

Principles of Network Plasticity in Neocortical Microcircuits

Présentée le 20 novembre 2023

À la Faculté des sciences de la vie
Projet Bluebrain
Programme doctoral en neurosciences

pour l'obtention du grade de Docteur ès Sciences

par

András ECKER

Acceptée sur proposition du jury

Prof. K. Hess Bellwald, présidente du jury
Prof. H. Markram, Dr M. Reimann, directeurs de thèse
Prof. S. Fusi, rapporteur
Prof. F. Zenke, rapporteur
Prof. W. Gerstner, rapporteur

Beauty is truth, truth (is) beauty, - that is all
— John Keats

To Donald O. Hebb, who was brave enough to link the mind to the brain,
in a time when no other psychologist dared to do so.

Acknowledgements

Firstly, I would like to thank my supervisors, Henry Markram and Michael W. Reimann. Henry, it was an honour to study plasticity as one of your last students. Michael, I am grateful to you for the weekly meetings, for constantly pushing things forward and getting them published, for even checking my code and for introducing me to climbing on the side. Secondly, I would like to thank Daniela Egas Santander, my closest collaborator. Your scientific rigour and attention to detail has shaped this thesis in countless ways. Thirdly, I would like to thank Eilif Muller and Giuseppe Chindemi. I will never forget that you invited me to write a grant proposal an hour after I finished my PhD interview. Even after leaving the group, you were always available for discussions about my PhD and about what to do next. Fourthly, I would like to thank my "*academic supervisor*" Wulfram Gerstner for always welcoming me at your lab meetings, and all members of the LCN for their kindness and feedback.

The Blue Brain Project is a group of more than 150 people, which makes every project a big collaboration. This thesis is composed of several articles, and I am thankful to all my co-authors. You taught me that clearly motivating and delivering results is just as important as the precise methodology behind them, and that splitting my sentences into two or three and using fewer pronouns would have facilitated that process . . . Special thanks to Sirio Bolaños-Puchet, my companion in misfortune, accelerator of any piece of code, and caretaker of my plants during Covid on the side. Furthermore, to James B. Isbister, for the countless disagreements and discussions we had and for always keeping these at the scientific level and never refusing a beer afterwards. I would also like to thank Cyrille Favreau, Elvis Boci, and Marwan Abdellah for their help with visualizing everything we did. A simulation neuroscience project of this scale would not be possible without professional software engineers. I might have only talked to you when things broke, but despite your sarcasm they always got fixed in no time. Mike Gevaert, Joni Herttuainen, Weina Ji, Fernando Pereira, M. Anil Tuncel, Thomas Delemontex, and Arseny V. Povolotsky huge thanks for this.

I would like to thank Florian Wyler and Amanda Klaeger, the only two Swiss people I actually got to know for countless hikes, "*Schpot breaks*", and always lifting me up. I would like to thank my family, for letting me study abroad and for always trying to give their best, despite getting the worst of me. I would also like to thank my grandmother Lúcia, for being so proud of me even if I did not have a single "*proper job*" so far. Lastly, I would like to thank you Csenge, for everything . . . but especially for the constant reminder of things more important in life than choosing colors for figures.

Abstract

Synaptic plasticity underlies our ability to learn and adapt to the constantly changing environment. The phenomenon of synapses changing their efficacy in an activity-dependent manner is often studied in small groups of neurons *in vitro* or indirectly through its effects on behavior *in vivo*. Investigating synaptic plasticity at an intermediate microcircuit level relies on simulation-based approaches, which offer a framework to reconcile fragmented and sparse experimental observations. Since Hebb's initial postulate, theoreticians have provided valuable insights about the role of cell assemblies, strongly interconnected groups of co-firing neurons, in learning and memory. However, most of these studies were limited in their scale, biological realism, and therefore generality. To overcome these limitations, we further improved and validated our previously published large-scale cortical network model featuring short-term plasticity and equipped it with a recently developed calcium-based model of long-term plasticity between excitatory cells. We calibrated the network to mimic an *in vivo* state characterized by low synaptic release probability and low-rate asynchronous firing and exposed it to ten different stimuli. By virtue of the model's non-random, biorealistic connectivity we could detect cell assemblies from the stimulus-evoked spiking activity of 186,665 neurons even in the naïve circuit, before the long-term plastic changes. This detection employed a combination of methods established by experimentalists. Leveraging the *in silico* nature of our setup, we then studied how the structure of synaptic connectivity underlies assembly composition ranging from feedforward thalamic innervation to intricate high-dimensional network motifs of the recurrent connectivity. Notably, we found that long-term plasticity sparsely and specifically strengthened synapses between cell assemblies: among 312 million synapses, only 5% experienced noticeable plasticity in 10 minutes of biological time and cross-assembly synapses underwent three times more changes than average. As our model neurons featured realistic morphologies and dendritic ion channels, we could also investigate how nonlinear dendritic processes influence assembly membership and the effects of long-term plasticity on synapses forming spatial clusters on postsynaptic dendrites. A comparative analysis of the network's responses to the different stimuli before and after the long-term changes revealed a network-level redistribution of efficacy from the superficial to the deep cortical layers. This shift led to prolonged stimulus-specific responses and more assemblies activating exclusively for a single pattern. In summary, using a state-of-the-art, bottom-up model of the cortical microcircuit we found sparse and specific plastic changes that reconfigured network dynamics while preserving its stability.

Abstract

Keywords: neocortex, microcircuit, simulation, connectivity, synaptic efficacy, short-term plasticity, long-term plasticity, cell assemblies, learning

Résumé

La plasticité synaptique est à la base de notre capacité à apprendre et à nous adapter à un environnement en constante évolution. La modification de l'efficacité de la transmission synaptique dépendante de l'activité neuronale est souvent étudiée dans de petits groupes de neurones *in vitro*, ou indirectement, par ses effets sur le comportement *in vivo*. L'étude de la plasticité synaptique au sein d'un microcircuit néocortical repose sur des approches basées sur la simulation, qui offrent un cadre permettant de réconcilier des observations expérimentales fragmentées et éparses. Depuis le postulat initial de Hebb, les théoriciens ont approfondi notre compréhension du rôle des assemblage de neurones (« cell assemblies »), des groupes de neurones activés conjointement et fortement interconnectés, dans l'apprentissage et la mémoire. Cependant, la plupart de ces études étaient limitées en terme d'échelle, de réalisme biologique et donc dans leur capacité à généraliser leurs résultats. Pour surmonter ces limites, nous avons amélioré et validé le modèle du microcircuit néocortical à grande échelle et équipé de plasticité à court terme que nous avons précédemment publié. Nous l'avons augmenté d'un modèle récemment développé de plasticité à long terme, entre les cellules excitatrices, basé sur le calcium. Nous avons calibré le réseau pour reproduire la faible probabilité de libération synaptique de neurotransmetteurs ainsi que le faible taux de décharge asynchrone observés *in vivo*, et nous l'avons exposé à dix stimuli différents. En raison de la connectivité structurelle, biologiquement réaliste du modèle, nous avons pu détecter des assemblage de neurones à partir de l'activité simulée de 186,665 neurones même au sein du circuit naïf. Notre méthode de détection se base sur une combinaison de méthodes établies par les expérimentateurs. Tirant parti de notre accès à l'entière connectivité du réseau *in silico*, nous avons ensuite étudié le lien entre l'organisation de la connectivité synaptique et la composition des assemblage de neurones, de l'innervation thalamique jusqu'aux motifs complexes à haute dimension de la connectivité récurrente du réseau. Nous avons notamment constaté que la plasticité à long terme renforçait de manière éparsée et spécifique les synapses entre les assemblage de neurones : sur 312 millions de synapses, seules 5% ont subi une plasticité notable en 10 minutes de temps biologique et les synapses inter-assemblées ont subi trois fois plus de changements que la moyenne. Comme nos neurones modèles présentaient des morphologies et des canaux ioniques dendritiques réalistes, nous avons également pu étudier l'influence des processus dendritiques non linéaires sur l'appartenance à une assemblée et les effets de la plasticité à long terme sur les synapses. Nous avons montré que celle-ci induit un regroupement des synapses sur les dendrites post-synaptiques. Une analyse comparative des réponses du circuit aux différents stimuli, avant et après les changements à long terme

Résumé

induit par la plasticité, a révélé une redistribution de l'efficacité synaptique, allant des couches corticales superficielles du réseau aux couches corticales profondes. Cette redistribution a eu pour effet de prolonger les réponses spécifiques au stimulus et d'augmenter le nombre d'assemblées s'activant exclusivement pour un seul motif. En résumé, en utilisant l'état de l'art de la modélisation du microcircuit néocortical, nous avons révélé des changements plastiques épars et spécifiques qui reconfigurent la dynamique du réseau tout en préservant sa stabilité.

Mots clefs : néocortex, microcircuit, simulation, connectivité, efficacité synaptique, plasticité à court-terme, plasticité à long-terme, assemblage de neurones, apprentissage

Contents

Acknowledgements	i
Abstract (English/Français)	iii
List of Figures	xi
List of Tables	xiii
Introduction	1
1 CA1 synapse physiology <i>in silico</i>	5
1.1 Introduction	6
1.2 Methods	7
1.2.1 Circuit building and synapse anatomy	7
1.2.2 Dendritic features of single cell models	8
1.2.3 Model of postsynaptic conductance and current	8
1.2.4 Short-term plasticity parameter fitting	9
1.2.5 Stochastic Tsodyks-Markram model with multi-vesicular release	10
1.2.6 Calibrating peak synaptic conductances through <i>in silico</i> paired recordings	11
1.2.7 Correcting for calcium ion concentration, temperature and liquid junction potential	11
1.2.8 Statistical analysis	13
1.3 Results	13
1.3.1 Literature curation	13
1.3.2 Synaptic model parameters	14
1.3.3 Validation of synaptic anatomy and dendritic attenuation	16
1.3.4 Short-term plasticity of synapses	19
1.3.5 Calibration of peak synaptic conductances to match PSP amplitudes	21
1.3.6 Parameter extrapolation	22
1.4 Discussion	25
1.5 Supplementary Methods	29
1.5.1 Single cell models	29
1.5.2 Different versions of the Tsodyks-Markram model	30
1.5.3 Membrane noise	32
	vii

Contents

1.6	Supplementary Figures and Tables	32
2	Cortical synapse physiology <i>in silico</i>	39
2.1	Results	40
2.1.1	Cortico-cortical synapse physiology	40
2.1.2	Thalamocortical synapse physiology	42
2.1.3	L4's contribution to L2/3 responses during whisker hold stimulus	46
2.2	Methods	48
2.2.1	Parametrization of spontaneous vesicle release	48
2.2.2	Reproducing Varani et al. (2022) <i>in silico</i>	48
2.3	Supplementary Tables	49
3	Cell assemblies and their underlying connectivity	53
3.1	Introduction	54
3.2	Results	56
3.2.1	Diverse set of assemblies can be detected from network simulations	56
3.2.2	Functional assemblies are determined by structural features	60
3.2.3	Assemblies are robust across simulation instances	66
3.3	Discussion	69
3.4	Methods	70
3.4.1	Network simulations	70
3.4.2	Distance metrics	71
3.4.3	Thalamic input stimuli	72
3.4.4	Assembly detection	72
3.4.5	Calculation of information theoretical measurements	73
3.4.6	Synaptic clustering coefficient	74
3.4.7	Determination of consensus assemblies	75
3.4.8	Calculation of spike time reliability	76
3.5	Supplementary Figures	77
4	Long-term plasticity induced sparse and specific synaptic changes	83
4.1	Introduction	84
4.2	Results	86
4.2.1	Calcium-based, biophysically detailed model of long-term plasticity	86
4.2.2	Achieving <i>in vivo</i> -like network activity	88
4.2.3	Sparse synaptic changes induced by long-term plasticity	89
4.2.4	More frequent plastic changes within and across cell assemblies	91
4.2.5	Synapse clustering contributes to the emergence of cell assemblies, and facilitates plasticity across them	94
4.2.6	Redistribution of assembly efficacies and prolonged stimulus-specific responses characterize the network after plasticity	97
4.2.7	Network topology changes are parametrized by input stimuli	100
4.3	Discussion	103

4.4	Methods	106
4.4.1	Calcium-based plasticity model	106
4.4.2	<i>In vivo</i> -like spontaneous and evoked activity	107
4.4.3	Network simulations	108
4.4.4	Evaluating control STDP rules	108
4.4.5	Cell assembly detection	108
4.4.6	Determination of consensus assemblies	109
4.4.7	Calculation of spike time reliability	110
4.4.8	Synaptic clustering coefficient and likelihood of plastic changes in synapse clusters	110
4.5	Supplementary Figures	112
	Conclusion	123
	Bibliography	149
	Curriculum Vitae	151

List of Figures

1.1	Synaptic data integration pipeline	13
1.2	Synapse model and parameter fitting	17
1.3	Hippocampal synapse anatomy	18
1.4	Hippocampal synapse physiology	20
1.5	Synapse diversity in the hippocampal CA1	25
1.6	Hippocampal single cell models	33
1.7	PSP attenuation in CA1 PCs	34
1.8	Calcium sensitivity of synaptic physiology	34
2.1	Overview of the physiology workflow (cortical network model)	41
2.2	Cortical synapse physiology	42
2.3	Thalamocortical synapse physiology	43
2.4	Reproducing and extending Varani et al. (2022) <i>in silico</i>	47
3.1	Cell assembly pipeline	55
3.2	<i>In vivo</i> -like activity <i>in silico</i>	57
3.3	Cell assembly detection	58
3.4	Connectivity determines cell assembly membership	62
3.5	Synapse clustering coefficient determines cell assembly membership	65
3.6	Consensus assemblies	67
3.7	Summary of <i>nIs</i> of all connectivity features and assembly membership	68
3.8	Spikes of significant time bins	77
3.9	Different number of clusters (of significant time bins)	78
3.10	Cell assemblies in L2/3	79
3.11	Anatomy of thalamocortical synapses	79
3.12	Synaptic clustering coefficient	80
3.13	Assemblies detected from averaged spike matrices	81
3.14	Efficiency of innervation	82
4.1	Overview of the cortical network model	87
4.2	Sparse synaptic changes in large-scale plastic simulations	90
4.3	More frequent plastic changes in cell assembly synapses	93
4.4	Changes in synapses participating in spatial clusters	96
4.5	Changes in cell assemblies after plasticity	99

List of Figures

4.6	Topology of single-pattern subnetworks	102
4.7	Physiology of excitatory cells and E to E connections	113
4.8	Synapse-specific parameters of the plasticity model	114
4.9	Calibration of the <i>in vivo</i> -like network state	115
4.10	Activity of the thalamic fibers	116
4.11	Changing connections in plastic control simulations	117
4.12	Individual changes in synaptic efficacy	118
4.13	Connectivity features underlying cell assembly membership	119
4.14	Comparison of cell assemblies before and after plasticity	120
4.15	Layer-wise propensity of changes per single pattern	121

List of Tables

1.1	Synaptic data inclusion and exclusion criteria	14
1.2	List of assumptions (CA1)	15
1.3	Synapse parameters (CA1)	24
1.4	Paired voltage-clamp recordings from rat CA1	35
1.5	Paired current-clamp recordings from rat CA1	36
1.6	Validation of hippocampal number of synapses per connection	37
1.7	Validation of (CA1) first PSP amplitude CVs	37
1.8	Validation of (CA1) first PSP amplitudes	37
2.1	Excitatory synapse parameters (cortex)	44
2.2	Inhibitory synapse parameters (cortex)	45
2.3	Thalamocortical synapse parameters	45
2.4	Validation of (cortical) first PSP amplitudes	50
2.5	Validation of (cortical) first PSP amplitude CVs	51
2.6	Validation of cortical mPSC frequencies	51
4.1	List of assumptions (cortex)	105

Introduction

All living organisms face an ongoing challenge to adapt to the ever-changing environment that surrounds them. Among higher life forms, mammals stand out for their remarkable ability to learn and recall memories when necessary. These essential cognitive processes find their roots in *synaptic plasticity*, the ability of synapses to change their *efficacy* in an activity-dependent manner. The term "*plastic*" comes from the Greek "*plastikós*" meaning "*fit for molding*" and its introduction to neuroscience dates back to the turn of the twentieth century (see Markram et al., 2011; Mateos-Aparicio and Rodríguez-Moreno, 2019 for reviews).

It became evident from early experiments at the neuromuscular junction that synaptic efficacy is indeed not static but changes over hundreds of milliseconds reflecting recent activity (Eccles et al., 1941). This form of synaptic plasticity was termed *short-term plasticity* (STP) and was later shown to be a prominent property of neocortical synapses as well (Thomson et al., 1993; Thomson and West, 1993; Deuchars and Thomson, 1996; Markram and Tsodyks, 1996). Its origin is presynaptic, i.e., it depends on the probability of transmitter release from the presynaptic axon (Dobrunz and Stevens, 1997). When the baseline release probability also known as *utilization of synaptic efficacy*¹ (U_{SE}) is high, consecutive postsynaptic potentials (PSPs) depress (the synaptic efficacy decreases). On the other hand, when U_{SE} is low PSPs facilitate. U_{SE} in turn depends on other factors, chief among them is the *extracellular calcium concentration* ($[Ca^{2+}]_o$). A decrease in $[Ca^{2+}]_o$ leads to a nonlinear decrease in U_{SE} (Ohana and Sakmann, 1998; Rozov et al., 2001; parametrized with a Hill coefficient of 4, Hill, 1910). Therefore, the difference between the $[Ca^{2+}]_o = 2 - 2.5$ mM used in slice preparations *in vitro* and the physiological $[Ca^{2+}]_o = 1 - 1.3$ mM *in vivo* (Jones and Keep, 1988; Massimini and Amzica, 2001) can cause drastic changes in STP profiles, turning depressing profiles into pseudo-linear², or even facilitating ones (Zucker and Regehr, 2002; Borst, 2010).

While providing evidence for the plasticity of synapses, the effect of STP vanishes after a second and thus cannot serve as a substrate for stable memories. One of the early theories about how learning and memory might manifest in the brain was formulated by Canadian psychologist Donald O. Hebb, who, in his influential book wrote: "*When an axon of cell A*

¹While release probability is perhaps easier to understand, in the articles presented in this thesis we follow the Tsodyks and Markram (1997) nomenclature and call it U_{SE} .

²Pseudo-linear STP profiles follows the nomenclature of Markram et al. (2015) and it describes an STP profile that is neither clearly depression, nor facilitating. Sometimes the same profile is referred to as "augmentation".

Introduction

is near enough to excite a cell B and repeatedly and 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 cell firing B, is increased." (Hebb, 1949). The highly quoted postulate of Hebb came to life from combining Golgi stainings, his experience in brain surgery, and psychology practice and was revolutionary in his time, as it stood in striking contrast to the views of Pavlov, Freud, and the Gestalt school of psychology. Its experimental validation came more than two decades later thanks to the technological advancement of slice electrophysiology and the discovery of *long-term potentiation* (LTP; an increase in synaptic efficacy that lasts for hours; Bliss and Lømo, 1973; Schwartzkroin and Wester, 1975). The molecular substrate of Hebbian coincidence detection is the N-methyl-D-aspartate (NMDA) receptor, which upon removal of the Mg^{2+} block conducts Ca^{2+} as well (Mayer et al., 1984, but see Nicoll and Malenka, 1995; Egger et al., 1999 for NMDA independent forms). The calcium-control hypothesis postulates, that prolonged, moderate amounts of Ca^{2+} lead to long-term depression (LTD; a long-lasting decrease in synaptic efficacy) while large transients of Ca^{2+} lead to LTP (Lisman, 1989). The entry of Ca^{2+} leads to the "*growth process . . . in one or both cells*" speculated by Hebb, and at the Schaffer collaterals between areas CA3 and CA1 of the hippocampus it is the insertion of α -amino-3-hydroxy-5-methyl-4-isoxazolepropionate (AMPA) receptors into the postsynaptic side of synapses (Malinow and Malenka, 2002). It took another two decades and another technological advancement, paired recordings from connected neurons, to prove the importance of not only the firing rates but the precise timing of the pre- and postsynaptic spikes (Markram et al., 1997b; Bi and Poo, 1998; Debanne et al., 1998). In these experiments, presynaptic spikes preceding the postsynaptic ones resulted in LTP, whereas LTD was observed when postsynaptic spikes occurred first, and this form of synaptic plasticity was christened spike-time dependent (STDP; Song et al., 2000). As Ca^{2+} plays a key role in long-term plasticity, this form of plasticity also depends on the $[Ca^{2+}]_o$. Low, physiological levels of $[Ca^{2+}]_o$ were recently shown to completely abolish STDP (Inglebert et al., 2020). Unlike the postsynaptically expressed plasticity of the Schaffer collaterals in the hippocampus, cortical plasticity affects the presynaptic U_{SE} as well, and therefore tightly couples short- and long-term plasticity (Markram and Tsodyks, 1996; Selig et al., 1999; Sjöström et al., 2003; but see Yasui et al., 2005; Enoki et al., 2009; Bliss and Collingridge, 2013). Markram and Tsodyks (1996) pointed out that cortical LTP is not necessary an increase in synaptic efficacy when multiple spikes are considered within a short interval, but rather a *redistribution of efficacy* towards the earlier spikes.

In parallel with neuroscientists, mathematicians and physicists also took up Hebb's postulate and formulated plenty of *learning rules*, by which artificial neural networks update the efficacy of connections based on the covariance of pre- and postsynaptic activity (Rochester et al., 1956; Sejnowski, 1977; Bienenstock et al., 1992; Hopfield, 1982; Oja, 1982). They coined the term "*computational neuroscience*" in a few years, and learning and memory are still among the main foci of the field. Rate-based models were translated to spike-based ones (Amit and Tsodyks, 1991; Gerstner et al., 1993) and a plethora of STDP rules were developed (Gerstner et al., 1996; Kempster et al., 1999; Song et al., 2000; Pfister and Gerstner, 2006). When STDP

proved to be synapse location- and therefore pathway-dependent (Sjöström and Häusser, 2006; Letzkus et al., 2006; Froemke et al., 2010), plasticity rules also evolved to feature voltage and calcium-dependence (Clopath et al., 2010; Graupner and Brunel, 2012) and a selected few even took into account the effects of the *in vivo* $[Ca^{2+}]_o$ (Higgins et al., 2014; Chindemi et al., 2022; Rodrigues et al., 2022). STP was also proposed to contribute to a range of functions, including temporal processing (Buonomano and Merzenich, 1995; Tsodyks and Markram, 1997; Markram et al., 1998; Goldman et al., 2002; Carvalho and Buonomano, 2011; Naud and Sprekeler, 2018), gain control, working memory, network stability, and compartment specific inhibition (Abbott et al., 1997; Maass and Markram, 2002; Sussillo et al., 2007; Keijsers and Sprekeler, 2022). Of particular interest is the biophysically detailed model of (Markram et al., 2015), in which modeling the pathway-specific $[Ca^{2+}]_o$ dependence of U_{SE} was crucial for understanding the transition from the *in vitro*-like synchronous activity, to the *in vivo*-like asynchronous firing regime. Combinations of diverse sets of plasticity rules (including, but not exclusively short- and long-term plasticity) have been recently used to model the formation and maintenance of *cell assemblies*, groups of neurons that Hebb hypothesized to develop a co-firing pattern after frequently repeated stimulation (Zenke et al., 2015; Fauth and Van Rossum, 2019).

While complex phenomena like plasticity are best studied under controlled laboratory experiments and in simplified models, the long-term goal of neuroscience is to characterize the rules governing plasticity under behaviorally relevant conditions. There was tremendous progress in the past decades in understanding cell assembly formation, or as the field calls it, in the identification and reactivation of "*engram cells*" in the hippocampus *in vivo* (Tonegawa et al., 2015). While these studies highlight the need to study Hebbian plasticity at the network level, most changes happen at the synapse level. The expression of synaptic proteins like PSD95 and AMPA subunit GluA1 are highly correlated with the efficacy of a synapse, and high-throughput methods tracking their levels are currently under development (Graves et al., 2021; Kim et al., 2023; Ray et al., 2023). While readily applicable to monitor the efficacy of millions of synapses *in vivo*, currently, these techniques cannot be supplemented with recordings of neural activity, thus the causes of the observed changes in efficacy remain unknown. Thus, while it is possible to control pre- and postsynaptic activity and measure subsequent changes in synaptic efficacy *in vitro*, no such tool exists for *in vivo* experiments. The bridge between them is often provided by complementary, *in silico* approaches (see above), and "*simulation neuroscience*" is emerging as a powerful tool to undertake tasks that require modeling high levels of detail (Markram et al., 2015; Fan and Markram, 2019). What distinguishes this approach from computational neuroscience is the scale and the attention to biophysical detail. By building models in a *bottom-up* manner, they can systematically integrate observations from hundreds of published articles and large, seemingly disconnected data sets (Markram et al., 2015; Fan and Markram, 2019; Billeh et al., 2020; Egger et al., 2020; Dura-Bernal et al., 2023). The present thesis aims to illustrate the integration of diverse set of experimental findings into a coherent biophysically detailed multi-scale model and then use it to study cell assemblies and the effects of plasticity under *in vivo* low $[Ca^{2+}]_o$.

Outline of the thesis

The rest of the thesis is structured as follows:

- Chapter 1 introduces the synaptic data-integration pipeline (including STP parameters) for the hippocampal area CA1 as a use case. It is a postprint of a "*Methods first*" publication, therefore offering a detailed view of the biophysical realism we use to model synapses in our group.
- Chapter 2 applies the methodology of Chapter 1 (not developed for but refined and published for the hippocampus) to parametrize the synaptic physiology of the latest release of our somatosensory cortex model. It also includes a high-level validation of synaptic pathways by reproducing an *in vivo* signal propagation experiment *in silico*.
- Chapter 3 describes in detail Hebbian cell assemblies, detected from spikes of cortical simulations under *in vivo*-like conditions. It mainly focuses on the structural connectivity rules underlying the assemblies' co-firing function and provides a foundation for Chapter 4.
- Chapter 4, in a true bottom-up fashion, is the main body of the thesis. It equips the cortical model with a calcium-based long-term plasticity model and studies the induced plastic changes in a network state that mimics *in vivo* activity. It emphasizes the stability of the network without explicitly modeling any homeostatic plasticity mechanism. Furthermore, it highlights that most of the sparse changes happen between cell assemblies and in synapses that form spatial clusters on the postsynaptic dendrites.

At the end, a general Conclusion section summarizes the key findings with respect to the network-level effects of short and long-term plasticity. This short summary is followed by a list of current limitations and possible future directions (not strictly restricted to plasticity).

1 CA1 synapse physiology *in silico*

This chapter is a postprint (with Wiley's permission and under the same CC-BY-NC license) of: **A. Ecker**, A. Romani, S. Sáray, S. Káli, M. Migliore, J. Falck, S. Lange, A. Mercer, A.M. Thomson, E. Muller, M.W. Reimann, S. Ramaswamy (2020) *Data-driven integration of hippocampal CA1 synaptic physiology in silico*. *Hippocampus* 30(11):1129–1145; doi: [10.1002/hipo.23220](https://doi.org/10.1002/hipo.23220).

Contribution: I reviewed the literature (based on A. Romani's initial sweep), helped A. Romani to derive and validate the structural connectome, fitted STP model parameters using open data sets, contributed to the soon to be open source software that calibrates and validates synaptic physiology parameters, created all the figures, and participated in writing the article. (A more detailed author contribution can be found at the end of the chapter.)

Abstract

The anatomy and physiology of monosynaptic connections in rodent hippocampal CA1 have been extensively studied in recent decades. Yet, the resulting knowledge remains disparate and difficult to reconcile. Here, we present a data-driven approach to integrate the current state-of-the-art knowledge on the synaptic anatomy and physiology of rodent hippocampal CA1, including axo-dendritic innervation patterns, number of synapses per connection, quantal conductances, neurotransmitter release probability, and short-term plasticity into a single coherent resource. First, we undertook an extensive literature review of paired-recordings of hippocampal neurons and compiled experimental data on their synaptic anatomy and physiology. The data collected in this manner is sparse and inhomogeneous due to the diversity of experimental techniques used by different groups, which necessitates the need for an integrative framework to unify these data. To this end, we extended a previously developed workflow for the neocortex to constrain a unifying *in silico* reconstruction of the synaptic physiology of CA1 connections. Our work identifies gaps in the existing knowledge

and provides a complementary resource towards a more complete quantification of synaptic anatomy and physiology in the rodent hippocampal CA1 region.

Keywords: hippocampus, CA1, data integration, *in silico* modeling, synapse

1.1 Introduction

The hippocampal formation, notably the CA1 region, is one of the most exhaustively studied regions in the mammalian brain and is thought to play a role, for example, in the acquisition of memory, recognition of place and language (Bliss and Collingridge, 2013; Buzsáki, 1989). Neuronal microcircuits in the hippocampal CA1 region process and store information through a myriad of cell-type-specific monosynaptic connections. Previous studies have shown that hippocampal cell-types are connected through multiple synaptic contacts, which are positioned across distinct axo-dendritic domains with a wide diversity in their physiology. Despite the wealth of data, we lack an integrative framework to reconcile the diversity of synaptic physiology, and therefore, identify knowledge gaps. There have been several noteworthy attempts to integrate knowledge on the cellular and synaptic components of hippocampal CA1 microcircuitry, which have provided a solid foundation to bring together anatomical properties and kinetic parameters of cell-type-specific connections - including the number of synapses per connection, connection probabilities, neurotransmitter release probabilities, and amplitudes of synaptic responses (Bezair and Soltesz, 2013; Wheeler et al., 2015; Moradi and Ascoli, 2019). As a complementary endeavor, we extended a previously developed framework to reconstruct neocortical microcircuitry at the cellular and synaptic levels of detail (Markram et al., 2015), by integrating disparate data on the physiology of short-term dynamics of depression and facilitation of cell-type-specific synaptic transmission in hippocampal CA1. Using this framework, we identified and extrapolated organizing principles to predict missing knowledge for a repertoire of connection types, for example, the short-term dynamics and peak conductance of synaptic connections between inhibitory interneurons (Klausberger and Somogyi, 2008; Pelkey et al., 2017), which remain largely uncharacterized, and could, therefore, require high-throughput strategies that employ multiple whole-cell patch-clamp recordings to surmount the relatively low yield obtained through conventional paired recordings (Perin et al., 2011; Jiang et al., 2015; Espinoza et al., 2018).

We accounted for the dynamic and probabilistic nature of synaptic transmission by fitting experimental traces using a stochastic generalization of the Tsodyks-Markram (TM) short-term plasticity (STP) model (Tsodyks and Markram, 1997; Markram et al., 1998; Fuhrmann et al., 2002), and also considered temperature and extracellular calcium concentration ($[Ca^{2+}]_o$) differences, which were adjusted using Q10 and Hill scaling factors, respectively.

Measuring peak quantal conductances directly at individual synaptic contacts remains very difficult, if not impossible with current experimental techniques. While theoretically the peak synaptic conductance can be calculated from voltage-clamp recordings by simply dividing the peak post-synaptic current (PSC) by the liquid junction potential (LJP)-corrected driving

force, this approach does not take into account the space-clamp artifact (Spruston et al., 1993; Williams and Mitchell, 2008; Gulyás et al., 2016). We have recently demonstrated that space-clamp corrected peak synaptic conductances in neocortical connections are at least two-three-fold higher than estimated previously (Markram et al., 2015). As a connection is formed by several synaptic contacts, each subject to a different space-clamp effect, a purely theoretical correction is challenging. We, therefore, used an alternative approach, where we calibrated peak synaptic conductances in the *in silico* model of connected pairs such that the resulting postsynaptic potential (PSP) amplitudes match *in vitro* recordings. This yielded an estimate of peak synaptic conductance since other factors affecting the PSP amplitude – such as number and location of synapses, release probability and reversal potential, depression, facilitation, and synaptic conductance rise and decay time constant - were independently validated beforehand.

The resulting models for a subset of hippocampal connection types were applied predictively to the remaining uncharacterized connection types by clustering them into nine groups based on synapse types and neuronal biomarkers and applying the estimated parameters within each group. Curated and predicted parameters presented here should serve as a resource to researchers aiming to model hippocampal synapses at any level, while the detailed methodology intends to give a guideline to utilize such a framework to integrate data from other brain regions or species.

1.2 Methods

1.2.1 Circuit building and synapse anatomy

A detailed model of the rat hippocampal CA1 area was built by adapting a previously described pipeline for reconstructing neocortical microcircuitry (Markram et al., 2015). In brief, detailed axo-dendritic morphological reconstructions and electrophysiological traces obtained from the dorsal part of hippocampal CA1 were used to build single cell-type-specific computational models (Migliore et al., 2018; see Supplementary Methods). The resulting single-cell models were assembled in an atlas-based volume corresponding to the dimensions of the hippocampal CA1 region (Ropireddy et al., 2012), cell-densities and proportions, which yielded a tissue model consisting about 400,000 cells, ~90% pyramidal cells (PCs) and ~10% interneurons comprising 11 distinct morphology types (m-types; see Supplementary Methods and Supplementary Figure 1.6; Bezaire and Soltesz, 2013). Structural appositions between axons and dendrites were detected based on touch distance criteria and subsequently pruned to yield a functional connectome through an algorithmic process, which was constrained with experimentally reported bouton density, number of synapses per connections and connection probability (Reimann et al., 2015). A previous study suggests targeted innervation of interneurons from PCs (Takács et al., 2012). Therefore, to recreate this tendency, touch distances from PCs to interneurons were set to 6 μm as against 1 μm for connections between PCs. Furthermore, touch distances of 6 μm for connections between all interneurons and 1 μm for

Chapter 1. CA1 synapse physiology *in silico*

connections between interneurons and PCs were assumed. In this manner, the number and location of synapses for each cell-types specific connection were derived in a data-driven manner. When reproducing paired recordings *in silico* (see below), monosynaptically connected pairs of neurons were sampled from this reconstructed circuit based on their inter-somatic distance as sampling criterion.

1.2.2 Dendritic features of single cell models

Detailed, multicompartmental morpho-electrical models with 3D reconstructed dendrites from Migliore et al. (2018) were used in the present study (see Supplementary Methods and Supplementary Figure 1.6). The attenuation of synaptic responses along the dendrites with varying diameters was validated against experimental data from Magee and Cook (2000) using the HippoUnit framework (see Supplementary Methods). To this end, excitatory postsynaptic current (EPSC) like currents were injected into the apical trunk of PCs with varying distance from the soma and PSPs were simultaneously measured at the local site of the injection and in the soma.

1.2.3 Model of postsynaptic conductance and current

Synaptic conductances were modeled with bi-exponential kinetics:

$$g(t) = \hat{g} A (e^{-t/\tau_{decay}} - e^{-t/\tau_{rise}}) \quad (1.1)$$

where \hat{g} (nS) is the peak synaptic conductance and τ_{rise} and τ_{decay} (ms) are PSC rise and decay time constants respectively. The $A = -e^{-t_p/\tau_{rise}} + e^{-t_p/\tau_{decay}}$ normalization constant ensures that synapses reach their peak conductance at $t_p = (\tau_{decay}\tau_{rise})/((\tau_{decay} - \tau_{rise})\log(\tau_{decay}/\tau_{rise}))$ (ms). (Equation (1.1) is modified below to take stochastic release of multiple vesicles into account.) AMPAR and GABAR synaptic currents are then computed as:

$$I(t) = g(t)(V_m(t) - E_{rev}) \quad (1.2)$$

where V_m (mV) is the membrane potential and E_{rev} (mV) is the reversal potential of the given synapse. NMDAR currents depend also on Mg^{2+} block:

$$I_{NMDA}(t) = g(t)mg(V_m(t))(V_m(t) - E_{rev}) \quad (1.3)$$

where $mg(V_m)$ is the LJP-corrected (see below) Jahr-Stevens nonlinearity (Jahr and Stevens, 1990):

$$mg(V_m) = \frac{1}{1 + e^{-c_1 V_m} (C_{Mg^{2+}} / c_2)} \quad (1.4)$$

where $C_{Mg^{2+}}$ (mM) is the extracellular magnesium concentration and $c_1 = 0.062$ (1/mV) and $c_2 = 2.62$ (mM) are constants (the difference from the original Jahr and Stevens, 1990 constant

is because the authors did not correct for the LJP offset of ~ 5 mV). PC to PC NMDAR rise and decay time constants are Q10 corrected (see below; Q10 = 2.2 ms for rise and 1.7 ms for decay time constants Hestrin et al., 1990; Korinek et al., 2010) values from Andrasfalvy and Magee (2001): $\tau_{rise} = 3.9$ ms, $\tau_{decay} = 148.5$ ms. All, but the CCK+ interneuron excitatory afferents have the same NMDAR time constants as the PC to PC ones, while the PC to CCK+ interneuron NMDAR conductance decays with a slower time constant: $\tau_{decay} = 298.75$ ms (Le Roux et al., 2013; Matta et al., 2013; Cornford et al., 2019). Peak NMDAR conductance \hat{g}_{NMDA} (nS) is calculated from the AMPAR one by multiplying it with NMDAR/AMPA peak conductance ratio. PC to PC NMDAR/AMPA peak conductance ratio = 1.22 was taken from Groc et al. (2002); Myne et al. (2003). PC to CCK+ interneuron NMDAR/AMPA \hat{g} ratio was set to 0.86, as against 0.28 for PC to other interneurons (Le Roux et al., 2013; Matta et al., 2013). Synaptic currents are individually delayed based on axonal path length and conduction velocity of $300 \mu\text{m}/\text{ms}$ (Stuart et al., 1997) and an additional 0.1 ms delay of neurotransmitter release (Ramaswamy et al., 2012).

1.2.4 Short-term plasticity parameter fitting

STP of synapse dynamics was fit by the TM model (Tsodyks and Markram, 1997; Markram et al., 1998). The model assumes that each synapse has a pool of available neurotransmitter resources (R) that is utilized by a presynaptic action potential (AP) with a release probability (U). The utilization of resources leads to postsynaptic conductance that is proportional to the amount utilized. R decreases and U increases after an AP and both R and U recover between spikes to a steady-state value. The speed of recovery is parameterized by time constants D and F (ms) that together determine the short-term dynamics of the synapse. This is described by the following differential equations:

$$\frac{dR(t)}{dt} = \frac{1 - R(t)}{D} - U(t)R(t)\delta(t - t_{spike}) \quad (1.5)$$

$$\frac{dU(t)}{dt} = \frac{U_{SE} - U(t)}{F} + U_{SE}(1 - U(t))\delta(t - t_{spike}) \quad (1.6)$$

where U_{SE} is the utilization of synaptic efficacy or absolute release probability (also known as the release probability in the absence of facilitation), $\delta(t)$ is the Dirac delta function and t_{spike} indicates the timing of a presynaptic spike. Each AP in a train elicits an $A_{SE}U(t_{spike})R(t_{spike})$ amplitude PSC, where A_{SE} is the absolute synaptic efficacy. $R = 1$ and $U = U_{SE}$ are assumed before the first spike.

The U_{SE} , D , F and A_{SE} free parameters of the model were fit to amplitudes of experimentally recorded trains of PSCs. In the case of Losonczy et al. (2002), amplitudes were already extracted by the authors, while in the case of Kohus et al. (2016) custom-written Python routines were used to extract them from the averaged postsynaptic traces. Fitting the 10 + 1 recovery spikes (Table 1.1) was done by using a multiobjective genetic algorithm from BluePyOpt (Van Geit et al., 2016). For Kohus et al. (2016) different frequency stimulations (10, 20 and 40 Hz) were fit

together for better generalization. Thus, the optimized error function contained 3 (frequencies) x 11 (peak amplitudes) points. For the event-based version of the equations above see Maass and Markram (2002) and Supplementary Methods. The Python source code fitting amplitudes from multiple frequencies is available on GitHub under [/BluePyOpt/examples](#).

1.2.5 Stochastic Tsodyks-Markram model with multi-vesicular release

For the simulation of synapses, the canonical TM model (introduced above and used for fitting experimental traces) was modified to include stochastic release of multiple vesicles, and connected to the model of postsynaptic conductance described above. To take multi-vesicular release (MVR) into account in the postsynaptic conductance model, the classical "quantal model" of Del Castillo and Katz (1954) was used. In this model, synapses are assumed to be composed of N_{RRP} (size of the readily releasable pool) release sites, each of which has a probability of release U (see deterministic TM model above) and contributes a $1/N_{RRP}$ quanta to the postsynaptic response (Loebel et al., 2009; Ramaswamy et al., 2012, 2015; Markram et al., 2015; Barros-Zulaica et al., 2019). Unlike in the deterministic TM model above, individual quanta were assumed to be released in an all-or-none fashion with probability $U(t)$ (Fuhrmann et al., 2002). Vesicle availability is also an all-or-none process where only available vesicles can be released. To this end, synaptic vesicles were implemented as 2-state (available:1 and unavailable:0) Markov processes. After release, the state is set to unavailable and the probability of staying in the unavailable state at time t was described as a survival process, with the time constant D . The state transitions are described by the following set of equations:

$$\begin{aligned}
 P_{1 \rightarrow 0} &= U(t) \quad \text{see equation (1.6)} \\
 P_{1 \rightarrow 1} &= 1 - P_{1 \rightarrow 0} \\
 P_{0 \rightarrow 0} &= e^{-(t-t_{spike})/D} \\
 P_{0 \rightarrow 1} &= 1 - P_{0 \rightarrow 0}
 \end{aligned} \tag{1.7}$$

The above-described model converges to the canonical TM model in the limit (number of trials $\rightarrow \infty$). In this formalism a presynaptic AP releases only a fraction $N_r \leq N_{RRP}$ fraction of vesicles, which follows a Bernoulli distribution. Equation (1.1) is thus updated as follows:

$$g(t) = \hat{g}(d(t) - r(t)) \tag{1.8}$$

$$\frac{dd(t)}{dt} = -\frac{d(t)}{\tau_{decay}} + A \frac{N_r}{N_{RRP}} \delta(t - t_{spike}) \tag{1.9}$$

$$\frac{dr(t)}{dt} = -\frac{r(t)}{\tau_{rise}} + A \frac{N_r}{N_{RRP}} \delta(t - t_{spike}) \tag{1.10}$$

where r and d are the rising and decaying components of the postsynaptic conductance respectively. The implementation of the above described stochastic synapse model is available at the open-access NMC portal (Ramaswamy et al., 2015).

These changes to the canonical TM model introduce variability of the postsynaptic traces, where the magnitude of the variability depends on the additional N_{RRP} parameter (Loebel et al., 2009; Barros-Zulaica et al., 2019). *In vitro* this variability is typically assessed by the coefficient of variation (CV, std/mean) of the peak PSC (or PSP) amplitudes. Therefore, the N_{RRP} was calibrated to match the CVs of the first PSCs extracted from the raw traces of Kohus et al. (2016). For a better comparison, artificial membrane noise was added to the simulated traces (see Barros-Zulaica et al., 2019 and Supplementary Methods).

1.2.6 Calibrating peak synaptic conductances through *in silico* paired recordings

Paired recordings were replicated *in silico* as follows: Firstly, pairs were selected from the circuit based on pathway specific distance criteria used by experimentalist ($100 \mu m^3$ for cells in the same layer and $200 \mu m^3$ for cell pairs from different layers). Secondly, postsynaptic cells were current-clamped to match the LJP-corrected (see below) steady-state potential specified in the experiments. It is important to note, that in the case of PCs sodium channels were blocked (*in silico* TTX application) when clamping above -58 mV to avoid spontaneous firing of the cell models (see Figure 5 in Migliore et al. (2018)), whereas sodium channels were not blocked in *in vitro* experiments. Next, the presynaptic cell was stimulated by somatic current injection, which resulted in a PSP recorded in the soma of the postsynaptic neuron. This protocol was repeated for 50 monosynaptic connections of the same pre-post combination with 35 repetitions for each neuron pair. Finally, the mean PSP amplitude was compared against experimentally data and the peak conductance value was calibrated using the formula:

$$\hat{g} = \hat{g} \frac{PSP_{exp}(1 - PSP_{model}/df)}{PSP_{model}(1 - PSP_{exp}/df)} \quad (1.11)$$

where PSP_{exp} (mV) and PSP_{model} (mV) are the experimental and modeled PSPs amplitudes respectively and $df = |E_{rev} - V_{SS}|$ (mV) is the driving force. For all the experiments we aimed to reproduce, $E_{rev} = -8.5$ mV was calculated for excitatory connections, while $E_{rev} = -73$ mV for inhibitory connections (Moradi and Ascoli, 2019). All simulations were run using the NEURON simulator as a core engine (Hines and Carnevale, 1997) with the Blue Brain Project's collection of hoc and NMODL (Hines and Carnevale, 2000) templates for parallel execution on supercomputers (Hines et al., 2008a,b). The default temperature in all simulations was set to 34 °C, and the integration time step to 0.025 ms.

1.2.7 Correcting for calcium ion concentration, temperature and liquid junction potential

Before integrating published parameters from different sources into the *in silico* synapse model, they were corrected for differences in experimental protocols. This included scalings for $[Ca^{2+}]_o$ levels different from 2 mM, temperatures different from 34 °C and the correction of holding and steady-state potentials by the theoretical LJP.

Chapter 1. CA1 synapse physiology *in silico*

Levels of $[Ca^{2+}]_o$ impact the neurotransmitter release probability. The corresponding *in silico* correction was applied by scaling the absolute release probability U_{SE} parameter (see above) of the synapses, using the Hill isotherm with $n = 4$ (Hill, 1910). The Hill equation below describes the non-linear increase in release probability as a function of increasing $[Ca^{2+}]_o$:

$$U_{SE} = U_{SE_{max}} \frac{[Ca^{2+}]_o^4}{K_{1/2}^4 + [Ca^{2+}]_o^4} \quad (1.12)$$

where $U_{SE_{max}}$ is the maximum value of the release probability (≤ 1) at high $[Ca^{2+}]_o$ and $K_{1/2}$ is the $[Ca^{2+}]_o$ at which U_{SE} is one-half of $U_{SE_{max}}$. $U_{SE_{max}}$ and $K_{1/2}$ parameters can be fit to data points (e.g. an indicator of release probability – the ratio between PSP amplitudes) measured at different $[Ca^{2+}]_o$ s. $K_{1/2}$ values were taken from Rozov et al. (2001), 2.79 (mM) for steep and 1.09 (mM) for shallow calcium dependence and were shown to generalize well for other characterized pathways of the neocortex (see Supplementary Figure S11 in Markram et al. (2015)). In the absence of hippocampus specific data, we followed the approach of Markram et al. (2015) and assumed a steep dependence in PC to PC and PC to distal dendrite targeting inhibitory (O-LM) cells, and a shallow dependence between PC to proximal targeting cells (PVBC (PV+ basket cell), CCKBC (CCK+ basket cell), and axo-axonic cell). For experimentally uncharacterized pathways an intermediate calcium dependence was used, as the average of the steep and shallow ones. This intermediate curve was in agreement with the few relevant data points for specific hippocampal synaptic connections (Price et al., 2008; Tyan et al., 2014).

The temperature dependence of kinetic parameters such as rise and decay time constants were corrected by dividing them with Q10 scaling factors:

$$\tau_{sim} = \tau_{exp} / Q10^{(T_{sim} - T_{exp})/10} \quad (1.13)$$

where τ is the time constant, Q10 is an empirically determined, receptor-specific parameter, $T_{sim} = 34^\circ\text{C}$ is the temperature used in the simulations, while $T_{exp} (< T_{sim})$ is the temperature of the experiment. The Q10 correction was only needed for the NMDA current between connected PCs (see above) because all other kinetic values that we used were recorded at near physiological temperature ($\sim 34^\circ\text{C}$).

Holding and steady-state potentials were corrected by the theoretical LJP (Neher, 1992). These potentials arise from the differences in solutions in the pipette and bath and are in 2-12 mV range for the standard solutions. Theoretical LJPs, calculated from the reported pipette and bath solutions were obtained from Moradi and Ascoli (2019).

1.2.8 Statistical analysis

R values for validating matching experimental and model values are Pearson correlations. Data are presented as mean \pm std to yield comparable values to the experimental ones. U_{SE} , D , F distributions from two different sources (e.g. found in the literature vs. fitted here) are said to be comparable if the mean of the second distributions is not further away than one-half of the std of the first distribution.

1.3 Results

1.3.1 Literature curation

First, we undertook an extensive literature review of paired recording experiments, and compiled data on the various parameters (Figure 1.1, step 1; Tables 1.1 and 1.2 for the data inclusion and exclusion criteria, and a list of data and modeling assumptions, respectively; see also Supplementary Tables 1.4 for voltage-clamp data from rat hippocampal CA1, and 1.5 for current-clamp recordings). The data collected in this manner is sparse and inhomogeneous, due to the disparate experimental conditions used by different groups and were, therefore, corrected for various aspects (Figure 1.1, step 2). For example, $[Ca^{2+}]_o$ is known to affect release probability and, therefore, an additional Hill scaling had to be considered while parametrizing short-term plasticity (STP) models (see Methods). Rise and decay time constants of synaptic currents are influenced by temperature differences but can be corrected with Q10 factors (see Methods). For electrophysiological recordings patch pipettes have become the method of choice over sharp electrodes, which necessitates applying an LJP correction for voltage traces (see Methods).




Figure 1.1: *In silico* data integration pipeline. **1:** 51 peer-reviewed papers, spanning 21 years were used to compile data on various parameters of connected neurons in rat CA1 including connection probability, number of synapses per connections, axo-dendritic innervation profile, kinetics, STP profiles, calcium and temperature sensitivity. **2:** Parameters were integrated into a common framework and experimental paradigm, including temperature, $[Ca^{2+}]_o$ and LJP corrections. TM models of STP were fit to publicly available raw traces. **3:** *In silico* paired recordings were run to correctly adjust the unitary peak conductance of connections with experimentally characterized PSP amplitudes. **4:** The resulting parameters were averaged within each of the 9 classes of synapses and used predictively to describe experimentally uncharacterized pathways.

Table 1.1: Data inclusion and exclusion criteria.

Data inclusion criteria	
1	For the characterization of number of the synapses per connections we used published values from anatomical studies employing electron and light microscopy; in rat CA1 slices.
2	For the validation of our axo-dendritic innervation patterns, we used figures from published studies with biocytin filled pairs; under light microscopy; in rat CA1 slices.
3	For the characterization of synaptic physiology, we prioritized data from: paired recordings from identified m-types; in rat CA1 slices; at 2 mM $[Ca^{2+}]_o$; and 34 °C; with reported holding/steady-state potential; and reported LJP or recording solutions.
4	For the parametrization the decay time constant of single PSCs we used published decay time constant fits (independent of the model e.g. single vs. bi-exponential fit).
5	For fitting the TM model we used average raw PSC traces as well as published peak PSC amplitudes; with 10 spikes at different frequencies plus a recovery spike.
6	For the validation of the TM model we used published fits from the neocortex (Markram et al., 2015) in order to compare U_{SE} , D , F values of the corresponding pathways.
7	For the estimation of the N_{RRP} we used raw PSC traces (all trials) to estimate the CV of the first peak PSC amplitude as well as published N_{RRP} estimates.
8	For the calibration of peak synaptic conductance amplitudes, we used published peak PSP amplitudes (see Supplementary Table 1.5).
9	For the validation of the peak synaptic conductances, we used single-receptor conductance and receptor number estimates.

Data exclusion criteria	
10	In the case of multiple reports of a single parameter or reference data, we prioritized publications which were already used for other parameters and excluded the others (see e.g. Pawelzik et al. (2002) in Supplementary Table 1.5).
11	When we had access to individual PSP amplitudes beyond the usually reported mean \pm std, we excluded outliers and used the updated mean \pm std as target PSP amplitude (see Supplementary Table 1.8).

1.3.2 Synaptic model parameters

We integrated the collected and corrected data into a model of synaptic transmission that includes STP and stochastic neurotransmitter release. We found that for some connection types the parameters of this model could be fully determined by employing *in silico* paired recordings (Figure 1.1, step 3). Yet, for the majority of connection types parameters had to be extrapolated (Figure 1.1, step 4). We use "*synapse*" to refer to a single anatomical synaptic contact and "*connection*" to indicate the collection of all synaptic contacts between a given presynaptic and postsynaptic neuron, comprising one or more synapses.

The underlying synapse model consisted of several parts, each with their associated parameters, which we determined in a 6 step procedure: We modeled synaptic connections with bi-exponential conductances requiring 8 parameters. Three parameters (E_{rev} , τ_{rise} , τ_{decay})

Table 1.2: List of assumptions. All the assumptions that were made to arrive at model parameters from a sparse set of raw data and published values.

-
- 1 We assume that after all the listed correction in this paper, all parameters coming from different sources can be used together to parameterize the synapse models.
 - 2 When using data from Kohus et al. (2016) we assumed that CCK+ DTIs (dendrite-targeting interneurons) are SCA cells in SR. Furthermore, we assumed that synaptic currents measured in mouse CA3 are representative of similar pathways in rat CA1.
 - 3 In the lack of representative data and our lack of neurogliaform cells we assumed that all inhibitory synapses are mediated purely by $GABA_A$ receptors.
 - 4 For calculating release probabilities at different $[Ca^{2+}]_o$, we assumed that Hill functions parameterized with cortical data generalize well for hippocampal connections.
 - 5 For modeling synaptic currents, we assumed that all CA1 synapses can be described with bi-exponential conductances, with vesicle release kinetics governed by the stochastic TM model. When modeling dendritic PSC decays, we assumed a single exponential function, parametrized with a time constant extracted from somatic recording.
 - 6 In the process of calibrating synaptic peak conductances we simulated only the synapses mediating the given connection and thus we assume that the background activity does not matter.
 - 7 Some of the biggest assumptions are inherited from the network model: In this work, we assumed that the published electrical models of single cells (Migliore et al., 2018) capture the behavior of different neurons in rat CA1. (The fact that unlike experimentalists, we cannot clamp PC models to potentials above -58 mV without blocking sodium channels seems to violate this assumption.) We also assumed that the cell composition and cell density within each layer are homogeneous and the constrained connectivity reflects the connectivity of rat CA1.
 - 8 Kinetic parameters for a given pathway are drawn from a distribution, but since (almost) all experimental data used to derive these parameters are representative for a given connection and not for individual synapses per se, we use the same parameters for all synapses mediating a single connection.
 - 9 The biggest assumption is that one can extrapolate parameters from experimentally characterized pathways, to fill in missing values. When generalizing our parameters for similar, experimentally uncharacterized pathways we group CA1 interneurons based on only one chemical marker. However, cells express many of these and the markers overlap (see hippocampome.org (Wheeler et al., 2015)). By PV+ cells we mean: SP_PVBCs, SP_BS cells, and SP_AA cells. By CCK+ cells we mean: SP_CCKBCs, SR_SCA cells and SLM_PPA cells. The only interneurons in our NOS+ class are SP_Ivy cells. (Neurogliaform cells would belong here as well.) We assume all neurons in SO: SO_OLM cells, SO_BS cells, SO_Tri cells, and SO_BP cells to be SOM+.
 - 10 A usually unspoken, implicit assumption on communication between neurons is used here as well, namely, we model only glutamatergic and GABAergic synapses between presynaptic axons and postsynaptic somata and dendrites. Thus, we leave out co-transmission and neuromodulators acting on different receptors, retrograde messengers, any kind of gap-junctions and any axonal receptors.
-

Chapter 1. CA1 synapse physiology *in silico*

were directly obtained from the literature (see Supplementary Table 1.4 for AMPAR and GABAR rise and decay time constants, Methods for NMDAR time constants, and Supplementary Table 1.5 for reversal potentials (Moradi and Ascoli, 2019)). In particular, for the τ_{decay} (Supplementary Table 1.4) with the exception of Maccaferri et al. (2000) who used either single or weighted bi-exponential fits, none of the other studies we considered explicitly reported how τ_{decay} was extracted. Therefore, we extrapolated single exponential fits for τ_{decay} of all pathways, which were measured through somatic voltage-clamp recordings. We used these measurements directly as dendritic PSC time constants without any correction for attenuation (Table 1.2). STP was modeled with the Tsodyks-Markram (TM) model, which added three parameters (U_{SE} , D , F) to a synaptic connection type. They were fit in conjunction to the experimentally observed STP behavior (Figure 1.2, step 4; see Methods). Stochastic synaptic transmission was modeled by extending the TM model to include quantal release from multiple sites. This added another parameter (N_{RRP}) that was fit to the observed variability of PSC amplitudes of experimental traces in terms of their coefficient of variation (CV std/mean; Figure 1.2, step 5; see Methods). Finally, the mean amplitude of PSPs depended on three of the parameters and thus could be fit to the peak synaptic conductance (\hat{g}) only after the other two parameters had been determined (Figure 1.2, step 6).

In addition to the parameters of synaptic models, the physiology of PSPs is also dependent on several anatomical parameters, which result from the single-cell and tissue modeling workflow (see Methods; Supplementary Figure 1.6). To ensure the accuracy of the fitted synaptic parameters we independently validated aspects of the modeled anatomy (Figure 1.2, steps 1, 2). In the following sections, we present the results of the anatomical validations, followed by the results of the various fits of synaptic parameters.

1.3.3 Validation of synaptic anatomy and dendritic attenuation

The anatomical properties of synaptic connections such as number of synapses per connection and axo-dendritic innervation patterns, along with the dendritic properties of single cell models were validated against experimental data (Figure 1.3). Pairs of synaptically connected neurons were sampled from a dense tissue-level reconstruction of the rat hippocampal CA1 region (see Methods; Supplementary Figure 1.6; Figures 1.3A and 1.4A). The number of synapses per connection for the handful of experimentally characterized pathways (Ali, 2011; Biro et al., 2005; Buhl et al., 1994a,b; Deuchars and Thomson, 1996; Földy et al., 2010; Maccaferri et al., 2000; Sik et al., 1995; Vida et al., 1998) were consistent with biological data ($r = 0.98$; Figure 1.3B and Supplementary Table 1.6). The mean number of synapses per connection for the *in silico* pathways that have been experimentally characterized are as follows: Excitatory to excitatory (E-E): 1.26 ± 0.6 ; inhibitory to excitatory (I-E): 8.2 ± 2.1 ; excitatory to inhibitory (E-I); only connections between PC to O-LM cells): 2.8 ± 1.2 ; inhibitory-inhibitory (I-I): 2.8 ± 0.2 (Supplementary Table 1.6). A systematic, quantitative characterization of axo-dendritic innervation profiles for hippocampal CA1 synaptic connections is largely lacking. Therefore, although we derived many predictions of axo-dendritic innervation profiles from *in silico*




Figure 1.2: *In silico* synapse model and parameter fitting: Properties of the network (left) and the parameters synapse model (right) determine certain features of the emergent PSPs (middle). (These PSP features are schematized at the bottom of the figure. Individual trials are shown in gray and their average postsynaptic voltage trace in black.) These dependencies between properties/parameters and PSP features (indicated by arrows, and dots where they join and continue as a single arrow) were used to fit the synapse model parameters to data in 6 steps. Left: Parts of the network model that affect these features such as biophysical and anatomical neuron models via dendritic attenuation (1) as well as dendritic innervation and the number of synapses per connection (2) are independently validated. Top right: Parameters of the model of postsynaptic conductance are taken from averaged experimental PSC traces (3). Middle right: The TM model of STP adds three parameters that are fit to observed STP behavior (4). Bottom right: The model of stochastic quantal release adds another parameter fit to the observed CV of PSP amplitudes (5). In the last step, peak synaptic conductances are calibrated to match PSP amplitudes from data (6). Numbers on arrows indicate that the given parameter was validated against - or fitted to data, while numbers on boxes indicate that the parameters were taken from literature and directly plugged in into the model.

synaptic pathways, these could, however, only be validated based on anecdotal evidence (Figures 1.3A and 1.4A). In addition, we sampled neuron pairs at inter-somatic distances of $0 - 200 \mu\text{m}$ to predict their connection probability and number of synapses per connection (Figure 1.3C, D). The upper bound of $200 \mu\text{m}$ ensured that we obtained a sufficient number ($100 \leq n \leq 5000$) of pairs for all connections, even where the pre-post neurons were in different layers e.g. Schaffer collateral-associated and O-LM cells to PC connections. Although the perforant path-associated cell to PC connections occur in our model, they were excluded in these analyses since their somata are farther apart than the general $200 \mu\text{m}$ distance criteria chosen for these predictions.

Finally, we also validated the dendritic attenuation profile of PSPs in single neuron models of PCs, which were also found to be consistent with experimental data (Magee and Cook, 2000) ($\tau_{model} = 235.2$, $\tau_{exp} = 155.6$; Supplementary Figure 1.7).

Chapter 1. CA1 synapse physiology *in silico*




Figure 1.3: *In silico* synapse anatomy. **A:** A representative *in silico* O-LM (purple) to PC (blue) pair, with synapses visualized in red. 3D morphologies were reconstructed with the NeuroLucida software by the members of the Thomson/Mercer lab (Migliore et al., 2018). **A1:** Branch order distribution ($n = 5000$ connections) of the presynaptic (O-LM) axons. **A2:** Branch order distribution of the postsynaptic (PC) tuft dendrites. **A3:** Distribution of the number of synapses per connection of the *in silico* O-LM to PC pathway. *In vitro* experimental data is indicated in red. **A4:** Distance dependent connection probability of the *in silico* O-LM to PC pathway. **B:** Validation of the number of synapses per connection against experimental data. (E: excitatory, I: inhibitory, e.g.: I-E: inhibitory to excitatory pathways.) Dashed gray line represents perfect correlation between experimental and model values. **C:** Predicted mean number of synapses per connections for all pathways in the full-scale CA1 network model. Only connections with $\leq 200 \mu\text{m}$ intersomatic distance were used to calculate the average. Averages were calculated from $100 \leq n \leq 5000$ pairs. White boxes represent connections that are not present in the circuit model due to the lack of axo-dendritic overlap (given the $\leq 200 \mu\text{m}$ intersomatic distance sampling criteria). Experimentally measured values (same as on its left) are highlighted with black rectangles. (Caption continues on the next page.)

Layer abbreviations: SR: stratum radiatum, SP: stratum pyramidale, SO: stratum oriens. M-type abbreviations: AA: axo-axonic cell, BP: back-projecting cell, BS: bistratified cell, CCKBC: CCK+ basket cell, Ivy: ivy cell, OLM: oriens-lacunosum moleculare cell, PC: pyramidal cell, PVBC: PV+ basket cell, PPA: performant path-associated cell, SCA: Schaffer collateral-associated cell, Tri: trilaminar cell (see Supplementary Methods). **D**: Predicted mean connection probability (within $200\ \mu\text{m}$ intersomatic distance) for all pathways in the CA1 network model. M-type abbreviations, white boxes, black rectangles and number of pairs are as in C.

1.3.4 Short-term plasticity of synapses

The synaptic physiology of hippocampal CA1 connections express a rich diversity of STP profiles in response to presynaptic AP trains at different stimulus frequencies (Ali et al., 1998, 1999; Ali and Thomson, 1998; Losonczy et al., 2002; Pouille and Scanziani, 2004; Kohus et al., 2016; Éltés et al., 2017). However, to the best of our knowledge, only Losonczy et al. (2002) reported TM model parameters for CA1 pathways and used an additional recovery spike elicited about 500 – 100 ms after the last spike in the train, which is crucial to characterize frequency-dependent STP profiles of depression and facilitation (Gupta et al., 2000). Published STP parameters from Losonczy et al. (2002) were used for PC to BC pathways, after refitting a subset of their data, and ensuring their consistency with our resulting U_{SE} , D , F values (see Methods). The dataset from Kohus et al. (2016) were obtained in the mouse CA3 region at $1.6\ \text{mM}\ [Ca^{2+}]_o$, which differs from the rest of the datasets we considered, we nevertheless made use of this resource due to the availability of their raw data, which was subsequently used in our procedure of fitting TM model parameters (see Methods; Table 1.1 for data inclusion and exclusion criteria; Table 1.2 for a list of data and modeling assumptions).

The resulting TM model parameters following the fitting procedure were consistent with those in the source dataset (Kohus et al., 2016). In addition, we were able to match the CVs of the first PSC amplitudes ($r = 0.8$; Figure 1.4B, Supplementary Table 1.7), by calibrating N_{RRP} (see Loebel et al., 2009; Barros-Zulaica et al., 2019 and Methods) with the resulting values of N_{RRP} in a biologically plausible range. An elegant study demonstrated that under experimental conditions to induce high neurotransmitter release probability (high Mg/Ca) CCKBC to PC connections in CA3 are characterized by MVR (with $N_{RRP} = 5 - 7$ vesicles) (Biró et al., 2006). However, uni-vesicular release (UVR, $N_{RRP} = 1$) is more prevalent under physiological conditions (Biró et al., 2006). The *in silico* CV of CCKBC to PC PSCs with $N_{RRP} = 1$ compared well against experimental data obtained under physiological conditions. In the cases of synaptic connections from PVBC to PC and PVBC a value larger than 1 ($N_{RRP} = 6$) vesicles were required (see Methods; Figure 1.4B). For pathways not present in the Kohus et al. (2016) dataset, the N_{RRP} could not be calibrated and was thus assumed. The assumption of MVR with $N_{RRP} = 2$ vesicles at each excitatory to excitatory connections was used in this study (Tong and Jahr, 1994; Conti and Lisman, 2003; Christie and Jahr, 2006; Barros-Zulaica et al., 2019), while UVR was assumed at all other non-calibrated pathways (see Biro et al. (2005); Gulyás et al. (1993) suggesting UVR for certain PC to interneuron connections).

Chapter 1. CA1 synapse physiology *in silico*




Figure 1.4: *In silico* synapse physiology. **A:** *In silico* paired recording experiment with the STP protocol used in Kohus et al. (2016). Presynaptic (PVBC) voltage trace is shown on top. *In silico* PVBC (green) to PC (blue) pair, with synapses visualized in red in the middle. 3D morphologies were reconstructed with the NeuroLucida software by the members of the Thomson/Mercer lab (Migliore et al., 2018). Postsynaptic (PC) experimental traces recorded *in vitro* (in gray) and their mean in red, as well as model traces recorded *in silico* (in gray) and their mean in blue, are presented at the bottom panel. Insets show the variance of the first IPSCs. **B:** Validation of the CV of the first PSC amplitudes (excluding failures) against experimental data. (E: excitatory, I: inhibitory, e.g.: I-E: inhibitory to excitatory pathways.) Dashed gray line represents perfect correlation between experimental and model values. **C:** Validation of the PSP amplitudes against experimental data. **D:** Predicted CVs of first PSC amplitudes (excluding failures) for all pathways in the CA1 network model after synapse parameter generalization. As in Figure 1.3C, only connections with $\leq 200 \mu m$ intersomatic distance were used to calculate the average postsynaptic response from $n = 20$ pairs with 35 repetitions for each pair. Postsynaptic cells were held at -65 mV in *in silico* voltage-clamp mode. M-type abbreviations, white boxes, and black rectangles are as in Figure 1.3C. **E:** Predicted PSP amplitudes for all pathways in the CA1 network model after synapse parameter generalization. 20 pairs with 35 repetitions for every possible connection. Postsynaptic cells were held at -65 mV steady-state potential in *in silico* current-clamp mode. Consistent with Gulyás et al. (1993), PC to interneurons are the strongest. M-type abbreviations, white boxes, black rectangles and number of pairs are as in D. (Caption continues on the next page.)

F: Properties of postsynaptic (PC) IPSPs from 100 PVBC to PC pairs with 35 repetitions each. **F1:** Distribution of *in silico* PSP amplitudes. *In vitro* experimental data from Pawelzik et al. (2002) is indicated in red. **F2:** Distribution of *in silico* PSP 10-90% rise times. (10-90% rise time constants of PSCs are fixed to 0.2 ms in the model, but the PSP rise times vary.) **F3:** Distribution of *in silico* PSP decay time constants (single exponential fit). **F4:** Distribution of *in silico* PSP latencies. **F5:** Distribution of the CVs of the first *in silico* PSP amplitudes (excluding failures). **F6:** Distribution of *in silico* failures (0 measurable PSP amplitude from 35 repetitions).

Based on the literature and our model-fitting we identified several rules to group STP profiles. The mapping of STP profiles for all pathways is as follows: PC to O-LM cells (Ali and Thomson, 1998; Biro et al., 2005; Losonczy et al., 2002; Pouille and Scanziani, 2004) and other interneurons in stratum oriens (Éltes et al., 2017) E1 (excitatory facilitating). PC to PC (Deuchars and Thomson, 1996), PC to all SOM- interneurons (Ali et al., 1998; Losonczy et al., 2002; Pouille and Scanziani, 2004) E2 (excitatory depressing). CCK+ interneurons to CCK+ interneurons (Ali, 2007, 2011; Kohus et al., 2016) I1 (inhibitory facilitating), PV+ and SOM+ interneurons to PC (Ali et al., 1998, 1999; Bartos et al., 2002; Buhl et al., 1995; Daw et al., 2009; Kohus et al., 2016; Maccaferri et al., 2000; Pawelzik et al., 2002) as well as interneurons to interneurons (except the CCK+ ones) (Bartos et al., 2002; Daw et al., 2009; Elfant et al., 2008; Karayannis et al., 2010; Kohus et al., 2016; Price et al., 2005) I2 (inhibitory depressing). CCK+ and NOS+ (only Ivy cells, since we lack NGF morphologies) to PC (Funtealba et al., 2008; Kohus et al., 2016; Price et al., 2008) I3 (inhibitory pseudo linear). The parameters of the groups and the resulting dynamics are summarized in Table 1.3 and Figure 1.5.

Neurotransmitter release probability and the STP profile are not only sensitive to the recording temperature and the developmental age but also $[Ca^{2+}]_o$ (Rozov et al., 2001; Williams and Atkinson, 2007; Guzman et al., 2016). Therefore, we modeled $[Ca^{2+}]_o$ sensitivity with a highly non-linear scaling of U_{SE} (baseline release probability) values (see Methods). As an exemplar result of this additional modeling detail, the PC to PC pathway exhibits an E3 (excitatory pseudo-linear) STP profile characterized by low PSP amplitudes with high trial-by-trial variability and failures at *in vivo* like $[Ca^{2+}]_o$ levels (1.1 – 1.3 mM) compared to the *in vitro* levels (2 – 2.5 mM) E2 (excitatory depressing) profile (Supplementary Figure 1.7B). U_{SE} values are scaled by a Hill isotherm (see Methods) parameterized with data from PSP amplitudes in neocortex (Supplementary Figure S11 in Markram et al. (2015)), which is an indirect measure of the release probability. Here, we have shown that applying this Hill isotherm directly to the U_{SE} values indeed results in the same scaling profile of PSP amplitudes in the case of PC to PC connection (Supplementary Figure 1.8A).

1.3.5 Calibration of peak synaptic conductances to match PSP amplitudes

There is a dearth of studies characterizing both the PSC and PSP amplitudes for the same connections in rat hippocampal CA1 (compare Supplementary Tables 1.4 and 1.5). Therefore, we only used PSP amplitudes that were measured experimentally to calibrate the *in silico* peak synaptic conductances in order to match the *in vitro* PSPs (Ali et al., 1998; Ali and Thomson,

Chapter 1. CA1 synapse physiology *in silico*

1998; Cobb et al., 1997; Deuchars and Thomson, 1996; Fuentealba et al., 2008; Pawelzik et al., 1999, 2002) (see Figure 1.4D and Table 1.3). Having parameterized all relevant anatomical and physiological synaptic properties including the number of synapses per connections, axo-dendritic innervation patterns, PSC rise and decay time constants, STP parameters, N_{RRP} , NMDA/AMPA peak conductance ratio, and reversal potential, we undertook *in silico* paired-recordings by following a sequence of steps. A connected pair of neurons within a pathway specific intersomatic distance (usually $\sim 100 \mu m$) for a given pathway was sampled from the hippocampal CA1 model, the postsynaptic neuron was current-clamped to a pathway-specific steady-state potential (see Supplementary Table 1.5), an AP was elicited in the presynaptic neuron, which caused a postsynaptic response, measured in the soma. After repeating this sequence for multiple pairs of the same pathway ($n = 50$) with many trials ($n = 35$), we derived the peak synaptic conductance value that yielded the reference mean experimental PSP amplitude (see Methods). Next, we repeated the same protocol on a set of 50 randomly selected pairs with the calibrated peak conductance values as a validation of our approach ($r = 0.99$; Figure 4C and Supplementary Table 1.8).

As an independent external validation of the peak conductances, we compared them against sparse published data estimating single-receptor conductance and receptor numbers in excitatory synapses on PCs. Hippocampal CA1 PCs receive most of their inputs from CA3 PCs through the Schaffer collaterals (Megías et al., 2001; Takács et al., 2012), whereas in this study we only considered intrinsic connections - e.g. excitatory connections between local CA1 PCs - and not long-range extrinsic projections. Thus, single-receptor conductances and receptor number estimates from the Schaffer collateral synapses were assumed to generalize for the intrinsic PC to PC connections. Using non-stationary fluctuation analysis on EPSCs recorded in outside-out dendritic membrane patches, Spruston et al. (1995) estimated peak single-receptor conductances of 10.2 pS and 43.5 pS for AMPARs and NMDARs, respectively. Based on these numbers, our calibrated values resulting in a peak AMPAR conductance of 0.6 ± 0.1 nS is the net result of ~ 59 AMPARs per synaptic contact. Based on an experimentally measured NMDAR/AMPA peak conductance ratio of 1.22 (Myme et al., 2003), we predict that there are about ~ 18 NMDARs constituting a single synaptic contact between CA1 PCs. Our *in silico* predictions are consistent with experimental studies that estimate $\sim 58-70$ AMPA and $\sim 5-30$ NMDA receptors (Jonas et al., 1993; Spruston et al., 1995; Nusser et al., 1998; Matsuzaki et al., 2001). Taken together, these experimental datasets enable an independent validation of the calibrated peak conductance of PC to PC connections in CA1. In addition, we also predict an average GABA peak conductance of 2 ± 1 nS at a single inhibitory synaptic contact comprising ~ 100 GABAergic receptors, which is also in good agreement with previous estimates (Mody and Pearce, 2004).

1.3.6 Parameter extrapolation

By integrating all the synaptic parameters and performing paired recordings *in silico*, we procured a dataset of 16 pathways (Table 1.3). The number of theoretically possible pathways

(based on 12 m-types) in our CA1 circuit model is 144, however, only 102 of these are biologically viable based on the extent of axo-dendritic overlap (Figure 1.3C, D). Therefore, the parameters of the remaining 90% of the pathways had to be extrapolated. We generalized the anatomical properties of synapses (number of synapses per connection, connection probability, bouton density, innervation profile) obtained from the fraction of characterized to the remainder of uncharacterized pathways as shown previously (Reimann et al., 2015; Markram et al., 2015). However, for STP profiles of hippocampal connections obtained from studies that reported measurements of paired-pulse ratios, but did not provide the raw experimental traces with ≥ 2 presynaptic spikes (Deuchars and Thomson, 1996; Ali and Thomson, 1998; Fuentealba et al., 2008), we applied analogous parameters from the somatosensory cortex (Markram et al., 2015). We performed a prior consistency check of the parameter ranges for similar connection types - perisomatic inhibitory (BCs) to PC, and inhibitory to inhibitory - that have been experimentally characterized in both somatosensory cortex and hippocampus and found them to be comparable. Therefore, our rationale to generalize 4 sets of U_{SE} , D , F values from the somatosensory cortex to the hippocampus (Tables 3) could be justified. Thereafter, we approximated the missing parameters with averaged values across specific connection types that were grouped according to neurochemical markers that appear to have similar STP parameters and peak conductances (Table 1.3). For example, it is known that excitatory synapses on distal dendrite targeting interneurons, which predominantly express SOM - such as PC to O-LM connections - are mostly facilitating, and on the contrary inhibitory synapses from SOM+ neurons to PCs are strongly depressing (Ali and Thomson, 1998). This exercise resulted in 9 synaptic classes, covering all connection types in the CA1 region (Table 1.3 and Figure 1.5). Most of these classes contain few experimentally characterized examples, especially between inhibitory interneurons (Table 1.3). We have previously shown that averaging STP parameters and peak conductances within synaptic classes is a valid method to extrapolate missing values (Markram et al., 2015; Ramaswamy et al., 2015).

With the integrated and calibrated, but mostly generalized set of parameters (\hat{g} , τ_{decay} , U_{SE} , D , F parameters of STP and N_{RRP} ; Figure 1.2) for all pathways in the CA1 model we predicted the CVs of the first PSCs (Figure 1.4D) and the first PSP amplitudes (Figure 1.4E), based on previously published cell models (Migliore et al., 2018) and statistically derived connectivity. In addition, we performed *in silico* paired recordings in all possible pre-post combinations of m-type-specific pathways ($n = 102$ biologically viable pathways) to generate detailed predictions of the physiological properties of synaptic transmission including PSP amplitudes, 10-90% rise times, decay time constants, latencies, CV of first PSP amplitude, and percentage of failures (Figure 1.4F). Although these predictions could provide preliminary insights into the organizing principles of synaptic transmission in hippocampal CA1 - in particular, inhibitory pathways, which remain mostly uncharacterized - they require further validation through targeted experiments e.g. employing state-of-the-art multiple whole-cell patch-clamp recordings (Perin et al., 2011; Guzman et al., 2016; Espinoza et al., 2018).

Chapter 1. CA1 synapse physiology *in silico*

Table 1.3: Parameters and generalization to 9 classes. Synapse parameters either taken from the literature (τ_{decay} (ms)), fitted directly to data (U_{SE} , D (ms), F (ms)), calibrated *in silico* (\hat{g} (nS), N_{RRP}) or taken from the somatosensory cortex ((Markram et al., 2015) marked with †). Values in the τ_{decay} column which neither appear in Supplementary Table S1 (summarizing rat PSCs) nor are taken from the somatosensory cortex, are from mouse recordings (Daw et al., 2009; Lee et al., 2014). Average class parameters are marked in bold and are used predictively for the remaining pathways belonging to the same class. The PC to Ivy \hat{g} was not taken into account for the PC to SOM- class average. For cells belonging into the same class see Table 1.2, assumption 9. M-type abbreviations are as in Figure 1.3C.

Pre	Post	\hat{g}	τ_{decay}	U_{SE}	D	F	N_{RRP}
PC to PC (E2)							
PC	PC	0.6±0.1	3±0.2	0.5±0.02 [†]	671±17 [†]	17±5 [†]	2
PC to SOM+ (E1)							
PC	O-LM	0.8±0.05	1.7±0.14[†]	0.09±0.12[†]	138±211[†]	670±830[†]	1
PC	SOM+	0.8±0.05	1.7±0.14[†]	0.09±0.12[†]	138±211[†]	670±830[†]	1
PC to SOM- (E2)							
PC	PVBC	2±0.05	4.12±0.5	0.23±0.09	410±190	10±11	1
PC	CCKBC	3.5±0.4	4.12±0.5	0.23±0.09	410±190	10±11	1
PC	BS	1.65±0.1	4.12±0.5	0.23±0.09	410±190	10±11	1
PC	Ivy	6.5±0.5	4.12±0.5	0.23±0.09	410±190	10±11	1
PC	SOM-	2.4±0.8	4.12±0.5	0.23±0.09	410±190	10±11	1
PV+ to PC (I2)							
PVBC	PC	2.15±0.2	5.94±0.5	0.16±0.02	965±185	8.6±4.3	6
AA	PC	2.4±0.1	11.2±0.9	0.1±0.01	1278±760	10±6.7	1
BS	PC	1.6±0.1	16.1±1.1	0.13±0.03	1122±156	9.3±0.7	1
PV+	PC	2±0.35	11.1±4.1	0.13±0.03	1122±156	9.3±0.7	1
CCK+ to PC (I3)							
CCKBC	PC	1.8±0.3	9.35±1	0.16±0.04	153±120	12±3.5	1
SCA	PC	2.15±0.3	8.3±0.44	0.15±0.03	185±32	14±5.8	1
CCK+	PC	2±0.15	8.8±0.25	0.16±0.01	168±15	13±0.5	1
SOM+ to PC (I2)							
Tri	PC	1.4±0.3	7.75±0.9	0.3±0.08[†]	1250±520[†]	2±4[†]	1
SOM+	PC	1.4±0.3	8.3±2.2[†]	0.3±0.08[†]	1250±520[†]	2±4[†]	1
NOS+ to PC (I3)							
Ivy	PC	0.48±0.05	16±2.5	0.32±0.14 [†]	144±80 [†]	62±31 [†]	1
CCK- to CCK- (I2)							
PVBC	PVBC	4.5±0.3	2.67±0.13	0.26±0.05	930±360	1.6±0.6	6
PVBC	AA	4.5±0.3	2.67±0.13	0.24±0.15	1730±530	3.5±1.5	1
CCK-	CCK-	4.5±0.3	2.67±0.13	0.26±0.05	930±360	1.6±0.6	1
CCK+ to CCK+ (I1)							
CCKBC	CCKBC	4.5±0.3	4.5±0.55	0.11±0.03	115±110	1542±700	1
CCK+	CCK+	4.5±0.3	4.5±0.55	0.11±0.03	115±110	1542±700	1




Figure 1.5: Summary of synapse diversity in the CA1 network model. Panels represent exemplar *in silico* pairs from the 9 generalized pathways (2 for PC to SOM- interneurons). Presynaptic voltage traces are shown in the upper traces of each panel (A-J), while the postsynaptic potentials elicited in 35 trials (in gray) and the average of these trials are superimposed in the lower traces of each panel. Postsynaptic cells were held at -65 mV steady-state potential in *in silico* current-clamp mode. Physical dimensions are as follows: decay time constant τ_d and D , F depression and facilitation time constants: ms, peak synaptic conductance \hat{g} : nS, while the absolute release probability U_{SE} and NMDA/AMPA conductance ratios are dimensionless. **A:** PC to PC (E2). **B:** PC to O-LM cell (E1). **C:** PC to (SP) bistratified cell (E2). **D:** PC to CCKBC (E2). **E:** O-LM cell to PC (I2). **F:** CCKBC to CCKBC (I1). **G:** Ivy cell to PC (I3). **H:** CCKBC to PC (I3). **I:** PVBC to PC (I2). **J:** PVBC to PVBC (I2). Vertical scale bars on each panel represent 0.25 mV . Connectivity in the schematic CA1 microcircuit in the middle is simplified for clarity (for example most of the interneuron to interneuron connections are missing). Simplified synapses of the pathways shown in the panels around are indicated with gray circles. M-type abbreviations are as in Figure 1.3C.

1.4 Discussion

Recent advances in high-performance computing have enabled biologically detailed, data-driven reconstructions and large-scale simulations of brain regions (Bezaire and Soltesz, 2013; Bezaire et al., 2016; Markram et al., 2015; Wheeler et al., 2015). Here, we demonstrate that a data-driven workflow grounded in biological first-principles, which was used to reconstruct a biologically detailed model of rat neocortical tissue digitally, can be extended to model other brain regions such as the hippocampal CA1, to reconcile disparate cellular and synap-

tic data, and to extrapolate from the sparse set of experimentally obtained parameters to predict those of synaptic connections not yet characterized experimentally. In this study, we chose a previous implementation of the phenomenological TM model of STP, which is based on the quantal model of neurotransmitter release. The approach was able not only to extract relevant parameters from raw experimental traces, but scaled well to simulate dynamic transmission (Ramaswamy et al., 2012; Markram et al., 2015; Ramaswamy et al., 2015). In addition, this version of the TM model also enabled us to simulate trial-to-trial fluctuations to recreate, validate and predict a broad spectrum of synaptic properties for cell-type-specific hippocampal connections including amplitudes, rise and decay times, latency, variability and response failures (Figure 1.4F). It is known that $[Ca^{2+}]_o$ regulates the neurotransmitter release probability, and therefore, the amplitudes of PSPs. In this study, we adapted the existing data-driven digital reconstruction workflow to reconcile differences in synaptic dynamics that were characterized at different $[Ca^{2+}]_o$ levels. Therefore, we scaled the neurotransmitter release probabilities for all pathways that were characterized at 1.6 – 2 mM $[Ca^{2+}]_o$ (Kohus et al., 2016; Losonczy et al., 2002; Markram et al., 2015) before calibrating peak conductances to match PSP amplitudes that were measured at 2.5 mM $[Ca^{2+}]_o$, which is more representative of baseline values for hippocampal slice experiments (Ali et al., 1998; Ali and Thomson, 1998; Deuchars and Thomson, 1996; Fuentealba et al., 2008; Pawelzik et al., 1999, 2002).

In the continuing spirit of bringing together hippocampal synaptic electrophysiology from published literature a recent complementary study leveraged text-mining techniques to extract the properties of synaptic connections in hippocampal CA1, including PSP amplitudes and peak conductances (Moradi and Ascoli, 2019). The authors have also open-sourced their collection of papers and parameters alongside useful cloud-based tools to calculate reversal potentials and LJPs, of which we took advantage for this paper. However, our approach to data integration from literature demonstrates that synaptic properties reported in the literature such as peak conductances should not be interpreted at face value but require further corrections to account for inadequate space-clamp errors, which could severely underestimate their value by two-three fold (Markram et al., 2015). Furthermore, when integrating data from whole-cell patch-clamp recordings, the interaction between the extracellular bath and intracellular pipette solutions, and their influence on the kinetics of ion channel mechanisms used in the *in silico* single-cell models becomes paramount. The presence of blockers such as TTX, QX314, cesium and gluconate among many others, alter the kinetics of dendritic ion channels, which are active in the subthreshold regime, and thus, are key factors in governing the attenuation of PSPs in active dendrites. However, in our study, the core experimental dataset that was used to calibrate the peak synaptic conductances (Supplementary Tables 1.5 and 1.8) were derived exclusively from sharp-electrode recordings where the intracellular medium is devoid of any of the above blockers, and therefore, the subthreshold regime of the single-cell models are not unduly influenced. Indeed, the effects of blockers on the subthreshold regime will not only become important for future refinements of single-cell models but also when more experimental data from whole-cell patch clamp recordings are available.

The results we report, to the best of our knowledge, probably constitute a comprehensive resource, not only for the anatomy but also the kinetic and short-term dynamic physiological properties of the rat hippocampal CA1 region. Consolidation of the state of the literature not only facilitates building detailed models but also highlights knowledge gaps and could help in prioritizing the identification of missing data on CA1 connections, such as PC to interneurons, and between interneurons, which could form diverse pre-postsynaptic combinations of potential CA1 connection types that are crucial in regulating hippocampal oscillations (Klausberger and Somogyi, 2008; Pelkey et al., 2017). Our modeling approach predicts relatively high connection probabilities for interneuron to interneuron connections, and low IPSP amplitudes (see Figure 1.3D and Figure 1.4E). However, these predictions need further experimental validation, probably through multiple patch-clamp recordings, which have enabled high-throughput mapping of inhibitory circuits not only in the neocortex (Jiang et al., 2015), but also in the dentate gyrus of the hippocampal formation (Espinoza et al., 2018). Indeed, the parameter set presented here should be considered a first draft, with many assumptions and limitations. For example, we assume somatically measured PSC decay time constants for dendritic synapses without any correction for attenuation, use U_{SE} , D , F values obtained in CA3, generalize NMDA/AMPA peak ratios characterized between PCs to all other excitatory pathways, and do not model $GABA_B$ receptors. We plan to refine these assumptions systematically in future versions of our model and overcome limitations by integrating new experimental data when available (see Table 1 for all data inclusion criteria and Table 2 for all explicit limitations).

By detailing all the integration steps in this study, we had two main objectives. First, we aimed to demonstrate that published parameters should not be taken at face value without rigorously checking their consistency within any modeling framework and the necessity of being abreast of the state-of-the-art experimental techniques. Second, we attempted to emphasize the fact that a growing diversity of experimental standards combined with published literature that provides access to only processed data sets but not raw experimental traces could lead to an inconsistent picture of a fundamental mechanism such as synaptic transmission. The bottom-up modeling framework presented as a resource in this article could facilitate the integration of disparate datasets and provide a platform within which a community-driven consensus of the synaptic organization of the hippocampal formation could develop.

Acknowledgements

We would like to thank Drs. Giuseppe Chindemi, Natali Barros-Zulaica, Julian M. L. Budd, Rodrigo Perin and Zoltán Nusser for fruitful discussions as well as Werner Van Geit, Michael Gevaert, Arseniy Povolotsky, Shailesh Appukuttan, Cyrille Favreau and Marwan Abdellah for technical assistance. This study was supported by funding to the Blue Brain Project, a research center of the École polytechnique fédérale de Lausanne, from the Swiss government's ETH Board of the Swiss Federal Institutes of Technology. Funding was also provided by The Human Brain Project through the European Union Seventh Framework Program (FP7/2007-

Chapter 1. CA1 synapse physiology *in silico*

2013) under grant agreement no. 604102 (HBP) and from the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreements No. 720270 (Human Brain Project SGA1) and No. 785907 (Human Brain Project SGA2). The Blue Brain 5 supercomputer (HPE SGI 8600 system) is financed by ETH Board Funding to the Blue Brain Project as a National Research Infrastructure and hosted at the Swiss National Supercomputing Center (CSCS) in Lugano. Paired recordings and reconstructions from A.T's lab were also supported by the Medical Research Council, the Wellcome Trust and Novartis Pharma. S.S was also supported by the ÚNKP-19-3-III new national excellence program of the ministry for innovation and technology, the European Union, co-financed by the European Social Fund (EFOP-3.6.3-VEKOP-16-2017-00002).

Author contribution

S.R, A.E, and E.M conceptualized the study. S.R supervised the study. J.F and S.L reconstructed single cells in *NeuroLucida*. A.M, and A.T provided experimental datasets. S.S, S.K, and M.M optimized and validated single cell models. A.R built the CA1 circuit with inputs from all authors. A.E performed literature curation, simulations, analyses and generated the figures with inputs from S.K, M.R and S.R. A.E, M.R, and S.R wrote the manuscript with inputs from all authors.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Data availability

The data that support the findings of this study are available on request from the corresponding author.

1.5 Supplementary Methods

1.5.1 Single cell models

Detailed biophysical models of PCs and interneurons of the CA1 region from Migliore et al. (2018) were used in the present study. The models are publicly available on [ModelDB:244688](https://modeldb.yale.edu/244688) and as a "live paper" of the Human Brain Project's (HBP's) Brain Simulation Platform (BSP), where they can be interrogated in a web browser without installing anything locally. Along the excitatory PCs, they modeled 11 inhibitory interneurons.

The full list of the interneurons - and their mapping to the morphological-types (m-types) defined on hippocampome.org (Wheeler et al., 2015) are as follows: stratum lacunosum moleculare (SLM): perforant path-associated (PPA) cell - Perforant Path-Associated; stratum radiatum (SR): Schaffer collateral-associated (SCA) cell - Schaffer Collateral-Associated; stratum pyramidale (SP): axo-axonic (AA) cell - Axo-axonic, bistratified (BS) cell - Bistratified, CCK+ basket cell (BC) - Basket CCK+, Ivy cell - Ivy, PV+ BC - Basket; stratum oriens (SO): back-projection (BP) cell - Back-Projection, BS - Oriens-Bistratified, O-LM cell - O-LM, trilaminar (Tri) cell - Trilaminar (Supplementary Figure S1A). Electrical types (e-types), based on the Petilla convention (Ascoli et al., 2008) were assigned to traces recorded *in vitro* and modeled accordingly. All PCs were classified as continuous accommodating cells (cAC). Interneurons were classified as cAC, bursting accommodating cells (bAC) and continuous non-accommodating cells (cNAC). Combining m- and e-types yielded 16 morpho-electrical types (me-types) (Supplementary Figure S1C) (Markram et al., 2015; Migliore et al., 2018).

Channel kinetics were based on those used in many previously published papers on hippocampal neurons (Migliore et al., 1999, 2005; Ascoli et al., 2010; Morse et al., 2010), and validated against a number of experimental findings on CA1 pyramidal neurons. Cell models were equipped with the following active membrane properties: transient sodium current (Na); A, D, and M types and a delayed rectifier potassium currents (K_A , K_D , K_M , and K_{DR}); L, N, and T types of calcium currents (Ca_L , Ca_N and Ca_T); the nonspecific I_h ; and two types of calcium-dependent potassium currents (slow: K_{Cas} and voltage-dependent: K_{Ca}). A simple calcium extrusion mechanism, with a single exponential decay of 100 ms, was also included in all compartments containing calcium channels.

All models were constrained with active dendritic conductances but were optimized using only somatic features. While the somatic responses to various step-current injections were correct, the dendrites of the single-cell models turned out to be too excitable, namely, single synaptic inputs ($g_{syn} = 1$ nS) were leading to spikelets and somatic spikes. For this reason, single-cell models were slightly re-optimized. The amplitude of the back-propagating action potential (in the apical trunk, 150 and 250 μm from the soma) as a dendritic feature was added to the list of objectives for PCs. As for the interneurons, homogeneous dendritic sodium channel densities were replaced with one that decays exponentially with distance from the soma (with a length constant of 50 μm) based on Hu et al. (2010). A-type potassium channels in the

dendrites of interneurons were also replaced with one that activates at a more hyperpolarized potential (see kinetics of "*kad*" for distal vs. "*kap*" for proximal A-type K^+ channels in Migliore et al. (1999)). Furthermore, the upper bounds (used by the multi-objective optimization algorithm) of dendritic sodium channel densities were reduced for all cell types. See re-optimized ion channel conductance in Supplementary Figure S1 B. After the re-optimization, single cells qualitatively reproduced the behavior presented in Migliore et al. (2018), assessed by HippoUnit, our single-cell model comparison framework.

The HippoUnit package contains standardized validations of biophysically detailed, multi-compartmental single hippocampal PC models and is publicly available on Github under [/KaliLab/hippounit](#). Moreover, a demo validation of the re-optimized cell models was added as a use case in the BSP and is publicly available (again in a web browser without downloading and installing anything) upon registration to the HBP collaboratory.

1.5.2 Different versions of the Tsodyks-Markram model

The Tsodyks-Markram (TM) model of short-term plasticity (STP) underwent many changes in the last twenty years. For a recent and consistent review see Hennig (2013). Furthermore, the equations are sometimes shown in the form of differential equations (Tsodyks and Markram, 1997; Tsodyks et al., 2000; Fuhrmann et al., 2002, 2004; Loebel et al., 2009; Hennig, 2013), while in other papers the iterative solution evaluated at spike arrivals is presented (Markram et al., 1998; Maass and Markram, 2002). The version used in this article follows the formalism presented in Hennig (2013):

$$\begin{aligned}\frac{dR(t)}{dt} &= \frac{1 - R(t)}{D} - U(t)R(t)\delta(t - t_{spike}) \\ \frac{dU(t)}{dt} &= \frac{U_{SE} - U(t)}{F} + U_{SE}(1 - U(t))\delta(t - t_{spike})\end{aligned}$$

where $R(t)$ is the fraction of available resources, $U(t)$ is the release probability, D , and F are depression and facilitation time constants respectively. U_{SE} is the utilization of synaptic efficacy or absolute release probability (also known as the release probability in the absence of facilitation). $\delta(t)$ is the Dirac delta function and t_{spike} indicates the timing of a presynaptic spike. Each action potential in a train elicits an $A_{SE}U(t_{spike})R(t_{spike})$ amplitude PSC, where A_{SE} is the absolute synaptic efficacy and is linked to the Nq part of the quantal model, where N is the number of release sites and q is the quantal amplitude. $R = 1$, and $U = U_{SE}$ are assumed before the first spike. In our simulations, we implement Fuhrmann et al. (2002) as the stochastic generalization of the model. (Where the value of $U(t)$ is actually used as a *probability*.) The equation of the release probability is slightly different in that article and it reads as follows:

$$\frac{dU(t)}{dt} = -\frac{U(t)}{F} + U_{SE}(1 - U(t))\delta(t - t_{spike})$$

According to this equation $U(t)$ decays to 0 (the wording of the articles suggest a decay to "*the baseline*"). To recover the definition of U_{SE} as the release probability in absence of

spikes (or U as the constant release probability in the first Tsodyks and Markram (1997) paper concentrating only on depressing connections) the $+U_{SE}(1 - U(t))$ has to be evaluated before the release happens. On the other hand, the $-U(t)R(t)$ jump in the equation of R still has to be evaluated after the event in order to be consistent with R being 1 in the absence of spikes. In this view $U(t)$ is mostly zero and at spike arrivals, before release happens it jumps to U_{SE} . From the biophysical point of view, this can be seen as a calcium-based model, where a quick calcium influx leads to release. On the other hand, in the Hennig (2013) version $U(t)$ decays to its baseline U_{SE} value and the $U_{SE}(1 - U(t))$ jump happens after the release. When fitting the deterministic TM model to experimental data as well as when simulating the stochastic version we use an event-based solution, meaning that the equations are only evaluated at spike times (as opposed to the ODE form). For the Fuhrmann et al. (2002) version the iterative update is:

$$\begin{aligned} R_{tmp} &= 1 + (R_n - 1) \exp\left(-\frac{\Delta t}{D}\right) \\ U_{tmp} &= U_n \exp\left(-\frac{\Delta t}{F}\right) \\ U_{n+1} &= U_{tmp} + U_{SE}(1 - U_{tmp}) \\ A_{n+1} &= A_{SE} U_{n+1} R_{tmp} \\ R_{n+1} &= R_{tmp} - U_{n+1} R_{tmp} \end{aligned}$$

where Δt is the time between the $(n + 1)$ th and n th spike and A_n is the n th amplitude. On the other hand, the Hennig (2013) version (used to fit models in Kohus et al. (2016)) is:

$$\begin{aligned} R_{tmp} &= 1 + (R_n - 1) \exp\left(-\frac{\Delta t}{D}\right) \\ U_{tmp} &= U_{SE} + (U_n - U_{SE}) \exp\left(-\frac{\Delta t}{F}\right) \\ A_{n+1} &= A_{SE} U_{tmp} R_{tmp} \\ R_{n+1} &= R_{tmp} - U_{tmp} R_{tmp} \\ U_{n+1} &= U_{tmp} + U_{SE}(1 - U_{tmp}) \end{aligned}$$

None of these forms are presented in the literature per se. Both Markram et al. (1998) and Maass and Markram (2002) integrate the ODEs in a single step:

$$\begin{aligned} R_{n+1} &= 1 + (R_n - 1 - U_n R_n) \exp\left(-\frac{\Delta t}{D}\right) \\ U_{n+1} &= U_{SE} + (U_n - U_{SE} + U_{SE}(1 - U_n)) \exp\left(-\frac{\Delta t}{F}\right) \\ &= U_{SE} + U_n(1 - U_{SE}) \exp\left(-\frac{\Delta t}{F}\right) \\ &= U_n \exp\left(-\frac{\Delta t}{F}\right) + U_{SE}(1 - U_n \exp\left(-\frac{\Delta t}{F}\right)) \\ A_{n+1} &= A_{SE} U_{n+1} R_{n+1} \end{aligned}$$

Chapter 1. CA1 synapse physiology *in silico*

Using the initialization $R_1 = 1$, $U_1 = U_{SE}$ and calculating the first two amplitudes with all 3 versions (Fuhrmann et al. (2002), Hennig (2013) and Maass and Markram (2002)) one gets:

$$A_1 = A_{SE}U_{SE}$$
$$A_2 = A_{SE}[U_{SE} + (U_{SE} - U_{SE}^2)\exp(-\frac{\Delta t}{F})](1 - U_{SE}\exp(-\frac{\Delta t}{D}))$$

With simulations, it is also possible to show that all the other amplitudes in response to a spike train will be the same for all versions. Thus, the three event-based models presented above are equivalent. We present the Hennig (2013) formalism in the article since we find it more intuitive that both Dirac deltas are evaluated at the same point (after the PSC amplitude is calculated) and is more in line with the wording of the papers, but emphasize that it is consistent with the other version Fuhrmann et al. (2002) and the fits presented in Markram et al. (2015).

1.5.3 Membrane noise


In order to correctly compare the coefficient of variation (CV, std/mean) of first PSC amplitudes, measurement noise was added to the simulated traces (Barros-Zulaica et al., 2019). To this end, noise parameters of *in vitro* traces were fitted and averaged for every different connection types and then stochastic noise generated with these extracted parameters was added to the corresponding *in silico* traces. Noise was described as an Ornstein-Uhlenbeck (OU) process. The OU process is a stationary Gauss-Markov process, which describes the velocity of the movement of a Brownian particle and is used in physics to describe noise relaxation (Bibbona et al., 2008). Mathematically it can be described with the following iterative equation:

$$X(i) = X(i-1) - \frac{X(i-1)}{\tau} dt + \sigma \sqrt{\frac{2dt}{\tau}} \mathcal{N}(0,1)$$

where dt is the time step of the signal, τ is the time constant fit to the exponential decay of the signal's autocorrelation function, σ is the standard deviation of the signal and $\mathcal{N}(0,1)$ is a draw from the normal distribution.


1.6 Supplementary Figures and Tables

1.6 Supplementary Figures and Tables




Supplementary Figure 1.6: Single cell models. **A:** Exemplar 3D reconstructions of the 12 m-types in the CA1 network model. 3D morphologies were reconstructed with the NeuroLucida software by the members of the Thomson/Mercer lab (Migliore et al., 2018). Axons are shown in blue, while dendrites in red. Rendering and visualization was done with NeuroMorphoVis (Abdellah et al., 2018). Diameters are scaled (x3) for better resolution. **B:** Re-optimized ion channel conductances for all e-types (6 bAC, 13 cAC (PC), 7 cAC (IN) and 13 cNAC). Where non-uniform channel distribution was used (e.g. h current in PC dendrites) the maximal values are shown. **C:** Fraction of e-types (4) recorded and modeled in each of the 12 m-types.

Chapter 1. CA1 synapse physiology *in silico*



Supplementary Figure 1.7: PSP attenuation. Validation of PSP attenuation against experimental data from Magee and Cook (2000). **A:** EPSC like currents were injected to the apical dendrites of the different pyramidal cell models from Migliore et al. (2018) and PSPs were measured at the injection site and at the soma. **B:** Summary of PSP attenuation (dendritic PSP/somatic PSP) in all PC models injected at different distances from the soma (in blue) and comparison to experimental data (in red).



Supplementary Figure 1.8: Calcium sensitivity of synaptic physiology. **A:** PC to PC PSP amplitudes at different extracellular calcium concentrations (normalized to 2 mM). Red curve indicates the experimentally measured scaling function which was applied to scale the U_{SE} parameter of the TM model. Shaded light blue area indicates the *in vivo* range 1.1 – 1.3 mM. **B:** Same *in silico* PC to PC pair at two different extracellular calcium concentrations. *In vitro* like is shown on top, while the *in vivo* one at the bottom. Single trials ($n = 35$) are shown in gray and their average in blue. Postsynaptic cells were held at -65 mV steady-state potential in *in silico* current-clamp mode. Vertical scale bar on the bottom panel represents the same value as the one on top.

Supplementary Table 1.4: Summary of paired recording experiments from rat CA1 in voltage-clamp mode (PSCs in pA). Liquid junction potentials (LJPs) and reversal potentials (E_{rev} s) are taken from Moradi and Ascoli (2019). Holding potentials (Hold.) are corrected for the indicated LJP with the correct sign. † in the rise time constant (τ_{rise}) column indicates 20-80% rise time, instead of 10-90%. M-type abbreviations are as in Figure 1.3C.

Pre	Post	Elect.	Ampl. (pA)	τ_{rise} (ms)	τ_{decay} (ms)	Hold. (mV)	[Ca ²⁺] _i (mM)	Temp. [°C]	E_{rev} (mV)	LJP (mV)	Reference
AA	PC	patch	308±103	0.8±0.1	11.2±0.9	-75.9	3	~30	-8.35	5.9	Maccaferri et al. (2000)
BS	PC	patch	Fig6) A,B	2±0.2	16.1±1.1	-75.9	3	~30	-8.35	5.9	Maccaferri et al. (2000)
CCKBC	PC	patch	118±13	0.73±0.05	6.8±0.2	-86.1	2	33	-95.1	16.1	Neu et al. (2007)
CCKBC	PC	patch	53.7±17.2	-	-	-86	2	33	-95.2	16	Földy et al. (2007)
CCKBC	PC	patch	23.6±8.25	2±0.8	9.35±1	-66.76	2.5	33±1	-94.45	16.76	Fuentealba et al. (2008)
CCKBC	PC	patch	115.4±10.8	0.63±0.04	6.47±0.27	-72.9±0.5	2	33	-90.3	14.6	Lee et al. (2010)
Ivy	PC	patch	8±2	2.8±0.2	16±2.5	-66.76	2.5	33±1	-94.45	16.76	Fuentealba et al. (2008)
O-LM	PC	patch	26±10	6.2±0.6	20.8±1.7	-75.9	3	~30	-8.35	5.9	Maccaferri et al. (2000)
PVBC	PC	patch	43.6±17.9	-	-	-86	2	33	-95.2	16	Földy et al. (2007)
SCA	PC	patch	60.2±8.1	1.43±0.12	8.3±0.44	-72.7±0.8	2	33	-90.3	14.6	Lee et al. (2010)
O-LM	PVBC	patch	11.7±1	2.6±1.3	16.5±3.9	-75.8	2	33±1	-112.2	15.8	Elfant et al. (2008)
O-LM	SCA	patch	19.5±4.7	1.9±0.4	31.2±4.5	-75.8	2	33±1	-112.2	15.8	Elfant et al. (2008)

Supplementary Table 1.5: Summary of paired recording experiments from rat CA1 in current-clamp mode (PSPs in mV). Liquid junction potentials (LJPs) and reversal potentials (E_{rev} s) are taken from Moradi and Ascoli (2019). M-type abbreviations are as in Figure 1.3C.

Pre	Post	Elect.	Ampl. (mV)	Rise (ms)	HalfW. (ms)	SS. (mV)	[Ca ²⁺] _i (mM)	Temp. [°C]	E_{rev} (mV)	LJP (mV)	Reference
PC	PC	sharp	0.7±0.5	2.7±0.19	16.8±4.1	-69.17-72.17	2.5	34-36	-8.5	-2.17	Deuchars and Thomson (1996)
AA	PC	sharp	0.51±0.07	5±0.2	45.6±2	-55.17±1	2.5	34-36	-73	-2.17	Pawelzik et al. (1999)
BS	PC	sharp	0.86±0.55	8.5±3.6	43.9±13.9	-59.77±4.4	2.5	34-36	-73	-2.17	Pawelzik et al. (1999)
BS	PC	sharp	0.55±0.15	7.4±1.4	54.6±4.2	-58.5±0.5	2	34-36	-73	-0.6	Pawelzik et al. (2002)
BS	PC	sharp	0.8±0.6	8.4±3.2	42.1±1.7	-55.17±5.1	2.5	34-36	-73	-2.17	Fuentealba et al. (2008)
CCKBC	PC	sharp	1.17±0.44	5.4±2.5	35.5±19.5	-57.17-67.17	2.5	34-36	-73	-2.17	Ali et al. (1999)
CCKBC	PC	sharp	1.47±1.06	6±2.2	47.6±13.3	-57.17-62.17	2.5?	34-36	-73	-2.17	Thomson (2000)
CCKBC	PC	sharp	0.7±0.5	6.5±1.5	44.2±10.1	-59.2±3.3	2.5	34-36	-73	-0.6	Pawelzik et al. (2002)
Ivy	PC	sharp	0.8±0.4	2.8±0.2	54.1±13.8	-59.17±2.6	2.5	34-36	-73	-2.17	Fuentealba et al. (2008)
PVBC	PC	sharp	0.45±0.24	4.6±3.2	32.4±18†	-60.95±4.6	2	34-35	-73	-3.36	Buhl et al. (1995)
PVBC	PC	sharp	1.17±0.57	4.5±2	30.4±11.6	-57.17-67.17	2.5	34-36	-73	-2.17	Ali et al. (1999)
PVBC	PC	sharp	0.81±0.92	6.8±2.7	47.2±16.9	-59.87±3.8	2.5	34-36	-73	-2.17	Pawelzik et al. (1999)
PVBC	PC	sharp	1.12±0.74	5.1±1.8	39.5±15.2	-57.17-62.17	2.5?	34-36	-73	-2.17	Thomson (2000)
PVBC	PC	sharp	0.83±0.37	5.13±2.06	38.32±12	-59±3	2.5	34-36	-73	-0.6	Pawelzik et al. (2002)
SCA	PC	sharp	0.38	10±2.8	45±2.2	-59.1±0.5	2.5	34-36	-73	-0.6	Pawelzik et al. (2002)
Tri	PC	sharp	0.8	5.6	48.8	-59.1±0.5	2.5	34-36	-73	-0.6	Pawelzik et al. (2002)
PC	BS	sharp	3.4±3.1	1.2±0.5	7.6±2.6	-68.17	2.5	34-35	-8.5	-2.17	Ali et al. (1998)
PC	BS	sharp	0.95±0.3	1.2±0.2	10.4±1.6	-66.6±1	2.5	34-36	-8.5	-0.6	Pawelzik et al. (2002)
PC	BS	sharp	1.8±2.3	1.5±0.3	6.4±2.7	-71.67±5.7	2.5	34-36	-8.5	-2.17	Fuentealba et al. (2008)
PC	CCKBC	sharp	2±2.1	1±0.4	6.1±1.5	-67.6±3	2.5	34-36	-8.5	-0.6	Pawelzik et al. (2002)
PC	Ivy	sharp	2.9±2.2	1.5±0.3	11.5±1.5	-67.97±5.4	2.5	34-36	-8.5	-2.17	Fuentealba et al. (2008)
PC	O-LM	sharp	0.93±1.06	1.2±0.5	7.5±0.7	-72.17±2.3	2.5	34-36	-8.5	-2.17	Ali and Thomson (1998)
PC	PVBC	sharp	1.4±1.05	0.88±0.44	5.4±2.2	-68.17	2.5	34-35	-8.5	-2.17	Ali et al. (1998)
PC	PVBC	sharp	3.51±2.9	1±0.3	5.74±1.78	-67.6	2.5	34-36	-8.5	-0.6	Pawelzik et al. (2002)
BC	BS	sharp	0.37	1	5.6	-58	2	34-35	-73	-3	Cobb et al. (1997)
BC	BS	sharp	1±0.4	1.65±0.5	15.6±2.8	-65.17±4.4	2.5	34-36	-73	-2.17	Pawelzik et al. (2003)
BC	BC	sharp	0.25	1.3	27	-62	2	34-35	-73	-3	Cobb et al. (1997)
BC	BC	sharp	1.1±0.47	2.5±0.9	18.7±9.1	-61.17±4	2.5	34-36	-73	-2.17	Pawelzik et al. (2003)
BS	BC	sharp	0.7±0.4	2.5±0.8	19.1±9.5	-61.87±2.7	2.5	34-36	-73	-2.17	Pawelzik et al. (2003)
SCA	SCA	sharp	0.5	5	34.3	-58.6	2.5	34-36	-73	-0.6	Pawelzik et al. (2002)
SCA	SCA	patch	0.6±0.41	7.0±1.38	41.1±12.5	-39.5	2	20-22	-79.5	15.5	Ali (2007)

1.6 Supplementary Figures and Tables

Supplementary Table 1.6: Validation of number of synapses per connections (see Figure 1.3B). M-type abbreviations are as in Figure 1.3C.

Pre	Post	Reference data	Model	Reference
PC	PC	1.2±0.4	1.26±0.6	Deuchars and Thomson (1996)
AA	PC	6.1	7±4.4	Buhl et al. (1994b)
BS	PC	6	6.5±3.2	Buhl et al. (1994a)
CCKBC	PC	8.3±0.8	8.6±3.9	Földy et al. (2010)
O-LM	PC	10±7	11±5.2	Maccaferri et al. (2000)
PVBC	PC	11±0.6	11.3±5.4	Földy et al. (2010)
SCA	PC	5.3±1.2	5±1.8	Vida et al. (1998)
PC	O-LM	2.8±0.8	2.8±1.2	Biro et al. (2005)
PVBC	PV+	1.54±1.08	2.6±1.3	Sik et al. (1995)
SCA	SCA	3.5±1.5	3±1.4	Ali (2011)

Supplementary Table 1.7: Validation of the CV of first PSC amplitudes (see Figure 1.4B). M-type abbreviations are as in Figure 1.3C.

Pre	Post	Reference data	Model	Reference
AA	PC	0.29±0.11	0.28±0.13	Kohus et al. (2016)
CCKBC	PC	0.43±0.14	0.36±0.1	Kohus et al. (2016)
PVBC	PC	0.26±0.06	0.28±0.07	Kohus et al. (2016)
SCA	PC	0.38±0.11	0.31±0.08	Kohus et al. (2016)
CCKBC	CCKBC	0.18±0.16	0.18±0.1	Kohus et al. (2016)
PVBC	AA	0.45±0.11	0.17±0.09	Kohus et al. (2016)
PVBC	PVBC	0.17±0.05	0.22±0.02	Kohus et al. (2016)

Supplementary Table 1.8: Validation of PSP amplitudes (see Figure 1.4C). PC to CCKBC and Ivy are not shown on the figure for visualization purpose. In some cases (indicated with †) outliers were removed from the reference data (see published reference data in Supplementary Table 1.5). M-type abbreviations are as in Figure 1.3C.

Pre	Post	Reference data (mV)	Model (mV)	Reference
PC	PC	0.7±0.5	0.68±0.43	Deuchars and Thomson (1996)
AA	PC	0.51±0.07	0.51±0.21	Pawelzik et al. (1999)
BS	PC	0.55±0.15	0.55±0.24	Pawelzik et al. (2002)
CCKBC	PC	0.7±0.5	0.68±0.26	Pawelzik et al. (2002)
Ivy	PC	0.8±0.4	0.82±0.35	Fuentealba et al. (2008)
PVBC	PC	0.83±0.37	0.83±0.23	Pawelzik et al. (2002)
SCA	PC	0.38	0.39±0.17	Pawelzik et al. (2002)
Tri	PC	0.8	0.81±0.36	Pawelzik et al. (2002)
PC	BS	0.95±0.3	0.96±0.54	Pawelzik et al. (2002)
PC	CCKBC	2±2.1	1.85±0.67	Pawelzik et al. (2002)
PC	Ivy	2.9±2.2	2.65±2	Fuentealba et al. (2008)
PC	O-LM	0.3±0.13 [†]	0.3±0.21	Ali and Thomson (1998)
PC	PVBC	1±0.4 [†]	1±0.75	Ali et al. (1998)
(PV)BC	(PV)BC	0.25	0.25±0.15	Cobb et al. (1997)

2 Cortical synapse physiology *in silico*

The data presented in this chapter is taken from our preprint: J.B. Isbister*, **A. Ecker***, C. Pokorny*, S. Bolaños-Puchet*, D. Egas Standander* et al. (2023) *Modeling and Simulation of Neocortical Micro- and Mesocircuitry. Part II: Physiology and Experimentation*. bioRxiv; doi: [10.1101/2023.05.17.541168](https://doi.org/10.1101/2023.05.17.541168)

Contribution: As the preprint covers diverse topics and has 40 authors, I did not include the entire manuscript, but cut out the parts I have contributed most to. (A detailed author contribution can be found at the end of the preprint.) Technically, I applied the pipeline presented in Chapter 1 to the synapses of an other brain region. This included a literature review, synapse physiology parameter calibration and validation, creation of the figures (except the first one shown in this Chapter) and tables and writing the first version of the corresponding sections in the manuscript. Since the pipeline is motivated, described, and discussed in the previous Chapter, this one only has a Results section (and a short one with the additional methodology). The Abstract is not the preprint's but newly written to reflect the aboves.

Abstract

Alongside their many strengths, detailed, large-scale models are excellent tools for integrating data from different sources. As the flow of experimental data is continuous, data integration has to be a regular exercise as well. Therefore, when building the latest version of our biophysically detailed, multi-scale model of the rat non-barrel primary somatosensory cortex, we took into account recently published experimental results about synaptic physiology and the propagation of activity across cortical layers. This process allowed us to update some of our general synapse parameters to pathway-specific ones and to provide a high-level validation of the anatomy and physiology of both the thalamocortical and the recurrent cortico-cortical synapses. By recreating a recent laboratory experiment *in silico*, we not only validated the flow of thalamus-evoked activity from layer 4 to layer 2/3 in our circuit, but also went beyond the

original experiment and provided predictions about how the other cortical layers influence the subthreshold activity of layer 2/3 excitatory cells. In summary, as many other elements of the new version of the model, the physiology of its synapses also improved. Furthermore, inspired by recent experiments, we set up a paradigm which led to new predictions about the interactions of cortical layers.

2.1 Results

The workflow for determining a dense parameter set for all synaptic pathways, starting with sparse data from the literature, is described for the use case of hippocampal CA1 in Chapter 1 (Figure 1.1).

2.1.1 Cortico-cortical synapse physiology

To parametrize the cortico-cortical synapses with the non-barrel primary somatosensory cortex (nbS1; Figure 2.1), we enriched the paired-recording data sources used in Markram et al. (2015) with recently published ones. To further constrain the variance of excitatory postsynaptic potential (EPSP) amplitudes in layer 5 thick-tufted pyramidal cells (L5 TTPCs), we used parameters from Barros-Zulaica et al. (2019). Furthermore, the dataset was also enriched with recent recordings from L6 (Qi and Feldmeyer, 2016; Yang et al., 2020, 2022). Compared to the hexagonal bounding boxes used in Markram et al. (2015), we built the new version of the model in an atlas-based manner (Figure 2.1; Reimann et al., 2022a). In order to follow the curvature of the cortical atlas, morphologies had to be tilted, which shifted the locations of synapses between them to an unknown degree. Furthermore, single cell models were re-optimized with additional ion channels (Figure 2.1 step 2, Supplementary Figure 4.7A for exemplary conductance densities; Reva et al., 2022). Although, the propagation of EPSPs along the dendrites of L5 TTPCs were validated as part of the re-optimization process (Supplementary Figure 4.7B2), PSP attenuation in other cell types might have changed. As somatic PSP amplitudes depend not only on the calibrated peak synaptic conductances, but on the location of synapses and on the physiology of the postsynaptic cells as well (Figure 1.2) not all synapse parameters from Markram et al. (2015) could be used as they were, but had to be re-adjusted (Figure 2.1 step 3). The resulting pathway-specific parameters are listed in Tables 2.1, 2.2, and 2.3, the most common short-term dynamics are depicted in Figure 2.2A1-2, and the assignment of STP profiles to different pathways are shown in Figure 2.2A3. PSP amplitudes and their coefficient of variation (CV; std/mean) closely matched their biological counterparts ($r = 0.99$, $n = 27$; Figure 2.2B1; Supplementary Table 2.4 and $r = 0.63$, $n = 10$; Figure 2.2C1; Supplementary Table 2.5, respectively). The dense parameter set also allowed prediction of PSP amplitudes and CVs for all cortical pathways (Figure 2.2B2, C2). The frequency of miniature postsynaptic currents (mPSCs) were also in line with *in vitro* measurements ($r = 0.92$, $n = 5$; Figure 2.2D; Supplementary Table 2.6; see Methods).

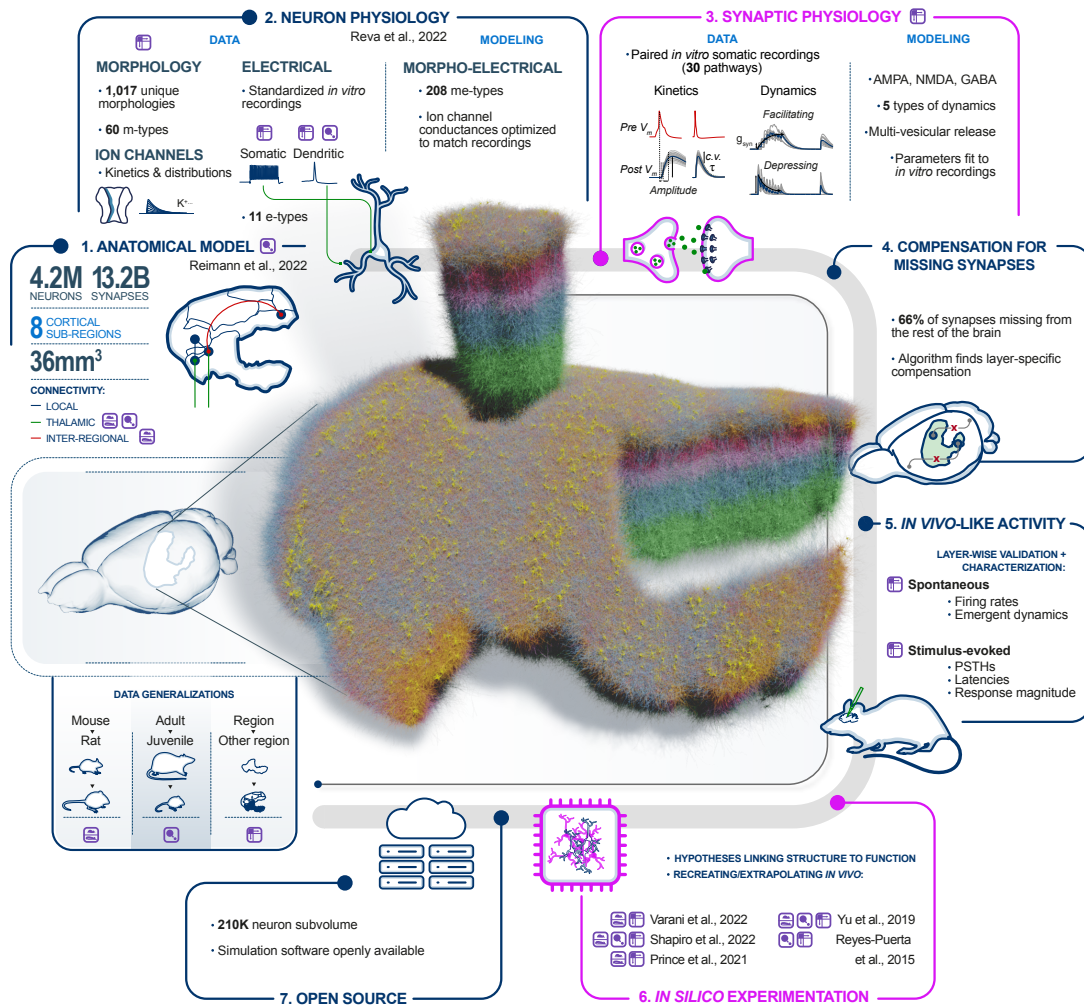


Figure 2.1: Overview of the physiology workflow of the cortical network model: The rendering of the whole nbS1 in the middle was done in Brayns. The "pulled out" seven column subvolume corresponds in size to our previous model (Markram et al., 2015). The boxes around it schematize the steps of the physiology workflow. Blue boxes correspond to the whole preprint (Isbister et al., 2023), while the pink ones to the parts that are described in detail in this Chapter. **1. Anatomical model:** Summary of the anatomical nbS1 model described in Reimann et al. (2022a). **2. Neuron physiology:** Neurons were modeled as multi-compartment models with ion channel densities optimised using previously established methods and data from somatic and dendritic recordings of membrane potentials *in vitro* (Reva et al., 2022). **3. Synaptic physiology:** Models of synapses were built using previously established methods and data from paired recordings *in vitro* (Markram et al., 2015; Ecker et al., 2020). **4. Compensation for missing synapses:** Excitatory synapses originating from outside nbS1 were compensated with noisy somatic conductance injection, parameterized by a novel algorithm. **5. In vivo-like activity:** We calibrated an *in silico* activity regime compatible with *in vivo* spontaneous and stimulus-evoked activity. **6. In silico experimentation:** Five laboratory experiments were recreated. Two were used for calibration and three of them were extended beyond their original scope. **7. Open Source:** Simulation software and a seven column subvolume of the model are available on Zenodo. **Data generalisations:** Three data generalisation strategies were employed to obtain the required data. Left: Mouse to rat, middle: Adult to juvenile rat, right: Hindlimb and barrel field subregions to the whole nbS1. Corresponding purple icons throughout the figure show where these strategies were used.

Chapter 2. Cortical synapse physiology *in silico*




Figure 2.2: *In silico* synapse physiology. **A:** Exemplary pair of L5 TTPCs (visualized with NeuroMorpho-Vis (Abdellah et al., 2018)). Presynaptic cell in gray, postsynaptic cell in red, synapses between them in purple. Neurite diameters are enlarged (x3) for visibility and axons were cut to fit into the figure. Pre- and postsynaptic voltage traces on the top right. **A2:** Exemplary postsynaptic traces with different STP profiles. **A3:** Assignment of STP profiles to viable pathways. (Pathways were considered viable if there were at least 10 connections in all eight subregions of the model.) **B1:** Validation of first PSP amplitudes. Dashed gray line represents perfect correlation between experimental and model values. Error bars show one standard deviation (also for C1 and D). **B2:** Predicted PSP amplitudes of all viable pathways in the circuit. Postsynaptic cells were held at -70 mV using an *in silico* voltage-clamp. Means were calculated over 100 pairs of neurons with 35 repetitions each. **C1** and **C2:** same as B1 and B2, but showing the CV of the first PSP amplitude. **D:** Validation of mPSC frequencies.

2.1.2 Thalamocortical synapse physiology

The model includes fibers from the thalamus, based on fibres projecting to the barrel cortex from the ventral posteromedial (VPM) and posteromedial (POm) thalamic nuclei (Reimann et al., 2022a). These fibres make synaptic contacts within a radius of the fiber probabilistically based on laminar innervation profiles (Figure 2.3A). To improve the physiology of VPM

synapses, synaptic peak conductances were constrained using EPSP amplitude measurements from thalamocortical slices (Figure 2.3B; Beierlein and Connors, 2002; Beierlein et al., 2003), instead of generalizing cortico-cortical L4 excitatory parametersets as in Markram et al. (2015). Latencies of layer-wise EPSPs increased with distance from the thalamus (Figure 2.3C). Additionally, thalamocortical EPSP amplitudes normalized relative to a single population were compared to normalized EPSPs in response to optogenetic stimulation targeting bundles of thalamic fibers in mice (Sermet et al., 2019). This provided contrasting insights, however. For example, whilst VPM to L6 inhibitory EPSPs match the initial validation data (Figure 2.3B), VPM to L6 parvalbumin (PV+) interneuron responses appear too strong relative to other populations (Figure 2.3D1). The results suggest that the model's POM to L5 excitatory pathway is too weak, when compared to other POM to excitatory and all POM to PV+ pathways (Figure 1.4H2, right).




Figure 2.3: *In silico* synapse physiology. **A:** Location of synapses from VPM fibers (purple) and POM fibers (red) on 38 neurons (dark gray) in a $5\ \mu\text{m}$ radius column (visualized with BioExplorer). **B:** Validation of thalamocortical EPSP amplitudes (as in Figure 2.2B1). The four pathways used for the validation are marked with a black rectangle on D1 to its right. **C:** EPSP latencies (time from presynaptic spike to the rise to 5% of peak EPSP amplitude in the postsynaptic trace). **D1** Left: mean VPM evoked EPSP amplitudes on postsynaptic cell types (over 50 pairs). Right: Comparison of normalized *in silico* amplitudes (normalized by L4 excitatory as in Sermet et al., 2019) to *in vitro* reference data from Sermet et al. (2019). Heatmap shows model minus reference values, thus positive values indicate a higher normalized EPSP amplitude in our model than in the reference experimental dataset. **D2:** same as D1 but for POM (normalized by L5 excitatory as in Sermet et al., 2019).

Chapter 2. Cortical synapse physiology *in silico*

Table 2.1: Excitatory synaptic pathways. Average class parameters are marked in bold and are used predictively (in lack of reference *in vitro* data) for the remaining pathways belonging to the same class. Physical dimensions are as follows: peak conductance \hat{g} : nS, depression and facilitation time constants D and F : ms, the release probability U_{SE} and the average number of vesicles in the release-ready pool N_{RRP} are dimensionless. All excitatory synaptic currents have a decay time constant of 1.74 ± 0.18 ms. (Proximal Targeting (PT) inhibitory mtypes: N/L/SBC, CHC. N/LBC etypes: cACint, cIR, bAC, bIR, cNAC have E1, while the rest of etypes; and SBC and CHC mtypes have E2 short-term dynamics. Distal Targeting (DT) inhibitory mtypes: MC, BP, DBC, BTC, of which MCs and the cACint e-type of BTC and DBC mtypes are Sst+. L1_GABAB- comprise all non NGC mtypes in L1)

Pre	Post	\hat{g}	U_{SE}	D	F	N_{RRP}
PC to Sst+IN (E1)						
PC	MC	0.2±0.1	0.09±0.12	138±211	670±830	1.5
PC	PT	0.6±0.5	0.02±0.00	194±18	507±37	4.5
PC	DT	0.2±0.1	0.02±0.00	194±18	507±37	1.5
PC to PC (E2)						
L23_PC	L23_PC	1.0±0.5	0.46±0.26	671±17	17±5	2.6
L4_PC	L4_PC	0.6±0.3	0.86±0.09	671±17	17±5	1.0
L4_SSC	L23_PC	0.2±0.1	0.79±0.04	671±17	17±5	1.8
L5_STPC	L5_STPC	0.9±0.3	0.39±0.03	690±90	44±21	1.0
L5_TTPC	L5_TTPC	1.9±1.0	0.38±0.10	365±100	25±45	2.8
L23_PC	L5_TTPC	0.5±0.2	0.50±0.02	671±17	17±5	1.5
L4_SSC	L5_STPC	0.6±0.3	0.50±0.02	671±17	17±5	1.2
L4_SSC	L6_PC	0.4±0.2	0.50±0.02	671±17	17±5	1.0
L6_TPC:A	L6_TPC:A	1.0±0.5	0.37±0.11	280±90	90±80	1.0
L6_TPC:C	L6_TPC:C	0.5±0.2	0.23±0.06	420±340	200±130	1.0
L6_IPC	L6_IPC	0.9±0.3	0.23±0.06	420±340	200±130	1.0
L6_PC to same L6_PC		0.8±0.2	0.23±0.06	420±340	200±130	1.0
L6_TPC:A	L6_TPC:C	1.2±0.5	0.23±0.06	420±340	200±130	1.0
L6_TPC:A	L6_BPC	0.3±0.1	0.23±0.06	420±340	200±130	1.0
L6_TPC:C	L6_IPC	0.2±0.1	0.23±0.06	420±340	200±130	1.0
L6_IPC	L6_BPC	0.4±0.1	0.23±0.06	420±340	200±130	1.0
L6_BPC	L6_TPC:A	0.2±0.1	0.23±0.06	420±340	200±130	1.0
L6_PC to diff. L6_PC		0.5±0.4	0.23±0.06	420±340	200±130	1.0
PC	PC	0.7±0.4	0.50±0.02	671±17	17±5	1.5
PC to Sst-IN (E2)						
PC	NBC	0.6±0.4	0.72±0.12	227±70	13±24	4.5
PC	PT	0.6±0.5	0.50±0.02	671±17	17±5	4.5
PC	L1_GABAB-	0.3±0.1	0.50±0.02	671±17	17±5	1.5
L6_TPC:A	L6_BC	0.4±0.1	0.58±0.13	240±80	70±90	1.5
L6_TPC:C	L6_BC	0.4±0.1	0.36±0.21	380±310	280±340	1.5
L6_IPC	L6_BC	0.3±0.1	0.51±0.20	440±300	100±50	1.5
L6_PC	L6_BC	0.4±0.1	0.47±0.21	370±290	155±215	1.5
PC	IN	0.4±0.1	0.50±0.02	671±17	17±5	1.5

2.1 Results

Table 2.2: Inhibitory synaptic pathways. Bold entries and physical dimensions as in Table 2.1 + τ_{decay} : ms. (Proximal Targeting (PT) inhibitory mtypes: N/L/SBC, CHC. N/LBC etypes: cNAC, dSTUT, cSTUT, bSTUT have I3, while the rest of etypes; and SBC (except: cACint and dNAC) and CHC mtypes have I2 short-term dynamics. DT and L1_GABAB- as in Table 2.1.

Pre	Post	\hat{g}	U_{SE}	D	F	N_{RRP}	τ_{decay}
IN to PC (I1)							
L6_BC	L6_BC	2.3±0.5	0.16±0.10	45±21	376±253	1.0	10.40±6.20
SBC (cACint)	PC	1.9±1.0	0.16±0.10	45±21	376±253	3.3	10.40±6.20
IN to PC and IN (I2)							
MC	PC	3.0±1.5	0.30±0.08	1250±520	2±4	1.0	8.30±2.20
DT	PC	3.0±1.5	0.25±0.13	706±405	21±9	1.0	8.30±2.20
NBC	PC	1.9±1.0	0.14±0.05	875±285	22±5	3.3	8.30±2.70
NGC	PC	0.2±0.1	0.25±0.13	706±405	21±9	1.0	36.50±1.30
L1_GABAB-	PC	0.3±0.1	0.25±0.13	706±405	21±9	1.0	8.30±2.20
SBC (dNAC)	PC	1.9±1.0	0.25±0.13	706±405	21±9	3.3	8.30±2.20
IN	*	2.3±0.5	0.25±0.13	706±405	21±9	1.0	8.30±2.20
IN to PC (I3)							
L6_BC	L6_PC	1.9±1.0	0.44±0.25	195±190	200±320	1.0	10.40±6.20
PT	PC	1.9±1.0	0.32±0.14	144±80	62±31	3.3	6.40±1.70

Table 2.3: Thalamocortical synaptic pathways. Values taken from the internal connectivity (Table 2.1) are marked in bold. Physical dimensions are the same as in Table 2.1. Sst+ inhibitory mtypes: MC and DBC, BTC (cACint etype only). PV+ inhibitory mtypes: N/LBC, CHC. 5HT3aR+ mtypes: the rest of the mtypes not listed above (e.g. NGC, SBC, and everything in L1).

Pre	Post	\hat{g}	U_{SE}	D	F	N_{RRP}
VPM, POm to Sst+IN (E1)						
*	Sst+	0.2±0.1	0.09±0.12	138±211	670±830	1.5
VPM, POm to PC (E2)						
VPM	L23_PC	1.7±0.6	0.75±0.1	671±17	17±5	1.5
VPM	L4_PC	1.1±0.4	0.75±0.1	671±17	17±5	1.5
VPM	L56_PC	2.4±0.9	0.75±0.1	671±17	17±5	1.5
POm	PC	1.7±0.6	0.75±0.1	671±17	17±5	1.5
VPM, POm to PV+IN (E2)						
VPM	L4_PV+	1.4±0.4	0.72±0.12	227±70	13±24	4.5
VPM	L6_PV+	3.1±1.0	0.72±0.12	227±70	13±24	4.5
VPM	L235_PV+	2.2±0.4	0.72±0.12	227±70	13±24	4.5
POm	PV+	2.2±0.4	0.72±0.12	227±70	13±24	4.5
VPM, POm to 5HT3aR+IN (E2)						
*	5HT3aR+	0.4±0.1	0.50±0.02	671±17	17±5	1.5

2.1.3 L4's contribution to L2/3 responses during whisker hold stimulus

In the canonical model of the cortex (reviewed e.g., in Lübke and Feldmeyer, 2007; Feldmeyer, 2012) information from the thalamus arrives to L4, then propagates to L2/3, from there to L5 (which serves as the main cortico-cortical output layer) and lastly to L6 (which projects back to the thalamus). The coordinated action of all layers of the primary somatosensory cortex is required for high-level behavioral tasks (Park et al., 2020). As the canonical model is based on the highest layer-wise density of axons, it cannot describe all interactions in the cortex. For example VPM does not only innervate L4, but also the border of L5 and L6, as well as the bottom of L2/3 (Figure 2.3A; Meyer et al., 2010; Constantinople and Bruno, 2013; Sermet et al., 2019).

To study how L4 contributes to the stimulus preference of L2/3 PCs, Varani et al. (2022) used optogenetic inactivation of L4 PCs during whisker stimulation and quantified the changes in the subthreshold response of L2/3 PCs. They found, that the early phase of the subthreshold response significantly differed from the control condition, if the whisker was deflected in either the most or the least preferred direction (see the top and bottom rows of their Figure 5B, C). From this they concluded that both L4 and VPM contribute to the direction tuning of L2/3 PCs. After reproducing their experimental conditions *in silico* (Figure 2.1 step 6; Figure 2.4A-D; see Methods) we confirmed that we can reproduce their results, i.e., subthreshold responses of L2/3 PCs decreased, when L4 PCs were inhibited (Figure 2.4E for preferred direction whisker stimulation; see Methods). This exercise can also be seen as a high-level, independent validation of the model's synaptic anatomy and physiology.

We then leveraged our *in silico* setup to study what Varani et al. (2022) could not, because of methodological limitations. In our reading, the authors aimed to test how direct excitatory connections from L4 PCs to L2/3 PCs influence the stimulus representation in L2/3. This connection can not specifically be blocked *in vivo*, instead (95% of) the L4 PC population is inhibited (as well as some lower L3 PCs). In our setup we could selectively block the connection and found almost the same result (compare Figure 2.4E and F, left). This extends the conclusion of Varani et al. (2022): L4 PCs contribute to the stimulus preference of L2/3 PCs via direct excitatory connections, and not via disinaptic inhibition.

The authors also discussed studying L5 PCs' contribution to L2/3 responses (as a large fraction of L5 PC axons terminate in L2/3), but this is infeasible with current mouse lines. Leveraging our model, we found that L5 contributes much less to subthreshold L2/3 traces than L4 (Figure 2.4F, right). Extending to other presynaptic layers, we found that the contribution of L2/3 is similar to that of L4, whereas inputs from L6 are negligible (Figure 2.4F, right). Whilst mouse lines targeting L5 PCs might arrive soon (which could validate our predictions), blocking L2/3 connections between L2/3 cells without hyperpolarizing the same L2/3 population seems only achievable *in silico*.




Figure 2.4: Reproducing and extending Varani et al. (2022) *in silico*. **A:** Schematics of whisker kinetics and VPM fiber rates during 500 ms long whisker hold stimulus. Fraction of VPM fibers coding for each kinetic feature are taken from Petersen et al. (2008). **B:** Mimicking the effect of activation of the Halo inhibitory opsin *in silico*. Injected somatic hyperpolarizing current mimicking opsin activation (top), and the resulting somatic voltage trace from a combination of injected conductance, current, and synaptic PSPs from the network (bottom). **C:** Raster plots of the microcircuit's activity and population firing rates below. **C1:** Control conditions, **C2:** 95% of L4 excitatory cells inhibited (by direct somatic current injection depicted in B above). **D:** Voltage traces of all L2/3 (top) and all L4 (bottom) excitatory cells. Panels show spiking traces (top), and subthreshold traces (bottom). **D1** and **D2** depict the same conditions as C above. **E:** Comparison of average traces from selected L2/3 PCs in control (black) and optogenetically inhibited (green) conditions. **F:** Same as E, but instead of mimicking the optogenetic inhibition of L4 excitatory cells, only the connections to L2/3 PCs are "cut" (compare inset with the one in E). The right part depicts connections systematically cut from PCs in all layers, while the left shows L4 only for a better visual comparison with the conditions of Varani et al. (2022) in E.

2.2 Methods

2.2.1 Parametrization of spontaneous vesicle release

As an additional source of biorealistic variability to the network simulations, single vesicles spontaneously released with a low frequency. To calibrate the mPSC frequencies of different pathways, single cell simulations with different values of the spontaneous release frequency for all synapses of a set of 1000 cells in a given pathway were run. In these simulations, *in silico* voltage-clamp recordings were performed to measure the resulting mPSC frequency at the soma. This data was then fitted with a logarithmic function and the value of the spontaneous release frequency matching the *in vitro* reference value for the mPSC frequency was interpolated. As *in vitro* paired recording data is sparse, all available sources to determine synaptic parameters were re-used for validation.

2.2.2 Reproducing Varani et al. (2022) *in silico*

To study how input from L4 contributes to L2/3 subthreshold responses Varani et al. (2022) used a 500 ms long whisker hold paradigm, while patch-clamping PCs in L2/3 in anesthetized and awake mice. The whisker hold stimulus was encoded as a step function (1, if $2000 \text{ ms} \leq t < 2500 \text{ ms}$, 0 otherwise) in 10% of the VPM fibers. One of four transfer functions were assigned to each fiber, based on the types of kinetic response properties of thalamic neurons identified in (Petersen et al., 2008). The types were selective for whisker position (v_{pos}), velocity (v_{vel}), acceleration (v_{acc}), or direction (v_{dir}), and were implemented as:

$$\begin{aligned}v_{pos}(t) &= r_{max} \cdot x(t) \\v_{vel}(t) &= r_{max} \cdot (x(t+1) - x(t)) \\v_{acc}(t) &= r_{max} \cdot (x(t+1) - 2x(t) + x(t-1)) \\v_{dir}(t) &= r_{max} \cdot |x(t+1) - x(t)|_+\end{aligned}\tag{2.1}$$

where $r_{max} = 150 \text{ Hz}$ is the firing rate of a thalamic fiber when its associated feature property is at the fiber's preferred value and $x(t)$ is the whisker position. Transfer functions were randomly assigned to fibers with the fractions identified in Petersen et al. (2008) (Figure 2.4A, 11% coding for position and acceleration, 58% for velocity and 20% for direction). The spiking process was an adapting Markov process (Muller et al., 2007) with an adaptation time constant of 100 ms evaluated at every 0.1 ms.

The optogenetic inhibition in Varani et al. (2022) targeted 95% of excitatory cells in L4. The authors found a few cells which also tested positive for the inhibitory opsin Halo at the bottom of L3 as well, but as they did not quantify it, lower L3 PCs were not targeted in the *in silico* version of the experiment. Optogenetic inhibition of the target L4 excitatory population was modeled through a current injection at the soma of these cells, with an intensity proportional to the cell's threshold current (see Reva et al., 2022). To mimic the conditions of surface

illumination, the dependence of effective depolarization strength on cortical depth using a modified Beer-Lambert law approximation was considered (Al-Juboory et al., 2013; Azimipour et al., 2014):

$$I(d) = I_0 \exp(-\mu_{eff} d) \quad (2.2)$$

where $I(d)$ describes the light intensity at depth d (in mm), with a maximum light intensity I_0 (on the surface of the cortex) and an effective attenuation coefficient given by:

$$\mu_{eff} = \sqrt{3\mu_a(\mu_a + \mu'_s)} \quad (2.3)$$

Based on the 595 nm wavelength (yellow light) the absorption coefficient $\mu_a \approx 0.49 \text{ mm}^{-1}$ and reduced scattering coefficient $\mu'_s \approx 4.12 \text{ mm}^{-1}$ were used (Mesradi et al., 2013). L4 excitatory cells were binned into 5 depth bins and for all cells belonging to the given bins the light intensity at the center of the bin was used. After scanning several values, I_0 was set to -200% as that reproduced the $\approx 10 \text{ mV}$ hyperpolarization of L4 PCs observed *in vivo*. In line with the *in vivo* experiment, the optogenetic stimulus ended in a (100 ms long) ramp to avoid rebound spikes (Figure 2.4B).

When going beyond reproducing the same experimental conditions and instead leveraging the *in silico* nature of our setup, synaptic pathways were lesioned by selecting the excitatory population in a given layer as the presynaptic population and the excitatory population in L2/3 as the postsynaptic population and not instantiating the connecting synapses during the simulation.

L2/3 PCs had to meet three criteria to be included in the subsequent analysis. Firstly, their activity was required to remain subthreshold during the 500 ms long whisker hold stimulus and in 200 ms long time windows before and after the stimulus, both in control and *in silico* optogenetic runs. Second, they had to be innervated by at least one (active) VPM fiber. Third, the derivative of their voltage trace had to cross the 1 mV/ms threshold in a 20 ms time window after stimulus onset in the control simulation. The last two were motivated by comparing subthreshold voltages to voltage traces from Varani et al. (2022) that showed large, stimulus evoked EPSPs. Around 8% of L2/3 PCs in the central column met all the above criteria and their voltages were averaged to arrive to the traces shown in Figure 2.4E-F. Thus, unlike in the original analysis, cells rather than trials were averaged. The motivation for this approach is that while *in vivo* it is easier to repeat the same paradigm after establishing stable recording conditions in a given cell, *in silico* it is quicker to record from all cells in a single simulation, instead of repeating the stimulus several times.

2.3 Supplementary Tables

Chapter 2. Cortical synapse physiology *in silico*

Supplementary Table 2.4: Validation of PSP amplitudes. Thick-tufted (TT) mtypes: L5_TPC:A, L5_TPC:B. Slender-tufted (ST) mtypes: L5_TPC:C, L5_UPC.

Pre	Post	<i>in vitro</i> (mV)	<i>in silico</i> (mV)	Reference
L23_PC	L23_PC	1.00±0.70	0.99±0.67	Feldmeyer et al. (2006)
L23_PC	L5_TTPC	0.30±0.30	0.30±0.24	Reyes and Sakmann (1999)
L4_EXC	L4_EXC	1.59±1.51	1.62±1.31	Feldmeyer et al. (1999)
L4_SSC	L23_PC	0.70±0.60	0.66±0.34	Feldmeyer et al. (2002)
L4_SSC	L5_STPC	0.60±0.40	0.59±0.33	Feldmeyer et al. (2005)
L4_SSC	L6_PC	0.29±0.16	0.30±0.30	Qi and Feldmeyer (2016)
L5_TTPC	L5_TTPC	1.30±1.10	1.24±0.73	Markram et al. (1997a)
L5_STPC	L5_STPC	0.80±0.20	0.75±0.40	Le Bé et al. (2007)
L6_BPC	L6_TPC:A	0.21±0.00	0.22±0.15	Berger (2009)
L6_IPC	L6_BPC	0.42±0.18	0.42±0.21	Berger (2009)
L6_IPC	L6_IPC	1.05±0.31	1.09±0.81	Berger (2009)
L6_TPC:C	L6_IPC	0.18±0.00	0.19±0.13	Berger (2009)
L6_TPC:C	L6_TPC:C	0.43±0.22	0.43±0.29	Berger (2009)
L6_TPC:A	L6_BPC	0.32±0.27	0.31±0.17	Berger (2009)
L6_TPC:A	L6_TPC:C	1.19±0.15	1.10±0.63	Berger (2009)
L6_TPC:A	L6_TPC:A	1.51±0.98	1.45±1.11	Berger (2009)
L23_PC	L1_GABAB-	1.10±0.30	1.09±0.83	Wozny and Williams (2011)
L4_EXC	L4_FS	2.20±2.20	2.17±2.46	Beierlein et al. (2003)
L5_TTPC	L5_MC	0.28±0.30	0.28±0.33	Silberberg and Markram (2007)
L6_IPC	L6_BC	1.59±1.60	1.56±1.48	Berger (2009)
L6_TPC:A	L6_BC	2.20±3.28	2.02±1.44	Berger (2009)
L6_TPC:C	L6_BC	1.29±1.65	1.28±0.84	Berger (2009)
L6_PC	L6_MC	0.20±0.12	0.15±0.14	Berger (2009)
L1_NGC	L23_PC	0.58±0.10	0.54±0.41	Wozny and Williams (2011)
L1_GABAB-	L23_PC	0.27±0.04	0.26±0.13	Wozny and Williams (2011)
L4_FS	L4_EXC	1.10±0.80	1.14±0.81	Beierlein et al. (2003)
L5_MC	L5_TTPC	0.50±0.40	0.47±0.25	Silberberg and Markram (2007)
VPM	L4_EXC	2.40±2.00	2.51±2.10	Beierlein et al. (2003)
VPM	L4_FS	4.10±3.20	4.09±2.52	Beierlein et al. (2003)
VPM	L6_EXC	1.20±0.80	1.28±1.96	Beierlein and Connors (2002)
VPM	L6_FS	3.90±3.50	3.11±2.57	Beierlein and Connors (2002)

2.3 Supplementary Tables

Supplementary Table 2.5: Validation of first PSP amplitudes' CVs. (TT and ST as in Supplementary Table 2.4)

Pre	Post	<i>in vitro</i> (mV)	<i>in silico</i> (mV)	Reference
L4_EXC	L4_EXC	0.37±0.16	0.25±0.06	Feldmeyer et al. (1999)
L4_EXC	L4_FS	0.27±0.13	0.38±0.25	Beierlein et al. (2003)
L4_FS	L4_EXC	0.25±0.11	0.28±0.09	Beierlein et al. (2003)
L4_SS	L5_STPC	0.33±0.20	0.43±0.06	Feldmeyer et al. (2005)
L4_SS	L6_PC	0.50±0.11	0.51±0.06	Qi and Feldmeyer (2016)
L4_SS	L23_PC	0.27±0.13	0.32±0.07	Feldmeyer et al. (2002)
L5_TTPC	L5_TTPC	0.31±0.14	0.39±0.09	Barros-Zulaica et al. (2019)
L5_STPC	L5_STPC	0.58±0.24	0.51±0.06	Le Bé et al. (2007)
L23_PC	L23_PC	0.33±0.18	0.43±0.15	Feldmeyer et al. (2006)
L234_PC	L234_NBC	0.32±0.08	0.21±0.08	Wang et al. (2002)

Supplementary Table 2.6: Validation of mPSC frequency. (Cortico-cortical (CC) mtypes: L6_UPC, L6_IPC, L6_HPC. Cortico-thalamic (CT) mtypes: L6_TPC:A, L6_TPC:C.)

Pre	Post	<i>in vitro</i> (Hz)	<i>in silico</i> (Hz)	Reference
E	L23_PC	8.20±2.90	9.36±4.38	Brasier and Feldman (2008)
E	L4_PC	11.90±2.40	15.64±7.84	Brasier and Feldman (2008)
E	L6_CC	2.80±0.80	3.87±2.14	Yang et al. (2020)
E	L6_CT	0.95±0.36	1.41±0.74	Yang et al. (2020)
I	L5_PC	21.10±4.80	16.06±6.74	Ling and Benardo (1999)

3 Cell assemblies and their underlying connectivity

This chapter is indential to our preprint: **A. Ecker***, D. Egas Standander*, S. Bolaños-Puchet, J.B. Isbister, M.W. Reimann (2023) *Cortical cell assemblies and their underlying connectivity: an in silico study*. bioRxiv; doi: [10.1101/2023.02.24.529863](https://doi.org/10.1101/2023.02.24.529863)

Contribution: I came up with the concept of detecting assemblies in our simulations after reviewing the literature, set up and ran all the simulations, wrote the biggest part of the open source analysis package, created all the figures and the open source data set, and participated in writing the manuscript. (A more detailed author contribution can be found at the end of the chapter.)

Abstract

Recent developments in experimental techniques have enabled simultaneous recordings from thousands of neurons, enabling the study of functional cell assemblies. However, determining the patterns of synaptic connectivity giving rise to these assemblies remains challenging. To address this, we developed a complementary, simulation-based approach, using a detailed, large-scale cortical network model. Using a combination of established methods we detected functional cell assemblies from the stimulus-evoked spiking activity of 186,665 neurons. We studied how the structure of synaptic connectivity underlies assembly composition, quantifying the effects of thalamic innervation, recurrent connectivity, and the spatial arrangement of synapses on dendrites. We determined that these features reduce up to 30%, 22%, and 10% of the uncertainty of a neuron belonging to an assembly. The detected assemblies were activated in a stimulus-specific sequence and were grouped based on their position in the sequence. We found that the different groups were affected to different degrees by the structural features we considered. Additionally, connectivity was more predictive of assembly membership if its direction aligned with the temporal order of assembly activation, if it originated from strongly interconnected populations, and if synapses clustered on dendritic branches. In summary, reversing Hebb's postulate, we showed how cells that are wired together, fire together, quanti-

Chapter 3. Cell assemblies and their underlying connectivity

fyng how connectivity patterns interact to shape the emergence of assemblies. This includes a qualitative aspect of connectivity: not just the amount, but also the local structure matters; from the subcellular level in the form of dendritic clustering to the presence of specific network motifs. This connectivity-based characterization of cell assemblies creates an opportunity to study plasticity at the assembly level, and beyond strictly pairwise interactions.

Keywords: Cell assemblies, simulation, connectivity, network topology, synapse clustering

3.1 Introduction

The formulation of the cell assemblies concept goes back to Hebb (1949), who defined them loosely as "*a diffuse structure comprising cells in the cortex*". In the past 70 years, the sequential activation of groups of neurons, the Hebbian "*phase sequence*" was linked to several complex cognitive processes, reviewed in Harris (2005) and Buzsáki (2010). Hebb's idea was later paraphrased as "*cells that fire together, wire together*" (Shatz, 1992), giving it both a structural, and a functional side. In this article we will concentrate on quantifying how the cortical structure underlies its neurons' co-firing function, but linking these groups of co-active neurons to cognitive processes is outside of our scope.

Cell assembly research rejuvenated in the hippocampus field when spikes could be reliably sorted from recordings with tetrodes and therefore neurons could be grouped to co-firing ensembles (Harris et al., 2003; Dragoi and Buzsáki, 2006; Sasaki et al., 2006; Lopes-dos Santos et al., 2013). The introduction of modern *in vivo* two-photon calcium imaging into the field, with its improved scalability and stability over time, allowed Bathellier et al. (2012) and Carrillo-Reid et al. (2015) to detect cell assemblies in auditory and visual cortices, where they showed how even a small set of them can serve as a backbone for cortical coding. These, and studies that followed (Miller et al., 2014; Montijn et al., 2016; Pérez-Ortega et al., 2021) contributed greatly to our understanding of the functional role of the Hebbian cell assemblies, but they could not make claims about the patterns of synaptic connectivity they originate from, as they could only predict functional connectivity from correlations in neuronal activity, but did not have access to the underlying structural connectivity of the neurons recorded. Additionally, results based on calcium imaging are limited to the superficial layers of the cortex, missing potential assemblies in the deeper layers, which would be of great interest as they serve as the output of the cortex (Feldmeyer, 2012; Harris and Shepherd, 2015).

Early theoretical work in the field explored the potential link between memories and cells that fire and therefore wire together, concentrating on the storage and retrieval of memories in strongly recurrent networks, such as the CA3 area of the hippocampus (Hopfield, 1982). Theories evolved and improved, but modeling studies about cell assemblies still concentrate on plasticity rules underlying the learning, storage and recall of various patterns (Fusi and Abbott, 2007; Zenke et al., 2015; Krotov and Hopfield, 2016; Fauth and Van Rossum, 2019; Kossio et al., 2021; Gastaldi et al., 2021). Thus, their focus lies on how function shapes structure, with little or no emphasis on the biologically accurate aspects of structural connectivity, such

as low connection probabilities and an abundance of directed motifs (Song et al., 2005; Perin et al., 2011; Reimann et al., 2017b).

On the other hand, the perspective can be reversed: how does a more bio-realistic structural connectivity influences a neuron's membership in one or more assemblies, or on a more general level: how does structure determine function? Additionally, how does innervation from different sources, such as local connectivity and various thalamic afferents, interact to shape assembly membership? Finally, are the afferent synapses from fellow assembly neurons scattered across the dendritic tree, or clustered on single branches, employing the nonlinear computational capabilities of dendrites (Poirazi et al., 2003; Stuart and Spruston, 2015; Kastellakis and Poirazi, 2019).




Figure 3.1: Pipeline summary. **A:** Schematics and quick facts about the detailed, large-scale cortical microcircuit that was used several times before to study the relationship of cortical structure and function (Reimann et al., 2017b, 2022b; Nolte et al., 2020). **B:** Schematics of the assembly detection pipeline from the spiking activity of 186,665 excitatory neurons in the circuit. **C:** Analysis of the connectivity of cell assembly neurons. **D:** Derivation of assembly membership probability based on different features of structural connectivity.

In order to provide insights into these questions, we employed an *in silico* approach, using an improved version of the detailed, large-scale (somatosensory) cortical circuit model of Markram et al. (2015) (Figure 3.1A), simulating the activity of tens of thousands neurons in response to a stream of thalamic input patterns. In the model, we have access not only to the spiking activity of every neuron, but also to the entire connectome, including dendritic locations of synapses. We then considered the established, purely functional definition of cell assemblies as neurons that fire together more than expected. Therefore, functional assemblies across all cortical layers were detected using a combination of previously published methods (Pérez-Ortega et al., 2021; Herzog et al., 2021; Figure 3.1B). We then analyzed their underlying structural connectivity, searching for rules that could explain assembly membership (Figure 3.1C, D). This analysis of the structure-function relation could be readily applied to assemblies detected with other methods (Lopes-dos Santos et al., 2013; Isbister et al., 2021; van der Plas et al., 2023).

Chapter 3. Cell assemblies and their underlying connectivity

We found that the structural strengths of afferents from various sources explained significant portions of the uncertainty of a neuron's membership in an assembly (Figure 3.1D). Specifically, innervation from VPM (ventral posteromedial nucleus of the thalamus) explained up to 30%, and up to 10% from POm (posteromedial nucleus of the thalamus). Strength of innervation through recurrent local connectivity explained up to 22%. The relative magnitudes of these effects differed between assemblies, with assemblies active immediately upon stimulus presentation being mostly determined by the structure of the thalamic innervation, and assemblies active 50 ms after stimulus onset being mostly determined by recurrent connectivity. Additionally, the effect of innervation strength on assembly membership can be much larger if the innervating population is highly interconnected within itself, and if the synapses are tightly clustered on the targeted dendrites (Figure 3.1D). The highly non-random structure of this connectome provides a more efficient wiring; implementing ensembles of reliably co-firing neurons with fewer synaptic connections than expected by chance.

3.2 Results

3.2.1 Diverse set of assemblies can be detected from network simulations

We simulated the electrical activity of a model of 2.4 mm^3 of cortical tissue, comprising 211,712 neurons in all cortical layers in an *in vivo*-like state. The model is a version of Markram et al. (2015) with anatomical improvements outlined in Reimann et al. (2022a) and physiological ones in Isbister et al. (2023) (Figure 3.1A). We consider the activity to be *in vivo*-like, based on a comparison of the ratios of spontaneous firing rates of sub-populations, and responses to brief thalamic inputs to *in vivo* results from Reyes-Puerta et al. (2015) (as described in Isbister et al., 2023). A stream of thalamic input patterns was applied to the model (see Methods), and the neuronal responses recorded (Figure 3.2A1). The circuit reliably responded to the brief stimuli with a transient increase in firing rate. This led to a slight shift to the right of the tail of the firing rate distribution from the spontaneous state (Figure 3.2A2), in line with experiments (Wohrer et al., 2013). The stream consisted of repeated presentations of ten different input patterns in random order (Reimann et al., 2022b). We designed the stimuli as 10 patterns with varying degrees of overlap (Figure 3.2B): 4 base patterns with no overlap (A, B, C, D), 3 patterns as combinations of two of the base ones (E, F, G), 2 patterns as combinations of three of the base ones (H, I), and 1 pattern as a combination of all four base ones (J). The overlap of these patterns can also be seen through the raster plots of their corresponding VPM fibers (Figure 3.2B bottom). For example the fibers corresponding to pattern A peak when stimulus A is presented, but also with 50% of the amplitude when E is presented.




Figure 3.2: *In vivo*-like activity *in silico*. **A1:** Raster plot of the microcircuit’s activity with 628,620 spikes from 98,059 individual neurons and the population firing rates below. The y-axis shows cortical depth. (As cortical layers do not have the same cell density, the visually densest layer is not necessarily the most active - see A2 bottom.) **A2:** Single cell firing rates (in excitatory and 3 classes of inhibitory cells) and layer-wise inhibitory and excitatory population firing rates in evoked (showed in A1) and spontaneous (not shown) activity. **B:** Top: pyramid-like overlap setup of VPM patterns, then the centers of the VPM fibers in flat map space. Bottom: raster plots of VPM fibers forming each of the patterns for the stimulus stream in A1 (i.e., from pattern A at 2000 ms to pattern J at 6500 ms). On the right: same for non-specific (POM) input.

Using a combination of the algorithms of Carrillo-Reid et al. (2015) and Herzog et al. (2021), we detected functional assemblies in 125 second-long recordings of simulated neuron activity while receiving the random input stream (25 repetitions of all 10 patterns with 500 ms inter-stimulus interval, see Methods). Briefly, the assembly detection algorithm first groups neuronal activity into 20 ms time bins, and identifies those with significantly increased firing rates (Sasaki et al., 2006; Carrillo-Reid et al., 2015; see Methods; Figure 3.3A). Then, these time bins are hierarchically clustered based on the cosine similarity of their activation vector, i.e., the vector of the number of spikes fired in the time bin for each neuron (Montijn et al., 2016; Pérez-Ortega et al., 2021; Supplementary Figure 3.8A, Figure 3.3B1). The threshold for cutting the clustering tree into clusters is determined by minimizing the resulting Davies-Bouldin index (Davies and Bouldin, 1979; see Methods; Supplementary Figure 3.9 and Supplementary Figure 3.8B for lower dimensional representations). Finally, these clusters correspond to the functional assemblies, with a neuron being considered a member if its spiking activity correlates with the activity of an assembly significantly more strongly than chance level (Montijn

Chapter 3. Cell assemblies and their underlying connectivity

et al., 2016; Herzog et al., 2021). This means that in each time bin only a single assembly is considered active, but neurons can be part of several assemblies (see Methods, Figure 3.1B, Figure 3.3B2).




Figure 3.3: Cell assembly detection. **A:** Population firing rate of excitatory neurons with the determined significance threshold. **B1:** Hierarchical clustering of the cosine similarity matrix of activation vectors of significant time bins (above threshold in A, see Methods). **B2:** Clustered significant time bins ordered by patterns presented. **C:** Number and location of neurons in each cell assembly: flat map view on top, depth-profile below. **D:** Jaccard similarity of cell assemblies and number of neurons participating in different number of assemblies. **E:** Input-output map: Input distance is calculated as the Earth mover's distance of the VPM fiber locations (see Figure 3.2B and Methods), while the output distance is the (normalized) Euclidean distance of pattern evoked time bin cluster counts (counts of different colors in the matrices above in B2, Methods).

We found that assemblies were activated in all stimulus repetitions and a series of two to three assemblies remained active for 110 ± 30 ms (Figure 3.3B2). Activation probability and duration depended on stimulus identity; the stimulus associated with the strongest response elicited 1.6 times as many significant time bins than the stimulus associated with the weakest

response. Not only were individual assemblies associated with only a subset of stimuli, but they also had a well-preserved temporal order, with some of them appearing early during a stimulus, and others later. Based on this, from now on we will refer to them as *early*, *middle*, and *late assemblies*, and will order them in the figures accordingly.

An assembly comprised on average $7 \pm 3.1\%$ of the simulated excitatory neurons, with late assemblies being about 3.7 times larger than early ones (Figure 3.3C). In terms of spatial distribution of assembly neurons, the early assemblies appear more "*patchy*" (small distinguishable clusters), and by visual inspection can be mapped back to the locations of VPM fibers corresponding to the stimuli that activated them (Figures 3.2B and 3.3C top). Moreover, their layer profile mimics that of VPM fiber innervation (Reimann et al., 2022a; Meyer et al., 2010; Supplementary Figure 3.11A), indicating that these assemblies may be determined by direct thalamic innervation. On the other hand, the late assembly neurons are more evenly distributed, and cover the entire surface of the simulated circuit, even beyond the range of VPM fiber centers, and are found mostly in deeper layers of the cortex. Middle assembly neurons are somewhat in-between, both in spatial distribution and depth profile. Although we found late assemblies to be nonspecific (at the chosen clustering threshold), early and middle assemblies belonging to the same stimulus occupy similar regions in space and can be shown to have a relatively high (25%) overlap of neurons (Figure 3.3D left). This also means that single neurons belong to several assemblies (up to 7 out of 11, Figure 3.3D right). In conclusion, the time course of stimulus responses are well-preserved and reliable enough (although with some temporal jitter) to be simplified into a sequence of distinct sets of functional cell assemblies.

Although stimulus-specific, the time courses are not unique (see the responses to patterns H and I in Figure 3.3B2). Thus, we wondered to what degree is the overlap of assemblies associated with the overlap of the input patterns given as stimuli? When we compared the distances between the locations of the VPM fibers making up an input pattern, and the assembly sequences detected from the network activity, we found a significant linear trend, i.e., patterns that are close in the input space (e.g. H and I) are close in the "*output space*" defined as the counts of individual assemblies popping up for a given stimuli across repetitions (Figure 3.3E). Thus, the activation sequence of cell assemblies can be seen as a low-dimensional representation of the complex, high-dimensional activity of the circuit's response to different stimuli. The data points are highly variable for mid-sized input distances, and the linear trend gets weaker with increasing number of assemblies (Supplementary Figure 3.9C). In summary, increasing the number of assemblies by cutting the clustering tree differently improves the separation of inputs at the cost of reducing the correlation between input and output distance. As our aim was not to build an ideal decoder of input patterns, in the following steps we analyze the assemblies resulting from a clustering that minimized the Davis-Bouldin index (Supplementary Figure 3.9A2).

Assemblies detected only in the superficial layers, are reminiscent of those detected across all layers of the cortex

Due to the limitations of traditional two-photon microscopy, *in vivo* experiments in the cortex can only detect cell assemblies in layer 2/3 (L2/3) (Bathellier et al., 2012; Carrillo-Reid et al., 2015; Miller et al., 2014; Montijn et al., 2016; Pérez-Ortega et al., 2021), whereas *in silico* we detect assemblies across all layers. This begs the question: are we detecting different assemblies, or do they just cover more depth but contain the L2/3 ones? To answer this question, we used the same data and methodology as before, but restricted our analysis to L2/3 neurons only (Supplementary Figure 3.10). When detecting assemblies exclusively from the spiking activity of L2/3 pyramidal cells, we got the same results overall, but with some specific differences. First, we found significant time bins for a shorter range (up to 100 ms, which corresponds to our definition of early and middle assembly time windows). And second, stimuli could be distinguished better, e.g. stimuli H and I correspond to different early assemblies in L2/3 (compare Figure 3.3B2 and Supplementary Figure 3.10A2). We compared the L2/3 assemblies to the original full assemblies and found that the early ones can get mapped to the L2/3 ones relatively well, in terms of the Jaccard similarity of their respective intersections with L2/3 neurons, and by visual inspection of the spatial locations of assembly neurons (Supplementary Figure 3.10D). Thus, we predict that assemblies detected *in vivo* are the superficial subset of full assemblies, and may not include late assemblies with only a small fraction of neurons in superficial layers.

3.2.2 Functional assemblies are determined by structural features

It appears as if the spatial structure of the thalamic input stimuli strongly determines assembly membership. At the same time, neuronal assemblies are thought to be strongly recurrently connected (Song et al., 2005; Perin et al., 2011).

We generally observe that some features of structural connectivity can be predictive of the probability of assembly membership. Here, we consider features related to thalamic innervation, recurrent connectivity and synaptic clustering. We quantify the strength of this effect by means of a thresholded and signed version of the mutual information which we call their *normalized mutual information* and denote by nI ; this allows us to compare the individual contributions (for details, see Methods). Its value has three basic properties: First, it is positive if the probability of assembly membership increases as the value of the structural feature increases, and is negative otherwise. Second, its absolute value is one if assembly membership can be completely predicted from the structural feature and it is only defined if the mutual information between the two processes is significantly larger than for randomly shuffled controls. Third, its absolute value does not require any assumptions about the shape of the dependency (e.g., linear, monotonic, etc.) between the structural feature and the probability of assembly membership.

Thalamic innervation explains early and middle assemblies

We began by considering the effect of direct thalamic innervation. Having confirmed that strong direct thalamic innervation facilitated assembly membership (Figure 3.4A, left), we formulated the following hypothesis: pairs of neurons are more likely to belong to the same assembly if they are innervated by overlapping sets of thalamic fibers. To test this, we first consider the common thalamic indegree of a pair of neurons, i.e., the number of thalamic fibers innervating both of them. Then, for each neuron, we use its *mean common thalamic indegree* over all cells in assembly A_n , as the structural feature to predict its membership in A_n . We performed this analysis separately for innervation from the VPM and POm nuclei.

In both cases, mean common thalamic indegree with an assembly increased the probability that a neuron is part of it (Figure 3.4A, second, olive curve for common thalamic indegree with the same assembly). The *nI* of the mean common thalamic indegree and membership in the same assembly was on average 0.165 for VPM and 0.054 for POm (Figure 3.4A, right, entries along the diagonals). More specifically, 0.157 and 0.034 for early assemblies, 0.198 and 0.087 for middle, and 0.082 and 0.037 for the late assembly. In addition, cross-assembly interactions were also observed, albeit at lower levels (Figure 3.4A, right, off-diagonal entries). The effect was strongest for pairs of early and middle assemblies that responded to the same stimuli, e.g. assemblies 1, and 2, responding to pattern A.

The lower *nI* values for the late assembly are expected, as it contains many neurons in layers not directly innervated by thalamus (Figure 3.3C, Supplementary Figure 3.11A), and its activity is largely restricted to time bins in which the thalamic input is only weakly active (80 – 140 ms after onset; Figure 3.3B2 vs. 3.2A1). Interestingly, common innervation with POm has the highest *nI* for middle assemblies, which seems related to the prevalence of L5 neurons in them (2.3 times more L5 cells, than in the early ones). POm targets the upper part of L5, and more importantly L1 (Reimann et al., 2022a; Meyer et al., 2010; Supplementary Figure 3.11A), where the apical tuft dendrites of thick-tufted L5 pyramidal cells reside (Harris and Shepherd, 2015; Ramaswamy et al., 2015). The delay caused by the long synapse to soma path distances (Supplementary Figure 3.11B) may explain the importance of common POm innervation 40 – 60 ms after stimulus onset.

Having confirmed that common innervation by thalamic fibers links pairs of neurons to the same assemblies, we then considered how much more assembly membership is determined by the identity of the specific patterns used. We hypothesized that direct innervation by fibers used in a pattern increases membership probability in assemblies associated with the same pattern. Specifically, we used as a structural feature the *pattern indegree* i.e., the total indegree of a neuron from VPM fibers used in each of the patterns.

Chapter 3. Cell assemblies and their underlying connectivity




Figure 3.4: Connectivity determines cell assembly membership. **A:** First: Effect of thalamic innervation (from VPM and POm nuclei) on participation in cell assemblies. Solid lines indicate the mean and the shaded areas indicate 95% confidence intervals. Second: probability of membership in an exemplary middle assembly against mean common POm indegrees with respect to all assemblies. Third and fourth: nI (normalized mutual information, see Methods) of mean common thalamic indegree and assembly membership. **B:** Probability of membership in exemplary early (first), middle (second), and late (third) assemblies against indegree with respect to all patterns. Fourth: nI of pattern indegree and assembly membership. **C:** Simplex counts within assemblies and random controls (same number of neurons with the same cell type distribution). (Caption continues on the next page.)

D: Probability of membership in exemplary early (first), middle (second) and late (third) assemblies against indegree with respect to all assemblies. Fourth: nI of indegree and assembly membership. **E:** First: Probability of membership in an exemplary early assembly against k -indegree with respect to the same assembly. Inset: k -indegree is given by the number of k -simplices within an assembly (orange nodes) completely innervating a given neuron (black). nI of k -indegree and assembly membership for $k = 1$ (second), $k = 2$ (third) and $k = 3$ (fourth). (White: nI not defined.)

As predicted from the previous results, probability of assembly membership grew rapidly with pattern indegree for early assemblies associated with the same pattern (Figure 3.4B, left). Every pattern had one early assembly strongly associated with it (except for pattern E, which only had middle and late assemblies Figure 3.3B2). On average, the nI of pattern indegree and assembly membership reached 0.26 (Figure 3.4B, right). A similar trend was observed for middle assemblies (mean: 0.215), while the late assembly was again an exception, with no value above 0.04, for reasons outlined above (Figure 3.4B third and fourth). In total, taking the stimulus patterns into account gives a 33% higher nI over the less specific mean common VPM innervation.

Recurrent connectivity explains late assemblies

Even with perfect knowledge of thalamic innervation and pattern identity the nI did not exceed values of 0.3, leading to the question: What other factors determine the rest? The most commonly accepted structural correlate of cell assemblies is the overexpression of recurrent connectivity motifs between participating neurons (Harris, 2005; Buzsáki, 2010; Song et al., 2005; Perin et al., 2011). One particular class of motifs that has been linked to neuronal function are directed simplices of dimension k (k -simplices Reimann et al., 2017b). A k -simplex is a motif on $k + 1$ neurons, which are all-to-all connected in a feed-forward fashion (Figure 3.4E left, inset), in particular 1-simplices are directed edges and 0-simplices are single cells. Indeed, we found a strong overexpression of directed simplices in the connectivity submatrices of cells within an assembly. In particular, the maximal simplex dimension found in assembly subgraphs is at least one higher than in the corresponding controls. Moreover, the peak of simplex counts in assembly graphs is in general one order of magnitude above the controls (Figure 3.4C).

Based on this we define the k -indegree with respect to an assembly of a neuron i , as the number of k -simplices in the assembly such that all the cells in the simplex innervate i (see Figure 3.4E left, inset). For the case $k = 0$, i.e., the number of cells in the assembly innervating i , we found that it is a good predictor of membership in the same assembly, with an average nI value of 0.143 (Figure 3.4D). This time, late assembly membership could be predicted the best, with an nI of 0.224, compared to an average of 0.118 and 0.16 for early and middle assemblies, respectively (Figure 3.4D). Additionally, probability of late assembly membership also increased with 1-indegree with respect to all other assemblies. Conversely, 0-indegree with respect to the late assembly decreases the membership probability for early, and most

Chapter 3. Cell assemblies and their underlying connectivity

of the middle assemblies (blue curves with negative slope on the first and second panels of Figure 3.4D). This reflects the temporal order of their activation and the fact that neurons in the deeper layers, that dominate the late assembly, mostly project outside the cortex and not back to superficial layers (Feldmeyer, 2012; Harris and Shepherd, 2015).

The nI values across the diagonals for $k = 1, 2$ are larger than for $k = 0$, for all assemblies except the late one. More precisely, we found nI averages of 0.152 and 0.161 for early, and 0.18 and 0.188 for middle assemblies for $k = 1$ and $k = 2$ respectively (Figure 3.4E second and third). On the other hand, off diagonal nI values drop for increasing k . This shows that not only the size of the presynaptic population has an effect on the activity of a neuron, but also the connectivity patterns between them. However, the effect of these non-local interactions are stronger within an assembly than across. For $k = 3$ the nI values drops, which can be explained by the narrower range of values the 3-indegree takes (Figure 3.4E first and fourth). As 0-indegree is the same as the general notion of indegree (the number of neurons in the afferent population) we will drop $k = 0$ and simply call it *indegree* in the following sections.

Synaptic clustering explains late assemblies

So far, we have only considered features which can be extracted from the connectivity matrix of the system. However, our model also offers subcellular resolution, specifically, the dendritic locations of all synapses (Markram et al., 2015; Kanari et al., 2019), which we have demonstrated to be crucial for recreating accurate post-synaptic potentials (Ramaswamy et al., 2012; Ecker et al., 2020), and long-term-plasticity (Chindemi et al., 2022). This allowed us to explore the impact of co-firing neurons potentially sending synapses to the same dendritic branch i.e., forming synapse clusters (Kastellakis and Poirazi, 2019; Wilson et al., 2016; Iacaruso et al., 2017; Ujfalussy and Makara, 2020). We hypothesized that innervation from an assembly is more effective at facilitating membership in an assembly if it targets nearby dendritic locations. We therefore defined the synaptic clustering coefficient (SCC) with respect to an assembly A_n , based on the path distances between synapses from A_n on a given neuron (see Methods and Supplementary Figure 3.12). The SCC is a parameter-free feature, centered at zero. It is positive for intersynaptic distances that are lower than expected (indicating clustering) and negative otherwise (indicating avoidance).

Overall, we found similar trends as for the recurrent connectivity, although with a lower impact on assembly membership. Late assembly membership was explained the best by the SCC with 0.114 nI , while early and middle assemblies had an average nI of only 0.048 and 0.061 (Figure 3.5A). Although SCC was not that powerful by itself, we found that a given value of indegree led to a higher assembly membership probability if the innervation was significantly clustered (Figure 3.5B, first). This lead us to explore the correlation between indegree and the SCC (Figure 3.5B second), finding a weak but significant correlation between these measures. Note that the SCC controls for the decrease in distance between synapses expected from a higher indegree (see Methods), therefore we conclude that this is non-trivial feature of the model. However, this means that the effects of indegree and SCC on assembly membership

are partially redundant. To dissociate these effects, we compute the nI between SCC and assembly membership conditioned by indegree (see Methods), which is on average 0.025, reaching values up to 0.045 (Figure 3.5B third and fourth). This shows that indegree and SCC affect assembly membership both independently but more so in conjunction.




Figure 3.5: Synapse clustering coefficient determines cell assembly membership. **A:** Probability of membership in exemplary early (first), middle (second), and late (third) assemblies against synapse clustering coefficient (SCC , see Methods) with respect to all assemblies. Fourth: nI (normalized mutual information, see Methods) of SCC and assembly membership. **B:** Combined effect of SCC and indegree (as in Figure 3.4D). First: Probability of membership in an exemplary early assembly, against indegree with respect to the same assembly, grouped by SCC significance (see Methods). Second: Joint distribution of SCC and indegree. Third: nI of SCC and assembly membership conditioned by indegree. (White: nI not defined.) Fourth: Relation of nI and conditional nI grouped by postsynaptic early, middle, and late assemblies (i.e., rows of the matrices in A fourth, and B third). **C:** Simulation results for 10 selected neurons per assembly (with the highest indegree and significant clustering; red arrow in B left) with modified physiological conditions. Left: correlation of spike times with the assembly. Right: Single cell firing rates.

The observed effect of SCC can only be explained by nonlinear dendritic integration of synaptic inputs (Stuart and Spruston, 2015; Kastellakis and Poirazi, 2019; Goetz et al., 2021) Specifically, our model has two sources of nonlinearity: First, active ion channels on the dendrites (Stuart and Spruston, 2015; Ujfalussy and Makara, 2020; Goetz et al., 2021), causing Na^+ and Ca^{2+} spikes. Second, NMDA conductances, which open in a voltage dependent manner,

Chapter 3. Cell assemblies and their underlying connectivity

leading to NMDA-spikes (Stuart and Spruston, 2015; Goetz et al., 2021). To show that these are the mechanisms by which the *SCC* acts, we studied how the removal of these non-linearities affects assembly membership of selected neurons. From each assembly we chose ten neurons whose probability of membership was most affected by their *SCC* i.e., those with highest indegree combined with significant clustering (red arrow in Figure 3.5B first). We subjected these neurons to the same input patterns they received in the simulations (see Methods), but with passive dendrites (Figure 3.5C green) or blocked NMDA receptors (Figure 3.5C orange). These modifications altered the neurons spiking activity, causing a non-negligible portion of them to drop out of their assembly, as their activity was no longer significantly correlated with it (Figure 3.5C first, see Methods). The manipulations resulted in a 45% drop in assembly membership for passive dendrites and 36% for blocking NMDA channels. Although, the manipulation resulted in an overall reduction of firing rate (Figure 3.5C second), this did not explain the drop in assembly membership (Figure 3.5C second: red dots with higher rate than black ones). We conclude that these nonlinearities contribute to the synchronization of activity within assemblies, underlying the observed effect of the *SCC*.

3.2.3 Assemblies are robust across simulation instances

The results of our simulations are stochastic (Nolte et al., 2019), leading to different outcomes for repetitions of the same experiment, as in biology. To assess the robustness of our results, we repeated our *in silico* experiment 10 times with the same thalamic inputs but different random seeds. Changing the seed mostly affects the stochastic release of synaptic vesicles (Markram et al., 2015; Nolte et al., 2019) especially at the low, *in vivo*-like extracellular Ca^{2+} concentration used. The assemblies detected in the repetitions were similar to the ones described so far, in term of member neurons (Figure 3.6A), temporal structure (early and middle groups and a non-specific late one, not shown), and their determination by connectivity features (not shown).

We hypothesize that cell assemblies in cortical circuits are inherently stochastic objects, partially determined by the structural connectivity, input stimuli, and neuronal composition. Thus, each repetition yields a different (but overlapping) set of assemblies and neurons contained in them. In order to get an approximation of these stochastic objects we pooled the assemblies detected in all repetitions and determined which best corresponded to each other by clustering them based on the Jaccard distance of their constituent neurons (see Methods, Figure 3.6A). According to this distance, nearby assemblies have a large intersection relative to their size. We called the resulting clusters *consensus assemblies*, and the assemblies contained in each its *instances*. We assigned to neurons different degrees of membership in a consensus assembly, based on the fraction of instances they were part of, normalized by a random control, and called it its *coreness* (see Methods). We found, that as coreness increased so did the neurons' spike time reliability (defined across the repetitions of the experiment (Nolte et al., 2019; Schreiber et al., 2003; see Methods), especially for the thalamus driven early assemblies (Figure 3.6B). We call the *union consensus assembly* the union of all its instances;

and the *core consensus assembly* the set of cells whose coreness is significantly higher than expected (see Methods and union vs. core distinction of Figure 3.6C).

When repeating our structural analysis on the core consensus assemblies, we found higher values of nI than before in all cases, except for the SCC (Figure 3.6D and Figure 3.7). This indicates that assembly membership in this model is not a binary property but exists on a spectrum from a highly reliable and structurally determined core to a more loosely associated and less connected periphery. The notion of consensus assemblies is a way of accessing this spectrum.




Figure 3.6: Consensus assemblies. **A:** Jaccard similarity based hierarchical clustering of assemblies from different simulation instances. **B:** Spike time reliability (see Methods) of neurons belonging to given fraction of assembly instances. **C:** Locations of neurons in the union (all instances) and core (at least 9/10 instances, see Methods) of exemplary (pattern A responsive) early consensus assembly. **D:** nI (normalized mutual information, see Methods) of connectivity features and consensus assemblies membership.

Another way to take all ten repetitions into account would be to average the time-binned spike trains of the simulated neurons. Thus, we averaged the input instead of the output of the assembly detection pipeline and we call these the *average assemblies*. We first compared these to the assemblies obtained in a single repetition. The similarity of significant time bins was higher for average assemblies (Supplementary Figure 3.13B1), and their sizes were larger (Supplementary Figure 3.13D second panel), with neurons belonging to up to 10 assemblies out of the 13 detected (Supplementary Figure 3.13D third panel). On the other hand, the nI of the structural features and membership remained the same as for a single repetition (compare matrices in Figures 3.4 and 3.5 to Supplementary Figure 3.13C).

Chapter 3. Cell assemblies and their underlying connectivity

When contrasting the average and consensus assemblies, we found pairs with high Jaccard similarity (Supplementary Figure 3.13D and E1). A detailed comparison showed that all neurons that are part of at least 6 out of 10 assembly instances were all contained in their matching average assembly. Further lowering the cutoff began admitting neurons that were not part of the corresponding average assembly (Supplementary Figure 3.13E2). On the other hand, there are barely any cells in the average that are not contained in the union consensus assembly. This demonstrates how assembly membership becomes less determined towards the periphery, mirroring the reduced nI with the structural connectivity features. Furthermore, a neuron in a consensus assembly will most likely belong to all instances (Supplementary Figure 3.13E2), unlike for the binomial distribution expected by chance.

In summary, while average assemblies give similar results to the union consensus assemblies, the coreness values used in the consensus assembly framework assign different degrees of membership to its neurons that can be taken into account in downstream analyses, e.g., by considering only the functionally reliable core. Furthermore, for this core the determination by most structural features, measured by nI , is stronger (Figure 3.7).




Figure 3.7: Summary of nI s of all connectivity features and assembly membership. Assemblies from single simulation on top and consensus assemblies from 10 simulations on the bottom. Only within-assembly interactions (diagonals of nI matrices) shown, except for the patterns, where for each column (postsynaptic assembly) we used the maximum value (strongest innervating pattern). Not only the colors, but the radii of the pies code for features: Red (VPM) longer one depicts common-innervation (as yellow for P0m) while the shorter one direct innervation from patterns. Green: Simplex dimension increases as length decreases (longest: indegree, shortest: 3-indegree). Blue: longer one is the SCC conditioned on indegree, while the shorter one is the same, but unconditioned (and that is why it always overlaps with indegree).

3.3 Discussion

Using a detailed, large-scale cortical network model we examined the link between cortical structure and function. Our principal findings from analyzing the connectivity of functional cell assemblies are as follows: (1) Different afferents dominated determination of membership in assemblies linked to different time windows: VPM innervation affected membership in early assemblies, POM innervation in middle assemblies, while recurrent innervation in the late assembly (Figure 3.7 red, yellow, green). (2) Recurrent innervation more strongly facilitated assembly membership the more the innervating neurons were wired among themselves, adding a non-local component to the structure-function relation. (3) Similarly, the innervation of a neuron by an assembly was significantly more powerful in facilitating membership when its synapses were clustered on the dendrites (Figure 3.7 blue). (4) In conjunction with the structure of cortical connectivity, features of subcellular physiology such as active dendritic channels and NMDA receptors, also influence assembly membership. (5) Interactions between assemblies emerged, where innervation by one assembly explained membership in another. Positive interactions (with increased membership probability) were found when the direction of innervation reflected the temporal order of assembly activation; otherwise, weaker or even negative interactions were found.

Point (1) above confirms our previous findings that, while the presence of an external stimulus makes the circuit much more reliable, this effect is not merely driven by direct innervation, but also requires recurrent connectivity (Nolte et al., 2019). Point (2) predicts a functional consequence for non-random features of neuronal connectivity, such as nodes with high centrality values and the presence of a rich club, that have been characterized in many species and regions, at various levels of resolution (Bassett and Bullmore, 2017). Points (3) and (4) link the theory of neuronal assemblies to the literature on active dendritic computation (Poirazi et al., 2003; Stuart and Spruston, 2015; Kastellakis and Poirazi, 2019). Point (5) strengthens the case that the topology of connectivity is best studied in a directed way, since undirected networks (although more amenable to network science methods) miss an essential part of the picture. In particular, the relationship between the structural directionality of the connection and the temporal direction of the flow of activity (Reimann et al., 2017b).

Our analysis supports the idea that neuronal activity revolves around activation of assemblies. We have shown that assembly membership is determined by certain structural prerequisites, mostly amounting to increased membership probability when more afferent synapses from various sources are formed on a neuron. These may be costly to fulfill, both in terms of energy (Harris et al., 2012) and space taken up by wiring. Chklovskii et al. (2002) considered efficient layouts of wiring, i.e., axons and dendrites, given a certain connectivity matrix. They found that the layout is tightly constrained by the available space and close to a theoretical optimum. Here, we expanded on that idea, demonstrating that on top of it, the structure of the connectivity matrix is efficient. Indeed, on the sub-cellular level, we have shown, that synaptic clustering on dendrites increases efficiency, in that the same probability of assembly membership can be obtained with about 50% of the indegree when the synapses are highly clustered (Figure

Chapter 3. Cell assemblies and their underlying connectivity

3.5B, Supplementary Figure 3.14A). Furthermore, on the connectivity level, another potential mechanism lies in the non-local interaction of the presynaptic population measured by the k -indegree. We have shown that for higher values of k , a given membership probability is attained for lower k -indegree values, suggesting higher efficiency (Figure 3.4E, left). However, as k grows, k -indegree counts the innervation by larger and larger motifs, potentially requiring more incoming synaptic connections (Supplementary Figure 3.14B2). Here, we control for this by repeating our analysis not with respect to the number of motifs but the size of the presynaptic subpopulation forming them. When going up from $k = 0$ to $k = 1$, $k = 2$, and $k = 3$ the same membership probability can be obtained for a given neuron with about 80%, 34% and 6% of the incoming connections (Supplementary Figure 3.14B1), confirming that non-local interactions in the presynaptic population make afferent innervation more efficient. This can be explained by our earlier finding that the simplicial motifs we considered increase the correlations and reliability of the spiking activity of participating neurons (Reimann et al., 2017b).

It is possible that *wiring efficiency* can be optimized further, based on experience, through structural and functional plasticity. Recent modelling studies investigated how starting from a random initial connectivity, plasticity rules and network activity lead to assembly formation, maintenance and *competition for member neurons* (Fauth and Van Rossum, 2019; Kossio et al., 2021; Gastaldi et al., 2021). Conversely, our network model has strongly non-random connectivity, constrained by neuronal morphology (Reimann et al., 2015, 2017a), and can thus be viewed as a circuit in a non-random plastic state, but unshaped by experience. The presence of assemblies in such a naive circuit is in line with the belief that the brain is not a tabula rasa (Buzsáki, 2019). It is also in line with the recent *in vivo* experiments of Bathellier et al. (2012) and Trägenap et al. (2022), who found endogenous cell assemblies in mouse auditory and ferret visual cortices. Furthermore, Trägenap et al. (2022) also found that these endogenous assemblies *solidify* and become more reliable after eye opening. These points lead to the question: How does long-term plasticity affect the three aspects discussed above, i.e., assembly wiring efficiency, competition for member neurons and assembly solidification? To address this question, we are integrating the functional plasticity model of Chindemi et al. (2022) into our network simulations. We can then analyze these with the methods we have introduced here, which provide new, quantitative ways to characterize assemblies, their temporal evolution and the connectivity underlying them.

3.4 Methods

3.4.1 Network simulations

The most recent version of the detailed, large-scale cortical microcircuit model of Markram et al. (2015) was used for the *in silico* experiments in this study. Updates on its anatomy, e.g., atlas-based cell densities are described in Reimann et al. (2022a), while updates on its physiology e.g., improved single cell models and missing input compensation in Isbister et al.

(2023). The 2.4 mm^3 subvolume of the juvenile rat somatosensory cortex, containing 211,712 neurons is freely available at: <https://zenodo.org/record/7930275>.

Although large-scale with bio-realistic counts of synapses originating from the local neurons, the neurons in the circuit still lacked most of their synapses (originating from other, non-modeled regions; Markram et al., 2015). In order to compensate for this missing input, layer and cell-type specific somatic conductances following an Ornstein-Uhlenbeck process were injected to the cell bodies of all neurons (Destexhe et al., 2001). The algorithm used to determine the mean and variance of the conductance needed to put the cells into an *in vivo*-like high-conductance state, and the network as whole into an *in vivo*-like asynchronous firing regime with low rates and realistic responses to short whisker stimuli is described in Isbister et al. (2023). The *in vivo*-like state used in this article is the same as $[Ca^{2+}]_o = 1.05 \text{ mM}$, percentage of reference firing rates = 50%, CV of the noise process = 0.4 from Isbister et al. (2023).

Simulations of selected cells with modified physiological conditions (active dendritic channels blocked or NMDA conductance blocked) used the activity replay paradigm of Nolte et al. (2019). In short, for each of these cells, spike times of its presynaptic population were recorded in the original network simulation. Then the selected cells were simulated in isolation by activating their afferent synapses according to the recorded spike times in the network simulation. Thus, the modified activity of the isolated cell did not affect the rest of the network.

Simulations were run using the NEURON simulator as a core engine with the Blue Brain Project's collection of hoc and NMODL templates for parallel execution on supercomputers (Hines and Carnevale, 1997; Kumbhar et al., 2019; Awile et al., 2022). Simulating 2 minutes of biological time took 100,000 core hours, on our HPE based supercomputer, installed at CSCS, Lugano.

3.4.2 Distance metrics

This section gives a brief overview and justification of the various distance metrics used below.

Population activity in time bins were compared using their cosine similarity Carrillo-Reid et al. (2015). Two time bins with high cosine similarity have similar sets of firing neurons, thus detecting co-firing. Note that there is an increasing relationship between firing rate and cosine similarity (Cutts and Eglen, 2014; Supplementary Figure 3.8A2).

Input patterns were defined using the Hamming distance between the sets of VPM fiber bundles involved to have specific values and thus specific sizes of intersections (Figure 3.2B). Input patterns were compared using Earth mover distances between the flat map locations of their contained fibers (Figure 3.3E).

Assemblies of neurons were compared using their Jaccard distances. Like Hamming, this also compares sizes of intersections but it is normalized with respect to the sizes of the sets involved.

Chapter 3. Cell assemblies and their underlying connectivity

Assembly sequences i.e., vectors of size the number of assemblies, with each entry counting the number of time bins the corresponding assembly was active in response to an input pattern, were compared using their normalized Euclidean distances. This is the Euclidean distance between vectors of normalized length.

Finally, afferent synapses were compared using their path distances. Specifically, the dendritic tree was represented as a graph with nodes being its branching points and edges between them weighted according to the length of sections connecting them. Distances between synapses was computed as the path distance in this graph.

3.4.3 Thalamic input stimuli

The VPM input spike trains were similar to the ones used in Reimann et al. (2022b). In detail, the 5388 VPM fibers innervating the simulated volume were first restricted to be $\leq 500 \mu\text{m}$ from the middle of the circuit in the horizontal plane to avoid boundary artefacts. To measure these distances we use a flat map, i.e., a two dimensional projection of the volume onto the horizontal plane, orthogonal to layer boundaries (Reimann et al., 2022a). Second, the flat map locations of the resulting 3017 fibers were clustered using k-means to form 100 bundles of fibers. The base patterns (A, B, C, and D) were formed by randomly selecting four non-overlapping groups of bundles, each containing 12% of them (corresponding to 366 fibers each). The remaining 6 patterns were derived from these base patterns with various degrees of overlap (see beginning of Results, Figure 3.2B). Third, the input stream was defined as a random presentation of these ten patterns with 500 ms inter-stimulus intervals, such that in every 30 second time intervals every pattern was presented exactly six times. Last, for each pattern presentation, unique spike times were generated for its corresponding fibers following an inhomogeneous adapting Markov process (Muller et al., 2007). When a pattern was presented, the rate of its fibers jumped to 30 Hz and decayed to 1 Hz over 100 ms. For the non-specific POm stimuli, a randomly selected 12% of the (unclustered) 3864 POm fibers were activated each time any pattern was presented (in every 500 ms). The spike trains were designed with the same temporal dynamics as described above for VPM, but with half the maximum rate (15 Hz). The implementation of spike time generation was based on `Elephant` (Denker et al., 2018).

3.4.4 Assembly detection

Our assembly detection pipeline was a mix of established techniques and consisted of five steps: binning of spike trains, selecting significant time bins, clustering of significant times bins via the cosine similarity of their activity, and determination of neurons corresponding to a time bin cluster and thus forming an assembly. Note, that time bins instead of neurons were clustered because this allows neurons to belong to several assemblies.

Spikes of excitatory cells were first binned using 20 ms time bins, based on Harris et al. (2003), which takes a postsynaptic reader neuron specific point of view. They suggest 10-30 ms as an ideal integration time window of the presynaptic (assembly) spikes (Harris et al., 2003; Buzsáki, 2010). Next, time bins with a significantly high level of activity were detected. The significance threshold was determined as the mean activity level plus the 95th percentile of the standard deviation of shuffled controls. The 100 random controls were rather strict, i.e., all spikes were shifted only by one time bin forward or backward, based on Sasaki et al. (2006); Carrillo-Reid et al. (2015). Next, a similarity matrix of significant time bins was built, based on the cosine similarity of activation vectors, i.e., vectors of spike counts of all neurons in the given time bins (Carrillo-Reid et al., 2015). The similarity matrix of significant time bins was then hierarchically clustered using Ward's linkage (Montijn et al., 2016; Pérez-Ortega et al., 2021). Potential number of clusters were scanned between five and twenty, and the one with the lowest Davis-Bouldin index was chosen, which maximizes the similarity within elements of the cluster while maximizing the between cluster similarity (Davies and Bouldin, 1979). These clusters corresponded to potential assemblies.

As the last step, neurons were associated to these clusters based on their spiking activity, and it was determined whether they formed a cell assembly or not. In detail, the correlations between the spike trains of all neurons and the activation sequences of all clusters were computed and the ones with significant correlation selected. Significance was determined based on exceeding the 95th percentile of correlations of shuffled controls (1000 controls with spikes of individual cells shifted by any amount Montijn et al., 2016; Herzog et al., 2021). In relation to Figure 3.5C left, it is important to note, that these correlation thresholds were specific to a pair of a neuron and an assembly. Finally, it is possible to have a group of neurons that is highly correlated with one part of the significant time bins in a cluster, and another that is highly correlated with the rest, while the two groups of neurons have uncorrelated activity. To filter out this scenario, it was required that the mean pairwise correlation of the spikes of the neurons with significant correlations was higher than the mean pairwise correlation of neurons in the whole dataset (Herzog et al., 2021). Clusters passing this test were considered to be functional assemblies and the neurons with significant correlations their constituent cells.

For a test of the methods on synthetic data and comparison with other ways of detecting cell assemblies please consult Herzog et al. (2021). Assembly detection was implemented in Python and is publicly available as `assemblyfire`.

3.4.5 Calculation of information theoretical measurements

To quantify the structural predictability of assemblies, *mutual information* of assembly membership (Y_n) and a structural feature of a neuron (X_m) was used, which is a measure of the mutual dependence between the two variables (Gray, 2011). More precisely, Y_n is a binary random variable that takes the value 1 if a neuron belongs to an assembly A_n , and 0 otherwise.

Chapter 3. Cell assemblies and their underlying connectivity

On the other hand, X_m is a random variable determined by a structural property of the neuron with respect to the assembly A_m , e.g., the number of afferent connections into the neuron from all neurons in A_m .

To assess the dependence between these two variables, first the dependence of the probability of a neuron belonging to assembly A_n given a specific value of the structural feature measured by X_m was studied. More precisely, the function $f_{n,m}$ whose domain is the values of X_m and is given by:

$$f_{n,m}(x) = P(Y_n = 1 | X_m = x) \quad (3.1)$$

was considered. If this function has an increasing or decreasing trend, then the random variables Y_n and X_m can not be independent.

Their dependence was quantified by means of their mutual information. The value of mutual information is always non-negative and it is zero when the random variables are independent. In order to restrict this value to $[0, 1]$ it was divided by the entropy of Y_n , which measures the level of inherent uncertainty of the possible outcome of the values of Y_n (Shannon, 1948). The calculation of mutual information is based on the probabilities of X_m and Y_n across all possible outcomes. If the number of possible values of X_m is large compared to the number of samples, there can be errors in determining these probabilities, possibly leading to inflated values of the mutual information. Therefore, the values of X_m were binned into 21 bins between the 1st and 99th percentile of all sampled values. The number of bins was determined such that the resulting value of mutual information in a shuffled control did not exceed 0.01. Shuffled controls (one per pair) were also used to threshold the mutual information values by considering only the pairs (n, m) whose mutual information was larger than the mean plus one standard deviation of all pairs in the shuffled controls. Finally, a negative sign was added to the significant mutual information value if the function $f_{n,m}$ was decreasing i.e., the probability of membership in A_n decreased as the values of X_m increased. This was assessed by the slope of a weighted (by the number of samples in each bins) linear fit of the function $f_{n,m}$. This normalized, thresholded and signed mutual information value was called *normalized mutual information* and denoted $nI(Y_n, X_m)$.

All the statements above can be made conditional with respect to a third random variable, yielding the *conditional normalized mutual information*, which was used when two structural features were inherently believed to be interacting as in the case of the SCC and indegree.

Calculations were done with the `pyitlib` package.

3.4.6 Synaptic clustering coefficient

To quantify the co-localization of synapses on the dendrites of a neuron i from its presynaptic population P_i with a single, parameter-free metric, synaptic clustering coefficient SCC was defined and calculated for all excitatory neurons in the circuit with respect to all assemblies. Based on these locations, D_i , the matrix of all pairwise path distances between synapses on i

from P_i were calculated. Let $D_{i,p}$ be the submatrix for pairs of synapses originating from a subpopulation $p \subseteq P_i$. Then the nearest neighbour distance for p can be written as:

$$nnd(i, p) = \text{mean} \left(\min_{\text{rows}} (D_{i,p}) \right) \quad (3.2)$$

In particular, for the subpopulation of P_i of neurons in the assembly A_n , denoted by p_n , $nnd(i, A_n) = nnd(i, p_n)$ was defined, where $p_n = P_i \cap A_n$. This value was normalized, using the nnd values of 20 random presynaptic populations from P_i of the same size as p_n . In summary, the SCC was defined as the negative z-score of $nnd(i, A_n)$ with respect to the distribution of control $nnds$ (Supplementary Figure 3.12). Additionally, the significance of the clustering or avoidance of the synapse locations was determined with a two-tailed t-test of $nnd(i, A_n)$ against the 20 random samples with an alpha level of 0.05. SCC was implemented using NeuroM and ConnectomeUtilities.

3.4.7 Determination of consensus assemblies

Consensus assemblies were defined over multiple repetitions of the same input stream. These were groups of assemblies with similar sets of neurons. Additionally, all assemblies in a group were required to originate from a different repetition; noted as the *repetition separation criterion*. The Jaccard distance matrix between all pairs of assemblies from all repetitions were computed and modified by setting the distances between pairs of assemblies from the same repetition to twice the maximum of the whole matrix. The matrix was then hierarchically clustered using Ward's linkage, and the lowest number of clusters that satisfied the repetition separation criterion was chosen. The resulting clusters were the *consensus assemblies*, and the assemblies within them their *instances*.

The *union consensus assembly* was defined as the set of neurons given by the union of all instances. Its member neurons were assigned a membership degree based on number of instances they were part of in two ways. First, by simply using the fraction of the instances a neuron was part of. Second, in what was called the *coreness* value of a neuron, which is the number of instances a neuron is part of normalized by its expected value, given the number and sizes of its instances. This was calculated as a binomial distribution with n set to the number of instances and p to the mean size of the instances, divided by the size of the union consensus assembly. Based on this, the coreness of a neuron contained in r instances was defined as $-\log_{10}(1 - B_{n,p}(r))$, where B is the cumulative binomial distribution. Neurons with a coreness value exceeding 4 were considered to be part of the corresponding *core consensus assembly*.

3.4.8 Calculation of spike time reliability

Spike time reliability was defined as the mean of the cosine similarities of a given neuron's mean centered, smoothed spike times across all pairs of repetitions (Schreiber et al., 2003; Cutts and Eglén, 2014). To smooth the spike times, they were first binned to 1 ms time bins, and then convolved with a Gaussian kernel with a standard deviation of 10 ms.

Acknowledgements

The authors thank Nicolas Ninin for his involvement in the early stage of this project, Elvis Boci and Cyrille Favreau for their help with visualizations and Alberto Antonietti, Christoph Pokorny, Kathryn B. Hess, Ran Levi, and Henry Markram for discussions.

Funding

This study was supported by funding to the Blue Brain Project, a research center of the École polytechnique fédérale de Lausanne (EPFL), from the Swiss government's ETH Board of the Swiss Federal Institutes of Technology.

Author contributions

- Conceptualization: A.E., D.E.S., M.W.R.
- Data curation: A.E., J.B.I., S.B.-P.
- Formal analysis: A.E., D.E.S., M.W.R.
- Investigation: A.E.
- Methodology: A.E., D.E.S., M.W.R.
- Project administration: M.W.R.
- Software: A.E., D.E.S., M.W.R., S.B.-P., J.B.I.
- Supervision: M.W.R.
- Validation: A.E., D.E.S., M.W.R., J.B.I., S.B.-P.
- Visualization: A.E.
- Writing - original draft: A.E., D.E.S., M.W.R.
- Writing - review & editing: A.E., D.E.S., M.W.R., J.B.I., S.B.-P.


Declaration of interests

The authors declare no competing interests.

Data and code availability


The 2.4 mm^3 subvolume of the juvenile rat somatosensory cortex (containing 211,712 neurons and their connectivity) used for the *in silico* experiments in this study has been deposited at Zenodo in SONATA format (Dai et al., 2020) and is publicly available at the following DOI: [10.5281/zenodo.7930275](https://doi.org/10.5281/zenodo.7930275). The simulator front-end that loads the SONATA model, and instantiates the simulation to be run in CoreNEURON (Kumbhar et al., 2019) is also publicly available at GitHub or under the following DOI: [10.5281/zenodo.8075202](https://doi.org/10.5281/zenodo.8075202). Assemblies from 10 simulation repetitions (with different random seeds), their consensus, and underlying significant spike times, and the whole excitatory connectivity matrix have been deposited at Zenodo and is publicly available at the following DOI: [10.5281/zenodo.8052721](https://doi.org/10.5281/zenodo.8052721). Analysis code that created (and can easily open the dataset above) is publicly available at GitHub or under the following DOI: [10.5281/zenodo.8112725](https://doi.org/10.5281/zenodo.8112725).

3.5 Supplementary Figures




Supplementary Figure 3.8: Spikes of (significant) time bins. **A1** Cosine similarity matrix (same as in Figure 3.2B1, but unsorted). **A2**: Joint distribution of pair-wise mean firing rate (of time bins) and cosine similarity. **B1**: 2D linear projection of mean centered and normalized spike matrix. **B2**: 2D nonlinear projection of spike matrix (using cosine distance).

Chapter 3. Cell assemblies and their underlying connectivity




Supplementary Figure 3.9: Different number of clusters (of significant time bins) tested. A1, B and C: as in Figure 3.3B, E. **A2:** Davis-Bouldin index (Davies and Bouldin, 1979) for the different number of clusters.

3.5 Supplementary Figures

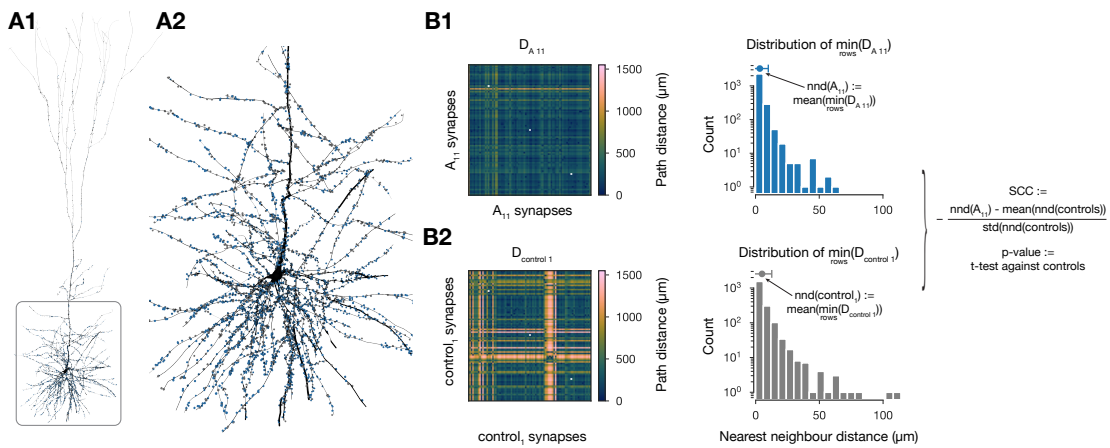


Supplementary Figure 3.10: Cell assembly detection in L2/3. A-C: as in Figure 3.3B, D, E. **B:** Jaccard similarity of assemblies (detected across layers, but restricted to L2/3 here) and the ones detected in L2/3 on the left, and number and location: flatmap view on top, depth-profile below of exemplary pairs (pattern A and B responsive ones) of assemblies with high similarity to its right.




Supplementary Figure 3.11: Anatomy of thalamocortical synapses. **A:** Density profile of VPM and POm synapses, digitized from Meyer et al. (2010). **B:** Synapse-to-soma path distances of different thalamocortical synapses on L5 neurons in an exemplary middle assembly (A_8). Ratios of box widths' represent the ratios of the number of synapses' from the inputs (e.g. most synapses on A_8 L5 pyramidal cells are coming from pattern G, in line with the indegree based A_8 membership probability on Figure 3.4B second).

Chapter 3. Cell assemblies and their underlying connectivity



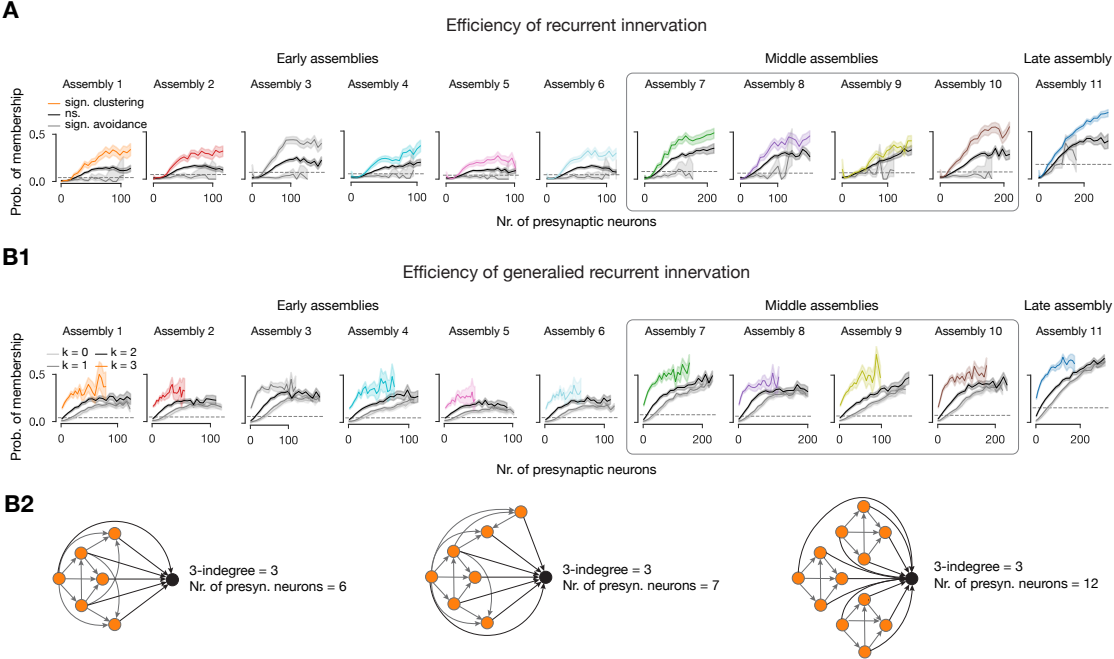
Supplementary Figure 3.12: Synaptic clustering coefficient. **A1:** Exemplary L5 pyramidal cell and all its afferent synapses from A_{11} (in blue) and from a control group (one out of the twenty) with the same number of presynaptic neurons (in gray). **A2:** Zoom in on A1. Soma, basal dendrites, and proximal apical dendrites are visible. Axon is not shown. The rendering was done with the BioExplorer package. **B1:** Distance matrix between all pairs of A_{11} synapses and distribution of nearest neighbour distances (minimum over the rows of the matrix) on its right (see Methods). **B2:** same as B1 but for the control group. The equations on the righthmost part of the figure are motivated and explained in the Methods.

3.5 Supplementary Figures



Supplementary Figure 3.13: Assemblies detected from averaged spike matrices. **A, B:** as in Figure 3.3A, B. **C:** as in Figure 3.6C. **D:** Left: Jaccard similarity of consensus assemblies and average assemblies. Middle: Number of neurons in consensus assemblies' union and core and average assemblies. Right: As middle, but number of neurons participating in given number of assemblies. **E1:** Jaccard similarity of consensus assemblies at different fraction thresholds (consensus assembly size grows to the right) and average assemblies. **E2:** Detailed comparison of the pair with the highest similarity in E1 at given number of assembly instances contained. (10 means that the consensus assembly neuron is part of 10/10 assembly instances, thus consensus assembly sizes grows to the right again.) Average consensus is negligible (419 neurons) and is not shown.

Chapter 3. Cell assemblies and their underlying connectivity



Supplementary Figure 3.14: Efficiency of innervation. Probabilities of within assembly memberships for all assemblies. **A:** As in Figure 3.5B left. **B1:** Similar to Figure 3.4E left, but with different x-axis (presynaptic population size instead of simplex counts, see B2). **B2:** Illustration of the difference between k -indegree and presynaptic population size.

4 Long-term plasticity induced sparse and specific synaptic changes

This chapter is an improved version (based on the thesis examiner's comments) of our recent preprint: **A. Ecker***, D. Egas Standander*, M. Abdellah, J. Blanco Alonso, S. Bolaños-Puchet, G. Chindemi, J.B. Isbister, J.G. King, P. Kumbhar, I. Magkanaris, E.B. Muller, M.W. Reimann (2023) *Long-term plasticity induces sparse and specific synaptic changes in a biophysically detailed cortical model*. bioRxiv; doi: [10.1101/2023.08.07.552264](https://doi.org/10.1101/2023.08.07.552264)

Contribution: I have parametrized the plastic synapses with the help of G. Chindemi, and helped J. Blanco Alonso, P. Kumbhar, J.G. King, and I. Magkanaris to make the plasticity model compatible with our simulator. I set up and ran all of the simulations, scripted most of the analysis based on input from D. Egas Standander, S. Bolaños-Puchet and M.W. Reimann, created all the figures and the open source data set, wrote the first version of the manuscript and participated in its rewriting. (A more detailed author contribution can be found at the end of the chapter.)

Abstract

Synaptic plasticity underlies the brain's ability to learn and adapt. This process is often studied in small groups of neurons *in vitro* or indirectly through its effects on behavior *in vivo*. Due to the limitations of available experimental techniques, investigating synaptic plasticity at the microcircuit level relies on simulation-based approaches. Although modeling studies provide valuable insights, they are usually limited in scale and generality. To overcome these limitations, we extended a previously published and validated large-scale cortical network model with a recently developed calcium-based model of functional plasticity between excitatory cells. We calibrated the network to mimic an *in vivo* state characterized by low synaptic release probability and low-rate asynchronous firing, and exposed it to 10 different stimuli. We found that synaptic plasticity sparsely and specifically strengthened synapses forming spatial clusters on postsynaptic dendrites and those between populations of co-firing neurons, also known as cell assemblies: among 312 million synapses, only 5% experienced noticeable

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes

plasticity and cross-assembly synapses underwent three times more changes than average. Furthermore, as occasional large-amplitude potentiation was counteracted by more frequent synaptic depression, the network remained stable without explicitly modeling homeostatic plasticity. When comparing the network's responses to the different stimuli before and after plasticity, we found that it became more stimulus-specific after plasticity, manifesting in prolonged activity after selected stimuli and more unique groups of neurons responding exclusively to a single pattern. Taken together, we present a plasticity rule that leads to sparse change and analyze the rules governing those changes.

4.1 Introduction

Learning and memory are orchestrated by synaptic plasticity, the ability of synapses to change their *efficacy* in an activity-dependent manner. Donald O. Hebb's postulate about how synaptic plasticity might manifest was paraphrased to the well known mantra: "*cells that fire together, wire together*" (Hebb, 1949; Shatz, 1992). The first proof of coincident pre- and postsynaptic population activity leading to *potentiation* (an increase in efficacy) came from pathway stimulation in hippocampal slices (Bliss and Lømo, 1973). It was later confirmed at the neuron pair level (Markram et al., 1997b; Bi and Poo, 1998), and spike-time dependent plasticity (STDP) became a quintessential protocol to study Hebbian plasticity *in vitro*. In the early 2000's a plethora of cortical pathways were studied and plasticity proved to be synapse location- and therefore pathway-dependent (Sjöström and Häusser, 2006; Letzkus et al., 2006; Froemke et al., 2010). The molecular substrate of Hebbian coincidence detection is the N-methyl-D-aspartate (NMDA) receptor, which upon removal of the Mg^{2+} block by depolarization, conducts Ca^{2+} as well (Mayer et al., 1984). The calcium-control hypothesis, put forward by Lisman (1989) postulates that prolonged, moderate amounts of Ca^{2+} lead to depression (a decrease in efficacy) while large transients of Ca^{2+} lead to potentiation. By putting these together, it became evident that it is not necessarily the timing of the postsynaptic spike, but the depolarization of the postsynaptic dendrite is important to evoke changes in synaptic efficacy (Goldberg et al., 2002; Lisman and Spruston, 2005).

In parallel with slice electrophysiology, Hebbian plasticity was also studied through its effect on behaviour via fear conditioning experiments (McKernan and Shinnick-Gallagher, 1997) and this line of research lead to a plethora of new techniques for tagging and re-activating cells that participate in newly formed memories (Tonegawa et al., 2015). While these studies highlighted the need to study plasticity at the network level, most changes are expected to happen at the synapse level. Therefore, high-throughput methods tracking synaptic proteins like PSD95 (Ray et al., 2023) and α -amino-3-hydroxy-5-methyl-4-isoxazolepropionate (AMPA) subunit GluA1 (Graves et al., 2021; Kim et al., 2023) are currently being developed. While readily applicable to monitor synaptic efficacy *in vivo*, currently, these techniques cannot be supplemented with recordings of neural activity thus the reason for the changes in efficacy can only be speculated.

The bridge between *in vitro* pairs of neurons and *in vivo* behavior is often provided by complementary simulation based-approaches. Early theoretical work explored the potential link between memories and cells that fire and therefore wire together, concentrating on the storage and retrieval of memories in strongly recurrent networks (Hopfield, 1982), which remained an active topic of research (Fusi and Abbott, 2007; Krotov and Hopfield, 2016; Widrich et al., 2020). In parallel with the STDP experiments, modelers developed plenty of *learning rules* that could explain the most recent pathway-specific findings (Gerstner et al., 1996; Kempter et al., 1999; Song et al., 2000; Pfister and Gerstner, 2006; Clopath et al., 2010). Of particular interest is the calcium-based model of Graupner and Brunel (2012), which models the evolution of intracellular calcium concentration ($[Ca^{2+}]_i$) given the pre- and postsynaptic spike trains and updates the efficacy of the synapse, upon $[Ca^{2+}]_i$ crossing thresholds for depression and potentiation. Linking memory storage, recall and bioplausible learning rules together, combinations of diverse sets of plasticity rules have been used to model the formation and maintenance of Hebbian *cell assemblies*, i.e., groups of neurons that fire together (Litwin-Kumar and Doiron, 2014; Zenke et al., 2015; Fauth and Van Rossum, 2019; Kossio et al., 2021). A common theme in these models is the necessity of fast homeostatic plasticity, that keeps the networks stable (Zenke et al., 2017a), however experimental evidence for those mechanisms is lacking (Turrigiano and Nelson, 2004). While these studies provided mechanistic explanation of learning and memory, they used point-neuron models, therefore neglecting the structural and functional importance of dendrites and other subcellular components (but see Bono et al., 2017; Kastellakis and Poirazi, 2019). The compartmentalized nature of dendritic trees gives rise to spatial clustering of synapses (Farinella et al., 2014; Iacaruso et al., 2017; Kastellakis and Poirazi, 2019) and local, non-linear voltage events (Poirazi et al., 2003; Stuart and Spruston, 2015) both of which are thought to contribute to removing the Mg^{2+} block from NMDA receptors and therefore gating plasticity.

To go beyond networks of point neurons stabilized with homeostatic plasticity, we equipped the biophysically detailed, large-scale cortical network model of Markram et al. (2015) with our recently developed, calcium-based model of functional plasticity (Chindemi et al., 2022) between excitatory cells (Figure 4.1). This way, we had access to more realistic pre- and postsynaptic activity and efficacy of millions of synapses and could characterize the rules governing plasticity at the microcircuit level. To make our predictions more relevant, we calibrated the circuit's activity to mimic an *in vivo* state, characterized by low synaptic release probability and low firing rates (Isbister et al., 2023). Thanks to the biophysical detail of the model, we could also take the effect of low extracellular calcium concentration ($[Ca^{2+}]_o$) into account (Chindemi et al., 2022), which was experimentally shown to reduce plasticity (Inglebert et al., 2020; Figure 4.1E). As we followed a bottom-up framework and did not model any specific task, we will refer to the effects of plasticity as *changes* in synaptic efficacy instead of learning. Changes in synaptic efficacy were sparse, affecting 5% of all synapses in 10 minutes of biological time. On the other hand, this was still enough to reorganize the network's dynamics, manifesting in more pattern-specificity after plasticity than before. We found an increased likelihood of changes within and across cell assemblies and in synapses

forming spatial clusters on postsynaptic dendrites. Among 312 million synapses, potentiation dominated in amplitude and depression counteracted it in frequency, which lead to stable firing rates without explicitly introducing any homeostatic terms (Turrigiano and Nelson, 2004; Zenke et al., 2017a). To support future, potentially more task-related studies of learning in the cortex, we made the model and the simulator available to the community.

4.2 Results

To achieve a continuous readout of plastic changes in synaptic efficacy of millions of excitatory synapses, we used a biophysically detailed, large-scale cortical model of the rat non-barrel somatosensory cortex (nbS1). The model improves on Markram et al. (2015) in terms of both anatomical, e.g., atlas based cell composition and placement (described in Reimann et al., 2022a), and physiological properties, e.g., improved single cell models, multi-vesicular synaptic release, and layer-wise compensation for missing synapses (described in Isbister et al., 2023). For this study, we used a seven column subvolume comprising 211,712 neurons in 2.4 mm^3 of tissue (Figure 4.1A) to keep the complexity of simulation and analysis manageable. In line with the biological variability, excitatory cells are modeled as a diverse set of morphologies (Kanari et al., 2019; Reimann et al., 2022a; Figure 4.1B) equipped with conductances distributed across all compartments (Reva et al., 2022; Supplementary Figure 4.7A). The connectivity and synaptic physiology of these cells were extensively validated (Reimann et al., 2022a; Isbister et al., 2023; Figure 4.1C; Supplementary Figure 4.7C). The model is also equipped with fibers from the ventral posteromedial nucleus of the thalamus (VPM) and the high-order posteromedial nucleus of the thalamus (POm; Figure 4.1D; Meyer et al., 2010). We use these fibers to deliver inputs with spatio-temporal precision.

4.2.1 Calcium-based, biophysically detailed model of long-term plasticity

In previous versions of the circuit model, synapses were only equipped with short-term plasticity (STP; Figure 4.1C). In the remainder of the manuscript we will call this the *non-plastic* version, as our scope here is long-term plasticity. To model long-term plasticity we integrated our recently published calcium-based plasticity model that was used to describe functional long-term potentiation and depression between pairs of pyramidal cells (PCs; Chindemi et al., 2022). In short, the model follows the formalism of Graupner and Brunel (2012), where pre- and postsynaptic spikes lead to changes in synaptic $[Ca^{2+}]_i$ (Figure 4.1E). Calcium entering through NMDA receptors and voltage-dependent calcium channels (VDCCs) contributes to $[Ca^{2+}]_i$ (equation (4.2) in Methods). When the integrated calcium trace of a synapse crosses the threshold for depression (θ_d) or the higher one for potentiation (θ_p), *synaptic efficacy* (ρ , exhibiting a bistable dynamics Lisman, 1985) is updated (Figure 4.1E left; equation (4.1) in Methods).




Figure 4.1: Overview of the network model. **A:** Visualisation of the seven column subvolume of rat nbS1. Rendering of 10% of the cells was done with Brayns. **B:** Representative morphologies for the 18 excitatory m-types and their typical firing pattern (e-type, top left). **C:** Exemplary connections to L5 TTPCs (top) and their STP profiles (bottom). Thin gray represent the 20 individual repetitions, while the thicker black ones their means. Renderings of morphologies (on B as well) were done with NeuroMorphoVis (Abdellah et al., 2018). Neurite diameters are scaled (x3) for better resolution. **D:** Bouton density profiles of thalamocortical fibers, and locations of VPM (black) and POM (purple) synapses on neurons (in a 5 μm radius subvolume). Rendering was done with BioExplorer. The scale bar on B applies to the whole figure. (Similar panels have been shown in Reimann et al., 2022a, Isbister et al., 2023, and Chindemi et al., 2022.) **E:** Variables of the plasticity model during coincident activation of the pre- and postsynaptic neurons. Left: under *in vitro*-like conditions (taken from Chindemi et al., 2022). Right: same pair of neurons under *in vivo*-like conditions. Schematics on their lefts illustrate the difference between *in vitro*- and *in vivo*-like conditions.

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes

As Graupner and Brunel (2012) modeled $[Ca^{2+}]_i$ of synapses on point neurons phenomenologically, they had to refit their plasticity model parameters to explain different experimental datasets. On the other hand, Chindemi et al. (2022) has shown that a generative model, optimized against STDP protocols from only two pathways can explain a large array of other experimentally measured pathways, thanks to the biophysically detailed model of $[Ca^{2+}]_i$ and the morphological complexity of the neurons. The finding of Chindemi et al. (2022), that one unique plasticity rule can rule them all, crucially depends on the location of synapses on the dendrites. The generative model converts location dependent pre- and postsynaptic $[Ca^{2+}]_i$ peaks into synapse-specific θ_d and θ_p parameters for all excitatory to excitatory (E to E) synapses in the circuit. In our model, we found presynaptically evoked $[Ca^{2+}]_i$ peaks to be three orders of magnitude larger, than the ones evoked by postsynaptic spikes (Supplementary Figure 4.8A). Postsynaptically evoked $[Ca^{2+}]_i$ peaks had a multimodal distribution in the apical dendrites (Supplementary Figure 4.8A right), in line with Landau et al. (2022).

Changes in ρ are then converted into changes in the utilization of synaptic efficacy (U_{SE}), a variable of the Tsodyks-Markram model of STP describing the baseline release probability (Tsodyks and Markram, 1997) and the peak AMPA receptor conductance (\hat{g}_{AMPA} ; equations (4.5) and (4.6) in Methods). As a result of updating U_{SE} as well, short- and long-term plasticity are tightly coupled in the model (Markram and Tsodyks, 1996; Costa et al., 2015; Deperrois and Graupner, 2020). In our network model U_{SE} is also modulated by $[Ca^{2+}]_o$, where a reduction in $[Ca^{2+}]_o$ leads to pathway-specific, non-linear reduction in U_{SE} (Figure 4.1E right; Markram et al., 2015; Ecker et al., 2020). At initiation, synapses are assumed to be at one of the two fixed points (fully depressed ones at $\rho = 0$ and fully potentiated ones at $\rho = 1$) and their assignment to these states is pathway-specific (Supplementary Figure 4.7C3).

4.2.2 Achieving *in vivo*-like network activity

After equipping the circuit with the extra parameters required for long-term plasticity, it was ready to be simulated. To drive network activity, we compensated for missing synaptic input arriving through long-range projections from other brain areas not included in the circuit model (Isbister et al., 2023) and provided inputs through the thalamocortical fibers. Complex phenomena like plasticity are traditionally studied under controlled laboratory conditions *in vitro*, but classical STPD protocols were shown to not induce any plastic changes under *in vivo*-like low $[Ca^{2+}]_o$ (Figure 4.1E, Inglebert et al., 2020; Chindemi et al., 2022). As our broad interest is understanding the rules governing plasticity in living brains, and our modeling pipeline is capable of taking the effects of low $[Ca^{2+}]_o$ into account (Markram et al., 2015), we calibrated the network's activity to mimic *in vivo* conditions. To that end, we calibrated layer-wise spontaneous firing rates and evoked activity to brief VPM inputs matching *in vivo* data from Reyes-Puerta et al. (2015). Spontaneous activity was driven by somatic injection of a layer- and cell type-specific noisy conductance signal (see Isbister et al., 2023 and Methods). By introducing plasticity at all E to E synapses, an additional depolarizing current from VDCCs was added to the model, which made the network more active than its

non-plastic counterpart (Supplementary Figure 4.9A). This required an algorithmic lowering of the amplitude of injected conductances from Isbister et al. (2023) to achieve the same *in vivo*-like layer-wise spontaneous firing rates (Supplementary Figure 4.9B).

Evoked activity was driven by a thalamocortical input stream already described in Ecker et al. (2023b). In short, ten VPM input patterns were repeatedly presented in random order with a 500 ms inter-stimulus interval, together with a non-specific POM input. The ten VPM patterns were defined with varying degrees of overlap in the sets of activated fibers (Figure 4.2A; see Methods). Spike trains delivered on the pattern fibers followed a 100 ms-long inhomogeneous adapting Markov process (Muller et al., 2007). The overlap of the patterns is clearly visible in the firing pattern of each group of fibers corresponding to them (Supplementary Figure 4.10). An exemplary raster plot, characterizing the evoked state of the plastic network is shown on Figure 4.2B.

4.2.3 Sparse synaptic changes induced by long-term plasticity

After achieving *in vivo*-like network activity, we simulated 10 minutes of biological time and measured the changes in synapses with respect to their initial states. The distribution of \hat{g}_{AMPA} remained lognormal, in line with biology (Buzsáki and Mizuseki, 2014; Rößler et al., 2023), and its mean shifted by only 0.07% (+0.5 pS, Figure 4.2C1). This minimal strengthening was achieved by less frequent, but stronger potentiation, and at the same time the network remained stable because of the more frequent, but weaker depression (Figure 4.2C2), without needing to model homeostatic plasticity (Turrigiano and Nelson, 2004; Zenke et al., 2017a). Changes in \hat{g}_{AMPA} are difficult to interpret, as the overall scale of its values is pathway-dependent (Supplementary Figure 4.7C3), i.e., the change associated with full potentiation in one pathway would indicate only partial potentiation in another. Therefore, in the rest of the article we will analyse ρ instead, as it always lies in the $[0, 1]$ interval. While ρ is changing on a faster time scale than \hat{g}_{AMPA} (see equations (4.5) and (4.6) in Methods), the propensity of changes at the end of a 10 minute-long simulation was virtually identical (Figure 4.2C3 vs. D3). When comparing the amount of changes in ρ across time steps, we found that most of the plastic changes happened in the first 1-2 minutes of the simulation, after which they stabilized (Figure 4.2D1). While small changes were still apparent towards the end of the simulation, by visualizing individual synaptic traces we confirmed that most of them oscillated around a dynamic fix point (Figure 4.12A) and the amount of changes in the second half of the simulation were negligible (Figure 4.12B). By splitting synapses at the end of the simulation based on their target neurite type and layer, we learned that an order of magnitude more synapses changed on basal dendrites compared to apical ones, although there are roughly the same amount of apical synapses in the circuit (Figure 4.2D2). Layer 5 (L5) PCs contributed mostly to changes on the basal dendrites, while apical changes happened mostly on L6 PCs.

In addition to looking at changes in individual synapses, we also performed analyses at the connection level. To do so, we averaged ρ values of all (4.1 ± 2.3 ; Supplementary Figure

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes




Figure 4.2: Synaptic changes in large-scale plastic simulations. **A:** Centers of the VPM fibers associated with the ten input patterns in flat map space. Bottom row 3rd: pyramid-like overlap setup of VPM patterns, 4th centers of POM fibers associated with all stimuli. **B:** Raster plot of the microcircuit’s activity and the population firing rates below. The y-axis shows cortical depth. (As cortical layers do not have the same cell density, the visually densest layer is not necessarily the most active. Similar panels have been shown in Ecker et al. (2023b)). **C:** Evolution of \hat{g}_{AMPA} during the 10 minute-long simulation. **C1:** Distribution of \hat{g}_{AMPA} in the beginning and end of the plastic simulation. **C2:** Plastic changes that lead to the shift in the distributions shown in C1 (blue: depression, red: potentiation throughout the figure). **C3:** Layer-wise propensity of changes. **D:** Evolution of synaptic efficacy (ρ). **D1:** L2 norm of changes in ρ across time. Similarly to C1, insert shows distribution of ρ values in the beginning and end of the simulation. **D2:** Layer- and neurite type-wise distribution of non-trivial (neither 0: totally depressed, nor 1: totally potentiated) ρ at the end of the simulation. **D3:** As C3. **E:** Evolution of mean ρ (aggregated over connections). **E1:** L2 norm of changes in mean ρ across time against STDP control (insert, see Methods). **E2:** Plastic changes (in mean ρ) vs. mean pairwise firing rates. **E3:** Same as C3.

4.7C1) synapses mediating a connection and analyzed the propensity of changes as before. As expected, changes at the connection level became more frequent than at the synapse level (Figure 4.2E3 vs. D3). By plotting the propensity of changes against the pairwise mean firing rates of the pre and postsynaptic neurons, we found that the percentage of changes increased as the pairwise firing rates increased (Figure 4.2E2), in line with previous modeling insights (Litwin-Kumar and Doiron, 2014; Graupner et al., 2016). Although, previous theoretical work has shown that embedding simple STDP rules in spiking networks without homeostatic plasticity leads to pathological behavior (Morrison et al., 2007), they relied on higher firing rates. To better understand if only the scale of the simulated network, its biorealistic connection probabilities and the low, *in vivo*-like rates contribute to the sparsity of changes observed in our simulation, we took the 36 M excitatory spikes from our simulation and characterized the propensity of changes resulting from a traditional spike pair-based STDP rule (Gerstner et al., 1996; Kempter et al., 1999; Song et al., 2000; see Methods). While our calcium-based rule stabilized in two minutes after the initial transient, the STDP rule kept inducing the same magnitude of changes throughout all 10 minutes of the simulation (Figure 4.2E1).

To test to what degree was presynaptic transmission required to trigger plasticity, we ran simulations without intrinsic connectivity between the neurons but keeping their extrinsic inputs, or parts thereof, intact (Supplementary Figure 4.11). When neurons received only the somatic conductance injection representing noisy background inputs, we did not observe any changes in mean ρ . When they additionally received the thalamic inputs patterns, we observed changes in ρ , albeit an order of magnitude fewer than in the baseline condition (Supplementary Figure 4.11). Therefore, while the calcium-based plasticity model of Chindemi et al. (2022) is not strictly Hebbian since the effect of postsynaptic firing alone could change synaptic efficacy, presynaptic release was required for most of the observed changes. Lastly, we ran control simulations in connected networks but instead of presenting the patterns, delivered random Poisson spikes on the same VPM fibers at a rate that resulted in the same thalamic spike count. This case was the closest in terms of changes to our baseline case, but still 25% fewer connections underwent plastic changes (Supplementary Figure 4.11), demonstrating the importance of the spatiotemporal structure of the stimuli over simply the firing of the pre- and postsynaptic neurons.

In summary, we observed that $\sim 5\%$ of synapses undergo long-term plasticity under realistic *in vivo*-like conditions in 10 minutes, and most of these synapses are on above-average firing rate L5 PC's basal dendrites. Potentiation dominated in amplitude, while depression counteracted it in frequency, keeping the network stable amidst ongoing plasticity without explicitly considering any homeostatic mechanisms.

4.2.4 More frequent plastic changes within and across cell assemblies

With 95% of synapses remaining unchanged, synaptic plasticity appears to be a highly specific mechanism. We therefore tried to understand the rules that determined which synapses

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes

changed. We specifically hypothesized that plasticity of connections may be structured by the membership of participating neurons in *Hebbian cell assemblies*, i.e., groups of neurons that fire together (Hebb, 1949; Harris, 2005). Our reasoning was as follows: from the parametrization of our plasticity model, we learned that presynaptic spikes contribute orders of magnitude higher calcium than postsynaptic ones (Supplementary Figure 4.8A) if the NMDA receptors are fully unblocked; thus, in order to effectively depolarize the dendrites and unblock NMDA receptors, spikes at low, *in vivo*-like rates must be synchronized in time, as in Hebbian assemblies. Thus, we detected cell assemblies from the *in silico* spiking activity of the 10 minute-long plastic simulation using methods established by experimentalists (Carrillo-Reid et al., 2015; Herzog et al., 2021). In modeling studies, assemblies are usually defined based on their strong internal connectivity, i.e., their structure (Litwin-Kumar and Doiron, 2014; Zenke et al., 2015; Fauth and Van Rossum, 2019; Kossio et al., 2021), but we wanted to use them to restrict our analysis of plastic changes and therefore detected them based on their co-firing function. The rationale for combining the methods above and the full pipeline is described in detail in our previous article, Ecker et al. (2023b) and briefly in the Methods. In short, spikes were binned and bins with significantly high firing rates (Figure 4.3A) were hierarchically clustered based on the cosine similarity of their activation vector (Figure 4.3B1). These clusters correspond to the functional assemblies, with a neuron being considered a member if its spiking activity correlates with the activity of an assembly significantly stronger than chance level (Figure 4.3C). Since time bins and not neurons, were clustered in the first place, this method yields one assembly per time bin and single neurons can be part of several assemblies (Figure 4.3B, D). Assemblies were activated in all stimulus repetitions and a series of three to four assemblies remained active for 190 ± 30 ms, similar to our previous results (Ecker et al., 2023b, Figure 4.3B2). Pattern A elicited the strongest response, while pattern B the weakest, and the responses of patterns H and I were the most similar to each other, as expected, since they share 66% of the VPM fibers (Figure 4.2A). Assembly activations had a well-preserved temporal order - with some of them always appearing early during a stimulus, while others later - and from now on we will refer to them as *early*, *middle*, and *late assemblies*, and will order them in the figures accordingly (Figure 4.3C-E and 4.4A, B).

In line with *in vivo* experiments, these assemblies were detected from functional activity (spikes). However, in our *in silico* approach we have access to the full biorealistic connectome (Reimann et al., 2022a) and can thus investigate how the underlying structure constrains function. In Ecker et al. (2023b) we presented an in-depth analysis of this question (in a non-plastic circuit), so here we will only give an overview of the findings important for this study. Dating back to Hebb (1949), the most commonly accepted structural correlate of cell assemblies is the abundance of recurrent connectivity motifs between participating neurons (Harris, 2005; Song et al., 2005; Perin et al., 2011). In our analysis we also observed that *assembly-indegree*, i.e., the number of afferent connections from an assembly, is a great predictor of a neuron's membership in an assembly (Supplementary Figure 4.13A1). Strong positive interactions were also found *across assemblies*, but only when the direction of innervation reflected the temporal order of assembly activation, e.g., assembly 8 to assembly 12 (A8 and A12 in Figure




Figure 4.3: More frequent changes in cell assembly synapses. **A:** Firing rate of excitatory cells with the determined significance threshold. **B1:** Hierarchical clustering of the cosine similarity matrix of activation vectors of significant time bins (i.e., above threshold in A). **B2:** Clustered significant time bins ordered by the patterns presented. **C:** Number and location of neurons in each cell assembly: flat view on top, depth-profile below. **D:** Jaccard similarity of cell assemblies. **E:** Propensity of changes in cell assemblies. **E1:** Initial mean efficacy (ρ) of within- and cross-assembly synapses. **E2:** Propensity of depression and potentiation of within- and cross-assembly synapses. As assemblies are overlapping (see D) single synapses are taken into account for many different pre- and postsynaptic assembly pairings. (Similar panels (except E) have been shown in Ecker et al., 2023b).

4.3B2 responding to patterns H and I). These results, combined with the biophysics of the plasticity model, suggest that connections within an assembly and the ones between temporarily ordered assemblies, are expected to undergo plastic changes with a higher probability.

When checking the propensity of changes within and across assemblies, we indeed found more synapses undergoing long-term plasticity (Figure 4.3E2). While only 3.5% of synapses depressed in the whole dataset, we found up to 10.5% when restricting the analysis to assemblies. Similarly, compared to 1.5% of all synapses potentiating, we observed up to 4.2% when

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes

restricting to assemblies. Interestingly, large values were found in the off-diagonal entries (Figure 4.3E2), i.e., synapses across assemblies underwent more plastic changes than the synapses within these assemblies. In our model, the initial ρ values are pathway-specific and highest in L4 pathways (Brémaud et al., 2007; Supplementary Figure 4.7C3). Therefore, early assemblies, with large number of L4 cells have a higher than average initial ρ (Figure 4.3C and E1 respectively), thus their synapses are more likely to depress (Figure 3.3E2). As early assemblies are stimulus specific, and thus not part of a Hebbian phase-sequence, synaptic depression between these cells can be seen as some kind of orthogonalization of the stimulus responses. On the other hand, late assemblies, that are predominantly composed of cells from the deep layers, have a low initial ρ (Figure 4.3E1; Supplementary Figure 4.7C3) and synapses towards them are more likely to potentiate. These assemblies are mostly non-specific and participate in all phase-sequences, thus the potentiation of their efferents means a general strengthening of the stimulus response as a whole.

Together these results indicate that, in line with 70 years old predictions, cells that fire together wire together (Hebb, 1949). Our contribution lies in making the qualitative statement above into a quantitative one: Under *in vivo*-like conditions cells that fire together more than expected have three times higher chances of changing the efficacy of their connections.

4.2.5 Synapse clustering contributes to the emergence of cell assemblies, and facilitates plasticity across them

In addition to co-firing, a group of innervating neurons is more effective in depolarizing a given dendritic branch if they all send synapses to the same branch, i.e., they form a spatial synapse cluster (Farinella et al., 2014; Iacaruso et al., 2017; Kastellakis and Poirazi, 2019). To quantify this trend, we previously defined the synaptic clustering coefficient (*SCC*) with respect to an assembly, based on the path distances between synapses from that assembly on to a given neuron (see Ecker et al., 2023b and Methods). For the assemblies detected in this study, we also found *SCC* to be a good predictor of a neuron's membership in an assembly (Supplementary Figure 4.13A2), although the effect was less than half as strong as that of assembly-indegree. We used assembly-indegree and *SCC* to select the 10 most innervated L5 TTPCs (thick-tufted pyramidal cells) within a cell assembly and then explicitly detected spatial clusters of synapses, defined as at least 10 synapses within a 20 μm stretch of a single dendritic branch (see Methods).

For our next analysis, we grouped all synapses on these 10 selected neurons per assembly into four categories based on assembly membership of the presynaptic neuron and whether the synapse was part of a cluster or not (see exemplary clustered assembly synapses on Figure 4.4E1). Then, we quantified the likelihood of plastic change in each category by contrasting the conditional probability of observing it in a given category with the probability of observing any change irrespective of the category (see equation (4.9) in Methods; Figure 4.4A2). Note that a nonzero value for one category always has to be compensated by a nonzero value with

opposite sign in another. Surprisingly, clustered within-assembly synapses were not likely to undergo any changes. When we repeated the analysis on the initial ρ values, we found that early and middle assembly synapses, especially the clustered ones, are very likely to be initialized as fully potentiated (Figure 4.4A1). On the other hand, synapses within the late assemblies were likely to be initialized in the fully depressed state, but were likely to change. Furthermore, when comparing the amplitude of changes across conditions with a 2-way ANOVA, we found that clustered within-assembly synapses depress to a smaller degree than the other ones (Figure 4.4C). When we checked the temporal evolution of within-assembly synapse cluster ρ values, we saw that while some of the synapses underwent small constant changes, most of them changed at the same time (vertical stripes on Figure 4.4E2). Thus the picture emerging is as follows: early and middle assemblies are partially defined by clustered (both spatial and functional) synapses that are initialized as fully potentiated. These synapses are unlikely to change, but when they do, they depress less than the others, and would converge back to $\rho = 1.0$ in absence of activity, as they do not cross the $\rho = 0.5$ unstable fix point. These stable early assemblies can therefore function as a stable backbone amid ongoing synaptic plasticity.

In our previous investigation, we found that most changes happened across assemblies, so we extended the analysis described above to cross-assembly synapses. Here, the picture was reversed: cross-assembly synapses that were part of a spatial cluster were likely to be initiated as fully depressed and then had a high chance of undergoing potentiation (Figure 4.4B). Interestingly, the amplitude of this potentiation was significantly less than that of the other groups' (Figure 4.4D), but on average, still enough to cross the $\rho = 0.5$ unstable fix point.

Together with the previous results, this suggests that synapses between assemblies are more likely to change, which is even more pronounced if these synapses form a cluster on the postsynaptic dendrite.

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes




Figure 4.4: Changes in synapses participating in spatial clusters. A: Michelson contrast (equation (4.9) in Methods) of probabilities of plastic changes within assembly synapses. Depression on top (blue colormap) and potentiation below (red colormap). Grey parts of the colormaps indicate lower than expected probabilities of observing synapses in a given state, given that it falls into the indicated category, while neon green means no synapses found in the given category. **A1:** Initial ρ , **A2:** plastic changes in ρ . **B:** Same as A, but for cross-assembly synapses (postsynaptic assembly fixed to A15). **C:** Distribution of within-assembly $\Delta\rho$ s across the four conditions. Boxes show all values, while black dots are 1000 samples from each. Significance test was run on the balanced samples (1000 each): 2-way ANOVA and post-hoc Tukey's test: *: $p \leq 0.05$, **: $p \leq 0.01$, ***: $p \leq 0.001$, ****: $p \leq 0.0001$. **D:** Same as C, but for cross-assembly synapses (data from several postsynaptic assemblies, not only A15 shown on B). (Caption continues on the next page.)

E1: Changes in within-assembly, clustered synapse on an exemplary A13 neuron. Rendering was done with `NeuroMorphoVis` Abdellah et al. (2018). Neurite diameters are scaled (x2) for better resolution. (Synapse diameters are arbitrary.) **E2:** Temporal evolution of the (~ 1000) synapses on basal dendrites shown on E1.

4.2.6 Redistribution of assembly efficacies and prolonged stimulus-specific responses characterize the network after plasticity

In the beginning of our study we used cell assemblies only as a powerful tool to restrict our analysis of plastic changes to biophysically motivated subpopulations of neurons. On the other hand, the evolution of assemblies in terms of their composition and association with stimuli is used to examine the functional consequences of plasticity and the stability of the neural code in contemporary literature (Fauth and Van Rossum, 2019; Kossio et al., 2021; Pérez-Ortega et al., 2021). From our investigation we have learned that the early assemblies are defined by clustered fully potentiated synapses at initialization. As ongoing plasticity strengthens their connections to the late assemblies we wondered what would happen to the assemblies if we detected them after the plastic changes.

To study this, we stabilized the network's state *after* our 10 minute-long plastic simulation, i.e., based on the ρ values in the last time step, assigned synapses to either fully potentiated (last $\rho \geq 0.5$) or fully depressed states and updated not only the U_{SE} and \hat{g}_{AMPA} values, but also the peak NMDA conductances (\hat{g}_{NMDA}) accordingly. Then we ran 2 minute-long, non-plastic simulations of this network and compared the resulting assemblies to the ones detected in a non-plastic simulation of the network *before* plasticity, i.e., in its naive state. Note that the stimulus streams presented were identical between the two cases. From a high level comparison of the network states before vs. after plasticity we learned that the firing rates increased (Supplementary Figure 4.14A1 left) but the pairwise spike correlations only slightly increased in line with recent findings (Oby et al., 2019; Feulner et al., 2022). Nonetheless spike time reliability of individual neurons increased (Figure 4.5A; see Methods). The observed increase in firing rate might explain the increase in spike time reliability after plasticity, as the two measures are correlated (Cutts and Eglén, 2014). Plotting pattern-specific peri-stimulus time histograms (PSTHs) before and after plasticity revealed a general lengthening of the late phases of the response and increased amplitudes for selected patterns (Figure 4.5B).

For a better comparison of assemblies, we ran five repetitions of both cases, and compared *consensus assemblies*, i.e., the sets of neurons that were reliably part of a given assembly across repetitions (Figure 4.5C1 and D1; see Ecker et al., 2023b and Methods). We found more consensus assemblies after plasticity than before (twelve vs. nine, compare Figure 4.5C2 and D2). The emergence of more consensus assemblies after plasticity is not an artefact of grouping assemblies together, as the optimal number of assemblies (assessed by Davis-Bouldin index (Davies and Bouldin, 1979); see Methods) was higher in four out of five repetitions after plasticity, and equal in one repetition (Supplementary Figure 4.14C). On the other hand, both the Davis-Bouldin index of the resulting clusters and the cosine similarity of consensus

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes

assembly counts across repetitions decreased after plasticity (Supplementary Figure 4.14D2). The sizes of consensus assemblies were similar before and after plasticity (Supplementary Figure 4.14E). Further comparing consensus assemblies before and after plasticity has revealed that corresponding pairs had more than 50% of their neurons shared (Figure 4.5E).

To gauge the functional consequences of plastic changes, we studied the functional and structural connectivity of consensus assemblies detected before and after plasticity. When comparing the mean pairwise spike correlations (a method usually applied to derive functional connectivity; see Methods) of neurons belonging to consensus assemblies, we observed a general decrease in the early ones and an increase in middle and late ones (Figure 4.5F). This was accompanied by the same arrangement of changes in the structural connectivity of the same consensus assemblies (assessed by comparing ρ values). We saw similar trends in the spike time reliability (see Methods) of individual neurons, i.e., a decrease for early, and an increase for middle and late consensus assembly neurons. More generally, we found that early assemblies grew less correlated with weaker internal connectivity through plasticity, which trend was weakened in the subset of neurons that remained part of the assembly (compare columns of Figure 4.5F). Conversely, in middle and late assemblies correlations and connections grew stronger, especially so in the neurons that were members of the consensus after plasticity.

As plasticity in the cortex changes not only \hat{g}_{AMPA} as in the hippocampus but also U_{SE} (Markram and Tsodyks, 1996; Selig et al., 1999; Sjöström et al., 2003; Costa et al., 2015; Chindemi et al., 2022), there is a *redistribution of synaptic efficacy* towards earlier spikes during high-frequency firing. This redistribution happens because the increased U_{SE} makes the STP profile of potentiated connections more depressing (Supplementary Figure 4.14B left; Markram and Tsodyks, 1996). However, in our simulation, we rarely observed high-frequency firing and also found the STP profile of potentiated connections to be facilitating at the low *in vivo* $[Ca^{2+}]_o$ (Supplementary Figure 4.14A2 and B right respectively). Thus, while Markram and Tsodyks (1996) showed a redistribution of synaptic efficacy after plasticity at the single connection level *in vitro*, we found a redistribution at the network level under *in vivo* like conditions: efficacy shifted towards synapses targeting the deeper layers of the cortex. Interestingly, while the firing rates only increased slightly in the significant time bins, there was a more pronounced increase when we compared them during the whole 2 minute-long simulation (compare Supplementary Figure 4.14A1 left to right). This strongly suggests that this network level redistribution of efficacy and the strengthening of late consensus assemblies lead to their reactivation during spontaneous activity, in line with experimental findings (Miller et al., 2014; Carrillo-Reid et al., 2015; Stringer et al., 2019; Herzog et al., 2021; Trägenap et al., 2022).




Figure 4.5: Changes in cell assemblies after plasticity. **A:** Functional network features extracted from spike times of non-plastic simulations *before* plasticity, i.e., in the naive circuit vs. *after* the 10 minute-long plastic simulation. Left: Pairwise spike correlation. Right: Spike time reliability (r_{spike}) measured over five repetitions of the same 2 minute-long simulations with the same input (see Methods). **B:** PSTHs by patterns before vs. after. (Only significant time bins are take into account, see assembly detection in Methods). **C:** Non-plastic consensus assemblies before plasticity. **C1:** Jaccard similarity based hierarchical clustering of assemblies from the five simulation instances. **C2:** Significant time bins from one of the repetitions, ordered by patterns presented, and colored by the consensus assemblies (not the ones detected from that instance). **D:** Same as B, but for non-plastic consensus assemblies after plasticity. (Caption continues on the next page.)

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes

E: Jaccard similarity of consensus assemblies detected before and after plasticity. **F:** From left to right: Changes in pairwise spike correlation, ρ , and r_{spike} of within consensus assembly neurons. Colors indicate changes (after - before), while columns indicate at which point the consensus assembly was detected. Right: Changes in total consensus assembly counts per pattern. Error bars are over the five repetitions. **G:** Input-output map. **G1:** Input distances as the Earth mover's distance of the VPM fiber locations (see Figure 4.2A). Insert shows the overlap (based on Hamming distance) of pattern fibers. **G2:** Output distances are calculated as the (normalized) Euclidean distance of pattern evoked consensus assembly cluster over repetitions (see Supplementary Figure 4.14B). **G3:** Correlation of distances from G1 and G2. (Similar panels (except A, B, and F) have been shown in Ecker et al., 2023b.)

Lastly, we further analyzed the total duration of consensus assembly responses to different patterns. In line with the prolonged PSTHs, we found a general increase (consensus assemblies active for 190 ± 45 ms before vs. 200 ± 60 ms after plasticity), and could trace it back to selected patterns A, E, H, and I (Figure 4.5D right). This provides the most likely explanation for the increased number of assemblies after plasticity: the higher number of significant time bins simply lead to a higher number of optimal clusters (given our metric). The activation sequence of consensus assemblies can be seen as a low-dimensional representation of the complex, high-dimensional activity of the network's response to different patterns. Following our previously established methods (Ecker et al., 2023b), we correlated the Earth mover's distances between the locations of the VPM fibers making up the input patterns (Figure 4.5C1), and the normalized Euclidean distances of output consensus assembly sequences across repetitions (Figure 4.5G2). We observed an increase in the input-output distance correlation after plasticity ($r = 0.443$ vs. $r = 0.357$; Figure 4.5G3). This increased input-output correlation after plasticity can partially be explained by the prolonged stimulus-specific assembly sequences.

In summary, when comparing assemblies before and after plasticity, we found that the network became more specific to the patterns it was exposed to. This manifested in assemblies splitting, weakening of early and strengthening of the late assemblies and the consequent prolonged assembly responses to specific patterns.

4.2.7 Network topology changes are parametrized by input stimuli

Increased pattern specificity after plasticity indicates that the plastic changes are indeed not random, but stimulus-driven. To better characterize this, we ran 2 minute-long plastic simulations in which we only presented a single pattern (with the same 500 ms inter-stimulus interval as before). We repeated this paradigm three times, for all 10 patterns, and analysed changes in mean ρ matrices as before. The propensity of changes in the connections was in line with the baseline ones (compare Supplementary Figures 4.11B and 4.15). While the number of connections changing was similar across patterns, we wondered if there is any pattern-specific information in them.

To investigate this, we used an input-output distance correlation analysis, similar to the one employed to compare consensus assemblies before and after plasticity. In detail, we

correlated the Earth mover's distance of the input patterns (Figure 4.5C1 as before) with the Euclidean distance of the steady state (last time step) mean ρ values and found a clear and strong correlation between them ($r = 0.666$, $p < 0.0001$; Figure 4.6A left). As controls we ran the analysis with other distance metrics as well. Neither Hamming distance (taking only the identity of changing connection into account) nor Earth mover's distance (taking only the distribution of steady state mean ρ values into account) showed a clear correlation with the distance of pattern fibers (Figure 4.6A middle and right respectively), indicating, that the whole network structure needs to be taken into account to describe the relationship between its steady state and the input patterns.

To further explore the structure of the changes, we focused on the plastic connections that evolved in the same direction (i.e., potentiation or depression) across all three repetitions for all patterns, which was around 40% of all changing connections (Supplementary Figure 4.15). As we found pattern-specific information in the changing connections, in the next step we analyzed to what degree the subnetworks they defined are determined only the neurons composing them. We did this by comparing them to random subnetworks of the entire circuit with the same pre- and post-synaptic populations and the same number of connections between them. We observed, that distributions of the changing pathways are different than expected from the network structure and the pre- and postsynaptic populations alone (Figure 4.6B1). To quantify this difference, we counted a particular class of motifs, directed simplices of dimension k , which are motifs on $k + 1$ neurons, which are all-to-all connected in a feed-forward fashion (Figure 4.6B2 inset). These motifs have previously shown to be linked to network function (Reimann et al., 2017b) as well as quantify complexity of the network's topology (Kahle, 2009; Bobrowski and Kahle, 2018). We found strong overexpression of these simplices in the subgraphs, compared to their random controls. In particular, the maximal simplex dimension found in the subgraphs was always one higher than in the corresponding controls (Figure 4.6B1).

While we learned that different connections change when different patterns are presented and the connections and the network topology they define are not entirely defined by the pre- and postsynaptic populations, so far have not linked the changes to individual patterns. To do so, we used methods developed in Ecker et al. (2023b) and first studied the propensity of changes against *pattern-indegree* i.e., the number of VPM fibers belonging to a pattern that innervate a neuron. The propensity of changes increased as either the pre- or the postsynaptic side of the connection's pattern-indegree increased (Figure 4.6C). Moreover, we confirmed that the effect of pattern-indegree of the pre- and postsynaptic neurons are not independent by computing the conditional mutual information between them and the probability of their connection to change. As expected, the mutual information is non-zero (between 0.0126 and 0.0291) and much larger than the one obtained for corresponding random controls (between 2.27×10^{-7} , 4.52×10^{-7}). Thus, in our last analysis we investigated the joint distribution and characterized the propensity of changes against pattern-indegree of both pre- and postsynaptic neurons. The propensity of both depression and potentiation grew rapidly with patten-indegree (Figure 4.6D1 and D2 respectively). While depressing and potentiating connections totaled to only

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes




Figure 4.6: Topology of changing subnetworks in response to single pattern presentations. A: Input-output distance correlations. **A1:** Different (from left to right: Euclidean, Hamming, Earth mover's) distances of mean ρ matrices. Three repetitions for each of the 10 input patterns. **A2:** Correlation of input distances (as in Figure 4.5C1) and distances of mean ρ matrices above on A1. **B1:** Layer-wise distribution of consistently changing (three out of three repetition) connections in response to presenting pattern A. Below its control, which was generated by taking the same number of connections between the same pre- and postsynaptic populations. **B2:** Directed simplex counts in subnetworks from B1 (but for all patterns). Colors correspond to simplex counts of the consistently changing subnetworks, while black to their controls. Note that by construction the controls must have the same number of 0- and 1-simplices which correspond to the number of cells and connections in the subnetwork. Insert illustrates simplex dimension. **C:** Propensity of changes vs. pattern-indegree of the presynaptic (left) or postsynaptic (right) neurons. **D:** Propensity of changes (split for depression **D1** and potentiation **D2**) against the pattern indegree of both pre- and postsynaptic neurons.

3% in the whole network, the amount of depressing connections was above 15%, and above 5% for potentiation when both sides of the connection were highly innervated by the VPM fibers. They reach the highest values for patterns A, H, and I, the same ones whose responses were strengthened by plasticity when all patterns were presented. Curiously, pattern-indegree of the presynaptic side was less important for potentiation, where pattern-indegree of the postsynaptic side was more predictive. The intuitive explanation for this is the following: if input from the VPM fibers depolarized the postsynaptic dendrites enough, then a successful release from any presynaptic neuron (independent of their pattern-indegree) caused Ca^{2+} entering through the at least partially opened NMDA receptors.

In summary, the network evolved differently when single patterns were presented, and the distance between the steady states achieved reflected the geometric distance between the patterns. Moreover, the steady state reached for each pattern could not be determined only by the pre- and postsynaptic populations of the changing connections, but by their precise location within the network as well as their pattern-indegree. The propensity of depression increased in connection in which both pre- and postsynaptic neurons was strongly innervated by pattern fibers, while for potentiation postsynaptic pattern-indegree dominated. Moreover, these effects are not independent.

4.3 Discussion

Using a detailed, large-scale cortical network model equipped with a calcium-based model of long-term functional plasticity, we have examined changes in synaptic efficacy in response to repeated presentation of 10 different stimuli over 10 minutes of biological time, under *in vivo*-like conditions. Our principal observations in this bottom-up simulation framework are as follows: (1) Plastic changes were sparse, affecting only 5% of the synapses. A balance between occasional large-amplitude potentiation and more frequent depression kept the network stable without explicitly modeling homeostatic plasticity. (2) Plastic changes were largely determined by the anatomical structure of synaptic connectivity and its relation to functional units, i.e., changes were most likely between co-firing cell assemblies and at clustered synapses. (3) Early-responding cell assemblies were defined by clustered synapses initialized as fully potentiated and remained fairly stable. In contrast, their synapses to late-responding assemblies underwent three times more changes than expected, resulting in prolonged and more reliable responses to selected patterns after plasticity. (4) Changes in the network evoked by the presentation of individual patterns reflected the geometric distance between the patterns themselves. The structure of these changes could be partially explained by the innervation of the pre- and postsynaptic neurons by the pattern fibers, though the populations alone are not enough to determine these changes, since the changing connections between them are not random.

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes

The first observation (1) is quite significant considering that we did not design the learning rule to be sparse and stable. In previous models of plastic networks, additional mechanisms were needed to keep the network's activity stable, not to mention the sparsity of changes (Turrigiano and Nelson, 2004; Litwin-Kumar and Doiron, 2014; Zenke et al., 2015, 2017a; Fauth and Van Rossum, 2019; Kossio et al., 2021). The machine learning community is also aware of the importance of sparse changes, as in continual learning one has to balance plasticity and stability to avoid catastrophic forgetting (McCloskey and Cohen, 1989; Ratcliff, 1990). In recent years, they have come up with impressive techniques that mask connections to improve the performance of deep neural networks (Zenke et al., 2017b; Mallya and Lazebnik, 2018; Frankle and Carbin, 2019), whereas in our model it emerged naturally from the high level of biophysical detail. Of course, the amount of data that deep networks are expected to store far exceeds the 10 patterns used here, and it is outside of our scope to find the maximal capacity of our network. On the other hand, we know from theoretical work that for bistable synapses operating on multiple time scales, capacity scales with the square root of the number of synapses (Crick, 1984; Fusi et al., 2005).

The second observation (2) can be explained from the biophysics of the plasticity model and links our results to the classical work of Hebb (1949) as well as the recent literature on synapse clustering (Farinella et al., 2014; Iacaruso et al., 2017; Kastellakis and Poirazi, 2019). With respect to the latter, we would highlight that our synapses are stochastic and the release probability between PCs is ~ 0.04 at the simulated low $[Ca^{2+}]_o = 1.05$ mM (Jones and Keep, 1988; Borst, 2010; Markram et al., 2015; Ecker et al., 2020). Therefore, care should be taken when comparing our results with either glutamate uncaging experiments, which bypass the presynaptic machinery (Pettit et al., 1997; Losonczy and Magee, 2006), or with other modeling studies that use deterministic synapses (Poirazi et al., 2003; Farinella et al., 2014; Ujfalussy and Makara, 2020). In relation to observations (2-4): While we were able to use cellular and subcellular features of the model's biorealistic structural connectivity (Reimann et al., 2022a) to predict plastic changes to a certain degree, this process also highlighted that many other rules govern plasticity at the network level. Further analysis considering the embedding of a connection in the entire network and thus the state of the whole network may be able to provide that explanation.

According to the contemporary view of L5 PCs, sensory bottom-up inputs target their basal dendrites, and top-down information arrives at the apical ones, and the coincidence activation of basal and apical inputs is encoded by bursts of action potentials (Larkum, 2013; Naud and Sprekeler, 2018). During bursts of action potentials, the bAPs propagate to the distal apical dendrites better (Williams and Stuart, 1999), enough to turn apical depression into potentiation (Letzkus et al., 2006). Therefore, bursts are not only important for coding, but for plasticity as well. L5 TTPC bursts were rare in our simulations, as the model is based on an early developmental stage (P14-16: juvenile rats) and burst firing only becomes prominent as the animals mature (Zhu, 2000). On the other hand, burst firing could probably be rescued with stronger top-down input. As the top-down input represents context/brain state and is thought to serve as an error/target signal for learning, it has to be highly specific (Makino,

2019). Although we added inputs from POm fibers as they were shown to gate plasticity in L2/3 PCs *in vivo* via dis-inhibiting the distal dendrites (Gambino et al., 2014; Williams and Holtmaat, 2019), we only used randomly distributed fibers, to keep our setup simple in this first investigation. For a profound understanding of the role of bursts in apical plasticity, more learning/task-related studies with more precise top-down input would be needed in the future.

We presented here what we believe to be a new way of studying unsupervised learning and plasticity in the cortex by taking the diversity of cell types and morphologies into account, modeling connections as multi-synaptic, validating synapse anatomy and physiology, modeling synapses with bistable dynamics, and simulating the network in an *in vivo*-like state. On the other hand, building a model of this scale and detail required gathering and systematic integration of a large amount of data over several years (Markram et al., 2015; Chindemi et al., 2022; Reimann et al., 2022a; Isbister et al., 2023). As the first of its kind, the work presented here did not exhaust all the additional understanding one could possibly gain from the high level of detail. To facilitate that process, we are open-sourcing our model alongside detailed instructions to launch simulations and analyze the results (see Data and code availability). As any other model, it has several assumptions (listed in Table 4.1) and limitations and can best be proven wrong and iteratively updated in a community-driven manner. Simulating the model requires a performant hardware and software infrastructure (e.g., we needed 16.5M core hours to run the simulation presented in this manuscript). With respect to the second part we are continuously improving the performance and efficiency of the simulator (Kumbhar et al., 2019).

Table 4.1: List of assumptions

-
- | | |
|---|--|
| 1 | As we combined the models of Isbister et al. (2023) and Chindemi et al. (2022) all assumptions therein are inherited. Of particular interest: |
| 2 | Extracellular recordings are assumed to have to same bias across layers and neuron populations. Furthermore it is assumed that different inhibitory subpopulations require the same ammount of input compensation. |
| 3 | The extracellular magnesium concentration of 1 mM used <i>in vitro</i> is assumed to be representative of the <i>in vivo</i> level. |
| 4 | As the plasticity model of Chindemi et al. (2022) is based on $[Ca^{2+}]_i$, by using it we assumed that other factors, like metabotropic glutamate receptors, endocannabinoid release, BDNF signaling are negligible for the network-level effects of plasticity that we investigated. |
| 5 | Spines are assumed to be separate biochemical compartments, i.e., $[Ca^{2+}]_i$ of the dendrites does not influence that of the synapses. |
| 6 | By detecting a single set of assemblies in the 10 minute-long plastic simulation we assumed that assemblies are stable on that time scale. |
-

4.4 Methods

4.4.1 Calcium-based plasticity model

The calcium-based plasticity model is fully described in Chindemi et al. (2022), but a minimal description of it can be found below. Synaptic efficacy (ρ) is based on the Graupner and Brunel (2012) formalism, which exhibits a bistable dynamics ($\rho = 0$ fully depressed, $\rho = 1$ fully potentiated, and $\rho = 0.5$ unstable fix point) described as:

$$\tau \frac{d\rho}{dt} = -\rho(1-\rho)(\rho_* - \rho) + \gamma_p(1-\rho)\Theta(Ca^*(t) - \theta_p) - \gamma_d\rho\Theta(Ca^*(t) - \theta_d) \quad (4.1)$$

where τ is the time constant of convergence, θ_d and θ_p are depression and potentiation thresholds, γ_p and γ_d are depression and potentiation rates and Θ is the Heaviside function. The dynamics of $[Ca^{2+}]_i$ in spines was modeled as:

$$\frac{d[Ca^{2+}]_i}{dt} = (I_{NMDAR}^* + I_{VDCC}) \frac{\eta}{2FX} - \frac{[Ca^{2+}]_i - [Ca^{2+}]_i^{(0)}}{\tau_{Ca}} \quad (4.2)$$

where I_{NMDAR}^* and I_{VDCC} are calcium currents through NMDA receptors and VDCCs, η is the fraction of unbuffered calcium, F is the Faraday constant, X is the spine volume, $[Ca^{2+}]_i^{(0)}$ is the resting value of $[Ca^{2+}]_i$, and τ_{Ca} is the time constant of free (unbuffered) calcium clearance. I_{NMDAR}^* depends on the state of the Mg^{2+} block. This nonlinear voltage dependence is described with the Jahr and Stevens (1990) formalism, with parameters fitted to cortical recordings from Vargas-Caballero and Robinson (2003).

Inspired by previous theoretical insights (Rubin et al., 2005), a leaky integrator of $[Ca^{2+}]_i$ was introduced (Ca^*) to slow down its time course instead of modeling enzymes downstream of calcium (e.g. CamKII as others did (Mäki-Marttunen et al., 2020; Rodrigues et al., 2022)):

$$\frac{dCa^*}{dt} = -\frac{Ca^*}{\tau^*} + [Ca^{2+}]_i - [Ca^{2+}]_i^{(0)} \quad (4.3)$$

where τ^* is the time constant of the integrator. Updates in ρ were done based on this Ca^* variable crossing θ_d and/or θ_p (see equation (4.1)). The two synapse-specific threshold were derived based on peaks in $[Ca^{2+}]_i$ caused by pre- and postsynaptic spikes, c_{pre} and c_{post} respectively. To measure these parameters for all 312,709,576 synapses, simulations of single cells were run, in which either the pre- or the postsynaptic cell was made to fire a single action potential and the local $[Ca^{2+}]_i$ was monitored in each synapse. Since 8% of L6 PCs could not be made to fire a single action potential (only bursts), synapses on those cells (10,995,513 in total) were assumed to be non-plastic, i.e., their thresholds were set to a negative value that could not be crossed. Similarly, as the plasticity of connections between L4 spiny stellate cells was shown to be non-NMDA dependent (Egger et al., 1999; Chindemi et al., 2022) those connections were made non-plastic. For the remainder of cells, θ_d and θ_p were derived as

follows:

$$\begin{bmatrix} \theta_d \\ \theta_p \end{bmatrix} = \begin{bmatrix} a_{00} & a_{01} \\ a_{10} & a_{11} \end{bmatrix} \times \begin{bmatrix} c_{pre} \\ c_{post} \end{bmatrix} \text{ or } \begin{bmatrix} b_{00} & b_{01} \\ b_{10} & b_{11} \end{bmatrix} \times \begin{bmatrix} c_{pre} \\ c_{post} \end{bmatrix} \quad (4.4)$$

where $a_{i,j}$ and $b_{i,j}$ are constants optimized during model fitting for apical and basal dendrites respectively. Changes in ρ were then converted by low-pass filtering into changes U_{SE} and \hat{g}_{AMPA} as follows:

$$\frac{dU_{SE}}{dt} = \frac{\bar{U}_{SE} - U_{SE}}{\tau_{change}} \quad \text{where} \quad \bar{U}_{SE} = U_{SE}^{(d)} + \rho \left(U_{SE}^{(p)} - U_{SE}^{(d)} \right) \quad (4.5)$$

$$\frac{d\hat{g}_{AMPA}}{dt} = \frac{\bar{g}_{AMPA} - \hat{g}_{AMPA}}{\tau_{change}} \quad \text{where} \quad \bar{g}_{AMPA} = \hat{g}_{AMPA}^{(d)} + \rho \left(\hat{g}_{AMPA}^{(p)} - \hat{g}_{AMPA}^{(d)} \right) \quad (4.6)$$

where $U_{SE}^{(d)}$, $U_{SE}^{(p)}$, $\hat{g}_{AMPA}^{(d)}$ and $\hat{g}_{AMPA}^{(p)}$ are the fully depressed (d) and fully potentiated (p) values of the given variables, in-between which they evolve. All values (fixed and optimized alike) are listed in Chindemi et al. (2022). Just to give a rough idea of time scales: $[Ca^{2+}]_i$ evolves at the timescale of tens of ms, Ca^* on the hundreds of ms, while changes in ρ are converted to changes in U_{SE} and \hat{g}_{AMPA} in seconds.

4.4.2 *In vivo*-like spontaneous and evoked activity

The calibration process that leads to the *in vivo*-like spontaneous activity is fully described in Isbister et al. (2023), but a minimal description and a list of the parameters used in this article can be found below. As extracellular recordings are known to overestimate firing rates (Wohrer et al., 2013), a spectrum of spontaneous states at fixed percentage of the rates reported in Reyes-Puerta et al. (2015) were calibrated (Isbister et al., 2023). Matching specific firing rates *in silico* was achieved by iterative adjustments of layer and cell-type (excitatory/inhibitory) specific somatic conductance injection (following an Ornstein-Uhlenbeck process Destexhe et al., 2001). The spontaneous state used in the article is characterized by the parameters: $[Ca^{2+}]_o = 1.05$ mM (Jones and Keep, 1988), percentage of reported firing rates = 40%, the coefficient of variation (CV; std/mean) of the noise process = 0.4.

The thalamic input patterns, and the spike trains delivered on them are fully described in Ecker et al. (2023b), but a minimal description, highlighting the changes applied in this study, can be found below. First, the flat map location of VPM fibers avoiding the boundaries of the network were clustered with k-means to form 100 bundles of fibers. Second, the four base patterns (A, B, C, and D) were formed by randomly selecting four non-overlapping groups of bundles, each containing 12% of them. The remaining six patterns were derived from these base patterns with various degrees of overlap: three patterns as combinations of two of the base ones (E, F, G), two patterns as combinations of three of the base ones (H, I), and one pattern as a combination of all four base ones (J). Third, the input stream was defined as a random presentation of these 10 patterns, in a balanced way. Last, for each pattern presentation, unique spike times were generated for its corresponding fibers following a 100 ms-long inhomogeneous adapting

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes

Markov process (Muller et al., 2007). The maximal rate of the VPM fibers was set to 17.5 Hz (compared to 30 Hz for the non-plastic circuit in Ecker et al., 2023b) and half of that for POM.

4.4.3 Network simulations

Simulations were run using the NEURON simulator as a core engine with the Blue Brain Project's collection of hoc and NMODL templates for parallel execution on supercomputers (Hines and Carnevale, 1997; Kumbhar et al., 2019; Awile et al., 2022; see Data and code availability). Simulating 10 minutes of biological time with reporting the state of all synapses (in every second) took 2,350,000 core hours ($\sim 4x$ more than the corresponding non-plastic circuit without reporting), on our HPE based supercomputer, installed at CSCS, Lugano. Simulations were always repeated at least three times to assess the consistency of the results.

4.4.4 Evaluating control STDP rules

To compare the amount of changes induced by Chindemi et al. (2022) with classical plasticity rules, the 36,573,737 excitatory spikes from the 10 minute-long simulation were evaluated with pair-based STDP rules (Gerstner et al., 1996; Kempter et al., 1999; Song et al., 2000). Synaptic weights evolved as follows under the STDP rule:

$$\Delta w_+ = A_+ \exp\left(-\frac{\Delta t}{\tau_+}\right) \text{ at } t_{post} \text{ if } t_{pre} < t_{post} \quad (4.7)$$

$$\Delta w_- = A_- \exp\left(\frac{\Delta t}{\tau_-}\right) \text{ at } t_{pre} \text{ if } t_{pre} > t_{post} \quad (4.8)$$

where t_{pre} and t_{post} are the times of pre- and postsynaptic spikes, $\Delta t = t_{post} - t_{pre}$ is the difference between them; $A_{\pm} = 0.05$ describe the weight update, which decayed exponentially with time constants $\tau_{\pm} = 20$ ms. The STDP rule was implemented in Brian2 (Stimberg et al., 2019).

4.4.5 Cell assembly detection

The combination of methods from Carrillo-Reid et al. (2015) and Herzog et al. (2021) yielding the assembly detection pipeline is fully described in Ecker et al. (2023b), but a minimal description, highlighting the changes applied in this study, can be found below. First, spikes of excitatory cells were binned using 20 ms time bins (Harris et al., 2003). Second, time bins with significantly high firing rates were determined as crossing a threshold defined as the mean activity level plus the 95th percentile of the standard deviation of 100 shuffled controls. These shuffled controls were less strict than in Ecker et al. (2023b). Unlike in the original study, where spikes were only shifted by one time bin forward or backward (Carrillo-Reid et al., 2015), spikes were shifted by any amount. This change was introduced because the network's response to the same patterns was more variable in the plastic simulations, and to not miss any of them, a lower threshold was more fitting. Third, based on the cosine similarity of activation

vectors, i.e., vectors of spike counts of all neurons in the given significant time bins, a similarity matrix was built (Carrillo-Reid et al., 2015). Fourth, this similarity matrix was hierarchically clustered using Ward's linkage (Montijn et al., 2016; Pérez-Ortega et al., 2021). Like for any other unsupervised clustering method, the number of optimal clusters cannot be known beforehand, thus potential number of clusters were scanned between five and twenty. In Ecker et al. (2023b), the one with the lowest Davis-Bouldin index was chosen, which maximizes the similarity within elements of the cluster while maximizing the the between cluster similarity (Davies and Bouldin, 1979). For assemblies detected over the 10 minutes-long plastic simulation, this optimal value was overwritten, to have at least one pattern-specific assembly for all 10 patterns. For the assemblies detected over the 2 minutes-long non-plastic simulation, the optimal value was chosen, to avoid biasing the before vs. after assembly comparisons. Fifth, neurons were associated to these clusters based on their spiking activity, and it was determined whether they formed a cell assembly or not. The correlations between the spike trains of all neurons and the activation sequences of all clusters were computed and the ones with significant correlation selected to be part of the potential assemblies. Significance was determined based on exceeding the 95th percentile of correlations of shuffled controls (1000 controls with spikes of individual cells shifted by any amount as above; Montijn et al., 2016; Herzog et al., 2021). Finally, it was required that the mean pairwise correlation of the spikes of the neurons with significant correlations was higher than the mean pairwise correlation of neurons in the whole dataset (Herzog et al., 2021). Clusters passing this last criterion were considered to be functional assemblies and the neurons with significant correlations their constituent cells. Assemblies of neurons were compared using their Jaccard distances. The `assemblyfire` package, developed for Ecker et al. (2023b) is publicly available on GitHub.

4.4.6 Determination of consensus assemblies

Consensus assemblies, resulting from the hierarchical clustering of the Jaccard similarity matrix of assemblies across repetitions of the same input stream, are fully described in (Ecker et al., 2023b), but a minimal description of them can be found below. It was ensured that assemblies from the same repetition did not cluster together, first by setting their distances to twice the maximum, and second, by cutting the tree in a way that resulted in the lowest number of cluster in which two assemblies from the same repetition did not cluster together. Membership of neurons in these consensus assemblies was based on the fraction of assembly instances they were part of, normalized by a binomial control and thresholded. As shown in Ecker et al. (2023b), consensus assemblies are similar to assemblies detected over the average spike trains across repetitions, but with the added benefit of the membership threshold. In rough terms, this threshold can be understood as follows: if a neuron was part of 80% of assembly instances that made up the consensus, then it was also a member of the consensus assembly.

In order to assess the functional connectivity of consensus assemblies before and after plasticity, the spike trains of their neurons across the five repetitions were first averaged and then

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes

binned (using the same 20 ms bins as above). Last, the Pearson correlation of all pairs of the preprocessed spike trains were calculated, and averaged across the population.

4.4.7 Calculation of spike time reliability

Spike time reliability, quantify the reliability of a single neuron across multiple presentations of the same input, is described in (Ecker et al., 2023b), but the same description of it can be found below. Spike time reliability was defined as the mean of the cosine similarities of a given neuron's mean centered, smoothed spike times across all pairs of repetitions (Schreiber et al., 2003; Cutts and Eglon, 2014). To smooth the spike times, they were first binned to 1 ms time bins, and then convolved with a Gaussian kernel with a standard deviation of 10 ms.

4.4.8 Synaptic clustering coefficient and likelihood of plastic changes in synapse clusters

Synaptic clustering coefficient (*SCC*), quantify the co-localization of synapses on the dendrites of a neuron from its presynaptic assembly with a single number, is fully described in (Ecker et al., 2023b), but a minimal description of it can be found below. First, the nearest neighbor distance (along the dendrites) between all pairs of synapses from the presynaptic assembly were computed and averaged (*mean nnd*). Second, 20 controls were generated by always selecting the same number of random presynaptic E cells from the circuit and *mean nnds* of the control populations were calculated. Last, *SCC* was defined as the negative z-score of assembly *mean nnd* with respect to the distribution of control *mean nnds*. *SCC* is thus a parameter-free metric, centered at zero, and is positive for intersynaptic distances that are lower than expected (indicating clustering) and negative otherwise (indicating dispersion). Additionally, the significance of the clustering or dispersion of the synapse locations was determined with a two-tailed t-test of assembly *mean nnd* against the 20 random samples with an alpha level of 0.05. *SCC* was implemented using NeuroM and ConnectomeUtilities.

Synapse clusters were also detected based on synapse neighbour distances. In order to be part of a spatial cluster, a synapse was required to have at least nine other synapses on the same dendritic branch, i.e., between two branching points of the dendrite, with $\leq 10 \mu\text{m}$ (Euclidean) distance. Significance of spatial clustering was determined similar to Druckmann et al. (2014). The distribution of synapse neighbour distances of the 10 selected synapses were compared with a Poisson model (assuming exponentially distributed inter-synapse distances) based on all (same branch) synapse neighbour distances on the given neuron. Clusters were merged in a subsequent step, thus synapse clusters with more than 10 synapses, spanning more than 20 μms were also feasible. As plastic changes in synapse clusters were only analyzed for a small subpopulation of assemblies (10 L5 PCs per assembly), *SCC* was used to select subpopulations with high probability of finding synapse clusters. To this end, assembly neurons with positive, significant *SSC* values with respect to an assembly (either the same assembly for within-

assembly analysis, or other ones for analysing cross-assembly interactions) were selected, and the ones with the 10 highest assembly indegree (with respect to the same assembly) selected (see Ecker et al., 2023b for the same selection method). Control synapse clusters, originating from non-assembly neurons were also detected on the same postsynaptic neurons.

The normalized likelihood of changes, conditioned on the four *categories* a synapse could fall into (assembly clustered, assembly non-clustered, non-assembly cluster, non-assembly non-clustered) were quantified using the Michaelson contrast, defined as:

$$\frac{P(\text{changed} | \text{category}) - P(\text{changed})}{P(\text{changed} | \text{category}) + P(\text{changed})} \quad (4.9)$$

where *changed* was split to be either potentiated or depressed.

Acknowledgements

The authors thank Nicolas Ninin for his involvement in the early stage of this project, Elvis Boci and Cyrille Favreau for their help with visualizations, Michael Gevaert, Joni Herttuainen and Thomas Delemontex for their assistance with software engineering, and Alberto Antonietti, Christoph Pokorny, Kathryn B. Hess, Ran Levi, Wulfram Gerstner and Henry Markram for discussions.

Funding

This study was supported by funding to the Blue Brain Project, a research center of the École polytechnique fédérale de Lausanne (EPFL), from the Swiss government's ETH Board of the Swiss Federal Institutes of Technology.


Author contributions

- Conceptualization: A.E., D.E.S., G.C., E.B.M., M.W.R.
- Formal analysis: A.E., D.E.S., M.W.R.
- Investigation: A.E., S.B.-P., G.C.
- Methodology: A.E., D.E.S., G.C., J.B.I., E.B.M., M.W.R.
- Project administration: M.W.R.
- Software: A.E., D.E.S., S.B.-P., G.C., J.K., P.K., M.W.R.
- Resources: J.B.A., I.M.
- Supervision: M.W.R.
- Validation: A.E., D.E.S., S.B.-P., J.B.I., M.W.R.
- Visualization: A.E., M.A.
- Writing - original draft: A.E., D.E.S., M.W.R.
- Writing - review & editing: A.E., D.E.S., S.B.-P., G.C., J.B.I., E.B.M., M.W.R.

Data and code availability


The 2.4 mm^3 subvolume of the juvenile rat somatosensory cortex, containing 211,712 neurons and 312,709,576 plastic synapses in SONATA format (Dai et al., 2020) is freely available at: <https://doi.org/10.5281/zenodo.8158471>. It can be loaded and instantiated in **CoreNEURON** (Kumbhar et al., 2019) with **neurodamus**. The circuit and the simulations can be analyzed using **Blue Brain SNAP** and **ConnectomeUtilities** built on top of it. Cell assemblies were detected and can be analyzed with **assemblyfire**. Exemplary Jupyter notebooks using the packages above were deposited in the same repository on Zenodo.

4.5 Supplementary Figures




Supplementary Figure 4.7: Physiology of excitatory cells and E to E connections. A: Distribution of ion-channel densities in the excitatory (cADpyr) electrical type (etype). B: Validation of dendritic physiology of the cADpyr e-type on L5 TTPC mtypes. B1: Validation of back-propagating action potential (bAP) amplitude for basal (teal) and apical (blue) dendrites. Reference data (in orange) comes from Stuart and Sakmann (1994); Larkum et al. (2001) (apical) and Nevian et al. (2007) (basal). Lines show exponential fits for the *in silico* (teal and blue) and *in vitro* (orange) data. Color bar indicates dendritic diameter. B2: Validation of EPSP attenuation. Reference data comes from Berger et al. (2001) (apical) and Nevian et al. (2007) (basal). Lines and color bar same as in B2. Data taken from (and partially shown in) Reva et al. (2022). (A similar panel has also been shown Isbister et al., 2023). C: Anatomy and physiology of E to E connections. C2: Connection probability and number of synapses per connections for all E to E connections. White boxes indicate non-feasible connections, or on the left panel: no pairs found within the 200 μm intersomatic distance used. C2: Mean (over 100 pairs) PSP amplitude (left) and CV (std/mean on the right) of all E to E connections. (Data taken from (and shown in) Isbister et al., 2023). C3: Initial synaptic physiology parameters. From left to right: ρ , \hat{g}_{AMPA} , and U_{SE} .

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes




Supplementary Figure 4.8: Synapse-specific parameters of the plasticity model. A: Layer- and neurite type-wise distribution of measured $[Ca^{2+}]_i$ peaks (used to derive parameters of the plasticity model shown in B). Synapses are grouped based on the soma location of the postsynaptic cell. 10% of all synapses are shown. Schematics on their lefts illustrate the measurement protocols. **B:** Layer- and neurite type-wise distribution of depression and potentiation thresholds (θ_p and θ_d) of the plasticity model. Synapses grouped and shown as in A. **C:** Correlations of the parameters shown in A and B.

4.5 Supplementary Figures




Supplementary Figure 4.9: Calibration of the *in vivo*-like network state. **A:** Same as 4.2B (i.e., raster plots of the microcircuit's activity) under different synapse setups. The microcircuit equipped with the plasticity model of Chindemi et al. (2022) only resembles that of the non-plastic network's of Isbister et al. (2023) when VDCCs (voltage-dependent calcium channels) are blocked (last row). **B:** Re-calibration of the *in vivo*-like state using the plasticity model. **B1:** Left: Euclidean distance of the measured percentages of firing rates (P_{FR} s) from the target ones in different iterations of the calibration process. Right: Validation of network states after the final (4th) iteration. Dashed gray line along the diagonal indicated perfect match. **B2** Left: Injected Ornstein-Uhlenbeck (OU) conductances in the non-plastic model of Isbister et al. (2023) vs. the plastic one for $P_{FR} = 40\%$ (the state used in the current article). Dashed gray line along the diagonal indicated perfect match. Right: Layer-wise (absolute) firing rates of excitatory (E) and inhibitory (I) subpopulations at $P_{FR} = 40\%$. Legend on the bottom applies to the last three panels in B.

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes




Supplementary Figure 4.10: Activity of the thalamic fibers. Raster plots of VPM fibers forming each of the ten input patterns (Figure 4.2A) for the stimulus stream in (i.e., from pattern A at 2000 ms to pattern J at 6500 ms). Bottom row shows the same for non-specific POM fibers. (A similar panel has been shown in Ecker et al., 2023b.)

4.5 Supplementary Figures




Supplementary Figure 4.11: Changing connections in plastic control simulations. A: Same as Figure 4.2B and E2 (i.e., raster plots of the microcircuit's activity and plastic changes in mean(ρ) vs. firing rates under different conditions. The last row of A2 is not an exact replica of Figure 4.2E2 as these simulations were run for 2 minutes. **B:** Similar to Figure 4.2E3 (i.e., layer-wise propensity of changes in mean(ρ)) but split across conditions. **C:** Similar to Figure 4.2D1 (i.e., L2 norm of changes) but for mean ρ values of connections (not synapses) for all conditions.

Chapter 4. Long-term plasticity induced sparse and specific synaptic changes




Supplementary Figure 4.12: Changes in synaptic efficacy during plasticity. A: Individual ρ traces (10 per panel) during plasticity. **B:** L2 norm of changes in ρ (similar to Figure 4.2D1, but) compared to $\rho_{t=5minutes}$. Insert shows the distribution of individual changes in each time steps (for 10 million samples).



Supplementary Figure 4.13: Connectivity features underlying cell assembly membership. **A:** Selected connectivity features of assemblies. **A1:** Left: Probability of membership in an exemplary assembly (A13) against assembly-indegree with respect to all assemblies. Solid lines indicate the mean and the shaded areas indicate 95% confidence interval. Right: nI (normalized mutual information, see Ecker et al., 2023b) of assembly-indegree and assembly membership (blue arrow indicates postsynaptic assembly A13, shown in detail on its left). **A2:** Probability of membership in the same exemplary assembly against synapse clustering coefficient (SCC , see Methods and Ecker et al., 2020) with respect to all assemblies; nI of SCC and assembly membership. White boxes indicate non-significant nI . **B:** Summary of within-assembly interactions (diagonals of nI matrices) for all connectivity features considered in Ecker et al. (2023b). (Similar panels have been shown in Ecker et al., 2023b.)


Chapter 4. Long-term plasticity induced sparse and specific synaptic changes



Supplementary Figure 4.14: Comparison of cell assemblies before and after plasticity. A1: Firing rates before and after plasticity. In all time bins on the left, and only in significant ones (see Methods) on the right. (Caption continues on the next page.)

4.5 Supplementary Figures

A2 Interspike interval (ISI) distribution (of all excitatory spikes) before and after plasticity on the left. On the right: Zoom in to low ISIs (≤ 5) ms split by layer. **B**: STP profiles before and after plasticity. At *in vitro* $[Ca^{2+}]_o$ on the left, and *in vivo* on the right. Thin lines represent the 20 individual repetitions, while the thicker ones their means. **C**: Davis-Bouldin index (see Ecker et al., 2023b and Methods) of different number of assemblies before and after plasticity across repetitions. (The index is to be minimized to achieve optimal number of clusters.) **D1**: Number of times a consensus assembly is active over repetitions before and after plasticity. E.g. the first rows per patterns are the counts of colored boxes from Figure 4.5A2 and B2. This representation can be used to judge the grouping of assemblies (see D2), and also for calculating their normalized Euclidean distance (see Figure 4.5G2). **D2**: Cosine similarity of rows of consensus assembly matrices (split by patterns before vs. after plasticity). **E**: Number and location of consensus assembly neurons before and after plasticity.



Supplementary Figure 4.15: Layer-wise propensity of changes per single pattern. As on Figure 4.2E3, layer corresponds to the soma location of the postsynaptic cells.

Conclusion

At the beginning of the thesis, we set out to use biophysically detailed, large-scale models of the cortex to study synaptic plasticity at the network level under physiological $[Ca^{2+}]_o$. Our group has already shown how the pathway-specific $[Ca^{2+}]_o$ dependence of U_{SE} , i.e., STP, contributes to the transition from *in vitro*-like synchronous activity, to *in vivo*-like asynchronous firing (Markram et al., 2015). In Romani et al. (2023), we extended the synapse parametrization process of Ecker et al. (2020) to the CA3 input of CA1 and described how parameters of those synapses (including the U_{SE} , D , and F parameters of STP) lead to feedforward inhibition, a hallmark of that pathway (Pouille and Scanziani, 2001; Pouille et al., 2009). These network-level effects of STP differ from previous theoretical insights (Abbott et al., 1997; Goldman et al., 2002; Maass and Markram, 2002; Sussillo et al., 2007; Naud and Sprekeler, 2018; Keijser and Sprekeler, 2022) because those were all based on short-term depression, which, due to the drastically decreased U_{SE} , is not so prevalent at *in vivo* low $[Ca^{2+}]_o$. Markram and Tsodyks (1996) showed how the interaction of short- and long-term plasticity leads to a redistribution of synaptic efficacy towards the early spikes during high-frequency firing. In Ecker et al. (2023a), we did not observe high frequency firing under *in vivo*-like conditions and could not study this effect in detail. On the other hand, long-term plasticity caused a network-level redistribution of efficacy from the superficial to the deep layers of the neocortex, which resulted in prolonged stimulus-specific responses. Compared to previous models (Zenke et al., 2015; Fauth and Van Rossum, 2019; Kossio et al., 2021), efficacy changes were sparse and our network remained stable without any homeostatic terms (Turrigiano and Nelson, 2004; Zenke et al., 2017a). This emerged from the combination of our calcium-based model (Chindemi et al., 2022) and the physiological levels of $[Ca^{2+}]_o$ used, as under these conditions, not every spikes contributed to the synaptic efficacy updates in our model, unlike in STDP rules, which take all spikes into account (Gerstner et al., 1996; Kempter et al., 1999; Song et al., 2000). The sparse changes were likely to be found at synapses forming spatial clusters of postsynaptic dendrites and between co-firing Hebbian cell assemblies (Ecker et al., 2023a). In Ecker et al. (2023b), we described these cell assemblies in detail and leveraged our *in silico* setup to study the structural connectivity underlying them.

Conclusion

Future directions

The remainder is structured into ten points that can be seen as current limitations that could be improved on or as feasible future avenues that one could pursue based on the results presented here and pipelines behind them.

1. As emphasized in Chapter 2, data-integration should be a continuous process and CA1 synaptic physiology presented in Chapter 1 is no exception. While in Ecker et al. (2020), we concentrated to the local, i.e., within CA1 synapses, in Romani et al. (2023), we parametrized the Schaffer collateral (projecting from CA3 to CA1) synapses as a step forward. However, there are several parameters in the current version of the CA1 model that are cortical "*placeholders*", like the U_{SE} , D , and F parameters of several pathways. Although hippocampal and cortical synapses have similar properties, there are several known differences. The difference between the NMDA nonlinearities is already taken into account (Jahr and Stevens, 1990; Vargas-Caballero and Robinson, 2003; Chindemi et al., 2022), but e.g., synaptic democracy¹, a feature of hippocampal, but not cortical connections, is not (Magee and Cook, 2000; Williams and Stuart, 2002). Thus, while new CA1 synaptic physiology data is sparse, every piece of it should be included to replace neocortical placeholder values in the future.
2. A clear next step would be introducing long-term plasticity at the Schaffer collaterals. This would require extra groundwork, as our model fitting pipeline is based on pairs of neurons in line with the cortical dual-patch experiments (Chindemi et al., 2022). On the other hand, as CA3 and CA1 are quite separate in space and the connection probability between them is low, long-term plasticity is studied using pathway stimulation (Schwartzkroin and Wester, 1975; Bliss and Collingridge, 1993; but see Debanne et al., 1995 for some heroic dual-patch effort in cultures). Thus, the presynaptic side is a population of (unknown number of) neurons, not a single one. Once the pipeline is operational, the model of Chindemi et al. (2022) would serve as an excellent tool for contributing to the postsynaptic only expression (Selig et al., 1999; Malinow and Malenka, 2002; Costa et al., 2015) vs. pre- and postsynaptic expression of plasticity debate (Yasui et al., 2005; Enoki et al., 2009; Bliss and Collingridge, 2013). By making this pathway plastic, we could also learn whether synaptic democracy results from plastic changes as suggested by modeling studies (Rumsey and Abbott, 2006) or if the answer lies elsewhere.

¹"*Synaptic democracy*" means that the activation of any synapse, independent of its location along the dendritic tree, results in the same somatic EPSP. At the Schaffer collateral synapses this is achieved by a distance dependent scaling of AMPA (but not NMDA) receptors (Andrasfalvy and Magee, 2001).

3. Once the plasticity of the Schaffer collaterals is well calibrated in our CA1 model, one could use it to study the mechanism behind behavioural time scale plasticity (BTSP; Bittner et al., 2015, 2017). BTSP is lauded as a new non-Hebbian learning rule, as it can lead to stable place cell² formation in a single shot (compared to the ~ 50 repetitions of spike pairings required in STDP protocols) and the presynaptic spikes can proceed (or follow) the prolonged postsynaptic depolarization, termed "*plateau potentials*", by seconds (compared to the millisecond precision required in STDP protocols). While patching cells deep in the brains of behaving animals that lead to this discovery is outstanding and so are the possible experiments one can do with this paradigm, calling it a profoundly new learning rule might be an overstatement. BTSP is known to require NMDA receptors (Bittner et al., 2017) and CamKII activation (Jain et al., 2023; Xiao et al., 2023), the same ingredients as our biophysically detailed plasticity model (Chindemi et al., 2022; Ecker et al., 2023a). In connected L5 TTPCs, a single *in vitro* spike pairing activates 5 – 7 synapses distributed across the basal dendrites (Markram et al., 1997a) into which the postsynaptic action potential does not propagate as well as into the apical ones (Williams and Stuart, 1999; Letzkus et al., 2006; Ecker et al., 2023a). On the other hand, CA1 PCs have $\sim 30k$ excitatory synapses, at least three times more than L5 TTPCs (Megías et al., 2001) and most of these synapses come from the CA3 and form clusters on CA1 dendrites (Druckmann et al., 2014). Based on the biophysics of Chindemi et al. (2022), pairing a huge, long-lasting dendritic depolarization with the activation of a large fraction of these clustered CA3 synapses will indeed lead to remarkable plastic changes, probably enough for single shot learning. The amount of Ca^{2+} entering during these events would be enhanced by the higher NMDA/AMPA ratio (Myme et al., 2003) and the steeper NMDA nonlinearity of the CA3-CA1 synapses (compared to the cortical L5-L5 ones). Thus, while BTSP is a faster **protocol** to induce plasticity than STDP is, the biophysical **rule** that governs the plastic changes seems to be the same.
4. A weak point of any assembly detection method, not just the one we presented in Chapter 3, is that the number of assemblies are not known beforehand. Since the rigorous work of Peyrache et al. (2010), the hippocampus community uses significant PCA scores³ (Lopes-dos Santos et al., 2013). However, when recording thousands of neurons, the activity of the sensory cortex is "*not so spontaneous*", and the explained variance increases only slowly with the number of PCA components (Stringer et al., 2019; Avitan and Stringer, 2022). Probably this was amongst the reasons why Carrillo-Reid et al. (2015) pioneered a new way of detecting cell assemblies from cortical data, but their method still involved a step of unsupervised clustering. (In the beginning of the cell assembly project we also used PCA for detecting assemblies, but only got unsatisfactory

²Some CA1 PCs have a spatial receptive field, i.e., they fire preferentially when the animal is at a given position in space (O'Keefe and Dostrovsky, 1971). CA1 PCs that exhibit this type of behavior are termed "*place cells*".

³Principal component analysis (PCA) is used to find the dimensions along which the data varies the most. Projecting the data to the first n components, that explain e.g. 95% of the variance is a commonly used dimensionality reduction technique. Contrary to this, Peyrache et al. (2010) suggested to use the Marčenko–Pastur distribution as a null hypothesis for testing the significance of each dimension, making the resulting number of PCA components, and thus the resulting number of cell assemblies less arbitrary.

Conclusion

results.) One of the strengths of our work is that contrary to other modeling studies that define assemblies based on their strong internal connectivity after plasticity (Litwin-Kumar and Doiron, 2014; Zenke et al., 2015; Fauth and Van Rossum, 2019; Kossio et al., 2021)), we used methods developed by experimentalists, thus we detect and talk about assemblies that are relevant to the experimental neuroscientist community. On the other hand, one could argue that if we were to take into account the structural connectivity as well, we could detect groups of neurons that are described by both their co-firing function, and strongly interconnected structure. We could do this the following way in the future: Instead of determining the number of clusters based on the distance between the leaves of the clustering tree beforehand, we could progressively cut the tree if the next cut would result in higher within-assembly indegree nI values, i.e., even better predicted assembly membership of a neuron based on the number of connections from assembly neurons.

5. At the beginning of Ecker et al. (2023b), we argue that we are in a good position to study the connectivity underlying cell assemblies as we have access to both functional and structural data, while to experimentalists, even if available, these modalities come from separate data sets. This picture seems to be changing with the release of the MICrONS data set (MICrONS Consortium, 2021). The data set contains functional (two-photon calcium imaging) and structural (electron microscopy, EM) data of mouse visual cortices. At the end of the preprint, the authors flash out possible experiments that take into account all connections of a functionally identified complex cell and could therefore test the postulate about their receptive field appearing complex because of inputs from several different simple cells (Hubel and Wiesel, 1962; Movshon et al., 1978). That would indeed be wonderful, but we think even more could be done, and a data set like this begs for population-level analysis. Thanks to our access to a large-scale model, we have already prototyped a toolchain that can detect cell assemblies and then describe how they emerge from the connectivity structure. The co-registration of functional and structural images requires manual work and is still ongoing (MICrONS Consortium, 2021). Once done, it could be used to test our predictions and provide further insights about how cell assemblies encode information in the sensory cortex.
6. One of the strong points of our bottom-up model of network plasticity presented in Chapter 4 is that we could put the network in an *in vivo*-like state, which we defined as low-rate, asynchronous firing and low release probability. The major driver of this is undoubtedly the low $[Ca^{2+}]_o$ (Borst, 2010), but there are several other differences between *in vitro* and *in vivo* conditions. Of special importance are: $[Mg^{2+}]_o$, high conductance state of the dendrites, and neuromodulation. The Mg^{2+} block of NMDA receptors can be removed by local depolarization, but the amount of depolarization needed depends on the $[Mg^{2+}]_o$, and hippocampal measurements suggest lower than 1 mM concentrations *in vivo* (Gonzalez et al., 2022). Once reliable cortical measurements become available, our simulation setup should be updated to feature *in vivo* $[Mg^{2+}]_o$, instead of the placeholder value taken from *in vitro* slice experiments. By in-

jecting noisy conductances to compensate for missing inputs from non-modeled brain regions we put the cell bodies into an *in vivo*-like high-conductance state (Destexhe et al., 2001; Isbister et al., 2023). On the other hand, as excitatory synapses are on the dendrites, they are mostly influenced by the activity of the local circuit. Modeling all those missing synapses would require tremendous resources (Cremonesi et al., 2020), but would put the dendrites into a high-conductance state as well and engage more dendritic nonlinearities, which might result in more plastic changes (Farinella et al., 2014; Stuart and Spruston, 2015; Kastellakis and Poirazi, 2019; Gonzalez et al., 2022). Neuromodulators, such as acetylcholine (ACh), are almost absent in slices, but are known to alter the shape of STDP kernels and plasticity in general (Brzosko et al., 2019). Most of these effects are mediated by metabotropic and G-protein coupled receptors, which would require substantial effort to model, but ACh in particular is known to reduce U_{SE} (Colangelo et al., 2019). This effect could easily be taken into account in our pipeline, and in an earlier, non-plastic version of the model it already was (Ramaswamy et al., 2018). Taken together, the state presented in Ecker et al. (2023a) resembles *in vivo* conditions more than any publication of which we know about, but could still be improved in the future.

7. According to many fellow modelers, an article describing a plastic network that is repeatedly exposed to different patterns is not complete without showing pattern completion from noisy or partial patterns (Hopfield, 1982). On the other hand, the patterns we present are noisy and "*partial*" by design. We only fixed the pattern fibers, but the spike trains delivered on them are stochastic. This stochasticity is further enhanced by the probabilistic release of vesicles. Moreover, as we designed overlapping patterns (in terms of the Hamming distance of the pattern fibers), presenting e.g. pattern E, which is half A and half B, can be seen as testing what would happen if we only presented half of pattern A. Furthermore, pattern completion requires a readout, i.e., a population of neurons whose firing rate would signal that the network "*converged to*" a pattern. Theoretical work implies that if a network is sufficiently large and has diverse units, then a simple linear readout, trained on all neurons, can become quite powerful even if the connections within the network are random (Maass and Markram, 2004). Therefore, when training the Allen Brain Institute's point neuron model (Billeh et al., 2020) for visual processing tasks, Chen et al. (2022) defined a randomly selected population of 30 L5 PCs as readout. We could do something similar, but as we model biorealistic plasticity and do not train the network with backpropagation through time (Bellec et al., 2020), I think this would not work. Another possibility is to use the pipeline of (Reimann et al., 2022b), which samples neurons based on their topological features and trains decoders based on the activity of these carefully selected cells. We could restrict it further to sample neurons from the late assemblies only, as we see those populations as the output of the cortex. Taken together, although it might raise more questions than provides answers, we could test pattern separation and pattern completion in the future.

Conclusion

8. At end of (Ecker et al., 2023a) we highlight the predictability of changes from the innervation of the pre- and postsynaptic neurons (which influence their firing rates), but also emphasize that these predictions are not perfect and other factors need to be investigated in the future. We can foresee three ways of doing this. The topologically inspired one would take into account the embedding of the changing connections into the whole network and e.g. look for motifs that these connections are part of instead of the simple pre- and postsynaptic indegree. The classical approach would be to assume an unknown function $\mathcal{F}(\rho_{ij}; v_i, v_j)$ that governs the changes in synaptic efficacy based on the initial efficacy and the firing rates of the pre- and postsynaptic neurons (v_i and v_j). If \mathcal{F} is well behaved, one can do a Taylor expansion (about $v_i = v_j = 0$) and fit its variables to the observed data (Gerstner et al., 2014). This exercise may even link the learning rule observed in our biophysically detailed model to the classical rate-based ones (Sejnowski, 1977; Bienenstock et al., 1992; Oja, 1982). The third, more modern way of doing it would be to identify a large set of features, like pre- and postsynaptic firing rate, pattern-indegree, assembly-indegree, participation in high-dimensional motifs etc. and feed it to a machine learning model that predicts either the direction or the precise amplitude of the plastic changes. Ideally, the model would order the features by importance, like XGBoost does (Chen and Guestrin, 2016).
9. At the end of (Ecker et al., 2023a) we propose future experiments with our model using more precise top-down inputs. Top-down inputs are seen as error/target signals and therefore boost specific plastic changes (Makino, 2019), but the precise pathways mediating this process are still under investigation. POM input was shown to gate plasticity in L2/3 PCs *in vivo* via activating vasoactive intestinal peptide (VIP; a subpopulation of our 5HT3aR+ population; Tremblay et al., 2016) interneurons, which in turn inhibit Sst+ interneurons and therefore dis-inhibit the distal dendrites of PCs (Gambino et al., 2014; Williams and Holtmaat, 2019). Sst+ interneurons have low firing rates in our simulations (thus, there is not much activity to be inhibited), probably because the calibration algorithm does not distinguish between inhibitory subpopulations and most of the spikes required to achieve the target layer-wise firing rates are coming from PV+ interneurons (Isbister et al., 2023). Our setup would be readily applicable to calibrate subpopulation specific firing rates, if those rates were available⁴. Although we know about available patch-clamp data from genetically labeled inhibitory neurons (Gentet et al., 2010, 2012; Yu et al., 2019), we decided not to use them because of the low number of cells recorded in mice⁵. If not as target firing rates, we could still have used those values to fix the ratios between the rates of the inhibitory subpopulations. Once the Sst+ rates are increased, they could be inhibited by the VIP+ interneurons, but that may require more targeted inhibitory connectivity (Reimann et al., 2022a). It has been known

⁴The target firing rates used in Isbister et al. (2023) are from spikes sorted from extracellular recordings (Reyes-Puerta et al., 2015) in which spikes of inhibitory subpopulations cannot be distinguished.

⁵The field in general seem to have shifted from rat to mouse, thanks to the recent advances in mouse genetics (Gurumurthy and Kent Lloyd, 2019; Brown, 2021; Azkona and Sanchez-Pernaute, 2022), which among other things made those inhibitory subpopulation-specific recordings possible.

that VIP+ cells inhibit qualitatively more interneurons than they do PCs (Pi et al., 2013; Karnani et al., 2016), but precise numbers were lacking. A recent preprint analysing EM data from the MICrONS data set provided those numbers (Schneider-Mizell et al., 2023) and when compared with our touch-based connectome, we found that we have the potential to make those VIP+ to interneuron connections, i.e., appositions between these cells exists (Reimann et al., 2015), but we pruned them in favour of VIP+ to PC synapses (Reimann et al., 2022a). In future releases of our connectome those precise targeting rules will be included which will increase the chances of reproducing the dis-inhibitory effect of VIP+ cells.

10. While this thesis focuses on short- and long-term plasticity, there are several other forms of synaptic plasticity. In Ecker et al. (2023a), we emphasize that unlike other modeling studies, we did not have to include homeostatic plasticity to keep our network stable. By doing so, we are not arguing against the existence of homeostatic plasticity, but highlight that it operates on slower time scales of hours - days, and we simulated "only" 10 minutes (Turrigiano and Nelson, 2004). Inhibitory plasticity is a faster form of synaptic plasticity that is also frequently used to keep the excitatory/inhibitory balance of network models (Vogels et al., 2011; Hennequin et al., 2017; Zenke et al., 2017a). Again, no questions about its existence in biology (D'amour and Froemke, 2015; Vickers et al., 2018; Field et al., 2020), but without a complete understanding of its biophysics, we cannot include this form of plasticity in our bottom-up model. Lastly, structural plasticity, in which new synapses are formed (instead of changing the efficacy of existing ones), was shown experimentally to be able to rewire complete networks in a non-NMDA receptor dependent manner (Le Bé and Markram, 2006), but just like homeostatic plasticity, on a slower time scale than what we can currently simulate. We are in a good position to model this type of rewiring in the future, because our connectome is "built" by pruning appositions, i.e., close contacts between presynaptic axons and postsynaptic dendrites (Reimann et al., 2015). The appositions that we do not keep are potential new synapses and after deleting existing weak ones (Le Bé and Markram, 2006) they could be added as new synapses, this way keeping the biorealistic bouton densities and forming new synapses at realistic locations (Reimann et al., 2015).

The motto of the thesis "*Beauty is truth, truth (is) beauty, - that is all*" is a line from Keats, often quoted by Hungarian neuroanatomist János Szentágothai. Szentágothai's ideology revolved around the concept that intricate nuances within nature bear profound beauty, a beauty that is uncovered through scientific discoveries (Gulyás and Somogyi, 2013). I think we also uncovered small, but important details and hope that our work will inspire others to continue to do so.

Bibliography

- Abbott, L. F., Varela, J. A., Sen, K., and Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science*, 275(5297):220–224.
- Abdellah, M., Hernando, J., Eilemann, S., Lapere, S., Antille, N., Markram, H., and Schürmann, F. (2018). NeuroMorphoVis: A collaborative framework for analysis and visualization of neuronal morphology skeletons reconstructed from microscopy stacks. *Bioinformatics*, 34(13):i574–i582.
- Al-Juboori, S. I., Dondzillo, A., Stubblefield, E. A., Felsen, G., Lei, T. C., and Klug, A. (2013). Light scattering properties vary across different regions of the adult mouse brain. *PLoS ONE*, 8(7):e67626.
- Ali, A. B. (2007). Presynaptic Inhibition of GABAA Receptor-Mediated Unitary IPSPs by Cannabinoid Receptors at Synapses Between CCK-Positive Interneurons in Rat Hippocampus. *Journal of Neurophysiology*, 98(2):861–869.
- Ali, A. B. (2011). CB1 modulation of temporally distinct synaptic facilitation among local circuit interneurons mediated by N-type calcium channels in CA1. *Journal of Physiology*, 105:1051–1062.
- Ali, A. B., Bannister, A. P., and Thomson, A. M. (1999). IPSPs elicited in CA1 pyramidal cells by putative basket cells in slices of adult rat hippocampus. *European Journal of Neuroscience*, 11(5):1741–1753.
- Ali, A. B., Deuchars, J., Pawelzik, H., and Thomson, A. M. (1998). CA1 pyramidal to basket and bistratified cell EPSPs: Dual intracellular recordings in rat hippocampal slices. *Journal of Physiology*, 507(1):201–217.
- Ali, A. B. and Thomson, A. M. (1998). Facilitating pyramid to horizontal oriens-alveus interneurone inputs: Dual intracellular recordings in slices of rat hippocampus. *Journal of Physiology*, 507(1):185–199.
- Amit, D. J. and Tsodyks, M. V. (1991). Quantitative study of attractor neural network retrieving at low spike rates. I. Substrate-spikes, rates and neuronal gain. *Network: Computation in neural systems*, 2(3):259.
- Andrasfalvy, B. K. and Magee, J. C. (2001). Distance-dependent increase in AMPA receptor number in the dendrites of adult hippocampal CA1 pyramidal neurons. *The Journal of Neuroscience*, 21(23):9151–9159.
- Ascoli, G. A. et al. (2008). Petilla terminology: Nomenclature of features of GABAergic interneurons of the cerebral cortex. *Nature Reviews Neuroscience*, 9(7):557–568.
- Ascoli, G. A., Gasparini, S., Medinilla, V., and Migliore, M. (2010). Local Control of Postinhibitory Rebound Spiking in CA1 Pyramidal Neuron Dendrites. *The Journal of neuroscience*, 30(18):6434–6442.
- Avitan, L. and Stringer, C. (2022). Not so spontaneous: Multi-dimensional representations of behaviors and context in sensory areas. *Neuron*, 110(19):3064–3075.
- Awile, O., Kumbhar, P., Cornu, N., Dura-Bernal, S., King, J. G., Lupton, O., Magkanaris, I., McDougal, R. A., Newton, A. J., Pereira, F., Săvulescu, A., Carnevale, N. T., Lytton, W. W., Hines, M. L., and

Bibliography

- Schürmann, F. (2022). Modernizing the NEURON Simulator for Sustainability, Portability, and Performance. *Frontiers in Neuroinformatics*, 16(884046).
- Azimipour, M., Baumgartner, R., Liu, Y., Jacques, S. L., Eliceiri, K. W., and Pashaie, R. (2014). Extraction of optical properties and prediction of light distribution in rat brain tissue. *Journal of biomedical optics*, 19(7).
- Azkona, G. and Sanchez-Pernaute, R. (2022). Mice in translational neuroscience: What R we doing? *Progress in Neurobiology*, 217(102330).
- Barros-Zulaica, N., Rahmon, J., Chindemi, G., Perin, R., Markram, H., Muller, E., and Ramaswamy, S. (2019). Estimating the Readily-Releasable Vesicle Pool Size at Synaptic Connections in a Neocortical Microcircuit. *Frontiers in Synaptic Neuroscience*, 11(29).
- Bartos, M., Vida, I., Frotscher, M., Meyer, A., Monyer, H., Geiger, J. R. P., and Jonas, P. (2002). Fast synaptic inhibition promotes synchronized gamma oscillations in hippocampal interneuron networks. *PNAS*, 99(20):13222–13227.
- Bassett, D. S. and Bullmore, E. T. (2017). Small-World Brain Networks Revisited. *Neuroscientist*, 23(5):499–516.
- Bathellier, B., Ushakova, L., and Rumpel, S. (2012). Discrete Neocortical Dynamics Predict Behavioral Categorization of Sounds. *Neuron*, 76(2):435–449.
- Beierlein, M. and Connors, B. W. (2002). Short-term dynamics of thalamocortical and intracortical synapses onto layer 6 neurons in neocortex. *Journal of Neurophysiology*, 88(4):1924–1932.
- Beierlein, M., Gibson, J. R., and Connors, B. W. (2003). Two Dynamically Distinct Inhibitory Networks in Layer 4 of the Neocortex. *Journal of Neurophysiology*, 90(5):2987–3000.
- Bellec, G., Scherr, E., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., and Maass, W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications*, 11(3625).
- Berger, T. (2009). *Properties of Neocortical Microcircuits*. PhD thesis, Ecole Polytechnique Fédérale de Lausanne.
- Berger, T., Larkum, M. E., and Lüscher, H. R. (2001). High Ih channel density in the distal apical dendrite of layer V pyramidal cells increases bidirectional attenuation of EPSPs. *Journal of Neurophysiology*, 85(2):855–868.
- Bezaire, M. J., Raikov, I., Burk, K., Vyas, D., and Soltesz, I. (2016). Interneuronal mechanisms of hippocampal theta oscillations in a full-scale model of the rodent CA1 circuit. *eLife*, 5(e18566):1–106.
- Bezaire, M. J. and Soltesz, I. (2013). Quantitative assessment of CA1 local circuits: Knowledge base for interneuron-pyramidal cell connectivity. *Hippocampus*, 23(9):751–785.
- Bi, G.-q. and Poo, M.-m. (1998). Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. *The Journal of Neuroscience*, 18(24):10464–10472.
- Bibbona, E., Panfilo, G., and Tavella, P. (2008). The Ornstein–Uhlenbeck process as a model of a low pass filtered white noise. *Meteorologia*, 45(6):S117.
- Bienenstock, E. E., Cooper, L. N., and Munro, P. W. (1992). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *The Journal of Neuroscience*, 2(1):32–48.
- Billeh, Y. N., Cai, B., Gratiy, S. L., Dai, K., Iyer, R., Gouwens, N. W., Abbasi-Asl, R., Jia, X., Siegle, J. H., Olsen, S. R., Koch, C., Mihalas, S., and Arkhipov, A. (2020). Systematic Integration of Structural and Functional Data into Multi-Scale Models of Mouse Primary Visual Cortex. *Neuron*, 106(3):388–403.
- Biro, A. A., Holderith, N. B., and Nusser, Z. (2005). Quantal Size Is Independent of the Release Probability at Hippocampal Excitatory Synapses. *Journal of Neuroscience*, 25(1):223–232.

- Biró, A. A., Holderith, N. B., and Nusser, Z. (2006). Release probability-dependent scaling of the postsynaptic responses at single hippocampal GABAergic synapses. *The Journal of neuroscience*, 26(48):12487–96.
- Bittner, K. C., Grienberger, C., Vaidya, S. P., Milstein, A. D., Macklin, J. J., Suh, J., Tonegawa, S., and Magee, J. C. (2015). Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. *Nature Neuroscience*, 18(8):1133–1142.
- Bittner, K. C., Milstein, A. D., Grienberger, C., Romani, S., and Magee, J. C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science*, 357(6355):1033–1036.
- Bliss, T. V. and Collingridge, G. L. (1993). A synaptic model of memory: long-term potentiation in the hippocampus. *Nature*, 361(6407):31–39.
- Bliss, T. V. and Collingridge, G. L. (2013). Expression of NMDA receptor-dependent LTP in the hippocampus: bridging the divide. *Molecular Brain*, 6(5).
- Bliss, T. V. P. and Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of Physiology*, 232(2):331–356.
- Bobrowski, O. and Kahle, M. (2018). Topology of random geometric complexes: a survey. *Journal of Applied and Computational Topology*, 1:331–364.
- Bono, J., Wilmes, K. A., and Clopath, C. (2017). Modelling plasticity in dendrites: from single cells to networks. *Current Opinion in Neurobiology*, 46:136–141.
- Borst, J. G. G. (2010). The low synaptic release probability in vivo. *Trends in Neurosciences*, 33(6):259–266.
- Brasier, D. J. and Feldman, D. E. (2008). Synapse-specific expression of functional presynaptic NMDA receptors in rat somatosensory cortex. *Journal of Neuroscience*, 28(9):2199–2211.
- Brémaud, A., West, D. C., and Thomson, A. M. (2007). Binomial parameters differ across neocortical layers and with different classes of connections in adult rat and cat neocortex. *PNAS*, 104(35):14134–14139.
- Brown, S. D. (2021). Advances in mouse genetics for the study of human disease. *Human Molecular Genetics*, 30(2):R274–R284.
- Brzosko, Z., Mierau, S. B., and Paulsen, O. (2019). Neuromodulation of Spike-Timing-Dependent Plasticity: Past, Present, and Future. *Neuron*, 103(4):563–581.
- Buhl, E. H., Cobb, S. R., Halasy, K., and Somogyi, P. (1995). Properties of unitary IPSPs evoked by anatomically identified basket cells in the rat hippocampus. *European Journal of Neuroscience*, 7(9):1989–2004.
- Buhl, E. H., Halasy, K., and Somogyi, P. (1994a). Diverse sources of hippocampal unitary inhibitory postsynaptic potentials and the number of synaptic release sites. *Nature*, 368:823–828.
- Buhl, E. H., Han, Z.-S., Lörinczi, Z., Stezhka, V. V., Karnup, S. V., and Somogyi, P. (1994b). Physiological Properties of Anatomically Identified AxoAxonic in the Rat Hippocampus. *Journal of Neurophysiology*, 71(4):1289–1307.
- Buonomano, D. V. and Merzenich, M. M. (1995). Temporal information transformed into a spatial code by a neural network with realistic properties. *Science*, 267(5200):1028–1030.
- Buzsáki, G. (1989). Two-stage model of memory trace formation: A role for "noisy" brain states. *Neuroscience*, 31(3):551–570.
- Buzsáki, G. (2010). Neural Syntax: Cell Assemblies, Synapsembles, and Readers. *Neuron*, 68(3):362–385.
- Buzsáki, G. (2019). *The brain from inside out*. Oxford University Press.
- Buzsáki, G. and Mizuseki, K. (2014). The log-dynamic brain: How skewed distributions affect network operations. *Nature Reviews Neuroscience*, 15(4):264–278.

Bibliography

- Carrillo-Reid, L., Miller, J. E. K., Hamm, J. P., Jackson, J., and Yuste, R. (2015). Endogenous sequential cortical activity evoked by visual stimuli. *Journal of Neuroscience*, 35(23):8813–8828.
- Carvalho, T. P. and Buonomano, D. V. (2011). A novel learning rule for long-term plasticity of short-term synaptic plasticity enhances temporal processing. *Frontiers in Integrative Neuroscience*, 5(20).
- Chen, G., Scherr, F., and Maass, W. (2022). A data-based large-scale model for primary visual cortex enables brain-like robust and versatile visual processing. *Science Advances*, 8(44).
- Chen, T. and Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794.
- Chindemi, G., Abdellah, M., Amsalem, O., Benavides-Piccione, R., Delattre, V., Doron, M., Ecker, A., Jaquier, A. T., King, J., Kumbhar, P., Monney, C., Perin, R., Rössert, C., Tuncel, M. A., van Geit, W., DeFelipe, J., Graupner, M., Segev, I., Markram, H., and Muller, E. B. (2022). A calcium-based plasticity model predicts long-term potentiation and depression in the neocortex. *Nature Communications*, 13(3038).
- Chklovskii, D. B., Schikorski, T., and Stevens, C. F. (2002). Wiring optimization in cortical circuits. *Neuron*, 34(3):341–347.
- Christie, J. M. and Jahr, C. E. (2006). Multivesicular Release at Schaffer Collateral-CA1 Hippocampal Synapses. *Journal of Neuroscience*, 26(1):210–216.
- Clopath, C., Büsing, L., Vasilaki, E., and Gerstner, W. (2010). Connectivity reflects coding: A model of voltage-based STDP with homeostasis. *Nature Neuroscience*, 13(3):344–352.
- Cobb, S. R., Halasy, K., Vida, I., Nyíri, G., Tamás, G., Buhl, E. H., and Somogyi, P. (1997). Synaptic effects of identified interneurons innervating both interneurons and pyramidal cells in the rat hippocampus. *Neuroscience*, 79(3):629–648.
- Colangelo, C., Shichkova, P., Keller, D., Markram, H., and Ramaswamy, S. (2019). Cellular, synaptic and network effects of acetylcholine in the neocortex. *Frontiers in Neural Circuits*, 13(24).
- Constantinople, C. M. and Bruno, R. M. (2013). Deep Cortical Layers Are Activated Directly by Thalamus. *Science*, 340(6140):1591–1594.
- Conti, R. and Lisman, J. (2003). The high variance of AMPA receptor- and NMDA receptor-mediated responses at single hippocampal synapses: Evidence for multiquantal release. *PNAS*, 100(8):4885–4890.
- Cornford, J., Mercier, M. S., Leite, M., Magloire, V., Häusser, M., and Kullmann, D. M. (2019). Dendritic NMDA receptors in parvalbumin neurons enable strong and stable neuronal assemblies. *eLife*, 8:e49872.
- Costa, R. P., Froemke, R. C., Sjöström, P. J., and van Rossum, M. C. (2015). Unified pre- and postsynaptic long-term plasticity enables reliable and flexible learning. *eLife*, 4(e09457).
- Cremonesi, F., Hager, G., Wellein, G., and Schürmann, F. (2020). Analytic performance modeling and analysis of detailed neuron simulations. *International Journal of High Performance Computing Applications*, 34(4):428–449.
- Crick, F. (1984). Memory and molecular turnover. *Nature*, 312:101.
- Cutts, C. S. and Eglon, X. S. J. (2014). Detecting pairwise correlations in spike trains: An objective comparison of methods and application to the study of retinal waves. *Journal of Neuroscience*, 34(43):14288–14303.
- Dai, K., Hernando, J., Billeh, Y. N., Gratiy, S. L., Planas, J., Davison, A. P., Dura-Bernal, S., Gleeson, P., Devresse, A., Dichter, B. K., Gevaert, M., King, J. G., van Geit, W. A., Povolotsky, A. V., Muller, E., Courcol, J. D., and Arhipov, A. (2020). The SONATA data format for efficient description of large-scale network models. *PLoS Computational Biology*, 16(2):e1007696.
- D’amour, J. A. and Froemke, R. C. (2015). Inhibitory and excitatory spike-timing-dependent plasticity in the auditory cortex. *Neuron*, 86(2):514–528.

- Davies, D. L. and Bouldin, D. W. (1979). A Cluster Separation Measure. *IEEE Transactions on Pattern Analysis and Machine Learning*, PAMI-1(2):224–227.
- Daw, M. I., Tricoire, L., Erdelyi, F., Szabo, G., and McBain, C. J. (2009). Asynchronous Transmitter Release from Cholecystinin-Containing Inhibitory Interneurons Is Widespread and Target-Cell Independent. *Journal of Neuroscience*, 29(36):11112–11122.
- Debanne, D., Gähwiler, B. H., and Thompson, S. M. (1998). Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *Journal of Physiology*, 507(1):237–247.
- Debanne, D., Guérineau, N. C., Gähwiler, B. H., and Thompson, S. M. (1995). Physiology and pharmacology of unitary synaptic connections between pairs of cells in areas CA3 and CA1 of rat hippocampal slice cultures. *Journal of Neurophysiology*, 73(3):1282–1294.
- Del Castillo, J. and Katz, B. (1954). Quantal components of the end-plate potential. *The Journal of Physiology*, 124(3):560–573.
- Denker, M., Yegenoglu, A., and Grün, S. (2018). Collaborative HPC-enabled workflows on the HBP Collaboratory using the Elephant framework. In *Neuroinformatics 2018*.
- Deperrois, N. and Graupner, M. (2020). Short-term depression and long-term plasticity together tune sensitive range of synaptic plasticity. *PLoS Computational Biology*, 16(9):e100826.
- Destexhe, A., Rudolph, M., Fellous, J. M., and Sejnowski, T. J. (2001). Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience*, 107(1):13–24.
- Deuchars, J. and Thomson, A. M. (1996). CA1 pyramid-pyramid connections in rat hippocampus in vitro: Dual intracellular recordings with biocytin filling. *Neuroscience*, 74(4):1009–1018.
- Dobrunz, L. E. and Stevens, C. F. (1997). Heterogeneity of release probability, facilitation, and depletion at central synapses. *Neuron*, 18(6):995–1008.
- Dragoi, G. and Buzsáki, G. (2006). Temporal Encoding of Place Sequences by Hippocampal Cell Assemblies. *Neuron*, 50(1):145–157.
- Druckmann, S., Feng, L., Lee, B., Yook, C., Zhao, T., Magee, J. C., and Kim, J. (2014). Structured Synaptic Connectivity between Hippocampal Regions. *Neuron*, 81(3):629–640.
- Dura-Bernal, S., Neymotin, S. A., Suter, B. A., Dacre, J., Moreira, J. V., Urdapilleta, E., Schiemann, J., Duguid, I., Shepherd, G. M., and Lytton, W. W. (2023). Multiscale model of primary motor cortex circuits predicts in vivo cell-type-specific, behavioral state-dependent dynamics. *Cell Reports*, 42(112574).
- Eccles, J. C., Katz, B., and Kuffler, S. W. (1941). Nature of the "endplate potential" in curarized muscle. *Journal of Physiology*, 4:362–387.
- Ecker, A., Romani, A., Sáray, S., Káli, S., Migliore, M., Falck, J., Lange, S., Mercer, A., Thomson, A. M., Muller, E., Reimann, M. W., and Ramaswamy, S. (2020). Data-driven integration of hippocampal CA1 synaptic physiology in silico. *Hippocampus*, 30(11):1129–1145.
- Ecker, A., Santander, D. E., Abdellah, M., Alonso, J. B., Bolaños-Puchet, S., Chindemi, G., Isbister, J. B., King, J. G., Kumbhar, P., Magkanaris, I., Muller, E. B., and Reimann, M. W. (2023a). Sparse and specific long-term plasticity emerge without homeostasis in a biophysically detailed cortical model. *bioRxiv*.
- Ecker, A., Santander, D. E., Bolaños-Puchet, S., Isbister, J. B., and Reimann, M. W. (2023b). Cortical cell assemblies and their underlying connectivity: an in silico study. *bioRxiv*.
- Egger, R., Narayanan, R. T., Guest, J. M., Bast, A., Udvardy, D., Messori, L. F., Das, S., de Kock, C. P., and Oberlaender, M. (2020). Cortical Output Is Gated by Horizontally Projecting Neurons in the Deep Layers. *Neuron*, 105:122–137.
- Egger, V., Feldmeyer, D., and Sakmann, B. (1999). Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. *Nature Neuroscience*, 2(12):1098–1105.

Bibliography

- Elfant, D., Pal, B. Z., Emptage, N., and Capogna, M. (2008). Specific inhibitory synapses shift the balance from feedforward to feedback inhibition of hippocampal CA1 pyramidal cells. *European Journal of Neuroscience*, 27(1):104–113.
- Éltes, T., Kirizs, T., Nusser, Z., and Holderith, N. (2017). Target Cell Type-Dependent Differences in Ca²⁺ Channel Function Underlie Distinct Release Probabilities at Hippocampal Glutamatergic Terminals. *The Journal of Neuroscience*, 37(7):1910–1924.
- Enoki, R., Ling Hu, Y., Hamilton, D., and Fine, A. (2009). Expression of Long-Term Plasticity at Individual Synapses in Hippocampus Is Graded, Bidirectional, and Mainly Presynaptic: Optical Quantal Analysis. *Neuron*, 62(2):242–253.
- Espinoza, C., Guzman, S. J., Zhang, X., and Jonas, P. (2018). Parvalbumin+ interneurons obey unique connectivity rules and establish a powerful lateral-inhibition microcircuit in dentate gyrus. *Nature Communications*, 9(4605).
- Fan, X. and Markram, H. (2019). A brief history of simulation neuroscience. *Frontiers in Neuroinformatics*, 13(32).
- Farinella, M., Ruedt, D. T., Gleeson, P., Lanore, E., and Silver, R. A. (2014). Glutamate-Bound NMDARs Arising from In Vivo-like Network Activity Extend Spatio-temporal Integration in a L5 Cortical Pyramidal Cell Model. *PLoS Computational Biology*, 10(4).
- Fauth, M. J. and Van Rossum, M. C. (2019). Self-organized reactivation maintains and reinforces memories despite synaptic turnover. *eLife*, 8:e43717.
- Feldmeyer, D. (2012). Excitatory neuronal connectivity in the barrel cortex. *Frontiers in Neuroanatomy*, 6(24).
- Feldmeyer, D., Egger, V., Lübke, J., and Sakmann, B. (1999). Reliable synaptic connections between pairs of excitatory layer 4 neurones within a single 'barrel' of developing rat somatosensory cortex. *Journal of Physiology*, 521(1):169–190.
- Feldmeyer, D., Lübke, J., and Sakmann, B. (2006). Efficacy and connectivity of intracolumnar pairs of layer 2/3 pyramidal cells in the barrel cortex of juvenile rats. *Journal of Physiology*, 575(2):583–602.
- Feldmeyer, D., Lübke, J., Silver, R. A., and Sakmann, B. (2002). Synaptic connections between layer 4 spiny neurone-layer 2/3 pyramidal cell pairs in juvenile rat barrel cortex: Physiology and anatomy of interlaminar signalling within a cortical column. *Journal of Physiology*, 538(3):803–822.
- Feldmeyer, D., Roth, A., and Sakmann, B. (2005). Monosynaptic connections between pairs of spiny stellate cells in layer 4 and pyramidal cells in layer 5A indicate that lemniscal and paralemniscal afferent pathways converge in the infragranular somatosensory cortex. *Journal of Neuroscience*, 25(13):3423–3431.
- Feulner, B., Perich, M. G., Chowdhury, R. H., Miller, L. E., Gallego, J. A., and Clopath, C. (2022). Small, correlated changes in synaptic connectivity may facilitate rapid motor learning. *Nature Communications*, 13(5163).
- Field, R. E., James, A. D., Tremblay, R., Miehl, C., Rudy, B., Gjorgjieva, J., Froemke, R. C., Field, R. E., James, A. D., Tremblay, R., Miehl, C., and Rudy, B. (2020). Heterosynaptic Plasticity Determines the Set Point for Cortical Excitatory-Inhibitory Balance Article Heterosynaptic Plasticity Determines the Set Point for Cortical Excitatory-Inhibitory Balance. *Neuron*, 106.
- Földy, C., Lee, S.-h., Morgan, R. J., and Soltesz, I. (2010). Regulation of fast-spiking basket cell synapses by the chloride channel ClC-2. *Nature Neuroscience*, 13(9):1047–1049.
- Földy, C., Lee, S. Y., Szabadics, J., Neu, A., and Soltesz, I. (2007). Cell type-specific gating of perisomatic inhibition by cholecystokinin. *Nature neuroscience*, 10(9):1128–1130.
- Frankle, J. and Carbin, M. (2019). The lottery ticket hypothesis: Finding sparse, trainable neural networks. *ICLR*.

- Froemke, R. C., Letzkus, J. J., Kampa, B. M., Hang, G. B., and Stuart, G. J. (2010). Dendritic synapse location and neocortical spike-timing-dependent plasticity. *Frontiers in Synaptic Neuroscience*, 2(29).
- Fuentealba, P., Begum, R., Capogna, M., Jinno, S., Márton, L. F., Csicsvari, J., Thomson, A., Somogyi, P., and Klausberger, T. (2008). Ivy Cells: A Population of Nitric-Oxide-Producing, Slow-Spiking GABAergic Neurons and Their Involvement in Hippocampal Network Activity. *Neuron*, 57(6):917–929.
- Fuhrmann, G., Cowan, A., Segev, I., Tsodyks, M., and Stricker, C. (2004). Multiple mechanisms govern the dynamics of depression at neocortical synapses of young rats. *Journal of Physiology*, 557(2):415–438.
- Fuhrmann, G., Segev, I., Markram, H., and Tsodyks, M. (2002). Coding of Temporal Information by Activity-Dependent Synapses. *Journal of Neurophysiology*, 87(1):140–148.
- Fusi, S. and Abbott, L. F. (2007). Limits on the memory storage capacity of bounded synapses. *Nature Neuroscience*, 10(4):485–493.
- Fusi, S., Drew, P. J., and Abbott, L. F. (2005). Cascade models of synaptically stored memories. *Neuron*, 45(4):599–611.
- Gambino, F., Pagès, S., Kehayas, V., Baptista, D., Tatti, R., Carleton, A., and Holtmaat, A. (2014). Sensory-evoked LTP driven by dendritic plateau potentials in vivo. *Nature*, 515(7525):116–119.
- Gastaldi, C., Schwalger, T., de Falco, E., Quiroga, R. Q., and Gerstner, W. (2021). When shared concept cells support associations: Theory of overlapping memory engrams. *PLoS Computational Biology*, 17(12):e1009691.
- Gentet, L. J., Avermann, M., Matyas, F., Staiger, J. F., and Petersen, C. C. (2010). Membrane Potential Dynamics of GABAergic Neurons in the Barrel Cortex of Behaving Mice. *Neuron*, 65(3):422–435.
- Gentet, L. J., Kremer, Y., Taniguchi, H., Huang, Z. J., Staiger, J. F., and Petersen, C. C. (2012). Unique functional properties of somatostatin-expressing GABAergic neurons in mouse barrel cortex. *Nature Neuroscience*, 15(4):607–612.
- Gerstner, W., Kempter, R., Hemmen, J. L. V., and Wagnert, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383:76–78.
- Gerstner, W., Kistler, W. M., Naud, R., and Paninski, L. (2014). *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge University Press.
- Gerstner, W., Ritz, R., and van Hemmen, J. L. (1993). Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. *Biological Cybernetics*, 69(5-6):503–515.
- Goetz, L., Roth, A., and Häusser, M. (2021). Active dendrites enable strong but sparse inputs to determine orientation selectivity. *PNAS*, 118(30):e2017339118.
- Goldberg, J., Holthoff, K., and Yuste, R. (2002). A problem with Hebb and local spikes. *Trends in Neurosciences*, 25(9):433–435.
- Goldman, M. S., Maldonado, P., and Abbott, L. F. (2002). Redundancy reduction and sustained firing with stochastic depressing synapses. *Journal of Neuroscience*, 22(2):584–591.
- Gonzalez, K. C., Losonczy, A., and Negrean, A. (2022). Dendritic Excitability and Synaptic Plasticity In Vitro and In Vivo. *Neuroscience*, 489:165–175.
- Graupner, M. and Brunel, N. (2012). Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. *PNAS*, 109(10):3991–3996.
- Graupner, M., Wallisch, P., and Ostojic, S. (2016). Natural firing patterns imply low sensitivity of synaptic plasticity to spike timing compared with firing rate. *Journal of Neuroscience*, 36(44):11238–11258.
- Graves, A. R., Roth, R. H., Tan, H. L., Zhu, Q., Bygrave, A. M., Lopez-Ortega, E., Hong, I., Spiegel, A. C., Johnson, R. C., Vogelstein, J. T., Tward, D. J., Miller, M. I., and Hugarir, R. L. (2021). Visualizing synaptic plasticity in vivo by large-scale imaging of endogenous AMPA receptors. *eLife*, 10:e66809.
- Gray, R. M. (2011). *Entropy and information theory*. Springer Science & Business Media.

Bibliography

- Groc, L., Gustafsson, B., and Hanse, E. (2002). Spontaneous unitary synaptic activity in CA1 pyramidal neurons during early postnatal development: constant contribution of AMPA and NMDA receptors. *The Journal of Neuroscience*, 22(13):5552–5562.
- Gulyás, A. I., Freund, T. F., and Káli, S. (2016). The Effects of Realistic Synaptic Distribution and 3D Geometry on Signal Integration and Extracellular Field Generation of Hippocampal Pyramidal Cells and Inhibitory Neurons. *Frontiers in Neural Circuits*, 10(88).
- Gulyás, A. I., Miles, R., Sík, A., Tóth, K., Tamamaki, N., and Freund, T. F. (1993). Hippocampal pyramidal cells excite inhibitory neurons through a single release site. *Nature*, 366:683–687.
- Gulyás, B. and Somogyi, P. (2013). János Szentágotthai. 31 October 1912—8 September 1994. *Biogr Mem Fellows R Soc*.
- Gupta, A., Wang, Y., and Markram, H. (2000). Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. *Science*, 287(5451):273–8.
- Gurumurthy, C. B. and Kent Lloyd, K. C. (2019). Generating mouse models for biomedical research: Technological advances. *Disease Models and Mechanisms*, 12:dmm029462.
- Guzman, S. J., Schlögl, A., Frotsher, M., and Jonas, P. (2016). Synaptic mechanisms of pattern completion in the hippocampal CA3 network. *Science*, 335(6304):11117–11123.
- Harris, J. J., Jolivet, R., and Attwell, D. (2012). Synaptic Energy Use and Supply. *Neuron*, 75(5):762–777.
- Harris, K. D. (2005). Neural signatures of cell assembly organization. *Nature Reviews Neuroscience*, 6(5):399–407.
- Harris, K. D., Csicsvari, J., Hirase, H., Dragoi, G., and Buzsáki, G. (2003). Organization of cell assemblies in the hippocampus. *Nature*, 424(6948):552–556.
- Harris, K. D. and Shepherd, G. M. (2015). The neocortical circuit: Themes and variations. *Nature Neuroscience*, 18(2):170–181.
- Hebb, D. O. (1949). *The Organization of Behavior; A Neuropsychological Theory*. John Wiley & Sons, Inc., New York.
- Hennequin, G., Agnes, E. J., and Vogels, T. P. (2017). Inhibitory Plasticity: Balance, Control, and Codependence. *Annual Review of Neuroscience*, 40:557–579.
- Hennig, M. H. (2013). Theoretical models of synaptic short term plasticity. *Frontiers in Computational Neuroscience*, 7(45).
- Herzog, R., Morales, A., Mora, S., Araya, J., Escobar, M. J., Palacios, A. G., and Cofré, R. (2021). Scalable and accurate method for neuronal ensemble detection in spiking neural networks. *PLoS ONE*, 16(7):e0251647.
- Hestrin, S., Sah, P., and Nicoll, R. A. (1990). Mechanisms Generating the Time Course of Dual Component Excitatory Synaptic Currents Recorded in Hippocampal Slices. *Neuron*, 5:247–253.
- Higgins, D., Graupner, M., and Brunel, N. (2014). Memory Maintenance in Synapses with Calcium-Based Plasticity in the Presence of Background Activity. *PLoS Computational Biology*, 10(10):e1003834.
- Hill, A. V. (1910). The possible effects of the aggregation of the molecules of haemoglobin on its dissociation curves. *Journal of Physiology*, 40:4–7.
- Hines, M. L. and Carnevale, N. T. (1997). The NEURON simulation environment. *Neural computation*, 9(6):1179–1209.
- Hines, M. L. and Carnevale, N. T. (2000). Expanding NEURON's repertoire of mechanisms with NMODL. *Neural Computation*, 12(5):995–1007.
- Hines, M. L., Eichner, H., and Schürmann, F. (2008a). Neuron splitting in compute-bound parallel network simulations enables runtime scaling with twice as many processors. *Journal of computational neuroscience*, 25(1):203–210.

- Hines, M. L., Markram, H., and Schürmann, F. (2008b). Fully implicit parallel simulation of single neurons. *Journal of computational neuroscience*, 25(3):439–448.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *PNAS*, 79(8):2554–2558.
- Hu, H., Martina, M., and Jonas, P. (2010). Dendritic mechanisms underlying rapid synaptic activation of fast-spiking hippocampal interneurons. *Science*, 327(5961):52–58.
- Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160:106–154.
- Iacaruso, M. F., Gasler, I. T., and Hofer, S. B. (2017). Synaptic organization of visual space in primary visual cortex. *Nature*, 547(7664):449–452.
- Inglebert, Y., Aljadeff, J., Brunel, N., and Debanne, D. (2020). Synaptic plasticity rules with physiological calcium levels. *PNAS*, 117(52):33639–33648.
- Isbister, J. B., Ecker, A., Pokorny, C., Bolanos-Puchet, S., Egas Santander, D., et al. (2023). Modeling and Simulation of Neocortical Micro- and Mesocircuitry. Part II: Physiology and Experimentation. *bioRxiv*.
- Isbister, J. B., Reyes-Puerta, V., Sun, J.-J., Horenko, I., and Luhmann, H. J. (2021). Clustering and control for adaptation uncovers time-warped spike time patterns in cortical networks in vivo. *Scientific Reports*, 11:15066.
- Jahr, C. E. and Stevens, C. F. (1990). Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics. *The Journal of neuroscience*, 10(9):3178–3182.
- Jain, A., Nakahata, Y., Watabe, T., Rusina, P., South, K., Adachi, K., Yan, L., Simorowski, N., Furukawa, H., and Yasuda, R. (2023). Dendritic, delayed, and stochastic CaMKII activation underlies behavioral time scale plasticity in CA1 synapses. *bioRxiv*.
- Jiang, X., Shen, S., Cadwell, C. R., Berens, P., Sinz, F., Ecker, A. S., Patel, S., and Tolias, A. S. (2015). Principles of connectivity among morphologically defined cell types in adult neocortex. *Science*, 350(6264):aac9462.
- Jonas, P., Major, G., and Sakmann, B. (1993). Quantal components of unitary EPSCs at the mossy fibre synapse on CA3 pyramidal cells of rat hippocampus. *The Journal of Physiology*, 472:615–663.
- Jones, H. and Keep, R. (1988). Brain Fluid Calcium Concentration and Response To Acute Hypercalcaemia During Development in the Rat. *Journal of Physiology*, 402:579–593.
- Kahle, M. (2009). Topology of random clique complexes. *Discrete Mathematics*, 309:1658–1671.
- Kanari, L., Ramaswamy, S., Shi, Y., Morand, S., Meystre, J., Perin, R., Abdellah, M., Wang, Y., Hess, K., and Markram, H. (2019). Objective Morphological Classification of Neocortical Pyramidal Cells. *Cerebral Cortex*, 29(4):1719–1735.
- Karayannis, T., Elfant, D., Huerta-Ocampo, I., Teki, S., Scott, R. S., Rusakov, D. A., Jones, M. V., and Capogna, M. (2010). Slow GABA transient and receptor desensitization shape synaptic responses evoked by hippocampal neurogliaform cells. *The Journal of neuroscience*, 30(29):9898–909.
- Karnani, M. M., Jackson, J., Ayzenshtat, I., Sichani, X. H., Manoocheri, K., Kim, S., and Yuste, R. (2016). Opening holes in the blanket of inhibition: Localized lateral disinhibition by vip interneurons. *Journal of Neuroscience*, 36(12):3471–3480.
- Kastellakis, G. and Poirazi, P. (2019). Synaptic Clustering and Memory Formation. *Frontiers in Molecular Neuroscience*, 12(300).
- Keijsers, J. and Sprekeler, H. (2022). Optimizing interneuron circuits for compartment-specific feedback inhibition. *PLoS Computational Biology*, 18(4):e1009933.
- Kempler, R., Gerstner, W., and van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Physical Review*, 59(4):4498–4514.

Bibliography

- Kim, D., Park, P., Li, X., Wong-Campos, J. D., Tian, H., Moulton, E. M., Grimm, J. B., Lavis, L., and Cohen, A. E. (2023). Mapping memories: pulse-chase labeling reveals AMPA receptor dynamics during memory formation. *bioRxiv*.
- Klausberger, T. and Somogyi, P. (2008). Neuronal diversity and temporal dynamics: The unity of hippocampal circuit operations. *Science*, 321(5885):53–57.
- Kohus, Z., Káli, S., Schlinghoff, D., Papp, O., Rovira-Esteban, L., Freund, T. F., Hájos, N., and Gulyás, A. I. (2016). Properties and dynamics of inhibitory synaptic communication within the CA3 microcircuits of pyramidal cells and interneurons expressing parvalbumin or cholecystokinin. *The Journal of physiology*, 594(13):3745–74.
- Korinek, M., Sedlacek, M., Cais, O., Dittert, I., and Vyklícký, L. (2010). Temperature dependence of N-methyl-d-aspartate receptor channels and N-methyl-d-aspartate receptor excitatory postsynaptic currents. *Neuroscience*, 165(3):736–748.
- Kossio, Y. F. K., Goedeke, S., Klos, C., and Memmesheimer, R. M. (2021). Drifting assemblies for persistent memory: Neuron transitions and unsupervised compensation. *PNAS*, 118(46):e2023832118.
- Krotov, D. and Hopfield, J. J. (2016). Dense associative memory for pattern recognition. *Advances in Neural Information Processing Systems*, 29:1172–1180.
- Kumbhar, P., Hines, M., Fouriaux, J., Ovcharenko, A., King, J., Delalondre, F., and Schürmann, F. (2019). CoreNEURON : An Optimized Compute Engine for the NEURON Simulator. *Frontiers in Neuroinformatics*, 13(63).
- Landau, A. T., Park, P., Wong-Campos, J. D., Tian, H., Cohen, A. E., and Sabatini, B. L. (2022). Dendritic branch structure compartmentalizes voltage-dependent calcium influx in cortical layer 2/3 pyramidal cells. *eLife*, 11(e76993).
- Larkum, M. E. (2013). A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends in Neurosciences*, 36(3):141–151.
- Larkum, M. E., Zhu, J. J., and Sakmann, B. (2001). Dendritic mechanisms underlying the coupling of the dendritic with the axonal action potential initiation zone of adult rat layer 5 pyramidal neurons. *Journal of Physiology*, 533(2):447–466.
- Le Bé, J.-V. and Markram, H. (2006). Spontaneous and evoked synaptic rewiring in the neonatal neocortex. *PNAS*, 103(35):13214–13219.
- Le Bé, J. V., Silberberg, G., Wang, Y., and Markram, H. (2007). Morphological, electrophysiological, and synaptic properties of corticocallosal pyramidal cells in the neonatal rat neocortex. *Cerebral Cortex*, 17(9):2204–2213.
- Le Roux, N., Cabezas, C., Böhm, U. L., and Poncer, J. C. (2013). Input-specific learning rules at excitatory synapses onto hippocampal parvalbumin-expressing interneurons. *Journal of Physiology*, 591(7):1809–1822.
- Lee, S.-H., Földy, C., and Soltesz, I. (2010). Distinct endocannabinoid control of GABA release at perisomatic and dendritic synapses in the hippocampus. *The Journal of neuroscience*, 30(23):7993–8000.
- Lee, S.-H., Marchionni, I., Bezaire, M., Varga, C., Danielson, N., Lovett-Barron, M., Losonczy, A., and Soltesz, I. (2014). Parvalbumin-Positive Basket Cells Differentiate among Hippocampal Pyramidal Cells. *Neuron*, 82(5):1129–1144.
- Letzkus, J. J., Kampa, B. M., and Stuart, G. J. (2006). Learning rules for spike timing-dependent plasticity depend on dendritic synapse location. *Journal of Neuroscience*, 26(41):10420–10429.
- Ling, D. S. and Benardo, L. S. (1999). Restrictions on inhibitory circuits contribute to limited recruitment of fast inhibition in rat neocortical pyramidal cells. *Journal of Neurophysiology*, 82(4):1793–1807.
- Lisman, J. (1989). A mechanism for the Hebb and the anti-Hebb processes underlying learning and memory. *PNAS*, 86(23):9574–9578.

- Lisman, J. E. (1985). A mechanism for memory storage insensitive to molecular turnover: A bistable autophosphorylating kinase. *PNAS*, 82:3055–3057.
- Lisman, J. E. and Spruston, N. (2005). Postsynaptic depolarization requirements for LTP and LTD: A critique of spike timing-dependent plasticity. *Nature neuroscience*, 8(7):839–841.
- Litwin-Kumar, A. and Doiron, B. (2014). Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nature Communications*, 5(5319).
- Loebel, A., Silberberg, G., Helbig, D., Markram, H., Tsodyks, M., and Richardson, M. J. E. (2009). Multiquantal release underlies the distribution of synaptic efficacies in the neocortex. *Frontiers in Cellular Neuroscience*, 3(27).
- Lopes-dos Santos, V., Ribeiro, S., and Tort, A. B. (2013). Detecting cell assemblies in large neuronal populations. *Journal of Neuroscience Methods*, 220(2):149–166.
- Losonczy, A. and Magee, J. C. (2006). Integrative Properties of Radial Oblique Dendrites in Hippocampal CA1 Pyramidal Neurons. *Neuron*, 50(2):291–307.
- Losonczy, A., Zhang, L., Shigemoto, R., Somogyi, P., and Nusser, Z. (2002). Cell type dependence and variability in the short-term plasticity of EPSCs in identified mouse hippocampal interneurons. *The Journal of physiology*, 542(1):193–210.
- Lübke, J. and Feldmeyer, D. (2007). Excitatory signal flow and connectivity in a cortical column: Focus on barrel cortex. *Brain Structure and Function*, 212:3–17.
- Maass, W. and Markram, H. (2002). Synapses as dynamic memory buffers. *Neural Networks*, 15(2):155–161.
- Maass, W. and Markram, H. (2004). On the computational power of circuits of spiking neurons. *Journal of Computer and System Sciences*, 69(4):593–616.
- Maccafferri, G., Roberts, J. D. B., Szucs, P., Cottingham, C. A., and Somogyi, P. (2000). Cell surface domain specific postsynaptic currents evoked by identified GABAergic neurons in rat hippocampus in vitro. *Journal of Physiology*, 524(1):91–116.
- Magee, J. C. and Cook, E. P. (2000). Somatic EPSP amplitude is independent of synapse location in hippocampal pyramidal neurons. *Nature neuroscience*, 3(9):895–903.
- Mäki-Marttunen, T., Iannella, N., Edwards, A. G., Einevoll, G. T., and Blackwell, K. T. (2020). A unified computational model for cortical post-synaptic plasticity. *eLife*, 9:e55714.
- Makino, H. (2019). Top-down control: A unified principle of cortical learning. *Neuroscience Research*, 141:23–28.
- Malinow, R. and Malenka, R. C. (2002). AMPA receptor trafficking and synaptic plasticity. *Annual review of neuroscience*, 25(1):103–126.
- Mallya, A. and Lazebnik, S. (2018). PackNet: Adding Multiple Tasks to a Single Network by Iterative Pruning. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.
- Markram, H., Gerstner, W., and Sjöström, P. J. (2011). A history of spike-timing-dependent plasticity. *Frontiers in Synaptic Neuroscience*, 3(4).
- Markram, H., Lübke, J., Frotscher, M., Roth, A., and Sakmann, B. (1997a). Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex. *Journal of Physiology*, 500(2):409–440.
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997b). Regulation of Synaptic Efficacy by Coincidence of Postsynaptic APs and EPSPs. *Science*, 275(5297):213–215.
- Markram, H., Müller, E. B., Ramaswamy, S., Reimann, M. W., et al. (2015). Reconstruction and Simulation of Neocortical Microcircuitry. *Cell*, 163:456–492.
- Markram, H. and Tsodyks, M. (1996). Neocortical Pyramidal Neurons. *Letters to Nature*, 382:807–810.

Bibliography

- Markram, H., Wang, Y., and Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *PNAS*, 95(9):5323–8.
- Massimini, M. and Amzica, F. (2001). Extracellular calcium fluctuations and intracellular potentials in the cortex during the slow sleep oscillation. *Journal of Neurophysiology*, 85(3):1346–1350.
- Mateos-Aparicio, P. and Rodríguez-Moreno, A. (2019). The impact of studying brain plasticity. *Frontiers in Cellular Neuroscience*, 13.
- Matsuzaki, M., Ellis-Davies, G. C., Nemoto, T., Miyashita, Y., Iino, M., and Kasai, H. (2001). Dendritic spine geometry is critical for AMPA receptor expression in hippocampal CA1 pyramidal neurons. *Nature Neuroscience*, 4(11):1086–1092.
- Matta, J. A., Pelkey, K. A., Craig, M. T., Chittajallu, R., Jeffries, B. W., and McBain, C. J. (2013). Developmental origin dictates interneuron AMPA and NMDA receptor subunit composition and plasticity. *Nature Neuroscience*, 16(8):1032–1041.
- Mayer, M. L., Westbrook, G. L., and Guthrie, P. B. (1984). Voltage-dependent block by Mg²⁺ of NMDA responses in spinal cord neurones. *Nature*, 309:261–263.
- McCloskey, M. and Cohen, N. J. (1989). Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. *The Psychology of Learning and Motivation*, 24:109–165.
- McKernan, M. G. and Shinnick-Gallagher, P. (1997). Fear conditioning induces a lasting potentiation of synaptic currents in vitro. *Nature*, 390:607–611.
- Megias, M., Emri, Z., Freund, T. F., and Gulyás, A. I. (2001). Total number and distribution of inhibitory and excitatory synapses on hippocampal CA1 pyramidal cells. *Neuroscience*, 102(3):527–540.
- Mesradi, M., Genoux, A., Cuplov, V., Abi Haidar, D., Jan, S., Buvat, I., and Pain, F. (2013). Experimental and analytical comparative study of optical coefficient of fresh and frozen rat tissues. *Journal of Biomedical Optics*, 18(11):117010.
- Meyer, H. S., Wimmer, V. C., Hemberger, M., Bruno, R. M., De Kock, C. P., Frick, A., Sakmann, B., and Helmstaedter, M. (2010). Cell type-specific thalamic innervation in a column of rat vibrissa cortex. *Cerebral Cortex*, 20(10):2287–2303.
- MICrONS Consortium (2021). Functional connectomics spanning multiple areas of mouse visual cortex. *bioRxiv*.
- Migliore, M., Ferrante, M., and Ascoli, G. A. (2005). Signal Propagation in oblique dendrites of CA1 pyramidal cells. *Journal of Neurophysiology*, 94(6):4145–4155.
- Migliore, M., Hoffman, D. A., Magee, J. C., and Johnston, D. (1999). Role of an A-Type K⁺ Conductance in the Back-Propagation of Action Potentials in the Dendrites of Hippocampal Pyramidal Neurons. *Journal of Computational Neuroscience*, 7:5–15.
- Migliore, R., Lupascu, C. A., Bologna, L. L., Romani, A., Courcol, J.-D., Antonel, S., Van Geit, W. A. H., Thomson, A. M., Mercer, A., Lange, S., Falck, J., Rössert, C. A., Shi, Y., Hagens, O., Pezzoli, M., Freund, T. F., Kali, S., Müller, E. B., Schürmann, F., Markram, H., and Migliore, M. (2018). The physiological variability of channel density in hippocampal CA1 pyramidal cells and interneurons explored using a unified data-driven modeling workflow. *PLoS Computational Biology*, 14(9):e1006423.
- Miller, J. E. K., Ayzenshtat, I., Carrillo-Reid, L., and Yuste, R. (2014). Visual stimuli recruit intrinsically generated cortical ensembles. *PNAS*, 111(38):E4053–E4061.
- Mody, I. and Pearce, R. A. (2004). Diversity of inhibitory neurotransmission through GABA A receptors. *Trends in Neurosciences*, 27(9):569–575.
- Montijn, X. J. S., Olcese, U., and Pennartz, X. C. M. A. (2016). Visual Stimulus Detection Correlates with the Consistency of Temporal Sequences within Stereotyped Events of V1 Neuronal Population Activity. *The Journal of Neuroscience*, 36(33):8624–8640.
- Moradi, K. and Ascoli, G. A. (2019). A comprehensive knowledge base of synaptic electrophysiology in the rodent hippocampal formation. *Hippocampus*.

- Morrison, A., Aertsen, A., and Diesmann, M. (2007). Spike-timing-dependent plasticity in balanced random networks. *Neural Computation*, 19(6):1437–1467.
- Morse, T., Carnevale, N., Mutalik, P., Migliore, M., and Shepherd, G. (2010). Abnormal excitability of oblique dendrites implicated in early Alzheimer's: a computational study. *Frontiers in Neural Circuits*, 4(16).
- Movshon, J. A., Thomson, I. D., and Tolhurst, D. J. (1978). Receptive field organization of complex cells in cat striate cortex. *Journal of Physiology*, 283:79–99.
- Muller, E., Buesing, L., Schemmel, J., and Meier, K. (2007). Spike-Frequency Adapting Neural Ensembles: Beyond Mean Adaptation and Renewal Theories. *Neural Computation*, 19:2958–3010.
- Myme, C. I. O., Sugino, K., Turrigiano, G. G., and Nelson, S. B. (2003). The NMDA-to-AMPA Ratio at Synapses Onto Layer 2/3 Pyramidal Neurons Is Conserved Across Prefrontal and Visual Cortices. *Journal of Neurophysiology*, 90(2):771–779.
- Naud, R. and Sprekeler, H. (2018). Sparse bursts optimize information transmission in a multiplexed neural code. *PNAS*, 115(27):E6329–E6338.
- Neher, E. (1992). Correction for liquid junction potentials in patch clamp experiments. *Methods in Enzymology*, 207:123–131.
- Neu, A., Földy, C., and Soltesz, I. (2007). Postsynaptic origin of CB1-dependent tonic inhibition of GABA release at cholecystokinin-positive basket cell to pyramidal cell synapses in the CA1 region of the rat hippocampus. *The Journal of physiology*, 578(1):233–247.
- Nevian, T., Larkum, M. E., Polsky, A., and Schiller, J. (2007). Properties of basal dendrites of layer 5 pyramidal neurons: A direct patch-clamp recording study. *Nature Neuroscience*, 10(2):206–214.
- Nicoll, R. A. and Malenka, R. C. (1995). Contrasting properties of two forms of long-term potentiation in the hippocampus. *Nature*, 377(6545):115–118.
- Nolte, M., Gal, E., Markram, H., and Reimann, M. W. (2020). Impact of higher order network structure on emergent cortical activity. *Network Neuroscience*, 4(1):292–314.
- Nolte, M., Reimann, M. W., King, J. G., Markram, H., and Muller, E. B. (2019). Cortical reliability amid noise and chaos. *Nature Communications*, 10(3792).
- Nusser, Z., Lujan, R., Laube, G., Roberts, J. D. B., Molnar, E., and Somogyi, P. (1998). Cell type and pathway dependence of synaptic AMPA receptor number and variability in the hippocampus. *Neuron*, 21(3):545–559.
- Oby, E. R., Golub, M. D., Hennig, J. A., Degenhart, A. D., Tyler-Kabara, E. C., Yu, B. M., Chase, S. M., and Batista, A. P. (2019). New neural activity patterns emerge with long-term learning. *PNAS*, 116(30):15210–15215.
- Ohana, O. and Sakmann, B. (1998). Transmitter release modulation in nerve terminals of rat neocortical pyramidal cells by intracellular calcium buffers. *Journal of Physiology*, 513(1):135–148.
- Oja, E. (1982). Simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, 15(3):267–273.
- O'Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34(1):171–175.
- Park, J. M., Hong, Y. K., Rodgers, C. C., Dahan, J. B., Schmidt, E. R., and Bruno, R. M. (2020). Deep and superficial layers of the primary somatosensory cortex are critical for whisker-based texture discrimination in mice. *bioRxiv*.
- Pawelzik, H., Bannister, A. P., Deuchars, J., Ilia, M., and Thomson, A. M. (1999). Modulation of bistratified cell IPSPs and basket cell IPSPs by pentobarbitone sodium, diazepam and Zn²⁺: Dual recordings in slices of adult rat hippocampus. *European Journal of Neuroscience*, 11(10):3552–3564.

Bibliography

- Pawelzik, H., Hughes, D. I., and Thomson, A. M. (2002). Physiological and morphological diversity of immunocytochemically defined parvalbumin- and cholecystokinin-positive interneurons in CA1 of the adult rat hippocampus. *Journal of Comparative Neurology*, 443(4):346–367.
- Pawelzik, H., Hughes, D. I., and Thomson, A. M. (2003). Modulation of inhibitory autapses and synapses on rat CA1 interneurons by GABAA receptor ligands. *Journal of Physiology*, 546(3):701–716.
- Pelkey, K. A., Chittajallu, R., Craig, M. T., Tricoire, L., Wester, J. C., and McBain, C. J. (2017). Hippocampal GABAergic inhibitory interneurons. *Physiological Reviews*, 97(4):1619–1747.
- Pérez-Ortega, J., Alejandre-García, T., and Yuste, R. (2021). Long-term stability of cortical ensembles. *eLife*, 10:e64449.
- Perin, R., Berger, T. K., and Markram, H. (2011). A synaptic organizing principle for cortical neuronal groups. *PNAS*, 108(13):5419–5424.
- Petersen, R. S., Brambilla, M., Bale, M. R., Alenda, A., Panzeri, S., Montemurro, M. A., and Maravall, M. (2008). Diverse and Temporally Precise Kinetic Feature Selectivity in the VPM Thalamic Nucleus. *Neuron*, 60(5):890–903.
- Pettit, D. L., Wang, S. S., Gee, K. R., and Augustine, G. J. (1997). Chemical two-photon uncaging: A novel approach to mapping glutamate receptors. *Neuron*, 19(3):465–471.
- Peyrache, A., Benchenane, K., Khamassi, M., Wiener, S. I., and Battaglia, F. P. (2010). Principal component analysis of ensemble recordings reveals cell assemblies at high temporal resolution. *Journal of Computational Neuroscience*, 29:309–325.
- Pfister, J.-P. and Gerstner, W. (2006). Triplets of Spikes in a Model of Spike Timing-Dependent Plasticity. *Journal of Neuroscience*, 26(38):9673–9682.
- Pi, H. J., Hangya, B., Kvitsiani, D., Sanders, J. I., Huang, Z. J., and Kepecs, A. (2013). Cortical interneurons that specialize in disinhibitory control. *Nature*, 503(7477):521–524.
- Poirazi, P., Brannon, T., and Mel, B. W. (2003). Pyramidal Neuron as Two-Layer Neural Network. *Neuron*, 37:989–999.
- Pouille, F., Marin-Burgin, A., Adesnik, H., Atallah, B. V., and Scanziani, M. (2009). Input normalization by global feedforward inhibition expands cortical dynamic range. *Nature Neuroscience*, 12(12):1577–1585.
- Pouille, F. and Scanziani, M. (2001). Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science*, 293(5532):1159–1163.
- Pouille, F. and Scanziani, M. (2004). Routing of spike series by dynamic circuits in the hippocampus. *Nature*, 429(6993):717–723.
- Price, C. J., Cauli, B., Kovács, E. R., Kukik, Á., Lambolez, B., Shigemoto, R., and Capogna, M. (2005). Neurogliaform Neurons Form a Novel Inhibitory Network in the Hippocampal CA1 Area. *Journal of Neuroscience*, 25(29):6775–6786.
- Price, C. J., Scott, R., Rusakov, D. A., and Capogna, M. (2008). GABAB Receptor Modulation of Feedforward Inhibition through Hippocampal Neurogliaform Cells. *The Journal of Neuroscience*, 28(27):6974–6982.
- Qi, G. and Feldmeyer, D. (2016). Dendritic Target Region-Specific Formation of Synapses between Excitatory Layer 4 Neurons and Layer 6 Pyramidal Cells. *Cerebral Cortex*, 26(4):1569–1579.
- Ramaswamy, S., Colangelo, C., and Markram, H. (2018). Data-Driven Modeling of Cholinergic Modulation of Neural Microcircuits: Bridging Neurons, Synapses and Network Activity. *Frontiers in Neural Circuits*, 12(77).
- Ramaswamy, S., Courcol, J.-D., et al. (2015). The neocortical microcircuit collaboration portal: a resource for rat somatosensory cortex. *Frontiers in Neural Circuits*, 9(44).

- Ramaswamy, S., Hill, S. L., King, J. G., Schürmann, F., Wang, Y., and Markram, H. (2012). Intrinsic morphological diversity of thick-tufted layer 5 pyramidal neurons ensures robust and invariant properties of in silico synaptic connections. *Journal of Physiology*, 590(4):737–752.
- Ratcliff, R. (1990). Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review*, 97(2):285–308.
- Ray, A., Christian, J. A., Mosso, M. B., Park, E., Wegner, W., Willig, K. I., and Barth, A. L. (2023). Quantitative Fluorescence Analysis Reveals Dendrite-Specific Thalamocortical Plasticity in L5 Pyramidal Neurons during Learning. *Journal of Neuroscience*, 43(4):584–600.
- Reimann, M. W., Bolaños-Puchet, S., Courcol, J.-D., Egas Santandre, D., et al. (2022a). Modeling and Simulation of Rat Non-Barrel Somatosensory Cortex. Part I: Modeling Anatomy. *bioRxiv*.
- Reimann, M. W., Horlemann, A. L., Ramaswamy, S., Muller, E. B., and Markram, H. (2017a). Morphological diversity strongly constrains synaptic connectivity and plasticity. *Cerebral Cortex*, 27(9):4570–4585.
- Reimann, M. W., King, J. G., Muller, E. B., Ramaswamy, S., and Markram, H. (2015). An algorithm to predict the connectome of neural microcircuits. *Frontiers in computational neuroscience*, 9(120).
- Reimann, M. W., Nolte, M., Scolamiero, M., Turner, K., Perin, R., Chindemi, G., Dłotko, P., Levi, R., Hess, K., and Markram, H. (2017b). Cliques of neurons bound into cavities provide a missing link between structure and function. *Frontiers in Computational Neuroscience*, 11(48).
- Reimann, M. W., Riihimäki, H., Smith, J. P., Lazovskis, J., Pokorny, C., and Levi, R. (2022b). Topology of synaptic connectivity constrains neuronal stimulus representation, predicting two complementary coding strategies. *PLoS ONE*, 17(1):e0261702.
- Reva, M., Rössert, C., Arnaudon, A., Damart, T., Mandge, D., Tuncel, A., Ramaswamy, S., Markram, H., and Werner, V. G. (2022). A universal workflow for creation, validation and generalization of detailed neuronal models. *bioRxiv*.
- Reyes, A. and Sakmann, B. (1999). Developmental switch in the short-term modification of unitary EPSPs evoked in layer 2/3 and layer 5 pyramidal neurons of rat neocortex. *Journal of Neuroscience*, 19(10):3827–3835.
- Reyes-Puerta, V., Sun, J. J., Kim, S., Kilb, W., and Luhmann, H. J. (2015). Laminar and Columnar Structure of Sensory-Evoked Multineuronal Spike Sequences in Adult Rat Barrel Cortex in Vivo. *Cerebral Cortex*, 25(8):2001–2021.
- Rochester, N., Holland, J. H., Haiht, L. H., and L, D. W. (1956). Tests on a cell assembly theory of the action of the brain, using a large digital computer. *IRE Trans. Inf. Theory*, 2:80–93.
- Rodrigues, Y. E., Tigaret, C., Marie, H., O'Donnell, C., and Veltz, R. (2022). A stochastic model of hippocampal synaptic plasticity with geometrical readout of enzyme dynamics. *bioRxiv*.
- Romani, A. et al. (2023). Community-based Reconstruction and Simulation of a Full-scale Model of Region CA1 of Rat Hippocampus. *bioRxiv*.
- Ropireddy, D., Bachus, S. E., and Ascoli, G. A. (2012). Non-homogeneous stereological properties of the rat hippocampus from high-resolution 3D serial reconstruction of thin histological sections. *Neuroscience*, 15(205):91–111.
- Rößler, N., Jungenitz, T., Sigler, A., Bird, A., Mittag, M., Rhee, J. S., Deller, T., Cuntz, H., Brose, N., Schwarzacher, S. W., and Jedlicka, P. (2023). Skewed distribution of spines is independent of presynaptic transmitter release and synaptic plasticity and emerges early during adult neurogenesis. *bioRxiv*.
- Rozov, A., Burnashev, N., Sakmann, B., and Neher, E. (2001). Transmitter release modulation by intracellular Ca²⁺ buffers in facilitating and depressing nerve terminals of pyramidal cells in layer 2/3 of the rat neocortex indicates a target cell-specific difference in presynaptic calcium dynamics. *Journal of Physiology*, 531(3):807–826.

Bibliography

- Rubin, J. E., Gerkin, R. C., Bi, G. Q., and Chow, C. C. (2005). Calcium time course as a signal for spike-timing-dependent plasticity. *Journal of Neurophysiology*, 93(5):2600–2613.
- Rumsey, C. C. and Abbott, L. F. (2006). Synaptic democracy in active dendrites. *Journal of Neurophysiology*, 96(5):2307–2318.
- Sasaki, T., Kimura, R., Tsukamoto, M., Matsuki, N., and Ikegaya, Y. (2006). Integrative spike dynamics of rat CA1 neurons: A multineuronal imaging study. *Journal of Physiology*, 574(1):195–208.
- Schneider-Mizell, C. M. et al. (2023). Cell-type-specific inhibitory circuitry from a connectomic census of mouse visual cortex. *bioRxiv*.
- Schreiber, S., Fellous, J. M., Whitmer, D., Tiesinga, P., and Sejnowski, T. J. (2003). A new correlation-based measure of spike timing reliability. *Neurocomputing*, 52–54:925–931.
- Schwartzkroin, P. A. and Wester, K. (1975). Long-lasting facilitation of a synaptic potential following tetanization in the in vitro hippocampal slice. *Brain Research*, 89:107–119.
- Sejnowski, T. J. (1977). Storing covariance with nonlinearly interacting neurons. *Journal of Mathematical Biology*, 4(4):303–321.
- Selig, D. K., Nicoll, R. A., and Malenka, R. C. (1999). Hippocampal long-term potentiation preserves the fidelity of postsynaptic responses to presynaptic bursts. *Journal of Neuroscience*, 19(4):1236–1246.
- Sermet, B. S., Truschow, P., Feyerabend, M., Mayrhofer, J. M., Oram, T. B., Yizhar, O., Staiger, J. F., and Petersen, C. C. (2019). Pathway-, layer- and cell-type-specific thalamic input to mouse barrel cortex. *eLife*, 8:e52665.
- Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27:623–656.
- Schatz, C. J. (1992). The Developing Brain. *Scientific American*, 267(3):60–67.
- Sik, A., Penttonen, M., Ylinen, A., and Buzsáki, G. (1995). Hippocampal CA1 Interneurons: An in vivo Intracellular Labeling Study. *Journal of Neuroscience*, 10(15):6651–6665.
- Silberberg, G. and Markram, H. (2007). Disynaptic Inhibition between Neocortical Pyramidal Cells Mediated by Martinotti Cells. *Neuron*, 53(5):735–746.
- Sjöström, P. J. and Häusser, M. (2006). A Cooperative Switch Determines the Sign of Synaptic Plasticity in Distal Dendrites of Neocortical Pyramidal Neurons. *Neuron*, 51(2):227–238.
- Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2003). Neocortical LTD via Coincident Activation of Presynaptic NMDA and Cannabinoid Receptors. *Neuron*, 39(4):641–654.
- Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, 3(9):919–926.
- Song, S., Sjöström, P. J., Reigl, M., Nelson, S., and Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biology*, 3(3):e68.
- Spruston, N., Jaffe, D. B., Williams, S. H., and Johnston, D. (1993). Voltage- and space-clamp errors associated with the measurement of electrotonically remote synaptic events. *Journal of Neurophysiology*, 70(2):781–802.
- Spruston, N., Jonas, P., and Sakmann, B. (1995). Dendritic glutamate receptor channels in rat hippocampal CA3 and CA1 pyramidal neurons. *Journal of Physiology*, 482(2):325–352.
- Stimberg, M., Brette, R., and Goodman, D. F. (2019). Brian 2, an intuitive and efficient neural simulator. *eLife*, 8(e47314).
- Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., and Harris, K. D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364:eaav7893.
- Stuart, G., Schiller, J., and Sakmann, B. (1997). Action potential initiation and propagation in rat neocortical pyramidal neurons. *Journal of Physiology*, 505(3):617–632.
- Stuart, G. J. and Sakmann, B. (1994). Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature*, 367(January):69–72.

- Stuart, G. J. and Spruston, N. (2015). Dendritic integration: 60 years of progress. *Nature Neuroscience*, 18(12):1713–1721.
- Sussillo, D., Toyozumi, T., and Maass, W. (2007). Self-tuning of neural circuits through short-term synaptic plasticity. *Journal of Neurophysiology*, 97(6):4079–4095.
- Takács, V. T., Klausberger, T., Somogyi, P., Freund, T. F., and Gulyás, A. I. (2012). Extrinsic and local glutamatergic inputs of the rat hippocampal CA1 area differentially innervate pyramidal cells and interneurons. *Hippocampus*, 22(6):1379–1391.
- Thomson, A. M. (2000). Facilitation, augmentation and potentiation at central synapses. *Trends in Neurosciences*, 23(7):305–312.
- Thomson, A. M., Deuchars, J., and West, D. C. (1993). Large, deep layer pyramid-pyramid single axon EPSPs in slices of rat motor cortex display paired pulse and frequency-dependent depression, mediated presynaptically and self-facilitation, mediated postsynaptically. *Journal of Neurophysiology*, 70(6):2354–2369.
- Thomson, A. M. and West, D. C. (1993). Fluctuations in pyramid-pyramid excitatory postsynaptic potentials modified by presynaptic firing pattern and postsynaptic membrane potential using paired intracellular recordings in rat neocortex. *Neuroscience*, 54(2):329–346.
- Tonegawa, S., Liu, X., Ramirez, S., and Redondo, R. (2015). Memory Engram Cells Have Come of Age. *Neuron*, 87(5):918–931.
- Tong, G. and Jahr, C. E. (1994). Multivesicular release from excitatory synapses of cultured hippocampal neurons. *Neuron*, 12(1):51–59.
- Trägenap, S., Whitney, D. E., Fitzpatrick, D., and Kaschube, M. (2022). Experience drives the development of novel, reliable cortical sensory representations from endogenously structured networks. *bioRxiv*.
- Tremblay, R., Lee, S., and Rudy, B. (2016). GABAergic Interneurons in the Neocortex: From Cellular Properties to Circuits. *Neuron*, 91(2):260–292.
- Tsodyks, M. and Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *PNAS*, 94(2):719–723.
- Tsodyks, M., Uziel, A., and Markram, H. (2000). Synchrony generation in recurrent networks with frequency-dependent synapses. *The Journal of neuroscience*, 20(RC50).
- Turrigiano, G. G. and Nelson, S. B. (2004). Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, 5(2):97–107.
- Tyan, L., Chamberland, S., Magnin, E., Camiré, O., Francavilla, R., Suzanne David, L., Deisseroth, K., and Topolnik, L. (2014). Dendritic inhibition provided by interneuron-specific cells controls the firing rate and timing of the hippocampal feedback inhibitory circuitry. *Journal of Neuroscience*, 34(13):4534–4547.
- Ujfalussy, B. B. and Makara, J. K. (2020). Impact of functional synapse clusters on neuronal response selectivity. *Nature Communications*, 11(1413).
- van der Plas, T. L., Tübiana, J., Le Goc, G., Migault, G., Kunst, M., Baier, H., Bormuth, V., Englitz, B., and Debrégeas, G. (2023). Neural assemblies uncovered by generative modeling explain whole-brain activity statistics and reflect structural connectivity. *eLife*, 11:e83139.
- Van Geit, W., Gevaert, M., Chindemi, G., Rössert, C., Courcol, J.-D., Müller, E., Schürmann, F., Segev, I., and Markram, H. (2016). BluePyOpt: Leveraging open source software and cloud infrastructure to optimise model parameters in neuroscience. *Frontiers in Neuroinformatics*, 10(17).
- Varani, S., Vecchia, D., Zucca, S., Forli, A., and Fellin, T. (2022). Stimulus Feature-Specific Control of Layer 2 / 3 Subthreshold Whisker Responses by Layer 4 in the Mouse Primary Somatosensory Cortex. *Cerebral Cortex*, 32(7):1419–1436.

Bibliography

- Vargas-Caballero, M. and Robinson, H. P. (2003). A slow fraction of Mg²⁺ unblock of NMDA receptors limits their contribution to spike generation in cortical pyramidal neurons. *Journal of Neurophysiology*, 89(5):2778–2783.
- Vickers, E. D., Clark, C., Osypenko, D., Fratzl, A., Kochubey, O., Bettler, B., and Schneggenburger, R. (2018). Parvalbumin-Interneuron Output Synapses Show Spike-Timing-Dependent Plasticity that Contributes to Auditory Map Remodeling. *Neuron*, 99(4):720–735.
- Vida, I., Halasy, K., Szinyei, C., Somogyi, P., and Buhl, E. H. (1998). Unitary IPSPs evoked by interneurons at the stratum radiatum — stratum lacunosum-moleculare border in the CA1. *Journal of Physiology*, 506(3):755–773.
- Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks. *Science*, 334(6062):1569–1573.
- Wang, Y., Gupta, A., Toledo-Rodriguez, M., Wu, C. Z., and Markram, H. (2002). Anatomical, physiological, molecular and circuit properties of nest basket cells in the developing somatosensory cortex. *Cerebral cortex*, 12(4):395–410.
- Wheeler, D. W., White, C. M., Rees, C. L., Komendantov, A. O., Hamilton, D. J., and Ascoli, G. A. (2015). Hippocampome.org: A knowledge base of neuron types in the rodent hippocampus. *eLife*, 4(e09960).
- Widrich, M., Schäfl, B., Pavlovic, M., Ramsauer, H., Gruber, L., Holzleitner, M., Brandstetter, J., Sandve, G. K., Greiff, V., Hochreiter, S., and Klambauer, G. (2020). Modern hopfield networks and attention for immune repertoire classification. *Advances in Neural Information Processing Systems*, 33.
- Williams, L. E. and Holtmaat, A. (2019). Higher-Order Thalamocortical Inputs Gate Synaptic Long-Term Potentiation via Disinhibition. *Neuron*, 101(1):91–102.
- Williams, S. R. and Atkinson, S. E. (2007). Pathway-specific use-dependent dynamics of excitatory synaptic transmission in rat intracortical circuits. *Journal of Physiology*, 585(3):759–777.
- Williams, S. R. and Mitchell, S. J. (2008). Direct measurement of somatic voltage clamp errors in central neurons. *Nature Neuroscience*, 11(7):790–798.
- Williams, S. R. and Stuart, G. J. (1999). Mechanisms and consequences of action potential burst firing in rat neocortical pyramidal neurons. *Journal of Physiology*, 521(2):467–482.
- Williams, S. R. and Stuart, G. J. (2002). Dependence of EPSP efficacy on synapse location in neocortical pyramidal neurons. *Science*, 295(5561):1907–1910.
- Wilson, D. E., Whitney, D. E., Scholl, B., and Fitzpatrick, D. (2016). Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex. *Nature Neuroscience*, 19(8):1003–1009.
- Wohrer, A., Humphries, M. D., and Machens, C. K. (2013). Population-wide distributions of neural activity during perceptual decision-making. *Progress in Neurobiology*, 103:156–193.
- Wozny, C. and Williams, S. R. (2011). Specificity of synaptic connectivity between layer 1 inhibitory interneurons and layer 2/3 pyramidal neurons in the rat neocortex. *Cerebral Cortex*, 21(8):1818–1826.
- Xiao, K., Li, Y., Chitwood, R. A., and Magee, J. C. (2023). A critical role for CaMKII in behavioral timescale synaptic plasticity in hippocampal CA1 pyramidal neurons. *bioRxiv*.
- Yang, D., Günter, R., Qi, G., Radnikow, G., and Feldmeyer, D. (2020). Muscarinic and Nicotinic Modulation of Neocortical Layer 6A Synaptic Microcircuits Is Cooperative and Cell-Specific. *Cerebral Cortex*, 30(6):3528–3542.
- Yang, D., Qi, G., Ding, C., and Feldmeyer, D. (2022). Layer 6A Pyramidal Cell Subtypes Form Synaptic Microcircuits with Distinct Functional and Structural Properties. *Cerebral Cortex*, 32(10):2095–2111.
- Yasui, T., Fujisawa, S., Tsukamoto, M., Matsuki, N., and Ikegaya, Y. (2005). Dynamic synapses as archives of synaptic history: State-dependent redistribution of synaptic efficacy in the rat hippocampal CA1. *Journal of Physiology*, 566(1):143–160.

- Yu, J., Hu, H., Agmon, A., and Svoboda, K. (2019). Recruitment of GABAergic Interneurons in the Barrel Cortex during Active Tactile Behavior. *Neuron*, 104(2):412–427.
- Zenke, F., Agnes, E. J., and Gerstner, W. (2015). Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nature Communications*, 6(6922).
- Zenke, F., Gerstner, W., and Ganguli, S. (2017a). The temporal paradox of Hebbian learning and homeostatic plasticity. *Current Opinion in Neurobiology*, 43:166–176.
- Zenke, F., Poole, B., and Ganguli, S. (2017b). Continual learning through synaptic intelligence. *ICML*.
- Zhu, J. J. (2000). Maturation of layer 5 neocortical pyramidal neurons: amplifying salient layer 1 and layer 4 inputs by Ca²⁺ action potentials in adult rat tuft dendrites. *Journal of Physiology*, 526(3):571–587.
- Zucker, R. S. and Regehr, W. G. (2002). Short-term synaptic plasticity. *Annual Review of Physiology*, 64:355–405.

András Ecker

Curriculum Vitæ

Bajcsy-Zsilinszky u. 32/2

Kőszeg 9730, Hungary

✉ andras.ecker@epfl.ch

📧 [andrisecker](#) [in](#) [András Ecker](#)
🐦 [andris_ecker](#) [s](#) [András Ecker](#)

Education

- 2018.10. - present **EPFL PhD Program in Neuroscience (EDNE)**, *Blue Brain Project*, EPFL (École polytechnique fédérale de Lausanne), Switzerland.
- IN SILICO NEUROSCIENCE PhD supervised by Henry Markram and Michael W. Reimann
 - *"Principles of Network Plasticity in Neocortical Microcircuits"*
 - Tutor at the EITN Computational Neuroscience Summer School in 2021 and 2022
 - Teaching assistant in Numerical analysis in 2020 and 2021
- 2016.09. - 2018.08. **Life Science and Technology MSc**, *EPFL (École polytechnique fédérale de Lausanne) SV Faculty*, Switzerland, GPA: 5.42/6.
- Minor in COMPUTATIONAL NEUROSCIENCE
 - Master research scholar of the faculty (2 days/week internship)
- 2012.09. - 2016.01. **Molecular Bionics Engineering BSc**, *Pázmány Péter Catholic University, Faculty of Information Technology and Bionics*, Hungary, GPA: 4.87/5.
- Teaching assistant in Calculus
 - Wrote official lecture notes for Neurobiology I and II
 - Member of the student association responsible for social events

Research/Work Experience

- 2016.09. - 2018.08. **Blue Brain Project**, *EPFL (École polytechnique fédérale de Lausanne)*, Switzerland.
- IN SILICO NEUROSCIENCE internship in Eilif B. Muller's group
- 2018.01. - 03. **GE Healthcare**, *Hungary*.
- DATA SCIENCE industrial intership with Levente Török
- 2017.05. - 08. **Google Summer of Code**, *INCF*, Sweden (remote working).
- *"Conversion of a large scale hippocampal network model to NeuroML"*
- 2016.03. - 07. **University College London**, *UK*.
- NEUROINFORMATICS intership with Pdraig Gleeson in Angus Silver's laboratory
- 2015.07. - 09. **Bernstein Center for Computational Neuroscience**, *Germany*.
- COMPUTATIONAL NEUROSCIENCE Amgen scholar in Thomas Wachtler's laboratory
- 2015.01. - 2016.01. **Institute of Experimental Medicine HAS**, *Hungary*.
- COMPUTATIONAL NEUROSCIENCE intership with Szabolcs Káli in Tamás Freud's laboratory

Preprints

- 2023 **A. Ecker**, D. Egas Santander, et al. *Long-term plasticity induces sparse and specific synaptic changes in a biophysically detailed cortical model*, <https://doi.org/10.1101/2023.08.07.552264>
- 2023 **A. Ecker**, D. Egas Santander, et al. *Cortical cell assemblies and their underlying connectivity: an in silico study*, <https://doi.org/10.1101/2023.02.24.529863>
- 2023 J.B. Isbister, **A. Ecker**, C. Pokorny, S. Bolaños-Puchet, D. Egas Santander, et al. *Modeling and Simulation of Neocortical Micro- and Mesocircuitry. Part II: Physiology and Experimentation*, <https://doi.org/10.1101/2023.05.17.541168>
- 2023 M.W. Reimann, S. Bolaños-Puchet, J-D. Courcol, D. Egas Santander, et al. *Modeling and Simulation of Neocortical Micro- and Mesocircuitry. Part I: Anatomy*, <https://doi.org/10.1101/2022.08.11.503144>
- 2023 A. Romani, et al. *Community-based Reconstruction and Simulation of a Full-scale Model of Region CA1 of Rat Hippocampus*, <https://doi.org/10.1101/2023.05.17.541167>

Publications

- 2023 E. Iavorne, et al. *Thalamic control of sensory processing and spindles in a biophysical somatosensory thalamoreticular circuit model of wakefulness and sleep*, <https://doi.org/10.1016/j.celrep.2023.112200>, Cell Reports, 42(3), 2023
- 2022 A. Ecker, et al. *"Hippocampal sharp wave-ripples and the associated sequence replay emerge from structured synaptic interactions in a network model of area CA3"*, <https://doi.org/10.7554/eLife.71850.sa0>, eLife, 11:e71850, 2022
- 2022 G. Chindemi, et al. *A calcium-based plasticity model predicts long-term potentiation and depression in the neocortex* <https://doi.org/10.1038/s41467-022-30214-w>, Nature Communications, 13(3038), 2022
- 2020 A. Ecker, et al. *"Data-driven integration of hippocampal CA1 synaptic physiology in silico"*, <https://doi.org/10.1002/hipo.23220>, Hippocampus, 30(11):1129-1145, 2020
- 2019 P. Gleeson, et al. *"Open Source Brain: a collaborative resource for visualizing, analyzing, simulating and developing standardized models of neurons and circuits"* <https://doi.org/10.1016/j.neuron.2019.05.019>, Neuron, 103(3):395-411, 2019

Awards and Scholarships

- 2018 Best master's thesis of the section (EPFL)
- 2017 GOOGLE SUMMER OF CODE
- 2017 Hungarian Scientific Students' Association Conference: 2nd place
- 2016 Honorous BSc degree (PPCU)
- 2015 Fellowship granted by the (Hungarian) Republic
- 2015 ERASMUS+ scholarship
- 2015 AMGEN SCHOLAR PROGRAM

Computer Skills

Programming languages: Python, MATLAB, (C++, SQL)

Neural simulators: NEURON, Brian, LEMS-NeuroML

Other skills: Adobe Illustrator, PhotoShop, InDesign; Unix, Git

Languages

Hungarian: mother tongue

English: fluent (C2)

German: intermediate (B2)

French: active learner (A1)