

A brain inspired sequence learning algorithm and foundations of a memristive hardware implementation

Younes Bouhadjar

Information

Band / Volume 95

ISBN 978-3-95806-693-9

Forschungszentrum Jülich GmbH
Institute of Neurosciences and Medicine (INM)
Computational and Systems Neuroscience (INM-6)

A brain inspired sequence learning algorithm and foundations of a memristive hardware implementation

Younes Bouhadjar

Schriften des Forschungszentrums Jülich
Reihe Information / Information

Band / Volume 95

ISSN 1866-1777

ISBN 978-3-95806-693-9

Bibliografische Information der Deutschen Nationalbibliothek.
Die Deutsche Nationalbibliothek verzeichnet diese Publikation in der
Deutschen Nationalbibliografie; detaillierte Bibliografische Daten
sind im Internet über <http://dnb.d-nb.de> abrufbar.

Herausgeber und Vertrieb: Forschungszentrum Jülich GmbH
Zentralbibliothek, Verlag
52425 Jülich
Tel.: +49 2461 61-5368
Fax: +49 2461 61-6103
zb-publikation@fz-juelich.de
www.fz-juelich.de/zb

Umschlaggestaltung: Grafische Medien, Forschungszentrum Jülich GmbH

Druck: Grafische Medien, Forschungszentrum Jülich GmbH

Copyright: Forschungszentrum Jülich 2023

Schriften des Forschungszentrums Jülich
Reihe Information / Information, Band / Volume 95

D 82 (Diss. RWTH Aachen University, 2023)

ISSN 1866-1777
ISBN 978-3-95806-693-9

Vollständig frei verfügbar über das Publikationsportal des Forschungszentrums Jülich (JuSER)
unter www.fz-juelich.de/zb/openaccess.



This is an Open Access publication distributed under the terms of the [Creative Commons Attribution License 4.0](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

“[...] To some degree we resemble the isolated tribe in New Guinea that in the 1940s encountered a crashed airplane and studied it without comprehending its primary function. Nevertheless, we can learn from the engineers: we should try to state the brain’s primary goal and basic performance specifications. We should try to intuit a role for each part. By placing the data in some framework, we can begin to evaluate how well our device works and begin to consider the why of its design. We will make this attempt, even though it will be incomplete, and sometimes wrong.”

— [Sterling and Laughlin \(2015\)](#), Principles of Neural Design.

SUMMARY

The brain uses intricate biological mechanisms and principles to solve a variety of tasks. These principles endow systems with self-learning capabilities, efficient energy usage, and high storage capacity. A core concept that lies at the heart of brain computation is sequence prediction and replay. This form of computation is essential for almost all our daily tasks such as movement generation, perception, and language. Understanding how the brain performs such a computation advances neuroscience and paves the way for new technological brain-inspired applications.

In the first part of this thesis, we propose a sequence learning model that explains how biological networks learn to predict upcoming elements, signal non-anticipated events, and recall sequences in response to a cue signal. The model accounts for anatomical and electrophysiological properties of cortical neuronal circuits, and learns complex sequences in an unsupervised manner by means of known biological plasticity and homeostatic control mechanisms. After learning, it self-organizes into a configuration characterized by a high degree of sparsity in connectivity and activity allowing for both high storage capacity and efficient energy usage.

In the second part, we extend the sequence learning model such that it permits probabilistic sequential memory recall in response to ambiguous cues. In the absence of noise, the model deterministically recalls the sequence shown most frequently during training. We investigate how different forms of noise give rise to more exploratory behavior. We show that uncorrelated noise averages out in population based encoding leading to non-exploratory dynamics. Locally coherent noise in the form of random stimulus locking to spatiotemporal oscillations addresses this issue. Our results show that depending on the amplitude and frequency of oscillation, the network can recall learned sequences according to different strategies: either always replay the most frequent sequence, or replay sequences according to their occurrence probability during training. The study contributes to an understanding of the neuronal mechanisms underlying different decision strategies in the face of ambiguity, and highlights the role of coherent network activity during sequential memory recall.

Finally, we investigate the feasibility of implementing the sequence learning model on dedicated hardware mimicking brain properties. Here, we focus on a type of hardware where synapses are emulated by memristive devices. As a first step in this direction, we replace the synapse dynamics of the original model with dynamics describing the phenomenological behavior of memristive elements, and demonstrate resilience with respect to different device characteristics. In this thesis, we further describe how the sequence learning model can be adapted at the algorithmic level to foster an implementation in a full electronic circuit centered around a memristive crossbar array.

Overall, this thesis sheds light on the key mechanisms underlying sequence learning, prediction, and replay in biological networks and demonstrates the feasibility of implementing this type of computation on neuromorphic hardware.

ZUSAMMENFASSUNG

Das Gehirn nutzt komplexe biologische Mechanismen und Prinzipien, um eine Vielzahl von Aufgaben zu lösen. Diese Prinzipien verleihen den Systemen selbstlernende Fähigkeiten, sowie eine hohe Energieeffizienz und Speicherkapazität. Ein zentrales Konzept, bei der Informationsverarbeitung im Gehirn, ist die Sequenzvorhersage und -wiedergabe. Diese Form der Berechnung ist essenziell für fast alle täglichen Aufgaben wie Sinneswahrnehmung, Bewegungsausführung und Sprache. Zu verstehen wie das Gehirn eine solche Berechnung durchführt, treibt die Neurowissenschaft voran und ebnet den Weg für neue technologische, vom Gehirn inspirierte Anwendungen.

Im ersten Teil dieser Arbeit schlagen wir ein Modell zum Sequenzlernen vor, das erklärt, wie biologische Netzwerke lernen können, nachfolgende Elemente vorherzusagen, nicht erwartete Ereignisse zu signalisieren und Sequenzen als Reaktion auf ein Hinweis-Signal abzurufen. Das Modell berücksichtigt anatomische und elektrophysiologische Eigenschaften kortikaler neuronaler Netzwerke und lernt komplexe Sequenzen auf unüberwachte Weise mittels bekannter biologischer Plastizität und homöostatischer Kontrollmechanismen. Während des Lernens organisiert es sich selbst in eine Konfiguration, die sich durch einen hohen Grad an Spärlichkeit in der Konnektivität und Aktivität auszeichnet, was sowohl eine hohe Speicherkapazität und Energieeffizienz zur Folge hat.

Im zweiten Teil erweitern wir das Modell so, dass es einen probabilistischen sequentiellen Gedächtnisabruf als Reaktion auf mehrdeutige Hinweise ermöglicht. In Abwesenheit von Rauschen ruft das Modell deterministisch die im Training am häufigsten gezeigte Sequenz ab. Wir untersuchen, wie verschiedene Formen von Rauschen zu einem explorativeren Verhalten führen. Wir zeigen, dass sich unkorreliertes Rauschen bei populationsbasierter Kodierung wegmittelt, was zu einer nicht-explorativen Dynamik führt. Dieses Problem wird durch lokal kohärentes Rauschen in Form einer zufälligen Stimuluskopplung an intrinsische, raum-zeitliche Oszillationen behoben. Unsere Ergebnisse zeigen, dass das Netzwerk, je nach Amplitude und Frequenz der Oszillation, gelernte Sequenzen gemäss verschiedener Strategien abrufen kann: entweder wird immer die häufigste Sequenz wiedergegeben, oder Sequenzen werden entsprechend ihrer Auftrittswahrscheinlichkeit beim Training abgespielt. Die Studie trägt zum Verständnis der neuronalen Mechanismen bei, die verschiedenen Entscheidungsstrategien zugrunde liegen, und unterstreicht die Rolle kohärenter Netzwerkaktivität bei sequenziellem Gedächtnisabruf.

Schließlich untersuchen wir die Realisierbarkeit der Implementierung des Sequenzverarbeitungsmodells auf einer speziellen Hardware, welche Eigenschaften des Gehirns imitiert. Hier konzentrieren wir uns auf eine Art von Hardware, bei der die Synapsen durch memristive Bauteile emuliert werden. Als ersten Schritt in diese Richtung ersetzen wir die Synapsen-Dynamik des ursprünglichen Modells durch eine Dynamik, die das phänomenologische Verhalten von memristiven Elementen beschreibt, und demonstrieren Robustheit der Vorhersagen bei Variation verschiedener Hardware-Eigenschaften. In dieser Arbeit beschreiben wir außerdem, wie das Sequenzlernmodell auf algorithmischer Ebene angepasst werden kann, um eine Implementierung in eine vollelektronische Schaltung basierend auf einem memristiven Crossbar-Array zu ermöglichen.

Insgesamt wirft diese Arbeit Licht auf die prinzipiellen Mechanismen, dem Lernen, der Vorhersage und dem Abruf von Sequenzen in biologischen Netzwerken zugrunde liegen, und demonstriert die Realisierbarkeit der Implementierung dieser Art von Berechnung auf neuromorpher Hardware.

DECLARATION ON PUBLICATIONS

The presented thesis is based on the publications and preprints below.

Constraints on sequence processing speed in biological neuronal networks
by Younes Bouhadjar, Markus Diesmann, Rainer Waser, Dirk J. Wouters, Tom Tetzlaff

Published in Proceedings of the International Conference on Neuromorphic Systems (Bouhadjar et al., 2019).

Sequence learning, prediction, and replay in networks of spiking neurons
by Younes Bouhadjar, Dirk J. Wouters, Markus Diesmann, Tom Tetzlaff

Published in PLOS Computational Biology (Bouhadjar et al., 2022b).

This publication is the basis of chapter 2.

Coherent noise enables probabilistic sequence replay in spiking neuronal networks

by Younes Bouhadjar, Dirk J. Wouters, Markus Diesmann, Tom Tetzlaff

In press at PLOS Computational Biology and published on arXiv (Bouhadjar et al., 2022a).

This publication is the basis of chapter 3.

Sequence learning in a spiking neuronal network with memristive synapses
by Younes Bouhadjar, Sebastian Siegel, Tom Tetzlaff, Markus Diesmann, Rainer Waser, Dirk J. Wouters

Manuscript under review and published on arXiv (Bouhadjar et al., 2022c).

This manuscript is the basis of chapter 4.

System model of neuromorphic sequence learning on a memristive crossbar array

by Sebastian Siegel, Younes Bouhadjar, Tom Tetzlaff, Rainer Waser, Regina Dittmann, Dirk J. Wouters

Manuscript under review (Siegel et al., 2023).

This manuscript is the basis of chapter 5.

The contribution of each author is indicated at the beginning of each chapter.

Additional related publications discussed only briefly in this thesis:

Learning and Replaying Spatiotemporal Sequences: A Replication Study

by Jette Oberländer, Younes Bouhadjar, Abigail Morrison

Published in Frontiers in Integrative Neuroscience (Oberländer et al., 2022).

Prediction: An Algorithmic Principle Meeting Neuroscience and Machine Learning Halfway

by Younes Bouhadjar*, Caterina Moruzzi*, Melika Payvand*

Published in CEUR Workshop Proceedings (Bouhadjar et al., 2022).

* Equal contribution.

Demonstration of neuromorphic sequence learning on a memristive array
by *Sebastian Siegel, Tobias Ziegler, Younes Bouhadjar, Tom Tetzlaff, Rainer Waser, Regina Dittmann, and Dirk J. Wouters*
In press at the Proceedings of the Neuro-inspired Computational Elements Workshop
([Siegel et al., 2023](#)).

CONTENTS

1	Introduction	1
1.1	Sequence prediction and replay	2
1.2	The neocortex	3
1.2.1	Organization of the neocortex	4
1.2.2	Neurons and synapses	6
1.2.3	Plasticity	7
1.2.4	Cortical activity	9
1.3	The Hierarchical Temporal Memory model	10
1.3.1	HTM regions and columns	10
1.3.2	HTM neurons and synapses	12
1.3.3	HTM cortical learning algorithm	13
1.3.4	Functional role of the cortical layers	13
1.4	Neuromorphic hardware	14
1.4.1	Memristive devices	15
1.4.2	Crossbar arrays	16
1.4.3	Redox-based random access memory devices	17
1.4.4	Valance change memory devices	18
1.5	Scope and structure of this thesis	18
2	Sequence learning in a network of spiking neurons	21
2.1	Introduction	22
2.2	Methods	22
2.2.1	Task and training protocol	22
2.2.2	Network model	24
2.2.3	Task performance measures	29
2.2.4	Model tables	34
2.3	Results	39
2.3.1	Sequence learning and prediction	39
2.3.2	Prediction performance	41
2.3.3	Dependence of prediction performance on the sequence speed	43
2.3.4	Sequence replay	45
2.4	Discussion	46
2.4.1	Summary	46
2.4.2	Relationship to other models	47
2.4.3	Conclusion	50
3	Probabilistic sequence replay in spiking neural networks	51
3.1	Introduction	52
3.2	Methods	53
3.2.1	Learning protocol and task	53
3.2.2	Network model	54
3.2.3	Task performance measures	56
3.2.4	Model tables	57
3.3	Results	62
3.3.1	A spiking neural network recalls sequences in response to ambiguous cues	62

3.3.2	Noise canceling cannot be overcome by increasing noise amplitude	67
3.3.3	Noise amplitude and level of correlation control replay strategy	68
3.3.4	Random stimulus locking to spatiotemporal oscillations as natural form of noise	70
3.4	Discussion	71
4	Sequence learning in a spiking neural network with memristive synapses	75
4.1	Introduction	76
4.2	Methods	78
4.2.1	Model tables	78
4.3	Results	82
4.3.1	ReRAM synapse model	82
4.3.2	A spiking neural networks with ReRAM synapses successful at sequence prediction	84
4.3.3	Influence of device characteristics on prediction performance	86
4.4	Discussion	93
4.4.1	Summary	93
4.4.2	Relationship to previous models	94
4.5	Conclusion	95
5	An algorithm for sequence learning in a memristive crossbar array	97
5.1	Introduction	98
5.2	Basic architectural concepts	98
5.2.1	Network structure	98
5.2.2	Neuron functionality	98
5.2.3	Synaptic functionality	98
5.3	Algorithmic adaptations	99
5.4	Memristive synapse model	101
5.5	Prediction performance results	101
5.6	Discussion	102
5.7	Conclusion	103
6	Discussion	105
6.1	Conclusion	106
6.2	Outlook	108
6.2.1	Testable predictions of the sequence learning model	108
6.2.2	Limitation and extensions of the sequence learning model	109
6.2.3	Prospects on neuromorphic computing	112
7	Supporting information	115
7.1	Supporting information chapter 2	115
7.1.1	Model and simulation parameters	115
7.2	Supporting information chapter 3	123
7.2.1	Model and simulation parameters	123
7.3	Supporting information chapter 4	126
7.3.1	Model and simulation parameters	126
	Bibliography	129

INTRODUCTION

Short excerpts of this chapter are based on the following publication:
Bouhadjar, Y., Diesmann, M., Wouters, D. J., & Tetzlaff, T. (2022). Sequence learning, prediction, and replay in networks of spiking neurons. *PLOS Computational Biology* 18 (6), e1010233.

Prediction, almost continually operative at conscious and reflex levels, is pervasive throughout most, if not all, levels of brain function.

[Llinás \(2002\)](#)

The brain is a complex network of nerve cells, which carries out cognitive computation in a distributed and energy-efficient manner. It uses intricate biological mechanisms and principles to solve a wide spectrum of problems ranging from perception and action to high-order cognitive functions such as planning and reasoning. Understanding the brain machinery is essential for driving new advances in neuroscience and for developing new real-world applications. A primary function performed by the brain is sequence prediction and recall. In this thesis, we aim at identifying the neural mechanisms and principles essential for this type of computation and determining hardware primitives that implement these concepts. To provide an answer to these questions, we follow a multidisciplinary approach following these four steps:

- Describe the task solved by the biological system using knowledge from behavioral neuroscience and psychophysics.
- Investigate the neuronal implementation of this task using knowledge from biology.
- Develop mathematical models describing the biological implementation and test these using computer simulations.
- Propose dedicated hardware implementation of the proposed models.

The introduction clarifies these steps and gives the reader an overview of the topics involved in this thesis. It starts with laying down a set of arguments of why sequence prediction and recall might be a central computation performed by the brain. Section 1.2 explains the organizational principles and components of a structure in the brain known as the neocortex. The following section 1.3 introduces the Hierarchical Temporal Memory (HTM) model, which links the working principles of the neocortex and sequence processing tasks. Section 1.4 discusses emerging neuromorphic hardware with a focus on that based on memristive devices. Finally, section 1.5 outlines the scope of the thesis and the questions addressed.

1.1 SEQUENCE PREDICTION AND REPLAY

Learning and processing sequences of events, objects, or percepts are fundamental computational building blocks of cognition ([Lashley, 1951](#); [Hawkins and Blakeslee, 2007](#); [Dehaene et al., 2015](#); [Clegg et al., 1998](#)). Data processed by mammals and many other living organisms is often sequential. This holds true for all types of sensory input data as well as motor output activity. Being able to form memories of such sequential data, predict future sequence elements, and replay learned sequences is a necessary prerequisite for survival.

Prediction was suggested to be the primary computation performed by the neocortex ([Hawkins and Blakeslee, 2007](#); [Bar, 2007](#); [Llinás, 2002](#)). As we move in our environment, we are constantly making predictions about what we are going to experience next. We effortlessly detect surprising, non-anticipated events and adjust our behavior

accordingly. When listening to a familiar melody, we can predict the next note before it has been played. We can even anticipate what someone is about to say before they have finished (Hawkins and Blakeslee, 2007). We can also make predictions across different sensory modalities. Listening to the voice of someone can trigger a rich repertoire of visual or even somatosensory memories. Often we are not aware of these predictions till some event violates our expectations and then we become immediately attentive. For instance, as we casually lift a cardboard box. Our brains predict what force is required to carry the box. If our prediction is wrong and the box suddenly becomes heavier, we notice the difference and correct immediately our behavior (Llinás and Roy, 2009; Hawkins and Blakeslee, 2007).

The ability to form predictions is not restricted only to sensory processing but also a component in higher cognitive functions such as planning and language (Bar, 2007). Often, these predictions are made in a context-dependent manner (Bar, 2004). In language, for example, we use context to disambiguate a sentence (Bod et al., 2003). Other cognitive processes use the context for fast and efficient interaction with the world. Take visual processing for instance: our environment dictates specific expectations of what we are about to see, for example, when we are in the library, we expect to see a specific set of objects such as bookshelves, desks, books, etc.

Predictions are often embedded and occur within sequences. Listening to a melody for instance evokes a series of predictions of what are the most likely next notes. In many situations, we cannot make predictions if the received information is non-sequential. For example, if we put our hands on an object, we would not recognize its texture only when we start moving our hands (Hawkins and Blakeslee, 2007). Our sensory organs are tuned to detect changes in the stimuli rather than capturing static inputs. When we look at a picture our eyes typically do not stand still but explore the picture following small movements known as saccades (Yarbus, 1967). They fixate on one position, and then suddenly jump to another one. They keep moving, even if we think they are fixating on a certain part of the image. The brain receives these sequences of information, makes sense of them, and then gives rise to the perception of the standing still image. If the eye saccades are artificially eliminated, the vision fades (Riggs and Ratliff, 1952; Martinez-Conde et al., 2004).

As we can carry out sequential predictions, we also manage to replay learned sequences, for example, when generating motor behavior or replaying sequential memories. These forms of processing have been studied extensively in a number of experimental works on sensory processing (Gavornik and Bear, 2014; Xu et al., 2012), motor production (Hahnloser et al., 2002), and decision making (Harvey et al., 2012).

The underlying mechanisms for sequence prediction and recall in the brain remain largely unknown. As pointed out by Lashley (1951), understanding how this type of processing is implemented by the brain is essential for understanding the working principles of the neural system.

1.2 THE NEOCORTEX

Sequence prediction and recall are principal computations performed by the brain and an integral part of cognition. These types of computations are implemented by a number of structures in the brain (Clegg et al., 1998). In this thesis, we investigate its neural circuitry in the neocortex.

The neocortex is believed to be the seat of intelligence in the brain. It allows for planning, thinking, and reasoning about the world. When hearing a tone, seeing an

object, or touching a cup, the sensory information is received by our sensory organs, preprocessed, packed, and sent via electrical pulses to the brain (Kandel et al., 2000; Bear et al., 2007). The sensory information is further processed by a number of structures in the brain till it reaches the neocortex. The neocortex then makes sense of the incoming sensory data giving rise to perception and other cognitive processes. Before discussing how sequence learning contributes to these abilities, we first dive into the organizing principles of the neocortex and its main components¹.

1.2.1 Organization of the neocortex

The neocortex is the outer region of the mammalian brain. In humans, it has a thickness of a few millimeters, a surface of around $\sim 2600 \text{ cm}^2$ (Mountcastle, 1997), and a total of 16 billion cortical neurons, each receiving up to 10,000 connections (Abeles, 1991). It is composed of two hemispheres, each subdivided into four lobes: frontal, parietal, temporal, and occipital (Fig. 1.1A). These lobes have been associated with different functions (Kandel et al., 2000). The frontal lobe is responsible for movement generation and high-order cognitive functions such as planning and decision-making. The other lobes are associated with sensory processing such as somatic sensation (parietal), vision (occipital), and hearing (temporal). In the early twentieth century, Korbinian Brodmann came up with a more detailed division of the neocortex. He divided it into 52 areas using its different morphological and cellular organization (Brodmann, 1909). These areas have been linked later on to specific functions (Payne and Peters, 2002). For example, areas 1, 2, and 3 are located in the parietal lobe and responsible for somatosensory processing, and area 22 is located in the temporal lobe and responsible for auditory processing.

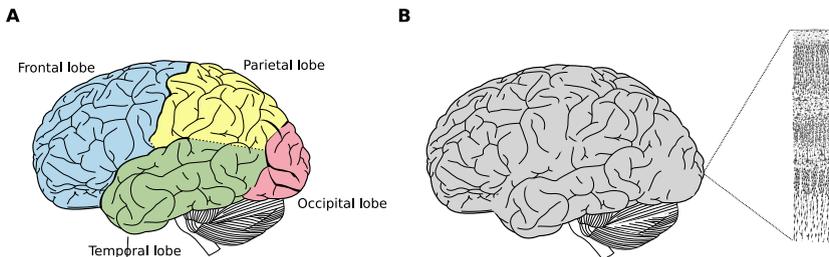


Figure 1.1: **Neocortical tissue.** The neocortex occupies the outer surface of the brain and is composed of four cortical lobes (A), each of which is constituted of repetitive structures known as macrocolumns or columns (B). Main figures in A and B are adapted from (Gray, 1918). The inset in B is adapted from (Ramon y Cajal, 1899).

HIERARCHIES In the neocortex, the areas are organized into a hierarchy. The processed information is communicated between lower and higher areas using feedforward and feedback connections (Felleman and Van Essen, 1991). The primary sensory areas first receive the information and then send it to higher areas in the hierarchy. Neurons across the different areas respond to different stimuli in the environment, i.e., a neuronal characteristic known as a receptive field (Hubel and Wiesel, 1962). Lower

¹ See also (Thiboust, 2020) for a brief and concise summary of key insights on the structure and function of the neocortex.

areas detect simple features in the sensory stimuli. Detected features become more abstract the higher the level in the hierarchy. For example, neurons in the primary visual cortex (or V1 cortex) receive information from the eyes and can detect simple features such as lines and edges (Hubel and Wiesel, 1962, 1968). Neurons in the next areas further process this information giving rise to more complex features (Milner and Goodale, 2006) such as colors and shapes (V2 area), visual motion (MT area), and object selectivity (IT area). Higher-level regions are known by their ability to form invariant representations (DiCarlo et al., 2012; Hawkins and Blakeslee, 2007). Unlike low-level regions, they can detect and recognize complex objects, regardless of environmental conditions such as lightning, position, or orientation.

MINICOLUMN The anatomical organization of the neocortex is very regular. Each one of the cortical areas is composed of a number of columnar structures, called minicolumns (Mountcastle, 1997; Braitenberg and Schüz, 1998). This structure repeats itself throughout the whole cortical sheet (Fig. 1.1B). Conventionally, the minicolumns are divided into six layers, with layers 2 and 3 often considered to constitute the same layer 2/3 (Fig. 1.2). However, layer thickness, neuronal types, and neuronal densities differ across the cortical areas. Adjacent minicolumns are innervated by the same axons. They form larger structures known as columns or macrocolumns (Horton and Adams, 2005; Thiboust, 2020). Each one of the minicolumns contains around 100 neurons (Mountcastle, 1997).

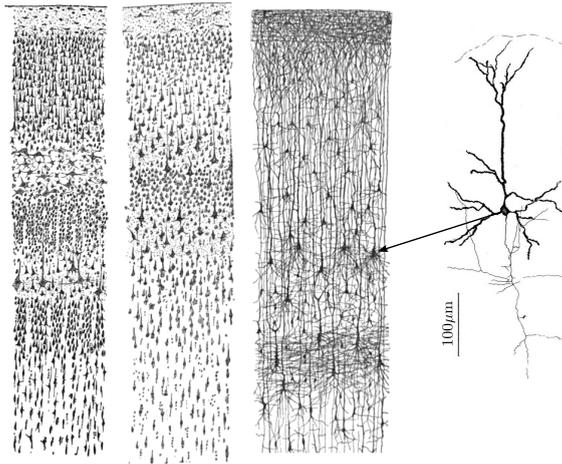


Figure 1.2: **Cortical layers.** Left: neurons in the visual, motor, and infant cortex, respectively, drawn by (Ramon y Cajal, 1899). Right: pyramidal neuron with its extended dendritic tree. Reprinted with permission from (Braitenberg and Schüz, 1991).

CONNECTIVITY A detailed dissection of the cortical connectivity reveals a unique connectivity pattern, both horizontally and vertically (Shipp, 2007). A variant of the same motif is reproduced in all cortical areas giving rise to the concept of the canonical circuit (Thomson et al., 2002; Douglas and Martin, 2004; Harris and Shepherd, 2015). Each cortical area interacts with other areas and subcortical structures with an

abundant fiber of coated myelin traveling underneath the neocortex in the white matter (Debanne, 2004). The neocortex is connected to the outside world through the thalamus, which acts as a gateway for both input and output. It innervates cortical areas via layers 4 and 3 and receives feedback connections from layer 5 (Thomson and Bannister, 2003). In the neocortex, the inputs from lower areas to higher areas first make contact with layer 4 and then project to layers 2 and 3 within the same cortical column. The information is subsequently transmitted to the next level of the cortical hierarchy. The feedback connections from higher-level areas to lower areas take a different path. The information is propagated from layer 6 in the upper cortical area to layer 1 in the lower areas and then sent to layers 2 and 3. This in turn excites cells in layers 5 and 6 (Hawkins and Blakeslee, 2007). Layer 5 closes the loop by innervating back the thalamus and other subcortical structures. Cortical neurons make also extensive lateral connections, especially in layers 2 and 3 (Binzegger et al., 2004). The connections appear to be mostly patchy in appearance, i.e., neurons with similar receptive fields connect to each other (Vanni et al., 2020) and are distance-dependent (Jiang et al., 2015; Boucsein et al., 2011). Distant cortical areas also connect to each other using long range myelinated fibers (Vanni et al., 2020).

1.2.2 *Neurons and synapses*

The neurons are electrically excitable cells and communicate with each other using voltage excursions called action potentials or spikes (Humphries, 2021, Fig. 1.3A). They are composed of a soma (or cell body), dendrites, and an axon. At reset, the soma maintains a gradient of ionic concentration between the intracellular and extracellular spaces resulting in a resting potential of about -70 mV (Dayan and Abbott, 2001). The soma collects inputs from other neurons using dendrites and generates a spike in case of sufficient stimulation. The spike travels across the axon till it reaches the synapses, where the spike from one neuron is transmitted to another neuron. The arrival of the spike at the synapse causes a release of neurotransmitters into the synaptic cleft. Most neurons in the cortex tend to release the same type of neurotransmitter on all their synapses with either an inhibitory or excitatory effect (see Dale’s law, Eccles et al., 1954). An excitatory neurotransmitter such as glutamate depolarizes the recipient neuron making the membrane potential less negative. The neuron generates a spike if it is sufficiently depolarized. An inhibitory neurotransmitter, on the other hand, such as gamma-aminobutyric acid (GABA) causes a hyperpolarization making the membrane potential of the recipient cell more negative (Kandel et al., 1991). The amount of this synaptic release and the density of the receptor on the postsynaptic neurons dictate the synaptic strength of the connection, which is often referred to as the synaptic weight. The majority of neurons in the cortex are excitatory (85% of which 75% are pyramidal neurons), while the remaining 15% are inhibitory (Abeles, 1991).

PYRAMIDAL NEURONS Pyramidal neurons are the primary source of excitation in the cerebral cortex. They can form not only connections within the same column but also with other columns or even other cortical areas. They populate all cortical layers (except layer 1), with the largest numbers in layers 6 and 5, and can exhibit different morphologies (Ledgergerber and Larkum, 2010). Their cell body is roughly pyramid shaped, hence their name. They are characterized by a distinct dendritic arbor composed of basal and apical dendrites (Spruston, 2008). Basal dendrites are close

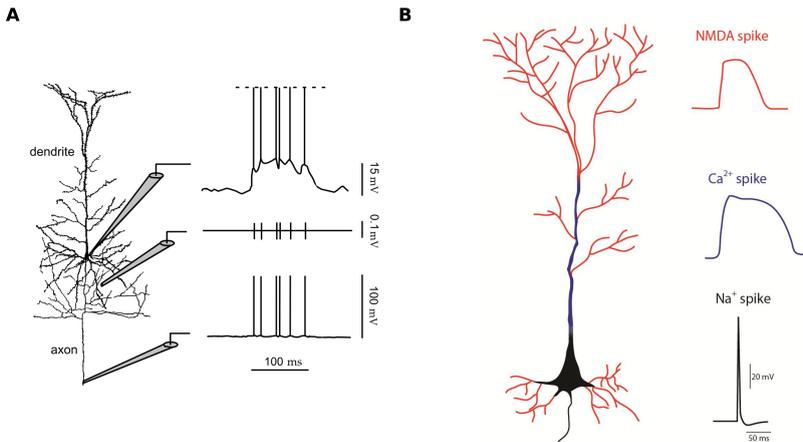


Figure 1.3: **Action potentials.** **A)** Simulated intracellular recording of the soma (top electrode) and the axon at a certain distance from the soma (bottom electrode) as well as an extracellular recording (middle electrode). Reprinted with permission from (Dayan and Abbott, 2001). **B)** The shape of three types of action potentials: sodium (black), calcium (blue), and NMDA spikes (red), as well as their initiation sites. Adapted figure from (Augusto and Gambino, 2019).

to the soma and apical dendrites can extend much further reaching the pial surface. Throughout the dendritic arbor, the pyramidal neuron collects its inputs coming from other neurons. Coincident activation of a cluster of 8–20 synapses localized on the same dendritic branch can trigger a dendritic action potential (dAP, for reviews, see Stuart and Spruston, 2015; Major et al., 2013; London and Häusser, 2005). There are three types of dendritic action potentials: sodium, calcium, and NMDA spikes. These active processes exhibit different temporal profiles. For example, the activation of the NMDA dAP results in strong and prolonged depolarization of the soma (~ 50 – 200 ms). Sodium and calcium spikes are less extended and can be activated for a duration of about 2 ms and 50 ms, respectively. They also occupy different regions within the dendritic arbor. NMDA and sodium spikes tend to be located on the basal and apical dendrites. Calcium spikes populate the apical trunk (Fig. 1.3B).

INHIBITORY NEURONS Inhibitory neurons populate all cortical layers and release the neurotransmitter GABA (Tremblay et al., 2016). Unlike pyramidal neurons, their activity remains local within the cortical column. They are highly diverse in morphology, firing patterns, and connectivity. The main function of inhibitory neurons is to balance excitation in cortical activity, coordinate the propagation of excitation in neuronal networks, and forge out stimulus-selective neuronal ensembles (Khan et al., 2018).

1.2.3 Plasticity

The neocortex rewires itself in the face of the ever-changing world, an ability known as neuronal plasticity (Buonomano and Merzenich, 1998). In an experiment, neuro-

scientists rewired the brain of a ferret such that the auditory region receives signals from the retina (von Melchner et al., 2000). This region was able to restructure itself to process the incoming information. Consequently, the ferret was able to perceive the visual data using the auditory cortex. Neuronal plasticity manifests itself also in almost everything we do. Our brains effortlessly retain information as we encounter new experiences, learn new concepts, or memorize a passage in a book.

The main mechanism behind plasticity is the brain’s ability to change the synaptic connections between the neurons (Martin et al., 2000). There are two forms of this synaptic change either structural or functional, realized via the creation and removal of synapses (Lamprecht and LeDoux, 2004) or the strengthening and depression of the synaptic strength (Citri and Malenka, 2007). Both structural and functional plasticity can be momentarily or long-lasting (Zucker and Regehr, 2002; Bear and Malenka, 1994). Both processes rely on morphological changes of the synapses (Stettler et al., 2006) or the density regulation of glutamate receptors such as NMDA and AMPA receptors (Anggono and Hugarir, 2012; Yashiro and Philpot, 2008). The synaptic plasticity is influenced by the presynaptic activity, the postsynaptic activity, neuromodulatory signals or combinations thereof (Magee and Grienberger, 2020). The simplest form of synaptic plasticity is known as Hebbian learning (Hebb, 1967). It postulates that when two neurons fire close to each other in time their synaptic strength increases. This form of learning is often summarized as “cells that fire together wire together” (Shatz, 1992). There is a plethora of synaptic learning rules and their role can never be underestimated. In this thesis, we employ two learning rules known as spike-timing-dependent plasticity and homeostatic control (Abbott and Nelson, 2000).

SPIKE-TIMING-DEPENDENT PLASTICITY. Since the first mention of Hebbian plasticity in 1949, a series of experimental works have been conducted to explore whether a similar mechanism exists in the neuronal system. The experiment by Markram et al. (1997) and Bi and Poo (1998) demonstrated a plasticity rule aligning with Hebb’s postulates called spike-timing-dependent plasticity (STDP). STDP posits that the timing between the generated spikes of the presynaptic and postsynaptic neurons influences the amount of synaptic change. It shows that if the postsynaptic spike follows after the presynaptic spike, the synaptic strength is increased, and decreased otherwise. Various types of STDP have been identified (Feldman, 2012), where the influence of the pre- and postsynaptic firing activity on the synaptic change can take different forms (Fig. 1.4). A structural form of STDP is also known from the literature (Liao et al., 1995; Wu et al., 1996).

HOMEOSTATIC PLASTICITY. Hebbian plasticity alone can either cause the strengthened synapses to become even more strengthened or depressed synapses to become more depressed (Abbott and Nelson, 2000). This positive feedback can result in a detrimental increase or decrease in the postsynaptic firing rates. To stabilize the learning and maintain a healthy working regime, the biological system employs additional processes known as homeostatic plasticity, which regulate the synaptic strength or neuronal excitability based on the postsynaptic firing rate or the total level of synaptic efficacy (Abbott and Nelson, 2000). Another benefit of homeostatic plasticity is that its combination with Hebbian plasticity can introduce competition between synapses (Tetzlaff et al., 2011). This combination causes the neurons to

develop different receptive fields, as neurons learn to become responsive to specific features of the input and unresponsive to others.

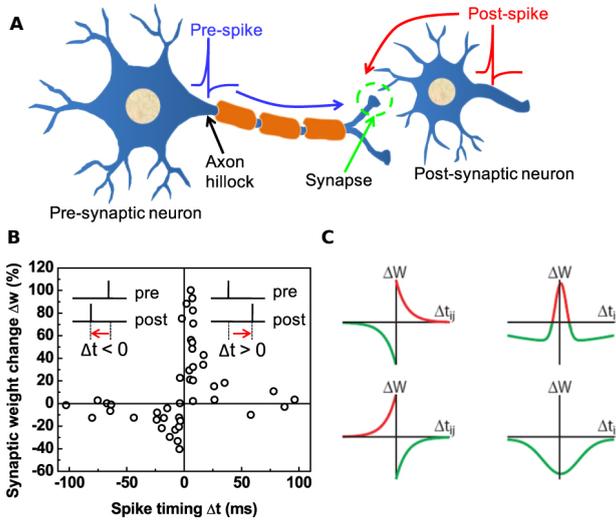


Figure 1.4: **Spike-timing-dependent plasticity (STDP) and its different forms.** **A**) Sketch of pre- and postsynaptic neurons with their corresponding spikes. **B**) The relative timing of pre- and postsynaptic spikes dictates the amount of the weight change. The data are adapted from (Bi and Poo, 1998). **C**) Different forms of STDP by Shouval, 2010. Reprinted with permission from (Kuzum et al., 2013).

1.2.4 Cortical activity

Neurons in cortical networks communicate with each other using sequences of spikes. This spiking activity is highly irregular (Shadlen and Newsome, 1998). In repeated experiments, the same stimulus can cause the same set of neurons to generate a varying number of spikes. Spiking activity in the cortex is not only irregular, but also asynchronous, i.e., the average level of synchrony for randomly chosen pairs of neurons is low (Ecker et al., 2010; Renart et al., 2010). Note, however, that as a response to behaviorally relevant events, neurons may synchronize their firing with millisecond precision (Riehle et al., 1997). Cortical activity is also characterized by a high degree of sparsity in space and time (Barth and Poulet, 2012; Shoham et al., 2002). The average firing rate, i.e., spike counts per unit time, of cortical neurons in different areas can range between < 1 and 30 spikes/s. A rate of 30 spikes/s is however rare. The bulk of the distribution is at smaller rates of 1 spikes/s or below (Swadlow, 1988; de Kock and Sakmann, 2009). This highly sparse code is thought to enhance the capacity of the biological network and the robustness against noise and failure (Olshausen and Field, 2004) as well as to optimize energy-usage (Attwell and Laughlin, 2001; Laughlin and Sejnowski, 2003).

The cooperative firing of neurons results in a rich repertoire of spatiotemporal patterns such as oscillations. They can be detected with non-invasive methods, like electroencephalography (EEG). Invasive methods recording electrocorticography (ECoG)

using electrodes at the surface of the cortical tissue or local field potentials (LFPs) using implanted electrodes can detect signals with a higher spatial resolution (Fig. 1.5, see also [Buzsáki et al., 2004](#)). These measurements reveal oscillatory activity over a broad range of frequencies ([Buzsáki and Draguhn, 2004](#)) including the theta (4–10 Hz), the beta (10–30 Hz), and the gamma bands (30–80 Hz). These oscillations have been linked to different behavioral and cognitive states ([Buzsáki, 2006](#)). For example, theta and gamma oscillations are involved in memory formation and retrieval ([Ward, 2003](#)), whereas beta oscillations are observed during movement initiation ([Denker et al., 2018](#)).

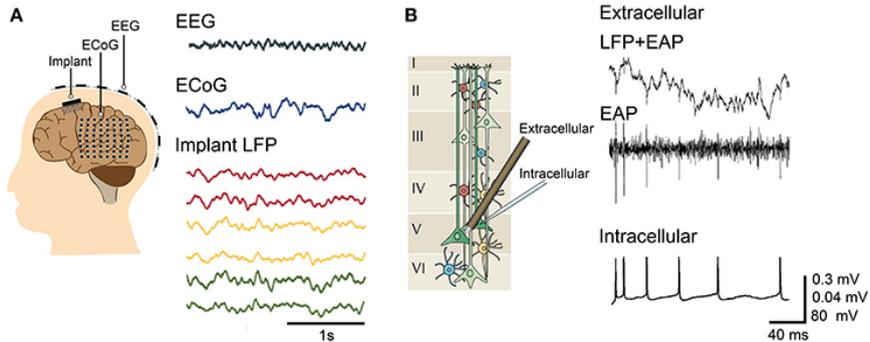


Figure 1.5: **Neural activity.** **A)** Electroencephalography (EEG; black), electrocorticography (ECoG; blue), local field potential (LFP) recordings. **B)** A combination of LFP and extracellular action potentials (EAPs), EAP, and intracellular recordings. For details on the experimental procedure see ([Obien et al., 2015](#)). Reprinted with permission from ([Obien et al., 2015](#)).

1.3 THE HIERARCHICAL TEMPORAL MEMORY MODEL

The remarkable uniformity of the neocortex is a well-established concept backed by a wealth of literature. As we have seen, the neocortical tissue is a collection of repetitive columns. [Mountcastle \(1978\)](#) suggested that each one of these columns might be implementing the same algorithm. He argued that understanding this algorithm would pave the way to discerning how the neocortex works. Inspired by these observations, [Hawkins and Blakeslee \(2007\)](#) proposed that learning and recalling sequences of patterns is the fundamental computation performed by the neocortex and laid down its cortical implementation. Later, Hawkins and his colleagues formulated the ideas of the book in a model called the Hierarchical Temporal Memory (HTM) model ([Hawkins et al., 2011](#)). In the following, we give a brief overview of the HTM model.

1.3.1 HTM regions and columns

Similar to the neocortex, an HTM network consists of regions organized in a hierarchy (Fig. 1.6). Each one of the regions comprises a sheet of interconnected cells arranged vertically in minicolumns. An HTM region implements a sequence memory and prediction algorithm. By combining past and current input, the region forms a prediction of what is likely to come next. Different regions receive different types of inputs and thus make specific predictions about the type of input they receive,

for example, visual regions make predictions about shapes, colors, etc, and auditory regions make predictions about tones, sounds, etc [Hawkins and Blakeslee \(2007\)](#). The hierarchical organization of the regions further promotes the formation of compositional sequences, which is critical for making predictions from sequence to sequence. To illustrate this, I borrow an example from ([Hawkins and Blakeslee, 2007](#)): vocal information reaching the auditory cortex is first decomposed to phonemes by the lower regions of the hierarchy and passed to the next level. The next higher region groups the sequences of phonemes to create words. Words are grouped into phrases by the next higher region and form our perception of the vocal information. The information can also flow downward from higher level regions to lower ones. This flow of information unfolds the phrases into sequences of phonemes, which can be passed to the motor cortex commanding the muscles to make sounds. The hierarchical organization of the processing makes a rich repertoire of predictions from either low-level sensory details to complex features or vice versa. For example, the top-down flow of information influences predictions in lower regions. As the information in higher regions is invariant with respect to environmental and external conditions, it can elucidate a more accurate expectation of what is likely to happen next. In an HTM network, the information does not move only within a particular sensory region, but can also travel to other sensory regions. For example, auditory information can make a prediction in the visual domain or other domains ([Fig. 1.6](#)).

The HTM model receives continuous inputs from the surrounding environment and learns in an unsupervised manner ([Hawkins et al., 2011](#)). It exploits the structure in the environment by grouping patterns that belong together. These patterns could for example reflect inputs from the same object or events unfolding in close proximity ([Hole, 2016](#)). Throughout the learning process, the model creates new representations and forgets old ones. The representations are invariant with respect to environmental conditions. A face is recognized as a face irrespective of external factors such as the lighting or the position. This way, the model is able to form predictions in novel situations. We discuss the learning algorithm in more detail in section [1.3.3](#).

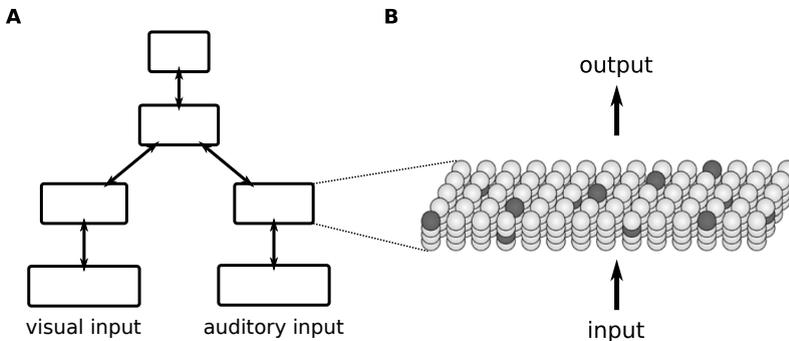


Figure 1.6: **HTM regions and columns.** **A)** HTM regions specialized for different sensory modalities (visual and auditory processing) are organized in a hierarchy. **B)** Each one of the HTM regions is composed of a set of minicolumns with active and silent cells shown in dark and light gray, respectively. Adapted figure from ([Hawkins et al., 2011](#)).

1.3.2 *HTM neurons and synapses*

In the HTM model (Hawkins and Ahmad, 2016), excitatory (pyramidal) neurons are described as abstract three-state systems that can assume an active, a predictive, or a non-active state (see Fig. 1.7). State updates are performed in discrete time. The current state is fully determined by the external input in the current time step and the network state in the previous step. Each HTM neuron is equipped with a number of dendrites (segments), modeled as coincidence detectors. The dendrites are grouped into distal, proximal, and apical dendrites. Distal dendrites receive inputs from other neurons in the local network, whereas proximal dendrites are activated by external sources. Inputs to proximal dendrites have a large effect on the soma and can easily trigger the generation of action potentials. Individual synaptic inputs to a distal dendrite, in contrast, have no direct effect on the soma. If the total synaptic input to a distal dendritic branch at a given time step is sufficiently large, the neuron becomes predictive. These dynamics mimic the generation of dendritic action potentials (dAPs), NMDA spikes (Antic et al., 2010; Schiller et al., 2000; Larkum et al., 2009), which result in a strong, long-lasting depolarization ($\sim 50\text{--}500$ ms) of the somata of neocortical pyramidal neurons. Apical dendrites are similar to distal dendrites but they receive inputs from higher-level regions and encode top-down expectations.

Synapses on any one of the dendrites have a binary weight and their maturity is characterized by an additional parameter called permanence. Throughout the learning process, the permanence value is incremented and decremented using a Hebbian-like rule. If the permanence is above a certain threshold the synapse is considered connected (or mature) and otherwise not connected (or immature). For a dendritic segment to be active, it requires a set of co-located mature synapses (typically 15–20) simultaneously receiving inputs from a subset of active presynaptic cells.

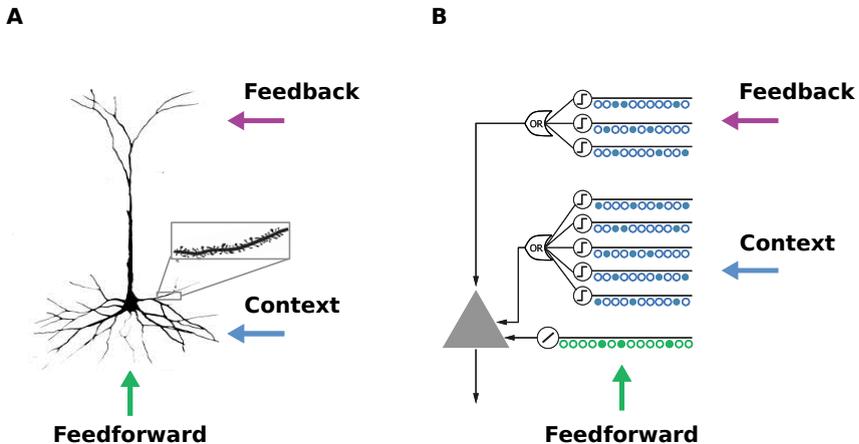


Figure 1.7: **HTM neuron.** **A**) A neocortical pyramidal neuron has three sources of inputs: feedforward (proximal dendrite), context (basal dendrite), and feedback (apical dendrite). Each dendrite is composed of thousands of excitatory synapses (inset). **B**) An HTM neuron with an array of coincident detectors each mimicking the different pyramidal dendrites. Adapted figure from (Hawkins and Ahmad, 2016).

1.3.3 *HTM cortical learning algorithm*

In a first step, an HTM network learns how to map inputs in the environment to the different minicolumns (see the Spatial Pooler (SP) algorithm in [Hawkins et al., 2011](#)). An external input is connected to proximal dendrites and causes a sparse set of minicolumns to become active. This set is selected by means of a winner take all mechanism, i.e., the minicolumns with the strongest input strength activate first and inhibit the minicolumns with smaller input strength ([Hole, 2016](#)). All cells in the corresponding minicolumns become active as they share the same receptive field. The algorithm strengthens the permanence between the activated minicolumns and the corresponding input pattern and decrements the permanence otherwise. This way similar input patterns get connected to the same set of minicolumns fostering an invariant representation.

In a second step, the algorithm learns the temporal order of the input patterns and associates the patterns that occur close to each other in time (see the Temporal Memory (TM) algorithm in [Hawkins et al., 2011](#)). The lateral connectivity between neurons belonging to the different minicolumns is subject to a form of Hebbian structural plasticity. Repetitive and consistent sequential presentation of sequence elements turns immature connections between successively activated minicolumns into mature connections and hence leads to the formation of sequence-specific subnetworks.

During the learning process, the number of mature connections grows to a point where the activation of a certain minicolumn by an external input activates the distal dendrites in a subset of cells in the subsequent minicolumn. The cells with active distal dendrites are referred to as predictive cells. When receiving an external input, only predictive cells become active. If no predictive cells exist within the minicolumn, then all the cells become active. The subset of activated cells is specific to the context, i.e., activity depends on the entire (or parts of the) stimulus history, and predicts by activating again a specific subset of distal dendrite in the next minicolumns. In the presence of a non-anticipated, non-predicted stimulus, the cells in the corresponding minicolumn fire collectively in a non-sparse manner, thereby signaling a “mismatch”.

HTM cortical algorithm also provides an explanation of how feedback connections influence the prediction ([Hawkins et al., 2011](#)). It proposes that inputs coming from higher-level cortical areas are integrated by the apical dendrite. These inputs implement a top-down prediction and constrain the network to adopt a particular interpretation ([Hawkins and Ahmad, 2016](#)). This is in line with a number of studies suggesting that feedback connections implement a form of expectation ([Lamme et al., 1998](#)).

1.3.4 *Functional role of the cortical layers*

[Hawkins et al. \(2011\)](#) hypothesize that a variation of the same algorithm is implemented by the different cortical layers. Here, we give a brief summary of these hypotheses. The algorithm implemented by layers 2 and 3 is the closest to the HTM cortical learning algorithm discussed in section 1.3.3. Similar to an HTM network, the receptive fields of neurons in layer 3 of the primary visual cortex (V1) are known to be context-specific and can differentiate by using the temporal context. For example, they exhibit a receptive field that is direction-selective, i.e., subsets of neurons become active when a line is moving toward the left, and other subsets become active as a response to a line moving right ([Hubel and Wiesel, 1965](#)). Other layers are thought to run a variant

of the same cortical learning algorithm. Layer 4 learns first-order sequences, i.e., does not account for the context of prior inputs. It learns to associate similar inputs to the same set of minicolumns. This form of sequence learning is critical for the formation of invariant representations, for example against spatial transformations. Neurons exhibiting such behavior are known in the neuroscience literature as complex cells (Hubel and Wiesel, 1968). Layer 5 possesses similar learning capabilities as layer 3, but in addition, it can predict not only “what” will happen next but also “when” it happens. This is in line with experimental data showing that layer 5 controls movement where the timing aspect is very critical. Learning the time is also crucial for other tasks such as replaying music or recognizing spoken words. In contrast to other layers, layer 6 is the origin of axons that feedback to lower regions.

1.4 NEUROMORPHIC HARDWARE

The growing knowledge of neuroscientific concepts and principles has led to a surge of brain-inspired algorithms able to solve complex problems with minimal energetic usage (Roy et al., 2019). For example, the Hierarchical Temporal Memory (HTM) model learns complex sequences and is able to detect anomalies in the data. The learning in HTM is continuous, unsupervised, and less prone to errors as compared to conventional computing paradigms (Cui et al., 2016).

This type of algorithm often relies on complex dynamics and is built using a large number of units with dense connectivity. Implementing these algorithms on traditional computers requires a substantial amount of energy and long simulation times (Mehonic and Kenyon, 2022). These issues are also hindering research in neuroscience (Kunkel et al., 2014). Nowadays, the modeling and the simulation of brain dynamics and functions are carried out on traditional computers. Simulations of large-scale networks with plastic, natural density connectivity run on even advanced supercomputers are substantially slower than biology. Studying biological processes such as learning and brain development that unfold on long timescales of hours, days, or more by means of computer simulations, therefore, remains challenging.

One of the main obstacles limiting the efficiency of traditional computers is the physical separation of memory and data processing. Conventional computers rely on the von-Neumann architecture, where the central processing unit (CPU) and the memory are physically separated (von Neumann, 1993). The CPU fetches the data from memory, executes the corresponding operations, and stores the results back in the memory. The transfer of the data between the CPU and the memory creates the so-called memory bottleneck (Mahapatra and Venkatrao, 1999). In the neuronal system, the memory and the computation are co-located on the same substrate allowing faster processing of the data and thus enabling superior performance compared to conventional computers.

An emerging type of hardware that holds promise for efficiently running brain-inspired algorithms is known as neuromorphic hardware. It mimics the way the brain processes information and borrows from its mechanisms and computational principles to realize highly efficient computers (Mead, 1990). This type of hardware allows for compact and low-power intelligent systems that can be integrated into smartphones or other devices such as brain-machine interfaces and it can aid neuroscience research by accelerating the simulation of large-scale networks. Current neuromorphic hardware is either designed in the digital domain such as SpiNNaker (Furber et al., 2014), TrueNorth (Merolla et al., 2014), and Loihi (Davies et al., 2018), or in the mixed

analog/digital domain such as BrainScaleS (Schemmel et al., 2010), Neurogrid (Benjamin et al., 2014), and DYNAPs (Moradi et al., 2018). See the reviews by Schuman et al. (2017); Shrestha et al. (2022). Most of these architectures rely on CMOS circuits. While these types of circuits are mature and robust, they are not particularly suited for neuromorphic computing. The hardware components implementing the neurons and synapses are often realized using complex circuitry and do not implement learning rules by themselves. To give an example, in contrast to a biological synapse, which occupies a surface of $0.001\mu\text{m}^2$, a CMOS-based synapse occupies 400 times larger surface (Xia and Yang, 2019).

To overcome these issues, memristive devices were proposed as bioinspired elements that can directly mimic neurons and synaptic functionalities (Yang et al., 2013; Ielmini and Wong, 2018; Wang et al., 2020). By integrating these devices into large and dense networks, neuromorphic hardware offers a possible solution for the emulation of the connectivity and processing capabilities of biological neural networks. In contrast to complex CMOS-based synaptic circuits, memristive devices offer a simpler hardware solution and can be substantially faster than their biological counterpart. Its programmable non-volatile conduction shows long term potentiation (LTP) and long term depression (LTD). Furthermore, this type of device is inherently noisy and heterogeneous. Developing neural network models of biological computation in the face of noise and heterogeneity enables us to better understand their role in nature.

1.4.1 Memristive devices

The term memristive device was introduced by Chua and Kang (1976) as a theoretical electrical component that exhibits resistive and memory effects. Later, Strukov et al. (2008) demonstrated that TiO₂-oxide based redox-based random access memory (ReRAM, Waser et al., 2009) devices show switching characteristics similar to the theoretically predicted memristor (Chua, 1971). Although the physical mechanism and the detailed switching mechanisms differ substantially (Valov and Kozicki, 2013), the term memristive device nowadays is used for a whole family of resisting switching devices, including ReRAM, phase change memory (PCM, Raoux et al., 2009), magnetoresistive random access memory (MRAM, Hosomi et al., 2005), and floating gate transistors (Hasler et al., 1994). The reviews by Wang et al. (2020); Zhu et al. (2020) give a summary of the different mechanisms exploited in these devices. ReRAM devices stand out at many aspects, they are highly scalable, energy-efficient, fast, and can be used to mimic learning mechanisms in the brain. In section 1.4.3, we give an overview of the different types of ReRAM devices and their working principles.

Memristive devices are, typically, composed of a metal-dielectric-metal stack (Wang et al., 2020). Similar to biological synapses, their resistance is plastic, non-volatile, and exhibits a response to the stimulus that is history-dependent. They can be programmed into different states (e.g., resistance), where the data is stored and represented. Due to their small size and scalability, these devices are highly suitable for integration into densely packed 2D or 3D crossbar arrays, which has resulted in their adoption in various neuromorphic applications (Pedretti et al., 2017; Wang et al., 2018).

Memristive devices are known, however, by their nonideal behavior including variability and nonlinearity in the update curves as well as limited resistance range and precision (Zhao et al., 2020; Fouda et al., 2020). These nonidealities can be detrimental if incorporated into networks (Fouda et al., 2020), such as Artificial Neural Networks (ANNs). A number of training algorithms have been developed to

mitigate these issues (Chen et al., 2015a; Gokmen and Haensch, 2020). However, since these imperfections are characteristics of biological synapses, current research is investigating how the biological system carries out reliable computation despite this latter. A prominent idea is the use of variability for probabilistic sampling (Nefci et al., 2016).

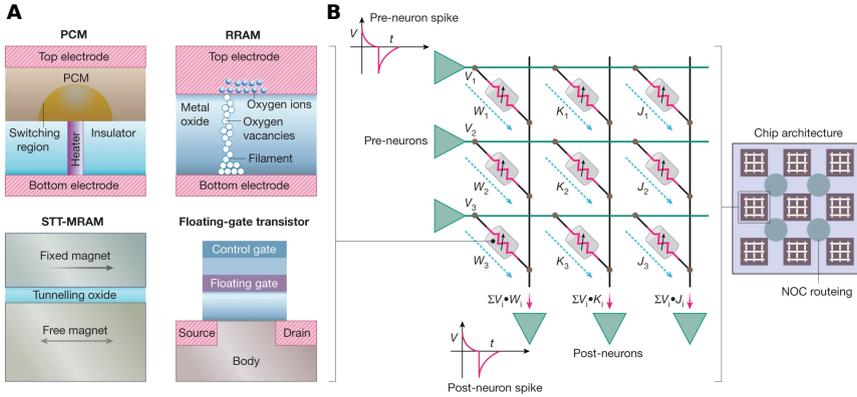


Figure 1.8: **Memristive technologies.** **A)** Sketches of different memristive devices: PCM, RRAM, S1T-MRAM, and floating-gate transistor. **B)** Crossbar array of memristive devices. Figures in panels A and B are reprinted with permission from (Roy et al., 2019).

1.4.2 Crossbar arrays

The neocortex is known by its massive and dense connectivity. Each neuron has on average 10,000 synapses (Abeles, 1991). The field of neuromorphic hardware proposes that such connectivity may be realized using a dense network of memristive devices, i.e., an architecture known as crossbar arrays (see Fig. 1.8B). Due to the energy efficiency of memristive devices, this approach has also the potential to close the gap between the energy consumption of computers and the brain (Burr et al., 2016; Xia and Yang, 2019).

In the biological network, the information travels from one neuron to the next via synapses. The crossbar array mimics this operation by a multiplication between the presynaptic input delivered as a voltage pulse and the conductance of the memristive device, which acts as a synapse. More specifically, crossbar arrays consist of multiple intersecting horizontal and vertical lines. The current flowing through the memristive device in each cross point, is determined by Ohm’s law expressed as a multiplication between the input voltage and the memristive conductance. By means of Kirchhoff’s law, the current induced at each vertical line is summed across all currents flowing through the memristive elements (Ielmini and Wong, 2018; Xia and Yang, 2019). Unlike von Neumann type of computing, a crossbar array can perform this computation in one pass without a transfer of any memory making them fast and energy efficient. Moreover, the small size of memristive elements allows for highly scalable hardware. In recent crossbar architectures, a CMOS transistor is connected in series with the memristive element (1T1R architecture). The additional transistor permits precise

operation and control of the memristive devices and helps suppress the sneak path effect, i.e., the current flowing through unselected cells (Xia and Yang, 2019).

The conception and development of crossbar arrays are still facing a number of challenges. Fabrication of large scale arrays is currently not possible due to device variations and technological limitations (Xia and Yang, 2019; Li et al., 2021). Moreover, the crossbar arrays rely on digital, peripheral circuits for updating and controlling the memristive synapses. This additional circuitry increases substantially the energy consumption and slows down the processing time. Future works aim at upscaling these arrays, while efficiently integrating the peripheral circuitry.

1.4.3 Redox-based random access memory devices

Redox-based random access memory (ReRAM) devices are composed of a resistive metal oxide “T” sandwiched between two metal conductors “M” functioning as electrodes (Waser, 2012b). After an initial electroforming cycle, the device can switch between two or more resistance states, either from high to low, i.e., SET operation, or from low to high, i.e., RESET operation. The switching procedure of the device can be either unipolar or bipolar. In bipolar type devices, the SET and RESET operations require different voltage polarities. In unipolar devices, the same voltage polarity can be used for both the SET and RESET operations. Typically, ReRAM plasticity is recorded as the change in resistance as a function of the amplitude of an applied voltage pulse. For example, the ReRAM device can be used to emulate STDP, by appropriate shaping of the pre- and postsynaptic action potentials, i.e., the applied pulse forms (Zamarreño-Ramos et al., 2011; Wang et al., 2015). Alternatively, the device resistance can be recorded based on the number of identical voltage pulses.

The mechanisms of this resistance modification are governed by electrochemical processes and the movement of metal ions or oxygen vacancies (Waser, 2012a; Abbaspour et al., 2018). The type of resistive switching mechanism in ReRAMs is influenced by the material composition of each layer in the stack and can be driven by electrochemical metallization memory effect (ECM), valence change memory effect (VCM), or thermochemical memory effect (TCM). In the following, we give an overview of the mechanism involved in these devices, for details see (Waser, 2012b; Valov, 2013). In ECM cells, the SET process causes the metal ions to move from an electrochemically active electrode (such as Ag) to the counter electrode. Upon arrival, the ions reduce into atoms forming a filament extending toward the active electrode. To RESET the device, a negative voltage is applied, which breaks the filament (Valov and Kozicki, 2013; Waser and Aono, 2007). On the other hand, in VCM ReRAM, the filaments are formed by extracting oxygen vacancies from the I layer. The conductivity of the device can be modulated by changing the electronic properties of the interface between the filament and the active electrode. The polarity of the SET or RESET pulses increases or decreases the concentration of oxygen vacancies, and thus, an increase or decrease of the conductivity (Waser, 2012b). In contrast to VCM and ECM cells, the switching of TCM cells is unipolar, i.e., the SET and RESET operations use the same voltage polarity. During the SET, the flow of current in the cell causes the cell to heat up. The current is further increased by the heat, which leads to a positive feedback loop between these two quantities leading to a thermoelectric breakdown event (Waser, 2012a). This results in local redox reactions and the formation of a conducting filament turning the cell into the ON-state. During RESET, the conducting filament is dissolved due to a substantial increase in the temperature caused by a high

current flow resulting from an application of a voltage (Telmini et al., 2011; Waser, 2012a).

In this thesis, we focus on VCM ReRAMs as synaptic elements in neuromorphic hardware and give more details on their switching mechanisms in the next section.

1.4.4 Valance change memory devices

The VCM ReRAM device consists of a thin (3–10 nm) insulating metal-oxide film sandwiched between a low workfunction, ohmic electrode (e.g., Zr), and a high workfunction, blocking Schottky-interface electrode (e.g., Pt). An initial electroforming step results in the formation of conducting filaments composed of oxygen vacancies. This step requires a high voltage, which extracts the oxygen atoms from the oxide layer creating oxygen vacancies. The filaments, referred in the literature as the *plug*, extend from the ohmic electrode toward the blocking interface, where they leave a gap called the *disc* (Fig. 1.9). The gap is initially fully oxidized creating an energy barrier, which gives rise to the OFF state (Waser, 2012a). During further operation, the movement of the oxygen vacancies in the disc area changes the device’s electrical properties. Applying a SET operation increases the concentration of oxygen vacancies at the blocking interface, which lowers the energetic barrier and increases the conductivity. On the other hand, the RESET process repels the oxygen vacancies from the interface and decreases the conductivity again. Recording the voltage at the active interface during the SET and RESET processes results in a switching polarity known as the counter-eight-wise shape (see Fig. 1.9).

In another type of VCM device, the switching occurs homogeneously over the whole electrode/oxide interface (Sawa, 2008). The oxide is initially conductive and forms an energetic barrier at the interface. The applied voltage causes the oxygen ions to drift toward or away from the interface. This results in changing the height of the energetic barrier and thus modulation of the resistance. This type of switching is observed in perovskite-type devices such as PCMO (Li et al., 2009).

1.5 SCOPE AND STRUCTURE OF THIS THESIS

The primary focus of this thesis is to uncover the potential mechanisms underlying sequence prediction and recall in the neocortex and develop a mapping of these mechanisms on neuromorphic hardware. Throughout the thesis, we answer the following questions.

Is it possible to implement the Temporal Memory (TM) component of the HTM model using biological ingredients in continuous time dynamics?

The TM model proposed by Hawkins and Ahmad (2016) constitutes a sequence processing algorithm inspired by the cortical circuitry. It learns high-order sequences in an unsupervised, continuous manner using local learning rules. Its capacity benefits from sparsity in the activity, and therefore provides a highly energy-efficient sequence learning and prediction mechanism. In its original implementation, however, the TM model is based on abstract neuron and synapse models updated in discrete time. Furthermore, some of its learning dynamics can hardly be reconciled with biology. A verification of this model by experimental data is therefore limited. **Chapter 2** seeks to fill this gap by providing a biologically inspired, continuous-time implementation of the TM model. We further assess whether we can replace its learning rules with more

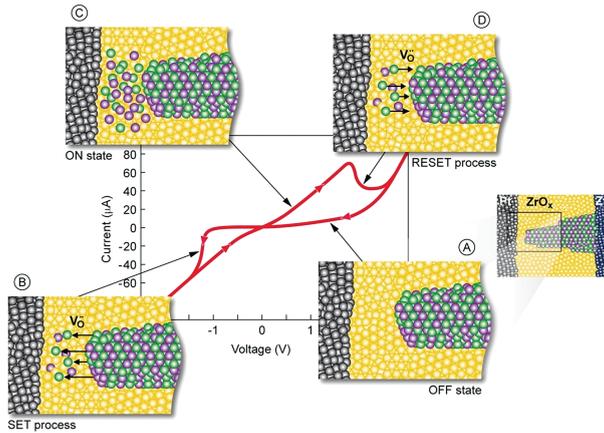


Figure 1.9: **The SET and RESET processes of VCM device.** The insets A to D depict the switching behavior of a Pt/ZrO_x/Zr stack, where Pt depicts the active electrode (AE) and Zr the ohmic electrode (OE). The oxygen vacancies and Zr ions are shown using green and purple spheres, respectively. Reprinted with permission from (Waser, 2012a).

biologically plausible plasticity mechanisms. By mapping the TM model ingredients and observables to those known from biology, we seek to foster its verifiability based on electrophysiological and behavioral data. This investigation identifies biological mechanisms and principles crucial for sequence processing. In the following, we refer to our reformulated model as the spiking Temporal Memory (spiking TM) model.

Can the spiking TM model replay learned sequences in response to an ambiguous cue using different strategies? Animals rely on different decision strategies when faced with ambiguous or uncertain cues. Depending on the context, decisions may be biased towards events that were most frequently experienced in the past, or be more explorative. A particular type of decision making central to cognition is sequential memory replay in response to ambiguous cues. In response to an ambiguous cue, the spiking TM model deterministically replays the sequence shown most frequently during training. **Chapter 3** investigates how different replay strategies can be achieved by adjusting the noise characteristics. We study two types of noise: random synaptic bombardment and random stimulus locking to spatiotemporal oscillations.

Is it possible to use memristive devices as a replacement of the biological synapses in the spiking TM model? Memristive devices have been identified as potential synaptic elements in neuromorphic hardware. They permit scalability, are energy efficient and fast, and can be used to implement biological plasticity rules. As a step towards a memristive hardware implementation of the spiking TM, **chapter 4** assesses the performance of the model by adapting its plasticity curves to those observed in memristive devices. The investigation includes two types of memristive devices: (i) an analog switching memristive device, where the conductance gradually changes between a low conductance state (LCS) and a high conductance state (HCS), and (ii) a binary switch memristive device, where the conductance abruptly changes

between the LCS and the HCS. We study the performance characteristics of the sequence learning model as a function of different device properties.

What algorithmic adaptations are needed for a full electronic circuit implementation of the spiking TM model, where the synaptic connectivity is realized by a memristive crossbar array? For a full electronic circuit implementation of the model, it is not sufficient to show only that synapses can be replaced by memristive devices (chapter 4). Memristive devices can emulate only certain aspects of the synapse dynamics. A full implementation of the plasticity dynamics, for example, requires additional hardware. A priori it is not clear whether the original model can be mapped to such hardware in a one-to-one manner. Certain hardware constraints may require certain adaptations. We address these points in **chapter 5**.

Chapter 6 provides short summarizing answers to each of the questions above and discusses future work. **Chapter 7** includes further experiments assessing the spiking TM model performance and provides model parameters.

SEQUENCE LEARNING IN A NETWORK OF SPIKING
NEURONS

This chapter is based on the following publication:

Bouhadjar, Y., Diesmann, M., Wouters, D. J., & Tetzlaff, T. (2022). Sequence learning, prediction, and replay in networks of spiking neurons. *PLOS Computational Biology* 18 (6), e1010233.

Author contributions: All authors contributed to the conception of the work. Younes Bouhadjar implemented the model, performed the simulations, and analyzed and visualized the data. All authors jointly wrote the paper, reviewed the manuscript, and approved it for publication. Younes Bouhadjar was supervised by Tom Tetzlaff and Dirk J Wouters.

2.1 INTRODUCTION

The majority of existing biologically motivated models of sequence learning addresses sequence replay (Maes et al., 2020; Klos et al., 2018; Cone and Shouval, 2021; Klampff and Maass, 2013). Sequence prediction and mismatch detection are rarely discussed. The Hierarchical Temporal Memory (HTM) (Hawkins et al., 2011) combines all three aspects: sequence prediction, mismatch detection and replay. Its Temporal Memory (TM) model (Hawkins and Ahmad, 2016) learns complex context-dependent sequences in a continuous and unsupervised manner using local learning rules (Cui et al., 2016), and is robust against noise and failure in system components. Furthermore, it explains the functional role of dendritic action potentials (dAPs) and proposes a mechanism of how mismatch signals can be generated in cortical circuits (Hawkins and Ahmad, 2016). Its capacity benefits from sparsity in the activity, and therefore provides a highly energy efficient sequence learning and prediction mechanism (Ahmad and Hawkins, 2015).

The original formulation of the TM model is based on abstract models of neurons and synapses with discrete-time dynamics. Moreover, the way the network forms synapses during learning is difficult to reconcile with biology. Here, we propose a continuous-time implementation of the TM model derived from known biological principles such as spiking neurons, dAPs, lateral inhibition, spike-timing-dependent structural plasticity, and homeostatic control of synapse growth. This model successfully learns, predicts and replays high-order sequences, where the prediction of the upcoming element is not only dependent on the current element, but also on the history. Bringing the model closer to biology allows for testing its hypotheses based on experimentally accessible quantities such as synaptic connectivity, synaptic currents, transmembrane potentials, or spike trains. Reformulating the model in terms of continuous-time dynamics moreover enables us to address timing-related questions, such as the role of the sequence speed for the prediction performance and the replay speed.

The study is organized as follows: the [Methods](#) describe the task, the network model, and the performance measures. The [Results](#) illustrate how the interaction of the model’s components gives rise to context dependent predictions and sequence replay, and evaluate the sequence processing speed and prediction performance. The [Discussion](#) finally compares the spiking TM model to other biologically motivated sequence learning models, summarizes limitations, and provides suggestions for future model extensions.

2.2 METHODS

In the following, we provide an overview of the task and the training protocol, the network model, and the task performance analysis. A detailed description of the model and parameter values can be found in [Tables 2.1](#) and [7.1](#).

2.2.1 Task and training protocol

In this study, we develop a neuronal architecture that can learn and process an ensemble of S sequences $s_i = \{\zeta_{i,1}, \zeta_{i,2}, \dots, \zeta_{i,C_i}\}$ of ordered discrete items $\zeta_{i,j}$ with $C_i \in \mathbb{N}^+$, $i \in [1, \dots, S]$. The length of sequence s_i is denoted by C_i . Throughout this study, the sequence elements $\zeta_{i,j} \in \{A, B, C, \dots\}$ are represented by Latin characters, serving as placeholders for arbitrary discrete objects or percepts, such as images,

numbers, words, musical notes, or movement primitives (Fig. 2.1A). The order of the sequence elements within a given sequence represents the temporal order of item occurrence.

The tasks to be solved by the network consist of

- i) predicting subsequent sequence elements in response to the presentation of other elements,
- ii) detecting unanticipated stimuli and generating a mismatch signal if the prediction is not met, and
- iii) autonomously replaying sequences in response to a cue signal after learning.

The architecture learns sequences in a continuous manner: the network is exposed to repeated presentations of a given ensemble of sequences (e.g., {A,D,B,E} and {F,D,B,C} in Fig. 2.1B). In the *prediction mode* (task i) and ii)), there is no distinction between a “training” and a “testing” phase. At the beginning of the learning process, all presented sequence elements are unanticipated and do not lead to a prediction (diffuse shades in Fig. 2.1B, left). As a consequence, the network generates mismatch signals (flash symbols in Fig. 2.1B, left). After successful learning, the presentation of some sequence element leads to a prediction of the subsequent stimulus (colored arrows in Fig. 2.1B). In case this subsequent stimulus does not match the prediction, the network generates a mismatch signal (red arrow and flash symbol in Fig. 2.1B, right). The learning process is entirely unsupervised, i.e., the prediction performance does not affect the learning. As described in section 2.3.4, the network can be configured into a *replay mode* where the network autonomously replays learned sequences in response to a cue signal (task iii)).

In general, the sequences in this study are “high-order” sequences, similar to those generated by a high-order Markov chain; the prediction of an upcoming sequence element requires accounting for not just the previous element, but for (parts of) the entire sequence history, i.e., the context. Sequences within a given set of training data can be partially overlapping; they may share certain elements or subsequences (such as in {A,D,B,E} and {F,D,B,C}). Similarly, the same sequence element (but not the first one, see section 6.2.2) may occur multiple times within the same sequence (such as in {A,D,B,D}). Throughout this work, we use two sequence sets:

SEQUENCE SET I: For an illustration of the learning process and the network dynamics in the prediction (section 2.3.1) and in the replay mode (section 2.3.4), as well as for the investigation of the sequence processing speed (section 2.3.3), we start with a simple set of two partially overlapping sequences $s_1 = \{A,D,B,E\}$ and $s_2 = \{F,D,B,C\}$ (see Fig. 2.1B).

SEQUENCE SET II: For a more rigorous evaluation of the sequence prediction performance (section 2.3.2), we consider a set of $S = 6$ high-order sequences: $s_1 = \{E,N,D,I,J\}$, $s_2 = \{L,N,D,I,K\}$, $s_3 = \{G,J,M,C,N\}$, $s_4 = \{F,J,M,C,I\}$, $s_5 = \{B,C,K,H,I\}$, $s_6 = \{A,C,K,H,F\}$, each consisting of $C = 5$ elements. The complexity of this sequence ensemble is comparable to the one used in (Hawkins and Ahmad, 2016), but is more demanding in terms of the high-order context dependence.

Results for two additional sequence sets are summarized in the Supporting information. The set used in Fig. 7.2 is composed of sequences with recurring first elements.

In Fig. 2.10, we show results for longer sequences with a larger number of overlapping elements.

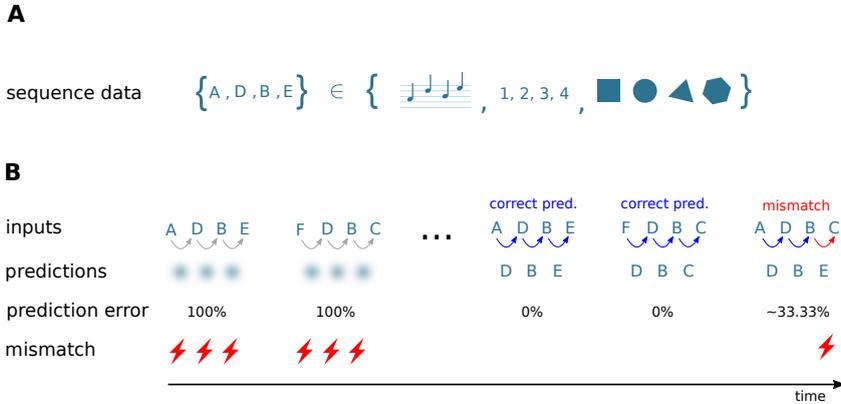


Figure 2.1: **Sketch of the task and the learning protocol.** **A)** The neuronal network model developed in this study learns and processes sequences of ordered discrete elements, here represented by characters “A”, “B”, “C”, Sequence elements may constitute arbitrary discrete items, such as musical notes, numbers, or images. The order of sequence elements represents the temporal order of item occurrence. **B)** After repeated, consistent presentation of sets of high-order sequences, i.e., sequences with overlapping characters (here, $\{A, D, B, E\}$ and $\{F, D, B, C\}$), the model learns to predict subsequent elements in response to the presentation of other elements (blue arrows) and to detect unanticipated elements by generating a mismatch signal if the prediction is not met (red arrows and flash symbols). The learning process is continuous and unsupervised. At the beginning of the learning process, all presented elements are unanticipated and hence trigger the generation of a mismatch signal. The learning progress is monitored and quantified by the prediction error (see section 2.2.3).

2.2.2 Network model

NETWORK STRUCTURE. The network consists of a population \mathcal{E} of N_E excitatory (“E”) and a population \mathcal{I} of N_I inhibitory (“I”) neurons. The neurons in \mathcal{E} are randomly and recurrently connected, such that each neuron in \mathcal{E} receives K_{EE} excitatory inputs from other randomly chosen neurons in \mathcal{E} . Note that these “EE” connections are potential connections in the sense that they can be either “mature” (“effective”) or “immature”. Immature connections have no effect on target neurons (see below). In the neocortex, the degree of potential connectivity depends on the distance between the neurons (Stepanyants et al., 2007). It can reach probabilities as high as 90% for neighboring neurons, and decays to 0% for neurons that are farther apart. In this work, the connection probability is chosen such that the connectivity is sufficiently dense, allowing for the formation of specific subnetworks, and sufficiently sparse for increasing the network capacity (see paragraph “Constraints on potential connectivity” below). The excitatory population \mathcal{E} is subdivided into M non-overlapping subpopulations $\mathcal{M}_1, \dots, \mathcal{M}_M$, each of them containing neurons with identical stimulus preference (“receptive field”; see below). Each subpopulation \mathcal{M}_k thereby represents a specific

element within a sequence (Figs 2.2A and 2.2B). In the original TM model (Hawkins and Ahmad, 2016), a single sequence element is represented by multiple (L) subpopulations (“minicolumns”). For simplicity, we identify the number M of subpopulations with the number of elements required for a specific set of sequences, such that each sequence element is encoded by just one subpopulation ($L = 1$). All neurons within a subpopulation \mathcal{M}_k are recurrently connected to a subpopulation-specific inhibitory neuron $k \in \mathcal{I}$. The inhibitory neurons in \mathcal{I} are mutually unconnected. The subdivision of excitatory neurons into stimulus-specific subpopulations defines how external inputs are fed to the network (see next paragraph), but does not affect the potential excitatory connectivity, which is homogeneous and not subpopulation specific.

EXTERNAL INPUTS. During the prediction mode, the network is driven by an ensemble $\mathcal{X} = \{x_1, \dots, x_M\}$ of M external inputs, representing inputs from other brain areas, such as thalamic sources or other cortical areas. Each of these external inputs x_k represents a specific sequence element (“A”, “B”, ...), and feeds all neurons in the subpopulation \mathcal{M}_k with the corresponding stimulus preference. The occurrence of a specific sequence element $\zeta_{i,j}$ at time $t_{i,j}$ is modeled by a single spike $x_k(t) = \delta(t - t_{i,j})$ generated by the corresponding external source x_k . Subsequent sequence elements $\zeta_{i,j}$ and $\zeta_{i,j+1}$ within a sequence s_i are presented with an inter-stimulus interval $\Delta T = t_{i,j+1} - t_{i,j}$. Subsequent sequences s_i and s_{i+1} are separated in time by an inter-sequence time interval $\Delta T_{\text{seq}} = t_{i+1,1} - t_{i,C_i}$. During the replay mode, we present only a cue signal encoding for first sequence elements $\zeta_{i,1}$ at times $t_{i,1}$. Subsequent cues are separated in time with an inter-cue time interval $\Delta T_{\text{cue}} = t_{i+1,1} - t_{i,1}$. In the absence of any other (inhibitory) inputs, each external input spike is strong enough to evoke an immediate response spike in all target neurons $i \in \mathcal{M}_k$. Sparse activation of the subpopulations in response to the external inputs is achieved by a winner-take-all mechanism implemented in the form of inhibitory feedback (see section 2.3.1).

NEURON AND SYNAPSE MODEL. In the original TM model (Hawkins and Ahmad, 2016), excitatory (pyramidal) neurons are described as abstract three-state systems that can assume an active, a predictive, or a non-active state. State updates are performed in discrete time. The current state is fully determined by the external input in the current time step and the network state in the previous step. Each TM neuron is equipped with a number of dendrites (segments), modeled as coincidence detectors. The dendrites are grouped into distal and proximal dendrites. Distal dendrites receive inputs from other neurons in the local network, whereas proximal dendrites are activated by external sources. Inputs to proximal dendrites have a large effect on the soma and trigger the generation of action potentials. Individual synaptic inputs to a distal dendrite, in contrast, have no direct effect on the soma. If the total synaptic input to a distal dendritic branch at a given time step is sufficiently large, the neuron becomes predictive. This dynamic mimics the generation of dendritic action potentials (dAPs), NMDA spikes (Antic et al., 2010; Schiller et al., 2000; Larkum et al., 2009), which result in a long-lasting depolarization (~ 50 - 500 ms) of the somata of neocortical pyramidal neurons.

In contrast to the original study (see section 1.3.2), the model proposed here employs neurons with continuous-time dynamics. For all types of neurons, the temporal evolution of the membrane potential is given by the leaky integrate-and-fire model Eq (2.10). The total synaptic input current of excitatory neurons is composed of currents in distal dendritic branches, inhibitory currents, and currents from external

sources. Inhibitory neurons receive only inputs from excitatory neurons in the same subpopulation. Individual spikes arriving at dendritic branches evoke alpha-shaped postsynaptic currents, see Eq (2.12). The dendritic current includes an additional nonlinearity describing the generation of dAPs: if the dendritic current I_{ED} exceeds a threshold θ_{dAP} , it is instantly set to a the dAP plateau current I_{dAP} , and clamped to this value for a period of duration τ_{dAP} , see Eq (2.16). This plateau current leads to a long lasting depolarization of the soma (see Fig. 2.3B). The dAP threshold θ_{dAP} is chosen such that the co-activation of γ neurons with mature connections to the target neuron reliably triggers a dAP. In this work, we use a single dendritic branch per neuron. However, the model could easily be extended to include multiple dendritic branches. External and inhibitory inputs to excitatory neurons as well as excitatory inputs to inhibitory neurons trigger exponential postsynaptic currents, see Eq (2.13–2.15). Similar to the original implementation, an external input strongly depolarizes the neurons and causes them to fire. To this end, the external weights J_{EX} are chosen to be supra-threshold (see Fig. 2.3A). Inhibitory interactions implement the WTA described in section 2.3.1. The weights J_{IE} of excitatory synapses on inhibitory neurons are chosen such that the collective firing of a subset of ρ excitatory neurons in the corresponding subpopulation causes the inhibitory neuron to fire. The weights J_{EI} of inhibitory synapses on excitatory neurons are strong such that each inhibitory spike prevents all excitatory neurons in the same subpopulation that have not generated a spike yet from firing. All synaptic time constants, delays and weights are connection-type specific (see Table 2.1).

PLASTICITY DYNAMICS. Both in the original (Hawkins and Ahmad, 2016) and in our model, the lateral excitatory connectivity between excitatory neurons (EE connectivity) is dynamic and shaped by a Hebbian structural plasticity mechanism mimicking principles known from the neuroscience literature (Liao et al., 1995; Wu et al., 1996; Lüscher et al., 2000; Nevian and Sakmann, 2004; Deger et al., 2012). All other connections are static. The dynamics of the EE connectivity is determined by the time evolution of the permanences P_{ij} ($i, j \in \mathcal{E}$), representing the synapse maturity, and the synaptic weights J_{ij} . Unless the permanence P_{ij} exceeds a threshold θ_P , the synapse $\{j \rightarrow i\}$ is immature, with zero synaptic weight $J_{ij} = 0$. Upon threshold crossing, $P_{ij} \geq \theta_P$, the synapse becomes mature, and its weight is assigned a fixed value $J_{ij} = W$ ($\forall i, j$). Overall, the permanences evolve according to a Hebbian plasticity rule: the synapse $\{j \rightarrow i\}$ is potentiated, i.e., P_{ij} is increased, if the activation of the postsynaptic cell i is immediately preceded by an activation of the presynaptic cell j . Otherwise, the synapse is depressed, i.e., P_{ij} is decreased. At the beginning of the learning process or during relearning, the activity in the individual subpopulations is non-sparse. Hebbian learning alone would therefore lead to the strengthening of all existing synapses between two subsequently activated subpopulations, irrespective of the context these two subpopulations participate in. After learning, the subsets of neurons that are activated by a sequence element recurring in different sequences would therefore largely overlap. As a consequence, it becomes harder to distinguish between different contexts (histories) based on the activation patterns of these subsets. The original TM model (Hawkins and Ahmad, 2016) avoids this loss of context sensitivity by restricting synaptic potentiation to a small subset of synapses between a given pair of source and target subpopulations: if there are no predictive target neurons, the original algorithm selects a “matched” neuron from the set of active postsynaptic cells as the one being closest to becoming predictive, i.e., the neuron receiving the

largest number of synaptic inputs on a given dendritic branch from the set of active presynaptic cells (provided this number is sufficiently large). Synapse potentiation is then restricted to this set of matched neurons. In case there are no immature synapses, the “least used” neuron or a randomly chosen neuron is selected as the “matched” cell, and connected to the winner cell of the previously active subpopulation. Restricting synaptic potentiation to synapses targeting such a subset of “matched” neurons is difficult to reconcile with biology. It is known that inhibitory inputs targeting the dendrites of pyramidal cells can locally suppress backpropagating action potentials and, hence, synaptic potentiation (Müllner et al., 2015). A selection mechanism based on such local inhibitory circuits would however involve extremely fast synapses and require fine-tuning of parameters. The model presented in this work circumvents the selection of “matched” neurons and replaces this with a homeostatic mechanism controlled by the postsynaptic dAP rate. In the following, the specifics of the plasticity dynamics used in this study are described in detail.

Within the interval $[P_{\min,ij}, P_{\max}]$, the dimensionless permanences $P_{ij}(t)$ evolve according to a combination of an additive spike-timing-dependent plasticity (STDP) rule (Morrison et al., 2008) and a homeostatic component (Abbott and Nelson, 2000; Tetzlaff et al., 2011):

$$\begin{aligned} P_{\max}^{-1} \frac{dP_{ij}}{dt} = & \lambda_+ \sum_{\{t_i^*\}'} x_j(t) \delta(t - [t_i^* + d_{\text{EE}}]) I(t_i^*, \Delta t_{\min}, \Delta t_{\max}) \\ & - \lambda_- \sum_{\{t_j^*\}} y_i \delta(t - t_j^*) \\ & + \lambda_h \sum_{\{t_i^*\}'} (z^* - z_i(t)) \delta(t - t_i^*) I(t_i^*, \Delta t_{\min}, \Delta t_{\max}). \end{aligned} \quad (2.1)$$

At the boundaries $P_{\min,ij}$ and P_{\max} , $P_{ij}(t)$ is clipped. While the maximum permanences P_{\max} are identical for all EE connections, the minimal permanences $P_{\min,ij}$ are uniformly distributed in the interval $[P_{0,\min}, P_{0,\max}]$ to introduce a form of persistent heterogeneity. The first term on the right-hand side of Eq (2.1) corresponds to the spike-timing-dependent synaptic potentiation triggered by the postsynaptic spikes at times $t_i^* \in \{t_i^*\}'$. Here, $\{t_i^*\}' = \{t_i^* | \forall t_j^* : t_i^* - t_j^* + d_{\text{EE}} \geq \Delta t_{\min}\}$ denotes the set of all postsynaptic spike times t_i^* for which the time lag $t_i^* - t_j^* + d_{\text{EE}}$ exceeds Δt_{\min} for all presynaptic spikes t_j^* . The indicator function $I(t_i^*, \Delta t_{\min}, \Delta t_{\max})$ ensures that the potentiation (and the homeostasis; see below) is restricted to time lags $t_i^* - t_j^* + d_{\text{EE}}$ in the interval $(\Delta t_{\min}, \Delta t_{\max})$ to avoid a growth of synapses between synchronously active neurons belonging to the same subpopulation, and between neurons encoding for the first elements in different sequences; see Eq (2.17). Note that the potentiation update times lag the somatic postsynaptic spike times by the delay d_{EE} , which is here interpreted as a purely dendritic delay (Morrison et al., 2007, 2008). The potentiation increment is determined by the dimensionless potentiation rate λ_+ , and the spike trace $x_j(t)$ of the presynaptic neuron j , which is updated according to

$$\frac{dx_j}{dt} = -\tau_+^{-1} x_j(t) + \sum_{t_j^*} \delta(t - t_j^*). \quad (2.2)$$

The trace $x_j(t)$ is incremented by unity at each spike time t_j^* , followed by an exponential decay with time constant τ_+ . The potentiation increment ΔP_{ij} at time t_i^* therefore depends on the temporal distance between the postsynaptic spike time t_i^* and all

presynaptic spike times $t_j^* \leq t_i^*$ (STDP with all-to-all spike pairing; (Morrison et al., 2008)). The second term in Eq (2.1) represents synaptic depression, and is triggered by each presynaptic spike at times $t_j^* \in \{t_j^*\}$. The depression decrement $y_i = 1$ is treated as a constant, independently of the postsynaptic spike history. The depression magnitude is parameterized by the dimensionless depression rate λ_- . The third term in Eq (2.1) corresponds to a homeostatic control triggered by postsynaptic spikes at times $t_i^* \in \{t_i^*\}'$. Its overall impact is parameterized by the dimensionless homeostasis rate λ_h . The homeostatic control enhances or reduces the synapse growth depending on the dAP trace $z_i(t)$ of neuron i , the low-pass filtered dAP activity updated according to

$$\frac{dz_i}{dt} = -\tau_h^{-1} z_i(t) + \sum_k \delta(t - t_{\text{dAP},i}^k). \quad (2.3)$$

Here, τ_h represents the homeostasis time constant, and $t_{\text{dAP},i}^k$ the onset time of the k th dAP in neuron i . According to Eq (2.1), synapse growth is boosted if the dAP activity $z_i(t)$ is below a target dAP activity z^* . Conversely, high dAP activity exceeding z^* reduces the synapse growth (Fig. 2.4). This homeostatic regulation of the synaptic maturity controlled by the postsynaptic dAP activity constitutes a variation of previous models (Abbott and Nelson, 2000; Tetzlaff et al., 2011) describing ‘synaptic scaling’ (Turrigiano et al., 1998; Turrigiano and Nelson, 2004; Turrigiano, 2008). It counteracts excessive synapse formation during learning driven by Hebbian structural plasticity. In addition, the combination of Hebbian plasticity and synaptic scaling can introduce a competition between synapses (Abbott and Nelson, 2000; Tetzlaff et al., 2011). Here, we exploit this effect to ensure that synapses are generated in a context specific manner, and thereby reduce the overlap between neuronal subpopulations activated by the same sequence element occurring in different sequences. To this end, the homeostasis parameters $z^* = 1$ and τ_h are chosen such that each neuron tends to become predictive, i.e., generate a dAP, at most once during the presentation of a single sequence ensemble of total duration $((C-1)\Delta T + \Delta T_{\text{seq}})S$ (see Network model). The time constant τ_h is hence adapted to the parameters of the task. For sequence sets I and II and the default inter-stimulus interval $\Delta T = 40$ ms, it is set to $\tau_h = 440$ ms and $\tau_h = 1560$ ms, respectively. In section section 2.3.3, we study the effect of the sequence speed (inter-stimulus interval ΔT) on the prediction performance for a given network parameterization. For these experiments, $\tau_h = 440$ ms is therefore fixed even though the inter-stimulus interval ΔT is varied.

The prefactor P_{max}^{-1} in Eq (2.1) ensures that all learning rates λ_+ , λ_- and λ_h are measured in units of the maximum permanence P_{max} .

CONSTRAINTS ON POTENTIAL CONNECTIVITY. The sequence processing capabilities of the proposed network model rely on its ability to form sequence specific subnetworks based on the skeleton provided by the random potential connectivity. On the one hand, the potential connectivity must not be too diluted to ensure that a subset of neurons representing a given sequence element can establish sufficiently many mature connections to a second subset of neurons representing the subsequent element. On the other hand, a dense potential connectivity would promote overlap between subnetworks representing different sequences, and thereby slow down the formation of context specific subnetworks during learning (see section 2.3.1). Here, we therefore identify the minimal potential connection probability p guaranteeing the existence of network motifs with a sufficient degree of divergent-convergent connectivity.

Consider the subset \mathcal{P}_{ij} of ρ excitatory neurons representing the j th sequence element ζ_{ij} in sequence s_i (see [Task and training protocol](#) and [Network model](#)). During the learning process, the plasticity dynamics needs to establish mature connections from \mathcal{P}_{ij} to a second subset $\mathcal{P}_{i,j+1}$ of neurons in another subpopulation representing the subsequent element $\zeta_{i,j+1}$. Each neuron in $\mathcal{P}_{i,j+1}$ must receive at least $c = \lceil \theta_{\text{dAP}}/W \rceil$ inputs from \mathcal{P}_{ij} to ensure that synchronous firing of the neurons in \mathcal{P}_{ij} can evoke a dAP after synapse maturing. For a random, homogeneous potential connectivity with connection probability p , the probability of finding these c potential connections for some arbitrary target neuron is given by

$$q(c; \rho, p) = \sum_{k=c}^{\rho} \binom{\rho}{k} p^k (1-p)^{\rho-k}. \quad (2.4)$$

For a successful formation of sequence specific subnetworks during learning, the sparse subset \mathcal{P}_{ij} of presynaptic neurons needs to recruit at least ρ targets in the set of n_E neurons representing the subsequent sequence element (Fig. 2.5A). The probability of observing such a divergent-convergent connectivity motif is given by

$$u(\rho; c, p, n_E) = \sum_{l=\rho}^{n_E} \binom{n_E}{l} q^l (1-q)^{n_E-l}. \quad (2.5)$$

Note that the above described motif does not require the size of the postsynaptic subset $\mathcal{P}_{i,j+1}$ to be exactly ρ . Eq (2.5) constrains the parameters p , c , n_E and ρ to ensure such motifs exist in a random network. Fig. 2.5B illustrates the dependence of the motif probability u on the connection probability p for our choice of parameters n_E , c , and ρ . For $p \geq 0.2$, the existence of the divergent-convergent connectivity motif is almost certain ($u \approx 1$). For smaller connection probabilities $p < 0.2$, the motif probability quickly vanishes. Hence, $p = 0.2$ constitutes a reasonable choice for the potential connection probability.

NETWORK REALIZATIONS AND INITIAL CONDITIONS. For every network realization, the potential connectivity and the initial permanences are drawn randomly and independently. All other parameters are identical for different network realizations. The initial values of all state variables are given in Table 2.1 and Table 7.1.

SIMULATION DETAILS. The network simulations are performed in the neural simulator NEST ([Gewaltig and Diesmann, 2007](#)) under version 3.0 ([Hahne et al., 2021](#)). The differential equations and state transitions defining the excitatory neuron dynamics are expressed in the domain specific language NESTML ([Plotnikov et al., 2016](#); [Nagendra Babu et al., 2021](#)) which generates the required C++ code for the dynamic loading into NEST. Network states are synchronously updated using exact integration of the system dynamics on a discrete-time grid with step size Δt ([Rotter and Diesmann, 1999](#)). The full source code for the implementation with a list of other software requirements is available at Zenodo: <https://doi.org/10.5281/zenodo.5578212>.

2.2.3 Task performance measures

To assess the network performance, we monitor the dendritic currents reporting predictions (dAPs) as well as the somatic spike times of excitatory neurons. To quantify the prediction error, we identify for each last element ζ_{i,C_i} in a sequence s_i all

excitatory neurons that have generated a dAP in the time interval $(t_{i,C_i} - \Delta T, t_{i,C_i})$, where t_{i,C_i} and ΔT denote the time of the external input corresponding to the last sequence element ζ_{i,C_i} and the inter-stimulus interval, respectively (see [Task and training protocol](#) and [Network model](#)). All subpopulations \mathcal{M}_k with at least $\rho/2$ neurons generating a dAP are considered “predictive”. The prediction state of the network is encoded in an M dimensional binary vector \mathbf{o} , where $o_k = 1$ if the k th subpopulation is predictive, and $o_k = 0$ else. The

$$\text{prediction error} = \frac{1}{L} \sqrt{\sum_{k=1}^M (o_k - v_k)^2} \quad (2.6)$$

is defined as the Euclidean distance between \mathbf{o} and the binary target vector \mathbf{v} representing the pattern of external inputs for each last element ζ_{i,C_i} , normalized by the number L of subpopulations per sequence element. Furthermore, we assess the

$$\text{false positive rate} = \frac{1}{L} \sum_{k=1}^M \Theta(o_k - v_k) \quad (2.7)$$

and the

$$\text{false negative rate} = \frac{1}{L} \sum_{k=1}^M \Theta(v_k - o_k), \quad (2.8)$$

where $\Theta(\cdot)$ denotes the Heaviside function. In addition to these performance measures, we monitor for each last sequence element the level of sparsity by measuring the ratio between the number of active neurons and the total number Ln_E of neurons representing this element. During learning, we expose the network repetitively to the same set $\{s_1, \dots, s_S\}$ of sequences for a number of training episodes K . To obtain the total prediction performance in each episode, we average the prediction error, the false negative and false positive rates, as well as the level of sparsity across the set of sequences.

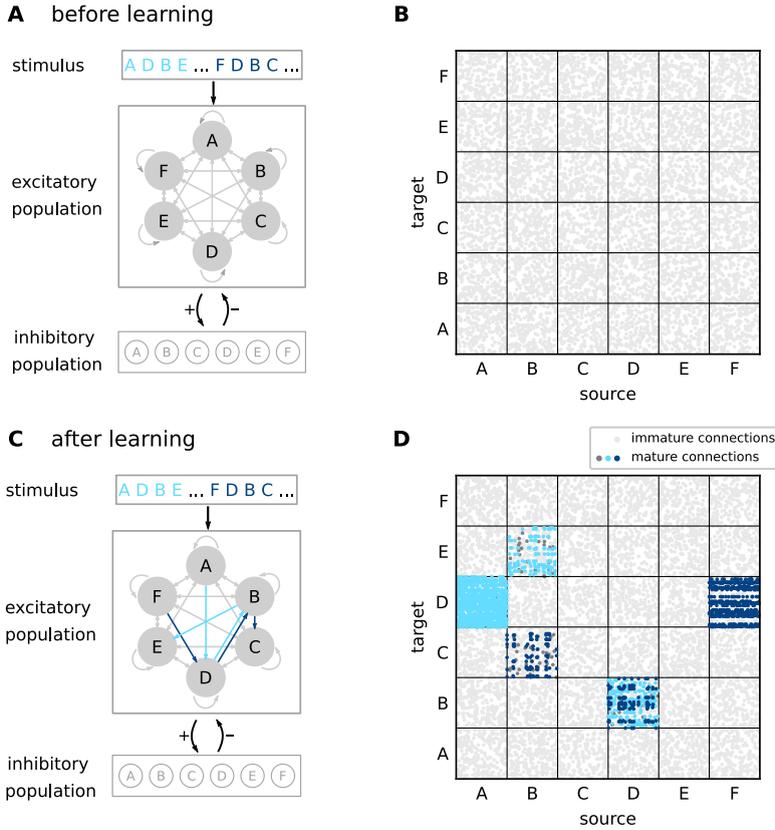


Figure 2.2: **Sketch of the network structure.** **A)** The architecture constitutes a recurrent network of excitatory and inhibitory neurons. Excitatory neurons are stimulated by external sources providing sequence-element specific inputs “A”; “D”, etc. The excitatory neuron population is composed of subpopulations containing neurons with identical stimulus preference (gray circles). Connections between and within the excitatory subpopulations are random and sparse. Inhibitory neurons are mutually unconnected. Each neuron in the inhibitory population is recurrently connected to a specific subpopulation of excitatory neurons. **B)** Initial connectivity matrix for excitatory connections to excitatory neurons (EE connections). Target and source neurons are grouped into stimulus-specific subpopulations (“A”, ..., “F”). Before learning, the excitatory neurons are sparsely and randomly connected via immature synapses (light gray dots). **C)** During learning, sequence specific, sparsely connected subnetworks with mature synapses are formed (light blue arrows: {A,D,B,E}, dark blue arrows: {F,D,B,C}). **D)** EE connectivity matrix after learning. During the learning process, subsets of connections between subpopulations corresponding to subsequent sequence elements become mature and effective (light and dark blue dots). Mature connections are context specific (see distinct connectivity between subpopulations “D” and “B” corresponding to different sequences), thereby providing the backbone for a reliable propagation of sequence-specific activity. In panels B and D, only 5% of sequence non-specific EE connections are shown for clarity. Dark gray dots in panel D correspond to mature connections between neurons that remain silent after learning. For details on the network structure, see Table 2.1 and Table 7.1.

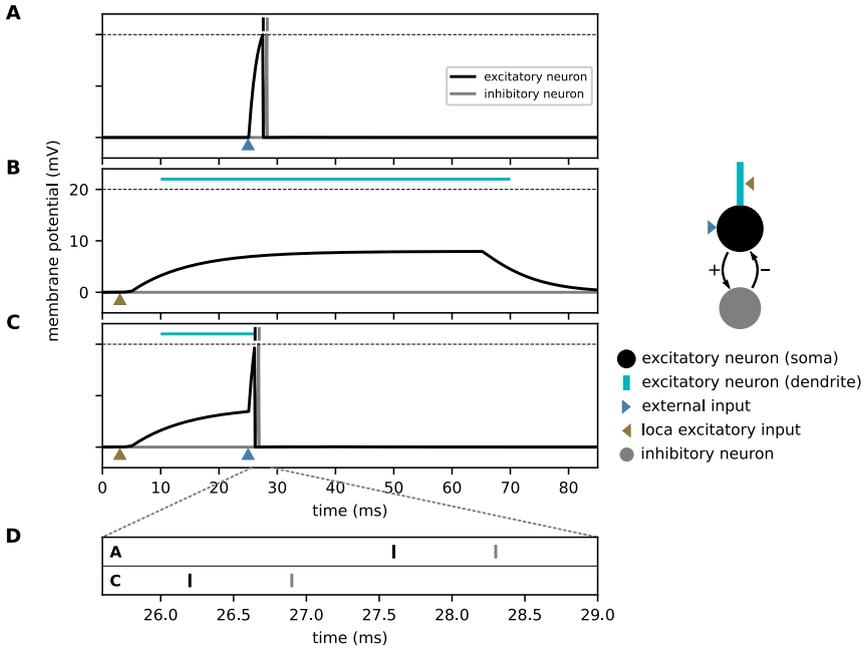


Figure 2.3: **Effect of dendritic action potentials (dAP) on the firing response to an external stimulus.** Membrane-potential responses to an external input (blue arrow, **A**), a strong dendritic input (brown arrow, **B**) triggering a dAP, and a combination of both (**C**). Black and gray vertical bars mark times of excitatory and inhibitory spikes, respectively. The horizontal dashed line marks the spike threshold θ_E . The horizontal light blue lines depict the dAP plateau. **D**) Magnified view of spike times from panels **A** and **C**. A dAP preceding the external input (as in panel **C**) can speed up somatic, and hence, inhibitory firing, provided the time interval between the dAP and the external input is in the right range. The excitatory neuron is connected bidirectionally to an inhibitory neuron (see sketch on the right).

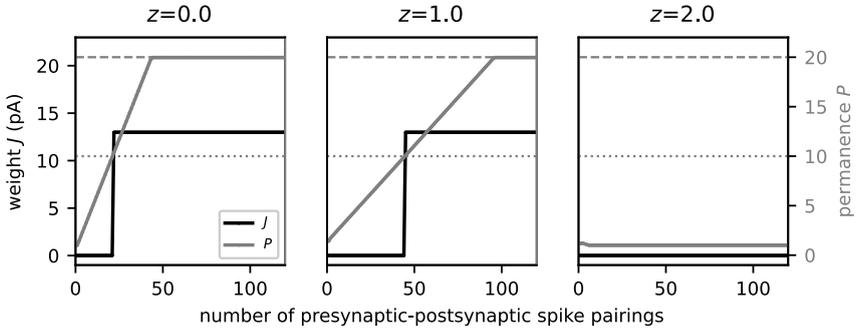


Figure 2.4: **Homeostatic regulation of the spike-timing-dependent structural plasticity by the dAP activity.** Evolution of the synaptic permanence (gray) and weight (black) during repetitive presynaptic-postsynaptic spike pairing for different levels of the dAP activity. In the depicted example, presynaptic spikes precede the postsynaptic spikes by 40 ms for each spike pairing. Consecutive spike pairs are separated by a 200 ms interval. In each panel, the postsynaptic dAP trace is clamped at a different value: $z = 0$ (left), $z = 1$ (middle), $z = 2$ (right). The dAP target activity is fixed at $z^* = 1$. The horizontal dashed and dotted lines mark the maximum permanence P_{\max} and the maturity threshold θ_P , respectively.

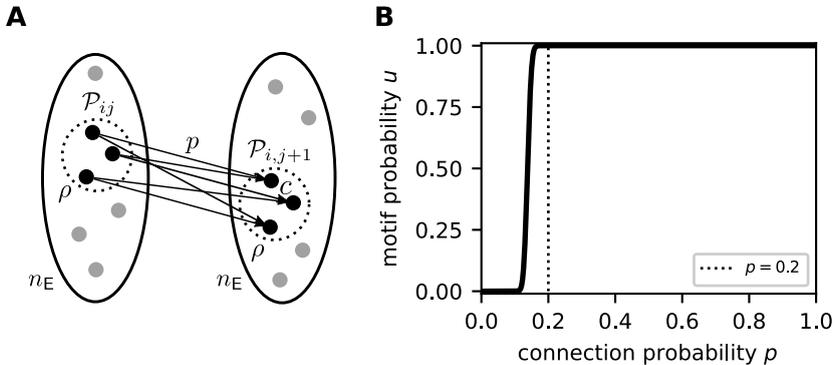


Figure 2.5: **Existence of divergent-convergent connectivity motifs in a random network.** **A)** Sketch of the divergent-convergent potential connectivity motif required for the formation of sequence specific subnetworks during learning. See main text for details. **B)** Dependence of the motif probability u on the connection probability p for $n_E = 150$, $c = 5$, and $\rho = 20$ (see Table 7.1). The dotted vertical line marks the potential connection probability $p = 0.2$ used in this study.

2.2.4 Model tables

Summary		
Populations	excitatory neurons (\mathcal{E}), inhibitory neurons (\mathcal{I}), external spike sources (\mathcal{X}); \mathcal{E} and \mathcal{I} composed of M disjoint subpopulations \mathcal{M}_k and \mathcal{I}_k ($k = 1, \dots, M$)	
Connectivity	<ul style="list-style-type: none"> • sparse random connectivity between excitatory neurons (plastic) • local recurrent connectivity between excitatory and inhibitory neurons (static) 	
Neuron model	<ul style="list-style-type: none"> • excitatory neurons: leaky integrate-and-fire (LIF) with nonlinear input integration (dendritic action potentials) • inhibitory neurons: leaky integrate-and-fire (LIF) 	
Synapse model	exponential or alpha-shaped postsynaptic currents (PSCs)	
Plasticity	homeostatic spike-timing-dependent structural plasticity in excitatory-to-excitatory connections	
Populations		
Name	Elements	Size
$\mathcal{E} = \cup_{i=k}^M \mathcal{M}_k$	excitatory (E) neurons	N_E
$\mathcal{I} = \cup_{i=k}^M \mathcal{I}_k$	inhibitory (I) neurons	N_I
\mathcal{M}_k	excitatory neurons in subpopulation k $\mathcal{M}_k \cap \mathcal{M}_l = \emptyset$ ($\forall k \neq l \in [1, M]$)	n_E
\mathcal{I}_k	inhibitory neurons in subpopulation k $\mathcal{I}_k \cap \mathcal{I}_l = \emptyset$ ($\forall k \neq l \in [1, M]$)	n_I
$\mathcal{X} = \{x_1, \dots, x_M\}$	external spike sources	M
Connectivity		
Source population	Target population	Pattern
\mathcal{E}	\mathcal{E}	random; fixed in-degrees $K_i = K_{EE}$, delays $d_{ij} = d_{EE}$, synaptic time constants $\tau_{ij} = \tau_{EE}$; plastic weights $J_{ij} \in \{0, J_{EE,ij}\}$ ($\forall i \in \mathcal{E}, \forall j \in \mathcal{E}$; "EE connections")
\mathcal{M}_k	\mathcal{I}_k	all-to-all; fixed delays $d_{ij} = d_{IE}$, synaptic time constants $\tau_{ij} = \tau_{IE}$, and weights $J_{ij} = J_{IE}$ ($\forall i \in \mathcal{I}_k, \forall j \in \mathcal{M}_k, \forall k \in [1, M]$; "IE connections")
\mathcal{I}_k	\mathcal{M}_k	all-to-all; fixed delays $d_{ij} = d_{EI}$, synaptic time constants $\tau_{ij} = \tau_{EI}$, and weights $J_{ij} = J_{EI}$ ($\forall i \in \mathcal{M}_k, \forall j \in \mathcal{I}_k, \forall k \in [1, M]$; "EI connections")
\mathcal{I}_k	\mathcal{I}_k	none ($\forall k \in [1, M]$; "II connections")
$\mathcal{X}_k = x_k$	\mathcal{M}_k	one-to-all; fixed delays $d_{ik} = d_{EX}$, synaptic time constants $\tau_{ij} = \tau_{EX}$, and weights $J_{ik} = J_{EX}$ ($\forall i \in \mathcal{M}_k, \forall k \in [1, M]$; "EX connections")
no self-connections ("autapses"), no multiple connections ("multapses")		
all unmentioned connections $\mathcal{M}_k \rightarrow \mathcal{I}_l, \mathcal{I}_k \rightarrow \mathcal{M}_l, \mathcal{I}_k \rightarrow \mathcal{I}_l, \mathcal{X}_k \rightarrow \mathcal{M}_l$ ($\forall k \neq l$) are absent		

Continued next page

Neuron and synapse	
Neuron	
Type	leaky integrate-and-fire (LIF) dynamics
Description	<p>dynamics of membrane potential $V_i(t)$ and spiking activity $s_i(t)$ of neuron i:</p> <ul style="list-style-type: none"> • emission of the kth spike of neuron i at time t_i^k if $V_i(t_i^k) \geq \theta_i \quad (2.9)$ <p style="text-align: center;">with somatic spike threshold θ_i</p> • spike train: $s_i(t) = \sum_k \delta(t - t_i^k)$ • reset and refractoriness: $V_i(t) = V_r \quad \forall k, \forall t \in (t_i^k, t_i^k + \tau_{\text{ref},i}]$ <p style="text-align: center;">with refractory time $\tau_{\text{ref},i}$ and reset potential V_r</p> • subthreshold dynamics: $\tau_{\text{m},i} \dot{V}_i(t) = -V_i(t) + R_{\text{m},i} I_i(t) \quad (2.10)$ <p style="text-align: center;">with membrane resistance $R_{\text{m},i} = \frac{\tau_{\text{m},i}}{C_{\text{m},i}}$, membrane time constant $\tau_{\text{m},i}$, and total synaptic input current $I_i(t)$ (see Synapse)</p> • excitatory neurons: $\tau_{\text{m},i} = \tau_{\text{m},\text{E}}$, $C_{\text{m},i} = C_{\text{m}}$, $\theta_i = \theta_{\text{E}}$, $\tau_{\text{ref},i} = \tau_{\text{ref},\text{E}}$ ($\forall i \in \mathcal{E}$) • inhibitory neurons: $\tau_{\text{m},i} = \tau_{\text{m},\text{I}}$, $C_{\text{m},i} = C_{\text{m}}$, $\theta_i = \theta_{\text{I}}$, $\tau_{\text{ref},i} = \tau_{\text{ref},\text{I}}$ ($\forall i \in \mathcal{I}$)

Continue next page.

Synapse	
Type	exponential or alpha-shaped postsynaptic currents (PSCs)
Description	<ul style="list-style-type: none"> • total synaptic input currents: <ul style="list-style-type: none"> excitatory neurons: $I_i(t) = I_{ED,i}(t) + I_{EX,i}(t) + I_{EI,i}(t), \forall i \in \mathcal{E}$ (2.11) inhibitory neurons: $I_i(t) = I_{IE,i}(t), \forall i \in \mathcal{I}$ with dendritic, external, inhibitory and excitatory input currents $I_{ED,i}(t), I_{EX,i}(t), I_{EI,i}(t), I_{IE,i}(t)$ evolving according to $I_{ED,i}(t) = \sum_{j \in \mathcal{E}} (\alpha_{ij} * s_j)(t - d_{ij}) \quad (2.12)$ with $\alpha_{ij}(t) = J_{ij} \frac{e}{\tau_{ED}} t e^{-t/\tau_{ED}} \Theta(t)$ and $\Theta(t) = \begin{cases} 1 & t \geq 0 \\ 0 & \text{else} \end{cases}$, $\tau_{EX} \dot{I}_{EX,i} = -I_{EX,i}(t) + \sum_{j \in \mathcal{X}} J_{ij} s_j(t - d_{ij}), \quad (2.13)$ $\tau_{EI} \dot{I}_{EI,i} = -I_{EI,i}(t) + \sum_{j \in \mathcal{I}} J_{ij} s_j(t - d_{ij}), \quad (2.14)$ $\tau_{IE} \dot{I}_{IE,i} = -I_{IE,i}(t) + \sum_{j \in \mathcal{E}} J_{ij} s_j(t - d_{ij}) \quad (2.15)$ with τ_{EX}, τ_{EI}, and τ_{IE} synaptic time constants of EX, EI, and IE connections, respectively, and J_{ij} the synaptic weight • suprathreshold dynamics of dendritic currents (dAP generation): <ul style="list-style-type: none"> – emission of kth dAP of neuron i at time $t_{dAP,i}^k$ if $I_{ED,i}(t_{dAP,i}^k) \geq \theta_{dAP}$ – dAP current plateau: $I_{ED,i}(t) = I_{dAP} \quad \forall k, \forall t \in (t_{dAP,i}^k, t_{dAP,i}^k + \tau_{dAP}) \quad (2.16)$ with dAP current plateau amplitude I_{dAP}, dAP current duration τ_{dAP}, and dAP activation threshold θ_{dAP} – reset: $I_{ED,i}(t_{dAP,i}^k + \tau_{dAP}) = 0$ ($\forall k$) – reset and refractoriness in response to emission of lth somatic spike of neuron i at time t_l^i: $I_{ED,i}(t) = 0 \quad \forall l, \forall t \in (t_l^i, t_l^i + \tau_{ref,i})$

Continued next page

Plasticity	
Type	spike-timing-dependent structural plasticity and dAP-rate homeostasis
EE synapses	<ul style="list-style-type: none"> • dynamics of synaptic permanence $P_{ij}(t)$ (synapse maturity): $P_{\max}^{-1} \frac{dP_{ij}}{dt} = \lambda_+ \sum_{\{t_j^*\}'} x_j(t) \delta(t - [t_i^* + d_{EE}]) I(t_i^*, \Delta t_{\min}, \Delta t_{\max})$ $- \lambda_- \sum_{\{t_j^*\}} y_i \delta(t - t_j^*)$ $+ \lambda_h \sum_{\{t_j^*\}'} (z^* - z_i(t)) \delta(t - t_i^*) I(t_i^*, \Delta t_{\min}, \Delta t_{\max})$ <p>with</p> <ul style="list-style-type: none"> – list of presynaptic spike times $\{t_j^*\}$, – list of postsynaptic spike times $\{t_i^*\}' = \{t_i^* \forall t_j^* : t_i^* - t_j^* + d_{EE} \geq \Delta t_{\min}\}$ – indicator function $I(t_i^*, \Delta t_{\min}, \Delta t_{\max}) = R(t_i^* - t_j^+ + d_{EE})$ $\text{with } R(\tau) = \begin{cases} 1 & \Delta t_{\min} < \tau < \Delta t_{\max} \\ 0 & \text{else,} \end{cases} \quad (2.17)$ – maximum permanence P_{\max}, potentiation and depression rates λ_+, λ_-, homeostasis rate λ_h, delay d_{EE}, depression decrement y_i, minimum Δt_{\min} and maximum Δt_{\max} time lags between pairs of pre- and postsynaptic spikes at which synapses are potentiated, t_j^+ is the nearest presynaptic spike time preceding t_i^*, – spike trace $x_j(t)$ of presynaptic neuron j, evolving according to $\frac{dx_j}{dt} = -\tau_+^{-1} x_j(t) + \sum_{t_j^*} \delta(t - t_j^*)$ with presynaptic spike times t_j^* and potentiation time constant τ_+, – dAP trace $z_i(t)$ of postsynaptic neuron i, evolving according to $\frac{dz_i}{dt} = -\tau_h^{-1} z_i(t) + \sum_k \delta(t - t_{\text{dAP},i}^k)$ with onset time $t_{\text{dAP},i}^k$ of the kth dAP, homeostasis time constant τ_h, and – target dAP activity z^* • dynamics of synaptic weights $J_{EE,ij}$: $J_{EE,ij}(t) = \begin{cases} W & \text{if } P_{ij}(t) \geq \theta_P \quad (\text{mature synapse}) \\ 0 & \text{if } P_{ij}(t) < \theta_P \quad (\text{immature synapse}) \end{cases}$ <p>with weight of mature EE connections W and synapse maturity threshold θ_P</p> <p>(for an algorithmic implementation of the plasticity dynamics, see Algorithm 1)</p>
all other synapses	non-plastic

Continued next page

Input
<ul style="list-style-type: none"> • prediction mode <ul style="list-style-type: none"> – repetitive stimulation with the same set $\mathcal{S} = \{s_1, \dots, s_S\}$ of sequences $s_i = \{\zeta_{i,1}, \zeta_{i,2}, \dots, \zeta_{i,C_i}\}$ of ordered discrete items $\zeta_{i,j}$ with number of sequences S and length C_i of ith sequence – presentation of sequence element $\zeta_{i,j}$ at time $t_{i,j}$ modeled by single spike $x_k(t) = \delta(t - t_{i,j})$, generated by the corresponding external source x_k – inter-stimulus interval $\Delta T = t_{i,j+1} - t_{i,j}$ between subsequent sequence elements $\zeta_{i,j}$ and $\zeta_{i,j+1}$ within a sequence s_i – inter-sequence time interval $\Delta T_{\text{seq}} = t_{i+1,1} - t_{i,C_i}$ between subsequent sequences s_i and s_{i+1} – example sequence sets: <ul style="list-style-type: none"> * sequence set I: $\mathcal{S} = \{\{A,D,B,E\}, \{F,D,B,C\}\}$ * sequence set II: $\mathcal{S} = \{\{E,N,D,I,J\}, \{L,N,D,I,K\}, \{G,J,M,C,N\}, \{F,J,M,C,I\}, \{B,C,K,H,I\}, \{A,C,K,H,F\}\}$ • replay mode <ul style="list-style-type: none"> – presentation of a cue encoding for first sequence elements $\zeta_{i,1}$ at $t_{i,1}$ – inter-cue time interval $\Delta T_{\text{cue}} = t_{i+1,1} - t_{i,1}$ between subsequent cues $\zeta_{i,1}$ and $\zeta_{i+1,1}$
Output
<ul style="list-style-type: none"> • somatic spike times $\{t_i^k \forall i \in \mathcal{E}, k = 1, 2, \dots\}$ • dendritic currents $I_{\text{ED},i}(t)$ ($\forall i \in \mathcal{E}$)
Initial conditions and network realizations
<ul style="list-style-type: none"> • membrane potentials: $V_i(0) = V_r$ ($\forall i \in \mathcal{E} \cup \mathcal{I}$) • dendritic currents: $I_{\text{ED},i}(0) = 0$ ($\forall i \in \mathcal{E}$) • external currents: $I_{\text{EX},i}(0) = 0$ ($\forall i \in \mathcal{E}$) • inhibitory currents: $I_{\text{EI},i}(0) = 0$ ($\forall i \in \mathcal{E}$) • excitatory currents: $I_{\text{IE},i}(0) = 0$ ($\forall i \in \mathcal{I}$) • synaptic permanences: $P_{i,j}(0) = P_{\min,i,j}$ with $P_{\min,i,j} \sim \mathcal{U}(P_{0,\min}, P_{0,\max})$ ($\forall i, j \in \mathcal{E}$) • synaptic weights: $J_{\text{EE},i,j}(0) = 0$ ($\forall i, j \in \mathcal{E}$) • spike traces: $x_i(0) = 0$ ($\forall i \in \mathcal{E}$) • dAP traces: $z_i(0) = 0$ ($\forall i \in \mathcal{E}$) • potential connectivity and initial permanences randomly and independently drawn for each network realization
Simulation details
<ul style="list-style-type: none"> • network simulations performed in NEST (Gewaltig and Diesmann, 2007) version 3.0 (Hahne et al., 2021) • definition of excitatory neuron model using NESTML (Plotnikov et al., 2016; Nagendra Babu et al., 2021) • synchronous update using exact integration of system dynamics on discrete-time grid with step size Δt (Rotter and Diesmann, 1999) • source code underlying this study: https://doi.org/10.5281/zenodo.5578212

Table 2.1: Description of the network model. Parameter values are given in Table 7.1.

Algorithm 1 Algorithmic description of the plasticity model, based on the algorithm proposed in (Morrison et al., 2007).

Update of permanence P_{ij} and synaptic weight $J_{EE,ij}$ at time t_j^k of the k th spike of presynaptic neuron j :

```

 $x_j \leftarrow$  get trace of presynaptic neuron  $j$  ▷ last update at time  $t_j^{k-1}$ 
 $L_{\text{post}}, Z_{\text{post}} \leftarrow$  get lists of postsynaptic spike times and corresponding dAP traces in the interval
 $(t_j^{k-1} - d_{EE}, t_j^k - d_{EE})$ 

for  $t_i, z_i$  in  $L_{\text{post}}, Z_{\text{post}}$  do
  if  $\Delta t_{\text{min}} < t_i - t_j^{k-1} + d_{EE} < \Delta t_{\text{max}}$  then
     $P_{ij} \leftarrow P_{ij} + \lambda_+ \cdot P_{\text{max}} \cdot x_j \cdot \exp(-(t_i - t_j^{k-1} + d_{EE})/\tau_+)$  ▷ potentiation
     $P_{ij} \leftarrow P_{ij} + \lambda_h \cdot P_{\text{max}} \cdot (z^* - z_i)$  ▷ homeostasis
  end if
end for

 $P_{ij} \leftarrow P_{ij} - \lambda_- \cdot P_{\text{max}} \cdot y_i$  ▷ depression

if  $P_{ij} > \theta_P$  then
   $J_{EE,ij} \leftarrow W$  ▷ mature synapse
else
   $J_{EE,ij} \leftarrow 0$  ▷ immature synapse
end if

 $x_j \leftarrow x_j \cdot \exp((t_j^{k-1} - t_j^k)/\tau_+) + 1$  ▷ update of presynaptic spike trace

```

Note: the clipping of the permanence P_{ij} at the boundaries of the interval $[P_{\text{min},ij}, P_{\text{max}}]$ is not included here for clarity.

2.3 RESULTS

2.3.1 Sequence learning and prediction

According to the Temporal Memory (TM) model, sequences are stored in the form of specific paths through the network. Prediction and replay of sequences correspond to a sequential sparse activation of small groups of neurons along these paths. Non-anticipated stimuli are signaled in the form of non-sparse firing of these groups. This subsection describes how the model components introduced in [Network model](#) interact and give rise to the network structure and behavior postulated by TM. For illustration, we here consider a simple set of two partly overlapping sequences {A,D,B,E} and {F,D,B,C} corresponding to the sequence set I (see Fig. 2.1B).

The initial sparse, random and immature network connectivity (Fig. 2.2A,B) constitutes the skeleton on which the sequence-specific paths will be carved out during the learning process. To guarantee a successful learning, this initial skeleton must be neither too sparse nor too dense (see [Methods](#)). Before learning, the presentation of a particular sequence element causes all neurons with the corresponding stimulus preference to reliably and synchronously fire a somatic action potential due to the strong, suprathreshold external stimulus (Fig. 2.3A). All other subpopulations remain silent (see Fig. 2.6A,B). The lateral connectivity between excitatory neurons belonging to the different subpopulations is subject to a form of Hebbian structural plasticity. Repetitive and consistent sequential presentation of sequence elements turns immature connections between successively activated subpopulations into mature connections, and hence leads to the formation of sequence-specific subnetworks (see Fig. 2.2C,D).

Synaptic depression prunes connections not supporting the learned pattern, thereby reducing the chance of predicting wrong sequence items (false positives).

During the learning process, the number of mature connections grows to a point where the activation of a certain subpopulation by an external input generates dendritic action potentials (dAPs), a “prediction”, in a subset of neurons in the subsequent subpopulation (blue neurons in Fig. 2.6C). The dAPs generate a long-lasting depolarization of the soma (Fig. 2.3B). When receiving an external input, these depolarized neurons fire slightly earlier as compared to non-depolarized (non-predictive) neurons (Fig. 2.3A,B,D). If the number of predictive neurons within a subpopulation is sufficiently large, their advanced spikes (Fig. 2.3C) initiate a fast and strong inhibitory feedback to the entire subpopulation, and thereby suppress subsequent firing of non-predictive neurons in this population (Fig. 2.6C,D). Owing to this winner-take-all dynamics, the network generates sparse spiking in response to predicted stimuli, i.e., if the external input coincides with a dAP-triggered somatic depolarization. In the presence of a non-anticipated, non-predicted stimulus, the neurons in the corresponding subpopulation fire collectively in a non-sparse manner, thereby signaling a “mismatch”.

In the model presented in this study, the initial synapse maturity levels, the permanences, are randomly chosen within certain bounds. During learning, connections with a higher initial permanence mature first. This heterogeneity in the initial permanences enables the generation of sequence specific sparse connectivity patterns between subsequently activated neuronal subpopulations (Fig. 2.2D). For each pair of sequence elements in a given sequence ensemble, there is a unique set of postsynaptic neurons generating dAPs (Fig. 2.6D). These different activation patterns capture the context specificity of predictions. When exposing a network that has learned the two sequences $\{A,D,B,E\}$ and $\{F,D,B,C\}$ to the elements “A” and “F”, different subsets of neurons are activated in “D” and “B”. By virtue of these sequence specific activation patterns, stimulation by $\{A,D,B\}$ or $\{F,D,B\}$ leads to correct predictions “E” or “C”, respectively (Fig. 2.6C–F).

Heterogeneity in the permanences alone, however, is not sufficient to guarantee context specificity. The subsets of neurons activated in different contexts may still exhibit a considerable overlap. This overlap is promoted by Hebbian plasticity in the face of the initial non-sparse activity, which leads to a strengthening of connections to neurons in the postsynaptic population in an unspecific manner (Fig. 2.7A,B). Moreover, the reoccurrence of the same sequence elements in different co-learned sequences initially causes higher firing rates of the neurons in the respective populations (“D” and “B” in Fig. 2.7). As a result, the formation of unspecific connections would even be accelerated if synapse formation was driven by Hebbian plasticity alone. The model in this study counteracts this loss of context specificity by supplementing the plasticity dynamics with a homeostatic component, which regulates synapse growth based on the rate of postsynaptic dAPs. This form of homeostasis prevents the same neuron from becoming predictive multiple times within the same set of sequences, and thereby reduces the overlap between subsets of neurons activated within different contexts (Fig. 2.7C, Fig. 7.3). To further aid the formation of context specific paths, the density of the initial potential connectivity skeleton is set close to the minimum value ensuring the existence of the connectivity motifs required for a faithful prediction (see [Methods](#)).

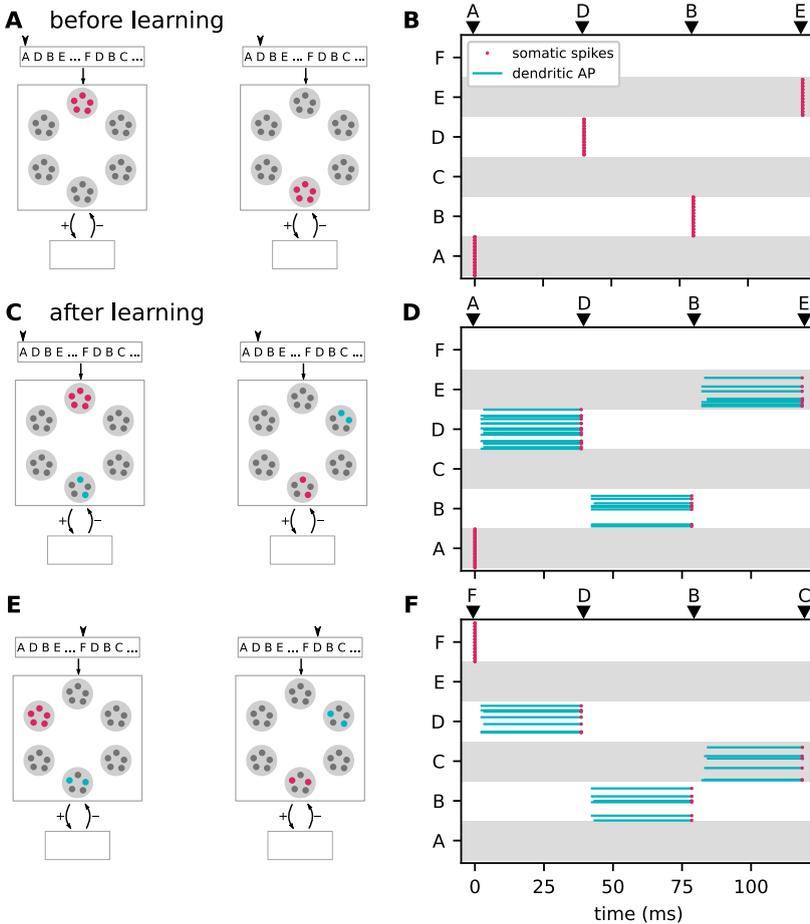


Figure 2.6: **Context specific predictions.** Sketches (left column) and raster plots of network activity (right column) before (top row) and after learning of the two sequences $\{A,D,B,E\}$ and $\{F,D,B,C\}$ (middle and bottom rows). In the left column, large light gray circles depict the excitatory subpopulations (same arrangement as in Fig. 2.2). Red, blue and gray circles mark active, predictive and silent neurons, respectively. In the right column, red dots and blue lines mark somatic spikes and dAP plateaus, respectively. Type and timing of presented stimuli are depicted by black arrows. **A,B**) Snapshots of network activity upon subsequent presentation of the sequence elements “A” and “D” (panel A), and network activity in response to presentation of the entire sequence $\{A,D,B,E\}$ (panel B) before learning. All neurons in the stimulated subpopulations become active.

2.3.2 Prediction performance

To quantify the sequence prediction performance, we repetitively stimulate the network with the sequences in sequence set I (see [Task and training protocol](#)), and continuously monitor the prediction error, the false-positive and false-negative rates, as well as

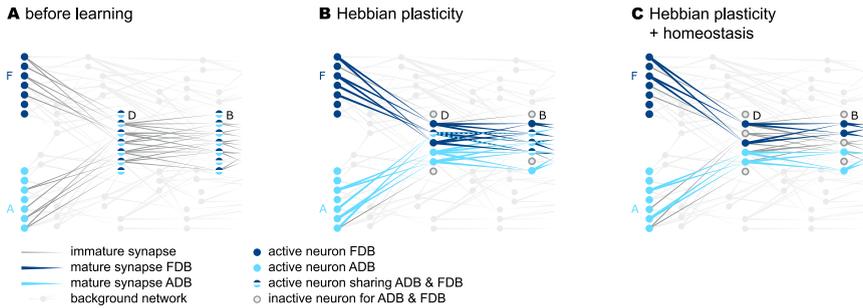


Figure 2.7: **dAP-rate homeostasis enhances context specificity.** **A)** Sketch of subpopulations of excitatory neurons representing the elements of the two sequences $\{F,D,B\}$ and $\{A,D,B\}$, depicted by light and dark blue colors, respectively. Before learning, the connections between the subpopulations are immature (gray lines). Hence, for each element presentation, all neurons in the respective subpopulations fire (filled circles). **B)** Hebbian plasticity drives the formation of mature connections between subpopulations representing successive sequence elements (colored lines), and leads to sparse firing. The sets of neurons contributing to the two sequences partly overlap. **C)** Incorporating dAP-rate homeostasis reduces this overlap in the activation patterns.

the fraction of active stimulated neurons as a measure of encoding sparsity (Fig. 2.8; section 2.2.3). To ensure the performance results are not specific to a single network, the evaluation is repeated for a number of randomly instantiated network realizations with different initial potential connectivities. At the beginning of the learning process, all neurons of a stimulated subpopulation collectively fire in response to the external input. Non-stimulated neurons remain silent. As the connectivity is still immature at this point, no dAPs are triggered in postsynaptic neurons, and, hence, no predictions are generated. As a consequence, the prediction error, the false-negative rate and the number of active neurons (in stimulated populations) are at their maximum, and the false positive rate is zero (Fig. 2.8). During the first training episodes, the consistent collective firing of subsequently activated populations leads to the formation of mature connections as a result of the Hebbian structural plasticity. Upon reaching of a critical number of mature synapse, first dAPs (predictions) are generated in postsynaptic cells (in Fig. 2.8, this happens after about 10 learning episodes). As a consequence, the false negative rate decreases, and the stimulus responses become more sparse. At this early phase of the learning, the predictions of upcoming sequence elements are not yet context-specific (for sequence set I, non-sparse activity in “B” triggers a prediction in both “E” and “C”, irrespective of the context). Hence, the false-positive rate transiently increases. As the context specific connectivity is not consolidated at this point, more and more presynaptic subpopulations fail at triggering dAPs in their postsynaptic targets when they switch to sparse firing. Therefore, the false-positive rate decreases again, and the false-negative rate increases. In other words, there exists a negative feedback loop in the interim learning dynamics where the generation of predictions leads to an increase in sparsity which, in turn, causes prediction failures (and, hence, non-sparse firing). With an increasing number of training episodes, synaptic depression and homeostatic regulation increase context selectivity and thereby break this loop. Eventually, sparse firing of presynaptic populations is sufficient to reliably trigger predictions in their postsynaptic targets. For sequence set I, the total prediction error

becomes zero and the stimulus responses maximally sparse after about 30 training episodes (Fig. 2.8). For a time resolved visualization of the learning dynamics, see Video 7.1.

Up to this point, we illustrated the model’s sequence learning dynamics and performance for a simple set of two sequences (sequence set I). In the following, we assess the network’s sequence prediction performance for a more complex sequence set (II) composed of five high-order sequences (see [Task and training protocol](#)), each consisting of five elements. This sequence set is comparable to the one used in ([Hawkins and Ahmad, 2016](#)), but contains a larger amount of overlap between sequences. The overall pattern of the learning dynamics resembles the one reported for sequence set I (Fig. 2.9). The prediction error, the false-positive and false-negative rates as well as the sparsity measure vary more smoothly, and eventually converge at minimal levels after about 40 training episodes. To compare the spiking TM model with the original, non-spiking TM model, we repeat the experiment based on the simulation code provided in ([Hawkins and Ahmad, 2016](#)), see Table 7.2. With our parameterization, the learning rates λ_+ and λ_- of the spiking model are by a factor of about 10 smaller than in the original model. As a consequence, learning sequence set II with the original model converges faster than with the spiking model (compare black and gray curves in Fig. 2.9). The ratio in learning speeds, however, is not larger than about 2. Increasing the learning rates, i.e., the permanence increments, would speed up the learning process in the spiking model, but bears the risk that a large fraction of connections mature simultaneously. This would effectively overwrite the permanence heterogeneity which is essential to form context specific connectivity patterns (see section 2.3.1). As a result, the network performance would decrease. The original model avoids this problem by limiting the number of potentiated synapses in each update step (see “Plasticity dynamics” in [Network model](#)).

In sequence sets I and II, the maximum sequence order is 2 and 3, respectively. For the two sequences $\{E,N,D,I,J\}$ and $\{L,N,D,I,K\}$ in sequence set II, for example, predicting element “J” after activation of “I” requires remembering the element “E”, which occurred three steps back into the past. The TM model can cope with sequences of much higher order. Each sequence element in a particular context activates a specific pattern, i.e., a specific subset of neurons. The number of such patterns that can be learned is determined by the size of each subpopulation and the sparsity ([Ahmad and Hawkins, 2016](#)). In a sequence with repeating elements, such as $\{ABBBBCC\}$, the maximum order is limited by this number. Without repeating elements, the order could be arbitrarily high provided the number of subpopulations matches or exceeds the number of distinct characters. In Fig. 2.10, we demonstrate successful learning of two sequences $\{A,D,B,G,H,I,J,K,L,M,N,E\}$, $\{F,D,B,G,H,I,J,K,L,M,N,C\}$ of order 10.

2.3.3 *Dependence of prediction performance on the sequence speed*

The reformulation of the original TM model in terms of continuous-time dynamics allows us to ask questions related to timing aspects. Here, we investigate the sequence processing speed by identifying the range of inter-stimulus intervals ΔT that permit a successful prediction performance (Fig. 2.11). The timing of the external inputs affects the dynamics of the network in two respects. First, reliable predictions of sequence elements can only be made if the time interval ΔT between two consecutive stimulus presentations is such that the second input coincides with the somatic depolarization

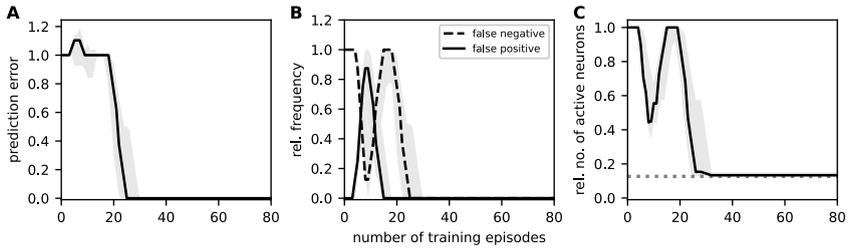


Figure 2.8: **Sequence prediction performance for sequence set I.** Dependence of the sequence prediction error (A), the false-positive and false-negative rates (B), and the number of active neurons relative to the subpopulation size (C) on the number of training episodes during repetitive stimulation with sequence set I (see [Task and training protocol](#)). Curves and error bands indicate the median as well as the 5% and 95% percentiles across an ensemble of 5 different network realizations, respectively. All prediction performance measures are calculated as a moving average over the last 4 training episodes. The dashed gray horizontal line in panel C depicts the target sparsity level $\rho/(Ln_E)$. Inter-stimulus interval $\Delta T = 40$ ms. See [Table 7.1](#) for remaining parameters.

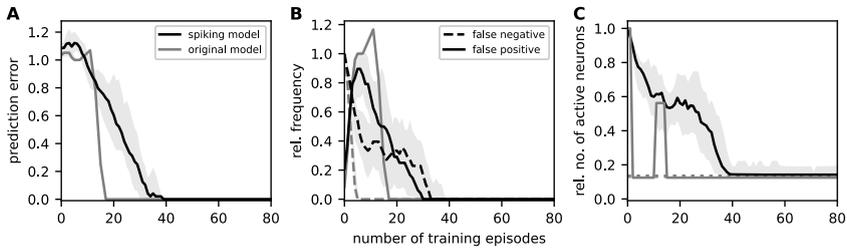


Figure 2.9: **Sequence prediction performance for sequence set II and comparison with original model.** Same figure arrangement, training and measurement protocol as in [Fig. 2.8](#). Data obtained during repetitive stimulation of the network with sequence set II (see [Task and training protocol](#)). Gray curves depict results obtained using the original (non-spiking) TM model from ([Hawkins and Ahmad, 2016](#)) with adapted parameters (see [Table 7.2](#)). The dashed gray horizontal line in panel C depicts the target sparsity level $\rho/(Ln_E)$.

caused by the dAP triggered by the first stimulus. Second, the formation of sequence specific connections by means of the spike-timing-dependent structural plasticity dynamics depends on ΔT .

If the external input does not coincide with the somatic dAP depolarization, i.e., if ΔT is too small or too large, the respective target population responds in a non-sparse, non-selective manner (mismatch signal; [Fig. 2.11C](#)), and in turn, generates false positives ([Fig. 2.11B](#)). For small ΔT , the external stimulus arrives before the dAP onset, i.e., before it is predicted. In consequence, the false negative rate is high. For large ΔT , the false negative rate remains low as the network is still generating predictions ([Fig. 2.11B](#)). The inter-stimulus interval ΔT in addition affects the formation of sequence specific connections due to the dependence of the plasticity dynamics on the timing of pre- and postsynaptic spikes, see [Eq \(2.1\)](#) and [Eq \(2.2\)](#). Larger ΔT results in smaller permanence increments, and thereby a slow-down of the learning process (red curve in [Fig. 2.11A](#)).

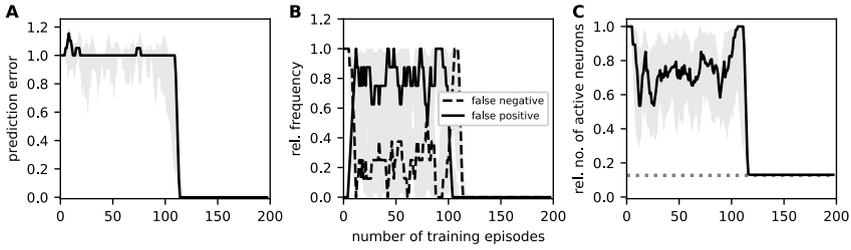


Figure 2.10: **Prediction performance for a sequence set with 10 overlapping elements.** Dependence of the sequence prediction error (A), the false positive frequency, the false negative frequency (B), and the number of active neurons relative to the subpopulation size (C) on the number of training episodes for a set of two sequences $s_1 = \{A, D, B, G, H, I, J, K, L, M, N, E\}$ and $s_2 = \{F, D, B, G, H, I, J, K, L, M, N, C\}$. Curves and error bands indicate the median as well as the 5% and 95% percentiles across 5 different network realizations, respectively. Inter-stimulus interval $\Delta T = 40$ ms. All prediction performance measures are calculated as a moving average over the last 4 training episodes. Parameters: $\Delta T = 40$ ms, $\lambda_+ = 0.39$, $\lambda_- = 0.0057$, $\lambda_h = 0.034$. See Table 7.1 for remaining parameters.

Taken together, the model predicts a range of optimal inter-stimulus interval ΔT (Fig. 2.11A). For our choice of network parameters, this range spans intervals between 10 ms and 75 ms. The lower bound depends primarily on the synaptic time constant τ_{EE} , the spike transmission delay d_{EE} , and the membrane time constant τ_m . The upper bound is mainly determined by the dAP plateau duration τ_{dAP} .

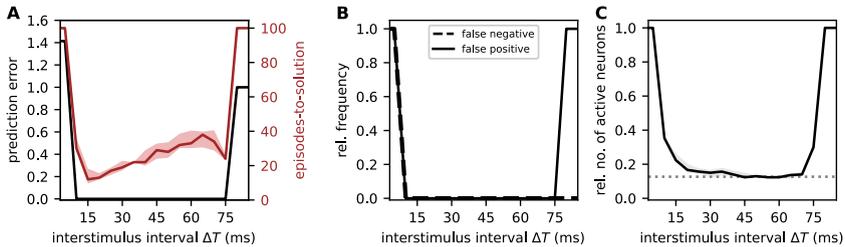


Figure 2.11: **Effect of sequence speed on network performance.** Dependence of the sequence prediction error, the learning speed (episodes-to-solution; A), the false-positive and false-negative rates (B), and the number of active neurons relative to the subpopulation size (C) on the inter-stimulus interval ΔT after 100 training episodes. Curves and error bands indicate the median as well as the 5% and 95% percentiles across an ensemble of 5 different network realizations, respectively. Same task and network as in Fig. 2.8.

2.3.4 Sequence replay

So far, we studied the network in the predictive mode, where the network is driven by external inputs and generates predictions of upcoming sequence elements. Another essential component of sequence processing is sequence replay, i.e., the autonomous generation of sequences in response to a cue signal (see [Task and training protocol](#)).

After successful learning, the network model presented in this study is easily configured into the replay mode by increasing the neuronal excitability, such that the somatic depolarization caused by a dAP alone makes the neuron fire a somatic spike. Here, this is implemented by lowering the somatic spike threshold θ_E of the excitatory neurons. In the biological system, this increase in excitability could, for example, be caused by the effect of neuromodulators (Atherton et al., 2015; Thomas, 2015), additional excitatory inputs from other brain regions implementing a top-down control, e.g. attention (Baluch and Itti, 2011; Noudoost et al., 2010), or propagating waves during sleep (Buzsáki, 2006; Grosmark et al., 2012).

The presentation of the first sequence element activates dAPs in the subpopulation corresponding to the expected next element in a previously learned sequence. Due to the reduced firing threshold in the replay mode, the somatic depolarization caused by these dAPs is sufficient to trigger somatic spikes during the rising phase of this depolarization. These spikes, in turn, activate the subsequent element. This process repeats, such that the network autonomously reactivates all sequence elements in the correct order, with the same context specificity and sparsity level as in the predictive mode (see Fig. 2.12A,B). The latency between the activation of subsequent sequence elements is determined by the spike transmission delay d_{EE} , the synaptic time constant τ_{EE} , the membrane time constant $\tau_{m,E}$, the synaptic weights $J_{EE,ij}$, the dAP current plateau amplitude I_{dAP} , and the somatic firing threshold θ_E . For sequences that can be successfully learned (see previous section), the time required for replaying the entire sequence is independent of the inter-stimulus interval ΔT employed during learning (Fig. 2.12C).

As shown in the previous section, sequences cannot be learned if the inter-stimulus interval ΔT is too small or too large. For small ΔT , connections between subpopulations corresponding to subsequent elements are strongly potentiated by the Hebbian plasticity due to the consistent firing of pre- and postsynaptic populations during the learning process. The network responses are, however, non-sparse, as the winner-take-all mechanism cannot be invoked during the learning (Fig. 2.11C). In the replay mode, sequences are therefore replayed in a non-sparse and non-context specific manner (left gray region in Fig. 2.12C). Similarly, connections between subsequent populations are slowly potentiated for very large ΔT . With sufficiently long learning, sequences can still be replayed in the right order, but the activity is non-sparse and therefore not context specific (right gray region in Fig. 2.12C).

2.4 DISCUSSION

2.4.1 Summary

In this work, we reformulate the Temporal Memory (TM) model (Hawkins and Ahmad, 2016) in terms of biophysical principles and parameters. We replace the original discrete-time neuronal and synaptic dynamics with continuous-time models with biologically interpretable parameters such as membrane and synaptic time constants and synaptic weights. We further substitute the original plasticity algorithm with a more biologically plausible mechanism, relying on a form of Hebbian structural plasticity, homeostatic control, and sparse random connectivity. Moreover, our model implements a winner-take-all dynamics based on lateral inhibition that is compatible with the continuous-time neuron and synapse models. We show that the revised TM

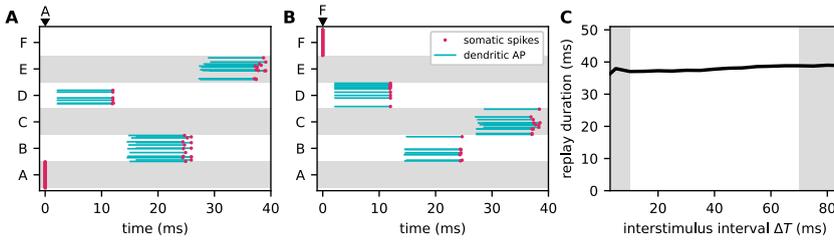


Figure 2.12: **Sequence replay dynamics and speed.** Autonomous replay of the sequences $\{A,D,B,E\}$ (A) and $\{F,D,B,C\}$ (B), initiated by stimulating the subpopulations “A” and “F”, respectively. Red dots and blue lines mark somatic spikes and dAP plateaus, respectively, for a fraction of neurons (30%) within each subpopulation. During learning, the inter-stimulus interval ΔT is set to 40 ms. C) Dependence of the sequence replay duration on the inter-stimulus interval ΔT during learning. Replay duration is measured as the difference between the mean firing times of the populations representing the first and last elements in a given sequence. Gray areas mark regions with low prediction performance (see section 2.3.3). Error bands represent the mean \pm standard deviation of the prediction error across 5 different network realizations. Same network and training set as in Fig. 2.8.

model supports successful learning and processing of high-order sequences with a performance similar to the one of the original model (Hawkins and Ahmad, 2016).

A new aspect that we investigated in the context of our work is sequence replay. After learning, the model is able to replay sequences in response to a cue signal. The duration of sequence replay is independent of the sequence speed during training, and determined by the intrinsic parameters of the network. In general, sequence replay is faster than the sequence presentation during learning, consistent with sequence compression and fast replay observed in hippocampus (Nádasy et al., 1999; Lee and Wilson, 2002; Davidson et al., 2009) and neocortex (Xu et al., 2012; Euston et al., 2007).

Finally, we identified the range of possible sequence speeds that guarantee a successful learning and prediction. Our model predicts an optimal range of processing speeds (inter-stimulus intervals) with lower and upper bounds constrained by neuronal and synaptic parameters (e.g., firing threshold, neuronal and synaptic time constants, coupling strengths, potentiation time constants). Within this range, the number of required training episodes is proportional to the inter-stimulus interval ΔT .

2.4.2 Relationship to other models

The model presented in this work constitutes a recurrent, randomly connected network of neurons with predefined stimulus preferences. The model learns sequences in an unsupervised manner using local learning rules. This is in essence similar to several other spiking neuronal network models for sequence learning (Lazar et al., 2009; Klampff and Maass, 2013; Klos et al., 2018; Maes et al., 2020; Cone and Shouval, 2021). The new components employed in this work are dendritic action potentials (dAPs) and Hebbian structural plasticity. We use structural plasticity to be as close as possible to the original model, and Hebbian forms of this are also known from the literature (Liao et al., 1995; Wu et al., 1996; Deger et al., 2012). However, preliminary results show that classical (non-structural) spike-timing-dependent plasticity (STDP) can yield

similar performance (see Fig. 7.1). Dendritic action potentials are instrumental for our model for two reasons. First, they effectively lower the threshold for coincidence detection and thereby permit a reliable and robust propagation of sparse activity (Jahnke et al., 2012; Breuer et al., 2014). In essence, our model bears similarities to the classical synfire chain (Abeles, 1991), one difference being that our mature network is not a simple feedforward network but has an abundance of recurrent connections. As shown in (Diesmann et al., 1999), a stable propagation of synchronous activity requires a minimal number of neurons in each synfire group. Without active dendrites, this minimal number is in the range of ~ 100 for plausible single-cell and synaptic parameters. In our (and in the original TM) model, coincidence detection happens in the dendrites. The number of presynaptic spikes needed to trigger a dAP is small, of the order of ~ 10 (Major et al., 2013; Mengual et al., 2020; Diesmann, 2002). This helps to reduce redundancy (only a small number of neurons needs to become active) and to increase the capacity of the network (the number of different patterns that can be learned is increased with pattern sparsity; (Ahmad and Hawkins, 2016)). Second, dAPs equip neurons with a third type of state (next to the quiescent and the firing state): the predictive state, i.e., a long lasting (~ 50 -200 ms) strong depolarization of the soma. Due to the prolonged depolarization of the soma, the inter-stimulus interval can be much larger than the synaptic time constants and delays. An additional benefit of dAPs, which is not exploited in the current version of our model, is that they equip individual neurons with more possible states if they comprise more than one dendritic branch. Each branch constitutes an independent pattern detector. The response of the soma may depend on the collective predictions in different dendritic branches. A single neuron could hence perform the types of computations that are usually assigned to multilayer perceptrons, i.e., small networks (Gidon et al., 2020; Poirazi et al., 2003).

Similar to a large class of other models in the literature, the TM network constitutes a recurrent network in the sense that the connectivity before and after learning forms loops at the subpopulation level. Recurrence in the immature connectivity permits the learning of arbitrary sequences without prior knowledge of the input data. In particular, recurrent connections enable the learning of sequences with repeating elements (such as in $\{A,B,B,C\}$ or $\{A,B,C,B\}$). Further, bidirectional connections between subpopulations are needed to learn sequences where pairs of elements occur in different orders (such as in $\{A,B,C\}$, $\{D,C,B\}$). Apart from providing the capability to learn sequences with all possible combinations of sequence elements, recurrent connections play no further functional role in the current version of the TM model. They may, however, become more important for future versions of the model enabling the learning of sequence timing and duration (see below).

Most of the existing models have been developed to replay learned sequences in response to a cue signal. The TM model can perform this type of pattern completion, too. In addition, it can act as a quiet, sparsely active observer of the world that becomes highly active only in the case of unforeseen, non-anticipated events. In this work, we didn't directly analyze the network's mismatch detection performance. However, this could be easily achieved by equipping each population with a "mismatch" neuron that fires if a certain fraction of neurons in the population fires (threshold detectors). In our model, predicted stimuli result in sparse firing due to inhibitory feedback (WTA). For unpredicted stimuli, this feedback is not effective, resulting in non-sparse firing indicating a mismatch. In (Schulz et al., 2021), a similar mechanism is employed to generate mismatch signals for novel stimuli. In this study, the strength of the inhibitory feedback needs to be learned by means of inhibitory synaptic plasticity.

In our model, the WTA mechanism is controlled by the predictions (dAPs) and implemented by static inhibitory connections. Furthermore, the model in (Schulz et al., 2021) can learn a set of elements, but not the order of these elements in the sequence.

In contrast to other sequence learning models (Maes et al., 2020; Cone and Shouval, 2021), our model is not able to learn an element specific timing and duration of sequence elements. The model in (Maes et al., 2020) relies on a clock network, which activates sequence elements in the correct order and with the correct timing. With this architecture, different sequences with different timings would require separate clock networks. Our model learns both sequence contents and order for a number of sequences without any auxiliary network. In an extension of our model, similar to (Maes et al., 2020) the timing of sequence element could be learned by additional plastic recurrent connections within each subpopulation. For the model by Maes et al. (2020) to serve as the basis of further research, in a recent Brief Report (Oberländer et al., 2022), we assessed the reproducibility of the model by giving a detailed description of its model description and pointed out missing elements, inconsistencies or errors in or between the original paper and its reference implementation. The model in (Cone and Shouval, 2021) can learn and recall higher-order sequences with limited history by means of an additional reservoir network with sparse readout. The TM model presents a more efficient way of learning and encoding the context in high-order sequences, without prior assignment of context specificity to individual neuron populations (Maes et al., 2020), and without additional network components (such as reservoir networks in (Cone and Shouval, 2021)).

An important sequence processing component that is not addressed in our work is the capability of identifying recurring sequences within a long stream of inputs. In the literature, this process is referred to as chunking, and constitutes a form of feature segmentation (Dehaene et al., 2015). Sequence chunking has been illustrated, for example, in (Asabuki and Fukai, 2020, 2021). Similar to our model, the network model in (Asabuki and Fukai, 2020, 2021) is composed of neurons with dendritic and somatic compartments, with the dendritic activity signaling a prediction of somatic spiking. Recurrent connections in their model improve the context specificity of neuronal responses, and thereby permit a context dependent feature segmentation. The model can learn high order sequence, but the history is limited. Although not explicitly tested here, our model is likely to be able to perform chunking if sequences are presented randomly across trials and without breaks. If the order of sequences is not systematic across trials, connections between neurons representing different sequences are not strengthened by spike-timing-dependent potentiation. Consecutive sequences are therefore not merged and remain distinct.

An earlier spiking neural network version of the HTM model has already been devised in (Billaudelle and Ahmad, 2015). It constitutes a proof-of-concept study demonstrating that the HTM model can be ported to an analog-digital neuromorphic hardware system. It is restricted to small simplistic sequences and does not address the biological plausibility of the TM model. In particular, it does not offer a solution to the question of how the model can perform online learning by known biological ingredients. Our study delivers a solution for this based on local plasticity rules and permits a direct implementation on a neuromorphic hardware system.

2.4.3 *Conclusion*

Our work demonstrates that the principal mechanisms underlying sequence learning, prediction, and replay in the TM model can be implemented using biologically plausible ingredients. By strengthening the link to biology, our implementation permits a more direct evaluation of the TM model predictions based on electrophysiological and behavioral data. Furthermore, this implementation allows for a direct mapping of the TM model on neuromorphic hardware systems.

PROBABILISTIC SEQUENCE REPLAY IN SPIKING
NEURAL NETWORKS

This chapter is based on the following preprint:

Bouhadjar, Y., Wouters, D. J., Diesmann, M., & Tetzlaff, T. (2022). Coherent noise enables probabilistic sequence replay in spiking neuronal networks. In press at PLOS Computational Biology, Preprint arXiv:2206.10538.

Author contributions: All authors contributed to the conception of the work. Younes Bouhadjar implemented the model, performed the simulations, analyzed and visualized the data, and wrote the first draft of the manuscript. All authors jointly reviewed the manuscript. Younes Bouhadjar was supervised by Tom Tetzlaff and Dirk J Wouters.

3.1 INTRODUCTION

The spiking Temporal Memory (spiking TM) model proposed in chapter 2 provides an energy-efficient sequence processing mechanism with high storage capacity by virtue of its sparse activity. It learns high-order sequences in an unsupervised, continuous manner using biological, local learning rules. After learning, the model successfully predicts upcoming sequence elements in a context-dependent manner and signals the occurrence of non-anticipated stimuli. Moreover, it can autonomously recall learned sequences in response to a cue stimulus.

In nature, cues are often incomplete or ambiguous, and it is not always clear what sequence to recall given the current context. Despite this ambiguity, we usually come to a clear decision on what sequence to recall. A key factor in decision-making is reward (Cohen et al., 2007; O’Doherty et al., 2017). In this regard, the optimal decision strategy is the one that maximizes the reward, and is hence referred to as the maximization or exploitation strategy. A number of studies demonstrate that decisions are often made in an apparently suboptimal manner, such as probability matching (Vulkan, 2000; Myers, 2014). In binary choice tasks, for example, where the probability of payoff is higher for one of the two possible choices, it appears most reasonable to always decide for this high-payoff option. Instead, however, humans and other animals often decide for each of the two choices with a probability that approximately matches the payoff probability. While this behavior appears unreasonable at first glance, it may in fact be optimal when taking into account previous (pre-experiment) experiences, such as prior knowledge of changing reward contingencies. In cases where the reward probability or amplitudes change over time, a more explorative behavior is beneficial (Cohen et al., 2007; Shanks et al., 2002). Previous studies suggest that decisions are not only determined by rewards, but also by the frequency of previously experienced input patterns (Bod et al., 2003; Hansen et al., 2012). Accordingly, suboptimal decision strategies may at least partly arise as a consequence of this additional influence of occurrence frequencies.

A number of previous studies propose neuronal network models of decision making in the face of ambiguous or incomplete stimuli. The majority of these models employ some form of intrinsic stochastic dynamics or uncorrelated noise to generate explorative behavior (Buesing et al., 2011; Legenstein and Maass, 2014; Hartmann et al., 2015; Neftci et al., 2016; Jordan et al., 2019). Noise has been introduced in the form of random or non-task-related synaptic background inputs (Jordan et al., 2019), or in the form of synaptic stochasticity (Neftci et al., 2016). An alternative solution is provided in the studies of Hartmann et al. (2015) and Dold et al. (2019), where the “noise” is generated by the complex but deterministic dynamics of the functional network itself, without any additional sources of stochasticity. In most models, the noise targeting different neurons or synapses is effectively uncorrelated. Supplying each element in a neuronal circuit with uncorrelated noise, however, does not necessarily lead to explorative dynamics: state variables arising from superpositions of many noisy, uncorrelated components become effectively deterministic as a result of noise averaging (Dold et al., 2019). The total input current of a neuron generated from superpositions of many synaptic inputs, for example, is largely unaffected by the variability in the individual synaptic responses. Similarly, in models where individual states are encoded by the activity of neuronal populations (Legenstein and Maass, 2014), the state representations become quasi deterministic if the single-neuron noise components are uncorrelated. Compensating this noise averaging effect by increasing

the noise amplitude appears to be an obvious strategy, but may be hard to realize by the biological system. Alternatively, noise averaging can be avoided by employing correlated noise. As an analogy, consider a particle in a still fluid: despite the constant bombardment by surrounding molecules, a large particle will hardly undergo any Brownian motion, because the momenta of the impinging molecules largely average out. Only if the molecules move in a coherent manner, such as in a turbulent fluid, they can have a substantial influence on the particle’s motion.

In biological neuronal networks, coherent noise may arise by different mechanisms: neighboring neurons typically receive inputs from partly overlapping presynaptic neuron populations. The synaptic input currents to these neurons are therefore correlated. In the literature, this type of correlation, which results from the anatomy of neurons and neuronal circuits, is referred to as shared-input correlation (Kriener et al., 2008; Tetzlaff et al., 2008). A second type of correlation in synaptic input currents arises from correlations in the presynaptic spiking activity (Renart et al., 2010; Tetzlaff et al., 2012; Helias et al., 2014). These dynamical correlations occur during stationary network states, or can be generated by different types of nonstationary activities, such as global oscillations in the population activity (Brunel and Hakim, 1999; Brunel, 2000) or traveling waves of activity propagating across the neuronal tissue (Sato et al., 2012; Takahashi et al., 2015; Roxin et al., 2005; Senk et al., 2020).

This study addresses the problem of sequential decision making in the face of ambiguity and the role of coherent noise in shaping decision strategies. We investigate how the spiking TM model recalls sequences in response to ambiguous cues in the presence of coherent noise, to what extent noise averaging can be overcome by increasing the noise amplitude, and how different recall strategies can be achieved by adjusting the noise characteristics. We further explore whether shared synaptic input and random stimulus locking to spatiotemporal oscillations can serve as appropriate, natural sources of coherent noise. In [Methods](#), we provide a detailed description of the task and the network model. For a first read, the main findings of the study can be understood without [Methods](#).

3.2 METHODS

In the following, we provide an overview of the task and the training protocol, the network model, and the task performance analysis. A detailed description of the model and parameter values is provided in [Tables 3.1](#) and [7.3](#).

3.2.1 *Learning protocol and task*

During learning, the network is continuously exposed to repeated presentations of an ensemble of S sequences $s_i = \{\zeta_{i1}, \zeta_{i2}, \dots, \zeta_{iC_i}\}$ ($C_i \in \mathbb{N}^+$, $i \in [1, \dots, S]$) of ordered discrete items ζ_{ij} . The order of the sequence elements within a given sequence represents the temporal order of the item occurrence. Each sequence s_i is presented with a relative frequency p_i in a given set of training data, where $\sum_{i=1}^S p_i = 1$. After successful learning, the presentation of some sequence element leads to a context dependent prediction of the subsequent stimulus. In case the prediction is wrong the network generates a mismatch signal. The network can also be configured into a replay mode where it autonomously replays learned sequences in response to a cue signal. We design the sequences such that they all start with the same two elements $\zeta_1 = \zeta_{11} = \zeta_{21} \dots \zeta_{i1}$ and $\zeta_2 = \zeta_{12} = \zeta_{22} \dots \zeta_{i2}$ ($i \in [1, \dots, S]$). As a consequence,

choosing the cue to be the first sequence element (ζ_1) results in an ambiguity. Here, we investigate the replay frequency of a given sequence s_i as a function of its training frequency p_i and study whether the network can choose between different replay strategies.

3.2.2 Network model

NETWORK STRUCTURE The network consists of a population \mathcal{E} of N_E excitatory (“E”) and a single inhibitory (“I”) neuron. The neurons in \mathcal{E} are randomly and recurrently connected, such that each neuron in \mathcal{E} receives K_{EE} excitatory inputs from other neurons in \mathcal{E} . Excitatory neurons are recurrently connected to the single inhibitory neuron. The excitatory population \mathcal{E} is subdivided into M non-overlapping subpopulations $\mathcal{M}_1, \dots, \mathcal{M}_M$, each of them containing neurons with identical stimulus preference (“receptive field”). Each subpopulation \mathcal{M}_k thereby represents a specific element within a sequence.

EXTERNAL INPUTS DURING LEARNING The network is driven by an ensemble $\mathcal{X} = \{x_1, \dots, x_{N_{\text{stim}}}\}$ of M external inputs. Each of these external inputs x_k represents a specific sequence element (“A”, “B”, ...), and feeds all neurons in the subpopulation \mathcal{M}_k with the corresponding stimulus preference. The occurrence of a specific sequence element $\zeta_{i,j}$ at time $t_{i,j}$ is modeled by a single spike $x_k(t) = \delta(t - t_{i,j})$ generated by the corresponding external source x_k . Subsequent sequence elements $\zeta_{i,j}$ and $\zeta_{i,j+1}$ within a sequence s_i are presented with an inter-stimulus interval $\Delta T = t_{i,j+1} - t_{i,j}$. Subsequent sequences s_i and s_{i+1} are separated in time by an inter-sequence time interval $\Delta T_{\text{seq}} = t_{i+1,1} - t_{i,C_i}$.

EXTERNAL INPUTS DURING REPLAY After learning the set of sequences S , we present a cue signal encoding for first sequence elements ζ_1 at time t^j for a number of trials N_t , where j denotes the trial id ($j \in [1, \dots, N_t]$). Subsequent cues are separated by an inter-trial interval $\Delta T_{\text{cue},j} = t^j - t^{j+1}$. In section 3.3.1, $\Delta T_{\text{cue},j}$ is constant and in section 3.3.4, $\Delta T_{\text{cue},j}$ is randomly and uniformly distributed between u_{min} ms and u_{max} .

During the replay, excitatory neurons are additionally driven by a background input implemented either in the form of asynchronous irregular synaptic bombardment (see section 3.3.1) or oscillatory inputs (see section 3.3.4). The first is realized using ensembles of excitatory and inhibitory spike sources \mathcal{Q}_k and \mathcal{V}_k ($k \in [1, \dots, M]$), each composed of n elements. Each source is an independent realization of a Poisson point process with a rate ν . Excitatory neurons in the same subpopulation \mathcal{M}_k receive K_{EQ} inputs with weight J_{EQ} from the ensemble \mathcal{Q}_k and K_{EV} inputs with weights $-J_{EV}$ from the ensemble \mathcal{V}_k . Spikes from \mathcal{Q}_k and \mathcal{V}_k give rise to a jump in the synaptic current of the postsynaptic cell followed by an exponential decay with a time constant τ_{EQ} and τ_{EV} , respectively. The time average input current of a neuron i is

$$\mu_i = 0 \quad (3.1)$$

and the variance across time

$$\sigma_i^2 = \frac{\tau_B}{2} \sum_{k=0}^{2K} J_{ik}^2 \nu, \quad (3.2)$$

where $J = J_{EQ} = -J_{EV}$, $\tau_B = \tau_{EQ} = \tau_{EV}$, and $K = K_{EQ} = K_{EV}$. Given that the populations of background sources are of a finite size, there is a probability that two

neurons in the same subpopulation to pick a certain number of identical sources, this gives rise to the so called shared input correlation. The correlation coefficient in the input current is governed by

$$c = \frac{K}{n}. \quad (3.3)$$

With this relationship, we can now vary the correlation coefficient by fixing K and varying n . For the special case where c is supposed to be zero, we assume that each neuron has its own set of independent Poissonian sources. The second type of background input is implemented using an ensemble \mathcal{G} of M sinusoidal current generators g_k , each with a frequency f , amplitude a , and a phase ϕ_k ($k \in [1, \dots, M]$). Excitatory neurons in the same subpopulation M_k receive oscillatory inputs from the same source g_k .

NEURON AND SYNAPSE MODEL For all types of neurons, the temporal evolution of the membrane potential is given by the leaky integrate-and-fire model Eq (3.6). The total synaptic input current of excitatory neurons is composed of currents in distal dendritic branches, inhibitory currents, and currents from external sources. The inhibitory neuron receives only inputs from excitatory neurons. Individual spikes arriving at dendritic branches evoke alpha-shaped postsynaptic currents, see Eq (3.8). The dendritic current includes an additional nonlinearity describing the generation of dendritic action potentials (dAPs): if the dendritic current I_{ED} exceeds a threshold θ_{dAP} , it is instantly set to the dAP plateau current I_{dAP} , and clamped to this value for a period of duration τ_{dAP} , see Eq (3.12). This plateau current leads to a long lasting depolarization of the soma. Inhibitory inputs to excitatory neurons as well as excitatory inputs to inhibitory neurons trigger exponential postsynaptic currents, see (3.9–3.10). The weights J_{IE} of excitatory synapses on inhibitory neurons are chosen such that the collective firing of a subset of ρ excitatory neurons in the corresponding subpopulation causes the inhibitory neuron to fire. The weights J_{EI} of inhibitory synapses on excitatory neurons are strong such that each inhibitory spike prevents all excitatory neurons in the network that have not generated a spike yet from firing. External inputs are composed of currents resulting from the presentation of the sequence elements or currents from background inputs (see Inputs in Table 3.1). All synaptic time constants, delays, and weights are connection-type specific.

PLASTICITY Only excitatory to excitatory (EE) synapses are plastic. All other connections are static. The dynamics of the EE synaptic weights J_{ij} evolve according to a combination of an additive spike-timing-dependent plasticity (STDP) rule (Morrison et al., 2008) and a homeostatic component (Abbott and Nelson, 2000; Tetzlaff et al., 2011). During the replay mode, the plasticity is disabled and the EE weights are kept constant (see Table 3.1 for details about the plasticity).

NETWORK REALIZATIONS AND INITIAL CONDITIONS. For every network realization, the connectivity and the initial weights are drawn randomly and independently. All other parameters are identical for different network realizations. The initial values of all state variables are given in Table 3.1 and Table 7.3.

SIMULATION DETAILS The network simulations are performed in the neural simulator NEST (Gewaltig and Diesmann, 2007) under version 3.0 (Hahne et al., 2021). The differential equations and state transitions defining the excitatory neuron dynamics are expressed in the domain specific language NESTML (Plotnikov et al., 2016;

(Nagendra Babu et al., 2021) which generates the required C++ code for the dynamic loading into NEST. Network states are synchronously updated using exact integration of the system dynamics on a discrete-time grid with step size Δt (Rotter and Diesmann, 1999). The full source code for the implementation with a list of other software requirements is available at Zenodo: <https://doi.org/10.5281/zenodo.6378376>.

3.2.3 Task performance measures

Consider the set $\mathcal{S} = \{s_1, s_2, \dots, s_S\}$ of S sequences learned by the network. Let

$$\mathcal{P} = \{\emptyset, \{s_1\}, \{s_2\}, \dots, \{s_1, s_2\}, \{s_1, s_3\}, \dots, \mathcal{S}\}$$

denote the power set of \mathcal{S} , i.e., the set of all subsets of \mathcal{S} , including the empty set and \mathcal{S} itself. We define the relative replay frequency $f_{\mathcal{P}_k}$ of each subset $\mathcal{P}_k \in \mathcal{P}$ of sequences as the normalized number of exclusive replays of this subset \mathcal{P}_k , such that

$$\sum_{\mathcal{P}_k} f_{\mathcal{P}_k} = 1. \quad (3.4)$$

For two sequences s_1 and s_2 , for example, we monitor the four different replay frequencies f_{\emptyset} (no sequence is replayed), $f_{\{s_1\}}$ (only s_1 is replayed), $f_{\{s_2\}}$ (only s_2 is replayed), and $f_{\{s_1, s_2\}}$ (both s_1 and s_2 are replayed). In this work, we refer to f_{\emptyset} as the “failure rate”. Simultaneous replay of both sequences ($f_{\{s_1, s_2\}}$) refers to cases where the network fails at coming to a unique decision.

3.2.4 Model tables

Summary		
Populations	excitatory neurons (\mathcal{E}), inhibitory neurons (\mathcal{I}), external spike sources (\mathcal{X}), background inputs in the form Poissonian sources (\mathcal{Q}_k and \mathcal{V}_k) or sinusoidal current generators (\mathcal{G}). \mathcal{E} composed of M disjoint subpopulations \mathcal{M}_k and ($k = 1, \dots, M$)	
Connectivity	<ul style="list-style-type: none"> • sparse random connectivity between excitatory neurons (plastic) • local recurrent connectivity between excitatory and inhibitory neurons (static) 	
Neuron model	<ul style="list-style-type: none"> • excitatory neurons: Leaky Integrate-and-Fire (LIF) with nonlinear input integration (dendritic action potentials) • inhibitory neurons: Leaky Integrate-and-Fire (LIF) 	
Synapse model	exponential or alpha-shaped postsynaptic currents (PSCs)	
Plasticity	homeostatic spike-timing-dependent plasticity in excitatory-to-excitatory connections	
Populations		
Name	Elements	Size
$\mathcal{E} = \cup_{i=k}^M \mathcal{M}_k$	excitatory (E) neurons	N_E
\mathcal{I}	inhibitory (I) neurons	N_I
\mathcal{M}_k	excitatory neurons in subpopulation k , $\mathcal{M}_k \cap \mathcal{M}_l = \emptyset$ ($\forall k \neq l \in [1, M]$)	n_E
\mathcal{Q}_k	excitatory Poisson generators, $\mathcal{Q}_k \cap \mathcal{Q}_l = \emptyset$ ($\forall k \neq l \in [1, M]$)	n
\mathcal{V}_k	inhibitory Poisson generators, $\mathcal{V}_k \cap \mathcal{V}_l = \emptyset$ ($\forall k \neq l \in [1, M]$)	n
$\mathcal{X} = \{x_1, \dots, x_M\}$	external spike sources	M
$\mathcal{G} = \{g_1, \dots, g_M\}$	sinusoidal current generators	M
Connectivity		
Source population	Target population	Pattern
\mathcal{E}	\mathcal{E}	random; fixed in-degrees $K_i = K_{EE}$, delays $d_{ij} = d_{EE}$, and synaptic time constants $\tau_{ij} = \tau_{EE}$, plastic synaptic weights J_{ij} ($\forall i \in \mathcal{E}, \forall j \in \mathcal{E}$; “EE connections”)
\mathcal{E}	\mathcal{I}	all-to-all; fixed delays $d_{ij} = d_{IE}$, synaptic time constants $\tau_{ij} = \tau_{IE}$, and weights $J_{ij} = J_{IE}$ ($\forall i \in \mathcal{E}, \forall j \in \mathcal{I}$; “IE connections”)
\mathcal{I}	\mathcal{E}	all-to-all; fixed delays $d_{ij} = d_{EI}$, synaptic time constants $\tau_{ij} = \tau_{EI}$, and weights $J_{ij} = J_{EI}$ ($\forall i \in \mathcal{I}, \forall j \in \mathcal{E}$; “EI connections”)
\mathcal{I}	\mathcal{I}	none (“II connections”)
\mathcal{Q}_k	\mathcal{M}_k	random; fixed in-degrees $K_i = K_{EQ}$, delays $d_{ij} = d_{EQ}$, synaptic time constants $\tau_{ij} = \tau_{EQ}$, and weights $J_{ij} \in \{0, J_{EQ}\}$ ($\forall i \in \mathcal{M}_k, j \in \mathcal{Q}_k, \forall k \in [1, M]$; “EQ connections”)
\mathcal{V}_k	\mathcal{M}_k	random; fixed in-degrees $K_i = K_{EV}$, delays $d_{ij} = d_{EV}$, synaptic time constants $\tau_{ij} = \tau_{EV}$, and weights $J_{ij} \in \{0, J_{EV}\}$ ($\forall i \in \mathcal{M}_k, j \in \mathcal{V}_k, \forall k \in [1, M]$; “EV connections”)
$\mathcal{X}_k = x_k$	\mathcal{M}_k	one-to-all; fixed delays $d_{ij} = d_{EX}$, synaptic time constants $\tau_{ij} = \tau_{EX}$, and weights $J_{ij} = J_{EX}$ ($\forall i \in \mathcal{M}_k, j \in \mathcal{X}_k, \forall k \in [1, M]$; “EX connections”)
$\mathcal{G}_k = g_k$	\mathcal{M}_k	one-to-all; fixed synaptic weights $J_{ij} = J_{EG}$ ($\forall i \in \mathcal{M}_k, j \in \mathcal{G}_k, \forall k \in [1, M]$; “EG connections”)
no self-connections (“autapses”), no multiple connections (“multapses”)		
all unmentioned connections $\mathcal{I} \rightarrow \mathcal{I}$, $\mathcal{V}_k \rightarrow \mathcal{V}_k$, $\mathcal{Q}_k \rightarrow \mathcal{Q}_k \dots \mathcal{X}_k \rightarrow \mathcal{M}_l$ ($\forall k \neq l$) are absent		

Continued next page

Neuron and synapse	
Neuron	
Type	Leaky Integrate-and-Fire (LIF) dynamics
Description	<p>dynamics of membrane potential $V_i(t)$ and spiking activity $s_i(t)$ of neuron i:</p> <ul style="list-style-type: none"> emission of the kth spike of neuron i at time t_i^k if $V_i(t_i^k) \geq \theta_i \quad (3.5)$ with somatic spike threshold θ_i spike train: $s_i(t) = \sum_k \delta(t - t_i^k)$ reset and refractoriness: $V_i(t) = V_r \quad \forall k, \forall t \in (t_i^k, t_i^k + \tau_{\text{ref},i}]$ with refractory time $\tau_{\text{ref},i}$ and reset potential V_r subthreshold dynamics: $\tau_{m,i} \dot{V}_i(t) = -V_i(t) + R_{m,i} I_i(t) \quad (3.6)$ with membrane resistance $R_{m,i} = \frac{\tau_{m,i}}{C_{m,i}}$, membrane time constant $\tau_{m,i}$, and total synaptic input current $I_i(t)$ (see Synapse) excitatory neurons: $\tau_{m,i} = \tau_{m,E}$, $C_{m,i} = C_m$, $\theta_i = \theta_E$, $\tau_{\text{ref},i} = \tau_{\text{ref},E}$ ($\forall i \in \mathcal{E}$) inhibitory neurons: $\tau_{m,i} = \tau_{m,I}$, $C_{m,i} = C_m$, $\theta_i = \theta_I$, $\tau_{\text{ref},i} = \tau_{\text{ref},I}$ ($\forall i \in \mathcal{I}$)
Synapse	
Type	continuous, exponential, or alpha-shaped postsynaptic currents (PSCs)
Description	<ul style="list-style-type: none"> total synaptic input current $\begin{aligned} \text{excitatory neurons: } I_i(t) &= I_{\text{ED},i}(t) + I_{\text{EX},i}(t) + I_{\text{EI},i}(t), \quad \forall i \in \mathcal{E} \\ \text{inhibitory neurons: } I_i(t) &= I_{\text{IE},i}(t), \quad \forall i \in \mathcal{I} \end{aligned} \quad (3.7)$ with dendritic, inhibitory, excitatory, and external input currents $I_{\text{ED},i}(t)$, $I_{\text{EI},i}(t)$, $I_{\text{IE},i}(t)$, $I_{\text{EX},i}(t)$ evolving according to $I_{\text{ED},i}(t) = \sum_{j \in \mathcal{E}} (\alpha_{ij} * s_j)(t - d_{ij}) \quad (3.8)$ with $\alpha_{ij}(t) = J_{ij} \frac{e}{\tau_{\text{ED}}} t e^{-t/\tau_{\text{ED}}} \Theta(t)$ and $\Theta(t) = \begin{cases} 1 & t \geq 0 \\ 0 & \text{else} \end{cases}$ $\tau_{\text{EI}} \dot{I}_{\text{EI},i} = -I_{\text{EI},i}(t) + \sum_{j \in \mathcal{I}} J_{ij} s_j(t - d_{ij}) \quad (3.9)$ $\tau_{\text{IE}} \dot{I}_{\text{IE},i} = -I_{\text{IE},i}(t) + \sum_{j \in \mathcal{E}} J_{ij} s_j(t - d_{ij}) \quad (3.10)$ $I_{\text{EX},i}(t) = I_{\text{S},i}(t) + I_{\text{B},i}(t) \quad (3.11)$ where $I_{\text{S},i}(t)$ and $I_{\text{B},i}(t)$ are the stimulus and the background input, respectively (see input). suprathreshold dynamics of dendritic currents (dAP generation): <ul style="list-style-type: none"> emission of kth dAP of neuron i at time $t_{\text{dAP},i}^k$ if $I_{\text{ED},i}(t_{\text{dAP},i}^k) \geq \theta_{\text{dAP}}$ dAP current plateau: $I_{\text{ED},i}(t) = I_{\text{dAP}} \quad \forall k, \forall t \in (t_{\text{dAP},i}^k, t_{\text{dAP},i}^k + \tau_{\text{dAP}}) \quad (3.12)$ with dAP current plateau amplitude I_{dAP}, dAP current duration τ_{dAP}, and dAP activation threshold θ_{dAP} reset: $I_{\text{ED},i}(t_{\text{dAP},i}^k + \tau_{\text{dAP}}) = 0$ ($\forall k$) reset and refractoriness in response to emission of lth somatic spike of neuron i at time t_i^l: $I_{\text{ED},i}(t) = 0 \quad \forall l, \forall t \in (t_i^l, t_i^l + \tau_{\text{ref},i}) \quad (3.13)$

	<p>– reset of $I_{ED,i}$ in case of a strong inhibitory current:</p> $I_{ED,i}(t_i^k) = 0, \text{ if } I_{EI,i}(t_i^k) < I_\theta, \quad (3.14)$ <p>where I_θ is the reset dAP current.</p>
Plasticity	
Type	spike-timing-dependent plasticity and dAP-rate homeostasis
EE synapses	<ul style="list-style-type: none"> • dynamics of synaptic weight $J_{ij}(t)$ (EE connections): $J_{\max}^{-1} \frac{dJ_{ij}}{dt} = \lambda_+ \sum_{\{t_i^*\}'} x_j(t) \delta(t - [t_i^* + d_{EE}]) - \lambda_- y_i J_{ij} \sum_{\{t_j^*\}} \delta(t - t_j^*) + \lambda_h \sum_{\{t_i^*\}'} (z^* - z_i(t)) \delta(t - t_i^*).$ <p>with</p> <ul style="list-style-type: none"> – list of presynaptic spike times $\{t_j^*\}$, – list of postsynaptic spike times $\{t_i^*\}' = \{t_i^* \exists t_j^* \Delta t_{\min} < t_i^* - t_j^* + d_{EE} < \Delta t_{\max}\}$, – maximum weight J_{\max}, potentiation and depression rates λ_+, λ_-, homeostasis rate λ_h, delay d_{EE}, depression decrement y_i, – spike trace of postsynaptic neuron i, evolving according to $\frac{dx_j}{dt} = -\tau_+^{-1} x_j(t) + \sum_{t_j^*} \delta(t - t_j^*)$ <p>with presynaptic spike times t_j^* and potentiation time constant τ_+,</p> <ul style="list-style-type: none"> – dAP trace of postsynaptic neuron i, evolving according to $\frac{dz_i}{dt} = -\tau_h^{-1} z_i(t) + \sum_k \delta(t - t_{dAP,i}^k)$ <p>with onset time $t_{dAP,i}^k$ of the kth dAP, homeostasis time constant τ_h, and</p> <ul style="list-style-type: none"> – target dAP activity z^*
all other synapses	non-plastic

Continued next page

Input

- prediction mode
 - stimulus
 - * repetitive stimulation of the network using the same set $\mathcal{S} = \{s_1, \dots, s_S\}$ of sequences $s_i = \{\zeta_{i,1}, \zeta_{i,2}, \dots, \zeta_{i,C_i}\}$ of ordered discrete items $\zeta_{i,j}$ with number of sequences S and length C_i of i th sequence
 - * presentation of sequence element $\zeta_{i,j}$ at time $t_{i,j}$ modeled by a single spike $x_k(t) = \delta(t - t_{i,j})$ generated by the corresponding external source x_k
 - * generated current as a response to the presentation of the sequence elements:

$$\tau_S \dot{I}_{S,i} = -I_{S,i}(t) + \sum_{j \in \mathcal{X}} J_{i,j} x_j(t - d_{ij}) \quad (3.15)$$
 - * inter-stimulus interval $\Delta T = t_{i,j+1} - t_{i,j}$ between subsequent sequence elements $\zeta_{i,j}$ and $\zeta_{i,j+1}$ within a sequence s_i
 - * inter-sequence time interval $\Delta T_{\text{seq}} = t_{i+1,1} - t_{i,C_i}$ between subsequent sequences s_i and s_{i+1}
 - * example sequence sets:
 - sequence set I: $\mathcal{S} = \{\{A,F,B,D\} \text{ and } \{A,F,C,E\}\}$
 - sequence set II: $\mathcal{S} = \{\{A,F,C,E\}, \{A,F,B,D\}, \{A,F,G,H\}\}$

- replay mode
 - stimulus
 - * presentation of a cue encoding for first sequence elements ζ_1 at time t^j , where j denotes the trial number ($j \in [1, \dots, N_t]$).
 - * inter-trial interval $\Delta T_{\text{cue}} = t^{j+1} - t^j$
 - background input
 - * in the form of correlated Poissonian inputs

$$\tau_B \dot{I}_{B,i}(t) = -I_{B,i}(t) + \sum_{j \in \mathcal{Q}} J_{i,j} s_j(t - d) + \sum_{j \in \mathcal{V}} J_{i,j} s_j(t - d) \quad (3.16)$$

with Poissonian spike trains $s_j(t)$ of rate ν , synaptic weight $J_{i,j} \in \{0, J\}$ where $J = J_{\text{EQ}} = -J_{\text{EV}}$, synaptic time constant $\tau_B = \tau_{\text{EQ}} = \tau_{\text{EV}}$, and delay $d = d_{\text{EQ}} = d_{\text{EV}}$

- the variance of $I_{B,i}(t)$ across time:

$$\sigma^2 = \text{Var}(I_{B,i}(t)) = J^2 K \nu \tau_B, \quad (3.17)$$

where $K = K_{\text{EQ}} = K_{\text{EV}}$ is the number of either excitatory or inhibitory Poissonian input per excitatory neuron

- the correlation coefficient of $I_{B,i}(t)$ and $I_{B,j}(t)$ across time:

$$c = \frac{\text{Cov}(I_{B,i}(t), I_{B,j}(t))}{\sigma^2} = \begin{cases} 0 & \forall i \in \mathcal{M}_k, \forall j \in \mathcal{M}_l (\forall k \neq l) \\ \frac{K}{n} & \forall i \in \mathcal{M}_k, \forall j \in \mathcal{M}_l (\forall k = l), \end{cases} \quad (3.18)$$

where n is the number of either excitatory or inhibitory Poissonian sources (see Connectivity)

- * or oscillatory current

$$I_{B,i}(t) = J_{\text{EG}} \cdot a \cdot \sin(2\pi f t + \phi_i) \quad (3.19)$$

with amplitude a , frequency f , and population specific phase $\phi_i = \phi_k (\forall i \in \mathcal{M}_k)$

Output

- somatic spike times $\{t_i^k | \forall i \in \mathcal{E}, k = 1, 2, \dots\}$
- dendritic currents $I_{\text{ED},i}(t)$ ($\forall i \in \mathcal{E}$)

Initial conditions and network realizations
<ul style="list-style-type: none"> • membrane potentials: $V_i(0) = V_r$ ($\forall i \in \mathcal{E} \cup \mathcal{I}$) • dendritic currents: $I_{ED,i}(0) = 0$ ($\forall i \in \mathcal{E}$) • external currents: $I_{S,i}(0) = 0$ and $I_{B,i}(0) = 0$ ($\forall i \in \mathcal{E}$) • inhibitory currents: $I_{EI,i}(0) = 0$ ($\forall i \in \mathcal{E}$) • excitatory currents: $I_{IE,i}(0) = 0$ ($\forall i \in \mathcal{I}$) • synaptic weights: $J_{ij}(0) \sim \mathcal{U}(J_{0,\min}, J_{0,\max})$ (uniform distribution; $\forall i, j \in \mathcal{E}$) • spike traces: $x_i(0) = 0$ ($\forall i \in \mathcal{E}$) • dAP traces: $z_i(0) = 0$ ($\forall i \in \mathcal{E}$) • initial weights randomly and independently drawn for each network realization
Simulation details
<ul style="list-style-type: none"> • network simulations performed in NEST (Gewaltig and Diesmann, 2007) version 3.0 (Hahne et al., 2021) • definition of excitatory neuron model using NESTML (Plotnikov et al., 2016; Nagendra Babu et al., 2021) • synchronous update using exact integration of system dynamics on discrete-time grid with step size Δt (Rotter and Diesmann, 1999) • source code underlying this study: https://doi.org/10.5281/zenodo.6378376

Table 3.1: Description of the network model. Parameter values are given in Table 7.3.

3.3 RESULTS

3.3.1 A spiking neural network recalls sequences in response to ambiguous cues

In this section, we provide a brief overview of the model and the task, illustrate how the network learns overlapping sequences occurring with different frequencies during the training, and show how these occurrence frequencies are encoded in the network. We then study the network responses to ambiguous cues and the influence of the occurrence frequencies on the recall behavior in the absence or presence of noise.

The model consists of a randomly and sparsely connected network of N_E excitatory neurons (population \mathcal{E}) and a single inhibitory neuron (Fig. 3.1A). Each excitatory neuron receives K_{EE} excitatory inputs from other randomly chosen neurons in \mathcal{E} . Excitatory neurons are subdivided into M subpopulations, each containing neurons with identical stimulus preference: in the absence of any additional connections, all neurons in a given subpopulation fire a spike upon presentation of a specific sequence element. The inhibitory neuron is recurrently connected to the excitatory neurons. In contrast to chapter 2 where each excitatory subpopulation is equipped with its own inhibitory neuron, we here use a single inhibitory neuron to implement competition between the subpopulations of excitatory neurons. The network is driven by external inputs, each representing a specific sequence element (“A”, “B”, ...), and feeds all neurons in the subpopulation \mathcal{M}_k with the corresponding stimulus preference. The model of the neurons and synaptic currents is similar to the one introduced in chapter 2. Synapses between excitatory neurons are plastic and subject to spike-timing-dependent plasticity and homeostatic control. Details on the network model are given in [Methods](#).

During the learning, the network is exposed to repeated presentations of two sequences $\{A, F, B, D\}$ and $\{A, F, C, E\}$, where the first sequence is shown with a relative frequency p and the second with $1 - p$ (e.g., $p = 0.2$ in Fig. 3.2A). In the following, we refer to $\{A, F, B, D\}$ as sequence 1 and to $\{A, F, C, E\}$ as sequence 2. Before learning, presenting a sequence element causes all neurons in the respective subpopulation to fire. During the learning process, the repetitive sequential presentation of sequence elements increases the strength of connections between the corresponding subpopulations to a point where the activation of a certain subpopulation by an external input generates dAPs in a specific subset of neurons in the subpopulation representing the subsequent element. The generation of the dAPs results in a long-lasting depolarization (~ 50 - 500 ms) of the soma. We refer to neurons that generate a dAP as predictive neurons. When receiving an external input, predictive neurons fire earlier as compared to non-predictive neurons. When enough predictive neurons are present within a certain subpopulation, their advanced spike initiates a fast and strong inhibitory feedback to all excitatory neurons, ultimately suppressing the firing of non predictive neurons. After learning, the model develops specific subnetworks representing the learned sequences (Fig. 3.1B), such that the presentation of a sequence element leads to a context dependent prediction of the subsequent element. As a result of Hebbian learning, the synaptic weights in the subnetwork corresponding to the most frequent sequence during learning are stronger than those for the less frequent sequence (Fig. 3.1B, Fig. 3.3A, Fig. 3.4A). In the prediction mode, this asymmetry in synaptic weights plays no role. For ambiguous stimuli, all potential outcomes are predicted, i.e., the network predicts both “C” and “E” simultaneously in response to stimuli “A” and “F”, irrespective of the training frequencies.

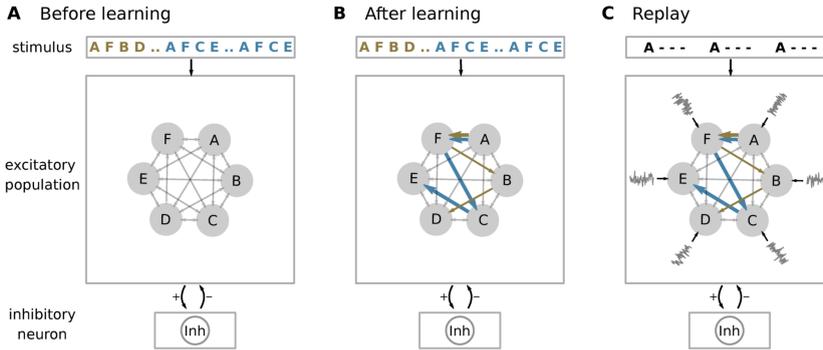


Figure 3.1: **Network structure.** **A)** The architecture constitutes a recurrent network of subpopulations of excitatory neurons (filled gray circles) and a single inhibitory neuron (Inh). Each excitatory subpopulation contains neurons with identical stimulus preferences. Excitatory neurons are stimulated by external sources providing sequence-element specific inputs “A”, “F”, “B”, etc. Connections between and within the excitatory subpopulations are random and sparse. The inhibitory neuron is recurrently connected to all excitatory neurons. In the depicted example, the network is repetitively presented with two sequences $\{A,F,C,E\}$ (brown) and $\{A,F,B,D\}$ (blue) during learning. The sequence $\{A,F,C,E\}$ occurs twice as often as $\{A,F,B,D\}$. **B)** After learning, the networks forms sequence specific subnetworks (blue and brown arrows representing $\{A,F,B,D\}$ and $\{A,F,C,E\}$, respectively). The connections between subpopulations representing the sequence shown more often are stronger (thick arrows). **C)** During the replay mode, the network is presented with a cue stimulus representing the first sequence element “A”. In addition, the excitatory subpopulations receive input from distinct sources of background noise (gray traces).

The model can be configured into a replay mode, where the network autonomously replays learned sequences in response to a cue stimulus. This is achieved by changing the excitability of the neurons such that the activation of a dAP alone can cause the neurons to fire. In the replay mode, we present ambiguous cues and study whether the network can replay sequences following different strategies (Fig. 3.2B). We refer to the “maximum probability” strategy (Fig. 3.2B, left) as the case where the network exclusively replays the sequence with the highest occurrence frequency. “Probability matching” is referred to when the replay frequency of a sequence matches its occurrence frequency during training (Fig. 3.2B, middle). We call the strategy “full exploration” when all sequences are randomly replayed with the same frequency, irrespective of the occurrence frequency during training (Fig. 3.2B, right). In Fig. 3.3, we illustrate the network’s decision behavior by presenting the ambiguous cue stimulus “A” three times. In the absence of noise, the network adopts the maximum probability strategy (Fig. 3.3B): as a result of the higher weights between the neurons representing the more frequent sequence, the dAPs get activated earlier in these neurons, which advances their somatic firing times with respect to the neurons of the less frequent sequence. This advanced response time quickly activates the inhibitory neuron, which suppresses the activity of the neurons representing the less frequent sequence.

To assess the replay performance quantitatively, we present the cue stimulus “A” for N_t trials and examine the replay frequency of the two sequences $\{A,F,B,D\}$ and

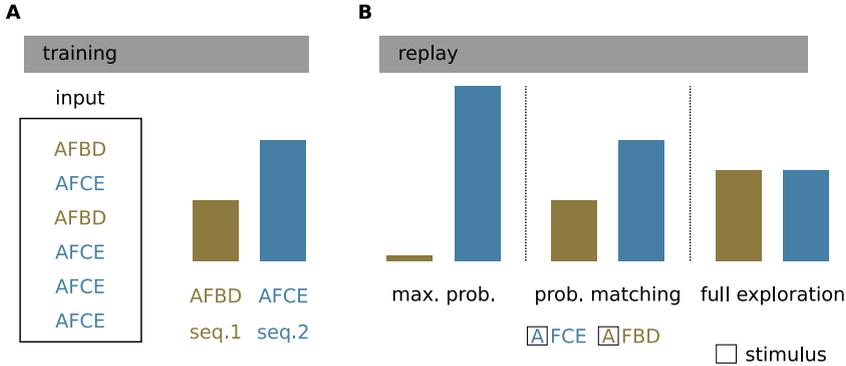


Figure 3.2: **Task.** **A**) During learning, the model is exposed to two (or more) competing sequences with different frequencies. Here, sequence 1 ($\{A,F,C,E\}$; blue) is shown twice as often as sequence 2 ($\{A,F,B,D\}$; brown). The respective normalized occurrence frequencies $1/3$ and $2/3$ are depicted by the histogram. **B**) During replay, the network autonomously recalls the sequences in response to an ambiguous cue signal (first sequence element; black squares) according to different strategies. Maximum probability (max-prob): Only the sequence with the highest occurrence frequency during training is replayed. Probability matching (prob. matching): The replay frequency of a sequence matches its occurrence frequency during training. Full exploration: All sequences are randomly replayed with the same frequency, irrespective of the occurrence frequency during training. Histograms represent the replay frequencies.

$\{A,F,C,E\}$ as a function of their occurrence frequencies during training. We define the sequences $\{A,F,B,D\}$ or $\{A,F,C,E\}$ to be successfully replayed if more than 50% of the neurons in the last subpopulations “E” or “D” have fired, respectively. In the absence of noise, the network replays only the sequence with the highest occurrence frequency p (Fig. 3.4E). For $p < 0.5$, the sequence $\{A,F,C,E\}$ is the only one replayed in all trials, and for $p > 0.5$ the sequence $\{A,F,B,D\}$ becomes the dominant one. To understand this behavior, we inspect the response latencies $t^{B/C}$ of the subpopulations “B” and “C” as a function of the occurrence frequencies of the sequences $\{A,F,B,D\}$ and $\{A,F,C,E\}$ (Fig. 3.4B). The mean of the response latency t^B or t^C is smaller for the subpopulation participating in the sequence with the higher frequency. The response latencies t^B and t^C decrease with increasing the respective occurrence frequencies. As the network is here operating in the absence of noise, the distribution of the response latencies $t^{B/C}$ across trials is narrow. Consequently, neurons representing the most frequent sequence fire earlier in all trials. This advanced response quickly activates the inhibitory neuron, which suppresses the activity of the neurons representing the less frequent sequence. For training frequencies between 0.4 and 0.6, the difference between t^B and t^C is small compared to the response latency of the WTA circuit. Hence, both sequences are occasionally replayed simultaneously (Fig. 3.4E).

To foster exploratory behavior, i.e., to enable occasional replay of the low-frequency sequence, we equip the excitatory neurons with noise. In this work, we investigate two different forms of noise. Here, we first consider noise provided in the form of stationary synaptic background input (see below for an alternative form of noise). Each

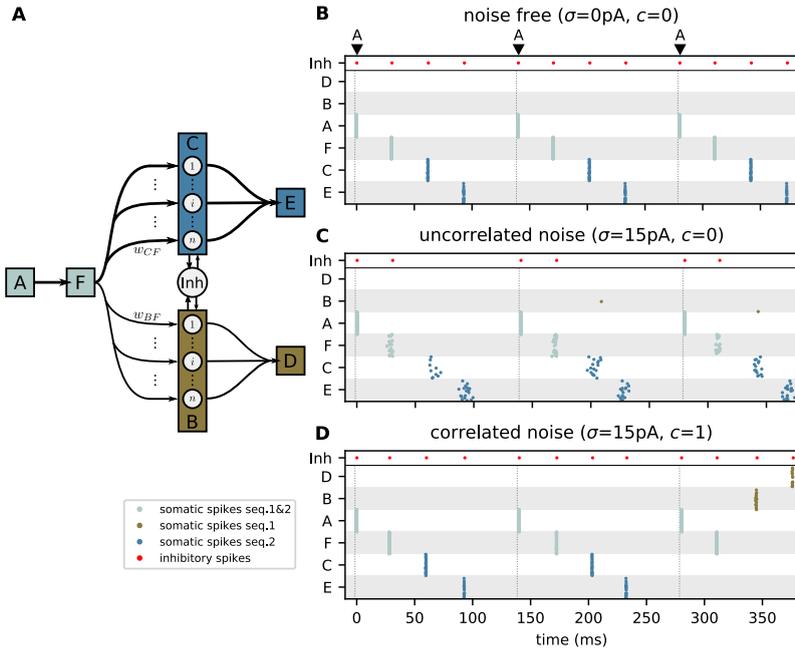


Figure 3.3: Correlated noise enhances exploratory behavior. **A**) Sketch of subpopulations of excitatory neurons (boxes) representing the elements of the two sequences $\{A, F, C, E\}$ (seq. 2) and $\{A, F, B, D\}$ (seq. 1). The subpopulations “C” and “B” are unfolded showing their respective neurons. The arrows depict the connections after learning the two sequences. The line thickness represents the synaptic weight. The presentation of the character “A” constitutes an ambiguous cue during replay. The inhibitory neuron (Inh) mediates competition between subpopulations (WTA). **B, C, D**) Spiking activity in the subpopulations depicted in panel A in response to three repetitions of the ambiguous cue “A” (black triangles at the top and vertical dotted lines) for three different noise configurations $\sigma = 0\text{ pA}$, $c = 0$ (**B**), $\sigma = 190\text{ pA}$, $c = 0$ (**C**), and $\sigma = 190\text{ pA}$, $c = 1$ (**D**). Brown, blue, and silver dots mark somatic spikes of excitatory neurons corresponding to sequence 1, sequence 2, and both, respectively. Red dots mark spikes of the inhibitory neuron. Network size: number of subpopulations $M = 8$ and number of neurons per subpopulation $n_E = 150$. See Table 7.3 for remaining parameters.

subpopulation of excitatory neurons receives input from its private pool of independent excitatory and inhibitory Poissonian spike sources (Fig. 3.1C). These background inputs are parameterized by the noise amplitude σ (standard deviation of the synaptic input current arising from these background inputs) and the noise correlation c (see Fig. 3.1C and Methods). Only inputs to neurons of the same subpopulation are correlated by an extent parameterized by c . Neurons in different subpopulations receive uncorrelated inputs. The noise amplitude σ is chosen such the subthreshold membrane potentials of the excitatory neurons are fluctuating without eliciting additional spikes. As a consequence, the distributions of response latencies $t^{B/C}$ across trials may be broadened and partly overlap (Fig. 3.4C,D). As we will show in the following, the network can adopt different replay strategies (Fig. 3.2B) depending on the amount of

this overlap. Note that noise is injected only during replay, but not during learning. During training, the weak noise employed here hardly affects the network behavior as the external inputs (stimulus) are strong and lead to a reliable, immediate responses.

With uncorrelated noise ($c = 0$), the replay behavior remains effectively non-explorative, i.e., only the high-frequency sequence is replayed in response to the cue (Fig. 3.3C). This is explained by the fact that each sequence element is represented by a subset of ρ neurons. The response latency t corresponds to the population average $t = \frac{1}{\rho} \sum_{i=1}^{\rho} t_{s,i}$ of the response latencies $t_{s,i}$ (time of first spike after the cue) for each individual neuron i within this subset. The across-trial variance

$$v = \text{Var}(t) = \frac{1}{\rho} v_s + \frac{\rho - 1}{\rho} c_s v_s \quad (3.20)$$

of this population measure t is determined by the population size ρ , the population averaged spike-time variance $v_s = \frac{1}{\rho} \sum_i \text{Var}(t_{s,i})$, and the population averaged spike-time correlation coefficient $c_s = \frac{1}{\rho(\rho-1)v_s} \sum_i \sum_{j \neq i} \text{Cov}(t_{s,i}, t_{s,j})$, with $\text{Cov}(t_{s,i}, t_{s,j})$ denoting the spike-time covariance for two neurons i and j . The response-latency statistics v_s and c_s depend on the input noise statistics σ and c in a unique and monotonous manner (Goedeke and Diesmann, 2008; De la Rocha et al., 2007). In the absence of correlations ($c = c_s = 0$), the across-trial variance v of t vanishes for large population sizes ρ . For finite population sizes, v is non-zero but small (Fig. 3.4C). The effect of the synaptic background noise on the variability of response latencies largely averages out. Hence, the average advance in the response of the population representing the high-frequency sequence cannot be overcome by noise; the network typically replays only the sequence with the higher occurrence frequency during training (Fig. 3.4F). For small differences in the occurrence frequencies (i.e., $p \approx 0.5$), the network occasionally fails to replay any sequence or replays both sequences. The mechanism underlying this behavior is explained below.

Noise averaging is efficiently avoided by introducing noise correlations. For perfectly correlated noise and, hence, perfectly synchronous spike responses ($c = c_s = 1$), the across-trial variance v of the response latency t is identical to the across-trial variance v_s of the individual spike responses, i.e., $v = v_s$, irrespective of the population size ρ ; see Eq (3.20). For smaller but non-zero spike correlations ($0 < c_s < 1$), the latency variance v is reduced but doesn't vanish as ρ becomes large. Hence, in the presence of correlated noise, the across-trial response latency distributions for two competing populations have a finite width and may overlap (Fig. 3.4D), thereby permitting an occasional replay of the sequence observed less often during training (Fig. 3.4G). Replay, therefore, becomes more exploratory, such that the occurrence frequencies during training are gradually mapped to the frequencies of the sequence replay. With an appropriate choice of the noise amplitude and correlation, even an almost perfect match between training and replay frequencies can be achieved (probability matching; Fig. 3.4G).

The results presented so far can be extended towards more than two competing sequences. As a demonstration, we train the network using three sequences $\{A,F,C,E\}$, $\{A,F,B,D\}$, and $\{A,F,G,H\}$ presented with different relative frequencies. By adjusting the noise amplitude σ and correlation c , the replay frequencies can approximate the training frequencies (Fig. 3.5).

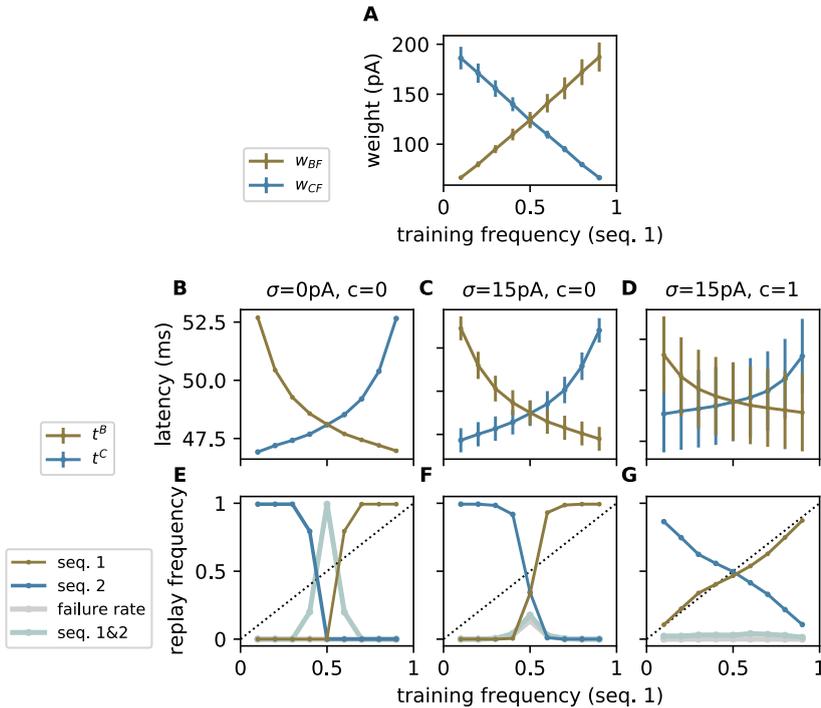


Figure 3.4: **Uncorrelated noise averages out in population based encoding.** Dependence of **A)** the compound weights (PSC amplitudes) w_{BF} (brown) and w_{CF} (blue; see Fig. 3.3A), **B–D)** the population averaged response latencies t^B and t^C (subpopulation averaged time of first spike after the cue) for subpopulations “B” (brown) and “C” (blue), and **E–G)** the relative replay frequencies of sequences 1 (brown) and 2 (blue) along with the failure rate (gray) and the joint probability of replaying both sequences (silver) on the occurrence frequency of sequence 1 during training. Panels B–G depict results for three different noise parameterization $\sigma = 0\text{pA}, c = 0$ (B,E), $\sigma = 150\text{pA}, c = 0$ (C,F), and $\sigma = 150\text{pA}, c = 1$ (D,G). In panel A, circles and error bars depict the mean and the standard deviation across different network realizations. In panels B–D, circles and error bars represent the mean and the standard deviation across $N_t = 150$ trials (cue repetitions), averaged across 5 different network realizations. In panels E–G, circles represent the mean across $N_t = 150$ trials, averaged across 5 different network realizations. See Table 7.3 for remaining parameters.

3.3.2 Noise canceling cannot be overcome by increasing noise amplitude

For subpopulations of finite size ρ , the variance v of the response latency t remains finite, and can be increased by scaling up the variance of the noise, even without correlation; see Eq (3.20). However, this solution cannot be applied to network models where a decision is mediated by a fast WTA circuit. In the presence of uncorrelated noise with high amplitude, the spikes from all neurons, in all competing subpopulations, are similarly dispersed. A large dispersion in spike times prohibits a fast and reliable activation of inhibition by one of the competing subpopulations. The WTA mechanism, therefore, fails at selecting a unique sequence. Consequently, both sequences run

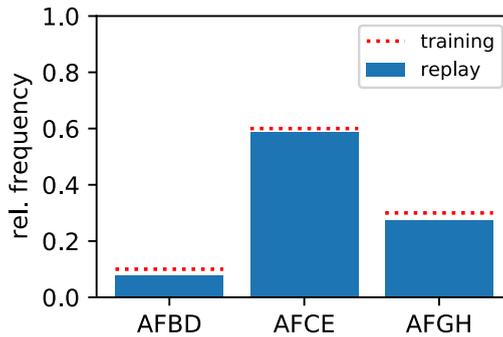


Figure 3.5: **Multiple competing sequences are learned and replayed according to their occurrence frequencies (probability matching).** During learning, three partly overlapping sequences $\{A,F,C,E\}$, $\{A,F,B,D\}$, and $\{A,F,G,H\}$ are repetitively presented with relative frequencies 0.2, 0.5, 0.3, respectively (dotted red lines). After learning, the network autonomously replays the learned sequences in response to the ambiguous cue “A” with frequencies depicted by the blue bars. Noise parameters: $\sigma = 190$ pA, $c = 1$. See Table 7.3 for remaining parameters.

through in most of the trials (Fig. 3.6A). An additional problem of the uncorrelated noise is that it impairs the propagation of the activity across the subpopulations of neurons. As our model relies on the propagation of synchronously firing neurons, the spike time dispersion resulting from the uncorrelated noise bears the risk that the spikes generated may be too dispersed to trigger dAPs in the next subpopulation (Fig. 3.6). As a result of these two problems, more explorative behavior cannot be achieved by increasing the amplitude of uncorrelated noise. Instead, the probability of simultaneous replay (no decision) and the failure rate increase (Fig. 3.6B).

Noise correlations lead to more synchronous responses, thereby reducing the overlap between the within-trial latency distributions of the two competing populations “B” and “C” (Fig. 3.3D). In each trial, the WTA dynamics is therefore triggered by just one of the two populations, rather than by both. Further, synchronous firing leads to a more robust activation of the subsequent subpopulation, and hence, a more robust replay. Hence, noise correlations help not only in generating more explorative behavior, but also in reducing replay failures and the chance of simultaneous activation of competing sequences (Fig. 3.4G).

3.3.3 Noise amplitude and level of correlation control replay strategy

We know from psychophysics experiments that humans and animals are able to flexibly change the strategy by which they select probabilistic outcomes (Shanks et al., 2002; Cohen et al., 2007). Our model is also able to express these different strategies. For small noise amplitudes and irrespective of the correlation level, the model replays deterministically the frequent sequence (max-prob, see Fig. 3.7A). With increasing the noise amplitude and sufficiently large correlation, the replay frequency approximately matches the training frequency (probability matching, Fig. 3.7B). Increasing the noise further leads to an even more explorative replay, where the replay frequencies become less dependent on the training frequencies (Fig. 3.7C). Full exploration requires a large amount of noise to overwrite the asymmetry in the synaptic weights. In the presence

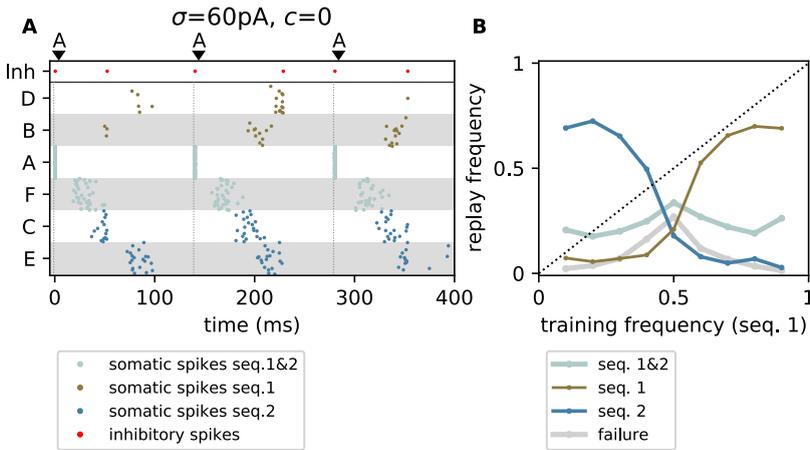


Figure 3.6: **Winner-take-all mechanism fails when increasing the amplitude of the uncorrelated noise.** **A)** Brown, blue, and silver dots mark somatic spikes of excitatory neurons belonging to sequence $\{A, F, B, D\}$ (seq. 1), sequence $\{A, F, C, E\}$ (seq. 2), or both, respectively. Red dots mark spikes from inhibitory neurons. Each trial is initiated by stimulating the first element in the sequence “A”, see dark arrows and vertical dashed lines). During training, the sequences 1 and 2 are shown with relative frequencies 0.2 and 0.8, respectively. **B)** The relative replay frequency of sequence 1 (brown) and sequence 2 (blue) along with the failure rate (gray) and the joint probability of replaying both sequences (silver) are plotted as a function of the relative training frequency of sequence 1. Circles represent the mean across $N_t = 150$ trials averaged across 5 network realizations. Parameters: $\sigma = 537$ pA and $c = 0$. See Table 7.3 for the remaining parameters.

of high-amplitude noise, neurons may become strongly hyperpolarized during stimulus arrival, and hence, fail to respond. Further, strong noise induced depolarizations may trigger spurious bursts of synchronous spikes.

In our work, it is also possible to change the replay strategy by changing the correlation level (see Fig. 3.8). If the noise amplitude is well adjusted, a low correlation level makes the dynamics non-explorative. With increasing the correlation level dynamics become more explorative. This suggests, therefore, that the replay strategy can be controlled by changing the noise amplitudes or the correlation level. Adjusting noise amplitudes can be achieved in biology by changing the rate of the background input or by changing the effective weights, which can be modulated by neuromodulation (Atherton et al., 2015) or attention (Baluch and Itti, 2011). So far, we studied correlations induced by shared presynaptic inputs. To achieve explorative dynamics, shared input correlations need to be high ($c \sim 1$). However, shared input correlations resulting from the cortical anatomy are rather low (Abeles, 1991; Braitenberg and Schüz, 1998; Shadlen and Newsome, 1998; Song et al., 2005). In the next section, we therefore propose an alternative form of noise, where correlations are generated by the network dynamics.

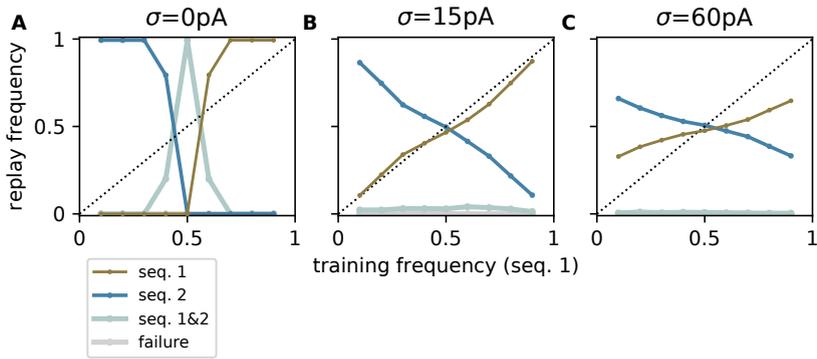


Figure 3.7: **Different replay strategies achieved by increasing the noise amplitude.**

The relative replay frequency of sequence 1 (brown) and sequence 2 (blue) along with the failure rate (gray) and the joint probability of replaying both sequences (silver) are plotted as a function of the relative training frequency of sequence 1 for different noise amplitudes $\sigma = 0\text{ pA}$ (A), $\sigma = 190\text{ pA}$ (B), and $\sigma = 537\text{ pA}$ (C) with correlation coefficient $c = 1$. Circles represent the mean across $N_t = 150$ trials, averaged across 5 different network realizations. See Table 7.3 for remaining parameters.

3.3.4 *Random stimulus locking to spatiotemporal oscillations as natural form of noise*

So far, we have used shared synaptic background input in the form of stationary Poissonian spikes as a source of correlated noise. Locally coherent noise may however also be provided in the form of randomness in stimulus timing with respect to spatiotemporal background oscillations. In the presence of traveling cortical waves, for example, nearby neurons in a given subpopulation share the same oscillation phase, whereas distant neurons belonging to different subpopulations experience different phases (Fig. 3.9). At the time of stimulus arrival, the neurons in the up phase are more excitable and tend to fire earlier than neurons in a down phase. Spatiotemporal oscillations in cortical activity are ubiquitous and occur in different forms (Nauhaus et al., 2009; Muller and Destexhe, 2012; Sato et al., 2012). In nature, external stimuli are usually not consistently locked to such oscillations. It is therefore reasonable to assume that the stimulus onset times are random with respect to the oscillation phase.

Here, we exploit this form of randomness to increase the trial-to-trial variability during replay. To investigate its effect on the replay performance, we first train the network in the absence of any background input using two sequences $\{A, F, C, E\}$ and $\{A, F, B, D\}$ with relative training frequencies p and $1 - p$, respectively. During replay, we inject an oscillating background current with amplitude a and frequency f into all excitatory neurons (see section 3.2). Neurons within a given subpopulation share the same oscillation phase. Phases for different subpopulations are randomly drawn from a uniform distribution between 0 and 2π . The replay performance of the network is assessed by monitoring the network responses to repetitive presentations of an external cue “A” with random, uniformly distributed inter-cue intervals $\Delta T_{\text{cue}} \sim \mathcal{U}(u_{\min} u_{\max})$. We define sequences $\{A, F, B, D\}$ or $\{A, F, C, E\}$ to be successfully replayed if more than 50% of the neurons in the last subpopulations “E” or “D” have fired, respectively. The

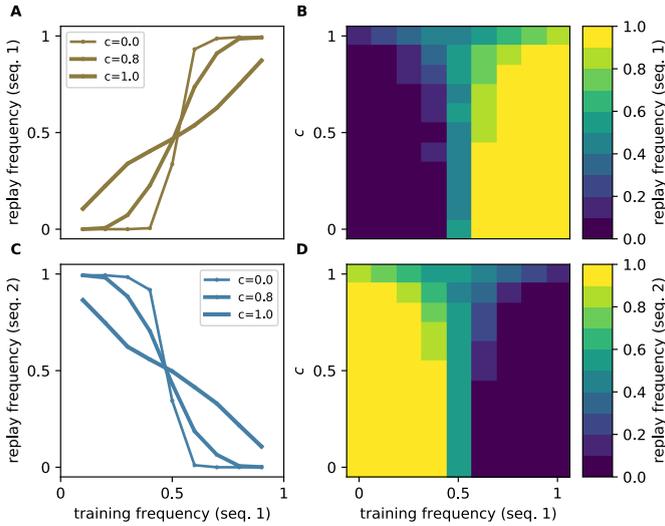


Figure 3.8: **Adjusting level of correlation permits different replay strategies.** The relative replay frequencies of sequence 1 (brown, **A**) and 2 (blue, **C**) as a function of the occurrence frequency of sequence 1 during training plotted for different correlation levels (different markers: $c = 0$, $c = 0.8$, and $c = 1$). Contour plot showing the dependence of the relative replay frequencies of sequences 1 (**B**) and 2 (**D**) on the training frequency of sequence 1 and different correlation levels. Noise amplitude $\sigma = 15$ pA. The replay frequencies are computed as the mean across $N_t = 150$ trials, averaged across 5 different network realizations. See Table 7.3 for remaining parameters.

analysis is repeated for a range of training frequencies p , oscillation amplitudes a , and frequencies f .

Depending on the choice of the oscillation amplitude a and frequency f , the network replicates different replay strategies (Fig. 3.10). For low-amplitude oscillations, the model replays only the sequence with the higher training frequency (max-prob). With increasing oscillation amplitude, the model becomes more explorative and occasionally replays the less frequent sequence. By adjusting the oscillation amplitude, the replay frequency can be closely matched to the training frequency. In line with experimental recording (Buzsáki, 2006; Buzsáki and Draguhn, 2004), this behavior is observed for a range of physiological oscillation frequencies: alpha (10 Hz), beta (30 Hz), gamma (70 Hz). Due to the low-pass filtering of the neuronal membranes and synapses, higher oscillation frequencies have a smaller effect. Consequently, increasing the oscillation frequency leads to a more reliable replay of the frequent sequence. For slow oscillations with long periods that are large compared to the average inter-cue interval, the network responses in subsequent trials are more correlated. For sufficiently many trials, however, the network can still explore different solutions.

3.4 DISCUSSION

This work proposes a spiking neural network model able to recall sequences in response to ambiguous cues following different strategies. In this model, explorative

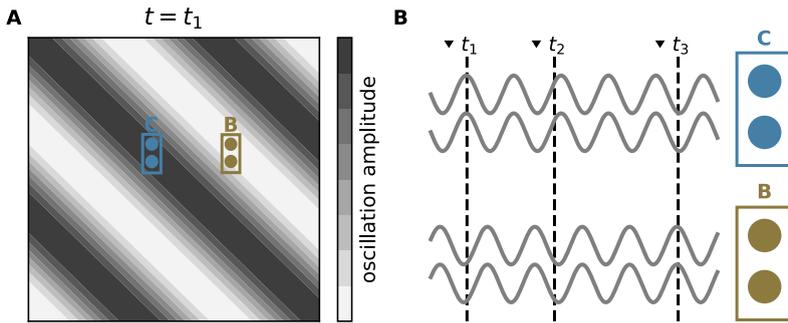


Figure 3.9: **Random locking of stimulus to global oscillations as a form of noise.**

A) Snapshot of a wave of activity traveling across a cortical region at time t_1 of the 1st stimulus onset. Grayscale depicts wave amplitudes in different regions. Brown and blue rectangles mark populations of neurons with stimulus preferences “B” and “C”, respectively. **B)** Background inputs to neurons in populations “B” and “C” at different times. Background inputs to each population “B” and “C” at different times. Background inputs to neurons within each population are in phase due to their spatial proximity. Background inputs to different populations are phase shifted. Arrows on the top depict stimulus onset times. The times t_1, t_2, \dots indicate input arrival to populations “B” and “C” (dashed vertical lines are random not locked to the background activity).

recall strategies are achieved by providing the network with external coherent noise. We explore two forms of coherent noise implemented either in the form of shared synaptic input or a random stimulus locking to global spatiotemporal oscillations in the neuronal activity. The model can switch between different replay strategies by adjusting the noise characteristics such as the noise or oscillation amplitude, as well as the noise correlation or oscillation frequency.

The sequence processing model proposed here relies on a form of population encoding. In the absence of correlations, noise injected to single neurons therefore largely averages out and leads to a quasi-deterministic and non-exploratory behavior. Locally correlated noise, in contrast, permits an explorative recall behavior where the sequence frequency during learning can be gradually mapped to the recall frequency. Furthermore, noise correlations foster synchronization between neurons within subpopulations, and thereby lead to a more robust context-specific activation of sequences during recall. The problem of noise averaging and the proposed solution are not unique to the model presented here, but are generic for all systems where relevant state variables arise from superpositions of many noisy, uncorrelated components. Fluctuations in the total input current of a single neuron resulting from superpositions of thousands of synaptic inputs, for example, can be efficiently controlled by the level of correlation in the presynaptic activity (Salinas and Sejnowski, 2001). Similarly, explorative behavior in other models of population based probabilistic computing (e.g. Legenstein and Maass, 2014) can be enhanced by equipping neurons within each population with correlated noise.

Correlation in neuronal firing can originate from both anatomical constraints or network dynamics (Tetzlaff et al., 2012; Helias et al., 2014). In this study, we investigate both types. The first type of noise is implemented in the form of irregular synaptic

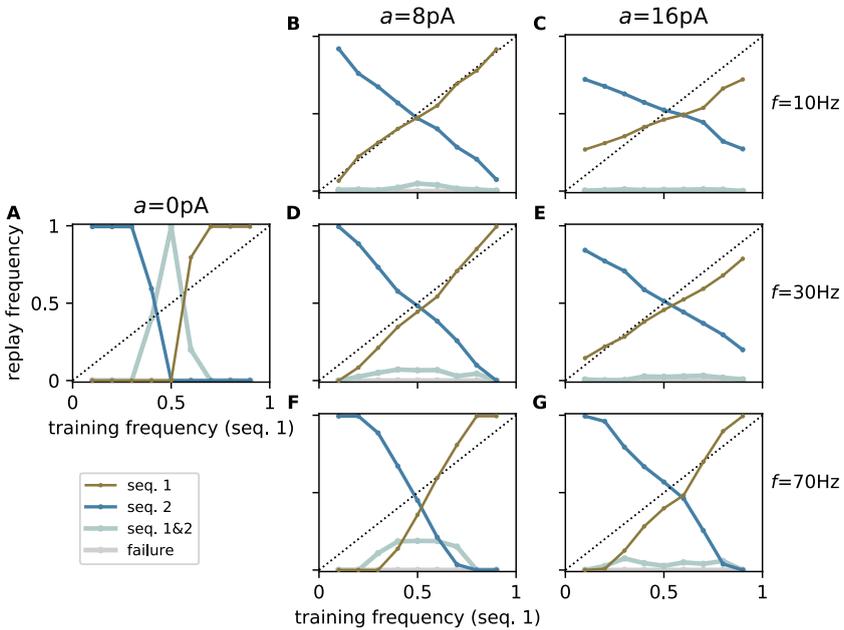


Figure 3.10: **Background oscillations as a mean for changing the replay strategies.** The relative replay frequency of sequence 1 (brown) and sequence 2 (blue) along with the failure rate (gray) and the joint probability of replaying both sequences (silver) are plotted as a function of the relative training frequency of sequence 1 for different amplitudes $a \in \{0, 8, 16\}$ and frequencies of the background oscillations: $f = 10$ Hz (B,C), $f = 30$ Hz (A,D,E), and $f = 70$ Hz (F,G). Circles represent the mean across $N_t = 150$ trials, averaged across 5 network realizations. See Table 7.3 for remaining parameters.

background input (Faisal et al., 2008; Fellous et al., 2004; Destexhe et al., 2001; Holt et al., 1996), where the correlation between neurons of the same subpopulation is resulting from shared presynaptic sources (Stroeve and Gielen, 2001; Kriener et al., 2008). From an anatomical perspective, this is reasonable as neighboring neurons indeed receive a considerable amount of inputs from identical presynaptic neurons. However, we show that the level of shared-input correlation required for an effective avoidance of noise averaging and maintenance of near synchronous activity is rather high, which contradicts anatomical studies reporting small connection probabilities in local cortical circuits, and hence, low levels of shared input correlation (Abeles, 1991; Braitenberg and Schüz, 1998; Song et al., 2005; Shadlen and Newsome, 1998). We therefore propose a second, biologically more plausible type of coherent noise resulting from a random stimulus locking to an intrinsic spatiotemporal coherent activity pattern on a large spatial scale, such as waves of cortical activity. Coherent spatiotemporal activity patterns in the cortex are observed in many different forms and under various conditions, including different sleep states, but also in awake behaving animals (Buzsáki, 2006; Buzsáki and Draguhn, 2004; Sato et al., 2012; Denker et al., 2018). In such states, neighboring neurons receive coherent input with identical phase, whereas distant neurons are exposed to different phases. Natural external stimuli such

as cue signals are usually not systematically locked to this type of intrinsic activity (unless stimuli are presented in closed-loop experiments). It is therefore reasonable to assume that the stimulus onset is random with respect to the internal state. In this study, we employ activity waves as a specific form of coherent spatiotemporal activity, and show that explorative behavior is generated for a range of plausible oscillation frequencies. We propose that a similar behavior can be achieved for other non-oscillatory forms of coherent activity, such as transient propagating wave fronts or bumps (Ermentrout, 1998; Coombes, 2005; Muller et al., 2018), as well as by other factors modulating the excitability of neighboring neurons in a coherent manner, such as transient neuromodulatory signals.

By changing the noise characteristics (such as the amplitude or frequency of the background activity, or the level of correlation), the model proposed in this study can replay competing sequences according to different strategies. For low levels of noise, the network systematically replays the sequence that occurred most often during learning (max-prob). For higher noise levels, it can match the replay frequency to the occurrence frequency during training (probability matching), or become even more explorative. This offers a potential mechanistic explanation of how animals can adjust their decision strategy based on environmental conditions (Cohen et al., 2007). In the living brain, the noise properties could be controlled by neuromodulatory signals or by inputs from other brain areas (e.g., during attention; Cohen and Maunsell, 2009). Our and many other studies predict that, in cases where the decision strategy is shifted towards exploration, more energy needs to be provided for noise generation. In line with this prediction, Daw et al. (2006) show that explorative behavior is accompanied by an increase in the BOLD signal amplitude in cortical areas associated with decision making.

A number of previous studies suggest that synaptic stochasticity, i.e., the variability in postsynaptic responses including synaptic failure (Branco and Staras, 2009), may constitute an efficient source of noise for probabilistic computations in neuronal circuits (Maass, 2014; Nefci et al., 2016). The total input to a neuron resulting from large ensembles of synapses, however, is likely to be subject to noise averaging, unless the variability of synaptic responses is correlated across synapses. To date, it remains unclear how such correlations could potentially arise. Localized neuromodulatory signals or shared presynaptic spike histories may play a role in this. Hartmann et al. (2015) demonstrate that a self-organizing recurrent neural network model can learn sequence probabilities in the absence of any form of stochasticity. Similar to our work, the model comprises recurrently coupled populations of excitatory and inhibitory threshold units and learns via a combination of spike-timing-dependent plasticity (STDP) and homeostatic plasticity mechanisms. The authors report that their model tends to overrepresent the more frequent stimuli, but they do not discuss mechanisms enabling a switching between different replay strategies. The ideas presented in our work could possibly be applied to this model.

Overall, our work ties together concepts from sequence processing and decision making in the face of ambiguity. It demonstrates that locally coherent noise is a potential mechanism underlying exploratory behavior, and shows that a random stimulus locking to macroscopic coherent activity patterns, e.g., traveling waves, can constitute such a form of noise.

SEQUENCE LEARNING IN A SPIKING NEURAL NETWORK WITH MEMRISTIVE SYNAPSES

This chapter is based on the following preprint:

Bouhadjar, Y., Siegel, S., Tetzlaff, T., Diesmann, M., Rainer W., & Wouters D. J. (2022). Sequence learning in a spiking neuronal network with memristive synapses. Preprint arXiv:2211.16592.

Author contributions: All authors contributed to the conception of the work. Younes Bouhadjar implemented the model, performed the simulations, analyzed and visualized the data, and wrote the first draft of the manuscript. All authors jointly reviewed the manuscript. Younes Bouhadjar was supervised by Tom Tetzlaff and Dirk J Wouters.

4.1 INTRODUCTION

In everyday's tasks such as learning, recognizing, or predicting objects in a noisy, ever-changing environment, brains outperform conventional computing systems and deep learning algorithms at many aspects: it has a higher capacity to generalize, can learn from small training examples, is robust with respect to perturbations and failure, and is highly resource and energy efficient. To achieve this performance, it uses intricate biological mechanisms and principles. Understanding these principles is essential for driving new advances in neuroscience and for developing new real-world applications. For instance, it is known that biological neural networks are highly sparse in activity and connectivity and they can self-organize in the face of the incoming sensory stimulus using unsupervised local learning rules. A number of biologically inspired algorithms relying on these principles have been developed for sequence prediction and replay (Lazar et al., 2009; Hawkins and Ahmad, 2016; Bouhadjar et al., 2022b), pattern recognition (Masquelier and Thorpe, 2007; Payeur et al., 2021), and decision making (Nefci and Averbeck, 2019). The spiking Temporal Memory (spiking TM) network proposed in chapter 2 learns high-order sequences in an unsupervised, continuous manner using local learning rules. Owing to its highly sparse activity and connectivity, it provides an energy-efficient sequence learning and prediction mechanism.

The spiking TM algorithm was implemented using the neural simulator NEST (Gewaltig and Diesmann, 2007). While NEST provides a simulation platform optimized for running large-scale networks efficiently in a reproducible manner, it is executed on standard von-Neumann-type computers, i.e., on hardware that is not specifically optimized for neuromorphic computing. This results in performance limitations as the simulation time and the energy dissipation become substantially high for brain-scale neural networks (Kunkel et al., 2014; Jordan et al., 2018). For using spiking TM in edge-computing applications, more efficient hardware is therefore required. Neuromorphic hardware, with dedicated solutions to the high demands imposed by the natural-density connectivity of the brain and the resulting communication load, as well as, specific circuit blocks emulating neuron and synapse functionalities, present a potential solution for that. The local learning rules and the sparse neuronal activation of the spiking TM model allow for efficient mapping of the algorithm on neuromorphic hardware.

Memristive devices were suggested as components in such a hardware (Yang et al., 2013; Jelmini and Wong, 2018; Yu, 2018). They can be used to emulate certain synaptic functionalities using only a single device, by replacing more complex CMOS-based circuits (Waser et al., 2009; Dittmann and Strachan, 2019). Their intrinsic dynamics capture similar characteristics as the biological synapses such as variability, weight dependence of the update, and nonvolatility. However, while single memristive devices may readily emulate the inference function, they cannot emulate on their own plasticity rules such as Spike-Timing-Dependent Plasticity (STDP) or homeostatic control. The change of the memristive conductivity depends on the momentary voltage difference between its two terminals, and the device has no memory of past spike events at either of its terminals nor of their relative timing. Hebbian learning such as STDP therefore can only be emulated using a memristive device by "reshaping" of the pre- and post-synaptic spike events using complex voltage pulses, so that the spike-time dependency is translated into a desired instantaneous voltage difference over the device (Zamarreño-Ramos et al., 2011; Wang et al., 2015). As a result, the learning

rule is controlled outside the actual device (see Fig. 4.1). As for implementing the learning, instead of using complex voltage pulse shapes, it is more efficient to use a controller to generate simple rectangular voltage pulses that can effectuate the desired change of the device conductance in a better, more energy efficient, and also more reliable way. The change of the device conductivity as a function of the number of applied voltage pulses can hereby be seen as an intrinsic plasticity curve of the device, where the actual pulse shape can be optimized towards desired potentiation and depression characteristics.

In this work, we investigate how the intrinsic potentiation and depression characteristics of memristive devices influence the learning of the model in chapter 2. Thereto, we adapt the original neuroscientific synapse model to accommodate memristive-type potentiation/depression characteristics. The performance of the system is assessed by varying device characteristics such as conductance values and ranges, granularity of conductance change, and device variability. In this work, we study a particular type of memristive device known as the valence change memory (VCM) ReRAM device (Waser, 2012b). We investigate its two operation modes (Cüppers et al., 2019): either the continuous, analog mode, where the conductivity changes gradually between a low conductance state (LCS) and a high conductance state (HCS), or the binary mode, where the conductivity changes abruptly between the LCS and the HCS. The binary switching is controlled by the value of an analog adaptable internal state variable (Doevenspeck et al., 2018; Zhao et al., 2019; Suri et al., 2013; Yu, 2018). It resembles the learning rule employed in chapter 2 and mimics a structural form of STDP known in the neuroscientific literature.

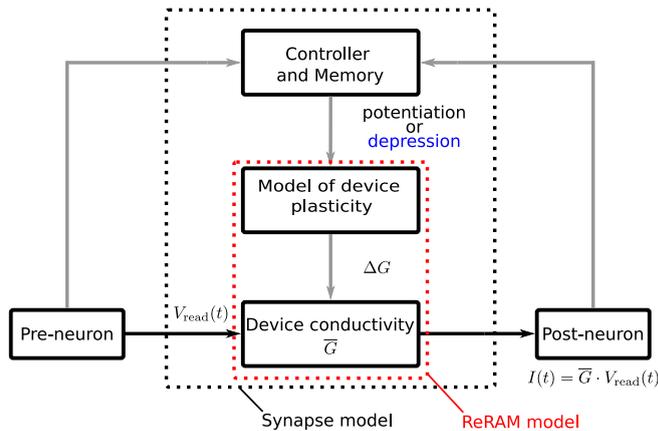


Figure 4.1: **ReRAM control circuit.** Sketch depicting the synapse model including the control circuit and the ReRAM model (red box). The circuit is composed of a read/inference path (black arrows) and a write/programming path (gray arrows). The device conductivity \bar{G} is read out whenever a presynaptic neuron emits a spike, which results in a postsynaptic current $I(t) = \bar{G} \cdot V_{\text{read}}$. The device conductivity is updated by the programming path. The controller receives pre- and postsynaptic spikes and decides on applying either a depression or a potentiation event (or both). In the next step, the model of device plasticity computes the conductance increment/decrement ΔG .

4.2 METHODS

The network structure, the external inputs, the neuron model, and the synaptic currents are similar to the ones introduced in chapter 2. We summarize them again in Table 4.1 and provide parameter values in Table 7.4. The main difference compared to chapter 2 lies in the plasticity dynamics, which is explained in section 4.3.1 and governs synapses from excitatory to excitatory neurons.

4.2.1 Model tables

Summary		
Populations	excitatory neurons (\mathcal{E}), inhibitory neurons (\mathcal{I}), external spike sources (\mathcal{X}); \mathcal{E} and \mathcal{I} composed of M disjoint subpopulations \mathcal{M}_k and \mathcal{I}_k ($k = 1, \dots, M$)	
Connectivity	<ul style="list-style-type: none"> • sparse random connectivity between excitatory neurons (plastic) • local recurrent connectivity between excitatory and inhibitory neurons (static) 	
Neuron model	<ul style="list-style-type: none"> • excitatory neurons: leaky integrate-and-fire (LIF) with nonlinear input integration (dendritic action potentials) • inhibitory neurons: leaky integrate-and-fire (LIF) 	
Synapse model	exponential postsynaptic currents (PSCs)	
Plasticity	homeostatic spike-timing-dependent plasticity in excitatory-to-excitatory connections	
Populations		
Name	Elements	Size
$\mathcal{E} = \cup_{i=k}^M \mathcal{M}_k$	excitatory (E) neurons	N_E
$\mathcal{I} = \cup_{i=k}^M \mathcal{I}_k$	inhibitory (I) neurons	N_I
\mathcal{M}_k	excitatory neurons in subpopulation k $\mathcal{M}_k \cap \mathcal{M}_l = \emptyset$ ($\forall k \neq l \in [1, M]$)	n_E
\mathcal{I}_k	inhibitory neurons in subpopulation k $\mathcal{I}_k \cap \mathcal{I}_l = \emptyset$ ($\forall k \neq l \in [1, M]$)	n_I
$\mathcal{X} = \{x_1, \dots, x_M\}$	external spike sources	M
Connectivity		
Source population	Target population	Pattern
\mathcal{E}	\mathcal{E}	random; fixed in-degrees $K_i = K_{EE}$, delays $d_{ij} = d_{EE}$, synaptic time constants $\tau_{ij} = d_{EE}$ plastic weights $G_{ij} \in \{0, \bar{G}_{ij}\}$ ($\forall i \in \mathcal{E}, \forall j \in \mathcal{E}$; “EE connections”)
\mathcal{M}_k	\mathcal{I}_k	all-to-all; fixed delays $d_{ij} = d_{IE}$, synaptic time constants $\tau_{ij} = \tau_{IE}$, and weights $G_{ij} = G_{IE}$ ($\forall i \in \mathcal{M}_k, \forall j \in \mathcal{I}_k, \forall k \in [1, M]$; “IE connections”)
\mathcal{I}_k	\mathcal{M}_k	all-to-all; fixed delays $d_{ij} = d_{EI}$, synaptic time constants $\tau_{ij} = \tau_{EI}$, and weights $G_{ij} = G_{EI}$ ($\forall i \in \mathcal{I}_k, \forall j \in \mathcal{M}_k, \forall k \in [1, M]$; “EI connections”)
\mathcal{I}_k	\mathcal{I}_k	none ($\forall k \in [1, M]$; “II connections”)
$\mathcal{X}_k = x_k$	\mathcal{M}_k	one-to-all; fixed delays $d_{ik} = d_{EX}$, synaptic time constants $\tau_{ij} = \tau_{EX}$, and weights $J_{ik} = G_{EX}$ ($\forall i \in \mathcal{M}_k, \forall k \in [1, M]$; “EX connections”)
no self-connections (“autapses”), no multiple connections (“multapses”)		
all unmentioned connections $\mathcal{M}_k \rightarrow \mathcal{I}_l, \mathcal{I}_k \rightarrow \mathcal{M}_l, \mathcal{I}_k \rightarrow \mathcal{I}_l, \mathcal{X}_k \rightarrow \mathcal{M}_l$ ($\forall k \neq l$) are absent		

Continued next page

Neuron and synapse	
Neuron	
Type	leaky integrate-and-fire (LIF) dynamics
Description	<p>dynamics of membrane potential $V_i(t)$ of neuron i:</p> <ul style="list-style-type: none"> • emission of the kth spike of neuron i at time t_i^k if $V_i(t_i^k) \geq \theta_i \quad (4.1)$ with somatic spike threshold θ_i • reset and refractoriness: $V_i(t) = V_r \quad \forall k, \forall t \in \left(t_i^k, t_i^k + \tau_{\text{ref},i} \right]$ with refractory time $\tau_{\text{ref},i}$ and reset potential V_r • subthreshold dynamics: $\tau_{m,i} \dot{V}_i(t) = -V_i(t) + R_{m,i} I_i(t) \quad (4.2)$ with membrane resistance $R_{m,i} = \frac{\tau_{m,i}}{C_{m,i}}$, membrane time constant $\tau_{m,i}$, and total synaptic input current $I_i(t)$ • $\tau_{m,i} = \tau_{m,E}$, $C_{m,i} = C_m$, $\theta_i = \theta_E$, $\tau_{\text{ref},i} = \tau_{\text{ref},E}$ ($\forall i \in \mathcal{E}$) • $\tau_{m,i} = \tau_{m,I}$, $C_{m,i} = C_m$, $\theta_i = \theta_I$, $\tau_{\text{ref},i} = \tau_{\text{ref},I}$ ($\forall i \in \mathcal{I}$)
Synapse	
Type	exponential or alpha-shaped postsynaptic currents (PSCs)
Description	<ul style="list-style-type: none"> • total synaptic input current $I_i(t) = I_{\text{ED},i}(t) + I_{\text{EX},i}(t) + I_{\text{EI},i}(t), \quad \forall i \in \mathcal{E} \quad (4.3)$ $I_i(t) = I_{\text{IE},i}(t), \quad \forall i \in \mathcal{I}$ with dendritic, inhibitory, external and excitatory input currents $I_{\text{ED},i}(t)$, $I_{\text{EI},i}(t)$, $I_{\text{EX},i}(t)$, $I_{\text{IE},i}(t)$ evolving according to $I_{\text{ED},i}(t) = \sum_{j \in \mathcal{E}} (\alpha_{ij} * s_j)(t - d_{ij}) \quad (4.4)$ with $\alpha_{ij}(t) = V_{\text{read}} G_{ij} \frac{e}{\tau_{\text{ED}}} t e^{-t/\tau_{\text{ED}}} \Theta(t)$ and $\Theta(t) = \begin{cases} 1 & t \geq 0 \\ 0 & \text{else} \end{cases}$ $\tau_{\text{EI}} \dot{I}_{\text{EI},i} = -I_{\text{EI},i}(t) + V_{\text{read}} \sum_{j \in \mathcal{I}} G_{ij} s_j(t - d_{ij}) \quad (4.5)$ $\tau_{\text{EX}} \dot{I}_{\text{EX},i} = -I_{\text{EX},i}(t) + V_{\text{read}} \sum_{j \in \mathcal{X}} G_{ij} s_j(t - d_{ij}) \quad (4.6)$ $\tau_{\text{IE}} \dot{I}_{\text{IE},i} = -I_{\text{IE},i}(t) + V_{\text{read}} \sum_{j \in \mathcal{E}} G_{ij} s_j(t - d_{ij}) \quad (4.7)$ with τ_{EX}, τ_{EI}, and τ_{IE} synaptic time constants of EX, EI, and IE connections, respectively, G_{ij} the synaptic weight, and V the read voltage • presynaptic spike trains $s_j(t) = \sum_k \delta(t - t_k^j)$ • dAP generation: <ul style="list-style-type: none"> – emission of lth dAP of neuron i at time t_i^l if $I_{\text{ED},i}(t_i^l) \geq \theta_{\text{dAP}}$ – dAP current plateau: $I_{\text{ED},i}(t) = I_{\text{dAP}} \quad \forall l, \forall t \in \left(t_i^l, t_i^l + \tau_{\text{dAP}} \right] \quad (4.8)$ with dAP current plateau amplitude I_{dAP}, dAP current duration τ_{dAP}, and dAP activation threshold θ_{dAP}.

Plasticity	
Type	Hebbian-type plasticity and dAP-rate homeostasis
EE synapses	<ul style="list-style-type: none"> • Hebbian plasticity described in section 4.3.1 • homeostatic control: <ul style="list-style-type: none"> – if $z_i(t) > z^*$: a depression pulse is applied (see Eq (4.9) or Eq (4.11)) – if $z_i(t) \leq z^*$: a potentiation pulse is applied (see Eq (4.9) or Eq (4.11)) with the dAP trace $z_i(t)$ and target dAP activity z^*. • dAP trace $z_i(t)$ of postsynaptic neuron i, evolving according to $\frac{dz_i}{dt} = -\tau_h^{-1} z_i(t) + \sum_k \delta(t - t_{\text{dAP},i}^k)$ with onset time $t_{\text{dAP},i}^k$ of the kth dAP, homeostasis time constant τ_h
all other synapses	non-plastic

Continued next page

Input
<ul style="list-style-type: none"> • repetitive stimulation of the network using the same set $\mathcal{S} = \{s_1, \dots, s_S\}$ of sequences $s_i = \{\zeta_{i,1}, \zeta_{i,2}, \dots, \zeta_{i,C_i}\}$ of ordered discrete items $\zeta_{i,j}$ with number of sequences S and length C_i of ith sequence • presentation of sequence element $\zeta_{i,j}$ at time $t_{i,j}$ modeled by a single spike $x_k(t) = \delta(t - t_{i,j})$ generated by the corresponding external source x_k • inter-stimulus interval $\Delta T = t_{i,j+1} - t_{i,j}$ between subsequent sequence elements $\zeta_{i,j}$ and $\zeta_{i,j+1}$ within a sequence s_i • inter-sequence time interval $\Delta T_{\text{seq}} = t_{i+1,1} - t_{i,C_i}$ between subsequent sequences s_i and s_{i+1} • example sequence sets: <ul style="list-style-type: none"> – sequence set I: $\mathcal{S} = \{\{A,D,B,E,I\}, \{F,D,B,E,C\}, \{H,L,J,K,D\}, \{G,L,J,K,E\}\}$
Output
<ul style="list-style-type: none"> • somatic spike times $\{t_i^k \forall i \in \mathcal{E}, k = 1, 2, \dots\}$ • dendritic currents $I_{\text{ED},i}(t)$ ($\forall i \in \mathcal{E}$)
Initial conditions and network realizations
<ul style="list-style-type: none"> • membrane potentials: $V_i(0) = V_r$ ($\forall i \in \mathcal{E} \cup \mathcal{I}$) • dendritic currents: $I_{\text{ED},i}(0) = 0$ ($\forall i \in \mathcal{E}$) • external currents: $I_{\text{EX},i}(0) = 0$ ($\forall i \in \mathcal{E}$) • inhibitory currents: $I_{\text{EI},i}(0) = 0$ ($\forall i \in \mathcal{E}$) • excitatory currents: $I_{\text{IE},i}(0) = 0$ ($\forall i \in \mathcal{I}$) • synaptic permanences: $P_{ij}(0) = P_{\min,ij}$ with $P_{\min,ij} \sim \mathcal{U}(P_{0,\min}, P_{0,\max})$ ($\forall i, j \in \mathcal{E}$) • synaptic weights: $\bar{G}_{ij}(0) = G_{\min,ij}$ with $G_{\min,ij} \sim \mathcal{U}(G_{0,\min}, G_{0,\max})$ ($\forall i, j \in \mathcal{E}$) (analog synapse) • synaptic weights: $\bar{G}_{ij}(0) = G_{\min}$ ($\forall i, j \in \mathcal{E}$) (binary synapse) • spike traces: $x_i(0) = 0$ ($\forall i \in \mathcal{E}$) • dAP traces: $z_i(0) = 0$ ($\forall i \in \mathcal{E}$) • potential connectivity and initial permanences randomly and independently drawn for each network realization
Simulation details
<ul style="list-style-type: none"> • network simulations performed in NEST (Gewaltig and Diesmann, 2007) version 3.0 (Hahne et al., 2021) • definition of excitatory neuron model using NESTML (Plotnikov et al., 2016; Nagendra Babu et al., 2021) • synchronous update using exact integration of system dynamics on discrete-time grid with step size Δt (Rotter and Diesmann, 1999) • source code underlying this study: 10.5281/zenodo.6754964

Table 4.1: Description of the network model. Parameter values are given in Table 7.4.

4.3 RESULTS

4.3.1 *ReRAM synapse model*

In this section, we briefly review the ReRAM device dynamics, introduce our model of the ReRAM device and its control circuitry (Fig. 4.1), and characterize the resulting model dynamics.

As explained in section 1.4.4, the VCM ReRAM device is composed of a stack of three materials: a metal-oxide film sandwiched between a low work function (WF) and a high WF, blocking Schottky-interface electrode (see Fig. 4.2). In a first step, a conductive filament is formed due to a high voltage applied across the device. This filament consists of oxygen vacancies that act as local dopant elements in the insulating metal-oxide matrix. During further operation, these oxygen vacancies can be moved by means of a high electric field and a local Joule heating (as a function of the polarity, either towards or away from the blocking electrode). The device conductance is controlled by the concentration of oxygen vacancies N_{VO} in a small region (gap) near the electrical blocking electrode: at low N_{VO} , the filament is “broken”, which gives rise to a high conduction barrier between the high WF electrode and the remaining filament (plug), i.e., the device is in the LCS. If N_{VO} is high, we have a “connecting” filament, where the high N_{VO} lowers the conduction potential barrier at the blocking electrode, i.e., the device is in the HCS. Hence, the concentration of the oxygen vacancies N_{VO} can be seen as an internal state variable of the device.

As mentioned above, depending on the initial resistance range and the voltage pulse amplitude and width, a VCM ReRAM device can operate in two different modes, i.e., binary or analog (Cüppers et al., 2019). In the analog mode, the applied pulses result in a gradual monotonous change of the device conductance, for both potentiation and depression. This operation mode is suitable for the implementation of STDP-type learning rules. It is, however, characterized by a limited conductivity range, and the device switching characteristics may slowly drift away from the analog behavior to a more abrupt conductivity change. In the binary mode, the conductivity can only be switched between two values, the LCS and the HCS state. The switching between these two states occurs abruptly. In previous works, the abrupt, binary switching is achieved using single program pulses with a sufficiently large amplitude (Cüppers et al., 2019). In contrast, here, we study the switching behavior of the device as a response to a certain number of pulses of smaller amplitudes. As a response to these pulses, an internal state variable N_{VO} gradually increases (Fleck et al., 2016). Only when this N_{VO} exceeds a certain threshold value, a thermal runaway condition is reached resulting in an abrupt switching event. Due to intrinsic ReRAM device variabilities (Fantini et al., 2013), the number of pulses to reach this thermal runaway conditions shows a strong device-to-device and cycle-to-cycle variation. During the depression, the switching is intrinsically more gradual, due to the lack of an internal runaway mechanism as present for the potentiation operation. Adding a series resistance (in or outside the device) can provide such runaway mechanism due to a voltage divider effect also in the RESET case (Hardtdegen et al., 2018). Hence, in both cases, the switching behavior can be summarized as follow: at first only a gradual change of the internal state variable N_{VO} is observed, associated with only a minor change of the device conductivity, followed by a strong switching effect (large change of N_{VO} as well as of the associated conductivity) when the internal state variable reaches a certain threshold. This operation mode is of particular interest for this study, as it is

similar to the structural STDP plasticity discussed and implemented in the original spiking TM model (see chapter 2).

Previous studies suggested both physics-based and phenomenological models for VCM type ReRAMs. Physics-based models such as the JART model (Bengel et al., 2020) capture detailed physical characteristics and predict their specific experimental behavior. They require however long simulation time and lead to convergence issues. On the other hand, the more phenomenological models give a high-level description of the operational characteristics, have good accuracy, are computationally less demanding, and can hence be combined with large-scale network models. In this study, we opt for a phenomenological model to implement both the analog and the binary ReRAM device.

The synapses are either potentiated or depressed by following learning rules similar to those outlined in the spiking TM model. The learning rules are implemented by the control circuit (Fig. 4.1) as follows: the synapse is depressed slightly at every presynaptic spike and potentiated if a postsynaptic spike follows after a presynaptic spike. In contrast to the original spiking TM model, synapses are potentiated by a fixed amount irrespective of the relative timing between the pre- and postsynaptic spikes. The potentiation is however disabled if these spikes occur very close to each other within the interval $[0, \Delta T_{\min}]$. This prohibits synchronously firing neurons from connecting to each other. The control circuit further implements a homeostatic control mechanism (see section 4.3.2).

In the analog mode, the increment

$$\Delta G_{i,j} = \begin{cases} G_{\max} \cdot \lambda_+ \cdot \left(1 - \frac{G_{i,j}}{G^*}\right)^{\mu_+} + X & \text{for potentiation} \\ -G_{\max} \cdot \lambda_- \cdot \left(\frac{G_{i,j}}{G^*}\right)^{\mu_-} + X & \text{for depression} \end{cases} \quad (4.9)$$

in the conductivity of the device (synapse) $j \rightarrow i$ following a potentiation or a depression event is modeled as in (Fusi and Abbott, 2007), but with an additional additive noise X . For each synapse and for each potentiation and depression step, the noise $X \sim \mathcal{N}(0, \sigma_w^2)$ is randomly and independently drawn from a normal distribution with zero mean and standard deviation σ_w . The conductance $G_{i,j}$ evolves between a lower and an upper bound G_{\min} and G_{\max} , and it is clipped at these boundaries, with learning rates λ_{plus} and λ_{minus} and weight dependence exponents μ_+ and μ_- . The conductance changes linearly with the internal state variable N_{VO} , thus no specification of the internal state variable is necessary. The initial conductance $G_{\min} = G_{i,j}(0)$ is drawn for every new device from a uniform distribution in the interval $[G_{0,\min}, G_{0,\max}]$.

For the binary switching behavior, we use a similar model as the structural STDP model proposed in chapter 2. The switching of the conductance between the LCS and the HCS is controlled by a permanence P . The permanence plays the role of the internal state variable N_{VO} . If it is above a certain threshold θ_P , the conductance $G_{i,j}$ is set to G_{\max} , otherwise it is set to G_{\min} :

$$G_{i,j}(t) = \begin{cases} G_{\max} & \text{if } P_{ij}(t) \geq \theta_P \\ G_{\min} & \text{if } P_{ij}(t) < \theta_P. \end{cases} \quad (4.10)$$

Similar to the analog synapse, the initial conductance G_{\min} is drawn for every new device from a uniform distribution in the interval $[G_{0,\min}, G_{0,\max}]$. At each potentiation

or depression step, the permanence P of the synapse $j \rightarrow i$ is incremented by an amount

$$\Delta P_{i,j} = \begin{cases} P_{\max} \cdot \lambda_+ \cdot \left(1 - \frac{P_{i,j}}{P_{\max}}\right)^{\mu_+} + X & \text{for potentiation} \\ -P_{\max} \cdot \lambda_- \cdot \left(\frac{P_{i,j}}{P_{\max}}\right)^{\mu_-} + X & \text{for depression,} \end{cases} \quad (4.11)$$

similar to the conductance increment of the analog synapse. It has a lower and an upper bound P_{\min} and P_{\max} and it is clipped at these boundaries. While the maximum permanences P_{\max} are identical for all EE connections, the minimal permanences $P_{\min,ij}$ are uniformly distributed in the interval $[P_{0,\min}, P_{0,\max}]$.

In addition to the write noise introduced by means of the variable X , both the analog and the binary synapse models incorporate a read noise. At each presynaptic spike of neuron j , a noisy component Z is added to the synaptic current

$$I_{i,j}(t) = (G_{i,j}(t) + Z) \cdot V_{\text{read}}(t) = \bar{G}_{i,j}(t) \cdot V_{\text{read}}(t), \quad (4.12)$$

of neuron i , where $Z \sim \mathcal{N}(0, \sigma_r^2)$ is randomly and independently drawn from a normal distribution with zero mean and standard deviation σ_r , and $V_{\text{read}}(t)$ is the applied voltage. In the course of this article, we use \bar{G} to denote the conductance incorporating both the read and the write noise.

Fig. 4.3 shows an exemplary switching behavior of the analog and binary synapse models for a specific set of parameters using 100 consecutive potentiation (i.e., SET) and depression (i.e., RESET) updates. We choose different learning rates (λ_+ and λ_-) for the two types of devices such that they switch from the LCS to the HCS state (and back) at about the same number of updates. Under normal operation of the spiking TM model, a potentiation update is always followed by a small depression (Fig. 4.4A). In the case of the analog synapse, the total synaptic growth in the absence of noise is therefore governed by

$$\Delta G_{i,j} = G_{\max} \left[\lambda_+ \cdot \left(1 - \frac{G_{i,j}}{G_{\max}}\right)^{\mu_+} - \lambda_- \cdot \left(\frac{G_{i,j}}{G_{\max}}\right)^{\mu_-} \right]. \quad (4.13)$$

The stationary solution of the device conductance (fixed point) G^* , obtained by setting $\Delta G_{i,j} = 0$, is always below the maximum conductance G_{\max} (Fig. 4.4B). The permanence of the binary synapse is subject to this effect, too. After a number of potentiation steps, it reaches a value P^* smaller than P_{\max} (see Fig. 4.4C). According to Eq (4.10), the conductance can however still assume G_{\max} . Only if the depression is too strong, the device may not reach the maturity threshold θ_P , and thus not switch to the HCS state.

In the next sections, we evaluate the effects of different characteristics of the analog and the binary switching dynamics such as the weight dependence of the device update (μ_+ , μ_-), the conductance range (G_{\min} , G_{\max}), the learning rates (λ_+ , λ_-), as well as the write and the read variability (σ_w , σ_r) on the learning process of the spiking TM model.

4.3.2 A spiking neural networks with ReRAM synapses successful at sequence prediction

Sequence learning and prediction are principal computations performed by the brain and have a number of potential technological applications. The study of chapter 2

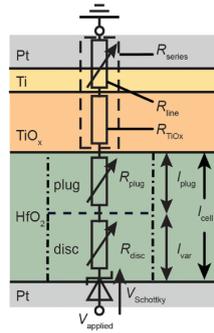


Figure 4.2: **ReRAM stack.** Equivalent circuit diagram for JART VCM model describing the Pt/HfO₂/TiO_x/Pt (HOTO) device. Figure by Bengel et al. (2020) licensed under a Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>).

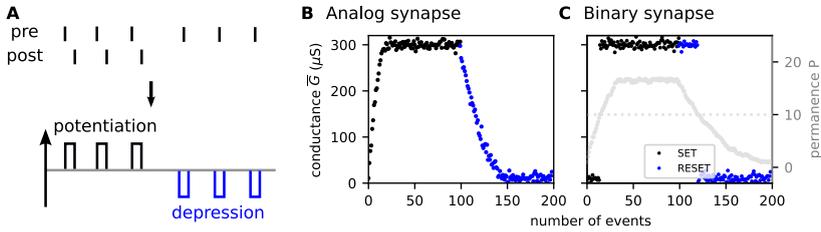


Figure 4.3: **Intrinsic dynamics of the ReRAM model (simulation results).** **A)** Sketch of the experimental protocol and mapping of pre- and postsynaptic spike timing (top) to the corresponding SET (potentiation; black) and RESET (depression; blue) operations (bottom). Evolution of the conductance \bar{G} in response to 100 SET (potentiation; black) updates, followed by 100 RESET (depression; blue) updates, for the analog (**B**) and the binary ReRAM model (**C**). In **C**, the permanence of the binary device is plotted in grey. Parameters: learning rates $\lambda_+ = 0.1$, $\lambda_- = \lambda_+/3$ (analog synapse), $\lambda_+ = 0.04$, $\lambda_- = \lambda_+/3$ (binary synapse), weight dependence exponents $\mu_+ = 0.5$, $\mu_- = 0.5$, and noise amplitudes $\sigma_r = 0.03$, $\sigma_w = 0.01$. For remaining parameters, see Table 7.4.

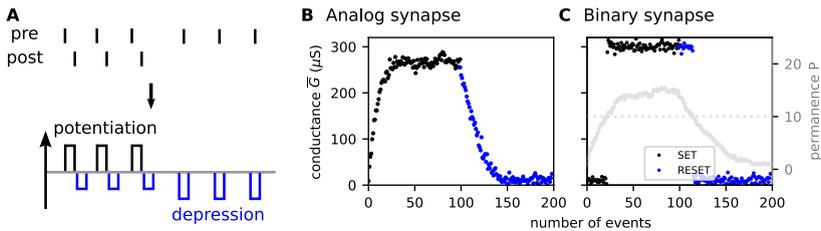


Figure 4.4: **Dynamics of the ReRAM model in the spiking TM context.** Same figure arrangement as in Fig. 4.3. In the context of the spiking TM model, each potentiation pulse is accompanied by a smaller depression pulse.

proposed a model of this type of computation known as the spiking Temporal Memory (spiking TM) model. It consists of a sparsely and recurrently connected network of spiking neurons and learns sequences continuously in an unsupervised manner by means of known biological plasticity mechanisms. After learning, the network successfully predicts and recalls complex sequences in a context-specific manner, and signals anomalies in the data.

Here, we study the prediction performance for the network with either the binary or the analog ReRAM synapses (Fig. 4.6). In contrast to chapter 2, The dAP threshold θ_{dAP} of excitatory neurons is chosen such that the co-activation of γ presynaptic neurons reliably triggers a dAP in the target neuron:

$$\theta_{\text{dAP}} = G_+ \cdot \gamma \cdot p. \quad (4.14)$$

In the case of the analog synapse, G_+ is taken to be the steady-state conductance G^* , and in the case of the binary synapse, it is taken to be G_{max} . We use the synaptic parameters fitted from the exemplary data discussed in section 4.3.1. To quantify the sequence prediction performance, we repetitively stimulate the network using the same set of sequence $\{A,D,B,E,I\}$, $\{F,D,B,E,C\}$, $\{H,L,J,K,D\}$, $\{G,L,J,K,E\}$ and assess the prediction error (see section 2.2.3). In the following, we refer to these sequences as sequence set I. To ensure the performance results are not specific to a single network, the evaluation is repeated for a number of randomly instantiated network realizations with different initial connectivities. After each new network instantiation, the initial prediction error is at 1 (Fig. 4.6). With an increasing number of training episodes, the prediction error for both networks with either the binary or the analog synapses decreases to zero as both networks learn the sequences and develop context-dependent pathways between successive sequence elements (see Fig. 4.5).

4.3.3 Influence of device characteristics on prediction performance

ReRAM devices reported in the literature exhibit different nonidealities, including 1) limited precision or the number of synaptic levels; 2) limited dynamic range; 3) dependence of the synaptic updates on the weight; 4) device variability, including read and write variability (see Zhao et al., 2020, for an overview). In this section, we study how these nonidealities affect the prediction performance in the spiking TM model.

4.3.3.1 Optimal prediction performance obtained for a broad range of on-off ratios and learning rates

The dynamic range is defined as the on-off ratio between the minimum (G_{min}) and the maximum conductance (G_{max}). Most ReRAM devices exhibit an on-off ratio in a range of $2x$ to $> 10^4x$ (Hong et al., 2018). Within the minimum and the maximum conductance, the synaptic precision or the number of synaptic steps is limited. In the synapse model proposed in section 4.3.1, we can influence the number of steps by changing parameters such as the learning rates (λ_+ , λ_-), weight dependence exponents (μ_+ , μ_-), and the on-off ratio. Given the difficulty to derive an analytical solution of the number of steps as a function of these parameters, we restrict the scope of the study in this section to investigating the influence of different learning rates and on-off ratios on the prediction performance.

We vary the on-off ratio between 5 and 40 by keeping G_{min} fixed and varying G_{max} . As G_{min} is drawn from a uniform distribution in the interval $[G_{0,\text{min}}, G_{0,\text{max}}]$, we

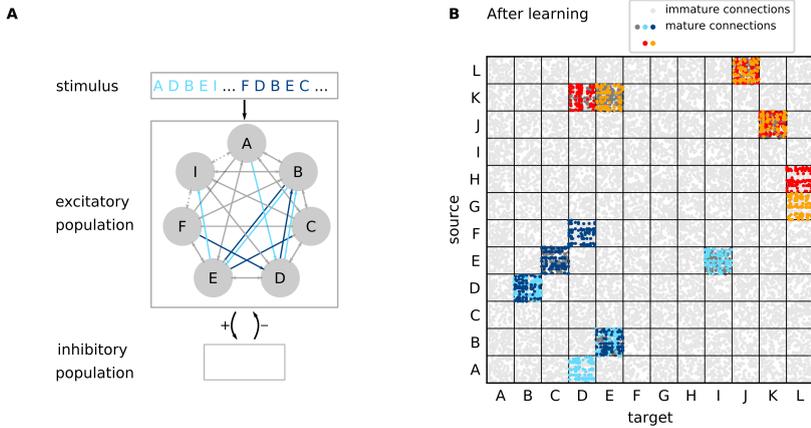


Figure 4.5: **Network structure.** **A)** Sketch of the model architecture composed of a randomly and sparsely connected recurrent network of excitatory and inhibitory neurons. The excitatory neuron population is subdivided into subpopulations according to stimulus preference (gray circles). During learning, sequence-specific, sparsely connected subnetworks with mature synapses are formed (light and dark blue arrows). For the example discussed in the main text and in panel B, the network learns four high-order sequences $\{A,D,B,E,I\}$, $\{F,D,B,E,C\}$, $\{H,L,J,K,D\}$ and $\{G,L,J,K,E\}$. In panel A, only two of them are depicted for clarity. **B)** Connectivity matrix of excitatory neurons after learning. Target and source neurons are grouped into stimulus-specific subpopulations (“A”, ..., “F”). During the learning process, subsets of connections between subpopulations corresponding to subsequent sequence elements become mature and effective ($\{A,D,B,E,I\}$: light blue, $\{F,D,B,E,C\}$: dark blue, $\{H,L,J,K,D\}$: red, $\{G,L,J,K,E\}$: orange). Immature synapses are marked by light gray dots. Dark gray dots in panel B correspond to mature connections between neurons that remain silent after learning. Only 1% of immature connections are shown for clarity.

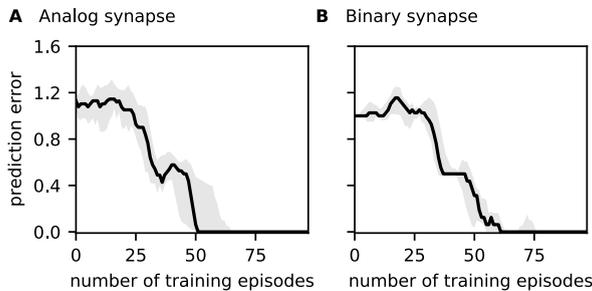


Figure 4.6: **Prediction error.** Dependence of the prediction error on the number of training episodes for the network with analog synapses (**A**) or binary synapses (**B**). Curves and error bands indicate the median as well as the 5% and 95% percentiles across an ensemble of 5 different network realizations, respectively. Same parameters as in Fig. 4.3.

compute the on-off ratio as G_{\max}/G_{\min}^* , where $G_{\min}^* = (G_{0,\max} + G_{0,\min})/2$. As we change G_{\max} , we modify the dAP threshold, see Eq (4.14). In addition, we vary the learning rate between 2% and 42% (Fig. 4.7). Parameters such as the read and write variability and the weight dependence exponents are taken from the exemplary data presented in section 4.3.1. We study the influence of the variability and the dependence of the synaptic updates on the weight more systematically in the upcoming sections. For the analog synapse, the prediction error converges to zero for an on-off ratio between 15 and 40 and for a learning rate between 2% and 18%. For the binary synapse, successful learning is obtained for an on-off ratio between 10 and 40 and for a learning rate between 2% and 18% (Fig. 4.7A,B). For learning rates above 18%, the prediction performance becomes less stable with sudden failures for some network realizations. While decreasing the learning rate yields minimum prediction error, the number of episodes-to-solution tends to increase as either the conductances or permanences need more learning steps to reach their maximum value (Fig. 4.7C,D). The learning in the network with binary synapses is faster due to the internal dynamics of binary synapses, which has faster switching dynamics compared to analog synapses: the permanence takes less number of update steps to reach the maturity threshold (θ_P) compared to the number of update steps the conductance of the analog synapses require to go from the LCS to the HCS.

In general, the on-off ratio in the spiking TM network is limited due to the following: the transition of the network activity from being initially non-sparse to becoming sparse after learning requires initial small conductances to avoid spurious activation of the dAPs, but high conductances after learning to allow the sparse set of active neurons to generate the dAP reliably. If the on-off ratio is too small this distinction between high and small conductances cannot be realized. Moreover, for successful learning, the network with analog synapses requires a higher on-off ratio compared to the network with binary synapses. This is due to the effect described in section 4.3.1 below equation Eq (4.13), which prohibits the conductance from reaching G_{\max} . Therefore, the effective on-off ratio is reduced. The learning mechanisms of the spiking TM also limit the range of possible learning rates. Increasing the learning rate bears the risk that a large fraction of neurons reaches the dAP threshold at the same time. The WTA mechanism selects then all neurons that generate dAP to become active. This leads to a loss of sparseness, which results in impairing the prediction performance. Decreasing the learning rate considerably is also not ideal as the network would learn very slowly.

4.3.3.2 Resilience of the model against weight dependence updates

The evolution of the conductance of realistic analog synapses grows in a nonlinear manner as a function of the potentiation and depression updates. The synapse model in section 4.3.1 captures this effect via the weight dependence factor controlled by the exponents (μ_+ , μ_-). During the potentiation process, the conductance tends to change rapidly at the beginning but saturates at the end of the process (see Fig. 4.8A). Similar behavior is also observed during the RESET. The potentiation and depression operations have, however, different dependencies on the device conductance. For high conductances, the potentiation increments are much smaller than the depression decrements. This asymmetry in the behavior can be further enhanced if the learning rates are different during the potentiation and depression operations. Similarly, it is reasonable to assume that for the binary synapses the evolution of the permanence

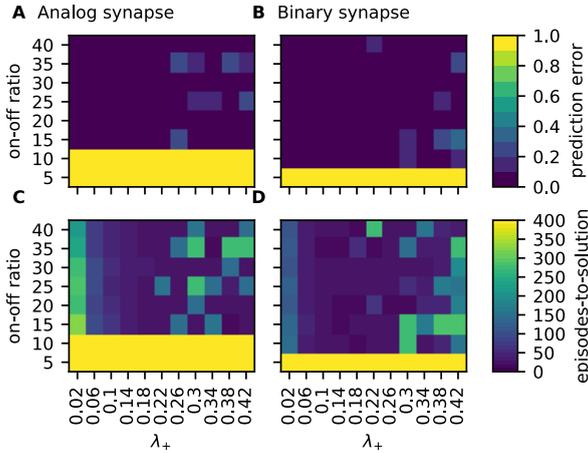


Figure 4.7: **Effect of the on-off ratio and the learning rate on the prediction performance.** Dependence of the prediction error and episodes-to-solution on the on-off ratio and the learning rate shown for the network with either analog (A,C) or binary synapses (B,D). Data depicts the median across an ensemble of 5 different network realizations. Parameters: depression learning rate $\lambda_- = \lambda_+/3$, weight dependence exponents $\mu_+ = 0.5$ $\mu_- = 0.5$, and variability amplitudes $\sigma_r = 0.03$, $\sigma_w = 0.01$. For remaining parameters see Table 7.4.

may exhibit a nonlinear dependence on the synaptic updates and an asymmetric behavior between the potentiation and depression dynamics (Fig. 4.8B).

Here, we first evaluate how the asymmetry in the learning rates between the potentiation and depression operations (λ_+ and λ_-) affects the prediction performance. To study this effect, we fix λ_+ and vary λ_- with the state dependence exponents μ_+ and μ_- being set to zero. The prediction error remains high if $\lambda_- \geq \lambda_+$ (see Fig. 4.9). In the spiking TM model, the potentiation operation is applied only when the postsynaptic spike follows after the presynaptic spike, in contrast, the RESET operation is applied every time the presynaptic neuron generates a spike. Therefore, for effective synaptic growth, the potentiation needs to be stronger than depression.

We assess, next, the prediction performance as a function of different weight dependence exponents for both potentiation and depression (μ_+ and μ_- , respectively, see Fig. 4.10). The results show that this latter has mild effects on the prediction error (see Fig. 4.10A,B). For larger values of μ_+ , the learning speed slows down as it takes longer for either the conductance or the permanence to reach their maximum values (see Fig. 4.10C,D). In the binary case, the steady-state permanence P^* may end up below the maturity threshold θ_P such that synapses can mature only due to the noise. The learning is therefore slowed down for large values of μ_+ or even not successful if the devices do not switch to the HCS. In the model, θ_P could be adjusted to P^* (similarly to adjusting θ_{dAP} to G^* in the analog synapse; see above). In this case, the learning in the analog and the binary networks may be similarly fast. In the physical device, however, the maturity threshold θ_P can hardly be changed.

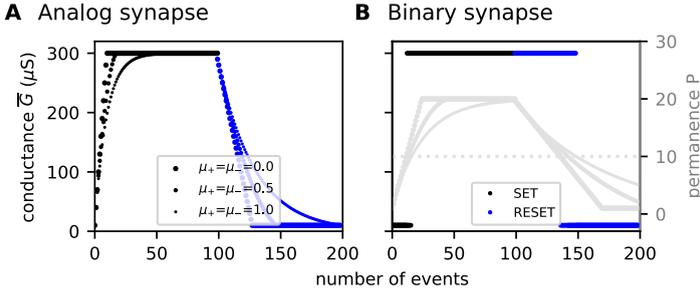


Figure 4.8: **An exemplary potentiation and depression curves depicting different levels of weight dependence updates.** Dependence of the conductance \bar{G} on the number of either SET (potentiation; black) or RESET (depression; blue) events as well as on different weight dependence exponents $\mu_+ = \mu_- = 0$ (large dot), $\mu_+ = \mu_- = 0.5$ (small dot), and $\mu_+ = \mu_- = 1$ (tiny dot) plotted for the analog (A) and the binary ReRAM models (B). In B, the permanence of the binary device is plotted in grey. Parameters: learning rates $\lambda_+ = 0.1$, $\lambda_- = \lambda_+/3$ (analog synapse), $\lambda_+ = 0.04$, $\lambda_- = \lambda_+/3$ (binary synapse), and variability amplitudes $\sigma_w = 0$, $\sigma_r = 0$. For remaining parameters see Table 7.4.

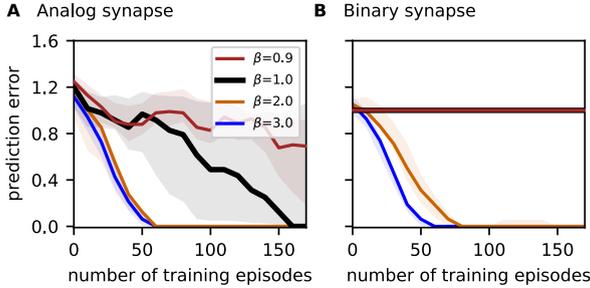


Figure 4.9: **Effect of the asymmetry in the learning rates on the prediction performance.** Dependence of the prediction error on the number of training episodes for different ratios (β) between the depression and potentiation learning rates ($\lambda_- = \lambda_+/\beta$), shown for the network with either analog (A) or binary (B) synapses. The potentiation learning rate λ_+ is fixed to 0.1 for the analog synapse and to 0.04 for the binary synapse. Curves and error bands indicate the median as well as the 5% and 95% percentiles across an ensemble of 5 different network realizations, respectively. For remaining parameters see Table 7.4.

4.3.3.3 Resilience of the model against variability

The resistive switching process of ReRAM devices involves the drift and diffusion of the ion vacancies. This phenomena is highly stochastic and shows considerable variation from device to device, and even from pulse to pulse within one device. In our work, we capture this effect by the read and write variability introduced in section 4.3.1. The influence of the read and write variability on the conductance curves are illustrated for both the analog and binary synapses in Fig. 4.11. For different trials, the write variability results in different conductance trajectories as a function of the

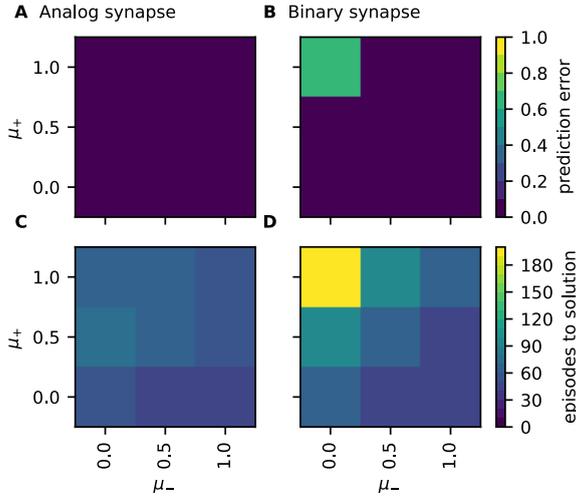


Figure 4.10: **Effect of the weight dependence update on the prediction performance.** Dependence of the prediction error and episodes-to-solution on the weight dependence exponents for both potentiation and depression (μ_+ and μ_-) shown for the networks with either analog (A,C) or binary synapses (B,D). Data depicts the median across an ensemble of 5 different network realizations. Parameters: learning rates $\lambda_+ = 0.1$, $\lambda_- = 0.03$ (analog synapse), $\lambda_+ = 0.04$, $\lambda_- = 0.01$ (binary synapse), and variability amplitudes $\sigma_w = 0.01$, $\sigma_r = 0.03$. For remaining parameters see Table 7.4.

applied potentiation or depression events. The read variability, on the other hand, causes only a jitter in the conductance curves.

To study how the variability influences the prediction performance, we assess the prediction error and episodes-to-solution for different magnitudes of the read and write noise σ_r and σ_w , respectively. Both networks with either analog or binary synapses allow similar read and write noise levels, with the binary synapse being slightly more resilient toward the read noise (Fig. 4.12A,B). In both cases, the write noise is more detrimental as it accumulates across the different learning episodes and can therefore have a higher impact on the learning performance. The read noise tends to average out as it is independent across the learning episodes. Concerning the learning speed, the number of episodes-to-solution is similar for the different noise levels where the learning is successful (Fig. 4.12C,D). Overall, increasing the read or write noise beyond what is acceptable leads to a spurious activation of the dAPs, i.e., predictions, and a decline in the prediction performance.

4.3.3.4 Robustness of the model against synaptic failure

When operating ReRAM devices, they risk failing by getting trapped in the HCS state even after applying voltage pulses with the appropriate magnitude across them (Kumar et al., 2017). To study how synaptic failure affects the prediction performance, we first train the network till it reaches zero prediction error (after 150 episodes in Fig. 4.13). Then, the conductance of a random fraction of synapses is set to the HCS. We quantify the level of synaptic failure by the ratio between the number of failed

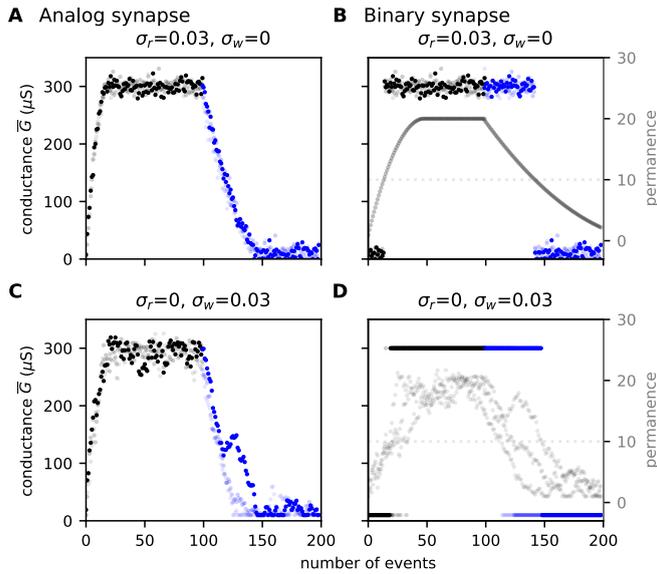


Figure 4.11: **An exemplary potentiation and depression curves in the presence of read and write variability.** Dependence of the conductance \bar{G} on the number of either potentiation (black) or depression (blue) events in the presence of read noise ($\sigma_r = 0.03$, **A,B**) or write noise ($\sigma_w = 0.01$, **C,D**) plotted for the analog (B) and the binary ReRAM models (C). In B and D, the permanence of the binary device is plotted in gray. Parameters: learning rates $\lambda_+ = 0.1$, $\lambda_- = \lambda_+/3$ (analog synapse), $\lambda_+ = 0.04$, $\lambda_- = \lambda_+/3$ (binary synapse), and weight dependence exponents $\mu_+ = 0.5$, $\mu_- = 0.5$. For remaining parameters see Table 7.4.

synapses and the total number of existing synapses. In the spiking TM model, a neuron may falsely generate a dAP if a sufficient number of its synapses are randomly switched to the HCS ($\sim \theta_{\text{dAP}}/G_{\text{max}}$, where W_{max} is the maximum conductance and θ_{dAP} is the dAP threshold). This may result in generating false positives and thus an increase in the prediction error. This is confirmed by our results presented in Fig. 4.13A,B. At up to 15% synaptic failure no impact is observed on the prediction performance (Fig. 4.13A,B). At greater than 15% synaptic failure the performance of the network declines and does not recover.

In a second experiment, instead of turning a selection of random synapses to the HCS state, we turn them to the LCS state. For the different levels of synaptic failures, the performance of the network initially declines. Due to the failing synapses, which are stuck at the LCS state, the neurons in certain subpopulations do not receive enough current and are thus not able to generate dAPs, i.e., make predictions. After further training episodes, the prediction errors converge back to zero as the network relearns using other synapses (Fig. 4.13C,D).

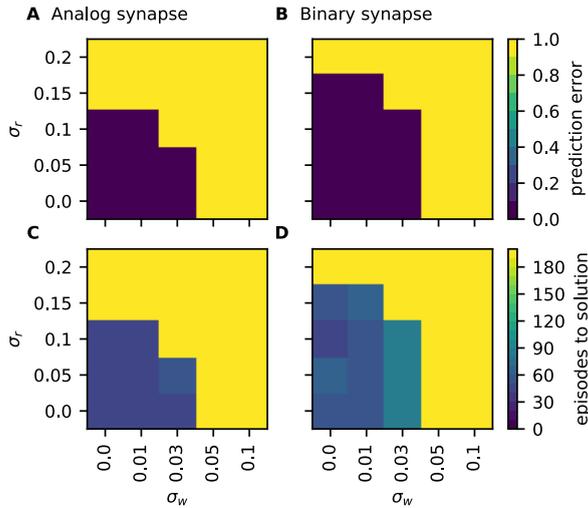


Figure 4.12: **Effect of the variability on the prediction performance.** Dependence of the prediction error and episodes-to-solution on the read and write variability σ_r , σ_w , shown for the networks with either analog (A,C) or binary synapse (B,D). Data depicts the median across an ensemble of 5 different network realizations. Parameters: learning rates $\lambda_+ = 0.1$, $\lambda_- = \lambda_+/3$ (analog synapse), $\lambda_+ = 0.04$, $\lambda_- = \lambda_+/3$ (binary synapse), and weight dependence exponents $\mu_+ = 0.5$, $\mu_- = 0.5$. For remaining parameters see Table 7.4.

4.4 DISCUSSION

4.4.1 Summary

In this work, we demonstrate that the learning rules of the spiking Temporal Memory (spiking TM) model proposed in chapter 2 can be realized using memristive dynamics. We investigate this for a particular type of memristive device known as VCM ReRAM. To this end, we show that the spiking TM retains high prediction performance for a broad range of on-off ratios and learning rates. The model is resilient toward the write and read variability as well as the dependence of the synaptic updates on the weight. Moreover, our results show that the VCM-type ReRAM device can be operated either in the binary or the gradual switching regime without performance loss. This is in line with the original spiking TM implementation, which shows that the learning rule can either be implemented using structural plasticity where the weight abruptly changes between two levels or a conventional form of STDP where the weight gradually changes until it saturates. This suggests that the intrinsic dynamics of the VCM ReRAM capture not only synaptic properties of biological synapses such as the variability and the dependence of the synaptic updates on the weight but also can implement known forms of plasticity in the neuroscientific literature. Our study, therefore, ties further principles from the biological and the memristive synapses.

Ultimately, the goal is to implement the spiking TM model on a standalone neuromorphic chip. In this chapter, we only investigate how the intrinsic properties of the memristive device affect the learning in the spiking TM. In a neuromorphic chip based on arrays of memristive devices, other factors should be taken into accounts

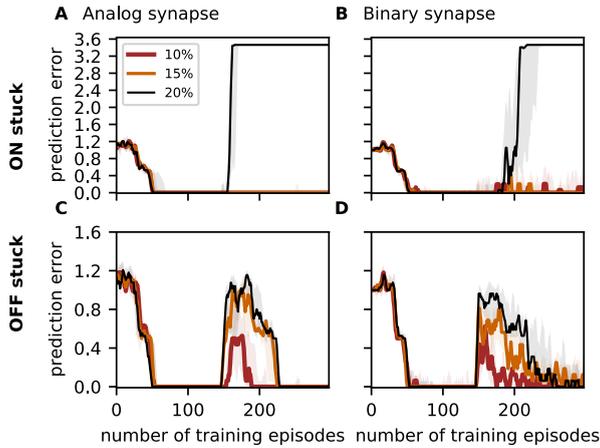


Figure 4.13: **Effect of synaptic failure on the prediction performance.** Dependence of the prediction error on the number of training episodes and different levels of synaptic failure (red 10%, orange 20%, black 30%) shown for both the analog (A, C) and the binary synapse (B, D). We implement the synaptic failure by fixing a random selection of synapses to be stuck at the HCS state (ON stuck; A and B) or stuck at the LCS state (OFF stuck; C and D). The synaptic failure is introduced after episode 150. Curves and error bands indicate the median as well as the 5% and 95% percentiles across an ensemble of 5 different network realizations, respectively. Same parameters as in Fig. 4.3.

such as the interaction between the crossbar array and the peripheral circuitry. In the study of the next chapter, we come up with a complete design of the neuromorphic circuit implementing the different components of the spiking TM.

4.4.2 Relationship to previous models

There are a number of biologically motivated sequence learning models that are closely related to the spiking TM, such as the self-organizing recurrent neural network model (SORN, Lazar et al., 2009). Recent work incorporated memristive dynamics into the synapses and neurons of the SORN model and showed that it retains successful performance (Payvand et al., 2022). The authors studied the role of variability and showed that it can improve prediction performance. However, the other memristive nonidealities were not investigated systematically.

The work by Doevenspeck et al. (2018) is another study that implements a sequence learning model using a memristive crossbar array. The memristive array consists of 1T1R elements and forms recurrent connections between the neurons. In contrast to our work, the model by Doevenspeck et al. (2018) employs a complex look-ahead algorithm to learn context-dependent sequences. Therefore, it disrupts the temporal locality of learning and makes the realization of online learning difficult.

4.5 CONCLUSION

ReRAM devices were suggested as promising synaptic elements in neuromorphic hardware due to their scalability and energy efficiency. In this work, we identified that a type of ReRAM device operated in either an analog or binary fashion can be used to implement well-known plasticity models from the neuroscientific literature. We demonstrate that both the analog or the binary ReRAM switching dynamics with their broad range of characteristics can be used as synaptic elements in the biologically inspired spiking TM model.

AN ALGORITHM FOR SEQUENCE LEARNING IN A MEMRISTIVE CROSSBAR ARRAY

This chapter is based on the following preprint:

Siegel, S., Bouhadjar, Y., Tetzlaff, T., Waser, R., Dittmann, R., & Wouters D. J. (2023). System model of neuromorphic sequence learning on a memristive crossbar array. Under review.

Author contributions: All authors contributed to the conception of the work. Sebastian Siegel conceived the overall circuit architecture and the detailed operation mechanism, performed the simulations, analyzed and visualized the data, and wrote the first draft of the manuscript. Younes Bouhadjar contributed to the discussions on the algorithmic adaptations for a neuromorphic hardware implementation of the spiking TM model. Tom Tetzlaff and Dirk Wouters overviewed the work and were involved in the scientific discussions. All authors reviewed the manuscript. In the following chapter, Younes Bouhadjar adapted and rewrote certain parts of the manuscript.

5.1 INTRODUCTION

In the previous chapter, we show that the spiking TM model retains successful operation when incorporating memristive dynamics in its synapse model. Here, we further assess the feasibility of implementing the model on memristive-based neuromorphic hardware. We shortly present a derived full circuit implementation centered around a memristive crossbar array. In the context of this thesis, we focus in particular on presenting the architectural and algorithmic differences and similarities between the original description of the spiking TM model and its hardware implementation (see Table 5.1). For practical details about the implementation, we refer the reader to (Siegel et al., 2023).

5.2 BASIC ARCHITECTURAL CONCEPTS

5.2.1 *Network structure*

The connection matrix of the spiking TM network is realized using a memristive crossbar array (Fig. 5.1A). Each synapse is represented by a memristive device (1R) in series with a transistor (1T). The neurons are subdivided into subpopulations, where each group of excitatory neurons shares an inhibitory neuron and the same source of external input. The terminals of each 1T1R circuit are attached to a horizontal input line and to a vertical output line. Neurons send pulses to the horizontal lines, receive back inputs from the vertical lines, and can further control the transistors gates. Recurrent connections of a neuron to itself are not realized.

5.2.2 *Neuron functionality*

In contrast to the original spiking TM algorithm, where the addition of the input currents is performed by the neuron model, in the model presented here, they are physically calculated in the array (using Kirchoff’s current law). Each vertical line can be seen as a single dendrite connected to the respective neuron, which is populated with a number of synaptic connections (i.e., memristive devices). This grouping of connections in the array, can result in spurious signals (e.g. so-called sneak currents Xia and Yang, 2019), but on the other hand, allows for efficient physical inference.

All other functionalities such as the generation of the spikes and the dendritic action potential (dAP) are implemented in the neuron blocks. The neurons also contribute to the implementation of the “synaptic” plasticity (see below). While it is possible to map their required functionalities into an analog electronic circuit, we here opt for implementation on a digital peripheral circuitry.

5.2.3 *Synaptic functionality*

As explained in section 4.3.1, a memristive device cannot emulate the complete synapse functionality. Especially, it cannot implement self-learning rules such as STDP, which is only enabled by a specific controller-memory function (see Fig. 4.2). However, it is not practically possible to include such a controller-memory function with each memristor device inside the array. A much more efficient implementation is to foresee that function outside the array, where it can be also combined for the different memristors in a row or column. In the proposed circuit, it is realized in

the functionality of the “neuron” blocks. The externalizing of this controller function, however, requires external row/column selection possibilities, and hence the use of a transistor selector for each memristor device resulting in a 1T1R array.

Summary		
	Spiking TM model	Crossbar implementation
Populations	excitatory neurons (\mathcal{E}), inhibitory neurons (\mathcal{I}), external spike sources (\mathcal{X}); \mathcal{E} and \mathcal{I} composed of M disjoint subpopulations \mathcal{M}_k and \mathcal{I}_k ($k = 1, \dots, M$)	same organization
Connectivity	<ul style="list-style-type: none"> • sparse random connectivity between excitatory neurons (plastic) • local recurrent connectivity between excitatory and inhibitory neurons (static) 	<ul style="list-style-type: none"> • all-to-all connectivity possible via memristive array, self-connections are not present (not formed) • inhibitory neuron circuit recurrently connected with the respective subpopulation
Neuron model	<ul style="list-style-type: none"> • excitatory neurons: leaky integrate-and-fire (LIF) with nonlinear input integration (dendritic action potentials) • inhibitory neurons: leaky integrate-and-fire (LIF) 	<ul style="list-style-type: none"> • excitatory neuron: linear integrate-and-fire dynamics without leak, reset after each time step • inhibitory neuron: digital counter emitting inhibition signal upon reaching a threshold
Synapse model	exponential postsynaptic currents (PSCs)	memristive conductances
Plasticity	homeostatic spike-timing-dependent plasticity in excitatory-to-excitatory connections	Hebbian learning and homeostasis controlled implemented by the peripheral neural circuitry

Table 5.1: Architectural and algorithmic differences between the original Spiking TM model and proposed hardware implementation. Modified from (Siegel et al., 2023).

5.3 ALGORITHMIC ADAPTATIONS

SEPARATION READ AND UPDATE PHASE The framework proposed here employs a global clock signal. The neurons receive alternating clock cycles. Each cycle of the clock encompasses a READ and an update phase (Fig. 5.1B,C), which corresponds to a single time step, i.e., a time step is equivalent to the presentation of a sequence element. As mentioned earlier, the neuron block contributes to implementing the learning rules. During the READ phase, the presentation of an external stimulus causes a neuron to emit a spike. This is translated into a rectangular voltage pulse, which is fed to the respective input of the horizontal line of the memristive array. Based on the current at the array outputs, the postsynaptic neurons generate a spike (or become predictive), and the controller decides on the update following the local learning rules:

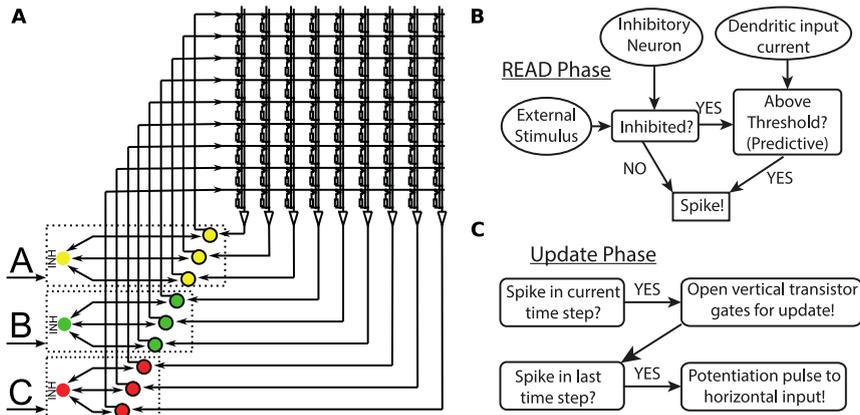


Figure 5.1: **System overview.** **A)** Sketch of the memristive array and the peripheral circuit. The peripheral circuit is composed of neurons ordered in subpopulations. Each subpopulation shares an inhibitory neuron and an external stimulus input. The neurons connect to the horizontal input lines of the memristive array and receive back connections through the vertical lines. **B)** Neuron functionality in the READ phase. Neurons generate spikes if they receive an external stimulus on the condition that they are not inhibited. The inhibition can be bypassed if the neuron is predicted. **C)** Neuron functionality in the Update phase. A presynaptic neuron that generates a spike in the current READ phase opens its vertically connected transistor gates for an update. If in addition, the postsynaptic neuron generates a spike in the next time step, it sends a potentiation pulse to its horizontal input line. A potentiation is caused only where a potentiation pulse coincides with an open transistor gate. Modified from (Siegel et al., 2023).

- **Hebbian learning:** in the hardware implementation, the read pulse polarity and amplitude are selected so that during the read phase all devices connected are slightly depressed. Only when there follows a postsynaptic spiking event, a (larger) potentiation pulse is applied.
- **Homeostasis:** in contrast to the original spiking TM algorithm, where the homeostatic function regulates the synaptic weight based on the activity of the neuron, here, in contrast, it allows each neuron to become predictive only once per episode.

GENERATION OF PREDICTION The neurons make predictions by generating dAPs. In contrast to the spiking TM algorithm, where the dAP is reset after a certain time period, here it is reset at the end of the READ phase in each step. This simplifies the circuit implementation as we no longer need an additional timer for resetting the dAP.

WINNER TAKE ALL MECHANISM If a certain number of neurons become predictive in the same subpopulation, they activate the inhibition, which hinders further neurons from becoming predictive. This is different compared to the original spiking TM algorithm, where the inhibitory neuron is activated by the somatic spikes and not by the subthreshold currents associated with the prediction, i.e., the dAP. This

allows for a more robust circuit implementation of the winner take all mechanism as the dAP is long-lasting and easier to detect compared to the somatic spikes.

TIMESCALES In contrast to the original spiking TM model, where the sequence processing speed is influenced by intrinsic timescales of neurons and synapses, which are in the range of milliseconds, in the hardware-adapted model, the speed of processing is determined by the period of the clock. The lower bound of the period depends on the intrinsic timescales of the electronic circuit, which is in the range of nanoseconds. The upper bound can be adjusted flexibly.

SIGNALS The signals in the hardware circuit are not spikes, but voltage pulses.

5.4 MEMRISTIVE SYNAPSE MODEL

In chapter 4, we investigate “extreme” cases of ReRAM switching behavior, i.e., analog or abrupt switching devices modeled by Eq (4.11) and Eq (4.9). In this work, ReRAM devices are simulated using a physically based device compact model. While this model reflects well the behavior of real devices, switching is typically intermediate between these two boundary cases (analog and binary). Chapter 4 shows that over a broad parameter range, both types of switching modes result in comparable successful learning results. It is expected therefore that the behavior of the device employed here to yield comparable performance (which is verified by the results below).

5.5 PREDICTION PERFORMANCE RESULTS

In this section, we assess whether the proposed hardware architecture can learn the high-order sequences $\{A,B,C,D\}$ and $\{E,B,C,F\}$ and compare results with the prediction performance of the original spiking TM model. The hardware architecture employs a 36-by-36 array using 36 excitatory neurons. These neurons are organized into six subpopulations with one shared inhibitory neuron each representing the characters “A” to “F”. In the array, 80% of the synapses are realized. Implementation and simulations of the hardware architecture are performed using the software Cadence. In this study, we downscaled the size of the original spiking TM model such that it has the same network size and level of sparsity as the hardware-adapted model.

To quantify the sequence prediction performance, we repetitively stimulate the network using the same set of sequences $\{A,B,C,D\}$, $\{E,B,C,F\}$ and assess the prediction performance (see Fig. 5.2). To ensure the performance results are not specific to a single network, the evaluation is repeated for a number of randomly instantiated network realizations with different initial connectivities. After each new network instantiation, the initial prediction error and false negatives are at 1, and the false positive is at zero (see Fig. 5.2). With an increasing number of training episodes, the prediction error as well as the false negative and positive of both the original spiking TM network and its hardware-adapted algorithm decrease and converge to zero. The original spiking TM algorithm converges slightly faster compared to the hardware-adapted algorithm.

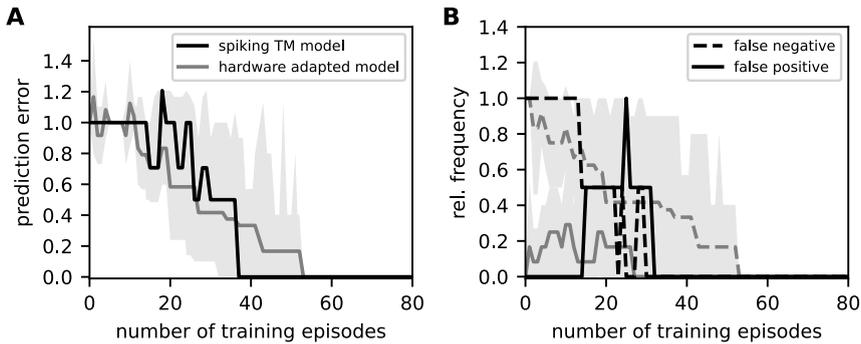


Figure 5.2: **Sequence prediction performance of spiking TM model and its hardware adapted model.** Dependence of the sequence prediction error (**A**), the false positive and false-negative rates (**B**) on the number of training episodes during repetitive stimulation with sequence $\{A,B,C,D\}$ and $\{E,B,C,F\}$. Curves indicate the median across an ensemble of 10 different network realizations. The parameters of the spiking TM model employed here are similar to the ones reported in Table 7.1 for sequence set I, except the following parameters: $n_E = 6$ (number of neurons per subpopulation), $M = 6$ (number of subpopulation), $\rho = 2$ (number of activate neurons per subpopulation), $p = 0.8$ (connection probability), $\lambda_+ = 0.02$, $\lambda_h = 0.004$ (potentiation and homeostatic learning rates). Data of the hardware-adapted model are from (Siegel et al., 2023).

5.6 DISCUSSION

In this work, we present the architectural and algorithmic adaptations of the spiking TM algorithm for implementation on neuromorphic hardware composed of a memristive crossbar array and peripheral circuitry. The crossbar array implements the synaptic weights and the peripheral circuitry implements the neuronal and the learning dynamics. We show that the proposed architecture can learn high-order sequences with similar prediction performance as the original spiking TM model.

The proposed hardware implementation around a memristive array offers interesting prospects as such arrays allow for efficient analog computation, e.g., during the read phase. It is expected that this hardware implementation can result in higher energy efficiency for obtaining the same results as the initial model. Interestingly, the array implementation enabled the incorporation of dendrites (vertical lines of the array), which were simplified to a point neuron in the original model. This gives the prospect for future work to expand this hardware model to multi-dendrite and possibly to hierarchical models.

Throughout this thesis, we first devise a continuous-time implementation of the TM algorithm by using biological ingredients such as spiking neurons. This translation facilitates porting the TM algorithm into neuromorphic hardware as a) it identified the biological ingredients that can implement the algorithm, which aided in finding its counterpart in the neuromorphic hardware, and b) it replaces the original learning rules, which are partially non-local, with hardware-friendly, fully local plasticity rules (see Plasticity dynamics in section 2.2.2). In true neuromorphic hardware, we have an ensemble of circuits each fully implementing the synapse or neuron functionalities, that are then connected in a particular network structure depending on the application. However, in our case, the hardware either does not emulate the complete functionality

(e.g. as memristive devices can emulate only part of the synapse functionality), or the function of different individual biological elements is grouped in a single (control) circuit. As a consequence, this chapter introduces the various algorithmic adaptations.

In a recent study, we demonstrate the algorithm's functionalities in a real hardware system (Siegel et al., 2023). The memristive crossbar array is fabricated using 130nm CMOS technology with HfO-based memristive devices (details on the process can be found in (Garbin et al., 2015)) and the peripheral circuitry is implemented in external digital hardware. The current implementation is limited to a small array size (8x8). Future studies could scale up the size and study the sequence learning capacity.

5.7 CONCLUSION

In this chapter, we present the algorithmic and architectural adaptations of the spiking TM algorithm for implementation on dedicated neuromorphic hardware employing memristive crossbar arrays to mimic the synaptic connectivity. The circuit-level simulations demonstrate the feasibility of our design choices. This work is a guide to future works aiming at integrating the architecture into a specialized electronic circuit, which can be embedded in systems at the edge.

The section [Limitation and extensions of the sequence learning model](#) is based on the following publication:

Bouhadjar, Y., Diesmann, M., Wouters, D. J., & Tetzlaff, T. (2022). Sequence learning, prediction, and replay in networks of spiking neurons. *PLOS Computational Biology* 18 (6), e1010233.

Other sections use short excerpts from the publication above ([Bouhadjar et al., 2022b](#)) or from ([Bouhadjar et al., 2022a](#), [Bouhadjar et al., 2022c](#)).

6.1 CONCLUSION

In this thesis, we propose a neuronal network architecture able to learn, predict, and replay sequences. This architecture provides an explanation of a number of biological mechanisms and principles observed in biology. Furthermore, we propose an efficient hardware implementation of these mechanisms using concepts from neuromorphic engineering. In the following, we summarize our key findings.

In chapter 2, we devise a continuous-time implementation of the temporal-memory (TM) component of the HTM algorithm (Hawkins and Ahmad, 2016), which is based on a recurrent network of spiking neurons with biophysically interpretable variables and parameters. The architecture learns high-order sequences by means of a structural Hebbian synaptic plasticity mechanism supplemented with a rate-based homeostatic control. In combination with nonlinear dendritic input integration and local inhibitory feedback, this type of plasticity leads to the dynamic self-organization of narrow sequence-specific subnetworks. These subnetworks provide the substrate for a faithful propagation of sparse, synchronous activity, and, thereby, for a robust, context-specific prediction of future sequence elements as well as for the autonomous replay of previously learned sequences. By strengthening the link to biology, our implementation facilitates the evaluation of the TM hypothesis based on experimentally accessible quantities. The continuous-time formulation in particular allows us to study how the sequence speed affects the prediction and replay performance. Furthermore, this study identifies biological mechanisms essential for sequence prediction and replay. It strengthens the concepts introduced in (Hawkins and Ahmad, 2016) and introduces new ones. In the following we provide a summary of these mechanisms:

- Learning and storage of sequences: sequences are represented by specific subnetworks embedded into the recurrent network. During the learning process, these subnetworks are carved out in an unsupervised manner by a form of structural Hebbian plasticity.
- Context specificity: learning of high-order sequences is enabled by a sparse, random potential connectivity, and by a homeostatic regulation of synaptic growth.
- Generation of predictions: neurons are equipped with a predictive state, implemented by a nonlinear synaptic integration mimicking the generation of dendritic action potentials (dAPs).
- Mismatch detection: only a few neurons become active if a prediction matches the stimulus. In our model, this sparsity is realized by winner-take-all (WTA) dynamics implemented in the form of inhibitory feedback. In the case of non-anticipated stimuli, the WTA dynamics cannot step in, thereby leading to a non-sparse activation of larger neuron populations.
- Sequence replay: an autonomous replay of learned sequences in response to a cue signal is enabled by increasing neuronal excitability.

In chapter 3, we present an extension of the spiking TM model able to replay sequences as a response to ambiguous cues following different strategies. The model in its original version recalls only the sequence shown most frequently during training. We introduce noise as a mechanism to achieve different replay strategies. To summarize our key findings:

- Uncorrelated noise, as used in many previous studies, averages out and is ineffective if sequence elements are represented by the activity of neuronal subpopulations, thereby leading to non-exploratory memory recall.
- Locally coherent noise permits explorative behavior without the need for unrealistically large noise amplitudes.
- Random stimulus locking to intrinsic spatiotemporal oscillations constitutes a natural form of such noise.
- Decision strategies can be adjusted by modulations in the amplitude or frequency of these oscillations.

In chapter 4, we investigate to what extent a type of memristive device known as ReRAM devices can replace biological synapses in the spiking TM model. In the first step, we implement a model of the ReRAM device in the neural simulator NEST. In the second step, we integrate the ReRAM device model into the spiking TM network. Our results demonstrate that successful learning is

- supported by both analog and binary ReRAM switching dynamics,
- obtained for a broad range of on-off ratios and learning rates,
- resilient with respect to device variability and synaptic failure, and
- robust with respect to weight dependence device updates.

In chapter 5, we describe what algorithmic and architectural changes of the spiking TM model are required to port it in a full electronic circuit architecture centered around a memristive crossbar array. The architecture is composed of a ReRAM crossbar array and peripheral circuitry. The ReRAM crossbar array implements the connectivity, where the synaptic elements are realized using a one-transistor-one-ReRAM-device (1T1R) structure. The peripheral circuitry implements the neuron model and facilitates the implementation of the learning rules. We assess our design using circuit-level simulations and show that the system can successfully learn context-dependent sequences.

6.2 OUTLOOK

6.2.1 Testable predictions of the sequence learning model

The spiking Temporal Memory (spiking TM) model developed in this thesis combines findings from both physiological-anatomical and behavioral data. It uses biological ingredients such as spiking neurons and local learning rules to explain how cortical networks predict and replay sequences in a context-dependent manner. It further explains how these networks replay sequences as a response to ambiguous cues. The model makes a number of testable predictions that can be verified in electrophysiological and behavioral experiments:

- Prediction and replay in the spiking TM model rely on the propagation of activity across subpopulations of neurons, where the number of spikes in each subpopulation is in the range ~ 10 – 20 . Detecting these patterns of activity using extracellular recording is not possible due to subsampling. It may be possible to detect such patterns using high-resolution calcium imaging (Grewe et al., 2010).
- The spiking TM network generates dendritic action potentials in a sparse subset of neurons to signal anticipated stimuli. The dendritic action potential results in a long-lasting depolarization of these neurons. In the case of correct predictions, the sparse subset of depolarized neurons generates somatic spikes. Otherwise, wrong predictions result in a non-sparse, synchronous somatic firing of neurons. These set of mechanisms and principles could be tested in an electrophysiological study using high-resolution calcium imaging.
- During the replay, the somatic depolarization due to a dendritic action potential can alone cause the neuron to generate a spike. We achieve this by increasing the excitability of the neurons, i.e., by decreasing the firing threshold. As suggested in chapter 2, the source of such an increase in excitability remains to be identified. Possible mechanisms could be a top-down signal delivered in the form of a neuromodulatory, attentional signal or propagating spatiotemporal oscillations (Atherton et al., 2015; Baluch and Itti, 2011; Noudoost et al., 2010; Buzsáki, 2006).
- The sequence replay speed in the spiking TM model is faster and independent of the sequence presentation speed during training. A number of studies have already demonstrated fast replay in both the neocortex and the hippocampus (Xu et al., 2012; Nádasdy et al., 1999). It remains to be determined whether the replay speed is completely independent of the training speed or if there exists a certain dependence.
- To foster context-dependent predictions, homeostatic control prevents the same neurons from generating predictions in different sequences. If it were possible to prevent the homeostasis mechanism, this would then increase the chance of overlap between the neuronal populations representing sequence elements in different contexts and may impair the prediction performance.
- In population-based encoding exploratory behavior is achieved by means of locally coherent noise. Random stimulus locking to spatiotemporal oscillations is a possible type of such noise: distant neurons see different phases, while close-by neurons see the same phase. At the stimulus onset, the neurons in the up

phase tend to fire earlier compared to neurons in the down phase. The stimulus onset is typically not locked to the oscillation and thus can be a mechanism for generating noise. In the biological system, if the neurons tuned to the different sequences are determined, it may be possible to identify the sequence the animal is about to recall by observing the phase of these neurons. This type of analysis is identical to the one used by Arieli et al. (1996).

- Sequence processing speed (i.e., interstimulus intervals) that permits successful learning has a lower and upper bound. Irrespective of the task complexity, the sequence prediction performance becomes poor for sequence speeds outside the optimal range. In addition, the number of training episodes required to reach optimal performance is proportional to the sequence speed, i.e., fast sequences with smaller interstimulus intervals can be learned faster, provided that the interstimulus interval is in the optimal range. The dependence of prediction performance on the processing speeds could be systematically tested in a behavioral study.

Additional testable predictions related to the original TM model are listed in (Hawkins and Ahmad, 2016).

6.2.2 *Limitation and extensions of the sequence learning model*

While the model is still fairly simplistic, it may provide the basis for a number of future extensions. Our results on the sequence processing speed revealed that the spiking TM model can process fast sequences with inter-stimulus intervals ΔT up to ~ 75 ms (see chapter 2). This range of processing speeds is relevant in many behavioral contexts such as motor generation, vision (saccades), music perception and generation, language, and many others (Mauk and Buonomano, 2004). However, slow sequences with inter-stimulus intervals beyond several hundreds of milliseconds cannot be learned by this model with biologically plausible parameters. This is problematic as behavioral time scales are often larger (Mauk and Buonomano, 2004; Paton and Buonomano, 2018). By increasing the duration τ_{dAP} of the dAP plateau, the upper bound for ΔT could be extended to 500 ms, and maybe beyond (Milojkovic, 2005). However, for such long intervals, the synaptic potentiation would be very slow, unless the time constant τ_+ of the structural STDP is increased and the depression rate λ_- is adapted accordingly. Furthermore, while our model explains the fast replay observed in the hippocampus and cortex, it is not able to learn an element-specific timing and duration of sequence elements (Dave and Margoliash, 2000; Louie and Wilson, 2001; Gavornik and Bear, 2014). This could be overcome by equipping the model with a working memory mechanism, which maintains the activity of the subpopulations for behaviorally relevant time scales (Maes et al., 2020; Tully et al., 2016).

In the current version of the model, the number of subpopulations, the number of neurons within each subpopulation, the number of dendritic branches per neuron, as well as the number of synapses per neuron are far from realistic (Hawkins and Ahmad, 2016). The number of sequences that can be successfully learned in this network is hence rather small. In addition, the current work is focusing on sequence processing at a single abstraction level, not accounting for a hierarchical network and task structure with both bottom-up and top-down projections. A further simplification in this work is that the lateral inhibition within a subpopulation is mediated by a single interneuron with unrealistically strong and fast connections to and from the pool of excitatory

neurons. In future versions of this model, this interneuron could be replaced by a recurrently connected network of inhibitory neurons, thereby permitting more realistic weights, and simultaneously speeding up the interaction between inhibitory and excitatory cells by virtue of the fast-tracking property of such networks (van Vreeswijk and Sompolinsky, 1998). Similarly, the external inputs in our model are represented by single spikes, which are passed to the corresponding target population by a strong connection, and thereby lead to an immediate synchronous spike response. Replacing each external input with a population of synchronously firing neurons would be a more realistic scenario without affecting the model dynamics. The external neurons could even fire in a non-synchronous, rate modulated fashion, provided the spike responses of the target populations remain nearly synchronous and can coincide with the dAP-triggered somatic depolarization (see Fig. 7.5). The current version of the model relies on a nearly synchronous immediate response to ensure that a small set of (~ 20) active neurons can reliably trigger postsynaptic dAPs, and that the predictive neurons (those depolarized by the dAPs) consistently fire earlier as compared to the non-predictive neurons, such that they can be selected by the WTA dynamics. Non-synchronous responses could possibly lead to a reliable generation of dAPs in postsynaptic neurons, but would require large active neuron populations (loss of sparsity) or unrealistically strong synaptic weights. The temporal separation between predictive and non-predictive neurons becomes harder for non-synchronous spiking. In future versions of the model, it could potentially be achieved by increasing the dAP plateau potential, and simultaneously equipping the excitatory neurons with a larger membrane time constant, such that non-depolarized neurons need substantially longer to reach the spike threshold. Increasing the dAP plateau potential, however, makes the model more sensitive to background noise (see below). Note that, in our model, only the immediate initial spike response needs to be synchronous. After successfully triggering the WTA circuit, the winning neurons could—in principle—continue firing in an asynchronous manner (for example, due the working-memory dynamics mentioned above). Similarly, long-lasting or tonic external inputs could lead to repetitive firing of the neurons in the TM network. As long as these repetitive responses remain nearly synchronous, the network performance is likely to be preserved.

In the predictive mode, the statistics of the spiking activity generated by our model are primarily determined by the temporal structure of the external inputs. Upon presentation of a sequence element, a specific subset of excitatory neurons fires a single volley of synchronous spikes. If the stimulus is predicted, this subset is small. The spike response is therefore highly sparse both in time and space, in line with experimental findings (Barth and Poulet, 2012). For simplicity and illustration, the sequences in this study are presented in a serial manner with a fixed order, and fixed inter-sequence and inter-element (inter-stimulus) intervals. As a consequence, the single-neuron spike responses are highly regular. The in-vivo spiking activity in cortical networks, in contrast, exhibits a high degree of irregularity (Shadlen and Newsome, 1998). A more natural presentation of sequences with irregular order and timing trivially leads to more irregular spike responses in our model. As long as the inter-stimulus intervals fall into the range depicted in Fig. 2.11, the model can learn and predict irregular sequences. Spiking activity in the cortex is not only irregular, but also fairly asynchronous in the sense that the average level of synchrony for randomly chosen pairs of neurons is low (Ecker et al., 2010; Renart et al., 2010). This, however, is not necessarily the case for any subset of neurons and at any point in time. It is well known that cortical neurons can systematically synchronize their

firing with millisecond precision in relation to behaviorally relevant events (see, e.g., (Riehle et al., 1997)). As demonstrated in (Schrader et al., 2008), synchronous firing of small subsets of neurons may easily go unnoticed in the presence of subsampling. The model proposed in this study relies on (near) synchronous firing of small subsets of neurons. In cases where the model processes large sets of sequences in parallel, this synchrony will hardly be observable if only a small fraction of neurons is monitored (see Fig. 7.4). After learning, different sequences are represented by distinct subnetworks with little overlap. Hence, the network can process multiple sequences at the same time with little interference between subnetworks. The model could even learn multiple sequences in parallel, provided there is no systematic across-trial dependency between the sequences presented simultaneously. We dedicate the task of testing these ideas to future studies. While the synchrony predicted by the TM model may hardly be observable in experimental data suffering from strong subsampling, the predicted patterns of spikes could be identified using methods accounting for both spatial and temporal dependencies in the spike data (Schrader et al., 2008; Quaglio et al., 2017, 2018). There are other factors that may contribute to a more natural spiking activity in extended versions of the model. First, equipping the model with a working memory mechanism enabling the learning of slow sequences and sequence timing (see above) would likely lead to sustained asynchronous irregular firing. Second, replacing the inhibitory neurons with recurrent networks of inhibitory neurons (see above) would generate asynchronous irregular activity in the populations of inhibitory neurons and thereby contribute to variability in the spike responses of the excitatory neurons. Third, the model proposed here may constitute a module embedded into a larger architecture and receive irregular inputs from other components. As shown in Fig. 7.5 and Fig. 7.6, the spiking activity and the prediction performance of the TM model are robust with respect to low levels of synaptic background activity, and, hence, membrane potential fluctuations reminiscent of those observed in vivo (DeWeese and Zador, 2006). For an increased level of noise, the learning speed decreases. For high noise levels leading to additional, non-task related background spikes, the dAP triggered plateau depolarization is overwritten, such that the WTA dynamics fail at selecting predictive neurons, ultimately leading to a loss of context specificity in the responses. Hence, the prediction performance degrades for large noise amplitudes. A potential application of introducing background noise is to allow the network to perform probabilistic computations (Jordan et al., 2019), such as replaying sequences in the presence of ambiguous cues presented in section 3.3.1.

Similar to the original TM model, the response of the population representing the first element in a sequence is non-sparse, indicating that the first sequence element is not anticipated and can therefore not be predicted. If a given first sequence element reoccurs within the same sequence (say, “A” in {A,B,A,C}) or in other sequences (e.g., in {D,E,A,F}), the non-sparse response of the respective population to a first sequence element leads to a simultaneous prediction of all possible subsequent elements, i.e., the generation of false positives. These false predictions would lead to a pruning of functional synapses as a response of the homeostatic regulation to the increased dAP activity. This could be overcome by three possible mechanisms: a) synaptic normalization avoiding excessive synapse growth (Turrigiano et al., 1998; Elliott, 2003), b) removing breaks between sequences, or c) sparse, sequence specific firing of subpopulations representing first elements. Results of applying the last mechanism are shown in Fig. 7.2, where dAPs are externally activated in random subsets of neurons in the populations representing first elements. In a more realistic hierarchical

network, a similar effect could be achieved by top-down projections from a higher level predicting sequences of sequences.

In the original model, synapses targeting silent postsynaptic cells are depressed, even if the presynaptic neuron is inactive. This pruning process, the freeing of unused synaptic resources, increases the network capacity while ensuring context sensitivity. According to the structural plasticity dynamics employed in our study, synapse depression is bound to presynaptic spiking, similar to other implementations of (non-structural) STDP (Morrison et al., 2007). As a consequence, strong connections originating from silent presynaptic neurons are not depressed (dark gray dots in Fig. 2.2D). This may complicate or slow down the learning of new sequences, and could be overcome by synaptic normalization.

For the dAP-rate homeostasis used in this study, the target dAP rate is set to one to make sure that each neuron contributes at most one dAP during each training episode. As a consequence, the time constant of the dAP-rate homeostasis is adapted to the duration of a training episode, which is in the range of a few seconds in this work. We are not aware of any biological mechanism that could account for such an adaptation. dAP-rate homeostasis is mediated by the intracellular calcium concentration, which, in turn, controls the synthesis of synaptic receptors, and hence, the synaptic strength. It is therefore known to be rather slow, acting on timescales of many minutes, hours, or days (Turrigiano and Nelson, 2004; Turrigiano, 2008). It is unclear to what extent the use of long homeostatic time constants and increased dAP target rates would alter the model performance. Alternatively, the dAP-rate homeostasis could be replaced by other mechanisms such as synaptic normalization.

6.2.3 *Prospects on neuromorphic computing*

The field of neuromorphic computing promises to provide efficient solutions to real world problems by taking inspiration from biology (Aimone, 2019). It strives to build systems with similar capabilities as the brain including continuous, online learning with efficient usage of energy and data resources. The progress on such an endeavor, however, has been slowed down due to a lack of understanding of the mechanisms underpinning neuronal computation. In this thesis, we follow a strategy similar to the one proposed by Bouhadjar et al. (2022) to tackle this issue. By combining knowledge of lower-level biological features such as dendritic computation and spiking neurons, with insights from the top cognitive function, we are able to explain a number of biological mechanisms and their role in computation. This allows us to devise a sequence learning model able to learn continuously in an energy-efficient manner. The model architecture is modular and allows the processing of data from different modalities. It may constitute a building block in future models aiming at learning tasks that require, for example, multisensory integration.

Brain-inspired algorithms run more efficiently on dedicated neuromorphic hardware. A prerequisite for such hardware's success is identifying hardware concepts and devices that can best implement neuroscientific mechanisms and principles. The type of neuromorphic hardware discussed in chapter 5 relies on analog components to implement biological neurons and synapses. These components can be faster than their biological counterparts, which foresee hardware that runs substantially faster than biology. Due to the shared similarities between this type of hardware and the brain, we could even use it as a platform to study neuroscientific questions. For example, to date, there are only speculations on how the brain contributes to sensible and

reliable behavior in the face of noisy and heterogeneous synapses. Throughout the work on neuromorphic systems, we can develop intuitions of how biological synapses counteract or even exploit these different characteristics for computation. Similar to biological synapses, ReRAM devices are characterized by a read and write variability. In the biological system, the read variability is analogous to the randomness in the vesicle release, and the write variability corresponds to the randomness in the density regulation of the postsynaptic receptors. So far, it is not clear how these different characteristics contribute to the learning dynamics in the biological system. Neuromorphic hardware can provide an environment where this question can be studied.

In this work, we show that the spiking TM model is resilient toward synaptic variability. Other works show that synaptic variability can even have a computational benefit. In probabilistic computing frameworks, variability is considered a prerequisite for probabilistic inference (Maass, 2014; Nefteci et al., 2016). It allows the system to explore the state space and come up with an estimate of how likely each solution is. Similarly, chapter 3 shows that the spiking TM model can replay sequences probabilistically by using noise. However, the study explores only additive noise. Future work can shed light on the possibility of using ReRAM type of noise, i.e., multiplicative noise, as an alternative. A major challenge in this respect, however, is to come up with a mechanism introducing correlated variability across neighboring synapses. From chapter 3, we know that this is necessary to avoid the averaging effect. Multiplicative noise suffers from averaging to the same extent as additive noise.

In artificial neural networks trained using gradient-based approaches, ReRAM nonidealities can severely undermine the overall performance (Fouda et al., 2020). Due to the ReRAM variability, devices can be hardly programmed to the desired state, and the asymmetry in the conductance change can affect the propagation of the gradient and lead to performance loss. Correcting for these nonidealities can be costly and may require additional circuitry (Chen et al., 2015b; Agarwal et al., 2016; Hong et al., 2018; Adnan et al., 2021). In the spiking TM and other brain-inspired self-organizing networks, the ReRAM characteristics are not a nuisance and can be even beneficial in certain applications. We know that the neuronal system carries out accurate computations despite its nonideal synaptic characteristics such as variability. This suggests the existence of biological principles accommodating for that, which we need to understand and port for a successful implementation of neuromorphic hardware. Spiking TM provides a step in this direction. It identifies a set of biological concepts that might be at the heart of brain processing capabilities. For instance, its highly sparse connectivity and activity are observed in biological networks and are essential for increasing the capacity of the system and decreasing energy consumption. Furthermore, it self-organizes by means of unsupervised learning rules without a teaching signal. It remains to be answered how these learning rules can be efficiently implemented in neuromorphic hardware. For example, in chapter 4, we stress that memristive devices are not able to mimic biological learning rules without the help of peripheral circuitry. Looking for ways to bring the learning rules closer to the hardware, i.e., implemented directly by the synaptic elements, would allow for energy efficient systems with online learning capabilities in edge applications.

SUPPORTING INFORMATION

7.1 SUPPORTING INFORMATION CHAPTER 2

7.1.1 *Model and simulation parameters*

Name	Value	Description
Network		
N_E	2100	total number of excitatory neurons
N_I	14	total number of inhibitory neurons
M	$A = 14$	number of excitatory subpopulations (= number of external spike sources)
n_E	$N_E/M = 150$	number of excitatory neurons per subpopulation
n_I	$N_I/M = 1$	number of inhibitory neurons per subpopulation
ρ	20	(target) number of active neurons per subpopulation after learning = minimal number of coincident excitatory inputs required to trigger a spike in postsynaptic inhibitory neurons
(Potential) Connectivity		
K_{EE}	420	number of excitatory inputs per excitatory neuron (EE in-degree)
p	$K_{EE}/N_E = 0.2$	probability of potential (excitatory) connections
K_{EI}	$n_I = 1$	number of inhibitory inputs per excitatory neuron (EI in-degree)
K_{IE}	n_E	number of excitatory inputs per inhibitory neuron (IE in-degree)
K_{II}	0	number of inhibitory inputs per inhibitory neuron (II in-degree)
Excitatory neurons		
$\tau_{m,E}$	10 ms	membrane time constant
$\tau_{ref,E}$	10 ms	absolute refractory period
C_m	250 pF	membrane capacitance
V_r	0.0 mV	reset potential
θ_E	20 mV (predictive mode), 5 mV (replay mode)	somatic spike threshold
I_{dAP}	200 pA	dAP current plateau amplitude
τ_{dAP}	60 ms	dAP duration
θ_{dAP}	59 pA (predictive mode), 41.3 pA (replay mode)	dAP threshold
Inhibitory neurons		
$\tau_{m,I}$	5 ms	membrane time constant
$\tau_{ref,I}$	2 ms	absolute refractory period
C_m	250 pF	membrane capacitance
V_r	0.0 mV	reset potential
θ_I	15 mV	spike threshold

Continued next page

Name	Value	Description
Synapse		
γ	5	number co-active presynaptic neurons required to trigger a dAP in the postsynaptic neuron
W	12.98 pA	weight of mature EE connections (EPSC amplitude)
\tilde{J}_{IE}	0.9 mV (predictive mode) 0.12 mV (replay mode)	weight of IE connections (EPSP amplitude)
J_{IE}	581.19 pA (predictive mode) 77.49 pA (replay mode)	weight of IE connections (EPSC amplitude)
\tilde{J}_{EI}	-40 mV	weight of EI connections (IPSP amplitude)
J_{EI}	-12915.49 pA	weight of EI connections (IPSC amplitude)
\tilde{J}_{EX}	22 mV	weight of EX connections (EPSP amplitude)
J_{EX}	4112.20 pA	weight of EX connections (EPSC amplitude)
τ_{EX}	2 ms	synaptic time constant of EX connection
τ_{EE}	5 ms	synaptic time constant of EE connections
τ_{EI}	1 ms	synaptic time constant of EI connections
τ_{IE}	0.5 ms	synaptic time constant of IE connections
d_{EE}	2 ms	delay of EE connections (dendritic)
d_{IE}	0.1 ms	delay of IE connections
d_{EI}	{0.1, 0.2} ms	delay of EI connections (non-default value used in Fig. 2.11 and Fig. 2.12)
d_{EX}	0.1 ms	delay of EX connections
Plasticity		
λ_+	0.08 (sequence set I), 0.28 (sequence set II)	potentiation rate
λ_-	0.0015 (sequence set I), 0.0061 (sequence set II)	depression rate
θ_P	20	synapse maturity threshold
$P_{\min,ij}$	$\sim \mathcal{U}(P_{0,\min}, P_{0,\max})$	minimum permanence
P_{\max}	20	maximum permanence
$P_{0,\min}$	0	minimal initial permanence
$P_{0,\max}$	8	maximal initial permanence
τ_+	20 ms	potentiation time constant
z^*	1	target dAP activity
λ_h	0.014 (sequence set I), 0.024 (sequence set II)	homeostasis rate
τ_h	440 ms (sequence set I), 1560 ms (sequence set II)	homeostasis time constant
y_i	1	depression decrement
Δt_{\min}	4 ms	minimum time lag between pairs of pre- and postsynaptic spikes at which synapses are potentiated
Δt_{\max}	$2\Delta T$	maximum time lag between pairs of pre- and postsynaptic spikes at which synapses are potentiated

Continued next page

Input		
L	1	number of subpopulations per sequence element = number of target subpopulations per spike source
S	2 (sequence set I), 6 (sequence set II)	number of sequences per set
C	4 (sequence set I), 5 (sequence set II)	number of elements per sequence
A	14	alphabet length (total number of distinct sequence elements)
ΔT	$\{2, \dots, \mathbf{40}, \dots, 90\}$ ms	inter-stimulus interval
ΔT_{seq}	$\max(2.5\Delta T, \tau_{\text{dAP}})$	inter-sequence interval
ΔT_{cue}	80 ms	inter-cue interval
Simulation		
Δt	0.1 ms	time resolution
K	$\{\mathbf{80}, 100\}$	number of training episodes

Table 7.1: Model and simulation parameters. Parameters derived from other parameters are marked in gray. Bold numbers depict default values.

Name	Value	Description
columnDimensions (M)	280	number of columns
numColumnsPerElement (L)	20	number of columns per element
cellsPerColumn (n_E)	8	number of cells per column
initialPermanence (P_0)	[0.1, 0.3]	initial permanence
connectedPermanence (θ_P)	0.5	threshold at which a synapses is considered connected
minThreshold	15	if the number of immature (potential) synapses active on a segment is at least this threshold, it is said to be “matching” and is eligible for learning.
maxNewSynapseCount	40	the maximum number of synapses added to a segment during learning
permanenceIncrement (λ_+ , λ_h)	0.1	amount by which permanences of synapses are incremented during learning.
permanenceDecrement (λ_-)	0.3	amount by which permanences of synapses are decremented during learning.
activationThreshold (θ_{dAP})	15	if the number of active connected synapses on a segment is at least this threshold, the segment is said to be active.
predictedSegmentDecrement	0.01	amount by which permanences of synapses are decremented during learning.

Table 7.2: **Adapted parameters of the original TM model used for Fig. 2.9.** Parameter names match those used in the original simulation code (https://github.com/numenta/htmpapers/tree/master/frontiers/why_neurons_have_thousands_of_synapses). Gray parameter names are those used in the spiking TM model.

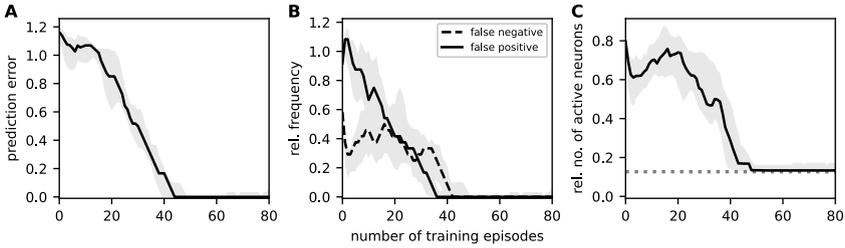


Figure 7.1: **Sequence prediction performance in the presence of conventional (non-structural) spike-timing-dependent plasticity (STDP).** Dependence of the sequence prediction error (A), the false-positive and false-negative rates (B), and the number of active neurons relative to the subpopulation size (C) on the number of training episodes for sequence set II. Curves and error bands indicate the median as well as the 5% and 95% percentiles across an ensemble of 5 different network realizations, respectively. All prediction performance measures are calculated as a moving average over the last 4 training episodes. In this experiment, structural STDP is replaced by conventional STDP, i.e., the permanences $P_{ij}(t)$ and P_{\max} in Eq (2.1) are replaced by the synaptic weights $J_{EE,ij}(t)$ and J_{\max} . The weights $J_{EE,ij}$ are restricted to the interval $[J_{\min,ij}, J_{\max}]$, and clipped at the boundaries. The minimal weights $J_{\min,ij}$ are randomly and independently drawn from a uniform distribution between $J_{0,\min}$ and $J_{0,\max}$. The performance characteristics shown here are comparable to those obtained with structural STDP (see Fig. 2.9 in Prediction performance). Parameters: $\Delta T = 40$ ms, $\lambda_+ = 0.43$, $\lambda_- = 0.0058$, $\lambda_h = 0.03$, $J_{0,\min} = 0$ pA, $J_{0,\max} = 2$ pA, $J_{\max} = 12.98$ pA. See Table 7.1 for remaining parameters.

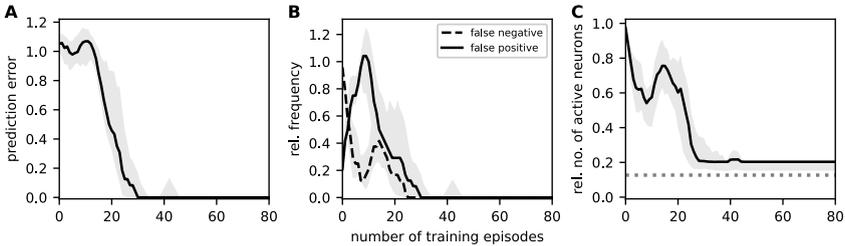


Figure 7.2: **Prediction performance for a sequence set with recurring first items.** Dependence of the sequence prediction error (A), the false positive frequency, the false negative frequency (B), and the number of active neurons relative to the subpopulation size (C) on the number of training episodes for a set of sequences $s_1 = \{B,D,I,C,H\}$, $s_2 = \{E,D,I,C,F\}$, $s_3 = \{F,B,C,A,H\}$, $s_4 = \{G,B,C,A,D\}$, $s_5 = \{E,C,I,H,A\}$, $s_6 = \{D,C,I,H,G\}$ with recurring first items. Curves and error bands indicate the median as well as the 5% and 95% percentiles across 5 different network realizations, respectively. As a solution to the issue discussed in Limitation and extensions of the sequence learning model concerning the recurring of first sequence elements in other sequences or within the same sequence, the dAPs are externally activated in a random subset of neurons in the populations representing first elements. Inter-stimulus interval $\Delta T = 40$ ms. All prediction performance measures are calculated as a moving average over the last 4 training episodes. Parameters: $\Delta T = 40$ ms, $\lambda_+ = 0.39$, $\lambda_- = 0.0057$, $\lambda_h = 0.034$. See Table 7.1 for remaining parameters

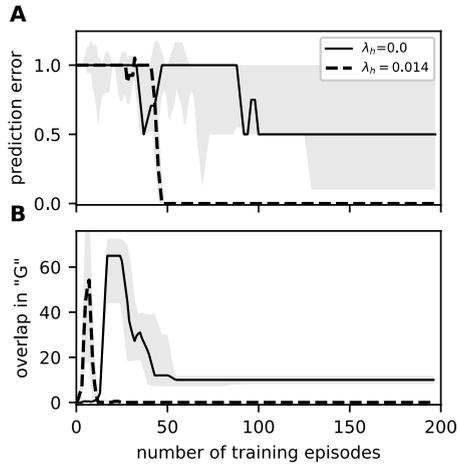


Figure 7.3: **Effect of the dAP-rate homeostasis on the prediction performance.** Dependence of the prediction error (A) and the overlap in the activation pattern between the neurons representing the sequence element “G” in the context of sequences {A,D,B,G,H,E} and {F,D,B,G,H,C} (B) on the number of training episodes explored for two values of the homeostasis rate (λ_h). Curves and error bands indicate the median as well as the 5% and 95% percentiles across 5 different network realizations, respectively. Disabling the homeostasis control ($\lambda_h = 0.0$) increases the overlap in the “G” activation pattern, which leads to a loss of context specificity and hence an increase in the prediction error (see [Sequence learning and prediction](#)). The parameters of the plasticity are similar to the ones reported in Table 7.1 for the sequence set I.

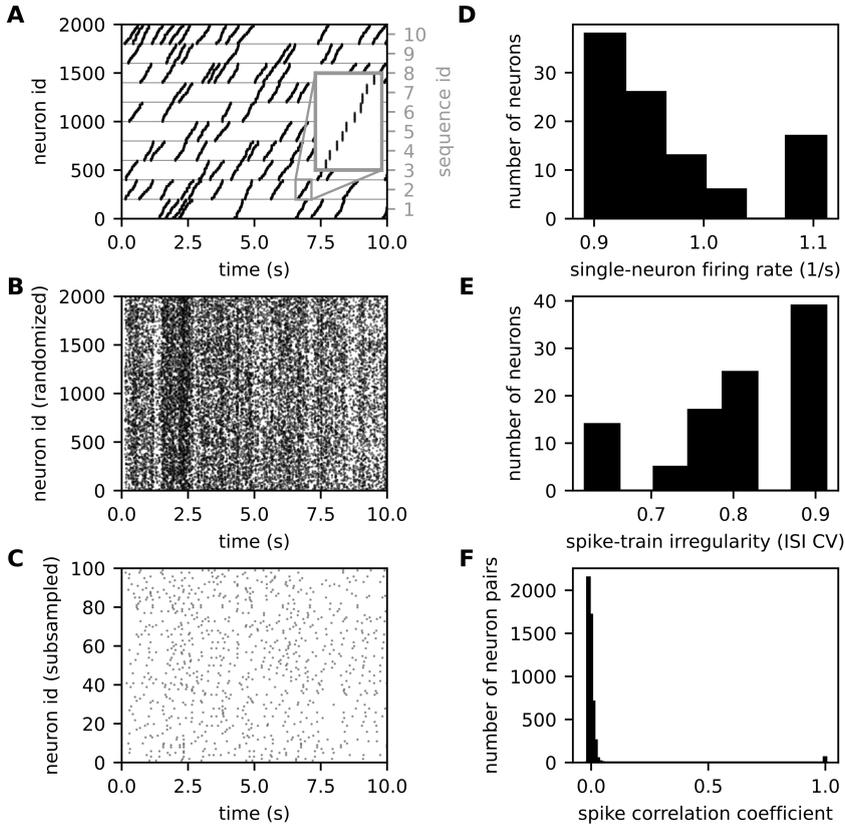


Figure 7.4: **Asynchronous irregular firing in a (hypothetical) network processing multiple sequences in parallel.** **A:** Artificial spike data mimicking activity of a TM network processing $S = 10$ sequences in parallel. Each sequence (right y-axis) is processed by a distinct subnetwork of 200 neurons, each composed of $C = 10$ subpopulations. The horizontal gray lines separate the different subnetworks. Upon activation of a sequence element, $\rho = 20$ neurons in the corresponding subpopulation synchronously fire a spike. Individual sequences are activated independently with rate 1 s^{-1} at random times (Poisson point process with 200 ms deadtime). Inter-element intervals $\Delta T \sim \mathcal{U}(10\text{ ms}, 80\text{ ms})$ are randomly drawn from a uniform distribution (cf. Fig. 2.11). The inset depicts a magnified view of a single activation of sequence 2. **B:** Same data as in A after random permutation of neuron identities. **C:** Spiking activity of a random subset of 100 neurons depicted in panel B. **D–F:** Distributions of single-neuron firing rates (D), inter-spike-interval variation coefficients (E; ISI CV), and spike-count correlation coefficients (F; binsize 10 ms) obtained from subsampled data shown in panel C for a total simulation time of 100 s (mean rate = 1 spikes/s, mean ISI CV = 0.8, mean correlation = 0.01). The data and analysis results shown here demonstrate that i) irregular sequence activation translates into irregular spiking, and ii) subsampling and the absence of prior knowledge of the network structure hide synchrony (but note the tiny peak at 1.0 in the distribution of correlation coefficients). The combination of both effects hence leads to asynchronous irregular firing, reminiscent of in-vivo cortical activity.

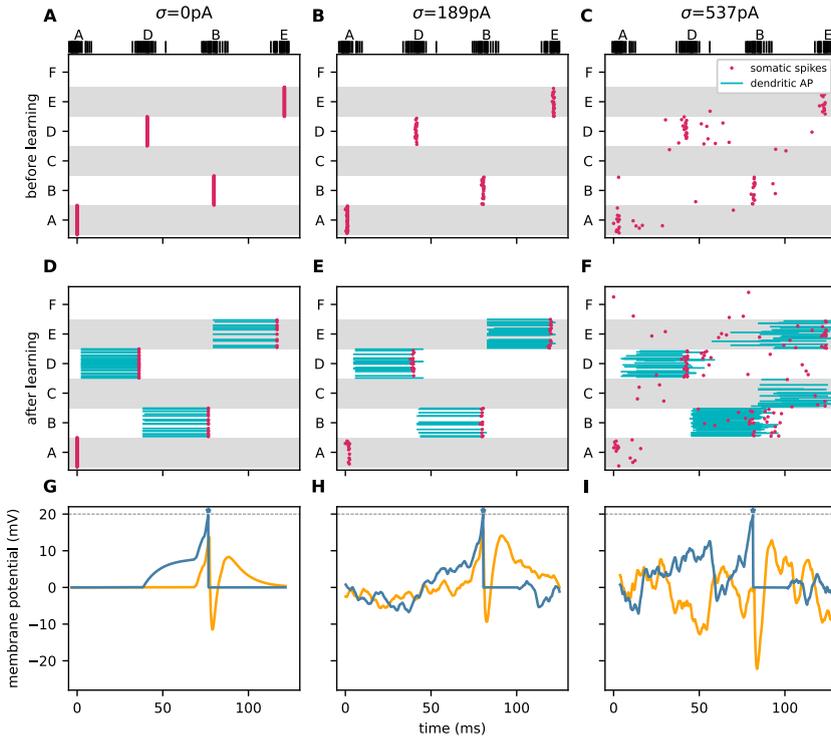


Figure 7.5: **Effects of background noise and non-synchronous stimulation on network activity.** **A–F)** Spiking activity before (panels A–C; 1st learning episode) and after learning sequence set I (panels D–F; 600th learning episode) in response to a presentation of sequence {A,D,B,E} without background noise (left) and in the presence of moderate (middle) or high synaptic background noise (right). External inputs are presented in the form of dispersed volleys of 50 spikes (black vertical bars at the top). Each of these spikes triggers an exponential synaptic input current in the target neurons with amplitude 134 pA and time constant 1 ms. Spike times in each spike volley are randomly drawn from a Gaussian distribution (width 4 ms), centered on the stimulus time. In each trial, all stimulated neurons receive the same realization of the Gaussian spike packet. Red dots and blue horizontal lines mark somatic spikes and dAPs, respectively. For clarity, only a fraction of 50% of excitatory neurons and external spikes are shown. Background noise to each excitatory neuron is provided in the form of balanced excitatory and inhibitory synaptic inputs, generated by distinct uncorrelated Poissonian spike sources (total rate per source $\nu = 10000 \text{ s}^{-1}$). Background synapses are modeled as exponential postsynaptic currents (time constant $\tau_B = 1 \text{ ms}$) with amplitudes $J = 0 \text{ pA}$ (left), 60 pA (middle), and 170 pA (right) for excitatory inputs, and $-J$ for inhibitory inputs, respectively. The mean background input $\mu = \tau_B \nu (J - J_-) = 0$ to each neuron vanishes due to the asymmetry in excitatory and inhibitory synaptic weights. The variance $\sigma^2 = \tau_B \nu J^2$ of the synaptic background current is modulated by adjusting the synaptic weight J (left: $\sigma = 0 \text{ pA}$, middle: $\sigma = 190 \text{ pA}$, right: $\sigma = 537 \text{ pA}$). **G,H,I)** Membrane potential traces of two neurons in the excitatory subpopulation “B” during the same time interval depicted in panels D–E for three noise levels $\sigma = 0 \text{ pA}$ (G), 190 pA (H), and 537 pA (I). One of the selected neurons (blue) is participating in the sequence, i.e., it generates a dAP and a somatic spike in response to sequence elements “D” and “B”. The other neuron (orange) is not part of the sequence. The horizontal dashed lines and blue stars mark the threshold θ_E and the times of somatic spikes, respectively. Parameters: $\Delta T = 40 \text{ ms}$, $\lambda_+ = 0.05$, $\lambda_- = 0.001$, $\lambda_h = 0.01$, $W = 23.6 \text{ pA}$, $\Delta t_{\min} = 20 \text{ ms}$, $\tau_{\text{dAP}} = 40 \text{ ms}$, $\tau_{\text{ref,I}} = 20 \text{ ms}$, $J_{\text{EI}} = -9686.62 \text{ pA}$. See Table 7.1 for remaining parameters.

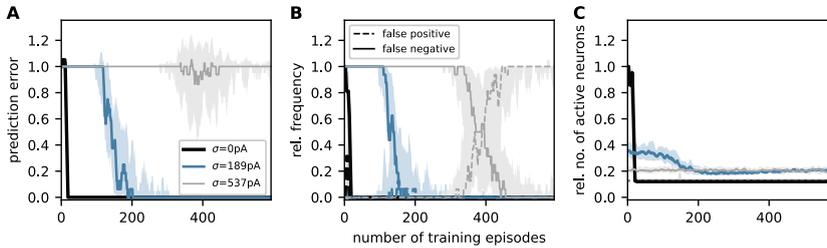


Figure 7.6: **Effects of background noise and non-synchronous stimulation on prediction performance and sparsity for sequence set I.** Dependence of the sequence prediction error (**A**), the false positive and false negative rate (**B**), and the sparsity (number of active neurons relative to the subpopulation size, **C**) on the number of training episodes for three different noise amplitudes $\sigma = 0$ pA (black), 190 pA (blue), and 537 pA (gray). See caption of Fig. 7.5 for details on the implementation of external inputs and background noise. Curves and error bands indicate the median as well as the 5% and 95% percentiles across 5 different network realizations, respectively. All prediction performance measures are calculated as a moving average over the last 4 training episodes. Same parameters as in Fig. 7.5.

Video 7.1: **Time resolved visualization of the learning dynamics:** Network activity (top) and connectivity (bottom) of the network during one learning episode. Each frame corresponds to a new training episode. In each learning episode, each of the two sequences {A,D,B,E} and {F,D,B,C} is presented once (black arrows in the top panel). **Top panel:** Red dots and blue bars mark spike and dAP times for each neuron. Neurons are sorted according to stimulus preference (vertical axis). **Bottom panel:** Network connectivity before learning (left) and during the current training episode (right). Light gray and black dots represent immature and mature connections, respectively, for each pair of source and target neurons (sorted according to stimulus preference; see [Sequence learning and prediction](#)).

7.2 SUPPORTING INFORMATION CHAPTER 3

7.2.1 Model and simulation parameters

Name	Value	Description
Network		
N_E	1200	total number of excitatory neurons
N_I	1	total number of inhibitory neurons
M	8	number of excitatory subpopulations (= number of external spike sources)
n_E	$N_E/M = 150$	number of excitatory neurons per subpopulation
ρ	20	(target) number of active neurons per subpopulation after learning = minimal number of coincident excitatory inputs required to trigger a spike in postsynaptic inhibitory neurons
n	{100, ..., 1000}	number of excitatory or inhibitory Poissonian sources
Connectivity		
K_{EE}	240	number of excitatory inputs per excitatory neuron (EE in-degree)
p	$K_{EE}/N_E = 0.2$	connection probability
K_{EI}	$N_I = 1$	number of inhibitory inputs per excitatory neuron (EI in-degree)
K_{IE}	$N_E = 1200$	number of excitatory inputs per inhibitory neuron (IE in-degree)
K_{II}	0	number of inhibitory inputs per inhibitory neuron (II in-degree)
K_{EQ}	100	number of excitatory Poissonian inputs per excitatory neuron (EQ)
K_{EV}	$K_{EQ} = 100$	number of inhibitory Poissonian inputs per excitatory neuron (EV)
Excitatory neurons		
$\tau_{m,E}$	10 ms	membrane time constant
$\tau_{ref,E}$	20 ms	absolute refractory period
C_m	250 pF	membrane capacity
V_r	0 mV	reset potential
θ_E	20 mV (training), 7 mV (replay)	somatic spike threshold
I_{dAP}	200 pA	dAP current plateau amplitude
τ_{dAP}	60 ms	dAP duration
θ_{dAP}	59 pA	dAP threshold
I_θ	-1000 pA	reset dAP current
Inhibitory neurons		
$\tau_{m,I}$	5 ms	membrane time constant
$\tau_{ref,I}$	2 ms	absolute refractory period
C_m	250 pF	membrane capacity
V_r	0 mV	reset potential
θ_I	15 mV (training), 12 mV (replay)	spike threshold

Continued next page

Name	Value	Description
Synapse		
J_{IE}	581.19 pA	weight of IE connections (EPSC amplitude)
J_{EI}	-12915.49 pA (training), -2260.21 pA (replay)	weight of EI connections (IPSC amplitude)
J_{EX}	4112.20 pA (training), 1962.64 pA (replay)	weight of EX connections (EPSC amplitude)
J_{EQ}	$\sigma/\sqrt{K_{EQ}\nu\tau_{EQ}}$	weight of EQ connections (EPSC amplitude)
J_{EV}	$-J_{EQ}$	weight of EV connections (EPSC amplitude)
J_{EG}	1 pA	weight of EG connections (EPSC amplitude)
τ_{EE}	17 ms	synaptic time constant of EE connections
τ_{EI}	1 ms	synaptic time constant of EI connections
τ_{EX}	2 ms	synaptic time constant of EX connection
τ_{IE}	0.5 ms	synaptic time constant of IE connections
τ_{EQ}	2 ms	synaptic time constant of EQ connections
τ_{EV}	$\tau_{EQ} = 2$ ms	synaptic time constant of EV connections
d_{EE}	2 ms	delay of EE connections (dendritic)
d_{IE}	0.1 ms	delay of IE connections
d_{EI}	0.1 ms	delay of EI connections
d_{EX}	0.1 ms	delay of EX connections
d_{EQ}	0.1 ms	delay of EQ connections
d_{EV}	$d_{EQ} = 0.1$ ms	delay of EV connections
Plasticity		
λ_+	0.001	potentiation rate
λ_-	0.000013	depression rate
λ_h	0.0008	homeostasis rate
J_{\min}	0 pA	minimum weight
J_{\max}	35 pA	maximum weight
$J_{0,\min}$	0 pA	minimal initial weight
$J_{0,\max}$	1 pA	maximal initial weight
τ_+	20 ms	potentiation time constant
z^*	10.35	target dAP activity
τ_h	2200 ms	homeostasis time constant
y_i	1	depression decrement

Continued next page

Input		
S	{2, 3}	number of sequences per set
C	4	number of characters per sequence
A	{6, 8}	alphabet length
ΔT	40 ms	inter-stimulus interval
ΔT_{seq}	100 ms	inter-sequence interval
ΔT_{cue}	140 ms or $\sim \mathcal{U}(u_{\text{min}} u_{\text{max}})$	inter-cue interval
u_{min}	140 ms	minimal inter-cue interval
u_{max}	280 ms	maximal inter-cue interval
σ	{0, 190, 537} pA	noise amplitude resulting from the Poissonian background inputs
c	$\frac{\mu}{K} = \{0, 1\}$	noise correlation
ν	1000 s^{-1}	rate of Poissonian background inputs
a	{0, 8, 16}	amplitude of the sinusoidal current generators
f	{10, 30, 70} Hz	frequency of the sinusoidal current generators
Simulation		
Δt	0.1 ms	time resolution
K	350	number of training episodes
N_{t}	150	number of trials

Table 7.3: Model and simulation parameters. Parameters derived from other parameters are marked in gray. Curly brackets depict a set of values corresponding to different experiments.

7.3 SUPPORTING INFORMATION CHAPTER 4

7.3.1 Model and simulation parameters

Name	Value	Description
Network		
N_E	1800	total number of excitatory neurons
N_I	12	total number of inhibitory neurons
M	14	number of excitatory subpopulations (= number of external spike sources)
n_E	$N_E/M = 150$	number of excitatory neurons per subpopulation
n_I	$N_I/M = 1$	number of inhibitory neurons per subpopulation
ρ	20	(target) number of active neurons per subpopulation after learning = minimal number of coincident excitatory inputs required to trigger a spike in postsynaptic inhibitory neurons
(Potential) Connectivity		
K_{EE}	450	number of excitatory inputs per excitatory neuron (EE in-degree)
p	$K_{EE}/N_E = 0.25$	probability of potential (excitatory) connections
K_{EI}	$n_I = 1$	number of inhibitory inputs per excitatory neuron (EI in-degree)
K_{IE}	n_E	number of excitatory inputs per inhibitory neuron (IE in-degree)
K_{II}	0	number of inhibitory inputs per inhibitory neuron (II in-degree)
Excitatory neurons		
$\tau_{m,E}$	10 ms	membrane time constant
$\tau_{ref,E}$	20 ms	absolute refractory period
C_m	$250 \mu F$	membrane capacity
V_r	0 mV	reset potential
θ_E	30 mV	somatic spike threshold
I_{dAP}	$200 \mu A$	dAP current plateau amplitude
τ_{dAP}	60 ms	dAP duration
θ_{dAP}	see Eq (4.14)	dAP threshold
Inhibitory neurons		
$\tau_{m,I}$	5 ms	membrane time constant
$\tau_{ref,I}$	2 ms	absolute refractory period
C_m	$250 \mu F$	membrane capacity
V_r	0. mV	reset potential
θ_I	15 mV	spike threshold

Continued next page

Name	Value	Description
Synapse		
\tilde{G}_{IE}	0.9 mV	weight of IE connections (EPSP amplitude)
G_{IE}	581.19 μ S	weight of IE connections (EPSC amplitude)
\tilde{G}_{EI}	-60 mV	weight of EI connections (IPSP amplitude)
G_{EI}	-19373.24 μ S	weight of EI connections (IPSC amplitude)
\tilde{G}_{EX}	33 mV	weight of EX connections (EPSP amplitude)
G_{EX}	6168.31 μ S	weight of EX connections (EPSC amplitude)
τ_{EE}	2 ms	synaptic time constant of EE connections
τ_{IE}	0.5 ms	synaptic time constant of IE connections
τ_{EI}	1 ms	synaptic time constant of EI connections
τ_{EX}	2 ms	synaptic time constant of EX connection
d_{EE}	2 ms	delay of EE connections (dendritic)
d_{IE}	0.1 ms	delay of IE connections
d_{EI}	0.1 ms	delay of EI connections
d_{EX}	0.1 ms	delay of EX connections
Plasticity		
λ_+	{0.02, ..., 0.1 , ..., 0.42} (analog synapse), {0.02, ..., 0.04 , ..., 0.42} (binary synapse)	potentiation learning rate
λ_-	λ_+/β	depression rate
β	{0.5, 1, 2, 3 }	ratio between depression and potentiation learning rates
λ_h	λ_-	homeostasis rate
μ_+	{0, 0.5 , 1}	weight dependence (potentiation) exponent (default parameter)
μ_-	{0, 0.5 , 1}	weight dependence (depression) exponent (default parameter)
θ_P	10	synapse maturity threshold
$P_{\min,ij}$	$\sim \mathcal{U}(P_{0,\min}, P_{0,\max})$	minimum permanence
$G_{\min,ij}$	$\sim \mathcal{U}(G_{0,\min}, G_{0,\max})$	minimum conductance
G_{\max}	{50, ..., 300 , ..., 400} μ S	maximum conductance
$G_{0,\min}$	7.5 μ S	minimal initial conductance
$G_{0,\max}$	12.5 μ S	maximal initial conductance
$P_{0,\max}$	8	maximal initial permanence
$P_{0,\min}$	0	minimal initial permanence
$P_{0,\max}$	8	maximal initial permanence
σ_r	{0, ..., 0.03 , ..., 0.1}	read noise
σ_w	{0, ..., 0.01 , ..., 0.25}	write noise
z^*	1.8	target dAP activity
τ_h	1040 ms	homeostasis time constant
Input		
S	4	number of sequences per set
C	5	number of characters per sequence
A	12	alphabet length
ΔT	40 ms	inter-stimulus interval
ΔT_{seq}	100 ms	inter-sequence interval
Simulation		
Δt	0.1 ms	time resolution
K	400	number of training episodes

Table 7.4: Model and simulation parameters. Parameters derived from other parameters are marked in gray. Bold numbers depict default values. Curly brackets depict a set of values corresponding to different experiments.

BIBLIOGRAPHY

- Abbaspour, E., S. Menzel, A. Hardtdegen, S. Hoffmann-Eifert, and C. Jungemann (2018). KMC simulation of the electroforming, set and reset processes in redox-based resistive switching devices. *IEEE Transactions on Nanotechnology* 17(6), 1181–1188.
- Abbott, L. F. and S. B. Nelson (2000). Synaptic plasticity: taming the beast. *Nature Neuroscience* 3, 1178–1183.
- Abeles, M. (1991). *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge: Cambridge University Press.
- Adnan, M. M., S. Sayyaparaju, S. D. Brown, M. S. A. Shawkat, C. D. Schuman, and G. S. Rose (2021). Design of a robust memristive spiking neuromorphic system with unsupervised learning in hardware. *ACM Journal on Emerging Technologies in Computing Systems* 17(4), 1–26.
- Agarwal, S., S. J. Plimpton, D. R. Hughart, A. H. Hsia, I. Richter, J. A. Cox, C. D. James, and M. J. Marinella (2016). Resistive memory device requirements for a neural algorithm accelerator. In *2016 International Joint Conference on Neural Networks (IJCNN)*. IEEE.
- Ahmad, S. and J. Hawkins (2015). Properties of sparse distributed representations and their application to Hierarchical Temporal Memory. *ArXiv*, 1503.07469.
- Ahmad, S. and J. Hawkins (2016). How do neurons operate on sparse distributed representations? a mathematical theory of sparsity, neurons and active dendrites. *ArXiv*, 1601.00720.
- Aimone, J. B. (2019). Neural algorithms and computing beyond moore’s law. *Communications of the ACM* 62(4), 110–110.
- Anggono, V. and R. L. Huganir (2012). Regulation of AMPA receptor trafficking and synaptic plasticity. *Current Opinion in Neurobiology* 22(3), 461–469.
- Antic, S. D., W.-L. Zhou, A. R. Moore, S. M. Short, and K. D. Ikonomu (2010). The decade of the dendritic nmda spike. *Journal of Neuroscience Research* 88(14), 2991–3001.
- Arieli, A., A. Sterkin, A. Grinvald, and A. Aertsen (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273(5283), 1868–1871.
- Asabuki, T. and T. Fukai (2020). Somatodendritic consistency check for temporal feature segmentation. *Nature Communications* 11(1), 1–13.
- Asabuki, T. and T. Fukai (2021). Neural mechanisms of context-dependent segmentation tested on large-scale recording data. *BioRxiv*, 2021.04.25.441363.

- Atherton, L. A., D. Dupret, and J. R. Mellor (2015). Memory trace replay: the shaping of memory consolidation by neuromodulation. *Trends in Neurosciences* 38(9), 560–570.
- Attwell, D. and S. B. Laughlin (2001). An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow & Metabolism* 21, 1133–1145.
- Augusto, E. and F. Gambino (2019). Can NMDA spikes dictate computations of local networks and behavior? *Frontiers in Molecular Neuroscience* 12, 238.
- Baluch, F. and L. Itti (2011). Mechanisms of top-down attention. *Trends in Neurosciences* 34(4), 210–224.
- Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience* 5(8), 617–629.
- Bar, M. (2007). The proactive brain: using analogies and associations to generate predictions. *Trends in Cognitive Sciences* 11(7), 280–289.
- Barth, A. L. and J. F. Poulet (2012). Experimental evidence for sparse firing in the neocortex. *Trends in Neurosciences* 35(6), 345–355.
- Bear, M. F., B. W. Connors, and M. A. Paradiso (2007). *Neuroscience: exploring the brain* (3 ed.). Lippincott Williams & Wilkins.
- Bear, M. F. and R. C. Malenka (1994). Synaptic plasticity: LTP and LTD. *Current Opinion in Neurobiology* 4(3), 389–399.
- Bengel, C., A. Siemon, F. Cuppers, S. Hoffmann-Eifert, A. Hardtdegen, M. von Witzleben, L. Hellmich, R. Waser, and S. Menzel (2020). Variability-aware modeling of filamentary oxide-based bipolar resistive switching cells using SPICE level compact models. *IEEE Transactions on Circuits and Systems I: Regular Papers* 67(12), 4618–4630.
- Benjamin, B., P. Gao, E. McQuinn, S. Choudhary, A. Chandrasekaran, J. Bussat, R. Alvarez-Icaza, J. V. Arthur, P. Merolla, and K. Boahen (2014). Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations. *Proceedings of the IEEE* 102(5), 699–716.
- Bi, G. and M. Poo (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience* 18, 10464–10472.
- Billaudelle, S. and S. Ahmad (2015). Porting HTM models to the heidelberg neuro-morphic computing platform. *ArXiv*, 1505.02142.
- Binzegger, T., R. J. Douglas, and K. A. C. Martin (2004). A quantitative map of the circuit of cat primary visual cortex. *Journal of Neuroscience* 39(24), 8441–8453.
- Bod, R., J. Hay, and S. Jannedy (2003). *Probabilistic linguistics*. MIT press.
- Boucein, C., M. Nawrot, P. Schnepel, and A. Aertsen (2011). Beyond the cortical column: abundance and physiology of horizontal connections imply a strong role for inputs from the surround. *Frontiers in Neuroscience* 5, 32.

- Bouhadjar, Y., M. Diesmann, R. Waser, D. J. Wouters, and T. Tetzlaff (2019). Constraints on sequence processing speed in biological neuronal networks. In *Proceedings of the International Conference on Neuromorphic Systems*, pp. 1–9.
- Bouhadjar, Y., C. Moruzzi, and M. Payvand (2022). Prediction: An algorithmic principle meeting neuroscience and machine learning halfway. In *CEUR Workshop Proceedings*, pp. 46–52.
- Bouhadjar, Y., S. Siegel, T. Tetzlaff, M. Diesmann, R. Waser, and D. J. Wouters (2022c). Sequence learning in a spiking neuronal network with memristive synapses. *ArXiv*, 2211.16592.
- Bouhadjar, Y., D. J. Wouters, M. Diesmann, and T. Tetzlaff (2022a). Coherent noise enables probabilistic sequence replay in spiking neuronal networks. *ArXiv*, 2206.10538.
- Bouhadjar, Y., D. J. Wouters, M. Diesmann, and T. Tetzlaff (2022b). Sequence learning, prediction, and replay in networks of spiking neurons. *PLOS Computational Biology* 18(6), e1010233.
- Braitenberg, V. and A. Schüz (1991). *Anatomy of the Cortex: Statistics and Geometry*. Berlin, Heidelberg, New York: Springer-Verlag.
- Braitenberg, V. and A. Schüz (1998). *Cortex: Statistics and Geometry of Neuronal Connectivity* (2nd ed.). Berlin: Springer-Verlag.
- Branco, T. and K. Staras (2009). The probability of neurotransmitter release: variability and feedback control at single synapses. *Nature Reviews Neuroscience* 10(5), 373–383.
- Breuer, D., M. Timme, and R.-M. Memmesheimer (2014). Statistical physics of neural systems with nonadditive dendritic coupling. *Physical Review X* 4(1), 011053.
- Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Großhirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Leipzig: Johann Ambrosius Barth.
- Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience* 8(3), 183–208.
- Brunel, N. and V. Hakim (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Computation* 11(7), 1621–1671.
- Buesing, L., J. Bill, B. Nessler, and W. Maass (2011). Neural dynamics as sampling: A model for stochastic computation in recurrent networks of spiking neurons. *PLOS Computational Biology* 7, e1002211.
- Buonomano, D. V. and M. M. Merzenich (1998). Cortical plasticity: From synapses to maps. *Annual Review of Neuroscience* 21(1), 149–186.
- Burr, G. W., R. M. Shelby, A. Sebastian, S. Kim, S. Kim, S. Sidler, K. Virwani, M. Ishii, P. Narayanan, A. Fumarola, L. L. Sanches, I. Boybat, M. L. Gallo, K. Moon, J. Woo, H. Hwang, and Y. Leblebici (2016). Neuromorphic computing using non-volatile memory. *Advances in Physics: X* 2(1), 89–124.
- Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford University Press.

- Buzsáki, G. and A. Draguhn (2004). Neuronal oscillations in cortical networks. *Science* 304, 1926–1929.
- Buzsaki, G., C. Geisler, D. Henze, and X.-J. Wang (2004). Interneuron diversity series: Circuit complexity and axon wiring economy of cortical interneurons. *Trends in Neurosciences* 27(4), 186–193.
- Chen, P.-Y., B. Lin, I.-T. Wang, T.-H. Hou, J. Ye, S. Vrudhula, J. sun Seo, Y. Cao, and S. Yu (2015a). Mitigating effects of non-ideal synaptic device characteristics for on-chip learning. In *2015 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pp. 194–199. IEEE.
- Chen, P.-Y., B. Lin, I.-T. Wang, T.-H. Hou, J. Ye, S. Vrudhula, J. sun Seo, Y. Cao, and S. Yu (2015b). Mitigating effects of non-ideal synaptic device characteristics for on-chip learning. In *2015 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE.
- Chua, L. (1971). Memristor—the missing circuit element. *IEEE Transactions on Circuit Theory* 18(5), 507–519.
- Chua, L. and S. M. Kang (1976). Memristive devices and systems. *Proceedings of the IEEE* 64(2), 209–223.
- Citri, A. and R. C. Malenka (2007). Synaptic plasticity: Multiple forms, functions, and mechanisms. *Neuropharmacology* 33(1), 18–41.
- Clegg, B. A., G. J. DiGirolamo, and S. W. Keele (1998). Sequence learning. *Trends in Cognitive Sciences* 2(8), 275–281.
- Cohen, J. D., S. M. McClure, and A. J. Yu (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B* 362(1481), 933–942.
- Cohen, M. R. and J. H. R. Maunsell (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience* 12, 1594–1600.
- Cone, I. and H. Z. Shouval (2021). Learning precise spatiotemporal sequences via biophysically realistic learning rules in a modular, spiking network. *eLife* 10, e63751.
- Coombes, S. (2005). Waves, bumps, and patterns in neural field theories. *Biological Cybernetics* 93, 91–108.
- Cui, Y., S. Ahmad, and J. Hawkins (2016). Continuous online sequence learning with an unsupervised neural network model. *Neural Computation* 28(11), 2474–2504.
- Cüppers, F., S. Menzel, C. Bengel, A. Hardtdegen, M. von Witzleben, U. Böttger, R. Waser, and S. Hoffmann-Eifert (2019). Exploiting the switching dynamics of HfO₂-based ReRAM devices for reliable analog memristive behavior. *APL Materials* 7(9), 091105.
- Dave, A. S. and D. Margoliash (2000). Song replay during sleep and computational rules for sensorimotor vocal learning. *Science* 290(5492), 812–816.
- Davidson, T. J., F. Kloosterman, and M. A. Wilson (2009). Hippocampal replay of extended experience. *Neuron* 63(4), 497–507.

- Davies, M., N. Srinivasa, T.-H. Lin, G. Chinya, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain, Y. Liao, C.-K. Lin, A. Lines, R. Liu, D. Mathaikutty, S. McCoy, A. Paul, J. Tse, G. Venkataramanan, Y.-H. Weng, A. Wild, Y. Yang, and H. Wang (2018). Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38(1), 82–99.
- Daw, N. D., J. P. O’Doherty, P. Dayan, B. Seymour, and R. J. Dolan (2006). Cortical substrates for exploratory decisions in humans. *Nature* 441(7095), 876–879.
- Dayan, P. and L. F. Abbott (2001). *Theoretical Neuroscience*. Cambridge: MIT Press.
- de Kock, C. P. J. and B. Sakmann (2009). Spiking in primary somatosensory cortex during natural whisking in awake head-restrained rats is cell-type specific. *Proceedings of the National Academy of Sciences of the United States of America* 106(38), 16446–16450.
- De la Rocha, J., B. Doiron, E. Shea-Brown, J. Kresimir, and A. Reyes (2007). Correlation between neural spike trains increases with firing rate. *Nature* 448(16), 802–807.
- Debanne, D. (2004). Information processing in the axon. *Nature Reviews Neuroscience* 5, 304–316.
- Deger, M., M. Helias, S. Rotter, and M. Diesmann (2012). Spike-timing dependence of structural plasticity explains cooperative synapse formation in the neocortex. *PLOS Computational Biology* 8(9), e1002689.
- Dehaene, S., F. Meyniel, C. Wacongne, L. Wang, and C. Pallier (2015). The neural representation of sequences: from transition probabilities to algebraic patterns and linguistic trees. *Neuron* 88(1), 2–19.
- Denker, M., L. Zehl, B. E. Kilavik, M. Diesmann, T. Brochier, A. Riehle, and S. Grün (2018). LFP beta amplitude is linked to mesoscopic spatio-temporal phase patterns. *Scientific Reports* 8(1), 1–21.
- Destexhe, A., M. Rudolph, J.-M. Fellous, and T. Sejnowski (2001). Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience* 107(1), 13–24.
- DeWeese, M. R. and A. M. Zador (2006). Non-gaussian membrane potential dynamics imply sparse, synchronous activity in auditory cortex. *Journal of Neuroscience* 26(47), 12206–12218.
- DiCarlo, J., D. Zoccolan, and N. Rust (2012). How does the brain solve visual object recognition? *Neuron* 73(3), 415–434.
- Diesmann, M. (2002). Conditions for stable propagation of synchronous spiking in cortical neural networks: Single neuron dynamics and network properties.
- Diesmann, M., M.-O. Gewaltig, and A. Aertsen (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402(6761), 529–533.
- Dittmann, R. and J. P. Strachan (2019). Redox-based memristive devices for new computing paradigm. *APL Materials* 7(11), 110903.

- Doevenspeck, J., R. Degraeve, A. Fantini, P. Debacker, D. Verkest, R. Lauwereins, and W. Dehaene (2018). Temporal sequence learning with a history-sensitive probabilistic learning rule intrinsic to oxygen vacancy-based ram. In *2018 IEEE International Electron Devices Meeting (IEDM)*, pp. 20–5. IEEE.
- Dold, D., I. Bytschok, A. F. Kungl, A. Baumbach, O. Breitwieser, W. Senn, J. Schemmel, K. Meier, and M. A. Petrovici (2019). Stochasticity from function—why the bayesian brain may need no noise. *Neural Networks 119*, 200–213.
- Douglas, R. J. and K. A. C. Martin (2004). Neuronal circuits of the neocortex. *Annual Review of Neuroscience 27*, 419–451.
- Eccles, J., F. P. and K. Koketsu (1954). Cholinergic and inhibitory synapses in a pathway from motor-axon collaterals to motoneurons. *Journal of Physiology 126*, 524–562.
- Ecker, A. S., P. Berens, G. A. Keliris, M. Bethge, and N. K. Logothetis (2010). Decorrelated neuronal firing in cortical microcircuits. *Science 327*(5965), 584–587.
- Elliott, T. (2003). An analysis of synaptic normalization in a general class of hebbian models. *Neural Computation 15*(4), 937–963.
- Ermentrout, B. (1998). Neural networks as spatio-temporal pattern-forming systems. *Reports on Progress in Physics 61*(4), 353–430.
- Euston, D. R., M. Tatsuno, and B. L. McNaughton (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science 318*(5853), 1147–1150.
- Faisal, A. A., L. P. Selen, and D. M. Wolpert (2008). Noise in the nervous system. *Nature Reviews Neuroscience 9*(4), 292–303.
- Fantini, A., L. Goux, R. Degraeve, D. Wouters, N. Raghavan, G. Kar, A. Belmonte, Y.-Y. Chen, B. Govoreanu, and M. Jurczak (2013). Intrinsic switching variability in HfO₂ RRAM. In *2013 5th IEEE International Memory Workshop*, pp. 30–33. IEEE.
- Feldman, D. E. (2012). The spike-timing dependence of plasticity. *Neuron 75*(4), 556–571.
- Felleman, D. J. and D. C. Van Essen (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex 1*, 1–47.
- Fellous, J., P. Tiesinga, P. Thomas, and T. Sejnowski (2004). Discovering spike patterns in neuronal responses. *Journal of Neuroscience 24*(24), 2989–3001.
- Fleck, K., C. La Torre, N. Aslam, S. Hoffmann-Eifert, U. Böttger, and S. Menzel (2016). Uniting gradual and abrupt set processes in resistive switching oxides. *Physical review applied 6*(6), 064015.
- Fouda, M. E., F. Kurdahi, A. Eltawil, and E. Neftci (2020). Spiking neural networks for inference and learning: a memristor-based design perspective. In *Memristive Devices for Brain-Inspired Computing*, Woodhead Publishing Series in Electronic and Optical Materials, pp. 499–530. Woodhead Publishing.

- Furber, S., F. Galluppi, S. Temple, and L. Plana (2014). The SpiNNaker Project. *Proceedings of the IEEE* 102(5), 652–665.
- Fusi, S. and L. F. Abbott (2007). Limits on the memory storage capacity of bounded synapses. *Nature Neuroscience* 10(4), 485–493.
- Garbin, D., E. Vianello, O. Bichler, Q. Raffay, C. Gamrat, G. Ghibaud, B. DeSalvo, and L. Perniola (2015). Hfo 2-based oxram devices as synapses for convolutional neural networks. *IEEE Transactions on Electron Devices* 62(8), 2494–2501.
- Gavornik, J. P. and M. F. Bear (2014). Learned spatiotemporal sequence recognition and prediction in primary visual cortex. *Nature Neuroscience* 17(5), 732.
- Gewaltig, M.-O. and M. Diesmann (2007). NEST (NEural Simulation Tool). *Scholarpedia Journal* 2(4), 1430.
- Gidon, A., T. A. Zolnik, P. Fidzinski, F. Bolduan, A. Papoutsi, P. Poirazi, M. Holtkamp, I. Vida, and M. E. Larkum (2020). Dendritic action potentials and computation in human layer 2/3 cortical neurons. *Science* 367(6473), 83–87.
- Goedeke, S. and M. Diesmann (2008). The mechanism of synchronization in feed-forward neuronal networks. *New Journal of Physics* 10, 015007.
- Gokmen, T. and W. Haensch (2020). Algorithm for training neural networks on resistive device arrays. *Frontiers in Neuroscience* 14, 103.
- Gray, H. (1918). *Anatomy of the human body (plate 728, labels removed)*. Wikimedia Commons, Public Domain Retrieved from https://commons.wikimedia.org/wiki/File:Lobes_of_the_brain_NL.svg on June 3, 2018.
- Grewe, B. F., D. Langer, H. Kasper, B. M. Kampa, and F. Helmchen (2010). High-speed in vivo calcium imaging reveals neuronal network activity with near-millisecond precision. *Nature Methods* 7(5), 399–405.
- Grosmark, A. D., K. Mizuseki, E. Pastalkova, K. Diba, and G. Buzsáki (2012). REM sleep reorganizes hippocampal excitability. *Neuron* 75(6), 1001–1007.
- Hahne, J., S. Diaz, A. Patronis, W. Schenck, A. Peyser, S. Graber, S. Spreizer, S. B. Vennemo, T. Ippen, H. Mørk, J. Jordan, J. Senk, S. Konradi, P. Weidel, T. Fardet, D. Dahmen, D. Terhorst, J. Stapmanns, G. Trensche, A. van Meegen, J. Pronold, J. M. Eppler, C. Linssen, A. Morrison, A. Sinha, J. Mitchell, S. Kunkel, R. Deepu, E. Hagen, T. Vierjahn, N. L. Kamiji, R. de Schepper, P. Machado, J. Albers, W. Klijjn, A. Myczko, W. Mayner, P. Nagendra Babu, H. Jiang, S. Billaudelle, B. S. Vogler, G. Miotto, L. Kusch, A. Antonietti, A. Morales-Gregorio, J. Dolderer, Y. Bouhadjar, and H. E. Plesser (2021). Nest 3.0.
- Hahnloser, R. H., A. A. Kozhevnikov, and M. S. Fee (2002). An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* 419(6902), 65–70.
- Hansen, K. A., S. F. Hillenbrand, and L. G. Ungerleider (2012). Effects of prior knowledge on decisions made under perceptual vs. categorical uncertainty. *Frontiers in Neuroscience* 6, 163.

- Hardtdegen, A., C. L. Torre, F. Cuppers, S. Menzel, R. Waser, and S. Hoffmann-Eifert (2018). Improved switching stability and the effect of an internal series resistor in HfO₂/TiO_x bilayer ReRAM cells. *IEEE Transactions on Electron Devices* 65(8), 3229–3236.
- Harris, K. D. and G. M. G. Shepherd (2015). The neocortical circuit: themes and variations. *Nature Neuroscience* 18(2), 170–181.
- Hartmann, C., A. Lazar, B. Nessler, and J. Triesch (2015). Where’s the noise? key features of spontaneous activity and neural variability arise through learning in a deterministic network. *PLOS Computational Biology* 11(12), e1004640.
- Harvey, C. D., P. Coen, and D. W. Tank (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484(7392), 62–68.
- Hasler, P., C. Diorio, B. Minch, and C. Mead (1994). Single transistor learning synapses. In *Advances in Neural Information Processing Systems*, Volume 7. MIT Press.
- Hawkins, J. and S. Ahmad (2016). Why neurons have thousands of synapses, a theory of sequence memory in neocortex. *Frontiers in Neural Circuits* 10, 23.
- Hawkins, J., S. Ahmad, D. Dubinsky, et al. (2011). Cortical learning algorithm and Hierarchical Temporal Memory. *Numenta Whitepaper* 1(68), 2.
- Hawkins, J. and S. Blakeslee (2007). *On intelligence: How a new understanding of the brain will lead to the creation of truly intelligent machines*. Macmillan.
- Hebb, D. O. (1967). *The Organization of Behaviour: A Neuropsychological Theory*. John Wiley & Sons.
- Helias, M., T. Tetzlaff, and M. Diesmann (2014). The correlation structure of local cortical networks intrinsically results from recurrent dynamics. *PLOS Computational Biology* 10(1), e1003428.
- Hole, K. J. (2016). *The HTM Learning Algorithm*, pp. 113–124. Cham: Springer International Publishing.
- Holt, G. R., W. R. Softky, C. Koch, and R. J. Douglas (1996). Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. *Journal of Neurophysiology* 75(5), 1806–1814.
- Hong, X., D. J. Loy, P. A. Dananjaya, F. Tan, C. Ng, and W. Lew (2018). Oxide-based RRAM materials for neuromorphic computing. *Journal of Materials Science* 53(12), 8720–8746.
- Horton, J. C. and D. L. Adams (2005). The cortical column: a structure without a function. *Philosophical Transactions of the Royal Society B* 360(1456), 837–862.
- Hosomi, M., H. Yamagishi, T. Yamamoto, K. Bessho, Y. Higo, K. Yamane, H. Yamada, M. Shoji, H. Hachino, C. Fukumoto, H. Nagao, and H. Kano (2005). A novel nonvolatile memory with spin torque transfer magnetization switching: spin-ram. In *IEEE International Electron Devices Meeting, 2005. IEDM Technical Digest.*, pp. 459–462. IEEE.

- Hubel, D. H. and T. N. Wiesel (1962). Receptive fields, binocular interaction, and functional architecture in the cat's visual cortex. *Journal of Physiology* 160, 106–154.
- Hubel, D. H. and T. N. Wiesel (1965). Receptive fields and functional architecture in two nonstriate areas (18 and 19) of the cat. *Journal of Neurophysiology* 28, 229–289.
- Hubel, D. H. and T. N. Wiesel (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Neurophysiology* 195(1), 215–243.
- Humphries, M. (2021). *The Spike*. Princeton University Press.
- Ielmini, D., R. Bruchhaus, and R. Waser (2011). Thermochemical resistive switching: materials, mechanisms, and scaling projections. *Phase Transitions* 84(7), 570–602.
- Ielmini, D. and H.-S. P. Wong (2018). In-memory computing with resistive switching devices. *Nature Electronics* 1(6), 333–343.
- Jahnke, S., M. Timme, and R.-M. Memmesheimer (2012). Guiding synchrony through random networks. *Physical Review X* 2(4), 041016.
- Jiang, X., S. Shen, C. R. Cadwell, P. Berens, F. Sinz, A. S. Ecker, S. Patel, and A. S. Tolias (2015). Principles of connectivity among morphologically defined cell types in adult neocortex. *Science* 350(6264), aac9462.
- Jordan, J., T. Ippen, M. Helias, I. Kitayama, M. Sato, J. Igarashi, M. Diesmann, and S. Kunkel (2018). Extremely scalable spiking neuronal network simulation code: From laptops to exascale computers. *Frontiers in Neuroinformatics* 12, 2.
- Jordan, J., M. A. Petrovici, O. Breitwieser, J. Schemmel, K. Meier, M. Diesmann, and T. Tetzlaff (2019). Deterministic networks for probabilistic computing. *Scientific Reports* 9, 18303.
- Kandel, E. R., J. H. Schwartz, and T. M. Jessel (1991). *Principles of neural science* (3 ed.). New York: Elsevier.
- Kandel, E. R., J. H. Schwartz, and T. M. Jessel (2000). *Principles of Neural Science* (4 ed.). New York: McGraw-Hill. ISBN 978-0838577011.
- Khan, A. G., J. Poort, A. Chadwick, A. Blot, M. Sahani, T. D. Mrsic-Flogel, and S. B. Hofer (2018). Distinct learning-induced changes in stimulus selectivity and interactions of GABAergic interneuron classes in visual cortex. *Nature Neuroscience* 21(6), 851–859.
- Klampff, S. and W. Maass (2013). Emergence of dynamic memory traces in cortical microcircuit models through stdp. *Journal of Neuroscience* 33(28), 11515–11529.
- Klos, C., D. Miner, and J. Triesch (2018). Bridging structure and function: A model of sequence learning and prediction in primary visual cortex. *PLOS Computational Biology* 14(6), e1006187.
- Kriener, B., T. Tetzlaff, A. Aertsen, M. Diesmann, and S. Rotter (2008). Correlations and population dynamics in cortical networks. *Neural Computation* 20, 2185–2226.

- Kumar, S., Z. Wang, X. Huang, N. Kumari, N. Davila, J. P. Strachan, D. Vine, A. L. D. Kilcoyne, Y. Nishi, and R. S. Williams (2017). Oxygen migration during resistance switching and failure of hafnium oxide memristors. *Applied Physics Letters* 110(10), 103503.
- Kunkel, S., M. Schmidt, J. M. Eppler, G. Masumoto, J. Igarashi, S. Ishii, T. Fukai, A. Morrison, M. Diesmann, and M. Helias (2014). Spiking network simulation code for petascale computers. *Frontiers in Neuroinformatics* 8, 78.
- Kuzum, D., S. Yu, and H.-S. P. Wong (2013). Synaptic electronics: materials, devices and applications. *Nanotechnology* 24(38), 382001.
- Lamme, V. A., H. Super, H. Spekreijse, et al. (1998). Feedforward, horizontal, and feedback processing in the visual cortex. *Current Opinion in Neurobiology* 8(4), 529–535.
- Lamprecht, R. and J. LeDoux (2004). Structural plasticity and memory. *Nature Reviews Neuroscience* 5, 45–54.
- Larkum, M. E., T. Nevian, M. Sandler, A. Polsky, and J. Schiller (2009). Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: A new unifying principle. *Science* 325(5941), 756–760.
- Lashley, K. S. (1951). *The problem of serial order in behavior*, Volume 21. Bobbs-Merrill.
- Laughlin, S. B. and T. J. Sejnowski (2003). Communication in neuronal networks. *Science* 301, 1870–1874.
- Lazar, A., G. Pipa, and J. Triesch (2009). SORN: a self-organizing recurrent neural network. *Frontiers in Computational Neuroscience* 3, 23.
- Ledergerber, D. and M. E. Larkum (2010). Properties of layer 6 pyramidal neuron apical dendrites. *Journal of Neuroscience* 30(39), 13031–13044.
- Lee, A. K. and M. A. Wilson (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* 36(6), 1183–1194.
- Legenstein, R. and W. Maass (2014). Ensembles of spiking neurons with noise support optimal probabilistic inference in a dynamically changing environment. *PLoS Computational Biology* 10(10), e1003859.
- Li, H., S. Wang, X. Zhang, W. Wang, R. Yang, Z. Sun, W. Feng, P. Lin, Z. Wang, L. Sun, and Y. Yao (2021). Memristive crossbar arrays for storage and computing applications. *Advanced Intelligent Systems* 3(9), 2100017.
- Li, S.-L., D. S. Shang, J. Li, J. L. Gang, and D. N. Zheng (2009). Resistive switching properties in oxygen-deficient pr_{0.7}ca_{0.3}mno₃ junctions with active al top electrodes. *Journal of Applied Physics* 105(3), 033710.
- Liao, D., N. A. Hessler, and R. Malinow (1995). Activation of postsynaptically silent synapses during pairing-induced ltp in ca1 region of hippocampal slice. *Nature* 375(6530), 400–404.
- Llinás, R. R. (2002). *I of the vortex: From neurons to self*. MIT press.

- Llinás, R. R. and S. Roy (2009). The ‘prediction imperative’ as the basis for self-awareness. *Philosophical Transactions of the Royal Society B* 364(1521), 1301–1307.
- London, M. and M. Häusser (2005). Dendritic computation. *Annual Review of Neuroscience* 28, 503–532.
- Louie, K. and M. A. Wilson (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* 29(1), 145–156.
- Lüscher, C., R. Noicoll, R. Malenka, and D. Muller (2000). Synaptic plasticity and dynamic modulation of the postsynaptic membrane. *Nature Neuroscience* 3(6), 545–550.
- Maass, W. (2014). Noise as a resource for computation and learning in networks of spiking neurons. *Proceedings of the IEEE* 102(5), 860–880.
- Maes, A., M. Barahona, and C. Clopath (2020). Learning spatiotemporal signals using a recurrent spiking network that discretizes time. *PLOS Computational Biology* 16(1), e1007606.
- Magee, J. C. and C. Grienberger (2020). Synaptic plasticity forms and functions. *Annual Review of Neuroscience* 43(1), 95–117.
- Mahapatra, N. R. and B. Venkatrao (1999). The processor-memory bottleneck. *XRDS: Crossroads, The ACM Magazine for Students* 5(3es), 2.
- Major, G., M. Larkum, and J. Schiller (2013). Active properties of neocortical pyramidal neuron dendrites. *Annual Review of Neuroscience* 36, 1–24.
- Markram, H., J. Lübke, M. Frotscher, and B. Sakmann (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–215.
- Martin, S. J., P. D. Grimwood, and R. G. M. Morris (2000). Synaptic plasticity and memory: An evaluation of the hypothesis. *Annual Review of Neuroscience* 23(1), 649–711.
- Martinez-Conde, S., S. L. Macknik, and D. H. Hubel (2004). The role of fixational eye movements in visual perception. *Nature Reviews Neuroscience* 5(3), 229–240.
- Masquelier, T. and S. J. Thorpe (2007). Unsupervised learning of visual features through spike timing dependent plasticity. *PLOS Computational Biology* 3(2), e31.
- Mauk, M. D. and D. V. Buonomano (2004). The neural basis of temporal processing. *Annual Review of Neuroscience* 27, 307–340.
- Mead, C. (1990). Neuromorphic electronic systems. *Proceedings of the IEEE* 78(10), 1629–1636.
- Mehonic, A. and A. J. Kenyon (2022). Brain-inspired computing needs a master plan. *Nature* 604(7905), 255–260.
- Mengual, U. M., W. A. Wybo, L. J. Spierenburg, M. Santello, W. Senn, and T. Nevian (2020). Efficient low-pass dendro-somatic coupling in the apical dendrite of layer 5 pyramidal neurons in the anterior cingulate cortex. *Journal of Neuroscience* 40(46), 8799–8815.

- Merolla, P. A., J. V. Arthur, R. Alvarez-Icaza, A. S. Cassidy, J. Sawada, F. Akopyan, B. L. Jackson, N. Imam, C. Guo, Y. Nakamura, B. Brezzo, I. Vo, S. K. Esser, R. Appuswamy, B. Taba, A. Amir, M. D. Flickner, W. P. Risk, R. Manohar, and D. S. Modha (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345(6197), 668–673.
- Milner, D. and M. Goodale (2006). *The visual brain in action*, Volume 27. OUP Oxford.
- Milojkovic, B. A. (2005). A strict correlation between dendritic and somatic plateau depolarizations in the rat prefrontal cortex pyramidal neurons. *Journal of Neuroscience* 25(15), 3940–3951.
- Moradi, S., N. Qiao, F. Stefanini, and G. Indiveri (2018). A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs). *IEEE Transactions on Biomedical Circuits and Systems* 12(1), 106–122.
- Morrison, A., A. Aertsen, and M. Diesmann (2007). Spike-timing dependent plasticity in balanced random networks. *Neural Computation* 19, 1437–1467.
- Morrison, A., M. Diesmann, and W. Gerstner (2008). Phenomenological models of synaptic plasticity based on spike-timing. *Biological Cybernetics* 98(6), 459–478.
- Mountcastle, V. B. (1978). *The mindful brain: an organizing principle for cerebral function: the unit module and the distributed system*. MIT Press.
- Mountcastle, V. B. (1997). The columnar organization of the neocortex. *Brain* 120, 701–722.
- Muller, L., F. Chavane, J. Reynolds, and T. J. Sejnowski (2018). Cortical travelling waves: mechanisms and computational principles. *Nature Reviews Neuroscience* 19(5), 255–268.
- Muller, L. and A. Destexhe (2012). Propagating waves in thalamus, cortex and the thalamocortical system: Experiments and models. *Journal of Physiology* 106(5-6), 222–238.
- Müllner, F. E., C. J. Wierenga, and T. Bonhoeffer (2015). Precision of inhibition: dendritic inhibition by individual gabaergic synapses on hippocampal pyramidal cells is confined in space and time. *Neuron* 87(3), 576–589.
- Myers, J. L. (2014). Probability learning and sequence learning. *Handbook of Learning and Cognitive Processes*, ed. WK Estes, 171–205.
- Nádasdy, Z., H. Hirase, A. Czurkó, J. Csicsvari, and G. Buzsáki (1999). Replay and time compression of recurring spike sequences in the hippocampus. *Journal of Neuroscience* 19(21), 9497–9507.
- Nagendra Babu, P., C. Linssen, J. M. Eppler, T. Schulte to Brinke, A. Ziaemehr, T. Fardet, Y. Bouhadjar, R. Duarte, B. Rumpe, and A. Morrison (2021). Nestml 4.0.
- Nauhaus, I., L. Busse, M. Carandini, and D. L. Ringach (2009). Stimulus contrast modulates functional connectivity in visual cortex. *Nature Neuroscience* 12, 70–76.

- Neftci, E. O. and B. B. Averbeck (2019). Reinforcement learning in artificial and biological systems. *Nature Machine Intelligence* 1(3), 133–143.
- Neftci, E. O., B. U. Pedroni, S. Joshi, M. Al-Shedivat, and G. Cauwenberghs (2016). Stochastic synapses enable efficient brain-inspired learning machines. *Frontiers in Neuroscience* 10, 241.
- Nevian, T. and B. Sakmann (2004). Single spine Ca^{2+} signals evoked by coincident epsps and backpropagating action potentials in spiny stellate cells of layer 4 in the juvenile rat somatosensory barrel cortex. *Journal of Neuroscience* 24(7), 1689–1699.
- Noudoost, B., M. H. Chang, N. A. Steinmetz, and T. Moore (2010). Top-down control of visual attention. *Current Opinion in Neurobiology* 20(2), 183–190.
- Oberländer, J., Y. Bouhadjar, and A. Morrison (2022). Learning and replaying spatiotemporal sequences: A replication study. *Frontiers in Integrative Neuroscience*, 113.
- Obien, M. E. J., K. Deligkaris, T. Bullmann, D. J. Bakkum, and U. Frey (2015). Revealing neuronal function through microelectrode array recordings. *Frontiers in Neuroscience* 8, 423.
- O’Doherty, J. P., J. Cockburn, and W. M. Pauli (2017). Learning, reward, and decision making. *Annual Review of Psychology* 68(1), 73–100.
- Olshausen, B. and D. Field (2004). Sparse coding of sensory inputs. *Current Opinion in Neurobiology* 14(4), 481–487.
- Paton, J. J. and D. V. Buonomano (2018). The neural basis of timing: Distributed mechanisms for diverse functions. *Neuron* 98(4), 687–705.
- Payeur, A., J. Guerguiev, F. Zenke, B. A. Richards, and R. Naud (2021). Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nature Neuroscience* 24(7), 1010–1019.
- Payne, B. R. and A. Peters (Eds.) (2002). *The Cat Primary Visual Cortex*. Elsevier.
- Payvand, M., F. Moro, K. Nomura, T. Dalgaty, E. Vianello, Y. Nishi, and G. Indiveri (2022). Self-organization of an inhomogeneous memristive hardware for sequence learning. *Nature Communications* 13(1), 5793.
- Pedretti, G., V. Milo, S. Ambrogio, R. Carboni, S. Bianchi, A. Calderoni, N. Ramaswamy, A. Spinelli, and D. Ielmini (2017). Memristive neural network for on-line learning and tracking with brain-inspired spike timing dependent plasticity. *Scientific Reports* 7(1), 1–10.
- Plotnikov, D., I. Blundell, T. Ippen, J. M. Eppler, B. Rumpe, and A. Morrison (2016). NESTML: a modeling language for spiking neurons. In A. Oberweis and R. Reussner (Eds.), *Modellierung 2016*, Volume P-254 of *Lecture Notes in Informatics (LNI)*, pp. 93–108. Modellierung 2016, Karlsruhe (Germany), 17 Mar 2016 - 19 Mar 2016: Gesellschaft für Informatik e.V. (GI).
- Poirazi, P., T. Brannon, and B. W. Mel (2003). Pyramidal neuron as two-layer neural network. *Neuron* 37(6), 989–999.

- Quaglio, P., V. Rostami, E. Torre, and S. Grün (2018). Methods for identification of spike patterns in massively parallel spike trains. *Biological Cybernetics*, 1–24.
- Quaglio, P., A. Yegenoglu, E. Torre, D. M. Endres, and S. Grün (2017). Detection and evaluation of spatio-temporal spike patterns in massively parallel spike train data with spade. *Frontiers in Computational Neuroscience* 11, 41.
- Ramon y Cajal, S. (1899). *Comparative study of the sensory areas of the human cortex*. Clark University.
- Raoux, S., W. Wehnich, and D. Ielmini (2009). Phase change materials and their application to nonvolatile memories. *Chemical Reviews* 110(1), 240–267.
- Renart, A., J. De La Rocha, P. Bartho, L. Hollender, N. Parga, A. Reyes, and K. D. Harris (2010). The asynchronous state in cortical circuits. *Science* 327, 587–590.
- Riehle, A., S. Grün, M. Diesmann, and A. Aertsen (1997). Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278(5345), 1950–1953.
- Riggs, L. and F. Ratliff (1952). The effects of counteracting the normal movements of the eye. In *Journal of the Optical Society of America*, Volume 42, pp. 872–873.
- Rotter, S. and M. Diesmann (1999). Exact digital simulation of time-invariant linear systems with applications to neuronal modeling. *Biological Cybernetics* 81(5-6), 381–402.
- Roxin, A., N. Brunel, and D. Hansel (2005). The role of delays in shaping spatio-temporal dynamics of neuronal activity in large networks. *Physical Review Letters* 94(23), 238103.
- Roy, K., A. Jaiswal, and P. Panda (2019). Towards spike-based machine intelligence with neuromorphic computing. *Nature* 575(7784), 607–617.
- Salinas, E. and T. J. Sejnowski (2001). Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience* 2(8), 539–550.
- Sato, T. K., I. Nauhaus, and M. Carandini (2012). Traveling waves in visual cortex. *Neuron* 75(2), 218–229.
- Sawa, A. (2008). Resistive switching in transition metal oxides. *Materials Today* 11(6), 28–36.
- Schemmel, J., D. Brüderle, A. Grübl, M. Hock, K. Meier, and S. Millner (2010). A wafer-scale neuromorphic hardware system for large-scale neural modeling. In *Proceedings of the 2010 International Symposium on Circuits and Systems (ISCAS)*, Paris, pp. 1947–1950. IEEE Press.
- Schiller, J., G. Major, H. J. Koester, and Y. Schiller (2000). Nmda spikes in basal dendrites of cortical pyramidal neurons. *Nature* 404(6775), 285–289.
- Schrader, S., S. Grün, M. Diesmann, and G. L. Gerstein (2008). Detecting synfire chain activity using massively parallel spike train recording. *Journal of Neurophysiology* 100, 2165–2176.

- Schulz, A., C. Miehl, M. J. Berry, and J. Gjorgjieva (2021). The generation of cortical novelty responses through inhibitory plasticity. *eLife* 10, e65309.
- Schuman, C. D., T. E. Potok, R. M. Patton, J. D. Birdwell, M. E. Dean, G. S. Rose, and J. S. Plank (2017). A survey of neuromorphic computing and neural networks in hardware. *ArXiv*, 1705.06963.
- Senk, J., K. Korvasová, J. Schuecker, E. Hagen, T. Tetzlaff, M. Diesmann, and M. Helias (2020). Conditions for wave trains in spiking neural networks. *Physical Review Research* 2(2).
- Shadlen, M. N. and W. T. Newsome (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *Journal of Neuroscience* 18(10), 3870–3896.
- Shanks, D. R., R. J. Tunney, and J. D. McCarthy (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making* 15(3), 233–250.
- Shatz, C. J. (1992). The developing brain. *Scientific American* 267(3), 60–67.
- Shipp, S. (2007). Structure and function of the cerebral cortex. *Current Biology* 17(12), R443–R449.
- Shoham, S., D. H. O’Connor, and R. Segev (2002). How silent is the brain: is there a “dark matter” problem in neuroscience? *Journal of Comparative Physiology* 192, 777–784.
- Shouval, H. (2010). Spike timing dependent plasticity: a consequence of more fundamental learning rules. *Frontiers in Computational Neuroscience* 4, 19.
- Shrestha, A., H. Fang, Z. Mei, D. P. Rider, Q. Wu, and Q. Qiu (2022). A survey on neuromorphic computing: Models and hardware. *IEEE Circuits and Systems Magazine* 22(2), 6–35.
- Siegel, S., Y. Bouhadjar, T. Tetzlaff, R. Waser, R. Dittmann, and D. J. Wouters (2023). System model of neuromorphic sequence learning on a memristive crossbar array. *Under review*.
- Siegel, S., T. Ziegler, Y. Bouhadjar, T. Tetzlaff, R. Waser, R. Dittmann, and D. J. Wouters (2023). Demonstration of neuromorphic sequence learning on a memristive array. *In press at the Proceedings of the Neuro-inspired Computational Elements Workshop*.
- Song, S., P. Sjöström, M. Reigl, S. Nelson, and D. Chklovskii (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLOS Biology* 3(3), e68.
- Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neuroscience* 9, 206–221.
- Stepanyants, A., J. A. Hirsch, L. M. Martinez, Z. F. Kisvárdy, A. S. Ferecskó, and D. B. Chklovskii (2007). Local potential connectivity in cat primary visual cortex. *Cerebral Cortex* 18(1), 13–28.
- Sterling, P. and S. Laughlin (2015). *Principles of neural design*. MIT press.

- Stettler, D. D., H. Yamahachi, W. Li, W. Denk, and C. D. Gilbert (2006). Axons and synaptic boutons are highly dynamic in adult visual cortex. *Neuron* 49, 877–887.
- Stroeve, S. and S. Gielen (2001). Correlation between uncoupled conductance-based integrate-and-fire neurons due to common and synchronous presynaptic firing. *Neural Computation* 13(9), 2005–2029.
- Strukov, D. B., G. S. Snider, D. R. Stewart, and R. S. Williams (2008). The missing memristor found. *Nature* 453(7191), 80–83.
- Stuart, G. J. and N. Spruston (2015). Dendritic integration: 60 years of progress. *Nature Neuroscience* 18(12), 1713–1721.
- Suri, M., D. Querlioz, O. Bichler, G. Palma, E. Vianello, D. Vuillaume, C. Gamrat, and B. DeSalvo (2013). Bio-inspired stochastic computing using binary CBRAM synapses. *IEEE Transactions on Electron Devices* 60(7), 2402–2409.
- Swadlow, H. A. (1988). Efferent neurons and suspected interneurons in binocular visual cortex of the awake rabbit: Receptive fields and binocular properties. *Journal of Neurophysiology* 59(4), 1162–1187.
- Takahashi, K., S. Kim, T. P. Coleman, K. A. Brown, A. J. Suminski, M. D. Best, and N. G. Hatsopoulos (2015). Large-scale spatiotemporal spike patterning consistent with wave propagation in motor cortex. *Nature Communications* 6(7169), 1–11.
- Tetzlaff, C., C. Kolodziejski, M. Timme, and F. Wörgötter (2011). Synaptic scaling in combination with many generic plasticity mechanisms stabilizes circuit connectivity. *Frontiers in Computational Neuroscience* 5, 47.
- Tetzlaff, T., M. Helias, G. T. Einevoll, and M. Diesmann (2012). Decorrelation of neural-network activity by inhibitory feedback. *PLOS Computational Biology* 8(8), e1002596.
- Tetzlaff, T., S. Rotter, E. Stark, M. Abeles, A. Aertsen, and M. Diesmann (2008). Dependence of neuronal correlations on filter characteristics and marginal spike-train statistics. *Neural Computation* 20(9), 2133–2184.
- Thiboust, M. (2020). Insights from the brain: The road towards machine intelligence. pp. 0–144.
- Thomas, S. A. (2015). Neuromodulatory signaling in hippocampus-dependent memory retrieval. *Hippocampus* 25(4), 415–431.
- Thomson, A. M. and P. Bannister (2003). Interlaminar connections in the neocortex. *Cerebral Cortex* 13, 5–14.
- Thomson, A. M., D. C. West, Y. Wang, and A. P. Bannister (2002). Synaptic connections and small circuits involving excitatory and inhibitory neurons in layer 2-5 of adult rat and cat neocortex: Triple intracellular recordings and biocytin labelling in vitro. *Cerebral Cortex* 12(9), 936–953.
- Tremblay, R., S. Lee, and B. Rudy (2016). GABAergic interneurons in the neocortex: from cellular properties to circuits. *Neuron* 91(2), 260–292.

- Tully, P. J., H. Lindén, M. H. Hennig, and A. Lansner (2016). Spike-based bayesian-hebbian learning of temporal sequences. *PLOS Computational Biology* 12(5), e1004954.
- Turrigiano, G. G. (2008). The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell* 135(3), 422–435.
- Turrigiano, G. G., K. R. Leslie, N. S. Desai, L. C. Rutherford, and S. B. Nelson (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391(6670), 892–896.
- Turrigiano, G. G. and S. B. Nelson (2004). Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience* 5, 97–107.
- Valov, I. (2013). Redox-based resistive switching memories (ReRAMs): Electrochemical systems at the atomic scale. *ChemElectroChem* 1(1), 26–36.
- Valov, I. and M. N. Kozicki (2013). Cation-based resistance change memory. *Journal of Physics D: Applied Physics* 46(7), 074005.
- van Vreeswijk, C. and H. Sompolinsky (1998). Chaotic balanced state in a model of cortical circuits. *Neural Computation* 10(6), 1321–1371.
- Vanni, S., H. Hokkanen, F. Werner, and A. Angelucci (2020). Anatomy and physiology of macaque visual cortical areas v1, v2, and v5/MT: Bases for biologically realistic models. *Cerebral Cortex* 30(6), 3483–3517.
- von Melchner, L., S. L. Pallas, and M. Sur (2000). Visual behavior mediated by retinal projections directed to the auditory pathway. *Nature* 404(6780), 871–876.
- von Neumann, J. (1993). First draft of a report on the EDVAC. *IEEE Annals of the History of Computing* 15(4), 27–75.
- Vulkan, N. (2000). An economist’s perspective on probability matching. *Journal of Economic Surveys* 14(1), 101–118.
- Wang, Z., S. Ambrogio, S. Balatti, and D. Ielmini (2015). A 2-transistor/1-resistor artificial synapse capable of communication and stochastic learning in neuromorphic systems. *Frontiers in Neuroscience* 8, 438.
- Wang, Z., S. Joshi, S. Savel’ev, W. Song, R. Midya, Y. Li, M. Rao, P. Yan, S. Asapu, Y. Zhuo, H. Jiang, P. Lin, C. Li, J. H. Yoon, N. K. Upadhyay, J. Zhang, M. Hu, J. P. Strachan, M. Barnell, Q. Wu, H. Wu, R. S. Williams, Q. Xia, and J. J. Yang (2018). Fully memristive neural networks for pattern classification with unsupervised learning. *Nature Electronics* 1(2), 137–145.
- Wang, Z., H. Wu, G. W. Burr, C. S. Hwang, K. L. Wang, Q. Xia, and J. J. Yang (2020). Resistive switching materials for information processing. *Nature Reviews Materials* 5(3), 173–195.
- Ward, L. M. (2003). Synchronous neural oscillations and cognitive processes. *Trends in cognitive sciences* 7(12), 553–559.
- Waser, R. (2012a). *Nanoelectronics and information technology: advanced electronic materials and novel devices*. John Wiley & Sons.

- Waser, R. (2012b). Redox-based resistive switching memories. *Journal of Nanoscience and Nanotechnology* 12(10), 7628–7640.
- Waser, R. and M. Aono (2007). Nanoionics-based resistive switching memories. *Nature Materials* 6(11), 833–840.
- Waser, R., R. Dittmann, G. Staikov, and K. Szot (2009). Redox-based resistive switching memories - nanoionic mechanisms, prospects, and challenges. *Advanced Materials* 21(25-26), 2632–2663.
- Wu, G.-Y., R. Malinow, and H. T. Cline (1996). Maturation of a central glutamatergic synapse. *Science* 274(5289), 972–976.
- Xia, Q. and J. J. Yang (2019). Memristive crossbar arrays for brain-inspired computing. *Nature Materials* 18(4), 309–323.
- Xu, S., W. Jiang, M. M. Poo, and Y. Dan (2012). Activity recall in a visual cortical ensemble. *Nature Neuroscience* 15(3), 449–455.
- Yang, J. J., D. B. Strukov, and D. R. Stewart (2013). Memristive devices for computing. *Nature nanotechnology* 8(1), 13–24.
- Yarbus, A. L. (1967). *Eye Movements and Vision*. Springer.
- Yashiro, K. and B. D. Philpot (2008). Regulation of NMDA receptor subunit expression and its implications for LTD, LTP, and metaplasticity. *Neuropharmacology* 55(7), 1081–1094.
- Yu, S. (2018). Neuro-inspired computing with emerging nonvolatile memorys. *Proceedings of the IEEE* 106(2), 260–285.
- Zamarreño-Ramos, C., L. A. Camuñas-Mesa, J. A. Pérez-Carrasco, T. Masquelier, T. Serrano-Gotarredona, and B. Linares-Barranco (2011). On spike-timing-dependent-plasticity, memristive devices, and building a self-learning visual cortex. *Frontiers in Neuroscience* 5, 26.
- Zhao, M., B. Gao, J. Tang, H. Qian, and H. Wu (2020). Reliability of analog resistive switching memory for neuromorphic computing. *Applied Physics Reviews* 7(1), 011301.
- Zhao, Y., W. Shen, P. Huang, W. Xu, M. Fan, X. Liu, and J. Kang (2019). A physics-based model of RRAM probabilistic switching for generating stable and accurate stochastic bit-streams. In *2019 IEEE International Electron Devices Meeting (IEDM)*. IEEE.
- Zhu, J., T. Zhang, Y. Yang, and R. Huang (2020). A comprehensive review on emerging artificial neuromorphic devices. *Applied Physics Reviews* 7(1), 011312.
- Zucker, R. S. and W. Regehr (2002). Short-term synaptic plasticity. *Annual Review of Physiology* 64, 355–405.

ACKNOWLEDGMENTS

During my stay at the Jülich research center, I was fortunate to be surrounded by many people who helped me answer the research questions in this thesis. I'm also grateful to all those who provided me with encouragement and motivation throughout this endeavor.

First and foremost, I would like to thank Markus Diesmann and Rainer Waser for giving me the opportunity to join their labs, for their efforts in creating a multi-disciplinary, open, and inspiring scientific environment, where I could effortlessly cultivate my scientific perspectives and conduct flexibly my research, and for their encouragement early on to participate in workshops and conferences.

I'm deeply grateful to have Tom Tetzlaff and Dirk J. Wouters as my supervisors. I would like to thank Tom Tetzlaff for devoting endless hours guiding me through this journey, for teaching me scientific writing, and for instilling in me a sense of responsibility to promote the scientific method, especially on questions surrounding the topic of reproducibility in science. You are a true inspiration. Dirk J. Wouters for teaching me how to mediate between the different disciplines involved in my thesis, for his thoroughness, and rigor, and for continually striving to produce excellent work. I also would like to thank him for the endless discussions on the projects of this thesis and for helping me navigate the hurdles of academia.

I would like to extend my thanks to Abigail Morrison for her feedback during the PhD plan meetings and for our short and successful project on reproducibility in science. I also would like to thank her for her efforts in promoting inclusive science. I would like to thank the computational neurophysics group in INM6 for the inspiring discussions, constructive criticism, and continuous feedback. Special thanks go to the NEST community including Charl Linssen, Dennis Terhorst, and Hans Ekkehard Plesser for their constant support.

Many thanks go to my collaborators and supervisee that helped shape this thesis: Sebastian Siegel, Barna Zajzon, Renato Duarte, Jette Oberländer, Hubertus Borsch, and Melissa Lober. Another thanks go to my colleagues for the diverse discussion and for never tiring teaching me science: Agnes Korcsak-Gorzo, Aitor Morales-Gregorio, Anno Kurth, Alexander van Meegen, Alexandre René, Christian Keup, Cristiano Köhler, Danylo Ulianych, and Johanna Senk, just to name a few.

I'm deeply thankful to my friends Praveen, Karim, Zakaria, Badr Eddine, and Samy for their companionship and encouragement during the whole period of this thesis, and to all my friends including Seif Eddine, Abderrahmane, Ayoub, Mohammed, and Anis for all their advice and motivation during the course of my education. Finally, I would like to express my gratitude to my parents for their boundless support, without you this thesis wouldn't have been possible. A big thank you to my siblings and my extended family for providing me with a prosperous, lively environment throughout the years.

CURRICULUM VITAE

Younes Bouhadjar

 orcid.org/0000-0003-4367-8236

EDUCATION

October 2018 - present

Doctoral candidate in Materials Science and Engineering, RWTH Aachen University, Germany

Institute of Neuroscience and Medicine (INM-6), & Institute for Advanced Simulation (IAS-6), & JARA BRAIN Institute Structure-Function Relationships (INM-10), & Peter Grünberg Institute (PGI-7,10), Jülich Research Centre, Jülich, Germany

September 2016 - June 2018

M.Sc. in Micro & Nanotechnologies for Integrated Systems, Phelma INP, Grenoble, France

Joint degree between EPFL, Switzerland; Phelma INP, Grenoble, France; and Politecnico di Torino, Italy

September 2013 – June 2016

B.Sc. in Physics and Electronics, Phelma INP, Grenoble, France

September 2010 – June 2013

Secondary Education Baccalaureate degree, Hireche Abdelmadjid, Chelghoum Laid, Algeria

COLOPHON

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede and Ivo Pletikosić. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". `classicthesis` is available for both L^AT_EX and L^yX:

<https://bitbucket.org/amiede/classicthesis/>

Band / Volume 83

Hybrid hydrogels promote physiological functionality of long-term cultured primary neuronal cells in vitro

C. Meeßen (2022), x, 120 pp

ISBN: 978-3-95806-643-4

Band / Volume 84

Surface states and Fermi-level pinning on non-polar binary and ternary (Al,Ga)N surfaces

L. Freter (2022), 137 pp

ISBN: 978-3-95806-644-1

Band / Volume 85

Dynamical and statistical structure of spatially organized neuronal networks

M. Layer (2022), xiii, 167 pp

ISBN: 978-3-95806-651-9

Band / Volume 86

Persistent firing and oscillations in the septo-hippocampal system and their relation to locomotion

K. Korvasová (2022), 111 pp

ISBN: 978-3-95806-654-0

Band / Volume 87

Sol-Gel-Synthese, Tintenstrahldruck und Blitzlampentemperatur von Tantaloxid-Dünnschichten zur pH-Messung

C. D. Beale (2022), xlix, 339 pp

ISBN: 978-3-95806-656-4

Band / Volume 88

Diversity of chiral magnetic solitons

V. Kuchkin (2022), xiv, 155 pp

ISBN: 978-3-95806-665-6

Band / Volume 89

Controlling the electrical properties of oxide heterointerfaces through their interface chemistry

M.-A. Rose (2022), vi, 162 pp

ISBN: 978-3-95806-667-0

Band / Volume 90

Modeling and Suppressing Unwanted Parasitic Interactions in Superconducting Circuits

X. Xu (2022), 123, XVIII pp

ISBN: 978-3-95806-671-7

Band / Volume 91

Activating molecular magnetism by controlled on-surface coordination

I. Cojocariu (2022), xi, 169 pp

ISBN: 978-3-95806-674-8

Band / Volume 92

Computational study of structural and optical properties of two-dimensional transition-metal dichalcogenides with implanted defects

S. H. Rost (2023), xviii, 198 pp

ISBN: 978-3-95806-682-3

Band / Volume 93

DC and RF characterization of bulk CMOS and FD-SOI devices at cryogenic temperatures with respect to quantum computing applications

A. Artanov (2023), xv, 80, xvii-lviii pp

ISBN: 978-3-95806-687-8

Band / Volume 94

HAXPES study of interface and bulk chemistry of ferroelectric HfO₂ capacitors

T. Szyjka (2023), viii, 120 pp

ISBN: 978-3-95806-692-2

Band / Volume 95

A brain inspired sequence learning algorithm and foundations of a memristive hardware implementation

Y. Bouhadjar (2023), xii, 149 pp

ISBN: 978-3-95806-693-9

Weitere **Schriften des Verlags im Forschungszentrum Jülich** unter
<http://www.zb1.fz-juelich.de/verlagextern1/index.asp>

Information

Band / Volume 95

ISBN 978-3-95806-693-9

Mitglied der Helmholtz-Gemeinschaft

