

Computing synaptic plasticity

En beräkningsmodell av synaptisk plasticitet

Lovisa Irberger

Supervisors: Anders Lansner^{1 2}, Pawel Herman¹, Lars Arvestad²

Examiner: Anders Mörtberg², Kristoffer Sahlin²

Reviewer: Chun-Biu Li²

Submission date: 2026-08-27

KTH, Department of Computational Science & Technology

Computational Brain Science Lab

SU, Department of Mathematics

"Men ought to know that from nothing else but the brain come joys, delights, laughter and sports, and sorrows, griefs, despondency, and lamentations. And by this, in an especial manner, we acquire wisdom and knowledge, and see and hear and know what are foul and what are fair, what are bad and what are good, what are sweet and what are unsavory. . . . And by the same organ we become mad and delirious, and fears and terrors assail us. . . . All these things we endure from the brain when it is not healthy. . . . In these ways I am of the opinion that the brain exercises the greatest power in the man.

—Hippocrates, *On the Sacred Disease* (Fourth century B.C.E.)

Abstract

Neuroplasticity has been of great interest for decades, and has been extensively studied through the lenses of biology, psychology and biochemistry, among others. Our brains are malleable, and we today know this remains true throughout adulthood and old age. We can learn, adapt and rewire neural pathways. The human brain can change in remarkable ways — both to its advantage and disadvantage. The changing of neural pathways, in relation to activity, is referred to as synaptic plasticity. This is a subtype of neuroplasticity which specifically refers to the ability of synapses (the microscopical connections between neurons) to strengthen and weaken in regards to activation. Synaptic plasticity is believed to play a vital role in the healthy and dysfunctional brain. It is hypothesised to impact learning, adaptation, habit forming and recovery, but also addiction, neuropsychiatric conditions (depression, Schizophrenia) and neurological diseases (Parkinson's disease, Alzheimer's disease). Neurons are found in large, complex networks in the brain, and communicate through sending spikes — which is why, in this thesis (and in many other similar models), a biologically plausible Spiking Neural Network (SNN) is used to mimic the role of neurons in the mammalian Central Nervous System (CNS). Through a computational framework, inspired from the tools of Machine Learning and employing a Bayesian learning rule (BCPNN), this thesis uses a model to compute how this specific strengthening and weakening of neuronal connections may behave. Characteristical behaviour found in simulations in a published paper are successfully recreated using a novel neural simulator (Brian2) and with a remarkably smaller (and less complex) network model. As a reference, additional simulations are done to match results of a second paper, which presented biological recordings from the cortex of mice. In this latter part, the focus is on the synaptic plasticity phenomenon of Long Term Depression (LTD).

Sammanfattning

Neuroplasticitet har varit ett relevant ämne i årtionde, och har studerats utförligt av forskare inom biologi, psykologi och biokemi, bland andra. Våra hjärnor är föränderliga, och fortsätter att vara det under hela vår livstid. Det här berör inlärning och modifiering av neurala kopplingar. Den mänskliga hjärnan kan förändras på anmärkningsvärda sätt — likväl till dess fördel som nackdel. Modifieringen av neurala kopplingar, i relation till deras aktivitet, benämns synaptisk plasticitet. Det här är en form av neuroplasticitet som specifikt refererar till förmågan av synapser (de mikroskopiska kopplingarna mellan neuroner) att förstärkas och försvagas i relation till aktivering. Synaptisk plasticitet tros spela en avgörande roll i den hälsosamma och dysfunktionella hjärnan. Många teorier idag pekar på att synaptisk plasticitet påverkar inlärning, beteendeförändringar och återhämtning, men också beroendeproblematik, neuropsykiatriska tillstånd (depression, Schizofreni) och neurologiska sjukdomar (Parkinsons sjukdom, Alzheimers sjukdom). Miljarderna av neuroner är organiserade i stora, komplexa nätverk i hjärnan, och kommunicerar genom att skicka spikar till varandra. I detta arbete (och i många andra liknande modeller) används ett biologiskt likt spikande neuralt nätverk (eng: SNN) för att efterlikna rollen av neuroner i

det centrala nervsystemet (CNS). Genom en beräkningsmodell, inspirerad från Maskininlärning och som använder en Bayesisk inlärningsregel (BCPNN), använder detta arbete en modell för att beräkna hur den här specifika modifieringen i styrka av neurala kopplingar kan fungera. Karakteristiskt beteende återfunnet i simuleringar från ett publicerat forskningsarbete återskapas med en ny neural simulator (Brian2) och med en anmärkningsvärt mycket mindre (och mindre komplex) nätverksmodell. Vidare simuleringar görs för att matcha resultat av ett andra forskningsarbete, där mätningar gjorda i cortex av möss redovisas, och detta används som en referens. I den här sista delen, är fokuset lagt på synaptiskt ångtidsdepression (eng: LTD), ett fenomen av synaptisk plasticitet.

Résumé

La neuroplasticité est un sujet qui depuis des décennies a suscité un grand intérêt et qui a été étudié de façon poussée au travers du prisme, entre autres, de la biologie, de la psychologie et de la bio-chimie. Nos cerveaux sont malléables, et nous savons aujourd'hui que cela reste vrai à l'âge adulte et à un âge avancé. Nous pouvons apprendre, nous adapter et changer nos connexions neuronales. Le cerveau humain peut changer de façon remarquable, à notre avantage ou à notre désavantage. Le changement des chemins de connections neuronales est appelé la plasticité synaptique. C'est un sous-domaine de la neuroplasticité qui se réfère à la possibilité pour les synapses (les connexions microscopiques entre les neurones) de se renforcer ou de s'affaiblir en fonction de leur activation. La plasticité synaptique semble jouer un rôle vital dans le cerveau fonctionnel ainsi que dans le cerveau dysfonctionnel. Selon certaines hypothèses, elle impacterait l'apprentissage, l'adaptation, et la récupération. Mais aussi les addictions, les états neuropsychiatriques tels que la dépression et la schizophrénie, ainsi que les maladies neurologiques (maladie de Parkinson, maladie d'Alzheimer). Les neurones forment de larges et complexes réseaux dans le cerveau et communiquent au moyen de l'envoi d'impulsions. C'est la raison pour laquelle dans cet essai (comme dans de nombreux autres), un Réseau de Neurones à Impulsions (RNI) biologiquement plausible est utilisé pour reproduire le rôle des neurones dans le Système Nerveux Central (SNC) des mammifères. En utilisant un cadre computationnel, inspiré des outils de l'apprentissage automatique (Machine Learning) et utilisant une règle d'apprentissage bayésien (BCPNN), ce travail utilise un modèle pour calculer comment se comportent spécifiquement ces renforcements et affaiblissements des connexions neuronales. Les comportements caractéristiques observés dans une étude antérieure publiée ont été reproduits avec succès à l'aide d'un nouveau simulateur neuronal (Brian 2) et d'un modèle de réseau remarquablement plus petit et moins complexe. Des simulations supplémentaires sont réalisées pour correspondre aux enregistrements biologiques du cortex des souris, d'un second article. Et dans cette partie, l'accent est mis sur le phénomène de dépression à long terme (LTD, Long Term Depression).

Acknowledgements

Above all, I want to dearly thank Drs. **Anders Lansner** and **Pawel Herman** (KTH), without whom this thesis would never have been realisable. Thank you, for welcoming me into your laboratory and research group in such a distinctly generous manner and for your constant kindness and light-heartedness. I owe a special thank you to Anders Lansner, who acted as my main supervisor and who offered weekly supervision and constant support throughout this process. Thank you to everyone in the Computational Cognitive Brain Science unit (Pawel Herman's research group), with whom I've had the pleasure to attend numerous weekly group meetings and share time in the office (*Angeliki, Naresh, Andrezej, Kascheng, Ferdinand, Florian, Martin — good luck with all of your future endeavours!*). My time at the laboratory has helped me gain valuable insights into the research field and academic world as a whole, and I am all the richer for the experience. Credit is certainly due to the amazing **Ferdinand Bujanowski** (UPCité, visiting Master's student), with whom I collaborated and discussed almost daily. Ferdinand created the framework in Brian2, from which the simulations in this thesis were created. Thank you, for your countless hours sharing your knowledge of the field, and for your kindness and friendship (*your future PhD awaits you, I know it will be both successful and gratifying*). I am thankful for the help of Dr. **Lars Arvestad** (SU), my official supervisor. Thank you, for initialising this collaboration and offering guidance and feedback along the way. Thank you, to Emil Bertlin and Angelica Goldring, for reading this thesis and providing useful comments. Thank you, to Christophe Lacour, for helping me translate the abstract to french. Thank you, to all of my friends at university, and all of my loved ones in life.

AI statement

For the purpose of structuring certain simulations and plots, generative AI (Claude) has been used in this thesis. No generative AI has been used in the writing of this thesis.

GitHub

Simulations were done in neural simulator Brian2, a Python3 library. The framework used can be found on github.com/irbergerlovisa/brian_BCPNN/tree/lovisa-branch/brian_bcpnn and the specific simulations used can be found on github.com/irbergerlovisa/brian_BCPNN/tree/lovisa-branch/brian_bcpnn/models/tully_2014.

Contents

1	Introduction	10
1.1	Background	10
1.2	Cognition & neural networks	11
1.3	Aim of Computational Neuroscience	12
1.4	Neurophysiology & synaptic transmission	12
1.5	Synaptic plasticity	13
1.6	Previous work	14
1.7	Spiking Neural Networks	14
1.8	Bayesian computations	15
1.9	Aim of thesis	17
2	Method	20
2.1	From Bayes' rule to spike based BCPNN	20
2.2	Local synaptic traces	22
2.3	Leaky Integrate And Fire Neuron Model	24
2.4	The Brian2 neural simulator	25
3	Results	27
3.1	Simulations from Tully et al.	27
3.1.1	Local synaptic traces	27
3.1.2	Dependency of spike timing interval	29
3.2	Simulation recreating Rodriguez et al.	33
3.2.1	Dependency of spiking frequency	33
3.3	Additional simulations of Long Term Depression	36
3.3.1	Long Term Depression	36
4	Discussion	39
5	Conclusions	42
5.1	Further research	43

Glossary

*This is the comprehensive list of important and field-specific terminology. Relevant words introduced in *italic* can be found here, with the purpose to aid the reader.*

Neuron, spike, action potential, activation, spiking: A **neuron** is a specialised cell, found in the mammalian Central Nervous System (CNS), and is synonymous to nerve cell. Neurons communicate with each other through the sending of electrical signals, these signals are called **spikes** or equally **action potentials**. When a neuron is sending a spike, we call this **activation**, or that the neuron is **spiking**.

Synapse, synaptic transmission, neurotransmitter, synaptic vesicles, synaptic strength, synaptic plasticity: Neurons are interconnected through **synapses**, which are microscopical connections between neurons through which information is transmitted. The procedure of an electrical signal (spike) being sent from one neuron to another (transmitting information) is called the **synaptic transmission**. Each synaptic transmission is facilitated by a **neurotransmitter** (in this thesis, most commonly Glutamate). In the neuron, it is the **synaptic vesicles** which contain the neurotransmitter. The strength, or efficacy, for which the synaptic transmission can occur between two neurons is referred to as the **synaptic strength**. The phenomenon of this synaptic strength between neurons adjusting in regards to activation is called **synaptic plasticity**.

Presynaptic neuron, postsynaptic neuron, excitatory neuron, inhibitory neuron, input: By convention, the neuron sending the spike is called the **presynaptic neuron**, and the neuron receiving the spike is called the **postsynaptic neuron**. However, in the model implemented in this thesis, it is not of importance whether it is the presynaptic or postsynaptic neuron that is activated first. An **excitatory neuron** is a neuron which sends spikes (excitatory spikes) which make the receiving neuron more likely to spike (pushing its membrane potential up). Reversely, **inhibitory neurons** send spikes (inhibitory spikes) which make the receiving neuron less likely to spike (dragging its membrane potential down). Although only excitatory neurons are used in this model, inhibitory spiking may still occur (whenever $w_{ij} < 0$), this is due to something called disinaptic inhibition. Incoming spikes will sometimes be referred to as **input**, and can be either excitatory or inhibitory.

Neuroplasticity, cortex, neocortex: The multiple different types of plasticity of the brain (structural plasticity, intrinsic plasticity, synaptic plasticity, and others) fall under the umbrella term of **neuroplasticity**. The **cortex**, short for cerebral cortex, is the outer layer of the brain. It governs high-level processing, voluntary movements,

hearing and the processing of visual information. The *neocortex* is the outermost part of the cerebral cortex, and is responsible for cognitive functions, including problem solving, reasoning, learning and memory.

Ion, Ca^{2+} , membrane, electrical current, conductances: An *ion* is an electrically charged atom (or molecule), with either a positive charge (surplus of protons) or negative charge (surplus of electrons). Ca^{2+} are calcium ions, and are vital in the calcium-governed synaptic transmission modelled in this thesis. The *membrane* of a neuron separates the internal environment of the neuron from the external environment. At rest, the neuron's membrane is at -70mV . A neuron's membrane is pushed up, towards the firing threshold of -55mV , whenever it receives an increase in *electrical current*, which is a flow of electrical charge generated from incoming spikes. The neuron's membrane potential is affected by *conductances*, which either pushes the potential up towards the firing threshold (when receiving excitatory input), or drags it down towards its resting potential (when receiving inhibitory input, or when at rest).

Hebbian learning, Hebbian, Hebb's postulate: The theory that simultaneous (or, close in time) activation leads to an increase in synaptic efficacy, or strength, is called *Hebbian learning*. Models and theories which follow this idea are commonly referred to as *Hebbian*. These theories are represented by *Hebb's postulate*, which is famously paraphrased as "*Neurons that fire together, wire together*". In this thesis, Hebbian synaptic plasticity is studied.

Artificial Neural Networks (ANNs), Spiking Neural Networks (SNNs), hypercolumns, minicolumns: From the field of Machine Learning comes the model called *Artificial Neural Networks (ANNs)*. These are inspired by the structure of mammalian neocortex, and the most biologically detailed subtype of ANNs is called *Spiking Neural Networks (SNNs)*. In a SNN, the strength of connection between neurons i, j is determined by the weight w_{ij} , and neurons communicate via sending spikes. A Spiking Neural Network is in a customary network implementation organised into *hypercolumns* and *minicolumns*, where each hypercolumn consists of a number of minicolumns.

Correlation, anti-correlation, muted: Synchronous, or near-synchronous, spiking of two neurons (one presynaptic, one postsynaptic) is called *correlation*. Reversely, non-synchronous spiking is called *anti-correlation*. A neuron that is inactive, meaning that it is not spiking and not currently receiving any excitatory input, is called *muted*.

Learning rule, Bayesian Confidence Propagation Neural Network (BCPNN), training data, spiking data, spike trains, local synaptic traces, Exponentially Weighted Moving Averages (EWMAs), weights: In a SNN, a *learning rule* is a mathematical formula which adjusts the strength of connection between neurons based on the timing of their spikes, and how this timing correlates to one another. In our model, the chosen learning rule is the *Bayesian Confidence Propagation*

Neural Network (BCPNN), which estimates the likelihood of two neurons firing together based on the number of occurrences of this in the **training data**, meaning the data of spiking which the model learns from. This data is also sometimes called **spiking data**, and comes in the form of **spike trains**, which are binary lists of zeroes (for a non-spike) and ones (for a spike). With this data as input, **local synaptic traces**, an adapted version of **Exponentially Weighted Moving Averages (EW-MAs)**, provide a low-pass filter and exponential integration over biologically inspired time scales. The output of this procedure is used to incrementally update the **weight**, the measure of synaptic strength between two neurons, at any given time.

Long Term Potentiation (LTP), Long Term Depression (LTD): A prolonged time of correlated activity between two neurons leads to an increase in synaptic strength (meaning, the weight), this is called **Long Term Potentiation (LTP)**. Reversely, the opposite phenomenon, where a prolonged time of anti-correlated activity leads to a decrease in synaptic strength, is called **Long Term Depression**. These are fundamental phenomena of synaptic plasticity, and seem to play important roles in learning and memory, among other things.

1 Introduction

1.1 Background

"When an axon of cell A is near enough to excite cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased."

Donald O. Hebb, 1949 (*The Organization of Behavior*) [1]

Have you ever wondered about the way in which the brain learns and adapts? How habits, memories and knowledge form and persist (or not) throughout time? This is, of course, interesting to reflect about on a personal level; to reflect on ones mind and personality and all the parts which they consist of. But ever so more intriguing, and far less well known still to this day, is what these processes resemble on a neuronal level.

Neurons, also known as nerve cells, are one of the two fundamental types of cells within the nervous system and can be described as the basic unit of the nervous system [2]. Research suggests that there are approximately 86 billion neurons in the average adult human brain [3] and that neurons play a crucial role in the majority of the unique functionality of the brain [4]. Neurons transmit information to other neurons through sending electrical signals or impulses [2], known as a *spike* or an *action potential*. The information is transported via the *synapse*, which can in an easy manner be described as the physical connection between two neurons. Neurons are interconnected in large networks in the Central Nervous System (CNS), which governs cognitive functions, motor control and the processing of sensory input.

When a neuron sends a spike, we say that the neuron is *spiking*. When one neuron, call it neuron *A*, is spiking, a chemical substance known as *neurotransmitter* is released from the end of neuron *A* (from its axon terminal) and travels via the synapse to the receiving neuron *B*, which electrically impacts the receiving neuron *B* [2]. Once the neuron *B* has been impacted by enough incoming spikes, this will cause it to emit a spike. We say that neuron *A* takes part in firing neuron *B*. This whole process is known as a *synaptic transmission*, and will be further explained in section 1.4. When we are talking about *synaptic strength*, a simple explanation is that it is the measure of the connection between neurons *A* and *B*, meaning the strength (or, equally, efficacy) of their synaptic transmission.

In his 1949 book *The Organization of Behavior*, Canadian psychologist Donald O. Hebb postulated what this interplay of neurons may look like. He introduced the theory that simultaneous *activation* of two neurons would lead to an increase in synaptic strength [1]. Meaning that, whenever neuron *A* spikes (activation) and shortly thereafter neuron *B* also spikes (due to the influence of other spiking neurons), the synaptic strength between neurons *A* and *B* is elevated [4]. This is commonly referred to as *Hebb's postulate* and is best represented by his above quote, but more famously paraphrased as "*Neurons that fire together, wire together*". Models or concepts aligning with this theory are referred to as *Hebbian learning*, or simply *Hebbian*.

To properly motivate why the concept of synaptic strength, i.e. the strength (or efficacy) of synaptic transmissions, is of great interest in this area of research, we must first introduce the concept of *synaptic plasticity*. In simple terms, synaptic plasticity refers to the modification (strengthening or weakening) of synaptic strength in regards to activity [5]. This particular kind of *neuroplasticity* (plasticity of the brain) has long intrigued neuroscientists, and will be the main topic of this thesis. Synaptic plasticity is of interest, since it seems to play a vital role in memory formation, learning, and several other cognitive functions.

When discussing synaptic plasticity, the distinction is often made between a *presynaptic neuron* and a *postsynaptic neuron*. By convention, the neuron *A* projecting the spike, and thus initiating the synaptic transmission, is called the presynaptic neuron, and the receiving neuron *B*, which is under the influence of (usually multiple) incoming spikes is called the postsynaptic neuron. A single neuron can receive thousands of spikes from many different neurons.

1.2 Cognition & neural networks

Every one of us uses our cognition every single day. It is what lets us think, reflect, learn, store and retrieve memories, analyse our surroundings and perceive sensory input. This encompasses mental processes such as reasoning, decision-making but also things such as knowing how to speak languages, recognising our neighbours and remembering which foods we like. Our cognition is formed in large, complex networks, known as neural networks. These are located in and between different regions of the mammalian brain, in what is known as the Central Nervous System (CNS, previously introduced). In this thesis, we are specifically studying neurons from the neocortex. The cortex, short for cerebral cortex, is the outer part of the mammalian brain, and the neocortex is its largest and most important component. The neocortex contains the parietal, temporal, occipital and frontal lobes, and is what governs all higher cognitive functions in humans, such as reasoning, language and problem solving. These neural networks consist of billions of neurons, with millions of connections between them. Every such connection is facilitated by a synapse, which is why synaptic plasticity refers to the ability of these connections to adapt in regards to activation.

1.3 Aim of Computational Neuroscience

Neuroscience is the scientific study of the organisation, functions and disorders of the mammalian nervous system. The aim is to understand complex mental and cognitive processes such as learning, reasoning, perception, memory and language in terms of the features (properties) of neural circuits and its components [6]. The sub-field of Computational Neuroscience (also sometimes called Theoretical Neuroscience) approaches this goal by means of constructing computational and mathematical models of the nervous system, at different levels of abstraction [6]. The ultimate goal is to, with the help of the computational models, reduce the mental phenomenon of cognition to properties of the nervous system [6], and measure these at a reasonably biologically plausible degree while maintaining computational feasibility [7].

1.4 Neurophysiology & synaptic transmission

Neurons consist of three main parts — the dendrites, the soma and the axon [4]. Figures 1.1 and 1.2 depicts the synapse between two neurons. The synapse consists of the axon terminal (the outermost part of the axon) of the presynaptic neuron, the opposite end of the postsynaptic neuron (which, in the case of neurons, is the dendrites) and the synaptic cleft separating the two. This full "area of contact" between the two neurons makes up what is referred to as the synapse, whereas the synaptic cleft specifically refers to only the small fluid-filled space between the two neurons [8]. The synapse is the physical junction between the axon terminal of one neuron and the dendrites of another [4]. Each synapse is highly specialised and it is along this microscopic physical distance that the synaptic transmission between two neurons occur.

Figure 1.1 shows the synapse before the arrival of the action potential. The light blue clusters in the axon terminal depicts *synaptic vesicles*, bound to the membrane [8]. Synaptic vesicles are filled with neurotransmitter molecules. Figure 1.2 shows the synapse once that the action potential has arrived. When an action potential arrives at the axon terminal, the membrane is depolarised (1) (made less negative). On its own, the membrane potential is at rest at around -70 mV, and so, when the (positively) electrically charged action potential arrives, it is increased. Because the membrane potential is depolarised, voltage-gated Ca^{2+} channels (see Figure 1.1, light green squares) in the membrane are activated and Ca^{2+} flows into the neuron (2) [8]. This results in (3), that the synaptic vesicles (light blue clusters) release neurotransmitter. Now, the neurotransmitter flows into the synaptic cleft and its molecules bind to receptor proteins on the postsynaptic neuron [8], on the dendrites more specifically.




Figure 1.1: Illustration of the synapse, before the arrival of an action potential




Figure 1.2: Illustration of the synapse, with an action potential

The Synapse. © Khan Academy. Licensed under CC BY-NC-SA 4.0. No changes made. Source: <https://www.khanacademy.org/science/biology/human-biology/neuron-nervous-system/a/the-synapse>. Note: All Khan Academy content is available for free at <http://www.khanacademy.org/>.

Simplified illustrations of the synaptic transmission between the axon terminal of a presynaptic neuron and the dendrites of a postsynaptic neuron.

The activation of the receptor proteins results in ion channels in the cell membrane opening and closing (4) and an inflow of positive ions (in our case, can be negative also in other cases). This generates an increase in current for the postsynaptic neuron and causes the postsynaptic neuron's membrane to depolarise (or, reversely, hyperpolarise if the ions would be of negative charge).

1.5 Synaptic plasticity

The ability of the mammalian brain to adapt is fundamental to learning and cognition. When we talk about the plasticity of the brain (neuroplasticity), multiple different types of neuroplasticity are usually referred to. Structural neuroplasticity or intrinsic plasticity are two such. Another one is synaptic plasticity, which specifically deals with the plasticity of synapses (synapses, see Section 1.4). This refers to the ability of synapses to strengthen or weaken, as a result of their activity level. Meaning, how strong the connection between two neurons grow, depending on how much they are spiking close in time to one another. Synaptic plasticity was first discovered biologically in 1973, when it was detected in the hippocampus [9, 10]. It was concluded that repeated, near-synchronous activation (co-activation, correlation) of presynaptic neurons and postsynaptic neurons resulted in an increase of synaptic strength [9, 10, 11]. This phenomenon soon became known as Long-term potentiation (LTP). The opposite phenomenon, where the absence of sufficient near-synchronous activation (rendering anti-correlation) leads to a decrease of synaptic strength, is known as Long-term depression (LTD). In more recent years, there has been extensive studies on both LTP

and LTD in a variety of brain areas [12, 13, 14, 10]. LTP and LTD are remarkable features of Hebbian synaptic plasticity, and LTD in particular will be of focus in this thesis. LTP and LTD alike are of interest, since they are hypothesised to impact how the brain learns, remembers and refines connections between neurons.

1.6 Previous work

Donald O. Hebb, whose work was presented in section 1.1, is today widely considered the father of neuropsychology. His ideas of how synaptic connections are strengthened and weakened, and how this relates to learning and memory, are still widely used today, and this thesis specifically investigates Hebbian synaptic plasticity. In Rodriguez et al. [15], characteristic properties of synaptic plasticity are measured in the cortex of mice. These biological recordings will be used as an important reference in this thesis. Today, much of Computational Neuroscience focuses on constructing probabilistic models of spiking behaviours in networks of neurons and biologically realistic models of neurons and their synaptic transmissions. The main goal is to compute how the spiking behaviour correlates to different synaptic plasticity characteristics. Popular learning rules include BCPNN, BCM and STDP. This thesis employs the Bayesian Confidence Propagation Neural Network (BCPNN) learning rule [16, 17], which will be introduced more extensively and explained further in sections 1.8, 2.1, 2.2. BCPNN has previously been used to model vision [18] and memory [19, 20, 21, 22], among other things. In particular, the spike-based extension of BCPNN from Tully et al. [23] is being used, which will be defined in section 2.1. Since one of the main goals of this thesis is to reproduce several results from Tully et al. [23], the method and theory used are largely based on the works of said paper's authors.

1.7 Spiking Neural Networks

When computing synaptic plasticity in Computational Neuroscience, a Spiking Neural Network (SNN) model is simulated. This is a subtype of the famous Machine Learning model Artificial Neural Network (ANN), and holds a higher biological resemblance than traditional ANN models. SNN models are modelled from the biological processes of the mammalian Central Nervous System (CNS), in particular by how neurons communicate by means of sending spikes [24]. There are multiple valid choices of SNN models for simulating synaptic plasticity, and generally speaking a model is chosen based on the trade-off between biological plausibility and computational feasibility [25].

In this thesis, a SNN is made up of a neuron model, a bayesian learning rule (the aforementioned BCPNN learning rule) and input poisson processes of 30 processes per neuron. For the neuron model, a version of the Integrate-and-Fire Neuron model is used. This model was first developed by French scientist Louis Lapicque in 1907 [26]

in his celebrated paper *Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation*. There are multiple versions of the Integrate-and-Fire neuron model, such as the more biologically detailed Adaptive Exponential Integrate-and-Fire neuron model [27], but for the sake of simplicity and computational feasibility the *Leaky Integrate-and-Fire (LIF)* neuron model will be used in this thesis. In particular, we will be using a slightly adapted model of the one in Tully et al. [23]. The adaptation is due to the fact that this thesis will use a minimal Spiking Neural Network model, meaning a model consisting of a single neuron pair. The pair will be made up of a presynaptic neuron i , and a postsynaptic neuron j (*this notation will remain throughout this thesis*). Our Leaky Integrate-and-Fire neuron model will be somewhat simplified to adapt to this minimal network.

The Leaky Integrate-and-Fire neuron model governs for which conditions the neurons will communicate with each other, meaning for which conditions they will send spikes to one another. As was explained in section 1.4, when a neuron will send a spike is decided by the voltage of the membrane potential. The main component of the Leaky Integrate-and-Fire neuron model, for this reason, is a differential equation which describes the rate of change of the neuronal membrane. This equation will be described in section 2.3 and takes into account excitatory and inhibitory input, either dragging the membrane voltage up (towards firing threshold) or down (towards resting potential), and has a leaky component which also drags the membrane voltage down whenever no recent input has been received. The theory is largely based on what has already been covered regarding the physiology of neurons and synapses in section 1.4. A key feature of the Leaky Integrate-and-Fire neuron model is that whenever the membrane potential V_m reaches (or surpasses) the threshold V_{th} (such that $V_m \geq V_{th}$) and a spike is emitted, the membrane potential is afterwards held to a reset potential for a given amount of time. The time is called the absolute refractory period and is represented by t_{ref} [23].

1.8 Bayesian computations

In addition to our neuron model (the Leaky Integrate-and-Fire neuron model, see section 1.7), the spike-based extension of the BCPNN learning rule from Tully et al. [23] (introduced in section 1.6) is used in this Spiking Neural Network (SNN) model. BCPNN is a learning rule, and is heuristically derived from Bayes' rule [28]. It computes posterior probabilities of spikes occurring, given already observed spiking events [23]. The formulas for BCPNN will be described in section 2.1. For now, it is sufficient to describe the relevant concepts. N.B. that this thesis implements a special case of this rule, since only a single neuron pair is simulated (see section 1.9), and not a whole network as is otherwise customary. The single neuron pair consists of the presynaptic neuron i and postsynaptic neuron j , as previously mentioned. However, in order to properly explain the model and the derivations of its formulas, it is necessary to explain the theory in terms of the customary network implementation. So, for this reason, the hereunder theory is explained in terms of this aforementioned customary

network implementation of spike-based BCPNN from Tully et al. [23]:

In BCPNN, the structural organisation used is inspired from that of mammalian neo-cortex, and groups neurons into minicolumns and hypercolumns. A unit is represented by a singular minicolumn, where each minicolumn models 30 highly connected and co-active spiking neurons. Each hypercolumn then consists of a number of minicolumns. The level of activation of a minicolumn is defined as the minicolumn’s average spike rate [23]. In a SNN, *weights* represent the strength of connection between two units (or, or in this thesis, between two neurons) — we will hereafter refer to this as a measure of synaptic strength. According to BCPNN, weights are based on the likelihood of units firing together (simultaneously, or, close in time), which is estimated based on the number of co-occurrences in the *training data* [28]. This means that, the strength of connection is determined by the estimated probability of simultaneous spiking. The model uses incrementally updated weights in a procedure bearing resemblance to proposed rules for biological synaptic plasticity [28, 29, 30]. As previously mentioned, the customary case of spike-based BCPNN computes strength of connection between different minicolumns, while our adapted model computes strength of connection between neurons i and j . A simple, but non-redundant, explanation of spike-based BCPNN is that for simultaneously active (*correlated*, or equally, co-active) units, the strength of connection is strengthened, while for *anti-correlated* or *muted* units, the strength of connection is decreased [28]. Correlated units are simultaneously active units, anti-correlated units are units that are not simultaneously active, meaning that they are not firing close enough in time, muted units are units that are not active (not firing).

The training data, or equally the *spiking data*, represents the observed spiking events, and is in the form of input *spike trains*. The spike trains are binary lists, where each non-spike is registered as a zero and each spike registered as a one. In our model, S_i represents the spike train of neuron i , and S_j represents the spike train of neuron j . The indexing of the lists is given by the time resolution (set to 0.1 ms). With spike trains of a given length as input, consisting of non-spikes and spikes, the probabilities of new spikes occurring is computed through a three-step procedure. The procedure is called *local synaptic traces*, and are exponential integrations over recent spiking activity at three different levels and time scales. The final level of the local synaptic traces consist of the P -traces, also known as the probability traces. It is for the sake of reaching these probability traces, through which the weight is incrementally computed, that the local synaptic traces are used in BCPNN. Meaning that, in our model, the local synaptic traces estimate probabilities that determine synaptic weights [23]. The probability traces are P_i , P_j and P_{ij} , where P_i is the probability of (presynaptic) neuron i spiking, P_j probability of (postsynaptic) neuron j spiking and P_{ij} probability of neurons i, j spiking simultaneously.

The local synaptic traces are implemented as *Exponentially Weighted Moving Averages (EWMAs)* [31] of the spike trains [23]. This is a filter which prioritises more recent data, and de-prioritised older data exponentially [32]. The integration of spikes is done over different time scales [23], and the time scales are inspired from what is

deemed biologically plausible for Hebbian plasticity [23] (with certain adjustments to our model). The three levels of synaptic traces in the procedure consists of the Z -traces (the *memory traces*), the E -traces (the *eligibility traces*) and the P -traces (the *probability traces*), the two former which will be explained further later. Each local synaptic trace is scaled by a time constant. The time constant determines how rapidly each trace will decay after recent activity. This makes so that the time constants also determine for how long time windows correlated, or anti-correlated, spiking between the neurons can be detected. The time constants are τ_{z_i} , τ_{z_j} , τ_e and τ_p , and will be explained later. The Z -trace time constants (τ_{z_i} , τ_{z_j}) are in this model chosen to match rapid Ca^{2+} influx via NMDA receptors, or voltage-gated Ca^{2+} channels [23, 33, 34].

1.9 Aim of thesis

Broadly speaking, this thesis aims to investigate the concept of synaptic plasticity, meaning the ability of interneural connections in the Central Nervous System (CNS) to adapt and regulate in relation to activity. More specifically, we wish to simulate (and to the extent possible, compute) the behaviour of synaptic strength between a presynaptic neuron and a postsynaptic neuron in mammalian neocortex by doing simulations of a model of this in neural simulator Brian2. Additional emphasis will be posed on the phenomenon of Long Term Depression (LTD; see 1.5), an important feature of synaptic plasticity.

Objectives

The goal is to simulate some of the characteristic behaviours of synaptic plasticity for a single neuron pair (consisting of one presynaptic and one postsynaptic neuron) in neural simulator Brian2. The results from the simulations will be compared to both previous results of computational models (from namely Tully et al.), and to biologically measured results (from namely Rodriguez et al.).

Research questions

From the objectives, two suitable research questions can be formalised:

1. *Can we reproduce some of the main findings from Tully et al. (Figures 2 and 5), for a network of a single neuron pair and in neural simulator Brian2 (instead of NEST), given the same setup and neuron model?*
2. *Can we induce LTD (Long Term Depression), given the same setup and simulator as in (1.), and get results comparable to biological recordings presented in Figure 2E of Rodriguez et al.?*

Anticipated outcome

Certain adaptations for the single neuron pair and different neural simulator will have to be made. We believe that it should be possible to make the model exhibit roughly the same behaviours as in Tully et al., but also expect some parameter adaptations to be needed. It could pose some difficulties comparing our model to biological recordings, which are oftentimes noisy.

Approach, methodology

The general methodology of the project is as follows: A spike-based extension of BCPNN from Tully et al. [23] (see 1.8, 2.1) with a setup (network) consisting of a single neuron pair (one being presynaptic, one being postsynaptic) will be used, together with the Leaky-Integrate-and-Fire neuron model (see 1.7, 2.3) of Tully et al. [23], adapted for our single neuron pair. Parameter changes will be done when necessary and will be explicitly stated. The simulations will be done in neural simulator Brian2 (see 2.4).

Parts 1 and 2 below relates to research questions 1 and 2, respectively.

1. Simulations will be done with the goal of reproducing qualitatively similar results as the ones from Tully et al. of Figures 2 and 5. The respective simulations will be performed in accordance with how they are described in said paper.
2. Simulations will be done with the goal of inducing and displaying Long Term Depression (LTD). One simulation will be done with the goal of reproducing results of biologically measured LTD (from Figure 2E of Rodriguez et al.), and two successive simulations will be done in order to further showcase examples of LTD in our model.

Possible indications

Why are the chosen research questions of interest, and what could they, and the methodology chosen, tell us? This thesis could lead to some interesting indications about, among other things:

How does the spike-based extension of the BCPNN model (see 1.8, 2.1) from Tully et al. hold for a single neuron pair?

How does the Brian2 neural simulator compare to the NEST simulator, for the spike-based extension of BCPNN mentioned above, and for simulations (Figures 2 and 5) from said paper

How does the spike-based extension of the BCPNN model mentioned above hold for simulations of LTD, additional to those already found in said paper?

Restrictions

For the sake of limited time and accounting for the time and memory demands of running simulations of larger networks, certain restrictions had to be imposed which

limit the scope of this thesis. The following restrictions were imposed:

We will only look at a single neuron pair, not a whole network. The neuron pair will consist of only excitatory neurons, no inhibitory neurons. We will only aim to reproduce some of the results of Tully et al.

Limitations

For the results reproducing figures from Tully et al. [23], an apparent limitation is the fact that we do not have access to the exact setup in NEST that was used for the paper, and so can not be sure that our implementation in Brian2 has an exactly corresponding setup. It is very likely that the two neural simulator (NEST and Brian2), will behave different to some degree, and since we are not familiar with NEST, this degree will be difficult to evaluate. A general limitation of the results of this thesis could be said to be the fact that we model specific neurons, and in a specific setting (excitatory, glutamatergic neurons in the neocortex of mammals).

2 Method

This chapter aims to introduce and explain all the relevant mathematical formulas and notation used in this thesis and its simulations. As well the neural simulator and setup used is explained. A table of parameters is included (Table 2.1). The methods used in this thesis are largely based on those already defined and explained in the paper which this thesis aims to reproduce its setup from, meaning Tully et al. [23].

2.1 From Bayes' rule to spike based BCPNN

As previously explained in Section 1.8, the BCPNN learning rule is used. This Bayesian computational rule is implemented in its spike-based form, which was adapted in Tully et al. [23], and is in this thesis adapted to our minimal model consisting of a single neuron pair. The BCPNN learning rule computes the measure of the synaptic strength, meaning the weight w_{ij} . The weight is defined in terms of the probability that both units (in our case; both neurons) will spike simultaneously. The equation of the weight will be adapted to our minimal model. However, all equations will first be explained in terms of the customary network implementation, where hypercolumns and presynaptic and postsynaptic minicolumns are used.

We are interested in how a neural unit can infer from its synaptic input what output it should produce, for instance if a spike should be generated or not in the next time interval. To determine the probability of such an event occurring, a computation is done given already observed events, where the observed events are represented by other active units. In BCPNN, each computational unit represents a stochastic event, and has an associated activation state. The activation state is represented by a value $[0, 1] \subseteq \mathbb{R}$ — this value corresponds to the probability [23]. In our model, this translates to: We are interested in the probability of both neurons spiking close in time to each other, and this probability is computed given already observed cases of simultaneous spiking. The spiking of the neurons is driven by Poisson processes, where, as input, each neuron is given 30 independent Poisson processes (n_{ex}), an input Poisson firing rate of 29 Hz (r_{ex}) and a weight per process of $w_{ex} = 10.75$ nS.

The aim is to compute the measure of the synaptic strength, meaning the weight. In the BCPNN learning rule, this is done by means of transforming input of spiking activity (the spike trains) into local synaptic traces (previously introduced), where the final form of the traces, the probability traces P_i , P_j and P_{ij} , representing estimations of the probabilities of neurons i , neurons j and both neurons simultaneously spiking, respectively, are used to incrementally calculate the weight. The customary

full network implementation of spike-based BCPNN calculates the weight as follows:

$$w_{i,j} = \log \left(\frac{P(x_i|y_j)}{P(x_i)} \right) = \log \left(\frac{P(x_i, y_j)}{P(x_i)P(y_j)} \right) \quad (1)$$

The weight symbolises the synaptic strength of connection between minicolumns x_i and y_j . In Equation 1, in the left-hand side of the equation, the numerator $P(x_i|y_j)$ is the conditional probability of the presynaptic minicolumn x_i being active, given that the postsynaptic minicolumn y_j already is active. The denominator, $P(x_i)$, is the probability of minicolumn x_i being active. This can be rewritten into the right-hand side of the equation, through the equality of $P(x_i|y_j) = P(x_i, y_j)/P(y_j)$, for discrete stochastic variables. On the right-hand side, the numerator $P(x_i, y_j)$ denotes the joint probability of minicolumns x_i and y_j being active at the same time, and the denominator $P(x_i)P(y_j)$ is the product of the probability of minicolumn x_i being active, and the probability of minicolumn y_j being active.

For our minimal model, the calculation is simplified from computing probabilities of minicolumns being active into calculating probabilities of neurons being active. The right-hand side of Equation 1 is rewritten:

$$w_{i,j} = \log \left(\frac{P_{ij}}{P_i P_j} \right) = \log \left(\frac{P_{syn}}{P_i P_j} \right) \quad (2)$$

The weight w_{ij} now symbolises the synaptic strength of connection between neurons i and j . In Equation 2, the estimated probabilities computed through the local synaptic traces are used. The numerator P_{ij} represents the probability of both neurons i and j spiking simultaneously, and the denominator $P_i P_j$ is the product of the two independent probabilities of neuron i spiking, and neuron j spiking, respectively. The P_{ij} -trace denotes simultaneous activity, and is also commonly referred to as the synaptic probability trace. P_{ij} can be rewritten (or, rather, renamed) as $P_{ij} = P_{syn}$, which gives the right-hand side of Equation 2.

The weight in the customary network implementation (Equation 1) is derived from Bayesian computations:

$$P(y_j|x_1, \dots, x_n) = P(y_j) \frac{P(x_1, \dots, x_n|y_j)}{P(x_1, \dots, x_n)} \quad (3)$$

Equation 3 uses Bayes' rule for conditional probabilities of random variables, and calculates the probabilities of neurons in the postsynaptic minicolumn y_j spiking, given already observed spiking rates from presynaptic minicolumns $x_1, \dots, x_n$. Through assuming conditional and unconditional independence for the presynaptic minicolumns $x_1, \dots, x_n$, this can be rewritten:

$$P(y_j|x_1, ..., x_n) = P(y_j) \frac{P(x_1|y_j)}{P(x_1)} \frac{P(x_2|y_j)}{P(x_2)} \dots \frac{P(x_n|y_j)}{P(x_n)} \quad (4)$$

The probabilities are recovered by normalising $P(y_j|x_1, ..., x_n)$ to sum to 1. From Equation 4, each x_h is defined as a discrete coded variable (such that, for example, x_h = "colour" and x_{h_i} = "pink", or "green"), and n_h minicolumns are seen as distributed into H hypercolumns:

$$P(y_j|x_1, ..., x_n) = P(y_j) \prod_{h=1}^H \sum_{i=1}^{n_h} \frac{P(x_{hi}|y_j)}{P(x_{hi})} \pi_{x_{hi}} \quad (5)$$

The term $\pi_{x_{hi}}$ represents relative activity (or uncertainty), of the attribute value x_{hi} , where $\pi_{x_{hi}} = 1$ indicates x_{hi} being observed with maximal certainty, or maximal activity. For computational purposes, Equation 5 is transformed into an equivalent expression of sums of logarithms:

$$\log(P(y_j|x_1, ..., x_n)) = \log(P(y_j)) + \sum_{h=1}^H \log \left[\sum_{i=1}^{n_h} \frac{P(x_{hi}|y_j)}{P(x_{hi})} \pi_{x_{hi}} \right] \quad (6)$$

The logarithms in Equation 6 refers to the natural logarithm $\ln$. Every hypercolumn is regarded as having normalised activity, such that $\sum_{i=1}^{n_h} \pi_{x_{hi}} = 1$. The activity statistics are gathered during learning, and the weight is incrementally calculated and updated. The corresponding weight w_{ij} (see Equation 1) is derived from Equation 6.

2.2 Local synaptic traces

As presented in Section 1.8, local synaptic traces are used to estimate probabilities P_i , P_j and P_{ij} , that determines the weight w_{ij} (in accordance with Equation 2). The weight is incrementally updated throughout the training process. Each trace is defined by a differential equation, describing their respective rate of change. The procedure of the local synaptic traces starts with the Z -traces: Z_i , filtering presynaptic spiking, and Z_j , filtering postsynaptic spiking:

$$\tau_{z_i} \frac{dZ_i}{dt} = \frac{S_i}{f_{max} t_{spike}} - Z_i + \varepsilon \quad (7)$$

$$\tau_{z_j} \frac{dZ_j}{dt} = \frac{S_j}{f_{max} t_{spike}} - Z_j + \varepsilon \quad (8)$$

The Z -traces filter spiking activity with time constants $\tau_{z_i}, \tau_{z_j} = 5$ ms. The time constants determine the time period for which correlated (or anti-correlated) spiking can be detected. For this reason, they are called memory traces. Equation 7 describes the Z_i -trace and Equation 8 describes the Z_j -trace. With their respective spike trains, S_i and S_j , as input, the Z -traces rise at each spike event to an amount that gets scaled by two constants: f_{max} and t_{spike} . These are, respectively, the maximal expected firing frequency (set to 30 Hz) and the time duration of a single spike (set to 0.1 ms). The smallest possible value that the traces can take on is decided by ε , which in our model is set to 0.0033. After each spike, the traces decay in accordance to their respective time constant. This models the leaky component, where older activity is deprioritised and exponentially forgotten by the system. When the respective presynaptic neuron and postsynaptic neuron fires at f_{max} , the mean over time converges towards 1. The next step in the procedure is the E -traces, also called eligibility traces. The Z -traces are fed into the E -traces:

$$\tau_e \frac{dE_i}{dt} = Z_i - E_i \quad (9)$$

$$\tau_e \frac{dE_j}{dt} = Z_j - E_j, \quad (10)$$

$$\tau_e \frac{dE_{ij}}{dt} = Z_i Z_j - E_{ij} \quad (11)$$

The traces E_i, E_j and E_{ij} are governed by a common time constant $\tau_e = 100$ ms. In every step of the procedure, a low-pass filter on the activity is provided, and so here this is done on the activity of the Z -traces. Each step provides exponential integration of activity, which acts as temporal smoothing thanks to the prolonged time constant that the traces in each subsequent step has. For the E -traces, defined by Equations 9, 10 and 11, each trace relaxes towards their respective driving signal, meaning the corresponding Z -trace(s) governing their behaviour, as can be seen by the Z_i, Z_j or $Z_i Z_j$ terms in each respective equation. Each E -trace mirrors its corresponding Z -trace(s) in an exponential manner, but with a considerable delay. The idea of this delay is to represent delayed learning that is hypothesised to happen in the brain. The E_{ij} -trace (Equation 11) tracks recent activity of $Z_i Z_j$, the product of the two respective Z -traces. This is called a synaptic trace, and is used to track correlated activity between the two neurons. The third and final step in the procedure is to translate the E -traces into the P -traces:

$$\tau_p \frac{dP_j}{dt} = \kappa(E_j - P_j) \quad (12)$$

$$\tau_p \frac{dP_i}{dt} = \kappa(E_i - P_i) \quad (13)$$

$$\tau_p \frac{dP_{ij}}{dt} = \kappa(E_{ij} - P_{ij}) \quad (14)$$

The P -traces are called probability traces. These traces, P_i, P_j and P_{ij} , represents the probabilities of neuron i , neuron j and neurons i, j simultaneously spiking. They are defined by Equations 12, 13 and 14, respectively. The P -traces provide a low-pass filter on the activity of the E -traces, and exponentially mirrors their behaviour with a substantial delay (specifically, by time constant $\tau_p = 1\,000$ ms). Furthermore, the equations (12, 13 and 14) are scaled by the parameter κ . This parameter is of note, since it moderates the plasticity during the training process. The plasticity parameter κ can take on any integer between zero and one, but is in normal cases set to $\kappa = 1$ (full plasticity).

2.3 Leaky Integrate And Fire Neuron Model

All parameters described in this section can be found in Table 2.1, along with their respective units and standard values. The parameters can be scaled to represent different types of synapses, among other things. However, apart from when else is stated (in certain simulations, specifically), the reader should assume all parameters to be of their standard values. As previously introduced, an adapted version of the Leaky Integrate-and-Fire neuron model used in Tully et al. [23] is used. The neuron model is represented by a differential equation describing the rate of change of the membrane potential V_m , since this governs the spiking activity of the neurons. As explained in 1.4, the membrane potential is the electrical voltage difference across the neuron's membrane, which is due to uneven distributions of ions outside of and inside the neuron. The membrane capacitance C_m measures the neuron's ability to store electrical charges. Receiving currents over the membrane capacitance, generated by incoming spikes, affects the membrane voltage [35]. The equation governing the rate of change of the neuronal membrane consists of three main parts:

$$g_L(V_m - E_L),$$

This first part is called the leak current and governs the membrane potential whenever the neuron is at rest. It pulls the membrane potential back toward the leak reversal potential E_L (the resting potential, -70mV), whenever no recent input has been received, since this makes the neuron less likely to fire. The strength is scaled by the leak conductance g_L (set to 16.67 nS), which describes the permeability of the membrane to let ions through when the neuron is at rest (whenever no recent input has been received).

$$\sum_{i=1}^n g_{ex,i}(V_m - E_{ex}),$$

The second part describes the difference between the membrane potential V_m and the excitatory reversal potential E_{ex} (set to 0 mV). When the neuron receives excitatory

input (spikes), it makes the neuron more likely to fire. In the equation, for every excitatory input i that is received, the membrane potential is driven up, towards the firing threshold ($V_{th} = -55$ mV). Each i contributes a current proportional to its conductance $g_{ex,i}$, which describes the permeability of the membrane to let ions through when the neuron receives excitatory input.

$$\sum_{i=1}^n g_{inh,i}(V_m - E_{inh}),$$

The third part is analogous to the second part, only for inhibitory input (meaning, for spikes that make the receiving neuron less likely to fire). Each inhibitory input i drags the membrane potential down, towards the inhibitory reversal potential E_{inh} (set to -75 mV). This is scaled by the inhibitory conductance $g_{inh,i}$, which describes the permeability of the membrane to let ions through when the neuron receives inhibitory input. These three parts are scaled by the membrane capacitance C_m , and this together describes the rate of change across the membrane potential V_m :

$$-C_m \frac{dV_m}{dt} = g_L(V_m - E_L) + \sum_{i=1}^n g_{ex,i}(V_m - E_{ex}) + \sum_{i=1}^n g_{inh,i}(V_m - E_{inh}). \quad (15)$$

Equation 15 serves as a modelling of biological neurons behaviours' within a neural network, and is implemented in its entirety in the neural simulator. The two latter conductances, g_{ex} and g_{inh} , are given by something called an *alpha-shaped PSC dynamics synapse model*, and are dependent on the timing of the receiving (excitatory/inhibitory) input, meaning they have no standard values.

2.4 The Brian2 neural simulator

In order to model our neuron pair and synapse, we are using the Brian2 neural simulator [36]. Brian2 is a library for Python and is a free, open source simulator for spiking neural networks available on almost all platforms [37]. It was developed by Marcel Stimberg, Romain Brette and Dan FM Goodman in 2019 [36]. The previously mentioned authors (Stimberg et al.) describe in their paper what they believe to be the largest problem among neural simulators: Either they are (1) able only to simulate a small range of models, but can do so very well, or they are (2) able to simulate a wide range of models, but suffer from either technical complexity, or poor performance [36]. Brian2's answer to this is letting the user express any given model in a few lines of high-level descriptions (maintaining high readability, low technical complexity), which in turn the simulator transforms into efficient low-level code (maintaining high performance) [36]. The Brian2 neural simulator supports multi-line strings of equations and equations defined as differential equations. All of the synapse equations and neuron equations are in the simulator defined as in their equations above.

Table 2.1: Model parameters

Neuron model			
$C_m = 250$ pF	Membrane capacitance	$g_L = 16.67$ nS	Leak conductance
$E_L = -70$ mV	Leak reversal potential	$E_{ex} = 0$ mV	Excitatory reversal potential
$E_{inh} = -75$ mV	Inhibitory reversal potential	$dt = 0.1$ ms	Time resolution
$V_{th} = -55$ mV	Voltage threshold	$V_{res} = -60$ mV	Reset potential
$t_{ref} = 2$ ms	Refractory period		
Channel model			
$\tau_{z_i} = 5$ ms	Z_i trace time constant	$f_{max} = 30$ Hz	Maximal firing rate
$\tau_{z_j} = 5$ ms	Z_j trace time constant	$f_{min} = 1$ Hz	Minimal firing rate
$\tau_e = 100$ ms	E trace time constant	$\varepsilon = 0.0033$	Baseline firing rate
$\tau_p = 1000$ ms	P trace time constant	$t_{spike} = 0.1$ ms	Spike duration
Synapse model			
$g_{max} = 2$ nS	Peak conductance	$d = 0.1$ ms	Transmission delay
$\tau_{ex} = 0.2$ ms	Rise time, excitatory input	$\kappa = 1$	Learning rate
$\tau_{inh} = 2$ ms	Rise time, inhibitory input	$g_{ex,i} : S$	Excitatory conductance
		$g_{inh,i} : S$	Inhibitory conductance
Input			
$n_{ex} = 30$	Poisson processes per neuron	$r_{ex} = 29$ Hz	Input firing rate, $0 < r_{ex} < f_{max}$
$w_{ex} = 10.75$ nS	Weight per process	$\tau_{input} = 0.7$ ms	Input EPSP time constant

Comprehensive table of parameters, mostly taken from Tully et al. [23] but with alterations of namely ε and the time constants of Z and P to accommodate for our choice of neural simulator (Brian2) and adapted (minimal) model. Certain simulations in this thesis use altered versions of parameters, this is then explicitly stated. The following intervals are used: $\tau_{z_i}, \tau_{z_j} : 2 - 10$ ms, $\tau_p : 1000 - 10000$ ms, $\kappa : 0.065 - 1$.

3 Results

3.1 Simulations from Tully et al.

As this thesis aims to reproduce some of the behavioural characteristics of simulations of spike-based BCPNN, for a single neuron pair, demonstrated in Tully et al. [23], and uses a close to identical setup and model to said paper, a substantial part of the simulations are reproductions of some of the first few figures found in Tully et al. [23]. In general, the results were found to correspond well to one another. Due to the different choice of neural simulator and a considerably smaller network (with the single neuron pair), several parameters had to be adjusted. Apart from where else is stated, the spiking activity is driven by Poisson processes (of 30 processes per neuron (n_{ex}), an input firing rate of 29 Hz (r_{ex}), and a weight per process of 10.75 nS (w_{ex}).

3.1.1 Local synaptic traces

As a first order of demonstration, we want to model the local synaptic traces from section 2.2. Crucially, this enables us to see that our model behaves as we expect it to, and that its local synaptic traces (Z -, E - and P -traces) correspond well to those in Figure 2 of Tully et al. [23]. In the interest of increasing clarity, the resulting change of weight w_{ij} is added as the bottom plot, as defined by Equation 2. In Figure 3.1, the three step EWMA (Exponentially Weighted Moving Average) procedure of the local synaptic traces is displayed. It illustrates how the input of two incoming trains of spikes take shape in the Z -traces, evolve through a low-pass filter to the E -traces and then a final time on to the P -traces. The P -traces representing the probabilities P_i of presynaptic spiking (red), P_j of postsynaptic spiking (blue) and P_{ij} of simultaneous pre- and postsynaptic spiking (black). The time constants ($\tau_{zi}, \tau_{zj}, \tau_e, \tau_p$) governing these traces, are chosen with essential plasticity-relevant changes that occur during learning in mind [23]. The two incoming spike trains are S_i (in red) and S_j (in blue), where S_i represents the spiking activity of the presynaptic neuron i and S_j represents the spiking activity of postsynaptic neuron j . The simulation is divided in five bins (of 100 ms each), marked by ticks on the x -axes. First, only the presynaptic neuron is stimulated. Then, only the postsynaptic neuron is stimulated. In the third and fifth bins, none are stimulated (both muted). In the fourth bin both neurons are stimulated (co-activation). They are stimulated as usual with Poisson processes. The spiking activity translates to the Z -traces through Equations 7, 8, where it is governed by the time constants τ_{zi}, τ_{zj} and scaled by the highest expected firing frequency f_{max} multiplied by the time duration of a single spike t_{spike} . Essentially, the time constants τ_{zi}, τ_{zj} determine for how long time periods co-activation between the neurons can be

detected. The traces enable to translate the spiking rates of neurons into probability values. The maximal probability is reached whenever the given neuron's firing rate is equal to f_{max} , and is equal to $\varepsilon = 0.0033$ whenever the firing rate is equal to f_{min} . This evolves into the E -traces, through Equations 9, 10 and 11, with a low-pass filter and higher time constant $\tau_e = 100$ ms (longer memory). Further to the P -traces, who have the highest time constant of them all (slower forgetting) and the slowest dynamics, given by $\tau_p = 1\,000$ ms. They are, as per Equations 12, 13 and 14, also scaled by κ : the parameter of plasticity, in standard cases set to 1 (full plasticity).




Figure 3.1: Local synaptic traces

At the top of the graph, a spike raster displays the input trains of spikes S_i and S_j . Five bins are simulated: First bin (0-100 ms) with (red) presynaptic spiking, second bin (100-200 ms) with (blue) postsynaptic spiking, third bin (200-300 ms) with no spikes, fourth bin (300-400 ms) with co-activation and fifth bin (400-500 ms) with no spikes. The traces decay according to their respective time constants: $\tau_{z_i} = \tau_{z_j} = 5$ ms, $\tau_e = 100$ ms and $\tau_p = 1\,000$ ms. The traces rise for each indexed spike; Z_i (red) for each presynaptic spike and Z_j (blue) for each postsynaptic spike. Red traces represent presynaptic spiking, blue traces represent postsynaptic spiking and black traces represent synaptic simultaneous spiking. The Z -traces are scaled by the maximum expected firing frequency f_{max} times the spike timing duration t_{spike} , and the P -traces are scaled by the plasticity parameter κ .

For each spike, the respective Z -trace increases by approximately 3 to 5 (Figure 3.1). Without any recent spiking activity, the traces decay to their minimum value ε . In this plot, the highest registered value for the Z -traces is approximately 7.5 for the

presynaptic trace (red, in the fourth bin) and 6 for the postsynaptic trace (blue, in the second bin). Following a low-pass filter, the E_i - (presynaptic) and E_j - (postsynaptic) traces take on values of approximately ε to 2.5 (at around 410 ms). The synaptic E_{syn} -trace, illustrating co-activation, grows to over 5 after repeated co-activation in the fourth bin. The higher time constant of $\tau_e = 100$ ms enables a longer period of detecting co-activation (in bins second to fifth) which allows for the high synaptic E_{syn} -trace. At the P -traces, representing probabilities and thus taking on values $\subseteq [0, 1]$, the time constant $\tau_p = 1\,000$ leads to a much slower decay of the traces compared to in Z and E . Presynaptic spikes (first bin) closely followed by postsynaptic spikes (second bin) leads to a steady increase in weight in the second bin. When the postsynaptic spikes (second bin) are not followed by any presynaptic spikes in neither the second nor third bin, the weight decreases and becomes negative, reaching its lowest weight of -0.5 slightly after the 300 ms mark. Repeated and close co-activation in the fourth bin again steadily increases the weight, until it reaches its highest value of around 1.3 around the second half of the fifth bin (see Figure 3.1).

The same patterns (bins) as in Figure 2 of Tully et al. [23] are simulated, but with remarkably less intense spiking activity. The amount of spikes in our figure is the result of the input Poisson processes of 30 processes per neuron. Important to point out that our model's parameters were tuned in order to achieve the most corresponding results to the ones in Tully et al [23]. This meant that τ_{z_i} and τ_{z_j} were lowered from 10 ms to 5 ms, τ_p lowered from 10 000 ms to 1 000 and ε set to 0.0033. Both Z -traces vary within the approximately same interval - meaning between ε and slightly above 6. The presynaptic trace however is slightly higher in our figure (with a maximum of approximately 7.5 compared to 6). The E -traces were found to correspond well, both in numerical values and behaviour. The notable exception is the less smooth and slightly higher E_{syn} -trace in our figure, of approximately 5.5 compared to 4.8. The P -traces were generally found to correspond well, importantly exhibiting the same behaviour. The overlap between the presynaptic (red) and postsynaptic (blue) traces was slightly delayed in time, occurring around the 350 ms mark in our figure compared to around the 300 ms mark in Tully et al.'s figure. The synaptic P_{syn} -trace (black) reaches a slightly higher value in our figure, approximately 0.6 compared to 0.55 in the figure from Tully et al. [23]. Figure 3.1 illustrates how the synaptic traces E_{syn} and P_{syn} are only increased in the fifth bin, where close correlation of spikes occur. The presynaptic spikes in the first bin followed by postsynaptic spikes in the second bin, which leads to a steady increase in weight, is not enough to alter the synaptic traces at all. Also only the first few postsynaptic spikes in the second bin increase the resulting weight, for the later postsynaptic spikes too much time has elapsed since the presynaptic spikes, and the weight is instead decreased.

3.1.2 Dependency of spike timing interval

For the second simulation, we aim to investigate the model's plasticity properties for different spike timing intervals. For spike-based BCPNN, spiking of the presynaptic neuron followed by spiking of the postsynaptic neuron (referred to as *pre-before-post*

spiking) within a certain time window should lead to an increase in synaptic strength (LTP, Long Term Potentiation) and for somewhat larger time windows should instead lead to a decrease in synaptic strength (LTD, Long Term Depression). The corresponding case holds for spiking of the postsynaptic neuron followed within a certain time window by spiking of the presynaptic neuron (referred to as *post-before-pre spiking*). This follows the Hebbian theory of that whenever two neurons (one being presynaptic, one being postsynaptic) fire (emit spikes) close enough in time to one another, the synaptic strength is increased. And, conversely, that whenever the time between the spikes exceed a certain limit, the strength is instead decreased. The time window between the first and the second spike is referred to as the spike timing interval. The goal is to determine for which spike timing intervals the synaptic strength is increased, and for which intervals it is instead decreased. The spike timing interval is given by Δt , and, in order to distinguish pre-before-post spiking from post-before-pre spiking, we let $\Delta t = t_j - t_i$ (ms), where t_i and t_j represent the timing of the spike of neuron i and neuron j , respectively. As a result, pre-before-post spiking will take on positive time intervals, whereas post-before-pre spiking will take on corresponding negative time intervals.

A spike from neuron i , followed by a spike from neuron j , or vice versa a spike from neuron j followed by a spike from neuron i , is referred to as a spike pair. For the aim of simulating spike pairs at specific time intervals the previously used method of stimulating neurons via the input of Poisson processes is no longer suitable. Instead we opt to manually simulate 101 spike pairs and assign the spike pairs spike timing intervals in the range of every integer between -50 and 50 . The change in synaptic strength is determined by measuring w_{ij} before the spike pair (represented by w_{before}), and then again following the spike pair (represented by w_{after}), and then subtracting w_{after} from w_{before} . Finally, every Δw_{ij} value is normalised by division with the maximal Δw_{ij} of all spike pairs:

$$\frac{\Delta w_{ij}}{\max(\Delta w_{ij})}$$

Figure 3.2 is simulated to reproduce Figures 5C and 5F of Tully et al. [23]. As previously stated, the standard time constant τ_p in this thesis is different to the one in [23] (1 000 instead of 10 000). The simulation was done for our corresponding standard $\tau_p = 1\,000$, for the doubled value $\tau_p = 2\,000$, and additionally for $\tau_p = 10\,000$.




Figure 3.2: Dependency of spike timing intervals

Plot of the relationship between spike timing intervals and the resulting (normalised) weight change for 101 manually simulated spike pairs $\in [-50, 50]$ ms. Dashed red lines as reference lines of $x = 0$ and $y = 0$ respectively. The spike timing interval between spikes i (presynaptic) and j postsynaptic is illustrated on the x -axis. $\Delta t = t_j - t_i$ yields *pre-before-post* spiking on the positive x -axis and *post-before-pre* spiking on the negative x -axis. The simulation was done for three different values of time constant τ_p : $\tau_p = 1\,000$ ms (dark blue, standard), $\tau_p = 2\,000$ ms (green), $\tau_p = 10\,000$ ms (cyan). For the standard τ_p (dark blue) synaptic strengths w_{ij} are increased for every spike timing interval $\in (-25, 25)$ ms and decreased for every spike timing interval $\in [-50, -26] \cup [26, 50]$ ms. Increasing τ_p leads to the weight change Δw_{ij} remaining positive for larger spike timing intervals: between $(-30, 30)$ ms for $\tau_p = 2\,000$ ms (green) and between $(-35, 35)$ ms for $\tau_p = 10\,000$ ms (cyan).

The positive weight change window is determined by the time constant τ_p , where increased τ_p values leads to an increase of the positive weight change window. The reason is that the value of τ_p defines the duration of the underlying P -trace (see Equations 12, 13, 14) which in its turn helps estimate weight w_{ij} (see Equation 2). Meaning that, an increase of τ_p induces a longer decay time for the P -traces [23] (see Figure 3.2). The longer decay time causes the overlap of P_i and P_j to increase which results in an increase of P_{ij} and subsequently of w_{ij} .

The positive weight change window represents the spike timing intervals where the synaptic strength was increased following the spike pair, and thus inducing LTP. The weight change window where LTP occurred was, for $\tau_p = 1\,000$ ms, $\tau_p = 2\,000$ ms and $\tau_p = 10\,000$ ms, respectively, $(-25, 25)$ ms, $(-30, 30)$ ms and $(-35, 35)$ ms. Meaning that, for our standard $\tau_p = 1\,000$ ms, the synaptic strength is increased whenever i (presynaptic) and j (postsynaptic) spike with less than 25 ms between each other. The

choice of τ_p directly correlates to the amount of LTP induced, and consequently to the amount of LTD. The spike timing intervals where the synaptic strength was decreased and LTD exhibited are represented by the negative weight change windows on either side of Figure 3.2. The simulation demonstrated Hebbian learning [23] and is almost perfectly symmetric in relation to the y -axis (due to the same values of τ_{z_i} and τ_{z_j}).

For the standard values of τ_p in the respective models (1 000 in ours, in dark blue, compared to 10 000 in Figure 5C of Tully et al.), the results are very close (see Figure 3.2). In our figure, the trace (dark blue) intersects the x -axis at -25 ms at 25 ms, respectively. In Figure 5C of Tully et al., the numerical value seems to be very close, at around -23 ms and 23 ms. However, since no exact numerical values are stated in the paper, it is not possible to know the exact values for which the trace intersects the x -axis. Comparing Figure 3.2 to Figure 5F of Tully et al., where their τ_p value is doubled from 10 000 ms to 20 000 ms, the closest matching trace for our model is achieved for $\tau_p = 10\,000$ (in cyan) compared to the standard τ_p value of 1 000 ms (in dark blue). However, the decay towards negative values is slightly slower in our figure.

Figures 3.3 and 3.4 are to be compared with Figure 5E and Figure 5D of Tully et al. [23], respectively. The same simulation and setup as before was used ($\tau_p = 1\,000$), with the exception being the time constants τ_{z_i} and τ_{z_j} governing the Z_i and Z_j traces. The standard values of the time constants are $\tau_{z_i} = \tau_{z_j} = 5$ ms. In Figure 3.3, τ_{z_j} was kept at 5 ms while τ_{z_i} was decreased to 2 ms. In Figure 3.4, τ_{z_i} was instead kept at 5 ms while τ_{z_j} was decreased to 2 ms.

The results show that the shape and width of the positive weight change window is determined by the time constants τ_{z_i} and τ_{z_j} (see Figures 3.3 and 3.4). When τ_{z_i} is decreased from 5 ms to 2 ms, the number of pre-before-post spike pairs where the resulting weight difference remained positive was severely reduced (see Figure 3.3). For standard $\tau_{z_i} = 5$ ms, all pre-before-post spike pairs with spike timing intervals up to 25 ms had a resulting positive weight difference (see dashed red reference line). Now, for $\tau_{z_i} = 2$ ms, only pre-before-post spike pairs with spike timing intervals up to 11 ms still had a positive resulting weight difference. When instead τ_{z_j} is decreased from 5 ms to 2 ms, the corresponding difference is the same for post-before-pre spike pairs (see Figure 3.4). Meaning that, for $\tau_{z_j} = 5$ ms, all post-before-pre spike pairs with spike timing intervals down to -25 had a positive resulting weight difference. For $\tau_{z_j} = 2$ ms, the corresponding value was instead only -11. The result is that the values of τ_{z_i} and τ_{z_j} determines for how large spike timing intervals the synaptic strength will be increased, and thus exhibit LTP, and not instead decreased and then exhibit LTD. This follows what was previously introduced, that the values of the τ_{z_i} and τ_{z_j} time constants determine the time window for which correlated (or rather, close-in-time) spiking can be detected.

Very similar results are obtained in our simulations as to the ones in Figures 5E and 5D of Tully et al. [23]. In our figures, the respective τ_{z_i} or τ_{z_j} time constant is decreased to 2 ms (from 5 ms), in Figures 5E and 5D of Tully et al., the respective time constant




Figure 3.3: Decreasing τ_{z_i}

$\tau_{z_i} = \tau_{z_j} = 5$ ms (red, reference line), $\tau_{z_i} = 2$ ms, $\tau_{z_j} = 5$ ms (blue). When time constant τ_{z_i} is decreased, the positive weight change window for pre-before-post spiking ($\Delta t > 0$) is severely reduced, as illustrated by the difference between the red reference line and the blue line on the left-hand side.




Figure 3.4: Decreasing τ_{z_j}

$\tau_{z_i} = \tau_{z_j} = 5$ ms (red, reference line), $\tau_{z_i} = 5$ ms, $\tau_{z_j} = 2$ ms (blue). When time constant τ_{z_j} is decreased, the positive weight change window for post-before-pre spiking ($\Delta t < 0$) is severely reduced, as illustrated by the difference between the red reference line and the blue line on the right-hand side.

is also decreased to 2 ms from 5 ms. For both figures, when the time constants are decreased to 2 ms, the trace (in dark blue) instead intersects the x -axis at around 10 ms and -10 ms respectively (see Figures 3.3 and 3.4). This seems to correspond well to what is displayed in Figures 5E and 5D of the paper, but again no exact comparison can be made since we do again not know the exact numerical results obtained in said paper. However, more importantly, the figures are qualitatively very similar.

3.2 Simulation recreating Rodriguez et al.

Having presented reproductions of several results from Tully et al. [23], the next aim of this thesis is to further model the synaptic plasticity characteristics of our model, with a focus on the phenomenon of Long Term Depression (LTD). This chapter aims to reproduce results presented in a paper by Rodriguez et al. [15], and will use this as an important reference in subsequent simulations.

3.2.1 Dependency of spiking frequency

Not only the amount of anti-correlated spiking is of importance for the induction of LTD, but likewise the frequency between the occurring spikes is important. This frequency is referred to as spiking frequency and is measured in Hertz (Hz) for the amount of spikes occurring every 1 s (1 000 ms). This part, section 3.2.1, is based on

the paper of Rodriguez et al. [15], where LTD is measured in an experimental setting. The paper is based on extracellular recordings of mouse cortex. Simulations were done to recreate the results from Figure 2E of Rodriguez et al. [15] (see Figure 3.5), which measured the resulting weight change of three presynaptic spikes occurring at four different spiking frequencies: 20, 50, 100 and 200 Hz. In Figure 2E of Rodriguez et al. [15], the three presynaptic spikes are followed by a single presynaptic spike. This final single spike is, in an experimental setting such as the one of Rodriguez et al. [15], a "tool" to measure the weight. In our figure however, since it is based on a computational model, the single spike at the end is removed. The resulting weight change after the three presynaptic spikes was in each case measured 100 ms after the final spike. In the figure by Rodriguez et al. the synaptic strength is represented by a normalised EPSP slope in percentage. EPSP, the Excitatory Postsynaptic Potential, is a common way to measure the strength of synapses in an experimental setting. Since our model and setup did not make it feasible to simulate values of EPSP, we instead opted for our closest corresponding value of measuring the strength of synapses. This was the synaptic strength (the weight) w_{ij} . Since we are comparing two very similar, but not identical, things, meaning the comparison between EPSP and w_{ij} , the nature of the comparison is such that an exact comparison is not possible. Instead, this becomes a qualitative comparison, where the main interest is whether or not the behaviour and percentage changes are consistent enough and comparable between the two figures.

For the sake of consistency (with earlier simulations, see Figures 3.2, 3.3, 3.4), the change in weight Δw_{ij} in Figure 3.5 was normalised by division with the maximal Δw_{ij} . To allow comparison between our figure and the one of Rodriguez et al., the two must be normalised in the same manner. The normalised in "normalised EPSP slope" in Figure 2E of Rodriguez et al. simply refers to the normal being 100 percent and the change measured from that baseline. In order to compare the two, the data is taken from the latter and normalised as in our figure. This meant dividing by the maximal total decrease, which was 42 percent and found at 200 Hz.

The results obtained were similar to those in Figure 2E of Rodriguez et al. [15], which measured biological LTD in the cortex of mice. For the most part, our model closely replicates the behaviour and percentage changes from Figure 2E of Rodriguez et al. [15], with especially the values at 20, 50 and 200 Hz being very close. As was stated before, the aim is not an exact comparison but rather a qualitative comparison, which can be made. The percentage decreases in Figure 2E of Rodriguez et al. [15] at 20, 50, 100 and 200 Hz are, approximately, 1%, 18%, 38 % and 42%, respectively. Normalised by the maximal total decrease (42 %), these become approximately -0.02, -0.43, -0.90 and -1.0, respectively. Comparing this with our results, the corresponding values are, approximately, -0.02, -0.45, -0.82 and -1.0, respectively (see Figure 3.5). The noticeable exception is at 100 Hz, where our Figure 3.5 "only" illustrates a normalised decrease of -0.82, compared to the normalised decrease of -0.90 in Figure 2E of Rodriguez et al. [15]. The closely corresponding results in this figure were achieved by adjusting the plasticity parameter κ all the way from 1 to 0.065, and setting the initial weight as slightly positive ($w_0 \approx 0.18$).




Figure 3.5: Dependency of spiking frequency

Spiking frequencies of 20, 50, 100 and 200 Hz (x -axis) were simulated and represented spiking intervals of 50, 20, 10 and 5 ms, respectively. The weight was initialised as slightly positive (approximately 0.18). For every spiking frequency (20, 50, 100, 200 Hz), three presynaptic spikes were simulated, adhering to the respective spiking frequency. The resulting weight difference was calculated as $\Delta w = w_{after} - w_{before}$, and then normalised by division with the maximal Δw in the y -axis. For spiking frequency of 20 Hz, the normalised weight change was approximately -0.02, for 50 Hz it was instead slightly lower than -0.4, for 100 Hz it was very close to -0.8 and for 200 Hz it was -1.0.

3.3 Additional simulations of Long Term Depression

With our model tuned to, as closely as possible, match the results of Figure 2E of Rodriguez et al. [15], the next step is to further investigate some of the LTD behaviours of our model.

3.3.1 Long Term Depression

Long Term Depression, LTD, the synaptic plasticity phenomenon where the weight w_{ij} , representing the measure of synaptic strength between neurons i (presynaptic) and j (postsynaptic), is decreased following some spike events. In this simulation, LTD is induced by letting the presynaptic neuron spike a given amount of times, and not letting the postsynaptic neuron spike at all. Specifically LTD is induced by letting a (slightly) positive initial weight be exposed to only multiple presynaptic spikes, or only multiple postsynaptic spikes. This causes anti-correlation, which decreases the synaptic strength. The synaptic weight w_{ij} is likewise a measure of the strength of the synapse between neurons i and j . We want to demonstrate how the strength of the synapse is decreased when LTD occurs. We want to demonstrate this for multiple different values of w_{ij} , and so we simulate weight traces for multiple different values of initial weights. We simulate weight traces for 8 different initial weights specifically, which are given from initialising the P_{syn} -trace to slightly increasing values. In the standard case, P_{syn} is initialised as $\varepsilon^2 = 0.0033^2$, but now we scale ε^2 by increasingly positive values. We initialise to a given weight x by letting P_{syn} be scaled by e^x , so that P_{syn} is initialised as $P_{syn} = e^x \cdot 0.0033^2 = P_{syn} = e^x \cdot \varepsilon^2$. See Equation 2 for more.

Spikes were simulated by manually setting the membrane voltage of the presynaptic neuron to 0 mV (importantly, above membrane voltage threshold $V_{th} = -55$ mV), which is registered as a spike in our neural simulator (Brian2). The plasticity parameter was set to $\kappa = 0.065$, as in accordance with the previous result (see Figure 3.5). The results of this simulations shows that the weights are clearly decreasing (see Figure 3.6).




Figure 3.6: Long Term Depression, 3 presynaptic spikes

8 different weights between 0.30 and 1.00 were initialised and then exposed to 3 spikes of the presynaptic neuron i , with spike frequency 100 Hz. The initial weights w_0 were 0.30 (blue), 0.40 (orange), 0.50 (green), 0.60 (red), 0.70 (purple), 0.80 (brown), 0.90 (pink) and 1.00 (grey). The resulting weight is measured at exactly 100 ms (dashed line, black). For $w_0 = 1.00$ (grey), after 100 ms the weight had decayed to 0.61, equivalent to a weight decrease of 39 percent. The spike raster indicates the specific times of the presynaptic spikes (5 ms, 15 ms and 25 ms, respectively).

Presynaptic spiking, without any postsynaptic spiking close in time, induces LTD (Long Term Depression), which causes the weight to decrease. The highest initial weight is 1.0 (see Figure 3.6, grey trace), and, at the 100 ms mark where the resulting weight is measured in our model, it has decreased by almost 39 percent. This result, which closely matches the result of 38 percent from Figure 2E of Rodriguez et al. [15], was achieved by setting the plasticity parameter to $\kappa = 0.065$. The initialised weights range from 0.30 to 1.0. Between 0 ms and 100 ms, the decrease is steep, to then after the 100 ms mark start decaying at a much slower rate. The reason is, that the presynaptic spikes only occur during the first 25 ms, so after some time has elapsed the model no longer detects as much (or, flags for) anti-correlation. For the trace initialised to 0.8 (brown), the corresponding weight decrease is also close to 39 percent (0.8 to 0.49, approximately). The same is, to varying degrees, also true for the remaining weights. Most of the other weights correlate within a close interval to this result, with the exception being the lower initialised weights. From initial weights of 0.5 (green) and under, the percentage decrease surpasses 40 percent, getting as high as 45 percent for the lowest initialised weight (0.3, blue). After the 200 ms mark, all the weights start to stabilise and plateau.




Figure 3.7: Long Term Depression, 6 presynaptic spikes

Simulation of LTD, as in Figure 3.6, but with 6 presynaptic spikes. The spiking frequency remained at 100 Hz (with spikes at 5, 15, 25, 35, 45 and 55 ms, as can be read from the spike raster). Initial weights w_0 of 0.30 (blue), 0.40 (orange), 0.50 (green), 0.60 (red), 0.70 (purple), 0.80 (brown), 0.90 (pink) and 1.00 (grey). The resulting weight is measured at 100 ms (dashed line, black). The highest initialised weight (gray trace) shows a weight decrease of 50 percent.

The highest initial weight of 1.0 (see Figure 3.7, grey trace) has after the 100 ms mark decreased to 0.5, equivalent to a decrease of 50 percent. In Figure 3.6, the corresponding value was 39 percent. The simulation shows the same behaviour as in Figure 3.6: all the initialised weights (0.3 to 1.0 with 0.1 intervals) are decreasing rapidly, with the steepest decrease occurring in the first 100 ms (before the dashed black line), followed by a considerably slower decrease, and after 200 ms they start to stabilise and plateau. Even after the 100 ms line, although the weights decrease at a lower rate, they are still clearly decreasing. From 0.5 ($w_0 = 1.0$, grey) at the 100 ms line, the weight continues to decrease to approximately 0.35 at 200 ms. The last presynaptic spike occurs at 0.55 ms in this figure (compared to at 0.25 ms in Figure 3.6), so it is only expected that the weights take longer time before the decrease is as noticeable. This is caused of the time it takes for the model before it starts detecting the anti-correlation, which in its turn caused LTD. The fact that the weights decrease considerably more than in the previous figure is expected from our model, since more numerous presynaptic spikes, combined with a lack of postsynaptic spikes, (or, vice versa) is what is referred to as anti-correlation and is exactly what caused LTD.

4 Discussion

To summarise our work, we implemented a Spiking Neural Network model using the Leaky Integrate-and-Fire neuron model, a Bayesian learning rule and input Poisson processes. The model was adapted to our one neuron pair, which was used instead of the customary full-network implementation. To compute probabilities of spikes occurring and the synaptic strength, an adapted version of the spike-based extension of BCPNN from Tully et al. [23] was used. Some relatively minor adjustments had to be done to this Bayesian learning rule to account for our minimal network model. Notably, the weight equation was rewritten from Equation 1 to Equation 2. The goal was to investigate how well this minimal network model would hold for simulating different synaptic plasticity characteristics. The interest was to see whether similar behaviours could be replicated in our model, rather than searching for exact numerical matches of results. The model and results were compared to, in the first part, Tully et al. [23], and in the second part, also to the results of biological recordings presented in Rodriguez et al. [15]. The general focus of the thesis was on the concept of synaptic plasticity, with a specific emphasis on the phenomenon of LTD (Long Term Depression) which is an important feature of Hebbian synaptic plasticity.

The research questions were:

1. *Can we reproduce some of the main findings from Tully et al. (Figures 2 and 5), for a network of a single neuron pair and in neural simulator Brian2 (instead of NEST), given the same setup and neuron model?*
2. *Can we induce LTD (Long Term Depression), given the same setup and simulator as before, and get results comparable to biological recordings from Figure 2E of Rodriguez et al.?*

Rather than getting exactly matching results, the more important task of this thesis has been to get qualitatively similar results, and, in particular, results exhibiting the same synaptic plasticity characteristics and behaviours as to the corresponding results from Tully et al. [23]. The results obtained from the simulations in this thesis were very qualitatively similar, when the relevant parameters were adjusted. It is natural, and frankly expected, that the numerical values of the results will not be identical. This is both due to the nature of the task of reproducing results in a completely different neural simulator, which may function quite differently, and due to the minimal size of our chosen model (containing only one neuron pair). In addition, the input spiking activity is not the same (see Figure 3.1). To account for these factors, several parameter changes had to be made; namely of τ_{z_i} , τ_{z_j} , τ_p , f_{max} and ε . The parameters were chosen within what was deemed reasonable for the model and in a biological

context, with the goal of making the model exhibit the same behaviours as the one in Tully et al. [23].

For the simulations of *the local synaptic traces* (see Figure 3.1), although not the same numerical values were obtained, the overall behaviour of, and between, the different traces and from the input spike trains onto the traces, were generally found to correspond well to the behaviour of the traces from Figure 2 of Tully et al. [23]. The most noticeable difference, between the two models, is the fact that much less spiking is registered in our figure. This difference is attributed to the minimality of our chosen model. Our model contains only one neuron pair, and although background noise is modelled by Poisson processes in both models, this does not seem to sufficiently make up for the large size difference between our model and the one of Tully et al., consisting of a full network. Multiple parameters had to be adjusted to make up for this difference, as mentioned above. The adjusted parameters did not importantly impact the amount of spiking, but did impact within which numerical values the traces were fluctuating and, crucially, the shape of the traces. Even with more limited spiking, the adjusted parameters helped replicate the behaviour of Figure 2 of Tully et al. in multiple ways. This made the goal of simulating qualitatively similar results possible.

For the simulations of *dependency of spike timing intervals* (see Figures 3.2, 3.3 and 3.4), the results were very similar to the ones in Tully et al. This was achieved mainly due to the time constant τ_p in our model being adjusted to 1 000 ms (compared to 10 000 ms in Tully et al. [23]). This is indeed a notably large change of an important time constant, which can be argued was set unusually high in Tully et al. The time constant τ_p can be of quite varying time scales and 1 000 is a more common choice.

With our model and the adjusted parameter values previously described, it works well to induce LTD and get corresponding results as to biological recordings shown in Figure 2E of Rodriguez et al. [15] (see Figure 3.5). For the simulations of *dependency of spiking frequency*, the results were generally close, although there were individual differences, but more importantly the shape and behaviour of the figures were very similar. However, the comparison is not ideal. Rodriguez et al. measures EPSP (Excitatory Postsynaptic Potential), and in our model the corresponding parameter is the weight. This is a good corresponding parameter, but EPSP and weight are, of course, not identical. Another aspect is that, in Rodriguez et al., the EPSP baseline was set after 10 minutes of measuring, this is not something that is possible to implement in our model. Furthermore, in our model the weight is measured 100 ms from the start. In Rodriguez et al., it is not stated after how long time the EPSP is measured. When comparing the two figures visually, it is also important to remember that the two models are not normalised in the same way. To compare percentage decrease between the two, one must first normalise Figure 2E of Rodriguez et al. by dividing with the highest measured weight change, as explained in the results. The qualitatively similar results were obtained by decreasing the parameter of plasticity κ all the way down to 0.065.

Once it was clear that the model could be adapted to match results of biologically measured LTD, it was of interest to further investigate these properties of LTD. In the simulations called *Long Term Depression*, LTD was successfully induced for all initial weights with presynaptic spikes and a spiking frequency of 100 Hz (Figures 3.6, 3.7). By adjusting the plasticity parameter κ from 1 to 0.065, the approximately same weight decrease as in Figure 2E of Rodriguez et al. [15] was obtained for several of the different initialised weights. The most clear example is for the weight initialised to 1.0 in Figure 3.6, which is after 100 ms measured to around 0.61, meaning a weight decrease of approximately 39 percent. This is compared to the 38 percent weight decrease for three presynaptic spikes at a 100 Hz frequency in Figure 2E of Rodriguez et al.

5 Conclusions

Although the exact results could not be reproduced, it was possible to reproduce results which followed the same behavioural patterns and were in the more important aspects close to the figures of Tully et al — and this was the main goal of the thesis. The minimal model of the single neuron pair naturally impacted the results, but did not in any important ways prohibit the replication of the chosen figures (Figures 2 and 5) from Tully et al. [23]. The minimal model of a single neuron pair most importantly affected the amount of spiking activity (which can be seen in Figure 3.1, by comparing it to Figure 2 of Tully et al. [23]). The notable adaptation to our model (and simulator) came in the form of adjusting the aforementioned parameters and time constants, in several cases the values had to be changed by quite a lot, compared to what was used in Tully et al. [23]. Given these adjusted parameters, the replication of results from Tully et al. was largely successful, meaning that although the exact results could not be replicated, we importantly manages to replicate the majority of behavioural patterns from Figures 2 and 5 of Tully et al.

The results of Figures 2 and 5 from Tully et al. could be fairly replicated, given our minimal model of a single neuron pair and in neural simulator Brian2 (instead of NEST). However, this was not possible given the same setup as in Tully et al. Several important parameters had to be adjusted in order to achieve similar results.

Although, again, no exact comparison can be made for the above stated reasons, with adapted plasticity parameter κ we can get an almost identical figure, and therefore get very corresponding results to the measured biological data. So we can induce LTD, and get results comparable to biological recordings from Figure 2E of Rodriguez et al. [15]. We can also perform additional simulations of LTD for our model, where the simulation with the same amount of spikes (Figure 3.6, 3 presynaptic spikes) correspond relatively close to the biological recordings obtained from Figure 2E of Rodriguez et al. [15]. The additional simulation in Figure 3.7 can not be compared with the aforementioned figure from Rodriguez et al., since the condition is not the same (the amount of spikes), however it still offers interesting and valuable results about the LTD characteristics of our model.

It was possible to get results corresponding to biological data from Rodriguez et al. Although not perfect, the replications were successful to a satisfiable degree. However, it was again needed to adjust several important parameters. Notably, the plasticity parameter of κ had to be severely adjusted. Meaning that, it was possible to reproduce the results (qualitatively), but not given the same setup as used in Tully et al.

The conclusion is that the spike-based extension of BCPNN, from Tully et al., adapted to our one neuron case (the minimal model) functions well, but certain restrictions are apparent. The adaptation to the one neuron case of spike-based BCPNN is uncommon, and it should be seen as positive how well the results of the larger network could be replicated. The model offers intriguing insights into how two singular neurons could communicate with each other, and how the strength of their connection could be strengthened and weakened, given different conditions. The computational models in this particular area of research (both the one used in this thesis, and its counterparts) seem to correlate well to what is today known about the neural networks in the mammalian Central Nervous System (CNS). The simplification of the modelling is needed, and enable to offer interesting insights into how particular components of the brain may function, in an accessible manner and while retaining computational feasibility. The minimal model used in this thesis, with the singular neuron pair, further lowers the time and memory demands of simulating behavioural characteristics of networks of neurons.

5.1 Further research

This thesis offers a contribution to the particular field of study in the way of recreating certain common simulations in a relatively novel neural simulator (Brian2), and in the way of investigating the role that down-scaling the spiking neural network (a "cortical network") to simply contain a single neuron pair have on these simulations. Importantly, the model is in the last part of the thesis adapted to match results of biologically measured Long Term Depression, which further strengthens the argument of the biological plausibility of the model. The results of this thesis are positive, and contribute with novel insights about a minimal version of the chosen SNN model (with notably Leaky Integrate-and-Fire neuron model, spike-based BCPNN learning rule). However, the results are limited and cannot be generalised to apply to all types of neurons and cortical areas (specifically, excitatory neurons in the neocortex is studied). Furthermore, more reliable and interesting results in terms of comparing the results of the full network implementation, versus the minimal network implementation used in this thesis, would be given by ensuring that the parameters are adapted to make up for the differences of the two neural simulators (NEST vs. Brian2). The comparison between the two neural simulators would more accurately be done with a full network implementation, where quantifiable comparisons could be possible (as opposed to in this thesis, where only qualitative comparisons were possible). An interesting expansion of the work done within this thesis is studying Brian2 simulations of a more complex cortical network and expanding the project to investigating different aspects of associative memory within the model. A version of this has been done by Ferdinand Bujanowski, in the thesis *The Effect of Scaling Down Minicolumn Size on SNN Associative Memory Performance* (see github.com/FerdinandBujanowski/brian_BCPNN for simulations).

Bibliography

- [1] Donald O. Hebb. *The Organization of Behavior: A Neuropsychological Theory*. Wiley, New York, 1949, pp. 62–63.
- [2] Maggie Norris and Donna Rae Siegfried. *Anatomy and Physiology For Dummies, 2nd Edition*. Wiley Publishing, Inc., 2011.
- [3] Frederico A.C. Azevedo et al. “Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain”. In: *The Journal of Comparative Neurology* (2009). DOI: <https://dx.doi.org/10.1002/cne.21974>.
- [4] Mark F. Bear, Barry W. Connors, and Michael A. Paradiso. *Neuroscience: Exploring the Brain, 4th Edition*. Jones and Bartlett Publishers, Inc, 2020, pp. 24–28.
- [5] Ami Citri and Robert C Malenka. “Synaptic Plasticity: Multiple Forms, Functions, and Mechanisms”. In: *Neuropsychopharmacology Reviews* 33 (2008), pp. 18–41. DOI: [doi:10.1038/sj.npp.1301559](https://doi.org/10.1038/sj.npp.1301559).
- [6] Ramón Heberto Martínez Mayorquin. “Sequence learning in the Bayesian Confidence Propagation Neural Network”. In: *KTH* (2022).
- [7] Peter Dayan and L. F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. Massachusetts Institute of Technology, 2001.
- [8] 2026. URL: <https://www.khanacademy.org/science/biology/human-biology/neuron-nervous-system/a/the-synapse>. (accessed 31.05.2026).
- [9] Terje Bliss and Tim Lømo. “Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path”. In: *The Journal of Physiology* (1973).
- [10] Jeffrey C. Magee and Christine Grienberger. “Synaptic Plasticity Forms and Functions”. In: *Annual Review Neuroscience* 43 (2020), pp. 95–117. DOI: <https://doi.org/10.1146/annurev-neuro-090919-022842>.
- [11] William B. Levy and Oswald Steward. “Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus”. In: *Neuroscience* (1983). DOI: [https://doi.org/10.1016/0306-4522\(83\)90010-6](https://doi.org/10.1016/0306-4522(83)90010-6).
- [12] Robert C. Malenka and Mark F. Bear. “LTP and LTD: An Embarrassment of Riches”. In: *Neuron* (2004). DOI: <https://doi.org/10.1016/j.neuron.2004.09.012>.

- [13] Natalia Caporale and Yang Dan. “Spike Timing–Dependent Plasticity: A Hebbian Learning Rule”. In: *Annual Review of Neuroscience* (2008). DOI: <https://doi.org/10.1146/annurev.neuro.31.060407.125639>.
- [14] Roger A. Nicoll. “A Brief History of Long-Term Potentiation”. In: *Neuron* (2017). DOI: <http://dx.doi.org/10.1016/j.neuron.2016.12.015>.
- [15] Antonio Rodriguez-Moreno et al. “Presynaptic Self-Depression at Developing Neocortical Synapses”. In: *Neuron* (2013). DOI: <http://dx.doi.org/10.1016/j.neuron.2012.10.035>.
- [16] Anders Lansner and Örjan Ekeberg. “A one-layer feedback artificial neural network with a bayesian learning rule.” In: *International Journal of Neural Systems* (1989). DOI: [10.1142/S0129065789000499](https://doi.org/10.1142/S0129065789000499).
- [17] Anders Lansner and Anders Holst. “A higher order bayesian neural network with spiking units.” In: *International Journal of Neural Systems* (1996). DOI: <https://doi.org/10.1142/S0129065796000816>.
- [18] Bernhard A. Kaplan et al. “Anisotropic connectivity implements motion-based prediction in a spiking neural network.” In: *Frontiers in computational neuroscience* (2013).
- [19] Anders Lansner et al. “Reactivation in working memory: an attractor network model of free recall.” In: *PLoS one* (2013).
- [20] Anders Lansner. “Associative memory models: from the cell-assembly theory to biophysically detailed cortex simulations.” In: *Trends in neurosciences* (2009).
- [21] Florian Fiebig and Anders Lansner. “Memory consolidation from seconds to weeks: a three-stage neural network model with autonomous reinstatement dynamics”. In: *Frontiers in computational neuroscience* (2014).
- [22] Florian Fiebig and Anders Lansner. “A spiking working memory model based on hebbian short-term potentiation.” In: *Journal of Neuroscience* (2017).
- [23] Philip J. Tully, Matthias H. Henning, and Anders Lansner. “Synaptic and nonsynaptic plasticity approximating probabilistic inference”. In: *Frontiers in Synaptic Neuroscience* vol. 6 (2014). DOI: [10.3389/fnsyn.2014.00008](https://doi.org/10.3389/fnsyn.2014.00008).
- [24] Khadeer Ahmed. “Brain-Inspired Spiking Neural Networks”. In: *IntechOpen* (2020).
- [25] Kashu Yamazaki et al. “Spiking Neural Networks and Their Applications: A Review”. In: *Brain Sciences* (2022). DOI: [10.3390/brainsci12070863](https://doi.org/10.3390/brainsci12070863).
- [26] Louis Lapique. “Recherches quantitatives sur l’excitation électrique des nerfs traitée comme une polarisation.” In: *Journal of Physiol Pathol Générale* vol. 9 (1907), pp. 620–635.
- [27] Romain Brette and Wulfram Gerstner. “Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity”. In: *Journal of Neurophysiology* vol. 94 (2005). DOI: [doi:10.1152/jn.00686.2005](https://doi.org/10.1152/jn.00686.2005).
- [28] A. Sandberg, A. Lansner, and K.M. Petersson & Ekeberg. “A Bayesian attractor network with incremental learning”. In: *Network: Computation in Neural Systems*, (2002), pp. 179–194. DOI: doi.org/10.1080/net.13.2.179.194.

- [29] William B. Levy, James A. Anderson, and Stephen Lehmkuhle. *Synaptic Modification, Neuron Selectivity, and Nervous System Organization*. Psychology Press, 1985.
- [30] Niklas Wahlgren and Anders Lansner. “Biological evaluation of a Hebbian–Bayesian learning rule”. In: *Neurocomputing* (2001).
- [31] S. W. Roberts. “Control Chart Tests Based on Geometric Moving Averages”. In: *Technometrics* vol. 1 (1959), pp. 239–250.
- [32] *Dealing with measurement noise*. 2010. URL: <https://web.archive.org/web/20100329135531/http://lorien.ncl.ac.uk/ming/filter/filewma.htm>. (accessed: 12.05.2026).
- [33] John Lisman. “A mechanism for the Hebb and the anti-Hebb processes underlying learning and memory.” In: *Proceedings of the National Academy of Sciences of the United States of America* (1989). DOI: <https://doi.org/10.1073/pnas.86.23.957>.
- [34] Tim Bliss and Graham Collingridge. “A synaptic model of memory: long-term potentiation in the hippocampus”. In: *Nature* (1993).
- [35] Florian Fiebig, Pawel Herman, and Anders Lansner. “An indexing theory for working memory”. In: *eNeuro.org* (2020). DOI: doi.org/10.1523/ENEURO.0374-19.2020.
- [36] Marcel Stimberg, Romain Brette, and Dan FM Goodman. “Brian 2, an intuitive and efficient neural simulator”. In: *eLIFE* (2019). DOI: doi.org/10.7554/eLife.47314.
- [37] 2026. URL: <https://briansimulator.org/>. (accessed 14.05.2026).



Stockholm
University

Bachelor thesis

Computer science

September 2026

Matematiska institutionen

Stockholms universitet

106 91 Stockholm