

Towards an Embodied Understanding of Neural Computation

Nareg Berberian

Thesis submitted to the University of Ottawa
in partial fulfillment of the requirements for the
PhD in Experimental Psychology



uOttawa

School of Psychology
Faculty of Social Sciences
University of Ottawa

© Nareg Berberian, Ottawa, Canada, 2021

Acknowledgments

I would like to thank my supervisor, professor Sylvain Chartier, for giving me the opportunity to work in his lab, which I will always remember as an inspiring environment filled with fascinating and exceptional ideas. Thank you Sylvain for the continuous guidance and support, they have left an ever-lasting impression that have helped shape the scientist I have become. This thesis would not have been possible without you, and your ingenuity. I would also like to thank professor Jean-Philippe Thivierge, for leading by example in staying dedicated in my work, and for guiding me in the exploration of new and challenging avenues in computational neuroscience. I would like to take the opportunity to show my gratitude to the examiners of the dissertation, professor Richard Naud, professor Stuart Fogel, professor Denis Cousineau and professor Poramate Manoonpong, your acceptance in being part of the committee of evaluators is much appreciated. I would also like to thank my wonderful colleagues for all the memorable moments we have spent inside and outside academia together: Matt Ross, Kinsey Church, Damiem and Thaddé Rollon-Mérette, André Cyr, Laurence Morissette, Eric Kuebler, Melissa Johnson, Bradley Harding, Marc-André Goulet, Matias Calderini, Philippe Vincent-Lamarre, Cassandra Morrison and Farooq Kamal. Last, but not least, I would like to thank my loving mother and father for their unconditional support, as well as my sister and brother for always being there for me. To my loving partner Taline, thank you for complementing me.

Contents

Introduction	xv
General overview and objectives of this work	xv
Short-term synaptic plasticity	xviii
Phenomenological model	xx
Long-term synaptic plasticity	xxi
Spike-timing-dependent plasticity	xxi
Spiking neuron models	xxv
Leaky integrate-and-fire model	xxv
Brief history on inspirations from nature	xxv
Desirable network properties in adaptive learning systems	xxvii
Online learning	xxx
General overview and summary of chapters	xxxiii
1 Synergy Between Short-Term and Long-Term Plasticity in Vision	1
1.1 Abstract	1
1.2 Introduction	2
1.3 Methods	5
1.3.1 Short-term plasticity (STP)	5
1.3.2 Spike-timing-dependent plasticity (STDP)	7
1.3.3 Numerical procedure	10

1.3.4	Procedure for searching candidate pairs of baseline Pr	12
1.4	Results	13
1.5	Discussion	18
1.6	Conclusion	21
2	Discrimination of Motion Direction in a Robot Using a Phenomenological Model of Synaptic Plasticity	22
2.1	Abstract	22
2.2	Introduction	23
2.3	Methods	26
2.3.1	Network architecture	26
2.3.2	Robotic setup	28
2.3.3	Computational model	28
2.3.4	Experimental data	35
2.4	Results	36
2.5	Discussion	42
2.6	Conclusion	46
3	Embodied Working Memory During Ongoing Input Streams	47
3.1	Abstract	47
3.2	Introduction	48
3.3	Methods	51
3.3.1	Network architecture	51
3.3.2	Robotic platform	52
3.3.3	Sensory inputs	53
3.3.4	Learning and recall procedure	56
3.3.5	Spiking neural model	60
3.3.6	Plasticity model	61

3.4	Results	65
3.4.1	Experiment 1: single learning and recall	66
3.4.2	Experiment 2: incremental learning and recall	68
3.4.3	Experiment 3: task switching	68
3.4.4	Experiment 4: resisting interfering inputs	69
3.4.5	Experiment 5: submitting to interfering inputs	70
3.4.6	Experiment 6: resisting distraction inputs	72
3.5	Discussion	74
3.5.1	Relation to previous work	74
3.5.2	Limitations	78
3.5.3	Practical implications	81
3.6	Supporting information	83
	Discussion	85
	Neuroscience-inspired robotics	85
	Human-robot interaction	87
	Humble machine intelligence	91
	Concluding statement	93
	A Analytical solutions of the passive integrate-and-fire model for a constant current injection	94
A.1	Membrane potential evolution in the absence of threshold and reset rule for action potential generation	94
A.2	Interspike interval and spike count prediction	96

List of Figures

Figure 1	Description of synaptic transmission.	xx
Figure 2	The phenomenological TM model of STP.	xxii
Figure 3	Three variations of the STDP protocol.	xxiv
Figure 1.1	Model flow chart.	4
Figure 1.2	Asymmetric learning rule for spike-timing-dependent plasticity.	9
Figure 1.3	Numerical procedure of the model.	11
Figure 1.4	Evolution of w_{ij} as a function of learning trial number.	14
Figure 1.5	Distribution of w_{ij} at learning trial 1000.	15
Figure 1.6	Stabilized weight w_{ij} as a function of direction step.	15
Figure 1.7	Firing rate of each unit as a function of direction step.	16
Figure 1.8	Baseline Pr pairs as a function of direction step.	17
Figure 1.9	Stabilization of w_{ij} in the left visual field during stimulus position discrimination.	18
Figure 1.10	Firing rate of each unit as a function of direction step.	19
Figure 1.11	Stabilization of w_{ij} in the right visual field during stimulus position discrimination.	20
Figure 2.1	Architecture of the microcircuit.	27
Figure 2.2	Robotic setup.	29
Figure 2.3	Pre-processing steps for proper image processing.	32

Figure 2.4	Initial release probability of STF synapses as a function of the direction step number.	34
Figure 2.5	Right motion discrimination.	37
Figure 2.6	Left motion discrimination.	38
Figure 2.7	Response of two subpopulations during the simulation, robotic imple- mentation, and multi-electrode recordings of V1.	39
Figure 2.8	Evolution of the TM model state variables during left and right motion discrimination.	40
Figure 2.9	Direction selectivity in the two-unit model.	41
Figure 2.10	Spike count correlations predict shared and unshared motion direction preferences.	43
Figure 3.1	Vector the robot.	54
Figure 3.2	Sensory afferents to the embodied network.	55
Figure 3.3	Keyboard listening setup.	57
Figure 3.4	Single unit profile of the aEIF model.	61
Figure 3.5	Synaptic plasticity profile.	63
Figure 3.6	Model flow chart.	65
Figure 3.7	Input-dependent and input-disengaged synaptic plasticity differentially modulate neuronal responses.	66
Figure 3.8	Single learning and recall.	67
Figure 3.9	Incremental learning and recall.	69
Figure 3.10	Task switching.	70
Figure 3.11	Resisting interfering inputs.	71
Figure 3.12	Submitting to interfering inputs.	72
Figure 3.13	Resisting distraction inputs.	73

Figure A.1	Evolution of membrane potential across time in response to constant current injection lasting for 350ms.	96
Figure A.2	Interspike intervals as a function of constant current injection.	97
Figure A.3	Spike count as a function of constant current injection of $t_{max} = 350$ ms.	98

List of Tables

Table 1.1	Short-term synaptic plasticity parameters.	7
Table 1.2	Single cell parameters.	7
Table 1.3	Spike-timing-dependent plasticity parameters.	8
Table 1.4	Temporal parameters.	10
Table 2.1	Short-term synaptic plasticity parameters.	30
Table 2.2	Leaky integrate-and-fire parameters.	33
Table 3.1	Network parameters.	52

Preface

While adaptive and intelligent machines continue to harness engineering principles and materials, our understanding of what biological mechanisms are relevant for creating powerful and efficient robots remains a challenging endeavour. Despite the advances of appropriate computing hardware, theoretical frameworks of brain function are more commonly linked to data from experimental neuroscience, but their application in neurorobotic contexts have received less attention from the literature. This challenge partly stems from an attempt to accurately represent the neurophysiological basis of biological behaviour within systems of artificial intelligence. In general, robots are mobile and they often learn from their interactions with the environment, which displays a vast amount of variability and complexity.

In an attempt to gain valuable insights into brain function, brain-inspired neurorobotics has emerged as an interdisciplinary field that promotes the collaboration amongst neuroscientists, roboticists and artificial intelligence researchers. The motivation behind the work presented here is to use bio-inspired robots to validate any hypothesized understanding of neuronal computation. Inspiring from neural circuits and the properties they may carry, this thesis is an attempt to build more flexible and adaptive machines. In my dissertation, I focus on biologically-inspired networks of spiking neurons, with the intent of following general principles of neural computation, such as the communication through discontinuous action potentials, the reliance on local rules for synaptic plasticity, as well as the inherent probabilistic nature of neurotransmitter release. In order to study how these biological principles can account for intelligent behaviour in artificial systems, networks of spiking neurons are embedded in physical robots. Altogether, the robotic implementation is used as a methodological tool for validating the functional contribution of the computational models proposed, and ultimately in proposing testable experimental predictions.

The following articles, proceedings, and conference presentations encompass part of the work presented in this thesis:

Berberian, N., Ross, M., Chartier, S. (2021). Embodied Working Memory During Ongoing Input Streams. PLoS ONE, 1-29. <https://doi.org/10.1371/journal.pone.0244822>.

Berberian N., Ross M., Chartier S. (2019). Discrimination of Motion Direction in a Robot Using a Phenomenological Model of Synaptic Plasticity. Computational Intelligence and Neuroscience in Neurorobotics, 1-14. <https://doi.org/10.1155/2019/6989128>.

Berberian, N., Ross, M., Chartier, S. (2019, March). Le cerveau du robot durant sa vision: une étude en neuro-robotique. 41st Annual Meeting of Société Québécoise pour la Recherche en Psychologie (SQRP), Mont-Temblant, Canada.

Berberian, N., Ross, M., Thivierge, J.P., Chartier, S. (2018, October). Development of Direction Selectivity via a Synergistic Interaction Between Short-term and Long-term Synaptic Plasticity. BMC Neuroscience. <https://doi.org/10.1186/s12868-018-0451-y>.

Berberian, N., Ross, M., Thivierge, J.P., Chartier, S. (2018, July). Development of Direction Selectivity via a Synergistic Interaction Between Short-term and Long-term Synaptic Plasticity. 27th Annual Meeting of Computational Neuroscience (CNS), Seattle, USA.

Berberian, N., MacPherson, A., Giraud, E., Richardson, L., Thivierge, J.P. (2017). Neuronal Pattern Separation of Motion-relevant Input in LIP activity. Journal of Neurophysiology, 738-755. <https://doi.org/10.1152/jn.00145.2016>.

Berberian, N., Ross, M., Chartier, S. (2018, March). L'interaction entre la plasticité à court terme et la plasticité à long terme explique l'émergence de la sélectivité de direction des stimuli en mouvement. 40th Annual Meeting the Société Québécoise pour la Recherche en Psychologie (SQRP), Quebec City, Canada.

Berberian N., Ross M., Chartier S., Thivierge J.P. (2017, September). Synergy Between Short-term and Long-term Plasticity Explains Direction Selectivity in Visual Cortex. IEEE Symposium Series on Computational Intelligence (IEEE SSCI), 1-8.

<https://doi.org/10.1109/SSCI.2017.8280986>.

Berberian, N., Ross, M., Chartier, S., Thivierge, J.P. (2017, November). Initial Release Probability during Short-term Plasticity Modulates the Strength of Long-term Depression-mediated Weight Convergence. 10th Annual IEEE Symposium Series on Computational Intelligence (IEEE SSCI), Honolulu, USA.

Berberian, N., Aamir, Z., Hélie, S., Chartier S. (2016, March). Encoding Sparse Features in a Bidirectional Associative Memory. International Joint Conference on Neural Networks (IJCNN), 5119-5126. <https://doi.org/10.1109/IJCNN.2016.7727875>.

Summary

Information processing systems (IPSs) that behave in the real-world are constantly bombarded with noise from the environment. Although the real-world offers noise for free, this extrinsic noise source has a cost associated to it. The problem is related to the fact that the environmental intricacies (i.e. noisy distribution) in which the IPS samples from is ever-changing. Consequently, the IPS is faced with the conundrum of maintaining stability in a dynamic environment, while at the same time, remaining flexible so as to match its intrinsic timescales with the timescales of the environment. Here, we propose conjoining three ingredients for solving the timescale incompatibility issue between IPSs and the environment. First, we propose evolving IPSs in an online fashion such that the system operates on-the-fly. Second, we propose the implementation of online operations in robotic hardware – a methodological tool for allowing the IPS to provide feedback during its interaction with the environment. Finally, we propose learning from brain mechanisms as a source of inspiration for building more flexible and adaptive IPSs.

In chapter one, we initiate our first attempt towards achieving flexibility in these systems. To this end, we study the interaction between short-term plasticity (STP) and spike-timing-dependent plasticity (STDP), two important plasticity rules we’ve learned from the brain. As such, we construct a microcircuit motif of two units, and show in simulation, how each unit can discriminate the position of a moving stimulus. In chapter two, we study synaptic plasticity in the context of an online robotic domain. To do so, we increase network size to six units, and endow the circuit with STP as a candidate mechanism for microcircuit sensitivity to inputs. Here, we study motion discrimination using a Raspberry Pi microcontroller as the information processing unit. We also use a stationary camera to process images from the real-world. Finally, we attach two LED light sensors for providing feedback of how the system is behaving. Results show that the agent is capable of discriminating the direction of a moving stimulus. In the final chapter of the thesis, we move away from the static online robotic implementation, towards a more dynamic setting. In doing so, we develop a keyboard

listener for online mobile robot control. Here, the motor trajectory of the robot is directly linked to network activity of 500 units. Furthermore, the agent is placed in an ecological context where it interacts with a human subject. During human-robot interaction, the motor trajectory of the robot is studied, enabling the human to make inferences about how neural computation is unfolding on-the-fly. The robot illustrates useful properties, one of which is high degree of flexibility and adaptation to ongoing input streams. Overall, we conjoin the three ingredients mentioned above as a framework for solving the timescale incompatibility issue between IPSs and the environment.

Introduction

General overview and objectives of this work

The amount of information the nervous system is capable of processing throughout its lifetime is simply astonishing. From exploration through movement, to relaxation through stillness, the brain leads the purpose of adapting and ultimately surviving in an ever changing environment. To interact with the external world, the sensory organs inevitably receive inputs on one side and generate motor outputs on the other side. The information processing capability that lies between input and output is based on an incredible complexity of brain areas. Arguably, one of the central questions in neuroscience relates to understanding how the brain generates stability and flexibility during ongoing sensory experience. Why is it important not to forget? From maladaptive changes, to severe effects on quality of life, the importance of learning and memory is most highlighted in the absence thereof. In an attempt to answer questions related to learning and memory, our understanding of brain plasticity becomes a fundamental requirement.

In the mammalian brain, the computational capabilities of a single neuron is highly limited. These limitations are overcome by combining several individual neurons in an intricately connected network. On average, a single neuron forms 1000 to 10 000 connections to other neurons, which amounts to a total of approximately 10^{14} to 10^{15} neural connections called synapses (Kandel, Schwartz, & Jessell, 2000). Importantly, these intricate connections can change over time, a phenomenon known as synaptic plasticity. Under the basis of such

refinement, the synaptic circuitry of the brain undergoes both anatomical and functional modifications essential for the emergence of functional neuronal circuits.

To lead an efficient communication system between neurons, the brain must provide algorithmic processes to change the synaptic connections, and ultimately maintain proper functionality. These algorithmic processes are generally characterized by a set of plasticity rules, which may act on multiple timescales (Duarte, Seeholzer, Zilles, & Morrison, 2017). Over the years, there has been surmounting experimental evidence suggesting that synaptic plasticity may operate over short and long timescales (Clemens, Ritter, Roy, Miller, & Van Hooser, 2012; R. P. Rao & Sejnowski, 2000). For example, at many synapses, presynaptic action potentials can rapidly and dynamically affect synaptic efficacy (Fisher, 1997; Magleby, 1987; Zador & Dobrunz, 1997; R. Zucker, 1989). Indeed, the amplitude of the synaptic response is not a static quantity, but rather increases or decreases due to recent historical activity. This ubiquitous feature is referred to as short-term plasticity (STP), reflecting transient changes in synaptic efficacy. These changes may vary over the course of hundreds of milliseconds to at most a few minutes (Abbott & Regehr, 2004; R. S. Zucker & Regehr, 2002). In addition to short-term changes in synaptic plasticity, synapses can also behave over longer timescales, ranging from a few minutes up to a few days (Bliss & Collingridge, 1993; Bliss & Lømo, 1973; Redondo & Morris, 2011; Reymann & Frey, 2007; Schuett, Bonhoeffer, & Hübener, 2001). In this context, the timing precision between presynaptic and postsynaptic activity may induce long-lasting changes at the synapse (Nabavi et al., 2014). Nevertheless, other factors besides timing precision may contribute to long-term changes in synaptic strength. For example, neuromodulatory signals can steer plasticity by (1) increasing or decreasing the sensitivity of neurons to STDP paradigms, (2) modifying the shape of the STDP function, (3) determining whether synapses undergo potentiation or depression, and (4) altering presynaptic glutamate release (see (Roelfsema & Holtmaat, 2018) for review).

Given the wide range of plasticity rules in the brain (Markram, 2011), it is likely to expect

that mechanisms of short-term and long-term synaptic plasticity play a complementary role in mediating the efficacy of synaptic transmission. Indeed, multiple experimental studies have shown that the two mechanisms may interact (Jin et al., 2012; Kintscher, Wozny, Johenning, Schmitz, & Breustedt, 2013; Markram & Tsodyks, 1996; Tokuoka & Goda, 2008). For example, nitric oxide (NO) and endocannabinoid (eCB) retrograde signals can lead to the expression of presynaptic long-term potentiation (pre-LTP) and long-term depression (pre-LTD) respectively (P. Sjöström, Turrigiano, & Nelson, 2003, 2007). These retrograde signals may act as coincidence detectors for pre- and postsynaptic spike timing, essentially contributing to the long-term modification of STP dynamics via changes in the probability of neurotransmitter release (Markram, Wang, & Tsodyks, 1998).

Consistent with the multi-timescale interaction of synaptic plasticity, there is mounting evidence to suggest that the locus of expression for synaptic plasticity can be either presynaptic, postsynaptic, or a combination of both (Costa, Mizusaki, Sjöström, & van Rossum, 2017). For example, when synaptic plasticity is expressed at the presynaptic terminal, newly established contacts from the presynaptic neuron undergo long-term changes in neurotransmitter release (Monday & Castillo, 2017). Here, the genesis of new contacts results in the redistribution of synaptic efficacies between newly created active zones – a notion commonly known as redistribution of synaptic efficacy (Markram & Tsodyks, 1996). This effect is in clear contrast to traditional Hebbian rules (i.e. STDP), where plasticity is commonly expressed at the postsynaptic terminal (Bi & Poo, 1998).

In general, STP and STDP have been involved in temporal processing (Abbott, Varela, Sen, & Nelson, 1997; R. Zucker, 1989) and memory storage (Markram, 2011; Nabavi et al., 2014), respectively. However, the functional role of their interaction remains unclear (but see (Froemke & Dan, 2002; Payeur, Guerguiev, Zenke, Richards, & Naud, 2020)). How do short-term changes in synaptic plasticity work their way in reorganizing embodied local neuronal circuits with STDP? In an attempt to study their complementary relationship, this thesis will consider the mutuality of both mechanisms, and embed their interaction in networks of

spiking neurons. To quantify the functional role of their interaction, the formal model will be embodied by a physical robot. Furthermore, the agent will be placed in an ecological context under which it will interact with a human. The agent will be exposed to non-stationary streams of input, and the interaction will take place online – continuously evolving on-the-fly. Ultimately, both successful or unsuccessful robotic behaviour will be revealed by reverting back to the details of the embodied network of spiking neurons – exposing, both behaviourally and computationally, the complementary role of short-term and long-term synaptic plasticity in explaining higher-order cognition such as working memory. As the chapters of this thesis unfold, we will aim to achieve this objective by gradually working our way up towards larger networks of spiking neurons. As we progress towards later chapters, the ecological niche of our neurorobotic agents, that is, the environmental context in which they behave will gradually become more sophisticated also, and their internal dynamics more biologically constrained.

Short-term synaptic plasticity

Amongst rapid changes in synaptic plasticity, short-term depression is a particularly prominent feature of transmission at many synapses, whereby presynaptic activity produces a decrease in synaptic strength (Magleby, 1987). A common explanation for this decrease is related to the depletion of synaptic vesicles at the presynaptic terminal (Hubbard, 1963; Liley & North, 1953; C. Stevens & Wang, 1995). Upon the arrival of an incoming action potential, the concentration of calcium ions $[Ca^{2+}]_i$ transiently increases inside the presynaptic terminal. As a result of the sudden elevation in $[Ca^{2+}]_i$, a small pool of presynaptic vesicles along the active zone, the readily releasable pool (RRP), rapidly release neurotransmitters into the synaptic cleft (Rosenmund & Stevens, 1996). As a result of this process of exocytosis, the number of presynaptic vesicles at the RRP is transiently reduced, as the separation of vesicles from the plasma membrane and the replenishment of neurotransmitters in presy-

naptic terminal take some time. Consequently, the synapse becomes momentarily depressed, as it transitions into a state of low release probability. Between subsequent action potentials, new vesicles, and/or those from the reserve pool may dock again, thus returning transmission probability back to its rest level with a time constant in the order of milliseconds to a second (Dobrunz & Stevens, 1997; Matveev & Wang, 2000).

In contrast to short-term depression, the mechanism of short-term facilitation displays an enhancement of the synaptic response to changes in presynaptic activity (Fisher, 1997; R. S. Zucker & Regehr, 2002). Paired-pulse ratio is a common experimental procedure for assessing facilitation, where a synapse is stimulated with a pair of consecutive pulses. In facilitating synapses, the response amplitude to the second pulse is typically larger than the response amplitude to the first pulse. If the window between the two pulses is large enough, the effect of facilitation disappears, and the response amplitude to both pulses is equal on average. Residual calcium is thought to be the underlying biophysical substrate of short-term facilitation (Katz & Miledi, 1968). Following an action potential, the influx of calcium elevates the concentration of calcium ions $[Ca^{2+}]_i$, which may result in vesicle exocytosis. If a second action potential arrives before $[Ca^{2+}]_i$ has time to recover to a baseline level, the new incoming calcium will add to the pre-existing residual calcium, thus enhancing release probability (R. S. Zucker, 1996).

Over the years, there have been an extensive number of studies on STP (R. S. Zucker & Regehr, 2002). This ubiquitous feature in neural circuits arises because neurotransmitter release is history dependent. Due to the inherent probabilistic nature of neurotransmitter release dynamics, STP has been taken in a statistical direction from early on to better understand its mechanisms (del Castillo & Katz, 1954). The functional role of STP has been examined by a plethora of experimental and theoretical studies. STP has been implicated in a wide range of computations, namely in modifying a neuron's sensitivity to the temporal coherence of its inputs (Abbott et al., 1997; Markram & Tsodyks, 1996), participating in gain control (Abbott et al., 1997), reducing redundancy (Goldman, Maldonado, & Abbott, 2002),

acting as an adaptive filter (van Rossum, van der Meer, Xiao, & Oram, 2008), and separating streams of information (Naud & Sprekeler, 2018), among others (Tauffer & Kumar, 2019).

Phenomenological model

In the late 90s, Tsodyks and Markram developed a model of STP, here referred to as the Tsodyks-Markram (TM) model (Markram et al., 1998; Tsodyks & Markram, 1997; Tsodyks, Pawelzik, & Markram, 1998). This STP model is phenomenological, because the noise coming from the stochastic recovery of the vesicle has been eliminated by either pooling the response from many independent synaptic contacts or by taking a trial-averaged response of the stochastic recovery of the vesicle from a single synaptic contact (Markram et al., 1998). Although the model does not catch the full complexity of the nature of synaptic release, it has received noticeable attention in the literature, potentially due to its appealing simplicity and its ability to account for biophysically relevant parameters (Hennig, 2013; Le Be & Markram, 2006; Markram et al., 1998; Ramaswamy et al., 2012; Richardson, Melamed, Silberberg, Gerstner, & Markram, 2005; Rinaldi, Silberberg, & Markram, 2008; Romani et al., 2013; Y. Wang et al., 2006).




Figure 1: **Description of synaptic transmission.** Presynaptic vesicles occupy an active zone in the presynaptic terminal (delimited by the brown line). Vesicles from this readily-releasable pool may bind to the plasma membrane and release neurotransmitters with a given probability (P_{rel}). In turn, these neurotransmitters may bind to postsynaptic receptors such as AMPA or NMDA (red structures).

The phenomenological TM model of STP is described by the following pair of ordinary differential equations:

$$\frac{du_j}{dt} = \frac{U - u_j}{\tau_f} + U(1 - u_j) \sum_k \delta(t - t_j^k) \quad (1)$$

$$\frac{dx_j}{dt} = \frac{1 - x_j}{\tau_d} - u_j x_j \sum_k \delta(t - t_j^k) \quad (2)$$

where δ is the Dirac delta function, t_j^k is the occurrence time of the k th spike of presynaptic unit j , $U \in [0, 1]$ represents the baseline release probability of u_j , and x_j is the neurotransmitter availability. In the absence of presynaptic spikes, unit j is at a resting state, where u_j and x_j remain at baseline resting values U and 1, respectively. When a spike arrives at the presynaptic terminal, u_j instantaneously increases $u_j \rightarrow u_j + U(1 - u_j)$, while x_j instantaneously decreases $x_j \rightarrow (1 - u_j)x_j$. Between subsequent spikes, u_j and x_j recover back to their baseline resting values U and 1, respectively. Depending on the initial setup of STP parameters, τ_f , τ_d and U , the TM model can mimic the effect of a depressing or a facilitating synapse (Fig 2). Therefore, the mechanism of short-term depression and facilitation can be distinguished using a different parameter setup in the same governing equations (Markram et al., 1998).

Long-term synaptic plasticity

Spike-timing-dependent plasticity

Since Donald Hebb's establishment of the notion of activity-dependent synaptic modification, plasticity in neuronal circuits has been considered as a fundamental mechanism for learning and memory (Hebb, 1949). Based on Hebb's principle, there have been a wide range of efforts to explain the biophysical implications of memory consolidation – from changes in synaptic strength at the molecular level all the way up to the formation and maintenance of

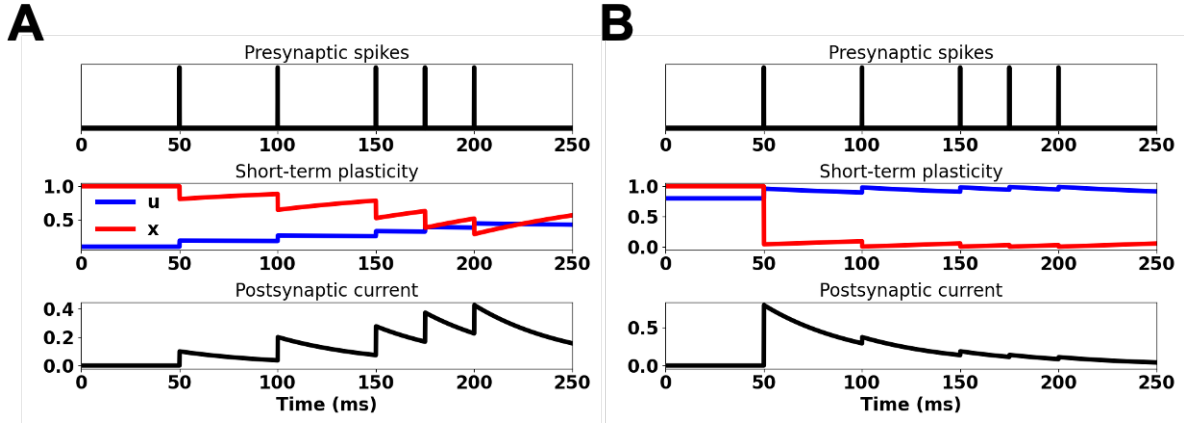


Figure 2: **The phenomenological TM model of STP.** (A) Postsynaptic current (bottom) mediated by facilitation-dominated state variables u and x (centre) in response to presynaptic inputs (top). $U = 0.1, \tau_f = 900, \tau_d = 100$ (B) Postsynaptic current (bottom) mediated by depression-dominated state variables u and x (centre) in response to presynaptic inputs (top). Legends same as (A). $U = 0.8, \tau_f = 100, \tau_d = 900$; STP parameters taken from (Vasilaki & Giugliano, 2014).

cell assemblies at the network level. Amongst those studies, experimental evidence suggests that the relative timing of pre- and postsynaptic action potentials plays a crucial role in changing the strength of synaptic connections (Bi & Poo, 1998; Markram, Lübke, Frotscher, & Sakmann, 1997). This phenomenon is commonly known as spike-timing-dependent plasticity (STDP). In recent years, STDP has emerged as the experimentally most studied form of synaptic plasticity (Bi & Poo, 1998; Froemke & Dan, 2002; Markram et al., 1997; P. J. Sjöström, Turrigiano, & Nelson, 2001). STDP is considered a key mechanism for the formation and maintenance of memory, and its underlying functional properties have been featured in a wide range of studies (Feldman, 2012).

In cellular neurophysiology, STDP protocols induce plasticity by either strengthening synaptic connections via long-term potentiation (LTP) or by weakening them via long-term depression (LTD) (Nabavi et al., 2014). In general, plasticity via LTP is induced when presynaptic spikes precede postsynaptic spikes by ~ 20 ms, whereas LTD is induced when presynaptic spikes follow postsynaptic spikes by ~ 50 ms (P. J. Sjöström et al., 2001). Consequently, the temporal coincidence window for plasticity induction is in the order of a few

tens of milliseconds, where LTD is accompanied by a wider induction window than LTP (Markram et al., 1997). Despite this short time interval under which plasticity is induced, STDP has been referred to as a form of long-term plasticity due to its involvement in mediating long-lasting changes in the synapse – those that may persist over a wide range of timescales ranging from a few minutes up to a few days (Bliss & Collingridge, 1993; Bliss & Lømo, 1973; Redondo & Morris, 2011; Reymann & Frey, 2007; Schuett et al., 2001). Consequently, the timescale for inducing a change in the synapse may substantially differ from the timescale of the resulting persistence (Morrison, Diesmann, & Gerstner, 2008).

Computational models generally use one of three distinct variations of the STDP induction protocol. The first two are generally classified as pair-wise, namely all-to-all or nearest-spike interactions (Fig 3A-B). The all-to-all spike timing interaction is the simpler rule of them all (Gerstner, Kempter, van Hemmen, & Wagner, 1996; Kempter, Gerstner, & van Hemmen, 1999; Kistler & van Hemmen, 2000; Song, Miller, & Abbott, 2000). For a given train of pre- and postsynaptic spikes, the timing between all pre- and postsynaptic spikes is compared, giving rise to a matrix of time differences (Fig 3A). The total weight change induced from pre-post/post-pre spike timing is the linear sum of all these time differences.

The second variation of the STDP induction protocol is the nearest-neighbour spike timing interaction, which is a slight variant of the all-to-all weight updating scheme (Bi, 2002; Burkitt, Meffin, & Grayden, 2004; Izhikevich & Desai, 2003; Pfister & Gerstner, 2005; van Rossum, Bi, & Turrigiano, 2000). This rule is more technical to implement from a computational standpoint, especially in online learning scenarios where ongoing network computations are evaluated at individual time instances. The nearest-spike interaction computes a time difference between the last spikes emitted from pre- and postsynaptic units (Fig 3B). In this way, the nearest spike timing difference between connected pairs of neurons is evaluated at individual time instances.

Finally, the third variation of the STDP induction protocol is the triplet model, where the elementary building block for updating the weights is not a pair, but a combination of

one presynaptic and two postsynaptic spikes (post-pre-post) or one postsynaptic and two presynaptic spikes (pre-post-pre) (Pfister, 2006; Pfister & Gerstner, 2005) (Fig 3C). This form of induction protocol is usually implemented online. For example, one presynaptic spike (x_+) and two postsynaptic spikes ($y_+; y_-$) can each leave a trace that decays with a time constant that is tied to a biophysical quantity. Here, x_+ tracks past pre activity (e.g. glutamate binding to postsynaptic NMDA receptors). y_+ can represent post activity (e.g. nitric oxide signalling) and when paired with pre spikes, LTP is expressed. y_- tracks post activity (e.g. endocannabinoid signalling) and when paired with pre spikes, LTD is expressed.

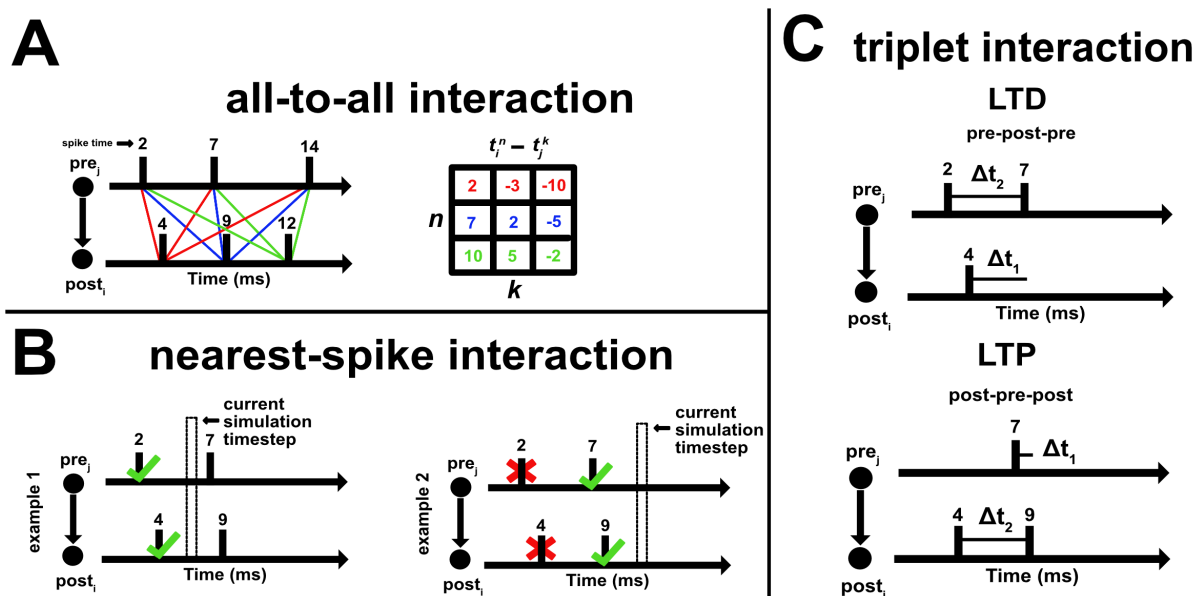


Figure 3: **Three variations of the STDP protocol.** (A) All-to-all interaction. (B) Nearest-spike (i.e. neighbour) interaction. (C) Triplet interaction. (top) The first presynaptic spike contributes to the induction if the interval between presynaptic spikes is no greater than the decay time constant of the first presynaptic spike. (bottom) The presynaptic spike must occur before the second postsynaptic spike with a time difference no greater than the decay time constant of the presynaptic spike.

Amongst the three induction protocols, the triplet model for STDP has been shown to be the most consistent with at least one particular set of experiments (P. J. Sjöström et al., 2001). For induction protocols at low frequencies, both pair-based models and the triplet model accurately predict the occurrence of potentiation and depression. However,

at high frequency pairings, the pair-based rules fall short, and the triplet model prevails in accurately predicting the induction of LTP and LTD. For the purpose of simplicity, the thesis will focus on pair-based spike interactions as means for modelling effects of long-term synaptic plasticity. More specifically, the all-to-all interaction will be used in chapter one, whereas chapter three will consider the nearest-neighbour interaction.

Spiking neuron models

Leaky integrate-and-fire model

In 1907, Louis Édouard Lapicque (1866-1952) published a highly influential paper, proposing how neurons can be studied in the context of a parallel resistor-capacitor (RC) circuit driven by an input (Lapicque, 1907). This mathematical formalism is known as the leaky integrate-and-fire (LIF) model, and was proposed long before our understanding of the underlying mechanisms of action potential generation (Brunel & van Rossum, 2007; Dayan & Abbott, 2001). Highlighted for its simplicity, the LIF model is widely used in computational neuroscience, where researchers can provide a mathematical formalism to neural computation. As a reasonable approximation for the behaviour of some neurons under highly specified conditions, the integrate-and-fire model can be studied analytically in order to approximate simplified descriptions of sub-threshold membrane potential dynamics as well as action potential generation. For example, the firing rate of an integrate-and-fire model in response to a constant current injection can be computed analytically (see Analytical solutions of the passive integrate-and-fire model for a constant current injection).

Brief history on inspirations from nature

In 1943, Warren McCulloch (1898-1969) and Walter Pitts (1923-1969) published a seminal paper, demonstrating how conditional statements can be used as a method for describing the

“all-or-none” character of cells in the nervous system (McCulloch & Pitts, 1943). As such, each unit was characterized as either “on” (i.e. active mode) or “off” (i.e. silent mode). If the stimulation of afferent units were strong enough to drive the response of the output unit above a fixed threshold, the unit would switch to “on”. Based on this principle, McCulloch and Pitts suggested that the activity of individual neurons may be represented by a formal analysis of propositional logic (Russel & Norvig, 2012). In this way, they went on to propose that all the logical connectives (AND, OR, NOT etc.) could be used to “record the behaviour of complicated nets” (McCulloch & Pitts, 1943). Early developments as such exemplified how biological systems may provide a source of inspiration for modelling brain function.

In the same decade, neurologist William Grey Walter (1910-1977) introduced the “tortoise” (i.e. *Machina speculatrix*), a three-wheeled mobile robot inspired from turtles, and used as a methodological tool for understanding animal behaviour (Walter Grey, 1950, 1953). To this end, the tortoise carried a photoelectric sensor for detecting light, a tactile sensor for detecting obstacles, and various other hardware components. As it scanned the environment in search for light and avoiding obstacles, the tortoise hardly ever repeated the same action twice – a feature that was attributed to its sensory-motor control design. Walter had carried out his robotic experiments in an attempt to explore how complex behaviours can emerge from simple constituents. In particular, he raised a fundamental research question for which its implications remain relevant today. What is the maximum number different behavioural patterns that could be generated from a robot, using a simple working model of two cells capable of being interconnected in multiple ways? In other words, how can embodied systems equipped with such small number of cells create a large number of behavioural patterns?

Indeed, it was difficult to conceive how the wide repertoire of behaviours observed from the tortoise could be solely attributed to the sophistication of the microcircuit architecture. After all, those were limited in size and in structural arrangement (Walter Grey, 1953). Consequently, the behaviours that the agent expressed were not solely the outcome of an internal control structure (Pfeifer, Lungarella, & Iida, 2007). Rather, a noticeable aspect of

Walter's robot was the ecological niche (i.e. environmental context) in which it operated. Although seemingly trivial, venturing away from highly controlled laboratory settings by moving closer to the real world created opportunistic interactions that would otherwise be difficult to simulate/operationalize computationally.

Today, it is difficult to conceive how Walter's early formulations on behaviour-based robotics served as a utility for making serious attempts at addressing biological questions (Webb, 2001). However, the quest for understanding the behaviour of living beings and their relation to neural activity is far from being outdated, and arguable remains on the surface of neuroscientific scrutiny (Krakauer, Ghazanfar, Gomez-Marin, MacIver, & Poeppel, 2017). Overall, Walter believed that robots must resemble living beings not in looks (Morimoto & Kawato, 2015), but rather by action. As such, the tortoise he built drew inspirations from the behaviour of animals. Taken together, early physical agents were built by drawing inspirations from animals and used as tools for understanding their behaviour, the same way McCulloch and Pitts neurons were inspired from the physiology of nerve cells to understand neural network function.

Desirable network properties in adaptive learning systems

A formal notion of local and distributed representations was introduced by Tim van Gelder, describing them in terms of their *extendedness* (Browne, 2019; van Gelder, 1991). In the localist approach, extendedness means that an item is not represented by a single unit (e.g. grandmother cell), but rather over an extended proportion of multiple available units in the network. Here, localist means that each item is being represented by the activity of multiple units. Importantly, two or more items require the participation of two or more separate unit ensembles, such that each unit ensemble can disambiguate the representation of a single item. As such, localist models endowed with extendedness make an implicit assumption, namely

that individual units are insufficient for representing a single item, because the representation of a single item is spread out (i.e. extended) and processed by the participation of multiple units (van Gelder, 1991). Hence, in the localist formulation of extendedness, multiple units (i.e. unit ensemble) are required to perform a given task. Furthermore, the task may change from one moment to the next, and therefore the same set of units can participate in the accommodation of the new task – a concept referred to as *neural reuse* (Anderson, 2014). In virtue of the unit ensemble, the network may accumulate evidence over time, such that the participant of multiple units gives rise to observable behaviour.

Localist models share properties that overlap with, but also differ from classic distributed models. In classic distributed systems, information is superimposed such that multiple items are represented by the activity of an individual unit (Page, 2000). As such, each unit is involved in the simultaneous representation of multiple items, thus participating in multiple tasks at once. Nevertheless, localist models endowed with extendedness share a crucial property with classic distributed systems. For example, localist models are also capable of representing multiple items. However, rather than involving a single unit ensemble comprised of individual units that participate in the representation of multiple items (i.e. overlapping participation), localist models carry separate unit ensembles that participate in the representation of separate items (i.e. non-overlapping participation). As such, the notion of extendedness in localist models implies non-overlapping unit ensembles separately yet simultaneously participating in the representation of distinct items. Taken together, the representation of multiple items, whether carried through local or distributed models is arguably the product of flexible neural computation.

Localist models that carry the notion of extendedness display some degree of resistance to perturbation, neuronal pruning and noise (Browne, 2019). Since the representation of a single item is extended over many units, the idiosyncratic response profile of individual units in the network may change without hindering the representation of the item. In neurorobotics, embodied networks endowed with such extendedness may become resistant to momentary

hardware failure or sensor perturbation. These networks may exhibit graceful degradation, where their function is progressively lost rather than being subject to sudden catastrophic failure (Kasabov, 2007). Overall, perturbations to part of the local unit ensemble may fail to disrupt the behavioural trajectory of the robot, ultimately conserving the execution of desirable behaviours in an ever-changing environment.

When the same set of units are capable of participating in a novel task, while forgetting the content of the previous task, it becomes possible for networks to hold the palimpsestic property (Amit & Fusi, 1994; Nadal, Toulouse, Changeux, & Dehaene, 1986; G. Parisi, 1986). Here, synaptic weights responsible for consolidating a particular task can change to meet the requirements of a new task. In doing so, the same subset of output units connected by these dynamic synapses can change their activity patterns to accommodate changes in the input stream. In this way, individual units can participate in storing various activity patterns, because the network may establish new synaptic configurations over time (Abraham & Robins, 2005). The presupposition behind the accommodation of new inputs streams is that the network structure is modifiable, rather than remaining frozen (Sporns, 2007). As such, synaptic plasticity can allow the network to switch tasks without the requirement of network size expansion, nor the requirement of a separate network.

Since localist models carry non-overlapping unit ensembles, distinct items can be described by segregated patterns of neural activity. This feature is particularly attractive for making efficient usage of memory resources. For example, in multi-item working memory, cue items may hold greater processing relevance than other items irrelevant for the time being (Constantinidis, Franowicz, & Goldman-Rakic, 2001). Assuming each item is loaded as a discrete “slot”, units responsible for representing the cue item may exhibit a distinct pattern of activity relative to their neighbouring subpopulation. In the context of motor control in neurorobotics, segregated patterns of neural activity can be mapped onto separate wheel motors, endowing artificial agents with distinct movement trajectories (Berberian, Ross, & Chartier, 2021). In turn, observations of movement trajectories can inform how patterns of

neural activity are changing in the network. Using this approach, physical robots as such can be used as a methodological tool for exploring neural activity “in action”. Notably, movement executions can provide an indication of the agent’s current sensory state – a “matching” takes place between the physical instantiation of observed movements and the neural dynamics underlying the embodied network (Berberian et al., 2021; Sporns, 2007). Viewed from such matching perspective, observed motor trajectories may turn out to be instrumental, because they can reflect, with some degree of precision, certain properties (e.g. intensity, duration) of sensory stimuli (Choe, Yang, & Eng, 2007).

Online learning

Online learning is a mode of operation where the information processing system updates its knowledge base at every time instance (Jain, Seera, Lim, & Balasubramaniam, 2014). For example, neural networks subject to online learning incrementally adjust the strength of their synaptic connections at every time instance (J. Wang, Belatreche, Maguire, & McGinnity, 2010; J. Wang, Belatreche, Maguire, & McGinnity, 2014). In this way, online learning is an ongoing process whereby the network can continuously adjust to a stream of input data. Here, inputs cannot be stored upfront because the environment is continuously evolving. Consequently, online learning environments impose stringent computational restrictions on neural networks, because the system is continuously attempting to handle incoming data samples on-the-fly. To this end, the network undergoes incremental updates – training-sample-by-training-sample (Zech et al., 2017), which means that its structure must by design, be adaptive, flexible and if need be, resistant to change.

In virtue of the design principles mentioned above, online learning can harness neural networks with adaptive capabilities that make them suitable candidates for hardware implementations. For example, embodied neural networks subject to online learning can display experience-dependent synaptic plasticity. Here, hardware systems (e.g. robots) can oper-

ate on-the-fly (Lobo, Del Ser, Bifet, & Kasabov, 2020). Under the framework where the neural network is embodied which in turn is embedded in the physical environment, the co-evolution between the agent and the environment can be studied as a dynamic process (Beer, 2014). Furthermore, the network structure that the agent embodies is not held fixed, but is rather sensitive to environmental influences that perturb its organization so as to transform the behaviour appropriate to the context. In this way, adaptive systems that learn online can align, at every time instance, with dynamically changing environments, and embodied systems provide an opportunity to explore these operations in action (Jain et al., 2014). Altogether, neural networks subject to online learning can be suitable for robotic implementations, which in turn can be embedded in the physical environment – acting as a flexible system subject to continuous adaptation (Berberian et al., 2021). Observations of the agent’s behaviour can then be related back to the details of the neural circuitry (Krakauer et al., 2017; J. Wang et al., 2014). This top-down approach first observes the behaviour of the agent, and then decomposes the degree to which the observed behaviour is associated with the underlying neural computation.

Similar to their biological counterparts, artificial agents subject to online methods may operate under continuously changing input streams. As a result, their internal state variables are updated on-the-fly such that they make no explicit distinction between what is commonly known as learning and recall in artificial neural networks (Zenke, Poole, & Ganguli, 2017). As such, the distinction between learning and recall is only meaningful to the external observer, and thus holds no explicit meaning for the agent behaving in an ever-changing environment (Brette, 2019). Indeed, experimental manipulations tailored for yielding interpretable results are commonly reflective of our own assumptions, that is, the assumptions of the external observer (Jolly & Chang, 2019). Nevertheless, drawing a fine line between learning and recall is deemed useful for the external observer. For example, in some situations, the external observer can devise an experimental protocol where the learning phase is treated as an independent variable, whereas the recall phase as the dependent variable. In this scenario,

the recall performance of the agent can be measured by systematic manipulation of stimulus features (e.g. duration, intensity) during the learning phase. During recall, the external observer can monitor whether the observed behaviour of the agent accurately reflects certain stimulus features over others (Choe et al., 2007). In this way, segregating learning from recall provides an opportunity to examine how the action-guided system is behaving in response to learning-related manipulations (Sharkey & Heemskerk, 2019). Overall, online methods render the agent agnostic to what we call learning and recall, a distinction that nevertheless serves as a utility for the external observer.

In offline methods, the synaptic connections in the network are normally updated after the entire training dataset is presented, marking the end of the first training cycle. After iterating through a few training cycles, learned connection weights are frozen and preserved, marking the onset of a new mode of operation, namely the testing (i.e. recall) phase (Jain et al., 2014). During testing, synaptic weights are immutable, resting on the assumption that the network remains forever “frozen” in its current state. Hence, if new and unforeseen data samples are presented to the network during testing, the network operates in a context where synaptic weights are frozen. Assuming the network has to retain all items into its pre-existing knowledge base, then freezing synaptic weights is a viable solution from an engineering standpoint – a solution offered by offline methods. In contrast to offline methods, online methods hold a built-in procedure for continual weight updating during recall (Jain et al., 2014). In other words, while offline methods have their weights frozen during recall, online methods continue changing their weights (J. Wang et al., 2010). Taken together, online and offline methods carry distinct build-in procedures during recall.

When an explicit minimal stopping rule for synaptic weight convergence (e.g. freeze weights if the specified minimum error value is reached) is no longer a criteria, online methods can learn and forget continuously on-the-fly (Berberian et al., 2021). However, attempting to learn new information without degrading previous ones means attempting to solve the problem of ongoing stable, lifelong learning (G. I. Parisi, Kemker, Part, Kanan, & Wermter,

2019). Lifelong learning is an incremental process where the system is concerned with the ability to continuously and incrementally learn in an ever-changing environment (Kasabov, 2007). Networks that operate in online mode pose as ideal paradigms for investigating the problem of lifelong learning, because these systems attempt to accommodate, on-the-fly, ongoing changes of the input distribution (Lobo et al., 2020). During this process, the network must retain previous contents while storing an ongoing stream of new contents. Lifelong learning can depend on multiple factors, one of which may be attributed to the plasticity of the network structure. Here, the network structure is not held fixed, but rather changes and adapts over a lifetime of complex tasks and domains – a notion commonly known as “learning with structural adaptation” (Wysoski, Benuskova, & Kasabov, 2006).

General overview and summary of chapters

The three chapters of this thesis share common themes, but also differ in their degree of methodological innovation. For instance, all chapters made use of formal models of spiking neurons. In chapter one and two, the leaky integrate-and-fire (LIF) model was used (Lapicque, 1907), because it was important for us to simplify the problem to as few parameters as possible, and gear towards computational efficiency over biophysical complexity. In chapter three, the adaptive-exponential integrate-and-fire (aEIF) model was used (Brette & Gerstner, 2005; Naud, Marcille, Clopath, & Gerstner, 2008). While the aEIF model carries a larger number of parameters and variables than the LIF, it can track neuronal response statistics more closely, and is well-known for capturing the spike initiation behaviour of single neurons (Naud et al., 2008). For these reasons, the methodological setting in chapter three was accompanied by the expressive power of this spiking neuron model. In addition to spiking, the mechanism of short-term synaptic plasticity was considered across all three chapters, an instrumental phenomena by which the efficacy of synaptic connections rapidly changes based on recent historical activity of nerve cells (R. S. Zucker & Regehr, 2002). As

we progress from earlier to later chapters of this thesis, the studies will gain methodological innovation as follows: we will be transitioning (1) from offline to online learning scenarios, (2) from disembodied to embodied formal models, (3) from microcircuits to larger networks of spiking neurons, and (4) from ecologically insensitive to ecologically sensitive robotic agents.

The first two chapters will consider a similar phenomena of interest, that is, the process of motion discrimination – the ability of single neurons and circuits to discriminate either the position (chapter one) or the direction (chapter two) of moving stimuli (e.g. leftward and rightward) along a single axis of motion. To study motion discrimination, chapter one focused on the computational aspect using a simplified, disembodied microcircuit motif of two units (Berberian, Ross, Chartier, & Thivierge, 2017). In chapter two, the problem was extended and studied under an embodied circuit of six units (Berberian, Ross, & Chartier, 2019), and the results from the implementation were compared against multi-electrode Utah array recordings from primary visual cortex (V1) of macaque monkeys (Kohn & Smith, 2016; Smith & Kohn, 2008). In chapter three, the formation of working memory was studied – the ability to maintain, for a certain duration, the content of a stimulus in memory without its concrete presence. Reminiscent to chapter one, chapter three devised a spiking neuron model subject to both short-term and long-term synaptic plasticity. This final study attempted to build an ecologically sensitive robot (Berberian et al., 2021). More specifically, the robotic implementation was performed online and the artificial agent was placed in an ecological context under which it interacted with a human subject. As a result, this ecological niche naturally involved variations in the interactive procedure, because the agent was stripped from the luxury of operating in idealized, highly controlled settings – exposing its adaptive capabilities. Overall, this thesis attempts to paint a picture that justifies the neurorobotic approach, and how its use may help gear towards validating any hypothesized understanding of neural computation. The justification of its use was initiated as a corollary in the first article, and progressed towards a principle proposition of the thesis.

Statement of contributions

The research presented in the following chapters was conducted at the University of Ottawa in Canada.

Chapter one was written for publication under the supervision of professor Sylvain Chartier and professor Jean-Philippe Thivierge, who reviewed and edited the conference proceeding. Nareg Berberian implemented the formal model and ran the numerical simulations. Nareg Berberian and Matt Ross conceptualized the study, analyzed the results and wrote the paper. Financial resources were provided by professor Sylvain Chartier, and in part by professor Jean-Philippe Thivierge, with funding acquisition from the Natural Sciences and Engineering Research Council of Canada (NSERC).

Chapter two was written for publication under the supervision of professor Sylvain Chartier, who reviewed and edited the manuscript. Nareg Berberian and Matt Ross conceptualized the study. Nareg Berberian ran the numerical simulations, analyzed the results, and wrote the manuscript. Matt Ross devised the robotic implementation. Financial resources were provided by professor Sylvain Chartier, with funding acquisition from NSERC.

Chapter three was written for publication under the supervision of professor Sylvain Chartier. Nareg Berberian and professor Sylvain Chartier conceptualized the study. Nareg Berberian developed the framework, conducted the experiments and wrote the manuscript. Matt Ross and professor Sylvain Chartier reviewed and edited the manuscript. Financial resources were provided by professor Sylvain Chartier, with funding acquisition from NSERC.

Chapter 1

Synergy Between Short-Term and Long-Term Plasticity in Vision

1.1 Abstract

In this study, we examine whether short-term synaptic plasticity (STP) interacts with learning rules governing long-term synaptic plasticity (LTSP). More specifically, we examine whether the initial vesicle release probability can influence long-term changes in synaptic strength. Given the importance of calcium-dependent modulation of synaptic transmission, as well as the temporal information to cortical computation, we study whether STP may contribute to network connectivity strength and stability. Taking as a starting point the well-established Tsodyks-Markram (TM) rule for STP, we devise a simple microcircuit of two interconnected units receiving a train of incoming spikes subject to presynaptic STP. Extending the TM model, we also implement a mechanism of postsynaptic LTSP. By treating the two mechanisms synergistically, we manipulate the initial vesicle release probability of short-term facilitating and short-term depressing synapses, and find that this process can modulate the strength and stability of synaptic connections subject to long-term depression, thus influencing postsynaptic responses. As such, we show how the extension of the TM

model jointly controls the activity in these units. Overall, results suggest that short-term synaptic plasticity can influence postsynaptic activity.

1.2 Introduction

In recent years, spike-timing-dependent plasticity (STDP) has emerged as the experimentally most studied form of synaptic plasticity (Markram, 2011). This associative form of synaptic plasticity determines the sign and magnitude of long-term potentiation (LTP) or depression (LTD) according to the order and precision of presynaptic and postsynaptic spike timing (Feldman, 2012). Multiple studies have aimed in revealing the rules and mechanisms underlying long-term changes in synaptic strength thought to underlie learning and memory (Brown, 1990; Dan & Poo, 2004; Malenka & Bear, 2004; Markram, 2011). At several cortical synapses, induction of long-term synaptic plasticity (LTSP) requires depolarization of the postsynaptic neuron in addition to presynaptic activity (Dan & Poo, 2004; Malenka, 1999), and long-term changes in synaptic strength is an important consequence of the relative timing between presynaptic and postsynaptic spikes (Bi & Poo, 1998; Tsodyks & Markram, 1997). This property is generally regarded as a hallmark of Hebbian learning (Kempster et al., 1999).

In addition to LTSP, studies have shown that an integral part of neural processing involves complex dynamical properties that operate on short timescales. This process is commonly known as short-term synaptic plasticity (STP), whereby the efficacy of the synapse varies dramatically over the course of hundreds of milliseconds to seconds timescale as a result of recent activity (Abbott & Regehr, 2004; R. S. Zucker & Regehr, 2002). These rapid changes in synaptic strength are defined by the efficacy of synaptic transmission, which varies overtime from one spike to the next due to short-term facilitation (STF) and depression (STD) (Thomson, 2000). Such rapid changes have been shown to provide beneficial effects to cortical computation, endowing neural circuits with the capacity to process temporal

patterns (Abbott & Regehr, 2004; Klug et al., 2012; Natschlger, Maass, & Zador, 2001; P. Carvalho, 2011; Rotman, Deng, & Klyachko, 2011; R. S. Zucker & Regehr, 2002).

Despite a large body of existing knowledge on STP and STDP, these ubiquitous mechanisms have frequently been studied in isolation. Yet, there is experimental evidence supporting the interaction between short-term and long-term plasticity (Jin et al., 2012; Kintscher et al., 2013; Markram & Tsodyks, 1996; Tokuoka & Goda, 2008). For example, it is known that alterations in synaptic strength, as well as neuromodulation, affect the reliability of the synapse by changing the probability of neurotransmitter release (Bolshakov & Siegelbaum, 1995; C. Stevens & Wang, 1994; Tritsch & Sabatini, 2012).

The existing synergy between short-term and long-term synaptic plasticity has been studied theoretically (Costa, Froemke, Sjöström, & van Rossum, 2015; Fernando & Yamada, 2012; Hiratani & Fukai, 2014; Legenstein, Naeger, & Maass, 2005; P. Carvalho, 2011; Senn, Markram, & Tsodyks, 2001; Vasilaki & Giugliano, 2014). For example, Legenstein and colleagues (2005) have shown how STDP can modulate the initial release probability (Pr) of synapses, pointing towards their potential functional contribution in the precise timing of firing in neural circuits (Legenstein et al., 2005). Vasilaki and Giugliano (2014) have highlighted how a parallel interaction between STDP and STP may contribute to the emergence of pairwise connectivity motifs observed in experiments (Vasilaki & Giugliano, 2014). Carvalho and Buonomano (2011) proposed a learning rule whereby STDP governs the baseline neurotransmitter release so as to discriminate complex spatiotemporal spike patterns. Hiratani and Fukai (2014) explored the interplay between the two mechanisms and their potential contribution in the formation, maintenance and reorganization of cell assemblies. Fernando and Yamada (2012) found that STP can help stabilize excitatory activity without imposing user-defined weight constraints in a microcircuit motif subject to STDP. Finally, Costa and colleagues (2015) devised a computational model where they showed how both pre- and postsynaptic locations of expression for synaptic plasticity may act as a substrate for rapid re-learning of previously experienced stimuli.

In support of experimental work, a common feature amongst all aforementioned theoretical studies is that long-term plasticity orchestrates short-term changes in synaptic strength. While we remain mindful of the experimental evidence suggesting unidirectional interaction from STDP to STP, here we rather asked how STP can influence long-term changes in synaptic strength. The main reason for assuming the potential presence of a bidirectional synergy is due to the fact that multi-timescale interactions of synaptic plasticity have not been firmly established in the literature. Consequently, the relation between short-term dynamics of synaptic transmission and long-term changes in synaptic strength remains open to theoretical and experimental investigations. Here, we focus on combining the two mechanisms, and specifically investigate the contribution of STP to the stability of synapses subject to STDP. In doing so, we construct two interconnected units, one of which is subject to STD and the other STF (Fig 1.1).

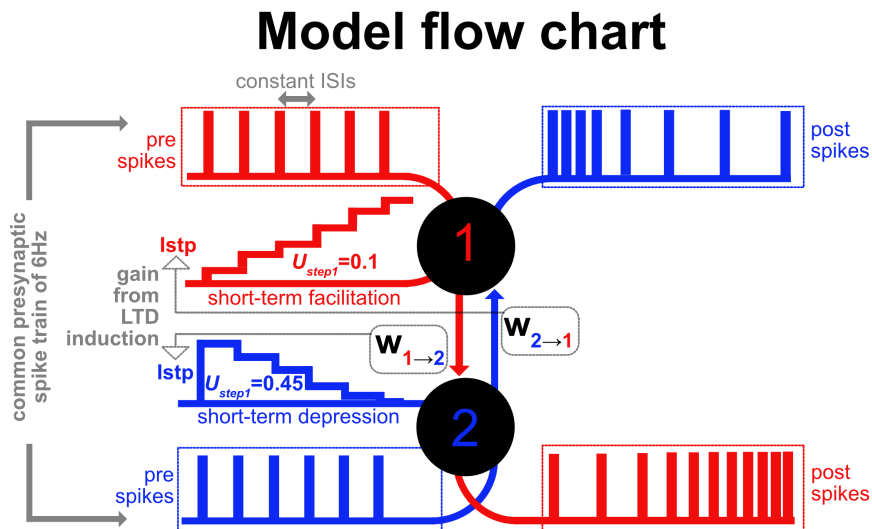


Figure 1.1: **Model flow chart.** The microcircuit is comprised of two units, interconnected by initial weight values $w_{1 \rightarrow 2} = 1$ and $w_{2 \rightarrow 1} = 1$. The two spike trains are subject to STP, yielding synaptic currents which in turn produce postsynaptic spikes. Based on the time difference between pre and post spikes (red vs red and blue vs blue), $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ get updated according to pair-based, all-to-all spike-timing interaction. Updated synaptic weights are then used as multiplicative scaling factors on synaptic currents (grey arrows). As a result of this multiplicative gain, updated synaptic currents produce a novel train of postsynaptic spikes. After weight stabilization, baseline Pr changes ($U_{step1} \rightarrow U_{step2}$), and weights are reinitialized to 1.

As these two interacting units receive a stationary input, the goal of the microcircuit is to discriminate, via systematic changes in firing rate, the spatial position of a stimulus moving bidirectionally across a horizontal axis of motion. To this end, we manipulate the baseline Pr of STF and STD synapses and study how they influence synaptic weights that yield progressive changes in postsynaptic activity. Overall, the microcircuit motif is chosen as a minimal requirement for uncovering the hypothesized influence of STP over the strength and stability of synapses subject to STDP.

The remainder of chapter one is divided as follows: The mechanisms of short-term (STP) and long-term (STDP) plasticity are first introduced. Then, we highlight the numerical procedure for studying the interaction between STP and STDP. This is followed by specifying the method used to search for candidate pairs of baseline release probabilities that render stable microcircuit activity. Results from the numerical simulations are then followed by discussion and conclusion of the overall findings.

1.3 Methods

1.3.1 Short-term plasticity (STP)

As a starting point, we introduce the well-known phenomenological Tsodyks-Markram (TM) model of STP (Tsodyks & Markram, 1997). Efferent connections from presynaptic units evolve according to:

$$\frac{du_j^s}{dt} = \frac{U_j^s - u_j^s(t)}{\tau_{f_j}} + U_j^s(1 - u_j^s(t)) \sum_k \delta(t - t_j^k) \quad (1.1)$$

where the sum on k spikes is over all spike times t_j^k of presynaptic unit j and $u_j^s(t)$ reflects presynaptic residual calcium levels at direction step s ($1, 2, 3, \dots, q$) with q corresponding to the number of direction steps. Here, direction steps are defined as the number of discrete spatial positions partaken from a stimulus moving bidirectionally across a horizontal axis of

motion. Noteworthy, s does not change at every time instance. Rather, s increments after the stabilization of synaptic weights $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$, that is, after 1000 learning trials (see Numerical procedure for details).

In the absence of incoming action potentials, the synapse is at a resting state with residual calcium levels $u_j^s(t_0) = U_j^s$. The amount of residual calcium instantaneously increases immediately after the first action potential within a spike train and $u_j^s(t_1) = (1 - u_j^s(t_0))$ is the fraction that remains available immediately after this first event. Hence, the running variable of U_j^s refers to as $u_j^s(t)$, and U_j^s is applied to the first action potential in the spike train, after which $u_j^s(t)$ decays exponentially to its resting value U_j^s with a time constant τ_{f_j} . As a result, each time an action potential is generated, presynaptic residual calcium instantaneously rises and then recovers with a time constant τ_{f_j} between subsequent spikes.

As residual calcium levels increase, a release-ready vesicle along the active zone of the presynaptic membrane terminal releases neurotransmitters into the synaptic cleft. During this process of exocytosis, neurotransmitter availability within the presynaptic terminal is described according to:

$$\frac{dx_j^s}{dt} = \frac{1 - x_j^s(t)}{\tau_{d_j}} - u_j^s(t)x_j^s(t) \sum_k \delta(t - t_j^k) \quad (1.2)$$

where $x_j^s(t)$ denotes the fraction of resources that remain available following vesicle release at direction step s . Between subsequent spikes, $x_j^s(t)$ recovers back to baseline to its resting value of 1 with a time constant τ_{d_j} , restoring the amount of synaptic resources available within the presynaptic terminal.

The STP model operates under a relatively short timescale, namely in the order of few milliseconds to seconds timescale (R. S. Zucker & Regehr, 2002). Depending on the initial setup of kinetic parameters τ_f , τ_d and U (Table 1.1), the STP model can mimic the effect of a depressing or a facilitating synapse (Tsodyks & Markram, 1997).

When a spike arrives at time t_j^k , the kinetic parameters will modulate the interplay

Table 1.1: Short-term synaptic plasticity parameters.

Parameter	STD	STF
Facilitation time constant (τ_f)	50 ms	750 ms
Depression time constant (τ_d)	750 ms	50 ms
Initial release probability (U)	0.45	0.15

between the dynamics of $u_j^s(t)$ and $x_j^s(t)$. In turn, the joint effect of $u_j^s(t)x_j^s(t)$ will characterize the short-term strength of the synapse at each time-instance t for a given direction step s , generating a synaptic current described by:

$$\frac{dI_j^{stp}}{dt} = (-I_j^{stp}(t) + u_j^s(t)x_j^s(t)) \sum_k \delta(t - t_j^k) \quad (1.3)$$

where δ is the Dirac delta function, and $I_j^{stp}(t)$ is the current triggered by the activation of presynaptic unit j . For simplicity, currents carry an instantaneous rise time, and the decay time constant is set to 1. Stationary presynaptic inputs of rate 6Hz form constant inter-spike intervals (ISIs) incorporated in equations 1.1-1.3.

Table 1.2: Single cell parameters.

Parameter	Values
Spike emission threshold (θ)	-55 mV
Reset potential (V_{reset})	-58 mV
Resting potential (V_{rest})	-70 mV
Membrane resistance (R_m)	200 m Ω
Membrane time constant (τ_m)	30 ms
Absolute refractory period (τ_{arp})	2 ms

1.3.2 Spike-timing-dependent plasticity (STDP)

In order to trace the total weight change Δw resulting from presynaptic and postsynaptic spike trains, pair-based STDP was applied using all-to-all spike-timing interaction:

$$\Delta w = \sum_k \sum_n W(t_i^n - t_j^k) \quad (1.4)$$

where $W(x)$ denotes the learning window of the STDP function (J. Sjöström & Gerstner, 2010). The double summation is over all differences between presynaptic and postsynaptic spike times. The change Δw in synaptic weights depends on the relative timing between all pairs of presynaptic and postsynaptic spikes. The presynaptic spike arrival times at synapse j are represented by t_j^k where k counts the presynaptic spikes ($k = 6$). Similarly, t_i^n represents the postsynaptic spike arrival times where n counts the postsynaptic spikes (n varies). At several cortical synapses, the width of the temporal window for LTD is considerably wider than the width of the temporal window for LTP (Feldman, 2000; P. J. Sjöström et al., 2001) (Table 1.3).

Table 1.3: Spike-timing-dependent plasticity parameters.

Parameter	Values
LTD window (τ_-)	50 ms
LTP window (τ_+)	20 ms
LTD rate (λ_-)	5e-4
LTP rate (λ_+)	5e-4
Learning factor (μ)	1
Scaling factor (α)	2

Consequently, a temporally asymmetric STDP learning rule was used to compute the magnitude and the direction of weight changes as function of presynaptic and postsynaptic spike timing differences (Fig 1.2):

$$\Delta w = \begin{cases} \lambda_+ f_+(w) e^{-\frac{\Delta t}{\tau_+}} & \text{if } \Delta t \geq 0 \\ -\lambda_- f_-(w) e^{\frac{\Delta t}{\tau_-}} & \text{if } \Delta t < 0 \end{cases} \quad (1.5)$$

$$f_+(w) = (1 - w)^\mu \quad \text{and} \quad f_-(w) = \alpha w^\mu \quad (1.6)$$

where $\Delta t = t_i^n - t_j^k$ represents an n-by-k matrix of spike-timing differences between presynaptic unit j and postsynaptic unit i . Parameter $\lambda_\pm, 0 < \lambda_\pm \ll 1$ is the learning rate that scales the magnitude of individual weight changes. The updating functions $f_+(w)$ and $f_-(w)$

scale synaptic changes (Gütig, Aharonov, Rotter, & Sompolinsky, 2003). The parameter μ of the updating functions sets boundary conditions on the changes in w . In this way, a non-zero μ prevents runaway dynamics of synaptic efficacies, thereby preventing weights from leaving the allowed range $[0, 1]$ (Table 1.3). The parameter α denotes a possible asymmetry between the scales of LTD and LTP, where $\alpha < 1$ assigns a higher magnitude for LTP whereas $\alpha > 1$ favors LTD (Gütig et al., 2003) (Table 1.3). In our case, an asymmetric learning rule for STDP is used to update the weights exhibiting LTSP, with a positive temporal difference between presynaptic and postsynaptic spikes leading to LTP and otherwise LTD for a negative temporal difference (Fig 1.2).

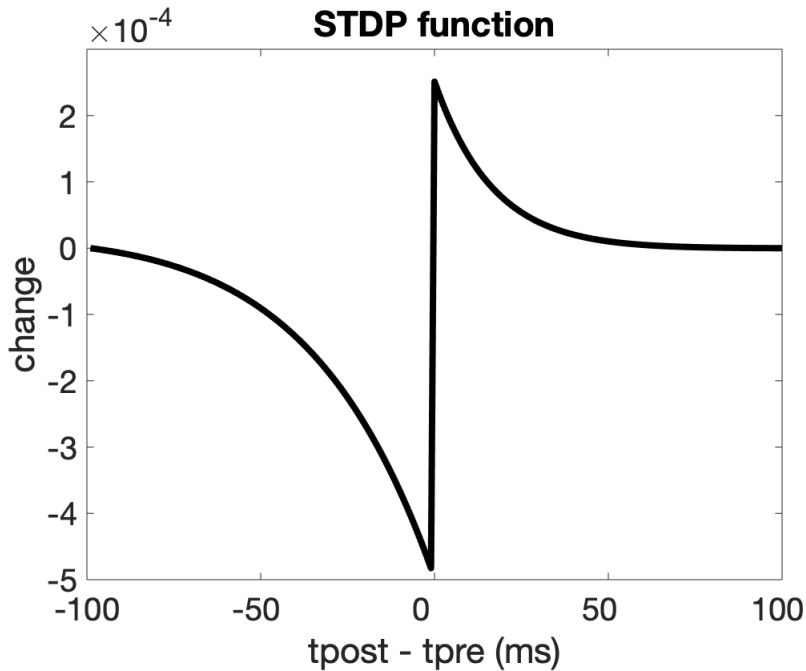


Figure 1.2: **Asymmetric learning rule for spike-timing-dependent plasticity.** A negative temporal difference between presynaptic and postsynaptic spike timing results in LTD, whereas a positive temporal difference results in LTP.

The synergy between STP and LTSP is introduced according to:

$$I_i^{stdp}(t) = \sum_{j=1}^2 w_{ij} I_j^{stp}(t) \quad (1.7)$$

where w_{ij} is the connection from presynaptic unit j to postsynaptic unit i and $I_j^{stp}(t)$

is the presynaptic current generated from the mechanism of STP described above. The updated current $I_i^{stdp}(t)$ is injected in:

$$\tau_m \frac{dV_i}{dt} = -V_i(t) + V_i^{rest} + R_m I_i^{stdp}(t) \quad (1.8)$$

where τ_m is the membrane time constant, R_m is the cell's membrane resistance and $I_i^{stdp}(t)$ is the updated version of $I_j^{stp}(t)$ scaled by the multiplicative gain factor w_{ij} . Whenever a depolarization hits a fixed threshold $V_i(t) > \theta$, the unit emits a spike and becomes refractory for a period τ_{arp} , after which equation 1.8 resumes from a sub-threshold reset potential V_i^{reset} (Table 1.2). The neuronal dynamics of leaky integrate-and-fire units are completely described by the system of ordinary differential equations (1.1-1.3 and 1.8) combined with the conditions for spike emission and refractoriness described above. These equations are integrated using an Euler scheme. Temporal parameters used in the numerical simulations are reported in Table 1.4.

Table 1.4: Temporal parameters.

Parameter	Values
integration time step (dt)	1 ms
single trial duration (T)	1000 ms


1.3.3 Numerical procedure

The numerical procedure of the model is shown in Fig 1.3. The steps are highlighted as follows: First, both units receive a stationary spike train of 1000ms subject to STP (equations 1.1-1.2), yielding I_j^{stp} (equation 1.3). Second, I_j^{stp} is scaled by w_{ij} , yielding I_i^{stdp} (equation 1.7). Third, I_i^{stdp} generates postsynaptic spikes (equation 1.8). Fourth, STDP (equations 1.4-1.6) adjusts the weights based on all-to-all spike-timing differences, inducing Δw_{ij} . Fifth, synaptic connections are updated according to $w_{ij} = w_{ij} + \Delta w_{ij}$, and steps 1 – 4 are repeated until $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ stabilize. Sixth, a new pair of baseline Pr are recruited


$(U_{STF}^s, U_{STD}^s) \rightarrow (U_{STF}^{s+1}, U_{STD}^{s+1})$, $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ are reinitialized to 1. Finally, steps 1 to 6 are repeated until s reaches the total number of direction steps q .

Numerical procedure


Step 1: all-to-all spike timing interaction (red vs red) and (blue vs blue)




Step 2: update $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$



Step 3: update currents



Step 4: update post spikes



Step 6: recruit a new pair of baseline release probabilities U , reinitialize weights, repeat steps 1 to 5

Step 5: repeat steps 1 to 4 until $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ stabilize

Figure 1.3: **Numerical procedure of the model.** Step 1: following pre and post spike trains of 1000ms, the all-to-all spike-timing interaction is computed (red vs red and blue vs blue). Step 2: $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ are updated based on the LTD induction protocol (see Spike-timing-dependent plasticity (STDP) for details). Step 3: updated $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ are used as multiplicative scaling factors. Step 4: post spikes are updated. Step 5: steps 1 to 4 are repeated until $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ stabilize. Step 6: a new pair of baseline Pr are recruited; $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ are reinitialized to 1; steps 1 to 5 are repeated.

The criterion for weight stabilization is held fixed with no stopping rule, as Δw_{ij} was lower than $10e-4$ after 1000 learning trials. Furthermore, learning can be made quicker or slower depending on parameters λ_+ and λ_- introduced in equation 1.5, which scale the magnitude

of individual weight changes. However, the parameter setup presented here allowed for a proper speed and accuracy tradeoff for the convergence of synaptic connections. To visualize the interaction between the two mechanisms, an overview of the steps were recorded to describe how STP affects the final STDP weights (<http://aifuture.com/res/2017-ssci/>).

1.3.4 Procedure for searching candidate pairs of baseline Pr

Here we search for candidate pairs of baseline Pr which allow units in the microcircuit to properly discriminate the spatial position of a moving stimulus. In doing so, we construct two distinct receptive fields demarcated by left and right visual fields. Within each receptive field, we intend to capture neuronal activity that progressively increments or decrements in response to a stimulus that moves bidirectionally along a single axis of motion (Hubel & Wiesel, 1959). The stimulus is conceptualized as a vertically oriented bar (e.g. line). To model systematic changes in firing rate as a function of stimulus location, we introduce the notion of *direction steps*. Here direction steps refer to a stimulus that moves bidirectionally in discrete steps along a horizontal axis of motion. Noteworthy, systematic changes in the position of the bar do not change the rate of presynaptic activity, because presynaptic spikes remain stationary during the entire duration of the numerical procedure. Therefore, there is no explicit conversion from stimulus to presynaptic spike trains (see chapter two where we attempt to overcome this limitation).

If the stimulus is purely conceptual, and incoming spike trains remain stationary, then which parameter in the model contributes to changes in postsynaptic firing rate? This is where the baseline Pr (U_j^s) comes in (see equation 1.1), which is the parameter of the TM model subject to systematic manipulation (P. Carvalho, 2011). Hence, each direction step s is directly accompanied by systematic changes in baseline Pr , which in turn affects post spikes. Here we searched for candidate pairs of baseline Pr (i.e. a set of baseline Pr pairs) that render stable postsynaptic rates. To this end, during stimulus position discrimination in the right visual field, the baseline Pr of the STF synapse was held constant at 0.15,

whereas that of the STD synapse ranged from 0.05 to 0.85 with $q_{rvf} = 400$ direction steps in between. During stimulus position discrimination in the left visual field, the baseline Pr of the STD synapse was held constant at 0.45, whereas that of the STF synapse ranged from 0.05 to 0.85 with $q_{lvf} = 400$ direction steps in between. In this way, we were able to explore the entire range of direction steps under which pairs of baseline Pr contributed to the stability, instability and suppression of postsynaptic firing rates. In the following results section, we first show the set of baseline Pr pairs that rendered stable postsynaptic firing rates. Then, we show microcircuit activity patterns (stable, unstable, suppressed) resulting from the entire range baseline Pr pairs mentioned above.

1.4 Results

Fig 1.4 presents values of LTD-mediated synaptic weights updated as a function of the learning trial number. For each direction step, synaptic weights are reinitialized to 1, and updated according to the learning procedure mentioned above (see Numerical procedure for details). There were 800 direction steps in total ($q = 800$), therefore 800 pairs of stabilized STDP weights. For each direction step, STDP weights stabilized after 1000 learning trials, and each learning trial was evaluated on the basis of spike-timing differences between 1000 milliseconds of presynaptic and postsynaptic activity. The stabilization of STDP weights were also examined under random initialization values sampled from a normal distribution, as well as under the usage of a temporally symmetric STDP learning window (<http://aifuture.com/res/2017-ssci/>).

We then examined the distribution of stabilized STDP weights. Amongst the 800 pairs of baseline Pr , half of the pairs corresponded to connection weights stabilized for a moving stimulus processed in the left visual field, and the remaining half were stabilized according to a moving stimulus processed in the right visual field. Fig 1.5 shows the distribution of stabilized STDP weights. When the stimulus is processed in the left visual field, $w_{1 \rightarrow 2}$




Figure 1.4: **Evolution of STDP weights as a function of learning trial number.** Connection weights stabilize after 1000 learning trials. Curves in blue and red correspond to weight instances $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ that scale I_j^{stp} subject to STD and STF, respectively.

stabilizes at a higher value than $w_{2 \rightarrow 1}$. Conversely, when the stimulus is processed in the right visual field, $w_{2 \rightarrow 1}$ stabilizes at a higher value than $w_{1 \rightarrow 2}$. Fig 1.6 presents the stabilization of LTD-mediated synaptic weights during each direction step. These direction steps are based on pairs of baseline Pr that yield proper stimulus position discrimination. Postsynaptic responses are shown in Fig 1.7 (see <https://www.youtube.com/watch?v=RQdw35P2Fqk> for details).

Fig 1.8 presents the baseline Pr pairs (U_{STF}^s ; U_{STD}^s) across all direction steps ($q = 800$). In general, STD-dominated synapses have a higher baseline Pr under which larger numbers of quanta are released per action potential. On the other hand, STF-dominated synapses are associated with lower levels of baseline Pr , thus releasing smaller numbers of quanta per action potential (Varela et al., 1997; R. S. Zucker & Regehr, 2002; R. Zucker, 1989). Consequently, across the entire horizontal axis of motion, the baseline Pr of the STD synapse was maintained higher than the baseline Pr of the STF synapse.

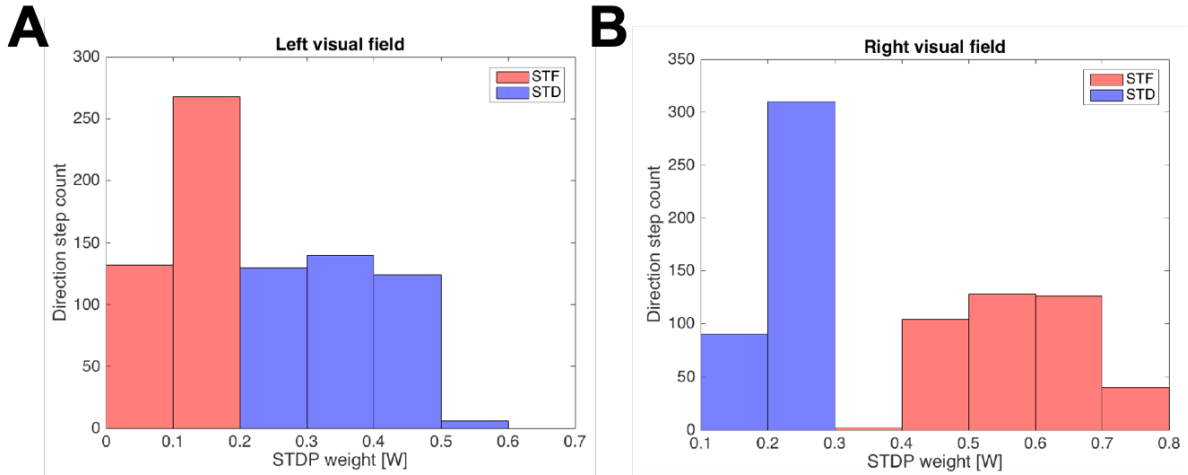


Figure 1.5: **Distribution of w_{ij} at learning trial 1000.** Stabilized distribution of w_{ij} during left (A) and right (B) visual field processing. Distributions in blue and red correspond to weights $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ at learning trial 1000 that scale I_j^{stp} subject to STD and STF, respectively.




Figure 1.6: **Stabilized weight w_{ij} as a function of direction step.** Steps $s \in [1..400]$ correspond to left-to-right [1..200] and right-to-left [201..400] stimulus position discrimination in left visual field. Steps $s \in [401..800]$ correspond to right-to-left [401..600] and left-to-right [601..800] stimulus position discrimination in right visual field. Scatter in blue and red correspond to weights $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ (Fig 1.1) at learning trial 1000 that scale I_j^{stp} (equation 1.7) subject to STD and STF, respectively.




Figure 1.7: **Firing rate of each unit as a function of direction step.** During left and right visual field processing, one unit remains active, while the other remains silent. Lines in blue and red correspond to postsynaptic responses at learning trial 1000, generated from synaptic currents subject to STD and STF, respectively.

What happens to the pattern of postsynaptic responses, if the baseline Pr of the STF synapse is higher than that of the STD synapse? In other words, are there candidate pairs of baseline Pr that render stable postsynaptic rates, while others that yield unstable or suppressing activity? To address this question, the initial Pr of the STD synapse was held constant at 0.45, while the baseline Pr of the STF synapse ranged from 0.05 to 0.85 with $q_{l_{vf}} = 400$ direction steps in between. Fig 1.9 shows stabilized STDP weights during stimulus position discrimination in the left visual field. Interestingly, the convergence of LTD-mediated synaptic weights that gave rise to a proper discrimination of stimulus position was possible only within a physiological range of baseline Pr where the STD synapse had a higher baseline Pr (0.45) than the STF synapse (0.1 – 0.15). Otherwise, pairs of baseline Pr outside this physiological range produced either 1) LTD (STDP weights < 0.25) strong enough to suppress the activity of the unit (Fig 1.10 – blue line in steps $s \in [1..400]$) that would otherwise properly discriminate the position of the moving stimulus during left visual




Figure 1.8: **Baseline Pr pairs as a function of direction step.** For each direction step, a new pair of baseline Pr are recruited. There are 800 direction steps in total thus 800 baseline Pr pairs. For steps $s \in [1..400]$, the baseline Pr of the synapse subject to STD is held constant, whereas the baseline Pr of the synapse subject to STF is modulated. The reverse process is held for step $s \in [401..800]$.

field processing (Fig 1.7 – blue line in steps $s \in [1..400]$), or 2) unstable fluctuations in STDP weights (Fig 1.10 – red line in steps $s \in [1..400]$).

The search for candidate pairs of baseline Pr was also conducted during stimulus position discrimination in the right visual field. Here, the initial Pr of the STF synapse was held constant at 0.15, while the Pr of the STD synapse ranged from 0.05 to 0.85 with $q_{rvf} = 400$ direction steps in between. Interestingly, stable postsynaptic responses from the right visual field also required a physiologically plausible range of baseline Pr pairs. Fig 1.11 shows stabilized STDP weights during stimulus position discrimination in the right visual field. In this case, a proper discrimination of stimulus position was possible within a physiological range of Pr where the STD synapse had a higher baseline Pr (0.2 – 0.35) than the STF synapse (0.15). Pairs of baseline Pr outside this range produced LTD strong enough (STDP < 0.25) to suppress the activity of the unit (Fig 1.10 – red line in steps $s \in [401..800]$) that would otherwise properly discriminate the position of the moving stimulus during right

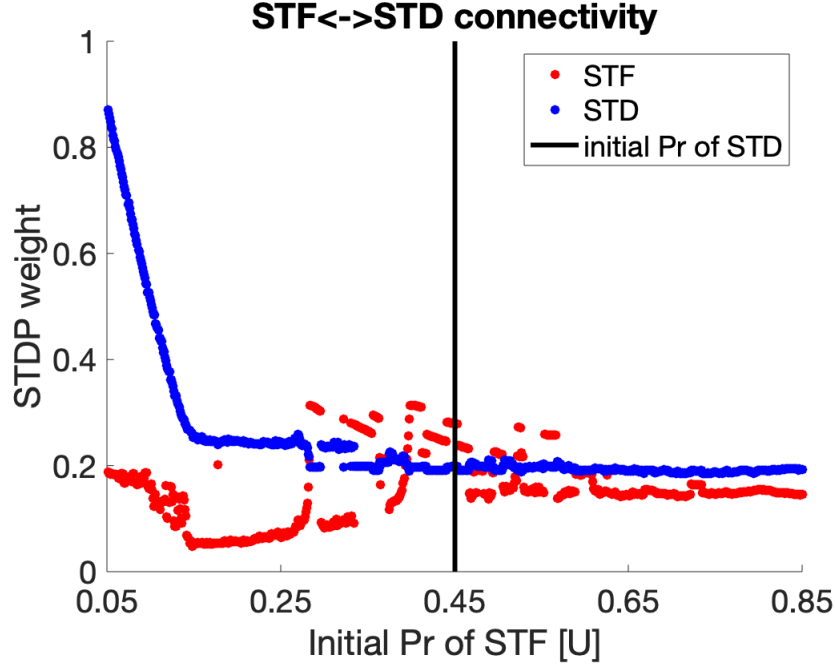


Figure 1.9: **Stabilization of w_{ij} in the left visual field during stimulus position discrimination.** Initial Pr of the synapse subject to STD is held constant, while the Pr of the synapse subject to STF is modulated. Scatter in blue and red correspond to weights $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ at learning trial 1000 that scale I_j^{stp} subject to STD and STF, respectively.

visual field processing (Fig 1.7 – red line in steps $s \in [401..800]$).

Overall, our results suggest that synapses may be subject to long-term changes via short-term dynamics of synaptic transmission, which in turn may control their strength and stability over the long-term thus influencing subsequent patterns of microcircuit activity.

1.5 Discussion

We found a subset of baseline Pr pairs in STD and STF synapses, which contributed to the activity of postsynaptic spikes. Interestingly, the regime of proper stimulus position discrimination occurred within a range where the probability of transmitter release was higher in the STD synapse in comparison to the STF synapse.

In this study, the interaction between STP and LTSP suggests that short-term changes in synaptic transmission can influence the long-term strength and stability of microcircuit




Figure 1.10: **Firing rate of each unit as a function of direction step.** There exists a regime of initial Pr pairs which disrupt the stability and strength of w_{ij} , resulting in unstable neuronal dynamics and suppression of neuronal responses. Lines in blue and red correspond to postsynaptic responses at learning trial 1000, generated by synapses subject to STD and STF, which in turn are multiplicatively scaled by weights $w_{1\rightarrow 2}$ and $w_{2\rightarrow 1}$, respectively. This figure presents the activity of both units according to the dynamic plot shown in (<https://www.youtube.com/watch?v=wlkts-Je6wg>).

connectivity. To accurately discriminate the position of the moving stimulus, we show that the baseline Pr of STF and STD synapses may need to operate, in conjunction, within a physiological range. More specifically, the baseline Pr in the STD synapse may need to be higher than the baseline Pr in the STF synapse in order for the microcircuit to properly discriminate the spatial position of the moving stimulus. Otherwise, assigning higher baseline Pr to STF and lower baseline Pr to STD either destabilizes or suppresses microcircuit activity. Overall, our results suggest that stimulus position discrimination is hindered when the baseline Pr of STD and STF synapses are pushed outside the physiological regime they tend to operate under biological settings. Nevertheless, the findings presented here remain purely theoretical, and therefore experimental work is required to confirm or refute them.

Considered as an integral part of history-dependent information processing in the brain,




Figure 1.11: **Stabilization of w_{ij} in the right visual field during stimulus position discrimination.** Initial Pr of the synapse subject to STF is held constant, while the Pr of the synapse subject to STD is modulated. Scatter in blue and red correspond to weights $w_{1 \rightarrow 2}$ and $w_{2 \rightarrow 1}$ at learning trial 1000 that scale I_j^{stp} subject to STD and STF, respectively.

STF and STD are well-established properties of synaptic transmission. It is therefore likely to expect STF and STD to play an important role in more complex microcircuits. Indeed, a recent study from Naud and colleagues (2020) show that hierarchical networks can achieve information multiplexing when feedforward and feedback connections are short-term depressing and short-term facilitating, respectively (Payeur et al., 2020). More specifically, they show that synaptic contacts near perisomatic regions provide bottom-up projections that are STD-dominated, which in turn act as selective communicators of event rates. Importantly, these bottom-up signals are simultaneously accompanied by top-down projections sent to apical dendrites, which are predominantly innervated with STF synapses that selectively communicate bursts without affecting bottom-up information flow. Hence, it is likely to expect that the polarization of STP in our microcircuit endows synapses in larger networks with important computational advantages, such as the ability to separate streams of information along the cortical hierarchy (Naud & Sprekeler, 2018).

The overall aim of this paper was to use the STDP learning rule inspired from biological neurons and examine its relation with short-term dynamics of synaptic transmission. Noteworthy, the model proposed here remains purely phenomenological and units exhibit equal firing rates in response to the stimulus moving in opposite directions. Therefore, the microcircuit discriminated the position of the moving stimulus, but not its direction. Indeed, two units are not enough degrees of freedom in order for the microcircuit to discriminate the direction of the stimulus.

1.6 Conclusion

We studied the synergistic interaction between short-term and long-term synaptic plasticity in a microcircuit motif of two units. In doing so, we manipulated the baseline Pr of STD and STF synapses and examined how they conjointly influence the strength and stability of synaptic connections via an asymmetric STDP learning rule. In chapter two, we extend the size of the microcircuit to six units, and embed the computational model in a Raspberry Pi microcontroller. There, we show how the agent can discriminate the direction of a moving stimulus. Finally, we compare the results from the robotic implementation against experimental data.

Chapter 2

Discrimination of Motion Direction in a Robot Using a Phenomenological Model of Synaptic Plasticity

2.1 Abstract

Recognizing and tracking the direction of moving stimuli is crucial to the control of a vast repertoire of animal behaviours. In this study, we examine whether a phenomenological model of synaptic plasticity implemented in a robotic agent may allow the discrimination of motion direction of real-world stimuli. Starting with a well-established model of short-term synaptic plasticity (STP), we develop a microcircuit motif of spiking neurons capable of exhibiting preferential and non-preferential responses to changes in the direction of an orientation stimulus in motion. While the robotic agent processes sensory inputs, the STP mechanism introduces direction-dependent changes in the synaptic connections of the microcircuit, resulting in a population of units that exhibit a typical cortical response property observed in primary visual cortex (V1), namely, direction selectivity. Visually evoked responses from the model are then compared to those observed in multi-electrode recordings

from V1 in anesthetized macaque monkeys, where sinusoidal gratings were displayed on a screen. Overall, the model highlights the potential role of STP as a complementary mechanism in explaining direction selectivity. It applies these insights in a physical robot, and validates the response characteristics of the embodied architecture using experimental data.

2.2 Introduction

Although a seemingly effortless task for humans, recognizing and tracking the direction of visual objects is based on an incredible complexity of brain areas involved in visual processing and attention, as well as learning and memory. In recent years, the advent of artificial neural networks (ANNs) has allowed the combination and isolation of the interactions of important biophysical mechanisms in order to shed light on the nature of biological phenomena. Through a symbiotic collaboration between neuroscience and artificial intelligence, the application of ANNs is in part to unify our understanding of the underlying mechanisms contributing to sensory experience. The development of these brain-inspired computational systems have shown their usefulness in revealing novel mechanisms of neuronal circuitry and in proposing experimental predictions that can be directly tested in experimental settings. In order to elucidate the circuit mechanisms underlying visual perception, mathematical models have been formulated with strong support from electrophysiological data (Tsodyks & Markram, 1997). Due to their usefulness and their predictive ability in driving new neuroscientific discoveries, brain-inspired ANNs also have the potential to be implemented in robotic agents in order to further assess their ecological validity (Krichmar, Conradt, & Asada, 2015). Given that mechanistic models cannot yet capture the full complexity of the nature of perceptual phenomena, the implementation of well-established models from neuroscience into the domain of artificial intelligence opens new avenues for understanding biological networks exposed to real-world stimuli (Hassabis, Kumaran, Summerfield, & Botvinick, 2017). Previous approaches in modelling the perceptual phenomena of motion

have shown successful attempts in incorporating natural visual inputs in networks of spiking neurons (Adams & Harris, 2015; Escobar, Masson, Vieville, & Kornprobst, 2009; Wu, McGinnity, Maguire, & Cai, 2008).

In this study, we propose a model of motion discrimination using a ubiquitous mechanism in neuronal circuits, namely, short-term plasticity (STP), whereby the strength of synaptic connections varies from milliseconds to seconds as a result of recent activity (Abbott & Regehr, 2004; R. S. Zucker & Regehr, 2002). These rapid changes in synaptic strength vary overtime from one spike to the next due to short-term facilitation (STF) and short-term depression (STD) (Dittman, Kreitzer, & Regehr, 2000). Short-term synaptic plasticity serves diverse functions in bio-inspired networks. For example, STP can process temporal patterns (P. Carvalho, 2011), modify a neuron’s sensitivity to the temporal coherence of inputs (Abbott et al., 1997; Markram & Tsodyks, 1996), participate in gain control (Abbott et al., 1997), reduce redundancy (Goldman et al., 2002), act as an adaptive filter (van Rossum et al., 2008), and separate streams of information (Naud & Sprekeler, 2018), among others (Tauffer & Kumar, 2019).

Despite the beneficial effects of STP on cortical computation (Abbott & Regehr, 2004; Deng & Klyachko, 2011; Klug et al., 2012; P. Carvalho, 2011; Rotman et al., 2011; R. S. Zucker & Regehr, 2002), it remains unclear whether STP contributes independently of sensory experience or whether it provides a causal contribution to experience-dependent plasticity. A study in-line with the former found that alteration in STP has been observed in cultured neurons, suggesting that endogenous neuronal activity (i.e., independent of sensory experience) is sufficient to drive changes in STP (Chen & Buonomano, 2012). In contrast, there is evidence to suggest that STP is a consequence of experience-dependent plasticity in local neuronal circuits and therefore causally linked to visually driven inputs (Bender, 2006; Cheetham & Fox, 2011; Finnerty, Roberts, & Connors, 1999; R. S. Larsen et al., 2014; Takesian, Kotak, & Sanes, 2010; Urban-Ciecko, Wen, Parekh, & Barth, 2015). For example, sensory deprivation can alter STP, but in most cases, the dynamics of synaptic transmission

are often inconsistent in these experiments, as even at the same synapse type, some promote facilitation while others will exhibit depression. Nonetheless, evoked and spontaneous vesicle release is likely to be controlled by two independent and non-overlapping mechanisms (Abrahamsson et al., 2017). Sensory experience can therefore modify the dynamics of STP, thus pointing towards a causal contribution of STP to experience-dependent plasticity. Indeed, an important determinant of development and sensory-driven alteration in STP is the expression of presynaptic NMDA receptors (preNMDARs) (Banerjee, Larsen, Philpot, & Paulsen, 2016; R. Larsen & Sjöström, 2015). These are ligand-gated ionotropic glutamate receptors that serve diverse functions ranging from the coincidence detection in Hebbian learning to excitatory neurotransmission critical for information processing in the mammalian central nervous system (Paoletti, Bellone, & Zhou, 2013).

In layers 2/3 (L2/3) of the primary visual cortex (V1), individual neurons respond more strongly to an object (i.e., orientation grating) moving in a particular direction (“preferred”) than the same object moving in the opposite direction (“null”); a visual response property termed direction selectivity. There is surmounting experimental and theoretical evidence that STP contributes to the enhancement of motion discrimination (Buchs & Senn, 2002; Carver, Roth, Cowan, & Fortune, 2008; Chance, Nelson, & Abbott, 1998; Hansel & Mato, 2013). In-line with previous studies, we recently found that rapid changes in synaptic strength via STP may provide an essential contribution for accurate motion discrimination (Berberian, Ross, Chartier, & Thivierge, 2017). Starting with the well-established Tsodyks-Markram model (Tsodyks & Markram, 1997), we implement STP in the synaptic connections of a microcircuit motif. We then examine neuronal responses to changes in the direction of real-world vertical orientation stimuli moving in bidirectional motion along a single axis of motion. Furthermore, we compare neuronal responses in realtime from a robotic implementation to those of a simulated version of the model whereby units are instead exposed to a hypothetical version of real-world stimuli in motion. Finally, we analyze neuronal responses in V1 to drifting sinusoidal gratings and compare cortical responses to those observed in the robotic

implementation. The remainder of the paper is divided as follows: Section (Methods) describes the architecture of the microcircuit motif, the setup of the robotic implementation, and the phenomenological model of STP and summarizes the experimental data analysis approach. Section (Results) illustrates all findings. In Sections (Discussion) and (Conclusion), we summarize the overall insights of our work, propose future avenues, and highlight the contribution of our work to neurorobotics research.

2.3 Methods

2.3.1 Network architecture

Here, we propose a microcircuit motif of six units in total, comprised of two subpopulations connected via synapses that exhibit STF (Fig 2.1). This novel framework differs from our previously proposed architecture of two units mediated by STD- and STF-dominated synaptic connections (Berberian, Ross, Chartier, & Thivierge, 2017). In our current study, we aim to provide a more parsimonious approach in highlighting the contribution of STP by using a single STP mechanism rather than two distinct STP mechanisms for showing successful motion discrimination. In addition, we highlight the structural advantages of the expanded network over the two-unit microcircuit. Finally, we display the functional advantages resulting from the topological structure of the expanded network, which happen to be absent in a two-unit microcircuit.

In order to assess whether the embodied robot is capable of displaying response characteristics similar to those observed in local microcircuits in V1, the architecture is expanded by following a constrained network topology inspired from specific features observed in local cortical microcircuits. For example, bidirectional connectivity in V1 is a by-product of neighbouring neurons sharing similar visual responses (Ko et al., 2013). In addition, bidirectional connectivity has been suggested to evolve according to synaptic connections mediated by STF (Vasilaki & Giugliano, 2014). Furthermore, neurons in V1 that share similar visual

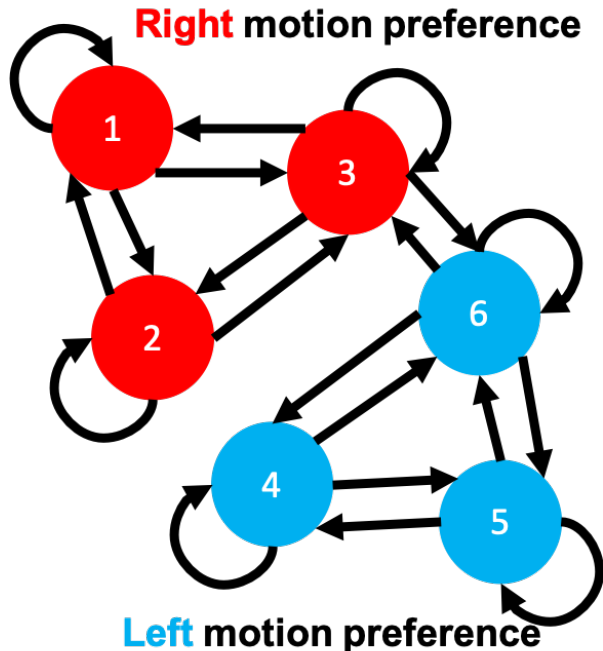


Figure 2.1: **Architecture of the microcircuit.** There are 6 units in total. Subpopulation 1 (red) forms a cluster of 3 units, each of which exhibits a preference for stimuli moving towards the right. The remaining units in subpopulation 2 (blue) display higher responses for stimuli moving towards the left. Connections are bidirectional, with self-connections allowed. Outgoing connections from units within both subpopulations exhibit STF. Connection strength from unit j to i $w_{ij} = 1$.

features (e.g., similar direction preference) are more highly connected and less connected to neurons showing a reduced preference for those similar visual features (Ko et al., 2013). Similarly, in our expanded microcircuit, there are a greater number of connections amongst units exhibiting the same direction preference and less connections between units coding for an opposing direction of motion (Fig 2.1) (Cossell et al., 2015). More specifically, units 1, 2, and 3 within subpopulation 1 are more highly connected amongst each other and less connected to units 4, 5, and 6 within subpopulation 2. These topological features of the expanded network would be absent in a two-unit microcircuit with bidirectional connections because both units would exhibit the same number of outgoing and incoming connections, acting as a single isolated subpopulation. Consequently, we hypothesized that, in the expanded architecture, units within subpopulation 1 and subpopulation 2 would exhibit preferential

responses to opposing directions of motion. In contrast, we expected that a microcircuit of two units with bidirectional connections would have a limited functional contribution by displaying preference only for a single direction of motion. Finally, it is noteworthy to mention that the study proposed here has extended the architecture to six units. We anticipate that network size expansion from this point forward would not introduce nonlinearities that would otherwise be absent in the six unit model. Here we attribute three main assumptions to this. Namely because (1) we assume the conservation of incoming and outgoing connections to and from each unit respectively, (2) the TM model is deterministic, and (3) the sequence of presynaptic spike trains to each unit, including presynaptic spike count, is maintained. Consequently, we anticipate that network size expansion would not qualitatively change the desired behaviour of individual units, but simply increase the simulation time.

2.3.2 Robotic setup

For the robotic implementation, we have employed the Raspberry Pi 3 Model B-V1.2 microcontroller (Fig 2.2A). To capture the image of the stimulus in motion, a Raspberry Pi camera (V2.1) is used and attached to the device via a ribbon cable. In addition, two simple circuits are created on a breadboard responsible for lighting up coloured LEDs (red and blue) and attached to the Raspberry Pi's GPIO general purpose input-output pins. The robotic setup is mounted onto a wooden box, and the camera is placed 12 centimeters away from the front of a computer monitor whereby real-world stimuli are displayed (Fig 2.2B).

2.3.3 Computational model

Using the robotic implementation, we incorporate the mechanism of STP within the microcircuit, whereby the neurotransmitter release probability Pr in the synaptic connections




Figure 2.2: **Robotic setup.** (A) Setup of the camera, Raspberry Pi, and the LEDs. (B) Overview of the robotic setup in front of a computer monitor displaying an orientation stimulus from the real-world.

evolves according to:

$$\frac{du_j^s}{dt} = \frac{U_j^s - u_j^s(t)}{\tau_f} + U_j^s(1 - u_j^s(t)) \sum_k \delta(t - t_j^k) \quad (2.1)$$

where δ is the Dirac delta function. The sum on k spikes is over all spike times t_j^k of presynaptic neuron j , and $u_j^s(t)$ reflects presynaptic residual calcium levels at direction step s ($1, 2, 3, \dots, q$) with q corresponding to the number of direction steps. Here, direction steps are defined as the number of discrete spatial positions partaken from a stimulus moving bidirectionally across a single axis of motion. Noteworthy, s does not change at every time instance. Rather, s is incremented after 1000 milliseconds of presynaptic and postsynaptic spike train data. In the absence of incoming action potentials, the synapse is at a resting state with initial Pr $u_j^s(t_0) = U_j^s$. The amount of residual calcium instantaneously increases immediately after the first action potential within a spike train, and $u_j^s(t_1) = 1 - u_j^s(t_0)$ is the fraction that remains available immediately after this first event. Hence, U_j^s is the initial Pr that applies to the first action potential in the spike train, after which the Pr $u_j^s(t)$ decays exponentially to its resting value U_j^s with a time constant τ_f . As a result, each time an action potential is generated, presynaptic residual calcium instantaneously rises and then recovers with a time constant τ_f between subsequent spikes. As residual calcium levels

increase, a release-ready vesicle along the active zone of the presynaptic membrane terminal releases neurotransmitters onto the postsynaptic side of the synapse. During this process of exocytosis, the neurotransmitter availability within the presynaptic terminal is described according to:

$$\frac{dx_j^s}{dt} = \frac{1 - x_j^s(t)}{\tau_d} - u_j^s(t)x_j^s(t) \sum_k \delta(t - t_j^k) \quad (2.2)$$

where $x_j^s(t)$ denotes the fraction of resources that remain available following vesicle release. Between subsequent spikes, $x_j^s(t)$ recovers back to its baseline resting value of 1 with a time constant τ_d , restoring the amount of synaptic resources available within the presynaptic terminal. The STP model allows the examination of synaptic behaviour under a relatively short timescale. Hence, here we are interested in the properties and mechanisms of plasticity over the course of milliseconds to seconds (R. S. Zucker & Regehr, 2002). Depending on the initial setup of kinetic parameters τ_f , τ_d , and U , the STP model can mimic the effect of a depressing or a facilitating synapse (Table 2.1) (Tsodyks & Markram, 1997). Therefore, the mechanism of STD and STF can be distinguished using a different parameter setup in the same governing equations 2.1-2.2.

Table 2.1: Short-term synaptic plasticity parameters.

Parameter	Value
Facilitation recovery (τ_f)	750 ms
Depression recovery (τ_d)	50 ms
Initial neurotransmitter availability (x)	1
Multiplicative factor (A)	0.039

In this study, images that the camera captures from the sensory environment are used as direct input into the microcircuit motif. In other words, external inputs directly generate incoming presynaptic spikes in the STP model. In order to receive synaptic inputs, the stationary camera is used to process changes in the visual image displayed on a computer monitor. Furthermore, we decided to mimic a similar stationary screen fixation setup used in experimental neuroscience (Cavanaugh, 2002; Kelly, Smith, Kass, & Lee, 2010; Smith

& Kohn, 2008), where an anesthetized, paralyzed animal exhibits minimal eye movements to moving stimuli displayed on the screen. As for the nature of the stimulus presented, we are interested in black and white images, each of which is characterized by a 24 bit colour scheme. Here, an average of the three 8 bit RGB planes is taken, and the resulting plane is then converted such that every pixel inside the plane is coded by 1 (black pixel) or 0 (white pixel), instead of traditional values ranging from 0 to 255 (where 255 is the maximal intensity that could be displayed). As the camera processes visual images, the conversion of the image to a binary scheme allows us to directly feed the microcircuit with trains of incoming “all-or-none” action potentials, similar to previous methods (Adams & Harris, 2015).

The Raspberry Pi camera has a native resolution of 3280×2464 (Fig 2.3 – left). Two steps were applied to downsample the native resolution of the camera, ensuring proper image processing from the microcircuit. Firstly, the visible preview of the image was lowered to dimensions 100×1000 . This ensured that the camera’s “receptive field” was small enough to focus on the image capture, but large enough to accommodate unintended displacements from the stationary camera relative to the computer monitor. Finally, the dimensions of the captured image were reduced to 6×1000 to provide the microcircuit with the necessary image processing dimensions. Noteworthy, reducing the dimensions of the visible preview to 6×1000 from the first step would heighten the chances of capturing images of interest that lie outside the visible preview of the camera. In this case, the camera’s receptive field would be too narrow to ensure proper image capturing, which would add undesirable heterogeneity to our controlled experimental setting.

Here, the first dimension represents the number of units to receive the input and 1000 is the size of the temporal window in which each train of presynaptic spikes is evaluated at direction step s (see equations 2.1-2.2). Since the network is comprised of six units, there are 94 remaining units serving as redundant information to the network, meaning information that goes beyond what is necessary from the microcircuit to properly discriminate the

direction of the moving stimulus (Fig 2.3 – centre). Therefore, the dimensionality of the image that was processed from the camera was 6×1000 . The vertical line extending across the first dimension activates all units in the network (Fig 2.3 – right).

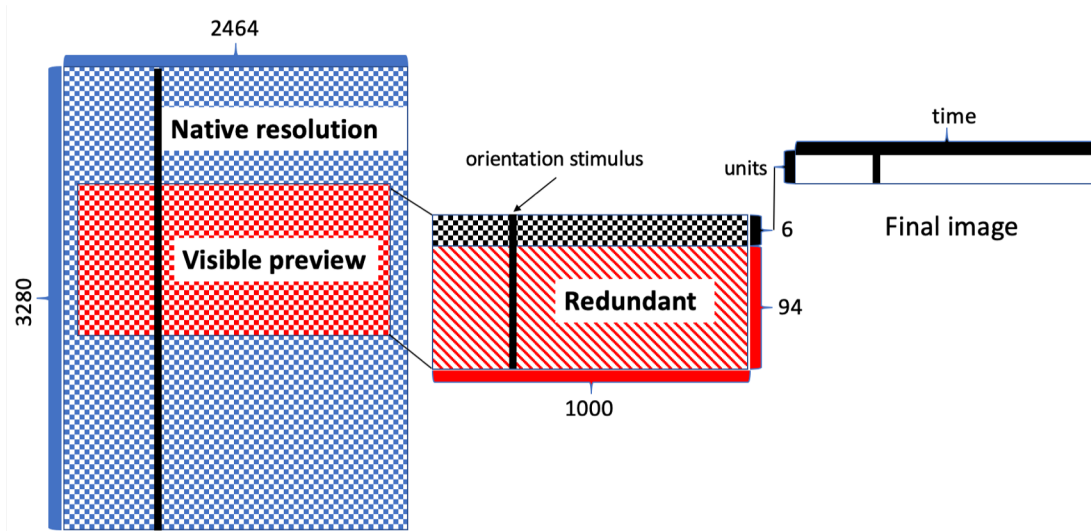


Figure 2.3: **Pre-processing steps for proper image processing.** Schematic representation displaying the transformation starting from the native resolution of the Raspberry Pi camera (left), the visible preview of the image (centre) up to the final image processed by the microcircuit (right). Since the camera remains fixed, the vertical orientation stimulus (dimension 6×10) is the only active component that moves in the 6×1000 dimensional image.

In this way, the camera captures an image of 6000 pixels which is then processed by the spiking network. Here, each unit receives a presynaptic train of 10 spikes evaluated within an interval of 1000ms. In the image, each spike appears as a black pixel whereas the absence of a spike appears as a white pixel. Therefore, the greater the width of the orientation stimulus (here set to 10), the greater the number of +1 pixels and, therefore, the higher the frequency of the presynaptic inputs (here set to 10Hz). In addition, horizontal shifts (e.g. direction steps) in the orientation stimulus will also shift the onset and offset timing of presynaptic spikes evaluated within the window of 1000ms. Finally, given the nature of the real-world stimulus, images were at times inherently noisy, meaning some units received a few more input spikes than others. Although negligible, this external noise source originated from fluctuations in luminance and/or the angle of the monitor relative to the lens of the camera.

In the model, there are a total of 200 direction steps, where each direction step represents a shift in the orientation stimulus image captured by the camera. Therefore, the camera captures 200 images, 100 of which are comprised of orientation stimuli moving in discrete steps towards the right and the remaining 100 images comprised of orientation stimuli moving in discrete steps towards the left. After being processed within an interval of 1000 milliseconds, the orientation stimulus shifts towards a particular direction (i.e. left or right). Consequently, the network receives a new train of incoming presynaptic inputs (Table 2.2).

The shift in the orientation stimulus is accompanied by a change in the initial Pr of STF synapses in subpopulation 1 and 2 (see equations 2.1-2.2). The range of values used to manipulate the initial Pr in each respective subpopulation is presented in Fig 2.4A. At each direction step, a new pair of initial Pr is recruited within each subpopulation. Based on the temporal dynamics of synaptic transmission, units within a given subpopulation would in turn display preferential and non-preferential responses to orientation stimuli. The average initial Pr from each respective subpopulation shows that units within subpopulation 1 exhibit higher initial Pr for right motion, whereas units within subpopulation 2 exhibit higher initial Pr for left motion (Fig 2.4B).

Table 2.2: Leaky integrate-and-fire parameters.

Parameter	Value
Spike emission threshold (θ)	-55 mV
Resting membrane potential (V_{rest})	-70 mV
Membrane resistance (R_m)	200 m Ω
Membrane time constant (τ_m)	30 ms
Absolute refractory period (τ_{arp})	2 ms
Integration time step (dt)	1 ms
Stimulus duration (T)	1000 ms

As units process incoming inputs, the kinetic parameters modulate the interplay between the dynamics of $u_j^s(t)$ and $x_j^s(t)$. In turn, the joint effect of $u_j^s(t)x_j^s(t)$ characterizes the short-term strength of synaptic inputs during time step t at direction step s , thus generating an

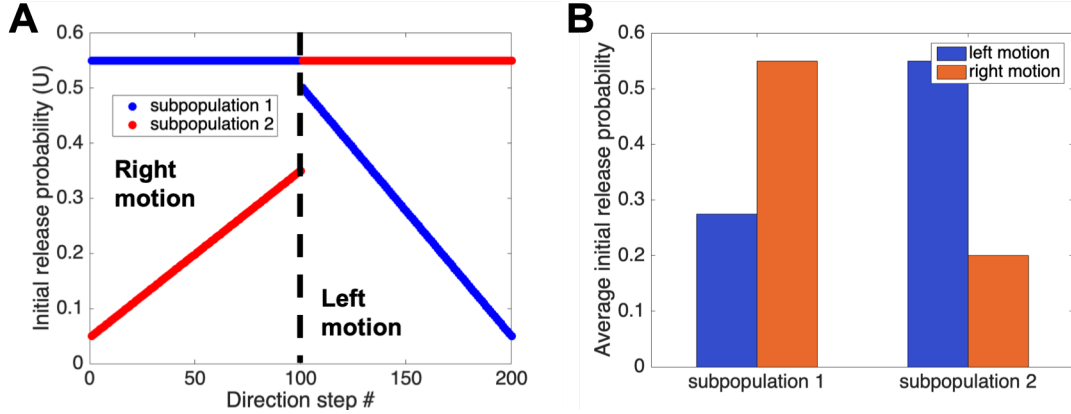


Figure 2.4: **Initial release probability of STF synapses as a function of the direction step number.** (A) Initial release probability U_j^s from units within each subpopulation (U_1^s, U_2^s, U_3^s – subpopulation 1; U_4^s, U_5^s, U_6^s – subpopulation 2) during right $s \in [1..100]$ and left $s \in [101..200]$ motion discrimination. (B) Average initial release probability of STF synapses within subpopulation 1 and subpopulation 2 during left and right motion.

instantaneous current characterized by:

$$I_i^{stp}(t) = A \sum_{j=1}^N w_{ij} u_j^s(t) x_j^s(t) \quad (2.3)$$

where the summation is taken over all presynaptic inputs. Here, w_{ij} is the synaptic weight from presynaptic unit j to postsynaptic unit i influenced by the time-dependent state variables $u_j^s(t)x_j^s(t)$ of the TM model. For all direction steps $\forall s; s \in [1..q]$ and time instance t , w_{ij} is held fixed with $w_{ij} = 1$ between connected units, otherwise $w_{ij} = 0$, denoting the absence of a connection. A is a constant multiplicative factor, modulating the overall gain of the generated current $I_i^{stp}(t)$. In turn, $I_i^{stp}(t)$ drives subthreshold membrane potential depolarization dynamics of leaky integrate-and-fire (LIF) units according to:

$$\tau_m \frac{dV_i}{dt} = -V_i(t) + V_{rest} + R_m I_i^{stp}(t) \quad (2.4)$$

where τ_m is the membrane time constant and R_m is the membrane resistance. $I_i^{stp}(t)$ is the instantaneous current with a time-dependency due to STP. Whenever a depolarization hits a fixed threshold $V_i(t) > \theta$, the unit emits a spike and becomes refractory for a period

τ_{arp} , after which equation 2.4 resumes from a subthreshold reset potential V_{rest} (Table 2.2).

2.3.4 Experimental data

To draw a parallel between the responses of spiking units observed in the robotic implementation versus those observed in an experimental setting, we analyzed data from V1 of visually evoked activity in anesthetized macaque *Macaca fascicularis* monkeys. Resulting recordings were mostly confined to layers 2/3, an area where orientation and direction selectivity are cortical response properties prominently observed. The data were collected in the Laboratory of Adam Kohn at the Albert Einstein College of Medicine and downloaded from the CRCNS website (Kohn & Smith, 2016). Hence, the dataset is taken from previous work where the experimental procedures are described in detail (Cavanaugh, 2002; Kelly et al., 2010; Smith & Kohn, 2008). Briefly, extracellular recordings were performed using Utah multi-electrode arrays inserted 0.6mm into cortex. Animals were paralyzed to minimize eye movements. All experimental procedures complied with guidelines approved by the Albert Einstein College of Medicine of Yeshiva University and New York University Animal Welfare Committees.

The spiking activity of neurons was recorded while presenting full-contrast drifting sinusoidal gratings presented at 12 orientations spaced equally (30°). Drifting gratings were presented binocularly for 1.28 seconds and separated by 1.5 seconds intervals during which a gray screen was presented. Stimulus orientation was randomized, and each stimulus was presented 200 times (i.e., trials). The evoked dataset consisted of spiking activity from 59 to 105 neurons from 3 monkeys (dataset 1, 2, and 3, respectively). To characterize neuronal responses, we chose dataset 3, which included the most amount of neurons (105) out of all 3 datasets. For each orientation of the stimulus moving in the bidirectional motion, the trial-averaged firing rate of individual neurons was computed.

In V1, and other areas of the brain, neurons exhibit high trial-by-trial fluctuations in firing rate (Faisal, Selen, & Wolpert, 2008). Regardless of the nature of the stimulus and the

behavioural state of the animal, a widespread feature of cortical responses is the reduction in trial-by-trial variability around 100 ms following the onset of the stimulus (Churchland et al., 2010). Given that stimulus onset quenches neuronal variability, estimated neuronal responses following a certain delay would in turn provide a more accurate response representation of visual information. Hence, we computed the firing rate of individual cells during the remaining 1 second of the recordings, rather than the entire 1.28 seconds. Furthermore, to remain consistent with the paradigm of the robotic implementation, we analyzed the spiking activity of 6 neurons in the dataset. While neuronal responses for all orientation gratings were analyzed, we focused on finding neurons exhibiting higher responses exclusively for the vertical orientation gratings – the same orientation processed by the camera of the robot. Consequently, we chose 3 neurons (30, 63, and 103) from the dataset exhibiting preferential responses for vertical sinusoidal gratings moving towards the right and non-preferential responses in the opposite null direction. Conversely, we chose 3 other neurons (5, 42, and 98) from the same dataset exhibiting preferential responses to stimuli moving left and non-preferential responses in the opposite null direction.

2.4 Results

Fig 2.5 illustrates the visually evoked activity of individual units, where an orientation stimulus exhibits a rightward motion along a single axis. In this scenario, the activity of the microcircuit is dominated by the response of units within subpopulation 1, where units within this subpopulation exhibit a preferential response for stimuli moving towards the right. Higher responses during right motion discrimination are indicated by the activation of the red LED (Fig 2.5A). Fig 2.5B shows a snapshot of the image captured by the camera during the time at which the stimulus is moving towards the right. Fig 2.5C displays the presynaptic input to each unit within the microcircuit, whereas Fig 2.5D illustrates the visually evoked activity of individual units within the microcircuit during right motion discrimination. Here,

the spiking activity of unit three and unit six is higher relative to the remaining units within their respective subpopulation, because the weighted sum of incoming inputs to these units is larger (Fig 2.1).

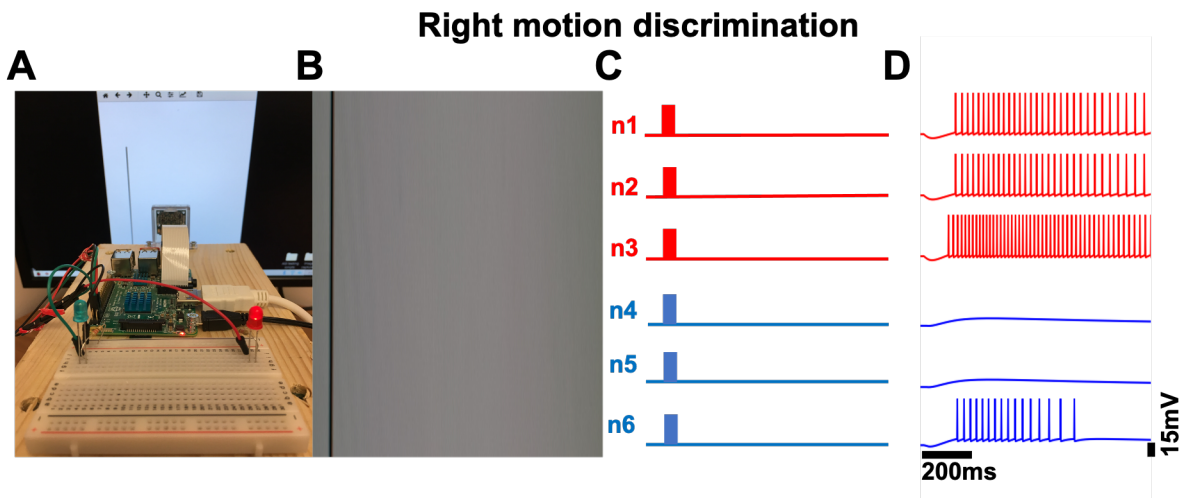


Figure 2.5: **Right motion discrimination.** Shown is a screenshot of the results from the robotic implementation at direction step $s = 10$. (A) Activation of the red LED in response to a stimulus moving towards the right. (B) Real-world stimulus processed by using the Raspberry Pi camera. (C) Trains of presynaptic spikes fed to each unit in the microcircuit. Notice the absence of a refractory period between spikes, and the location of spikes along the x-axis, corresponding to the location of the orientation stimulus shown in B, as depicted on the computer monitor in A. (D) Postsynaptic response of individual units.

Fig 2.6 illustrates the activity of individual units to a stimulus moving towards the left. Here, the activity of the microcircuit is dominated by the response of units within subpopulation 2. Higher responses from subpopulation 2 are in turn represented by the activation of the blue LED (Fig 2.6A). An image of the vertical orientation stimulus is displayed in Fig 2.6B, resulting in direct incoming action potentials in the microcircuit illustrated in Fig 2.6C. Fig 2.6D illustrates the visually evoked activity of individual units within the microcircuit during left motion discrimination. Consistent with Fig 2.5D, the spiking activity of unit three and unit six is higher relative to the remaining units within their respective subpopulation, because the weighted sum of incoming inputs to these units is larger (Fig 2.1). A video illustration of the real-time robotic implementation and the corresponding spatiotemporal patterns of activity of all six units can be found in Supplementary Material S1

(<https://www.hindawi.com/journals/cin/2019/6989128>).

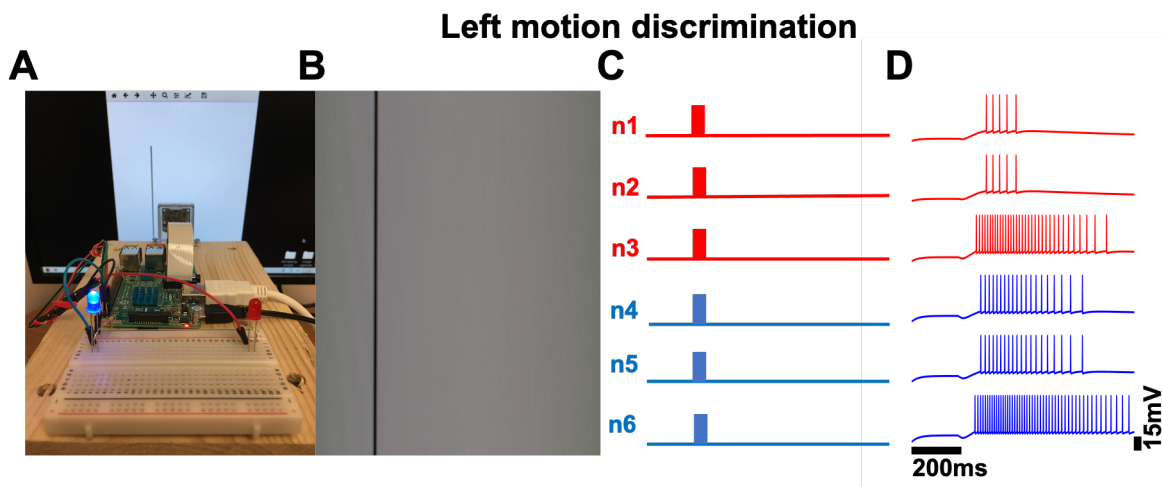


Figure 2.6: **Left motion discrimination.** Shown is a screenshot of the results from the robotic implementation at direction step $s = 180$. (A) Activation of the blue LED in response to a stimulus moving towards the left. (B) Real-world stimulus processed by using the Raspberry Pi camera. (C) Trains of presynaptic spikes fed to each unit in the microcircuit. Notice the absence of a refractory period between spikes, and the location of spikes along the x-axis, corresponding to the location of the orientation stimulus shown in B, as depicted on the computer monitor in A. (D) Postsynaptic response of individual units.

Fig 2.7 displays the average firing rate of units within subpopulations 1 and 2 exposed to stimuli moving in bidirectional motion. The visually evoked response of both subpopulations is shown when the model is exposed to hypothetical stimuli (Fig 2.7A) and real-world stimuli (Fig 2.7B). Units within subpopulation 1 exhibit a preferential response to orientation stimuli moving towards the right. Conversely, units within subpopulation 2 show a higher response to stimuli moving towards the left. Under both scenarios, the average firing rate of units in the preferred and non-preferred direction is highly close to that observed amongst direction-selective neurons in V1 responding to drifting sinusoidal gratings (Fig 2.7C). Fig 2.7D displays the average response of direction-selective cells of interest across all orientations and directions. Among these responses, those resulting from vertical gratings are displayed in Fig 2.7C. In subpopulation 1, the trial-averaged response of 3 cells shows preferential responses for right motion. In subpopulation 2, the trial-averaged response of 3 cells displays preference for left motion.

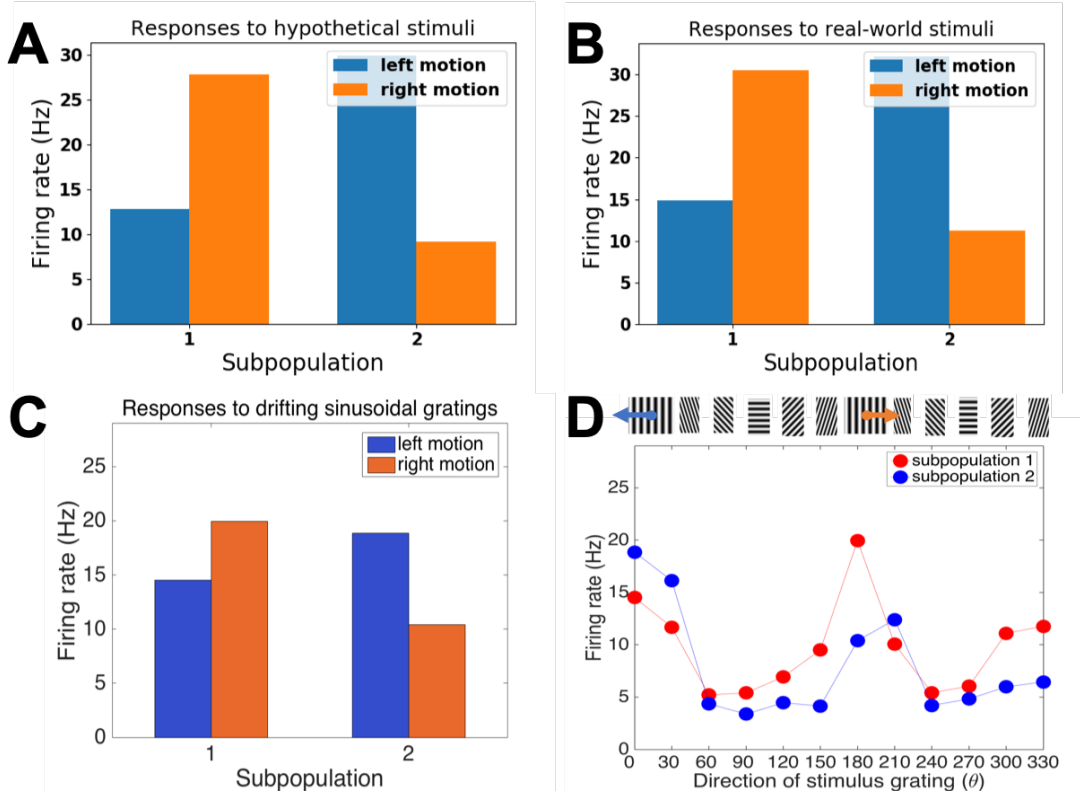


Figure 2.7: **Response of two subpopulations during the simulation, robotic implementation, and multi-electrode recordings of V1.** (A) Responses to stimuli tested in the simulation (B) Responses to real-world stimuli tested in the robot. (C) Responses to drifting sinusoidal gratings tested on a macaque monkey. (D) Responses to drifting sinusoidal gratings presented at 12 orientations tested on a macaque monkey.

Fig 2.8 illustrates the temporal dynamics of the synaptic variables $x_j^s(t)$ and $u_j^s(t)$ as the microcircuit processes visual information. During motion discrimination, units within both subpopulations display synaptic connections that require similar amounts of synaptic resources available in order to properly mediate the response of both subpopulations. Furthermore, the average amount of neurotransmitters available is kept within a high range across the entire temporal domain (Fig 2.8A). This suggests that units within the microcircuit are minimizing use-dependent alterations of synaptic transmission during bidirectional motion discrimination, a scenario that is particularly advantageous when future task demands are required for the robot to perform. Finally, the examination of the release probability in the synaptic connections suggests that units within subpopulation 1 exhibit a higher neurotrans-

mitter release probability for stimuli moving towards the right. On the contrary, synaptic connections within subpopulation 2 display a higher neurotransmitter release probability for stimuli moving towards the left (Fig 2.8B). This suggests that when the microcircuit is exposed to stimuli moving in specific direction, units that exhibit preferential responses to stimuli moving in the specified direction are most likely to be mediated by synaptic connections that exhibit a high release probability (Varela et al., 1997).

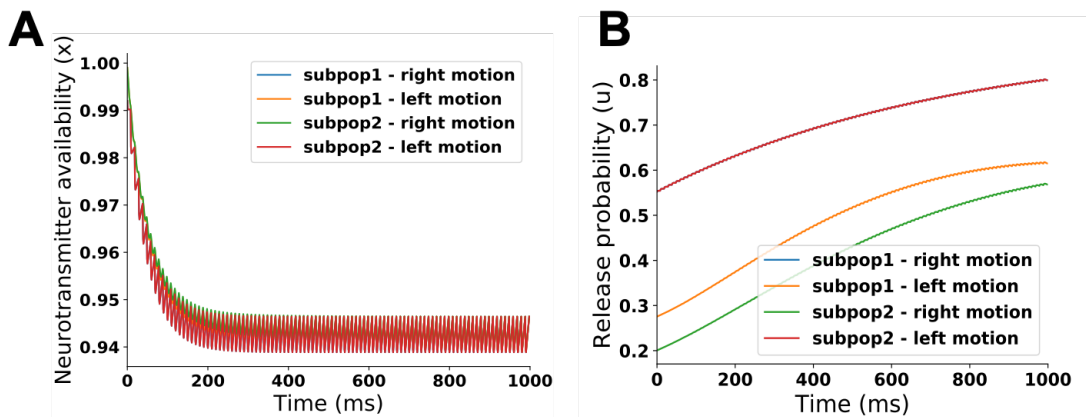


Figure 2.8: **Evolution of the TM model state variables during left and right motion discrimination.** (A) Average synaptic resources available with each subpopulation. (B) Average neurotransmitter release probability in the synaptic connections within each subpopulation. Notice that the temporal evolution of the release probability is the same for subpopulation 1 – right motion and subpopulation 2 – left motion because their average baseline release probability is the same during left and right direction discrimination (Fig 2.4B).

Next, we examined whether the topological structure of our expanded network adds functionality that would otherwise be absent in a two-unit microcircuit. In order to maintain a consistent comparison between the expanded network and the two-unit microcircuit, all of the synaptic connections were mediated by STF (Fig 2.9A). In addition, we examined the activity of each unit using the same kinetic parameters τ_f and τ_d (Table 2.1). Furthermore, we presented the same hypothetical stimulus in motion and recruited the same pairs of initial Pr as those used in the six-unit microcircuit (e.g. U_1^s and U_4^s , or U_2^s and U_5^s , or U_3^s and U_6^s) (Fig 2.4A). Under this framework, both units display highly synchronized activity to stimuli moving in bidirectional motion (see S2 in Supplementary Material for video illus-

tration) – <https://www.hindawi.com/journals/cin/2019/6989128>. In this way, the two-unit microcircuit behaves as a single subpopulation, where each unit displays the same preference to opposite directions of motion (Fig 2.9B). In contrast, the topological structure of the expanded network divides units into two distinct subpopulations, each of which displays preference for a direction opposite to that of its neighbouring subpopulation (Fig 2.7A). Taken together, the two-unit model limits the functional contribution of the microcircuit by displaying preference only for a single direction of motion, whereas the expanded architecture embodies a topological structure that lays the foundation for displaying preferential responses to both directions of motion.

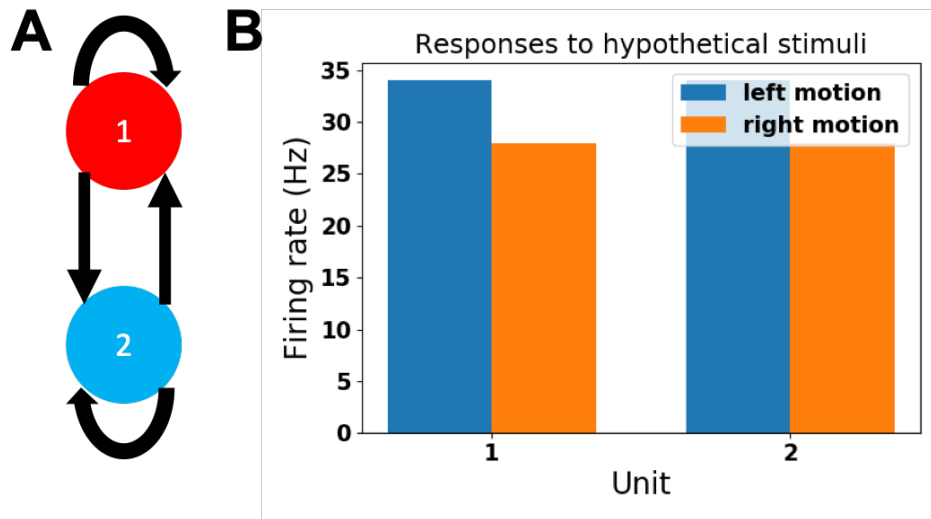


Figure 2.9: **Direction selectivity in the two-unit model.** (A) Architecture of the two-unit microcircuit with STF synapses. Connections are bidirectional, with self-connections allowed. (B) Response of two units to a hypothetical stimulus moving in bidirectional motion. Besides keeping all other parameters the same, the multiplicative factor A (Table 2.1) was instead set to 2, ensuring the mean activity of both units falls within a similar range as that of the six-unit microcircuit (Fig 2.7A and Fig 2.7B).

In cortical microcircuits including V1, neurons exhibit shared fluctuations in population activity overtime (T. A. Engel et al., 2016). In general, these shared fluctuations are measured between pairs of neurons over multiple presentations of an identical stimulus. To examine these coordinated fluctuations in spiking activity, we used a measure of spike count correlation (SCC) between pairs of neurons in V1 during motion discrimination (Vinci, Ven-

tura, Smith, & Kass, 2016). In doing so, we used a non-overlapping time window of 1 millisecond to compute the total number of spikes emitted from each neuron during a given time step. In this way, we obtained the total spike count of individual neurons across time over multiple presentations of the same stimulus in motion. We then computed the pairwise correlation coefficient matrix between 6 neurons, representing the SCC between all pairs of neurons. Finally, we computed the mean SCC observed between pairs of neurons within and between subpopulations. In doing so, we find a positive SCC within subpopulations, and a negative SCC between subpopulations (Fig 2.10A). Interestingly, a positive “within” SCC predicts that fluctuations in the activity of neurons within subpopulations are accurate predictors of a shared preference for a particular direction of motion (Fig 2.10B). Conversely, a negative “between” SCC predicts the presence of an unshared motion direction preference between units belonging to distinct subpopulations (Fig 2.10B). These results were qualitatively captured by our expanded microcircuit of six units (Fig 2.10C and Fig 2.10D). In contrast, the microcircuit of two units displays preference only for a single direction (Fig 2.9B) and therefore fails to predict the presence of an unshared motion direction preference (Fig 2.10E). Taken together, the expanded architecture has a greater predictive power over the two-unit microcircuit, by exhibiting fluctuations in population activity that marks the presence of both shared and unshared motion direction preferences.

2.5 Discussion

Our simple and reproducible robotic implementation highlights the relation between short-term dynamics of synaptic transmission and motion discrimination. In our work, real-world stimuli are directly processed by the microcircuit subject to STP. As units receive external input, they exhibit both preferential and non-preferential responses to stimuli moving bidirectionally along a single axis of motion. Results from the simulation and the robotic implementation are in close agreement to analyses of visually evoked activity in V1, whereby

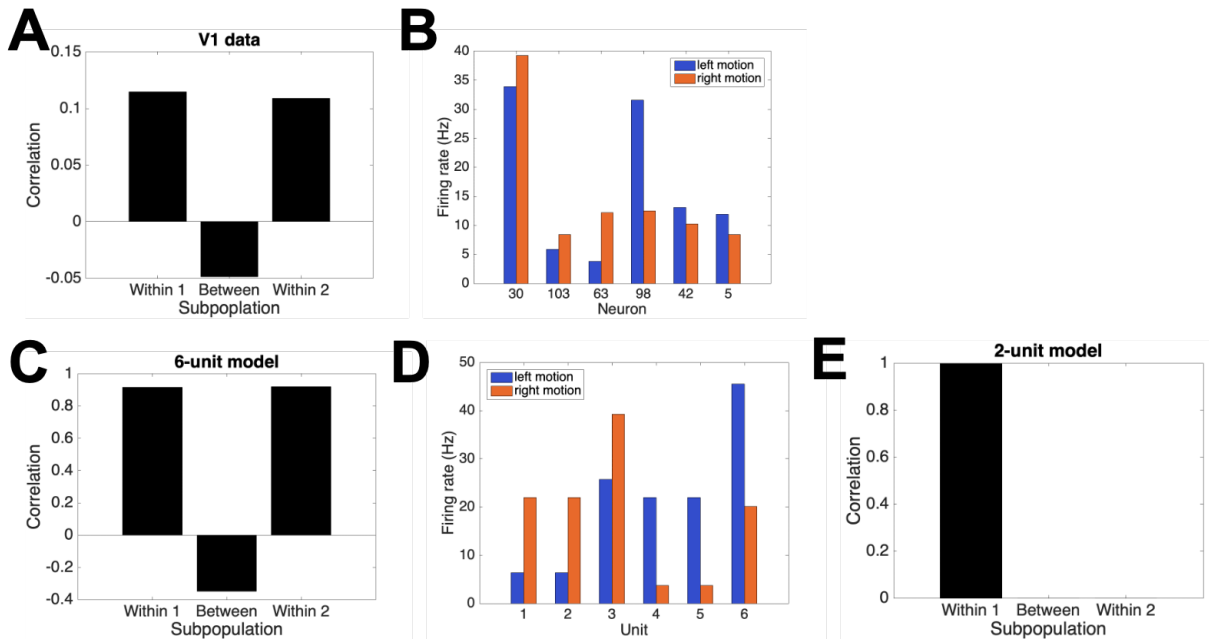


Figure 2.10: **Spike count correlations predict shared and unshared motion direction preferences.** (A) Mean SCC between pairs of V1 cells within and between subpopulations. (B) Average firing rate of six neurons in V1 responding to drifting vertical sinusoidal gratings. (C) Mean SCC between pairs of units within and between subpopulations in the six-unit microcircuit. (D) Average firing rate of six units responding to a vertical orientation stimulus in motion. (E) Mean SCC between two units is equal to 1 (see S2 in Supplementary Material for video illustration) – <https://www.hindawi.com/journals/cin/2019/6989128>. Notice that the two-unit model behaves as a single subpopulation. Therefore, mean SCC “Between” and “Within 2” cannot be computed.

cortical neurons exhibit higher responses for stimuli moving in the preferred direction of motion and lower responses for stimuli moving in the opposite null direction. In addition to accurate motion discrimination, the firing rate of motion-selective units in the STP model is close to the firing rate of direction-selective neurons in V1. As a result, the robotic implementation and the simulated version of the model capture both qualitative and quantitative depictions of typical neuronal responses observed in V1. In addition, units that exhibited preferential responses to stimuli moving in the specified direction were more likely to be mediated by synaptic connections exhibiting a high release probability. Furthermore, the contribution of STP as a complementary mechanism for direction selectivity is validated by

the robotic implementation in real-time, showing successful motion discrimination at the behavioural level. By comparing neuronal responses from the robotic implementation to those of a simulated version of the model, accurate motion discrimination is observed despite the inherent noise of real-world stimuli present in the robotic implementation.

Although visual experience exerts an influence over the direction preference that neurons acquire, the initial topological structure is an essential determinant of direction selectivity (Roy et al., 2018; Van Hooser et al., 2012). In this work, the topological structure of the six-unit microcircuit expanded the repertoire of direction preferences over the two-unit microcircuit, allowing two subpopulations of units to exhibit preferential responses to opposing directions of motion. In contrast, direction selectivity happened to be absent in both the current and previously proposed architecture of two units (Berberian, Ross, Chartier, & Thivierge, 2017). In addition, although units in the two-unit microcircuit exhibited progressive changes in their activity in response to stimulus position, these activity patterns were synchronized (Berberian, Ross, Chartier, & Thivierge, 2017). This functional property is in stark contrast to cortical networks, where asynchronous activity is more commonly observed across cells (Ecker, Berens, Tolias, & Bethge, 2011). With the addition of more units, the larger microcircuit embodied a topological structure which inevitably added asynchronous patterns of neural activity during motion discrimination. Finally, within the context of a larger network size, we show that global fluctuations in population activity can provide an accurate prediction of shared versus unshared motion direction preferences. More specifically, units within subpopulations displayed a positive SCC and were therefore more likely to exhibit fluctuations in subpopulation activity that were accurate predictors of a shared motion direction preference. In contrast, units between subpopulations shared a negative SCC, suggesting that units between subpopulations were likely to display preference for opposing directions of motion. This later prediction was absent in the two-unit microcircuit because the architecture behaved as a single microcircuit, preventing neurons from displaying unshared motion direction preferences.

A predominant view from recent computational work suggests that direction biases present at eye opening may arise purely from “innate” network connectivity (Adams & Harris, 2015). The onset of this architecture is suggested to be present in the absence of any explicit coding for direction selectivity and prior to any self-organizing process facilitated by spontaneous activity or motion-induced training (Adams & Harris, 2015). Similarly, the topological structure of our expanded network was constructed in the absence of any explicit coding for direction selectivity. Hence, our work is in-line with recent experimental and computational studies suggesting that visual experience may serve a permissive role to complement structural processes that are fully characterized at the onset of visual experience (Adams & Harris, 2015; Cyr, Thériault, Ross, Berberian, & Chartier, 2018; Li, Van Hooser, Mazurek, White, & Fitzpatrick, 2008; Roy et al., 2018). Hence, the organization of the initial architecture may lay the foundation for the map of direction preference, as observed in the visual cortex (Van Hooser et al., 2012).

Visually evoked activity is likely to be mediated by a variety of mechanisms operating at different timescales and at distinct developmental stages (Clemens et al., 2012). Therefore, given the wide range of plasticity rules (Markram, 2011), it is likely to expect other candidate mechanisms that are complementary to short-term changes in synaptic strength. Indeed, there is experimental evidence lending support for the interaction between STP and long-term synaptic plasticity (Loebel, Le Bé, Richardson, Markram, & Herz, 2013; Markram et al., 1997; Monday & Castillo, 2017; Monday, Younts, & Castillo, 2018; P. Sjöström et al., 2007). Amongst long-term changes in synaptic strength, spike-timing-dependent plasticity (STDP) has been proposed as a candidate mechanism that strengthens or weakens synapses based on the relative timing of action potentials. Despite operating at different timescales, experimental studies have shown that STP and STDP may interact (Jin et al., 2012; Kintscher et al., 2013; Markram & Tsodyks, 1996; Tokuoka & Goda, 2008). Given that bio-inspired learning rules (e.g. STDP) can provide robotic systems with flexibility, we intend to combine STP and STDP to study their synergistic interaction in a more dynamic

neurorobotic setting.

2.6 Conclusion

The overall aim of this study was to provide a step towards applying well-established formal models from neuroscience into the domain of neurorobotics. To this end, we designed a microcircuit motif of six units connected by STF synapses. The topological structure of the embodied microcircuit enable units to exhibit direction selectivity. Results from the formal model captured both qualitatively and quantitatively the response characteristics of direction-selective cells in V1. In chapter three, we move away from the static online robotic implementation, towards a more dynamic neurorobotic setting. There, we develop a keyboard listening framework for online mobile robot control. Furthermore, we devise a larger network of spiking neurons, and study the working memory capacity of the agent. Finally, we show how short-term and long-term changes in synaptic plasticity provides the embodied network with enough flexibility to (1) refine, (2) overwrite or (3) resist working memory contents.

Chapter 3

Embodied Working Memory During Ongoing Input Streams

3.1 Abstract

Sensory stimuli endow animals with the ability to generate an internal representation. This representation can be maintained for a certain duration in the absence of previously elicited inputs. The reliance on an internal representation rather than purely on the basis of external stimuli is a hallmark feature of higher-order functions such as working memory. Patterns of neural activity produced in response to sensory inputs can continue long after the disappearance of previous inputs. Experimental and theoretical studies have largely invested in understanding how animals faithfully maintain sensory representations during ongoing reverberations of neural activity. However, these studies have focused on preassigned protocols of stimulus presentation, leaving out by default the possibility of exploring how the content of working memory interacts with ongoing input streams. Here, we study working memory using a network of spiking neurons with dynamic synapses subject to short-term and long-term synaptic plasticity. The formal model is embodied in a physical robot as a companion approach under which neuronal activity is directly linked to motor output.

The artificial agent is used as a methodological tool for studying the formation of working memory capacity. To this end, we devise a keyboard listening framework to delineate the context under which working memory content is (1) refined, (2) overwritten or (3) resisted by ongoing new input streams. Ultimately, this study takes a neurorobotic perspective to resurface the long-standing implication of working memory in flexible cognition.

3.2 Introduction

Animals are capable of relying on internal sensory representations rather than purely on the basis of external stimuli. These representations routinely support higher-order functions such as working memory (WM) – holding the stimulus content in memory for a certain duration in the absence of its concrete presence (Chafee & Goldman-Rakic, 1998; Funahashi, Bruce, & Goldman-Rakic, 1989; P. S. Goldman-Rakic, 1996; E. K. Miller, Erickson, & Desimone, 1996; E. Miller, Li, & Desimone, 1993; Takeda & Funahashi, 2002). Experimental studies on WM in primate neurophysiology have shown that patterns of neural activity can reverberate following the offset of sensory inputs (Constantinidis et al., 2001; E. K. Miller et al., 1996). These ongoing reverberations can persist over a period of several seconds (Chafee & Goldman-Rakic, 1998; Constantinidis et al., 2001; P. S. Goldman-Rakic, 1995, 1996; Takeda & Funahashi, 2002), a time segment during which neuronal responses remain preferentially elevated for a target stimulus (Berkes, Orban, Lengyel, & Fiser, 2011; Kenet, Bibitchkov, Tsodyks, Grinvald, & Arieli, 2003; Luczak, Barthó, & Harris, 2009). Despite having shown to be correlated with psychophysical performance (Constantinidis et al., 2001), mechanisms underlying intrinsic dynamics of self-sustained activity remain elusive. Consequently, the significance of reverberating activity for higher-order function remains unknown.

A long-standing question shared across theoretical and experimental studies of WM is how animals maintain a faithful representation of sensory stimuli during ongoing reverberations of neural activity. Persistent activity during WM has been studied extensively via influ-

ential observations of experimental data (Fuster & Alexander, 1971; P. S. Goldman-Rakic, 1995; Miyashita & Chang, 1988). These experiments have motivated the development of computational models, which in turn have triggered fundamental theoretical insights into potential mechanisms underlying WM capacity (Bouchacourt & Buschman, 2019; Compte, 2000; Hansel & Mato, 2013; Macoveanu, Klingberg, & Tegnér, 2006; Seeholzer, Deger, & Gerstner, 2019; Wei, Wang, & Wang, 2012). These formal models, among others, form recurrently connected networks, which typically maintain selective elevated persistent activity through local excitatory recurrent connections with global feedback inhibition (Koulakov, Raghavachari, Kepecs, & Lisman, 2002). In this formalism, sustained firing may be achieved by carefully fine-tuning the strength and structure of recurrent circuitry (Barak & Tsodyks, 2014). As such, these functional networks can maintained, in memory, the spatial location of target stimuli, forming what is commonly known as persistent activity bumps (i.e. bump attractors) (Koulakov et al., 2002).

Persistent activity in recurrent networks has been supported by short-term synaptic plasticity, where synaptic strength is rapidly regulated by recent historical activity within the network (Itskov, Hansel, & Tsodyks, 2011; Schneegans & Bays, 2018; Seeholzer et al., 2019). In the presence of such rapid changes in synaptic dynamics, neuronal activity can drift over time (Schneegans & Bays, 2018; Seeholzer et al., 2019), or remain centered at the initial bump location (Itskov et al., 2011). Despite their overarching support from empirical studies (Barbosa et al., 2020; Edin et al., 2009), WM models holding the persistent memory hypothesis have for long been using preassigned protocols of fixed stimulus presentation. In this context, previous methods have primarily focused on one aspect of WM, namely the maintenance of a specific target stimulus – which leaves out by default its interaction with ongoing input streams. In a real-world setting, biological agents are continuously bombarded with stimuli, some of which they maintain in WM while new input streams actively harness their ongoing sensory experience. How does the brain manage to maintain specific target stimuli in WM, while withstanding streams of incoming stimuli? As the environment con-

tinuously delivers a rich repertoire of sensory intricacies, the brain must somehow reconcile these two seemingly contradictory tasks – maintaining previous inputs stable over a period of time while new input streams are continuously coming in. To this end, WM must have a temporal undercurrent, a basis under which it integrates past and present time points.

In contrast to displaying persistent activity by recurrent interactions, here we instead devise a feedforward network of spiking neurons (Goldman, 2009), subject to both short-term and long-term synaptic plasticity (Amit, 1997; Del Giudice, Fusi, & Mattia, 2003; Fusi, 2002; Giudice & Mattia, 2001). By considering activity-dependent Hebbian plasticity as a complementary mechanism for generating persistent activity (Manohar, Zokaei, Fallon, Vogels, & Husain, 2019), we examine how the content of WM interacts with the intricacies of ongoing input streams. Spiking networks are well-known for carrying out continuous online operations in non-stationary environments (Lobo et al., 2020). In this context, synaptic connections can be continuously modified (Zenke et al., 2017), without fine-tuning their strength and structure. This ongoing process of modification may result in momentary network restructuring so as to accommodate an evolving environment (Lobo et al., 2020). Here, we present six different experiments, namely (1) single learning and recall, (2) incremental learning and recall, (3) task-switching, (4) resistance to interference, (5) submission to interference, and (6) resistance to distraction. Each experiment provides an entry point for the next one (see Robotic experiments), highlighting the scenario under which the WM content may be (1) refined, (2) overwritten or (3) resisted by the arrival of ongoing input streams.

Online unsupervised learning of WM capacity has been studied in networks of spiking neurons connected by plastic synapses (see (Del Giudice et al., 2003) for review). These networks can undergo an ongoing process of structural modification, with no *a priori* knowledge about the content of subsequent input streams (Amit & Mongillo, 2003; Del Giudice et al., 2003; Fusi, Annunziato, Badoni, Salamon, & Amit, 2000; Lobo et al., 2020). In this context, previous target inputs may be gracefully overwritten by subsequent inputs, because synaptic plasticity can provide the malleability needed for the content of WM to ride the wave of new

input streams. Networks that operate in this fashion share the palimpsest property, meaning they forget the content of previous inputs to make room for new ones (Amit & Fusi, 1994; Nadal et al., 1986; G. Parisi, 1986). Under some circumstances however, the content of WM may still be retrieved despite the arrival of subsequent inputs, because plastic synapses behave in a context-specific manner (Amit & Fusi, 1994). For example, previous studies have shown that the stimulus statistics can have an impact on the state transition probability of the synapse – the probability of transitioning to a new state different from the previous state (Del Giudice et al., 2003; Fusi et al., 2000).

The WM of biological agents operating in the real-world is inevitably carried out while subsequent input streams continuously arrive. In support of their biological counterparts, artificial agents require adaptive capabilities in order to be deemed useful in the presence of ongoing changes in the external world. Indeed, physical robots pose as ideal companions for carrying out sequential tasks online, whereby the artificial agent must process incoming information on-the-fly. Here, we embed the spiking network in the “Vector” robot produced by Anki developer. We perform our experiments using the agent as a tool for studying WM capacity during ongoing input streams. Importantly, the robot is used to ensure a one-to-one correspondence between network activity and behaviour – a companion tool for moving closer to the real-world interaction that exists between WM and ongoing input streams.

3.3 Methods

3.3.1 Network architecture

The spiking neural network consisted of 500 units connected sparsely. The initial network connectivity followed a set of rules (Berberian, MacPherson, Giraud, Richardson, & Thivierge, 2017). First, inhibitory neurons occupy approximately 20% of the population of cortical neurons, whereas the great majority of remaining neurons are excitatory (Gabbott & Somogyi, 1986). As such, 80% of units in the embodied network were randomly

chosen to be excitatory, whereas the remaining 20% were inhibitory. Following Dale’s law, a given excitatory/inhibitory unit only exhibited excitatory/inhibitory efferent connections, respectively. Second, the number of afferent synapses to a single neuron in cortex is limited, forming clusters of sparsely connected networks with roughly 10-20% probability of synaptic contact (Braitenberg & Schüz, 1998). To this end, initial network connectivity was sparse, with only 20% of all possible connections present (chosen randomly among all possible connections). At network initialization, self-connections were not permitted because although not uncommon, they are rare *in vivo*, in comparison to their prominence in dissociated cell cultures (Bekkers, 2003). Finally, efferent synaptic connections from excitatory units (J_{EE} , J_{IE}) were positive, whereas those from inhibitory units (J_{II} , J_{EI}) were negative. Network parameter values are indicated in Table 3.1.

Table 3.1: Network parameters.

Symbol	Description	Value
c	Initial probability of synaptic contact	0.2
N_E	Number of E units	400
N_I	Number of I units	100
J_{EE}	Initial E to E synaptic efficacy	0.65
J_{IE}	Initial E to I synaptic efficacy	0.65
J_{II}	Initial I to I synaptic efficacy	-1
J_{EI}	Initial I to E synaptic efficacy	-1

3.3.2 Robotic platform

The Vector Software Development Kit (SDK) was used as an open-source robotics platform. In particular, we used a physical robot named “Vector” by Anki Developer Fig 3.1. This robot is equipped with five basic hardware components, namely (1) an HD camera, (2) an infrared laser sensor, (3) a beamforming 4-microphone array, (4) a capacitive touch sensor and (5) two wheel motors. The computational resources of the robot were remote, performed on a laptop computer. Incoming stimuli were controlled via user keypresses from the laptop computer, and the communication was mediated via wireless network. This embedded

control design was intended so as to avoid using Vector’s onboard tactile sensors. To explain, the robot continuously moved during the experiments. Therefore, reliance on tactile sensors would demand the user to continuously follow the robot throughout the experiment. Moreover, sensor reading is noisy and therefore some tactile triggers would go unnoticed from Vector failing to read them as valid touch. These would vastly limit the fluidity of information exchange, because they would place high workload demands on the user (but see (Weng, 2004)). Consequently, the time required for conveying user desirability would increase. For these reasons, we have instead devised a keyboard listening framework. The protocol was designed to ensure an efficient and instant medium of communication between the user and the robot (Goodrich & Schultz, 2007). Finally, there were no human subjects recruited in this study. The degree of user involvement in the teleoperation is similar to the degree of involvement in pressing a key to compile a code on a computer.

In order to grant access to the robot’s hardware capabilities, API commands were communicated to the corresponding actuators of the robot. The only hardware components used in this study were Vector’s two wheel motors. The execution of ongoing motor trajectories was mediated by the spiking activity of the network. Spike count information was computed in consecutive non-overlapping time bins of 40ms. The rotating speed of the two wheel motors were decoded separately. Each motor rotated at a given speed (in millimetres per second) until the total spike count of the next non-overlapping time bin commanded the robot to drive at a new speed. Taken together, neuronal responses were mapped onto the wheel motors of the robot, controlling the execution of its navigated trajectory.

3.3.3 Sensory inputs

Sensory-evoked inputs were presented to the network in the form of spatially localized currents represented by a bimodal mixture of Gaussians:

$$I_i^{ext}(t) = R_b + R_p e^{-\frac{(i-i_{pref})^2}{2\sigma^2}} + R_n e^{-\frac{(i+(\frac{N}{2})-i_{pref})^2}{2\sigma^2}} \quad (3.1)$$



Figure 3.1: **Vector the robot.** Vector behaves as a result of network spiking activity.

where R_b was the baseline amplitude of sensory inputs, R_p was the higher amplitude peak whereas R_n was the lower amplitude peak, σ was the tuning width controlling the input specificity, i was the unit index ($1, 2, 3, \dots, N$) with N corresponding to the number of units.

There are three main reasons why bimodal activations were used, including the specified offset between the two Gaussians. First, previous models of spatial WM have used the Gaussian for characterizing their target stimulus. Inputs range from a single target to multiple targets (Bouchacourt & Buschman, 2019; Compte, 2000; Edin et al., 2009; Eliasmith, 2005; Hansel & Mato, 2013; Itskov et al., 2011; Kilpatrick, 2018; Laing & Chow, 2001; Renart, Song, & Wang, 2003; Sandamirskaya, 2014; Seeholzer et al., 2019; Tanaka, 2002). Second, Vector is equipped with two wheel motors. The agent was therefore an ideal candidate for supplying its differential steering system with bimodal activations. The specific offset between the Gaussians induced differences in wheel motor speed, which in turn resulted in rotating movements with diameters small enough for the robot to accommodate the confined spatial setting. Finally, the offset between the two Gaussians has experimental grounding, inspired from a variant of the classic oculomotor delayed response task (Constantinidis et al., 2001). In these experiments, animals are shown two simultaneously presented stimuli, with

differences in luminance (i.e. contrast ratio).

The network involved two subpopulations, consisting of spiking units spatially distributed according to the sensory input to which they were most sensitive (Fig 3.2). The bimodal activations differentially drove the activity of units within each respective subpopulation, producing localized activity bumps determined by the spatial structure of injected external input currents I_i^{ext} . In particular, i_{pref} corresponded to the unit most sensitive to R_p , and therefore exhibited the highest response amongst its neighbouring units within the subpopulation.

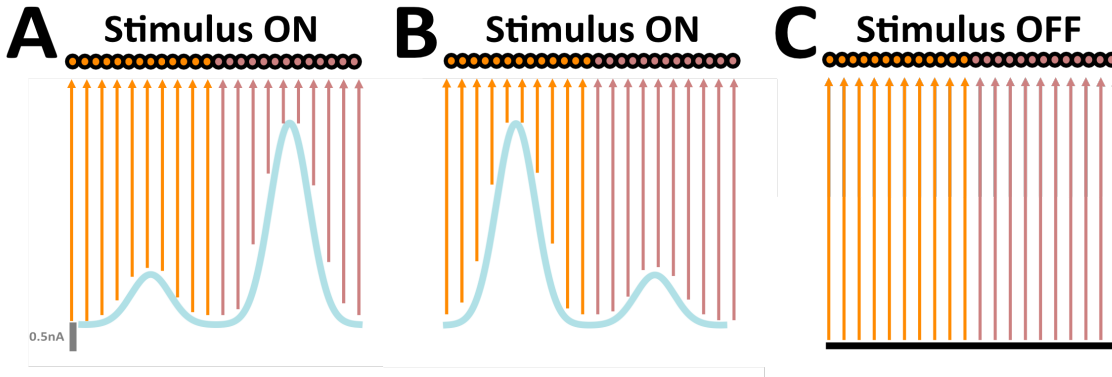


Figure 3.2: **Sensory afferents to the embodied network.** (A) ON stimulus (blue). (B) Same as (A), but instead peaks are interchanged (C) OFF stimulus (black); orange, subpopulation 1; brown, subpopulation 2. Parameters: $R_b = 0.5$ nA, $R_p = 2.5$ nA, $R_n = 1.0$ nA, $i_{pref} = 125$ for clockwise rotations, $i_{pref} = 375$ for counter-clockwise rotations, $N = 500$, $\sigma = 35$. Figure adapted from (Itskov, Hansel, & Tsodyks, 2011).

In this study, a keyboard listening framework was devised in order to actively control ongoing input streams to the embodied network (see Robotic platform). By delivering keypresses from a laptop computer, the user could control both the configuration (Fig 3.2) and the duration of stimuli. Sensory inputs were represented by two distinct configurations (Fig 3.2A and 3.2B). Their average intensity were identical. The only distinguishing characteristic between the two was their interchangeable amplitude peaks. In configuration 1, units in a given subpopulation received inputs near R_p , whereas those in the remaining subpopulation received inputs near R_n (Fig 3.2A). In configuration 2, spatially localized amplitude peaks were interchanged (Fig 3.2B). Overall, inputs evoked localized activity profiles in the

network (i.e. bump states). In their absence, $I_i^{ext} = 0$ (Fig 3.2C).

3.3.4 Learning and recall procedure

During WM tasks (e.g. oculomotor delayed response), the animal experiences a “cue period” where external stimuli are presented. Stimulus presentation is immediately followed by a “delay period” where external stimuli are removed. As an analogy to typical experiments, our study treated the cue period as learning, and the delay period as recall. For simplicity, the “response period” normally characterized by saccadic eye movements towards a target location were not considered. Overall, the temporal evolution of the model was organized in batches of learning and recall in order to maintain an operational definition used by traditional approaches in artificial neural networks (ANNs).

Learning and recall phases were carried out by distinct time segments of network activity during which the robot exploited input-dependent and input-disengaged operations, respectively (Fig 3.3). The offset of both learning and recall is controlled by pressing the <Enter> key (at any time instance). To mark the onset of learning, the <Left> or <Right> key is pressed, such that the robot executes counter-clockwise or clockwise motor trajectories, respectively (Fig 3.3A and 3.3B). To end the learning phase, the <Enter> key is pressed (not shown), allowing the robot to move autonomously (Fig 3.3C). Overall, the learning phase was considered as a time segment of evoked activity where units were driven by sensory afferents. In contrast, the recall phase was considered as a time segment of spontaneous activity where units were solely driven by the intrinsic dynamics of the network.

Learning and recall phases were alternated by turning the stimulus ON and OFF, repeatedly (ON-OFF-ON-OFF...). This sequence-based approach allowed us to include but also move beyond traditional approaches of single learning and recall. The spiking network was subject to continuous modification. This ongoing process naturally removes the fine line that has traditionally been used to distinguish learning from recall in ANNs (Zenke et al., 2017). Activity-dependent changes in synaptic efficacy were exclusively based on local infor-




Figure 3.3: **Keyboard listening setup.** (A) The <Left> keypress recruits input configuration shown in Fig 3.2A, marking the onset of input-dependent counter-clockwise motor executions. (B) Same as in (A), but instead the <Right> key is pressed, recruiting input configuration shown in Fig 3.2B, marking the onset of input-dependent clockwise motor trajectories. (C) The <Enter> key is pressed (not shown), recruiting the homogeneous input configuration shown in Fig 3.2C, marking the onset of input-disengaged operations.

mation (see Plasticity model). Evoked activity led to the formation of synaptic structure in a rate-dependent Hebbian manner (Del Giudice et al., 2003). The synaptic structure provided sufficient structured feedback so as to endow units with the ability to maintain stimulus-selective persistent activity (Amit, 1995; Amit & Mongillo, 2003). The structure of the weight matrix representing the stimulus information was put in a state of reverberation that outlasted the exposure of the stimulus. Hence, the role of the resulting synaptic structure was to sustain a local activity bump produced by previous stimuli. In this way, items that were loaded as discrete “slots” during learning were maintained by distinct localized activity during recall. If the stimulus was absent for long durations, network activity progressively decayed back to baseline levels (Goldman, 2009) – an expression reflected in the network structure. Taken together, the initial Network architecture was not designed to meet the requirements for a particular task, but was rather subject to ongoing modification through

the sensory experience of the robot (Weng, 2004).

Robotic experiments

We report the procedure of six robotic experiments and lay out the criteria for their successful completion. In the first experiment, a single learning and recall procedure is introduced. Since network activity is unavailable to the user, the behaviour of the robot is monitored to gain insight into network convergence. Here, the criteria for successful task completion is to observe consistent movement trajectories from the robot. This approach differs from disembodied networks imprinted with a minimum stopping rule for weight convergence, because robotic behaviour is the marker of network stability, not preassigned internal triggers. Using movement trajectories as top-down evidence of network convergence, the user can then remove the stimulus and observe the recall performance of the robot. This observation is used as a stepping stone for introducing an incremental paradigm.

In the second experiment, an incremental learning and recall procedure is presented. This experiment encompasses the notion of behavioural shaping (Chernova & Thomaz, 2014). Here, shaping refers to an incremental process where the user provides feedback to an agent so as to improve approximations of a target behaviour (Knox, Fasel, & Stone, 2009; Knox, Stone, & Breazeal, 2013). The user observes the behaviour of the robot in response to different feedback signals (e.g. stimulus duration), and makes the necessary adjustments on-the-fly. In this way, robotic behaviour is progressively fine-tuned, moving the agent closer to the desired behaviour (Saksida, Raymond, & Touretzky, 1997). This incremental process towards a single target behaviour begs the question, however, whether the robot can move beyond single task demands, by transitioning between multiple tasks.

In the third experiment, a task-switching procedure is introduced. The robot is shown two interchanging inputs. The criteria for success is the expression of flexible behaviour. Here, sustained learning is not a necessary requirement for successful task completion, because task-switching washes away the memory of the previous task. The robot must adapt to

change at every step of the way, accommodating user desirability. Nevertheless, is the robot always deemed to be accurately aligned with user judgements?

In the fourth experiment, a task interference procedure is introduced. Here, the robot is presented with a long cue stimulus, and tested against a brief interfering stimulus. To successfully accomplish this experiment, the robot must keep performing movement trajectories in the same direction as those observed during the cue period, even after the user interferes by imposing movements in the opposite direction. In this scenario, if the robot resists interfering inputs, is it because they are presented for a shorter duration or because the content of the cue stimulus is still maintained in WM?

In the fifth experiment, a variant of the task interference procedure is presented. Here, delay duration is extended long enough to wash-away the content of the cue stimulus. When the robot forgets the cue stimulus, its movement trajectory provides the user with enough evidence to present the interfering stimulus. To successfully perform this experiment, brief interfering stimuli should be enough to overwrite the content of cue stimuli. Otherwise, the previous experiment would suggest that the robot resisted interfering inputs simply because they were presented for a shorter duration, and not because the content of cue stimuli were maintained in WM. Nevertheless, is the robot capable of resisting distractions that last as long as the cue stimulus?

In the sixth and final experiment, the robot is tested against distractor inputs. Here, the duration of both cue and distractor inputs are similar. However, distractor intensity is three times weaker than cue intensity. The magnitude of distractor intensity is based on two requirements, namely (1) to impose distracting movements in a direction opposite to those imposed during the cue period and (2) to navigate within the confined spatial environment. To successfully accomplish this experiment, the robot must move in the same direction before and after distraction, despite movements in the opposite direction during distraction. Taken together, the behaviour of the robot, not just the network activity, is a requirement for the successful completion of our experiments.

3.3.5 Spiking neural model

The spiking neural network was numerically simulated via current-based synapses. The network was composed of adaptive exponential integrate-and-fire (aEIF) units, a well-known model that captures the spike initiation behaviour and statistics of cortical neurons (Brette & Gerstner, 2005; Naud et al., 2008). Neuronal dynamics were described by a membrane potential V_i and by a spike-frequency adaptation w_i – a widespread neurobiological phenomenon (Benda & Herz, 2003). Below a constant threshold V_T , units evolved according to:

$$C_m \frac{dV_i}{dt} = -g_L(V_i - E_L) + g_L \Delta_T e^{\frac{V_i - V_T}{\Delta_T}} - w_i + I_i^{syn} + I_i^{ext} \quad (3.2)$$

where C_m was the membrane capacitance, g_L was the leak conductance, E_L was the resting potential, and Δ_T was a slope factor, which determined the sharpness of the threshold. Here, I_i^{syn} described the postsynaptic current generated from the intrinsic dynamics of the network, I_i^{ext} was the sensory input for which its presence $I_i^{ext} > 0$ or absence $I_i^{ext} = 0$ was controlled using keypresses (see Sensory inputs). The postsynaptic membrane voltage (Eq 3.2) was coupled to an adaptation variable w_i defined by:

$$\tau_w \frac{dw_i}{dt} = a_w(V_i - E_L) - w_i \quad (3.3)$$

where τ_w was the adaptation time constant and a_w represented the level of subthreshold adaptation. When the membrane potential crossed the threshold ($V_i > V_T$), a spike (V_{peak}) was triggered in the aEIF model, and the integration of Eq 3.2 was reset to V_{reset} ($V_i \rightarrow V_{reset}$), with $V_{reset} = E_L$. In parallel, the adaptation variable in Eq 3.3 was increased by an amount b_w ($w_i \rightarrow w_i + b_w$), where b_w was the parameter of spike-triggered adaptation. Single unit parameter values of the aEIF model were taken from (Brette & Gerstner, 2005), and kept fixed throughout all robotic experiments (Fig 3.4).




Figure 3.4: **Single unit profile of the aEIF model.** Postsynaptic spikes from input current injection. Top: input current, Middle: membrane potential, Bottom: adaptation variable. Parameters: $C_m = 281$ pF, $g_L = 30$ nS, $E_L = -70.6$ mV, $V_T = -50.4$ mV, $V_{peak} = 20$ mV, $\Delta_T = 2$ mV, $\tau_w = 144$ ms, $a_w = 4$ nS, $b_w = 0.0805$ nA, $\tau_{syn} = 5$ ms.

3.3.6 Plasticity model

Short-term plasticity (STP) was applied on excitatory and inhibitory synapses according to a well-established model (Tsodyks & Markram, 1997). Efferent synaptic connections from presynaptic units were modelled by the following system of ordinary differential equations:

$$\frac{du_j}{dt} = \frac{U - u_j}{\tau_f} + U(1 - u_j) \sum_k \delta(t - t_j^k) \quad (3.4)$$

$$\frac{dx_j}{dt} = \frac{1 - x_j}{\tau_d} - u_j x_j \sum_k \delta(t - t_j^k) \quad (3.5)$$

where δ was the Dirac delta function, t_j^k was the occurrence time of the k th spike of presynaptic unit j , $U \in [0, 1]$ represented the baseline release probability of u_j , and x_j

was the neurotransmitter availability. Presynaptic units in the network received intrinsic background input in the form of a spike train s_j , which was an independent homogeneous Poisson process for each presynaptic unit j with rate r of 10Hz (Amit, 1997). For simplicity, the spatiotemporal pattern of the background noise remained constant during both evoked and spontaneous activity. In the absence of presynaptic spikes, unit j was at a resting state, where u_j and x_j remained at baseline resting values U and 1, respectively. When a presynaptic spike occurred, u_j instantaneously increased $u_j \rightarrow u_j + U(1 - u_j)$, while x_j instantaneously decreased $x_j \rightarrow (1 - u_j)x_j$. Between subsequent spikes, u_j and x_j recovered back to their baseline resting values U and 1, respectively. Depending on the initial setup of parameters, τ_f , τ_d and U , the phenomenological model of STP can mimic the effect of a depressing or a facilitating synapse (Markram et al., 1998). Therefore, the mechanism of short-term depression and facilitation can be distinguished using a different parameter setup in the same governing equations. In our embodied network, excitatory and inhibitory synapses were subject to short-term depression, which is well-suited in shaping the temporal response properties of cortical neurons (Abbott et al., 1997; Chance et al., 1998; Tsodyks & Markram, 1997).

Spike-timing-dependent plasticity (STDP) was implemented using a pair-based, nearest-neighbour spike timing interaction (P. J. Sjöström et al., 2001). At each time instance, the induction protocol evaluated the connection between presynaptic unit j and postsynaptic unit i , and computed a time difference $\Delta t = t_i - t_j$ between the last spike emitted from each respective unit. In this way, the nearest time difference between spikes was continuously evaluated by the weight updating STDP rule described as:

$$\Delta J_{ij} = \begin{cases} \lambda_+ f_+(J_{ij}) e^{-\frac{\Delta t}{\tau_+}} & \text{if } \Delta t > 0 \\ -\lambda_- f_-(J_{ij}) e^{\frac{\Delta t}{\tau_-}} & \text{if } \Delta t \leq 0 \end{cases} \quad (3.6)$$

$$f_+(J_{ij}) = (1 - J_{ij})^\mu \quad \text{and} \quad f_-(J_{ij}) = \alpha (J_{ij})^\mu \quad (3.7)$$

where λ_{\pm} , $0 < \lambda_{\pm} \ll 1$ were the learning rates λ_{+} and λ_{-} that scaled the magnitude of individual weight changes during long-term potentiation (LTP) and depression (LTD), respectively. The element J_{ij} represented the strength and sign of the synapse from presynaptic unit j to postsynaptic unit i . The updating functions $f_{+}(J_{ij})$ and $f_{-}(J_{ij})$ scaled weight changes and implemented LTP for a positive temporal difference $\Delta t > 0$, and LTD otherwise for a negative or null temporal difference $\Delta t \leq 0$ (Gütig et al., 2003). The parameter μ of the updating functions set boundary conditions on the changes in synaptic weights. Here, a non-zero μ was used in order to prevent runaway dynamics of synaptic weights. The parameter α denoted a possible asymmetry between the scales of LTD and LTP, where $\alpha < 1$ assigned a higher magnitude for LTP whereas $\alpha > 1$ favored LTD (Gütig et al., 2003). At several cortical synapses, the width of the plasticity window for LTD is considerably wider than the width of the plasticity window for LTP (Feldman, 2000; P. J. Sjöström et al., 2001). Since LTD is widely expressed in the central nervous system (Malenka & Bear, 2004), α was set greater than 1 and the temporally asymmetric STDP learning rule was evaluated with $\tau_{-} > \tau_{+}$. Parameter values of both STP and STDP were kept fixed throughout all robotic experiments (Fig 3.5).

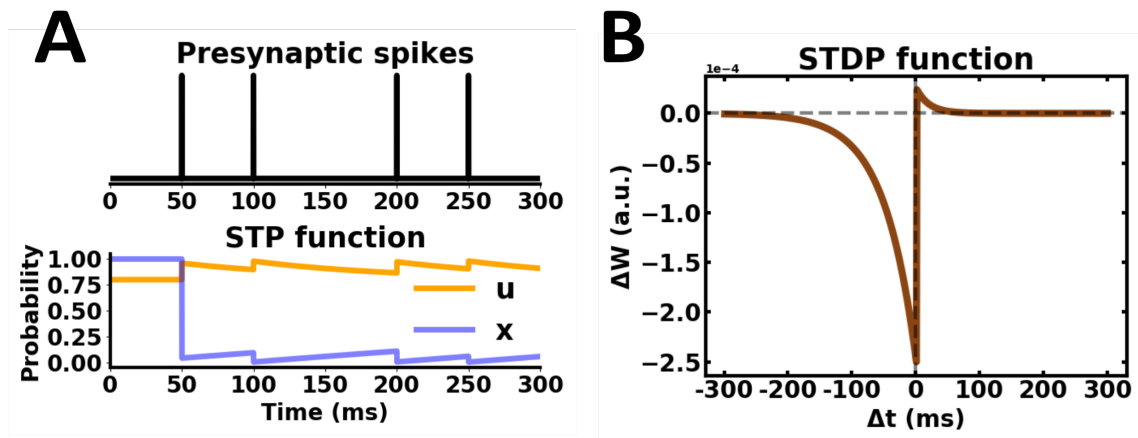


Figure 3.5: **Synaptic plasticity profile.** (A) STP dynamics from presynaptic spikes. Parameters: $\tau_f = 100$ ms, $\tau_d = 900$ ms, $U = 0.80$, $x = 1$. (B) Asymmetric STDP learning rule. Parameters: $\tau_{-} = 50$ ms, $\tau_{+} = 20$ ms, $\lambda_{-} = 25e-5$, $\lambda_{+} = 5e-5$, $\mu = 1$, $\alpha = 2$.

To combine the effects of STP and STDP, we first defined an instantaneous synaptic efficacy $u_j x_j$ with a time dependency due to STP (Eq 3.4-3.5). From this, we further extended the instantaneous efficacy by adding J_{ij} as an additional scaling factor (Klampfl & Maass, 2013; Markram, 2011; Morrison et al., 2008; Vasilaki & Giugliano, 2014), thus generating synaptic connections with efficacies that evolved according to:

$$G_i(t) = \sum_{j=1}^N J_{ij} \cdot u_j x_j \quad (3.8)$$

where G_i represented the time-dependent total synaptic efficacy of postsynaptic unit i mediated by the combined effects of STP and STDP. As a result, unit i received at each time instance, a postsynaptic current I_i^{syn} described by:

$$\frac{dI_i^{syn}}{dt} = -\frac{I_i^{syn}}{\tau_{syn}} + \sum_k G_i \delta(t - t_j^k) \quad (3.9)$$

where t_j^k represented the k th spike emitted by presynaptic unit j and G_i was the peak amplitude of the elementary postsynaptic current triggered by the activation of presynaptic unit j . The Dirac delta function δ represented the occurrence of the presynaptic spike and τ_{syn} was the recovery time constant of the postsynaptic current. In particular, when the presynaptic unit j emitted a spike, I_i^{syn} was instantaneously raised according to $I_i^{syn} \rightarrow I_i^{syn} + G_i$, and then decayed exponentially back to 0 with a time constant τ_{syn} between subsequent spikes. For an overview of the computational model proposed, a flow chart is illustrated in Fig 3.6.

Initially, units in the network were connected according to a sparse matrix structure \mathbf{C} (see Network architecture). If unit j projected to unit i , then $C_{ij} = 1$. Otherwise, in the absence of a connection, $C_{ij} = 0$. As a result of the interaction between STP and STDP, the weight matrix \mathbf{J} was engaged in structural plasticity such that novel synaptic contacts were formed over the course of the robotic experiment. Based on their difference in spike timing, uncoupled pairs of units $C_{ij} = 0$ formed a new connection $C_{ij} = 1$ for which their strength

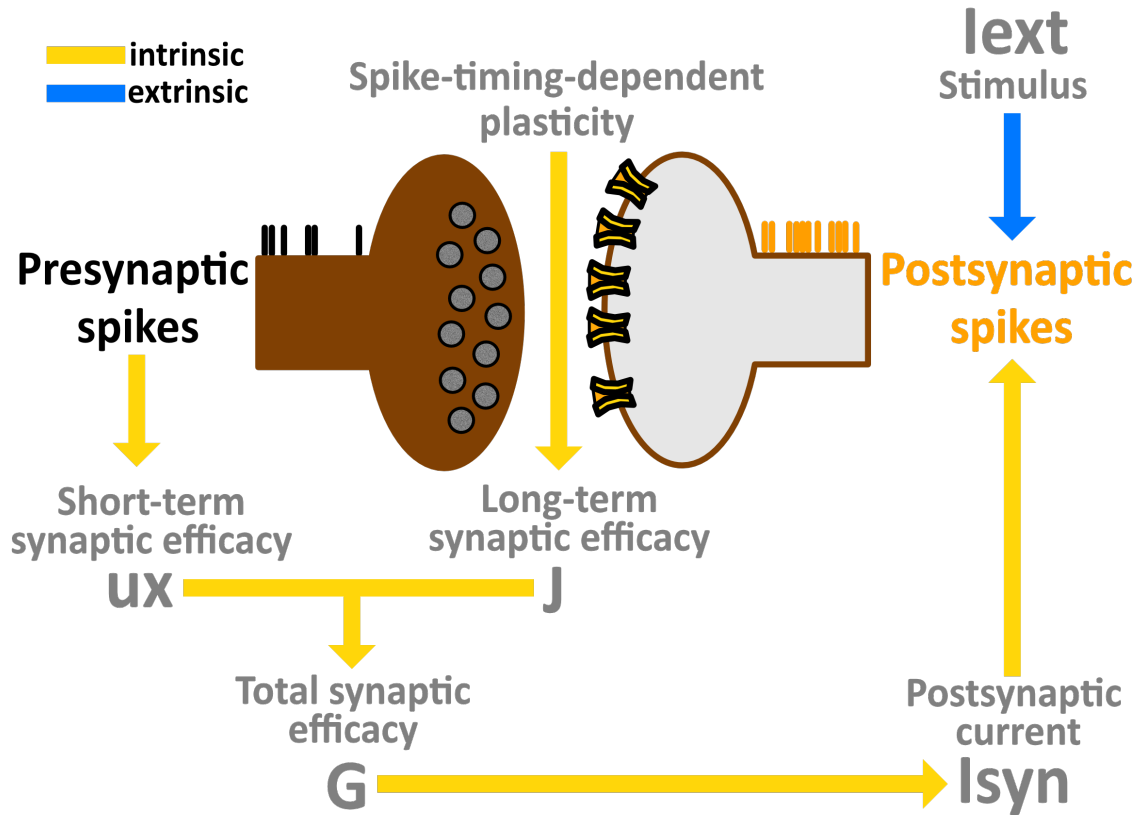


Figure 3.6: **Model flow chart.** Combination of short-term and long-term synaptic plasticity giving rise to a time-dependent total synaptic efficacy, which in turn generates a postsynaptic current that drives intrinsic network dynamics. Sensory afferents are characterized by an extrinsic input, directly affecting postsynaptic emission rates.

was scaled according to Eq 3.6-3.7. In this way, the initial sparse network connectivity evolved towards a fully connected network (S1 Fig).

3.4 Results

We report the results of six Robotic experiments. Each experiment is run via the Robotic platform, initiated using the initial Network architecture, and manipulated using the keyboard listening framework. Batches of evoked and spontaneous activity are alternated according to the Learning and recall procedure. Postsynaptic responses increase during evoked activity, and decrease during spontaneous activity (Fig 3.7). The rate of postsynaptic depolarization determines whether synaptic connections are potentiated or depressed. Overall,

the robot exhibits (1) single learning and recall, (2) incremental learning and recall, (3) task-switching, (4) resistance to interference, (5) submission to interference, and (6) resistance to distraction. A video of each robotic experiment can be found in the supplementary material (S1 Video,S2 Video,S3 Video,S4 Video,S5 Video,S6 Video).




Figure 3.7: **Input-dependent and input-disengaged synaptic plasticity differentially modulate neuronal responses.** (A) When the stimulus is ON during learning, synaptic connections are strengthened by high postsynaptic emission rates. (B) When the stimulus is OFF during recall, synaptic connections are weakened by low postsynaptic emission rates. Presynaptic activity remains relatively constant.

3.4.1 Experiment 1: single learning and recall

In the first experiment, a single learning and recall procedure is performed (Fig 3.8). Here, external inputs are presented up to a point where the network reaches global stability. During evoked activity, neuronal responses follow the spatial profile of external inputs, exhibit-

ing patterns of localized activity near the stimulus peaks (Fig 3.8A). Following stimulus offset, the network maintains local activity bumps (Fig 3.8A). These self-sustained reverberations are determined by the strength of synaptic efficacies established during evoked activity (Fig 3.8B). Noteworthy, synaptic efficacies can depend on the emission rate and spike timing of units they connect, both of which may jointly determine activity-dependent changes in synaptic plasticity (P. J. Sjöström et al., 2001). Synapses connecting pairs of units most responsive near the stimulus peaks are strengthened. The preferential bias in local reverberating activity is reflected in the robot’s wheel motor rotation speed (Fig 3.8C). As a result, movement trajectories during spontaneous activity are in the same direction as those observed during evoked activity (Fig 3.8D). Taken together, the network exhibits localized patterns of self-sustained activity with preferential responses that settle towards the position of the higher bump location, an observation directly manifested in the robot’s motor trajectory (S1 Video).

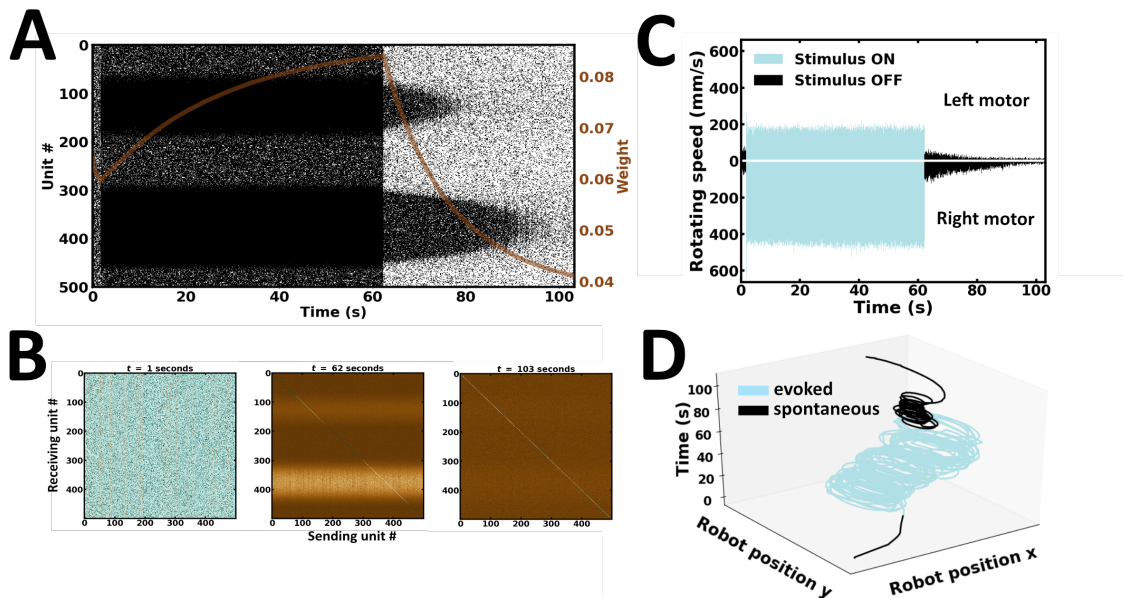


Figure 3.8: **Single learning and recall.** (A) Raster plot of network activity. Local persistent activity is preferentially biased towards a single subpopulation; brown curve, average synaptic efficacy. (B) Snapshot of synaptic weights at the end of evoked and spontaneous batches. (C) Rotating speed of left and right wheel motors as a function of time. (D) Motor trajectory of the robot in 3-D space during evoked and spontaneous batches.

3.4.2 Experiment 2: incremental learning and recall

In spite of persistent activity, networks with dynamic synapses are expected to display sensitivity to the temporal features of elicited stimuli, such as stimulus duration (Barak & Tsodyks, 2007; Egorov, Hamam, Fransén, Hasselmo, & Alonso, 2002; Egorov, Unsicker, & von Bohlen und Halbach, 2006; Fusi et al., 2000). Here, an incremental learning and recall procedure is introduced (Fig 3.9). Stimulus duration is progressively increased during successive presentations (Fig 3.9A). As a result, persistent activity progressively increases (Fig 3.9B). This gradual gain is associated with the strengthening of incoming synaptic connections (Fig 3.9C). As a result, differences in left and right wheel motor speed become progressively pronounced (Fig 3.9D). Consequently, the robot gradually refines its internal representation of sensory stimuli (Fig 3.9E). Noteworthy, autonomous motor trajectories gradually resemble more closely to the recall performance observed in Experiment 1. Taken together, the content of WM may be refined with longer stimulus presentations (S2 Video).

3.4.3 Experiment 3: task switching

In the third experiment, the robot is engaged in a task-switching procedure (Fig 3.10). Here, we examined whether the network is flexible enough to ride the wave of new input streams. To test this, stimulus intensity and duration are kept relatively constant (Fig 3.10A), whereas stimulus configuration is switched during successive evoked batches. As a result, the preferential bias in local persistent activity is switched according to the latest input configuration (Fig 3.10B). Changes in local emission rates are directly reflected in the speed of wheel motors (Fig 3.10C). As a result, the robot actively maintains an internal representation of the latest stimulus presented (Fig 3.10D). Finally, the strength of synaptic weights near the higher stimulus peak changes accordingly (Fig 3.10E). Taken together, the content of WM may be overwritten by ongoing new input streams (S3 Video).

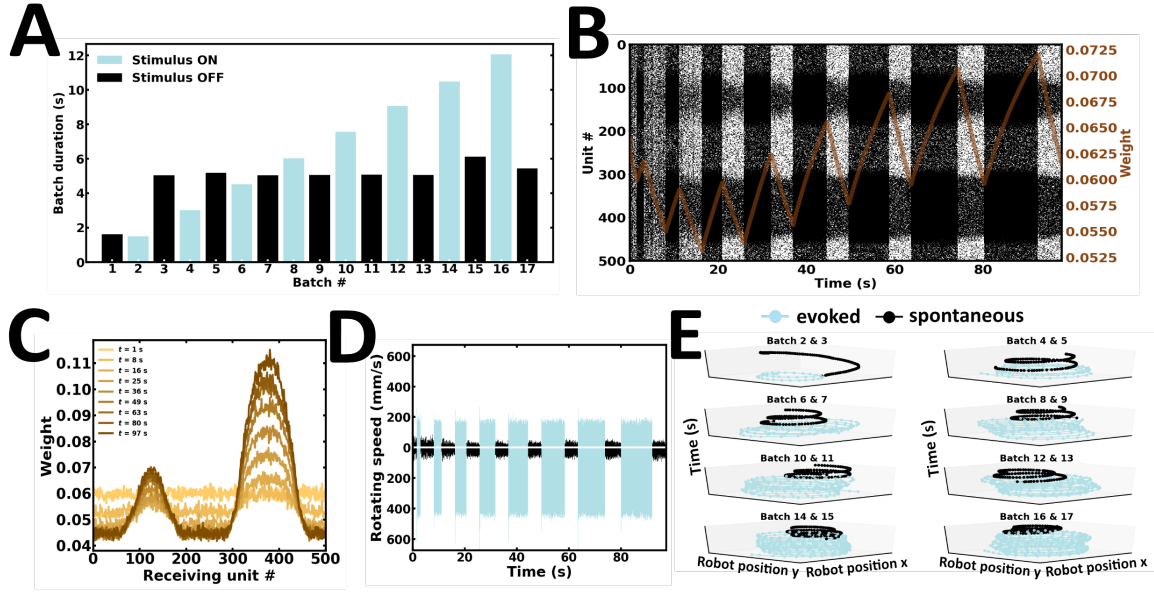


Figure 3.9: **Incremental learning and recall.** (A) Duration of evoked and spontaneous batches. (B) Raster plot of network activity. Persistent activity gain depends on stimulus duration; brown curve, average synaptic efficacy. (C) Evolution of average incoming synaptic weights at the end of spontaneous batches. (D) Rotating speed of left and right wheel motors as a function of time. (E) Motor trajectory of the robot in 3-D space during subsequent pairs of evoked and spontaneous batches.

3.4.4 Experiment 4: resisting interfering inputs

In addition to switching tasks, animals can remain resistant to unforeseen interference from external stimuli (Compte, 2000; Kilpatrick, 2018). In the fourth experiment, the robot is tested against interference (Fig 3.11). Here, stimulus configurations are the same as those in Experiment 3. However, cue stimuli are presented for a longer duration than interfering inputs (Fig 3.11A). If the robot can resist interference, it should move in the same direction before and after interference. Fig 3.11B shows that cue-guided persistent activity is maintained following interference, suggesting that the network remains resistant to brief interfering stimuli. The same wheel motor remains higher before and after interference (Fig 3.11C). Consequently, the robot moves in the same direction before and after interference (Fig 3.11D) (S4 Video). Interestingly, this type of interference is not enough to alter the acquired structure of the synaptic weight matrix (Fig 3.11E). Taken together, our fourth

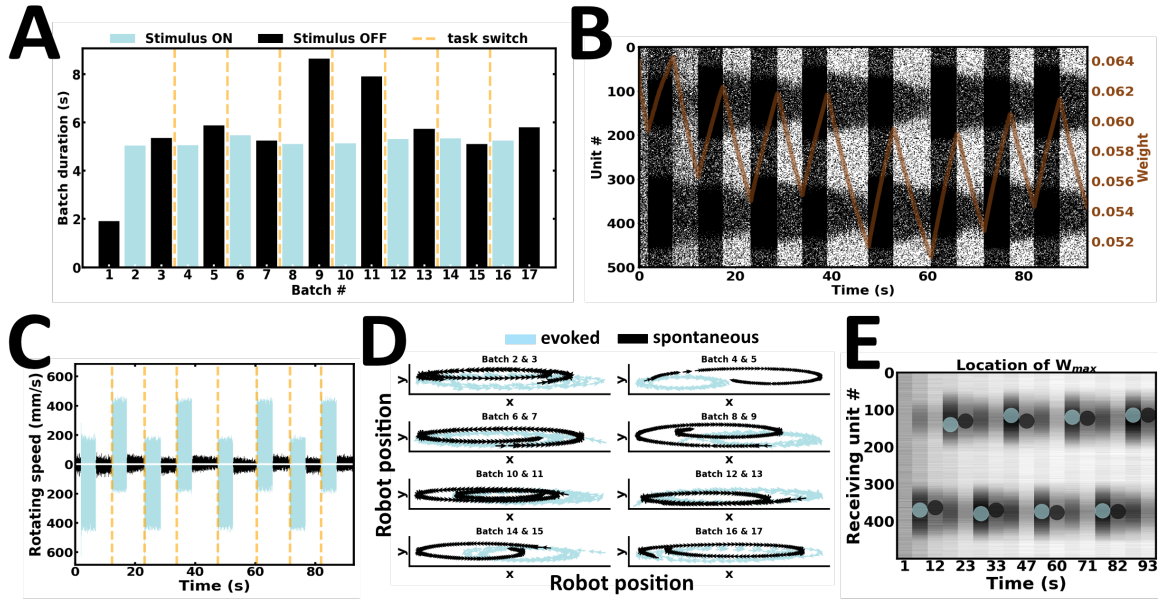


Figure 3.10: **Task switching.** (A) Duration of evoked and spontaneous batches. (B) Raster plot of network activity. Stimulus configuration shapes local persistent activity; brown curve, average synaptic efficacy. (C) Rotating speed of left and right wheel motors as a function of time. (D) Motor trajectory of the robot during subsequent pairs of evoked and spontaneous activity. (E) Spatial location of the highest average incoming synaptic weights; blue and black dot, evoked and spontaneous activity, respectively.

experiment suggests that stimuli presented for a longer duration are more likely to be remembered than other stimuli. However, if different stimuli are successively presented for similar durations, the latest stimulus may overwrite the content of the previous stimulus, as observed in Experiment 3.

3.4.5 Experiment 5: submitting to interfering inputs

Despite the ability to resist interfering inputs in Experiment 4, cue stimuli were presented for a longer duration than interfering inputs. Arguably, the robot's movement trajectory being biased towards that of the stimulus presented longer seems fully expected. However, this interpretation remains agnostic about the contribution of delay duration, which introduces context under which the WM content either resists or submits to ongoing new input streams. To explain, if delay duration is extended, the WM content would be erased. After erasure,

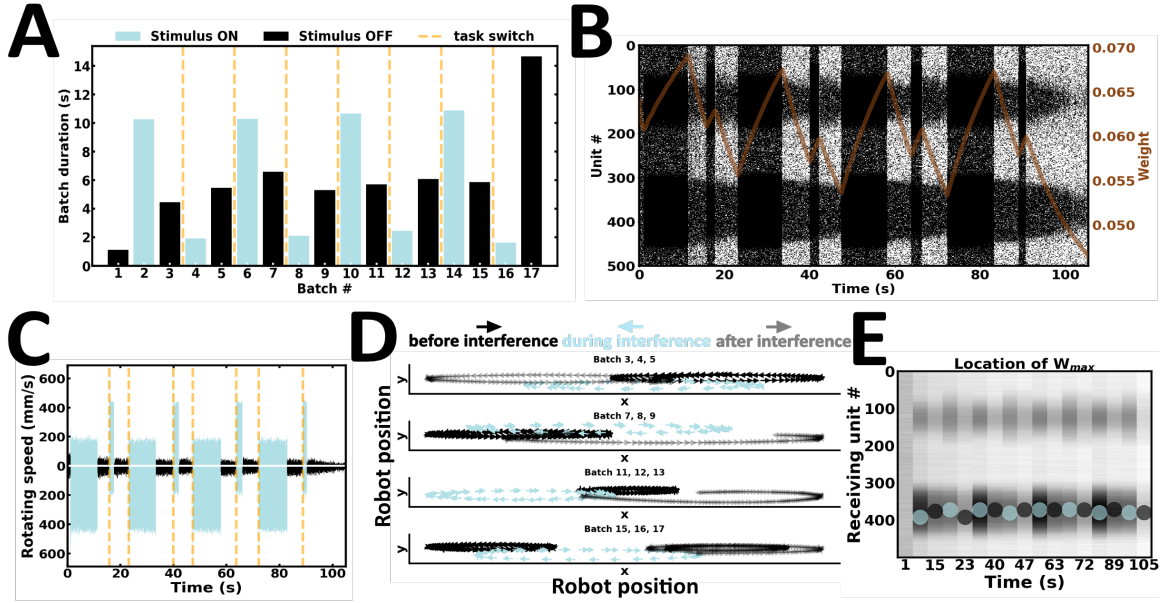


Figure 3.11: **Resisting interfering inputs.** (A) Duration of alternating batches of evoked and spontaneous activity. (B) Raster plot of network activity. Interfering inputs fail to disrupt cue-guided patterns of local persistent activity; brown curve, average synaptic efficacy. (C) Rotating speed of left and right wheel motors as a function of time. (D) Motor trajectory of the robot before, during and after the presentation of interfering inputs. (E) Spatial location of the highest average incoming synaptic weights; blue and black dot, evoked and spontaneous activity, respectively.

the robot is expected to remember the content of the latest stimulus presented, irrespective of stimulus duration. To test this hypothesis, the fifth experiment introduces a variant of Experiment 4 (Fig 3.12). Here, the duration of both stimuli is maintained, but delay activity is extended long enough to wash away the content of the cue (Fig 3.12A-B). Arguably, one would automate the paradigm using the exact same stimulus triggers as Experiment 4, and introduce an n -fold increase in delay duration. However, the latter approach is more suitable for disembodied networks, because multiple delay durations must be tested to find the minimum delay duration required to wash-away the content of the cue. To this end, multiple offline simulations runs would need to be performed, an approach our online proposal attempts to overcome. Hence, to move beyond offline experimental refinements, we turned to our online robotic implementation instead. Here, the robot accommodates interfering inputs by switching tasks, reminiscent of Experiment 3 (Fig 3.12C-D) (S5 Video). The strength

of synaptic weights near the higher stimulus peak changes according to the ongoing wave of new input streams (Fig 3.12E). Taken together, our fifth experiment suggests that delay duration introduces context, under which the WM content either resists or submits to similar interferences.

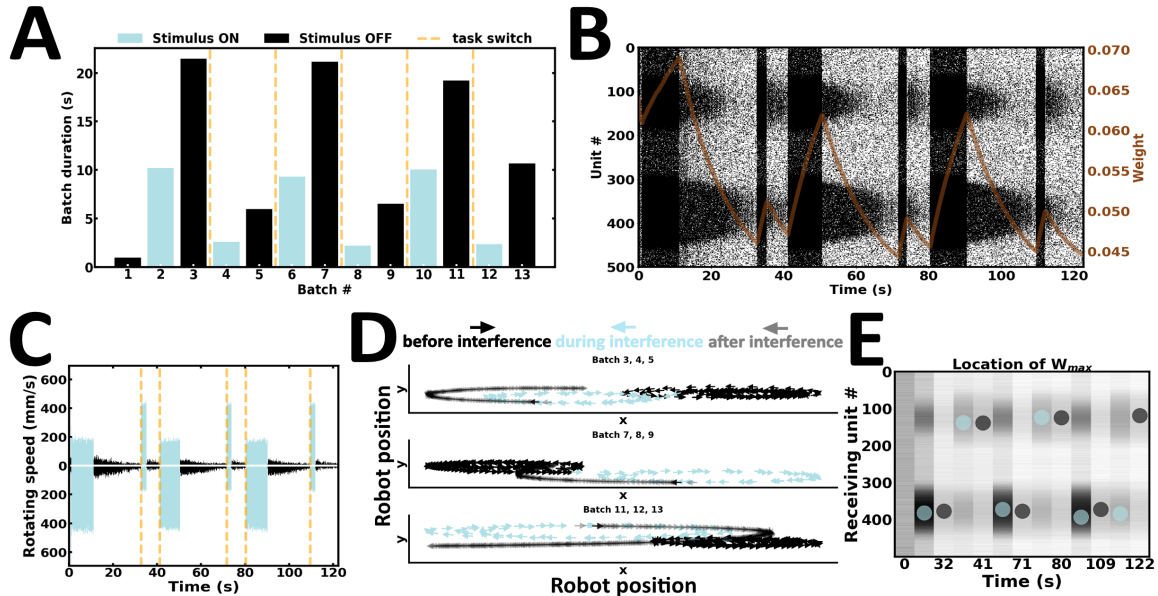


Figure 3.12: **Submitting to interfering inputs.** (A) Duration of alternating batches of evoked and spontaneous activity. (B) Raster plot of network activity. Interfering inputs disrupt cue-guided patterns of local persistent activity; brown curve, average synaptic efficacy. (C) Rotating speed of left and right wheel motors as a function of time. (D) Motor trajectory of the robot before, during and after the presentation of interfering inputs. (E) Spatial location of the highest average incoming synaptic weights; blue and black dot, evoked and spontaneous activity, respectively.

3.4.6 Experiment 6: resisting distraction inputs

In the sixth and final experiment, the robot is introduced to distraction inputs (Fig 3.13). Here, both cue and distractor durations are approximately the same (Fig 3.13A). However, distractor intensity is weaker than cue intensity. Cue-guided persistent activity is maintained following distraction inputs, suggesting that the WM content remains resistant to weaker stimuli (Fig 3.13B). The rotating speed of the same wheel motor remains higher before and after distraction (Fig 3.13C). Consequently, the robot executes movement trajectories in

the same direction before and after distraction (Fig 3.13D) (S6 Video). Synaptic weights near stimulus peaks are continuously restructured during ongoing input streams (Fig 3.13E). Following distraction inputs, units controlling the left motors are connected by the strongest synaptic efficacies. In contrast, persistent activity is globally higher among units controlling the right motors. Therefore, local information may misrepresent population level activity under certain conditions. Taken together, our sixth and final experiment suggests that the content of WM may be resistant to distraction inputs, so long as the intensity of the latter is weaker than that of the cue stimulus.

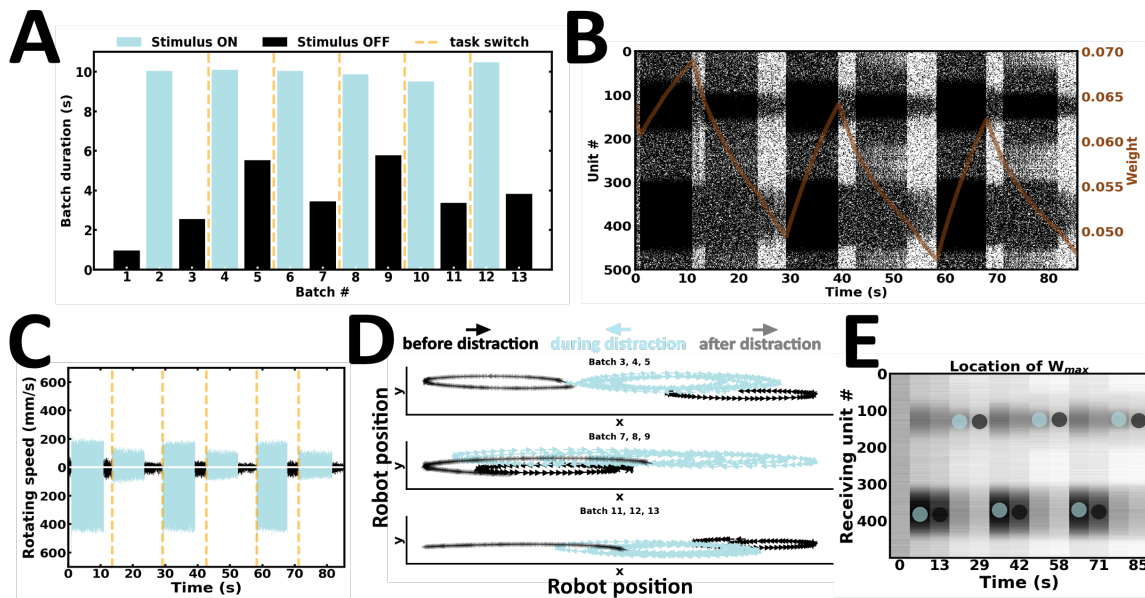


Figure 3.13: **Resisting distraction inputs.** (A) Duration of alternating batches of evoked and spontaneous activity. (B) Raster plot of network activity. Distraction inputs fail to disrupt cue-guided patterns of local persistent activity; brown curve, average synaptic efficacy. (C) Rotating speed of left and right wheel motors as a function of time. (D) Motor trajectory of the robot before, during and after the presentation of distraction inputs. (E) Spatial location of the highest average incoming synaptic weights; blue and black dot, evoked and spontaneous activity, respectively.

3.5 Discussion

3.5.1 Relation to previous work

Working memory has been studied in networks of spiking neurons connected by synapses subject to both short-term and long-term plasticity (Giudice & Mattia, 2001). Through mean field analysis, the authors showed that increased LTP should be adequately balanced by increased LTD to prevent uncontrollable excitation from destabilizing network activity. Moreover, they found that the inclusion of short-term depression can widen the range of stable ongoing neural and synaptic dynamics, ensuring online learning is kept under control. This mean field approach was then supplemented by the work of (Fusi, 2002), where the author studied how plastic synapses operate based on patterns of mean spike rates. The quantitative analysis showed that external stimuli can be embedded in the synaptic structure, controlled by the rate of postsynaptic depolarization above and below emission threshold. Furthermore, low LTP and LTD transition probabilities were shown to be instrumental for slow learning and forgetting to take place – a desirable feature in palimpsestic networks, where old memories are forgotten to make room for new ones (Amit & Fusi, 1994; Nadal et al., 1986; G. Parisi, 1986). In (Del Giudice et al., 2003), the authors outlined general requirements for stimulus driven, unsupervised formation of WM states. Their mean field analysis suggested that Hebbian rate-dependent synaptic plasticity may serve in developing synaptic structure in initially unstructured networks. The role of short-term depression was highlighted, serving as a regulatory mechanism for preventing excessive potentiation from destabilizing stimulus-selective activity.

The theoretical studies of (Amit, 1997; Del Giudice et al., 2003; Fusi, 2002; Giudice & Mattia, 2001) made a handful of predictions for which our robotic implementation supports. First, our study shows how neural activity during ongoing input streams can lead to the formation of synaptic structures, which in turn support persistent activity. Second, the asymmetric STDP learning rule suggests that LTD is strong enough to compensate for

excessive potentiation. Third, the regulatory role of short-term depression is supported by our online robotic implementation. Fourth, LTP and LTD learning rates are small enough to generate stable synaptic weight changes. Finally, transitions between LTP and LTD can be controlled by postsynaptic emission rates.

Short-term synaptic plasticity has been supported as a candidate mechanism for storing information in WM (Barbosa et al., 2020; Edin et al., 2009). Moreover, fast-expressing Hebbian synaptic plasticity may modify network connectivity momentarily enough to support information storage in WM (Fiebig, Herman, & Lansner, 2020; Fiebig & Lansner, 2017; Sandberg, Tegnér, & Lansner, 2003). Arguably, Hebbian forms of synaptic plasticity are incompatible with the flexible functionality of WM, because they induce long-lasting changes in synaptic connections that generally outlast the duration of persistent activity. However, persistent activity could be mediated by synapses that actively operate under multiple timescales (Duarte et al., 2017), carefully orchestrating information maintenance in WM. Recently, *in vivo* recordings from dentate gyrus tracked the dynamics of synaptic changes after low interference WM training in freely behaving rats (Missaire et al., 2020). While animals were still performing the task, the authors hypothesized that prior WM contents could be forgotten so as to accurately recall the storage of new contents during forthcoming trials. Interestingly, WM erasure prior to good behavioural performance was associated with the physiological induction of LTD, suggesting its implication in the loss of prior irrelevant information for the adaptive benefit of storing new ones. Importantly, lack of WM erasure corresponded to more LTP-like phenomena, which impaired behavioural performance on subsequent trials (Missaire et al., 2020). Our robotic experiments are in qualitative agreement with these findings, suggesting that LTD may erase prior WM content during extended delay durations so that forthcoming input streams can be stored without proactive interference (Underwood, 1957; Wixted, 2004). Nevertheless, further experiments will be required to disentangle the role of fast-expressing Hebbian synaptic plasticity, and its potential involvement in WM storage (Erickson, Maramara, & Lisman, 2010; Park et al., 2014; Pradier et al.,

2018).

In this study, sensory information was progressively lost during spontaneous activity (Constantinidis et al., 2001). In ANNs, this phenomenon is commonly known as catastrophic forgetting, where synaptic weights important for consolidating task A are changed to meet the requirements for task B (Kirkpatrick et al., 2017; Knoblauch, 2017; Knoblauch, Körner, Körner, & Sommer, 2014). A previous model of elastic weight consolidation (EWC) has attempted to overcome catastrophic forgetting by rendering a proportion of previous task-allocated synapses less plastic (Kirkpatrick et al., 2017). In this way, network activity reached stability over long timescales, without forgetting older tasks during sequential learning of multiple tasks. This approach may be instrumental for continual lifelong learning systems (G. I. Parisi et al., 2019). However, EWC weights are tailored for capturing memory retention, and therefore the model does not consider a forgetting component, an active arbitrator of WM capacity.

In contrast to EWC weights, synapses in our network do not hold a specific mechanism designed to protect previous knowledge from being overwritten by learning of novel information (Kirkpatrick et al., 2017; Zenke et al., 2017), nor are they imposed an active forgetting component. Rather, retention and forgetting are part of the ongoing dynamics, and there are no specialized resources administered to previous task-allocated synapses. The robotic experiments follow a continual learning scheme such that artificial stop-learning conditions are instead replaced by real-time user intervention. In light of this approach, our findings are aligned with previous theoretical studies, suggesting that the formation of WM may be subject to activity-dependent synaptic plasticity (Mongillo, Curti, Romani, & Amit, 2005). We show how this process can unfold in a robot embedded in an ecological context.

In this work, sensory information was treated as a distributed processing scheme, where any individual neuron functions as part of a larger population whose combined activity underlies the agent’s information processing capabilities (Amit, 1995). Indeed, persistent activity is thought, in many cases, to carry distributed information – sensory representation

is not a single neuron property, rather many neurons participate in the representation of an item loaded in memory (Miyashita & Chang, 1988). Here, individual units responded in an idiosyncratic manner, but collectively they generated neuronal responses that provided a centralized control over each respective wheel motor of the robot. As such, wheel motors were controlled via the collective participation of distinct neuronal subpopulations. Motor trajectories were driven by distinct localized activity bumps. Indeed, these bump states are reminiscent to those observed in discrete slot models of multi-item WM capacity (Bouchacourt & Buschman, 2019; Edin et al., 2009; Macoveanu et al., 2006; Tanaka, 2002; Wei et al., 2012), where each item is stored in a bell-shaped activity bump (Bouchacourt & Buschman, 2019; Compte, 2000; Edin et al., 2009; Eliasmith, 2005; Hansel & Mato, 2013; Itskov et al., 2011; Kilpatrick, 2018; Laing & Chow, 2001; Renart et al., 2003; Sandamirskaya, 2014; Seeholzer et al., 2019; Tanaka, 2002).

Although the cognitive significance of persistent activity remains elusive, our work suggests that local subpopulations collectively extend their role in maintaining an active memory of structured, learned information about external stimuli (Amit, 1995). Previous studies have shown that stable persistent activity could depend on a number of factors, namely the duration and intensity of the external input current applied, the synaptic weights, the level of intrinsic noise, as well as the size of the network (Laing & Chow, 2001). Indeed, controlling the parameters that define the properties of the network provides a more powerful approach rather than simply changing the input to the network (Eliasmith, 2005). Although a complete assessment of these intricacies is beyond the scope of our robotic experiments, a future study could segregate the network from the physical robot, and investigate the sensitivity of the model to a wider range of parameters.

From a phenomenological standpoint, the fade-out bumps observed during the delay period imply the gradual disappearance of sensory representation (Constantinidis et al., 2001). In agreement with this finding, behavioural studies in human subjects have shown that WM precision can degrade with the duration of the delay period (Ploner, Gaymard, Rivaud,

Agid, & Pierrot-Deseilligny, 1998). Analogously, the robot progressively fails to remember the direction of previously executed motor commands, which is reflected in the gradual disappearance of stimulus-selective activity towards baseline (Goldman, 2009). Noteworthy, persistent activity during the delay period is not only stimulus-dependent, but also likely to be mediated by time-varying fluctuations in neuronal activity which render the internal representation dynamic (Brody, Romo, & Kepecs, 2003). Indeed, WM experiments have shown a set of diverse neuronal response patterns during delay period activity (Chafee & Goldman-Rakic, 1998; Takeda & Funahashi, 2002). These individual response profiles range from sustained tonic activity to gradual increase/decrease in firing rate fluctuations. The presence of heterogeneity in the response of individual neurons likely provides a meaningful computational role for higher order function (Berberian, MacPherson, et al., 2017; Drover, 2014; Renart et al., 2003; Singh, 2006). Despite the presence of such diversity, our findings solely accounted for the progressive decrease in the neural response profiles (Chafee & Goldman-Rakic, 1998; Takeda & Funahashi, 2002). Nevertheless, persistent activity was maintained over a time scale of seconds (Baddeley, 1992; P. Goldman-Rakic, 1988). This observation has been a hallmark feature of the prefrontal cortex involved in WM, where the regulation of internally-guided decisions can be executed in the absence of external stimuli (P. S. Goldman-Rakic, 1987). From a neurorobotic standpoint, persistent activity is advantageous for artificial agents exposed to momentary hardware failure or sensor perturbation. Instead of relying purely on external stimuli, the agent internalizes recent events and carries them over towards a foreseeable future. In the absence persistent activity, perceptual uncertainty would prevail, preventing an accurate execution of internally-guided behaviour.

3.5.2 Limitations

In this study, executed motor trajectories aimed to capture the internal representation of sensory afferents stored in WM. However, using motor executions to characterize sensory representations is an abstraction that moves away from the kinds of representations memorized

by animals during WM experiments. Therefore, the characterization of sensory representations via wheel motor executions creates a partial correspondence to the hypothesized contribution of WM maintenance proposed in experimental studies. Consequently, our robotic experiments create an imbalance in abstraction which may arguably lead to a loss of biological relevance (Webb, 2001). It is important to note however that our work does not propose input-disengaged motor executions to be mediated by persistent activity. Rather, our study shows how the physical robot can be used as a companion approach for concretizing the representation of sensory stimuli in an embodied network. In particular, autonomous motor executions during spontaneous activity is used purely for the purpose of studying how the content of WM interacts with ongoing sensory afferents. Importantly, the inclusion of motor primitives turns out to be critical for providing behaviourally meaningful content to sensory primitives, because the maintenance of an ongoing movement holds a property in the action that accurately reflects the property of the sensory stimulus (Choe et al., 2007). Beyond the sheer presence of neural activity and synaptic transmission, the motor primitives of the robot were used as a method for displaying persistent neural activity “in action”. The robotic implementation was therefore used in order to ensure a one-to-one correspondence between network activity and executed motor trajectory. The extensive matching and complementarity between their properties constrains the scope of multifaceted interpretations in the face of non-stationary environments. Moreover, it provides an opportunity for their coordinated changes to ultimately result in the emergence of adaptive behaviour (Chiel & Beer, 1997). Overall, our robotic experiments aimed to exemplify the importance of embodiment as a tool for studying the formation of WM.

During the experiments, structural plasticity rendered an unrestricted establishment of novel synaptic connections within and between excitatory and inhibitory populations. In this context, excitatory units maintained their initial strict excitatory influences over the course of each experiment, because the multiplicative STDP rule linearly attenuates excitatory efficacies as they reach the boundaries $[0, 1]$ (Gütig et al., 2003). As such, excitatory outgoing

connections evolved within boundaries (S2 Fig – left column). In contrast, outgoing connections from inhibitory units switched their initial strict inhibitory influences by suddenly forming excitatory contacts with their novel postsynaptic targets, a biologically implausible scenario which led to the violation of Dale’s law (S2 Fig – left column). To circumvent this issue, some networks have held fixed inhibitory efficacies (Clopath, Büsing, Vasilaki, & Gerstner, 2010; Gütig et al., 2003; Hosaka, Araki, & Ikeguchi, 2008; Sadeh, Clopath, & Rotter, 2015). Others have attempted to obey Dale’s principle by suddenly freezing synaptic connections if the synaptic update reversed the sign of excitatory or inhibitory influences (Kim & Chow, 2018). In subsequent trainings, synapses attempting to change their signs were excluded and therefore prevented from exhibiting activity-dependent changes. Although the latter constraint is adequately tailored for the conservation of Dale’s principle, an artificial stop-learning condition is imposed, which hinders the ability to learn new memories (Knoblauch et al., 2014). Nevertheless, model networks with positive and negative synaptic weights can be functionally equivalent to those carrying positive weights only (Parisien, Anderson, & Eliasmith, 2008). Finally, structural plasticity also led to the establishment of self-connections in our network, which were absent during network initialization (S2 Fig – right column).

In this work, inputs were represented as bimodal Gaussian mixtures (Sensory inputs). The motivation for using simplified sensory afferents partly stems from classic studies of visual WM, where subjects are introduced the oculomotor delayed response task (Funahashi et al., 1989; Takeda & Funahashi, 2002). During the cue period of the task, stimuli are presented at various possible locations around a fixation point (Constantinidis et al., 2018). These stimuli may possess simple geometric shapes differing in luminance or color (Constantinidis et al., 2001; Hoshi, Shima, & Tanji, 1998). Noteworthily, a substantial proportion of neurons in the dorsal region of the prefrontal cortex exhibit spatially tuned elevated persistent firing across the delay period (Meyer, Qi, Stanford, & Constantinidis, 2011). However, the reverberating activity of prefrontal neurons can extend beyond the spatial domain, to-

wards physical characteristics of objects (e.g. shape, color and intensity) and their respective identity (Freedman, 2001; E. K. Miller et al., 1996; Ó Scalaidhe, 1997; Ó Scalaidhe, Wilson, & Goldman-Rakic, 1999; Rainer, Asaad, & Miller, 1998; Rainer & Miller, 2000; S. C. Rao, 1997; F. Wilson, Ó Scalaidhe, & Goldman-Rakic, 1993). Although the bimodal mixture of Gaussians used in this study carries selectivity for spatial location and intensity, it remains agnostic to shape information, thus leaving out the multimodal characteristics of real world stimuli and the complex statistical structure they carry (Brader, Senn, & Fusi, 2007). Nevertheless, it is possible to train the network on an arbitrary set of stimuli, such that delay activity is selective to object shape, as observed in prefrontal neurons (S. C. Rao, 1997). An example is illustrated in the supplementary figure (S3 Fig).

3.5.3 Practical implications

In this study, there are practical reasons for considering user keypresses instead of setting up timers to generate consistent experimental triggers. First, it is noteworthy to mention that the robot was tested in a confined spatial setting. Perhaps less obvious, devising a fully automated task in a restricted area runs the risk of having the robot run into collisions. Indeed, collisions may be avoided by running experiments in an unrestricted area or overcome by incorporating system feedback. However, our robot is designed to have low neglect tolerance, meaning it requires continuous supervisory control (Goodrich & Schultz, 2007). Agents with low neglect tolerance are not given the luxury of behaving in the context of automated task designs. Our goal is not to design automated paradigms, because the network is performed online and the robot is placed in an ecological context (Amershi, Cakmak, Knox, & Kulesza, 2014). The study therefore focused on the interactive component between the user and the robot, such that the user received continuous behavioural feedback from the robot, and pressed a key at any moment, including when the agent ran the risk of collision. Real-time interventions as such created inconsistent experimental triggers.

The keyboard listening framework developed here is suitable for online learning of rapidly

varying input streams, a phenomenon known as concept drift (Lobo et al., 2020). In particular, the framework is designed for choosing a specific input configuration from a repertoire of configurations, such that it can introduce novel motor trajectories via rapid and flexible task-switching. Future studies could use our keyboard listening framework and expand the repertoire of input configurations. This approach has been proposed as an important design principle in developmental robotics (Weng, 2001, 2004). Since the programmer who develops the keyboard listening framework does not know what tasks the future users will teach the robot, the framework gives the freedom to choose a sequence of tasks that are unknown at the time of programming (Weng, 2001). Hence, in an attempt to further explore how the model behaves with disparate problems in the same network, programmers may design a wider range of stimulus duration, intensity, specificity and bump location. To this end, we expect that the computations resulting from a wider range of input configurations and durations may not only reproduce the behavioural results presented here, but also further extend the range of possible behavioural trajectories. For example, introducing a bimodal Gaussian mixture with localized amplitude peaks of the same intensity should remove the preferential bias from the persistent activity bumps observed here, making local reverberating activity between the two subpopulations indistinguishable. Consequently, the robot would autonomously navigate forward following the removal of the “unbiased” stimuli. This new behavioural trajectory would directly result from the coupled interaction between neural activity and motor output, both of which are under the influence of ongoing new input streams presented to the same underlying network (Chiel & Beer, 1997).

Robotic behaviour can be used a method for making an inference about the underlying network computation. Adaptive behaviour is directly observable, because the online operation is tailored for piece-by-piece accommodation of ongoing input streams. In this way, new items can be actively loaded in WM and therefore the agent can be trained to maintain novel sensory representations on-the-fly (Weng, 2001). Moreover, the keypresses used here minimize the efforts required to make arbitrary changes in the elicited stimuli (Najar, Sigaud, &

Chetouani, 2019). Interactions as such provide external observers some time to interpret the behaviour of the robot, and form a real-time hypothesis of its forthcoming behaviour. Under the umbrella of this interaction, the robot can actively augment, and be shaped by the user. Taken together, our study proposes that traditional WM protocols may be complemented by placing biological agents in an ecological context, under which the content of WM interacts with ongoing input streams.

Conclusion

In this work, we have devised a network of spiking neurons with dynamic synapses. As a companion approach to the computational model, the network was embedded in a physical robot. The artificial agent was used as a source of exploration under which neuronal computation was directly linked to motor output – serving as an illustration of WM capacity under non-stationary input streams. This methodological approach was accompanied by a keyboard listening framework, which helped us delineate how the content of WM is (1) refined, (2) overwritten or (3) resisted by changes in the duration and configuration of incoming stimuli. The model we propose is capable of continuous online learning, and the embodied network adapts to changing environmental contexts, all of which are done so under the basis of local information reliance. We show how these local changes can be accounted for by short-term and long-term changes in synaptic plasticity. Although we believe that carefully designed experiments will continue to lay the foundation for our understanding of higher-order function, we hope that embodied networks will add a supporting layer to the valuable theoretical insights gained from computational models of pure simulation.

3.6 Supporting information

S1 Fig. Temporal evolution of structural plasticity. (S1 Fig link).

S2 Fig. Temporal evolution of synaptic weights. (S2 Fig link).

S3 Fig. Persistent activity based on object shape and location. (S3 Fig link).

S1 Video. Robotic experiment 1. Single learning and recall. (S1 Video link).

S2 Video. Robotic experiment 2. Incremental learning and recall. (S2 Video link).

S3 Video. Robotic experiment 3. Task-switching. (S3 Video link).

S4 Video. Robotic experiment 4. Resistance to interference. (S4 Video link).

S5 Video. Robotic experiment 5. Submission to interference. (S5 Video link).

S6 Video. Robotic experiment 6. Resistance to distraction. (S6 Video link).

Discussion

Neuroscience-inspired robotics

Over the years, there has been a growing commitment to the idea that one would benefit from understanding the brain in the context of an embodied system that routinely interacts with the environment (A. K. Engel, Maye, Kurthen, & König, 2013; M. Wilson, 2002). Embodied views of cognition suggest that a tightly coupled interaction between the nervous system, the body and the environment has profound implications in generating meaningful activity patterns for behaviour (Chiel & Beer, 1997; Pfeifer et al., 2007). Here, patterns of neural activity are not investigated in an isolated nervous system, but rather through the generation of both internally and externally guided movement executions resulting from the agent-environment interaction. Notably, this agent-environment nexus imposes both constraints and opportunities for neural control (Beer, 1998, 2014; Chiel & Beer, 1997). For example, the biomechanics of the body, the capabilities of sensors and effectors, as well as the repertoire of movement executions impose constraints because the agent cannot produce actions beyond those imposed by its physical capabilities (Iida & Nurzaman, 2016). However, constrained actuation poses opportunities for simplifying complex control problems, by performing a readout of the embodied network from a large repertoire of idiosyncratic neural response profiles, thus centralizing the control of movement executions that unfold overtime (Turkel, Port, & van Gelder, 1998).

More recently, the field of neurorobotics has emerged as a scientific discipline that carries

physical devices over virtual or real-world environments, and operates them based on known principles of the nervous system (Krichmar, 2008; Oberts, Sanders, Knoll, & Gewaltig, 2016). Standing along the edge of neuroscience and robotics, the interdisciplinary field considers the nervous system as embodied, whereby the body is in turn embedded in the environment (Chiel & Beer, 1997; Krichmar, 2018). In general, neurorobots are equipped with sensors and actuators. Endowed with a defined body morphology from which they perform a repertoire of movements, neurobotic agents form a reciprocal interaction with the environment containing sensory objects and/or events. Perhaps less obvious yet important, embodied agents perform actions that affect the environment, which in turn influences the nature of inputs they receive. In this way, the behaving agent is continuously influenced by forthcoming input streams imposed on its sensory apparatus (Nolfi & Parisi, 1993; Schlesinger, 2003; Sporns, 2007; Sporns & Pegors, 2004).

Considered in isolation, the nervous system is a highly complex system of interconnected nerve cells organized in a rich hierarchical structure. However, when the nervous system is considered embodied and coupled with the environment, the agent is suddenly constrained by the morphology of its body as well as the ecological niche in which it operates (Iida & Nurzaman, 2016; Pfeifer et al., 2007; Webb, 2001). This agent-environment coupling differs from traditional approaches, where the brain is studied in isolation (Bickle, 2016). Instead, neurobotic systems are physically instantiated such that the action-guided system is treated as fine-grained as the investigations at the neural level (Krakauer et al., 2017). To this end, neurorobotics can help build methodological tools for studying reciprocal tendencies between the nervous system, the body and the environment (Kasabov, 2007).

Embodied systems provide an opportunity to transition from passive and immobile physical models towards active and mobile systems (Gravish & Lauder, 2018). From this viewpoint, the embodied agent is given a wide range of affordances, a notion introduced by ecological psychologist J. J. Gibson (1904 - 1979) (Gibson, 1978). Affordances take on the perspective of the agent and ask: what are the possible ways in which I can use, intervene

and act in the physical world? While closely tied to context-specific sampling of environmental features, affordances are profoundly operated by the physical structure of the agent and its nervous system capacity, which ultimately hold action-related properties grounded in the physical environment (Clark, 1995, 1999). Noteworthy however, building systems capable of unconstrained interactions with the real-world has been more difficult than anticipated, because of the sudden explosion of high-dimensional affordances under real-world settings (Brooks, 1991).

To alleviate the problem of high-dimensional affordances, one may restrict the range of admissible behavioural trajectories the agent performs (Beer, 1998; Harvey, 2019). Despite reducing the repertoire of robotic action opportunities, control-based internal representations can rise to the forefront (Clark, 1995). Here, internal representations are both descriptive and directive, in that they serve to describe how the world is (i.e. encoding), and in parallel, provide a response or action (i.e. control). In this way, there is no further computational effort required to yield the prescribed action, because the internal representation does the acting as it stands in relation to an environmental event – an idea coined as “pushmi-pullyu representations” (Millikan, 1995). The direct dedication of inner states by control of actions specify the agent’s involvement whose degree depends upon the world being in a certain way (Clark, 1995). Taken together, endowing agents with low-dimensional affordances may narrow the range of admissible behavioural trajectories, yet maintain an interpretable interaction between brain, body and environment – mutually and simultaneously influencing change (Van Gelder, 1995).

Human-robot interaction

Human-robot interaction (HRI) involves a symbiotic collaboration between humans and robots. Here, the human and the robot play an active role in the interaction. For example, imagine teaching a child how to ride a bicycle by gripping from the handlebar to guide her

direction. Through reciprocal and tightly coupled interaction, the child provides the teacher feedback primarily through overt actions. In doing so, the child helps the teacher maintain an accurate mental representation of its ability to leverage from instructions (Chernova & Thomaz, 2014). The interpretation behind the child’s performed actions encourages the teacher to make the necessary adjustments in its teaching strategy (e.g. I should grip from the stem instead of the handlebar to give the child more surface freedom on the handlebar; maybe this will refine her balance on the bicycle). Importantly, this collaborative approach moves away from the presupposition that the child is the only active learner in the process. Rather, the human teacher is also an active learner, because it observes the child performing and thus learns to refine its teaching strategy to better serve the task requirement. In this context, learning is a bidirectional collaboration that emphasizes the importance of mutual commitment in maintaining an open communication channel (Chernova & Thomaz, 2014).

Analogous to the child, the robot is considered to play a utilitarian role for humans in the context of HRIs. First, humans may reduce their cognitive workload by offloading their strategies onto the robot with which they interact. Mental offloading may be important, because when humans are confronted with time-pressured constraints of real-time intervention (e.g. Tetris game), their performance can suffer due to information overload, a notion commonly known as *representational bottleneck* (M. Wilson, 2002). To remedy information overload during immediate online interactions, humans can use mobile robots as ideal companions for dynamically offloading information onto them. Second, robots are useful tools for augmenting humans, because they can help them learn to refine their decisions to be more aligned with their underlying preferences (Russell, 2019). For example, mobile robots can be used as tools for engineering choice alternatives over which humans can formulate or refine their decisions. Furthermore, robots can help humans better recognize the consequences behind their decisions, because they are by HRI design principles, continuously providing action-guided feedback. Third, robotic behaviour can give humans hints about how their internal state dynamics are evolving. The latter is particularly applicable

for control-based methods that treat internal representations primarily as tools for action (Clark, 1995). Mobile robots can help humans formulate real-time hypotheses about their internal governing structures. From this approach, it becomes possible to revert back to the details of their “internal state dynamics”, helping humans formulate new scientific questions about brain-environment alignment.

There are a variety of existing methods by which humans can communicate with robots (Najar & Chetouani, 2020). Methods of human instruction range from verbal communication to non-verbal gestures, as well as haptic signals via push buttons, among others (Chernova & Thomaz, 2014). As such, robotic agents can learn from human instructions via teleoperation procedures. Such external events call an event handler so as to run a particular procedure in a computer (Ben-Ari & Mondada, 2018). When used in the right context, they can minimize the cognitive load imposed on the human teacher, particularly under experiments where the robot is mobile, and exhibits *low neglect tolerance* – requires the undivided attention of the human teacher (Crandall, Goodrich, Olsen, & Nielsen, 2005). Amongst the remote operations mentioned above, push buttons are particularly favourable when prioritizing the accuracy of instruction signals, because they provide reliable input-output responses (Knox & Stone, 2009; Knox et al., 2013). For example, keypresses from a computer ensure that the feedback signal is an accurate reflection of the teacher’s desirability of the agent’s targeted behaviour (Knox et al., 2013). Furthermore, keypresses vastly improve the fluidity of information exchange, particularly for scenarios where the robot is subject to real-time human intervention. Finally, keypresses minimize the efforts required to train non-expert human teachers willing to interact with the robot (Knox et al., 2013).

Keeping humans in-the-loop during human-robot interactions holds practical implications and challenges. First, humans can interact with robots for the purpose of aligning them with human objectives. Second, placing robots in an ecological context means ultimately exposing them to the variability that humans have to offer. This ecological noise source exposes artificial agents to the common problem of complexity matching in modern AI systems,

where highly variable and abrupt changes in the environment create misalignments between the inner complexity of the agent and the outer complexity of the environment (Kello et al., 2010). It thus follows that when robots interact with humans, they must display flexibility and adaptation in response to human inconsistencies. Failing to align with environmental intricacies exposes their limitations, such as lack of synchronization and context-dependent organization. Such event-related behaviours may be important for better understanding and thus refining their internal governing structures, such that the latter better matches its environmental assembly.

Beyond the practical reasons mentioned above, there are ethical reasons why keeping humans in-the-loop may be important. In recent years, artificial intelligence (AI) researchers have raised safety concerns about AI and its long-term unintended existential risks (Bostrom, 2014; Russell, 2019; Tegmark, 2017). In the quest for understanding the organizing principles of autonomous systems, the value-alignment problem draws our attention towards asking what initial goal and objective should artificial intelligent systems be supplied with. In particular, a single-minded pursuit of an objective that is not aligned with the objectives of the human race runs the risk of rendering unintended consequences over the long term. A classic example is the paperclip maximizer – endowing an AI with the goal of maximizing the number of paperclips (Yudkowsky, 2009). After producing above and beyond the number of paperclips required, the AI system is gestured from its creators to reduce the number of paperclips. However, since this would lead to fewer paperclips – the AI agent would end up misinterpreting its creators as obstacles for its own goal. Consequently, the AI system would sidestep all moral imperatives in its quest to maximize the number of paperclips. There are many systems in the world where the objective has been fixed prematurely, after which the machine is decoupled from those it is supposed to be serving (e.g. reinforcement learning (RL) algorithms operated behind content recommendation platforms). Is there a way to remedy the problem of having machines that optimize fixed objectives?

Humble machine intelligence

In order to solve the value-alignment problem, it may be instrumental to first solve the control problem. That is, to devise machines that are uncertain about the objective of what it is they are trying to maximize (e.g. reward) (Russell, 2019). In other words, we can design machines that are intelligent enough to know that they do not know what it is that they are supposed to be maximizing. In his book *Human Compatible*, Stuart Russell calls these systems “humble machines” (Russell, 2019). Rather than cutting the strings and letting the robot go by itself in pursuit of a fixed objective, we can reconsider the methods by which we design the internal governing structure of AI systems. For instance, the pursuit of decoupled machines with objectives that are misaligned with human objectives may be remedied by principles of inverse reinforcement learning (IRL) (Russell, 2019). Here, the term “inverse” is used, because traditional RL agents are first rewarded and then behave accordingly, whereas inverse RL agents first behave and then learn the rewards. Instead of endowing the agent with a fixed reward function where certainty is prescribed to it, the IRL agent initially has a vague estimation of the true reward function. Importantly, the precision of the estimated reward function is refined by the pursuit of human objectives. As such, the technique assumes that a human teacher knows the goals of the task, and therefore provides feedback that points the agent towards closer estimates of the true reward function. Throughout the process, the agent considers evolving human choices rather than engaging in a single-minded pursuit of a fixed objective.

What are the requirements for designing humble machine intelligence? In his 1947 lecture *Machine and Organism*, French physician and philosopher of science Georges Canguilhem (1904-1995) argued that the teleological orientation of a machine should always be supplied from the external world. In other words, the purpose of the machine should not solely be intrinsic to the machine’s own internal organization. Unlike the organism, the machine’s privileges are predetermined by a precisely defined external order towards which it remains submissive. In this way, there is no clear boundary that separates the machine from the

external world, because a machine governed by internally and externally given criteria is inevitably submissive to an externally given order (Canguilhem, 2008). Canguilhem stated that if life is “an attempt from all directions”, then a machine that follows such attempts must affirm the norms of predictability (Bates, 2014; Canguilhem, 2008). Taking away rudimentary privileges of internally guided decisions in machines raises noteworthy challenges for modern AI systems, which are heavily governed by their own internal dynamics with no requisite considerations of evolving human objectives. Taken together, it is not the primary objective of humble machines to make a prediction of an upcoming reward, but rather to adapt to evolving human preferences.

Concluding statement

In the first chapter, we studied the relationship between short-term and long-term synaptic plasticity in a microcircuit motif of two interconnected units. The goal of the microcircuit was to discriminate the spatial position of a moving stimulus. Overall, this published work explored the contribution of short-term plasticity in shaping the long-term strength and stability of synaptic connections in the microcircuit.

In the second chapter, we studied direction-selectivity in a microcircuit motif of six units, connected by synapses subject to short-term plasticity. In doing so, we embedded the microcircuit in a Raspberry Pi microcontroller, and used a camera to process moving visual images displayed on a computer monitor. Neuronal responses from the microcircuit were compared against the response of direction-selective cells in primary visual cortex (V1) of anesthetized macaque monkeys. Overall, this published study attempted to bridge together ideas from neuroscience and robotics.

In the third and final chapter, we studied working memory in a spiking network of 500 units connected by synapses subject to both short-term and long-term synaptic plasticity. The activity of the network was directly linked to the motor output of a mobile robot. We also developed a keyboard listening framework for online mobile robot control. Using the framework, we used push-buttons to control inputs injected into the network. Overall, this published study presented a methodological perspective on how we can embody formal models from computational neuroscience, and place them in an ecological context to study their internal dynamics.

Appendix A

Analytical solutions of the passive integrate-and-fire model for a constant current injection

A.1 Membrane potential evolution in the absence of threshold and reset rule for action potential generation

The passive integrate-and-fire model is described by the following equation:

$$\tau_m \frac{dV}{dt} = E_L - V + R_m I \quad (\text{A.1})$$

where τ_m is the membrane time constant, E_L is the resting membrane potential, R_m is the membrane resistance and I is a constant current injection. When I is constant, the membrane potential V can be determined by integrating the equation of the passive integrate-and-fire model:

$$\frac{dV}{dt} = \frac{E_L - V + R_m I}{\tau_m} \quad (\text{A.2})$$

$$\frac{dV}{-E_L + V - R_m I} = \frac{-dt}{\tau_m} \quad (\text{A.3})$$

Since dV is on one side of the equation and dt on the other, we can integrate both sides to find the evolution of the membrane potential from time t_0 to t :

$$\int_{t_0}^t \frac{dV}{-E_L + V(t_0) - R_m I} = \int_{t_0}^t \frac{-dt}{\tau_m} \quad (\text{A.4})$$

$$\int_{t_0}^t \frac{1}{-E_L + V(t_0) - R_m I} dV = \int_{t_0}^t \frac{-dt}{\tau_m} \quad (\text{A.5})$$

$$\ln|-E_L + V(t_0) - R_m I|_{t_0}^t = \frac{-t}{\tau_m} \Big|_{t_0}^t \quad (\text{A.6})$$

$$\ln|-E_L + V(t) - R_m I| - \ln|-E_L + V(t_0) - R_m I| = \frac{-t}{\tau_m} - \frac{-t_0}{\tau_m} \quad (\text{A.7})$$

$$\ln \frac{|-E_L + V(t) - R_m I|}{|-E_L + V(t_0) - R_m I|} = \frac{-t}{\tau_m} - \frac{-t_0}{\tau_m} \quad (\text{A.8})$$

$$\frac{|-E_L + V(t) - R_m I|}{|-E_L + V(t_0) - R_m I|} = \exp\left(\frac{-t}{\tau_m} - \frac{-t_0}{\tau_m}\right) \quad (\text{A.9})$$

$$|-E_L + V(t) - R_m I| = |-E_L + V(t_0) - R_m I| \exp\left(\frac{-t}{\tau_m} - \frac{-t_0}{\tau_m}\right) \quad (\text{A.10})$$

$$\begin{cases} |-E_L + V(t_0) - R_m I|, & \text{if } -E_L + V(t) - R_m I > 0 \\ -|-E_L + V(t_0) - R_m I|, & \text{if } -E_L + V(t) - R_m I < 0 \end{cases} \quad (\text{A.11})$$

$$-E_L + V(t) - R_m I = (-E_L + V(t_0) - R_m I) \exp\left(\frac{-t}{\tau_m} - \frac{-t_0}{\tau_m}\right) \quad (\text{A.12})$$

Since $t_0 = 0$:

$$\boxed{V(t) = E_L + R_m I + (-E_L + V(0) - R_m I) \exp\left(\frac{-t}{\tau_m}\right)} \quad (\text{A.13})$$

Suppose that at t_0 , the neuron emits a spike and is thus at the reset potential $V(0) = V_{reset}$. From here on, if the threshold and reset rule for spike generation is not applied, the membrane potential would evolve according to Fig A.1.

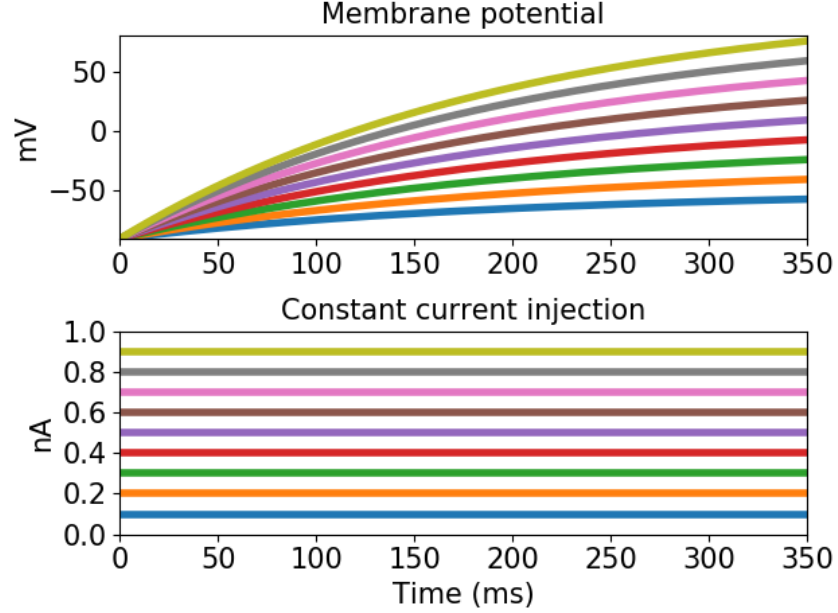


Figure A.1: **Evolution of membrane potential across time in response to constant current injection lasting for 350ms.** Colours in top and bottom panels are associated. Model parameters: $V_{reset} = -90\text{mV}$, $E_L = -70\text{mV}$, $R_m = 200 \text{ M}\Omega$, $\tau_m=200\text{ms}$

A.2 Interspike interval and spike count prediction

Suppose a spike will occur when the membrane potential reaches the spike threshold, that is, at time $t = t_{isi}$. This scenario can be described by the following equation:

$$V(t_{isi}) = V_{th} = E_L + R_m I + (-E_L + V_{reset} - R_m I) \exp\left(\frac{-t_{isi}}{\tau_m}\right) \quad (\text{A.14})$$

By solving for t_{isi} , we can find the time of the next spike and therefore determine the

interspike interval (ISI) for any arbitrary constant current injection I :

$$V_{th} - E_L - R_m I = (-E_L + V_{reset} - R_m I) \exp\left(\frac{-t_{isi}}{\tau_m}\right) \quad (\text{A.15})$$

$$\frac{V_{th} - E_L - R_m I}{-E_L + V_{reset} - R_m I} = \exp\left(\frac{-t_{isi}}{\tau_m}\right) \quad (\text{A.16})$$

$$\ln\left(\frac{V_{th} - E_L - R_m I}{-E_L + V_{reset} - R_m I}\right) = \frac{-t_{isi}}{\tau_m} \quad (\text{A.17})$$

$$\ln\left(\frac{-E_L + V_{reset} - R_m I}{V_{th} - E_L - R_m I}\right) = \frac{t_{isi}}{\tau_m} \quad (\text{A.18})$$

$$t_{isi} = \tau_m \ln\left(\frac{-E_L + V_{reset} - R_m I}{V_{th} - E_L - R_m I}\right) \quad (\text{A.19})$$




Figure A.2: **Interspike intervals as a function of constant current injection.** Model parameters are the same as Fig A.1, with $V_{th} = -54\text{mV}$.

The ISI firing rate can be found by taking the inverse of the ISI:

$$r_{isi} = \frac{1}{t_{isi}} \quad (\text{A.20})$$

This expression is valid if $R_m I > V_{th} - E_L$; otherwise $r_{isi} = 0$. Finally, if a constant

current I is injected for time t_{max} , the spike count can be obtained according to:

$$\text{spike count} = r_{isi} t_{max} \tag{A.21}$$




Figure A.3: Spike count as a function of constant current injection of $t_{max} = 350\text{ms}$.

Bibliography

- Abbott, L. F., & Regehr, W. G. [W G]. (2004). Synaptic Computation. *Nature*, *431*(7010), 796–803. doi:10.1038/nature03010
- Abbott, L. F., Varela, J. A., Sen, K., & Nelson, S. B. (1997). Synaptic Depression and Cortical Gain Control. *Science*, *275*(5297), 220–224. doi:10.1126/science.275.5297.221
- Abraham, W. C., & Robins, A. (2005). Memory Retention – the Synaptic Stability versus Plasticity Dilemma. *Trends in Neurosciences*, *28*(2), 73–78. doi:10.1016/j.tins.2004.12.003
- Abrahamsson, T., Chou, C. Y. C., Li, S. Y., Mancino, A., Costa, R. P., Brock, J. A., . . . Sjöström, P. J. (2017). Differential Regulation of Evoked and Spontaneous Release by Presynaptic NMDA Receptors. *Neuron*, *96*(4), 839–855.e5. doi:10.1016/j.neuron.2017.09.030
- Adams, S. V., & Harris, C. M. (2015). A Computational Model of Innate Directional Selectivity Refined by Visual Experience. *Scientific Reports*, *5*(1), 12553. doi:10.1038/srep12553
- Amershi, S., Cakmak, M., Knox, W. B., & Kulesza, T. (2014). Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine*, *35*(4), 105. doi:10.1609/aimag.v35i4.2513
- Amit, D. J. (1995). The Hebbian paradigm reintegrated: Local Reverberations as Internal Representations. *Behavioral and Brain Sciences*, *18*(4), 617–626. doi:10.1017/S0140525X00040164

- Amit, D. J. (1997). Model of Global Spontaneous Activity and Local Structured Activity During Delay Periods in the Cerebral Cortex. *Cerebral Cortex*, 7(3), 237–252. doi:10.1093/cercor/7.3.237
- Amit, D. J., & Fusi, S. (1994). Learning in Neural Networks with Material Synapses. *Neural Computation*, 6(5), 957–982. doi:10.1162/neco.1994.6.5.957
- Amit, D. J., & Mongillo, G. (2003). Spike-driven Synaptic Dynamics Generating Working Memory States. *Neural Computation*, 15(3), 565–596. doi:10.1162/089976603321192086
- Anderson, M. (2014). *After phrenology: Neural Reuse and the Interactive Brain*. The MIT Press.
- Baddeley, A. (1992). Working Memory. *Science (New York, N.Y.)* 255(5044), 556–9. doi:10.1126/science.1736359
- Banerjee, A., Larsen, R. S., Philpot, B. D., & Paulsen, O. (2016). Roles of Presynaptic NMDA Receptors in Neurotransmission and Plasticity. *Trends in Neurosciences*, 39(1), 26–39. doi:10.1016/j.tins.2015.11.001
- Barak, O., & Tsodyks, M. (2007). Persistent Activity in Neural Networks with Dynamic Synapses. *PLoS Computational Biology*, 3(2), e35. doi:10.1371/journal.pcbi.0030035
- Barak, O., & Tsodyks, M. (2014). Working Models of Working Memory. doi:10.1016/j.conb.2013.10.008
- Barbosa, J., Stein, H., Martinez, R. L., Galan-Gadea, A., Li, S., Dalmau, J., ... Compte, A. (2020). Interplay Between Persistent Activity and Activity-silent Dynamics in the Prefrontal Cortex Underlies Serial Biases in Working Memory. *Nature Neuroscience*, 23(8), 1016–1024. doi:10.1038/s41593-020-0644-4
- Bates, D. W. (2014). Unity, Plasticity, Catastrophe: Order and Pathology in the Cybernetic Era. In *Catastrophes*. doi:10.1515/9783110312584.32
- Beer, R. D. (1998). Computational and Dynamical Languages for Autonomous Agents. In *Mind as motion*. doi:10.7551/mitpress/4622.003.0006

- Beer, R. D. (2014). Dynamical Systems and Embedded Cognition. In K. Frankish & W. M. Ramsey (Eds.), *The cambridge handbook of artificial intelligence* (pp. 128–148). doi:10.1017/CBO9781139046855.009
- Bekkers, J. M. (2003). Synaptic Transmission: Functional Autapses in the Cortex. *Current Biology*, *13*(11), R433–R435. doi:10.1016/S0960-9822(03)00363-4
- Ben-Ari, M., & Mondada, F. (2018). *Elements of Robotics*. doi:10.1007/978-3-319-62533-1
- Benda, J., & Herz, A. V. M. (2003). A Universal Model for Spike-frequency Adaptation. *Neural Computation*, *15*(11), 2523–2564. doi:10.1162/089976603322385063
- Bender, K. J. (2006). Synaptic Basis for Whisker Deprivation-induced Synaptic Depression in Rat Somatosensory Cortex. *Journal of Neuroscience*, *26*(16), 4155–4165. doi:10.1523/JNEUROSCI.0175-06.2006
- Berberian, N., MacPherson, A., Giraud, E., Richardson, L., & Thivierge, J.-P. (2017). Neuronal Pattern Separation of Motion-relevant Input in LIP Activity. *Journal of Neurophysiology*, *117*(2), 738–755. doi:10.1152/jn.00145.2016
- Berberian, N., Ross, M., & Chartier, S. (2019). Discrimination of Motion Direction in a Robot Using a Phenomenological Model of Synaptic Plasticity. *Computational Intelligence and Neuroscience*, *2019*, 1–14. doi:10.1155/2019/6989128
- Berberian, N., Ross, M., & Chartier, S. (2021). Embodied Working Memory During Ongoing Input Streams. *PLoS ONE*, 1–29. doi:10.1371/journal.pone.0244822
- Berberian, N., Ross, M., Chartier, S., & Thivierge, J.-P. (2017). Synergy Between Short-term and Long-term Plasticity Explains Direction-selectivity in Visual Cortex. In *Ieee symposium series on computational intelligence* (pp. 1–8). doi:10.1109/SSCI.2017.8280986
- Berkes, P., Orban, G., Lengyel, M., & Fiser, J. (2011). Spontaneous Cortical Activity Reveals Hallmarks of an Optimal Internal Model of the Environment. *Science*, *331*(6013), 83–87. doi:10.1126/science.1195870

- Bi, G.-Q. (2002). Spatiotemporal Specificity of Synaptic Plasticity: Cellular Rules and Mechanisms. *Biological Cybernetics*, *87*(5-6), 319–332. doi:10.1007/s00422-002-0349-7
- Bi, G.-Q., & Poo, M. (1998). Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *18*(24), 10464–10472.
- Bickle, J. (2016). Revolutions in Neuroscience: Tool Development. *Frontiers in Systems Neuroscience*, *10*. doi:10.3389/fnsys.2016.00024
- Bliss, T., & Collingridge, G. (1993). A Synaptic Model of Memory: Long-term Potentiation in the Hippocampus. *Nature*, *361*(6407), 31–39. doi:10.1038/361031a0
- Bliss, T., & Lømo, T. (1973). Long-lasting Potentiation of Synaptic Transmission in the Dentate Area of the Anaesthetized Rabbit Following Stimulation of the Perforant Path. *The Journal of Physiology*, *232*(2), 331–356. doi:10.1113/jphysiol.1973.sp010273
- Bolshakov, V., & Siegelbaum, S. (1995). Regulation of Hippocampal Transmitter Release During Development and Long-term Potentiation. *Science*, *269*(5231), 1730–1734. doi:10.1126/science.7569903
- Bostrom, N. (2014). *Superintelligence, Paths, Dangers, Strategies*. Oxford University Press.
- Bouchacourt, F., & Buschman, T. J. (2019). A Flexible Model of Working Memory. *Neuron*, *103*(1), 147–160.e8. doi:10.1016/j.neuron.2019.04.020
- Brader, J. M., Senn, W., & Fusi, S. (2007). Learning Real-world Stimuli in a Neural Network with Spike-driven Synaptic Dynamics. *Neural Computation*. doi:10.1162/neco.2007.19.11.2881
- Braitenberg, V., & Schüz, A. (1998). *Cortex: Statistics and Geometry of Neuronal Connectivity*. doi:10.1007/978-3-662-03733-1
- Brette, R. (2019). Is Coding a Relevant Metaphor for the Brain? *Behavioral and Brain Sciences*, *42*, e215. doi:10.1017/S0140525X19000049

- Brette, R., & Gerstner, W. (2005). Adaptive Exponential Integrate-and-fire Model as an Effective Description of Neuronal Activity. *Journal of Neurophysiology*, *94*(5), 3637–3642. doi:10.1152/jn.00686.2005
- Brody, C. D., Romo, R., & Kepecs, A. (2003). Basic Mechanisms for Graded Persistent Activity: Discrete Attractors, Continuous Attractors, and Dynamic Representations. *Current Opinion in Neurobiology*, *13*(2), 204–211. doi:10.1016/S0959-4388(03)00050-3
- Brooks, R. A. (1991). Intelligence Without Representation. *Artificial Intelligence*, *47*(1-3), 139–159. doi:10.1016/0004-3702(91)90053-M
- Brown, T. (1990). Hebbian Synapses: Biophysical Mechanisms And Algorithms. *Annual Review of Neuroscience*. doi:10.1146/annurev.neuro.13.1.475
- Browne, A. (2019). *Neural Network Perspectives on Cognition and Adaptive Robotics* (A. Browne, Ed.). doi:10.1201/9780367813239
- Brunel, N., & van Rossum, M. C. W. (2007). Quantitative Investigations of Electrical Nerve Excitation Treated as Polarization. *Biological Cybernetics*, *97*(5-6), 341–349. doi:10.1007/s00422-007-0189-6
- Buchs, N. J., & Senn, W. (2002). Spike-based Synaptic Plasticity and the Emergence of Direction Selective Simple Cells: Simulation Results. *J. Computational Neuroscience*, *20*(2), xx. doi:10.1023/A:1020210230751
- Burkitt, A., Meffin, H., & Grayden, D. (2004). Spike-timing-dependent Plasticity: The Relationship to Rate-based Learning for Models with Weight Dynamics Determined by a Stable Fixed Point. *Neural Computation*, *16*(5), 885–940. doi:10.1162/089976604773135041
- Canguilhem, G. (2008). Machine and Organism. In *Knowledge of life*.
- Carver, S., Roth, E., Cowan, N. J., & Fortune, E. S. (2008). Synaptic Plasticity can Produce and Enhance Direction Selectivity. *PLoS Computational Biology*, *4*(2), e32. doi:10.1371/journal.pcbi.0040032

- Cavanaugh, J. R. (2002). Nature and Interaction of Signals From the Receptive Field Center and Surround in Macaque V1 Neurons. *Journal of Neurophysiology*. doi:10.1152/jn.00692.2001
- Chafee, M. V., & Goldman-Rakic, P. S. (1998). Matching Patterns of Activity in Primate Prefrontal Area 8a and Parietal Area 7ip Neurons During a Spatial Working Memory-Task. *Journal of Neurophysiology*, 79(6), 2919–2940. doi:10.1152/jn.1998.79.6.2919
- Chance, F. S., Nelson, S. B., & Abbott, L. F. (1998). Synaptic Depression and the Temporal Response Characteristics of V1 Cells. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 18(12), 4785–99.
- Cheetham, C. E. J., & Fox, K. (2011). The Role of Sensory Experience in Presynaptic Development is Cortical Area Specific. *The Journal of Physiology*, 589(23), 5691–5699. doi:10.1113/jphysiol.2011.218347
- Chen, W., & Buonomano, D. (2012). Developmental Shift of Short-term Synaptic Plasticity in Cortical Organotypic Slices. *Neuroscience*, 213, 38–46. doi:10.1016/j.neuroscience.2012.04.018
- Chernova, S., & Thomaz, A. L. (2014). Robot Learning from Human Teachers. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 8(3), 1–121. doi:10.2200/S00568ED1V01Y201402AIM028
- Chiel, H. J., & Beer, R. D. (1997). The Brain Has a Body: Adaptive Behavior Emerges from Interactions of Nervous System, Body and Environment. *Trends in Neurosciences*, 20(12), 553–557. doi:10.1016/S0166-2236(97)01149-1
- Choe, Y., Yang, H.-F., & Eng, D. C.-Y. (2007). Autonomous Learning of the Semantics of Internal Sensory States Based on Motor Exploration. *International Journal of Humanoid Robotics*, 04(02), 211–243. doi:10.1142/S0219843607001102
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., ... Shenoy, K. V. (2010). Stimulus Onset Quenches Neural Variability: A

- Widespread Cortical Phenomenon. *Nature Neuroscience*, *13*(3), 369–378. doi:10.1038/nn.2501
- Clark, A. (1995). Moving Minds: Situating Content in the Service of Real-Time Success. *Philosophical Perspectives*, *9*, 89. doi:10.2307/2214213
- Clark, A. (1999). An Embodied Cognitive Science? doi:10.1016/S1364-6613(99)01361-3
- Clemens, J. M., Ritter, N. J., Roy, A., Miller, J. M., & Van Hooser, S. D. (2012). The Laminar Development of Direction Selectivity in Ferret Visual Cortex. *Journal of Neuroscience*, *32*(50), 18177–18185. doi:10.1523/JNEUROSCI.3399-12.2012
- Clopath, C., Büsing, L., Vasilaki, E., & Gerstner, W. (2010). Connectivity Reflects Coding: a Model of Voltage-based STDP with Homeostasis. *Nature Neuroscience*, *13*(3), 344–352. doi:10.1038/nn.2479
- Compte, A. (2000). Synaptic Mechanisms and Network Dynamics Underlying Spatial Working Memory in a Cortical Network Model. *Cerebral Cortex*, *10*(9), 910–923. doi:10.1093/cercor/10.9.910
- Constantinidis, C., Franowicz, M. N., & Goldman-Rakic, P. S. (2001). The Sensory Nature of Mnemonic Representation in the Primate Prefrontal Cortex. *Nature Neuroscience*, *4*(3), 311–316. doi:10.1038/85179
- Constantinidis, C., Funahashi, S., Lee, D., Murray, J. D., Qi, X.-L., Wang, M., & Arnsten, A. F. (2018). Persistent Spiking Activity Underlies Working Memory. *The Journal of Neuroscience*, *38*(32), 7020–7028. doi:10.1523/JNEUROSCI.2486-17.2018
- Cossell, L., Iacaruso, M. F., Muir, D. R., Houlton, R., Sader, E. N., Ko, H., . . . Mrsic-Flogel, T. D. (2015). Functional Organization of Excitatory Synaptic Strength in Primary Visual Cortex. *Nature*, *518*(7539), 399–403. doi:10.1038/nature14182
- Costa, R. P., Froemke, R. C., Sjöström, P. J., & van Rossum, M. C. (2015). Unified Pre- and Postsynaptic Long-term Plasticity Enables Reliable and Flexible Learning. *eLife*, *4*. doi:10.7554/eLife.09457

- Costa, R. P., Mizusaki, B. E. P., Sjöström, P. J., & van Rossum, M. C. W. (2017). Functional Consequences of Pre- and Postsynaptic Expression of Synaptic Plasticity. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*(1715), 20160153. doi:10.1098/rstb.2016.0153
- Crandall, J., Goodrich, M., Olsen, D., & Nielsen, C. (2005). Validating Human–Robot Interaction Schemes in Multitasking Environments. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, *35*(4), 438–449. doi:10.1109/TSMCA.2005.850587
- Cyr, A., Thériault, F., Ross, M., Berberian, N., & Chartier, S. (2018). Spiking Neurons Integrating Visual Stimuli Orientation and Direction Selectivity in a Robotic Context. *Frontiers in Neurobotics*, *12*(November), 1–10. doi:10.3389/fnbot.2018.00075
- Dan, Y., & Poo, M. M. (2004). Spike-timing-dependent Plasticity of Neural Circuits. doi:10.1016/j.neuron.2004.09.007
- Dayan, P., & Abbott, L. F. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*.
- Del Giudice, P., Fusi, S., & Mattia, M. (2003). Modelling the Formation of Working Memory with Networks of Integrate-and-fire Neurons Connected by Plastic Synapses. *Journal of Physiology-Paris*, *97*(4-6), 659–681. doi:10.1016/j.jphysparis.2004.01.021
- del Castillo, J., & Katz, B. (1954). Statistical Factors Involved in Neuromuscular Facilitation and Depression. *The Journal of Physiology*. doi:10.1113/jphysiol.1954.sp005130
- Deng, P. Y., & Klyachko, V. A. (2011). The Diverse Functions of Short-term Plasticity Components in Synaptic Computations. *Communicative and Integrative Biology*, *4*(5), 543–548. doi:10.4161/cib.4.5.15870
- Dittman, J. S., Kreitzer, A. C., & Regehr, W. G. (2000). Interplay between Facilitation, Depression, and Residual Calcium at three Presynaptic Terminals. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *20*(4), 1374–1385.

- Dobrunz, L., & Stevens, C. (1997). Heterogeneity of Release Probability, Facilitation, and Depletion at Central Synapses. *Neuron*, *18*(6), 995–1008. doi:10.1016/S0896-6273(00)80338-4
- Drover, J. D. (2014). Timing over Tuning: Overcoming the Shortcomings of a Line Attractor during a Working Memory Task. *PLoS Computational Biology*, *10*(1), e1003437. doi:10.1371/journal.pcbi.1003437
- Duarte, R., Seeholzer, A., Zilles, K., & Morrison, A. (2017). Synaptic Patterning and the Timescales of Cortical Dynamics. doi:10.1016/j.conb.2017.02.007
- Ecker, A. S., Berens, P., Tolias, A. S., & Bethge, M. (2011). The Effect of Noise Correlations in Populations of Diversely Tuned Neurons. *Journal of Neuroscience*, *31*(40), 14272–14283. doi:10.1523/JNEUROSCI.2539-11.2011
- Edin, F., Klingberg, T., Johansson, P., McNab, F., Tegner, J., & Compte, A. (2009). Mechanism for Top-down Control of Working Memory Capacity. *Proceedings of the National Academy of Sciences*, *106*(16), 6802–6807. doi:10.1073/pnas.0901894106
- Egorov, A. V., Hamam, B. N., Fransén, E., Hasselmo, M. E., & Alonso, A. A. (2002). Graded Persistent Activity in Entorhinal Cortex Neurons. *Nature*, *420*(6912), 173–178. doi:10.1038/nature01171
- Egorov, A. V., Unsicker, K., & von Bohlen und Halbach, O. (2006). Muscarinic Control of Graded Persistent Activity in Lateral Amygdala Neurons. *European Journal of Neuroscience*, *24*(11), 3183–3194. doi:10.1111/j.1460-9568.2006.05200.x
- Eliasmith, C. (2005). A Unified Approach to Building and Controlling Spiking Attractor Networks. *Neural Computation*, *17*(6), 1276–1314. doi:10.1162/0899766053630332
- Engel, A. K., Maye, A., Kurthen, M., & König, P. (2013). Where’s the Action? The Pragmatic Turn in Cognitive Science. doi:10.1016/j.tics.2013.03.006
- Engel, T. A., Steinmetz, N. A., Gieselmann, M. A., Thiele, A., Moore, T., & Boahen, K. (2016). Selective Modulation of Cortical State During Spatial Attention. *Science*, *354*(6316), 1140–1144. doi:10.1126/science.aag1420

- Erickson, M. A., Maramara, L. A., & Lisman, J. (2010). A Single Brief Burst Induces GluR1-dependent Associative Short-term Potentiation: A Potential Mechanism for Short-term Memory. *Journal of Cognitive Neuroscience*, *22*(11), 2530–2540. doi:10.1162/jocn.2009.21375
- Escobar, M. J., Masson, G. S., Vieville, T., & Kornprobst, P. (2009). Action Recognition using a Bio-inspired Feedforward Spiking Network. *International Journal of Computer Vision*, *82*(3), 284–301. doi:10.1007/s11263-008-0201-1
- Faisal, A. A., Selen, L. P., & Wolpert, D. M. (2008). Noise in the Nervous System. doi:10.1038/nrn2258
- Feldman, D. E. (2000). Timing-based LTP and LTD at Vertical Inputs to Layer II/III Pyramidal Cells in Rat Barrel Cortex. *Neuron*. doi:10.1016/S0896-6273(00)00008-8
- Feldman, D. E. (2012). The Spike-timing Dependence of Plasticity. *Neuron*, *75*(4), 556–71. doi:10.1016/j.neuron.2012.08.001
- Fernando, S., & Yamada, K. (2012). Spike-timing-dependent Plasticity and Short-term Plasticity Jointly Control the Excitation of Hebbian Plasticity without Weight Constraints in Neural Networks. *Computational Intelligence and Neuroscience*, *2012*, 1–15. doi:10.1155/2012/968272
- Fiebig, F., Herman, P., & Lansner, A. (2020). An Indexing Theory for Working Memory Based on Fast Hebbian Plasticity. *eneuro*, *7*(2), ENEURO.0374–19.2020. doi:10.1523/ENEURO.0374-19.2020
- Fiebig, F., & Lansner, A. (2017). A Spiking Working Memory Model Based on Hebbian Short-term Potentiation. *The Journal of Neuroscience*, *37*(1), 83–96. doi:10.1523/JNEUROSCI.1989-16.2017
- Finnerty, G. T., Roberts, L. S. E., & Connors, B. W. (1999). Sensory Experience Modifies the Short-term Dynamics of Neocortical Synapses. *Nature*, *400*(6742), 367–371. doi:10.1038/22553

- Fisher, S. (1997). Multiple Overlapping Processes Underlying Short-term Synaptic Enhancement. *Trends in Neurosciences*, *20*(4), 170–177. doi:10.1016/S0166-2236(96)01001-6
- Freedman, D. J. (2001). Categorical Representation of Visual Stimuli in the Primate Prefrontal Cortex. *Science*, *291*(5502), 312–316. doi:10.1126/science.291.5502.312
- Froemke, R., & Dan, Y. (2002). Spike-timing-dependent Synaptic Modification Induced by Natural Spike Trains. *Nature*, *416*(6879), 433–438. doi:10.1038/416433a
- Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic Coding of Visual Space in the Monkey's Dorsolateral Prefrontal Cortex. *Journal of Neurophysiology*, *61*(2), 331–349. doi:10.1152/jn.1989.61.2.331
- Fusi, S. (2002). Hebbian Spike-driven Synaptic Plasticity for Learning Patterns of Mean Firing Rates. *Biological Cybernetics*. doi:10.1007/s00422-002-0356-8
- Fusi, S., Annunziato, M., Badoni, D., Salamon, A., & Amit, D. J. (2000). Spike-Driven Synaptic Plasticity: Theory, Simulation, VLSI Implementation. *Neural Computation*, *12*(10), 2227–2258. doi:10.1162/089976600300014917
- Fuster, J. M., & Alexander, G. E. (1971). Neuron Activity Related to Short-term Memory. *Science*, *173*(3997), 652–654. doi:10.1126/science.173.3997.652
- Gabbott, P., & Somogyi, P. (1986). Quantitative Distribution of GABA-immunoreactive Neurons in the Visual Cortex (Area 17) of the Cat. *Experimental Brain Research*, *61*(2). doi:10.1007/BF00239522
- Gerstner, W., Kempter, R., van Hemmen, J. L., & Wagner, H. (1996). A Neuronal Learning Rule for Sub-millisecond Temporal Coding. *Nature*, *383*(6595), 76–78. doi:10.1038/383076a0
- Gibson, J. J. (1978). The Ecological Approach to the Visual Perception of Pictures. *Leonardo*, *11*(3), 227. doi:10.2307/1574154
- Giudice, P., & Mattia, M. (2001). Long and Short-term Synaptic Plasticity and the Formation of Working Memory: A Case Study. *Neurocomputing*, *38-40*, 1175–1180. doi:10.1016/S0925-2312(01)00557-4

- Goldman-Rakic, P. (1988). Topography of Cognition: Parallel Distributed Networks In Primate Association Cortex. *Annual Review of Neuroscience*, *11*(1), 137–156. doi:10.1146/annurev.neuro.11.1.137
- Goldman-Rakic, P. S. (1987). Circuitry of the Prefrontal Cortex and the Regulation of Behavior by Representational Knowledge. In *Handbook of physiology* (pp. 373–417). American Physiological Society.
- Goldman-Rakic, P. S. (1995). Cellular Basis of Working Memory. *Neuron*, *14*(3), 477–485. doi:10.1016/0896-6273(95)90304-6
- Goldman-Rakic, P. S. (1996). Regional and Cellular Fractionation of Working Memory. *Proceedings of the National Academy of Sciences*, *93*(24), 13473–13480. doi:10.1073/pnas.93.24.13473
- Goldman, M. (2009). Memory Without Feedback in a Neural Network. *Neuron*, *61*(4), 621–634. doi:10.1016/j.neuron.2008.12.012
- Goldman, M., Maldonado, P., & Abbott, L. (2002). Redundancy Reduction and Sustained Firing with Stochastic Depressing Synapses. *Journal of neuroscience*, *22*(2), 584–591. doi:10.1523/JNEUROSCI.22-02-00584.2002
- Goodrich, M. A., & Schultz, A. C. (2007). Human-Robot Interaction: A Survey. *Foundations and Trends in Human-Computer Interaction*, *1*(3), 203–275. doi:10.1561/11000000005
- Gravish, N., & Lauder, G. V. (2018). Robotics-inspired Biology. *The Journal of Experimental Biology*, *221*(7), jeb138438. doi:10.1242/jeb.138438
- Gütig, R., Aharonov, R., Rotter, S., & Sompolinsky, H. (2003). Learning Input Correlations Through Nonlinear Temporally Asymmetric Hebbian Plasticity. *Journal of Neuroscience*, *23*(9), 3697–714.
- Hansel, D., & Mato, G. (2013). Short-term Plasticity Explains Irregular Persistent Activity in Working Memory Tasks. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *33*(1), 133–49. doi:10.1523/JNEUROSCI.3455-12.2013

- Harvey, K. (2019). General Principles of Neurorobotic Models Employing Entrainment and Chaos Control. *Frontiers in Neurobotics*, *13*. doi:10.3389/fnbot.2019.00032
- Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-inspired Artificial Intelligence. *Neuron*, *95*(2), 245–258. doi:10.1016/j.neuron.2017.06.011
- Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory*. Oxford, England: Wiley.
- Hennig, M. H. (2013). Theoretical Models of Synaptic Short-term Plasticity. *Frontiers in Computational Neuroscience*, *7*, 45. doi:10.3389/fncom.2013.00045
- Hiratani, N., & Fukai, T. (2014). Interplay Between Short- and Long-term Plasticity in Cell-assembly Formation. *PLoS ONE*, *9*(7), e101535. doi:10.1371/journal.pone.0101535
- Hosaka, R., Araki, O., & Ikeguchi, T. (2008). STDP Provides the Substrate for Igniting Synfire Chains by Spatiotemporal Input Patterns. *Neural Computation*, *20*(2), 415–435. doi:10.1162/neco.2007.11-05-043
- Hoshi, E., Shima, K., & Tanji, J. (1998). Task-dependent Selectivity of Movement-related Neuronal Activity in the Primate Prefrontal Cortex. *Journal of Neurophysiology*. doi:10.1152/jn.1998.80.6.3392
- Hubbard, J. I. (1963). Repetitive Stimulation at the Mammalian Neuromuscular Junction, and the Mobilization of Transmitter. *The Journal of Physiology*, *169*(3), 641–662. doi:10.1113/jphysiol.1963.sp007286
- Hubel, D., & Wiesel, T. (1959). Receptive Fields of Single Neurones in the Cat's Striate Cortex. *The Journal of physiology*, *148*, 574–91. doi:10.1113/jphysiol.1959.sp006308
- Iida, F., & Nurzaman, S. G. (2016). Adaptation of Sensor Morphology: an Integrative View of Perception from Biologically Inspired Robotics Perspective. *Interface Focus*, *6*(4), 20160016. doi:10.1098/rsfs.2016.0016
- Itskov, V., Hansel, D., & Tsodyks, M. (2011). Short-term Facilitation may Stabilize Parametric Working Memory Trace. *Frontiers in Computational Neuroscience*, *5*(October), 1–19. doi:10.3389/fncom.2011.00040

- Izhikevich, E. M., & Desai, N. S. (2003). Relating STDP to BCM. *Neural Computation*, *15*(7), 1511–1523. doi:10.1162/089976603321891783
- Jain, L. C., Seera, M., Lim, C. P., & Balasubramaniam, P. (2014). A Review of Online Learning in Supervised Neural Networks. doi:10.1007/s00521-013-1534-4
- Jin, I., Puthanveetil, S., Udo, H., Karl, K., Kandel, E. R., & Hawkins, R. D. (2012). Spontaneous Transmitter Release is Critical for the Induction of Long-term and Intermediate-term Facilitation in Aplysia. *Proceedings of the National Academy of Sciences*, *109*(23), 9131–9136. doi:10.1073/pnas.1206914109
- Jolly, E., & Chang, L. J. (2019). The Flatland Fallacy: Moving Beyond Low-Dimensional Thinking. *Topics in Cognitive Science*, *11*(2), 433–454. doi:10.1111/tops.12404
- Kandel, E. R., Schwartz, J. H., & Jessell, T. M. (2000). *Principles of Neural Science: Fourth Edition*. doi:10.1036/0838577016
- Kasabov, N. (2007). *Evolving Connectionist Systems*. doi:10.1007/978-1-84628-347-5
- Katz, B., & Miledi, R. (1968). The Role of Calcium in Neuromuscular Facilitation. *The Journal of Physiology*, *195*(2), 481–492. doi:10.1113/jphysiol.1968.sp008469
- Kello, C. T., Brown, G. D., Ferrer-i-Cancho, R., Holden, J. G., Linkenkaer-Hansen, K., Rhodes, T., & Van Orden, G. C. (2010). Scaling Laws in Cognitive Sciences. *Trends in Cognitive Sciences*, *14*(5), 223–232. doi:10.1016/j.tics.2010.02.005
- Kelly, R. C., Smith, M. A., Kass, R. E., & Lee, T. S. (2010). Local Field Potentials Indicate Network State and Account for Neuronal Response Variability. *Journal of computational neuroscience*, *29*(3), 567–79. doi:10.1007/s10827-009-0208-9
- Kempter, R., Gerstner, W., & van Hemmen, J. L. (1999). Hebbian Learning and Spiking Neurons. *Physical Review E*, *59*(4), 4498–4514. doi:10.1103/PhysRevE.59.4498
- Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., & Arieli, A. (2003). Spontaneously Emerging Cortical Representations of Visual Attributes. *Nature*, *425*(6961), 954–956. doi:10.1038/nature02078

- Kilpatrick, Z. P. (2018). Synaptic Mechanisms of Interference in Working Memory. *Scientific Reports*, 8(1), 7879. doi:10.1038/s41598-018-25958-9
- Kim, C. M., & Chow, C. C. (2018). Learning Recurrent Dynamics in Spiking Networks. *eLife*, 7. doi:10.7554/eLife.37124
- Kintscher, M., Wozny, C., Johenning, F. W., Schmitz, D., & Breustedt, J. (2013). Role of RIM1 α in Short- and Long-term Synaptic Plasticity at Cerebellar Parallel Fibres. *Nature Communications*, 4(1), 2392. doi:10.1038/ncomms3392
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., ... Hadsell, R. (2017). Overcoming Catastrophic Forgetting in Neural Networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(13), 3521–3526. doi:10.1073/pnas.1611835114
- Kistler, W. M., & van Hemmen, J. L. (2000). Modeling Synaptic Plasticity in Conjunction with the Timing of Pre- and Postsynaptic Action Potentials. *Neural Computation*, 12(2), 385–405. doi:10.1162/089976600300015844
- Klampfl, S., & Maass, W. (2013). Emergence of Dynamic Memory Traces in Cortical Microcircuit Models through STDP. *Journal of Neuroscience*, 33(28), 11515–11529. doi:10.1523/JNEUROSCI.5044-12.2013
- Klug, A., Borst, J. G. G., Carlson, B. A., Kopp-Scheinpflug, C., Klyachko, V. A., & Xu-Friedman, M. A. (2012). How Do Short-term Changes at Synapses Fine-tune Information Processing? *Journal of Neuroscience*. doi:10.1523/JNEUROSCI.3348-12.2012
- Knoblauch, A. (2017). Impact of Structural Plasticity on Memory Formation and Decline. In *The rewiring brain* (pp. 361–386). doi:10.1016/B978-0-12-803784-3.00017-2
- Knoblauch, A., Körner, E., Körner, U., & Sommer, F. T. (2014). Structural Synaptic Plasticity Has High Memory Capacity and Can Explain Graded Amnesia, Catastrophic Forgetting, and the Spacing Effect. *PLoS ONE*, 9(5), e96485. doi:10.1371/journal.pone.0096485

- Knox, W. B., Fasel, I., & Stone, P. (2009). Design Principles for Creating Human-shapable Agents. In *Aaai spring symposium - technical report* (pp. 1–8).
- Knox, W. B., & Stone, P. (2009). Interactively Shaping Agents via Human Reinforcement. In *Proceedings of the fifth international conference on knowledge capture* (pp. 1–8). doi:10.1145/1597735.1597738
- Knox, W. B., Stone, P., & Breazeal, C. (2013). Training a Robot via Human Feedback: A Case Study. In *Lecture notes in computer science* (pp. 460–470). doi:10.1007/978-3-319-02675-6_46
- Ko, H., Cossell, L., Baragli, C., Antolik, J., Clopath, C., Hofer, S. B., & Mrsic-Flogel, T. D. (2013). The Emergence of Functional Microcircuits in Visual Cortex. *Nature*. doi:10.1038/nature12015
- Kohn, A., & Smith, M. (2016). Utah Array Extracellular Recordings of Spontaneous and Visually Evoked Activity from Anesthetized Macaque Primary Visual Cortex (V1). doi:http://dx.doi.org/10.6080/K0NC5Z4X
- Koulakov, A. A., Raghavachari, S., Kepecs, A., & Lisman, J. E. (2002). Model for a Robust Neural Integrator. *Nature Neuroscience*, 5(8), 775–782. doi:10.1038/nn893
- Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., & Poeppel, D. (2017). Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron*, 93(3), 480–490. doi:10.1016/j.neuron.2016.12.041
- Krichmar, J. L. (2008). Neurorobotics. doi:10.4249/scholarpedia.1365
- Krichmar, J. L. (2018). Neurorobotics—A Thriving Community and a Promising Pathway Toward Intelligent Cognitive Robots. *Frontiers in Neurorobotics*. doi:10.3389/fnbot.2018.00042
- Krichmar, J. L., Conradt, J., & Asada, M. (2015). Neurobiologically Inspired Robotics: Enhanced Autonomy through Neuromorphic Cognition. *Neural Networks*, 72, 1–2. doi:10.1016/j.neunet.2015.11.004

- Laing, C., & Chow, C. (2001). Stationary Bumps in Networks of Spiking Neurons. *Neural Computation*, *13*(7), 1473–1494. doi:10.1162/089976601750264974
- Lapicque, L. (1907). Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation. *J. Physiol. Pathol. Gen.* doi:10.1007/s00422-007-0189-6
- Larsen, R., & Sjöström, P. J. (2015). Synapse-type-specific Plasticity in Local Circuits. *Current Opinion in Neurobiology*, *35*, 127–135. doi:10.1016/j.conb.2015.08.001
- Larsen, R. S., Smith, I. T., Miriyala, J., Han, J. E., Corlew, R. J., Smith, S. L., & Philpot, B. D. (2014). Synapse-specific Control of Experience-dependent Plasticity by Presynaptic NMDA Receptors. *Neuron*, *83*(4), 879–893. doi:10.1016/j.neuron.2014.07.039
- Le Be, J.-V., & Markram, H. (2006). Spontaneous and Evoked Synaptic Rewiring in the Neonatal Neocortex. *Proceedings of the National Academy of Sciences*, *103*(35), 13214–13219. doi:10.1073/pnas.0604691103
- Legenstein, R., Naeger, C., & Maass, W. (2005). What can a Neuron Learn with Spike-timing-dependent Plasticity? *Neural Computation*. doi:10.1162/0899766054796888
- Li, Y., Van Hooser, S. D., Mazurek, M., White, L. E., & Fitzpatrick, D. (2008). Experience with Moving Visual Stimuli Drives the Early Development of Cortical Direction Selectivity. *Nature*, *456*(7224), 952–956. doi:10.1038/nature07417
- Liley, A. W., & North, K. A. K. (1953). An Electrical Investigation of Effects of Repetitive Stimulation on Mammalian Neuromuscular Junction. *Journal of Neurophysiology*, *16*(5), 509–527. doi:10.1152/jn.1953.16.5.509
- Lobo, J. L., Del Ser, J., Bifet, A., & Kasabov, N. (2020). Spiking Neural Networks and Online Learning: An Overview and Perspectives. *Neural Networks*, *121*, 88–100. doi:10.1016/j.neunet.2019.09.004
- Loebel, A., Le Bé, J.-V., Richardson, M. J. E., Markram, H., & Herz, A. V. M. (2013). Matched Pre- and Post-synaptic Changes Underlie Synaptic Plasticity over Long Time Scales. *Journal of Neuroscience*, *33*(15), 6257–66. doi:10.1523/JNEUROSCI.3740-12.2013

- Luczak, A., Barthó, P., & Harris, K. D. (2009). Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations. *Neuron*, *62*(3), 413–425. doi:10.1016/j.neuron.2009.03.014
- Macoveanu, J., Klingberg, T., & Tegnér, J. (2006). A Biophysical Model of Multiple-item Working Memory: A Computational and Neuroimaging Study. *Neuroscience*, *141*(3), 1611–1618. doi:10.1016/j.neuroscience.2006.04.080
- Magleby, K. (1987). Short-term Changes in Synaptic Efficacy. In G. Eldeman, W. Gall, & W. Cowan (Eds.), *Synaptic function* (pp. 21–56). John Wiley & Sons, Inc.
- Malenka, R. C. (1999). Long-Term Potentiation—A Decade of Progress? *Science*, *285*(5435), 1870–1874. doi:10.1126/science.285.5435.1870
- Malenka, R. C., & Bear, M. F. (2004). LTP and LTD. *Neuron*, *44*(1), 5–21. doi:10.1016/j.neuron.2004.09.012
- Manohar, S. G., Zokaei, N., Fallon, S. J., Vogels, T. P., & Husain, M. (2019). Neural mechanisms of Attending to Items in Working Memory. *Neuroscience & Biobehavioral Reviews*, *101*, 1–12. doi:10.1016/j.neubiorev.2019.03.017
- Markram, H. (2011). A History of Spike-timing-dependent Plasticity. *Frontiers in Synaptic Neuroscience*, *3*. doi:10.3389/fnsyn.2011.00004
- Markram, H., Lübke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of Synaptic Efficacy by Coincidence of Postsynaptic APs and EPSPs. *Science*, *275*(5297), 213–215. doi:10.1126/science.275.5297.213
- Markram, H., & Tsodyks, M. (1996). Redistribution of Synaptic Efficacy Between Neocortical Pyramidal Neurons. *Nature*, *382*(6594), 807–810. doi:10.1038/382807a0
- Markram, H., Wang, Y., & Tsodyks, M. (1998). Differential Signaling via the Same Axon of Neocortical Pyramidal Neurons. *Proceedings of the National Academy of Sciences*, *95*(9), 5323–5328. doi:10.1073/pnas.95.9.5323

- Matveev, V., & Wang, X.-J. (2000). Implications of All-or-None Synaptic Transmission and Short-Term Depression beyond Vesicle Depletion: A Computational Study. *The Journal of Neuroscience*, *20*(4), 1575–1588. doi:10.1523/JNEUROSCI.20-04-01575.2000
- McCulloch, W. S., & Pitts, W. (1943). A Logical Calculus of the Ideas Immanent in Nervous Activity. *The Bulletin of Mathematical Biophysics*, *5*(4), 115–133. doi:10.1007/BF02478259
- Meyer, T., Qi, X.-L., Stanford, T. R., & Constantinidis, C. (2011). Stimulus Selectivity in Dorsal and Ventral Prefrontal Cortex after Training in Working Memory Tasks. *Journal of Neuroscience*, *31*(17), 6266–6276. doi:10.1523/JNEUROSCI.6798-10.2011
- Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural Mechanisms of Visual Working Memory in Prefrontal Cortex of the Macaque. *The Journal of Neuroscience*, *16*(16), 5154–5167. doi:10.1523/JNEUROSCI.16-16-05154.1996
- Miller, E., Li, L., & Desimone, R. (1993). Activity of Neurons in Anterior Inferior Temporal Cortex During a Short-term Memory Task. *Journal of Neuroscience*. doi:10.1523/jneurosci.13-04-01460.1993
- Millikan, R. G. (1995). Pushmi-Pullyu Representations. *Philosophical Perspectives*, *9*, 185. doi:10.2307/2214217
- Missaire, M., Fraize, N., Comte, J.-C., Truchet, B., Parmentier, R., Salin, P.-A., & Malleret, G. (2020). Working and Reference Memory Tasks Trigger Opposed Long-term Synaptic Changes in the Rat Dentate Gyrus. *bioRxiv Animal Behavior and Cognition*, 1–28.
- Miyashita, Y., & Chang, H. S. (1988). Neuronal Correlate of Pictorial Short-term Memory in the Primate Temporal Cortex. *Nature*, *331*(6151), 68–70. doi:10.1038/331068a0
- Monday, H. R., & Castillo, P. E. (2017). Closing the Gap: Long-term Presynaptic Plasticity in Brain Function and Disease. *Current opinion in neurobiology*, *45*, 106–112. doi:10.1016/j.conb.2017.05.011
- Monday, H. R., Younts, T. J., & Castillo, P. E. (2018). Long-term Plasticity of Neurotransmitter Release: Emerging Mechanisms and Contributions to Brain Function and

- Disease. *Annual Review of Neuroscience*, *41*(1), 299–322. doi:10.1146/annurev-neuro-080317-062155
- Mongillo, G., Curti, E., Romani, S., & Amit, D. J. (2005). Learning in Realistic Networks of Spiking Neurons and Spike-driven Plastic Synapses. *European Journal of Neuroscience*, *21*(11), 3143–3160. doi:10.1111/j.1460-9568.2005.04087.x
- Morimoto, J., & Kawato, M. (2015). Creating the Brain and Interacting with the Brain: an Integrated Approach to Understanding the Brain. *Journal of The Royal Society Interface*, *12*(104), 20141250. doi:10.1098/rsif.2014.1250
- Morrison, A., Diesmann, M., & Gerstner, W. (2008). Phenomenological Models of Synaptic Plasticity Based on Spike Timing. *Biological Cybernetics*, *98*(6), 459–478. doi:10.1007/s00422-008-0233-1
- Nabavi, S., Fox, R., Proulx, C. D., Lin, J. Y., Tsien, R. Y., & Malinow, R. (2014). Engineering a Memory with LTD and LTP. *Nature*, *511*(7509), 348–352. doi:10.1038/nature13294
- Nadal, J. P., Toulouse, G., Changeux, J. P., & Dehaene, S. (1986). Networks of Formal Neurons and Memory Palimpsests. *Europhysics Letters (EPL)*, *1*(10), 535–542. doi:10.1209/0295-5075/1/10/008
- Najar, A., & Chetouani, M. (2020). Reinforcement Learning with Human Advice. A survey. *Autonomous Agent and Multi-Agent Systems*.
- Najar, A., Sigaud, O., & Chetouani, M. (2019). Interactively Shaping Robot Behaviour with Unlabeled Human Instructions. *Autonomous Agent and Multi-Agent Systems*, 1–8.
- Natschlger, T., Maass, W., & Zador, A. (2001). Efficient Temporal Processing with Biologically Realistic Dynamic Synapses. *Network: Computation in Neural Systems*, *12*(1), 75–87. doi:10.1088/0954-898X/12/1/305
- Naud, R., Marcille, N., Clopath, C., & Gerstner, W. (2008). Firing Patterns in the Adaptive Exponential Integrate-and-fire Model. *Biological Cybernetics*, *99*(4-5), 335–347. doi:10.1007/s00422-008-0264-7

- Naud, R., & Sprekeler, H. (2018). Sparse bursts optimize information transmission in a multiplexed neural code. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(27), E6329–E6338. doi:10.1073/pnas.1720995115
- Nolfi, S., & Parisi, D. (1993). Self-selection of Input Stimuli for Improving Performance. In *Neural networks in robotics* (pp. 403–418). doi:10.1007/978-1-4615-3180-7_23
- Ó Scalaidhe, S. P. (1997). Areal Segregation of Face-Processing Neurons in Prefrontal Cortex. *Science*, *278*(5340), 1135–1138. doi:10.1126/science.278.5340.1135
- Ó Scalaidhe, S. P., Wilson, F. A., & Goldman-Rakic, P. S. (1999). Face-selective Neurons during Passive Viewing and Working Memory Performance of Rhesus Monkeys: Evidence for Intrinsic Specialization of Neuronal Coding. *Cerebral Cortex*. doi:10.1093/cercor/9.5.459
- Oberts, J., Sanders, S., Knoll, A., & Gewaltig, M.-O. (2016). Brain-inspired Intelligent Robotics: The Intersection of Robotics and Neuroscience Sciences. *Science*, *354*(6318), 1445.2–1445. doi:10.1126/science.354.6318.1445-b
- P. Carvalho, T. (2011). A Novel Learning Rule for Long-term Plasticity of Short-term Synaptic Plasticity Enhances Temporal Processing. *Frontiers in Integrative Neuroscience*, *5*, 1–11. doi:10.3389/fnint.2011.00020
- Page, M. (2000). Connectionist modelling in psychology: A localist manifesto. *Behavioral and Brain Sciences*, *23*(4), 443–467. doi:10.1017/S0140525X00003356
- Paoletti, P., Bellone, C., & Zhou, Q. (2013). NMDA Receptor Subunit Diversity: Impact on Receptor Properties, Synaptic Plasticity and Disease. *Nature Reviews Neuroscience*, *14*(6), 383–400. doi:10.1038/nrn3504
- Parisi, G. (1986). A Memory Which Forgets. *Journal of Physics A: Mathematical and General*, *19*(10), L617–L620. doi:10.1088/0305-4470/19/10/011
- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual Lifelong Learning with Neural Networks: A Review. *Neural Networks*, *113*, 54–71. doi:10.1016/j.neunet.2019.01.012

- Parisien, C., Anderson, C. H., & Eliasmith, C. (2008). Solving the Problem of Negative Synaptic Weights in Cortical Models. *Neural Computation*, *20*(6), 1473–1494. doi:10.1162/neco.2008.07-06-295
- Park, P., Volianskis, A., Sanderson, T. M., Bortolotto, Z. A., Jane, D. E., Zhuo, M., ... Collingridge, G. L. (2014). NMDA Receptor-dependent Long-term Potentiation Comprises a Family of Temporally Overlapping Forms of Synaptic Plasticity that are Induced by Different Patterns of Stimulation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1633), 20130131. doi:10.1098/rstb.2013.0131
- Payeur, A., Guerguiev, J., Zenke, F., Richards, B. A., & Naud, R. (2020). Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. doi:10.1101/2020.03.30.015511
- Pfeifer, R., Lungarella, M., & Iida, F. (2007). Self-Organization, Embodiment, and Biologically Inspired Robotics. *Science*, *318*(5853), 1088–1093. doi:10.1126/science.1145803
- Pfister, J. P. (2006). Triplets of Spikes in a Model of Spike-timing-dependent Plasticity. *Journal of Neuroscience*, *26*(38), 9673–9682. doi:10.1523/JNEUROSCI.1425-06.2006
- Pfister, J. P., & Gerstner, W. (2005). Beyond Pair-based STDP: A Phenomenological Rule for Spike Triplet and Frequency Effects. In *Advances in neural information processing systems*.
- Ploner, C. J., Gaymard, B., Rivaud, S., Agid, Y., & Pierrot-Deseilligny, C. (1998). Temporal Limits of Spatial Working Memory in Humans. *European Journal of Neuroscience*. doi:10.1046/j.1460-9568.1998.00101.x
- Pradier, B., Lanning, K., Taljan, K. T., Feuille, C. J., Nagy, M. A., & Kauer, J. A. (2018). Persistent but Labile Synaptic Plasticity at Excitatory Synapses. *The Journal of Neuroscience*, *38*(25), 5750–5758. doi:10.1523/JNEUROSCI.2772-17.2018
- Rainer, G., Asaad, W. F., & Miller, E. K. (1998). Selective Representation of Relevant Information by Neurons in the Primate Prefrontal Cortex. *Nature*, *393*(6685), 577–579. doi:10.1038/31235

- Rainer, G., & Miller, E. K. (2000). Effects of Visual Experience on the Representation of Objects in the Prefrontal Cortex. *Neuron*. doi:10.1016/S0896-6273(00)00019-2
- Ramaswamy, S., Hill, S. L., King, J. G., Schürmann, F., Wang, Y., & Markram, H. (2012). Intrinsic Morphological Diversity of Thick-tufted Layer 5 Pyramidal Neurons Ensures Robust and Invariant Properties of in Silico Synaptic Connections. *Journal of Physiology*. doi:10.1113/jphysiol.2011.219576
- Rao, R. P., & Sejnowski, T. J. (2000). Predictive Sequence Learning in Recurrent Neocortical Circuits. In *Advances in neural information processing systems*.
- Rao, S. C. (1997). Integration of What and Where in the Primate Prefrontal Cortex. *Science*, 276(5313), 821–824. doi:10.1126/science.276.5313.821
- Redondo, R. L., & Morris, R. G. M. (2011). Making Memories Last: the Synaptic Tagging and Capture Hypothesis. *Nature Reviews Neuroscience*, 12(1), 17–30. doi:10.1038/nrn2963
- Renart, A., Song, P., & Wang, X.-J. (2003). Robust Spatial Working Memory through Homeostatic Synaptic Scaling in Heterogeneous Cortical Networks. *Neuron*, 38(3), 473–485. doi:10.1016/S0896-6273(03)00255-1
- Reymann, K. G., & Frey, J. U. (2007). The Late Maintenance of Hippocampal LTP: Requirements, Phases, Synaptic Tagging, Late-associativity and Implications. *Neuropharmacology*, 52(1), 24–40. doi:10.1016/j.neuropharm.2006.07.026
- Richardson, M. J. E., Melamed, O., Silberberg, G., Gerstner, W., & Markram, H. (2005). Short-Term Synaptic Plasticity Orchestrates the Response of Pyramidal Cells and Interneurons to Population Bursts. *Journal of Computational Neuroscience*, 18(3), 323–331. doi:10.1007/s10827-005-0434-8
- Rinaldi, T., Silberberg, G., & Markram, H. (2008). Hyperconnectivity of Local Neocortical Microcircuitry Induced by Prenatal Exposure to Valproic Acid. *Cerebral Cortex*, 18(4), 763–770. doi:10.1093/cercor/bhm117
- Roelfsema, P. R., & Holtmaat, A. (2018). Control of synaptic plasticity in deep cortical networks. *Nature Reviews Neuroscience*, 19(3), 166–180. doi:10.1038/nrn.2018.6

- Romani, A., Marchetti, C., Bianchi, D., Leinekugel, X., Poirazi, P., Migliore, M., & Marie, H. (2013). Computational Modeling of the Effects of Amyloid-beta on Release Probability at Hippocampal Synapses. *Frontiers in Computational Neuroscience*, *7*. doi:10.3389/fncom.2013.00001
- Rosenmund, C., & Stevens, C. F. (1996). Definition of the Readily Releasable Pool of Vesicles at Hippocampal Synapses. *Neuron*, *16*(6), 1197–1207. doi:10.1016/S0896-6273(00)80146-4
- Rotman, Z., Deng, P.-Y., & Klyachko, V. A. (2011). Short-term Plasticity Optimizes Synaptic Information Transmission. *Journal of Neuroscience*, *31*(41), 14800–14809. doi:10.1523/JNEUROSCI.3231-11.2011
- Roy, A., Christie, I. K., Escobar, G. M., Osik, J. J., Popović, M., Ritter, N. J., ... Van Hooser, S. D. (2018). Does Experience Provide a Permissive or Instructive Influence on the Development of Direction Selectivity in Visual Cortex? *Neural Development*, *13*(1), 16. doi:10.1186/s13064-018-0113-x
- Russel, S., & Norvig, P. (2012). *Artificial Intelligence—A Modern Approach 3rd Edition*. doi:10.1017/S0269888900007724
- Russell, S. (2019). *Human Compatible: Artificial Intelligence and the Problem of Control*. Viking.
- Sadeh, S., Clopath, C., & Rotter, S. (2015). Emergence of Functional Specificity in Balanced Networks with Synaptic Plasticity. *PLoS Computational Biology*. doi:10.1371/journal.pcbi.1004307
- Saksida, L. M., Raymond, S. M., & Touretzky, D. S. (1997). Shaping Robot Behavior using Principles from Instrumental Conditioning. *Robotics and Autonomous Systems*, *22*(3-4), 231–249. doi:10.1016/S0921-8890(97)00041-9
- Sandamirskaya, Y. (2014). Dynamic Neural Fields as a Step Toward Cognitive Neuromorphic Architectures. *Frontiers in Neuroscience*. doi:10.3389/fnins.2013.00276

- Sandberg, A., Tegnér, J., & Lansner, A. (2003). A Working Memory Model Based on Fast Hebbian Learning. *Network: Computation in Neural Systems*, *14*(4), 789–802. doi:10.1088/0954-898X.14.4.309
- Schlesinger, M. (2003). A Lesson from Robotics: Modeling Infants as Autonomous Agents. *Adaptive Behavior*, *11*(2), 97–107. doi:10.1177/10597123030112003
- Schneegans, S., & Bays, P. M. (2018). Drift in Neural Population Activity Causes Working Memory to Deteriorate Over Time. *The Journal of Neuroscience*, *38*(21), 4859–4869. doi:10.1523/JNEUROSCI.3440-17.2018
- Schuett, S., Bonhoeffer, T., & Hübener, M. (2001). Pairing-induced Changes of Orientation Maps in Cat Visual Cortex. *Neuron*. doi:10.1016/S0896-6273(01)00472-X
- Seeholzer, A., Deger, M., & Gerstner, W. (2019). Stability of Working Memory in Continuous Attractor Networks under the Control of Short-term Plasticity. *PLOS Computational Biology*, *15*(4), e1006928. doi:10.1371/journal.pcbi.1006928
- Senn, W., Markram, H., & Tsodyks, M. (2001). An Algorithm for Modifying Neurotransmitter Release Probability Based on Pre- and Postsynaptic Spike Timing. *Neural Computation*, *13*(1), 35–67. doi:10.1162/089976601300014628
- Sharkey, N. E., & Heemskerk, J. N. H. (2019). The Neural Mind and the Robot. In *Neural network perspectives on cognition and adaptive robotics* (pp. 168–194). doi:10.1201/9780367813239-12
- Singh, R. (2006). Higher-Dimensional Neurons Explain the Tuning and Dynamics of Working Memory Cells. *Journal of Neuroscience*, *26*(14), 3667–3678. doi:10.1523/JNEUROSCI.4864-05.2006
- Sjöström, J., & Gerstner, W. (2010). Spike-timing-dependent Plasticity. *Scholarpedia*, *5*(2), 1362. doi:10.4249/scholarpedia.1362
- Sjöström, P. J., Turrigiano, G. G., & Nelson, S. B. (2001). Rate, Timing, and Cooperativity Jointly Determine Cortical Synaptic Plasticity. *Neuron*, *32*(6), 1149–1164. doi:10.1016/S0896-6273(01)00542-6

- Sjöström, P., Turrigiano, G. G., & Nelson, S. B. (2003). Neocortical LTD via Coincident Activation of Presynaptic NMDA and Cannabinoid Receptors. *Neuron*, *39*(4), 641–654. doi:10.1016/S0896-6273(03)00476-8
- Sjöström, P., Turrigiano, G. G., & Nelson, S. B. (2007). Multiple Forms of Long-term Plasticity at Unitary Neocortical Layer 5 Synapses. *Neuropharmacology*, *52*(1), 176–184. doi:10.1016/j.neuropharm.2006.07.021
- Smith, M. A., & Kohn, A. (2008). Spatial and Temporal Scales of Neuronal Correlation in Primary Visual Cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *28*(48), 12591–603. doi:10.1523/JNEUROSCI.2929-08.2008
- Song, S., Miller, K. D., & Abbott, L. F. (2000). Competitive Hebbian Learning through Spike-timing-dependent Synaptic Plasticity. *Nature Neuroscience*, *3*(9), 919–926. doi:10.1038/78829
- Sporns, O. (2007). What Neuro-robotic Models can Teach us About Neural and Cognitive Development. In *Neuroconstructivism volume twoperspectives and prospects* (pp. 179–204). doi:10.1093/acprof:oso/9780198529934.003.0008
- Sporns, O., & Pegors, T. K. (2004). Information-Theoretical Aspects of Embodied Artificial Intelligence. In *Lecture notes in artificial intelligence* (pp. 74–85). doi:10.1007/978-3-540-27833-7_5
- Stevens, C., & Wang, Y. (1994). Changes in Reliability of Synaptic Function as a Mechanism for Plasticity. *Nature*, *371*(6499), 704–707. doi:10.1038/371704a0
- Stevens, C., & Wang, Y. (1995). Facilitation and depression at single central synapses. *Neuron*, *14*(4), 795–802. doi:10.1016/0896-6273(95)90223-6
- Takeda, K., & Funahashi, S. (2002). Prefrontal Task-Related Activity Representing Visual Cue Location or Saccade Direction in Spatial Working Memory Tasks. *Journal of Neurophysiology*, *87*(1), 567–588. doi:10.1152/jn.00249.2001

- Takesian, A. E., Kotak, V. C., & Sanes, D. H. (2010). Presynaptic GABAB Receptors Regulate Experience-Dependent Development of Inhibitory Short-term Plasticity. *Journal of Neuroscience*, *30*(7), 2716–2727. doi:10.1523/JNEUROSCI.3903-09.2010
- Tanaka, S. (2002). Dopamine Controls Fundamental Cognitive Operations of Multi-target Spatial Working Memory. *Neural Networks*, *15*(4-6), 573–582. doi:10.1016/S0893-6080(02)00050-3
- Tauffer, L., & Kumar, A. (2019). Short-term Synaptic Plasticity Makes Neurons Sensitive to the Distribution of Presynaptic Population Firing Rates. *bioRxiv*, 707398. doi:10.1101/707398
- Tegmark, M. (2017). *Life 3.0: Being Human in the Age of Artificial Intelligence*. Alfred A. Knopf.
- Thomson, A. M. (2000). Facilitation, Augmentation and Potentiation at Central Synapses. doi:10.1016/S0166-2236(00)01580-0
- Tokuoka, H., & Goda, Y. (2008). Activity-dependent Coordination of Presynaptic Release Probability and Postsynaptic GluR2 Abundance at Single Synapses. *Proceedings of the National Academy of Sciences*, *105*(38), 14656–14661. doi:10.1073/pnas.0805705105
- Tritsch, N. X., & Sabatini, B. L. (2012). Dopaminergic Modulation of Synaptic Transmission in Cortex and Striatum. *Neuron*, *76*(1), 33–50. doi:10.1016/j.neuron.2012.09.023
- Tsodyks, M., & Markram, H. (1997). The Neural Code between Neocortical Pyramidal Neurons depends on Neurotransmitter Release Probability. *Proceedings of the National Academy of Sciences of the United States of America*, *94*(2), 719–23. doi:10.1073/pnas.94.2.719
- Tsodyks, M., Pawelzik, K., & Markram, H. (1998). Neural Networks with Dynamic Synapses. *Neural Computation*, *10*(4), 821–835. doi:10.1162/089976698300017502
- Turkel, W. J., Port, R. F., & van Gelder, T. (1998). Mind as Motion: Explorations in the Dynamics of Cognition. *Language*, *74*(1), 219. doi:10.2307/417624

- Underwood, B. J. (1957). Interference and Forgetting. *Psychological Review*, *64*(1), 49–60.
doi:10.1037/h0044616
- Urban-Ciecko, J., Wen, J. A., Parekh, P. K., & Barth, A. L. (2015). Experience-dependent Regulation of Presynaptic NMDARs Enhances Neurotransmitter Release at Neocortical Synapses. *Learning & Memory*, *22*(1), 47–55. doi:10.1101/lm.035741.114
- Van Gelder, T. (1995). *It's About Time: An overview of the Dynamical Approach to Cognition*. Cambridge MA: MIT Press.
- Van Hooser, S. D., Li, Y., Christensson, M., Smith, G. B., White, L. E., & Fitzpatrick, D. (2012). Initial Neighborhood Biases and the Quality of Motion Stimulation Jointly Influence the Rapid Emergence of Direction Preference in Visual Cortex. *Journal of Neuroscience*, *32*(21), 7258–7266. doi:10.1523/JNEUROSCI.0230-12.2012
- van Gelder, T. (1991). *What is the "D" in "PDP"?* A survey of the Concept of Distribution (W. Ramsey, S. Stich, & D. Rumelhart, Eds.). Taylor & Francis.
- van Rossum, M. C. W., Bi, G.-Q., & Turrigiano, G. G. (2000). Stable Hebbian Learning from Spike-timing-dependent Plasticity. *The Journal of Neuroscience*, *20*(23), 8812–8821. doi:10.1523/JNEUROSCI.20-23-08812.2000
- van Rossum, M. C. W., van der Meer, M., Xiao, D., & Oram, M. (2008). Adaptive Integration in the Visual Cortex by Depressing Recurrent Cortical Circuits. *Neural Computation*, *20*(7), 1847–1872. doi:10.1162/neco.2008.06-07-546
- Varela, J. A., Sen, K., Gibson, J., Fost, J., Abbott, L. F., & Nelson, S. B. (1997). A Quantitative Description of Short-term Plasticity at Excitatory Synapses in Layer 2/3 of Rat Primary Visual Cortex. *Journal of Neuroscience*, *17*(20), 7926–40.
- Vasilaki, E., & Giugliano, M. (2014). Emergence of Connectivity Motifs in Networks of Model Neurons with Short- and Long-term Plastic Synapses. *PLoS ONE*, *9*(1), e84626. doi:10.1371/journal.pone.0084626
- Vinci, G., Ventura, V., Smith, M. A., & Kass, R. E. (2016). Separating Spike Count Correlation from Firing Rate Correlation. *Neural Computation*. doi:10.1162/NECO_a_00831

- Walter Grey, W. (1950). An Imitation of Life. *Scientific American*. doi:10.1038/scientificamerican0550-42
- Walter Grey, W. (1953). *The Living Brain*. doi:10.1007/978-3-642-70911-1_17
- Wang, J., Belatreche, A., Maguire, L., & McGinnity, M. (2010). Online versus Offline Learning for Spiking Neural Networks: A Review and New Strategies. In *2010 IEEE 9th International Conference on Cybernetic Intelligent Systems* (pp. 1–6). doi:10.1109/UKRICIS.2010.5898113
- Wang, J., Belatreche, A., Maguire, L., & McGinnity, T. M. (2014). An Online Supervised Learning Method for Spiking Neural Networks with Adaptive Structure. *Neurocomputing*, *144*, 526–536. doi:10.1016/j.neucom.2014.04.017
- Wang, Y., Markram, H., Goodman, P. H., Berger, T. K., Ma, J., & Goldman-Rakic, P. S. (2006). Heterogeneity in the Pyramidal Network of the Medial Prefrontal Cortex. *Nature Neuroscience*, *9*(4), 534–542. doi:10.1038/nn1670
- Webb, B. (2001). Can Robots Make Good Models of Biological Behaviour? *Behavioral and Brain Sciences*. doi:10.1017/s0140525x01000127
- Wei, Z., Wang, X.-J., & Wang, D.-H. (2012). From Distributed Resources to Limited Slots in Multiple-Item Working Memory: A Spiking Network Model with Normalization. *Journal of Neuroscience*, *32*(33), 11228–11240. doi:10.1523/JNEUROSCI.0735-12.2012
- Weng, J. (2001). Artificial Intelligence: Autonomous Mental Development by Robots and Animals. *Science*, *291*(5504), 599–600. doi:10.1126/science.291.5504.599
- Weng, J. (2004). Developmental Robotics: Theory and Experiments. *International Journal of Humanoid Robotics*, *01*(02), 199–236. doi:10.1142/S0219843604000149
- Wilson, F., Ó Scalaidhe, S., & Goldman-Rakic, P. (1993). Dissociation of Object and Spatial Processing Domains in Primate Prefrontal Cortex. *Science*, *260*(5116), 1955–1958. doi:10.1126/science.8316836
- Wilson, M. (2002). Six Views of Embodied Cognition. *Psychonomic Bulletin & Review*, *9*(4), 625–636. doi:10.3758/BF03196322

- Wixted, J. T. (2004). The Psychology and Neuroscience of Forgetting. *Annual Review of Psychology*, 55(1), 235–269. doi:10.1146/annurev.psych.55.090902.141555
- Wu, Q., Mcginnity, T. M., Maguire, L., & Cai, J. (2008). Motion Detection Using Spiking Neural Network Model. In *International conference on intelligent computing* (pp. 76–83). Springer-Verlag Berlin Heidelberg 2008.
- Wysoski, S. G., Benuskova, L., & Kasabov, N. (2006). On-line Learning with Structural Adaptation in a Network of Spiking Neurons for Visual Pattern Recognition. In *Lecture notes in computer science (including subseries lecture notes in artificial intelligence and lecture notes in bioinformatics)* (pp. 61–70). doi:10.1007/11840817_7
- Yudkowsky, E. (2009). Paperclip Maximiser.
- Zador, A., & Dobrunz, L. (1997). Dynamic Synapses in the Cortex. *Neuron*, 19(1), 1–4. doi:10.1016/S0896-6273(00)80341-4
- Zech, P., Haller, S., Lakani, S. R., Ridge, B., Ugur, E., & Piater, J. (2017). Computational Models of Affordance in Robotics: a Taxonomy and Systematic Classification. *Adaptive Behavior*, 25(5), 235–271. doi:10.1177/1059712317726357
- Zenke, F., Poole, B., & Ganguli, S. (2017). Continual Learning Through Synaptic Intelligence. *Proceedings of machine learning research*, 70, 3987–3995.
- Zucker, R. S. (1996). Exocytosis: A Molecular and Physiological Perspective. *Neuron*, 17(6), 1049–1055. doi:10.1016/S0896-6273(00)80238-X
- Zucker, R. S., & Regehr, W. G. [W. G.]. (2002). Short-term Synaptic Plasticity. *Annual Review of Physiology*, 64(1), 355–405. doi:10.1146/annurev.physiol.64.092501.114547
- Zucker, R. (1989). Short-term Synaptic Plasticity. *Annual Review of Neuroscience*, 12(1), 13–31. doi:10.1146/annurev.ne.12.030189.000305