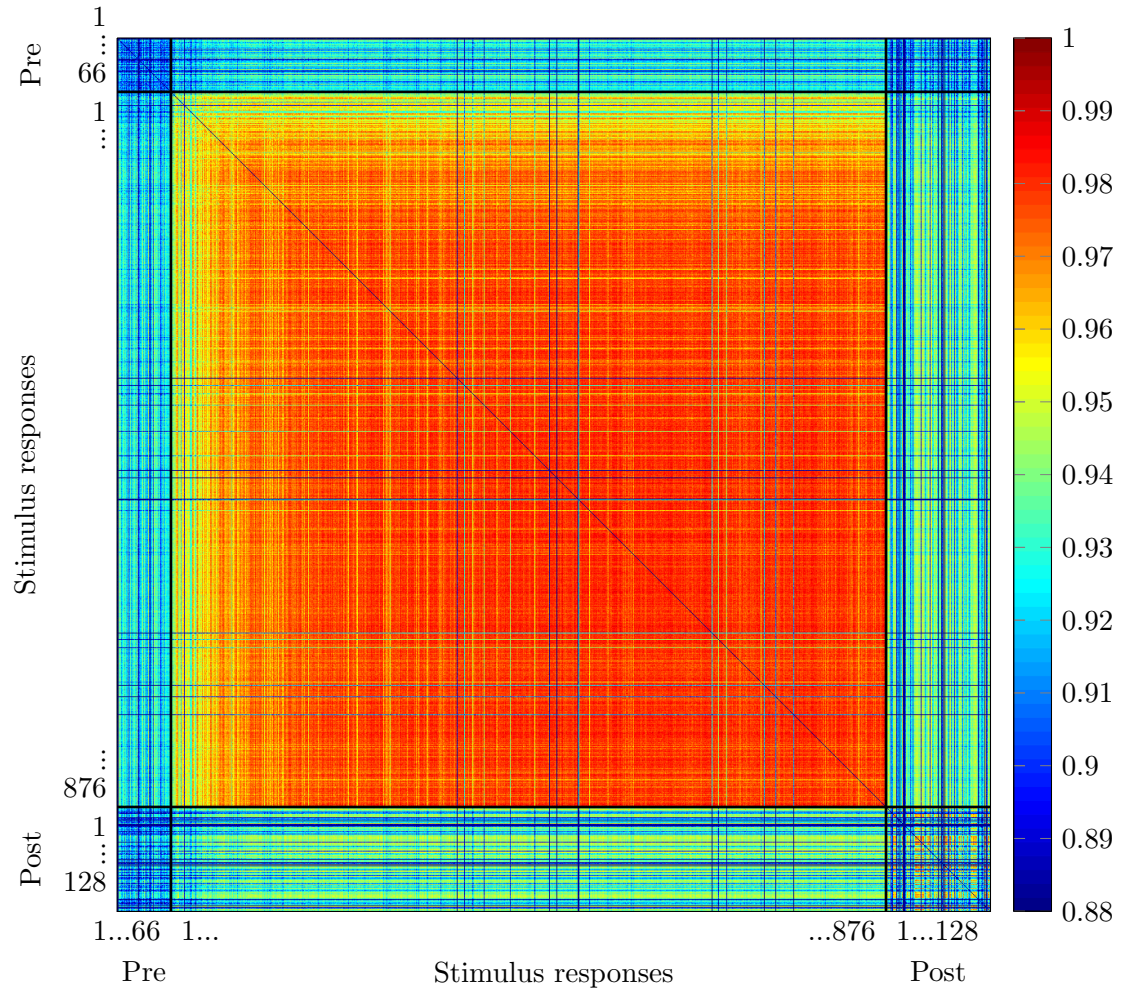

Figure 3.14.: Matrix with the mutual similarities between all the synchronous bursting events (SBEs) in simulation 7. The SBEs are sorted by time, i.e. the first SBE is located at the top/left, and categorized according to prestimulus bursts (Pre), stimulus responses and poststimulus bursts (Post)

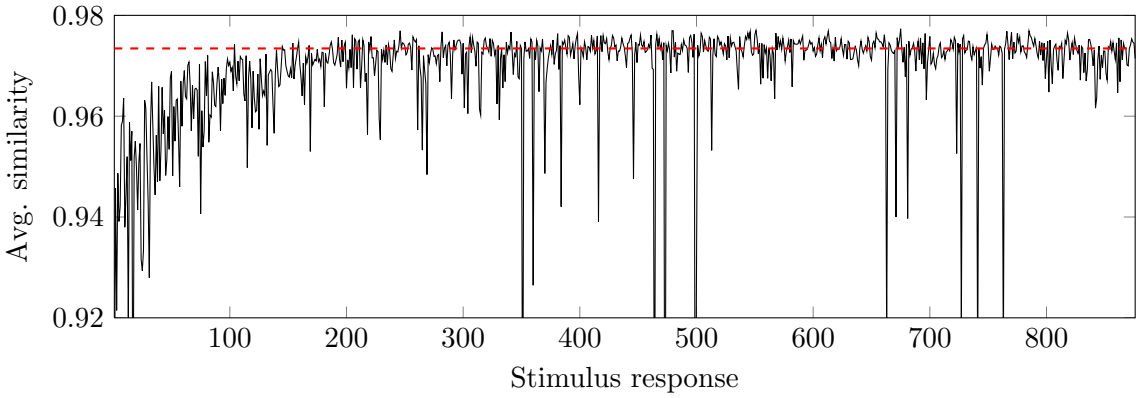

Figure 3.15.: The average similarity between a stimulus response and all other stimulus responses, for all the stimulus responses of simulation 7. The dashed red line indicates the median of the last 300 values.

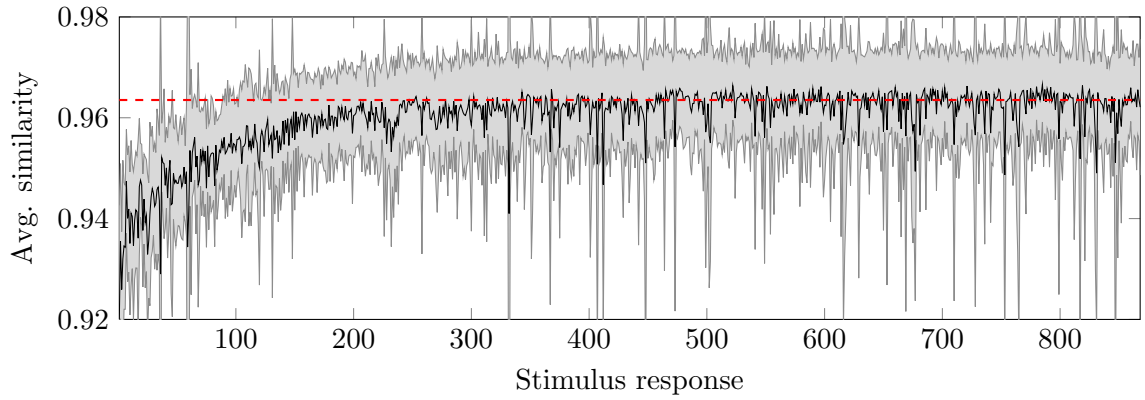

Figure 3.16.: This plot shows the mean similarity between a stimulus response and all other stimulus responses, averaged over all simulations. The shaded areas indicate the unbiased estimation of the standard deviation. The dashed red line indicates the median of the last 300 values.

with the poststimulus bursts is low as well and vice versa. Figure 3.18a shows the same quantity, averaged over all simulations.

The difference between the mean similarity of a stimulus response with all the bursts before and all bursts after stimulation has been averaged over all simulations and is depicted in figure 3.18b. This difference is positive for all stimulus responses but significantly increases during the first 700 stimulus responses. The smoothness of this difference, compared with the individual similarities shown in figure 3.18a (note the different scales), affirm the correlation described in the previous paragraph: apparently, the difference in similarity between a stimulus response and the pre- / poststimulus bursts is rather independent of the absolute similarity.

Burst leaders

The burst leader distributions before and after stimulation are compared for every simulation, using the similarity measure described in section 2.3.4, which yields a similarity

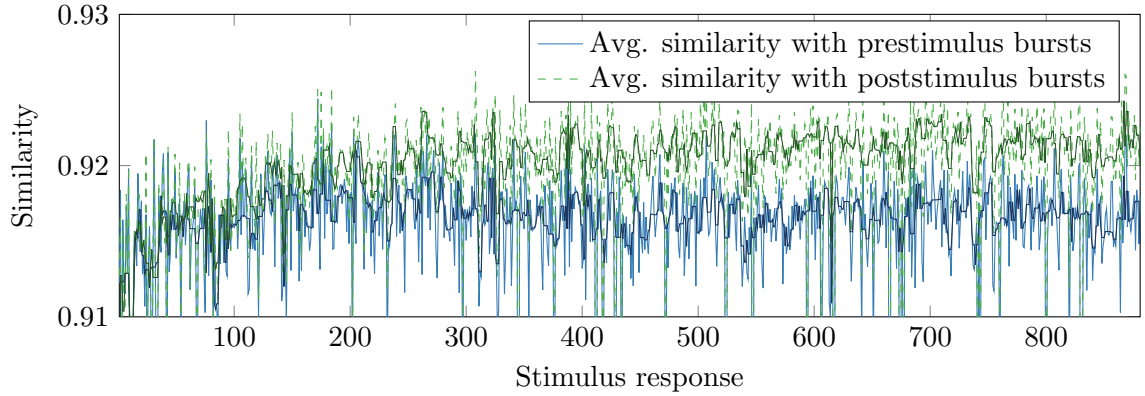

Figure 3.17.: The average similarity between a stimulus response and all pre-/poststimulus bursts for simulation 1. The solid lines depict the smoothed values using a median filter with a window size of 5.

between 0 and 1. Since three different initial networks are used, we also create three control networks. In the control networks, no stimulation is applied during the stimulation phase and we therefore estimate the change in leader distribution similarity due to intrinsic changes in the network and chance. The changes in the burst leader distributions are denoted in table 3.3.

Hypothesis test and summary

We now test the hypothesis that the stimulus responses in a simulation have a higher similarity with the poststimulus bursts than with the prestimulus bursts. We do this individually per simulation and for all the simulations combined. On the individual level, we average the similarity of every burst in a simulation with all the stimulus responses and denote by $S_{\text{pre},m,i}$ and $S_{\text{post},n,i}$ this average similarity for the mth pre- / nth poststimulus

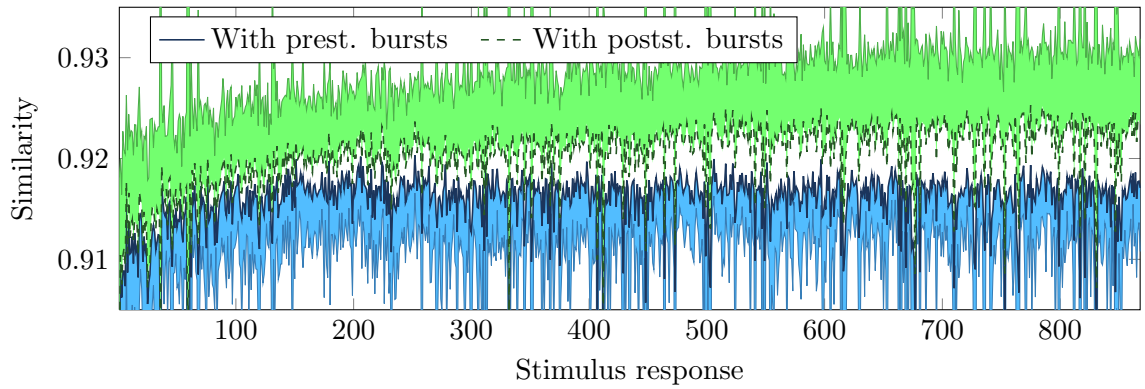

(a)

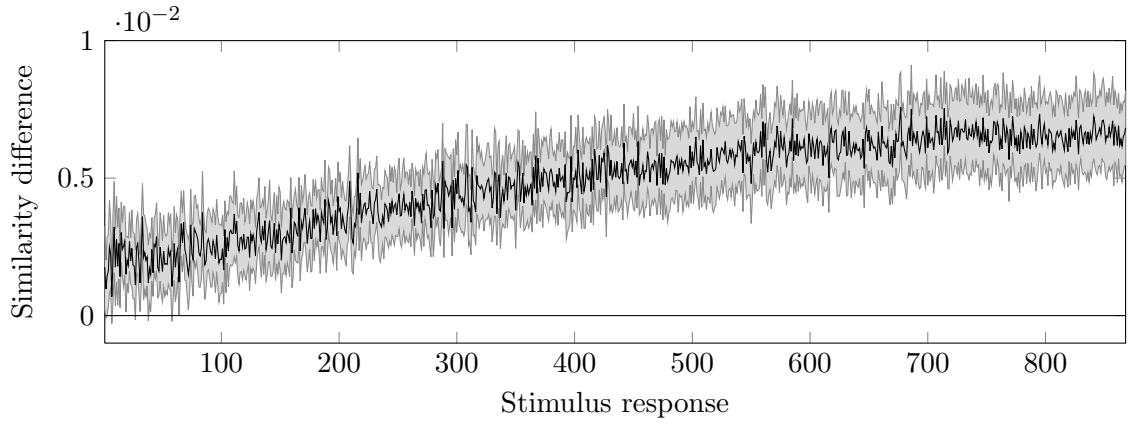

(b)

Figure 3.18.: (a) shows the mean similarities between a stimulus response and all pre- / poststimulus network bursts of the corresponding simulation, averaged over all simulations. The shaded areas indicate the unbiased estimation of the standard deviation. Only half of the shaded area is shown to preserve clarity. (b) shows the difference between the mean similarity with poststimulus bursts and prestimulus bursts per stimulus response, averaged over all simulations. The shaded area indicates the standard error of the mean, calculated using the unbiased estimation of the standard deviation.

burst in the ith simulation. For every simulation i we test the hypothesis

$$\begin{aligned} \mathbf{H_0}: &\quad \overline{S_{\text{post},i}} - \overline{S_{\text{pre},i}} = 0 \\ \mathbf{H_1}: &\quad \overline{S_{\text{post},i}} - \overline{S_{\text{pre},i}} \neq 0 \end{aligned}$$

i.e. whether the average similarity of a *poststimulus* burst with all stimulus responses in a simulation is significantly higher/lower than the average similarity of a *prestimulus* burst with all stimulus responses. The underlying distributions of $S_{\text{pre},m,i}$ and $S_{\text{post},n,i}$ are unknown, but the number of pre- / poststimulus bursts is large (≥ 39). We therefore use the two-tailed, two-sample t-test with a significance level of 5 % to test the significance of the difference between the two averages. The results are listed in table 3.3. In 11 of the 12 simulations, the average similarity of the stimulus responses is higher with bursts after stimulation than before, but only in three simulations this difference is significant. In simulation 6, the average similarity of the stimulus responses with the bursts after stimulation is *smaller* than the average similarity with bursts before stimulation. It is interesting to note that this is the only simulation in which the burst leader distribution before the stimulation has a higher similarity with the burst leader distribution after stimulation than in the control network.

For the overall hypothesis check, we denote by $S^*_{\text{pre},i}$ and $S^*_{\text{post},i}$ the average similarity of all pre- / postimulus bursts with all stimulus responses for every simulation i, which leads to a total of 12 similarity pairs as listed in table 3.3. We then want to check whether similarities are on average higher for poststimulus bursts than for prestimulus bursts. Because of the small sample size and the unknown distributions for $S^*_{\text{pre},i}$ and $S^*_{\text{post},i}$, we have to use a nonparametric test and can therefore only test the medians of the distributions instead of the means

$$\begin{aligned} \mathbf{H_0}: &\quad \widetilde{S^*}_{\text{post}} - \widetilde{S^*}_{\text{pre}} \leq 0 \\ \mathbf{H_1}: &\quad \widetilde{S^*}_{\text{post}} - \widetilde{S^*}_{\text{pre}} > 0 \end{aligned}$$

Using the Wilcoxon signed rank test, we find a p-value of 7.3242×10^{-4}. This means that if the median of the similarity after stimulation were equal to or less than the median of the similarity before stimulation, the chance of getting our data would be less than one tenth of a percent. This makes the alternative hypothesis, that the similarity after stimulation has a higher median, very likely. For a discussion whether the premises for this test are fulfilled, particularly the independence of the value pairs, see section 4.1.2.

Table 3.3.: Summary of the effect of stimulation on three different computational network realizations with four sets of stimulated neurons each.

Sim.	Netw.[1]	Stim. neurons	Pre-/poststim. bursts[2]	Burst leader similarity/control	Pre./poststim. similarity[3]	Similarity difference[4]
1	7	1-3	47/140	0.296/0.613	0.9155/0.9191 ↑	○ (0.135)
2	7	4-6	47/153	0.372/0.613	0.9143/0.9206 ↑	+ (0.017)
3	7	7-9	47/187	0.134/0.613	0.9168/0.9191 ↑	○ (0.226)
4	7	10-12	47/167	0.164/0.613	0.9157/0.9177 ↑	○ (0.257)
5	8	1-3	66/77	0.300/0.525	0.9120/0.9151 ↑	○ (0.340)
6	8	4-6	66/97	0.566/0.525	0.9099/0.9081 ↓	○ (0.415)
7	8	7-9	66/128	0.224/0.525	0.9141/0.9226 ↑	○ (0.127)
8	8	10-12	66/170	0.238/0.525	0.9111/0.9220 ↑	○ (0.063)
9	9	1-3	39/113	0.185/0.292	0.9198/0.9217 ↑	○ (0.264)
10	9	4-6	39/232	0.188/0.292	0.9207/0.9274 ↑	+ (0.008)
11	9	7-9	39/156	0.073/0.292	0.9234/0.9321 ↑	+ (0.000)
12	9	10-12	39/97	0.055/0.292	0.9165/0.9183 ↑	○ (0.220)

[1] All networks after 10 days of spontaneous activity.
[2] Per 3 h.
[3] An uparrow indicates an increase in similarity, a downarrow a decrease in similarity.
[4] +/- = similarities poststimulus bursts with the stimulus responses are significantly larger/smaller than the similarities of the prestimulus bursts with the stimulus responses, ○ = no significant difference (α-value of 5 %). The value in brackets is the p-value.

3.2. In vitro results

We performed nine experiments on four different cultures as summarized in table 3.4. Three of the four cultures were in vitro between three and four weeks, the fourth culture was in vitro for approximately two weeks. In all 9 experiments, we applied the measurement protocol as described in section 2.2.3.

We first checked for every experiment whether the culture showed a stable response during the full stimulation phase. Appendix C shows the number of spikes and the length of every stimulus response for all nine experiments. There is a large deviation in the development of these parameters over time. In experiments 1 and 5, the stimuli do not yield a reliable stimulus response. In experiments 2, 4 and 8, the parameters are stable for approximately the first half of the responses and then show larger variations in the second half. For experiments 3 and 9, exactly the opposite can be observed. For the remaining experiments 6 and 7, the two plotted parameters appear to be stable over time. Every stimulus response was verified using the algorithm described in section 2.3.2. In experiments 1 and 5, less than 70 % of the stimulus responses were valid, **we therefore discarded the data obtained from experiments 1 and 5.**

Development of the stimulus responses and network bursts

We will present here the same three quantities that were described in the stimulation section of the simulation results. These are:

- The mean similarity of a stimulus response with all other stimulus responses in the experiment.

Table 3.4.: Summary of the cultures and stimulation parameters of the nine experiments.

Exp.	DIV[1]	Stim. electrode[2]	Stim. amplitude	Stim. responses[3]	Internal ID
1	21	78 (53)	16 µA	401	20379
2	22	66 (43)	12 µA	810	16659
3	25	51 (30)	16 µA	749	20379
4	28	28 (14)	12 µA	897	16659
5	27	34 (17)	16 µA	600	20379
6	26	31 (13)	16 µA	880	16659
7	12	45 (29)	16 µA	899	20380
8	26	66 (43)	12 µA	889	20382
9	13	48 (26)	16 µA	885	20380

[1] Days in vitro.
[2] The first value denotes the electrode as its column and row (e.g. 35 is the electrode in the third column and fifth row). The number in brackets is the internal electrode number which ranges between 0 and 59.
[3] The number of valid stimulus responses according to the algorithm in section 2.3.2.

- The mean similarity of a stimulus response with all pre- / poststimulus bursts.
- The difference between the mean similarity of a stimulus response with all poststimulus bursts and the similarity of the stimulus response with all prestimulus bursts.

Contrary to the simulations, the quantities do not follow a common trend among the experiments and therefore, we cannot present exemplary plots for either of them. Instead, we will show here only the values averaged over all experiments. The individual plots can be found in appendix D.

The mean similarity of every stimulus response with every other stimulus response in the experiment, averaged over all experiments (equivalent to figure 3.16) has been plotted in figure 3.19. The equivalents to figures 3.18a and 3.18b, i.e. the averages of the second and third quantity over all experiments, are shown in figure 3.20.

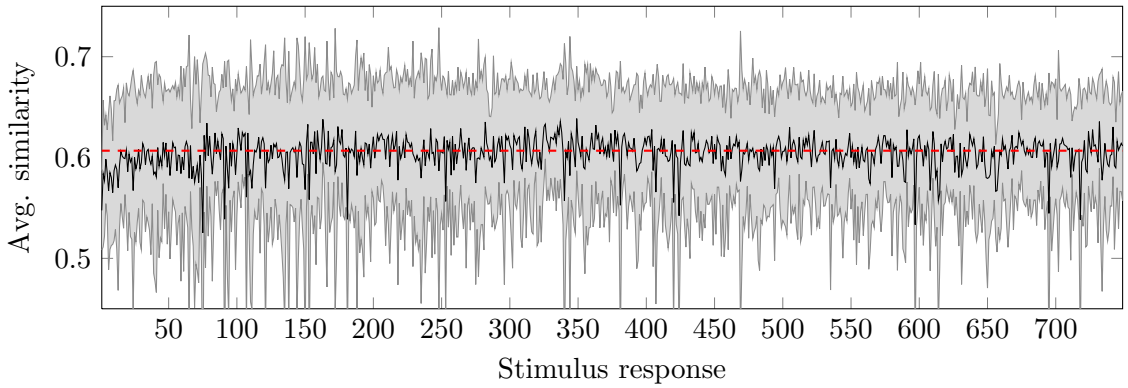

Figure 3.19.: This plot shows the mean similarity between a stimulus response and all other stimulus responses, averaged over all experiments. The shaded areas indicate the unbiased estimation of the standard deviation. The dashed red line indicates the median of the last 300 values.

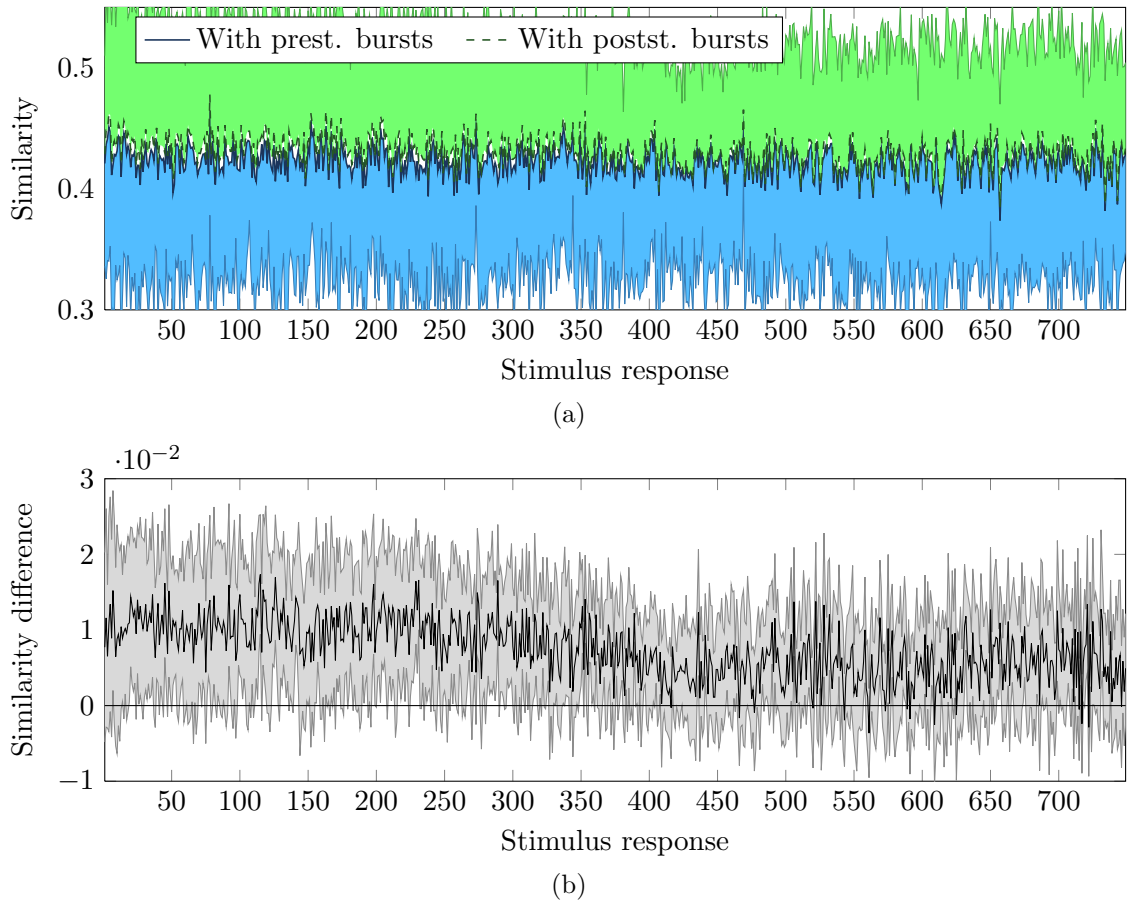

Figure 3.20.: (a) shows the mean similarities between a stimulus response and all pre- / poststimulus network bursts of the corresponding experiment, averaged over all experiments. Only half of the shaded area is shown to preserve clarity. (b) shows the difference between the mean similarity with poststimulus bursts and prestimulus bursts per stimulus response, averaged over all simulations. The shaded area indicates the standard error of the mean, calculated using the unbiased estimation of the standard deviation.

Stimulation induced changes in burst leader distributions

In some of the experiments, the measurement was continued for up to 17 additional hours after the protocol was finished. For these cultures, control values in the similarity between burst leader distributions were obtained by measuring the similarity between the burst leader distributions of the third to last and last hour of measurement data, such that two hours of spontaneous activity elapsed between the two data blocks.

Hypothesis test and summary

The stimulation effects per experiment are summarized in table 3.5. In four of the seven experiments, the average similarity of the poststimulus bursts with all the stimulus responses was higher than the average similarity of the prestimulus bursts with all the stimulus responses. In three experiments, this difference was significant at a 5 % level. In three of the seven experiments, the similarity with the poststimulus bursts was lower and in one experiment, this was significant.

Table 3.5.: Summary of the effect of stimulation for 7 different experiments.

Exp.	DIV[1]	Stim. electrode	Pre-/poststim. bursts[2]	Burst leader similarity/control	Pre./poststim. similarity[3]	Similarity difference[4]
2	22	66 (43)	112/94	0.944/∼	0.4836/0.4800↓	○ (0.309)
3	25	51 (30)	602/402	0.712/0.920	0.3690/0.3651↓	○ (0.074)
4	28	28 (14)	98/82	0.852/0.931	0.4700/0.5063↑	+ (0.000)
6	26	31 (13)	161/159	0.937/0.894	0.4927/0.4947↑	○ (0.380)
7	12	45 (29)	141/127	0.882/∼	0.4791/0.4569↓	- (0.000)
8	26	66 (43)	185/194	0.898/0.884	0.4264/0.4461↑	+ (0.007)
9	13	48 (26)	175/140	0.929/∼	0.2319/0.2544↑	+ (0.000)

[1] Days in vitro.
[2] Per 1 h.
[3] An uparrow indicates an increase in similarity, a downarrow a decrease in similarity.
[4] +/- = similarities poststimulus bursts with the stimulus responses are significantly larger/smaller than the similarities of the prestimulus bursts with the stimulus responses, ○ = no significant difference (α-value of 2.5 %). The value in brackets is the p-value.

We did the same overall hypothesis check as for the computational data. We denote by $S^*_{\text{pre},i}$ and $S^*_{\text{post},i}$ the average similarity of all pre- / postimulus bursts with all stimulus responses for every experiment i, which leads to a total of seven similarity pairs as listed in table 3.5. We then want check whether the medians of the similarities are on average higher for poststimulus bursts than for prestimulus bursts

$$\mathbf{H_0}: \quad \widetilde{S^*}_{\text{post}} - \widetilde{S^*}_{\text{pre}} \leq 0$$
$$\mathbf{H_1}: \quad \widetilde{S^*}_{\text{post}} - \widetilde{S^*}_{\text{pre}} > 0$$

Using the Wilcoxon signed rank test, we find a p-value of 0.289, which is far from being significant. If we only average over the first half of the stimulus responses instead of all stimulus response, the p-value decreases to 0.1875, while it increases to 0.4062 when considering only the second half of the stimulus responses.

4. Discussion

We will start this chapter by summarizing and interpreting the results gained by the simulations. We will then discuss how well the data gained by the MEA experiments agrees with these results and comment on the differences. The chapter concludes with a discussion on our methodology.

4.1. Simulation results

We started our research with a small-scale computational network, adapted from the model used in [39]. After approximately one day of simulation time, this network reaches a weight distribution such that weights in the network have a strong negative correlation (< -0.8) with the axonal delay of the corresponding synapse. This distribution appears to maintain itself and it can be stable for at least 10 days of network time. Depending on the parameter settings, neurons that we call leader neurons may emerge spontaneously. These are neurons that loose virtually all their synaptic input and develop very strong outgoing synapses. This happens if a neuron, that has little excitatory input, initiates bursts without firing again in the same burst. According to the spike-timing-dependent plasticity mechanism, all outgoing synapses become stronger while the incoming synapses decrease in strength. Having stronger outgoing synapses increases the chance of initiating bursts, while lower excitatory input decreases the chance of the neuron firing again in the same burst, a vicious circle that ends in all the outgoing synaptic weights of that neuron becoming one and all incoming synaptic weights becoming zero. Because the outgoing synapses of such a neuron are so strong, the neuron starts to initiate a significant amount of the bursts in the network, hence the name.

Reducing the amount of inhibitory activity, either by removing the low-threshold spiking neurons or by reducing the amount of noise applied to the inhibitory neurons, reduces the chance of the emergence of such leader neurons because it is then less likely that a neuron fires only once in a burst. In one of the simulated networks, one of the inhibitory neurons loses all its input, illustrating that the development of a leader neuron is initiated by its lack of input and not its strong outgoing synapses, since an inhibitory neuron cannot initiate a burst. It is unclear whether these leader neurons are only an artifact of the simulation or whether they emerge in cultured networks as well. In cultured neuronal networks, there are so-called burst leaders, neurons that fire first in a disproportionate number of the bursts [8]. While the leader neurons initiate a lot of bursts in the simulation and are therefore burst leaders, the opposite is not necessarily true: simulated networks without leader neurons still showed burst leaders.

4.1.1. Analysis of the equilibrium

In the equilibrium, the synaptic weights still show some fluctuations as illustrated in figure 3.6, but it appears that these are fluctuations around a constant mean. In addition, the fluctuations in the synapses seem to be independent of each other, because after the network has settled, the Euclidean distance between two randomly chosen weight matrices that are more than half a day apart shows little variation (see figure 3.7). Networks that

were re-simulated with the very same settings except for different initial weights and/or a noise realization reached the same equilibrium, making it unlikely that multiple equilibria exist in the simulated networks.

Looking at the weight distribution of a network without leader neurons in its equilibrium state (illustrated in figure 3.5b), we see that it is unimodal and positively skewed, which agrees with the theoretical predictions made by [30]. A remarkable property of the weight distribution in the equilibrium is that the average synaptic weight between two excitatory neurons drops to circa 0.3. Our synaptic plasticity model STDP measures the cross-correlation between the arrival of a presynaptic spike and the spiking of the postsynaptic neuron. If both events were uncorrelated, the equilibrium weight can be calculated at which the average weight increase is equal to the average weight decrease:

$$\begin{aligned} & A_+ &=& \ 2 \cdot S_{eq} \cdot A_- \\ \Leftrightarrow \quad & A_+ &=& \ 2 \cdot S_{eq} \cdot A_+ \cdot 1.05 \\ \Leftrightarrow \quad & S_{eq} &=& \ \tfrac{1}{2.1} \approx 0.488 \end{aligned}$$

Because of our slightly asymmetric STDP implementation, this value is just below 0.5. The equilibrium is self-regulating since the decrease is proportional to the synaptic weight while the increase is absolute, see section 2.1.3. Therefore, even if there was no correlation between the two events, we would expect the weights to average at approximately 0.5. Since the arrival of a pre-synaptic spike can cause the post-synaptic neuron to fire, we actually expect a positive correlation, making the low synaptic weight average counterintuitive. We will now give a possible explanation for this observation.

We showed that the weights in the equilibrium show a strong negative correlation with the associated axonal delay, which is illustrated in figure 3.11. Neurons that are connected by an axon with a relatively short delay are on average strong and vice versa. The explanation for this lies in the shape of the bursts realized by a network. Synaptic weights change predominantly during bursts due to the large amount of temporally close spikes. When looking at the burst pattern shown in figure 3.1, it can be seen that bursts start with some individual spikes, followed by a shorter period of high-intense spiking, in which all excitatory neurons fire with a high frequency within a window of approximately 20 ms.

We show in appendix E that if every neuron fired only once within this window with a uniform chance distribution and independent of each other, the expected value of the synapse would be a function of its axonal delay as plotted in figure E.2. The equilibrium value of ≈ 0.5 is only reached if the axonal delay us 0 and for increasing axonal delays, the equilibrium decreases.

Assuming that the neurons fire independently is not too far fetched. Figure 4.1 shows the input profile of a neuron during a burst. In the beginning, a couple of individual spikes arriving at the neuron can be identified. Here, the precise timing is crucial in determining whether the input leads to a spike. At approximately $t = 40$ ms, the input increases strongly, due to the large number of excitatory neurons in the network that spike simultaneously. Because of the strong input, the effect of a single post-synaptic potential is rather irrelevant and it is mostly the internal neuron dynamics that determine when the neuron spikes.

Assuming that both neurons fire only once is a strong simplification. The calculation for multiple spikes during a burst is much more complicated, especially if the neurons fire an unequal amount of spikes, but qualitatively, the effect is probably the same: When an action potential arrives at a synapse that is connected by a long axon, it is more likely that there was a recent postsynaptic action potential than for synapses connected by short axons. Since axonal delays are relatively long compared to the length of the bursts, the

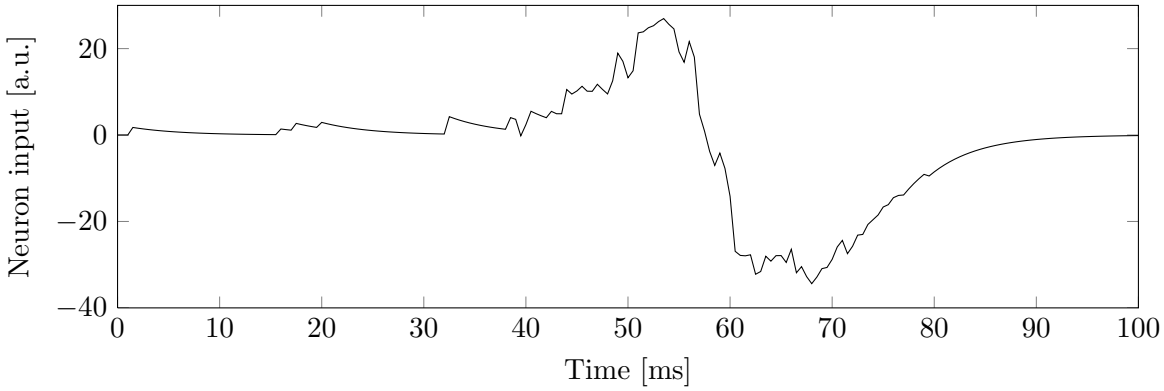

Figure 4.1.: Exemplary neuron input of one of the neurons in the simulation during a spontaneous network burst.

average synaptic weight is significantly lower than the equilibrium value of 0.488 in case of uncorrelated firing. This becomes clear when we think of a network with only two neurons that are bilaterally connected with axonal delays that are larger than delays caused by the dendritic tree. If both neurons fire simultaneously, the strength of both synapses decrease.

We simulated three networks with longer axonal delays. The bursts in these networks were wider, but the relative ratio between axonal delays and burst width was larger than in the other simulations. In agreement with the previous thoughts, the synaptic weights of the excitatory synapses were on average lower than for the other networks.

Connectivity and activity

The activity of the neurons during a burst is limited to a narrow time window. Synapses that match that activity pattern are those with short latencies, which will be relatively strong on average. Synapses with long axonal latencies do not fit into this activity pattern and therefore decrease in strength. This effect may increase itself, since it is likely that due to this change in weights, the bursts will become even narrower. This supports the idea of a balance between activity and connectivity and it would therefore be interesting to look at the development of the burst length over time in an unstimulated network.

The simple burst pattern also gives rise to the question whether our network structure is too simple. All excitatory/inhibitory neurons are statistically the same and the connection density among the neurons is very high. Therefore, activity during a burst snowballs and reaches all neurons more or less at the same time. More versatile and probably realistic burst profiles and thereby weight development may be obtained by reducing the connection density, reducing the synaptic efficacy or adding positions to the neurons in the network that affect the connections between the neurons. The latter is discussed in more detail in section 4.3.1.

4.1.2. Effects of stimulation

As the initial networks for our stimulation protocol, we chose the three networks with less inhibitory noise after 10 days of spontaneous activity. By then, they had not developed leader neurons and therefore did not have extreme synapses. We chose these networks because of two reasons. First of all, we do not know whether leader neurons occur in cultured networks or whether they are only an artifact of the network. Second, synapses can undergo stronger changes if they are not bounded at their maximum/minimum value.

Stimulation led to a change in synaptic weights that is strongest in the beginning of the stimulation and settles after approximately 2 h. Conversely, this change in synaptic weights affects the stimulus responses. Since most of the activity during the stimulation phase is composed of stimulus responses, an equilibrium may establish between the connectivity and the stimulus responses, similar to the equilibrium between the connectivity and spontaneous activity before stimulation.

Since the stimulus response is affected by the synaptic weights, its development may give an idea of the development of the synaptic weights. The temporal development of the stimulus response was quantified by averaging the similarity of a stimulus response with all other stimulus responses. If this average similarity is relatively low, this means that the stimulus response still changes, since it is not very similar to previous or future stimulus responses. A relatively high similarity indicates stability of the stimulus response since it has a relatively high occurrence among the stimulus responses. In the simulations, the mutual similarity is relatively low at first and then asymptotically approaches its maximum value, which is reached after circa 300 stimuli. Interestingly, the stimulus responses appear to reach their final state faster than synaptic weights.

We averaged every stimulus response with all the pre- / poststimulus bursts to get an impression of the development of the stimulus response relative to the intrinsic pre- and poststimulus activity in the network. These similarities are plotted in figure 3.18a. Noteworthy, their similarity with both the bursts before and after the stimulation increase for approximately the first 150 stimulus responses. This means that not only does the network adapt to the stimuli, but also the stimulus responses adapt to the internal network dynamics. This supports the idea that by stimulation, a new balance between stimulus responses and connectivity develops.

After the first 150 stimuli, the similarity of the stimulus responses with all the prestimulus bursts does not further increase. It seems that there is in fact even a small decrease before the similarity reaches its final value. The similarity with the poststimulus bursts on the other hand, and thereby the difference between the two similarities, continue to increase. For the first stimulus responses, the similarity with the bursts after the stimulation is just slightly higher than the similarity with the bursts before stimulation. The difference increases by a factor of 5 within approximately 700 stimulus responses. The relative differences in the similarities are very small. This is due to the high resemblance between all the network bursts and stimulus responses in the simulated networks as described in the previous section. All the similarities are therefore close to their maximum value of 1. The significance of the similarity measure is discussed in section 4.3.3.

To test the statistical significance of the difference, we averaged the similarity of every burst with all the stimulus responses. The reason for this is that we cannot treat every stimulus response / burst pair as an independent observation, since the stimulus responses are not independent of each other, i.e. if a burst has a high similarity with one of the stimulus responses, it is likely that it has a strong similarity with all the stimulus responses. For every simulation, we calculated the average similarity of every individual burst with all the stimulus responses. Per simulation, it was then tested whether this similarity was significantly higher for poststimulus bursts than for prestimulus bursts. Because there is a relatively large amount of bursts, a two-sample t-test was used even though we do not know the underlying distributions of the similarities.

To test the overall hypothesis, the average similarities per burst were averaged again over all pre- / poststimulus bursts in the network, leading to a pair of similarities for every simulation, describing the similarity between the induced activity with the spontaneous activity before and after stimulation. We tested using the Wilcoxon signed rank test,

whether the mean of the similarity with the poststimulus activity was significantly higher than the mean of the similarity with the prestimulus activity. For this test, the individual observations need to be independent. Since we only used three distinct networks for all 12 networks, there were only three distinct groups of prestimulus bursts, so it can be argued that the average similarities of the prestimulus bursts with the stimulus responses are not fully independent. These similarities however are calculated using not only the prestimulus bursts but also the stimulus responses, for which different stimulation neurons are used.

The bursting frequency of the networks increased strongly after stimulation and the burst leader distributions changed significantly. We therefore assume that there is a connection between a change in bursts leaders and a change in spontaneous activity. This is supported by the fact that the only simulation in which the average similarity of the stimulus responses with the spontaneous activity was lower after stimulation than before, showed the smallest change in the burst leader distribution.

4.1.3. Interpretation of the simulation results

The phenomenon that induced patterns become visible in the spontaneous activity after stimulation is very interesting, since it has been observed in the hippocampi and visual cortices of mammalian brains as well. One of the experiments described in [43] was similar to this study but with natural stimuli and biological networks. A cat was repeatedly exposed to a natural movie for a couple of minutes in total while the activity of multiple units in the primary visual cortex was recorded. Before and after the natural movie, the spontaneous activity was measured by presenting a blank movie. The spontaneous activity after the exposure to the natural movie contained traces of the activity evoked by the natural movie that were not present in the activity before exposure. Similar observations are made for special pyramidal cells in the hippocampus of some mammalian brains, known as place cells. Place cells show activity patterns that represent the current environment of the mammal. Even after external influences have been removed, these patterns reoccur.

The last phenomenon is called replay. Interestingly however, the reoccurring patterns in replay can be much faster than the original patterns and the temporal order may be reversed. Our model is unable to explain these changes in timing. In addition, it seems plausible that a network can incorporate and thereby "learn" multiple stimulation patterns. In our model, it is very likely that stimulating the network again at different neurons will eventually lead to an equilibrium between the new stimulus responses and the new connectivity that is independent of the previous stimuli. In between however, a mixed equilibrium that incorporates the equilibria of both stimuli may emerge. A more complicated network structure as discussed in 4.3.1 may allow multiple equilibria in the network. This way, multiple stimuli may be remembered on a longer timescale and independent of each other.

4.2. Experimental results

The main results gained by the simulations are that the stimulus response stabilizes after repeated stimulation and that the spontaneous activity after stimulation has a higher similarity with the stimulus responses than the spontaneous activity before the stimulation. When looking at the averaged development of the stimulus response (figure 3.19), no clear trend can be observed. Looking at the individual plots as shown in appendix D, it can be seen that for most experiments, the average similarity of a stimulus response with all

other stimulus responses increases at first but the time-scale of this development is very different among the experiments.

For all successful experiments, we calculated the mean similarity between the induced activity with the pre- and poststimulus activity. Averaged over the successful experiments, the similarity with the poststimulus bursts is higher but contrary to the simulations, this difference is not statistically significant. We could not find any differences between the properties of the cultures that could explain the large variation in the results. In fact, one of the cultures that was used for two experiments on two consecutive days showed very different results in either experiment (experiments 7 and 9).

We believe that the biggest flaw in the experiments, that may explain the large variation among the data, lies in the stimulation. This is discussed further in section 4.3.2.

4.3. Methodology

4.3.1. Modeling choices

Structure

Very little is known about the structure of cultured neuronal networks, i.e. the connections between neurons and the axonal delays. It seems plausible that the chance of two neurons being connected as well as the corresponding axonal delay are functions of the distance between and/or the location of the neurons. Neurons with a smaller distance between them may be more likely to be connected and have on average shorter axons. Indeed, the authors of [5] showed that the mutual positions of electrodes affect the mutual activity, since their center-of-trajectory (CAT) measure lost significance when the positions of the neurons were removed: "By applying a shuffling method to the CAT analysis to erase spatial information about recording location in its calculation, we found that changes in activity patterns recorded from neighboring electrodes were not independent and contributed to the better performance of CAT to detect plasticity. The network plasticity was *region specific*: despite the apparent random connectivity of neurons, plasticity was not symmetrically distributed, and the location of neurons played a role in stimulus-induced plasticity." For their study, mature network cultures that were in vitro between one and three months were used. The authors also state that: "A common misconception regarding dissociated cultures is that they are random, homogeneous and lack structure, and thus cannot support stable changes to synaptic weights associated with memory formation. While plated from a random cell suspension, microscopic observation reveals that a heterogeneous arrangement develops over time[.]"

In order to implement advanced structures in the network, it may be necessary to increase the number of neurons, since the strong synapses and high connection density were required for the network to show bursts.

Spike-timing-dependent plasticity

STDP is the only long-term plasticity mechanism included in our model and is therefore crucial for this research. We did not find any evidence for changes that should be made to this method. We do however want to comment on some of our choices. First of all, our model uses multiplicative decrease and additive increase of synaptic strengths. Biological evidence suggests that the increase may be a mixed form between multiplicative and additive, see section 2.1.3. [30] added noise to their STDP model. This is very plausible

from a biological standpoint but we do not know whether it has any effect in the simulations. Lastly we do not model dendrites. With dendrites it takes some time for the action potentials of postsynaptic neurons to backpropagate into the synapse. Since postsynaptic spikes will therefore arrive at the synapse later, it is likely that the synapses will become stronger on average.

Homeostatic plasticity

Due the positive feedback associated with STDP, i.e. strong synapses tend to become even stronger and weak synapses even weaker, it is often suggested that neuronal networks must feature homeostatic plasticity to avoid a bimodal weight distribution. Evidence for homeostatic plasticity has been found in cultured networks [1] and several ideas of homeostatic plasticity on local and global levels exist. Usually, activity is controlled by globally or locally sensing the momentary activity and scaling the incoming or outgoing synapses accordingly. The precise underlying mechanisms are still unknown however and implementations of activity control in simulations often appear artificial. It is therefore noteworthy that the multiplicative decrease in our STDP model (see section 2.1.3), a local mechanism that depends solely on the activity of the two neurons directly connected to it, leads to an unimodal weight distribution, as long as a neuron does not loose all of its input. We therefore expect that homeostatic plasticity on the incoming synapses counteracts the emergence of leader neurons.

4.3.2. Experimental setup

The experiment design was based on the results gained by the simulations. In the simulations, the synaptic weights in the network changed for approximately 2 h of stimulation, which is the reason for choosing a stimulation period of 2 h in the MEA experiments as well. We do not know however whether the change in synaptic weights per stimulus is the same in cultured and and simulated networks. The magnitude of change in the STDP model has been chosen according to biological data, but the change in synaptic weights due to a stimulus response strongly depends on the shape of the response. Similar to bursts, the stimulus responses in the simulated networks are narrower and the neurons fire with a higher density. A synapse may therefore experience multiple changes in strength with opposite signs during one burst that neutralize each other.

It turned out that not all cultures showed a consistent response during the full stimulation phase. This suggests that cultured networks adapt to stimuli not only by changes in synaptic weights. In addition, unknown effects may take place in the culture that reduce its susceptibility to stimuli. For instance, the set of activated neurons/axons may vary due to internal adaption or electrochemical effects at the stimulation electrode. This is supported by the development of the mutual similarities among the stimulus responses. In the simulations, this similarity increases and than stabilizes. For most of the experiments, we see an increase of the similarity in the beginning as well, but it decreases again at the end. We therefore suggest to examine whether shorter stimulation periods and/or higher stimulation amplitudes lead to more consistent responses. That less stimulus responses can be sufficient to see significant changes in the poststimulus activity is demonstrated in [21, 42].

4.3.3. Data analysis

Extraction of bursts and stimulus responses

In this research, we compared the similarity of stimulus responses with spontaneous network bursts. While stimulus responses are very easy to extract, this is much more difficult for bursts since there is no common definition. In the simulation, bursts can be unambiguously identified most of the time, but this is less clear in experimental data. Whether a group of spikes is denotes as a burst may eventually boil down to a threshold setting. Our analysis may therefore be sensitive to the burst detection algorithm, since we reduce the spontaneous activity to bursts. It should therefore either be verified that our analysis is independent of the burst detection or, even better, it may be possible to examine the hypothesis independently of bursts and take all the activity before/after stimulation into account.

Comparison of bursts and stimulus responses

Several different methods can be found in literature to compare synchronous bursting patterns realized by (computational) neuronal networks. In a basic method, all the temporal information is discarded and only the number of spikes measured per electrode are compared [8]. Other methods, such as time histograms, ignore the spatial information and do not differentiate between the neurons. We chose the cross-correlation method by [33] since it combines both temporal and spatial information. The method of the "center of activity trajectory" (CAT) [5] also combines temporal and spatial data, CAT however depends on the location of the electrodes. The neurons in our simulation do not have positions and therefore we decided not to use this method.

From our results, we do not see any reason to deviate from the used similarity measure. On the contrary, it shows its significance as well as ability to detect very subtle changes when comparing the stimulus responses in the simulations with each other. These stimulus responses show a very high mutual similarity, with similarities between 0.93 and 0.98[1]. Yet, a clear trend in the temporal development of the mutual similarity was observable. In our further analysis however, we average over all bursts before or after the stimulation which may reduce the significance. This is especially true because of the imprecise definition of a network burst as mentioned in the previous section. In [33], a clustering method is used to make subgroups of synchronous bursting events according to their mutual similarities. It would be interesting to cluster the bursts and then see if new subgroups emerge after stimulation. We could then average similarities per subgroup instead of over all bursts. Another interesting study would then be whether these groups show any correlation with burst leaders.

[1] The similarity is 1 only if the same action potentials are fired with the exact same relative timing.

5. Conclusion

This research originated from the idea that the activity and connectivity in a cultured neuronal network form a balance: The connectivity of a network determines how spontaneous activity spreads through the network, while this activity leads to changes in the connectivity according to long-term plasticity effects. If the intrinsic activity in a network due to a certain connectivity maintains this connectivity, an equilibrium between activity and connectivity is reached. It was further hypothesized that if stimulating the network induces activity that does not agree with this equilibrium, the equilibrium may be disturbed and incorporate the induced activity. The induced activity patterns should then be visible in the intrinsic activity after stimulation.

The computational network of this study perfectly supports the hypothesis. Without stimulation, the synaptic weights and the activity in the network develop a balance, such that the weights do not significantly change anymore. When applying external input through stimulation, the synaptic weights asymptotically reach new values such that the induced activity and connectivity are again in balance. The new synaptic weights temporary affect the intrinsic activity in the network after stimulation, such that it has a higher similarity with the stimulus responses than the intrinsic activity before stimulation. This can be interpreted as the network showing traces of the induced activity, due to the disturbance of the "natural" equilibrium between activity and connectivity. This is a very interesting concept, since traces of stimuli have also been found in the poststimulus activity of the hippocampi and visual cortices of mammalian brains.

The results in the experiments were less clear. Although for most experiments, the mutual similarity between the stimulus responses increased at first, we did not observe a uniform development of the stimulus responses among the experiments. Averaged over all experiments, the induced activity did have a higher similarity with the poststimulus activity than with the prestimulus activity but this difference was not statistically significant since there was a large inconsistency in the data. Some of the experimental networks showed trends very similar to their simulated counterparts, while others behaved quite differently. Even two experiments performed with the same culture on two consecutive days produced contradictory data with respect to the hypothesis.

Now the only valid conclusion from the discrepancy between simulated and experimental data is that either the simulation does not accurately mimic the behavior of a cultured network or that our experimental setup is flawed. But we already knew that our simulation is not accurately modeling reality, which leaves us with two options. Either the hypothesis is wrong for cultured networks, or the experimental setup fails to properly examine it. Due to the unambiguous results in the simulations as well as the positive trend in the experimental data, it is tempting to simply perform more experiments since the trend may become significant with more data. Its diversity however would remain unexplained. Reproducibility of the experimental data, whether supporting or contradicting the hypothesis, should therefore be addressed first.

A. Simulation verification

A.1. Neuron model & Time step

Both the neuron and synapse model are based on differential equations, which are solved using the forward Euler method. The accuracy of this method depends on the time step relative to the speed at which the variables change. A smaller time step leads to a higher accuracy but also a slower simulation speed. We therefore want to use the largest time step for which any decrease does not significantly improve the accuracy of the simulation.

The neuron model includes two variables, v and u, of which v is changing much faster than u. Izhikevich therefore updates v twice as often with half the time step used for u, in the order v-v-u, see section 2.1.2. We will now demonstrate that changing this updating sequence to v-u-v shows better results for the same time step. Figures A.1a to A.1d show the four different neuron types used in the simulation with a constant input of $I = 14$ a.u. starting at $t = 10$ ms and four different time steps from 0.125 ms to 1 ms. When looking at the chattering neuron, there is a significant change in the waveform between a time step of 1 ms and 0.5 ms. But even with the smaller time step, the firing of the fast spiking neuron is irregular. The firing becomes regular when the time step is further decreased to 0.25 ms and another decrease to 0.125 ms does not affect the waveforms of the different neuron types anymore. Figures A.2a and A.2b show the very same simulation except for the alternative updating sequence, for the time steps 0.25 ms and 0.5 ms. Even at the larger time step, the firing of the fast spiking neuron is regular and a decrease of the time step does not substantially change the waveforms.

This result seems odd at first, since repeatedly applying the update sequence leads to v-u-v-v-u-v-... either way. The difference must therefore lie in the moments of reset after an action potential was fired.

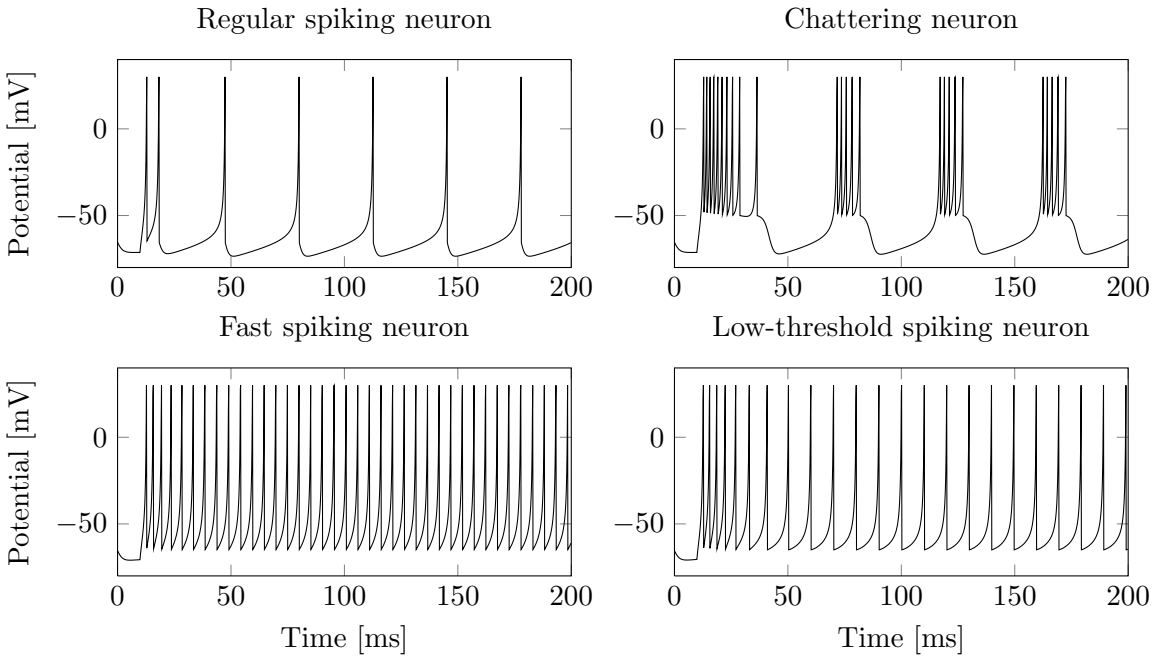

(a) Original updating sequence, $dt = 0.125\,\text{ms}$

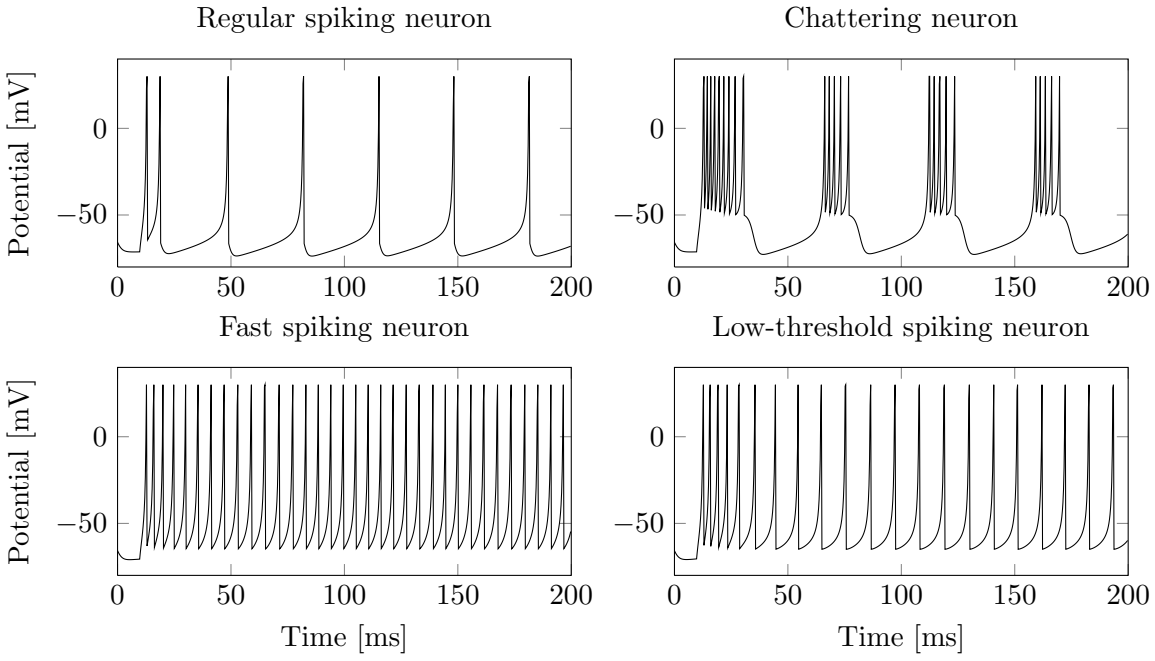

(b) Original updating sequence, $dt = 0.25\,\text{ms}$

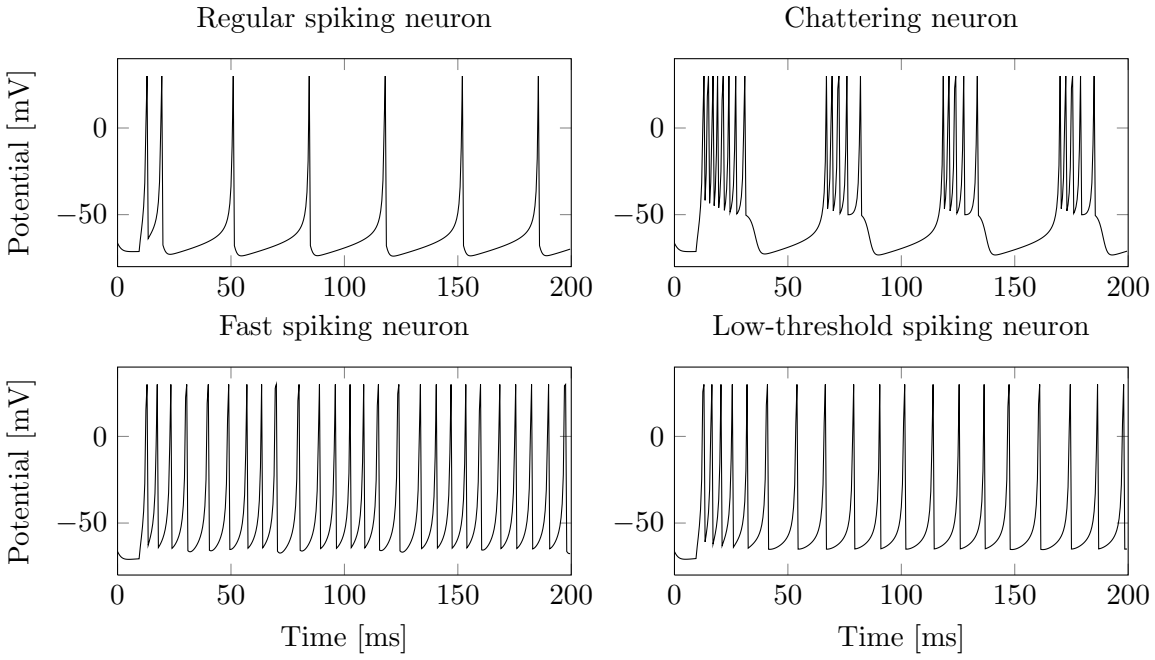

(c) Original updating sequence, $dt = 0.5\,\text{ms}$

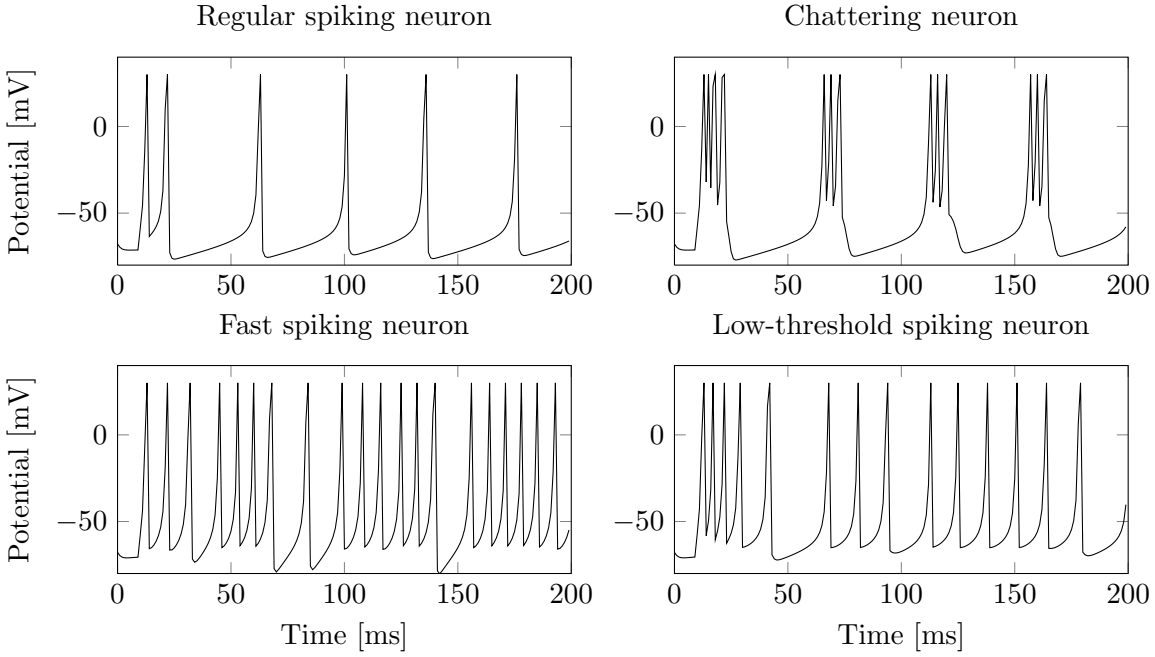

(d) Original updating sequence, $dt = 1\,\text{ms}$

Figure A.1.: The four neuron types, receiving a constant input $I = 14$ starting at $t = 10\,\text{ms}$. The time steps are (a) 0.125 ms, (b) 0.25 ms, (c) 0.5 ms and (d) 1 ms (d) and the original updating sequence v-v-u is used.

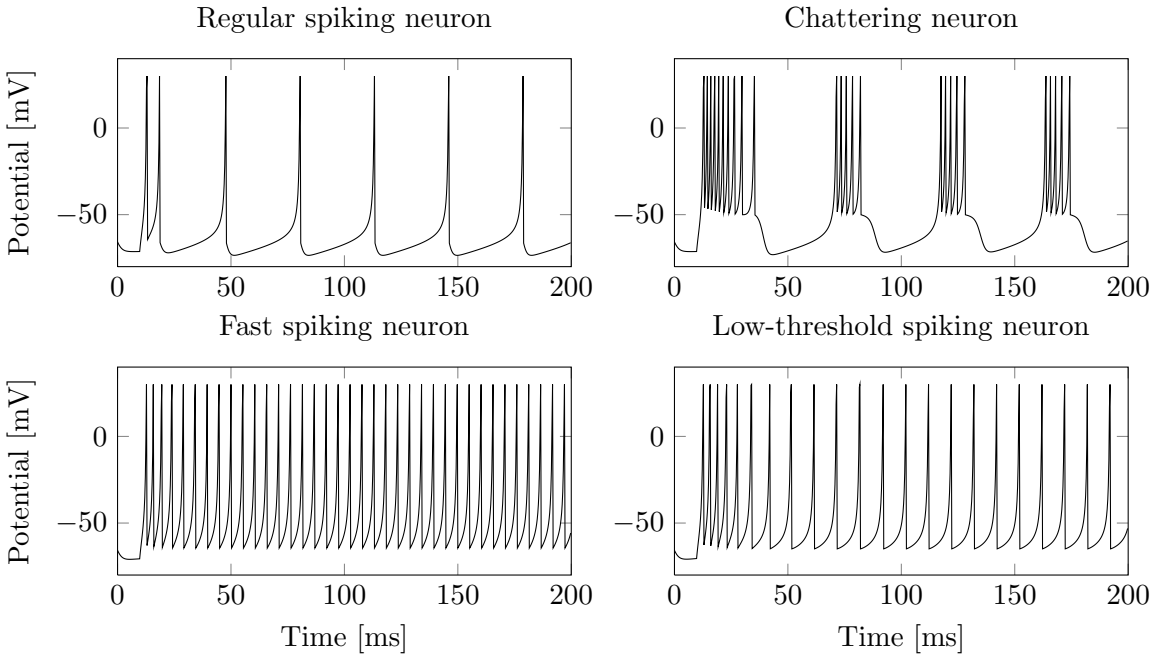

(a) Alternative updating sequence, $dt = 0.25\,\text{ms}$

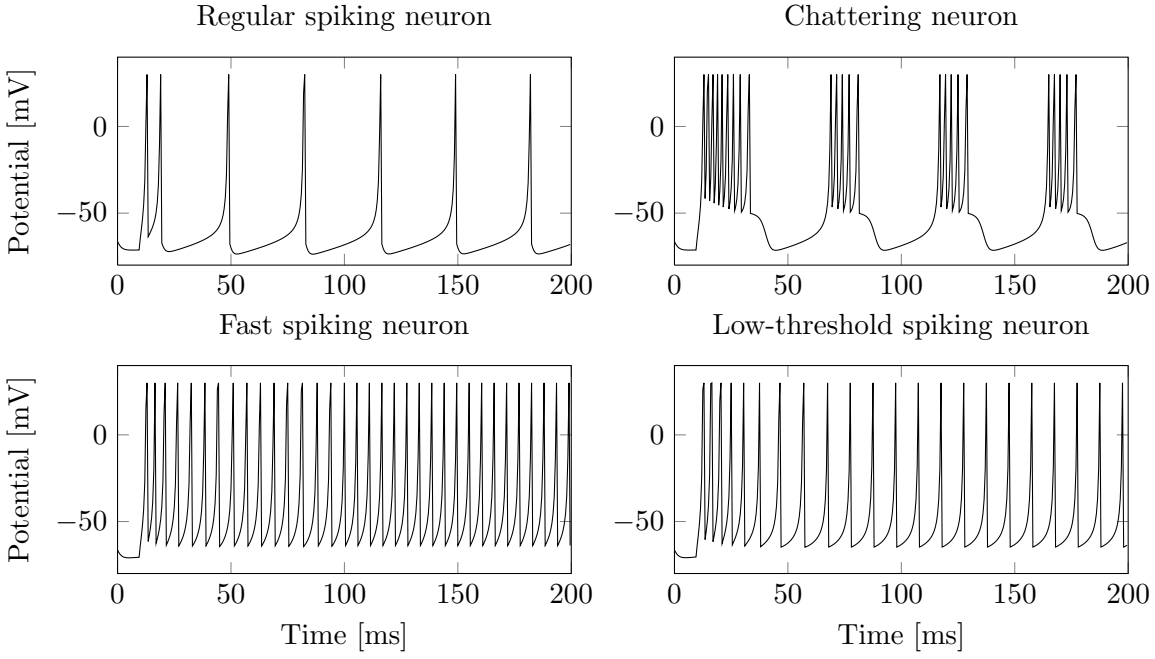

(b) Alternative updating sequence, $dt = 0.5\,\text{ms}$

Figure A.2.: The four neuron types, receiving a constant input $I = 14$ starting at $t = 10\,\text{ms}$. The time steps are (a) $0.25\,\text{ms}$ and (b) $0.5\,\text{ms}$ and the alternative updating sequence v-u-v is used.

A.2. Synapses

The model distinguishes between four different synapse types since there are four different combinations between excitatory and inhibitory synapses: EE, EI, IE and II. To test the proper working of the synapses, we set up a minimal network with all four synapse types. Two presynaptic neurons of either type were connected to an excitatory and inhibitory neuron respectively with an axonal delay of 5 ms to the excitatory neuron and 10 ms to the inhibitory neuron. The synaptic strengths were set to ± 1. The presynaptic neurons were stimulated with a frequency of 20 Hz. Figure A.3a shows the postsynaptic currents for the EE and EI synapses. The EE synapses clearly shows short-term depletion, while the EI synapse exhibits short-term facilitation. Figure A.3b shows the same plot for an inhibitory presynaptic neuron. Now the IE synapse shows short-term facilitation while the II is characterized by short-term depletion. The synaptic currents generated by the first spike are defined by the A and U values of the synapse type (see section 2.1.3) and are summarized in table A.1.

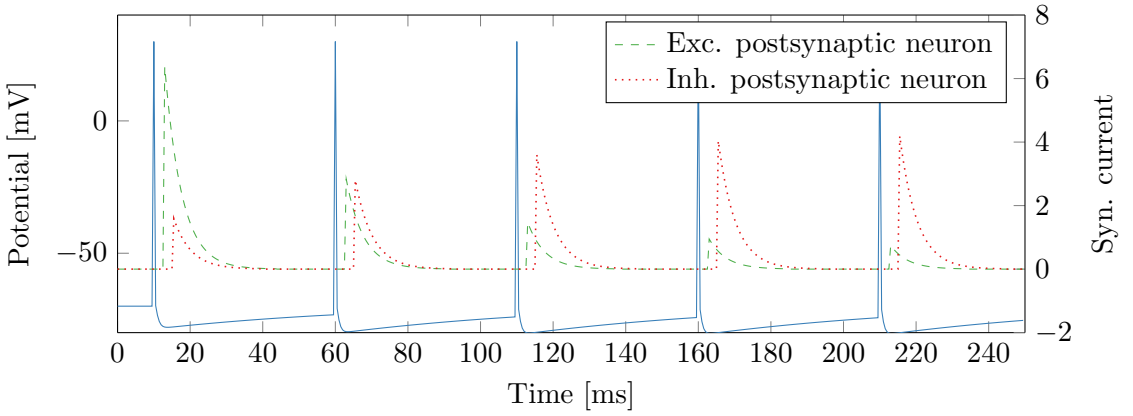

(a) Excitatory presynaptic neuron.

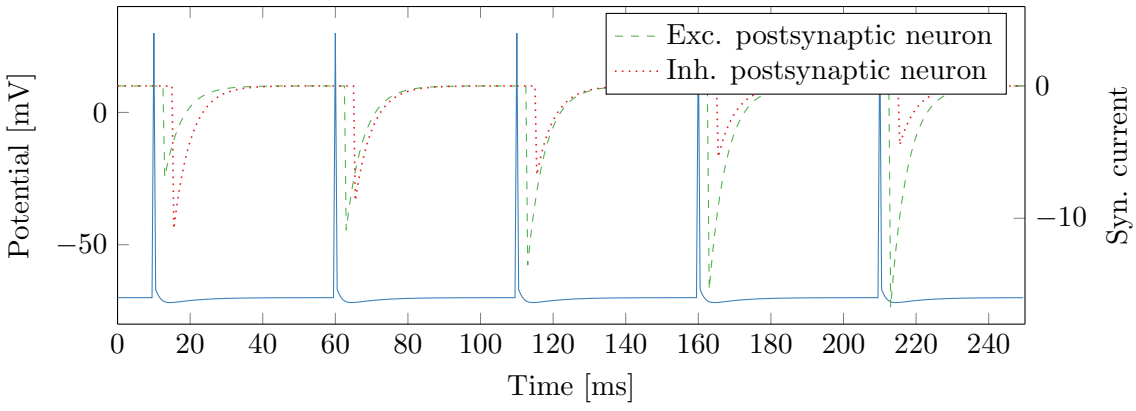

(b) Inhibitory presynaptic neuron.

Figure A.3.: A excitatory/inhibitory presynaptic neuron, whose membrane potential is depicted with a solid line, is stimulated with a frequency of 20 Hz. The dashed line shows the current generated by a synapse that is connected to an excitatory postsynaptic neuron by an axon with a delay of 5 ms. The dotted line shows the current generated by a synapse which is connected to an inhibitory postsynaptic neuron by an axon with a delay of 10 ms.

Table A.1.: Currents generated by synapses in their steady state and having a maximal synaptic strength.

Synapse (Pre → Post)	U[1]	A[2]	$U \cdot A$
Excitatory → Excitatory	0.59	10.8	6.37[3]
Excitatory → Inhibitory	0.049	32.4	1.59
Inhibitory → Excitatory	0.16	43.2	6.91
Inhibitory → Inhibitory	0.25	43.2	10.8

[1] Initial neurotransmitter fraction used, see synapse model 2.1.3.
[2] Synaptic efficacy, see synapse model 2.1.3.
[3] Sufficient to trigger an action potential at a resting target neuron.

A.3. STDP

The STDP model was tested using a network with two mutually connected neurons with axons having a delay of 1 ms. Both neurons are stimulated every 5 s with a relative timing difference of 2 ms as shown in figure A.4a. This leads to a change in synaptic weights as depicted in figure A.4b. A clear difference between the additive and multiplicative model can be seen.

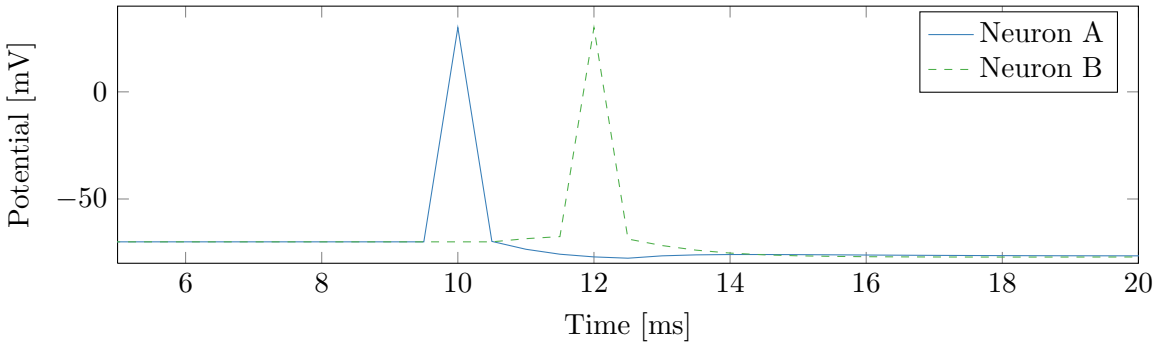

(a) Two mutually connected neurons are stimulated every 5 s with a relative delay of 2 ms.

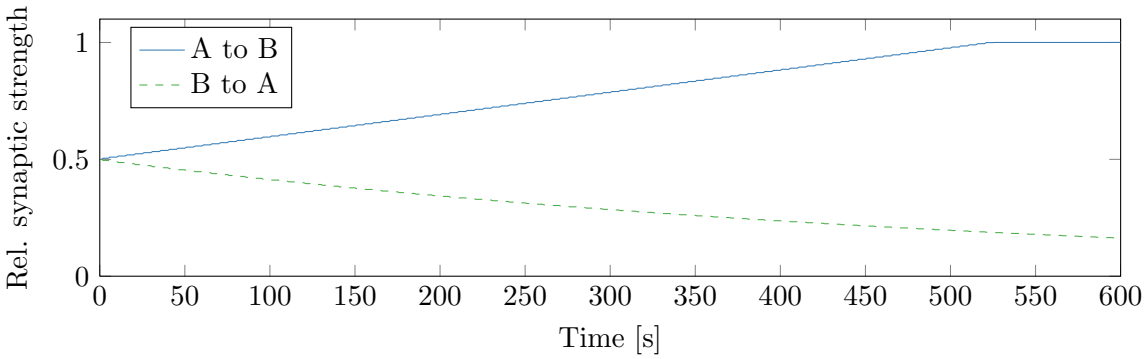

(b) Strength development of the synapses of the neurons. The strength of the synapse from the neuron that is stimulated first to the neuron that is stimulated second increases linearly while the other synapse decreases exponentially.

Figure A.4.

B. SBE similarity algorithm

Listing B.1: `ECmatrix.m`

```
function [ SimMat ] = ECmatrix( BurstsTs, BurstsCs )
%ECMATRIX Returns a matrix with the mutual correlations between
%synchronous bursting events
%   BurstsTs and BurstsCs are cell arrays containing an array with spike
%   timings and electrode numbers for every SBE.

    len = 300; % maximum time window of the SBEs
    nElectrodes = 60;

    nBursts = length(BurstsTs);
    ActivityVectors = zeros(nBursts, nElectrodes, len+1);
    SimMat = zeros(nBursts);
    NSpikes = zeros(nBursts, nElectrodes);

    % create Gaussian filter
    sigma = 3;
    siz = ceil(6 * sigma);
    if mod(siz, 2) == 1
        siz = siz + 1;
    end
    linS = linspace(-siz/2, siz/2, siz+1);
    gFilter = exp(-linS .^ 2 / (2 * sigma^2));
    gFilter = gFilter ./ sum(gFilter);

    for b = 1:nBursts
        BurstTs = BurstsTs{b};
        BurstCs = BurstsCs{b};

        if length(BurstTs) <= 1
            continue
        end

        % start at t = 20ms to fit the smoothed vector
        BurstTs = BurstTs - BurstTs(1) + 20;

        for e = 1:nElectrodes
            T = BurstTs(BurstCs == (e-1));

            nSpikes = length(T);
            NSpikes(b, e) = nSpikes;
            if nSpikes == 0
                continue;
            end

            Activity = histc(T, 0:len);
            Activity = conv(Activity, gFilter, 'same');
            ActivityVectors(b, e, :) = Activity;
        end
    end
```

```
    for b1 = 1:nBursts
        SimMat(b1, b1) = 1;
        for b2 = (b1+1):nBursts
            S = 0;

            % create vector with the average amount of spikes per
            % electrode, used for normalization
            NActive = zeros(nElectrodes, 1);
            for e = 1:nElectrodes
                NActive(e) = (NSpikes(b1, e) + NSpikes(b2, e)) / 2;
            end
            nActiveTotal = sum(NActive);
            for e = 1:nElectrodes
                Activity1 = ActivityVectors(b1, e, :);
                Activity2 = ActivityVectors(b2, e, :);

                if sum(Activity1) == 0 || sum(Activity2) == 0
                    continue;
                end

                Cor = xcorr(Activity1, Activity2, 'coeff');
                Cor = Cor(:);
                S = S + Cor .* NActive(e);
            end

            SimMat(b1, b2) = max(S) / nActiveTotal;
            SimMat(b2, b1) = SimMat(b1, b2);
        end
    end
end
```

C. Stimulus responses

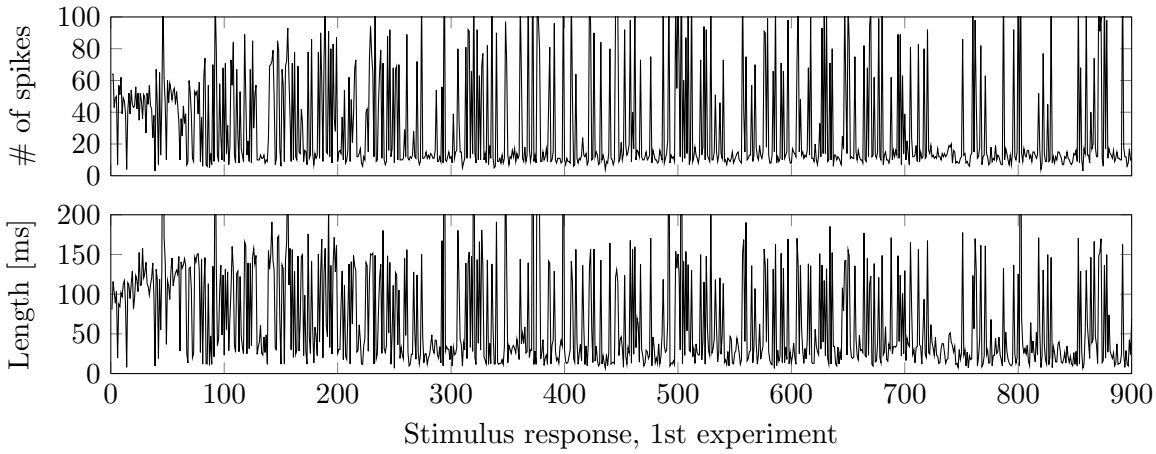

Stimulus response, 1st experiment

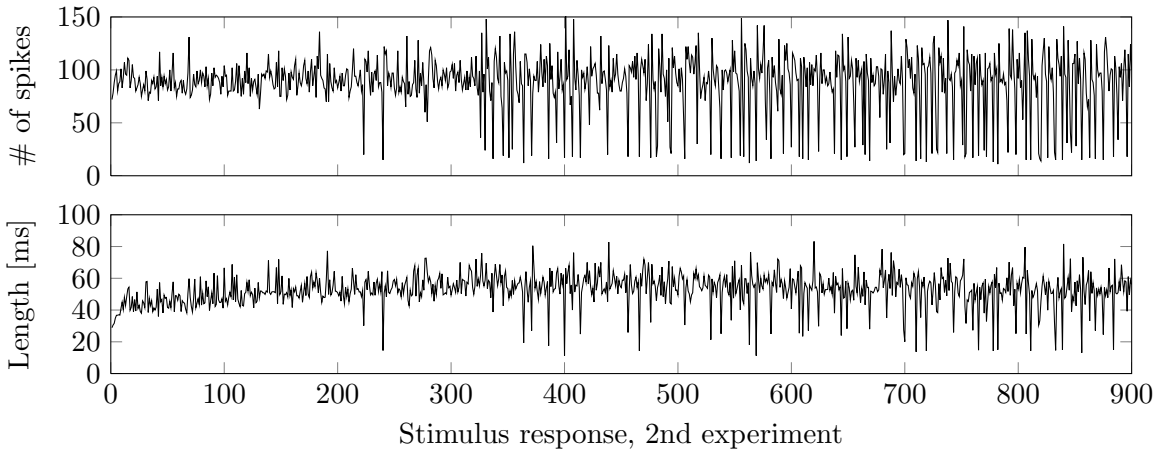

Stimulus response, 2nd experiment

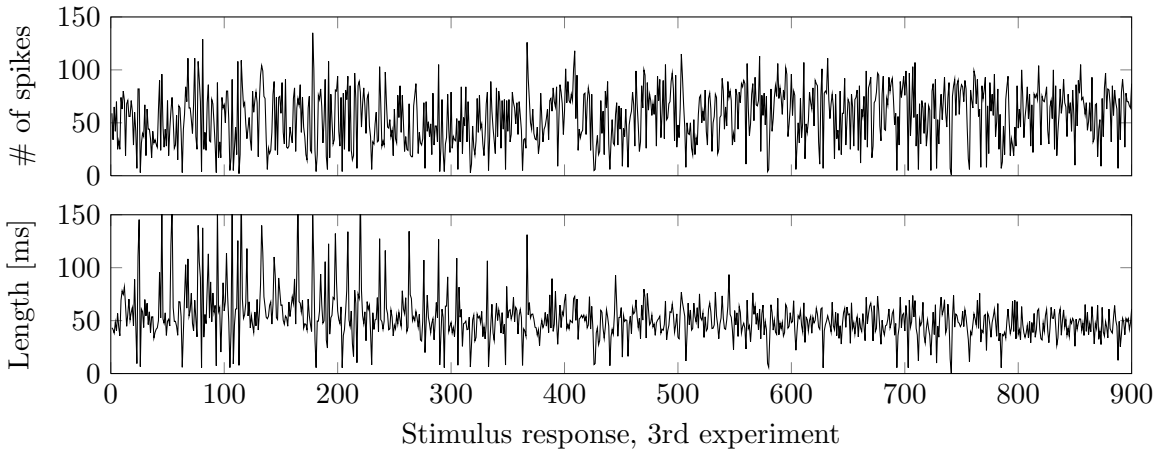

Stimulus response, 3rd experiment

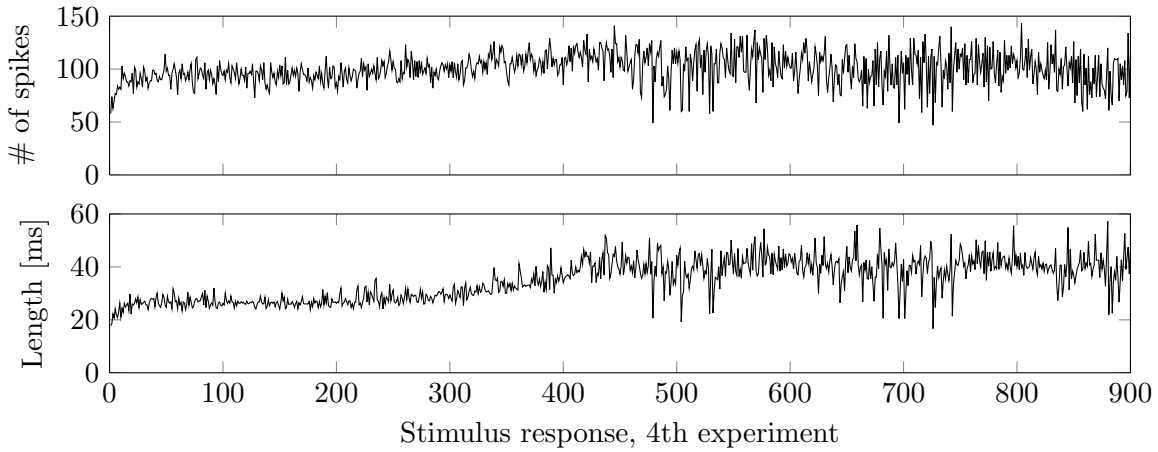

Stimulus response, 4th experiment

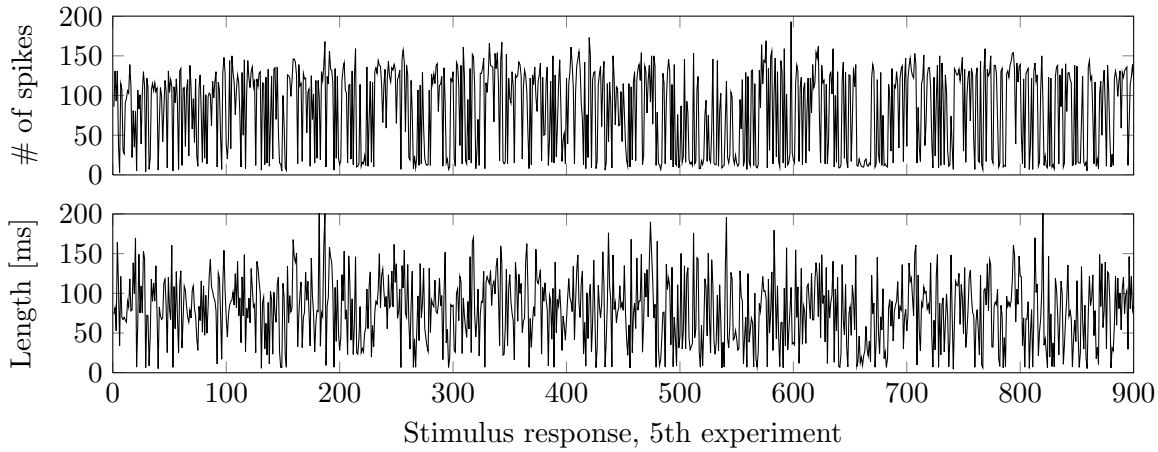

Stimulus response, 5th experiment

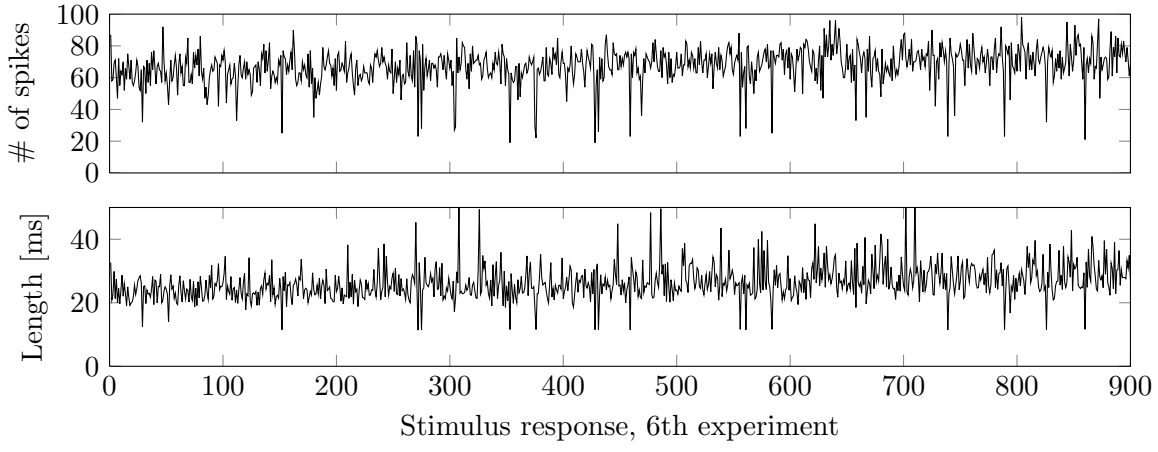

Stimulus response, 6th experiment

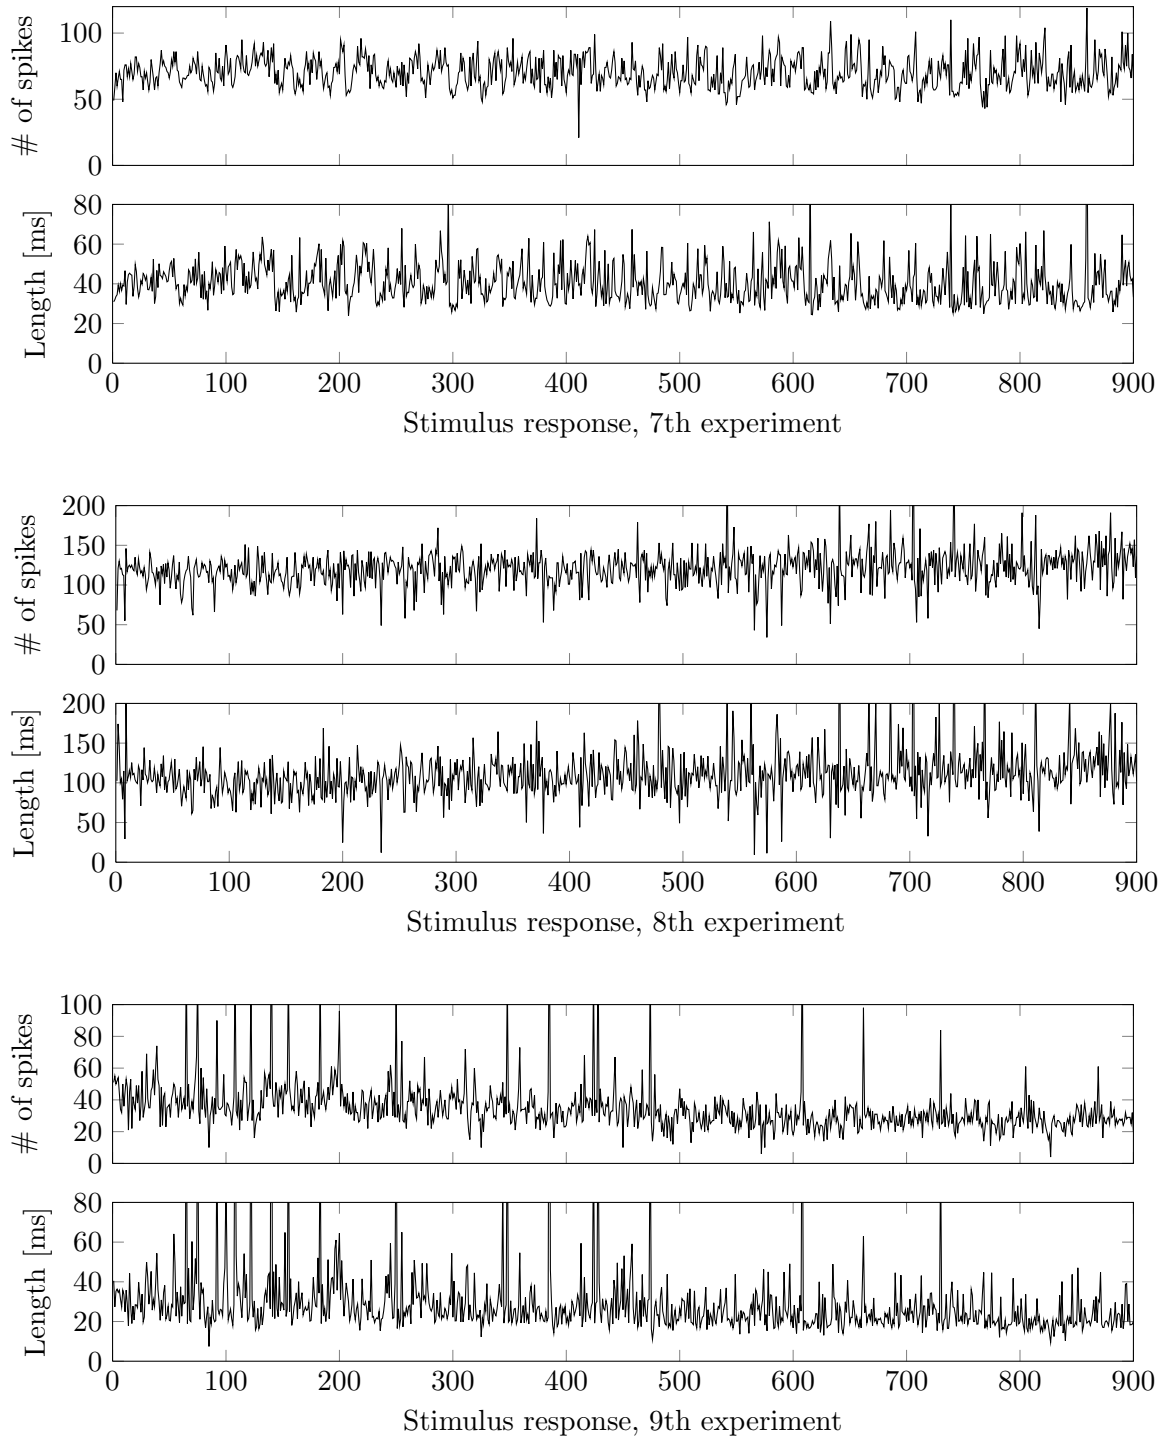

Figure C.1.: Plots showing the number of spikes and the length of every stimulus response for all nine experiments with 900 stimuli each.

D. Complementary plots for the experimental results

For every successful, the following three quantities are plotted.

- The mean similarity of a stimulus response with all other stimulus responses in the experiment.
- The mean similarity of a stimulus response with all pre- / poststimulus bursts.
- The relative change between the mean similarity of a stimulus response with all poststimulus bursts compared with the similarity of the stimulus response with all prestimulus bursts.

In the first plot, the red line indicates the filtered graph using a moving average filter with a width of 50. In the second plot the data is filtered using a median filter of width 5.

Figure D.1.: Complementary plots for the seven successful experiments on cultured networks. See paragraph above for a description of the plots.

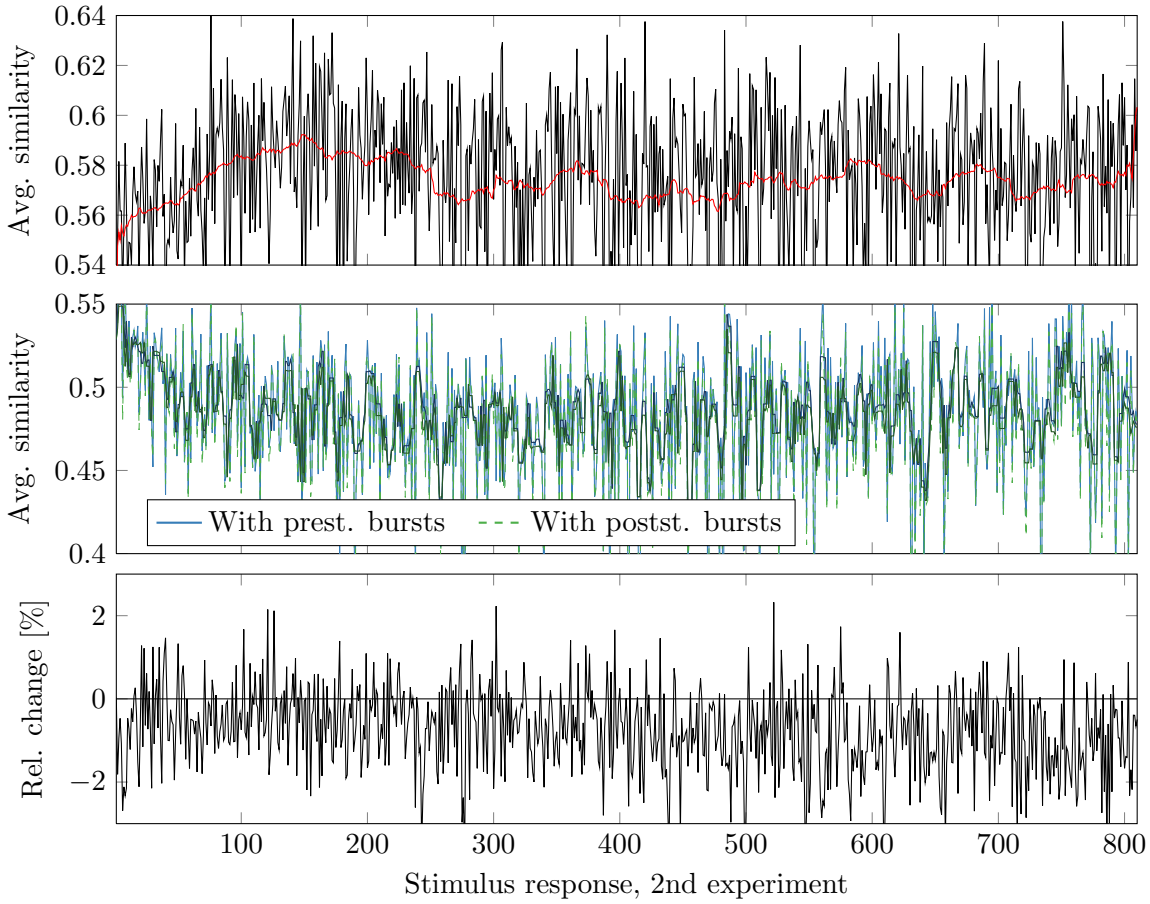

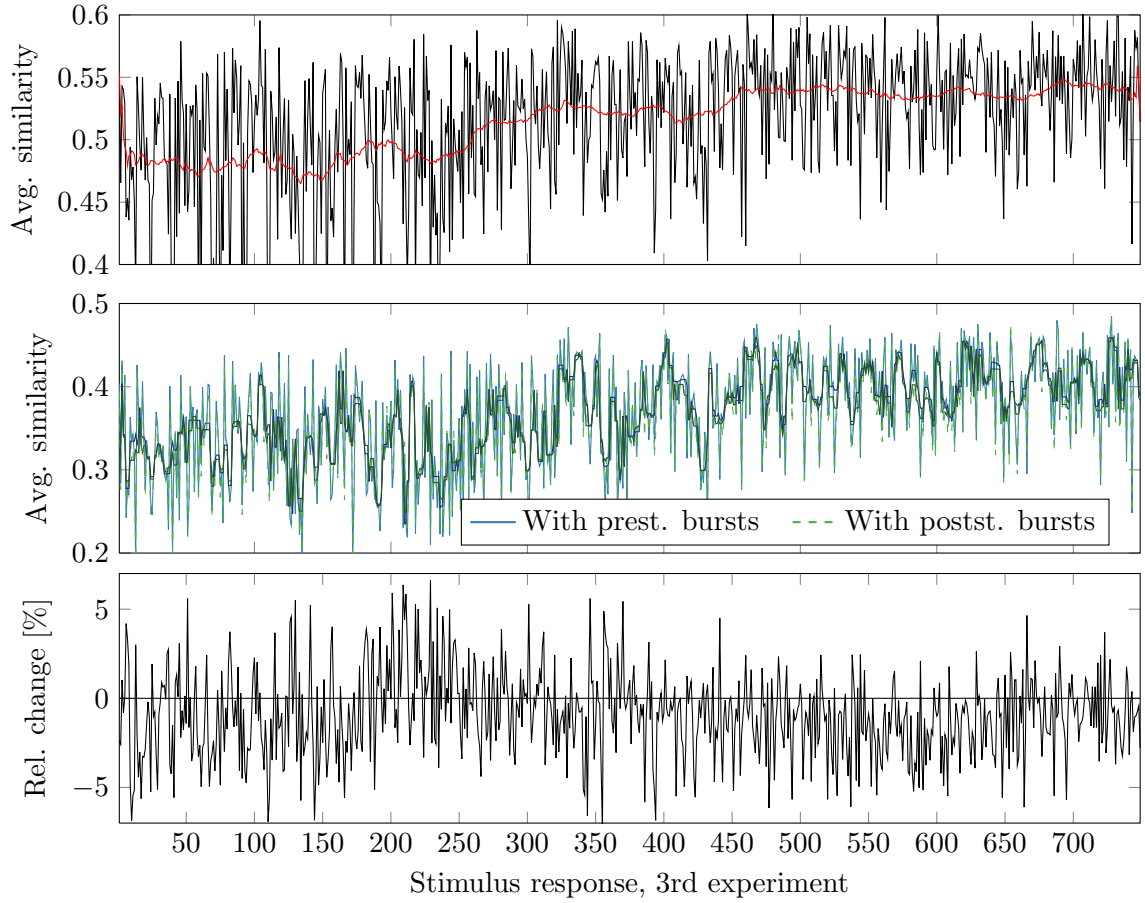
Avg. similarity
0.6
0.55
0.5
0.45
0.4
Avg. similarity
0.5
0.4
0.3
0.2
With prest. bursts
With poststt. bursts
Rel. change [%]
5
0
−5
50
100
150
200
250
300
350
400
450
500
550
600
650
700
Stimulus response, 3rd experiment

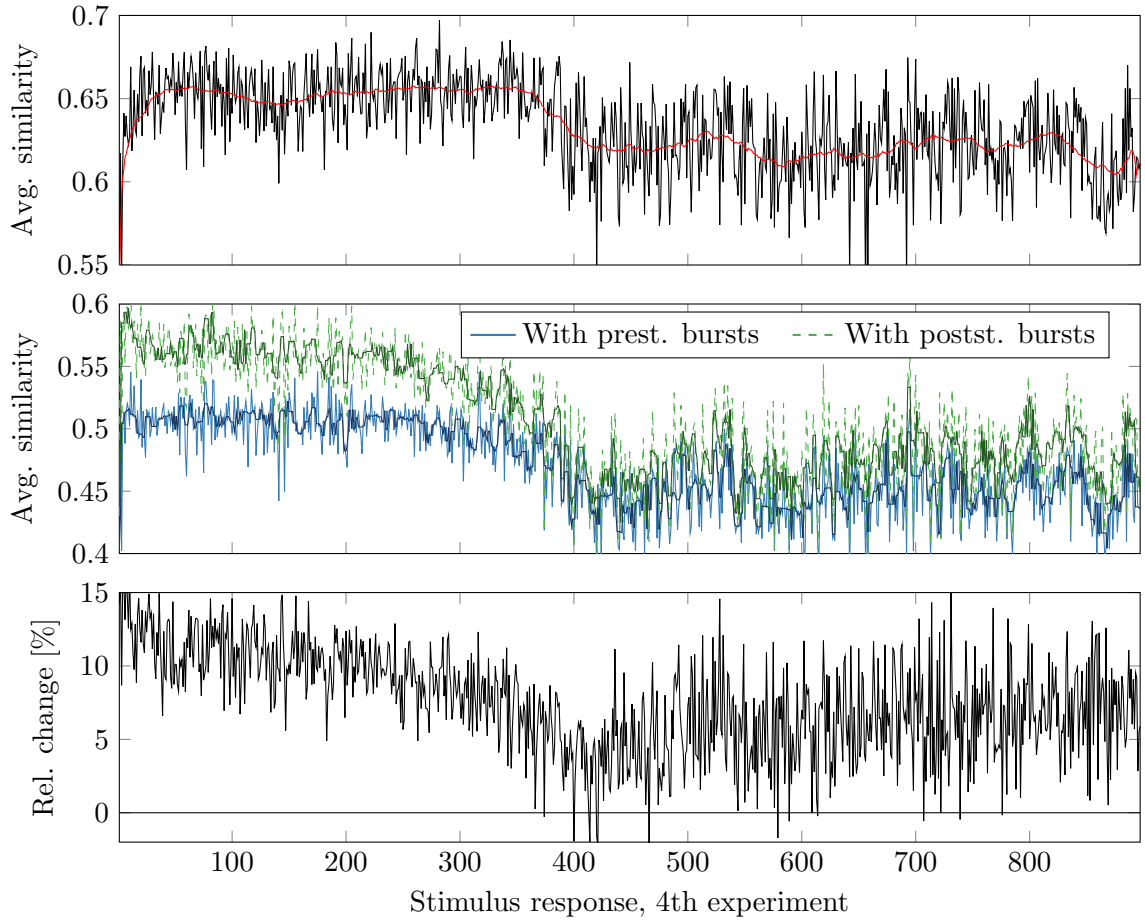
Avg. similarity
0.7
0.65
0.6
0.55
Avg. similarity
0.6
0.55
0.5
0.45
0.4
With prest. bursts
With poststt. bursts
Rel. change [%]
15
10
5
0
100
200
300
400
500
600
700
800
Stimulus response, 4th experiment

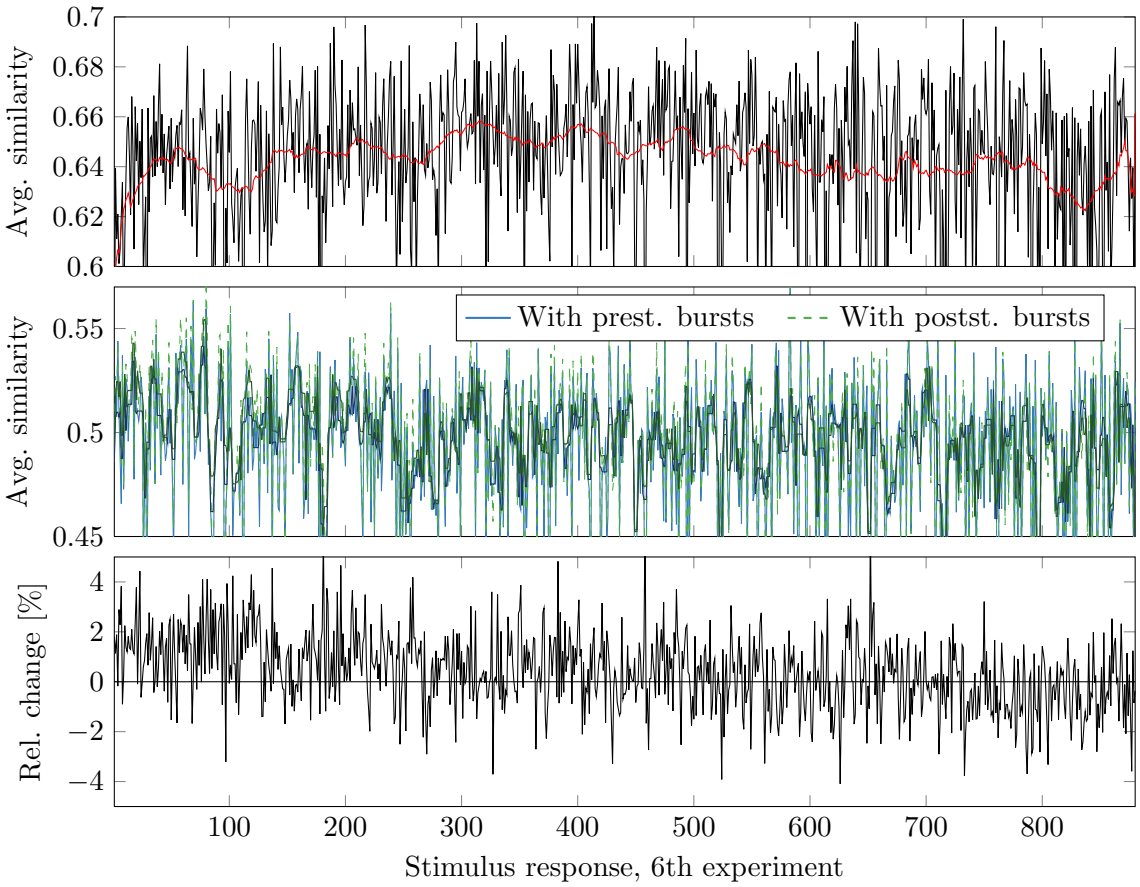
Avg. similarity
0.7
0.68
0.66
0.64
0.62
0.6
Avg. similarity
With prest. bursts
With poststt. bursts
0.55
0.5
0.45
Rel. change [%]
4
2
0
−2
−4
100
200
300
400
500
600
700
800
Stimulus response, 6th experiment

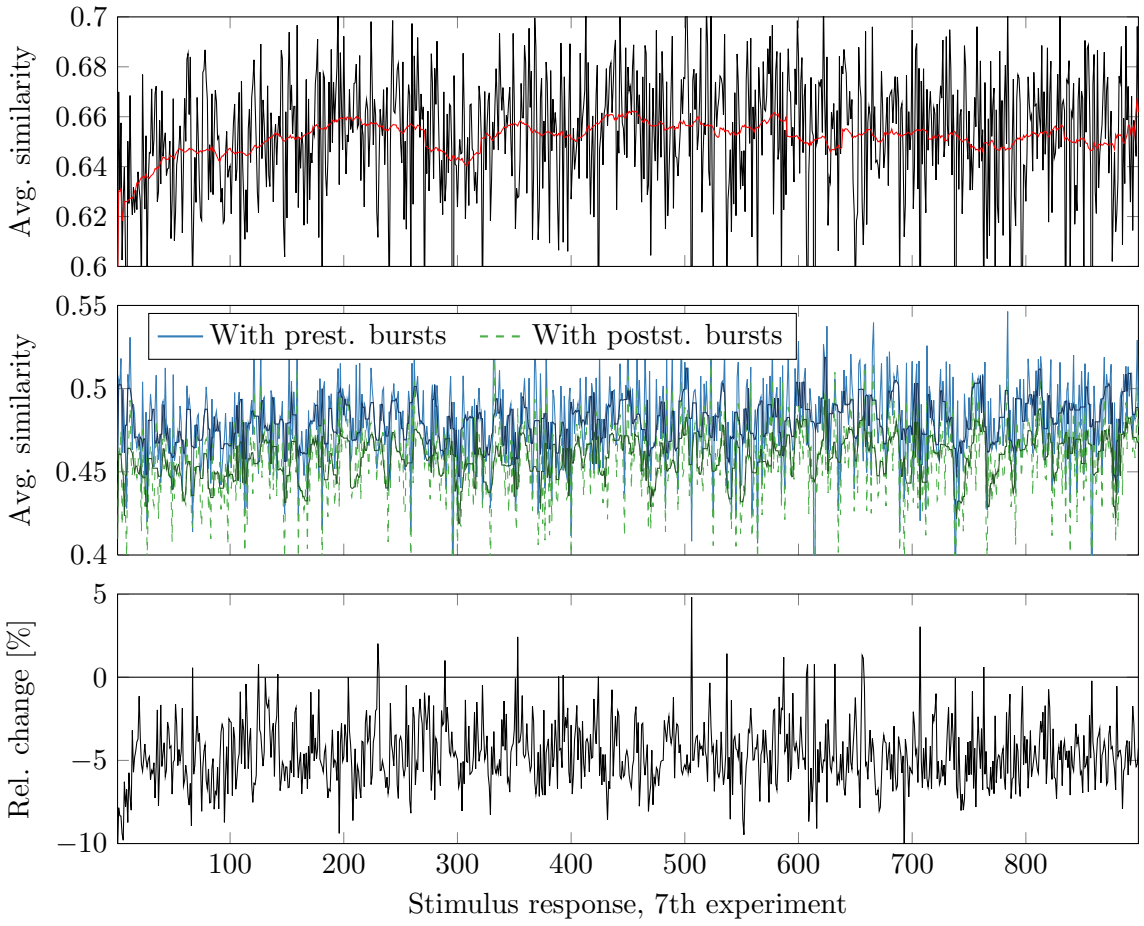
Avg. similarity
0.7
0.68
0.66
0.64
0.62
0.6
Avg. similarity
With prest. bursts
With poststt. bursts
0.55
0.5
0.45
0.4
Rel. change [%]
5
0
−5
−10
100
200
300
400
500
600
700
800
Stimulus response, 7th experiment

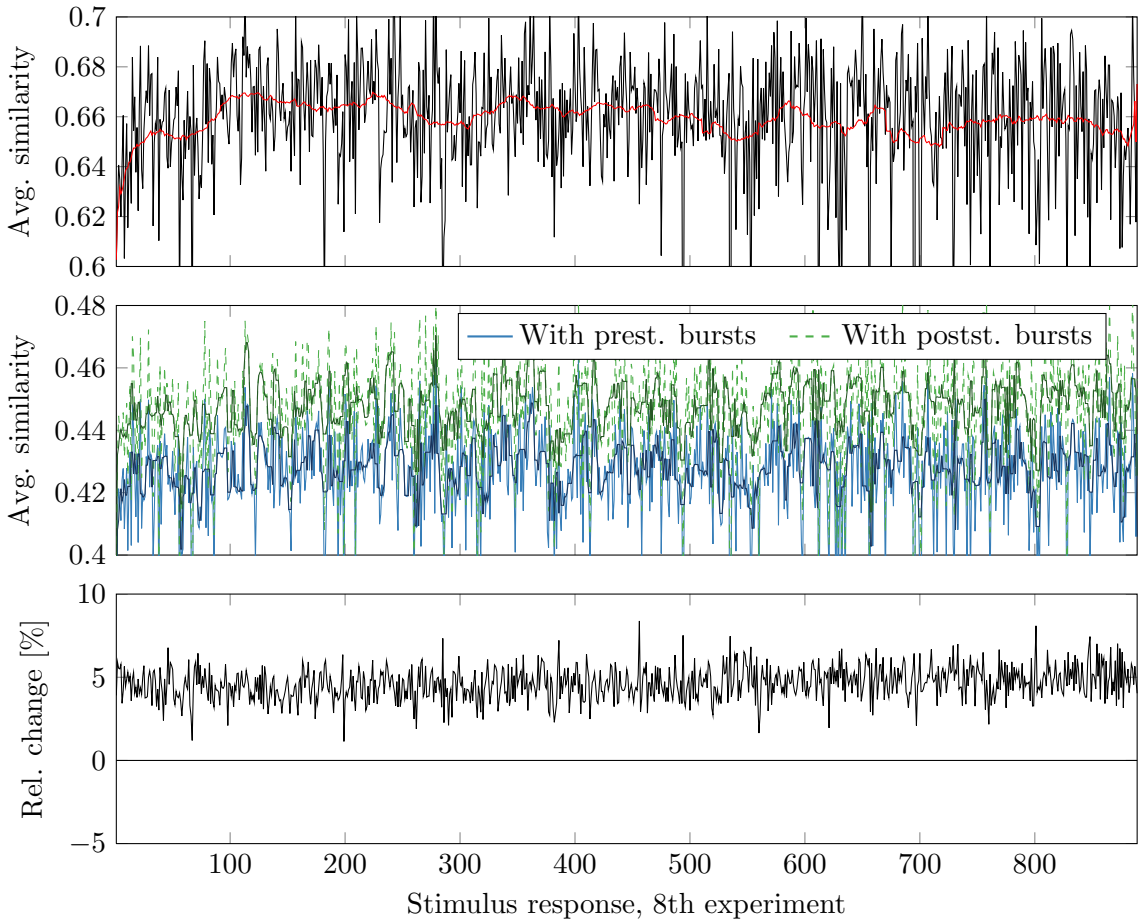
Avg. similarity
0.7
0.68
0.66
0.64
0.62
0.6
Avg. similarity
0.48
0.46
0.44
0.42
0.4
With prest. bursts
With poststt. bursts
Rel. change [%]
10
5
0
−5
100
200
300
400
500
600
700
800
Stimulus response, 8th experiment

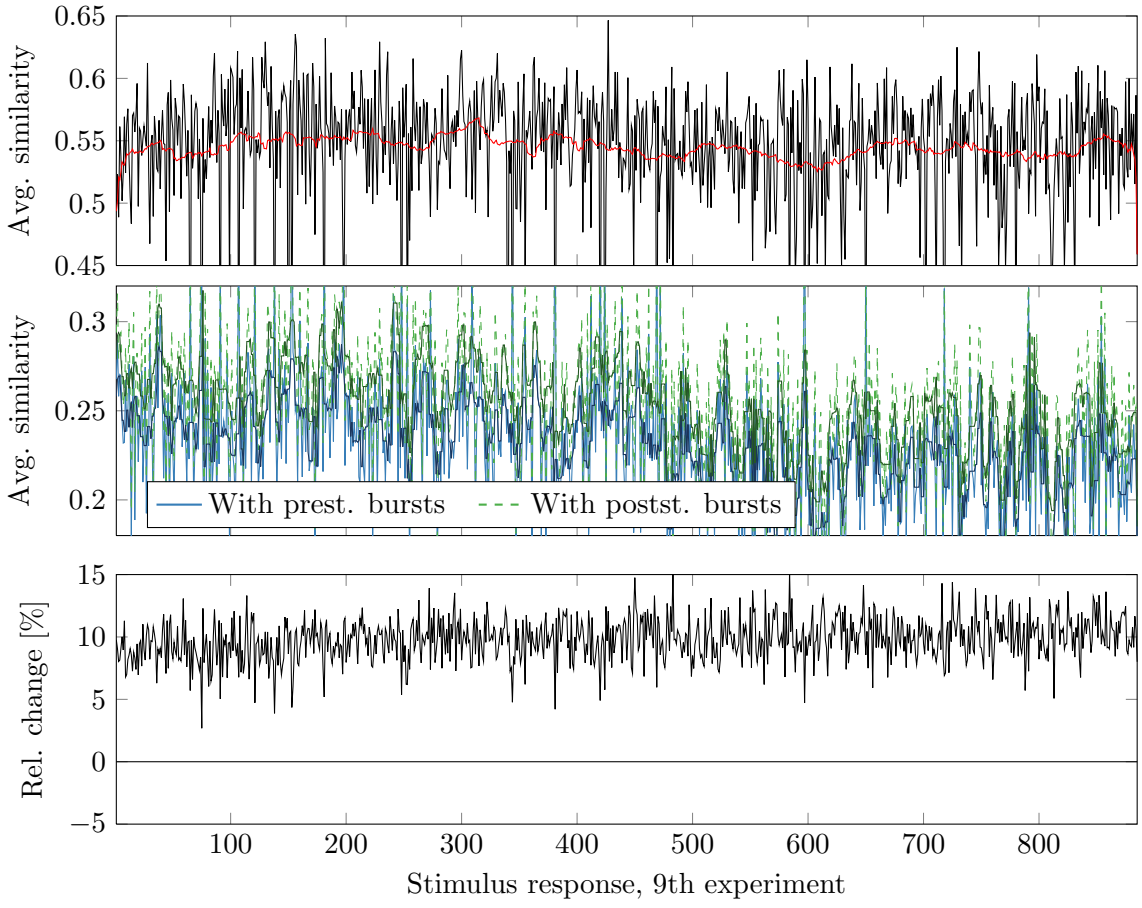
Avg. similarity
0.65
0.6
0.55
0.5
0.45
Avg. similarity
0.3
0.25
0.2
With prest. bursts
With poststt. bursts
Rel. change [%]
15
10
5
0
−5
100
200
300
400
500
600
700
800
Stimulus response, 9th experiment

E. Theoretical STDP analysis

Since the STDP model in the simulation is fully deterministic, the average synaptic weight change can be predicted if the firing probabilities of two neurons as well as the axonal delays between them are known. To start, let us assume that we have two neurons, A and B, with a synaptic connection from A $\rightarrow$ B. The axonal delay is d and both neurons fire independently of each other once within an interval $[0, T]$. If we assume a uniform chance distribution for both neurons, the chance of neuron A or B firing at $t = t'$ equals

$$p_A(t') = p_B(t') = \begin{cases} \frac{1}{T}, & 0 \leq t' \leq T \\ 0, & \text{else} \end{cases}$$

If neuron A fires at $t = t'$, the spike will arrive at the pre-synaptic terminal at $t = t' + d$. The probability density function (pdf) for the arrival of a spike at the presynaptic terminal therefore is

$$p_{\text{pre}}(t') = \begin{cases} \frac{1}{T}, & d \leq t' \leq T + d \\ 0, & \text{else} \end{cases}$$

Since we assumed that the spikings are independent and therefore uncorrelated, we can now calculate the chance of neuron B firing an interval t after the arrival of a presynaptic spike, where t can be negative to denote that B has spiked before the arrival of the pre-synaptic spike.

$$p_{\text{int}}(t) = \int_{-\infty}^{\infty} p_{\text{pre}}(t') \cdot p_B(t' + t) \cdot \mathrm{d}t'$$

If we assume $d \leq T$ and consider the two cases $t \geq -d$ and $t < -d$ we get

$$t \geq -d\text{:} \quad p_{\text{int}}(t) = \int_{d}^{T-t} \frac{1}{T} \cdot \frac{1}{T} \cdot \mathrm{d}t' = \frac{T - t - d}{T^2}$$

$$t < -d\text{:} \quad p_{\text{int}}(t) = \int_{-t}^{T+d} \frac{1}{T} \cdot \frac{1}{T} \cdot \mathrm{d}t' = \frac{T + d + t}{T^2}$$

and can combine these to

$$p_{\text{int}}(t) = \frac{T - |t + d|}{T^2}$$

This function is plotted in figure E.1. It shows a maximum for a delta time $t = -d$, i.e. neuron B spikes an interval d *before* the spike of neuron A arrives at the presynaptic terminal.

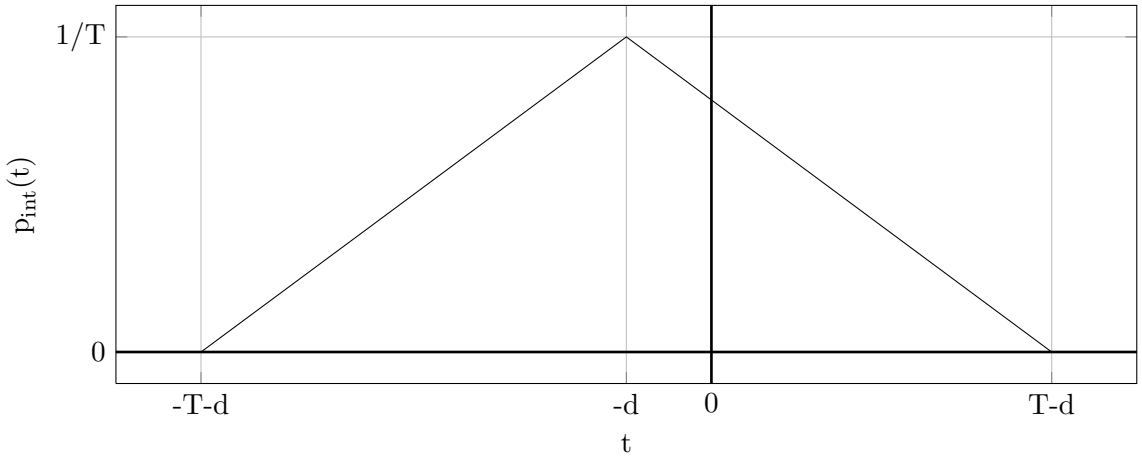

Figure E.1.: The probability density function for neuron B spiking an interval t after a presynaptic spike from neuron A arrives. It is assumed that both neurons fire independently once in the interval $[0, T]$ with a uniform pdf and that the axonal delay is d.

We can now calculate the expected synaptic weight change due to our STDP model:

$$\begin{aligned} dw &= \int_{-\infty}^{\infty} p_{\text{int}}(t) \cdot \text{STDP}(t) \cdot \text{d}t \\ &= \int_{-\infty}^{0} p_{\text{int}}(t) \cdot -A_- \cdot S \cdot \exp\left(\frac{t}{\tau_{STDP}}\right) \cdot \text{d}t + \int_{0}^{\infty} p_{\text{int}}(t) \cdot A_+ \cdot \exp\left(\frac{-t}{\tau_{STDP}}\right) \cdot \text{d}t \\ &= dw_- + dw_+ \end{aligned}$$

where S denotes the current synaptic weight between the two. Both integrals will be solved separately, starting with dw_-. For clarity, τ_{STDP} will be reduced to τ.

$$\begin{aligned} dw_- &= -A_- \cdot S \cdot \left(\int_{-T-d}^{-d} p_{\text{int}}(t) \cdot \exp\left(\frac{t}{\tau}\right) \cdot \text{d}t + \int_{-d}^{0} p_{\text{int}}(t) \cdot \exp\left(\frac{t}{\tau}\right) \cdot \text{d}t \right) \\ &= -A_- \cdot S \cdot \left(\int_{-T-d}^{-d} \left(\frac{1}{T} + \frac{t+d}{T^2}\right) \cdot \exp\left(\frac{t}{\tau}\right) \cdot \text{d}t + \int_{-d}^{0} \left(\frac{1}{T} + \frac{-t-d}{T^2}\right) \cdot \exp\left(\frac{t}{\tau}\right) \cdot \text{d}t \right) \\ &= -A_- \cdot S \cdot \frac{\tau}{T^2} \left(T + \tau - d - 2 \cdot \tau \cdot \exp\left(\frac{-d}{\tau}\right) + \tau \cdot \exp\left(\frac{-T-d}{\tau}\right) \right) \end{aligned}$$

And similarly we find

$$dw_+ = A_+ \cdot \frac{\tau}{T^2} \left(T - \tau - d + \tau \cdot \exp\left(\frac{-T+d}{\tau}\right) \right)$$

In our implementation, $A_- = 1.05 \cdot A_+$. For any d and T we can find the equilibrium S for which there is on average no change in synaptic weight. This equilibrium for $T = 30\,\text{ms}$ has been plotted in figure E.2.

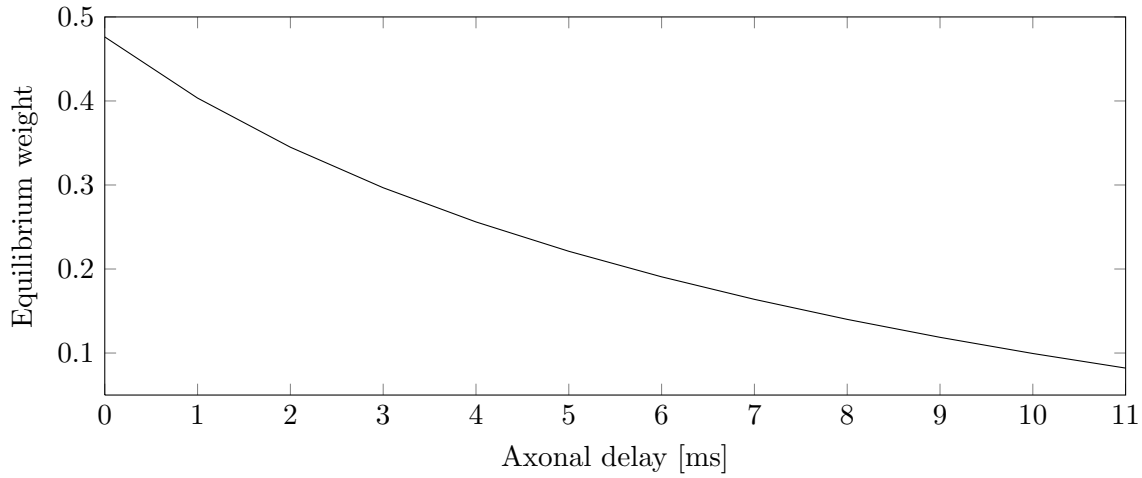

Figure E.2.: The weight between two neurons for which the average synaptic change is zero if both neurons fire once in an interval $T = 30\,\text{ms}$, per axonal delay.

Bibliography

[1] T.E. Bartlett and Y.T. Wang. Illuminating synapse-specific homeostatic plasticity. *Neuron*, 72:682–685, 2011.

[2] G.Q. Bi and M.M. Poo. Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of Neuroscience*, 18(24), 1998.

[3] G.Q. Bi and M.M. Poo. Distributed synaptic modification in neural networks induced by patterned stimulation. *Nature*, 401:792–796, 1999.

[4] P.C. Bressloff. Lectures in mathematical neuroscience. *Mathematical Biology*, 14:293–398, 2009.

[5] Z.C. Chao, D.J. Bakkum, and S.M. Potter. Region-specific network plasticity in simulated and living cortical networks. *Journal of Neural Engineering*, 4(3):294–308, 2007.

[6] C.E. Collins, D.C. Airey, N.A. Young, D.B. Leitch, and J.H. Kaas. Neuron densities vary across and within cortical areas in primates. *Proceedings of the National Academy of Sciences of the United States of America*, 107(35):15927–15932, 2010.

[7] A. Destexhe and Z.F. Mainen. Synthesis of models for excitable membranes, synaptic transmission and neuromodulation using a common kinetic formalism. *Journal of Computational Neuroscience*, 1:195–230, 1994.

[8] J-P. Eckmann, S. Jacobi, S. Marom, E. Moses, and C. Zbinden. Leader neurons in population bursts of 2d living neural networks. *New Journal of Physics*, 10, 2008.

[9] R.C. Froemke, I.A. Tsay, M. Raad, J.D. Long, and Y. Dan. Contribution of individual spikes in burst-induced long-term synaptic modification. *Journal of Neurophysiology*, 95:1620–1629, 2006.

[10] T.A. Gritsun, J. Le Feber, and J. Stegenga. Network bursts in cortical cultures are best simulated using pacemaker neurons and adaptive synapses. *Biological Cybernetics*, 102:293–310, 2010.

[11] A. Gupta, Y. Wang, and H. Markram. Organizing principles for a diversity of gabaergic interneurons and synapses in the neocortex. *Science*, 286:273–278, 2000.

[12] A.L. Hodgkin and A.F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544, 1952.

[13] M. Ichikawi, K. Muramoto, K. Kobayashi, M. Kawahara, and Y. Kuroda. Formation and maturation of synapses in primary cultures of rat cerebral cortical cells: an electron microscopic study. *Neuroscience Research*, 16:95–103, 1993.

[14] E.M. Izhikevich. Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14(6):1569–1572, 2003.

[15] E.M. Izhikevich. Which model to use for cortical spiking neurons? *IEEE Transactions on Neural Networks*, 15(5):1063–1070, 2004.

[16] E.M. Izhikevich and N.S. Desai. Relating stdp to bcm. *Neural Computation*, 15(7):1511–1523, 2003.

[17] E.M. Izhikevich and G.M. Edelman. Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences of the United States of America*, 105(9):3593–3598, 2008.

[18] Y. Jimbo, A. Kawana, P. Parodi, and V. Torre. The dynamics of a neuronal culture of dissociated cortical neurons of neonatal rats. *Biological Cybernetics*, 83, 2000.

[19] H. Kamioka, W. Maeda, Y. Jimbo, H.P.C. Robinson, and A. Kawana. Spontaneous periodic synchronized bursting during formation of mature patterns of connections in cortical cultures. *Neuroscience Letters*, 206:109–112, 1996.

[20] J. Le Feber, W.L.C. Rutten, J. Stegenga, P.S. Wolters, G.J.A. Ramakers, and J. van Pelt. Conditional firing probabilities in cultured neuronal networks: a stable underlying structure in widely varying spontaneous activity patterns. *Journal of Neural Engineering*, 4:54–67, 2007.

[21] J. Le Feber, J. Stegenga, and W.L.C. Rutten. The effect of slow electrical stimuli to achieve learning in cultured networks of rat cortical neurons. *PLoS ONE*, 5(1), 2010.

[22] E. Maeda, H.P.C. Robinson, and A. Kawana. The mechanisms of generation and propagation of synchronized bursting in developing networks of cortical neurons. *The Journal of Neuroscience*, 15(10):6834–6845, 1996.

[23] H. Markram, Y. Wang, and M. Tsodyks. Differential signaling via the same axon of neocortical pyramidal neurons. *Proceedings of the National Academy of Sciences of the United States of America*, 95:5323 – 5328, 1989.

[24] S. Marom and G. Shahaf. Development, learning and memory in large random networks of cortical neurons: lessons beyond anatomy. *Quarterly Reviews of Biophysics*, 2002.

[25] R.G.M. Morris. D.o. hebb: The organization of behavior, wiley: New york; 1949. *Brain Research Bulletin*, 50:437, 1999.

[26] A. Morrison, A. Aertsen, and M. Diesmann. Spike-timing-dependent plasticity in balanced random networks. *Neural Computation*, 19:1437–1467, 2007.

[27] A. Morrison, M. Diesmann, and W. Gerstner. Phenomenological models of synaptic plasticity based on spike timing. *Biological Cybernetics*, 98:459–478, 2008.

[28] E.A. Neale, W.H. Oertel, L.M. Bowers, and V.K. Weise. Glutamate decarboxylse immunoreactivity and $\gamma - [^3h]$ aminobutyric acid accumulation within the same neurons in dissociated cell cultures of cerebral cortex. *Journal of Neuroscience*, 3(2):376–382, 1983.

[29] H.J. Romijn, F. Van Huizen, and P.S. Wolters. Towards an improved serum-free, chemically defined medium for long-term culturing of cerebral cortex tissue. *Neuroscience & Biobehavioral Reviews*, 8:301–334, 1983.

[30] M.C.W. Rossum van, G.Q. Bi, and G.G. Turrigiano. Stable hebbian learning from spike timing-dependent plasticity. *The Journal of Neuroscience*, 20:8812–8821, December 2000.

[31] M.E. Ruaro, P. Bonifazi, and V. Torre. Toward the neurocomputer: Image processing and pattern recognition with neuronal cultures. *IEEE Transactions on Biomedical Engineering*, 52(3):371–383, 2005.

[32] J. Rubin, D.D. Lee, and H. Sompolinsky. Equilibrium properties of temporally asymmetric hebbian plasticity. *Physical Review Letters*, 86:364–367, Jan 2001.

[33] R. Sevev, I. Baruchi, E. Hulata, and E. Ben-Jacob. Hidden neuronal correlations in cultured networks. *Physical Review Letters*, 92(11), 2004.

[34] G. Shahaf and S. Marom. Learning in networks of cortical neurons. *The Journal of Neuroscience*, 21:8782–8788, November 2001.

[35] P.J. Sjöström, G.G. Turrigiano, and S.B. Nelson. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32:1149–1164, 2001.

[36] S. Song, K.D. Miller, and L.F. Abbott. Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, 3(9):919–926, 2000.

[37] J. Stegenga, J. Le Feber, E. Marani, and W.L.C. Rutten. Analysis of cultured neuronal networks using intraburst firing characteristics. *IEEE Transactions on Biomedical Engineering*, 55(4):1382–1390, April 2008.

[38] M. Tsodyks, A. Uziel, and H. Markram. Synchrony generation in recurrent networks with frequency-dependent synapses. *The Journal of Neuroscience*, 20, 2000.

[39] T.M. Veenendaal van. Computational modeling of memory in neuronal networks. Master's thesis, Twente University, 2012.

[40] D.A. Wagenaar, T.B. Demarse, and S.M. Potter. Meabench: A toolset for multi-electrode data acquisition and on-line analysis. *2nd International IEEE EMBS Conference on Neural Engineering*, 2005:518–521, 2005.

[41] D.A. Wagenaar, J. Pine, and S.M. Potter. Effective parameters for stimulation of dissociated cultures using multi-electrode arrays. *Journal of Neuroscience Methods*, 138:27–37, 2004.

[42] T. Witteveen. External stimuli and connectivity in cultured neuronal networks. Master's thesis, Twente University, 2012.

[43] H. Yao, L. Shi, F. Han, H. Gao, and Y. Dan. Rapid learning in cortical coding of visual scenes. *Nature Neuroscience*, 10(6):772–778, 2007.

[44] L.I. Zhang, W.H. Tao, C.E. Holt, W.A. Harris, and M.M. Poo. A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, 395:37–44, 1998.