

Modeling the impact of protein statistics on spontaneous and plasticity-induced synaptic dynamics

Dissertation
zur
Erlangung des Doktorgrades (Dr. rer nat.)
der
Mathematisch-Naturwissenschaftlichen Fakultät
der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von
Janko Petkovic
aus Belgrad, Serbien

Bonn, 2025

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

Gutachterin/Betreuerin: Prof. Dr. Tatjana Tchumatchenko

Gutachter: Prof. Dr. Jan Hasenauer

Tag der Promotion: 19.12.2025

Erscheinungsjahr: 2026

Table of contents

Summary	7
Acknowledgements	11
1. Introduction	13
1.1. The experimental foundations of neuroscience	13
1.2. Synaptic plasticity: a possible implementation of memory and learning	15
1.3. Molecular basis of synaptic plasticity	18
1.4. Research goal	22
1.5. Contribution statement	24
2. Stochastic modeling of synaptic fluctuations	25
2.1. State of the art and open challenges	25
2.2. Preliminary experimental observations	27
2.3. Designing the model	30
2.3.1. Bounded Wiener, and Ornstein-Uhlenbeck processes	30
2.3.2. Negative momentum and non-Markovianity	32
2.3.3. The LN-OU Model	34
2.4. The impact of long-term potentiation on synaptic distributions	36
2.5. Extending the LN-OU model to stimulated spines	39
2.6. Summary and remarks	42
3. Modeling the impact of protein distributions on multi-spine synaptic plasticity .	46
3.1. State of the art and open challenges	46
3.2. The optimality of the multi-spine plasticity profile	48
3.3. Designing the model	50
3.3.1. Quasi-steady-state approximation and closed form solution	53
3.4. Optimization	57
3.4.1. Structural identifiability	57
3.4.2. Definition and pre-processing of the target values	60
3.4.3. Likelihood maximization	61
3.4.4. Goodness of fit	62
3.5. Inferred and observed log-normality of synaptic features	64
3.6. Log-normality emergence from molecular dynamics	66

3.7. The effect of inter-stimulus distance on the plasticity profile	73
3.8. Parameter degeneracy and bidirectional effects of FK506	76
3.9. Variability of synaptic response and effect of basal synaptic size	79
3.10. Summary and remarks	84
4. Conclusion	88
4.1. Log-normal characteristics of Synaptic Fluctuations	88
4.2. Possible origin of log-normality	89
4.3. Size-dependent plasticity mechanisms	89
4.4. Catalytic degeneracy in potentiation mechanisms	89
4.5. Future Directions	90
Bibliography	93
Appendices	105
Appendix A: Shifted log-normal distribution	105
Appendix B: Inter-spine distance estimation	107
Appendix C: FK506 induced parameter changes	109
Supplementary figures	111
List of peer-reviewed publications	117

Summary

The fundamental question of how the brain encodes, processes, and retrieves stored information remains a central, unresolved challenge in neuroscientific research. This question holds great importance for a number of reasons, from the purest philosophical and scientific charm, to the plethora of technological applications its answer could bring. Maybe one of the most noteworthy of such applications concerns the medical field, where knowledge about behaviour and neuronal functioning *in toto*, represents the prime scientific instrument towards treatment. Understanding on a mechanistic, molecular level the specific role a certain biochemical factor plays, is the first step towards understanding how to control this factor and ultimately, how to correct for its malfunction.

Synaptic plasticity is regarded as one of the central mechanisms responsible for neuronal information encoding and, consequently, the emergence of anything that can be considered repeatable and investigable behaviour. The degree to which a presynaptic neuron can excite its postsynaptic partner upon spike arrival determines the synapse’s “strength”. Synaptic strength is governed by an elaborate interplay of biochemical components, and it exhibits a wide range of dynamical changes both across different spines and over time — a possible signature of memory formation. One of the most significant neuroscientific breakthroughs of the past century was the recognition that alterations in the strength of a synapse are not predetermined but instead depend critically on the activity patterns of the neurons it connects.

To understand the causes and implications of this mechanism, a number of experimental and theoretical studies have been conducted in the last decades. This body of work encompasses a variety of approaches, from normative principles proposing information-theoretic optimality frameworks, to ad-hoc mechanistic models focusing on specific plasticity features, to increasingly sophisticated descriptive experiments. This research has provided remarkable insights into the phenomenon, slowly but steadily leading to more robust connections between theoretical principles and observations. Nevertheless, despite these advances, the precise mechanisms underlying synaptic plasticity remain poorly understood and continue to represent one of the field’s most challenging questions.

A particularly challenging aspect concerns how exactly the spatio-temporal dynamics of synaptic plasticity emerge from the thousands of molecular components that regulate it. Resolving this mechanistic puzzle is essential, as it would ultimately pave the way to a consistent theory bridging the microscopic, causally sound, biochemical level to the macroscopic

(potentially optimal?) emergent level, providing a crucial key to understand the “rules” governing behaviour, pathology, and, potentially, treatment. This work addresses this formidable question by harnessing the capabilities of dynamical modeling and optimization techniques, proposing and examining various mathematical frameworks that enable not only a precise description of recent experimental observations but also allow for novel interpretations of previous, apparently contradictory findings.

This thesis is organized as follows:

Chapter 1 provides a comprehensive overview of synaptic plasticity. Following a brief historical introduction – necessary for understanding our current position within the neuroscientific landscape – I will introduce the concept of synaptic plasticity, focusing on its strong compatibility with a biological basis of memory formation and learning. Subsequently, I will present the primary biochemical factors driving synaptic plasticity, with brief discussions of their individual functions and mutual interactions. Finally, I will conclude the chapter with the specific scientific questions that this thesis aims to answer.

In Chapter 2, I introduce my first original contribution, which consists of a stochastic descriptive model of synaptic size statistics. Building upon experimental observations, a model describing synaptic fluctuations is constructed from fundamental principles and developed incrementally in close adherence with the existing literature. After validation, the model is then utilized to identify possible “governing principles” of synaptic fluctuations, such as the tendency for large spines to shrink and small spines to grow. The focus then shifts to examining how synaptic stimulation affects the synaptic ensemble dynamics, utilizing an appropriately extended version of the baseline model. Lastly, the model is used to put forward possible optimality principles driving synaptic response to potentiation, focusing, in particular, on the effects of synaptic simulation on the synaptic population size entropy.

In Chapter 3, I present my second and main original contribution, consisting of another modeling endeavour formulated within a different mathematical framework: reaction-diffusion dynamics. Starting from the data kindly provided by our experimental collaborators (T.E. Chater, Y. Goda), a reaction-diffusion model is introduced, describing synaptic dynamics in relation to two main molecular processes: protein diffusion and phosphorylation. The resulting equations, which permit a remarkably intuitive closed-form solution, are subsequently fitted to the experimental data using a novel optimization toolkit (PyPesto), in close collaboration with the Hasenauer Group for data-driven inference (specifically E. Raimundez and D. Pathirana). Synaptic plasticity is then analyzed through the statistical and biochemical lens,

yielding surprising insights into how initial synaptic statistics, and in particular their log-normal distribution, emerge and influence the outcome of a plasticity experiment. The model's adherence to the underlying biochemistry also enables the formulation of a novel interpretation of FK506's bidirectional effects on synaptic potentiation, an observation reported across multiple experimental studies but not yet fully explained.

The thesis concludes with Chapter 4, which presents a comprehensive summary of results alongside the primary research directions that could extend these findings in future investigations, subdivided into broadly two main categories. From a theoretical standpoint, possible generalizations of the proposed models are considered, leading, among other options, to a "combined model" able to reconcile in a joint fashion the spontaneous, fast ensemble fluctuations with the synapse-specific, slower plasticity dynamics induced by the stimulus. From a methodological standpoint, alternative optimization choices are considered, focusing in particular on Bayesian hierarchical optimization - one of the most natural frameworks for a wide class of biological problems. A final remark ends the thesis, examining the possible applications of this work in terms of basic science and, crucially, clinical and therapeutic research.

Acknowledgements

As most of the significant achievements that one has the fortune to achieve throughout their career, and more importantly, life, this Ph.D. is far from being the result of the exclusive efforts of a single person, and truly, many are the people that I am grateful to for helping and supporting me throughout these three years. To some, I can express this gratitude here, but many others I will not report, promising that I will find another way to give them my thanks.

First and foremost, I want to express my deepest gratitude to my supervisor, Prof. Dr. Tatjana Tchumatchenko, whose experience and wisdom both in and beyond the neuroscience field have taught me a lot about what being a researcher, an academic and a group leader is. For your availability, your kindness, and your eagerness to push your students to travel, meet, and start shaping their own scientific identity, my most sincere thanks.

I would then like to express my gratitude to my second supervisor, Prof. Dr. Jan Hasenauer, for the invaluable scientific assistance he provided concerning my main doctoral project, and which, beyond any doubt, was among the decisive factors of its final success. I would also like to extend this gratitude to two members of his lab in particular, Dr. E. Raimundez, for her openness and her enthusiasm to assist, and pass on her expertise, and D. Pathirana, who's strikingly sharp and curious questions, and principled approach to any field of discussion, I will carry with me throughout my career as great inspiration.

I would also like to thank Prof. Dr. Haruhiko Bito for including my research in what he considers “the science” that “needs time”, even if this time is taken from late-night hours.

Science is but a part of a researcher's job, and we could not devote to it without the invaluable support of a group of inspired and competent people dedicating their time and efforts to taking care of the inevitable logistical and organisational burden that accompanies every research lab. In this regard, I want to express my deepest thanks to the secretaries of the Tchumatchenko Lab, who, through the years, have enabled us to do our job without having to worry (too much) about the sometimes overwhelming bureaucratic tempests. To Sabine Eichholz, for going above and beyond to support us, and trying her best to make our group a harmonious, familiar and happy environment in which everyone would feel welcome and considered; anything I could write would not come close to acknowledging the merits she has. To Sascha, who, despite the darkness that regrettably life carries with itself, has always been ready to offer his smile, as well as his kind and reassuring help, inside and outside of work.

I want to thank the members of the lab I had the luck to spend my doctoral studies in. Beyond thanking, I want to remark how invaluable and irreplaceable they were, not only as brilliant researchers with which science could be discussed creatively and insightfully, but also as companions in a life experience that is not always free of challenges. To Surbhit, Ulzii, Lorenzo, Kanaan, Pietro, Icaro, Dario, I can only say that it has been a true joy, and a fortune, for me, to spend these years in your company, and that I hope I will be able to enjoy this pleasure in the future.

Finally, I want to thank my family, for everything. To my sister, to her courage and determination, I dedicate this thesis.

1. Introduction

1.1. The experimental foundations of neuroscience

“How the brain works” has been one of the central questions for humanity over millennia. The answer to this question, representing not only a scientific endeavor but also a medical and, maybe even more significantly, a religious one, has been dramatically shaped by shifts in philosophical thought, technological innovation, and technical discovery. Ancient civilizations such as the Egyptians recognized the brain’s existence, as well as its central connection to some of the symptoms arising from head trauma (e.g., aphasia and seizure, 17th century BC). The connection between brain and higher-order behavioral functions, however, was hypothesized only around the 5th century BC, with Alcmaeon of Croton and Hippocrates of Kos being among the first to propose that the brain was involved in vision, sensation, and intelligence, giving rise to the understanding that separated humans from animals. Remarkably, during the Hellenistic period, Herophilus and Erasistratus of Alexandria (among others) engaged in studies that involved dissecting human bodies, providing evidence for the primacy of the brain. They affirmed the distinction between the cerebrum and the cerebellum, and identified the ventricles and the dura mater. Their works are now mostly lost, and some of their discoveries had to be rediscovered a millennium after their death.

In the following centuries, one of the main drivers of neuroscientific research was Islamic medicine, led by brilliant pioneers such as Ibn Sina (latinized Avicenna). With striking modernity, “mental health” became, for the first time, a subject of study, seeding neuroanatomical investigation (e.g., the cerebellar vermis and the caudate nucleus) and most strikingly, structural hypotheses about the connection between brain and behaviour (mental deficits as a result of frontal lobe alterations). In the European continent, this remarkable push towards discovery arrived with the Scientific Revolution, as a direct consequence of the methodological advances (generally referred to as the “Scientific Method”) pioneered by, among others, G. Galilei and F. Bacon. The role of electricity in nerves was first observed in dissected frogs by Luigi Galvani, Lucia Galeazzi Galvani, and Giovanni Aldini in the second half of the 18th century. In the same period, dissection studies were able for the first time to link a specific nervous region to a specific behavioural function, identifying the medulla oblongata as the respiratory center, as well as characterizing the effects of different encephalic and medullary lesions on motoricity, sensibility, and behavior. Other pioneering observations were conducted in the next century, with an increasingly quantitative characterization of medical manifestations. This is the case

1. Introduction

for the role of the frontal lobe in behavior (the case of Phineas Gage studied by J. Harlow in 1848), as well as the speech-related functions of the Broca and Wernicke's motor and temporal regions (P. Broca 1861, and C. Wernicke 1875). Ultimately, in 1909, K. Brodmann published a full cytoarchitectural characterization of the human brain, defining the 52 Brodmann areas still used in clinical and theoretical neuroscience to this day.

This characterization, representing the first instance of its kind, was one of the pinnacles of the technical revolution that involved microscopy, the prime instrument guiding histological research and that, with the introduction of controlled electrical lighting (in particular the Köhler lighting) was for the first time able to achieve resolution compatible with the theoretical limits of light microscopy. This technical advancement could be considered the prime instrumental factor allowing neuroscientific research to move from a macroscopic, behavioral and correlational level to a microscopic, structural one.

The refinement of the microscope and the development of staining procedures (C. Golgi, S. Ramón y Cajal) led to the formation of the neuron doctrine, the hypothesis that the functional unit of the brain is the neuron. Golgi and Ramón y Cajal shared the Nobel Prize in Physiology or Medicine in 1906 for their extensive observations, descriptions, and categorizations of neurons throughout the brain. The hallmark of the 20th century, however, can be considered the exponential development of novel technological solutions. The invention of the electroencephalogram (EEG) by Hans Berger in 1929 enabled scientists to measure the brain's electrical activity in real time. The development of imaging technologies, such as magnetic resonance imaging (MRI) and positron emission tomography (PET), revolutionized the ability to study the living brain's structure and function. The groundbreaking introduction of electron microscopy (attributed among others to the work of Reinhold Rüdenberg, and, independently, Max Knoll and Ernst Ruska, 1931) allowed for the first time the description and characterization of subcellular structures, while immunostaining (first described by A. Coons in 1941) allowed researchers to specifically label (i.e., "color") these structures with unprecedented molecular precision. A number of techniques sprouted from these methods, leveraging the fundamental biological and, especially, biochemical discoveries of the same century, and combining them into the methodological state of the art that can be observed today, where observation of the specimen can be coupled with direct, pinpoint modification (optogenetics [1], CRISPR-Cas9 [2], among others).

1.2. Synaptic plasticity: a possible implementation of memory and learning

The initial interest in synaptic plasticity emerged from a fundamental constraint in neuroscience: the apparent stability of the number of total neurons throughout adult life. In 1894, the Spanish neuroanatomist Santiago Ramón y Cajal confronted this puzzle and proposed a revolutionary hypothesis. Given that the number of neurons in the brain remained relatively constant after development, Cajal reasoned that memories must be formed through the strengthening of existing neuronal connections rather than through the generation of new neurons. This insight laid the groundwork for what would become one of the most influential concepts in neuroscience.

Cajal's farsighted hypothesis gained substantial theoretical support fifty years later when Donald O. Hebb formalized this learning mechanism in his seminal work "The Organization of Behavior" [3]. Hebb's postulate provided a clear mechanistic framework for synaptic modification: "When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased." This principle, now known as Hebb's rule, suggested that activities that are causally related should be reinforced through synaptic strengthening. The implications of Hebb's postulate extended far beyond simple pairwise connections. In recurrent neural networks, where neurons are extensively interconnected, the consistent application of Hebbian learning could theoretically lead to the emergence of stable cell assemblies. These assemblies would represent coherent patterns of neural activity that could persist over time, potentially serving as the neural substrate for memories and learned behaviors.

Despite the elegance of the Hebbian framework, experimental validation proved extraordinarily challenging. The difficulties were multiple, and intrinsic to the nature of neural systems. First, the experimental resolution required to confirm synaptic strengthening between specific neurons demanded technological capabilities that were not available in the early-mid-20th century. Second, the timescales over which neuronal connections needed to remain stable to support long-term memory were extremely long, often corresponding to the entire lifespan of an organism. Third, with trillions of neurons in complex brains, identifying the specific connections responsible for particular memories seemed an insurmountable task. Beyond these practical obstacles lay deeper theoretical issues concerning the very nature of synaptic connections and information processing in *spiking* neurons. The spiking nature of neural activity — which likely first emerged during the Ediacaran period in association with the evolution of

1. Introduction

predatory behavior — presents itself computational and theoretical challenges. Unlike systems with continuous responses, neurons generate discrete “spikes,” creating instantaneous and transitory changes in intracellular voltage with specific spectral properties. This discontinuous nature of neural signaling made it difficult to develop a rigorous definition of what “strengthening” a connection between two neurons means.

The technological advances achieved in the second half of the 20th century finally made it possible to design experiments capable of observing signs of synaptic plasticity in simpler organisms. A crucial breakthrough came from the studies of E. Kandel on *Aplysia californica*, a giant sea slug that later became a model organism for learning and memory. Researchers demonstrated that as the animal learned to withdraw its gills in response to noxious stimuli, the strength of the synapses involved in this defensive behavior became stronger [4]. Synaptic strength was operationally defined as the slope of the induced post-synaptic action potential, providing a quantifiable measure of connection efficacy. Conversely, as the *Aplysia* became desensitized to innocuous stimuli through repeated exposure, those same synapses became weaker.

A pivotal moment in the field came in 1973 when Terje Lømo and Timothy Bliss published groundbreaking research revealing long-lasting changes in synaptic transmission [5,6]. Their study showed that high-frequency electrical stimulation — termed tetanic bursting — of the perforant path in the hippocampus of anesthetized rabbits led to sustained enhancement in synaptic strength, as measured by recordings of field potentials. In the same issue of *The Journal of Physiology*, Bliss and Tony Gardner - Medwin reported analogous findings from recordings in freely behaving, chronically implanted rabbits, demonstrating that this phenomenon was not an artifact of anesthesia. These seminal papers outlined a form of synaptic potentiation that was both activity-dependent and remarkably persistent, lasting up to three days after the plasticity-inducing stimulation protocol. The phenomenon was initially termed “long-lasting potentiation,” but Douglas and Goddard soon coined the now-famous appellation “long-term potentiation” (LTP), which became the standard terminology for this fundamental form of synaptic plasticity [7].

In subsequent years, researchers not only further characterized LTP but also its counterpart: (long term) synaptic depression, or LTD. In various experimental contexts, they observed that dendritic spines could undergo depression both homosynaptically (at the stimulated spine itself) and heterosynaptically (at synapses in the neighborhood of the stimulated one). This led to the development of a general descriptive framework that broadly characterized poten-

1.2. Synaptic plasticity: a possible implementation of memory and learning

itation as the synaptic response to strong input — such as when a neuron represents the next encoding step downstream from an active neuron — and depression as the response to weak input or, in the heterosynaptic case, to the echo of input received from a neighboring active synapse. The apparent property of heterosynaptic depression to conserve some form of “total” synaptic weight, or “mass”, sparked investigations into “homeostatic plasticity”. This concept proposed that a general activity invariance principle might prompt neurons to preserve their global activity levels by coupling homosynaptic potentiation with heterosynaptic depression, potentially maintaining the network stability while allowing for selective strengthening of specific connections.

The confirmation of synaptic plasticity as a strong biological correlate of behavioral learning provided a candidate mechanism for understanding how memories are stored in neural networks. However, this framing revealed the inherent complexity of the question itself, as the very definition of “memory” required careful consideration. If memory is operationally defined as an organism’s ability to consistently reproduce an investigable behavior, then the “how” of memory storage becomes closely intertwined with the “what” a memory is. From a reductionist perspective, a neuronal system can be conceptualized as a deterministic recurrent network—essentially a graph—with at most a certain degree of uncertainty (noise) in the activity of its nodes. Under this framework, a specific configuration of synaptic weights uniquely defines the average activity landscape and, consequently, the observed behavior in response to particular inputs. This raises a fundamental question: what is the fundamental difference between a “random” weight configuration from one that encodes a specific memory?

The condition of behavioral reproducibility provides one answer to this question. A weight configuration pertaining to an encoded memory must reliably produce the same output given the same input. Formally, this requirement can be understood as the necessity for the weight configuration to support attractors in the network’s configuration space. If these attractors — rather than specific solutions — can be mapped to recognizable behaviors, they could be experimentally interpreted as memories. This minimal, arguably Occamian definition has generated an impressive corpus of work in recent decades, in line with the computational revolution of the 21st century. Researchers have demonstrated how *attractor network* dynamics in various settings can be recognized in neuronal recordings, laying a potential foundation for a theoretical understanding of biological learning and information processing.

The extreme simplification that provides the generality and power of the attractor network approach also constitutes its main limitation. After understanding the mathematical structure

1. Introduction

that the biological brain hypothetically implements, researchers are confronted with another “how”: how does the unfathomably complex network of biological and biochemical events taking place in each neuron — unfolding over multiple spatial and temporal scales — lead to mean behavior compatible with abstract network dynamics? If synaptic weights represent the “how” at the network level, then the molecular factors and their interactions driving synaptic plasticity constitute the “how” at the level immediately below. This multi-scale nature of neural function — from molecules to networks to behavior — constitutes one of the most significant challenges in contemporary neuroscience, requiring integration across vastly different levels of biological organization to achieve a complete understanding of memory formation and storage.

1.3. Molecular basis of synaptic plasticity

To understand the molecular basis of synaptic plasticity, it is good to begin with the definition of synaptic strength: the slope of the excitatory post-synaptic potential (EPSP) measured at the time of pre-synaptic spike arrival. Changing synaptic strength, or equivalently the weight of the synapse, involves modifying the amount of ionic current that enters the post-synaptic side of the neuronal synapse following neurotransmitter release from the pre-synaptic terminal upon arrival of an action potential.

In the following discussion, I will specifically focus on excitatory synapses, where the released neurotransmitter induces the influx of cationic current, thereby increasing and exciting the post-synaptic potential. These synapses are predominantly glutamatergic, meaning that glutamate serves as the primary neurotransmitter released from the presynaptic terminal. Once glutamate is released into the synaptic cleft, it rapidly diffuses and binds to glutamate receptors, a family of receptors located on the postsynaptic cellular membrane. In particular, two elements of this family, each possessing distinct kinetic properties, play a remarkably important role as plasticity drivers.

AMPA (α -amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid) receptors (AMPARs) represent ionotropic receptors responsible for fast excitatory synaptic transmission in the central nervous system. These receptors form tetrameric complexes composed of various combinations of four subunits—GluA1, GluA2, GluA3, and GluA4—which collectively influence their ion selectivity and kinetic properties. When glutamate binds to AMPARs, they rapidly open to

permit the entry of sodium (Na^+) ions, and in certain configurations calcium (Ca^{2+}) ions, into the postsynaptic neuron, resulting in fast signal transmission.

NMDA (N-methyl-D-aspartate) receptors (NMDARs) constitute another class of ionotropic glutamate receptors, playing a crucial role in excitatory neurotransmission within the central nervous system. Structurally, these receptors are large macromolecular complexes formed from various combinations of three types of subunits: GluN1, GluN2 (A-D), and GluN3 (A-B). The precise subunit composition determines the receptor's kinetic properties and impact on synaptic plasticity. NMDA receptors are distinguished from AMPARs by their unique activation requirements: both glutamate and either glycine or D-serine must bind for the channel to open. Additionally, their ion channel is blocked by magnesium ions at resting membrane potential, with this block only relieved through sufficient depolarization. This dual requirement makes NMDARs effective "coincidence detectors," as they require simultaneous presynaptic glutamate release and postsynaptic depolarization for activation. Once opened, NMDA receptors conduct calcium (Ca^{2+}), sodium (Na^+), and potassium (K^+) ions. The high permeability to Ca^{2+} is particularly significant, as calcium influx through NMDARs triggers intracellular signaling cascades that ultimately mediate synaptic plasticity, both potentiating and depressing.

The opening of these channels allows cationic current influx, generating an EPSP that, if capable of traversing the dendrite and reaching the soma, can induce action potential generation in the post-synaptic neuron, effectively relaying the input to downstream portions of the neuronal circuit. The temporal profile of the synaptic EPSP depends on the quantity of (glutamatergic) receptors present in the post-synaptic terminal, more specifically in the post-synaptic density (PSD). To modify synaptic weight, therefore, one primary mechanism involves changing the number of glutamatergic receptors, particularly AMPA receptors, exposed on the membrane. While the relationship between receptor number and synaptic strength is not entirely bijective — additional factors, for example, include the exposed AMPAR subunit composition or their fraction included in PSD nanodomains — this relationship holds true to a reasonable degree of approximation.

Understanding how synaptic weight changes, and therefore how synaptic plasticity operates, requires comprehension of the processes that drive receptor integration and removal from the post-synaptic membrane. As is common in biological systems, the answer to this question is complex, involving hundreds of processes that shape the final outcome across multiple timescales. However, two processes can be considered the main drivers of synaptic change.

1. Introduction

Protein synthesis is the first of these processes, producing new receptor copies from mRNAs. This can occur either in the neuronal cell body, with receptors subsequently migrating to spines of interest through diffusion within the dendritic volume or along its surface, or locally within the spines themselves. Increasing the pool of receptors available for exocytosis and exposure on the synaptic surface translates to an increase in effectively exposed receptors, as their number depends on the dynamic equilibrium between the exocytosis and endocytosis rates. This relationship implies that when a general increase in receptor resources occurs throughout the entire neuron, and these resources can diffuse and enter all available spines in an unbiased manner, an overall increase in synaptic weight is observed—a process sometimes referred to as multiplicative potentiation. Conversely, if the total amount of resources decreases, a proportional decrease in synaptic strength should be observed throughout the entire neuron. Importantly, protein translation is not a rapid process and is classically associated with synaptic changes unfolding over hours, primarily linked to long-term and structural plasticity.

Protein phosphorylation, and its inverse, dephosphorylation, represent the second major mechanism of synaptic weight change. Within dendritic spines, a rich molecular network regulates the phosphorylation of various plasticity-related factors, particularly glutamatergic receptors and protein translation mediators. Unlike protein translation, phosphorylation acts on much shorter timescales, typically minutes, with increased phosphorylation generally associated with synaptic potentiation promotion through mechanisms such as increased externalization of AMPA receptors or enhanced polymerization of the scaffolding protein beta-actin, leading to increased synaptic size. Conversely, dephosphorylation is associated with synaptic depression.

The phosphorylation biochemical network implemented in each spine consists of dozens of different molecular types, each modulated by hundreds of targeting proteins. However, two enzymes have been extensively studied and are traditionally considered the central drivers of this process, potentially representing the core of the bow-tie architecture regulating fast synaptic plasticity.

Calcium/calmodulin-dependent protein kinase II (CaMKII) is a multifunctional serine/threonine kinase highly concentrated in excitatory synapses throughout the brain. CaMKII is uniquely positioned to regulate structural and functional adaptations at synapses. Upon synaptic activity, calcium influx through NMDA receptors activates CaMKII via binding to calmodulin, followed by autophosphorylation that can sustain CaMKII activity even after

calcium levels diminish. CaMKII directly modulates synaptic machinery by phosphorylating key targets such as AMPA receptor subunits and scaffolding proteins, promoting enhanced AMPA receptor conductance and increased insertion into the postsynaptic membrane. This process strengthens synaptic connections and transforms silent synapses into active ones. Interestingly, it has been reported that CaMKII is also involved in long-term depression (LTD), mediating activity-dependent synaptic weakening in cerebellar Purkinje cells by altering its substrate specificity in response to different signaling conditions. Different CaMKII isoforms (α , β , δ , γ) and their subunit composition further contribute to the precision and diversity of synaptic plasticity modulation. Notably, the α and β isoforms have distinct affinities for calmodulin and actin, allowing fine regulation of synaptic strength and spine morphology.

Calcineurin (CaN), also known as protein phosphatase 2B, is a calcium/calmodulin-dependent serine/threonine phosphatase that acts as a negative regulator, counteracting kinase activity to modulate synaptic efficacy by dephosphorylating synaptic proteins such as glutamate receptors and ion channels. Under normal conditions, active calcineurin reduces synaptic strength by promoting removal of AMPA receptors. Calcineurin is implicated in both Hebbian plasticity mechanisms, such as long-term depression (LTD), and in the modulation of ion channel phosphorylation states that govern synaptic transmission and neuronal excitability. Moreover, calcineurin acts as a critical Ca^{2+} -sensor, linking postsynaptic calcium dynamics to pathways such as retinoic acid-dependent signaling that further mediate synaptic homeostatic regulation.

These two enzymes act in opposing ways: CaMKII promotes synaptic potentiation while calcineurin hinders it and promotes depression. To understand how synaptic activity links to synaptic change, it is necessary to comprehend how synaptic activity connects to these two molecular players. Following release into the synaptic cleft, glutamate binds to both AMPA and NMDA receptors. AMPA receptors are primarily permeable to sodium ions, whose influx induces EPSP formation. NMDA receptors, however, are primarily permeable to calcium ions, whose influx impacts post-synaptic potential but is mainly responsible for initiating the calcium cascade.

One principal branch of this cascade begins with the second messenger calmodulin (CALcium-MODulated proteIN, CaM), a multifunctional intermediate calcium-binding messenger protein. After binding calcium, CaM undergoes activation through conformational change and proceeds to activate downstream effector molecules, including CaN and CaMKII. This activation occurs through reversible, non-covalent binding and, most importantly, with active

1. Introduction

competition between these two enzymes for CaM. The final phosphorylation/dephosphorylation regime depends on the amount of available calmodulin as well as the different binding affinities. It is strongly suspected that the higher affinity of calcineurin for CaM, coupled with the higher synaptic abundance of CaMKII, leads to the frequency-dependent plasticity phenomenon. Indeed, experimental observations have revealed that low-frequency activation of the post-synaptic terminal, with consequently lesser calcium influx and CaM activation, leads to synaptic depression, while high activity frequency produces the opposite effect. This result can be explained, at least approximately, by an ordered activation of CaN and CaMKII, where calcineurin activates first at low CaM concentrations, while CaMKII activates later but in higher amounts, shifting plasticity toward potentiation.

Understanding how this elementary principle relates to the plethora of other plasticity drivers, as well as to the ability of all these molecular players to diffuse outside stimulated spines into neighboring dendritic and synaptic regions, represents a significant challenge. The complex interplay between these mechanisms across multiple spatial and temporal scales continues to be an active area of investigation, requiring sophisticated experimental and theoretical approaches to fully elucidate the molecular basis of synaptic plasticity and its role in learning and memory formation.

1.4. Research goal

Learning and memory formation depend fundamentally on synaptic plasticity. At the biological level, synaptic plasticity emerges from complex molecular dynamics that operate across multiple spatial and temporal scales. While contemporary experimental approaches provide unprecedented insights into the remarkable intricacy of molecular factors within neuronal dendrites and axons, our theoretical understanding of the complete activity landscape of these factors, as well as their relation to synaptic stimulation, remains incomplete.

Techniques that can capture the spatial and temporal dynamics of molecular movement are typically limited to examining only a small number of molecules simultaneously. Conversely, high-throughput approaches like omics technologies, while capable of analyzing thousands of molecular species, provide merely static snapshots of neuronal states at discrete time points, lacking the temporal resolution necessary to understand dynamic processes. The investigator is, therefore, constrained to trade off the breadth of the observed species for the spatio-temporal precision of their observations, accepting in both cases to work with only a partially

observable ground truth. This lack of observability is particularly critical within neuronal compartments like dendrites, where the behavior of numerous molecular species becomes interconnected across spatial scales spanning from tens to hundreds of microns. When an individual molecular component responds to stimulation, it frequently initiates elaborate downstream signaling networks that produce a multitude of effects throughout the neuron. Moreover, molecular responses to identical stimuli exhibit significant context dependency, influenced by both the neuron's present state and its previous activation history.

Experimental neuroscience faces an additional major obstacle in disentangling the relative contributions of diverse mechanisms that govern protein dynamics. Although researchers can examine specific processes such as local protein synthesis with considerable precision, separating and quantifying the influences of concurrent mechanisms, including diffusion, degradation, active transport, and vesicular trafficking events such as endocytosis and exocytosis, remains a formidable challenge. These interconnected and overlapping processes work together to regulate synaptic plasticity, making it difficult to attribute observed changes to particular underlying mechanisms.

The role of a consistent theory is exactly to address these challenges. Numerous models of synaptic plasticity have been developed across a multitude of spatial and temporal scales, shedding light and reuniting disparate, apparently contradictory experimental observations under a single common theoretical framework. In this work, we follow this approach, proposing two mathematical models expressed in two different (although compatible) mathematical languages. Both these works address, from different perspectives, three very open questions concerning synaptic plasticity and, more in general, synaptic dynamics:

1. what are the statistical properties of a synaptic ensemble, and what are these statistical properties driven by;
2. how do these basal statistical properties drive the synaptic response to a stimulus, sometimes even switching the response from potentiation to depression;
3. what are the spatio-temporal dynamics of a synapto-dendritic system as a whole, and how can these predict homo- and heterosynaptic plasticity in a joint, multi-spine plasticity framework.

1.5. Contribution statement

Research is a fundamentally collective endeavor, and this is particularly evident in computational neuroscience, where discoveries emerge through constant interaction between scientists from diverse fields, each contributing their unique perspective to the question under consideration. Like an architectural endeavor, the final result does not have a single true author but represents the collective achievement of individuals who, by sharing their respective expertise, enable the construction of knowledge that hopefully will belong *openly* to everyone.

This principle applies equally to the present thesis. Although I have carried out the theoretical and computational work presented here, I have never worked in isolation. To better clarify my specific contribution to each project while rightfully acknowledging the efforts of my colleagues, I provided detailed attribution statements in gray boxes at the beginning of Chapters 2 and 3, clearly delineating the role each collaborator has played in the piece of work being reported.

2. Stochastic modeling of synaptic fluctuations

This chapter is adapted from the published work (star indicates shared first authorship):

M.F. Eggli*, T.E. Chater*, J. Petkovic*, Y. Goda, & T. Tchumatchenko, (2023). *Linking spontaneous and stimulated spine dynamics*. Communications biology, 6(1), 930. <https://doi.org/10.1038/s42003-023-05303-1>.

The authors' contributions are as follows: M.F.E. and J.P. analyzed the dataset and developed the model; T.C. conducted the experiments; M.F.E., J.P., T.C., Y.G., and T.T. prepared the manuscript; M.F.E. and T.T. conceived the study. Y.G. and T.T. supervised the project.

Specifically, my contributions include the data analysis of the spontaneous synaptic dynamics, the design and validation of the proposed models on experimental and simulated synaptic data, the analysis of the synaptic change autocorrelation function, as well as the final framing and interpretation of the obtained results.

2.1. State of the art and open challenges

Learning and memory are believed to depend on modifications in synaptic strength, involving the enhancement and weakening of particular synaptic connections [8–13]. Multiple investigations have examined the molecular mechanisms underlying synaptic plasticity across minutes [14] as well as over periods spanning hours to days [6,15,16]. Although synaptic plasticity typically targets particular synaptic locations, synapses can also demonstrate dynamic behavior without directed plasticity, making it difficult to distinguish spontaneous changes from activity-induced modifications [17]. Synapses experience substantial size variations across hours and days, likely resulting from spontaneous molecular dynamics at synaptic sites [18–24].

Although individual synapses may undergo considerable fluctuations across time, average characteristics of the population demonstrate remarkable temporal stability [19,21,25–30]. Numerous experimentally observed synaptic size distributions display asymmetry and possess an extended right tail, which has been proposed to relate to optimality regarding information storage capacity, neuronal firing patterns, and long-range information transmission [31,32]. Although these distributions are typically assumed to result from the combined effects of spine shrinkage and growth [29,33], the relationship between activity-independent and activity-

2. Stochastic modeling of synaptic fluctuations

dependent elements remains incompletely understood [30]. Furthermore, modeling studies frequently rely on one key assumption: synapses maintain their characteristics indefinitely unless actively stimulated to change. This assumption is critical since otherwise, spontaneous changes would result in alterations to network function or the loss of newly learned abilities. Nevertheless, the reality that synaptic modifications are governed by inherently noisy molecular processes (such as lateral diffusion, active trafficking, endocytosis, and exocytosis [23,34]) suggests that spontaneous changes are unavoidable. Therefore, investigating how fundamental properties of synapse populations are preserved (including probability of release, total receptor conductance, size, morphology, ultrastructure, and composition) across extended timescales represents another essential component of memory research.

This ability of synapses to maintain their characteristics is termed synaptic tenacity in the scientific literature [17,35]. Models connecting these observations to individual spine dynamics through different methodologies have been developed [18,21,28,36,37]. In the present chapter, we introduce a model capable of reproducing both LTP-induced spine modifications and activity-independent spine fluctuations using a unified framework. It is proposed that LTP affects smaller spines more significantly due to their greater capacity for growth [38], whereas larger spines may represent stable long-term memory storage [39,40]. In the activity-independent context, research has demonstrated that larger spines exhibit greater variability [18,21,30]. Our model, drawing inspiration from the Kesten process and multiplicative dynamics from prior research, enabled us to replicate experimental findings related to spontaneous spine fluctuations while accounting for their log-normal characteristics. We were also able to apply our model to describe post-LTP spine behavior and document a notable increase in entropy (representing a measure of dendritic information storage capacity). Our findings regarding spontaneous spine fluctuations align with previously documented phenomena, including large spine variance, stable population distributions, and oscillatory spine behavior resulting from negative correlations between timesteps [18,19,37], and can clarify how LTP signals influence spontaneous spine distributions.

We propose that a fundamental process generating spontaneous spine distributions (activity-independent spine plasticity) is altered by plasticity induction, allowing both spontaneous and induced spine distributions to be characterized using the same model under different model conditions. Consequently, before examining stimulation effects, we sought to identify the model mechanisms required to capture activity-independent, spontaneous spine fluctuations.

2.2. Preliminary experimental observations

We started by turning our attention to the dataset kindly shared by our experimental collaborators, who conducted confocal imaging of spines on apical oblique dendrites of GFP-expressing CA1 pyramidal neurons in cultured hippocampal organotypic slices. In one experimental setup, quasi-simultaneous potentiation of a subset of spines was performed through glutamate uncaging, in order to induce structural LTP (sLTP). In a separate independent experimental set, the caged glutamate molecule was excluded from the bath, preventing spines from undergoing sLTP following laser illumination. This sham stimulation dataset served as our activity-independent group. For both conditions, across 55 minutes (15 minutes pre- and 40 minutes post-stimulation), spine sizes are recorded at eight time points (at $-15, -10, -5, 2, 10, 20, 30, 40$ minutes, where negative values indicate pre-stimulation) to examine spine dynamics. This dataset includes three baseline measurements, followed by glutamate uncaging or sham-uncaging, then five additional time points. This approach enabled us to directly observe LTP induction effects on spine populations and incorporate how newly potentiated synapses and their unstimulated neighbors develop within a unified model.

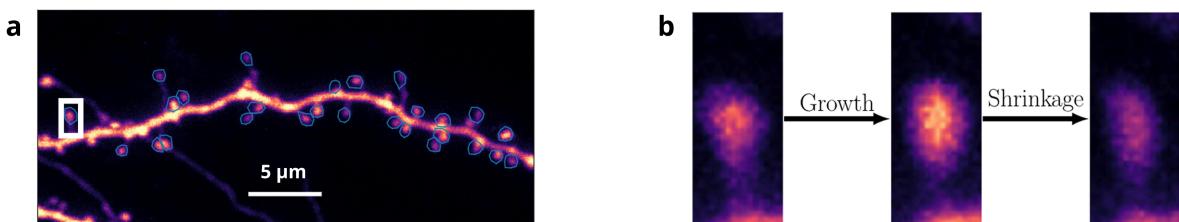


Figure 2.1: **Example image from the utilized experimental dataset.** **a** a GFP-expressing CA1 neuron whose spine dynamics we analyze and model. **b** spontaneous dynamics at the single spine level. The spine (marked by a white rectangle in a) exhibits both growth and shrinkage in the observed time frame.

We assessed synaptic strength at each time point by measuring spine head size [38,41,42] since numerous synaptic parameters correlate with head volume [43,44]. We present an example image, including semi-automatically generated ROIs used for spine head size measurement in Figure 2.1. We have marked a synapse with a white rectangle in and illustrated its varying sizes at different time points (Figure 2.1 b) to highlight the variable dynamics spines experience. These recordings are conducted in an imaging solution containing tetrodotoxin (TTX), picrotoxin, and with nominally 0 mM Mg^{2+} . Under these conditions, without neuronal spiking and experimentally imposed stimulations, spines continuously fluctuate spontaneously in size across time.

2. Stochastic modeling of synaptic fluctuations

Despite this variability, the spine size distribution (Figure 2.2 a) remains remarkably stable over time. Its shape is right-skewed and displays an extended right tail, consistent with findings reported previously across various experimental studies [21,30]. Importantly, we observed that the spine population mean is also remarkably stable, contrasting with individual spine dynamics (see inset of Figure 2.2 a). We note that spine size change distributions (Figure 2.2 c) demonstrate Gaussian behavior with no significant differences between time points (Kolmogorov-Smirnov (KS) test not significant except for the change from +10 minutes to +20 minutes, which is indicated by an *).

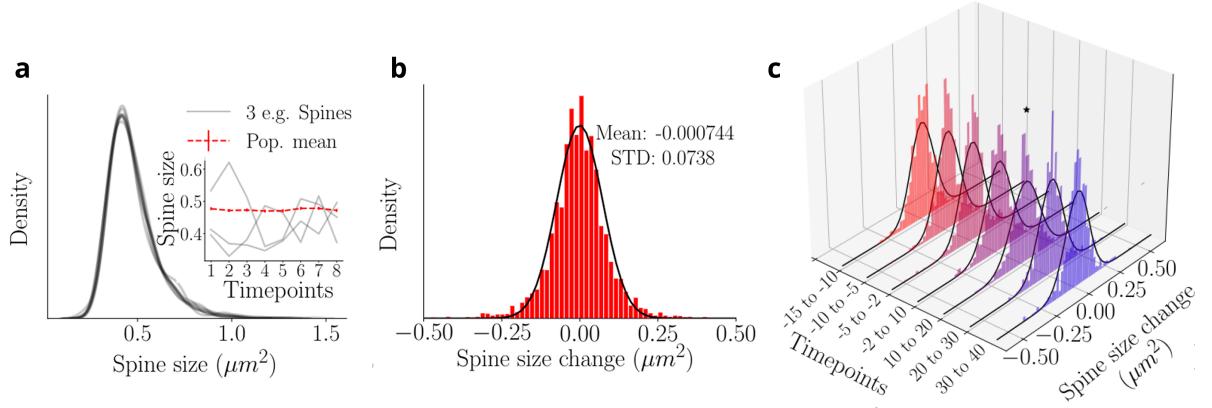


Figure 2.2: **Temporal stability of synaptic population statistics.** **a** The spine sizes follow a temporally stable right-skewed distribution with a long tail. Each gray line refers to a different snapshot distribution, which shows significant overlap. Inset: Temporal evolution of 3 randomly selected spines (gray) and population mean (red). **b** The collection of all spine changes across all time points follows a zero mean Gaussian distribution and a standard deviation $\sigma \approx 0.07 \mu\text{m}^2$. **c** Collective distributions of the spine size changes from time point to time point follow a Gaussian distribution. The black lines denote the corresponding Gaussian fits. The * denotes the single distribution that is significantly different ($p < 0.05$ when tested with KS-test).

We can also combine all these changes into a single distribution and calculate the sample mean, μ , and sample standard deviation, σ . The resulting distribution and sample statistics are shown in Figure 2.2 b. Spine size changes are consistently negatively correlated between adjacent time steps (see darker red colors in Figure 2.3 a). This effect occurs on the scale of tens of minutes in our data, which is considerably shorter than the day-long spine correlations (which also have smaller values) documented by previous studies [18,36]. We note that this correlation also continues throughout our experiments, provided the timesteps directly follow each other, e.g., computing the correlation of timestep 4 - timestep 2 and timestep 6 - timestep 4.. Finally, spines of different sizes display distinct spine change distributions (Figure 2.3 c,

2.2. Preliminary experimental observations

KS-test performed between samples yielded p values all under 0.05), which are all well-characterized by log-normal distributions (black lines).

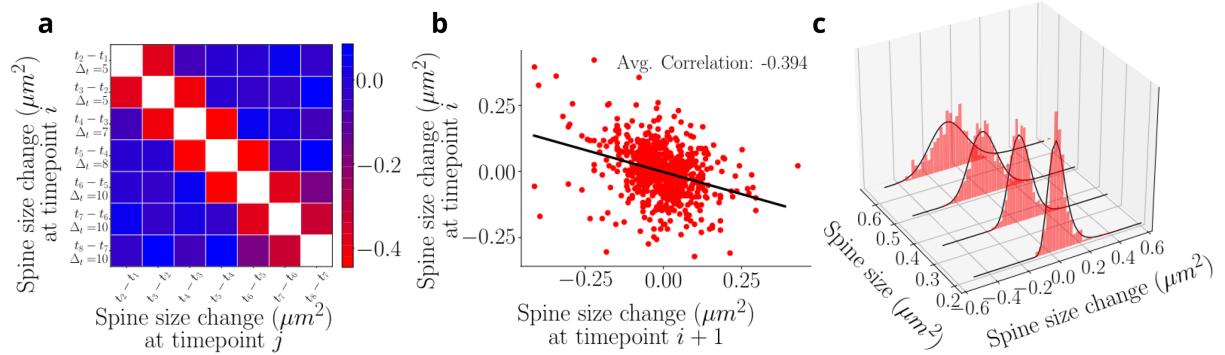


Figure 2.3: **Temporal features of synaptic fluctuation process.** **a** Spine sizes display correlations across time, whereby the neighboring time points are negatively correlated (negative off-diagonal values). **b** Correlation of two time points. **c** Evaluating spine size changes as a function of the spine size across time points shows that small spines exhibit a narrow distribution of changes, while larger spines show larger variability. Black lines represent the corresponding log-normal fits of the data.

These experimentally observed findings in our data lead us to the following question: given individual spine dynamics (oscillatory, small vs large), how is the steady size distribution maintained? We address this question by introducing an abstract stochastic model that incorporates the minimum number of parameters to preserve tractability while capturing the following key characteristics of our experimental data:

1. the temporal spine dynamics must remain stable around the distribution observed in the dataset (Figure 2.2 a). As a result, the distribution mean must remain stable through time (Figure 2.2 a, inset);
2. if we begin with another distribution, e.g., a uniform or delta distribution, the model should return to the original stable distribution. This assumption does not emerge directly from the observed data but rather from the fact that, as neurons develop and change, the initial spines could start small and still reach the distribution of (Figure 2.2 a), which is stable over the timescales we consider (approximately 10 minutes). Therefore, to maintain biological realism, we will incorporate this feature;
3. the global dynamics of spine changes and their distribution from one time point to the next should follow a Gaussian distribution, (Figure 2.2 c);
4. time points directly following each other should be negatively correlated with one another (Figure 2.3 b), suggesting an oscillatory dynamic component.

2.3. Designing the model

Before proceeding with the modeling, we introduce some helpful notation. Throughout this manuscript, we used the absolute change in spine areas, which is defined as follows:

$$\Delta V_i = V_i - V_{i-1} \quad (2.1)$$

To compare distributions against each other, the populations were taken (in the case where these samples were very large, randomly subsampled), and a Kolmogorov–Smirnov test was performed. Single asterisks indicate $p < 0.05$. Fits of probability distribution functions were performed using SciPy. Correlations report the Pearson linear correlation coefficients. Unless reported otherwise, error bars in line plots refer to the standard error and in box-and-whisker plots refer to the inter-quartile range.

2.3.1. Bounded Wiener, and Ornstein-Uhlenbeck processes

We begin by examining the Gaussian distribution of the experimentally observed spine changes in Figure 2.2. Therefore, a purely Gaussian model (Wiener process) for the spine changes appears as a natural initial choice. This model has the form:

$$\Delta V_{i+1} = \eta_i \quad (2.2)$$

where $\eta_i \sim \mathcal{N}(\mu, \sigma)$ and ΔV_i is the size change between time points i and $i - 1$. While this model is straightforward and captures the experimentally observed statistics of spine changes, it displays an inherent incompatibility with other experimental results. Since a Gaussian distribution is naturally unbounded, this model allows infinitely large (negative and positive) spine size values. Historically, the absence of bounds in a Gaussian distribution has been addressed through the introduction of bounding walls W_l, W_r (e.g., in [18]): at each time step, the value V_{i+1} is reset to be within the range $[W_l, W_r]$, where $W_l < W_r$. This can be accomplished, for example, by using either a bounce-back mechanism (i.e., a change in the opposite direction) or imposing no change, i.e., $V_{i+1} = V_i$. To examine whether the introduction of walls can allow us to proceed with the Gaussian model, we implemented two walls (W_l and W_r) which we set to equal the fifth percentile and the largest experimentally observed spine size, respectively.

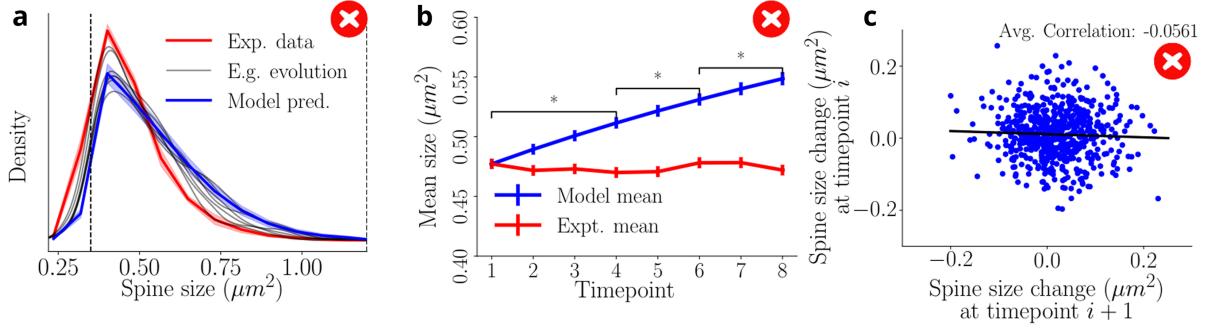


Figure 2.4: **Statistical properties of the bounded Brownian process (2.2)** **a** Synaptic size probability distribution. Starting from the (experimentally observed) initial condition (red), the system evolves (gray lines), arriving at the last simulated state (blue). **b** The ensemble mean of the process (blue) does not reach stationarity, in disagreement with the experimental observation (red). **c** Correlation of subsequent size changes, considerably smaller than the experimental observation.

The model simulations resulting from (2.2) are shown in Figure 2.4, where the dashed lines represent the walls. Despite good agreement with the collective spine distribution, three conceptual issues eliminate this model:

1. the left wall enforces an accumulation of smaller sizes that leads to the desired asymmetry, but also results in a complete drop-off in spines smaller than this size (Figure 2.4 a);
2. spines are free to grow until they reach the right wall value, causing an overall increase in the population mean and a biologically implausible growth at the right tail of the size distribution (Figure 2.4 a,b).
3. the negative correlation between subsequent size changes is lost due to the memory-less additive Gaussian noise (Figure 2.4 c).

Therefore, we will modify our model further to include a negative temporal correlation and achieve a biologically plausible spine size distribution. To this end, we will replace the purely diffusive process with an Ornstein-Uhlenbeck process. This approach was previously also used in [21] to model activity-independent plasticity in a framework with multiplicative noise. Here we will be applying it in an additive manner:

$$\Delta V_{i+1} = -\theta(V_i - \mu) + \eta_i \quad (2.3)$$

where θ, μ are the drift terms and η_i is as above. We observe that this process, characterized by the deterministic drift toward the long-term average μ , can reproduce the experimental mean-reverting behavior shown in Figure 2.3 if θ is large enough. However, if we choose θ to be too large, all the spine sizes will eventually stabilize in a narrow neighborhood of μ , which is inconsistent with the experimental observation that even after hours and days, there was

2. Stochastic modeling of synaptic fluctuations

a stable and diverse set of different spine sizes [18,23]. Adopting a set of different values of constant μ for the different spines while maintaining a high value of θ allows the recovery of this phenomenon, but inevitably locks the spines into their stable size and prevents them from changing from one size to the other.

2.3.2. Negative momentum and non-Markovianity

To avoid these pitfalls, we introduce a drift μ that is (i) unique to each spine and (ii) time-dependent. Thus we avoid both the global stable size as well as the local stable size. The simplest implementation of this principle is the introduction of a “negative-momentum” term, obtained by setting $\mu = V_{i-1}$:

$$\Delta V_{i+1} = -\theta (V_i - V_{i-1}) + \eta_i \quad (2.4)$$

This non-Markovian process contains a bounce-back mechanism that induces the spines that have grown in the previous step to have a higher probability of shrinking in the next one. Importantly, this effect vanishes at longer timescales. We implement this model by setting θ to achieve the experimentally observed correlation. The results of the simulations can overcome two of the three issues illustrated above: the population mean remains stable over time (Figure 2.5 a, inset), and the oscillatory behavior reappears in agreement with the experimental observations (Figure 2.5 b). However, the additive Gaussian term is still responsible for improper tail-flattening and, ultimately, for improper symmetrization of the spine size distribution. This fact and the observation that the different spine sizes exhibit different variation profiles (see Figure 2.3 c) demonstrate that more complicated noise-generating models are required to reproduce activity-independent plasticity.

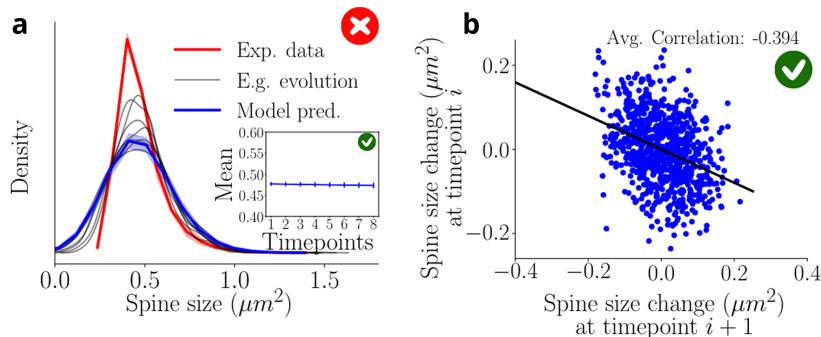


Figure 2.5: **Statistical properties of the negative-momentum model**
(2.4). Panels are presented as in Figure 2.4.

To comprehend the required noise profile driving spine size modifications, we begin with a more detailed examination of the experimental data. The overall distribution of spine size changes across time appears to be Gaussian (Figure 2.2 b), which may suggest a model based on Gaussian dynamics. However, when we tested such a model, we found that there were fundamental issues that did not align with the experimental results.

In fact, we observe that the overall profile of the spine size population is a skewed, log-normal-like profile (Figure 2.2 a). Additionally, when we examine the changes in spines with different initial values separately, the distribution of changes also displays a skewed profile (Figure 2.3 c). Furthermore, we observe that these distributions differ from one another, indicating that spines belonging to different size intervals behave in a fundamentally different manner. Thus, we introduce a model with a noise profile, η_i , which is sampled from a set of log-normal distributions such that

$$\Delta V_{i+1} = \eta_i, \quad \eta_i \sim \text{Lognorm}\left(\mu_{log}(V_i), \sigma_{log}(V_i), \hat{\delta}\right) \quad (2.5)$$

where μ_{log} and σ_{log} are parameters that depend on the spine size V_i and determine the shape of the log-normal sampling distribution, and $\hat{\delta}$ is a location parameter (see Appendix A for more details). To establish the dependence of μ_{log} and σ_{log} on the size of the spine V_i we assume, following observations seen in [18,21,30], that there exist two linear functions f_μ and f_σ that map spine sizes onto the corresponding log-normal change parameters. However, rather than identifying the linear functions that are optimal for all spines which (i) becomes computationally expensive, (ii) can lead to overfitting, or (iii) leads to difficulty inferring the underlying distribution due to insufficient data, we simplify the above model by binning spines in equal-size bins and then evaluating the sample means and standard deviations of those bins. This provides exactly the linear functions f_μ and f_σ which allow us to estimate the sample means and deviations for all spine sizes (identified by the pedix \cdot_s), i.e.,

$$\mu_s(V) = f_\mu(V), \quad \sigma_s(V) = f_\sigma(V) \quad (2.6)$$

These values can be used to estimate the parameters of the underlying normal distribution, which can then be transformed into the parameters to define the log-normal distribution (μ_{log} and σ_{log}) using equations (2.39), generating our noise profile. We note that previous work (including that of [30]) found linear relations between the spine size squared and the variance and mean. We observed that such fits were equally effective as the fits presented here and led to similar results. The fits for f_μ and f_σ (Figure 2.6 a) result in the following interesting outcomes: (i) small spines have a positive mean change and have smaller standard deviation,

2. Stochastic modeling of synaptic fluctuations

so they tend to grow but are less variable and (ii) large spines have a negative mean change and larger standard deviation, so they tend to shrink and are more variable.

After optimization, we can simulate the synaptic size process (2.5) (which we refer to as Best fit LN Mode), and examine the properties of arising size dynamics (Figure 2.6 b,c). The generated results do not recreate the desired experimental characteristics, i.e., the mean of the simulated distributions (inset of the same figure) decreases, and the negative correlation is too small (compare Figure 2.6 c and Figure 2.3 a).

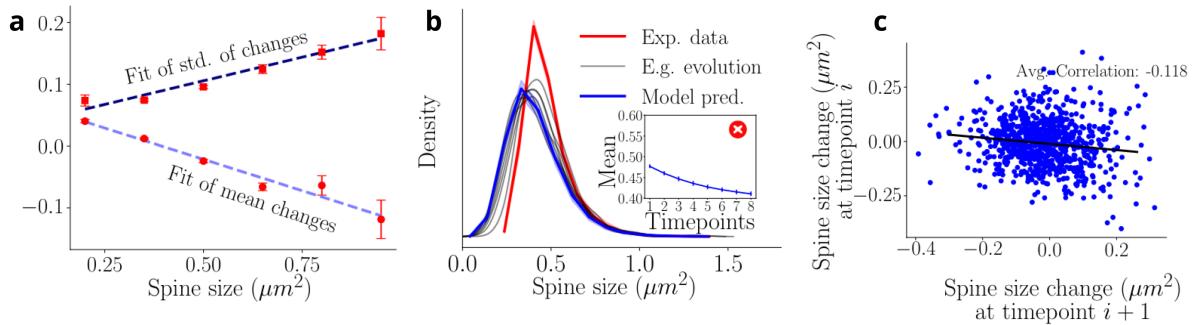


Figure 2.6: **Parameter optimization and dynamics obtained with model (2.5)** **a** Fit to data of the size dependent parameters $\mu_{log}(V)$ and σ_{log} . **b** Deriving ensemble simulation does not result in a stable distribution compatible with the experimental observations. In the inset, the simulated mean is shown, which decreases significantly. **c** Correlation obtained from one example step of the best-fit log-normal simulations. The value of the correlation is $r \approx 0.1$, drastically smaller than what was observed experimentally.

2.3.3. The LN-OU Model

We observe, however, one critical fact: by slightly modifying the “best” linear fits of the means and standard deviations (raising the mean and lowering the standard deviation, Figure 2.7 a), we obtain an alternative parametrization (Alt fit LN Model) which allows for the emergence a spine size distribution in excellent agreement with the distribution which was experimentally observed (Figure 2.7 b). Still, however, the correlation between subsequent size changes is still significantly underestimated (Figure 2.7 c).

We can alleviate this discrepancy by implementing the negative momentum term (cfr. (2.4)) together with the altered parametrization. Despite the excellent agreement with the experimental results, we found it necessary to use the manually tuned fits for obtaining the mean and the standard deviation. As such, when implementing the Alt. Fit LN model, we were not using the optimal fits shown in Figure 2.6. We assume that the discrepancy in using the optimal fits is not due to any noise arising from the experimental setup, but, instead, because we are

missing a crucial aspect that the “altered” fits are accounting for. These observations lead us to introduce two key modifications in model (2.5):

1. to recover the negative correlation between subsequent size changes, we introduce the negative momentum term;
2. by observing that the manual changes applied to the fits are equal across all spine bins (Figure 2.7 a), we propose that an additional global drift term can recover the experimentally reported dynamics of the spine while allowing the differential analysis of spine dynamics in different size groups. Therefore, we also add a global Ornstein-Uhlenbeck (OU) drift term (referred to as *drift* below).

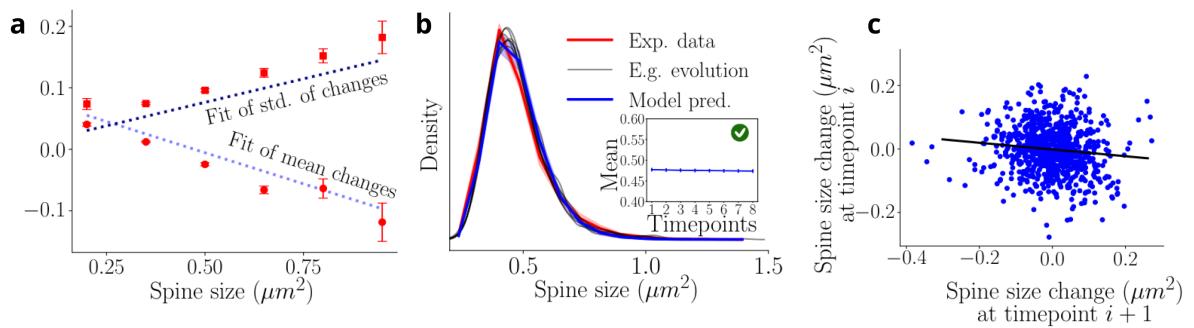


Figure 2.7: **Effect of parameter alteration on model’s (2.5) dynamics.** **a** Altered linear fits are used to achieve modeling goals. **b** Deriving ensemble probability distribution, achieving stability compatible with experimental observations. **c** The obtained correlation between subsequent size changes is, however, still smaller than required.

The parameters of our final model, referred to as the Lognormal-Ornstein-Uhlenbeck model, or *LN-OU model*

$$\Delta V_{i+1} = \underbrace{\text{Lognorm}\left(\mu_{log}(V_i), \sigma_{log}(V_i), \hat{\delta}\right)}_{\text{Long-term stochasticity}} - \underbrace{\tilde{\theta}(V_i - \tilde{\mu})}_{\text{Drift}} - \underbrace{\theta(V_i - V_{i-1})}_{\text{Negative momentum}} \quad (2.7)$$

are fitted to achieve the best match to the experimental data. The resulting simulation is illustrated in Figure 2.8 and demonstrates that we correctly reproduced all the experimental data we started out with in Figure 2.2. Both the size distribution and the collective size change distributions are captured accurately and maintain a correct degree of negative correlation between subsequent size changes.

2. Stochastic modeling of synaptic fluctuations

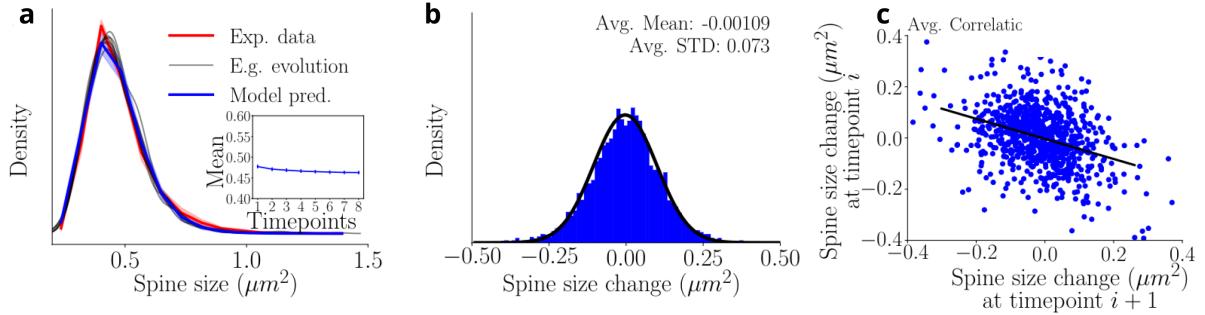


Figure 2.8: **Synaptic dynamics obtained with model (2.7).** **a** Synaptic size distribution emerging from the best linear fits (Figure 2.6 a). Stability is reached compatibly with the experimental observations. **b** Simulated activity-independent synaptic changes, distributed compatibly with Figure 2.2 b. **c** Substantial negative correlation between subsequent size variations, compatible with the experimental observations.

2.4. The impact of long-term potentiation on synaptic distributions

In previous analyses, all spines along the imaged dendritic branch were grouped into a single set, since no unambiguous criterion existed to distinguish them. However, since we intentionally induced plasticity through glutamate uncaging at specific spine groups, we can now establish two separate spine categories: those that received stimulation (homosynaptic, i.e., synaptic targets that have been specifically selected for LTP) and those that remained unstimulated (heterosynaptic, i.e., spines located on the same dendritic segment that are not directly stimulated). We stress that the heterosynaptic spines, which did not receive laser targeting for glutamate uncaging despite being on the same dendritic branch as homosynaptic spines, differ from the spines in the previous sham stimulation experiments, which received laser targeting but lacked glutamate and therefore did not undergo potentiation. We limit our definition of heterosynaptic spines to those within $4\ \mu\text{m}$ of stimulation sites and consider them as one separate group. Finally, to ensure adequate numbers of homosynaptic spines, we focus on the experimental protocol where plasticity induction was carried out on 15 distinct spines on the same dendritic branch. Prior to applying the previously established LN-OU model to this dataset, we must examine how stimulation affects the activity-independent spine dynamics.

2.4. The impact of long-term potentiation on synaptic distributions

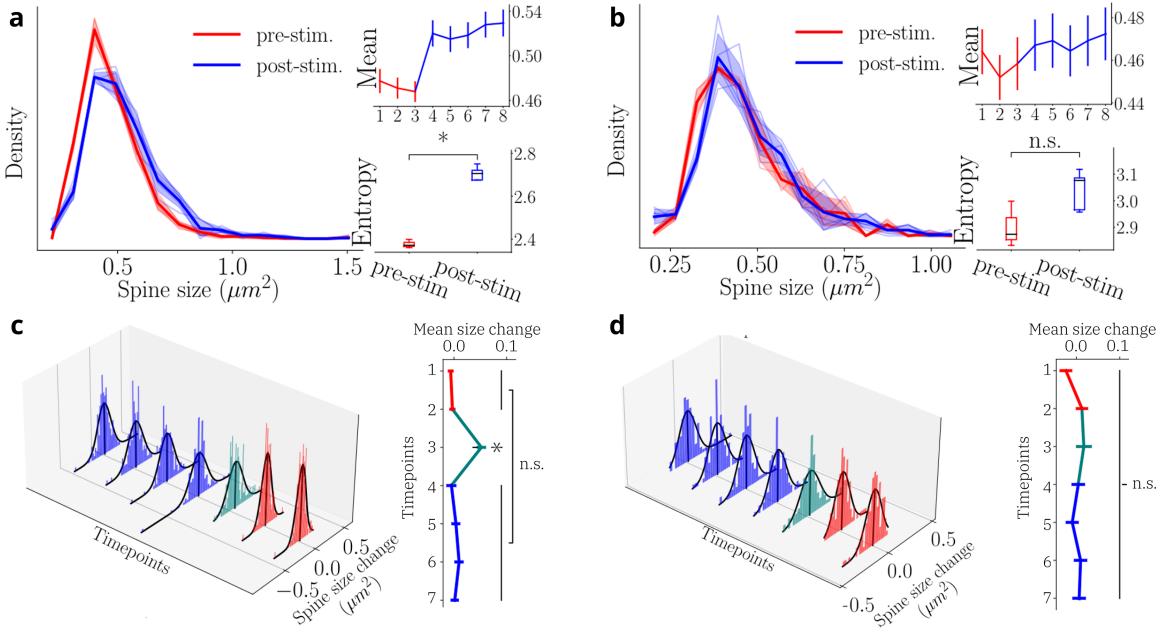


Figure 2.9: Synaptic size dynamics arising from the LTP experiment. **a-b** Homo- and heterosynaptic spine size distribution at different time points, with red and blue referring to pre- and post-stimulation, respectively. **c** Homosynaptic change distributions at each timepoint (red before, teal at, and blue after LTP induction). On the right, the corresponding mean and standard error are shown. Significant difference in change distribution is present only at the stimulation timepoint. **d** As in the previous panel, but relative to heterosynaptic spines. This time no statistical difference is found between size variations.

Starting with the homosynaptic spine distribution (Figure 2.9 a), we observe that the pre-(red) and post-stimulation (blue) stationary distributions differ significantly. This difference is also evident in the time point means (Figure 2.9 a, top inset). This indicates that spine dynamics before and after stimulation can be characterized as activity-independent plasticity around their respective stable distributions, but the stimulation event acts instantaneously (at our temporal resolution) to shift the spine size distribution. To further quantify this distributional change, we measured the information content or “uncertainty” within the spine size distributions [45]. Here, we employ Shannon entropy, which, given a random variable X with probability mass function $p(x)$ is defined by

$$\mathcal{H}(X) \triangleq \mathbb{E}[-\log p(x)] = - \sum_{x \in \mathcal{X}} p(x) \log p(x) \quad (2.8)$$

In our application, \mathcal{X} represents the complete set of possible (discretized) sizes, and $p(x)$ is the probability of these sizes. We compute the information capacity of spine size distributions before and after stimulation. The lower inset in Figure 2.9 a demonstrates a significant increase in the neuron’s potential information-storage capacity following stimulation. Conversely,

2. Stochastic modeling of synaptic fluctuations

Figure 2.9 b reveals that heterosynaptic spine (with distance from stimulation $< 4 \mu m$) size distributions and sample means show no significant shift during stimulation. While entropy does increase, this change is not significant. The time-to-time changes in both homosynaptic and heterosynaptic spines (Figure 2.9 c and d) resemble those of activity-independent plasticity (Gaussian distributions). In Figure 2.9 c, we observe that the stimulation protocol (shown in teal) produces a significant shift in the change distribution location, but no significant alteration in shape. In contrast, heterosynaptic spines show no significant differences in shape or location compared to other time points (Figure 2.9 d). Additionally, when examining the average changes at each time point (vertical black lines in the 3D plots), we find that the stimulation time point for homosynaptic spines is significantly elevated above other time points. This elevation confirms the “shift” event observed in Figure 2.9 a. We also note no significant differences among all other time points. Since we assume that pre-stimulation time points resemble activity-independent plasticity (i.e., no prior knowledge of impending stimulation exists), we can assume that post-stimulation spine change distributions are also governed by activity-independent plasticity. Heterosynaptic spines do not show such elevation, so we assume these spines predominantly undergo activity-independent plasticity.

We subsequently categorized the spine population by size into bins of $0.15 \mu m^2$, shown in Figure 2.10. Since we assume all non-stimulation time points represent activity-independent plasticity, we combine these and plot spine size changes in the left figures. The right figures display only the immediate post-stimulation time point. We note that all are approximated by log-normal distributions (black fits) (Figure 2.10 b and d). We can also compare distributions for each bin. The inset p values represent KS-test results between the two datasets. Significant differences occurred for homosynaptic spines only below $0.5 \mu m^2$, and for heterosynaptic spines, only below $0.35 \mu m^2$. This finding aligns with results from [38], suggesting that small spines are proportionally more influenced by glutamate uncaging events and play more important roles in new memory acquisition. Conversely, larger spines remain more stable and do not change significantly from baseline activity-independent plasticity.

2.4. The impact of long-term potentiation on synaptic distributions

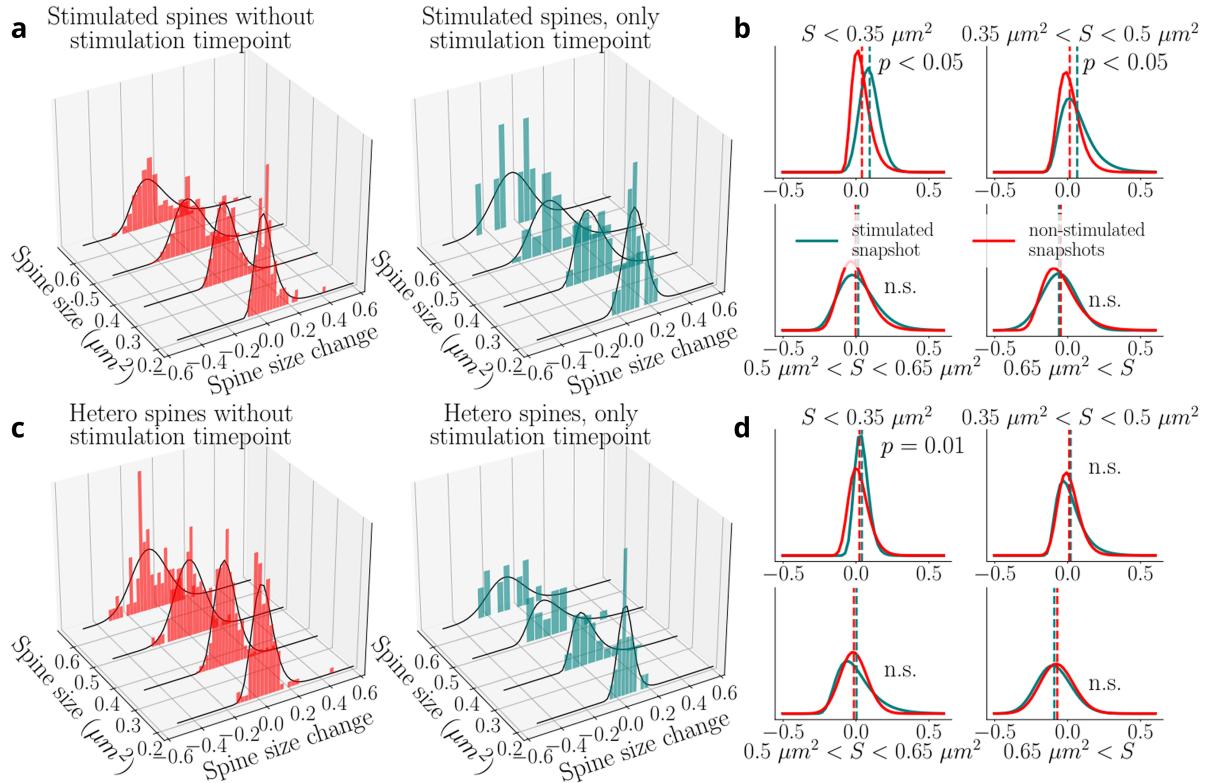


Figure 2.10: **Synaptic variations, spontaneous and stimulated, relative to each size bin (bucket).**

a Homosynaptic spines, considering the non-stimulated timesteps (left) and the stimulated timestep (right). Black lines show a log-normal fit to the data. **b** Log-normal fit comparison between non-stimulated and stimulated timepoints, for each of the four size buckets. Red refers to the non-stimulated time point, and the teal to the stimulated ones. The p-value in the figure refers to a KS test performed on the data in panel a. **c** and **d** Same as in a and b, for heterosynaptic spines.

Finally, we observe that the stimulated spine change distribution becomes narrower for small ($< 0.35 \mu\text{m}^2$) heterosynaptic spines (Figure 2.10 d, teal vs red). This narrowing appears right-skewed, suggesting that reduced activity-independent fluctuations could be preferentially linked to small spine shrinkage. Unlike stimulated small spines that experience growth, neighboring small spines encounter stimulation only peripherally. In this scenario, growth-inducing components may not reach levels sufficient for actual growth while remaining present at levels that could still oppose (or compete with) activity-independent shrinkage.

2.5. Extending the LN-OU model to stimulated spines

When applying our model to the stimulation scenario, we must establish the new linear dependencies on spine size and lognormal statistics that emerge. Initially, we examine the sample means and standard deviations for homo- (Figure 2.11 a) and heterosynaptic spines

2. Stochastic modeling of synaptic fluctuations

(Figure 2.11 b), excluding the stimulation snapshot. The resulting behaviour shows good agreement with previous fits (shown in gray), validating our observation that the pre-stimulation baseline model remains applicable.

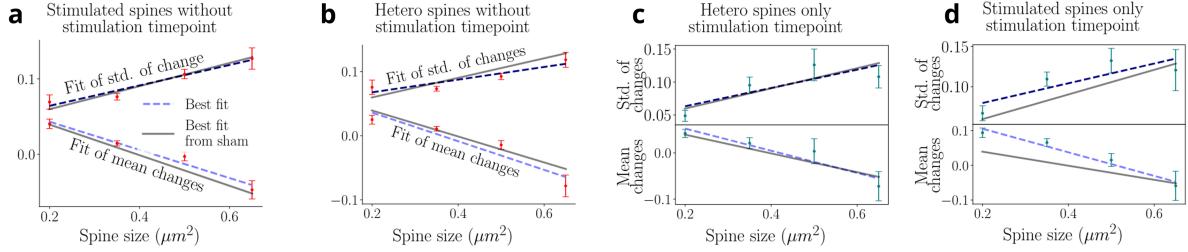


Figure 2.11: Fits for the linear functions μ_{log} and σ_{log} in model (2.7) on dataset with 15 stimulations. Subsets of homo- and heterosynaptic spines were split according to size, and linear fits were carried out for the sample mean and standard deviation of the spine variations. **a-c** of the non-stimulation snapshot of the homosynaptic spine and all snapshots of the heterosynaptic spines show good agreement with the activity-independent plasticity fits (gray). **d** Stimulation snapshot of the homosynaptic spine shows a substantial difference in the fit for smaller spines.

Subsequently, we examine the stimulation snapshot and find that the model fits for heterosynaptic spines in Figure 2.11 c show only minimal deviation in the smallest spines from the activity-independent baseline. Thus, for simplicity, we assume that heterosynaptic spines experience activity-independent baseline dynamics across all time points. In contrast, homosynaptic spines (Figure 2.11 d) exhibit different behavior. The standard deviation shows an upward tilt, indicating that the resulting log-normal distribution has enhanced its standard deviation and spines became more variable during stimulation. This increase is intuitive for several reasons: as spines undergo rapid enlargement through the potentiation protocol, their variance also increases because (i) they have expanded beyond the typical size range of activity-independent plasticity and (ii) they now represent large spines, which have been shown to display greater variance than small spines.

This standard deviation increase occurs only in medium-sized bins, not in small or large spines. This pattern may be explained by the fact that medium spines, capable of growing to large spine dimensions, now display the characteristics typical of large spines, including enhanced variance. Moreover, in the present study, medium-sized spines were preferentially selected for stimulation, since prior studies have demonstrated that this population shows the greatest lability regarding potentiation (for instance, [38]). Lastly, the mean spine change shows a clear linear trend, where smaller spines exhibit larger mean increments relative to the activity-independent plasticity fit.

To identify which model parameters require modification to replicate the stimulation time point for homosynaptic spines (Figure 2.11 a), we modify each component (long-term stochasticity and drift) of the log-normal model separately. We also assume that the negative momentum term represents an inherent feature of activity-independent plasticity, functioning as a stabilization mechanism that counteracts previous stochastic changes. Since stimulation constitutes directed activity, negative momentum would impede post-stimulation spine growth by promoting shrinkage and suggest that prior stochastic activity-independent plasticity directly influences subsequent activity-dependent changes. Therefore, we choose to disable this term in the model during stimulation to prevent this scenario. Nevertheless, future investigations could incorporate this or a generalized negative momentum term and examine its influence on the resulting synaptic size distribution.

Initially, we modified the long-term stochasticity component by implementing the linear fits for the stimulation time point (Figure 2.12 a). While the rapid component of stimulation is reproduced, maintaining constant drift causes a gradual return to the original distribution. This differs from our experimental observations with 15 stimulated spines (Figure 2.9 a). However, we do observe decay back to baseline in a separate case involving only seven stimulated spines (not shown).

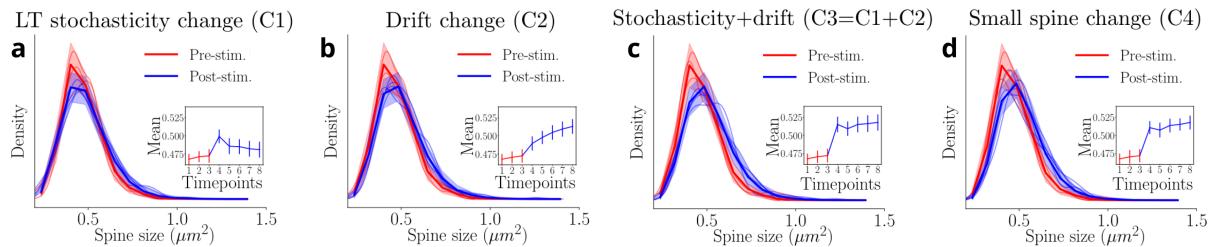


Figure 2.12: Fits for the linear functions μ_{log} and σ_{log} in model (2.7) on dataset with 15 stimulations. Subsets of homo- and heterosynaptic spines were split according to size, and linear fits were carried out for the sample mean and standard deviation of the spine activity. **a-c** of the non-stimulation snapshot of the homosynaptic spine and all snapshots of the heterosynaptic spines show good agreement with the activity-independent plasticity fits (gray). **d** Stimulation snapshot of the homosynaptic spine shows a difference in the fit for smaller spines.

This suggests that sustained LTP response correlates with increased drift and indicates that the long-term stochasticity component replicates immediate potentiation while drift produces the sustained effect. Additional support for this interpretation appears in Figure 2.12 b, where only the drift term is modified at all post-stimulation points, using linear fits from activity-independent plasticity. The changes in mean and distribution occur more slowly without

2. Stochastic modeling of synaptic fluctuations

instantaneous potentiation. Previously, long-term stochasticity and drift components operated on comparable time scales. During stimulation, we observe that long-term stochasticity produces instantaneous structural changes in spines over our considered timescale, while drift toward the new steady state occurs subsequently over longer time scales. Finally, we modify both components by adjusting linear fits at post-stimulation time points and the drift parameter μ following stimulation. Figure 2.12 c shows that we reproduce distinct stable distributions before and after stimulation on appropriate timescales (cfr. Figure 2.12 a). Therefore, the LN-OU model successfully reproduces experimental results for both plasticity types.

To generate the distributional jump observed in the data (Figure 2.9), we employed the complete linear fits (Figure 2.11 d) for the long-term stochasticity. Additionally, given that small spines show the greatest stimulation response, we investigated the effect of modifying only the smallest spine parameters in the model (Figure 2.12 d). We therefore increased only the sample mean of spines with initial area $< 0.35 \mu\text{m}^2$ during stimulation and treated the stochastic component of all other spines as undergoing activity-independent plasticity. Drift parameters were applied as described above, since they influence all spines equally. In essence, we modified the slow component for all spines but altered the fast potentiating component only for the smallest spines. This modification reproduces experimental results with no discernible difference from using complete linear fits (compare Figure 2.12 c and d).

We additionally computed the Shannon entropy of simulated distributions [45–47]. In all cases, we substantially enhance the information encoding capacity of synaptic weight distributions following stimulation. However, modifying only long-term stochasticity (representing the short-term stimulation effect) produces a smaller entropy increase that could potentially decay to pre-stimulation levels after the observed period. Other model modifications emulating the stimulation event show considerably higher entropy values (similar to experimental observations). Notably, entropy values show no significant difference between fully altered and small spine models. We conclude that models aiming to reproduce spine size population dynamics can concentrate on the smallest spines to streamline the simulation process while maintaining good results, at least over the time scales examined in this study.

2.6. Summary and remarks

In this study, we examined experimentally recorded population dynamics of both stimulated and unstimulated spines on the same dendrite. Drawing inspiration from previous work

2.6. Summary and remarks

[18,21,30], we have developed a model framework that incorporates the dynamics for spontaneous and plasticity-driven spine changes observed in our data. Specifically, we identified a stable right-skewed distribution of spine sizes where the dynamics of small and large spines appear to follow distinct computational rules.

We constructed a model that operates at the synaptic population level and can be described by a single stochastic differential equation, avoiding detailed molecular principles such as those described in [36,48], following the traditional approach of previous abstract models [18,21,28]. Through this high-level perspective, we gained an understanding of activity-independent plasticity and incorporated multiple time scales. Previous studies (e.g., [18,21,30,48]) have provided a comprehensive set of stochastic modeling approaches to describe specific individual effects present in spine dynamics. Consistent with previous literature, we applied the ergodic hypothesis for our modeling. However, given the 55-minute recording window in our dataset, we could not directly test ergodicity or demonstrate that each spine explores the complete phase space.

Our model includes both fast and slow mechanisms that have implications for neuronal synaptic stability. The rapid spine changes that are anti-correlated with previous size changes may prevent a winner-take-all system by distinguishing between small and large spines. Ensuring that large spines shrink on average guarantees bounded spine size and aligns with the long tail of the spine size distribution. Moreover, in our analysis, we found that small spines preferentially exhibited positive size changes (Figure 2.12 d), and therefore they could serve as sites of information acquisition during plasticity induction. Conversely, large spines did not significantly alter their dynamics following stimulation, suggesting that large spines could help preserve the stability of previous states. In the absence of direct plasticity signals, we found that large spines were more variable and, on average, susceptible to shrinkage. Large spines, with their complex structural organization, require greater numbers of proteins, membrane traffic, and actin filaments for maintenance, resulting in higher energy costs. This would justify favoring size reduction for large spines in accordance with an energy-efficient (homeostatic) system that preferentially degrades large spines (older memories that became obsolete) to optimize storage and energy in the brain. We note that our results do not explain how a small subset of spines (e.g., large spines representing selected memories) can be preserved over timescales of days or months.

Our model builds upon and extends several modeling studies that address the differences in small and large spine dynamics. An early study by [18] separated small and large spines

2. Stochastic modeling of synaptic fluctuations

into different categories based on manual group assignment to model activity-independent plasticity. Our work proposes a plausible mechanism for activity-independent plasticity that avoids such rigid categorization. Another study by [21] demonstrated that temporal changes in spine size could be approximated by a model incorporating two timescales using multiplicative dynamics and Ornstein-Uhlenbeck processes, consistent with the fast and slow components of our model. A different model by [28,30] employed the Kesten process to describe synaptic remodeling dynamics. In [48], the authors developed a molecular model that explained how size fluctuations and distributional shapes can emerge from stochastic assimilation and removal of synaptic molecules at synaptic sites. Finally, [36,37] used actin dynamics to model rapid, spontaneous shape fluctuations of dendritic spines, predicting that these polymerization dynamics self-organize into a critical state that generates negative correlations in spine dynamics on short time scales.

Furthermore, a crucial aspect of our study is the examination of both stimulation and activity-independent plasticity within a single experimental paradigm and single mathematical model. Previous imaging studies have either restricted glutamate-uncaging to individual spines [38,49] or small spine clusters [50,51] and did not monitor population-level changes in synaptic sizes. Others monitored multiple spines while applying global chemicals to induce plasticity (e.g., [52]). Here, we could validate one of the findings of [38] that small spines are the primary targets for growth and, therefore, may be the substrates for new memory acquisition and, consequently, that large spines are likely to be the reservoirs for long-term memories [39,40,53]. Model justifications for distinct dynamics in small and large spines have been discussed in [54], which proposed a mechanism based on clusters of interacting receptors in the synaptic membrane, in [55], who considered a reaction-diffusion model of calcium dynamics, and in [56], which showed that discrete, stochastic reactions and macroscopic reactions could be exploited for size-dependent regulation. Remarkably, we observed that the distribution of spine sizes differed post compared to pre-stimulation. In contrast, we found that the changes in spine size, when viewed as a population across all time points (longer than 2 minutes away from plasticity induction), were indistinguishable from activity-independent, spontaneous changes.

Therefore, our model provides a unified stochastic framework that helps understand spine plasticity operating spontaneously after stimulation. Finally, we examined the entropy and information content of the synaptic populations. Entropy is a measure of disorder in a system and can be assessed by observing the diversity of synaptic sizes in a neural network. Higher entropy implies a more disordered system that allows for more diverse information encoding

capabilities. Following LTP stimulation, we observed an increase in the range of synaptic sizes and, thus, a larger set of possible states consistent with higher entropy. This higher entropy could facilitate learning by enabling the network to differentially encode a wider range of inputs. Secondly, entropy can also reflect the stability and robustness of synaptic connections. A higher entropy, reflected by a more diverse distribution of synaptic strengths, could make a network less sensitive to changes in individual synapses. This increased ability to buffer against noise or disruptions, such as the loss or weakening of specific synapses, helps promote the overall robustness of the network.

Our study provides not only a unified framework for understanding spontaneous versus evoked dynamics across spines but also helps establish a coherent view of various features related to spontaneous spine dynamics that align with prior reports obtained in different experimental preparations. Since spontaneous spine dynamics is often studied across both *in vivo* and *in vitro* preparations, slices, hippocampal or organotypic cultures, and across different brain regions, confirming or differentiating these reports within a common model framework is an ongoing challenge. While our experiments are conducted in slices, (e.g., [30]) has utilized primary culture models to image spines over hours to days while monitoring fluorescently-tagged PSD components. Similarly, [18] worked in hippocampal slices and employed different blockers to silence neuronal activity while observing several spine dynamics features compatible with our model and data. Other studies, such as [21], imaged dendritic spines *in vivo* in the auditory cortex, measuring populations of spines over days to weeks. During imaging sessions, the mice were lightly anesthetized, but activity at these synapses evolved spontaneously between sessions, leading to synaptic strength changes. Remarkably, despite these differences in experimental preparations, many reported features align with our experimental data, including the right-skewed spine distributions and size-dependent statistics consistent with our model.

In conclusion, this study established a connection between activity-independent spontaneous spine dynamics and directed synaptic plasticity. Within this modeling framework, we were able to unite new and previously reported synaptic features such as stable distribution of spine sizes, [18,21,28], higher variability of larger spines vs. small spines [18,19,21,30], the oscillatory behavior of the spines [18,36], and incorporate plasticity-induced dynamics into the same model. This framework can open avenues for interpreting specific experimentally reported synaptic changes relative to spontaneous activity and help constrain plasticity models operating at the circuit level.

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

This chapter is adapted from the following article, which is currently under review at Nature Communications:

J. Petkovic, M.F. Eggl, D. Pathirana, J.P. Hasenauer, S.O. Rizzoli, & T. Tchumatchenko, *Push-and-pull protein dynamics leads to log-normal synapse1 distributions and probabilistic multi-spine plasticity*. Preprint link: <https://doi.org/10.64898/2026.01.29.702571>

The authors' contributions are as follows: J.P. and T.T. designed the study; S.O.R. provided the experimental data and contributed data analysis ideas; J.P. and M.F.E. analyzed the data; J.P. designed the model in collaboration with M.F.E., and optimized the parameterization under the supervision of D.P and J.P.H.; J.P.-M.F.E. and D.P. prepared the paper. All the authors contributed to editing and reviewing the paper.

Specifically, my contributions include the analysis of the experimental data, the model design, the mathematical derivation of the closed form approximation, its identifiability analysis and subsequent optimization, as well as the conception and validation of the final results.

3.1. State of the art and open challenges

Synaptic plasticity is a highly complex phenomenon, driven by hundreds of molecular factors whose dynamics occur on a wide variety of timescales [57–60]. These dynamics, moreover, arise from different physical processes, from molecular diffusion and transport [61], to phospho-dephosphorylation [57,58,62], to protein translation [63,64], to calcium and voltage-dependent signalling [65,66]. Recent technical and experimental advances have been able to shed light on the interplay of these different factors, collecting exceptionally rich datasets that describe synaptic features and dynamics at high spatial and temporal resolutions [67–71]. The complexity of these datasets was accompanied by the development of more elaborated computational models [68,72–75], able to navigate the vast observational landscape in a principled fashion, and to provide experimentally testable hypotheses.

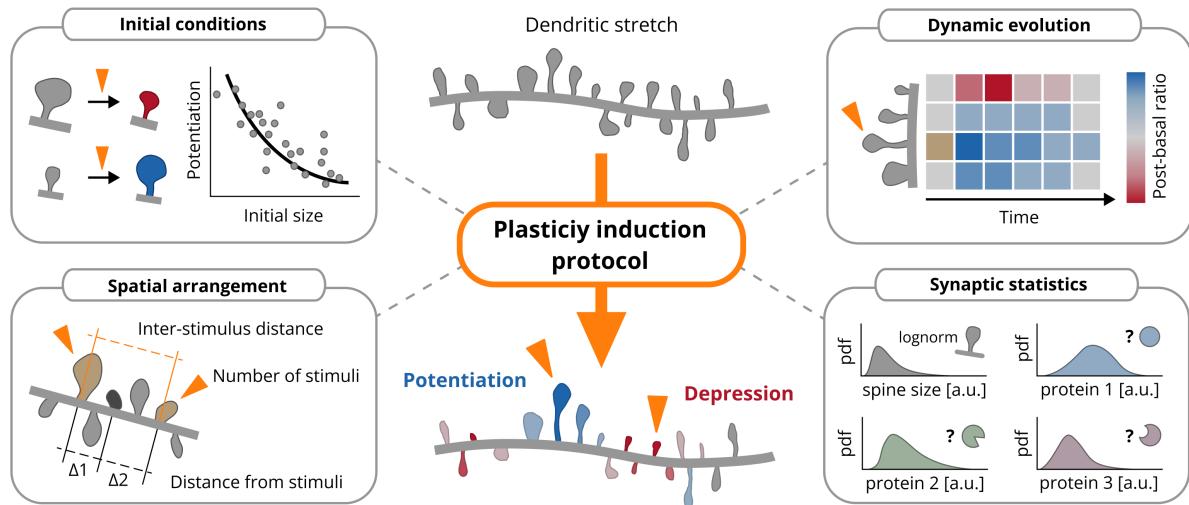


Figure 3.1: Overview of the multi-spine plasticity phenomenon. In a synaptic plasticity experiment, a subset of spines is stimulated, according to the same stimulation protocol. The outcome of synaptic changes, however, is not uniquely defined, but rather depends on a multitude of additional, confounding factors, including initial synaptic sizes, the synaptic arrangement across along the dendrite, the specific times of observation, and the synaptic content of different protein species.

Introducing model complexity, however, leads to a reduction in the model's tractability, imposing a choice between the detail of the mathematical description and the possibility to meaningfully interpret the prediction it provides. To tackle this issue, modelers have resorted to different strategies, constraining their description to the dynamics of a specific set of molecules [64,73,75], recurring to different levels of mechanistic abstraction [68,74,76], or exploring general optimality principles from a normative perspective [77,78]. Many of these models, however, still do not allow for a closed form solution, hindering the optimization of the model's parameters and, ultimately, the interpretation of the model's predictions. Crucially, the lack of a closed-form solution also complicates a clear statistical framing of these predictions, constraining the models to address only the average properties of plasticity response. This is in stark contrast with the recent theories regarding basal synaptic properties, where a substantial corpus of work [71,79–82] has been able to study the distributional properties of various synaptic metrics and, in particular, the log-normal compatible distribution followed by spine sizes and their weights. Despite strong experimental hints [83–85], whether and, especially, how these distribution impacts the synaptic response to stimulus is still an open question.

In this chapter, we will see how this question can be explored focusing on short-term changes in excitatory spine size following glutamate uncaging, examining the behaviour of a minimal

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

set of equations describing these dynamics through two fundamental processes: diffusion and phosphorylation. Despite its essential nature, this framework is able to capture several hallmarks of synaptic plasticity, aligning with established functional principles such as synaptic competition, cooperation, or the protein "tug-of-war" theory. Importantly, the model's analytical tractability and its direct connection from the underlying biochemical processes, allow it to account for synaptic variability while offering a direct link between a spine's size and its response to induction, paving the way towards a more unified understanding of synaptic function.

3.2. The optimality of the multi-spine plasticity profile

This section is adapted from the following published work:

Petkovic, J., & Fioresi, R. (2024). *Spontaneous emergence of robustness to light variation in CNNs with a precortically inspired module*. Neural Computation, 36(9), 1832-1853. https://doi.org/10.1162/neco_a_01691

The authors' contributions are as follows: J.P. and R.F. designed the study, designed the modeling framework, and proved the deriving theorems; J.P. implemented the model and conducted the experiments; J.P. and R.F. interpreted the results and prepared the paper.

Specifically, my contributions include the translation of the existing fiber bundle framework to the biology of the first stages of the low visual pathway, the proofs regarding the smoothness of the different visual activation fields, and their lifts, as well as the PyTorch code for the proposed RetiLeNet CNN model.

Synaptic plasticity has been experimentally studied throughout a number of different cerebral regions, experimental conditions, and, importantly, stimulation regimes. Focusing on the response of the stimulated spine, two different behaviours can be identified, characterized by either an increase or a reduction in synaptic size. Historically, this has led to the splitting of the stimulation protocols into the two reciprocally excluding categories: potentiating and depressing. However, once the spines located in the neighbourhood of the stimulated one (the heterosynaptic spines) are considered, a much more intriguing behaviour emerges. Not only does the stimulus effect spread to this neighbourhood, inducing heterosynaptic plasticity, but the change it induces in heterosynaptic spines can also mimic or antagonize the change induced at the stimulus location. In total, four different plasticity profiles appear [86], going

3.2. The optimality of the multi-spine plasticity profile

from a fully potentiating to a fully depressing one, and passing through two “center-surround” scenarios where the homo and heterosynaptic changes have different signs. A crucial modeling question becomes, therefore, finding a molecular mechanism that is able to jointly support these four different scenarios.

The existence of the two center-surround plasticity profiles is particularly important for this question. Structurally, it suggests that a single mechanism could underlie these different observations, with the four plasticity profiles corresponding to four different degrees of a continuous, cumulative, multi-spine potentiation spectrum. Functionally, moreover, these profiles follow a spatial motif that has been observed across multiple regions of the brain, and, in particular, in relation to the low portion of the visual pathway, in the form of center-surround receptive fields. Many theories have been put forward relating this specific spatial arrangement to a decoding optimality principle, e.g., showing that Gaussian Laplacian (or, almost equally, sums of Gaussians) minimizes the uncertainty when inferring the position and the orientation of a border in a perceived image. In [87], we have also shown that center-surround filter arrangements are learned in the first layers of an extended LeNet 5 CNN model (which we called RetiLeNet), rendering the model invariant, for example, to contrast and brightness variation. This shape was learned via gradient descent, without external constraints being imposed on the learning process. Importantly, while this shape was clearly emerging in the first layers of RetiLeNet, it was not universally present, with deeper convolutional layers showing different arrangements.

Taken together, these observations confirm that if a single biochemical mechanism underlies multi-spine plasticity, it needs to possess *the right amount* of flexibility. If neural computations follow optimal encoding principles, center-surround organizations have to be able to emerge, for example, from the reiterated, additive action of a center-surround plasticity rule (the “Mexican hat” shape described in, e.g., [86] and [88]). For neural computations occurring in a different brain region, however, the optimal synaptic configuration could be different, and the same rule, possibly parametrized by different, region-specific stimulation features, has to be able to allow for its stabilization. In the next section, we will see how these structural insights can be implemented in a rigorous model, leading, ultimately, to recover a number of apparently contradicting results observed throughout the experimental literature.

3.3. Designing the model

After being exposed to a stimulus, spines can show different plasticity responses, potentiating, depressing, or exhibiting no change in their size [89] (Figure 3.1). These responses are dictated by a variety of factors, related to the stimulus (e.g., its intensity, duration, the number of stimulated spines [70,83]), to the spine in consideration (e.g., its distance from the stimulus or its basal size [69,85,90]), and by the dendritic system as a whole [68,91]. Importantly, these responses show a large degree of stochastic variability, with a significant inverse correlation between the synaptic tendency to potentiate and its initial size or weight [83,84,92]. A suitable model for multi-spine plasticity, therefore, needs to implement three main functional components: the features of the stimulus applied to one (or more) spines, the spatiotemporal relationship between the spines and the stimulus, and some information regarding basal synaptic conditions. These three factors can be naturally brought together in a reaction-diffusion framework.

Let's start by considering the synaptodendritic system and, in particular, the protein dynamics that take place throughout its domain (Figure 3.2). Several different molecular processes have been shown to underpin synaptic plasticity, giving rise to spine size dynamics that occur on a range of different spatial and temporal scales [88,93–95]. Among these, the phospho-dephosphorylation process has been classically associated with synaptic plasticity occurring on the minute-to-hour timescale [96–100], and has been successfully used by previous works to gain computational insight on various types of synaptic plasticity [75,101,102].

Consider, therefore, an abstract unphosphorylated synaptic protein resource, referred to as U . This resource is able to diffuse throughout the dendrite (dendritic fraction U_d) and fill the dendritic spines, indexed i (synaptic fraction $U_s^{(i)}$). Within the dendritic spine, two general families of active catalysts are present: kinases (referred to as $K^{(i)}$) and phosphatases (referred to as $N^{(i)}$). These two families regulate the conversion of the synaptic resource $U_s^{(i)}$ into its phosphorylated counterpart, denominated $P^{(i)}$, integrating it into the synaptic structure and leading to the experimentally observed change in synaptic size [62,103].

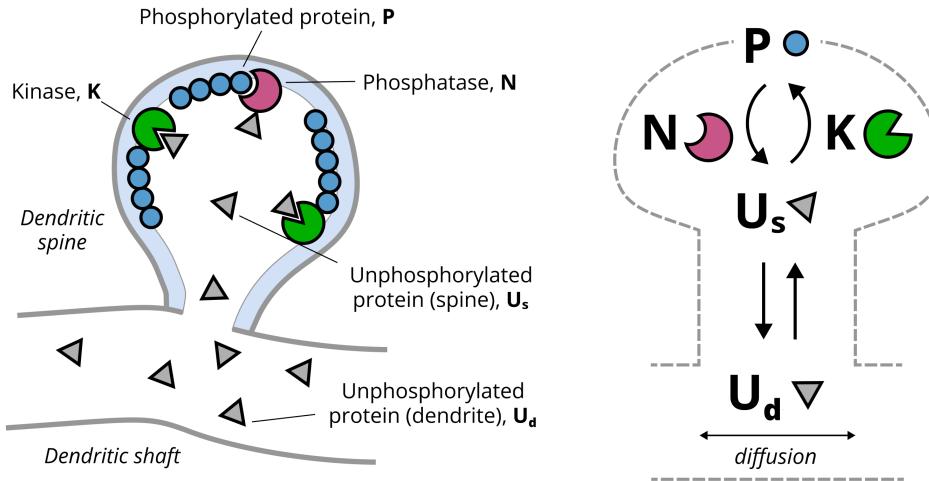


Figure 3.2: **Model overview detailing the molecular dynamics affecting plasticity.** A freely diffusive synaptic resource is able to permeate the spines, where it gets converted to its phosphorylated counterpart, inducing synaptic strengthening and enlargement. The opposite process is also admitted, leading to synaptic depression and increase of the shared unphosphorylated pool.

Following the established reaction-diffusion formalism (e.g. [104]), we now give a rigorous definition of the above-introduced elements. We define the dendrite as the set $D = [0, L] \subset \mathbb{R}$ with length L , with n spines located at positions $x_i \in D, i = 1, \dots, n$. We also consider time as the real valued domain $T = [t_{min}, t_{max}] \subset \mathbb{R}$. On these domains, we define five different molecular families as functions of time and space. We have

- the *dendritic unphosphorylated resource*, defined as positive function $U_d : D \times T \rightarrow \mathbb{R}$;
- the *synaptic unphosphorylated resource*, defined as n positive functions, $U_s^{(i)} : T \rightarrow \mathbb{R}^+$, each describing the temporal evolution of the resource in the spine located at position x_i . Notice that U_d and U_s represent the same molecular species (unphosphorylated resource) but describe its behaviour in different spatial compartments (dendrite and spines);
- the remaining synaptic quantities, i.e. *activated kinases*, *activated phosphatases*, and *phosphorylated synaptic resource* are defined in line with U_s , giving raise, respectively, to the three function families $K^{(i)}$, $N^{(i)}$, and $P^{(i)}$.

We assume the functions defined above present a sufficient degree of smoothness to undergo usual differentiation, as they represent classical, well-behaved physical quantities. Considering a first-order approximation of the Michaelis-Menten catalytic kinetics, the final reaction-diffusion process can be written as

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

$$\begin{cases} \frac{dP^{(i)}}{dt}(t) = k_K K^{(i)}(t) U_s^{(i)}(t) - k_N N^{(i)}(t) P^{(i)}(t) \\ \frac{dU_s^{(i)}}{dt}(t) = k_{in} U_d(t, x_i) - k_{out} U_s^{(i)}(t) - \frac{dP^{(i)}}{dt}(t) \\ \frac{\partial_t U_d}{\partial t}(x, t) = D_U \frac{\partial^2 U_d}{\partial x^2}(x, t) - \sum_{i=1}^n \delta(x - x_i) [k_{in} U_d(x, t) + k_{out} U_s^{(i)}(t)] \end{cases} \quad (3.1)$$

where

- k_K and k_N are the k_{cat}/k_M constant ratios for kinases and phosphatases respectively;
- k_{in} and k_{out} are the exchange constants between spine and dendrite;
- D_U is the diffusion coefficient for the unphosphorylated resource.

The boundary conditions for the partial differential portion of this system will be discussed in the next section. In order to be able to use this model to describe synaptic changes induced by spine stimulation, it is now necessary to introduce the dynamics of the activated catalysts $K^{(i)}$ and $N^{(i)}$. These, evolving in time, drive the observed changes in synaptic size. In accordance with the established literature [57,58,61,68,75], we assume that (i) immediately after induction, the quota of activated catalysts spikes to a new higher amount, (ii) that it then decays exponentially with time to its basal value, and (iii) that the catalyst activation effect depends on the distance from stimulation in a Gaussian-shaped fashion (Figure 3.2.e). The equations describing such dynamics are

$$\begin{cases} K^{(i)}(t) = K_b^{(i)} + \Theta(t - t_s) e^{-\frac{t-t_s}{\tau_K}} \sum_{x_s} K_s e^{-\left(\frac{x_i - x_s}{\sigma_K}\right)^2} \\ N^{(i)}(t) = N_b^{(i)} + \Theta(t - t_s) e^{-\frac{t-t_s}{\tau_N}} \sum_{x_s} N_s e^{-\left(\frac{x_i - x_s}{\sigma_N}\right)^2} \end{cases} \quad (3.2)$$

where t_s and x_s are the time and synaptic locations of plasticity induction, Θ is the Heaviside theta function, τ_K and τ_N are the active catalyst decay timescales, $K_b^{(i)}$ and $N_b^{(i)}$ are the basal values for K and N at each spine i , K_s and N_s are the homosynaptically induced catalyst amounts, and σ_K and σ_N are the spatial induction decay scales (Figure 3.3).

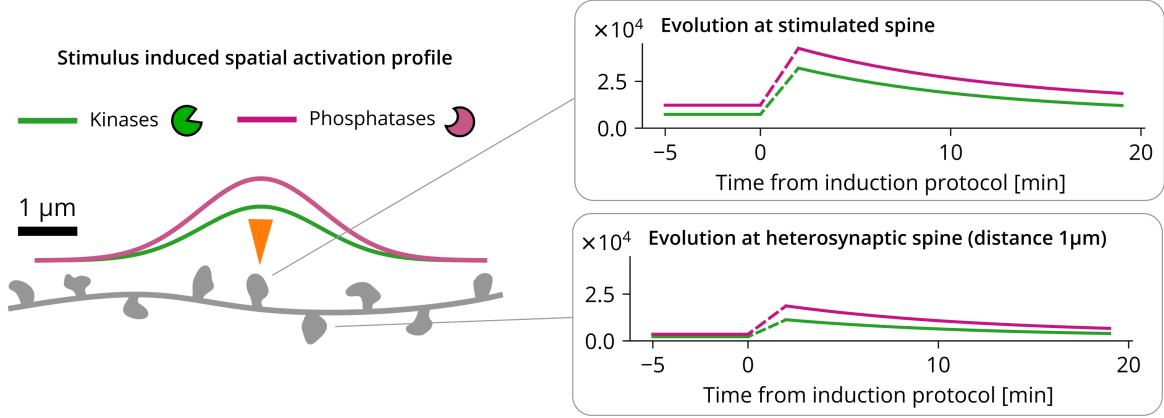


Figure 3.3: **Sketched spatio-temporal dynamics of activated catalyst.** At each spine, starting from a spine-specific basal value, kinases and phosphatases are increased by the stimulation in a distance-dependent fashion. Their dynamics then exponentially decay to the initial value with constants τ_K and τ_N .

For a synaptodendritic system composed of n spines, this model, in its full form, is composed of one partial differential equation and $4n$ ordinary differential equations, counting $9 + 2n$ total parameters. This represents a considerable issue, especially since, in order to validate the model on experimental data, it is necessary to find its optimal parametrization.

3.3.1. Quasi-steady-state approximation and closed form solution

In order to obtain a treatable closed form for the model, we start by noticing that there are two different processes that are taking place simultaneously in the system. One is the pure diffusion of U_d , happening along spine-free regions of the dendritic shaft ($x \neq x_i$), while the other is a combination of diffusion of U_d , its exchange with the synapse and its phosphorylation into and from $P^{(i)}$, happening at $x = x_i$. Since the unbiased influx and outflux of protein resources into the dendrite can be considered extremely low during the time span of the studied phenomenon [63], and since the degradation timescale of dendritic proteins occurs on the timescale of days [63,64], we adopt no-flux boundary conditions for U_d :

$$\partial_x U_d(0, t) = \partial_x U_d(L, t) = 0 \quad (3.3)$$

These boundary conditions allow for the conservation of the total protein resources in time. Indeed, if we define the total resource quantity as the sum of phosphorylated and dephosphorylated families

$$\Pi(t) = \int_D U_d(t, x) dx + \sum_{i=1}^n [U_s^{(i)}(t) + P^{(i)}(t)] \quad (3.4)$$

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

we can see that it does not vary in time

$$\begin{aligned}
\frac{d}{dt} \Pi(t) &= \frac{d}{dt} \left\{ \int_D U_d(t, x) dx + \sum_{i=1}^n \left[U_s^{(i)}(t) + P^{(i)}(t) \right] \right\} \\
&= \int_D \frac{\partial U_d}{\partial t}(t, x) dx + \sum_{i=1}^n \left[\frac{dU_s^{(i)}}{dt}(t) + \frac{dP^{(i)}}{dt}(t) \right] \\
&= \int_{D \setminus \{x=x_i\}} \frac{\partial U_d}{\partial t}(x, t) dx \\
&= [\partial_x U_d(x, t)]_0^L = 0
\end{aligned} \tag{3.5}$$

where the last step is justified by the finite cardinality of the possible discontinuity points $\{x = x_i\}$. This is an important property that will allow us later to obtain a closed-form expression for $P^{(i)}(t)$.

We now notice that multiple timescales are involved in the process. The phosphor-dephosphorylation of U_s and P , as well as the diffusion of U_d , can be considered fast compared to the deactivation times of the activated catalysts K and N [57,105,106]. Conversely, the timescale at which the amount of activated catalysts decays corresponds to the timescale at which the experiment used in this work was conducted (tenths of minutes, [57]) and can be considered the leading timescale. We can, therefore, conduct a quasi-steady-state approximation of the initial equations, obtaining

$$\begin{cases} P^{(i)}(t) = \frac{k_K K^{(i)}(t)}{k_N N^{(i)}(t)} U_s^{(i)} \\ U_s^{(i)}(t) = \frac{k_{in}}{k_{out}} U_d(t) \\ \frac{\partial^2 U_d}{\partial x^2}(x, t) = 0 \end{cases} \tag{3.6}$$

The equation for $U_d(x, t)$ has become a one-dimensional Laplace equation which, considering the no-flux boundary conditions and the continuity of U_d itself, has constant solution in space, for a given time:

$$U_d(x, t) = \bar{U}_d(t) \tag{3.7}$$

Therefore, the amount of unphosphorylated dendritic resource at a given time will be equal to

$$\int_D \bar{U}_d(t) dx = L \bar{U}_d(t) \tag{3.8}$$

The exchange dynamics of unphosphorylated resources between spines and dendrite depend only on the geometric features of the spine neck. Since these can only favour or hinder molecular passage symmetrically, we can set $k_{in} = k_{out}$, and therefore $U_s^{(i)}(t) = \bar{U}_d(t)$.

We can now set about finding the closed form for $P^{(i)}(t)$ following a procedure similar to the one exposed in [107]. We start by recalling the definition of the total resources (3.4), and noticing that we can explicitly separate it into its dendritic and synaptic component

$$\Pi = \int_D U_d(t, x) dx + \sum_{i=1}^n [U_s^{(i)}(t) + P^{(i)}(t)] \quad (3.9)$$

By substituting (3.6) and (3.8) we obtain

$$\begin{aligned} \Pi &= L \bar{U}_d(t) + n \bar{U}_d(t) + \sum_{i=1}^n \frac{k_K}{k_N} \frac{K^{(i)}(t)}{N^{(i)}(t)} \bar{U}_d \\ &= \left(L + n + \frac{k_K}{k_N} \sum_{i=1}^n \frac{K^{(i)}(t)}{N^{(i)}(t)} \right) \bar{U}_d(t) \end{aligned} \quad (3.10)$$

which yields an expression of U_d in function of the total synaptic catalysts $K^{(i)}$ and $N^{(i)}$:

$$\bar{U}_d(t) = \frac{\Pi}{L + n + \sum_{i=1}^n \frac{k_K}{k_N} \frac{K^{(i)}(t)}{N^{(i)}(t)}} \quad (3.11)$$

This expression can be now substituted into the equation (3.6) for $P^{(i)}$, obtaining

$$P^{(i)}(t) = \Pi \frac{\frac{K^{(i)}(t)}{N^{(i)}(t)}}{\frac{k_N}{k_K} (L + n) + \sum_j \frac{K^{(j)}(t)}{N^{(j)}(t)}} \quad (3.12)$$

As a final step, we define the linear spine density $\lambda = n/L$ and the subsequent “geometric” constant $\Omega = \frac{k_N}{k_K} (1 + \lambda)L$, and split the sum over j , obtaining

$$P^{(i)}(t) = \Pi \frac{\frac{K^{(i)}(t)}{N^{(i)}(t)}}{\Omega + \frac{K^{(i)}(t)}{N^{(i)}(t)} + \sum_{j \neq i} \frac{K^{(j)}(t)}{N^{(j)}(t)}} \quad (3.13)$$

It is important to remark that the obtained equation is a quasi-steady-state solution to the general system (3.1). This implies that the predicted synaptic trajectories will be sufficiently accurate only after the quick processes have had time to relax. For a dendritic stretch of roughly $L = 70 \mu m$, this time is approximately $L/2D_u \sim 2 \text{ min}$. This fact, together with our focus on dynamics happening on the timescale of roughly one hour, allow us to consider a plasticity-inducing protocol lasting one minute as a “single stimulation” at $t = t_s$. An example

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

representation of the model's (3.13) dynamics in its validity regime is shown Figure 3.4, where plasticity induction is carried out on the central spine of a small dendritic stretch. We can see that both homosynaptic potentiation and heterosynaptic depression emerge in a Mexican-hat-like fashion [69,89], due to the non-linear interplay between basal synaptic catalyst distributions and the deterministic distance-dependent induction provided by the stimulus. Ultimately, progressing with time, the induced changes vanish, and the system returns to its basal configuration.

In every spine, a dynamical equilibrium is present between the synaptic $P^{(i)}$ and $U^{(i)}$, and this equilibrium is strictly dependent on the ratio of synaptic kinases $K^{(i)}$ and phosphatases $N^{(i)}$. This gives rise to important structural consequences, which can be immediately recognized in the (3.13). First, different basal contents of activated kinases and phosphatases directly translate to different spine sizes. Second, changes of the amounts of synaptic kinases and phosphatases induce a *local* change of synaptic size, while a change in the total amount of resources Π translates to a *global*, multiplicative change of all synaptic sizes. Thirdly, since (3.13) is a homographic function of the ratio $K_b^{(i)}/N_b^{(i)}$ for every spine, synaptic sizes are prevented from undergoing runaway dynamics, as $P^{(i)}$ can at most reach the value of Π , when the ratio $K^{(i)}/N^{(i)}$ tends to infinity. From a biochemical standpoint, this corresponds to a scenario where all the available resources have been segregated in the i -th spine, depriving the rest of the system.

In order to utilize this model, it is necessary to find the optimal value for its parameters via model optimization. To this end, we follow the procedure described in Section 3.4, using the synaptic size dynamics studied in [68]. This choice derives from the widely observed correlation between a spine's volume, its (phosphorylated) protein content and, ultimately, its weight [67,108,109]. Since the model admits the closed form (3.13) an explicit likelihood can be computed and minimized, thus avoiding the necessity of numerical integration and computationally expensive sampling techniques.

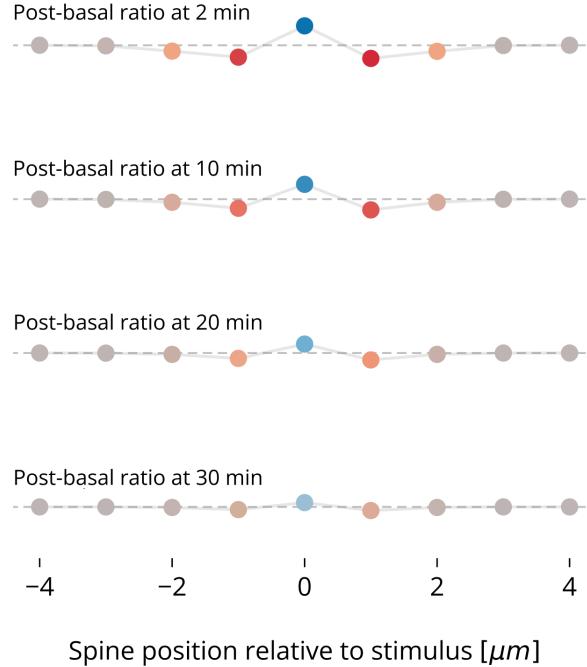


Figure 3.4: **Plasticity dynamics simulated in space and time.** The ratio $P^{(i)}(t)/P_{\text{base}}^{(i)}$ is plotted for each spine at different time snapshots. According to our model, this value resembles the evolution of synaptic sizes (cfr. next sections).

3.4. Optimization

3.4.1. Structural identifiability

The model (3.13) describes the evolution of each spine on a dendrite subject to a specific plasticity induction protocol. Consider now a scenario where $m = 1, \dots, M$ different induction protocols have been carried out on different neurons. Each of these scenarios entails different stimulation locations and different dendrites, with different geometric features, different quantities of available resources and different basal initial conditions. For each condition, our model yields

$$P_m^{(i)}(t) = \Pi_m \frac{\alpha_m^{(i)}(t)}{\Omega_m + \sum_{j=1}^{n_m} \alpha_m^{(j)}(t)}, \quad m = 1, \dots, M \quad (3.14)$$

where, assuming that the stimulation protocol happens at $t = 0$

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

$$\alpha_m^{(i)}(t) = \frac{K_{b,m}^{(i)} + \Theta(t) K_s e^{-\frac{t}{\tau_K}} \sum_{\bar{x} \in \mathcal{X}_m} e^{-\left(x_i - \frac{\bar{x}^2}{\sigma_K^2}\right)}}{N_{b,m}^{(i)} + \Theta(t) N_s e^{-\frac{t}{\tau_N}} \sum_{\bar{x} \in \mathcal{X}_m} e^{-\left(x_i - \frac{\bar{x}^2}{\sigma_N^2}\right)}} \quad (3.15)$$

The number of observed spines for each protocol (n_m) and the locations of the stimulations ($\mathcal{X} = \{\bar{x}_1, \dots, \bar{x}_1\}$) are available a priori from the experimental procedure. The other parameters used in the equations, however, are not known and need to be recovered via data fitting. In particular, we observe that two families of parameters have to be inferred (Table 3.1):

1. *global parameters*, driving the plasticity dynamics independently of the specific experimental realization. These are τ_K, τ_N, σ_K , and σ_N ;
2. *specific parameters*, which depend on a specific experimental setup. These are the spine initial conditions ($K_{b,m}^{(i)}, N_{b,m}^{(i)}$), the available resources Π_m , and the dendritic geometric factor Ω_m . This parameter category is indexed by the protocol index m .

Global parameters	Number
Timescales	τ_N, τ_K
Length scales	σ_K, σ_N
Stimulus contribution	K_s, N_s

Specific parameters	Number
Total resources	Π_m
Dendritic geometric factor	Ω_m
Initial kinases per spine	$K_{b,m}^{(i)}$
Initial phosphatase per spine	$N_{b,m}^{(i)}$

Table 3.1: **Model parameters necessary for the simultaneous optimization of M protocols with n_m spines each.** The specific parameters are indexed by the protocol index m , and describe quantities that can differ between experiments, like the total available resources in the dendrite, the geometry of the dendrite and the synaptic catalyst distributions at stimulus time.

The total number of parameters that have to be inferred when fitting data arising from M experimental protocols amounts to $6 + 2M + 2\sum_m n_m$. It is not possible, however, to fit all these parameters together as the proposed model is not structurally identifiable due to the rescaling invariance for $P_m^{(i)}$ and $\alpha_m^{(i)}$ in (3.14) and (3.15):

$$\begin{aligned} P_m^{(i)}(\Pi_m, \alpha_m^{(1)}(t), \dots, \alpha_m^{(n_m)}(t), \Omega_m) &= P_m^{(i)}(\Pi_m, s \alpha_m^{(1)}(t), \dots, s \alpha_m^{(n_m)}(t), s \Omega_m), \quad s \in \mathbb{R} \\ \alpha^{(i)}(t; K_{b,m}^{(i)}, K_s, N_{b,m}^{(i)}, N_s) &= \alpha^{(i)}(t; s K_{b,m}^{(i)}, s K_s, s N_{b,m}^{(i)}, s N_s), \quad s \in \mathbb{R} \end{aligned} \quad (3.16)$$

In other words, the full optimization problem is underdetermined. This can be solved by introducing two additional constraints, one for each of the two degrees of freedom. Defining, for clarity, the “fractionary index” symbols for the catalysts

$$K_{b,m/M}^{(i)} = \frac{K_{b,m}^{(i)}}{\Omega_M N_{b,M}^{(n_M)}}, \quad N_{b,m/M}^{(i)} = \frac{N_{b,m}^{(i)}}{N_{b,M}^{(n_M)}}, \quad K_{s/M} = \frac{K_s}{\Omega_M N_{b,M}^{(n_M)}}, \quad N_{s/M} = \frac{N_s}{N_{b,M}^{(n_M)}} \quad (3.17)$$

and for the dendritic parameter

$$\Omega_{m/M} = \frac{\Omega_m}{\Omega_M} \quad (3.18)$$

we can rewrite the problem equations (3.14)

$$\begin{cases} P_m^{(i)}(t) = \Pi_m \frac{\alpha_m^{(i)}(t)}{\Omega_{m/M} + \sum_{j=1}^{n_m} \alpha_m^{(j)}(t)}, & m = 1, \dots, M-1 \\ P_M^{(i)}(t) = \Pi_M \frac{\alpha_M^{(i)}(t)}{1 + \sum_{j=1}^{n_M} \alpha_M^{(j)}(t)}, & m = M \end{cases} \quad (3.19)$$

with

$$\begin{aligned} \alpha_m^{(i)}(t) &= \frac{K_{b,m/M}^{(i)} + K_{s/M} e^{-\frac{t}{\tau_K}} \sum_{x^- \in \mathcal{X}_m} e^{-\frac{(x_i - x^-)^2}{\sigma_K^2}}}{N_{b,m/M}^{(i)} + N_{s/M} e^{-\frac{t}{\tau_N}} \sum_{x^- \in \mathcal{X}_m} e^{-\frac{(x_i - x^-)^2}{\sigma_N^2}}}, \quad m = 1, \dots, M-1 \\ \alpha_M^{(i)}(t) &= \frac{K_{b,m/M}^{(i)} + K_{s/M} e^{-\frac{t}{\tau_K}} \sum_{x^- \in \mathcal{X}_m} e^{-\frac{(x_i - x^-)^2}{\sigma_K^2}}}{1 + N_{s/M} e^{-\frac{t}{\tau_N}} \sum_{x^- \in \mathcal{X}_m} e^{-\frac{(x_i - x^-)^2}{\sigma_N^2}}}, \quad m = M \end{aligned} \quad (3.20)$$

In this fashion, we set Ω_M and $N_{b,M}^{(n_M)}$ to 1 and introduce the two necessary constraints to render the optimization problem well posed. We can then proceed to identify a biologically plausible value range for the remainder of the parameters (Table 3.2) and carry out the inference for (3.19). As a final step, in accordance with the previous literature [67], we select a reasonable value for the two constraints Ω_M and $N_{b,M}^{(n_M)}$, and recover in respect to that the optimal values for the full optimization problem (3.14).

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

Parameter	Unit	BVs	Notes
$K_{b,m/M}^{(i)}$	-	0.002 – 0.12	Supposing 4% of CaM is active basally, and equally split between K and N [57,66,67]
$K_{s/M}$	-	0.07 – 3.71	Assuming that CaM is the limiting activation factor [57,66,67,101]
$N_{b,m/M}^{(i)}$	-	0.28 – 3.5	Using the same argument as in $K_{b,m/M}^{(i)}$ [57,66,67]
$N_{s/M}$	-	0.07 – 3.71	Assuming that calcineurin is the limiting factor [57,66,67]
τ_K	min	0 – 100	Covering bulk, local and reciprocal kinase-effector activation timescales [57]
τ_N	min	0 – 100	Range as in τ_K
σ_K	μm	0 – 100	Wide parameter range
σ_N	μm	0 – 100	Wide parameter range
$\Omega_{m/M}$	-	0.3 – 3.33	Observations in [67] and general dendritic statistics

Table 3.2: **Parameter boundary values (BVs) used in the optimization**

3.4.2. Definition and pre-processing of the target values

The equations (3.14) describe, for every protocol $m = 1, \dots, M$, the dynamics of the n_m spines located on one dendrite. Since in our experimental data [68] each protocol is reproduced multiple times on different cells, we have to build a dataset containing one statistically representative dendrite for each protocol. For a given experimental protocol m , uniquely defined by the stimulus \mathcal{X}_m , we proceed as follows (a graphical summary is shown in S 4.1):

1. **image preprocessing:** utilizing in-house developed software [110], we segment the images, extracting for each spine belonging to every cell the “raw integrated density” values (a proxy for synaptic sizes), as well as the relative temporal and spatial coordinates. We then conduct statistical analysis (baseline correction) on the obtained values to verify that the initial conditions are not significantly different between different cells as well as between stimulated and non-stimulated spines;
2. **distance assignment:** for each cell, we assign to each synapse its distance value from the nearest stimulation. Importantly, if the spine is located between two stimulations, we assign it a negative distance value. We assign a distance value of 0 to the stimulated spines;
3. **average spine density:** for the model (3.14) to reproduce the studied dynamics, we have to consider the correct number of spines located on the dendrite. We infer this parameter by

estimating the mean inter-spine distance across all the cells used for each protocol. Details and benchmarking of the estimation procedure are reported in Appendix B;

4. **binning:** for each protocol m , we generate a statistically representative dendrite with spines located regularly at the inferred inter-spine distance, and size values equal to the luminosity average across different cells of the corresponding spatial bin. To avoid overlapping, we choose a bin width equal to the inferred inter-spine distance.

In this fashion, we obtain a dataset consisting of 6 raw integrated density fields, one for each protocol, effectively describing the synaptic size evolution before and after plasticity induction. We use 4 of these datasets for model optimization (1, 3, 5, and 7 clustered stimulation protocols under control conditions), and leave out one (7 distributed stimulation protocols under control conditions and 7 clustered stimulation protocols under FK506 2 μM condition) for model validation.

One final factor that needs to be considered before proceeding with the optimization, is that the datasets obtained with this procedure contain a reduced amount of spines (Fig. S 4.1). This would lead, during the fitting, to an underestimation of the factor $\sum_j \alpha^{(j)}(t)$, in the denominator of (3.14). To correctly account for all the spines in the experimental dataset, we introduce a differential weighting of the spines depending on their type (stimulated, inside of the stimulation cluster and outside of the stimulation cluster), and compute the sum as

$$\sum_{j=1}^{n_m} \alpha^{(j)} \simeq |\mathcal{X}_m| \alpha^{(0)} + 2 \sum_{j \in \text{inside}} \alpha^{(j)} + (2 - |\mathcal{X}_m|) \sum_{j \in \text{outside}} \alpha^{(j)} \quad (3.21)$$

where $|\mathcal{X}_m|$ is the cardinality of \mathcal{X}_m , i.e. the number of stimulations for the protocol m .

3.4.3. Likelihood maximization

To perform parameter estimation, we used the maximum likelihood estimate (MLE) approach. The optimization problem is given by

$$\theta_{MLE} = \arg \min_{\theta} -\log \mathcal{L}(\mathcal{D}|\theta) \quad (3.22)$$

where \mathcal{L} is the probability of observing measurements $\bar{y} \in \mathcal{D}$ given some system parameterized by θ , which contains the unknown parameters given in Table Table 3.2. We minimize the negative logarithm of \mathcal{L} for numerical efficiency.

The likelihood requires an observation model $\mathbf{h} : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_y}$ that maps the model space to the data space. The spine measurements are noisy, hence the observables $\mathbf{y} = \mathbf{h}(x(t, \theta), \theta)$ are re-

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

lated to data by a noise model $\bar{y}(t_i, y_i) = y_{t_i, y_i} + \epsilon_{t_i, y_i}$, where $t_i \in \{1, \dots, n_t\}$ and $y_i \in \{1, \dots, n_y\}$ index the time points observables, respectively, and ϵ_{t_i, y_i} is the measurement-specific noise.

Assuming i.i.d. Student's t-distributed noise ϵ_{t_i, y_i} with standard deviations σ_{t_i, y_i} computed from the standard error of the mean in the data and with d degrees of freedom, and ignoring constants that do not affect the location of optima in the parameter estimation problem, the negative log-likelihood is given by

$$\mathcal{J}(\theta) = -\log \mathcal{L}(\mathcal{D} \mid \theta) = \sum_{t_i, y_i} \frac{d_i + 1}{2} \log \left\{ 1 + \frac{1}{d_i} \left(\frac{\bar{y} - y(\theta)}{\sigma_{t_i, y_i}} \right)^2 \right\} \quad (3.23)$$

We used d_i equal to the number of data points available for each spatial bin x_i . The model and \mathcal{J} were implemented in PyTorch [111], and the Python Parameter Estimation Toolbox (pyPESTO) was used with the Fides optimizer to perform parameter estimation [112,113]. We log-transformed parameters for efficiency, and bounds are given in Table 3.2. We used multi-start, gradient-based optimization, with 1200 starts and with gradients computed in PyTorch.

3.4.4. Goodness of fit

We start investigating the quality of the optimization by looking at the final convergence values of each of the runs (Figure 3.5). As shown in the waterfall plot (Figure 3.5.a), there are two wide local minima (LM1 and LM2) in the parameter space, with 437 and 656 converging runs, respectively. The minima show almost identical final criterion values ($fval_1 = 108.88$, $fval_2 = 108.96$), as well as very similar inferred values for the parameters. Interestingly, the only substantial difference seems to reside in the estimates for the decay constants τ_K and τ_N , with LM2 converging on significantly higher values.

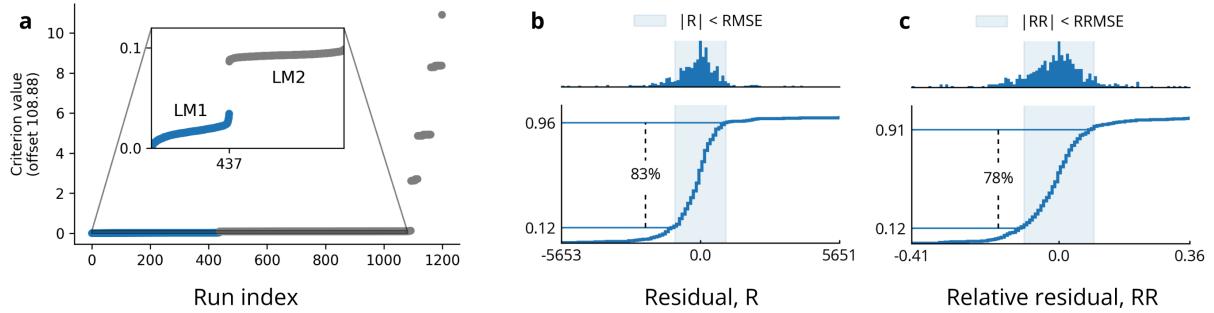


Figure 3.5: Model optimization summary. **a** Cascade plot concerning the 1200 optimization runs. Two (almost equivalent) local minima are found (LM1 and LM2) with convergence basins covering the majority of the available parameter space. **b** Absolute residual (R) probability density (upper graph) and cumulative distribution functions (lower graph) for the best optimization run; 83% of the residuals fall, in absolute value, under the root mean squared error (RM SE = 1034.51). **c** Same analysis concerning relative residuals (RR). Importantly, a value smaller than 10% is achieved for the relative root mean squared error (RRM SE = 9.72%), a hallmark of good model convergence.

Following standard practice, we pick the best run converging to the best local minimum as providing the best estimate for the model parameters. The resulting fits to data, shown in S 4.2, display remarkable qualitative accuracy, correctly reproducing both homosynaptic potentiation at $x^{(i)} = 0.00 \mu m$ and heterosynaptic depression, when present. Quantitatively, the model achieves a 9.72% relative root mean squared error (RRMSE), with narrow and symmetric residual and relative residual distributions (Figure 3.5). Remarkably, the same metrics were achieved using values for τ_K and τ_N deriving from the best run of LM2. Given their better compatibility with the investigated timescales, we used these values in our subsequent simulations.

As a final step in the optimization evaluation, we estimate the posterior distributions of the inferred parameters. Due to the number of parameters involved, we carry out a full Markov chain Monte Carlo sampling only for the six global parameters (S 4.3, adaptive parallel tempered sampling [114]). For the remaining parameters corresponding to the rescaled basal values $K_{b,m/M}^{(i)}$ and $N_{b,m/M}^{(i)}$, we inspect the convergence values of the 437 runs belonging to LM1, and verify that indeed, with very few exceptions, they follow a well behaved distribution (S 4.4).

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

3.5. Inferred and observed log-normality of synaptic features

In the previous section, we carried out the parameter optimization for the model (3.19), and with the aid of the experimental literature, we were able to fully determine the specific parametrization of the problem (3.13). In order to be able to utilize the general model for synaptic plasticity simulation, it is first necessary to find a suitable statistical description for the *spine-specific parameters*, as these correspond to stochastic quantities of which each neuron represents just one specific instantiation. We posit that the variability of these components plays a crucial role both in synaptic size distribution and in the outcome of the plasticity experiment, as it inevitably impacts the strength with which a spine is able to compete for resources.

To this end, we start by analyzing the distributions of the inferred values for $K_b^{(i)}$ and $N_b^{(i)}$, both independently and in a joint fashion. Log-normal patterns have been observed in relation to a number of synaptic metrics [115], suggesting that log-normality can arise from a fundamental process driving the general dynamics taking place inside the neuron. Importantly, [71,81] also show that log-normality can be uncoupled from neuronal activity, arguing that its emergence could be driven not to optimal information encoding but caused mechanistically by a different process.

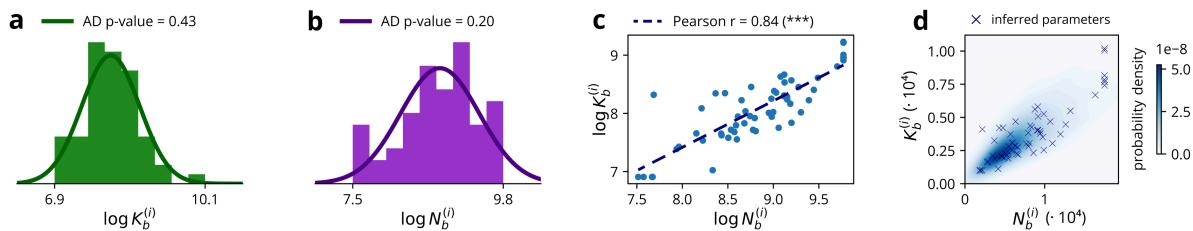


Figure 3.6: Joint and marginal lognormality of inferred synaptic catalytic distribution. **a,b** The inferred optimal values of synaptic $K_b^{(i)}$ ($N = 59$) and $N_b^{(i)}$ ($N = 58$) each follow a log-normal compatible distribution. **c** The logarithms of the inferred synaptic catalyst amounts show a strong, significant linear correlation. **d** Graphical representation of the bivariate log-normal describing the distribution of the estimates of $K_b^{(i)}$ and $N_b^{(i)}$.

Given that $K_b^{(i)}$ and $N_b^{(i)}$ represent the synaptic amount of molecular species undergoing a plethora of dynamical diffusive and transport processes, a natural probability density that could describe their synaptic distribution is the log-normal distribution. We test this hypothesis and find that, indeed, the inferred basal catalytic values do not show statistically significant difference from log-normal distributions (Figure 3.6 a,b). Moreover, the respective joint distribution shows a substantial degree of correlation (Figure 3.6 c), with a highly signif-

3.5. Inferred and observed log-normality of synaptic features

icant Pearson r value of 0.84. These two findings (log-normality and correlation between $K_b^{(i)}$ and $N_b^{(i)}$), together with the absence of evident pathological features (visual evaluation) allow us to assume that the overall distribution of these two catalyst families can be modeled as a bivariate log-normal distribution. From our model's perspective, each spine can therefore identified with a pair of values sampled from this bivariate distribution.

In order to validate the reliability of the log-normal function for describing synaptic molecular distributions, we turn our attention to the experimental data presented in [67], where confocal and super-resolution microscopy is conducted on several synaptic targets in cultured hippocampal neurons. Encouragingly, we find that a number of different plasticity-related proteins show a log-normal compatible distribution across spines, together with a high correlation with the scaffolding protein Homer1 (S 4.5). These features are also displayed by two critical synaptic catalysts, i.e. Ca²⁺/calmodulin-dependent protein kinase II (CaMKII) and Calcineurin (CaN) (Figure 3.7), which can be considered a major counterpart of the model's kinase and phosphatase families.

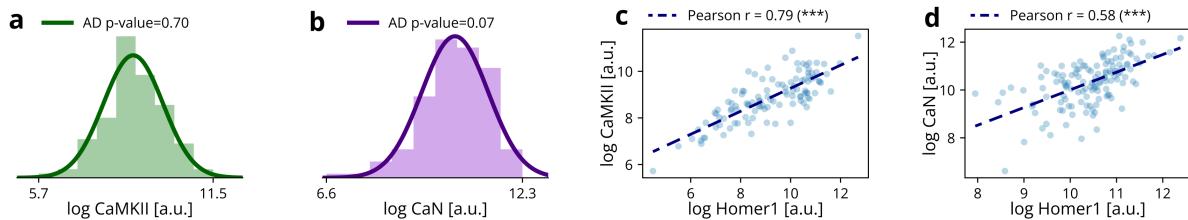


Figure 3.7: Experimentally observed catalyst statistical features. **a,b** Ca²⁺/calmodulin-dependent protein kinase II (CaMKII, $N = 110$) and Calcineurin (CaN, $N = 133$) distribution across spines. Both result compatible with a log-normal distribution. **c,d** Correlation between synaptic Homer1 and synaptic CaMKII/CaN signals respectively.

Importantly, in accordance with [116], the correlation observed between CaMKII/CaN and Homer1 allows us to derive theoretical bounds for the Pearson correlation between CaMKII and CaN. These bounds, evaluating to [0.05, 0.92], result compatible with the correlation inferred for the values of $K_b^{(i)}$ and $N_b^{(i)}$.

So far, we have observed that our optimization routine has identified the log-normality of the distribution of basal amounts of synaptic catalysts and phosphatases. Recurring to experimental data [67], we have confirmed that this distribution is indeed retrievable for a number of plasticity-related proteins, including, in particular, major catalyst families (CaMKII and Calcineurin). To provide a definitive corroboration for our choice of distribution, we now ask ourselves if log-normality (which, so far, was just not excluded, but never confirmed as a “true

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

distribution") can be mechanistically explained with a fundamental mechanism underpinning the dynamics of a wide family of synaptic proteins.

3.6. Log-normality emergence from molecular dynamics

To investigate how log-normality can emerge from a fundamental dynamic process, we resort to the stochastic simulation of a linear dendrite where an abstract molecule is able to diffuse throughout the dendritic shaft, exiting and entering the dendritic spines. Under general conditions, we find that the resulting synaptic distribution of this diffusive resource is undistinguishable from a log-normal distribution. Consider a one-dimensional dendrite of length L with spines located at positions $x = x_i, i = 1, \dots, N$. Consider now an abstract molecular family able to diffuse throughout the dendrite, entering and exiting the dendritic spines. Let now $m_d(x, t)$ be the number of molecules at time t in the dendritic section $[x, x + dx]$, and let $m_s^{(i)}(t)$ be this number in the spine connected to the dendrite at x_i . We can write this system as

$$\begin{cases} \frac{dm_s^{(i)}}{dt}(t) = k_{in}^{(i)} m_d(x_i, t) - k_{out}^{(i)} m_s^{(i)}(t) \\ \frac{\partial m_d}{\partial t}(x, t) = D_m \frac{\partial^2 m_d}{\partial x^2}(x, t) - \sum_{i=1}^N \delta(x - x_i) \frac{dm_s^{(i)}}{dt}(t) \end{cases} \quad (3.24)$$

where D_m is the diffusion constant of the considered molecule and $k_{in}^{(i)}$ and $k_{out}^{(i)}$ are the spine-dendrite exchange rate constants pertaining to the i -th spine, depending on a number of synaptic and neck features (e.g., width, shape, synaptic crowding and confinement). We associate to this system the no flux boundary conditions $\partial_x m_d(0, t) = \partial_x m_d(L, t) = 0$, as we assume that the influx rate in and outside of the dendrite is low enough compared to the diffusion constant D_m . Under these conditions, the system admits the steady-state solution

$$\begin{cases} m_s^{(i)} = \frac{k_{in}^{(i)}}{k_{out}^{(i)}} m_d \\ m_d(x_i) = m_d \end{cases} \quad (3.25)$$

with a constant value of m_d throughout the dendrite, and $m_s^{(i)}$ proportional to this value through the ratio of the exchange constants at the i -th spine. The solution of our original problem is hereby represented by the histogram of the values $m_s^{(i)}$, corresponding to the distribution density of the ratio $k_{in}^{(i)}/k_{out}^{(i)}$. A crucial step of this derivation is, therefore, the choice of the probability distribution from which the exchange constants are sampled. We examine, in the next paragraphs, the effect of different plausible choices of this distribution.

Log-normal distribution A number of different synaptic and neck features have been found to follow a log-normal distribution, so this could represent a first reasonable choice for the statistics of $k_{in}^{(i)}$ and $k_{out}^{(i)}$. From such a scenario, however, it trivially follows that $m_s^{(i)}$ would also follow a log-normal distribution, as the ratio of two log-normally distributed variables is also described by a log-normal distribution. This fact highlights the stability of log-normality in systems with multiplicative noise (as ours), but does not provide insight into how log-normality emerges in the first place from a less generous, possibly additive, noise. We choose, therefore, to move on to more general distribution choice, focusing in particular on two *symmetric* families.

Normal distribution This distribution represents an appealing choice, due to its ability to naturally emerge in distributions subject to general additive noise. The ratio of normal variables, however, does not have a universally defined, well-behaved distribution and strongly depends on the parameters driving the dividend and the divisor. There is, however, one physical constraint that leads to interesting properties. By definition, $k_{in}^{(i)}$ and $k_{out}^{(i)}$ have to be positive quantities, as they describe positively defined flux rates, into and from the spines. To respect this condition, their distributions will have to have positive means and small enough variances to render the probability of a negative sample negligible. In other words, they have to have a small enough coefficient of variation (CV). Under these conditions, it can be shown that the ratio of two Gaussian variables can be well approximated by a log-normal distribution, with the quality of the approximation decreasing with the magnitude of the coefficients of variation. This approximation, known in the statistical literature as *delta method*, is valid in general for every function of a narrowly distributed normal random variable. In our case, the derivation starts by considering the variable $Z = \log X/Y$, and approximating it via the Taylor expansion up to the second order

$$\log\left(\frac{X}{Y}\right) \approx \log\left(\frac{\mu_X}{\mu_Y}\right) + \frac{X - \mu_X}{\mu_X} + \frac{Y - \mu_Y}{\mu_Y} - \frac{1}{2} \left[\left(\frac{X - \mu_X}{\mu_X}\right)^2 + \left(\frac{Y - \mu_Y}{\mu_Y}\right)^2 \right] \quad (3.26)$$

Remembering that $(X - \mu_X)/\mu_X$ and $(Y - \mu_Y)/\mu_Y$ follow independent normal distributions with mean 0 and variances CV_X^2 and CV_Y^2 respectively, we can find the first central momenta of Z up to the leading order

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

$$\begin{aligned}
\mathbb{E}[Z] &\approx \log\left(\frac{\mu_X}{\mu_Y}\right) \\
\mathbb{E}[(Z - \mathbb{E}[Z])^2] &\approx CV_X^2 + CV_Y^2 \\
\mathbb{E}[(Z - \mathbb{E}[Z])^3] &\approx CV_X^3 + CV_Y^3 \\
&\dots
\end{aligned} \tag{3.27}$$

the general formula for the n -th moment being

$$\begin{aligned}
\mathbb{E}[(Z - \mathbb{E}[Z])^n] &\approx \mathbb{E}\left[\left(\tilde{X} - \frac{\tilde{X}^2}{2} + \tilde{Y} - \frac{\tilde{Y}^2}{2}\right)^n\right] \\
&= \sum_{k=0}^n \binom{n}{k} \left\{ \sum_{j=0}^k \binom{k}{j} \left(-\frac{1}{2}\right)^j \mathbb{E}[X^{k+j}] \sum_{l=0}^{n-k} \binom{n-k}{l} \left(-\frac{1}{2}\right)^l \mathbb{E}[Y^{n-k+l}] \right\}
\end{aligned} \tag{3.28}$$

and allowing to show that in general the n -th moment is of the order $\mathcal{O}(\max\{CV_X, CV_Y\}^n)$. For small values of CV_X and CV_Y the distribution Z will, therefore, be well approximated by a normal distribution, and the ratio X/Y will consequently be compatible with a log-normal distribution - independently of the mean for $k_{in}^{(i)}$ and $k_{out}^{(i)}$. To give a quantitative evaluation of this compatibility, we resort to computational sampling. For every pair of coefficients of variation CV_{in} and CV_{out} taken in the range $[0.01, 0.22]$, we extract $n_{spines} = 1000$ “synaptic” values of $k_{in} \sim \mathcal{N}(0.5, 0.5 CV_{in})$ and $k_{out} \sim \mathcal{N}(0.2, 0.2 CV_{out})$, arbitrarily picking the mean values as they do not impact the approximation. We then test the resulting log-ratio $\log k_{in}/k_{out}$ for normality using the Anderson-Darling test, rejecting compatibility for $p < 0.05$. We repeat this procedure $n = 1000$ times and report the final ratio of log-normal compatible samples for each value of CV_{in} and CV_{out} (Figure 3.8).

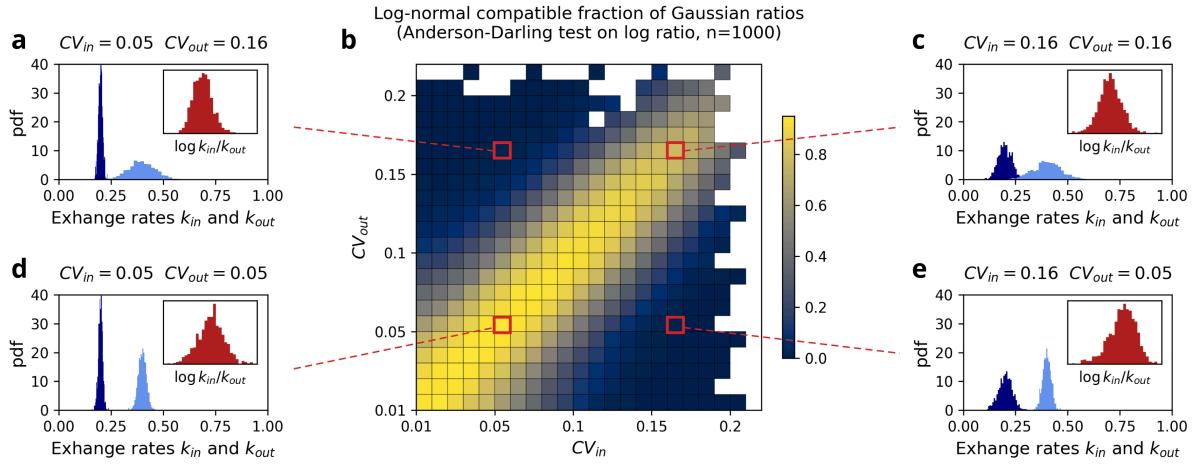


Figure 3.8: **Log-normal compatibility of ratio distribution generated by normally distributed k_{in} and k_{out} .** For each pair of coefficients of variation of k_{in} and k_{out} (referred to as CV_{in} and CV_{out}) we sample $n = 1000$ ratios k_{in}/k_{out} and compute the log-normal compatible fraction by testing the log-ratio $\log k_{in}^{(i)}/k_{out}^{(i)}$ for normality with the Anderson-Darling test (panel **b**). As a significance threshold, we choose $p = 0.05$. To give better intuition of the results, we also plot for realizations of different CV values (panels **a,c,d,e**). The white squares in panel **b** represent instances where sampling of k_{in} and k_{out} produced values smaller than zero, automatically falsifying log-normal compatibility due to ratios not being positive.

As expected, in this range of coefficients of variation, the log-normal compatible fraction results remarkably high (Figure 3.8 b,d,c). We do notice, however, that if the two CVs differ by more than 5%, this compatibility drops considerably, with log-ratio distributions acquiring pronounced left or right tails (Figure 3.8.a,e). One final source of non compatibility emerges when the values of the variation coefficients increase beyond ~ 0.15 : the sampled k_{in} and k_{out} start including negative values, rendering the ratio distribution structurally incompatible with a log-normal distribution (which by definition has positive support). This leads to conclude that, to model a scenario where the exchange rates are driven by high coefficients of variation, the Gaussian distribution is not a valid modeling choice. To complete our analysis, we therefore switch to another symmetric distribution, the Beta distributions, which is known to be a reasonable approximation of the Gaussian distribution for low CVs, but can also behave as a high CV probability density while maintaining symmetry and a positive, compact support.

Beta distribution In order to understand how the (log-)ratio behaves when k_{in} and k_{out} show a higher coefficient of variation, we model their probability density as a Beta distribution

$$\text{Beta}(x; \alpha, \beta) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{\text{B}(\alpha, \beta)}, \quad x \in [0, 1] \quad (3.29)$$

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

where $B(\alpha, \beta)$ is the beta function

$$B(\alpha, \beta) = \int_0^1 x^{\alpha-1} (1-x)^{\beta-1} dx \quad (3.30)$$

This distribution, due to the compactness of its support, can take arbitrary variation coefficient values, while giving raise to positive and symmetrically distributed values. Using the same procedure illustrated in the previous paragraph, we explore log-normal compatibility of the ratio distribution k_{in}/k_{out} extending it to a broader range of CV values, confirming that compatibility is reliably retrievable until roughly $CV \sim 0.35$, and decays quickly thereafter due to an increase in the log-ratio kurtosis (Figure 3.9, in particular panel c). As before, a high degree of log-normal compatibility is found only when CV_{in} and CV_{out} take similar values, with the log-ratio exhibiting left or right skewness otherwise (Figure 3.9.a,d).

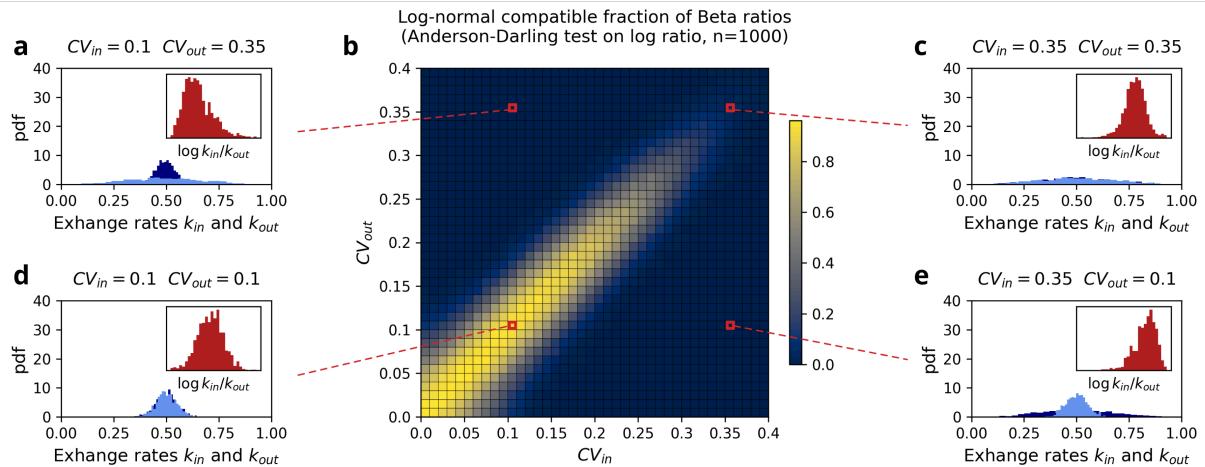


Figure 3.9: **Log-normal compatibility of ratio distribution generated by Beta distributed k_{in} and k_{out} .** The same procedure as in Figure 3.8 is carried out. Notice that due to the compact support of the Beta distribution, we are able to explore a wider range of coefficients of variation, finding a considerable degree of log-normal compatibility even in the case of $CV > 30\%$.

Stochastic simulation To confirm the theoretical predictions derived so far, we simulate the system (3.24) as a discrete-time Markov chain. In this framework, the dendrite is subdivided into a finite number of small volumes (of the same size), with or without a connection to a spine. At every time step, each compartment (dendritic volume, synapse) will be able to exchange the resource m with its neighbours, following the transition equations

$$\begin{cases} \Delta m_s^{(i)}(t+1) = \epsilon_{in}^{(i)} m_d(x_i, t) - \epsilon_{out}^{(i)} m_s^{(i)}(t) \\ \Delta m_d(x_j, t+1) = \epsilon_D [m_d(x_{j+1}, t) + m_d(x_{j-1}, t) - 2m_d(x_j, t)] - \delta_{ij} \Delta m_s^{(i)}(t+1) \end{cases} \quad (3.31)$$

3.6. Log-normality emergence from molecular dynamics

where the different ϵ represent the different transition probabilities deriving from the respective rate constants. The associated Markov transition matrix is

$$\begin{bmatrix} \vdots \\ \Delta m_d(x_{i-1}) \\ \Delta m_d(x_i) \\ \Delta m_s^{(i)} \\ \Delta m_d(x_{j+1}) \\ \vdots \end{bmatrix} = \begin{bmatrix} \ddots & & & & & & \\ \dots & \epsilon_D & -2\epsilon_D & \epsilon_D & 0 & 0 & 0 & \dots \\ \dots & 0 & \epsilon_D & -2\epsilon_D - \epsilon_{in}^{(i)} & \epsilon_{out}^{(i)} & \epsilon_D & 0 & \dots \\ \dots & 0 & 0 & \epsilon_{in} & -\epsilon_{out}^{(i)} & 0 & 0 & \dots \\ \dots & 0 & 0 & \epsilon_D & 0 & -2\epsilon_D & \epsilon_D & \dots \\ \dots & \ddots & & & & & & \end{bmatrix} \begin{bmatrix} m_d(x_{i-1}) \\ m_d(x_i) \\ m_s^{(i)} \\ m_d(x_{j+1}) \\ \vdots \end{bmatrix} \quad (3.32)$$

In the simulation, the exchange step is implemented as a multinomial sampling for each compartment, with the number of partitions equal to the number of possible transitions. By repeating the same CV sampling procedure carried out for Figure 3.8 and Figure 3.9, we confirm that indeed $m_s^{(i)}$ shows a synaptic distribution compatible with a log-normal density under the hypothesized conditions on the exchange rate coefficients of variation (Figure 3.10 a,b).

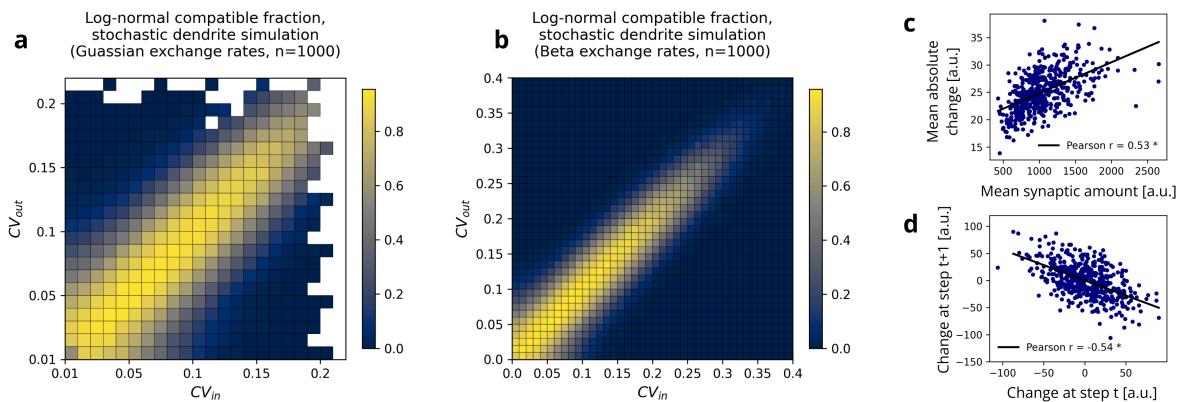


Figure 3.10: Log-normal compatibility of stochastically generated synaptic molecular distributions. The same procedure as in Figure 3.8 and Figure 3.9 is carried out. As hypothesized, stochasticity does not impact the equilibrium synaptic molecular distributions, which follow the theoretically predicted degree of log-normal compatibility.

Moreover, the simulated process is able to show two additional hallmarks which have been experimentally observed in relation to synaptic size dynamics [117,118], i.e. the proportionality between the average synaptic molecular content and its average absolute change, as well as the anticorrelation between changes in synaptic molecular amounts between subsequent observation timesteps (Figure 3.10 c,d).

Having corroborated the inferred log-normality of the basal parameters of (3.13) both through experimental observation and mechanistic theory and simulation, we accept it as a reliable

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

distribution for the model's catalytic distribution. Methodologically, this allows us to complete the implementation of our model, structuring our next stimulations as follows:

1. first, we define a dendrite with a given number of spines, and for each of the spine we draw a value for $K_b^{(i)}$ and $N_b^{(i)}$ from the general bivariate log-normal catalytic distribution (Figure 3.6 d);
2. second, we define which spines will be selected for plasticity induction, and consider the following computations, taking them as reference;
3. third, using the model (3.13) we compute for each spine position x_i and time t the value of the synaptic size $P^{(i)}(t)$.

Importantly, since each time we instantiate a dendrite, its spines contain different amounts of $K_b^{(i)}$ and $N_b^{(i)}$, for each experiment we conduct multiple model runs, and evaluate the final result in terms of summary statistics. We assume that this is equivalent to repeating the same plasticity experiment in different neurons.

As a final consequence, we notice that given the log-normality of the marginals $K_b^{(i)}$ and $N_b^{(i)}$ immediately implies the log-normality of the ratio $K_b^{(i)}/N_b^{(i)}$ and, through (3.13), the widely observed log-normality of synaptic sizes (Figure 3.11). In this regard, one starts by observing that the quasi-steady state approximation introduced in Section 3.3.1 leads to the intermediate expression for spine sizes

$$P^{(i)}(t) = \frac{k_K}{k_N} \frac{K^{(i)}}{N^{(i)}} U_d \quad (3.33)$$

as $U_s^{(i)}$ has balanced out with the dendritic quota U_d , which itself is constant along the dendrite. For a fixed time t , the probability density of $P^{(i)}$ will therefore be equal to the ratio of the probability densities of $K^{(i)}$ and $N^{(i)}$. Having found that these two densities are jointly log-normal, their ratio also follows a log-normal distribution.

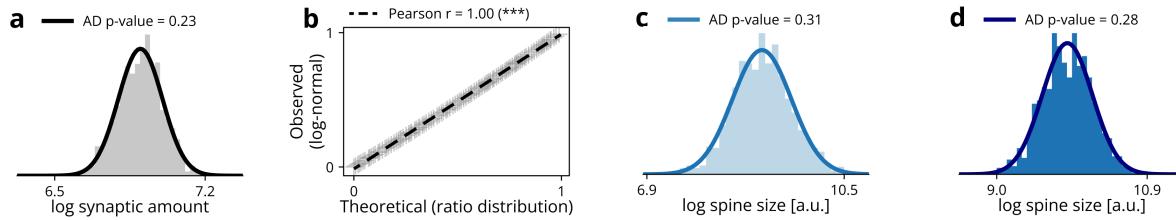


Figure 3.11: Statistical features of the synaptic plasticity-related proteins. When presented, log-normality is assessed by testing the logarithms of the data for normality with the Anderson-Darling (AD) test. A resulting p-value higher than 0.05 denotes a positive test. **a** Synaptic steady-state distribution of a freely diffusive abstract molecule, showing strong compatibility with a log-normal distribution ($N = 1000$). **b** Quantile-quantile plot (QQ plot) comparing the known theoretical distribution of the diffusive molecule (ratio distribution) with a fitted log-normal. From both the visual inspection and the linear correlation coefficient, the two distributions are experimentally indistinguishable. **c,d** Experimentally observed ($N = 1105$) and model simulated ($N = 1000$) distributions of synaptic sizes. Both are compatible with a log-normal distribution.

In summary, we set out to find a suitable probability density describing the fitted values of $K_b^{(i)}$ and $N_b^{(i)}$. We start by exploring the synaptic distribution of a freely diffusing molecule under general dynamical conditions, and find that it is statistically undistinguishable from a log-normal distribution. We verify that a number of synaptic proteins, and specifically plasticity-related catalysts, show a log-normal compatible distribution across spines, and a high degree of correlation with the structural protein Homer1. We confirm that the optimal parameter values inferred for our model's $K_b^{(i)}$ and $N_b^{(i)}$ follow a bivariate log-normal distribution, with correlations compatible with the ones observed experimentally. Finally, we show that our model, endowed with this bivariate distribution, is able to reproduce the experimentally observed log-normal distribution of synaptic sizes, a feature that directly emerges from the properties of log-normal distributions and their ratios.

3.7. The effect of inter-stimulus distance on the plasticity profile

After understanding the statistical properties of the molecular factors underpinning our model, we move on to understanding its predictive power. In particular, we are interested in exploring whether and how the protein variability that we observed for the model's basal conditions can impact the synaptic response to an induction protocol, possibly challenging result interpretation in an experimental setting.

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

Also, more importantly, we are interested in testing the reliability of the optimized model in reproducing synaptic plasticity experiments that were not used for the fitting.

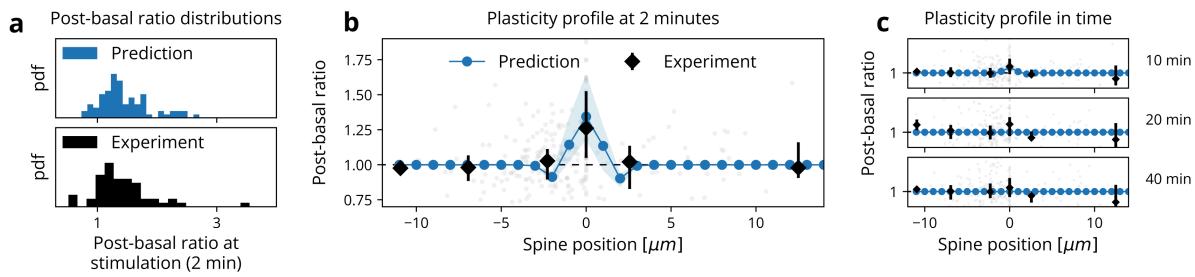


Figure 3.12: Model validation with plasticity prediction. **a** The model is able to predict the skewed distribution of the post-pre size ratios at the stimulation locations (Kolmogorov-Smirnov $p = 0.13$). **b,c** Model prediction tested on the left-out dataset (seven distributed stimulations). The stimulated spatio-temporal profile of synaptic sizes shows good accordance with the observed data at all timepoints.

To this end, we start by validating our model to predict outcomes of one of the experiments reported in [68]. We focus on one of the datasets that has not been used for model fitting, where glutamate uncaging is performed on 7 spines distributed along a dendritic stretch, at an average inter-stimulation distance of $16 \mu\text{m}$. The simulation predicts very closely the dynamics observed in the experiment (Figure 3.12 b,c), with synaptic change being particularly evident at 2 minutes after the induction. Three qualitatively different behaviours can be observed at different distances from stimulation (Figure 3.12 b): the stimulated, as well as the spines very close to the stimulation ($\sim 2 \mu\text{m}$) undergo potentiation, while spines located at an intermediate distance ($\sim 2 - 4 \mu\text{m}$) undergo depression. Finally, spines that are further than $4 \mu\text{m}$ do not show a significant response to the stimulus. These effects decay in time, with synaptic sizes returning to baseline at the final mark of 40 min (Figure 3.12 c). In addition to the average spatio-temporal plasticity dynamics, the model is also able to predict plasticity statistics for the stimulated spine, correctly reproducing the distribution of post-pre size ratios at 2 minutes after the induction (Figure 3.12 a, Kolmogorov-Smirnov test $p = 0.13$). Crucially, we are able to see the presence of a small fraction of stimulated spines undergoing depression (Figure 3.12 a, post-pre ratios smaller than 1). This observation strongly agrees with the hypothesis that synaptic-specific protein variability plays a crucial role in determining the outcome of a plasticity experiment, allowing for diverse behaviours to emerge in response to a single stimulation type.

To further disentangle the role of the induction protocol from the role of synaptic variability, we simulate a specific set of plasticity induction experiments (Figure 3.13), with stimulations

3.7. The effect of inter-stimulus distance on the plasticity profile

applied at 2 locations with increasing distances. This allows us to characterize the effect of the stimulus geometry on synaptic catalyst activation, ultimately, synaptic change.

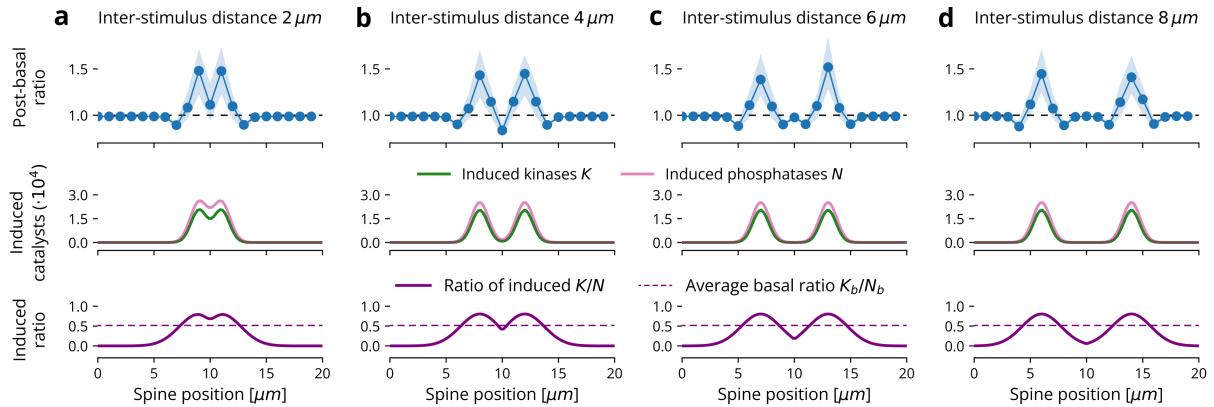


Figure 3.13: Plasticity response profiles when stimulating two spines at different distances (2, 4, 6 and 8 μm). Each column compares the post-pre spine size ratio with the newly induced catalysts and the ratio of newly induced K and N . (Post-pre size ratio, both experimental and simulated, is shown as median and inter-quartile range, given the non-normality of the underlying distributions).

As expected, we find qualitatively distinct behaviours between stimulated and non-stimulated spines. The first show distinct potentiation, the amount of which does not depend on the distance between the stimuli. This agrees with the fact that the induction protocol the model has been fitted with (glutamate uncaging, 60 pulses at 0.5 Hz), is considered a potentiation protocol, and elicits *on average* synaptic potentiation. Non-stimulated spines, on the other hand, exhibit a much more nuanced behaviour, and strongly depend on the stimulation geometry. In particular, the spines located between the two stimulations can undergo potentiation, depression or no plasticity change, depending on the distance between the stimuli. This is the result of the differential catalyst activation that happens at different distances and strongly depends on the inter-stimulus proximity (Figure 3.13 middle rows). Different catalyst activations and, in particular, different resulting ratios $K^{(i)}/N^{(i)}$ then lead to different plasticity outcomes. When the stimulations are close to each other (Figure 3.13 d), the newly induced ratio is high, leading to spine potentiation. Increasing the inter-stimulus distance reduces this ratio, leading first to stronger depression (Figure 3.13 e), and then to the absence of plasticity. This non-linear relationship response and distance between stimuli becomes clear when, together with the ratio, we consider the absolute amounts of newly activated $K^{(i)}$ and $N^{(i)}$. When a spine is far from both stimuli, these absolute amounts decay to zero, leaving the spine unaltered.

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

In summary, our model is able to reproduce plasticity dynamics not only by correctly predicting the average synaptic change, but by recovering the statistical distribution of changes observed across different experiment trials. This stochasticity is directly linked to the variability of synaptic protein distributions and accounts for apparently contradictory observations like spines undergoing depression after being exposed to a potentiation protocol. We also show how stimuli cooperating together can lead to an increase of both potentiation and depression of hetero-synaptic spines, depending on the interplay between their distance and the protein activation profile they induce.

3.8. Parameter degeneracy and bidirectional effects of FK506

In the previous section, we have shown that different stimulus configurations can lead to different plasticity outcomes depending on the balance between the baseline synaptic catalysts and the stimulus activated quota. This prompts us to ask how a change of the baseline protein distributions alone can lead to observable changes in synaptic response, and if this link can be quantitatively characterized in our framework.

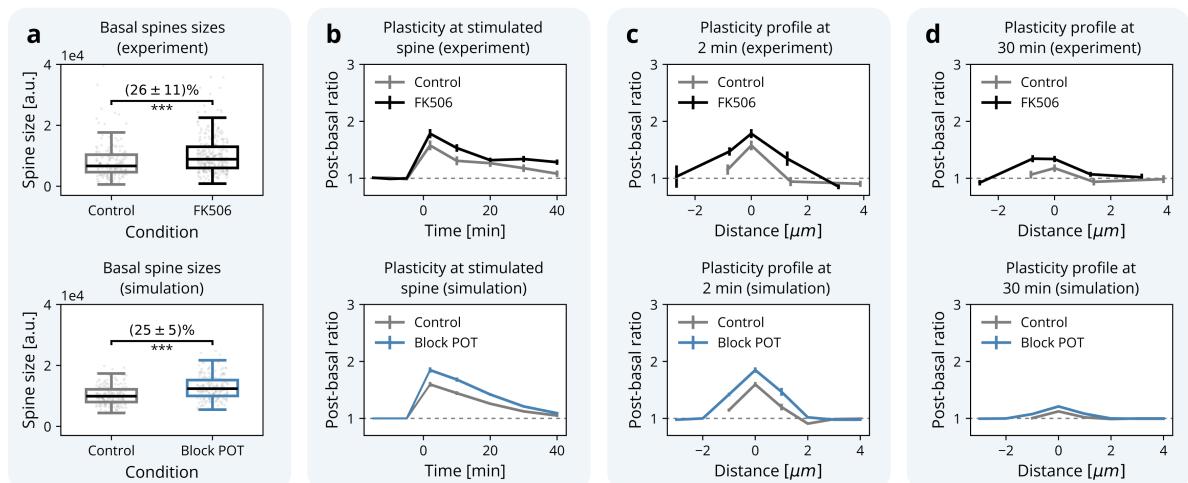


Figure 3.14: **Dose-dependent effect of FK506 (2 μ M).** Data is presented as mean \pm SEM for compatibility with the considered experimental work. In each panel, the upper row reports the experimental observation, while the lower shows the model simulation. **a** The addition of FK506 induces a significant increase in average basal synaptic size (Kolmogorov-Smirnov p -value < 0.01), as well as an overall increment of synaptic potentiation after protocol induction, both temporally (panel **b**) and spatially (panels **c** and **d**).

To address these questions, we again focus our attention on experimental observations reported in the literature, and, in particular, on the effect of Calcineurin inhibition via FK506.

3.8. Parameter degeneracy and bidirectional effects of FK506

In their work, [68] show that this inhibition leads to an overall increase of uncaging-induced potentiation, with higher and more lasting synaptic strengthening occurring both at stimulated and hetero-synaptic spines (Figure 3.14 a-d, upper plots). This result follows intuition, since a reduced amount of spine phosphatases should lead to an increase in calcium-induced protein phosphorylation and, consequently, overall synaptic potentiation. An opposing result, however, is presented in [98] (Figure 3.15 a, upper plot). Here, application of FK506 leads to a reduction in potentiation in a dose-dependent fashion. The authors explore several possible reasons that could explain this counterintuitive effect, considering, among others, qualitative differences in calcium signalling elicited by NMDA receptors and voltage-gated calcium channels.

We propose a unified interpretation, able to generate these conflicting observations from a common dynamical framework. The increase or the reduction in potentiation derives from an increase or a reduction of the ratio between basal and activated catalysts in respect to the control condition. Two factors drive this change. First, the addition of FK506 impacts the basal amounts of active CaMKII and Calcineurin, increasing the first and reducing the second [119,120]. Secondly, both FK506 and the difference in the calcium channels driving plasticity (NMDARs in [68] and VGCCs in [98]) modulate the stimulus-induced quota of calcineurin, due to inhibition and due to a quantitative difference in the emerging calcium dynamics. All these effects can be transparently implemented in the model by changing the value of four parameters describing the corresponding catalytic features (Table 4.3 in Appendix C).

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

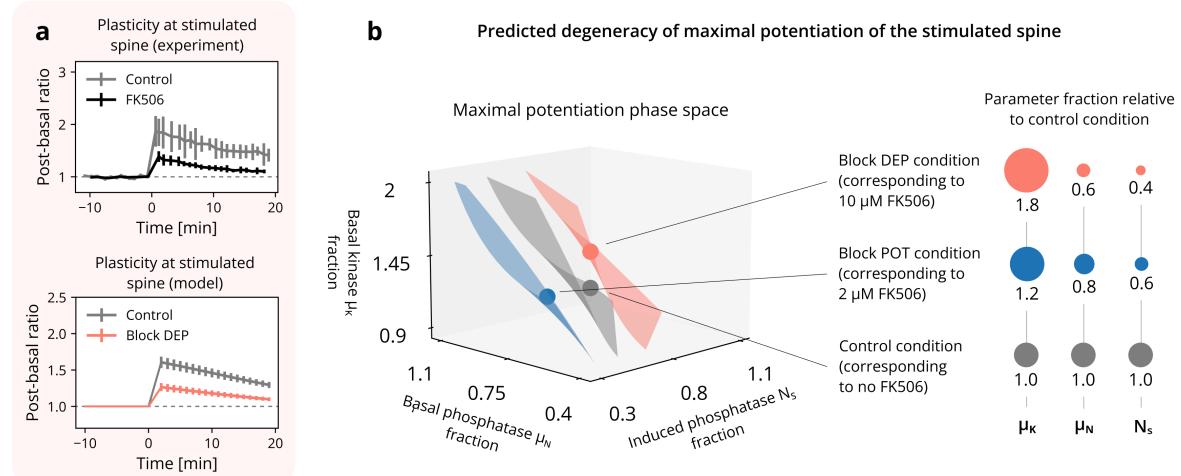


Figure 3.15: **Dose-dependent effect of FK506 (10 μM).** Data is shown as in Figure 3.14. **a** Upper panel: experimental observations reported [98]. In this case, the addition of FK506 reduces synaptic potentiation following plasticity induction. Lower panel: model reproduction after blocking the phosphatase-associated parameters (Block DEP). **b** Average maximum potentiation level surfaces (isosurfaces) in function of the average basal phosphatases μ_N , average basal kinases μ_K , and the induced phosphatase quota N_s . These surfaces identify parameter values that lead to the same average maximum potentiation. Three model parameter points, corresponding to the previously used control, Block POT, and Block DEP conditions are shown on the respective isosurface.

We start by reproducing the observations in [68], and modify the optimal model parameters to account for the effects of the addition FK506, in accordance with the previous literature [119–121]. This strategy not only solves the technical challenge of fitting *de novo* a new optimal parameter set (unachievable due to the reduced number of points in the dataset), but also corroborates the claim that our model is able to give a transparent link between its parameters and their biochemical counterpart. Encouragingly, we find that the blocked model (Block POT) correctly reproduces the observed outcomes in [68], with higher and longer lasting potentiation both at stimulated and surrounding spines (Figure 3.14, lower panels). We are also able to reproduce the slight increase in basal spine sizes observed after the FK506 has been applied to the culture (Figure 3.14 a). In agreement with [119], this arises from an increase in average basally active kinases μ_K combined with a reduction of basally active phosphatases μ_N . After accounting for the difference in applied FK506 (2 μM in [68], 10 μM in [98]), we can repeat the procedure described above and obtain the results presented in [98]. As before, we are able to reproduce the correct plasticity behaviour, this time with the addition of FK506 inducing a reduction of potentiation at the stimulated spines (Figure 3.15, lower plot). In fact, our model indicates that potentiation change shows a continuous, albeit nonlinear, dependence

3.8. Parameter degeneracy and bidirectional effects of FK506

on three coacting factors (Figure 3.15 b): the average basal phosphatase and kinase content μ_N and μ_K , directly modulated by FK506, and the stimulus induced quota N_s , determined by FK506 and the stimulus features. The emerging average plasticity behaviour, in particular of the stimulated spine, is determined by specific values of these parameters, and, importantly, we see that there is a whole set of different values leading to no change in comparison to the control condition (Figure 3.15 b, grey surface). Moreover, due to the continuity of the change in potentiation with respect to the these parameters, very low dosages of FK506 ($1 \mu M$) would not substantially alter the basal/induced ratios, leading to statistically non-significant effects on synaptic plasticity [98].

In summary, our model is able to quantitatively characterize the interplay between the basal synaptic protein distributions and the molecular action of a synaptic stimulus, correctly reproducing the effect of catalyst inhibitors on synaptic plasticity. Moreover, it is able to provide a unified interpretation for the antagonistic effects on synaptic potentiation observed for the calcineurin inhibitor FK506, giving a clear interpretation of its dose-dependent effect.

3.9. Variability of synaptic response and effect of basal synaptic size

In the previous sections we show that the outcome of a plasticity experiment is determined by a variety of factors and, in particular, by a nuanced interplay between the synaptic protein statistics and the activation effect induced by the stimulus. We characterized this interplay in relation to different stimulus features and different basal catalytic conditions, obtaining a description that provides multiple experimentally testable aspects. As a final step, we ask ourselves if our model can provide a clear hypothesis for the widely observed inverse relationship between synaptic size and its tendency to potentiate [83].

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

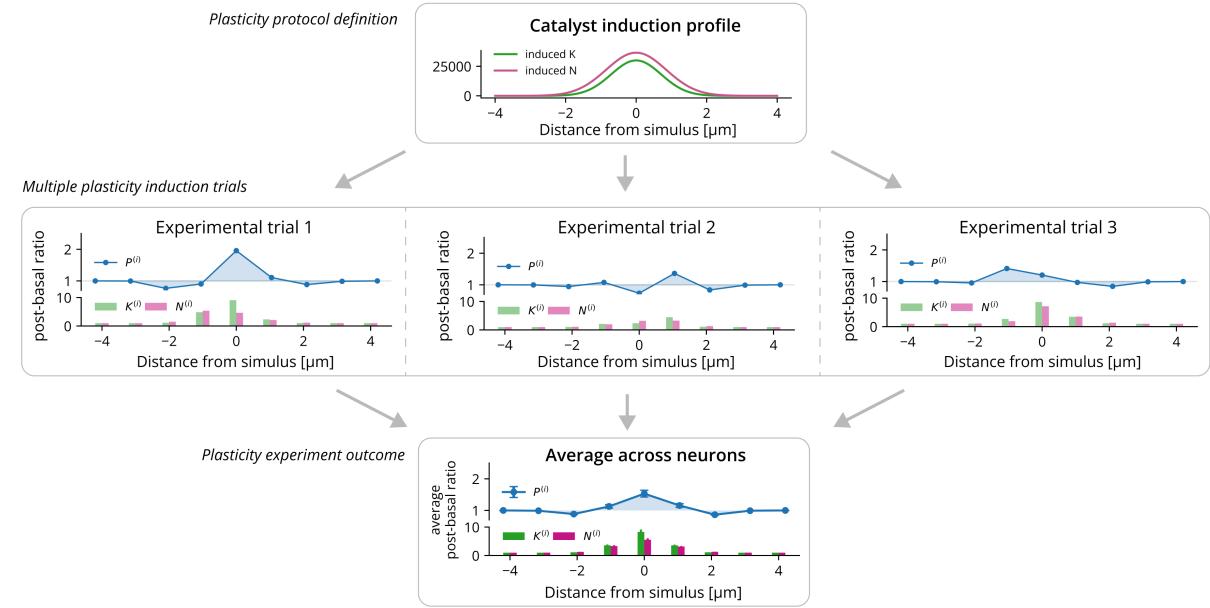


Figure 3.16: The variability of plasticity response across trials The application of a synaptic potentiation protocol provides a fixed contribution of active synaptic kinases and phosphatases (upper panel). The synaptic response, however, is determined by the sum of this deterministic quota with the stochastic basal amount, and varies significantly across experiment repetitions (middle panels). On average, however, a well-defined spatial plasticity profile emerges, uniquely determined by the stimulus type and the synaptic protein statistics (lower panel).

In previous works, the variability of synaptic plasticity response has been attributed to a number of causes, from biological stochasticity to a degree of methodological and observational noise [84,85,122]. In our model, this variability is present as well, with different instances of the same experimental protocol leading to different plasticity behaviours (Figure 3.16). For example, while the average synaptic response profile shows a Mexican-hat spatial structure, single experimental instances can considerably detach from it, even admitting depression at the site where the LTP protocol is applied. In the presented framework, this variability can be directly attributed to the stochastic nature of synaptic basal conditions. The connection can be clearly understood from a phase space representation of these basal conditions (Figure 3.17).

3.9. Variability of synaptic response and effect of basal synaptic size

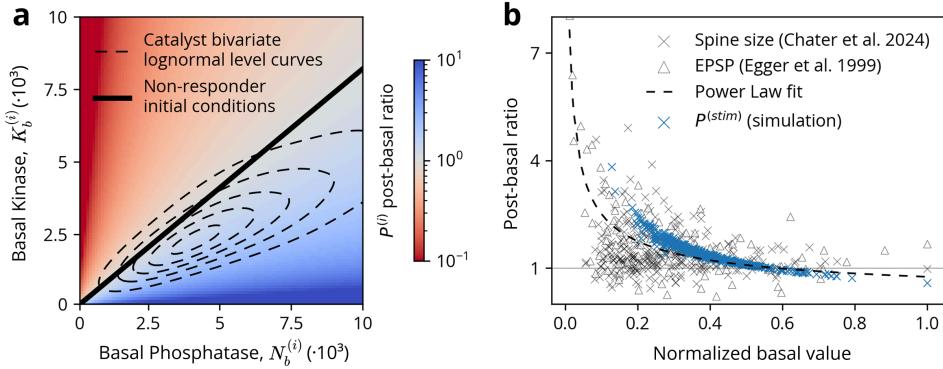


Figure 3.17: **Size dependence of plasticity response: potentiating protocol.**

a For each basal synaptic condition (a pair $K_b^{(i)}$ and $N_b^{(i)}$), the model allows for finding the average synaptic response ($P^{(i)}$ post-basal ratio) to a given stimulus protocol. This response depends on the synaptic catalytic distribution, with a clear distinction between potentiating (blue) and depressing (red) regions. Between these regions, a set of initial catalytic values lead to no change after stimulus (black line). **c** Several synaptic measures (spine size, EPSP) show an overlapping trend between their relative variation (post-pre ratio) and their normalized basal value. Our simulations also lead to a very similar behaviour, closely following a power-law fit to data proposed in [72].

In this representation, a spine is uniquely defined by a pair of values $(K_b^{(i)}, N_b^{(i)})$ corresponding to its catalytic content before a protocol is applied. As discussed in Section 3.5, biological synapses do not uniformly fill this space but distribute themselves following a bivariate log-normal distribution (Figure 3.6 d and dashed lines in Figure 3.17 a). For a defined stimulus (a specific value for the model's catalyst activation parameters) our model can predict the average (maximal) response for each point of the phase plane (red-blue color code in Figure 3.17 a), showing that, at least in theory, the same stimulus can lead to both potentiation and depression. Moreover, it shows that the transition between potentiating and depressing behaviours is smooth and, in fact, that they can be separated by a smooth line corresponding to *non-responder* initial conditions, i.e. conditions identifying spines that will exhibit no change in response to the stimulus. To derive the explicit equation for this line, which represents the 0 level-set of average synaptic change, we proceed as follows.

We start by considering the solution (3.13) and rewrite it for convenience as

$$P^{(i)} = \Pi \frac{\frac{K^{(i)}}{N^{(i)}}}{R^{(i)} + \frac{K^{(i)}}{N^{(i)}}} \quad (3.34)$$

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

where $R^{(i)} = \Omega + \sum_{i \neq j} \frac{K^{(j)}}{N^{(j)}}$ represents the heterosynaptic competitive portion of the system. As described in the previous sections, each spine starts from some basal catalytic values $K_b^{(i)}$ and $N_b^{(i)}$, and a stimulus acts by modifying these values by some amount Δ_K and Δ_N depending on various factors (e.g., distance from the considered spine, time elapsed from induction). Importantly, this action takes place in every spine in the system, so that $R^{(i)}$ is also modified by some amount Δ_R . We focus our attention on stimulated spines, i.e. spines where $\Delta_K = K_s$ and $\Delta_N = N_s$. Indicating with $P_b^{(i)}$ the basal size and with $P_{post}^{(i)}$ the post-induction size of a spine, we are interested in finding the relation between $K_b^{(i)}$ and $N_b^{(i)}$ satisfying $P_b^{(i)} = P_{post}^{(i)}$.

By applying (3.34) to both sides of the equation, one obtains

$$\frac{K_b^{(i)}}{R^{(i)} N_b^{(i)} + K_b^{(i)}} = \frac{K_b^{(i)} + K_s}{(R^{(i)} + \Delta_R)(N^{(i)} + N_s) + K_b^{(i)} + K_s} \quad (3.35)$$

which can then be solved to obtain the final solution

$$K_b^{(i)} = N_b^{(i)} \frac{K_s}{N_s} \left[1 + \frac{\Delta_R}{R^{(i)}} \left(1 + \frac{N_b^{(i)}}{N_s} \right) \right]^{-1} \quad (3.36)$$

describing the curve in the K and N phase plane corresponding to the spines that will not change size in response to a stimulus inducing K_s and N_s amounts. This expression is stochastic, as for a fixed value of $K_b^{(i)}$ and $N_b^{(i)}$, the heterosynaptic factor $R^{(i)}$, as well as its variation, depend on the specific dendritic instance. We resorted to exploring this stochasticity by running multiple simulations using our optimal parameter set, and found that the value of $\Delta_R/R^{(i)}$ was extremely small in all cases. This effectively allows us to approximate (3.36) to the linear equation $K_b^{(i)} = \frac{K_s}{N_s} N_b^{(i)}$, providing a much clearer intuition on the initial conditions that characterize a spine showing no change in size in response to a stimulus.

From a practical standpoint, conducting multiple trials of a synaptic plasticity experiment consists precisely of fixing the stimulus parameters (the color-code) and sampling the catalytic phase space following the bivariate synaptic distribution; the average response (the color corresponding to the mean of the bivariate log-normal) is then what defines the protocol as potentiating or depressing.

Despite being experimentally inaccessible, the catalytic phase space has a well-defined relation to the observable $P^{(i)}$. The considerations illustrated in the previous paragraph can be promptly translated to synaptic sizes via (3.13), predicting the existence of an inverse relationship between basal values and the synaptic tendency to potentiate. Moreover, the same protocol could induce potentiation in smaller spines and depression in bigger ones, with a

3.9. Variability of synaptic response and effect of basal synaptic size

family of “intermediate” spines showing no change after induction. In order to validate this prediction, we again turn our attention to experimental observations, describing synaptic size [68] and synaptic weight evolution [84] after plasticity induction (Figure 3.17 b). After normalizing the basal values of each dataset (dividing them by their maximum), we see that both observations and the simulated variations of $P^{(i)}$ show good agreement, and, in particular, we see that the simulation closely follows the power-law fit proposed in [72].

So far, we have assumed that the difference between plasticity protocols resides only in the value of the parameters driving catalyst dynamics. This would imply that the same model can also be used to describe the synaptic response to an LTD protocol, the only difference from before being a rotation in the phase space color-code (cfr. panels a in Figure 3.17 and Figure 3.18).

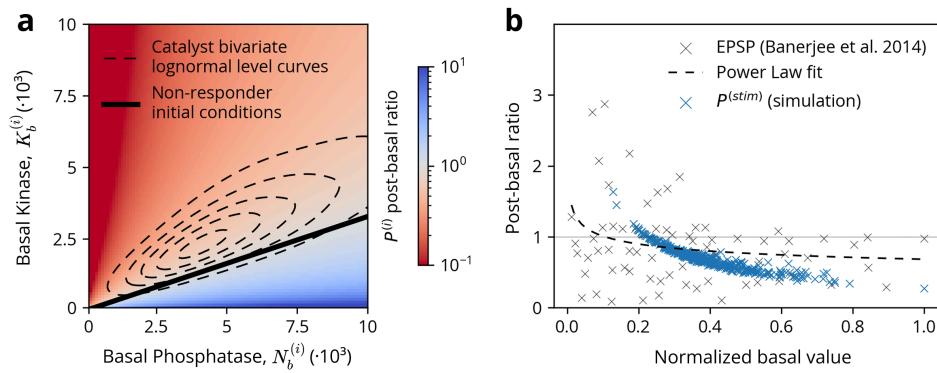


Figure 3.18: **Size dependence of plasticity response: depressing protocol.**

The panels are presented as in Figure 3.17, after scaling at roughly half the value ($K_s^{dep} = 0.4 K_s$). **a** same as panel Figure 3.17 a. Notice that the majority of the spines falls in the red region, allowing the protocol to be considered depressing in an experimental setting. **b** Same as Figure 3.17 b in a depressing setting. The simulation does not obey the power-law fit to data, but qualitatively recovers the inverse relationship between spine size and potentiation/depression, admitting the potentiation of small spines.

After adapting the stimulation parameters to reflect depression induction, we test our predictions against experimental data coming from an LTD experiment [85] (Figure 3.18 b). Despite showing less adherence to the experimental data, the model is still able to capture the dependence of the observed synaptic response on their initial value, with pure depression happening only for bigger spines, and potentiation being possible for the smaller ones. We do not see, however, a remarkably good compatibility between predictions and data, having to conclude

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

that there could be additional qualitative differences between potentiating and depressing protocols that are not accounted for in our framework.

In summary, our model shows that the behaviour of a spine undergoing plasticity strictly depends on the interplay between its basal protein distributions and the inducing features of the stimulus. This interplay translates to a smooth and monotonically decreasing dependence of synaptic size variation and synaptic basal size, recovering counterintuitive experimental findings like depressing response to potentiating stimuli, or non-responding spines. The model also predicts that a hallmark of this dependence can be observed in the synaptic response to any stimulus protocol, but further investigation is necessary to understand its descriptive power for induced synaptic depression.

3.10. Summary and remarks

In this work, we have presented a minimal model describing multi-spine plasticity occurring at the minute-to-hour timescale along a linear dendritic branch. We have derived its governing equations from two fundamental biophysical processes, compatible with these spatial and temporal scales [57,58,61,62,72]: molecular diffusion and phosphorylation. To the best of our knowledge, this is the first study that characterized synaptic plasticity at the single observation level, and not only as an average response across trials.

Previous studies have employed a similar approach [68,74,123], with a notable example represented by [107], where the authors demonstrate how the general principle of resource-sharing can account for non-linear synaptic properties like multiplicative scaling and runaway dynamics prevention. Our model is able to extend this framework with explicit biochemical meaning, as well as a clear spatial structure for the underlying dendritic system. This specification is essential to understanding how an induction protocol influences the plasticity dynamics of a multi-spine system. It enables, for example, a clear distinction between a “passive” depression induced by competition for resources [68,107] and the “active” depression, mediated by the induction of the phosphatase family N in both stimulated and neighbouring spines. This latter form, in particular, is directly connected to the Mexican-hat shaped plasticity profile observed for single stimulations [69], and derives from the higher activation spread of phosphatases in comparison to kinases [119,124]. A strong validation of this hypothesis could come, among other possibilities, by investigating the heterosynaptic plasticity induced by two close stimuli as a function of their distance from each other (as in Figure 3.13 d-g).

The specific characteristics of the induction protocol are not, however, the only factors determining plasticity response. In fact, our model strongly indicates that they determine only the average behaviour of a much more diversified and variable dynamic landscape [68,72,84,85,98]. Guided by our model, we propose the hypothesis that a considerable portion of this variability is encoded in the synaptic basal catalytic distributions. A substantial corpus of modeling work has characterized the features and the origins of synaptic statistical properties, focusing in particular on what appears to be a generalized compatibility of a number of synaptic quantities with a log-normal distribution [115]. Several mechanisms have been proposed to explain this observation, ranging from the most fundamental multiplicative noise [115], to more sophisticated models [82] grounded on general stochastic processes [80,118], or local binding mechanisms [79]. Despite providing an accurate characterization of log-normality, these models often do not provide an immediate mapping between their driving parameters and the biochemical machinery underlying synaptic distributions. Moreover, an assumption on which all of these models implicitly rely is that dendritic spines represent a *statistical ensemble*, i.e., they can be considered a different instantiation of the same, stationary random process. By taking a different approach, and constraining our model to depend only on elementary molecular dynamics, we have showed that log-normal compatibility could emerge spontaneously as the result of the elementary diffusive dynamics of synaptic proteins. This hypothesis is not only able to directly connect the spine-size log-normality to the underlying catalytic log-normality, but is also able to avoid the ensemble assumption. Moreover, it is able to provide a robust, minimal mechanism for the observations in [71,81], where log-normal compatible synaptic distributions are shown to emerge independently of neuronal activity and, therefore, possibly not driven by an information encoding optimality principle.

Our models also allows for a translation of the statistical properties of synaptic catalyst to the probabilistic description of the synaptic response to plasticity induction. Several experimental and theoretical works have observed an inverse relationship between synaptic size (or weight) and its tendency to potentiate when stimulated, under a variety of plasticity protocols [83,85,98,118]. In some instances, this inverse relationship is not restricted to a specific plasticity direction, with the same induction being able to elicit both depression and potentiation, depending on the initial synaptic strength [125]. Our model is able to support this observation, in strong similarity with the power-law dependence proposed in [72]. Following this, we propose that the synaptic response profile is quantitatively related to the calcium-induced catalytic dynamics occurring at every spine, and directly linked to the kinase-to-phosphatase

3. Modeling the impact of protein distributions on multi-spine synaptic plasticity

ratio before and after stimulation. For small spines, starting from a low K_b^i/N_b^i , this ratio would on average increase, while the opposite would occur for big spines. Lastly, a class of intermediate spines, the size of which would depend on the stimulation features, would not show a change in its catalytic ratio and therefore would appear not to respond to the induction protocol (as reported in, e.g. [125]). This mechanism is consistent with the observations in [100], where partial inhibition of Protein Phosphatase 1 (one component of the phosphatase family N) induces an overall shift towards potentiation, switching the response to a 10 Hz stimulus protocol from LTD to LTP. It is important to note that our model, despite being able to predict the existence of potentiation in very small spines undergoing LTD-inducing stimulation [85], does show a lower quantitative predictive power in this latter case. This discrepancy could arise from several reasons, one of the most likely being the difference in biochemical pathways leading to potentiation or depression [92], [93]. Despite characterizing two possible general depression mechanisms (passive resource subtraction and active heterosynaptic phosphatase induction), we have ultimately fitted the model on data obtained under a potentiation protocol. The emerging predictions could, therefore, not be optimal for describing dynamics mediated by other types of receptors like, for example, metabotropic glutamate receptors [92].

One final question our model allows us to explore is the dependence of synaptic plasticity on catalyst perturbation and, in particular, on the differential block of calcineurin. The effect of FK506 on synaptic potentiation has been studied under a plethora of experimental conditions (see [119,126,127] and references therein) with observations supporting both a facilitating and a hindering action. Multiple reasons have been proposed for the contradictory nature of these results, focusing strongly on the qualitative differences in the evoked calcium second-messenger cascade [128,129] under different experimental conditions.

Our model suggests the possibility that these qualitatively different results could emerge from a quantitative feature, i.e, a degeneracy in the phospho-dephosphorylation dynamics with respect to their driving parameters. This degeneracy corresponds to the invariance of the observed synaptic potentiation when a change of initial and stimulus-induced catalyst amounts takes place. It has been shown that FK506 impacts both basal activity of CaMKII and calcineurin [119,120], as well as the efficacy of the newly activated calcineurin quota in response to a plasticity protocol. These three components correspond directly to three model parameters (μ_N , μ_K , and N_s) and, consequently, to three degrees of freedom. We show that this 3-dimensional space is foliated by the smooth maximal potentiation isosurfaces an average spine can show after stimulation. Different experimental conditions would correspond to dif-

3.10. Summary and remarks

ferent points in this space, mapping to different surfaces and, consequently, different degrees of synaptic potentiation. Moreover, the dose-dependent effect of FK506 would correspond to a smooth line traversing these isosurfaces. In order to characterize this line, however, a precise titration of the impact of FK506 on basal and newly activated CaMKII and CaN is necessary.

4. Conclusion

Memory encoding and learning processes are fundamentally dependent on synaptic plasticity mechanisms. Elucidating how synaptic weights are regulated could provide crucial insights into these cognitive functions while also offering groundbreaking applications. Clinically, epilepsy represents a prime therapeutic target, as many epileptic manifestations arising from compromised network properties stem directly from disrupted neuronal protein turnover, misfolding, and, generally, dynamical alterations.

In this thesis, I have investigated synaptic plasticity from a spine-centric modeling standpoint, dissecting how individual synaptic variability leads to the final, average plasticity outcome. My first approach characterized spontaneous synaptic dynamics as a stochastic process, interpreting statistical patterns as manifestations of the stationary distribution within a synaptic statistical ensemble. The second approach addressed the phenomenon at a more mechanistic level, directly modeling molecular redistribution and phosphorylation through reaction-diffusion equations, and carefully dissecting the role of different protein families on the emerging plasticity phenomenon. Together, these complementary frameworks have yielded several significant insights into the nature of synaptic plasticity, and, to the best of my knowledge, they were the first work able to describe synaptic plasticity at the single observation level, and not only as an average response across trials.

4.1. Log-normal characteristics of Synaptic Fluctuations

Log-normality extends beyond previously observed synaptic structural quantities, such as weights and sizes, and includes also changes between discrete time intervals. This finding represents a particularly intriguing result, as it reveals a dynamical feature rarely encountered in conventional exchange stochastic dynamics. In classical multi-compartment systems, the steady-state distribution of quantities of interest typically follows binomial distributions, with temporal changes exhibiting symmetry around average values. The log-normal distribution deviates significantly from this pattern. Furthermore the covariance structure of synaptic processes cannot be attributed solely to noise, as the anticorrelation observed between successive changes extends to temporally non-adjacent modifications. This observation provides a critical validation criterion for models attempting to describe synaptic dynamics through stochastic processes.

4.2. Possible origin of log-normality

While log-normality in synaptic structural quantities has been interpreted through various lenses and linked to information-theoretic optimality principles, an alternative mechanistic explanation is possible. Rather than being selected in a normative fashion, log-normality may emerge as an approximate stationary distribution arising from the diffusive dynamics that synaptic structural proteins undergo throughout their lifetime. This hypothesis demonstrates remarkable robustness by connecting synaptic protein dynamics to the well-established central limit theorem scenario, where log-normality emerges from multiplicative processes. In the dendrite, multiplicativity arises not from random noise but from the combined effects of hundreds of multiplicative reactions that determine the final distribution of proteins across spines. Ultimately, since synaptic size can be described as the ratio of these proteins, its distribution is necessarily constrained to follow a log-normal-compatible distribution.

4.3. Size-dependent plasticity mechanisms

The anticorrelation between initial synaptic weight or size and a spine's propensity for potentiation has been consistently observed across the literature. Additionally, basal synaptic size has been shown to determine the directionality of synaptic plasticity under fixed stimulation protocols, with small spines exhibiting potentiation and larger spines showing depression. The presented modeling framework successfully reproduces this behavior, directly linking both deterministic (power-law) and stochastic profiles to catalytic ratio dynamics. Importantly, this approach also predicts the existence of non-responder synapses, interpreting them as spines with specific, stimulus-dependent catalytic content whose lack of weight change reflects an inherent invariance in the relationship between synaptic weights and protein content rather than experimental failure.

4.4. Catalytic degeneracy in potentiation mechanisms

Synaptic size is a function of catalytic content, incorporating both basal and stimulus-induced components through four free parameters. Specifying a particular maximum potentiation level introduces a single constraint within this four-dimensional parameter space, creating a three-dimensional hyperplane where diverse synaptic conditions can produce identical maximum potentiation levels. The application of FK506 in synaptic potentiation protocols has revealed nonlinear effects on the final experimental outcomes, sometimes enhancing and sometimes

4. Conclusion

diminishing maximum potentiation. This complex behavior can be explained through FK506's action within the four-dimensional catalytic space, particularly given its documented effects on both calcineurin and CaMKII basal and induced activities. By modifying these components in a dose-dependent manner, this formalism naturally accounts for all observed plasticity modifications, including the statistically insignificant effects observed at very low FK506 concentrations.

4.5. Future Directions

This work contributes to the extensive literature of synaptic modeling, representing a focused exploration of how two mathematical frameworks can illuminate synaptic dynamics and plasticity mechanisms. While not providing definitive answers to the fundamental questions in the field, these findings offer reasonable hypotheses upon which future investigations can build.

The most immediate investigation should reconsider the model presented in Chapter 3 in its full form, and reperform the optimization without relying on quasi-steady-state approximations. Although this approximation enabled successful parameter fitting and closed-form treatment of size-dependent potentiation, it introduces two structural limitations that impede the correct interpretation of synaptic plasticity transients. The steady-state assumption for resource diffusion necessarily trivializes predicted spatiotemporal dynamics, precluding fascinating phenomena such as potentiation/depression waves propagating throughout dendrites. Similarly, the steady-state approximation for catalytic dynamics reduces model expressivity and prevents accounting for differences in plasticity kinetics, which depend on absolute kinase and phosphatase amounts rather than their ratios. These limitations inevitably compromise parameter estimation reliability and reduce predictive power. Future work should therefore repeat optimization procedures, potentially within a hierarchical Bayesian framework.

Additional investigations could extend the current minute-to-hour timescale model to encompass the multi-hour characteristics of structural plasticity. While short-term and structural plasticity are traditionally considered distinct phenomena — partially supported by their different biochemical processes — all the involved mechanisms belong to a single, general reaction network. Depending on stimulation strength, short-term reversible plasticity can combine with stable structural changes. A unified, potentially perturbative model could be developed and fitted to existing datasets, better describing transient and lasting changes across these timescales. Mathematical treatment through matched asymptotic expansions

could prove valuable to allow for insightful closed-form expressions and feasible optimization, in line with what I presented here. Importantly, when incorporating longer timescales, mass conservation constraints become invalid, and closed-form approximations lose reliability. Future work must address this critical limitation by developing quasi-steady-state approximations that support transient phenomena like traveling plasticity waves while maintaining analytical tractability.

Another potential extension could operate at shorter rather than longer timescales, investigating how our deterministic reaction-diffusion model can incorporate meaningful stochastic structure (such as the one I introduced in Chapter 2). While theoretically feasible, this represents a non-trivial challenge, as direct approaches using Chapman-Kolmogorov equations become intractable due to the overwhelming number of biochemical reactions underlying synaptic plasticity. An intriguing possibility could come from recurring to renormalization group theory, potentially generating appropriate fluctuation and noise profiles while identifying general synaptic quantities relevant from normative optimality perspectives.

These research directions are the natural following steps towards understanding synaptic plasticity at the appropriate spatial and temporal scales, and could directly translate to comprehending the origins of corresponding learning and behavioral phenomena, bridging the gap between molecular mechanisms and cognitive function, and, ultimately, understanding how to recover the latter from pathology.

Bibliography

- [1] Nagel G, Brauner M, Liewald JF, Adeishvili N, Bamberg E, Gottschalk A. Light activation of channelrhodopsin-2 in excitable cells of *Caenorhabditis elegans* triggers rapid behavioral responses. *Current Biology* 2005;15:2279–84.
- [2] Jinek M, Chylinski K, Fonfara I, Hauer M, Doudna JA, Charpentier E. A programmable dual-RNA-guided DNA endonuclease in adaptive bacterial immunity. *Science* 2012;337:816–21.
- [3] Hebb DO. *The organization of behavior: A neuropsychological theory*. Psychology press; 2005.
- [4] Castellucci V, Pinsker H, Kupfermann I, Kandel ER. Neuronal mechanisms of habituation and dishabituation of the gill-withdrawal reflex in *Aplysia*. *Science* 1970;167:1745–8.
- [5] Caya-Bissonnette L, Béïque J-C. Half a century legacy of long-term potentiation. *Current Biology* 2024;34:R640–62.
- [6] Bliss TV, Lømo T. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of Physiology* 1973;232:331–56.
- [7] Douglas RM, Goddard GV. Long-term potentiation of the perforant path-granule cell synapse in the rat hippocampus. *Brain Research* 1975;86:205–15.
- [8] Stevens CF, Sullivan J. Synaptic plasticity. *Current Biology* 1998;8:R151–3.
- [9] Maletic-Savatic M, Malinow R, Svoboda K. Rapid dendritic morphogenesis in CA1 hippocampal dendrites induced by synaptic activity. *Science* 1999;283:1923–7.
- [10] Engert F, Bonhoeffer T. Dendritic spine changes associated with hippocampal long-term synaptic plasticity. *Nature* 1999;399:66–70.
- [11] Trachtenberg JT, Chen BE, Knott GW, Feng G, Sanes JR, Welker E, et al. Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. *Nature* 2002;420:788–94.
- [12] Yuste R, Bonhoeffer T. Morphological changes in dendritic spines associated with long-term synaptic plasticity. *Annual Review of Neuroscience* 2001;24:1071–89.
- [13] Magee JC, Grienberger C. Synaptic plasticity forms and functions. *Annual Review of Neuroscience* 2020;43:95–117.

- [14] Zucker RS, Regehr WG. Short-term synaptic plasticity. *Annual Review of Physiology* 2002;64:355–405.
- [15] Redondo RL, Okuno H, Spooner PA, Frenguelli BG, Bito H, Morris RG. Synaptic tagging and capture: differential role of distinct calcium/calmodulin kinases in protein synthesis-dependent long-term potentiation. *Journal of Neuroscience* 2010;30:4981–9.
- [16] Smolen P, Baxter DA, Byrne JH. How can memories last for days, years, or a lifetime? Proposed mechanisms for maintaining synaptic potentiation and memory. *Learning & Memory* 2019;26:133–50.
- [17] Ziv NE, Brenner N. Synaptic tenacity or lack thereof: spontaneous remodeling of synapses. *Trends in Neurosciences* 2018;41:89–99.
- [18] Yasumatsu N, Matsuzaki M, Miyazaki T, Noguchi J, Kasai H. Principles of long-term dynamics of dendritic spines. *Journal of Neuroscience* 2008;28:13592–608.
- [19] Minerbi A, Kahana R, Goldfeld L, Kaufman M, Marom S, Ziv NE. Long-term relationships between synaptic tenacity, synaptic remodeling, and network activity. *Plos Biology* 2009;7:e1000136.
- [20] Kasai H, Fukuda M, Watanabe S, Hayashi-Takagi A, Noguchi J. Structural dynamics of dendritic spines in memory and cognition. *Trends in Neurosciences* 2010;33:121–9.
- [21] Loewenstein Y, Kuras A, Rumpel S. Multiplicative dynamics underlie the emergence of the log-normal distribution of spine sizes in the neocortex *in vivo*. *Journal of Neuroscience* 2011;31:9481–8.
- [22] Cane M, Maco B, Knott G, Holtmaat A. The relationship between PSD-95 clustering and spine stability *in vivo*. *Journal of Neuroscience* 2014;34:2075–86.
- [23] Ziv NE, Fisher-Lavie A. Presynaptic and postsynaptic scaffolds: dynamics fast and slow. *The Neuroscientist* 2014;20:439–52.
- [24] Ishii K, Nagaoka A, Kishida Y, Okazaki H, Yagishita S, Ucar H, et al. In vivo volume dynamics of dendritic spines in the neocortex of wild-type and Fmr1 KO mice. *Eneuro* 2018;5.
- [25] Murthy VN, Sejnowski TJ, Stevens CF. Heterogeneous release properties of visualized individual hippocampal synapses. *Neuron* 1997;18:599–612.

- [26] Harms KJ, Tovar KR, Craig AM. Synapse-specific regulation of AMPA receptor subunit composition by activity. *Journal of Neuroscience* 2005;25:6379–88.
- [27] Zeidan A, Ziv NE. Neuroligin-1 loss is associated with reduced tenacity of excitatory synapses. *Plos ONE* 2012;7. <https://doi.org/10.1371/journal.pone.0042314>.
- [28] Statman A, Kaufman M, Minerbi A, Ziv NE, Brenner N. Synaptic size dynamics as an effectively stochastic process. *Plos Computational Biology* 2014;10:e1003846.
- [29] Scheler G. Logarithmic distributions prove that intrinsic learning is Hebbian. *F1000research* 2017;6.
- [30] Hazan L, Ziv NE. Activity dependent and independent determinants of synaptic size diversity. *Journal of Neuroscience* 2020;40:2828–48.
- [31] Buzsáki G, Mizuseki K. The log-dynamic brain: how skewed distributions affect network operations. *Nature Reviews Neuroscience* 2014;15:264–78.
- [32] Humble J, Hiratsuka K, Kasai H, Toyoizumi T. Intrinsic spine dynamics are critical for recurrent network learning in models with and without autism spectrum disorder. *Frontiers in Computational Neuroscience* 2019;13:38.
- [33] Van Rossum MC, Bi GQ, Turrigiano GG. Stable Hebbian learning from spike timing-dependent plasticity. *Journal of Neuroscience* 2000;20:8812–21.
- [34] Choquet D, Triller A. The dynamic synapse. *Neuron* 2013;80:691–703.
- [35] Fisher-Lavie A, Zeidan A, Stern M, Garner CC, Ziv NE. Use dependence of presynaptic tenacity. *Journal of Neuroscience* 2011;31:16770–80.
- [36] Bonilla-Quintana M, Wörgötter F, Tetzlaff C, Fauth M. Modeling the shape of synaptic spines by their actin dynamics. *Frontiers in Synaptic Neuroscience* 2020;12:9.
- [37] Bonilla-Quintana M, Wörgötter F, D'Este E, Tetzlaff C, Fauth M. Reproducing asymmetrical spine shape fluctuations in a model of actin dynamics predicts self-organized criticality. *Scientific Reports* 2021;11:1–17.
- [38] Matsuzaki M, Honkura N, Ellis-Davies GC, Kasai H. Structural basis of long-term potentiation in single dendritic spines. *Nature* 2004;429:761–6.
- [39] Grutzendler J, Kasthuri N, Gan W-B. Long-term dendritic spine stability in the adult cortex. *Nature* 2002;420:812–6.

[40] Holtmaat AJ, Trachtenberg JT, Wilbrecht L, Shepherd GM, Zhang X, Knott GW, et al. Transient and persistent dendritic spines in the neocortex in vivo. *Neuron* 2005;45:279–91.

[41] Yang G, Pan F, Gan W-B. Stably maintained dendritic spines are associated with lifelong memories. *Nature* 2009;462:920–4.

[42] Hayashi-Takagi A, Yagishita S, Nakamura M, Shirai F, Wu YI, Loshbaugh AL, et al. Labelling and optical erasure of synaptic memory traces in the motor cortex. *Nature* 2015;525:333–8.

[43] Harris KM, Stevens JK. Dendritic spines of CA 1 pyramidal cells in the rat hippocampus: serial electron microscopy with reference to their biophysical characteristics. *Journal of Neuroscience* 1989;9:2982–97.

[44] Bourne JN, Harris KM. Balancing structure and function at hippocampal dendritic spines. *Annu Rev Neurosci* 2008;31:47–67.

[45] Samavat M, Bartol TM, Bromer C, Bowden JB, Hubbard DD, Hanka DC, et al. Regional and LTP-Dependent Variation of Synaptic Information Storage Capacity in Rat Hippocampus. *Biorxiv* 2022:2022–8.

[46] Strong SP, Koberle R, Van Steveninck RRDR, Bialek W. Entropy and information in neural spike trains. *Physical Review Letters* 1998;80:197.

[47] Granot-Atedgi E, Tkačik G, Segev R, Schneidman E. Stimulus-dependent maximum entropy models of neural population codes. *Plos Computational Biology* 2013;9:e1002922.

[48] Shomar A, Geyrhofer L, Ziv NE, Brenner N. Cooperative stochastic binding and unbinding explain synaptic size dynamics and statistics. *Plos Computational Biology* 2017;13:e1005668.

[49] Murakoshi H, Wang H, Yasuda R. Local, persistent activation of Rho GTPases during plasticity of single dendritic spines. *Nature* 2011;472:100–4.

[50] Govindarajan A, Israely I, Huang S-Y, Tonegawa S. The dendritic branch is the preferred integrative unit for protein synthesis-dependent LTP. *Neuron* 2011;69:132–46.

[51] Oh WC, Parajuli LK, Zito K. Heterosynaptic structural plasticity on local dendritic segments of hippocampal CA1 neurons. *Cell Reports* 2015;10:162–9.

[52] Kopec CD, Li B, Wei W, Boehm J, Malinow R. Glutamate receptor exocytosis and spine enlargement during chemically induced long-term potentiation. *Journal of Neuroscience* 2006;26:2000–9.

[53] O'Donnell C, Nolan MF, Rossum MC van. Dendritic spine dynamics regulate the long-term stability of synaptic plasticity. *Journal of Neuroscience* 2011;31:16142–56.

[54] Shouval HZ. Clusters of interacting receptors can stabilize synaptic efficacies. *Proceedings of the National Academy of Sciences* 2005;102:14440–5.

[55] Bell MK, Holst MV, Lee CT, Rangamani P. Dendritic spine morphology regulates calcium-dependent synaptic weight change. *Journal of General Physiology* 2022;154:e202112980.

[56] Jozsa M, Donchev TI, Sepulchre R, O'Leary T. Autoregulation of switching behavior by cellular compartment size. *Proceedings of the National Academy of Sciences* 2022;119:e2116054119.

[57] Yasuda R, Hayashi Y, Hell JW. CaMKII: a central molecular organizer of synaptic plasticity, learning and memory. *Nature Reviews Neuroscience* 2022;23:666–82.

[58] Jain A, Nakahata Y, Pancani T, Watabe T, Rusina P, South K, et al. Dendritic, delayed, stochastic CaMKII activation in behavioural time scale plasticity. *Nature* 2024;635:151–9.

[59] Frey U, Huang Y -Y, Kandel ER. Effects of cAMP simulate a late stage of LTP in hippocampal CA1 neurons. *Science* 1993;1661–4.

[60] Citri A, Malenka RC. Synaptic Plasticity: Multiple Forms, Functions, and Mechanisms. *Neuropsychopharmacology* 2008;33:18–41.

[61] Earnshaw BA, Bressloff PC. A diffusion-activation model of CaMKII translocation waves in dendrites. *Journal of Computational Neuroscience* 2010;28:77–89.

[62] Penzes P, Rafalovich I. Regulation of the actin cytoskeleton in dendritic spines. *Advances in Experimental Medicine and Biology* 2012;970:81–95.

[63] Fonkeu Y, Kraynyukova N, Hafner A-S, Kochen L, Sartori F, Schuman EM, et al. How mRNA Localization and Protein Synthesis Sites Influence Dendritic Protein Distribution and Dynamics. *Neuron* 2019;103:1109–1122.e7.

[64] Wagle S, Kracht MK, Bührke A, Acker-Palmer A, Kraynyukova N, Hafner A-S, et al. An integrative model of AMPA receptor trafficking reveals the central contribution of local translation in subtype-specific kinetics. *Biorxiv* 2025.

[65] Zucker RS. Calcium- and activity-dependent synaptic plasticity. *Current Opinion in Neurobiology* 1999;9:305–13.

[66] Kubota Y, Waxham MN. Lobe Specific Ca²⁺-Calmodulin Nano-Domain in Neuronal Spines: A Single Molecule Level Analysis. *Plos Computational Biology* 2010;6.

[67] Helm MS, Dankovich TM, Mandad S, Rammner B, Jähne S, Salimi V, et al. A large-scale nanoscopy and biochemistry analysis of postsynaptic dendritic spines. *Nature Neuroscience* 2021;24:1151–62.

[68] Chater TE, Eggl MF, Goda Y, Tchumatchenko T. Competitive processes shape multi-synapse plasticity along dendritic segments. *Nature Communications* 2024;15.

[69] Tong R, Chater TE, Emptage NJ, Goda Y. Heterosynaptic cross-talk of pre- and postsynaptic strengths along segments of dendrites. *Cell Reports* 2021;108693.

[70] Oh WC, Parajuli LK, Zito K. Heterosynaptic structural plasticity on local dendritic segments of hippocampal CA1 neurons. *Cell Reports* 2015;162–9.

[71] Rössler N, Jungeńitz T, Sigler A, Bird AD, Mittag M, Rhee J-S, et al. Skewed distribution of spines is independent of presynaptic transmitter release and synaptic plasticity, and emerges early during adult neurogenesis. *Open Biology* 2023;13.

[72] Chindemi G, Abdellah M, Amsalem O, Benavides-Piccione R, Delattre V, Doron M, et al. A calcium-based plasticity model for predicting long-term potentiation and depression in the neocortex. *Nature Communications* 2022;13.

[73] Kirchner JH, Gjorgjieva J. Emergence of local and global synaptic organization on cortical dendrites. *Nature Communications* 2021;12.

[74] Shafiee S, Schmitt S, Tetzlaff C. Calcium-based input timing learning. *Biorxiv* 2025.

[75] Li L, Stefan MI, Novère NGL. Calcium Input Frequency, Duration and Amplitude Differentially Modulate the Relative Activation of Calcineurin and CaMKII. *Plos ONE* 2012;7.

[76] Graupner M, Brunel NJ-B. Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. *Proceedings of the National Academy of Sciences* 2012;109:3991–6.

[77] Chen H, Xie L, Wang Y, Zhang H. Postsynaptic Potential Energy as Determinant of Synaptic Plasticity. *Frontiers in Computational Neuroscience* 2020;16.

[78] Pache A, Rossum MCW van. Energetically efficient learning in neuronal networks. *Current Opinion in Neurobiology* 2023;83.

[79] Shomar A, Geyrhofer L, Ziv NE, Brenner N. Cooperative stochastic binding and unbinding explain synaptic size dynamics and statistics. *Plos Computational Biology* 2017;13.

[80] Statman A, Kaufman M, Minerbi A, Ziv NE, Brenner N. Synaptic Size Dynamics as an Effectively Stochastic Process. *Plos Computational Biology* 2014;10.

[81] Rössler N, Smilović D, Vuksić M, Jedlicka P, Deller T. Maintenance of Lognormal-Like Skewed Dendritic Spine Size Distributions in Dentate Granule Cells of TNF, TNF-R1, TNF-R2, and TNF-R1/2-Deficient Mice. *The Journal of Comparative Neurology* 2024:e25645.

[82] Matsuzaki M, Honkura N, Ellis-Davies GCR, Kasai H. Structural basis of long-term potentiation in single dendritic spines. *Nature* 2004;429:761–6.

[83] Bi G, Poo M-m. Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. *The Journal of Neuroscience* 1998;18:10464–72.

[84] Egger V, Feldmeyer D, Sakmann BJ. Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. *Nature Neuroscience* 1999;2:1098–105.

[85] Banerjee A, González-Rueda A, Sampaio-Baptista C, Paulsen O, Rodríguez-Moreno A. Distinct mechanisms of spike timing-dependent LTD at vertical and horizontal inputs onto L2/3 pyramidal neurons in mouse barrel cortex. *Physiological Reports* 2014;2.

[86] Chater TE, Goda Y. My neighbour hetero—deconstructing the mechanisms underlying heterosynaptic plasticity. *Current Opinion in Neurobiology* 2021;67:106–14.

[87] Petkovic J, Fioresi R. Spontaneous emergence of robustness to light variation in cnns with a precortically inspired module. *Neural Computation* 2024;36:1832–53.

[88] Royer S, Paré D. Conservation of total synaptic weight through balanced synaptic depression and potentiation. *Nature* 2003;422:518–22.

[89] Chater TE, Goda Y. My Neighbour Hetero—deconstructing the mechanisms underlying heterosynaptic plasticity. *Current Opinion in Neurobiology* 2020;67:106–14.

[90] Tazerart S, Mitchell DE, Miranda-Rottmann S, Araya R. A spike-timing-dependent plasticity rule for dendritic spines. *Nature Communications* 2020;11.

[91] d'Aquin S, Szőnyi A, Mahn M, Krabbe S, Gründemann J, Lüthi A. Compartmentalized dendritic plasticity during associative learning. *Science* 2022;376.

[92] Stein IS, Park DK, Claiborne N, Zito K. Non-ionotropic NMDA receptor signaling gates bidirectional structural plasticity of dendritic spines. *Cell Reports* 2020;34:108664.

[93] Scanziani M, Malenka RC, Nicoll RA. Role of intercellular interactions in heterosynaptic long-term depression. *Nature* 1996;380:446–50.

[94] Kastellakis G, Poirazi P. Synaptic Clustering and Memory Formation. *Frontiers in Molecular Neuroscience* 2019;12.

[95] Kuriu T, Inoue A, Bito H, Sobue K, Okabe S. Differential Control of Postsynaptic Density Scaffolds via Actin-Dependent and -Independent Mechanisms. *The Journal of Neuroscience* 2006;26:7693–706.

[96] Lisman JE, Yasuda R, Raghavachari S. Mechanisms of CaMKII action in long-term potentiation. *Nature Reviews Neuroscience* 2012;13:169–82.

[97] Purkey AM, Dell'Acqua ML. Phosphorylation-Dependent Regulation of Ca²⁺-Permeable AMPA Receptors During Hippocampal Synaptic Plasticity. *Frontiers in Synaptic Neuroscience* 2020;12.

[98] Onuma H, Lu Y-f, Tomizawa K, Moriwaki A, Tokuda M, Hatase O, et al. A calcineurin inhibitor, FK506, blocks voltage-gated calcium channel-dependent LTP in the hippocampus. *Neuroscience Research* 1998;30:313–9.

[99] Foley K, McKee C, Nairn AC, Xia H. Regulation of Synaptic Transmission and Plasticity by Protein Phosphatase 1. *The Journal of Neuroscience* 2021;41:3040–50.

[100] Jouvenceau A, Hédou G, Potier B, Kollen M, Dutar P, Mansuy IM. Partial inhibition of PP1 alters bidirectional synaptic plasticity in the hippocampus. *European Journal of Neuroscience* 2006;24.

[101] Romano DR, Pharris MC, Patel NM, Kinzer-Ursem TL. Competitive tuning: Competition's role in setting the frequency-dependence of Ca²⁺-dependent proteins. *Plos Computational Biology* 2017;13.

[102] Pinto TM. Computational Models of Intracellular Signalling and Synaptic Plasticity Induction in the Cerebellum. In:; 2013.

[103] Bertling E, Englund J, Minkeviciene R, Koskinen M, Segerstråle M, Castrén E, et al. Actin Tyrosine-53-Phosphorylation in Neuronal Maturation and Synaptic Plasticity. *The Journal of Neuroscience* 2016;36:5299–313.

[104] Bressloff PC, Xu B. Stochastic Active-Transport Model of Cell Polarization. *SIAM J Appl Math* 2015;75:652–78.

[105] Konietzny A, Bär J, Mikhaylova M. Dendritic Actin Cytoskeleton: Structure, Functions, and Regulations. *Frontiers in Cellular Neuroscience* 2017;11.

[106] Obashi K, Matsuda A, Inoue Y, Okabe S. Precise Temporal Regulation of Molecular Diffusion within Dendritic Spines by Actin Polymers during Structural Plasticity. *Cell Reports* 2019;1503–1515.e8.

[107] Triesch J, Vo AD, Hafner A-S. Competition for synaptic building blocks shapes synaptic plasticity. *Elife* 2018;7.

[108] Liu KKL, Hagan MF, Lisman JE. Gradation (approx. 10 size states) of synaptic strength by quantal addition of structural modules. *Philosophical Transactions of the Royal Society B: Biological Sciences* 2017;372.

[109] Gürth C-M, Dankovich TM, Rizzoli SO, D'Este E. Synaptic activity and strength are reflected by changes in the post-synaptic secretory pathway. *Scientific Reports* 2020;10.

[110] Eggli MF, Wagle S, Filling JP, Chater TE, Goda Y, Tchumatchenko T. SpyDen: Automating molecular and structural analysis across spines and dendrites. *Biorxiv* 2024:2024–6.

[111] Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *Arxiv* 2019.

[112] Schälte Y, Fröhlich F, Jost PJ, Vanhoefer J, Pathirana D, Stapor P, et al. pyPESTO: a modular and scalable tool for parameter estimation for dynamic models. *Bioinformatics* 2023;39.

[113] Fröhlich F, Sorger PK. Fides: Reliable trust-region optimization for parameter estimation of ordinary differential equation models. *Plos Computational Biology* 2021;18.

[114] Vosden WD, Farr WM, Mandel I. Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations. *Monthly Notices of the Royal Astronomical Society* 2015;455:1919–37.

[115] Buzsáki G, Mizuseki K. The log-dynamic brain: how skewed distributions affect network operations. *Nature Reviews Neuroscience* 2014;15:264–78.

[116] Glass GV, Collins JR. Geometric Proof of the Restriction On the Possible Values of r_{xy} When r_{xz} and r_{yz} Are Fixed. *Educational and Psychological Measurement* 1970;30:37–9.

[117] Loewenstein Y, Kuras A, Rumpel S. Multiplicative Dynamics Underlie the Emergence of the Log-Normal Distribution of Spine Sizes in the Neocortex In Vivo. *The Journal of Neuroscience* 2011;31:9481–8.

[118] Eggli M, Chater TE, Petkovic J, Goda Y, Tchumatchenko T. Linking spontaneous and stimulated spine dynamics. *Communications Biology* 2023;6.

[119] Fujii H, Inoue M, Okuno H, Sano Y, Takemoto-Kimura S, Kitamura K, et al. Nonlinear decoding and asymmetric representation of neuronal input information by CaMKII α and calcineurin. *Cell Reports* 2013;978–87.

[120] Dumont FJ. FK506, an immunosuppressant targeting calcineurin function. *Current Medicinal Chemistry* 2000;731–48.

[121] Ye Q, Feng Y-d, Yin Y-x, Faucher F, Currie MA, Rahman MN, et al. Structural basis of calcineurin activation by calmodulin. *Cellular Signalling* 2013;2661–7.

[122] Jang J, Anisimova M, Oh WC, Zito K. Induction of input-specific spine shrinkage on dendrites of rodent hippocampal CA1 neurons using two-photon glutamate uncaging. *STAR Protocols* 2021;2.

[123] Gupta R, O'Donnell C. Dendritic spine neck plasticity controls synaptic expression of long-term potentiation. *Biorxiv* 2023.

[124] Yasuda R. Biophysics of Biochemical Signaling in Dendritic Spines: Implications in Synaptic Plasticity. *Biophysical Journal* 2017;2152–9.

[125] Milstein AD, Li Y, Bittner KC, Grienberger C, Soltesz I, Magee JC, et al. Bidirectional synaptic plasticity rapidly modifies hippocampal representations. *Elife* 2020;10.

[126] Funauchi M, Haruta H, Tsumoto T. Effects of an inhibitor for calcium/calmodulin-dependent protein phosphatase, calcineurin, on induction of long-term potentiation in rat visual cortex. *Neuroscience Research* 1994;19:269–78.

[127] Torii N, Kamishita T, Otsu Y, Tsumoto T. An inhibitor for calcineurin, FK506, blocks induction of long-term depression in rat visual cortex. *Neuroscience Letters* 1995;185:1–4.

[128] Lu Y-f, Hayashi Y, Moriwaki A, Tomizawa K, Matsui H. FK506, a Ca²⁺/calmodulin-dependent phosphatase inhibitor, inhibits the induction of long-term potentiation in the rat hippocampus. *Neuroscience Letters* 1996;205:103–6.

[129] Stelzer A, Shi HG. Impairment of gabaa receptor function by N-methyl-D-aspartate-mediated calcium influx in isolated CA1 pyramidal cells. *Neuroscience* 1994;62:813–28.

[130] Chen M, Qi J, Poo M-c, Yang Y. Stability and dynamics of dendritic spines in macaque prefrontal cortex. *National Science Review* 2022;9.

Appendices

Appendix A: Shifted log-normal distribution

In probability theory and statistics, the log-normal distribution is a continuous probability distribution of a random variable whose logarithm is normally distributed. That is, if the random variable X is log-normally distributed, then $Y = \ln(X)$ is normally distributed. The log-normal distribution is parameterized by the mean, μ , and standard deviation, σ , of the underlying normal distribution. The probability density function of the log-normal distribution is given by

$$p(x) = \frac{1}{x \sigma \sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{\log(x) - \mu}{\sigma}\right)^2\right], \quad x \in \mathbb{R}^+ \quad (4.37)$$

where x is the value of the log-normally distributed variable. As we will be modeling data that can take negative values (the spines can shrink) and the standard log-normal is only defined for positive values, $x > 0$, we also need one additional parameter to characterize our distribution: the shift parameter. This parameter shifts the distribution so that $x > \delta$, where δ can be positive (shifted to the right) or negative (shifted to the left). The probability distribution is then

$$p(x) = \frac{1}{(x - \delta) \sigma \sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{\log(x - \delta) - \mu}{\sigma}\right)^2\right], \quad x \in \mathbb{R}^+ \quad (4.38)$$

Given access to the entire population of spine size changes, the parameters that define the log-normal distribution can be found by transforming the sample means and standard deviations (μ_s and σ_s) of the spine size changes as follows:

$$\mu_{log} = \log\left[\frac{(\mu_x + \hat{\delta})^2}{\sqrt{\sigma_s^2 + (\mu_s + \hat{\delta})^2}}\right], \quad \sigma_{log} = \sqrt{\log\left[\left(\frac{\sigma_s}{\mu_s} + \hat{\delta}\right)^2 + 1\right]} \quad (4.39)$$

where we have introduced the positive term $\hat{\delta}$, which shifts the sample mean towards positive values. The choice of $\hat{\delta}$ is relatively trivial as long as all the values of the dataset are positive after the shift. This ensures that μ_{log} is also positive, thus avoiding the log-normal distribution transformation accumulating values around $x = 0$.

Once the parameters of the lognormal have been estimated, the model uses the log-normal distribution to generate the subsequent time points. This model then takes the form

$$\Delta V_{i+1} = \eta_i, \quad \eta_i \sim \text{Lognormal}(\mu_{log}, \sigma_{log}, \hat{\delta}) \quad (4.40)$$

which mirrors the form of the conventional Wiener process. Here, we include the $-\hat{\delta}$ term to be able to map our log-normal back to the original range of values that we observe in the data. We emphasize here that the change ΔV_{i+1} for each spine is log-normal but that the population change, i.e., the collection of all changes should still be normally distributed (in accordance with Figure 2.2 c). By the central limit theorem and the assumption that the activity-independent plasticity of the spines is independent of each other, we will obtain this Gaussian nature as long as we have sufficiently many spines.

Appendix B: Inter-spine distance estimation

We are interested in estimating the mean inter-synaptic distance ($MISD$) from the data describing a linear dendritic stretch. As stated in the Methods section, each spine is given a distance value corresponding to the distance from the closest stimulation. Moreover, this value is negative in case the considered spine lies between different stimulations and positive otherwise.

We start by assuming that the linear spine density is homogeneous, i.e. it does not change with the absolute position along the dendrite at the considered length scales ($10 - 100 \mu m$). From the array of distances, we then consider only the positive values, deriving from the spines located outside the stimulation cluster; we then sort this array in ascending order. This sorted array now contains roughly double the amount of spines located on an average dendritic stretch, as we have mapped two different stretches onto one. In order to estimate the average inter-spine distance, led by this heuristic, we subsample this array, taking every second distance value. As a final step, we compute the differences between consecutive distances and use their average value as an estimate of the true inter-spine distance.

We quantify the goodness of this estimator using a Monte-Carlo sampling. We generate $N = 10^5$ different dendrites with 100 spines each. In accordance with [130], we use a Weibull distribution

$$f(x; c, s) = \frac{c}{s} \left(\frac{x}{s} \right)^{c-1} e^{-(x/s)^c} \quad (4.41)$$

with fixed scale and shape parameters (s and c) to generate random inter-synaptic distances for each dendrite. We then carry out our estimation using the procedure described above, and evaluate its performance in terms of the mean relative error

$$MRE = \left\langle \frac{\overline{MISD}}{MISD} - 1 \right\rangle \quad (4.42)$$

where \overline{MISD} is the estimated mean inter-spine distance and $MISD$ is the true theoretical value $MISD = s\Gamma(1 + 1/c)$. We quantify this metric on a set of different values of c and s , focusing on parameter ranges giving rise to mean inter-spine distances between 1 and a few microns.

The results, reported in Fig. Figure 4.19, show that the described estimator has a well-behaved, bell-shaped distribution, with an average bias of $\sim 3\%$, confined under a 30% error. We consider

this error acceptable since its absolute value ($\sim 0.8 \mu\text{m}$) is comparable with the spatial resolution of the experimental setup used to collect the data.

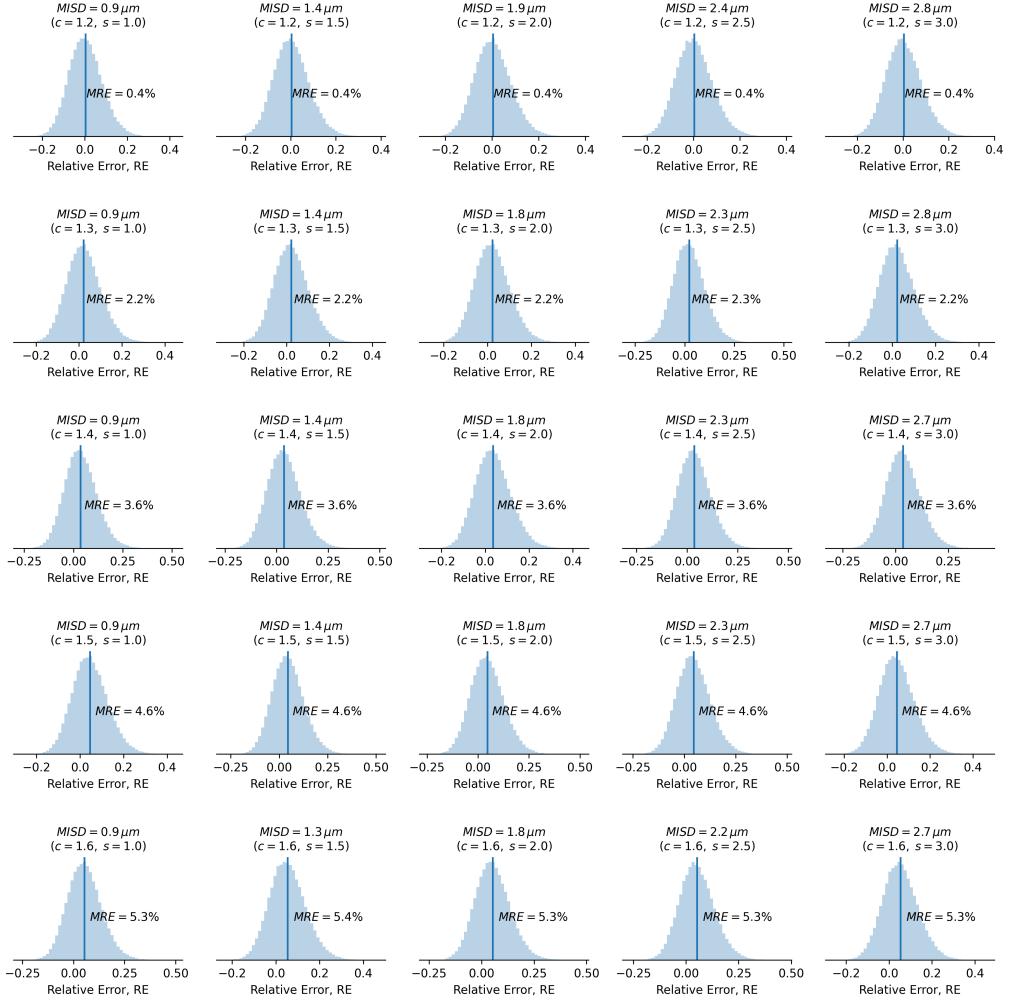


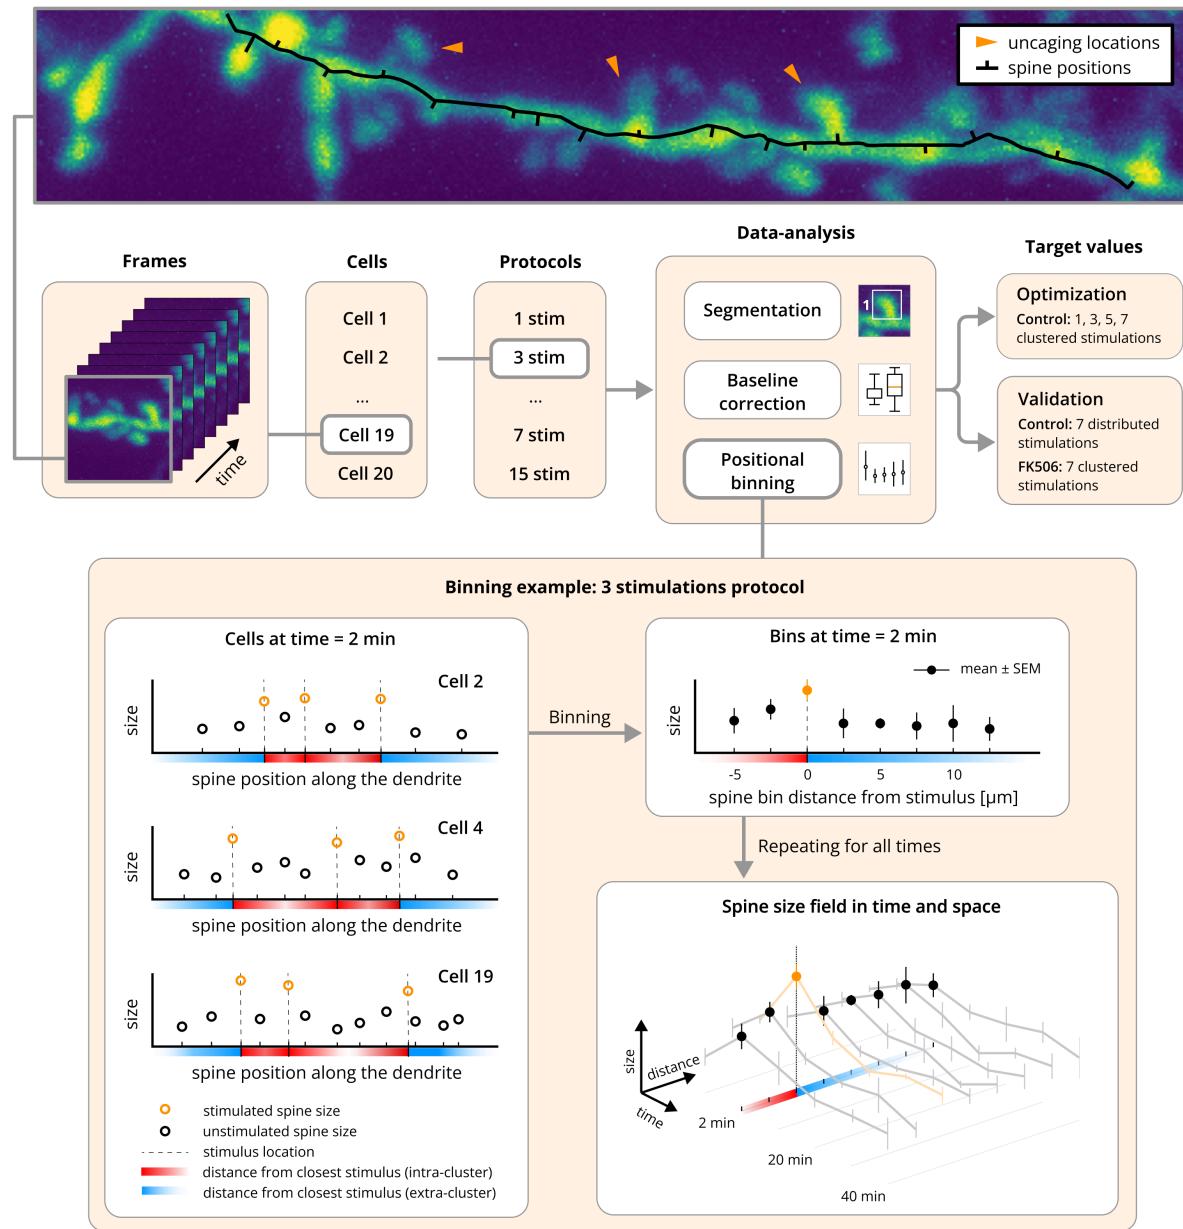
Figure 4.19: Relative error sampling of the inter-spine distance Weibull estimator. The sampling is focused on inter-spine distances compatible with hippocampal spine densities. Notice that the same average inter-spine distance (MISD) can be obtained with different values of the parameters c, s . The vertical bar shows the mean relative error (MRE) of the estimator.

Appendix C: FK506 induced parameter changes

The parameter changes are reported for the two works described in Section 3.8. In [68], a concentration of $2 \mu M$ FK506 was used, while in [98] the concentration amounts to $10 \mu M$.

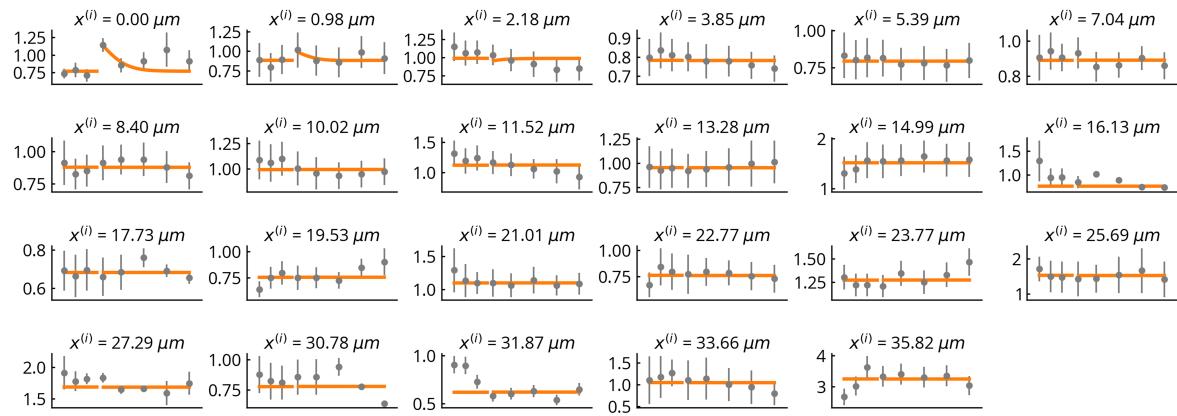
Work	Parameter		Factor	Notes
[68]	Average basally active kinases	μ_K	1.2	In accordance with [119]
	Average basally active phosphatases	μ_N	0.83	FK506 hindering effect on CaN binding to substrate [120]
	Stimulus induced phosphatases	N_s	0.55	Qualitatively following the reasoning for μ_N
	Phosphatase activation spread	σ_N	0.91	Posited effect of FK506 on calcineurin rate of binding to calcium/calmodulin
[98]	Average basally active kinases	μ_K	1.8	Considering [119] and the higher concentration of FK506 used in experiment compared to [68]
	Average basally active phosphatases	μ_N	0.53	FK506 hindering effect on CaN binding to substrate [120], higher FK506 concentration compared to [68]
	Stimulus induced phosphatases	N_s	0.40	Qualitatively following the reasoning for μ_N

Supplementary figures

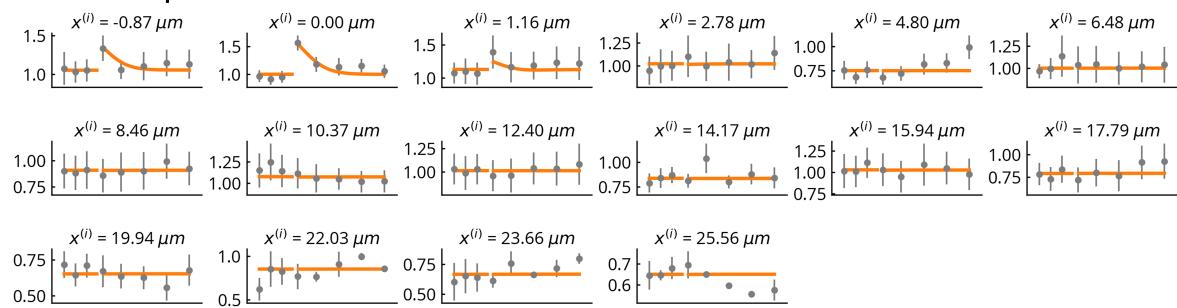


S 4.1: Graphical representation of the data analysis and pre-processing procedure. In the lower panel, an example of the binning procedure is shown: Considering a specific protocol (e.g., 3 stimulations), for each neuron, at a fixed time point, each spine gets assigned a distance value from the closest stimulus. This value is positive for spines lying outside of the stimulation cluster (blue color code) or negative (red color code). Using a bins of size equal to the inferred mean inter-spine (MISD), the spine sizes are binned, and the result reported in terms mean and standard error of the mean. Notice that the resulting “summary” dendrite contains a reduced number of spines. The procedure is repeated for all the considered time points, giving raise to a synaptic size field in function of time and distance from the closest stimulus.

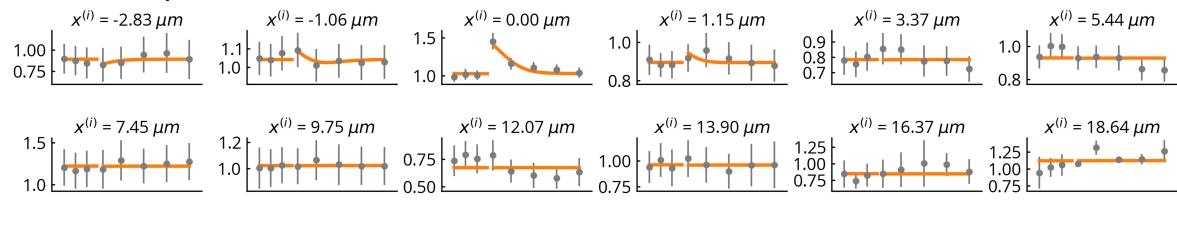
Stimulation protocol: 1 stimulation



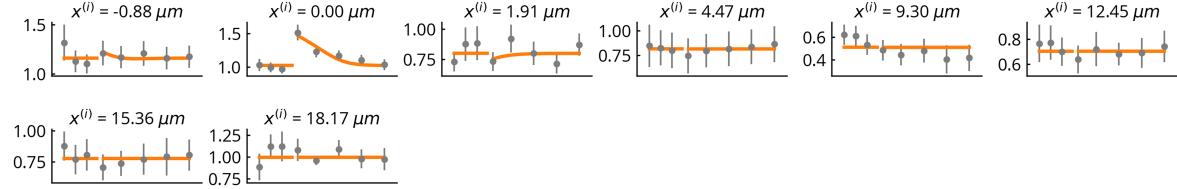
Stimulation protocol: 3 stimulations



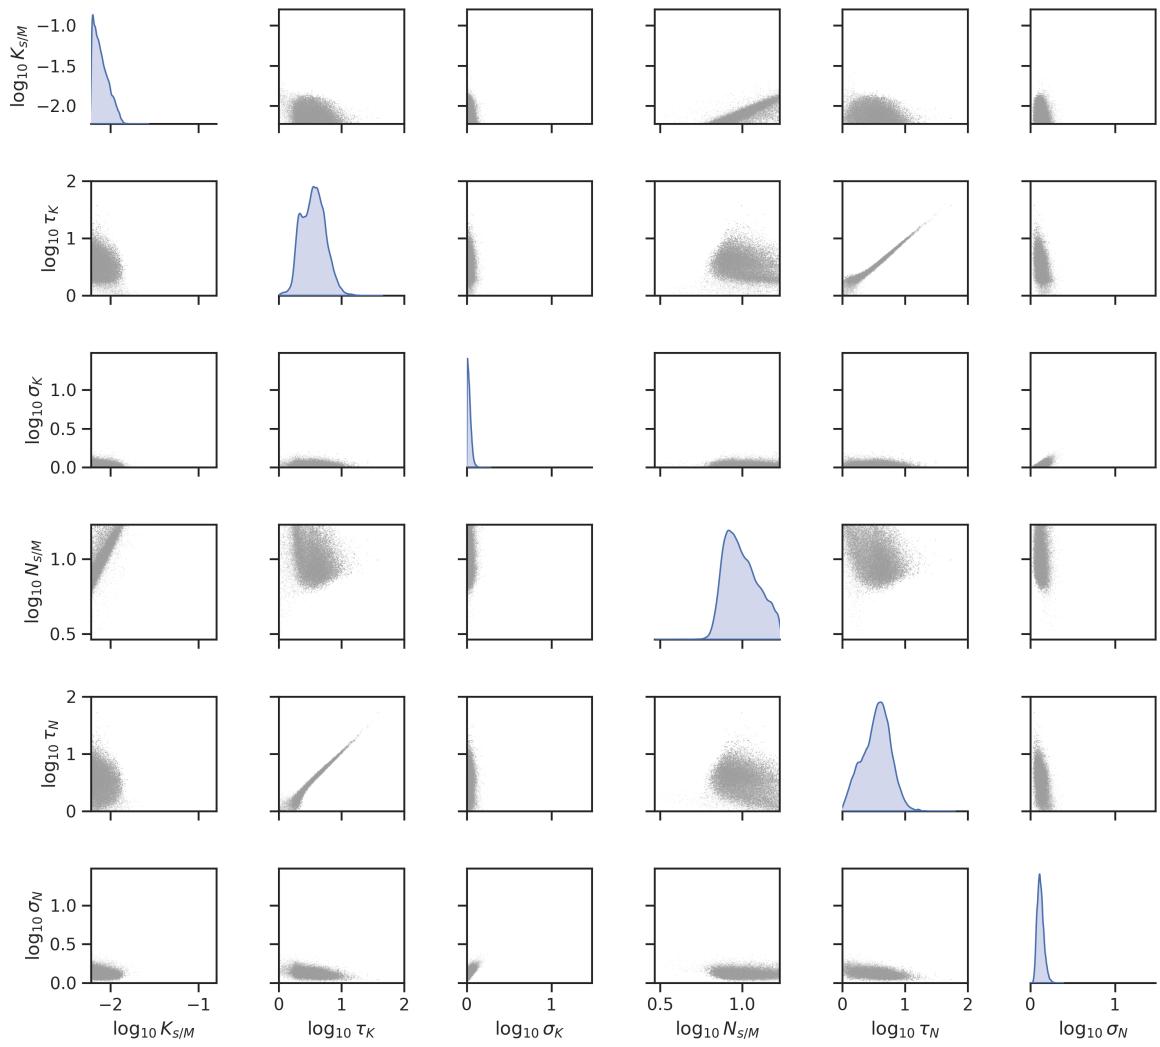
Stimulation protocol: 5 stimulations



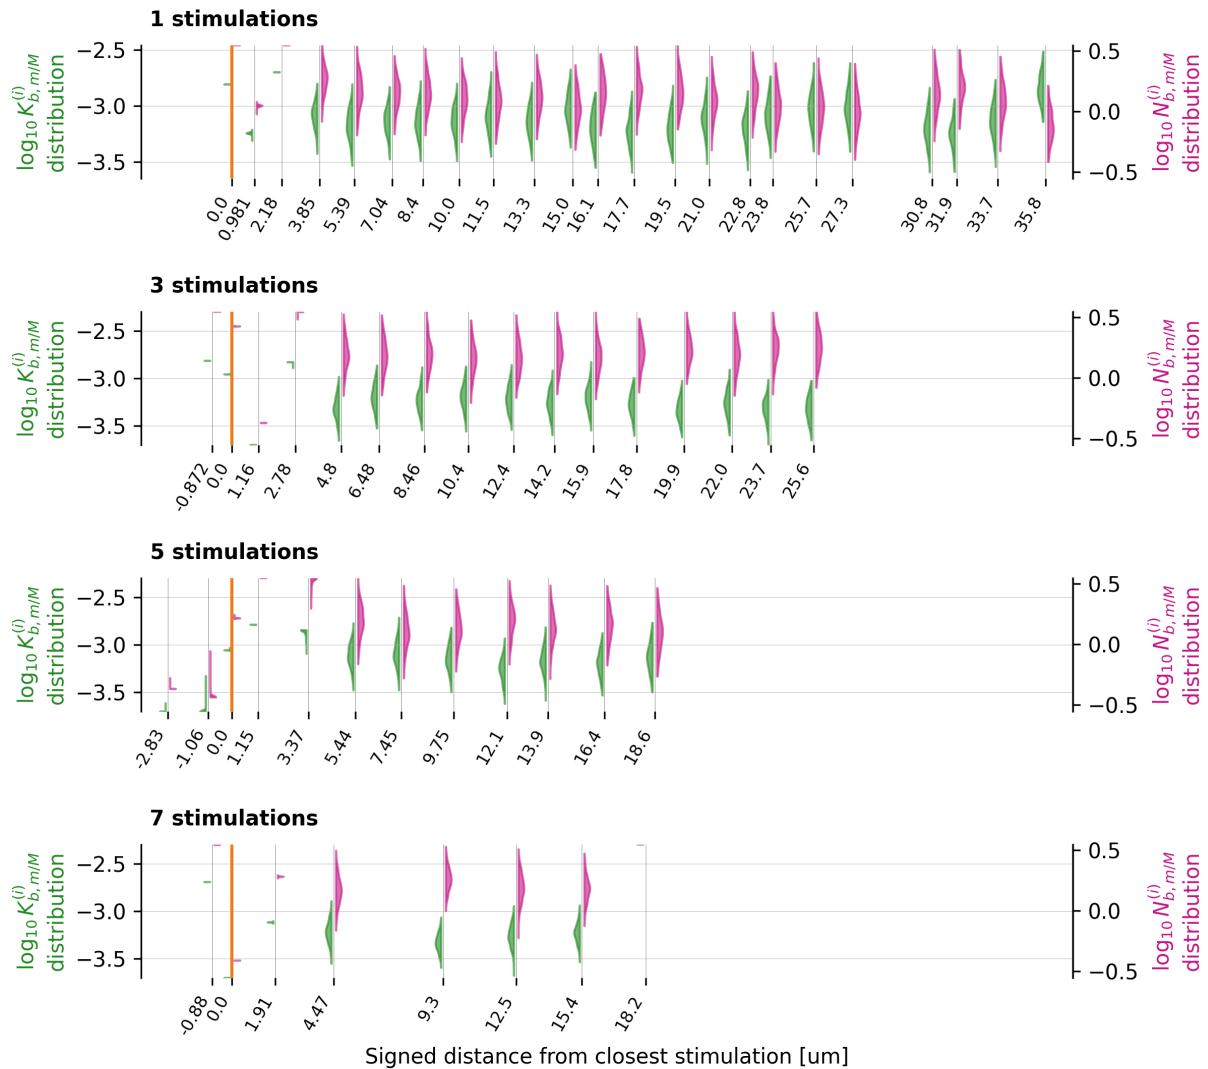
Stimulation protocol: 7 stimulations



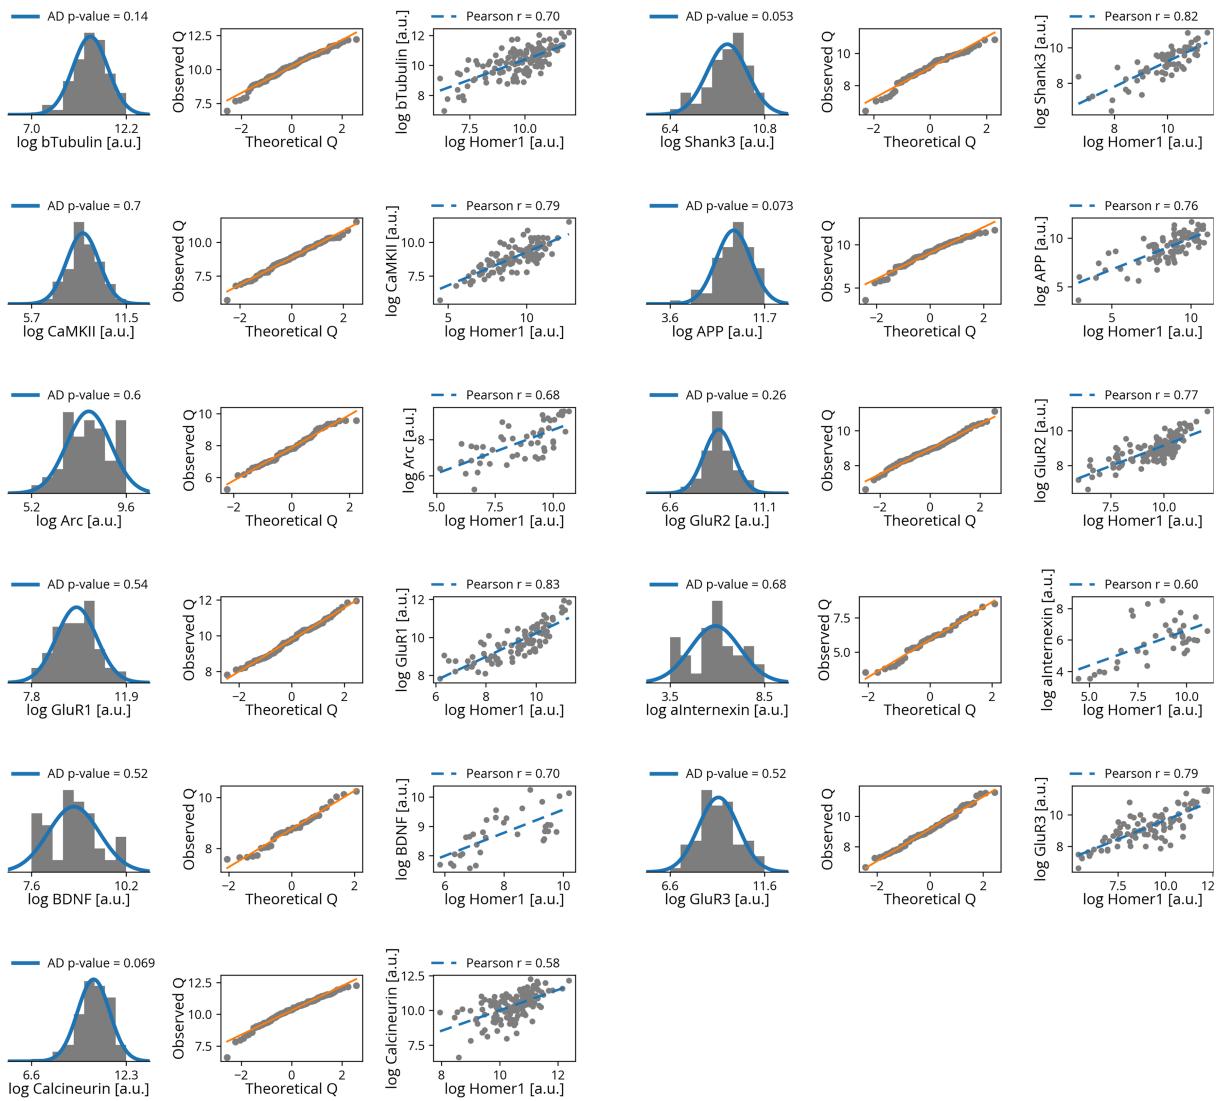
S 4.2: Fit quality of the optimized model for the experimental data For each plot, the x axis represents the time in minutes, with data located at $-15, -10, -5, 2, 10, 20, 30$ and 40 minutes (not drawn for graphical improvement).



S 4.3: Global parameters' posterior distribution obtained with adaptive parallel tempered MCMC sampling.



S 4.4: Convergence values of synaptic specific parameters across runs converging to the best local minimum LM1. Each row shows the data obtained from the corresponding protocol. The orange vertical line represents the stimulus location ($0\mu\text{m}$ distance).



S 4.5: Distribution, log-normal compatibility and correlation with Homer 1 of several synaptic proteins.

List of peer-reviewed publications

During my PhD research work, I published the following works:

- Coutinho C., Mehrab A., **Petkovic J.**, Zanchetta F., Fioresi R. (2025). Enhancing CNNs robustness to occlusions with bioinspired filters for border completion, Geometric Science of Information: 7th International Conference, GSI 2025 Proceedings, Part I. Geometric Science of Information. https://doi.org/10.1007/978-3-032-03918-7_28
- **Petkovic, J.**, & Fioresi, R. (2024). Spontaneous Emergence of Robustness to Light Variation in CNNs With a Precortically Inspired Module. *Neural Computation*, 36, 1832-1853. <https://doi.org/10.64898/2026.01.29.702571>
- **Petkovic, J.***, Fioresi, R.* , Marraffa, A.* (2023). A new perspective on border completion in visual cortex as bicycle rear wheel geodesics paths via sub Riemannian Hamiltonian formalism. *Differential Geometry and its Applications*. <https://doi.org/10.1016/j.difgeo.2024.102125>
- **Petkovic, J.***, Eggli, M.* , Chater, T.E.* , Goda, Y., & Tchumatchenko, T. (2023). Linking spontaneous and stimulated spine dynamics. *Communications Biology*, 6. <https://doi.org/10.1038/s42003-023-05303-1>