

The Impact of Structural Constraints on Learning in Biological and Artificial Neural Networks

Dissertation

der Mathematisch-Naturwissenschaftlichen Fakultät
der Eberhard Karls Universität Tübingen
zur Erlangung des Grades eines
Doktors der Naturwissenschaften
(Dr. rer. nat.)

vorgelegt von
Emmanouil Giannakakis
aus Athen/Griechenland

Tübingen
2024

Gedruckt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der
Eberhard Karls Universität Tübingen.

Tag der mündlichen Qualifikation:	15.09.2025
Dekan:	Prof. Dr. Thilo Stehle
1. Berichterstatter/-in:	Prof. Dr. Anna Levina
2. Berichterstatter/-in:	Prof. Dr. Martin Giese

For all the friends I love and might someday forget

Abstract

Biological neural networks are characterized by vast complexity, which manifests itself in the highly intricate and specific structures that arise during development. From the numerous ion channels and complex subcellular biophysical processes that drive the dynamics of single neurons to the intricate structure of the dendritic tree and the highly specific connectivity patterns between hundreds of distinct neuron types, biological networks are highly optimized for information processing by millions of years of evolution. On top of these intricate structures, a vast array of adaptive mechanisms modifies the neuronal connection strengths with precision and efficiency to incorporate new information and adapt behaviour to changing conditions. However, unlike artificial networks, for whom the learning process itself is the sole origin of structure and function, learning in biological networks only happens within the constraints imposed by these intricate cellular and network level structures.

This thesis investigates the influence that these structural constraints have on the learning process of biological neural networks. We begin by studying the evolutionary origins of synaptic plasticity and the extent to which different aspects of the network and task structure can influence the form of evolved plasticity rules. Continuing with a greater focus on biological detail, we study how structural features of biological networks, including a network's topology as well as biophysical properties of individual neurons such as complex dendritic structures, exert a very significant influence on the ability of local synaptic plasticity mechanisms to perform simple unsupervised learning tasks. Subsequently, we turn to artificial networks performing more challenging tasks. Specifically, we investigate how modular structures in a balanced network of excitatory and inhibitory neurons affect population dynamics and how networks in different dynamical states can be used as reservoirs in a time series prediction task. We then study how long timescales in recurrent networks learning memory tasks can be created via distinct (cellular and network level) mechanisms. We identify different training curricula that drive the network to utilize these different mechanisms and find differences in performance. Finally, using the same memory tasks, we show how a neural growth curriculum that imposes a task-specific, modular structure to the network outperforms conventional training methods.

In summary, our findings suggest that structural constraints on biological and artificial networks can significantly affect their ability to learn. Using a variety of settings, tasks and learning mechanisms, we have demonstrated that the evolution and function of local learning mechanisms can be effectively driven by biophysical constraints on the network structure. Moreover, we have shown how insights drawn from the impact of structural constraints on local learning can generalize to artificial systems, leading to improvements in performance, robustness and generalizability.

Kurzfassung

Biologische neuronale Netze zeichnen sich durch eine enorme Komplexität aus, die sich in hochkomplexen und spezifischen Strukturen manifestiert, die während der Entwicklung entstehen. Von den zahlreichen Ionenkanälen und komplexen subzellulären biophysikalischen Prozessen, die die Dynamik einzelner Neuronen steuern, über die komplexe Struktur des Dendritenbaums bis hin zu den hochspezifischen Konnektivitätsmustern zwischen Hunderten verschiedener Neuronentypen wurden biologische Netze im Laufe von Millionen Jahren der Evolution für die Informationsverarbeitung optimiert. Zusätzlich zu diesen komplexen Strukturen verändert eine Vielzahl von Anpassungsmechanismen mit Präzision und Effizienz die Stärke der neuronalen Verbindungen, um neue Informationen aufzunehmen und das Verhalten an veränderte Bedingungen anzupassen. Im Gegensatz zu künstlichen Netzen, bei denen der Lernprozess selbst der einzige Ursprung von Struktur und Funktion ist, findet das Lernen in biologischen Netzen jedoch nur innerhalb der Beschränkungen statt, die durch diese komplexen Strukturen auf Zell- und Netzebene vorgegeben sind.

In dieser Arbeit wird untersucht, welchen Einfluss diese strukturellen Beschränkungen auf den Lernprozess biologischer neuronaler Netze haben. Wir beginnen mit der Untersuchung der evolutionären Ursprünge der synaptischen Plastizität und der Frage, inwieweit verschiedene Aspekte der Netzwerk- und Aufgabenstruktur die Form der entwickelten Plastizitätsregeln beeinflussen können. Mit einem stärkeren Fokus auf biologische Details untersuchen wir, wie strukturelle Merkmale biologischer Netzwerke, einschließlich der Topologie eines Netzwerks sowie biophysikalischer Eigenschaften einzelner Neuronen wie komplexer dendritischer Strukturen, einen bedeutenden Einfluss auf die Fähigkeit lokaler synaptischer Plastizitätsmechanismen zur Durchführung einfacher unbeaufsichtigter Lernaufgaben ausüben. Anschließend wenden wir uns künstlichen Netzwerken zu, die anspruchsvollere Aufgaben erfüllen. Insbesondere untersuchen wir, wie modulare Strukturen in einem ausgewogenen Netzwerk aus exzitatorischen und inhibitorischen Neuronen die Populationsdynamik beeinflussen und wie Netzwerke in verschiedenen dynamischen Zuständen als Reservoir für eine Zeitreihenvorhersage verwendet werden können. Weiterhin untersuchen wir, wie lange Zeitskalen in rekurrenten Netzwerken, die Gedächtnisaufgaben lernen, durch unterschiedliche Mechanismen (auf zellulärer und Netzwerkebene) erzeugt werden können. Wir identifizieren verschiedene Trainingscurricula, die das Netzwerk dazu bringen, diese verschiedenen Mechanismen zu nutzen, und finden Unterschiede in der Leistung. Schließlich zeigen wir anhand derselben Gedächtnisaufgaben, wie ein neuronales Wachstumscurriculum, das dem Netzwerk eine aufgabenspezifische, modulare Struktur auferlegt, herkömmliche Trainingsmethoden übertrifft.

Zusammenfassend deuten unsere Ergebnisse darauf hin, dass strukturelle Beschränkungen biologischer und künstlicher Netzwerke deren Lernfähigkeit erheblich beeinflussen können. Anhand einer Vielzahl von Einstellungen, Aufgaben und Lernmechanismen

zeigen wir, dass die Entwicklung und Funktion lokaler Lernmechanismen effektiv durch biophysikalische Beschränkungen der Netzwerkstruktur gesteuert werden kann. Darüber hinaus zeigen wir, wie sich die Erkenntnisse aus den Auswirkungen struktureller Beschränkungen auf lokales Lernen auf künstliche Systeme übertragen lassen, was zu Verbesserungen bei Leistung, Robustheit und Verallgemeinerbarkeit führt.

Acknowledgements

The years I have spent working in Tübingen have been some of the happiest of my life, largely thanks to wonderful people who have supported me throughout my PhD. I was extremely fortunate to have a supportive, erudite, funny and kind supervisor who encouraged and motivated me through good and difficult times. Anna Levina has been a great influence on me, both scientifically and personally, and I am grateful for her unwavering support over the last years.

Oleg Vinogradov is one of the very few true polymaths I know and one of the nicest, most empathetic and fun people I had the pleasure to interact with. I would like to thank him for many interesting discussions on science, politics and literature. Likewise, I am grateful to Victor Buendía for his mentorship and friendship, as well as for his - occasionally genuinely superhuman - efforts in finishing our common projects. Sina Khajehabdollahi has been my friend, scientific collaborator and travel companion during the last four years. I have learned a lot from him and have benefited from his company both scientifically and personally. Tim Schäfer has been an ideal colleague and collaborator, and I am grateful to him for his help and support over the last years. Finally, I would like to thank Roxana Zeraati, the true Muse of Neuroscience, whose kindness, resolve and dedication to science has been a great inspiration for me. Working with her has been an honour and a pleasure.

I am grateful to Mani Hamidi for many great discussions, introducing me to Persian food and encouraging me to pursue diverse interests. I aspire to match his energy, motivation and scientific curiosity. Furthermore, I would like to thank Aaron Spieler for a wonderful collaboration and many interesting conversations on science, life and everything else. Additionally, despite our short overlap in Tübingen, I am happy to have met Tanguy Fardet, whose moral conviction has been an inspiration. I am also grateful to Miriam Bautista, Alina Krause, Daniel Gramling and Joao Barretto-Bittar for making the last months of my PhD much more interesting and fun than I ever hoped this time would be.

During the last years Zina Stefanidi has been a constant source of positivity, an occasional therapist and great friend. I am grateful for her support. I would also like to thank Auguste Schulz, whose encouragement and advice motivated me in critical moments. Xin Sui, David Nagy, Csenge Fráter, Ryutaro Uchiyama and Alex Witt made my time in Tübingen much more joyful and interesting than it would have been without them. The time spent in their company has been stimulating, memorable and fun. I would also like to thank Rama Krishna Kandukuri for introducing me to everything Indian and for being an ideal flatmate for the last three years.

I would like to thank Charley Wu for his support, advice and encouragement. Also, I have greatly enjoyed my collaboration with Alex Bird, Hermann Cuntz and Peter Jedlicka, and I am very grateful for having the opportunity to work with them.

Throughout this time, Laura Ritter was my compass, my best friend and greatest

source of inspiration. I cannot thank her enough for supporting me, keeping me sane and making me happy over the past years. Finally, I would like to thank my mother for her endless support throughout my studies. Nothing I have done would be possible without her patience and encouragement.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Thesis Structure	2
1.3	Publications and Contributions	3
1.4	Publications not included in this thesis	5
2	Background on Biological and Artificial Learning	6
2.1	Long-term Synaptic Plasticity	6
2.2	Biologically-inspired Learning in Artificial Systems	7
2.3	Structure in Biological Neural Networks	9
3	Outline of Methods and Results	12
3.1	Emergence of diverse synaptic learning rules	12
3.2	Biological structures as inductive biases for synaptic plasticity	15
3.3	Biological Structure as an inductive bias in Artificial Networks	19
3.4	Conclusion	23
4	When is it useful to learn ?	24
4.1	Motivation	24
4.2	Summary of results	24
4.3	Discussion	25
5	Co-evolution of network structure and synaptic plasticity	26
5.1	Motivation	26
5.2	Summary of results	26
5.3	Discussion	28
6	Dendritic structures control synaptic learning	29
6.1	Motivation	29
6.2	Summary of results	29
6.3	Discussion	30
7	Structural connectivity influences synaptic plasticity	31
7.1	Motivation	31
7.2	Summary of results	31
7.3	Discussion	32
8	Reservoir computing with structured E/I networks	33
8.1	Motivation	33
8.2	Summary of results	33

8.3	Discussion	34
9	Task structure drives the emergence of long timescales in RNNs	35
9.1	Motivation	35
9.2	Summary of results	35
9.3	Discussion	36
10	Modular Growth of Hierarchical Networks	37
10.1	Motivation	37
10.2	Summary of results	37
10.3	Discussion	38
11	Outlook	39
11.1	Meta-learning Plasticity Rules and Network Structure	39
11.2	Introducing greater Biological Plausibility in the modeling of Learning .	40
11.3	Developing Artificial Learning Systems with Biological Priors	41
	Bibliography	44
A	Environmental variability and network structure determine the optimal plasticity mechanisms in embodied agents	68
B	Network bottlenecks and task structure control the evolution of interpretable learning rules in a foraging agent	79
C	Dendritic nonlinearities and synapse-type specific input clustering enable the development of input selectivity in diverse settings	93
D	Structural influences on synaptic plasticity: The role of presynaptic connectivity in the emergence of E/I co-tuning	103
E	Modularity in E/I networks controls the dynamical regime and optimizes computational capabilities	144
F	Emergent mechanisms for long timescales depend on training curriculum and affect performance in memory tasks	172
G	Modular Growth of Hierarchical Networks: Efficient, General, and Robust Curriculum Learning	201

1 Introduction

1.1 Motivation

The concept of an inductive bias - the set of built-in assumptions or structural constraints that drive an agent's learning process - is fundamental to the design of artificial neural networks (ANNs) and machine learning systems in general (Mitchell, 1980). The identification of structural constraints that could provide an artificial neural network with the correct inductive biases is a major focus of AI research and often takes direct inspiration from different structures found in biological neural networks, including the morphology and dynamics of single neurons (Chavlis and Poirazi, 2021; Fan *et al.*, 2024) as well as different connectivity patterns found in biological networks (Zhang *et al.*, 2024).

However, while identifying biological structures that could be introduced in ANNs to direct the learning process as much as possible is a well established and rapidly growing



Figure 1.1: *Unlike untrained ANNs, newborn animal brains are optimized to perform complex computations prior to any learning.*¹

area of research, the mechanisms by which such structures influence learning in the original biological systems, remains largely obscure and has only recently become a topic of study (Conwell *et al.*, 2024; Dorrell *et al.*, 2023). Unlike ANNs, which are optimized via global, gradient-based methods (overwhelmingly backpropagation of errors for deep or recurrent networks), biological networks learn via complex learning rules that are largely local, act over a broad range of timescales and often depend on a variety of modulatory signals (Citri and Malenka, 2008; Speranza *et al.*, 2021). The extent to which structures that positively bias ANNs towards successful learning can also drive biologically plausible learning mechanisms in the same way (and inversely) is unclear.

A further difference between biological and artificial learning systems is the ultimate scope of their ability to learn. ANNs and other artificial learning models are primarily optimized to *learn*, that is, a well performing artificial learning agent is one that can go from being utterly unable to perform a task to being able to perform that task as efficiently and accurately as possible. On the contrary, a well performing biological agent is one that manages to *survive* (Fig. 1.1). Learning is clearly useful and ubiquitous among animals (Shettleworth, 1993), suggesting a strong correlation between the ability to learn and the ability to survive, but it is not an end in itself. Indeed, most animals are born with highly

¹Image source: Wikimedia Commons, public domain. Retrieved from https://commons.wikimedia.org/wiki/File:Baby_loggerhead_turtle_detailed_close_up.jpg. Credit: Steve Hillebrand, U.S. Fish and Wildlife Service.

complex nervous systems (Kaas, 2011) that are extremely well optimized to perform a vast amount of very intricate computations before any learning can occur (Tinbergen, 2020), including avoiding predators (Yilmaz and Meister, 2013) and performing other complex behaviours (Metz *et al.*, 2017; Weber and Hoekstra, 2009). Still, the interaction of these innate complex structures and the different learning processes taking place after birth remain rather poorly understood.

Thus, so far, we have argued that:

1. Simple biologically inspired constraints can significantly affect the learning process in ANNs.
2. Biological neural networks are characterized by vastly more complex (genetically encoded) structure than ANNs.
3. It is unclear how much of this structure found in biological neural networks can act as inductive bias for local learning mechanisms.

The goal of this thesis is to try to elucidate the last point by studying how network structure can influence learning in biological neural networks. We investigate the circumstances under which biological learning might arise in an evolving agent, the influence of network structure on the evolved plasticity rules and the extent to which cellular and network level constraints can influence learning in detailed simulations of biological neural networks performing unsupervised tasks. Finally, we ask whether different biological structures we identify can be introduced to artificial systems and improve their performance.

1.2 Thesis Structure

A detailed statement of contributions for each paper immediately follows this section. Then, the main part of the thesis begins in chapter 2, which gives a short introduction to relevant neurobiology and prior research in the interaction between network structure and learning mechanisms in biological and artificial networks. The next section (chapter 3) provides an outline of methods, relevant related work and the main results from the different papers, and is intended as a general description of the research forming the basis of this thesis, connecting the individual projects into a single thread.

Each of the subsequent chapters presents a summary of the motivations, main results and conclusions of the individual research papers forming the basis of the thesis. This part starts with chapters 4 and 5, which present the results published in (Giannakakis *et al.*, 2023) and (Giannakakis *et al.*, 2024a), studying the evolution of synaptic plasticity in embodied foraging agents.

The more biologically focused part of the thesis is presented in the next two chapters. Chapter 6 outlines our findings on the detailed control exerted over local learning dynamics by the shape of dendritic nonlinearities and the specificity of excitatory and

inhibitory inputs on the dendritic tree. The closely related work of (Giannakakis *et al.*, 2024b), studying how the structure of E and I recurrent connectivity can impact population dynamics and drive the emergence of input selectivity via local learning rules, is presented in chapter 7.

The last overarching theme of the thesis is covered in the next three chapters, dealing with the impact that biologically inspired constraints can have on the performance of artificial networks solving complex tasks. The impact of excitatory and inhibitory modular connectivity on a reservoir computing network is presented in chapter 8, followed by a presentation of the findings of Khajehabdollahi *et al.* (2024) on different mechanisms for developing long timescales in RNNs in chapter 9. The last research contribution of the thesis (Hamidi *et al.*, 2024), presenting a novel developmental curriculum for RNNs, is summarized in chapter 10.

Finally, the thesis closes with a discussion of general conclusions and an outlook for future research on related topics in chapter 11. The complete text of the published papers and unpublished manuscripts described in chapters 4–10 are attached after the bibliography as appendices.

1.3 Publications and Contributions

The thesis is based on the publications and manuscripts listed below. The contributions of the different authors for each project were as follows:

1. Giannakakis, E., Khajehabdollahi, S., and Levina, A. (2023). Environmental variability and network structure determine the optimal plasticity mechanisms in embodied agents. In *Artificial Life Conference Proceedings 35*, number 1, page 22
Contributions: **EG:** project conceptualization, implementation of static agents, analysis of results, literature review, writing and editing the text. **SK:** project conceptualization, implementation of moving agents, analysis of results, writing and editing the text. **AL:** project conceptualization, supervision, writing and editing the text
2. Giannakakis, E., Khajehabdollahi, S., and Levina, A. (2024). Network Bottlenecks and Task Structure Control the Evolution of Interpretable Learning Rules in a Foraging Agent. *Artificial Life*, pages 1–18.
Contributions: **EG:** project conceptualization, implementation of static agents, analysis of results, literature review, writing and editing the text. **SK:** project conceptualization, implementation of moving agents, analysis of results, editing the text. **AL:** project conceptualization, supervision, writing and editing the text
3. Giannakakis, E., Bird, A., Jedlicka, P., Cuntz, H., and Levina, A. (In Preparation). Dendritic nonlinearities and synapse-type specific input clustering enable the development of input selectivity in diverse settings.

- Contributions:** **EG:** project conceptualization, implementation of main simulations, analysis of results, literature review, writing and editing the text. **AB:** project conceptualization, compartmental model simulations, writing and editing the text. **HC:** project conceptualization, supervision, editing the text. **PJ:** project conceptualization, supervision, editing the text. **AL:** project conceptualization, supervision, writing and editing the text.
4. Giannakakis, E., Vinogradov, O., Buendía, V., and Levina, A. (2024). Structural influences on synaptic plasticity: The role of presynaptic connectivity in the emergence of e/i co-tuning. *PLOS Computational Biology*, 20(10), 1–22.
Contributions: **EG:** project conceptualization, implementation of main simulations, analysis of results, literature review, writing and editing the text. **OV:** project conceptualization, coding and implementation of Bayesian inference, writing and editing the text. **VB:** analytics for neural mass model, writing and editing the text. **AL:** project conceptualization, supervision, writing and editing the text.
 5. Giannakakis, E., Buendía, V., Vinogradov, O., and Levina, A. (In Preparation). Modularity in E/I networks controls the dynamical regime and optimizes computational capabilities.
Contributions: **EG:** project conceptualization, training of Reservoirs, writing and editing the text. **VB:** project conceptualization, analytics, writing and editing the text. **OV:** coding and implementation of simulation-based inference, writing and editing the text. **AL:** project conceptualization, supervision, writing and editing the text.
 6. Khajehabdollahi, S., Zeraati, R., Giannakakis, E., Schäfer, T. J., Martius, G., and Levina, A. (2024). Emergent mechanisms for long timescales depend on training curriculum and affect performance in memory tasks. In *The Twelfth International Conference on Learning Representations*.
Contributions: **SK:** project conceptualization, training of RNNs, writing and editing the text. **RZ:** project conceptualization, training of RNNs, timescales analysis, literature review, writing and editing the text. **EG:** project conceptualization, literature review, connectivity and timescales analysis, writing and editing the text. **TS:** training of RNNs, perturbation analysis, writing and editing the text. **GM:** writing and editing the text. **AL:** project conceptualization, supervision, writing and editing the text.
 7. Hamidi, M., Khajehabdollahi, S., Giannakakis, E., Schäfer, T. J., Levina, A., and Wu, C. M. (2024). Modular growth of hierarchical networks: Efficient, general, and robust curriculum learning. In *ALIFE 2024: Proceedings of the 2024 Artificial Life Conference*. MIT Press.

Contributions: **MH:** project conceptualization, developing and training modular RNNs, literature review, writing and editing the text. **SK:** project conceptualization, training modular RNNs, writing and editing the text. **EG:** project conceptualization, connectivity and timescales analysis, writing and editing the text. **TS:** perturbation analysis, writing and editing the text. **AL:** project conceptualization, supervision, editing the text. **CW:** project conceptualization, making figures, supervision, writing and editing the text.

1.4 Publications not included in this thesis

During the course of my doctoral studies, in addition to the different projects which form the basis for this thesis, I worked on a number of collaborative projects (Khajehabdollahi *et al.*, 2022; Fardet *et al.*, 2024; Vinogradov *et al.*, 2024; Khajehabdollahi *et al.*, 2023). While these projects are on topics largely unrelated to this thesis and my overall contribution to them was minor, working on them was both enjoyable and important in terms of learning new concepts, techniques and skills. Most importantly, working on these projects initiated the fruitful collaborations which supported the completion of many of the papers forming part of the thesis. Thus, the resulting publications are listed here for completeness:

1. Khajehabdollahi, S., Prosi, J., Giannakakis, E., Martius, G., and Levina, A. (2022). When to be critical? performance and evolvability in different regimes of neuralising agents. *Artificial Life*, 28(4), 458–478.
2. Fardet, T., Giannakakis, E., Paulun, L., and Levina, A. (2024). Revising clustering and small-worldness in brain networks. arXiv preprint arXiv:2401.15630.
3. Vinogradov, O., Giannakakis, E., Buendia, V., Uysal, B., Ron, S., Weinreb, E., Schwarz, N., Lerche, H., Moses, E., and Levina, A. (2024). Effective excitability captures network dynamics across development and phenotypes. *bioRxiv*, pages 2024–08.
4. Khajehabdollahi, S., Giannakakis, E., Buendia, V., Martius, G., and Levina, A. (2023). Locally adaptive cellular automata for goal-oriented self-organization. In *Artificial Life Conference Proceedings 35*, volume 2023, page 59.

2 Background on Biological and Artificial Learning

This chapter presents a concise introduction to the historical development of our current understanding of long-term synaptic plasticity, along with efforts to relate biological local learning mechanisms with more powerful learning algorithms like backpropagation of errors. This is followed by a brief outline of the complexity of brain structures and their potential effects on how different learning mechanisms behave.

2.1 Long-term Synaptic Plasticity

Biological learning is a complex phenomenon, involving multiple interacting processes over a diverse range of timescales. However, it is generally agreed that virtually for every organism with a nervous system, learning involves some form of *long-term synaptic plasticity*, the activity-dependent modification of the synaptic strengths between their neurons, which produces stable synaptic configurations over time (Kandel, 2001). While the basic idea of synaptic modification as the basis for biological learning dates back more than a century (Bain, 1873), the modern idea of synaptic learning originates with Donald Hebb who stipulated that a causal relation between pre- and postsynaptic firing would lead to the long-term strengthening the synaptic strength (Hebb, 2005).

Hebb's theoretical insight became the basis for experimental investigation, starting with the discovery of long-term potentiation (Bliss and Lømo, 1973) (the activity dependent strengthening of synapses) and followed by the experimental demonstration of long-term depression (Barrionuevo *et al.*, 1980; Ito *et al.*, 1982) (the activity dependent reduction of synaptic strength following a prolonged stimulus). The discovery of these fundamental components of long-term plasticity (Fig. 2.1) was followed by an explosion of experimental studies that documented synaptic long term plasticity in a variety of circuits and settings in the brain (Malenka and Bear, 2004; Nicoll, 2017; Citri and Malenka, 2008).

This increased focus on plasticity quickly revealed a previously unimaginable complexity of long-term plasticity mechanisms in the brain. In particular, observations of a variety of different learning rules, including anti-Hebbian (Koch *et al.*, 2013; Roberts and Leen, 2010), symmetric (Mishra *et al.*, 2016) and bidirectional learning rules (Fino *et al.*, 2005), acting over a variety of timescales (Bittner *et al.*, 2017; O'Donnell, 2023) revealed a much more intricate collection of learning mechanisms than classic Hebbian modification. The discovery of inhibitory plasticity (Castillo *et al.*, 2011;




Figure 2.1: An illustration of an STDP window. The timing of pre and postsynaptic spikes determines the depression or potentiation of the synaptic strength.

Hennequin *et al.*, 2017) as well as the effect of third factors, such as different neuromodulators, (Gelinas and Nguyen, 2007; Speranza *et al.*, 2021) on synaptic learning added a novel dimension that further complicated the picture.

In parallel to the explosion of experimental work revealing this unimaginable complexity of local learning mechanisms in the brain, a similar build-up of theoretical work added new insights into how these mechanisms could function. Starting with models of the development of input selectivity in the cortex (Bienenstock *et al.*, 1982; Intrator and Cooper, 1992), multiple studies have proposed different models of plasticity mechanisms, including spike timing dependent plasticity (STDP)¹ (Abbott and Nelson, 2000; Gerstner *et al.*, 1996) and reward modulated plasticity (Frémaux *et al.*, 2010; Frémaux and Gerstner, 2016) as well as theoretical results on how different plasticity mechanisms might coordinate to achieve complex learning objectives (Zenke *et al.*, 2015; Schulz *et al.*, 2021; Eckmann *et al.*, 2024).

In summary, as our understanding of long-term synaptic plasticity increases, it becomes clear that it is a collection of extremely diverse processes, which differ across brain areas and neuron types (Citri and Malenka, 2008). Moreover, complex interaction between synaptic plasticity and different regulatory and modulatory mechanisms have been shown to learn complex tasks (Wu *et al.*, 2020; Illing *et al.*, 2021; Halvagal and Zenke, 2023). However, we are still very far from a complete theory of biological learning or even a detailed understanding of how many of the tasks that even simpler organisms routinely learn can be tackled via local synaptic plasticity.

Other forms of Plasticity: In addition to synaptic long-term plasticity, biological neural networks are characterized by a wide range of additional plasticity mechanisms, including short-term plasticity (Zucker and Regehr, 2002; Regehr, 2012; Hennig, 2013) and structural plasticity (Holtmaat and Svoboda, 2009; Butz *et al.*, 2009). These mechanisms are very important for the function of biological neural networks, but since they were not studied in any of the publications forming the basis of this thesis, they will not be discussed further.

2.2 Biologically-inspired Learning in Artificial Systems

The ability of the biological neural networks to efficiently process vast amounts of information and direct well-tuned responses to complex stimuli has inspired multiple artificial imitations of neural information processing, many of whom have been hugely successful in terms of producing functional, efficient and, quite often, powerful artificial intelligence systems. These include the idea of an ANN itself (McCulloch and Pitts, 1943; Rosenblatt, 1958) which was explicitly inspired by the function of the nervous system, including recurrent (Hopfield, 1982) and deep networks (Rumelhart *et al.*, 1986; Hinton

¹A wonderfully written and detailed history of how our understanding of STDP developed can be found in Markram *et al.* (2011)

and Salakhutdinov, 2006) as well as convolutional neural networks (Fukushima, 1980; LeCun *et al.*, 1989), which drew inspiration from the way the visual cortex processes information.

However, despite there being a direct influence of the *structure* of biological neural networks on the design of artificial systems, attempts to emulate biological learning mechanisms (local, activity dependent learning rules) have usually fallen short of achieving the same performance as the standard method of training ANNs which is backpropagation of errors (Linnainmaa, 1970; Rumelhart *et al.*, 1986; Werbos, 1994). Backpropagation relies on there being a global error signal for optimizing a network’s connectivity, and modifies network connectivity according to the exact contribution of each neuron to the global error by computing the gradient of the error with respect to the weight of each connection (LeCun *et al.*, 2015). No biological learning mechanism known so far can modify network connectivity as efficiently and accurately as backpropagation. Attempts to bridge this gap in performance between backpropagation and biologically plausible local mechanisms, usually fall into one of two categories (Schmidgall *et al.*, 2024):

A first line of tackling this problem is to try and identify learning mechanisms that can achieve high performance in complex tasks, while at the same time maintaining the biological constraints under which biological learning operates (locality in space and time). The identification of such learning rules can happen via different methods, including gradient descent (Miconi *et al.*, 2018; Schmidgall *et al.*, 2021) or evolutionary optimization (Najarro and Risi, 2020; Berg Palm *et al.*, 2020). Despite the high performance such methods can achieve in different tasks (Duan *et al.*, 2023; Tyulmankov *et al.*, 2022), their biological plausibility has been called into question, largely due to the fact that they usually optimize distinct learning rules for individual synapses, something that is unlikely to happen in biological networks.

A different approach to the problem of biologically plausible learning that compares with backpropagation in terms of performance is to develop learning algorithms that emulate the basic properties of backpropagation or other gradient-based learning algorithms within the constraints of biological structure (Betti *et al.*, 2018; Richards and Kording, 2023). The most prominent of these methods are feedback alignment (Lillicrap *et al.*, 2016; Nøkland, 2016; Moskovitz *et al.*, 2018), which used a random feedback weight matrix for credit assignment and eligibility propagation and its extension in spiking networks, eligibility propagation (Bellec *et al.*, 2020). More biologically plausible versions of these methods, incorporating features of biological networks such as bursting activity (Payeur *et al.*, 2021), dendritic structures (Sacramento *et al.*, 2018; Max *et al.*, 2024) or local neuromodulators (Liu *et al.*, 2021, 2022) can perform even better in benchmark tasks, although they usually fall short of achieving the same performance as classical backpropagation.

In summary, biologically plausible training mechanisms tend to underperform the gradient-based optimization methods used for training ANNs. However, the development of learning algorithms that use various structural features of biological networks (dendrites, neuromodulators, etc) can partially improve performance.

2.3 Structure in Biological Neural Networks

A potential explanation of why artificial emulations of biological learning tend to underperform gradient-based methods, hinted at by the boost in performance in biologically plausible learning algorithms that incorporate biological structures, may be found in the nature of the medium in which the two kinds of learning take place. Specifically, unlike artificial neural networks, which consist of simple, identical units with largely uniform connectivity structures, biological networks are made up of hundreds of highly distinct neuron types, interacting via intractably complex connectivity patterns and receiving multiple modulatory signals over a range of timescales. The potential impact of these factors on learning is discussed in this section.

2.3.1 Single Neuron Structure and Diversity

The human brain contains approximately 85 billion neurons, which perform the bulk of information processing required for survival (Herculano-Houzel, 2012). Unlike the simple processing units of artificial neural networks, usually performing a weighted non-linear summation of incoming inputs (Nwankpa *et al.*, 2018), biological neurons are extremely complex structures with highly non-trivial dynamics (Gerstner and Kistler, 2002).




Figure 2.2: An outline of the highly complex structure of a typical cortical pyramidal neuron.

The membrane dynamics of individual neurons are known to depend on a variety of ion channels, acting on different timescales, (Hodgkin and Huxley, 1952; Bean, 2007) in addition to energetic constraints (Skou, 1990; Howarth *et al.*, 2012) and other subcellular processes (Verkhatsky and Petersen, 1998; Kilpatrick *et al.*, 2016). Moreover, biological neurons are characterized by a complex spatial structure, the most prominent feature of which is the dendritic tree (Fig. 2.2), often extending far away from the soma and enabling the neuron to receive inputs from multiple sources (Larkman, 1991). As our understanding of cellular function increases, it becomes clearer that dendrites are not simply passive extensions of the soma, but are characterized by complex dynamics (Hausser *et al.*, 2000), such as dendritic plateaus and spikes (Augustinaite *et al.*, 2014; Palmer *et al.*, 2014). This complexity of single neuron dynamics had led to a realization that biological neurons are not really analogous to the processing units of ANNs but can be better approximated as small neural networks that perform highly nonlinear processing of incoming information (Poirazi *et al.*, 2003; Spieler *et al.*, 2023).

In addition to the complexity of each individual neuron, biological neural networks

are also characterized by a great diversity of neuron types. It is estimated that there are hundreds of neuron types in the brain, each with its own unique structure, dynamics and connectivity patterns (Stevens, 1998; Masland, 2004). These can be broadly separated into excitatory and inhibitory neurons (having a depolarizing and hyperpolarizing effect on postsynaptic neurons, respectively), but these two groups can be additionally divided into multiple neuron subtypes with significant morphological and functional variability (Hattori *et al.*, 2017b).

All this complexity has a significant impact on the learning capabilities of biological neural networks. Dendritic structures and their interaction with the soma have long been associated with efficient learning in experimental (Yaeger *et al.*, 2024) and theoretical studies (Mikulasch *et al.*, 2021; Wilmes and Clopath, 2023). Additionally, different neuron types are usually characterized by distinct learning rules (Larsen and Sjöström, 2015; Hennequin *et al.*, 2017), which allows for much more sophisticated learning dynamics than a single update rule for all connections.

2.3.2 Specific Connectivity

Perhaps even more that the complexity of individual neurons, the most prominent difference in structure between biological and artificial neural networks is the staggering difference between the relatively uniform connectivity of the latter, which contrasts with the highly inhomogeneous and complex connectivity patterns found in biological brain networks.

Biological neural networks are physical objects, whose connectivity is largely determined by the constraints of the three-dimensional space they inhabit. Thus, neurons tend to connect with other neurons that are close to them in space (Rossi *et al.*, 2020; Ding *et al.*, 2023), share similar inputs (Yoshimura and Callaway, 2005) or targets (Brown and Hestrin, 2009). At the same time, long range connections are also well documented (Markov *et al.*, 2013) and important for neural information processing (Deco *et al.*, 2021; Wang *et al.*, 2023). Additionally, neurons tend to connect with other neurons that encode for similar features, forming functional assemblies of strongly interconnected neurons with similar activation patterns (Miehl *et al.*, 2023).

Different neuron types are also well known to follow distinct connectivity patterns within and between brain areas (Udvary *et al.*, 2022; Rossi *et al.*, 2020). For example, excitatory connections tend to be more far-reaching, transmitting information between different areas while inhibitory connections




Figure 2.3: An illustration of the connectivity of a cortical microcircuit of pyramidal neurons (Pyr) and different interneurons (PV, SOM and VIP).

largely remain within the same area (Fino *et al.*, 2013; Watts and Thomson, 2005), albeit with significant exceptions (Tomioka *et al.*, 2005; Higo *et al.*, 2009). In an example of even greater specificity, different types of cortical interneurons tend to be very particular with respect to their targets (Fig. 2.3), with Somatostatin (SOM)-positive interneurons targeting the distal dendrites and Parvalbumin (PV)-positive interneurons targeting the soma and proximal dendrites of pyramidal neurons, while Vasoactive intestinal peptide (VIP)-expressing interneurons target other interneurons (Tremblay *et al.*, 2016; Rudy *et al.*, 2011). Other brain areas are also characterized by highly specific connectivity between different neuron types (Booker and Vida, 2018; Jörntell, 2017).

These highly non-trivial connectivity structures that characterize biological neural networks are known to generate rich dynamics (Litwin-Kumar and Doiron, 2012) and enable complex computations (Lagzi and Rotter, 2015). Moreover, neuron-type specific connectivity structures have been shown to interact with different plasticity mechanisms and enabling the learning of complex tasks (Hertäg and Clopath, 2022; Wilmes and Clopath, 2019; Lakshminarasimhan *et al.*, 2024).

2.3.3 Modulation

Finally, the signalling mechanisms in biological networks tend to be much more complex and diverse than a single error signal that is used to train artificial networks. The mammalian brain is modulated by dozens of different neuromodulators (Nadim and Bucher, 2014), many of whom are associated with distinct aspects of learning. Among these, dopamine is most prominently known to encode reward prediction errors (Schultz *et al.*, 1997), encouraging analogies between the brain and artificial reinforcement learning algorithms (Glimcher, 2011). Still, other neuromodulators have also been associated with learning. In particular, acetylcholine is known to regulate the initiation of plasticity in neural circuits, (Richards and Kording, 2023; Partridge *et al.*, 2002) while modulators such as noradrenaline and serotonin also appear to regulate the learning of some complex tasks (Doya, 2002).

Although biological networks clearly rely significantly on such modulatory mechanisms for regulating neural activity and learning, their exact functions are still poorly understood. Several studies using neuromodulation-inspired mechanisms in artificial networks have recently shown how they could potentially improve the performance of biologically plausible learning algorithms (Liu *et al.*, 2021, 2022; Lee *et al.*, 2024b), suggesting that this type of diverse signaling regulating different aspects of the learning process potentially plays a significant role in boosting the learning capabilities of biological systems.

3 Outline of Methods and Results

The main purpose of artificial neural networks is to learn, as efficiently and accurately as possible, to perform a specific task *de novo*. On the contrary, biological neural networks develop into well-structured, almost fully functional systems prior to any exposure to sensory inputs, and only use different learning mechanisms to improve their performance in tasks they are generally already well-tuned to solve.

The main focus of this thesis is the interaction between such biological structures that emerge early during development and remain largely constant throughout an organism's lifetime and the various learning processes that ubiquitously take place in biological systems. We investigate how synaptic plasticity mechanisms can emerge and co-evolve with structured motor networks, and how cellular and network structures found in biological circuits can regulate the effects of synaptic plasticity. Finally, we study how many of these biological structures can be incorporated in artificial networks solving complex tasks and boost their performance.

In this section, we will give an outline of the motivations behind the main research directions presented in this thesis, a general introduction into the most important methods we used and a summary of our main results. A more detailed description of each individual project is presented in the next chapter.

3.1 Emergence of diverse synaptic learning rules

Despite huge morphological and behavioural variability among living organisms, virtually all living creatures display some form of adaptation to changing environmental conditions. In simple organisms, this ability might be limited to simple avoidance behaviours or growing near developing food sources (Dussutour, 2021; Wong and Rankin, 2019; Gourgou *et al.*, 2021). For more complex life forms, including for animals that possess a simple nervous system (Roberts and Glanzman, 2003), learning via modification of the synaptic strengths between neurons, i.e., via synaptic plasticity, is the primary form of adapting behaviour to changing environmental conditions (Citri and Malenka, 2008).

Still, despite its near-universal presence across the animal kingdom, the evolutionary origins of synaptic plasticity remain poorly elucidated (Ovsepian *et al.*, 2020). There is a consensus on it partially arising as a mechanism for dealing with unpredictable environmental variability (Papini and Torres, 2017), but the exact evolutionary pressures that different types of environmental variability exert on the emergence of synaptic learning remain largely speculative and have been only studied in relatively high-level settings (Lange and Sprekeler, 2022).

Additionally, the interaction of different forms of plasticity with brain structures that overwhelmingly emerge during development and remain largely constant during an or-

ganism's lifetime is poorly understood (Ecker *et al.*, 2023). It is known that most complex organisms are born with highly optimized brains, that can perform a large amount of intricate computations before any plasticity mechanism could optimize them. This suggests that a two-way adaptation is taking place during evolution, with plasticity mechanisms evolving to effectively optimize these pre-existing structures during an organism's lifetime and the network structure also adapting to accommodate efficient learning via local synaptic modification.

The first part of this thesis is thus dedicated to the investigation of the environmental conditions that would produce the appropriate evolutionary pressures for synaptic plasticity to develop in a simple nervous system as well as the impact that different structural features of that nervous system would have on the evolved plasticity rules. This section is based on two published papers (Giannakakis *et al.*, 2023, 2024a), which are presented in more detail in chapters 4 and 5.

Taking inspiration from the olfactory system of *Drosophila Melanogaster*, which is known to be responsible for reward-modulated associative learning (De Belle and Heisenberg, 1994; McGuire *et al.*, 2001; Aso *et al.*, 2012; Rajagopalan *et al.*, 2023), we developed simple embodied agents consisting of a plastic sensory network and a non-plastic motor network (Fig. 3.1a). The idea behind this choice was to capture, in as simple a model system as possible, the interaction between the more conserved¹ motor structures (Clark *et al.*, 2018) of drosophila and the highly adaptable olfactory areas that use reward-modulated synaptic plasticity (Bennett *et al.*, 2021) to learn novel environmental stimuli (Masuda-Nakagawa *et al.*, 2009; Hattori *et al.*, 2017a).

These simple learning agents are embodied and embedded in a 2D surface scattered with food particles of different values (Fig. 3.1b). The agents' capacity to learn the values of the food particles and consume those with a positive value while avoiding the ones with a negative value is used to evaluate their genetic fitness and evolve the connectivity of the motor network as well as the parameters of the plasticity of the sensory network (Fig. 3.1c). While different aspects of our setting have been used in past studies (Yerushalmi and Teicher, 2008; Khajehabdollahi *et al.*, 2022; Yaman *et al.*, 2021a), the combination of embodied plastic agents, evolutionary optimization, and a complex task (foraging in a changing environment) is novel and makes the setting highly relevant as an approximation of the evolutionary pressures that organisms face in natural settings.

Using this experimental setup, we study how different environmental conditions and types of variability promote or suppress the emergence of synaptic learning. Interestingly, we identify some settings in which the best-performing agents do not evolve the ability to learn at all, relying on their structural priors for interacting with their environment. This suggests a complementary function for learning, endowing an already well-functioning motor/sensory system with adaptive abilities only when environmental

¹There is some evidence of plasticity in the drosophila motor system (Fushiki *et al.*, 2013; Fernandes *et al.*, 2023), making the underlying biology somewhat complicated and our study a rather rough approximation of the nervous system of any specific organism.

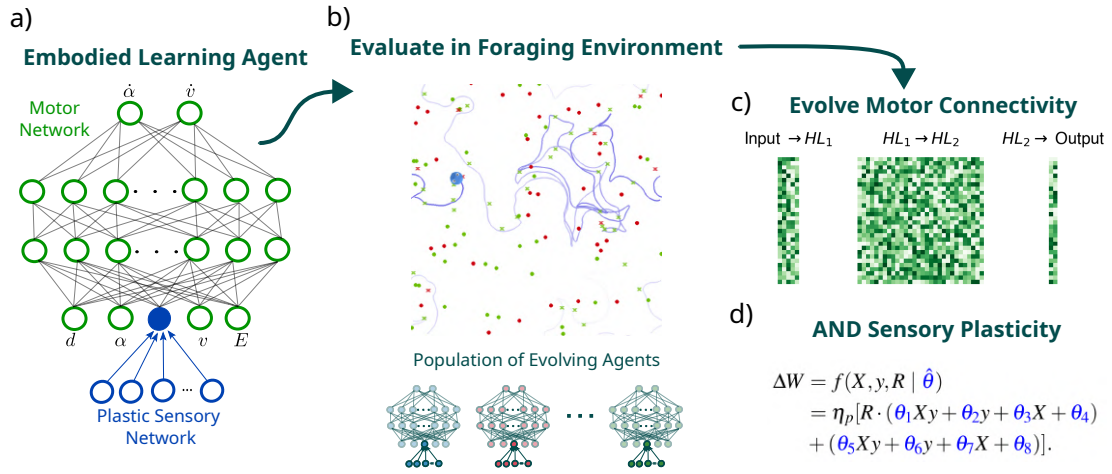


Figure 3.1: *Evolving Plasticity Rules in Embodied Agents.* **a.** The embodied agents consist of a single-layer plastic sensory network and a non-plastic motor network with 2 hidden layers. **b.** A population of agents is placed in a 2D environment and evaluated on their ability to learn the distribution of food particles and navigate towards the ones with the higher values. The agents’ performance is then used to evolve both (c.) the connectivity of the motor network and (d.) the parameters of the plasticity rule of the sensory network.

conditions make such an adaptation beneficial (See Chapter 4 for details).

Multiple studies have used meta-learning methods (Ramesh *et al.*, 2024; Confavreux *et al.*, 2023; Yaman *et al.*, 2021b), including evolutionary algorithms, (Confavreux *et al.*, 2020; Jordan *et al.*, 2021; Confavreux *et al.*, 2024) to optimize local plasticity rules for specific tasks, often finding a huge diversity of forms that is hard to interpret. Thus, we then turn to the study of the learning rules that evolve in our embodied agents. After identifying a simple regularization mechanism that reduces the initial diversity of evolved rules, allowing us to extract interpretable forms, we find that the evolutionary trajectory of the learning rules is highly dependent on different details of the network. In particular, we find that the readout neuron’s activation function, the plasticity’s normalization mechanism as well as the structure of the task itself, all affect the form of the evolved learning rules. Moreover, we find that the different rules that evolve under these different conditions achieve significantly different levels of performance and generalizability (See Chapter 5 for details).

These findings suggest that structural features of a network can significantly impact its ability to learn via synaptic plasticity. This interdependence between these two aspects (learning rules and structural features) motivates a holistic approach to the study of biological learning that does not exclusively focus on the form of the plasticity mechanism but also takes into account the (potentially critically important) local and network-level structures that may control the learning process in non-trivial ways.

3.2 Biological structures as inductive biases for synaptic plasticity

Having gained an intuition of the extent to which different aspects of a network's structure can drive the evolution of distinct plasticity rules in our simple embodied agents, we next investigate how these principles might manifest in more realistic biological settings. The embodied agents from chapters 4 and 5 are useful as minimal models of a simple learning system and can help us with deriving general, proof of concept, insights into learning, but they are far from even approximating the complex dynamics and intricacies of an actual biological neural network.

Thus, for the next section of this dissertation, we turn to biological networks and focus on the interaction between structural constraints and synaptic plasticity. However, unlike the previous section, where we examined how optimal plasticity rules might evolve to fit different structural network features, we now adopt the opposite approach and examine how a given plasticity protocol can be affected by changes in neuronal and network-level structures. This change in focus was in part motivated by experimental evidence on the high degree of evolutionary conservation of synaptic structure (and consequently synaptic plasticity mechanisms) in biological neural networks (Glanzman, 2010), combined with faster evolutionary change in their specific circuitry (Finlay *et al.*, 1998).

General synaptic morphology and plasticity mechanisms seem to be extremely well conserved across species³, with many basic functions conserved even between vertebrates and invertebrates (Ryan and Grant, 2009; Kandel, 2001; Glanzman, 2010), albeit with some significant differences between the two (Emes *et al.*, 2008; Emes and Grant, 2012). Moreover, when looking at related organisms with vast differences in behavioural sophistication (including the ability to learn complex tasks), it becomes clear that the underlying synaptic structures are too similar to reliably account for different learning capabilities. For example, the fundamental properties of LTP and LTD along with several other aspects of synaptic plasticity are highly conserved across rodents and primates



Figure 3.2: *The charming but relatively unsophisticated mollusc Aplysia is characterized by a similar level of synaptic protein diversity as the taxonomically related but behaviourally much more complex Octopus (Orvis et al., 2022).*²

²Image source: Wikimedia Commons, Public domain. Retrieved from https://commons.wikimedia.org/wiki/File:Aplysia_californica_NHGRI-79108.jpg.

³As with most such principles in biology, there are some exceptions to this rule (Beed *et al.*, 2020). Still despite examples to the contrary, the scientific consensus seems to support a relatively high degree of evolutionary conservation in plasticity mechanisms across species

(Huang *et al.*, 2014). On the same direction, a recent study (Orvis *et al.*, 2022) examining the synaptic proteins of the rather simple (20,000 neurons) mollusc *Aplysia* (Fig. 3.2) found a similar diversity of synaptic proteins to the *Octopus*, by far the most intelligent invertebrate, known for its remarkable capability to learn complex tasks (Richter *et al.*, 2016; Grasso and Basil, 2009).

On the contrary, when looking at structural features such as the cellular properties of individual neurons (Jacobs *et al.*, 2014; Nomura *et al.*, 2018; Marchetto *et al.*, 2019) and even more prominently, specific connectivity structures (Tierney, 1995; Katz and Harris-Warrick, 1999; Balaban, 1997), including structures largely responsible for learning (Ellis *et al.*, 2024), we often find a significant diversity even among closely related species. This suggests that, notwithstanding the generally high degree of evolutionary conservation of neural structures (Kavanau, 1990), it is more plausible that structural features of a network would adapt to accommodate the function of a learning rule rather than the other way around. With this in mind, we opt to study how, given a fixed plasticity protocol, changes in different biophysical structures can affect synaptic learning in biological networks performing an unsupervised task.

3.2.1 A Classic Unsupervised Task for Biological Networks: Developing Input Selectivity via E & I Plasticity

Neurons in the visual cortex have been long known to display selectivity to particular inputs. A classic early experimental study showed that upon presenting an animal with gratings of different orientations, different neurons show strong preference (i.e. significantly increase their firing rate) for specific orientations (Hubel and Wiesel, 1962). The development of this selectivity in neurons that do not receive direct sensory input was assumed to develop via synaptic plasticity.

The celebrated experiments of Wiesel and Hubel (1963) and Blakemore and Cooper (1970), demonstrated that the development of orientation selectivity happens in the first weeks of life and is dependent on neural activity, conclusively demonstrating that some form of synaptic plasticity is involved. A plasticity protocol that could create stable orientation selectivity was later proposed by Bienenstock *et al.* (1982) with improvements introduced in later years (Intrator and Cooper, 1992; Law and Cooper, 1994).

The discovery of input selectivity in cortical inhibitory neurons (Niell and Stryker, 2008; Sohya *et al.*, 2007; Najafi *et al.*, 2020), as well as the experimental demonstration of inhibitory plasticity (Komatsu and Iwakiri, 1993; Woodin *et al.*, 2003; Holmgren and Zilberter, 2001), complicated this picture and initiated a search for a plasticity protocol that simultaneously create both inhibitory selectivity and detailed⁴ E/I balance (Vogels

⁴A neuron is balanced when the incoming E and I currents cancel out on average (Okun and Lampl, 2008). The time resolution on which this cancelling out takes place determines the tightness of the balance, with global or loose balance indicating larger time windows while detailed/tight balance indicating very small time windows. The relative benefits of loose/global and tight/detailed E/I balance

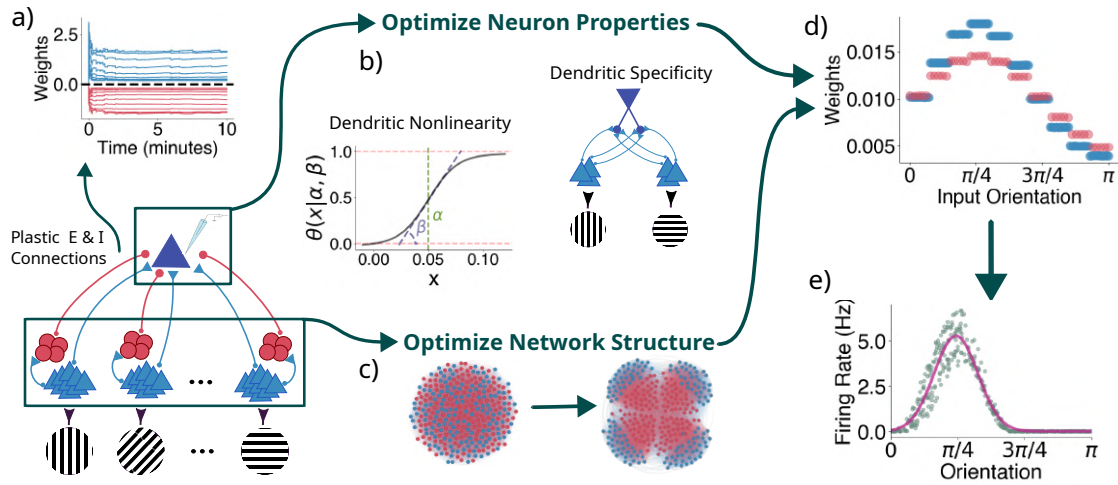


Figure 3.3: *Inferring biological structures that drive the emergence of input selectivity via E/I plasticity. a.* A network of excitatory and inhibitory neurons, grouped into multiple populations, each of whom is selective to a different stimulus, is feedforwardly connected to a single postsynaptic neuron with plastic E and I connections. We study how the optimization of either (b.) the dendritic structures of the postsynaptic neuron or (c.) the connectivity structure of the presynaptic network, can enable the development of tuned E&I feedforward connectivity (d.) that generates input selectivity (e.) in the postsynaptic neuron.

and Abbott, 2009) in a neuron receiving tuned E and I inputs.

The first step into resolving this problem was made by Vogels *et al.* (2011), who demonstrated how a simple homeostatic inhibitory learning rule can produce detailed E/I balance by matching the inhibitory connectivity to non-plastic, tuned excitatory connections. This was followed by a study of interacting E and I plasticity (Clopath *et al.*, 2016) creating stable input selectivity in certain settings. Finally, a more general protocol was proposed by (Eckmann *et al.*, 2024), who demonstrated that simultaneous, synapse-type specific competitive normalization of the E and I synapses, enables the stable learning of orientation selectivity along with the creation of detailed E/I balance (Fig. 3.3a).

3.2.2 Inferring structural inductive biases in biological networks

Our contribution to this line of research is to study how structural biophysical and topological features of the network can interact with the plasticity protocol and influence

have been a topic of extensive study (Denève and Machens, 2016) and remain a topic of debate, but overall, detailed/tight balance seems to be beneficial for gating multiple signals (Vogels and Abbott, 2009) and efficient coding (Zhou and Yu, 2018).

learning outcomes. The state-of-the-art plasticity protocol for simultaneously creating E/I balance and input selectivity (Eckmann *et al.*, 2024) heavily relies on a beneficial statistical structure of the inputs to the post-synaptic neuron. Our aim is to identify cellular and network structures that could enable these same plasticity mechanisms to function well in settings with potentially suboptimal input structures.

The experience of co-evolving a network’s structure and learning rules from chapters 4 and 5 suggests that while the influence of structure on learning is strong, it is also clearly stochastic. This, in combination with the well-known noisiness of synaptic learning (Stein *et al.*, 2005; Faisal *et al.*, 2008), motivated us to opt for probabilistic methods for identifying structures that can drive synaptic plasticity. Specifically, we used approximate Bayesian computation (ABC) (Beaumont, 2019) and simulation-based inference (SBI) (Tejero-Cantero *et al.*, 2020) to estimate the probability distributions of different structural features generating the desired plasticity outcomes. This approach provides us with a full posterior distribution of the way different structural parameters influence learning and the interactions between them. Thus, we gain a more complete understanding of the influence structure has on the dynamics of plasticity than by merely optimizing the structure, since we unveil dependencies and correlations between different parameters that would otherwise have remained obscure.

We separately investigated how changes in the postsynaptic neuron’s dendritic structure (Fig. 3.3b) and the presynaptic network’s connectivity (Fig. 3.3c) can impact learning. At first, inspired by findings on the influence of dendritic nonlinearities and the distribution of synapses along the dendritic tree on orientation selectivity (Wilson *et al.*, 2016; Iacaruso *et al.*, 2017; Weiler *et al.*, 2022), we extended the model of Eckmann *et al.* (2024) to include a simple dendritic structure and investigated how the shape of the dendritic nonlinearity and the specificity of E/I synapses on the dendritic tree influences learning outcomes. We found that the introduction of dendrites makes the original model much more robust to noise. We also demonstrated how the shape of the nonlinearity and the dendritic specificity can reliably control the sharpness of the resulting orientation tuning. Details of these findings are presented in chapter 6.

Subsequently, our attention turned to the structure of the presynaptic network and in particular on how connectivity can impact the statistics of the population activity that influence learning outcomes. Using a spiking network with excitatory (Pfister and Gerstner, 2006) and inhibitory (Vogels *et al.*, 2011) plasticity along with competitive normalization (Eckmann *et al.*, 2024), we show that the introduction of noise and recurrent connectivity among the presynaptic neurons destroys the ability of plasticity to generate input selectivity. We then demonstrate how changes in the connectivity structure of the presynaptic network, in particular the introduction of synapse-type specific neuronal assemblies (Miehl *et al.*, 2023) of varying strengths can restore the optimal statistical structure in the population activity and enable the plasticity protocol to generate E/I weight balance (Fig. 3.3d) and input selectivity (Fig. 3.3e) in the post-synaptic neuron. These findings, presented in the publication Giannakakis *et al.* (2024b), are summarized in chapter 7.

Our results suggest that simple biological structures can exert a significant influence on the activity of neural networks and strongly impact the behaviour of local learning mechanisms. Using low-dimensional parametrizations of cellular and network structure, we have demonstrated the precise control such structures can exert on learning dynamics, effectively regulating the emergence of input selectivity. These findings suggest that the study of specific plasticity protocols that ignores the surrounding biophysical structures is unlikely to provide a complete understanding of learning in biological system and argues for a greater focus on structural features as essential regulators of synaptic plasticity.

3.3 Biological Structure as an inductive bias in Artificial Networks

While unsupervised learning via Hebbian-like synaptic plasticity is well documented across different brain areas in multiple species (Citri and Malenka, 2008) and it undoubtedly performs a crucial function, it is to some extent constrained by its simplicity. Hebbian rules essentially learn the underlying statistical structure of inputs, often performing an approximation of principal component analysis (Oja, 1992; Pehlevan and Chklovskii, 2015). While this is a necessary component of many learning algorithms, it is insufficient to explain the full range of learning processes taking place in biological neural networks. Thus, for the next and final part of the thesis, we seek to study whether the structural features that we identified as crucial in regulating unsupervised Hebbian learning can also impact the ability of neural networks to perform complex, supervised learning tasks.

Supervised learning in the brain has been studied extensively (Knudsen, 1994). This includes neuromodulatory reward-based reinforcement learning (Schultz *et al.*, 1997), cortical learning driven by feedback signals originating in subcortical areas, (Roelfsema *et al.*, 2010) as well as error-driven learning in the cerebellum (Guo *et al.*, 2014; Popa and Ebner, 2019). However, while these learning functions are almost certainly performed via local learning rules, the exact plasticity protocols that can learn complex tasks using biologically plausible learning are still somewhat unclear. Thus, for this final part of the thesis, we eschew strict biological plausibility when it comes to the learning mechanisms and study how the introduction of biological structures in artificial networks trained via supervised, error-gradient-based methods, to perform complex temporal tasks can enhance or hamper their performance.

3.3.1 Reservoir computing with modular E/I networks

In chapter 7, we show how the introduction of modular, synapse-type specific connectivity in a balanced E/I network can shape the statistical structure of the network's population activity and enable synaptic plasticity to produce input selectivity in a readout neuron (Giannakakis *et al.*, 2024b). Motivated by this result, we then turn to investigate

the extent to which such a simple, low-dimensional parametrization of network connectivity can affect population dynamics in a way that impacts a network’s capacity to solve complex tasks.

To this end, after analytically investigating how different modular E/I structures generate distinct dynamical states in RNNs obeying Dale’s law, we use a reservoir computing framework to evaluate how such structures can impact the networks’ computational capabilities. Reservoir computing (Nakajima, 2018) is a computational paradigm that utilizes the complex intrinsic dynamics of a non-trainable dynamical system for processing complex inputs. Reservoir computing approaches are usually utilized for temporal tasks, such as time series prediction, due to the high cost of end-to-end training networks solving memory-demanding tasks (Shahi *et al.*, 2022).

In the original formulation of reservoir computing, a fixed recurrent neural network is used as a reservoir (Fig. 3.4). A given input (usually a time series) is presented to the reservoir, whose complex temporal dynamics project the input onto a higher dimensional space, making it more interpretable. Then, a simple readout layer (trained via linear or Ridge regression) can be used to map the reservoir’s state onto the target output of the task at hand (Jaeger and Haas, 2004).

While the most common type of dynamical system used in reservoir computing is a recurrent neural network, there has been a wide range of studies on how different dynamical systems can be used as reservoirs, including simulated as well as physical systems (Tanaka *et al.*, 2019). Spiking networks with varying degrees of biological plausibility have been commonly used as reservoirs for a wide range of tasks (Maass, 2011; Schliebs *et al.*, 2011). The use of spiking or other biological networks in reservoir computing allows the reliable assessment of their computational capabilities in a comparable and computationally efficient manner, which has made reservoir computing systems a popular setting for evaluating the dynamics and computational capacity of biological networks (Lazar *et al.*, 2009; Cramer *et al.*, 2020).

In our experiments, presented with greater detail in chapter 8, we used the modular E/I networks as reservoirs in a challenging time series prediction task. Our findings indicate that different levels of E and I modularity significantly affect network dynamics and control the network’s distance from a chaotic transition point. Moreover, we found that the different connectivity-driven dynamical states can have a significant impact on the network’s performance, indicating that biological structures, such as Dale’s law and modularity in network connectivity can exert a significant influence in the ability of recurrent networks to perform complex prediction tasks.




Figure 3.4: A sketch of a reservoir computing network. The complex dynamics of a non-trainable recurrent network transform a given input and make it more interpretable to a trainable readout layer.

3.3.2 Cellular Mechanisms and Network Structures for Memory tasks

Tasks that require memory of past states are particularly challenging for both biological and artificial agents (Bengio *et al.*, 1994; Squire and Wixted, 2011; Carruthers, 2013). Specifically, networks solving temporal tasks, are required to maintain traces of past states in their activity, a process that can be quantified by the decay rate of the neurons' autocorrelation function (Zeraati *et al.*, 2023), which gives the neurons' intrinsic timescales (Fig. 3.5a). A tentative association between longer intrinsic timescales and computational complexity has been observed in the mammalian brain, where brain areas dealing with more cognitively demanding tasks are characterized by longer timescales (Murray *et al.*, 2014). In biological networks, such long timescales can be created either via the biophysical properties of individual cells (Gjorgjieva *et al.*, 2016; Duarte *et al.*, 2017) or via network interactions (Litwin-Kumar and Doiron, 2012; Zeraati *et al.*, 2023; Shi *et al.*, 2023). Still, the exact computational implications of these different mechanisms in artificial networks have not been fully investigated. Thus, for this last section of the thesis, using an RNN with trainable single neuron timescales (representing the biophysical membrane timescale of biological neurons), we investigated how these different biological mechanisms for creating long timescales can be introduced in artificial recurrent neural networks and their impact on their performance.

Experimenting with different training mechanisms, we discovered that different training curricula can drive the RNN to develop long timescales using distinct mechanisms. Curriculum learning is a training paradigm that introduces progressively harder versions of a task to a learning agent, allowing gradual adjustment to the task demands. Curriculum training has been shown to lead to higher performance in both artificial agents (Wang *et al.*, 2021; Soviany *et al.*, 2022) and biological organisms (Lee *et al.*, 2024a). For our experiments, we found that the choice of training curriculum can drive a network to utilize cellular (trainable single-neuron timescale) or network-based (recurrent connectivity) mechanisms (Fig. 3.5b & c) for developing long timescales. We then compare the performance of the two types of networks and find that relying on the network connectivity for developing long timescales appears to be a far superior solution in terms of both performance and robustness.

Our results, published in (Khajehabdollahi *et al.*, 2024) and summarized here in chapter 9, suggest that different learning processes might recruit distinct mechanisms for solving a complex task. Moreover, our results indicate that different biological mechanisms for generating the necessary dynamics needed to solve complex temporal tasks may offer distinct benefits and disadvantages. In the case of the tasks we examined, we found that network-based mechanisms for creating long timescales lead to superior performance than relying on the time constants of individual cells. The fact that network connectivity appears to be the preferred mechanism by which biological networks generate long timescales (Zeraati *et al.*, 2023), suggests that its benefits for computation may translate across learning mediums to both artificial and biological networks.

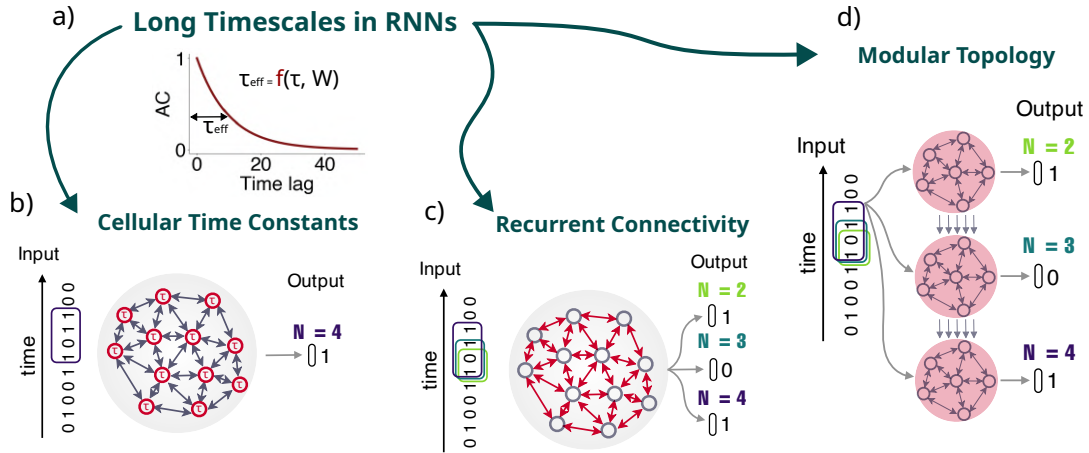


Figure 3.5: *Biophysical mechanisms for creating long timescales in RNNs solving temporal tasks. a.* For a neural network to solve temporal tasks, individual neurons need to have long effective timescales (measured as the decay rate of the autocorrelation function). This effective timescale arises via the interaction of the cellular time constant and network interactions. Thus, in order to create such timescales a network may rely on *b.* large cellular membrane time constants, *c.* tuned recurrent network connectivity, or *d.* a high-level topological structure that encourages the emergence of long timescales (a special case of (c)).

Our findings on the overwhelming benefit of using the network connectivity for developing long timescales, led us to wonder whether there is a way to further tune the connectivity to a specific task’s needs. To this end, questioning whether the classic all-to-all network connectivity of RNNs is optimal, we opted to study whether following a neurodevelopmental approach (Stanley and Miikkulainen, 2002) which introduces topological features specific to the task structure could improve performance.

Although training minimally structured networks *de novo* has overwhelmingly been the most prominent approach on developing artificial learning systems, there is an increasing number of experts shifting towards developmental approaches that introduce hard-coded features in network structure prior to learning. This approach, voiced most prominently in Zador (2019), argues that biological networks rely much more on hard-coded features that are either genetically encoded or organically emerge during development, than on learning during an organism’s lifetime and proposes a similar approach for the design of artificial networks. This can take the form of genome-to-phenotype mappings, where an agent is trained to output the structure of a network (Stanley, 2007; Ha et al., 2016; Shuvaev et al., 2024) or of developmental approaches, where instead of starting with a fixed number of neurons and training the connectivity, the network grows during the learning process (Risi, 2021; Plantec et al., 2024; Najarro et al., 2023), a mechanism akin to the embryogenesis of a biological organism (Hiesinger, 2018).

Inspired by such approaches and the computational benefits of structural modularity and hierarchical processing (Hihi and Bengio, 1995; Mohajerin and Waslander, 2014) in ANNs, we propose a growing developmental curriculum that builds up a hierarchical modular network (Fig. 3.5d). Comparing the performance of this modular network with equivalent networks from our original timescales study (Khajehabdollahi *et al.*, 2024), we found a significant boost in performance, training efficiency, and robustness. Our findings are presented in detail in Hamidi *et al.* (2024) and are summarized here in chapter 10.

Our results in this last study, add to the growing evidence that developmental and structural approaches may offer an avenue for innovation in the design of artificial learning systems. Moreover, our findings on the benefits of hierarchical and modular structures, connectivity patterns that are extremely widespread in biological networks (Meunier *et al.*, 2010; Redies and Puelles, 2001; Hilgetag and Kaiser, 2004) and are known to boost computational capacity (Zhang *et al.*, 2024), further demonstrates the importance of structural connectivity for biological learning and the potential benefits of incorporating biological structures in artificial systems.

3.4 Conclusion

In summary, this thesis focused on the different ways biological structures can influence learning in biological and artificial systems. We showed how cellular and network structures can influence the evolutionary dynamics of distinct synaptic plasticity rules, and how different biophysical features of neuronal networks can influence the learning outcomes of Hebbian plasticity. Finally, we demonstrated how the incorporation of different biological structures into artificial networks can establish strong inductive biases, boosting performance in a variety of complex learning tasks.

Our findings establish a link between network structure and learning outcomes across a variety of settings, and motivate further study into the potential influence of diverse biological structures on the ability of biological and artificial networks to learn different tasks. The following seven chapters present the individual publications forming the basis of this work in greater detail. Finally, the thesis closes with an outlook chapter (chapter 11), proposing potential future research directions.

4 When is it useful to learn ?

This chapter is based on the publication (Giannakakis *et al.*, 2023) titled *Environmental variability and network structure determine the optimal plasticity mechanisms in embodied agents*, which was published at the proceedings of the 2023 conference on Artificial Life (ALIFE 2023).

4.1 Motivation

While the ability to adapt to their environment is a defining feature of nearly all living organisms, the evolutionary origins of this ability remain unclear (Papini and Torres, 2017). The answer at least partially lies in the fact that the amount of information required to encode all information required for creating a functional nervous system is simply too large to be transmitted genetically (Zador, 2019; Shuvaev *et al.*, 2024). However, an even more crucial factor driving the development of learning, is that natural environments are very volatile and thus expected to change during an organisms's lifetime in unpredictable ways (Ellefsen, 2014; Dunlap and Stephens, 2016). The relation between environmental volatility and the emergence of learning has been studied in both biological and in artificial settings (Kerr and Feldman, 2003; Lange and Sprekeler, 2022).

In this work, we examined the effects of environmental volatility on the emergence of biologically plausible learning, namely reward-modulated synaptic plasticity. We aimed at classifying the effects that different types of environmental variability have on the evolution of synaptic learning in simple artificial organisms. Additionally, we examine the extent to which structural features of the learning networks influence the evolved learning rules.

4.2 Summary of results

We simulated both static agents (consisting in a single-layer plastic neural network) and embodied moving agents (consisting in a single-layer plastic network attached to a multi-layered non-plastic network controlling the agents' movement in a 2D surface with scattered food particles). The agents switch between two environments with different food particle distributions and need to adjust their sensory systems to the distribution of each environment in order to identify which foods to consume. We used a genetic algorithm to meta-learn the parameters of a reward-modulated plasticity rule, with the aim of enabling the agents to learn the distribution of food particle values. The plasticity was optimized on a number of environmental settings, introducing different types of variability. We further simulated sensory networks with linear and binary readouts, to examine the effects the readout nonlinearity can have on the resulting learning rules. Our main findings were:

- **Environmental variability controls the learning rate of the plasticity:** After simulating the agents (static and moving) in a variety of environmental conditions and recording the evolved learning rate of the resulting plasticity mechanisms, we find that in order for synaptic plasticity to emerge (i.e. for the learning rates to be sufficiently high to lead to substantial modification of synaptic strengths during an agent’s lifetime), the following conditions must be met:
 1. The environment must change over an agent’s lifetime. A fully static environment prevents the evolution of synaptic learning, and the variability of environmental conditions an agent is exposed to is directly correlated with its evolved learning rate.
 2. The change in environmental conditions should be slow enough that a learning process has time to at least partially converge. Too rapid environmental change prevents learning.
 3. The agent should have a reliable enough sensory system to collect information about its environment. A reduction in the reliability of an agent’s sensory system (via the introduction of noise) leads to a decrease in its evolved learning rate.
- **Network structure impacts the form of the evolved learning rule:** We simulated both static and moving agents with linear and binary sensory readouts. In the static agents, we found that the evolved learning rule consistently converged to a different form for each readout function. Moreover, for the moving agents, we noted a large diversity among the evolved learning rules, suggesting the development of specialized learning rules for different motor networks.

4.3 Discussion

Our findings suggest that different forms of environmental variability can reliably affect the optimal learning rates of synaptic plasticity and by extension the extent to which an agent relies on its initial “genetically encoded” network structure vs the learned connectivity that arises via interaction with the environment (Shuvaev *et al.*, 2024). Our results indicate that synaptic learning is not always a desirable strategy, extending earlier findings (Lange and Sprekeler, 2022; Pontes *et al.*, 2020) in a mechanistic setting where synaptic plasticity is explicitly modelled.

Moreover, we find that small changes in network structure (in our case, the readout function of the sensory network) can significantly impact the form of the evolved learning rule. Finally, in the case of co-evolved learning rules and motor networks for the moving agents, we observe a large variability in the resulting learning rules. Both these findings suggest that the form of learning mechanisms can be significantly impacted by structural features of the learning networks. We further explore this topic in chapter 5.

5 Co-evolution of network structure and synaptic plasticity

This chapter is based on the publication (Giannakakis *et al.*, 2024a) titled *Network bottlenecks and task structure control the evolution of interpretable learning rules in a foraging agent*, which was published in *Artificial Life*.

5.1 Motivation

The meta-learning of plasticity rules, i.e. the optimization of a synaptic update function in order to enable a network to learn a particular task, has been widely used, both to introduce continuous learning capabilities in artificial networks (Najarro and Risi, 2020; Pedersen and Risi, 2021a; Yaman *et al.*, 2021b) and to study synaptic learning in simulations of biological networks (Shervani-Tabar and Rosenbaum, 2023; Confavreux *et al.*, 2023). Recent studies (Ramesh *et al.*, 2024; Confavreux *et al.*, 2024) have found that the optimization of plasticity for specific tasks results in a large diversity of rules that converge equally well to the target state. This diversity of equally fitting learning rules for any given task makes it hard to interpret how different aspects of a network’s structure or various task details can influence the form of meta-learned plasticity rules.

Here, following on the preliminary results presented in chapter 4, we identify a simple regularization mechanism which significantly reduces the apparent variability of evolved learning rules. We use this to study how different network structures, normalization mechanisms for the synaptic plasticity (Zenke and Gerstner, 2017) as well as details of the task structure and loss function can drive the evolution of distinct learning rules.

5.2 Summary of results

We evolve moving and static agents following the structure described in chapter 4. In this study, we keep the environmental variability constant (to a level that encourages the development of synaptic plasticity) and focus on the form of the evolved learning rules.

- **Information bottlenecks promote generalizability:** We simulate moving agents with different readout neurons (binary and linear) in their sensory network. As in chapter 4, we originally find that there is a great diversity in the resulting rules, with no discernible pattern. However, we test how well these diverse rules can generalize when combined with motor networks co-evolved with a different learning rule. To this end, we swap the sensory and motor networks of different agents and test their performance. We find that, while the networks with a binary readout retain their performance when combined with different motor networks, those with

a linear readout decline significantly. This suggests that the introduction of an information bottleneck (i.e. a binary readout) between sensory and motor networks enables better generalization.

- **Regularization of the plasticity parameter reduces redundancy:** Introducing a simple L_1 regularization on the plasticity parameters leads to a very clear convergence of the learning rules into consistent sparse forms, regardless of the readout of the sensory network. This promotion of sparsity in the plasticity parameters avoids the confusing diversity we observe in the evolved rules of the moving agents, and was also found in other meta-learning studies (Ramesh *et al.*, 2024; Confavreux *et al.*, 2024). Consequently, the resulting rules become ubiquitously interpretable, enabling us to analytically study their dynamics and convergence points.
- **Choice of readout nonlinearity affects the evolved learning rule:** We repeat the experiments of chapter 4 with linear and binary readouts for the sensory networks of the moving agents, with the addition of the L_1 regularization leading to a convergence of the evolved learning rules in interpretable sparse forms. We observe that depending on the activation function of the readout network, different rules evolve, which suggests that rules are specific to the readout function of the sensory network. To further test this assumption, we simulate agents with a sigmoid readout activation function whose slope is evolved along with the other network parameters, and we observe that the resulting sigmoid tends to approach a step function. As expected, the corresponding evolved rules converge to the same general form as those evolved in networks with binary readouts.
- **Choice of normalization mechanism affects the evolved learning rule:** We further examine the impact that different normalization mechanisms have on the evolution of plasticity. Hebbian learning rules require compensatory processes in order to be stable (Zenke and Gerstner, 2017), the most prominent of which are subtractive and divisive normalization of synaptic weights (Elliott, 2003). We test how different normalization mechanisms affect the form of the resulting rule, and find that the evolved learning rule changes depending on the type of normalization that is applied. The impact of the different normalization mechanisms is most prominent in networks with scalar readouts, but also visible in networks with binary or sigmoid readouts.
- **Task structure affects the evolved rule:** Finally, we simulate static agents with L_1 regularization of the plasticity parameters, aiming to test the conditions under which the learning rules that evolve in the moving agents can emerge in the static agents. We find that the details of the objective function, as well as assumptions about the accuracy of the moving agent during its learning phase, significantly affect the form of the evolved rule. Specifically, we find that static agents that are exposed to both positive and negative food particles, regardless of their readout

value, converge to the same rule as the moving agents. This suggests that the plasticity rules that evolve in the moving agents take into account the unreliability of the agents' motor networks (i.e. the fact that negative food particles may be consumed by accident despite a correct sensory prediction).

5.3 Discussion

We have shown that in simple systems, the redundancy of meta-learned rules found in (Ramesh *et al.*, 2024; Confavreux *et al.*, 2023, 2024) can potentially be reduced by introducing L_1 regularization of the plasticity parameters. While this regularization does not make the emerging rules more biologically plausible (there is no biophysical reason favouring simpler rather than more complex synaptic plasticity rules), it does enable the analytical treatment of the resulting interpretable learning rules.

The analysis of the learning rules, evolved under different readout functions and normalization mechanisms, reveals a significant diversity in emerging forms and highlights the intricate influence of different structural factors on the development of optimal learning rules. Thus, our findings suggest that plasticity mechanisms can be very specific to various structural details of the networks they are acting upon. This motivates a more holistic study of biological learning, where different synaptic plasticity mechanisms are examined in the context of the specific circuits they are acting upon.

Finally, using the simplified static agents, we demonstrated that important details about the behaviour of the learning agents (such as the accuracy of communication between the motor and sensory systems) can be inferred by simulating the evolution of learning mechanisms under different conditions. Although our system is extremely simple compared to biological organisms, this setting can serve as a minimal proof of concept for using meta-learning techniques to study learning in biological networks. We argue that such methods can offer insights not only into the function of specific circuits, but also potentially provide intuitions about higher level structures and behaviours.

6 Dendritic structures control synaptic learning

This chapter is based on the manuscript titled *Dendritic nonlinearities and synapse-type specific input clustering enable the development of input selectivity in diverse settings*, which is currently in preparation for submission.

6.1 Motivation

Cortical pyramidal neurons have an extremely complex structure, characterized by rich dynamics driven by a multitude of ion channels (Hodgkin and Huxley, 1952), as well as a highly intricate dendritic tree (Larkman, 1991). The dendritic structures of cortical neurons are known to contribute, among other things, to memory storage (Kastellakis *et al.*, 2015, 2016) and efficient information processing (Mel, 1994). Moreover, recent experimental studies have shown that the distribution of E/I synapses along the dendritic tree is linked to the sharpness of orientation selectivity of neurons in the visual cortex (Wilson *et al.*, 2016; Iacaruso *et al.*, 2017; Weiler *et al.*, 2022).

Here, taking inspiration from a long line of research using Hebbian plasticity models to study the emergence of input selectivity (Bienenstock *et al.*, 1982; Clopath *et al.*, 2016; Eckmann *et al.*, 2024), we propose that these experimental findings linking dendritic structures with a neuron's orientation selectivity, could be potentially explained by the influence that dendritic nonlinearities and the specificity of inputs on the dendritic tree have on local E/I synaptic plasticity mechanisms.

6.2 Summary of results

Starting from an established protocol for creating orientation selectivity via interacting E and I plasticity (Eckmann *et al.*, 2024), we modify the original setting of a single compartment postsynaptic rate neuron by randomly grouping inputs and passing their sum through a sigmoid function before adding it to the somatic voltage. These external nonlinearities are meant to approximate the effects of dendritic summation of inputs in biological neurons.

- **Dendritic nonlinearities encourage the emergence of orientation selectivity:** The original E/I plasticity protocol (Eckmann *et al.*, 2024) is fairly sensitive to the presence of noise, which damages the correlation structure of the inputs. We find that the introduction of a sigmoid dendritic nonlinearity significantly improves the capacity of the plasticity to produce input selectivity in the presence of noise.
- **The shape of the dendritic nonlinearity controls the emerging neuron tuning:** The sigmoid dendritic nonlinearity is parametrized by its slope and centre. We use simulation-based inference (Tejero-Cantero *et al.*, 2020) to estimate the impact that

these parameters have on the sharpness of the learned orientation selectivity. We find that the shape of the nonlinearity can effectively control the resulting tuning, with steep sigmoids with a high threshold encouraging sharper, while sigmoids that are less steep and shifted to the left encouraging a broader tuning.

- **E/I synaptic specificity controls tuning:** We finally test how the specificity of inputs to a given dendrite (whether similar inputs are more likely to target the same dendrite) affects the emerging tuning. We find that, as was experimentally demonstrated in [Weiler *et al.* \(2022\)](#), the specificity of *E* and *I* synapses control different aspects of the emerging orientation tuning, with *E* specificity regulating the sharpness of tuning and *I* specificity the stability of the learned connectivity.

6.3 Discussion

The impact of the dendritic tree on single neuron activity ([Bastian and Nguyenkim, 2001](#); [Gollo *et al.*, 2009](#)) and information processing ([Smith *et al.*, 2013](#); [Eyal *et al.*, 2014](#); [Branco *et al.*, 2010](#)) has been a topic of intense study, demonstrating that dendritic structures enable single neurons to perform highly non-linear computations ([Poirazi *et al.*, 2003](#); [Spieler *et al.*, 2023](#)). Additionally, experimental studies have shown an impact of dendritic structures on orientation tuning ([Wilson *et al.*, 2016](#); [Iacaruso *et al.*, 2017](#); [Weiler *et al.*, 2022](#)) as well as an influence of dendritic dynamics on synaptic plasticity ([Golding *et al.*, 2002](#); [Kampa *et al.*, 2007](#)).

Our findings suggest a link between these two findings, by demonstrating how details of the dendritic structures can drive the learning dynamics of a Hebbian plasticity protocol ([Eckmann *et al.*, 2024](#)) and control the sharpness of the emerging orientation selectivity. We show how distinct aspects of the dendritic tree structure (namely the shape of the nonlinearity and the synaptic specificity) can act as strong inductive biases, pushing the development of specific orientation tunings, which indicates that structural features of single cells can effectively control the trajectories of synaptic learning rules. These findings motivate the further study of such synergistic effects between cellular structure and learning mechanisms that could potentially offer additional insights into biological learning.

7 Structural connectivity influences synaptic plasticity

This chapter is based on the publication (Giannakakis *et al.*, 2024b) titled *Structural influences on synaptic plasticity: the role of presynaptic connectivity in the emergence of E/I co-tuning*, which has been published in *PLOS Computational Biology*.

7.1 Motivation

Biological neural networks are characterized by highly intricate connectivity patterns, including clustering (Song *et al.*, 2005; Perin *et al.*, 2011), neuronal assemblies (Miehl *et al.*, 2023) and highly specific connectivity profiles for different neuron types (Jiang *et al.*, 2015). Many of these connectivity structures, such as clusters (Litwin-Kumar and Doiron, 2012; Lagzi and Rotter, 2015) and different spatial distributions of excitatory and inhibitory connectivity, (Rosenbaum and Doiron, 2014; Rosenbaum *et al.*, 2017) have been associated with distinct network dynamics. However, the extent to which structured connectivities and the associated population dynamics affect synaptic learning remains poorly understood.

In this paper, we investigated whether the introduction of structured E/I recurrent connectivity in a spiking network, can generate population activity with a desired statistical structure which enables synaptic plasticity to produce input selectivity and E/I weight co-tuning in a postsynaptic readout neuron.

7.2 Summary of results

We simulated a large spiking network of excitatory and inhibitory neurons that project with feedforward plastic connections to a single readout neuron. The network was divided in different subpopulations, each of whom received the same spike train as external input. We tested the ability of the synaptic plasticity to produce input selectivity and matching E/I connectivity in the feedforward neuron under different conditions.

- **Noise and random recurrent connectivity prevent the development of input selectivity and E/I co-tuning:** Synaptic plasticity is able to produce E/I co-tuning and input selectivity on the postsynaptic neuron in the case of feedforward networks with low noise. However, we find that increasing the noise level in the input or introducing random, fixed recurrent connections between the neurons of the presynaptic network, changes the statistics of the population activity and prevents the emergence of co-tuned and orientation selective feedforward connectivity.
- **The introduction of structured fixed recurrent connectivity ameliorates the effects of noise and recurrence:** We use Approximate Bayesian Computation to identify network structures that might restore the population statistics required for

the plasticity to produce E/I co-tuning and input selectivity. Specifically, we vary the strength of synapse-type specific assemblies, formed between E and I neurons of the same subpopulation (neurons receiving the same input signal). We find that a structure of strong E assemblies and intermediate strength I assemblies, can restore the desired population statistics even in recurrent networks with high levels of noise, which allows the synaptic plasticity protocol to produce co-tuning and input selectivity in the postsynaptic neuron.

- **The optimal connectivity structure depends on sparsity levels:** We perform Bayesian inference to identify the optimal assembly structure in networks with varying sparsity levels, and we find that there is a shift in the inferred assemblies strengths as the networks become sparser. Specifically, we find that for sparser networks, the optimal inhibitory connectivity becomes even less specific, while the excitatory connectivity remains at similar levels.
- **Structured connectivity is beneficial even in fully plastic networks:** We finally, introduce plasticity in the recurrent connectivity of the presynaptic network and test how the inferred connectivity structures from the static network impact the development of input selectivity in a fully plastic network. We verify that the structure that we inferred for static networks can still influence the population statistics of fully plastic networks, which boosts the ability of plasticity to produce input selectivity.

7.3 Discussion

While the ability of synaptic plasticity to shape the connectivity of neural networks has been the focus of multiple computational studies (Vogels *et al.*, 2011; Litwin-Kumar and Doiron, 2014; Zenke *et al.*, 2015), the inverse influence, i.e. the ways in which structural network features might impact the behaviour of different plasticity protocols, has been largely neglected as a topic of study. However, biological neural networks are characterized by highly complex connectivity patterns (Song *et al.*, 2005) that largely emerge during early development and change very little during an organism’s lifetime. Thus, synaptic plasticity has to act within the constraint these highly intricate connectivity structures impose on the network activity and information flow.

In this paper, we demonstrated how the introduction of simple, biologically plausible (Miehl *et al.*, 2023), structured connectivity can have a massive impact on the statistics of population activity that spike-timing dependent synaptic plasticity relies on to produce input selectivity. In essence, we provide a setting in which structural features of a network that largely arise during early development can act as a very strong inductive bias for synaptic learning, effectively regulating the emergence of input selectivity and balanced E/I connectivity.

8 Reservoir computing with structured E/I networks

This chapter is based on the manuscript titled *Modularity in E/I networks controls the dynamical regime and optimizes computational capabilities*, which is currently in preparation for submission.

8.1 Motivation

In chapter 7 we showed how varying the strengths of E and I assemblies in a recurrent network can modify the statistics of the population activity so that synaptic plasticity manages to produce input selectivity in a post-synaptic neuron (Giannakakis *et al.*, 2024b). We then wondered whether similar E/I connectivity patterns can impact the population activity in ways that could boost a network’s ability to perform much more complex tasks. Modular connectivity patterns are known to create distinct dynamics in neural networks (Litwin-Kumar and Doiron, 2012; Lagzi and Rotter, 2015) and multiple studies have demonstrated that the computational capabilities of recurrent networks are associated with their dynamical state (Boedecker *et al.*, 2012; Legenstein and Maass, 2007).

Thus, in this study, we use dynamic mean-field theory (Kadmon and Sompolinsky, 2015) to systematically analyse the impact that different levels of E/I modularity have on the dynamics of a recurrent neural network. Then, using a reservoir computing framework (Schrauwen *et al.*, 2007), we evaluate how the connectivity-induced dynamics affect the network’s computational capabilities in a complex time series prediction task.

8.2 Summary of results

We use a rate network obeying Dale’s law which we train as an Echo State network (Jaeger and Haas, 2004) to predict the trajectories of chaotic attractors.

- **DMFT of E/I modular networks:** Using dynamic mean-field theory, we show that the network’s dynamics are controlled by its E/I modularity and its spectral radius. We find that for sufficiently large spectral radius, the difference in E and I modularity controls the distance from the edge of chaos.
- **The network’s dynamical state affects the computational capabilities of the ESN:** We examine how networks with different E/I modularity levels behave as reservoirs. We find that the analytically derived dynamical states can predict the loss map of the ESN reasonable well, suggesting that the recurrent dynamics of the network (and consequently the E/I modular structure that controls the dynamics) determine the network’s computational capabilities.

- **Network size and time series dimensionality shift the critical point:** Subsequently, we study the performance of ESNs which are tasked with reconstructing multiple chaotic attractors simultaneously. We note that as the number of input attractors increases, the loss pattern shifts, suggesting that the critical transition to the chaotic state happens for smaller network spectral radius. We find that this effect can be ameliorated by increasing network size proportionally to the number of attractors.
- **Input structure has a minor impact on the network’s performance:** We finally test whether the structure of the input has any impact on the performance of the ESN. In particular, using an ESN with four modules which is reconstructing four different chaotic attractors, we test whether different levels of E and I input specificity of each attractor targeting a different module would make a difference in performance. Using simulation-based inference (Tejero-Cantero *et al.*, 2020) to better estimate this effect, we find that depending on the recurrent network’s dynamical state, different input structures may have a small but significant impact on performance. This suggests that despite the overwhelming influence of recurrent dynamics (and the structural connectivity that controls it), structured input connectivity may also play a small role in controlling computational capabilities.

8.3 Discussion

Our work suggests that in networks obeying Dale’s law, different modular structures for the E and I connectivity may offer precise control of network dynamics, enabling the network to tune its distance from the chaotic state. These findings complement our conclusions about how modularity/clustering can influence synaptic learning (Giannakakis *et al.*, 2024b) and suggest that neuron type-specific inhomogeneous connectivity structures - a commonly observed feature of biological neural networks (Lund, 2002; Callaway, 2002; Fino *et al.*, 2013) - may impact a network’s capacity to learn and process information in a variety of different ways.

Additionally, our results suggest that reservoir computing (Schrauwen *et al.*, 2007; Nakajima and Fischer, 2021), despite potential shortcoming in terms of performance compared to state-of-the-art methods (Schrauwen *et al.*, 2007; Gauthier *et al.*, 2021), can perform quite well on challenging temporal tasks. These computational capabilities along with the extreme efficiency of reservoir networks in terms of training resources make reservoir computing not only a valuable framework for evaluating the computational capabilities of different biological neural network architectures, but also highlight its potential as a component of high-performing artificial systems (Lukoševičius and Jaeger, 2009). Specifically, as our understanding of ways to reliably control network dynamics and the relationship between dynamical state and computational capacity increases, a greater use of reservoir computing methods in different applications where training cost needs to be kept at a minimum should be considered.

9 Task structure drives the emergence of long timescales in RNNs

This chapter is based on the publication (Khajehabdollahi *et al.*, 2024) titled *Emergent mechanisms for long timescales depend on training curriculum and affect performance in memory tasks*, which was published in the proceedings of the Twelfth International Conference on Learning Representations (*ICLR 2024*).

9.1 Motivation

Biological networks often face the challenge of solving tasks with complex temporal dependencies that may span a wide range of timescales (Zeraati *et al.*, 2024). In order to effectively perform such computations, the underlying timescale of the task needs to be reflected in the population activity of the network, leading to a tuning of intrinsic network timescales across brain areas (Murray *et al.*, 2014). Long timescales can emerge via complex recurrent network interactions (Litwin-Kumar and Doiron, 2012; Zeraati *et al.*, 2023) or via the intrinsic properties of individual neurons (Gjorgjieva *et al.*, 2016), with both of these mechanisms having been observed in different settings. However, the relative benefits and disadvantages of these two mechanisms for computation have not been systematically studied.

Here, we demonstrate how the task structure itself can drive a network to utilize either biophysical or network mediated mechanisms for developing long timescales. We moreover provide a direct comparison of the two mechanisms in terms of task performance, training speed and robustness.

9.2 Summary of results

We use a recurrent neural network (RNN) with trainable recurrent weights and intrinsic timescales for each neuron. The network is presented with a binary sequence and trained to solve one of two tasks (N -delayed match-to-sample or N -parity).

All RNNs are trained using a curriculum learning approach (Bengio *et al.*, 2009). Specifically, we develop two distinct training curricula.

1. **Single-head:** The network is first trained to solve the task for $N = 2$. The trained network for each N is used the initial condition for training a network to solve the task for $N + 1$.
2. **Multi-head:** At each curriculum step, a new readout head is added, which is trained to solve the task for increasing values of N . Thus, on the K -th step of the curriculum, the network is simultaneously trained to solve the task for $N = 2, \dots, K$.

Our main results are the following:

- **The multi-head curriculum boosts performance:** The multi-head curriculum appears to be superior to the single-head one in all metrics considered. Specifically, networks trained with the multi-head curriculum reach higher N , train faster and are more robust to perturbations compared to networks trained with the single-head curriculum
- **Different training curricula lead to different timescale creation mechanisms:** While the average effective timescale (neuron autocorrelation time) increases with N for all networks, we find that networks trained with the single-head curriculum, rely on the trainable single-neuron (membrane) timescale of individual neurons to create long timescales. On the contrary, networks trained with the multi-head curriculum develop very fast single-neuron timescales and rely on network interactions for developing longer timescales as N increases.
- **The connectivity is tuned to match the membrane dynamics:** The learned connectivity of trained networks appears to be linked to the different timescale creation mechanisms. Multi-head networks maintain a balanced connectivity (average weight around 0) as N increases, while the connectivity of single-head networks becomes negatively skewed for higher N s. The balanced E/I connectivity of the multi-head networks is associated with the creation of long timescales (Bhatia *et al.*, 2019). On the contrary, the inhibition-dominated connectivity in the single-head networks is presumably necessary to create stable dynamics in the presence of long single-neuron timescales.

9.3 Discussion

While curriculum learning is well known to improve learning performance in both artificial and biological agents (Bengio *et al.*, 2009; Lee *et al.*, 2024a), the impact of different curricula on the learned network structure has only recently become the subject of systematic study (Kepple *et al.*, 2022).

In this work, we showed how different curricula can fundamentally alter the mechanisms by which a learning objective is achieved. We found that a network’s reliance on single neuron biophysical properties vs network structure for developing long timescales can be effectively determined by the structure of the learning task. The massive difference in performance between the two network types indicates that structural properties of neural networks (in this case the presence of large membrane time constants and the associated inhibition-dominated connectivity vs fast single-neuron time constants combined with E/I balance for achieving long timescales) can have a significant impact on the ability of networks to learn complex tasks.

10 Modular Growth of Hierarchical Networks

This chapter is based on the publication (Hamidi *et al.*, 2024), titled *Modular Growth of Hierarchical Networks: Efficient, General, and Robust Curriculum Learning*, which was published at the proceedings of the 2024 conference on Artificial Life (ALIFE 2024).

10.1 Motivation

Structural modularity is a unifying characteristic of complex brains. The division into functionally specialized subunits allows the compartmentalization of information processing, increases efficiency (Clune *et al.*, 2013) and enables mosaic evolution (Fong *et al.*, 2021; Schumacher and Carlson, 2022). Additionally, modular networks are associated with distinct dynamics (Litwin-Kumar and Doiron, 2012; Giannakakis *et al.*, 2024b) that are beneficial for temporal computations and learning. These features have been frequently exploited for the design of artificial networks solving complex tasks (Hihi and Bengio, 1995; Mohajerin and Waslander, 2014; Yuan *et al.*, 2024).

Here, we apply these insights into the problem of creating long timescales in networks solving temporal tasks. Following on our findings presented in chapter 9 and recent research on self-assembling neural networks (Najarro *et al.*, 2023), we propose a developmental curriculum (Bengio *et al.*, 2009), which gradually assembles a network from replicable modules, creating a hierarchical structure that is tuned to the demands of the task at hand.

10.2 Summary of results

Following on the methodology of (Khajehabdollahi *et al.*, 2024), which was presented in chapter 9, we train both the connectivity and the single-neuron timescales of our models. The curriculum we develop, starts with a very small RNN (5 - 15 neurons), which is trained to solve a temporal task (either N -delayed match-to-sample or N -parity) for $N = 2$. Once a target accuracy has been achieved, the RNN is replicated and connected with feedforward connections to the original RNN. The resulting structure is now trained (the recurrent connections of the old and new modules, as well as the feedforward layer connecting them) to solve the task for $N = 2$ and $N = 3$, with the readout layer from the original module being tasked with solving it for $N = 2$ and the readout from the new module solving it for $N = 3$. Once this network is trained, a new module (a copy of the previous last module), solving the task for $N = 4$ is added creating a hierarchy of modules solving the task for increasingly high N_{max} .

- **Modular vs. Non-Modular Networks:** We compare a set of networks trained with our growing curriculum with a set of equivalent non-modular networks (trained

with the multi-head curriculum from chapter 9). We find that the modular networks outperform the equivalent non-modular networks (with a comparable number of trainable parameters) in terms of training time and the highest N_{max} reached. Additionally, we find that modular networks are more robust to perturbations of the recurrent connectivity than their non-modular counterparts.

- **Analysis of modular networks:** The highly specific structure of the modular networks allows dissecting their function to a degree that it usually impossible with classic RNNs. We first find that, similar to the multi-head networks from chapter 9, modular networks converge to fast single-neuron timescales and rely on their connectivity for the development of long intrinsic timescales. We further study the contribution of different connections types (feed-forward between modules vs recurrent within modules). Via perturbation experiments, we find that the feed-forward connections are much more sensitive to perturbation than the recurrent ones. Moreover, we note that during training, the feed-forward connectivity changes significantly more than the recurrent connectivity across modules. These findings suggest that the details of the recurrent connectivity matter much less than the specific feed-forward connections. We verify this by comparing performance when only one type of weights is trained while the other is replicated across modules. We find that freezing or duplicating recurrent connectivity barely affects performance, while doing the same in the feed-forward connections essentially destroys the network’s learning capabilities. These findings suggest that modular networks function as deep reservoirs, where the complex dynamics of the recurrently connected modules are utilized by the specific feed-forward connections to solve the task. Thus, by replicating the recurrent modules instead of training them on every step, we can effectively half the number of trainable parameters without any significant loss in performance.

10.3 Discussion

Biological networks are highly structured (Sporns, 2015), with modular topologies in particular being a prominent feature of brain connectivity (Bertolero *et al.*, 2015; Sporns and Betzel, 2016). Modular architectures have been used in artificial networks, (Happel and Murre, 1994; Mohajerin and Waslander, 2014) increasing performance and efficiency across a variety of tasks.

In this study, we have demonstrated how a developmental curriculum that imposes a modular, task-specific structure to a network can significantly boost performance and robustness while simultaneously decreasing training time and number of trainable parameters. Our findings suggest that the presence of replicable connectivity structures may play an important role in biological learning and highlight how biologically plausible structural constraints can be utilized for efficient learning in artificial networks.

11 Outlook

In this work, we have demonstrated a strong link between cellular and network structure and the capacity of different learning mechanisms to perform various tasks. We have shown how the form of meta-learned plasticity rules can be strongly affected by network and task constraints, (Giannakakis *et al.*, 2023, 2024a) and how biological structure on both the cellular (dendritic structures) and network levels (E/I connectivity) can control the learning dynamics of Hebbian plasticity rules (Giannakakis *et al.*, 2024b). Finally, we have demonstrated how the introduction of biological structural constraints on artificial learning systems performing complex temporal tasks (time series prediction and short-term memory tasks) can impact their performance (Khajehabdollahi *et al.*, 2024; Hamidi *et al.*, 2024).

However, despite the novel insights these works bring on the relationship between network structure and learning, our contributions remain limited by the proof of concept approach we adopted for most of these works. Specifically, in all the works presented here, we have used simplified neuron models and network structures that are far from capturing the full complexity of biological systems in either the cellular or the network level. Moreover, we have focused on relatively simple supervised and unsupervised tasks that were largely selected for their interpretability but fail to capture the intricacies of the complex behavioural tasks facing even the simplest biological organisms. These shortcomings can be improved upon by future work focusing on greater biological detail as well as more challenging applications. Thus, in this last section of the thesis, we will outline some potential future research directions through which the main themes of this thesis can be expanded upon.

11.1 Meta-learning Plasticity Rules and Network Structure

One of the unifying themes of the work presented in this thesis is the use of various optimization methods for discovering cellular and network structures, as well as plasticity mechanisms that optimize the learning capabilities of a given system. This included using evolutionary algorithms (Beyer and Schwefel, 2002) for optimizing both the network connectivity and the parameters of the plasticity rules in chapters 4 and 5, approximate Bayesian computation (Beaumont, 2019) to infer beneficial E/I network topologies in chapter 7 as well as simulation-based inference (Tejero-Cantero *et al.*, 2020) to identify optimal dendritic configurations in chapter 6 and input connectivity structures in chapter 8.

Such meta-learning approaches have been extensively used in the past to identify plasticity rules that could learn a given task. To this end, a variety of methods has been used, including evolutionary optimization (Jordan *et al.*, 2021; Confavreux *et al.*, 2020), simulation-based inference (Confavreux *et al.*, 2023, 2024), generative adversarial networks (Ramesh *et al.*, 2024) and direct gradient-based optimization (Mehta *et al.*, 2023),

with varying degrees of success. Additionally, despite the focus of meta-learning studies being overwhelmingly on the form of the plasticity rules, similar approaches have been used on identifying specific connectivity structures that can influence learning (Lakshminarasimhan *et al.*, 2024).

Our work argues for a greater focus on structure as a driver of synaptic learning and proposes a synergistic study of plasticity mechanisms and network structure using meta-learning methods. In particular, our approach could be extended by a simultaneous optimization of structure and plasticity (following a similar logic to the work presented in chapters 4 and 5) in more realistic biological settings. This could involve the simultaneous meta-learning of network architectures or dendritic structures along with the parameters of synaptic plasticity rules, focusing, for example, on the interaction between varying levels of structural E/I clustering and the parameters of synaptic learning rules active in a plastic recurrent setting. The increased dimensionality of this type of parametrization would require an amount of data that could put a strain on computational resources, but as inference techniques improve (Confavreux *et al.*, 2023) such computations are becoming more tractable.

11.2 Introducing greater Biological Plausibility in the modeling of Learning

This thesis examined how biological structures can influence learning in neural networks. Using dendritic nonlinearities in chapter 6 and clustered E/I connectivity patterns in chapters 7 and 8, we demonstrated how such features of a biological network can act as inductive biases for different learning mechanisms. However, despite the biological plausibility of such structures (Hilgetag and Kaiser, 2004), our models remain very far from even approaching the complexity of real biological cellular and network structures since it only captures high-level features of biological systems. Future work could build in this direction by focusing on integrating detailed experimental data about the structure of biological neural networks in studies of biological learning.

At first, the complexity of individual neurons had been until very recently almost impossible to model, except by using very large and computationally intractable multi-compartment models (Hodgkin and Huxley, 1952). However, as more flexible and efficient phenomenological neuron models - that can nevertheless capture the full complexity of biological neurons - are developed (Beniaguev *et al.*, 2021; Spieler *et al.*, 2023), studying learning in networks of biologically plausible neurons becomes an increasingly plausible research direction. This could involve the modeling of more detailed learning rules, that take into account neuron morphology and other biophysical parameters.

Most importantly, the last years have seen an explosion in the volume of available connectomic data, which has made the detailed study of large biological circuits a distinct possibility. In particular, the recent publications of detailed connectomic data from the *Drosophila Melanogaster* larva (Winding *et al.*, 2023) as well as the mouse cortex

(Consortium *et al.*, 2021; Eisenstein, 2024), has opened up new possibilities for data driven studies of network structure. Using connectomic data to constraint the structure of neural networks performing visual tasks has already revealed interesting insights on the function of different neuron types in the fruit fly’s visual system (Lappalainen *et al.*, 2024). A similar approach for studying plasticity mechanisms, focusing on well mapped plastic circuits such as the mushroom body (Smith *et al.*, 2008; Menzel, 2014; Modi *et al.*, 2020), could provide useful insights on the form and function of plasticity rules in biological circuits as well as the impact that detailed connectivity structures may have on the function of these learning mechanisms.

11.3 Developing Artificial Learning Systems with Biological Priors

The final part of the thesis dealt with the benefits of introducing biological structures in artificial learning systems. We showed how different levels of E/I modularity can control the dynamical regime and, consequently, the computational capabilities of recurrent networks in chapter 8. Furthermore, in chapters 9 and 10 we showed how different biological mechanisms for creating long timescales (Murray *et al.*, 2014) can impact the performance and robustness of RNNs solving temporal tasks. In all these cases, we showed how structural features of these neural networks can impact their learning process and affect performance. Still, our examples were largely focused on the biological aspects of the networks and used relatively simple tasks as a means to prove a general concept. Nevertheless, integrating biological structures into artificial networks, in particular via the use of developmental approaches like the growing curriculum of chapter 10, is a promising direction in artificial intelligence that could increase the efficiency and generalizability of artificial systems.

The reliance of biological organisms on inherent structures, that emerge during the early stages of development, as opposed to sensory-dependent learning via their lifetime, has led to suggestions that artificial systems should also adopt an approach of optimizing in-built properties as opposed to learning capabilities (Zador, 2019). This approach emphasizes the ability of genetic encoding to compress vast amounts of information that generate functional networks prior to any learning. Several applications of this principle have produced promising results (Stanley and Miikkulainen, 2002; Ha *et al.*, 2016; Najarro *et al.*, 2023; Shuvaev *et al.*, 2024) by using different mechanisms to develop fully functional networks that are not explicitly trained.

Our results on the interaction between synaptic plasticity and network structure, suggest a complimentary approach that emphasizes both genetically encoded structure and lifelong learning via local learning mechanisms. Several studies have demonstrated how optimized local learning rules can produce state-of-the-art results in artificial networks (Najarro and Risi, 2020; Pedersen and Risi, 2021b) and developmental approaches that take into account both structure and local plasticity have shown promise in proof of

concept tasks (Montero *et al.*, 2024). Our findings encourage further research in this direction, potentially with the inclusion of additional biological features such as Dale's law, diverse neuron time constants and explicit modular architectures that could interact with different learning mechanisms.

List of abbreviations

The following abbreviations are used throughout the text of the thesis for brevity. Their meanings are given in the table below:

<i>E</i>	Excitatory
<i>I</i>	Inhibitory
STDP	Spike Timing Dependent Plasticity
LTP	Long-Term Potentiation
LTD	Long-Term Depression
ANN	Artificial Neural Network
RNN	Recurrent Neural Network
SNN	Spiking Neural Network
ABC	Approximate Bayesian Computation
SBI	Simulation-Based Inference
DMFT	Dynamic Mean-Field Theory

Bibliography

- Abbott, L. F. and Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nature neuroscience*, **3**(11), 1178–1183.
- Aso, Y., Herb, A., Ogueta, M., Siwanowicz, I., Templier, T., Friedrich, A. B., Ito, K., Scholz, H., and Tanimoto, H. (2012). Three dopamine pathways induce aversive odor memories with different stability. *PLoS genetics*, **8**(7), e1002768.
- Augustinaite, S., Kuhn, B., Helm, P. J., and Heggelund, P. (2014). Nmda spike/plateau potentials in dendrites of thalamocortical neurons. *Journal of Neuroscience*, **34**(33), 10892–10905.
- Bain, A. (1873). *Mind and body: The theories of their relation*, volume 4. HS King & Company 1873 Farnborough Eng. Gregg International.
- Balaban, E. (1997). Changes in multiple brain regions underlie species differences in a complex, congenital behavior. *Proceedings of the National Academy of Sciences*, **94**(5), 2001–2006.
- Barrionuevo, G., Schottler, F., and Lynch, G. (1980). The effects of repetitive low frequency stimulation on control and “potentiated” synaptic responses in the hippocampus. *Life sciences*, **27**(24), 2385–2391.
- Bastian, J. and Nguyenkim, J. (2001). Dendritic modulation of burst-like firing in sensory neurons. *Journal of neurophysiology*, **85**(1), 10–22.
- Bean, B. P. (2007). The action potential in mammalian central neurons. *Nature Reviews Neuroscience*, **8**(6), 451–465.
- Beaumont, M. A. (2019). Approximate bayesian computation. *Annual review of statistics and its application*, **6**(1), 379–403.
- Beed, P., Ray, S., Velasquez, L., Stumpf, A., Parthier, D., Swaminathan, A., Nitzan, N., *et al.* (2020). Species-specific differences in synaptic transmission and plasticity. *sci rep* 10: 16557.
- Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., and Maass, W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature communications*, **11**(1), 3625.
- Bengio, Y., Simard, P., and Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, **5**(2), 157–166.

- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. (2009). Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48.
- Beniaguev, D., Segev, I., and London, M. (2021). Single cortical neurons as deep artificial neural networks. *Neuron*, **109**(17), 2727–2739.
- Bennett, J. E., Philippides, A., and Nowotny, T. (2021). Learning with reinforcement prediction errors in a model of the drosophila mushroom body. *Nature communications*, **12**(1), 2569.
- Berg Palm, R., Najarro, E., and Risi, S. (2020). Testing the genomic bottleneck hypothesis in hebbian meta-learning. *arXiv e-prints*, pages arXiv–2011.
- Bertolero, M. A., Yeo, B. T., and D’Esposito, M. (2015). The modular and integrative functional architecture of the human brain. *Proceedings of the National Academy of Sciences*, **112**(49), E6798–E6807.
- Betti, A., Gori, M., and Marra, G. (2018). Backpropagation and biological plausibility. *arXiv preprint arXiv:1808.06934*.
- Beyer, H.-G. and Schwefel, H.-P. (2002). Evolution strategies—a comprehensive introduction. *Natural computing*, **1**, 3–52.
- Bhatia, A., Moza, S., and Bhalla, U. S. (2019). Precise excitation-inhibition balance controls gain and timing in the hippocampus. *Elife*, **8**, e43415.
- Bienenstock, E. L., Cooper, L. N., and Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, **2**(1), 32–48.
- Bittner, K. C., Milstein, A. D., Grienberger, C., Romani, S., and Magee, J. C. (2017). Behavioral time scale synaptic plasticity underlies ca1 place fields. *Science*, **357**(6355), 1033–1036.
- Blakemore, C. and Cooper, G. F. (1970). Development of the brain depends on the visual environment. *Nature*, **228**(5270), 477–478.
- Bliss, T. V. and 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.
- Boedecker, J., Obst, O., Lizier, J. T., Mayer, N. M., and Asada, M. (2012). Information processing in echo state networks at the edge of chaos. *Theory in Biosciences*, **131**, 205–213.

- Booker, S. A. and Vida, I. (2018). Morphological diversity and connectivity of hippocampal interneurons. *Cell and tissue research*, **373**(3), 619–641.
- Branco, T., Clark, B. A., and Häusser, M. (2010). Dendritic discrimination of temporal input sequences in cortical neurons. *Science*, **329**(5999), 1671–1675.
- Brown, S. P. and Hestrin, S. (2009). Intracortical circuits of pyramidal neurons reflect their long-range axonal targets. *Nature*, **457**(7233), 1133–1136.
- Butz, M., Wörgötter, F., and Van Ooyen, A. (2009). Activity-dependent structural plasticity. *Brain research reviews*, **60**(2), 287–305.
- Callaway, E. M. (2002). Cell type specificity of local cortical connections. *Journal of neurocytology*, **31**(3), 231–237.
- Carruthers, P. (2013). Evolution of working memory. *Proceedings of the National Academy of Sciences*, **110**(supplement_2), 10371–10378.
- Castillo, P. E., Chiu, C. Q., and Carroll, R. C. (2011). Long-term plasticity at inhibitory synapses. *Current opinion in neurobiology*, **21**(2), 328–338.
- Chavlis, S. and Poirazi, P. (2021). Drawing inspiration from biological dendrites to empower artificial neural networks. *Current opinion in neurobiology*, **70**, 1–10.
- Citri, A. and Malenka, R. C. (2008). Synaptic plasticity: multiple forms, functions, and mechanisms. *Neuropsychopharmacology*, **33**(1), 18–41.
- Clark, M. Q., Zarin, A. A., Carreira-Rosario, A., and Doe, C. Q. (2018). Neural circuits driving larval locomotion in drosophila. *Neural development*, **13**, 1–10.
- Clopath, C., Vogels, T. P., Froemke, R. C., and Sprekeler, H. (2016). Receptive field formation by interacting excitatory and inhibitory synaptic plasticity. *BioRxiv*, page 066589.
- Clune, J., Mouret, J.-B., and Lipson, H. (2013). The evolutionary origins of modularity. *Proceedings of the Royal Society b: Biological sciences*, **280**(1755), 20122863.
- Confavreux, B., Zenke, F., Agnes, E., Lillicrap, T., and Vogels, T. (2020). A meta-learning approach to (re) discover plasticity rules that carve a desired function into a neural network. *Advances in Neural Information Processing Systems*, **33**, 16398–16408.
- Confavreux, B., Ramesh, P., Goncalves, P. J., Macke, J. H., and Vogels, T. P. (2023). Meta-learning families of plasticity rules in recurrent spiking networks using simulation-based inference. In *Thirty-seventh Conference on Neural Information Processing Systems*.

- Confavreux, B., Agnes, E. J., Zenke, F., Sprekeler, H., and Vogels, T. P. (2024). Balancing complexity, performance and plausibility to meta learn plasticity rules in recurrent spiking networks. *bioRxiv*, pages 2024–06.
- Consortium, M., Bae, J. A., Baptiste, M., Bishop, C. A., Bodor, A. L., Brittain, D., Buchanan, J., Bumbarger, D. J., Castro, M. A., Celi, B., *et al.* (2021). Functional connectomics spanning multiple areas of mouse visual cortex. *BioRxiv*, pages 2021–07.
- Conwell, C., Prince, J. S., Kay, K. N., Alvarez, G. A., and Konkle, T. (2024). A large-scale examination of inductive biases shaping high-level visual representation in brains and machines. *Nature Communications*, **15**(1), 9383.
- Cramer, B., Stöckel, D., Kreft, M., Wibrals, M., Schemmel, J., Meier, K., and Priesemann, V. (2020). Control of criticality and computation in spiking neuromorphic networks with plasticity. *Nature communications*, **11**(1), 2853.
- De Belle, J. S. and Heisenberg, M. (1994). Associative odor learning in drosophila abolished by chemical ablation of mushroom bodies. *Science*, **263**(5147), 692–695.
- Deco, G., Perl, Y. S., Vuust, P., Tagliazucchi, E., Kennedy, H., and Kringelbach, M. L. (2021). Rare long-range cortical connections enhance human information processing. *Current Biology*, **31**(20), 4436–4448.
- Denève, S. and Machens, C. K. (2016). Efficient codes and balanced networks. *Nature neuroscience*, **19**(3), 375–382.
- Ding, Z., Fahey, P. G., Papadopoulos, S., Wang, E. Y., Celi, B., Papadopoulos, C., Kunin, A. B., Chang, A., Fu, J., Ding, Z., *et al.* (2023). Functional connectomics reveals general wiring rule in mouse visual cortex. *bioRxiv*.
- Dorrell, W., Yuffa, M., and Latham, P. E. (2023). Meta-learning the inductive bias of simple neural circuits. In *International Conference on Machine Learning*, pages 8389–8402. PMLR.
- Doya, K. (2002). Metalearning and neuromodulation. *Neural networks*, **15**(4–6), 495–506.
- Duan, Y., Jia, Z., Li, Q., Zhong, Y., and Ma, K. (2023). Hebbian and gradient-based plasticity enables robust memory and rapid learning in RNNs. In *The Eleventh International Conference on Learning Representations*.
- Duarte, R., Seeholzer, A., Zilles, K., and Morrison, A. (2017). Synaptic patterning and the timescales of cortical dynamics. *Current Opinion in Neurobiology*, **43**, 156–165.

- Dunlap, A. S. and Stephens, D. W. (2016). Reliability, uncertainty, and costs in the evolution of animal learning. *Current Opinion in Behavioral Sciences*, **12**, 73–79. Behavioral ecology.
- Dussutour, A. (2021). Learning in single cell organisms. *Biochemical and Biophysical Research Communications*, **564**, 92–102.
- Ecker, A., Egas Santander, D., Abdellah, M., Blanco Alonso, J., Bolaños-Puchet, S., Chindemi, G., Gowri Mariyappan, D. P., Isbister, J. B., King, J. G., Kumbhar, P., *et al.* (2023). Assemblies, synapse clustering and network topology interact with plasticity to explain structure-function relationships of the cortical connectome. *bioRxiv*, pages 2023–08.
- Eckmann, S., Young, E. J., and Gjorgjieva, J. (2024). Synapse-type-specific competitive hebbian learning forms functional recurrent networks. *Proceedings of the National Academy of Sciences*, **121**(25), e2305326121.
- Eisenstein, M. (2024). A milestone map of mouse-brain connectivity reveals challenging new terrain for scientists. *Nature*, **628**(8008), 677–679.
- Ellefsen, K. O. (2014). The evolution of learning under environmental variability. pages 649–656.
- Elliott, T. (2003). An analysis of synaptic normalization in a general class of hebbian models. *Neural computation*, **15**(4), 937–963.
- Ellis, K. E., Bervoets, S., Smihula, H., Ganguly, I., Vigato, E., Auer, T. O., Benton, R., Litwin-Kumar, A., and Caron, S. J. C. (2024). Evolution of connectivity architecture in the drosophila mushroom body. *Nature Communications*, **15**(1), 4872.
- Emes, R. D. and Grant, S. G. (2012). Evolution of synapse complexity and diversity. *Annual review of neuroscience*, **35**(1), 111–131.
- Emes, R. D., Pocklington, A. J., Anderson, C. N., Bayes, A., Collins, M. O., Vickers, C. A., Croning, M. D., Malik, B. R., Choudhary, J. S., Armstrong, J. D., *et al.* (2008). Evolutionary expansion and anatomical specialization of synapse proteome complexity. *Nature neuroscience*, **11**(7), 799–806.
- Eyal, G., Mansvelder, H. D., de Kock, C. P., and Segev, I. (2014). Dendrites impact the encoding capabilities of the axon. *Journal of Neuroscience*, **34**(24), 8063–8071.
- Faisal, A. A., Selen, L. P., and Wolpert, D. M. (2008). Noise in the nervous system. *Nature reviews neuroscience*, **9**(4), 292–303.

- Fan, F.-L., Wang, M., Dong, H.-C., Ma, J., and Zeng, T. (2024). No one-size-fits-all neurons: Task-based neurons for artificial neural networks. *arXiv preprint arXiv:2405.02369*.
- Fardet, T., Giannakakis, E., Paulun, L., and Levina, A. (2024). Revising clustering and small-worldness in brain networks. *arXiv preprint arXiv:2401.15630*.
- Fernandes, A. R., Martins, J. P., Gomes, E. R., Mendes, C. S., and Teodoro, R. O. (2023). *Drosophila* motor neuron boutons remodel through membrane blebbing coupled with muscle contraction. *Nature Communications*, **14**(1), 3352.
- Finlay, B. L., Hersman, M. N., and Darlington, R. B. (1998). Patterns of vertebrate neurogenesis and the paths of vertebrate evolution. *Brain behavior and evolution*, **52**(4-5), 232–242.
- Fino, E., Glowinski, J., and Venance, L. (2005). Bidirectional activity-dependent plasticity at corticostriatal synapses. *Journal of Neuroscience*, **25**(49), 11279–11287.
- Fino, E., Packer, A. M., and Yuste, R. (2013). The logic of inhibitory connectivity in the neocortex. *The Neuroscientist*, **19**(3), 228–237.
- Fong, S., Rogell, B., Amcoff, M., Kotrschal, A., van der Bijl, W., Buechel, S. D., and Kolm, N. (2021). Rapid mosaic brain evolution under artificial selection for relative telencephalon size in the guppy (*poecilia reticulata*). *Science Advances*, **7**(46), eabj4314.
- Frémaux, N. and Gerstner, W. (2016). Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. *Frontiers in neural circuits*, **9**, 85.
- Frémaux, N., Sprekeler, H., and Gerstner, W. (2010). Functional requirements for reward-modulated spike-timing-dependent plasticity. *Journal of Neuroscience*, **30**(40), 13326–13337.
- Fukushima, K. (1980). Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological cybernetics*, **36**(4), 193–202.
- Fushiki, A., Kohsaka, H., and Nose, A. (2013). Role of sensory experience in functional development of *drosophila* motor circuits. *PLoS One*, **8**(4), e62199.
- Gauthier, D. J., Bollt, E., Griffith, A., and Barbosa, W. A. (2021). Next generation reservoir computing. *Nature communications*, **12**(1), 1–8.
- Gelinas, J. and Nguyen, P. (2007). Neuromodulation of hippocampal synaptic plasticity, learning, and memory by noradrenaline. *Central Nervous System Agents in Medicinal Chemistry (Formerly Current Medicinal Chemistry-Central Nervous System Agents)*, **7**(1), 17–33.

- Gerstner, W. and Kistler, W. M. (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press.
- Gerstner, W., Kempter, R., Van Hemmen, J. L., and Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature*, **383**(6595), 76–78.
- Giannakakis, E., Khajehabdollahi, S., and Levina, A. (2023). Environmental variability and network structure determine the optimal plasticity mechanisms in embodied agents. In *Artificial Life Conference Proceedings 35*, number 1, page 22.
- Giannakakis, E., Khajehabdollahi, S., and Levina, A. (2024a). Network Bottlenecks and Task Structure Control the Evolution of Interpretable Learning Rules in a Foraging Agent. *Artificial Life*, pages 1–18.
- Giannakakis, E., Vinogradov, O., Buendía, V., and Levina, A. (2024b). Structural influences on synaptic plasticity: The role of presynaptic connectivity in the emergence of *e/i* co-tuning. *PLOS Computational Biology*, **20**(10), 1–22.
- Gjorgjieva, J., Drion, G., and Marder, E. (2016). Computational implications of biophysical diversity and multiple timescales in neurons and synapses for circuit performance. *Current opinion in neurobiology*, **37**, 44–52.
- Glanzman, D. L. (2010). Common mechanisms of synaptic plasticity in vertebrates and invertebrates. *Current Biology*, **20**(1), R31–R36.
- Glimcher, P. W. (2011). Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, **108**(supplement_3), 15647–15654.
- Golding, N. L., Staff, N. P., and Spruston, N. (2002). Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature*, **418**(6895), 326–331.
- Gollo, L. L., Kinouchi, O., and Copelli, M. (2009). Active dendrites enhance neuronal dynamic range. *PLoS computational biology*, **5**(6), e1000402.
- Gourgou, E., Adiga, K., Goettemoeller, A., Chen, C., and Hsu, A. (2021). *Caenorhabditis elegans* learning in a structured maze is a multisensory behavior. *iscience*, 24: 102284.
- Grasso, F. W. and Basil, J. A. (2009). The evolution of flexible behavioral repertoires in cephalopod molluscs. *Brain, behavior and evolution*, **74**(3), 231–245.
- Guo, C. C., Ke, M. C., and Raymond, J. L. (2014). Cerebellar encoding of multiple candidate error cues in the service of motor learning. *Journal of Neuroscience*, **34**(30), 9880–9890.
- Ha, D., Dai, A., and Le, Q. V. (2016). Hypernetworks. *arXiv preprint arXiv:1609.09106*.

- Halvagal, M. S. and Zenke, F. (2023). The combination of hebbian and predictive plasticity learns invariant object representations in deep sensory networks. *Nature neuroscience*, **26**(11), 1906–1915.
- Hamidi, M., Khajehabdollahi, S., Giannakakis, E., Schäfer, T. J., Levina, A., and Wu, C. M. (2024). Modular growth of hierarchical networks: Efficient, general, and robust curriculum learning. In *ALIFE 2024: Proceedings of the 2024 Artificial Life Conference*. MIT Press.
- Happel, B. L. and Murre, J. M. (1994). Design and evolution of modular neural network architectures. *Neural networks*, **7**(6-7), 985–1004.
- Hattori, D., Aso, Y., Swartz, K. J., Rubin, G. M., Abbott, L., and Axel, R. (2017a). Representations of novelty and familiarity in a mushroom body compartment. *Cell*, **169**(5), 956–969.
- Hattori, R., Kuchibhotla, K. V., Froemke, R. C., and Komiyama, T. (2017b). Functions and dysfunctions of neocortical inhibitory neuron subtypes. *Nature neuroscience*, **20**(9), 1199–1208.
- Hausser, M., Spruston, N., and Stuart, G. J. (2000). Diversity and dynamics of dendritic signaling. *Science*, **290**(5492), 739–744.
- Hebb, D. O. (2005). *The organization of behavior: A neuropsychological theory*. Psychology press.
- Hennequin, G., Agnes, E. J., and Vogels, T. P. (2017). Inhibitory plasticity: balance, control, and codependence. *Annual review of neuroscience*, **40**(1), 557–579.
- Hennig, M. H. (2013). Theoretical models of synaptic short term plasticity. *Frontiers in computational neuroscience*, **7**, 45.
- Herculano-Houzel, S. (2012). The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost. *Proceedings of the National Academy of Sciences*, **109**(supplement_1), 10661–10668.
- Hertäg, L. and Clopath, C. (2022). Prediction-error neurons in circuits with multiple neuron types: Formation, refinement, and functional implications. *Proceedings of the National Academy of Sciences*, **119**(13), e2115699119.
- Hiesinger, P. R. (2018). The self-assembling brain: Contributions to morphogenetic engineering from a self-organizing neural network in the insect brain. In *Artificial Life Conference Proceedings*, pages 202–203. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info . . .

- Higo, S., Akashi, K., Sakimura, K., and Tamamaki, N. (2009). Subtypes of gabaergic neurons project axons in the neocortex. *Frontiers in neuroanatomy*, **3**, 842.
- Hihi, S. and Bengio, Y. (1995). Hierarchical recurrent neural networks for long-term dependencies. *Advances in neural information processing systems*, **8**.
- Hilgetag, C. C. and Kaiser, M. (2004). Clustered organization of cortical connectivity. *Neuroinformatics*, **2**, 353–360.
- Hinton, G. E. and Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *science*, **313**(5786), 504–507.
- Hodgkin, A. L. and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, **117**(4), 500.
- Holmgren, C. D. and Zilberter, Y. (2001). Coincident spiking activity induces long-term changes in inhibition of neocortical pyramidal cells. *Journal of Neuroscience*, **21**(20), 8270–8277.
- Holtmaat, A. and Svoboda, K. (2009). Experience-dependent structural synaptic plasticity in the mammalian brain. *Nature Reviews Neuroscience*, **10**(9), 647–658.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, **79**(8), 2554–2558.
- Howarth, C., Gleeson, P., and Attwell, D. (2012). Updated energy budgets for neural computation in the neocortex and cerebellum. *Journal of Cerebral Blood Flow & Metabolism*, **32**(7), 1222–1232.
- Huang, S., Rozas, C., Treviño, M., Contreras, J., Yang, S., Song, L., Yoshioka, T., Lee, H.-K., and Kirkwood, A. (2014). Associative hebbian synaptic plasticity in primate visual cortex. *Journal of Neuroscience*, **34**(22), 7575–7579.
- Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, **160**(1), 106.
- Iacaruso, M. F., Gasler, I. T., and Hofer, S. B. (2017). Synaptic organization of visual space in primary visual cortex. *Nature*, **547**(7664), 449–452.
- Illing, B., Ventura, J., Bellec, G., and Gerstner, W. (2021). Local plasticity rules can learn deep representations using self-supervised contrastive predictions. *Advances in neural information processing systems*, **34**, 30365–30379.

- Intrator, N. and Cooper, L. N. (1992). Objective function formulation of the bcm theory of visual cortical plasticity: Statistical connections, stability conditions. *Neural Networks*, **5**(1), 3–17.
- Ito, M., Sakurai, M., and Tongroach, P. (1982). Climbing fibre induced depression of both mossy fibre responsiveness and glutamate sensitivity of cerebellar purkinje cells. *The Journal of physiology*, **324**(1), 113–134.
- Jacobs, B., Johnson, N. L., Wahl, D., Schall, M., Maseko, B. C., Lewandowski, A., Raghanti, M. A., Wicinski, B., Butti, C., Hopkins, W. D., Bertelsen, M. F., Walsh, T., Roberts, J. R., Reep, R. L., Hof, P. R., Sherwood, C. C., and Manger, P. R. (2014). Comparative neuronal morphology of the cerebellar cortex in afrotherians, carnivores, cetartiodactyls, and primates. *Frontiers in Neuroanatomy*, **8**.
- Jaeger, H. and Haas, H. (2004). Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *science*, **304**(5667), 78–80.
- Jiang, X., Shen, S., Cadwell, C. R., Berens, P., Sinz, F., Ecker, A. S., Patel, S., and Tolias, A. S. (2015). Principles of connectivity among morphologically defined cell types in adult neocortex. *Science*, **350**(6264), aac9462.
- Jordan, J., Schmidt, M., Senn, W., and Petrovici, M. A. (2021). Evolving interpretable plasticity for spiking networks. *Elife*, **10**, e66273.
- Jörntell, H. (2017). Cerebellar physiology: links between microcircuitry properties and sensorimotor functions. *The Journal of physiology*, **595**(1), 11–27.
- Kaas, J. H. (2011). Neocortex in early mammals and its subsequent variations. *Annals of the New York Academy of Sciences*, **1225**(1), 28–36.
- Kadmon, J. and Sompolinsky, H. (2015). Transition to chaos in random neuronal networks. *Physical Review X*, **5**(4), 041030.
- Kampa, B. M., Letzkus, J. J., and Stuart, G. J. (2007). Dendritic mechanisms controlling spike-timing-dependent synaptic plasticity. *Trends in neurosciences*, **30**(9), 456–463.
- Kandel, E. R. (2001). The molecular biology of memory storage: a dialogue between genes and synapses. *Science*, **294**(5544), 1030–1038.
- Kastellakis, G., Cai, D. J., Mednick, S. C., Silva, A. J., and Poirazi, P. (2015). Synaptic clustering within dendrites: an emerging theory of memory formation. *Progress in neurobiology*, **126**, 19–35.
- Kastellakis, G., Silva, A. J., and Poirazi, P. (2016). Linking memories across time via neuronal and dendritic overlaps in model neurons with active dendrites. *Cell reports*, **17**(6), 1491–1504.

- Katz, P. S. and Harris-Warrick, R. M. (1999). The evolution of neuronal circuits underlying species-specific behavior. *Current Opinion in Neurobiology*, **9**(5), 628–633.
- Kavanau, J. L. (1990). Conservative behavioural evolution, the neural substrate. *Animal behaviour*, **39**(4), 758–767.
- Kepple, D., Engelken, R., and Rajan, K. (2022). Curriculum learning as a tool to uncover learning principles in the brain. In *International Conference on Learning Representations*.
- Kerr, B. and Feldman, M. (2003). Carving the cognitive niche: Optimal learning strategies in homogeneous and heterogeneous environments. *Journal of Theoretical Biology*, **220**(2), 169–188.
- Khajehabdollahi, S., Prosi, J., Giannakakis, E., Martius, G., and Levina, A. (2022). When to be critical? performance and evolvability in different regimes of neural ising agents. *Artificial Life*, **28**(4), 458–478.
- Khajehabdollahi, S., Giannakakis, E., Buendía, V., Martius, G., and Levina, A. (2023). Locally adaptive cellular automata for goal-oriented self-organization. In *Artificial Life Conference Proceedings 35*, volume 2023, page 59. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info
- Khajehabdollahi, S., Zeraati, R., Giannakakis, E., Schäfer, T. J., Martius, G., and Levina, A. (2024). Emergent mechanisms for long timescales depend on training curriculum and affect performance in memory tasks. In *The Twelfth International Conference on Learning Representations*.
- Kilpatrick, B. S., Magalhaes, J., Beavan, M. S., McNeill, A., Gegg, M. E., Cleeter, M. W., Bloor-Young, D., Churchill, G. C., Duchon, M. R., Schapira, A. H., *et al.* (2016). Endoplasmic reticulum and lysosomal ca²⁺ stores are remodelled in gba1-linked parkinson disease patient fibroblasts. *Cell Calcium*, **59**(1), 12–20.
- Knudsen, E. I. (1994). Supervised learning in the brain. *The Journal of Neuroscience*, **14**(7), 3985.
- Koch, G., Ponzo, V., Di Lorenzo, F., Caltagirone, C., and Veniero, D. (2013). Hebbian and anti-hebbian spike-timing-dependent plasticity of human cortico-cortical connections. *Journal of Neuroscience*, **33**(23), 9725–9733.
- Komatsu, Y. and Iwakiri, M. (1993). Long-term modification of inhibitory synaptic transmission in developing visual cortex. *Neuroreport*, **4**(7), 907–910.
- Lagzi, F. and Rotter, S. (2015). Dynamics of competition between subnetworks of spiking neuronal networks in the balanced state. *PloS one*, **10**(9), e0138947.

- Lakshminarasimhan, K. J., Xie, M., Cohen, J. D., Sauerbrei, B. A., Hantman, A. W., Litwin-Kumar, A., and Escola, S. (2024). Specific connectivity optimizes learning in thalamocortical loops. *Cell reports*, **43**(4).
- Lange, R. T. and Sprekeler, H. (2022). Learning not to learn: Nature versus nurture in silico. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 7290–7299.
- Lappalainen, J. K., Tschopp, F. D., Prakhya, S., McGill, M., Nern, A., Shinomiya, K., Takemura, S.-y., Gruntman, E., Macke, J. H., and Turaga, S. C. (2024). Connectome-constrained networks predict neural activity across the fly visual system. *Nature*, pages 1–9.
- Larkman, A. U. (1991). Dendritic morphology of pyramidal neurones of the visual cortex of the rat: I. branching patterns. *Journal of Comparative Neurology*, **306**(2), 307–319.
- Larsen, R. S. and Sjöström, P. J. (2015). Synapse-type-specific plasticity in local circuits. *Current opinion in neurobiology*, **35**, 127–135.
- Law, C. C. and Cooper, L. N. (1994). Formation of receptive fields in realistic visual environments according to the bienenstock, cooper, and munro (bcm) theory. *Proceedings of the National Academy of Sciences*, **91**(16), 7797–7801.
- Lazar, A., Pipa, G., and Triesch, J. (2009). Sorn: a self-organizing recurrent neural network. *Frontiers in computational neuroscience*, **3**, 800.
- LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., and Jackel, L. D. (1989). Backpropagation applied to handwritten zip code recognition. *Neural computation*, **1**(4), 541–551.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *nature*, **521**(7553), 436–444.
- Lee, J. H., Mannelli, S. S., and Saxe, A. (2024a). Why do animals need shaping? a theory of task composition and curriculum learning. *arXiv preprint arXiv:2402.18361*.
- Lee, S., Garcia, S. L., Clopath, C., and Dabney, W. (2024b). Lifelong reinforcement learning via neuromodulation. *arXiv preprint arXiv:2408.08446*.
- Legenstein, R. and Maass, W. (2007). Edge of chaos and prediction of computational performance for neural circuit models. *Neural networks*, **20**(3), 323–334.
- Lillicrap, T. P., Cownden, D., Tweed, D. B., and Akerman, C. J. (2016). Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, **7**(1), 13276.

- Linnainmaa, S. (1970). *The representation of the cumulative rounding error of an algorithm as a Taylor expansion of the local rounding errors*. Ph.D. thesis, Master's Thesis (in Finnish), Univ. Helsinki.
- Litwin-Kumar, A. and Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature neuroscience*, **15**(11), 1498–1505.
- Litwin-Kumar, A. and Doiron, B. (2014). Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nature communications*, **5**(1), 5319.
- Liu, Y. H., Smith, S., Mihalas, S., Shea-Brown, E., and Sümbül, U. (2021). Cell-type-specific neuromodulation guides synaptic credit assignment in a spiking neural network. *Proceedings of the National Academy of Sciences*, **118**(51), e2111821118.
- Liu, Y. H., Smith, S., Mihalas, S., Shea-Brown, E., and Sümbül, U. (2022). Biologically-plausible backpropagation through arbitrary timespans via local neuromodulators. *Advances in Neural Information Processing Systems*, **35**, 17528–17542.
- Lukoševičius, M. and Jaeger, H. (2009). Reservoir computing approaches to recurrent neural network training. *Computer science review*, **3**(3), 127–149.
- Lund, J. S. (2002). Specificity and non-specificity of synaptic connections within mammalian visual cortex. *Journal of neurocytology*, **31**(3), 203–209.
- Maass, W. (2011). Liquid state machines: motivation, theory, and applications. *Computability in context: computation and logic in the real world*, pages 275–296.
- Malenka, R. C. and Bear, M. F. (2004). Ltp and ltd: an embarrassment of riches. *Neuron*, **44**(1), 5–21.
- Marchetto, M. C., Hrvoj-Mihic, B., Kerman, B. E., Yu, D. X., Vadodaria, K. C., Linker, S. B., Narvaiza, I., Santos, R., Denli, A. M., Mendes, A. P., Oefner, R., Cook, J., McHenry, L., Grasmick, J. M., Heard, K., Fredlender, C., Randolph-Moore, L., Kshirsagar, R., Xenitopoulos, R., Chou, G., Hah, N., Muotri, A. R., Padmanabhan, K., Semendeferi, K., and Gage, F. H. (2019). Species-specific maturation profiles of human, chimpanzee and bonobo neural cells. *eLife*, **8**, e37527.
- Markov, N. T., Ercsey-Ravasz, M., Lamy, C., Ribeiro Gomes, A. R., Magrou, L., Misery, P., Giroud, P., Barone, P., Dehay, C., Toroczka, Z., *et al.* (2013). The role of long-range connections on the specificity of the macaque interareal cortical network. *Proceedings of the National Academy of Sciences*, **110**(13), 5187–5192.
- Markram, H., Gerstner, W., and Sjöström, P. J. (2011). A history of spike-timing-dependent plasticity. *Frontiers in synaptic neuroscience*, **3**, 4.

- Masland, R. H. (2004). Neuronal cell types. *Current Biology*, **14**(13), R497–R500.
- Masuda-Nakagawa, L. M., Gendre, N., O’Kane, C. J., and Stocker, R. F. (2009). Localized olfactory representation in mushroom bodies of drosophila larvae. *Proceedings of the National Academy of Sciences*, **106**(25), 10314–10319.
- Max, K., Kriener, L., Pineda García, G., Nowotny, T., Jaras, I., Senn, W., and Petrovici, M. A. (2024). Learning efficient backprojections across cortical hierarchies in real time. *Nature Machine Intelligence*, pages 1–12.
- McCulloch, W. S. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, **5**, 115–133.
- McGuire, S. E., Le, P. T., and Davis, R. L. (2001). The role of drosophila mushroom body signaling in olfactory memory. *Science*, **293**(5533), 1330–1333.
- Mehta, Y., Tyulmankov, D., Rajagopalan, A. E., Turner, G. C., Fitzgerald, J. E., and Funke, J. (2023). Model based inference of synaptic plasticity rules. *bioRxiv*, pages 2023–12.
- Mel, B. W. (1994). Information processing in dendritic trees. *Neural computation*, **6**(6), 1031–1085.
- Menzel, R. (2014). The insect mushroom body, an experience-dependent recoding device. *Journal of Physiology-Paris*, **108**(2-3), 84–95.
- Metz, H. C., Bedford, N. L., Pan, Y. L., and Hoekstra, H. E. (2017). Evolution and genetics of precocious burrowing behavior in peromyscus mice. *Current Biology*, **27**(24), 3837–3845.
- Meunier, D., Lambiotte, R., and Bullmore, E. T. (2010). Modular and hierarchically modular organization of brain networks. *Frontiers in neuroscience*, **4**, 200.
- Miconi, T., Stanley, K., and Clune, J. (2018). Differentiable plasticity: training plastic neural networks with backpropagation. In *International Conference on Machine Learning*, pages 3559–3568. PMLR.
- Miehl, C., Onasch, S., Festa, D., and Gjorgjieva, J. (2023). Formation and computational implications of assemblies in neural circuits. *The Journal of Physiology*, **601**(15), 3071–3090.
- Mikulasch, F. A., Rudelt, L., and Priesemann, V. (2021). Local dendritic balance enables learning of efficient representations in networks of spiking neurons. *Proceedings of the National Academy of Sciences*, **118**(50), e2021925118.

- Mishra, R. K., Kim, S., Guzman, S. J., and Jonas, P. (2016). Symmetric spike timing-dependent plasticity at ca3–ca3 synapses optimizes storage and recall in autoassociative networks. *Nature communications*, **7**(1), 1–11.
- Mitchell, T. M. (1980). The need for biases in learning generalizations (rutgers computer science tech. rept. cbm-tr-117). *Rutgers University*.
- Modi, M. N., Shuai, Y., and Turner, G. C. (2020). The drosophila mushroom body: from architecture to algorithm in a learning circuit. *Annual review of neuroscience*, **43**(1), 465–484.
- Mohajerin, N. and Waslander, S. L. (2014). Modular deep recurrent neural network: Application to quadrotors. In *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 1374–1379. IEEE.
- Montero, M. L., Plantec, E., Nisioti, E., Pedersen, J. W., and Risi, S. (2024). Meta-learning an evolvable developmental encoding. *arXiv preprint arXiv:2406.09020*.
- Moskovitz, T. H., Litwin-Kumar, A., and Abbott, L. (2018). Feedback alignment in deep convolutional networks. *arXiv preprint arXiv:1812.06488*.
- Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., *et al.* (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature neuroscience*, **17**(12), 1661–1663.
- Nadim, F. and Bucher, D. (2014). Neuromodulation of neurons and synapses. *Current opinion in neurobiology*, **29**, 48–56.
- Najafi, F., Elsayed, G. F., Cao, R., Pnevmatikakis, E., Latham, P. E., Cunningham, J. P., and Churchland, A. K. (2020). Excitatory and inhibitory subnetworks are equally selective during decision-making and emerge simultaneously during learning. *Neuron*, **105**(1), 165–179.
- Najarro, E. and Risi, S. (2020). Meta-learning through hebbian plasticity in random networks. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20719–20731. Curran Associates, Inc.
- Najarro, E., Sudhakaran, S., and Risi, S. (2023). Towards self-assembling artificial neural networks through neural developmental programs. In *Artificial Life Conference Proceedings 35*, volume 2023, page 80. MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info
- Nakajima, K. (2018). Reservoir computing: Theory, physical implementations, and applications. *IEICE Technical Report; IEICE Tech. Rep.*, **118**(220), 149–154.

- Nakajima, K. and Fischer, I. (2021). *Reservoir computing*. Springer.
- Nicoll, R. A. (2017). A brief history of long-term potentiation. *neuron*, **93**(2), 281–290.
- Niell, C. M. and Stryker, M. P. (2008). Highly selective receptive fields in mouse visual cortex. *Journal of Neuroscience*, **28**(30), 7520–7536.
- Nøkland, A. (2016). Direct feedback alignment provides learning in deep neural networks. *Advances in neural information processing systems*, **29**.
- Nomura, T., Yamashita, W., Gotoh, H., and Ono, K. (2018). Species-specific mechanisms of neuron subtype specification reveal evolutionary plasticity of amniote brain development. *Cell Reports*, **22**(12), 3142–3151.
- Nwankpa, C., Ijomah, W., Gachagan, A., and Marshall, S. (2018). Activation functions: Comparison of trends in practice and research for deep learning.
- O’Donnell, C. (2023). Nonlinear slow-timescale mechanisms in synaptic plasticity. *Current opinion in neurobiology*, **82**, 102778.
- Oja, E. (1992). Principal components, minor components, and linear neural networks. *Neural networks*, **5**(6), 927–935.
- Okun, M. and Lampl, I. (2008). Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nature neuroscience*, **11**(5), 535–537.
- Orvis, J., Albertin, C. B., Shrestha, P., Chen, S., Zheng, M., Rodriguez, C. J., Tallon, L. J., Mahurkar, A., Zimin, A. V., Kim, M., *et al.* (2022). The evolution of synaptic and cognitive capacity: Insights from the nervous system transcriptome of aplysia. *Proceedings of the National Academy of Sciences*, **119**(28), e2122301119.
- Ovsepián, S. V., O’Leary, V. B., and Vesselkin, N. P. (2020). Evolutionary origins of chemical synapses. *Vitamins and hormones*, **114**, 1–21.
- Palmer, L. M., Shai, A. S., Reeve, J. E., Anderson, H. L., Paulsen, O., and Larkum, M. E. (2014). Nmda spikes enhance action potential generation during sensory input. *Nature neuroscience*, **17**(3), 383–390.
- Papini, M. R. and Torres, C. (2017). Comparative learning and evolution.
- Partridge, J. G., Apparsundaram, S., Gerhardt, G. A., Ronesi, J., and Lovinger, D. M. (2002). Nicotinic acetylcholine receptors interact with dopamine in induction of striatal long-term depression. *Journal of Neuroscience*, **22**(7), 2541–2549.
- Payeur, A., Guerguiev, J., Zenke, F., Richards, B. A., and Naud, R. (2021). Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nature neuroscience*, **24**(7), 1010–1019.

- Pedersen, J. W. and Risi, S. (2021a). Evolving and merging hebbian learning rules: Increasing generalization by decreasing the number of rules. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '21*, page 892–900, New York, NY, USA. Association for Computing Machinery.
- Pedersen, J. W. and Risi, S. (2021b). Evolving and merging hebbian learning rules: increasing generalization by decreasing the number of rules. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 892–900.
- Pehlevan, C. and Chklovskii, D. (2015). A normative theory of adaptive dimensionality reduction in neural networks. *Advances in neural information processing systems*, **28**.
- Perin, R., Berger, T. K., and Markram, H. (2011). A synaptic organizing principle for cortical neuronal groups. *Proceedings of the National Academy of Sciences*, **108**(13), 5419–5424.
- Pfister, J.-P. and Gerstner, W. (2006). Triplets of spikes in a model of spike timing-dependent plasticity. *Journal of Neuroscience*, **26**(38), 9673–9682.
- Plantec, E., Pedersen, J. W., Montero, M. L., Nisioti, E., and Risi, S. (2024). Evolving self-assembling neural networks: From spontaneous activity to experience-dependent learning. In *ALIFE 2024: Proceedings of the 2024 Artificial Life Conference*. MIT Press.
- Poirazi, P., Brannon, T., and Mel, B. W. (2003). Pyramidal neuron as two-layer neural network. *Neuron*, **37**(6), 989–999.
- Pontes, A. C., Mobley, R. B., Ofria, C., Adami, C., and Dyer, F. C. (2020). The evolutionary origin of associative learning. *The American Naturalist*, **195**(1), E1–E19.
- Popa, L. S. and Ebner, T. J. (2019). Cerebellum, predictions and errors. *Frontiers in cellular neuroscience*, **12**, 524.
- Rajagopalan, A. E., Darshan, R., Hibbard, K. L., Fitzgerald, J. E., and Turner, G. C. (2023). Reward expectations direct learning and drive operant matching in drosophila. *Proceedings of the National Academy of Sciences*, **120**(39), e2221415120.
- Ramesh, P., Confavreux, B., Gonçalves, P. J., Vogels, T. P., and Macke, J. H. (2024). Indistinguishable network dynamics can emerge from unlike plasticity rules.
- Redies, C. and Puelles, L. (2001). Modularity in vertebrate brain development and evolution. *Bioessays*, **23**(12), 1100–1111.
- Regehr, W. G. (2012). Short-term presynaptic plasticity. *Cold Spring Harbor perspectives in biology*, **4**(7), a005702.

- Richards, B. A. and Kording, K. P. (2023). The study of plasticity has always been about gradients. *The Journal of Physiology*, **601**(15), 3141–3149.
- Richter, J. N., Hochner, B., and Kuba, M. J. (2016). Pull or push? octopuses solve a puzzle problem. *PloS one*, **11**(3), e0152048.
- Risi, S. (2021). The future of artificial intelligence is self-organizing and self-assembling. *sebastianrisi.com*, page 14.
- Roberts, A. C. and Glanzman, D. L. (2003). Learning in aplysia: looking at synaptic plasticity from both sides. *Trends in neurosciences*, **26**(12), 662–670.
- Roberts, P. D. and Leen, T. K. (2010). Anti-hebbian spike-timing-dependent plasticity and adaptive sensory processing. *Frontiers in computational neuroscience*, **4**, 156.
- Roelfsema, P. R., van Ooyen, A., and Watanabe, T. (2010). Perceptual learning rules based on reinforcers and attention. *Trends in cognitive sciences*, **14**(2), 64–71.
- Rosenbaum, R. and Doiron, B. (2014). Balanced networks of spiking neurons with spatially dependent recurrent connections. *Physical Review X*, **4**(2), 021039.
- Rosenbaum, R., Smith, M. A., Kohn, A., Rubin, J. E., and Doiron, B. (2017). The spatial structure of correlated neuronal variability. *Nature neuroscience*, **20**(1), 107–114.
- Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, **65**(6), 386.
- Rossi, L. F., Harris, K. D., and Carandini, M. (2020). Spatial connectivity matches direction selectivity in visual cortex. *Nature*, **588**(7839), 648–652.
- Rudy, B., Fishell, G., Lee, S., and Hjerling-Leffler, J. (2011). Three groups of interneurons account for nearly 100% of neocortical gabaergic neurons. *Developmental neurobiology*, **71**(1), 45–61.
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning representations by back-propagating errors. *nature*, **323**(6088), 533–536.
- Ryan, T. J. and Grant, S. G. (2009). The origin and evolution of synapses. *Nature Reviews Neuroscience*, **10**(10), 701–712.
- Sacramento, J., Ponte Costa, R., Bengio, Y., and Senn, W. (2018). Dendritic cortical microcircuits approximate the backpropagation algorithm. *Advances in neural information processing systems*, **31**.
- Schliebs, S., Mohemmed, A., and Kasabov, N. (2011). Are probabilistic spiking neural networks suitable for reservoir computing? In *The 2011 International Joint Conference on Neural Networks*, pages 3156–3163. IEEE.

- Schmidgall, S., Ashkanazy, J., Lawson, W., and Hays, J. (2021). Spikepropamine: Differentiable plasticity in spiking neural networks. *Frontiers in neurorobotics*, **15**, 629210.
- Schmidgall, S., Ziaei, R., Achterberg, J., Kirsch, L., Hajiseyedrazi, S., and Eshraghian, J. (2024). Brain-inspired learning in artificial neural networks: a review. *APL Machine Learning*, **2**(2).
- Schrauwen, B., Verstraeten, D., and Van Campenhout, J. (2007). An overview of reservoir computing: theory, applications and implementations. In *Proceedings of the 15th european symposium on artificial neural networks*. p. 471-482 2007, pages 471–482.
- Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, **275**(5306), 1593–1599.
- Schulz, A., Miehl, C., Berry II, M. J., and Gjorgjieva, J. (2021). The generation of cortical novelty responses through inhibitory plasticity. *Elife*, **10**, e65309.
- Schumacher, E. L. and Carlson, B. A. (2022). Convergent mosaic brain evolution is associated with the evolution of novel electrosensory systems in teleost fishes. *Elife*, **11**, e74159.
- Shahi, S., Fenton, F. H., and Cherry, E. M. (2022). Prediction of chaotic time series using recurrent neural networks and reservoir computing techniques: A comparative study. *Machine learning with applications*, **8**, 100300.
- Shervani-Tabar, N. and Rosenbaum, R. (2023). Meta-learning biologically plausible plasticity rules with random feedback pathways. *Nature Communications*, **14**.
- Shettleworth, S. J. (1993). Varieties of learning and memory in animals. *Journal of Experimental Psychology: Animal Behavior Processes*, **19**(1), 5.
- Shi, Y.-L., Zeraati, R., Levina, A., and Engel, T. A. (2023). Spatial and temporal correlations in neural networks with structured connectivity. *Physical Review Research*, **5**(1), 013005.
- Shuvaev, S., Lachi, D., Koulakov, A., and Zador, A. (2024). Encoding innate ability through a genomic bottleneck. *Proceedings of the National Academy of Sciences*, **121**(38), e2409160121.
- Skou, J. (1990). The energy coupled exchange of na⁺ for k⁺ across the cell membrane: The na⁺, k⁺-pump. *FEBS letters*, **268**(2), 314–324.
- Smith, D., Wessnitzer, J., and Webb, B. (2008). A model of associative learning in the mushroom body. *Biological cybernetics*, **99**, 89–103.

- Smith, S. L., Smith, I. T., Branco, T., and Häusser, M. (2013). Dendritic spikes enhance stimulus selectivity in cortical neurons in vivo. *Nature*, **503**(7474), 115–120.
- Sohya, K., Kameyama, K., Yanagawa, Y., Obata, K., and Tsumoto, T. (2007). Gabaergic neurons are less selective to stimulus orientation than excitatory neurons in layer ii/iii of visual cortex, as revealed by in vivo functional ca²⁺ imaging in transgenic mice. *Journal of Neuroscience*, **27**(8), 2145–2149.
- Song, S., Sjöström, P. J., Reigl, M., Nelson, S., and Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS biology*, **3**(3), e68.
- Soviany, P., Ionescu, R. T., Rota, P., and Sebe, N. (2022). Curriculum learning: A survey. *International Journal of Computer Vision*, **130**(6), 1526–1565.
- Speranza, L., Di Porzio, U., Viggiano, D., de Donato, A., and Volpicelli, F. (2021). Dopamine: the neuromodulator of long-term synaptic plasticity, reward and movement control. *Cells*, **10**(4), 735.
- Spieler, A., Rahaman, N., Martius, G., Schölkopf, B., and Levina, A. (2023). The expressive leaky memory neuron: an efficient and expressive phenomenological neuron model can solve long-horizon tasks. In *The Twelfth International Conference on Learning Representations*.
- Sporns, O. (2015). Cerebral cartography and connectomics. *Philosophical Transactions of the Royal Society B: Biological Sciences*, **370**(1668), 20140173.
- Sporns, O. and Betzel, R. F. (2016). Modular brain networks. *Annual review of psychology*, **67**(1), 613–640.
- Squire, L. R. and Zola-Morgan, J. (1991). The cognitive neuroscience of human memory since hm. *Annual review of neuroscience*, **34**(1), 259–288.
- Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, **8**, 131–162.
- Stanley, K. O. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary computation*, **10**(2), 99–127.
- Stein, R. B., Gossen, E. R., and Jones, K. E. (2005). Neuronal variability: noise or part of the signal? *Nature Reviews Neuroscience*, **6**(5), 389–397.
- Stevens, C. F. (1998). Neuronal diversity: too many cell types for comfort? *Current Biology*, **8**(20), R708–R710.

- Tanaka, G., Yamane, T., Héroux, J. B., Nakane, R., Kanazawa, N., Takeda, S., Numata, H., Nakano, D., and Hirose, A. (2019). Recent advances in physical reservoir computing: A review. *Neural Networks*, **115**, 100–123.
- Tejero-Cantero, A., Boelts, J., Deistler, M., Lueckmann, J.-M., Durkan, C., Gonçalves, P. J., Greenberg, D. S., and Macke, J. H. (2020). sbi: A toolkit for simulation-based inference. *Journal of Open Source Software*, **5**(52), 2505.
- Tierney, A. (1995). Evolutionary implications of neural circuit structure and function. *Behavioural processes*, **35**(1-3), 173–182.
- Tinbergen, N. (2020). *The study of instinct*. Pygmalion Press, an imprint of Plunkett Lake Press.
- Tomioka, R., Okamoto, K., Furuta, T., Fujiyama, F., Iwasato, T., Yanagawa, Y., Obata, K., Kaneko, T., and Tamamaki, N. (2005). Demonstration of long-range gabaergic connections distributed throughout the mouse neocortex. *European Journal of Neuroscience*, **21**(6), 1587–1600.
- Tremblay, R., Lee, S., and Rudy, B. (2016). Gabaergic interneurons in the neocortex: from cellular properties to circuits. *Neuron*, **91**(2), 260–292.
- Tyulmankov, D., Yang, G. R., and Abbott, L. (2022). Meta-learning synaptic plasticity and memory addressing for continual familiarity detection. *Neuron*, **110**(3), 544–557.
- Udvary, D., Harth, P., Macke, J. H., Hege, H.-C., de Kock, C. P., Sakmann, B., and Oberlaender, M. (2022). The impact of neuron morphology on cortical network architecture. *Cell Reports*, **39**(2).
- Verkhatsky, A. J. and Petersen, O. H. (1998). Neuronal calcium stores. *Cell calcium*, **24**(5-6), 333–343.
- Vinogradov, O., Giannakakis, E., Buendia, V., Uysal, B., Ron, S., Weinreb, E., Schwarz, N., Lerche, H., Moses, E., and Levina, A. (2024). Effective excitability captures network dynamics across development and phenotypes. *bioRxiv*, pages 2024–08.
- Vogels, T. P. and Abbott, L. (2009). Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nature neuroscience*, **12**(4), 483–491.
- Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science*, **334**(6062), 1569–1573.
- Wang, X., Chen, Y., and Zhu, W. (2021). A survey on curriculum learning. *IEEE transactions on pattern analysis and machine intelligence*, **44**(9), 4555–4576.

- Wang, Y., Royer, J., Park, B.-y., Vos de Wael, R., Larivière, S., Tavakol, S., Rodriguez-Cruces, R., Paquola, C., Hong, S.-J., Margulies, D. S., *et al.* (2023). Long-range functional connections mirror and link microarchitectural and cognitive hierarchies in the human brain. *Cerebral Cortex*, **33**(5), 1782–1798.
- Watts, J. and Thomson, A. M. (2005). Excitatory and inhibitory connections show selectivity in the neocortex. *The Journal of physiology*, **562**(1), 89–97.
- Weber, J. N. and Hoekstra, H. E. (2009). The evolution of burrowing behaviour in deer mice (genus *Peromyscus*). *Animal Behaviour*, **77**(3), 603–609.
- Weiler, S., Nilo, D. G., Bonhoeffer, T., Hübener, M., Rose, T., and Scheuss, V. (2022). Orientation and direction tuning align with dendritic morphology and spatial connectivity in mouse visual cortex. *Current Biology*, **32**(8), 1743–1753.
- Werbos, P. J. (1994). *The roots of backpropagation: from ordered derivatives to neural networks and political forecasting*, volume 1. John Wiley & Sons.
- Wiesel, T. N. and Hubel, D. H. (1963). Effects of visual deprivation on morphology and physiology of cells in the cat's lateral geniculate body. *Journal of neurophysiology*, **26**(6), 978–993.
- Wilmes, K. A. and Clopath, C. (2019). Inhibitory microcircuits for top-down plasticity of sensory representations. *Nature communications*, **10**(1), 5055.
- Wilmes, K. A. and Clopath, C. (2023). Dendrites help mitigate the plasticity-stability dilemma. *Scientific Reports*, **13**(1), 6543.
- Wilson, D. E., Whitney, D. E., Scholl, B., and Fitzpatrick, D. (2016). Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex. *Nature neuroscience*, **19**(8), 1003–1009.
- Winding, M., Pedigo, B. D., Barnes, C. L., Patsolic, H. G., Park, Y., Kazimiers, T., Fushiki, A., Andrade, I. V., Khandelwal, A., Valdes-Aleman, J., *et al.* (2023). The connectome of an insect brain. *Science*, **379**(6636), eadd9330.
- Wong, J. S. and Rankin, C. H. (2019). *Caenorhabditis elegans* learning and memory. In *Oxford Research Encyclopedia of Neuroscience*.
- Woodin, M. A., Ganguly, K., and Poo, M.-m. (2003). Coincident pre- and postsynaptic activity modifies GABAergic synapses by postsynaptic changes in Cl⁻ transporter activity. *Neuron*, **39**(5), 807–820.
- Wu, Y. K., Hengen, K. B., Turrigiano, G. G., and Gjorgjieva, J. (2020). Homeostatic mechanisms regulate distinct aspects of cortical circuit dynamics. *Proceedings of the National Academy of Sciences*, **117**(39), 24514–24525.

- Yaeger, C. E., Vardalaki, D., Zhang, Q., Pham, T. L., Brown, N. J., Ji, N., and Harnett, M. T. (2024). A dendritic mechanism for balancing synaptic flexibility and stability. *Cell reports*, **43**(8).
- Yaman, A., Iacca, G., Mocanu, D. C., Coler, M., Fletcher, G., and Pechenizkiy, M. (2021a). Evolving plasticity for autonomous learning under changing environmental conditions. *Evolutionary computation*, **29**(3), 391–414.
- Yaman, A., Iacca, G., Mocanu, D. C., Coler, M., Fletcher, G., and Pechenizkiy, M. (2021b). Evolving Plasticity for Autonomous Learning under Changing Environmental Conditions. *Evolutionary Computation*, **29**(3), 391–414.
- Yerushalmi, U. and Teicher, M. (2008). Evolving synaptic plasticity with an evolutionary cellular development model. *Plos one*, **3**(11), e3697.
- Yilmaz, M. and Meister, M. (2013). Rapid innate defensive responses of mice to looming visual stimuli. *Current Biology*, **23**(20), 2011–2015.
- Yoshimura, Y. and Callaway, E. M. (2005). Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity. *Nature neuroscience*, **8**(11), 1552–1559.
- Yuan, X., Savarese, P., and Maire, M. (2024). Accelerated training via incrementally growing neural networks using variance transfer and learning rate adaptation. *Advances in Neural Information Processing Systems*, **36**.
- Zador, A. M. (2019). A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature communications*, **10**(1), 3770.
- Zenke, F. and Gerstner, W. (2017). Hebbian plasticity requires compensatory processes on multiple timescales. *Philosophical Transactions of the Royal Society B: Biological Sciences*, **372**(1715), 20160259.
- Zenke, F., Agnes, E. J., and Gerstner, W. (2015). Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nature communications*, **6**(1), 6922.
- Zeraati, R., Shi, Y.-L., Steinmetz, N. A., Gieselmann, M. A., Thiele, A., Moore, T., Levina, A., and Engel, T. A. (2023). Intrinsic timescales in the visual cortex change with selective attention and reflect spatial connectivity. *Nature communications*, **14**(1), 1858.
- Zeraati, R., Levina, A., Macke, J. H., and Gao, R. (2024). Neural timescales from a computational perspective. *arXiv preprint arXiv:2409.02684*.

- Zhang, R., Pitkow, X., and Angelaki, D. E. (2024). Inductive biases of neural network modularity in spatial navigation. *Science Advances*, **10**(29), eadk1256.
- Zhou, S. and Yu, Y. (2018). Synaptic ei balance underlies efficient neural coding. *Frontiers in neuroscience*, **12**, 46.
- Zucker, R. S. and Regehr, W. G. (2002). Short-term synaptic plasticity. *Annual review of physiology*, **64**(1), 355–405.

Appendix A Environmental variability and network structure determine the optimal plasticity mechanisms in embodied agents

Contains the full text of the publication: Giannakakis, E., Khajehabdollahi, S., and Levina, A. (2023). Environmental variability and network structure determine the optimal plasticity mechanisms in embodied agents. In *Artificial Life Conference Proceedings* 35, number 1, page 22

Environmental variability and network structure determine the optimal plasticity mechanisms in embodied agents

Emmanouil Giannakakis^{1,2}, Sina Khajehabdollahi¹ and Anna Levina^{1,2,3}

¹Department of Computer Science, University of Tübingen, Tübingen, Germany

²Max Planck Institute for Biological Cybernetics, Tübingen, Germany

³Bernstein Center for Computational Neuroscience Tübingen, Tübingen, Germany

Abstract

The evolutionary balance between innate and learned behaviors is highly intricate, and different organisms have found different solutions to this problem. We hypothesize that the emergence and exact form of learning behaviors is naturally connected with the statistics of environmental fluctuations and tasks an organism needs to solve. Here, we study how different aspects of simulated environments shape an evolved synaptic plasticity rule in static and moving artificial agents. We demonstrate that environmental fluctuation and uncertainty control the reliance of artificial organisms on plasticity. Interestingly, the form of the emerging plasticity rule is additionally determined by the details of the task the artificial organisms are aiming to solve. Moreover, we show that co-evolution between static connectivity and interacting plasticity mechanisms in distinct sub-networks changes the function and form of the emerging plasticity rules in embodied agents performing a foraging task.

Introduction

One of the defining features of living organisms is their ability to adapt to their environment and incorporate new information to modify their behavior. It is unclear how the ability to learn first evolved (Papini, 2012), but its utility appears evident. Natural environments are too complex for all the necessary information to be hardcoded genetically (Snell-Rood, 2013) and more importantly, they keep changing during an organism's lifetime in ways that cannot be anticipated (Ellefsen, 2014; Dunlap and Stephens, 2016). The link between learning and environmental uncertainty and fluctuation has been extensively demonstrated in both natural (Kerr and Feldman, 2003; Snell-Rood and Steck, 2019), and artificial environments (Nolfi and Parisi, 1996).

Nevertheless, the ability to learn does not come without costs. For the capacity to learn to be beneficial in evolutionary terms, a costly nurturing period is often required, a phenomenon observed in both biological (Thornton and Clutton-Brock, 2011), and artificial organisms (Eskridge and Hougen, 2012). Additionally, it has been shown that in some complex environments, hardcoded behaviors may be superior to learned ones given limits in the agent's lifetime

and environmental uncertainty (Dunlap and Stephens, 2009; Fawcett et al., 2012; Lange and Sprekeler, 2020).

The theoretical investigation of the optimal balance between learned and innate behaviors in natural and artificial systems goes back several decades. However, it has recently found also a wide range of applications in applied AI systems (Lee and Lee, 2020; Biesialska et al., 2020). Most AI systems are trained for specific tasks, and have no need for modification after their training has been completed. Still, technological advances and the necessity to solve broad families of tasks make discussions about life-like AI systems relevant to a wide range of potential application areas. Thus the idea of open-ended AI agents (Open Ended Learning Team et al., 2021) that can continually interact with and adapt to changing environments has become particularly appealing.

Many different approaches for introducing lifelong learning in artificial agents have been proposed. Some of them draw direct inspiration from actual biological systems (Schmidhuber, 1987; Parisi et al., 2019). Among them, the most biologically plausible solution is to equip artificial neural networks with some local neural plasticity (Thangarasa et al., 2020), similar to the large variety of synaptic plasticity mechanisms (Citri and Malenka, 2008; Feldman, 2009; Caroni et al., 2012) that performs the bulk of the learning in the brains of living organisms (Magee and Grienberger, 2020). The artificial plasticity mechanisms can be optimized to modify the connectivity of the artificial neural networks toward solving a particular task. The optimization can use a variety of approaches, most commonly evolutionary computation.

The idea of meta-learning or optimizing synaptic plasticity rules to perform specific functions has been recently established as an engineering tool that can compete with state-of-the-art machine learning algorithms on various complex tasks (Burms et al., 2015; Najarro and Risi, 2020; Pedersen and Risi, 2021; Yaman et al., 2021). Additionally, it can be used to reverse engineer actual plasticity mechanisms found in biological neural networks and uncover their functions (Confavreux et al., 2020; Jordan et al., 2021).

Here, we study the effect that different factors (environ-

mental fluctuation and reliability, task complexity) have on the form of evolved functional reward-modulated plasticity rules. We investigate the evolution of plasticity rules in static, single-layer simple networks. Then we increase the complexity by switching to moving agents performing a complex foraging task. In both cases, we study the impact of different environmental parameters on the form of the evolved plasticity mechanisms and the interaction of learned and static network connectivity. Interestingly, we find that different environmental conditions and different combinations of static and plastic connectivity have a very large impact on the resulting plasticity rules.

Methods

Environment

We imagine an agent who must forage to survive in an environment presenting various types of complex food particles. Each food particle is composed of various amounts and combinations of N ingredients that can have positive (food) or negative (poison) values. The value of a food particle is a weighted sum of its ingredients. To predict the reward value of a given resource, the agent must learn the values of these ingredients by interacting with the environment. The priors could be generated by genetic memory, but the exact values are subject to change.

To introduce environmental variability, we stochastically change the values of the ingredients. More precisely, we define two ingredient-value distributions E_1 and E_2 (Guttenberg, 2019) and switch between them, with probability p_{tr} for every time step. We control how (dis)similar the environments are by parametrically setting $E_2 = (1 - 2d_e)E_1$, with $d_e \in [0, 1]$ serving as a distance proxy for the environments; when $d_e = 0$, the environment remains unchanged, and when $d_e = 1$ the value of each ingredient fully reverses when the environmental transition happens. For simplicity, we take values of the ingredients in E_1 equally spaced between -1 and 1 (for the visualization, see Fig. 3a, b).

Static agent

The static agent receives passively presented food as a vector of ingredients and can assess its compound value using the linear summation of its sensors with the (learned or evolved) weights, see Fig. 1. The network consists of N sensory neurons that are projecting to a single post-synaptic neuron. At each time step, an input $X_t = (x_1, \dots, x_N)$ is presented, where the value x_i , $i \in \{1, \dots, N\}$ represents the quantity of the ingredient i . We draw x_i independently from a uniform distribution on the $[0, 1]$ interval ($x_i \sim U(0, 1)$). The value of each ingredient w_i^c is determined by the environment (E_1 or E_2).

The postsynaptic neuron outputs a prediction of the food X_t value as $y_t = g(WX_t^T)$. Throughout the paper, g will be either the identity function, in which case the prediction neuron is linear, or a step-function; however, it could be any

other nonlinearity, such as a sigmoid or ReLU. After outputting the prediction, the neuron receives feedback in the form of the real value of the input R_t . The real value is computed as $R_t = W^c X_t^T + \xi$, where $W^c = (w_1^c, \dots, w_N^c)$ is the actual value of the ingredients, and ξ is a term summarizing the noise of reward and sensing system $\xi \sim \mathcal{N}(0, \sigma)$.




Figure 1: An outline of the static agent's network. The sensor layer receives inputs representing the quantity of each ingredient of a given food at each time step. The agent computes the prediction of the food's value y_t and is then given the true value R_t ; it finally uses this information in the plasticity rule to update the weight matrix.

For the evolutionary adjustment of the agent's parameters, the loss of the static agent is the sum of the mean squared errors (MSE) between its prediction y_t and the reward R_t over the lifetime of the agent. The agent's initial weights are set to the average of the two ingredient value distributions, which is the optimal initial value for the case of symmetric switching of environments that we consider here.

Moving Agent

As a next step, we incorporate the sensory network of static agents into embodied agents that can move around in an environment scattered with food. To this end, we merge the static agent's network with a second, non-plastic motor network that is responsible for controlling the motion of the agent in the environment. Specifically, the original plastic network now provides the agent with information about the value of the nearest food. The embodied agent has additional sensors for the distance from the nearest food, the angle between the current velocity and the nearest food direction, its own velocity, and its own energy level (sum of consumed food values). These inputs are processed by two hidden layers (of 30 and 15 neurons) with tanh activation. The network's outputs are angular and linear acceleration, Fig. 2.

The embodied agents spawn in a 2D space with periodic boundary conditions along with a number of food particles that are selected such that the mean of the food value distribution is ~ 0 . An agent can eat food by approaching it sufficiently closely, and each time a food particle is eaten,




Figure 2: An outline of the network controlling the foraging agent. The sensor layer receives inputs at each time step (the ingredients of the nearest food), which are processed by the plastic layer in the same way as the static sensory network, Fig. 1. The output of that network is given as input to the motor network, along with the distance d and angle α to the nearest food, the current velocity v , and energy E of the agent. These signals are processed through two hidden layers to the final output of motor commands as the linear and angular acceleration of the agent

it is re-spawned with the same value somewhere randomly on the grid (following the setup of (Khajehabdollahi et al., 2022)). After 5000 time steps, the cumulative reward of the agent (the sum of the values of all the food it consumed) is taken as its fitness, at which point the best agents are selected by the genetic algorithm and used to initialize the next generation. The environment (food and agents' positions) is re-initialized at the start of each generation. During the evolutionary optimization, the parameters for both the motor network (connections) and plastic network (learning rule parameters) are evolved simultaneously (the genotype includes both motor weights and plasticity parameters), and so agents must learn to move and discriminate good/bad food at the same time.

Plasticity rule parametrization

Reward-modulated plasticity is one of the most promising explanations for biological credit assignment (Legenstein et al., 2008). In our network, the plasticity rule that updates the weights of the linear sensor network is a reward-modulated rule which is parameterized as a linear combination of the input, the output, and the reward at each time

step:

$$\Delta W_t = \eta_p [R_t \cdot \underbrace{(\theta_1 X_t y_t + \theta_2 y_t + \theta_3 X_t + \theta_4)}_{\text{Reward Modulated}} + \underbrace{(\theta_5 X_t y_t + \theta_6 y_t + \theta_7 X_t + \theta_8)}_{\text{Hebbian}}]. \quad (1)$$

Additionally, after each plasticity step, the weights are normalized by mean subtraction, an important step for the stabilization of Hebbian-like plasticity rules (Zenke and Gerstner, 2017).

We use a genetic algorithm to optimize the learning rate η_p and amplitudes of different terms $\theta = (\theta_1, \dots, \theta_8)$. The successful plasticity rule after many food presentations must converge to a weight vector that predicts the correct food values (or allows the agent to correctly decide whether to eat a food or avoid it).

To have comparable results, we divide $\theta = (\theta_1, \dots, \theta_8)$ by $\theta_{\max} = \max_k |\theta_k|$. So that $\theta/\theta_{\max} = \theta^{\text{norm}} \in [-1, 1]^8$. We then multiply the learning rate η_p with θ_{\max} to maintain the rule's evolved form unchanged, $\eta_p^{\text{norm}} = \eta_p \cdot \theta_{\max}$. In the following, we always use normalized η_p and θ , omitting norm .

Evolutionary Algorithm

To evolve the plasticity rule and the moving agents' motor networks, we use a simple genetic algorithm with elitism (Deb, 2011). The agents' parameters are initialized at random (drawn from a Gaussian distribution), then the sensory network is trained by the plasticity rule and finally, the agents are evaluated. After each generation, the best-performing agents (top 10 % of the population size) are selected and copied into the next generation. The remaining 90 % of the generation is repopulated with mutated copies of the best-performing agents. We mutate agents by adding independent Gaussian noise ($\sigma = 0.1$) to its parameters. Unless specified otherwise, we train a population of 100 agents for 200 generations.

Results

Environmental and reward variability control the evolved learning rates of the static agents

To start with, we consider a static agent whose goal is to identify the value of presented food correctly. The static reward-prediction network quickly evolves the parameters of the learning rule, successfully solving the prediction task. We first look at the evolved learning rate η_p , which determines how fast (if at all) the network's weight vector is updated during the lifetime of the agents. We identify three factors that control the learning rate parameter the EA converges to: the distance between the environments, the noisiness of the reward, and the rate of environmental transition.

The first natural factor is the distance d_e between the two environments, with a larger distance requiring a higher

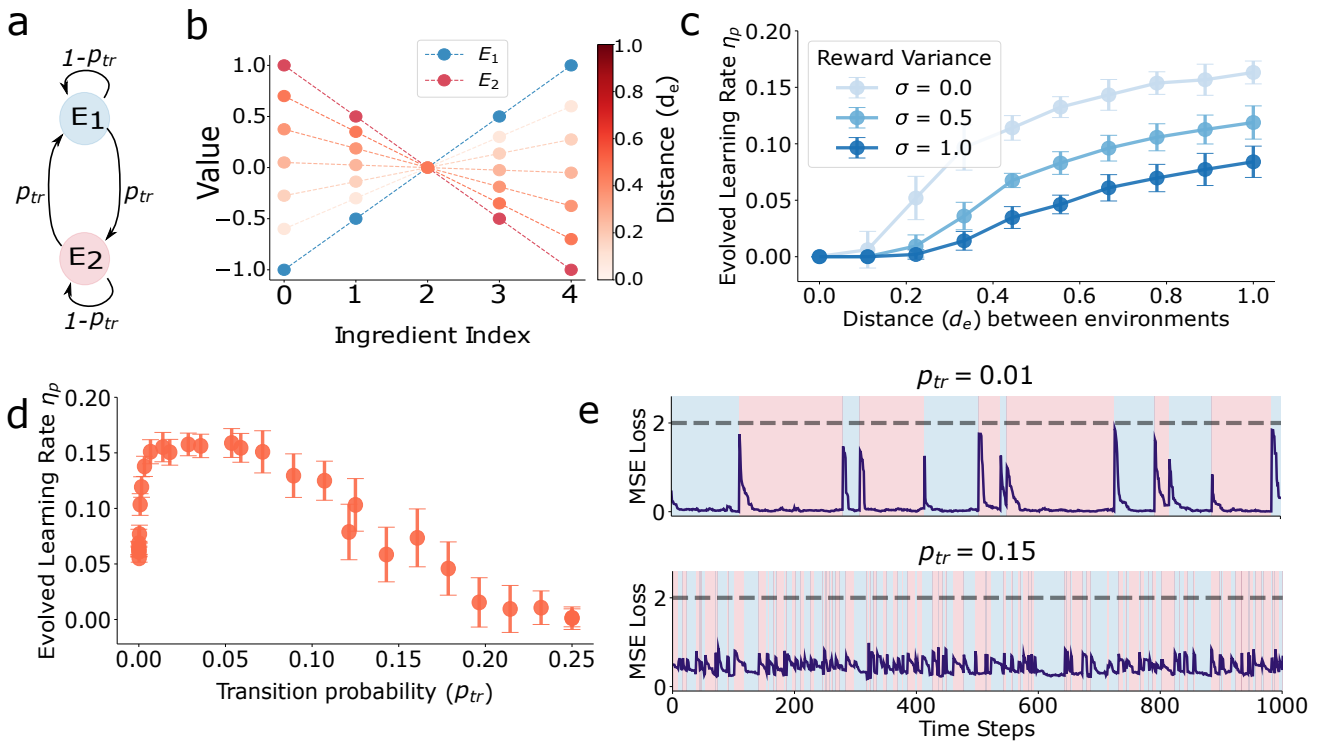


Figure 3: **a.** Schematic representation of two-states Markov model with transition probability p_{tr} between two environments E_1 and E_2 defined by the ingredient value distributions. **b.** We vary the E_2 environment by changing the ingredient values linearly $E_2 = (1 - 2d_e)E_1$, the d_e is indicated by the color. **c.** The evolved learning rate η_p grows with the distance d_e between the environments and decreases with the reward variance σ . **d.** The environment transition probability p_{tr} (here for $d_e = 1$ and $\sigma = 0.25$) has a non-monotonous relationship with the evolved learning rate η_p . Up to a certain point, more rapid transitions lead to faster learning, but too rapid environmental transition leads to a reduction of the evolved learning rate. **e.** For slow environmental transition (top), the agent fully adapts to the environment after each transition. If the transitions happen fast (bottom), the agent maintains an intermediate position between the two environments and never fully adapts to either of them.

learning rate, Fig. 3c. This is an expected result since the convergence time to the “correct” weights is highly dependent on the initial conditions. If an agent is born at a point very close to optimality, which naturally happens if the environments are similar, the distance it needs to traverse on the fitness landscape is small. Therefore it can afford to have a small learning rate, which leads to a more stable convergence and is not affected by noise.

A second parameter that impacts the learning rate is the variance of the rewards. The reward an agent receives for the plasticity step contains a noise term ξ that is drawn from a zero mean Gaussian distribution with standard deviation σ . This parameter controls the unreliability of the agent’s sensory system, i.e., higher σ means that the information the agent gets about the value of the foods it consumes cannot be fully trusted to reflect the actual value of the foods. As σ increases, the learning rate η_p decreases, which means that the more unreliable an environment becomes, the less an agent relies on plasticity to update its weights, Fig. 3c.

Indeed for some combinations of relatively small distance d_e and high reward variance σ , the EA converges to a learning rate of $\eta_p \approx 0$. This means that the agent opts to have no adaptation during its lifetime and remain at the mean of the two environments. It is an optimal solution when the expected loss due to ignoring the environmental transitions is, on average, lower than the loss the plastic network will incur by learning via the (often misleading because of the high σ) environmental cues.

A final factor that affects the learning rate the EA will converge to is the frequency of environmental change during an agent’s lifetime. Since the environmental change is modeled as a simple, two-state Markov process (Fig. 3a), the control parameter is the transition probability p_{tr} .

When keeping everything else the same, the learning rate rapidly rises as we increase the transition probability from 0, and after reaching a peak, it begins to decline slowly, eventually reaching zero (Fig. 3d). This means that when environmental transition is very rare, agents opt for a very

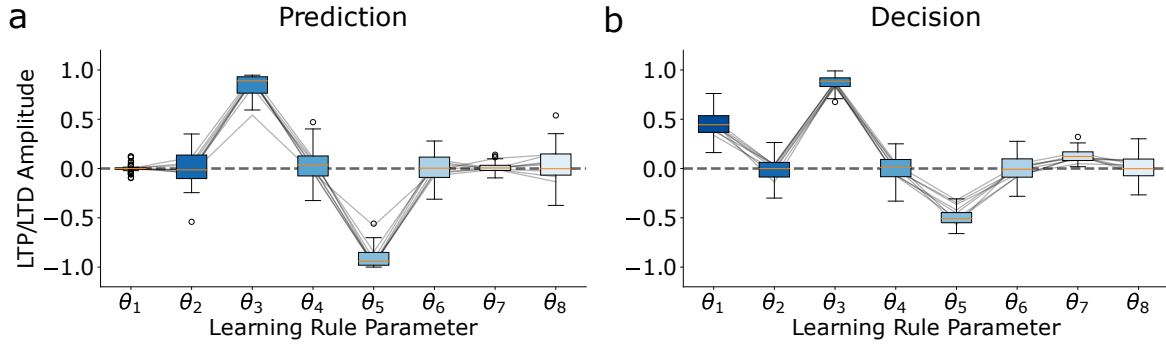


Figure 4: The evolved parameters $\theta = (\theta_1, \dots, \theta_8)$ of the plasticity rule for the reward prediction (a.) and the decision (b.) tasks, for a variety of parameters ($p_{tr} = 0.01$, $d_e \in 0, 0.1, \dots, 1$, and $\sigma \in 0, 0.1, \dots, 1$ in all 100 combinations). Despite the relatively small difference between the tasks, the evolved learning rules differ considerably. For visual guidance, the lines connect θ s from the same run.

low learning rate, allowing a slow and stable convergence to an environment-appropriate weight vector that leads to very low losses while the agent remains in that environment. As the rate of environmental transition increases, faster learning is required to speed up convergence in order to exploit the (comparatively shorter) stays in each environment. Finally, as the environmental transition becomes too fast, the agents opt for slower or even no learning, which keeps them near the middle of the two environments, ensuring that the average loss of the two environments is minimal (Fig. 3d).

The form of the evolved learning rule depends on the task: Decision vs. Prediction

The plasticity parameters $\theta = (\theta_1, \dots, \theta_8)$ for the reward-prediction task converge on approximately the same point, regardless of the environmental parameters (Fig. 4a). In particular, $\theta_3 \rightarrow 1$, $\theta_5 \rightarrow -1$, $\theta_i \rightarrow 0$ for all other i , and thus the learning rule converges to:

$$\Delta W_t = \eta_p [\theta_3 X_t R_t + \theta_5 X_t y_t] \approx \eta_p X_t (R_t - y_t). \quad (2)$$

Since by definition $y_t = g(W_t X_t^T) = W_t X_t^T (g(x) = x$ in this experiment) and $R_t = W^c X_t^T + \xi$ we get:

$$\Delta W_t = \eta_p X_t (W^c - W_t) X_t^T + \eta_p \xi X_t^T. \quad (3)$$

Thus the distribution of ΔW_t converges to a distribution with mean 0 and variance depending on η_p and σ and W converges to W^c . So this learning rule will match the agent's weight vector with the vector of ingredient values in the environment.

We examine the robustness of the learning rule the EA discovers by considering a slight modification of our task. Instead of predicting the expected food value, the agent now needs to decide whether to eat the presented food or not. This is done by introducing a step-function nonlinearity ($g(x) = 1$ if $x \geq 1$ and 0 otherwise). Then the output $y(t)$

is computed as:

$$y_t = \begin{cases} 1, & \text{if } W_t X_t^T \geq 0, \\ 0, & \text{if } W_t X_t^T < 0. \end{cases} \quad (4)$$

Instead of the MSE loss between prediction and actual value, the fitness of the agent is now defined as the sum of the food values it chose to consume (by giving $y_t = 1$). Besides these two changes, the setup of the experiments remains exactly the same.

The qualitative relation between η_p and parameters of environment d_e, σ and p_{tr} is preserved in the changed experiment. However, the resulting learning rule is significantly different (Fig. 4). The evolution converges to the following learning rule:

$$\Delta W_t = \begin{cases} \eta_p X_t [\theta_3 R_t + \theta_7], & y_t = 0, \\ \eta_p X_t [(\theta_1 + \theta_3) R_t + (\theta_5 + \theta_7)], & y_t = 1. \end{cases} \quad (5)$$

In both cases, the rule has the form $\Delta W_t = \eta_p X_t [\alpha_y R_t + \beta_y]$. Thus, the ΔW_t is positive or negative depending on whether the reward R_t is above or below a threshold ($\gamma = -\beta_y / \alpha_y$) that depends on the output decision of the network ($y_t = 0$ or 1).

Both learning rules have a clear Hebbian form (coordination of pre- and post-synaptic activity) and use the incoming reward signal as a threshold. These similarities indicate some common organizing principles of reward-modulated learning rules, but their significant differences highlight the sensitivity of the optimization process to task details.

The learning rate of embodied agents depends on environmental variability

We now turn to the moving embodied agents in the 2D environment. To optimize these agents, both the motor network's connections and the sensory network's plasticity parameters

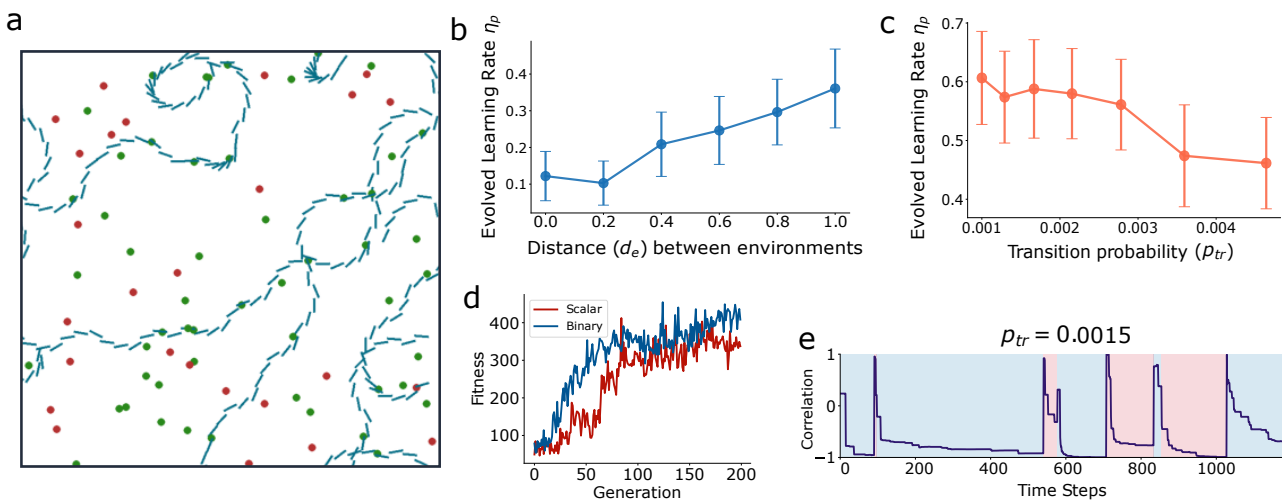


Figure 5: **a.** The trajectory of an agent (blue line) in the 2D environment. A well-trained agent will approach and consume food with positive values (green dots) and avoid negative food (red dots). **b.** The learning rate of the plastic sensory network η_p grows with the distance between environments d_e **c.** and decreases with the frequency of environmental change. **d.** The fitness of an agent (measured as the total food consumed over its lifetime) increases over generations of the EA for both the scalar and binary readouts in the sensory network. **e.** The Pearson correlation coefficient of an evolved agent’s weights with the ingredient value vector of the current environment (E_1 - blue, E_2 - red). In this example, the agent’s weights are anti-correlated with its environment, which is not an issue for performance since the motor network can interpret the inverted signs of food.

evolve simultaneously. Since the motor network is initially random and the agent has to move to find food, the number of interactions an agent experiences in its lifetime can be small, slowing down the learning. However, having the larger motor network also has benefits for evolution because it allows the output of the plastic network to be read out and transformed in different ways, resulting in a broad set of solutions.

The agents can solve the task effectively by evolving a functional motor network and a plasticity rule that converges to interpretable weights (Fig. 5a). After ~ 100 evolutionary steps (Fig. 5d), the agents can learn the ingredient value distribution using the plastic network and reliably move towards foods with positive values while avoiding the ones with negative values.

We compare the dependence of the moving and the static agents on the parameters of the environment: d_e and the state transition probability p_{tr} . At first, in order to simplify the experiment, we set the transition probability to 0, but fixed the initial weights to be the average of E_1 and E_2 , while the real state is E_2 . In this experiment, the distance between states d_e indicates twice the distance between the agent’s initial weights and the optimal weights (the environment’s ingredient values) since the agent is initialized at the mean of the two environment distributions. Same as for the static agent, the learning rate increases with the distance d_e (Fig. 5b).

Then, we examine the effect of the environmental tran-

sition probability p_{tr} on the evolved learning rate η_p . In order for an agent to get sufficient exposure to each environment, we scale down the probability p_{tr} from the equivalent experiment for the static agents. We find that as the probability of transition increases, the evolved learning rate η_p decreases (Fig. 5c). This fits with the larger trend for the static agent, although there is a clear difference when it comes to the increase for very small transition probabilities that were clearly identifiable in the static but not the moving agents. This could be due to much sparser data and possibly the insufficiently long lifetime of the moving agent (the necessity of scaling makes direct comparisons difficult). Nevertheless, overall we see that the associations observed in the static agents between environmental distance d_e and transition probability p_{tr} and the evolved learning rate η_p are largely maintained in the moving agents. Still, more data would be needed to make any conclusive assertions about the exact effect of these environmental parameters on the emerging plasticity mechanisms.

Rule redundancy in the embodied agents

A crucial difference between the static and the moving agents is the function the plasticity has to perform. While in the static agents, the plasticity has to effectively identify the exact value distribution of the environment in order to produce accurate predictions, in the embodied agents, the plasticity has to merely produce a representation of the environment that the motor network can evolve to interpret

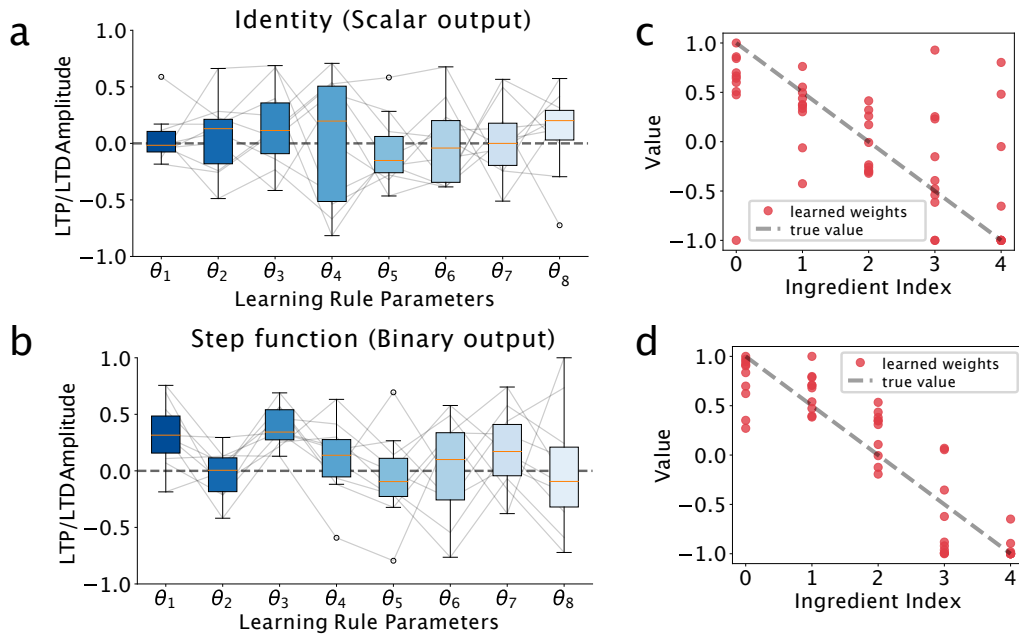


Figure 6: The evolved parameters of moving agents' plasticity rule for the $g(s) = x$, identity (a.) and the step function (Eq. 4) (b.) sensory networks (the environmental parameters here are $d_e \in [0, 1]$, $\sigma = 0$ and $p_{tr} = 0.001$). The step function (binary output) network evolved a more structured plasticity rule (e.g., $\theta_3 > 0$ for all realizations) than the linear network. Moreover, the learned weights for the identity network (c.) have higher variance and correlate significantly less with the environment's ingredient distribution compared to the learned weights for the thresholded network (d.)

adequately enough to make decisions about which food to consume.

To illustrate the difference, we plot the Pearson correlation coefficient between an agent's weights and the ingredient values of the environment it is moving in (Fig. 5e). We use the correlation instead of the MSE loss (which we used for the static agents in Fig. 3e) because the amplitude of the weight vector varies a lot for different agents and meaningful conclusions cannot be drawn from the MSE loss. For many agents, the learned weights are consistently anti-correlated with the actual ingredient values (an example of such an agent is shown in Fig. 5e). This means that the output of the sensory network will have the opposite sign from the actual food value. While in the static network, this would lead to very bad predictions and high loss, in the foraging task, these agents perform exactly as well as the ones where the weights and ingredients values are positively correlated, since the motor network can simply learn to move towards food for which it gets a negative instead of a positive sensory input.

This additional step of the output of the plastic network going through the motor network before producing any behavior has a strong effect on the plasticity rules that the embodied agents evolve. Specifically, if we look at the emerging rules the top performing agents have evolved (Fig. 6a), it becomes clear that, unlike the very well-structured rules

of the static agents (Fig. 4a), there is now virtually no discernible pattern or structure. The difference becomes even clearer if we look at the learned weights (at the end of a simulation) of the best-performing agents (Fig. 6c). While there is some correlation with the environment's ingredient value distribution, the variance is very large, and they do not seem to converge on the "correct" values in any way. This is to some extent expected since, unlike the static agents where the network's output has to be exactly correct, driving the evolution of rules that converge to the precise environmental distribution, in the embodied networks, the bulk of the processing is done by the motor network which can evolve to interpret the scalar value of the sensory network's output in a variety of ways. Thus, as long as the sensory network's plasticity rule co-evolves with the motor network, any plasticity rule that learns to produce consistent information about the value of encountered food can potentially be selected.

To further test this assumption, we introduce a bottleneck of information propagation between the sensory and motor networks by using a step-function nonlinearity on the output of the sensory network (Eq. 4). Similarly to the decision task of the static network, the output of the sensory network now becomes binary. This effectively reduces the flow of information from the sensory to the motor network, forcing the sensory network to consistently decide whether food should be consumed (with the caveat that the motor network can

still interpret the binary sign in either of two ways, either consuming food marked with 1 or the ones marked with 0 by the sensory network). The agents perform equally well in this variation of the task as before (Fig. 5d), but now, the evolved plasticity rules seem to be more structured (Fig. 6b). Moreover, the variance of the learned weights in the best-performing agents is significantly reduced (Fig. 6d), which indicates that the bottleneck in the sensory network is increasing selection pressure for rules that learn the environment's food distribution accurately.

Discussion

We find that different sources of variability have a strong impact on the extent to which evolving agents will develop neuronal plasticity mechanisms for adapting to their environment. A diverse environment, a reliable sensory system, and a rate of environmental change that is neither too large nor too small are necessary conditions for an agent to be able to effectively adapt via synaptic plasticity. Additionally, we find that minor variations of the task an agent has to solve or the parametrization of the network can give rise to significantly different plasticity rules.

Our results partially extend to embodied artificial agents performing a foraging task. We show that environmental variability also pushes the development of plasticity in such agents. Still, in contrast to the static agents, we find that the interaction of a static motor network with a plastic sensory network gives rise to a much greater variety of well-functioning learning rules. We propose a potential cause of this degeneracy; as the relatively complex motor network is allowed to read out and process the outputs from the plastic network, any consistent information coming out of these outputs can be potentially interpreted in a behaviorally useful way. Reducing the information the motor network can extract from the sensory system significantly limits learning rule variability.

Our findings on the effect of environmental variability concur with the findings of previous studies (Lange and Sprekeler, 2020) that have identified the constraints that environmental variability places on the evolutionary viability of learning behaviors. We extend these findings in a mechanistic model which uses a biologically plausible learning mechanism (synaptic plasticity). We show how a simple evolutionary algorithm can optimize the different parameters of a simple reward-modulated plasticity rule for solving simple prediction and decision tasks. Reward-modulated plasticity has been extensively studied as a plausible mechanism for credit assignment in the brain (Florian, 2007; Baras and Meir, 2007; Legenstein et al., 2008) and has found several applications in artificial intelligence and robotics tasks (Burms et al., 2015; Bing et al., 2019). Here, we demonstrate how such rules can be very well-tuned to take into account different environmental parameters and produce optimal behavior in simple systems.

Additionally, we demonstrate how the co-evolution of plasticity and static functional connectivity in different sub-networks fundamentally changes the evolutionary pressures on the resulting plasticity rules, allowing for greater diversity in the form of the learning rule and the resulting learned connectivity. Several studies have demonstrated how, in biological networks, synaptic plasticity heavily interacts with (Butz et al., 2014; Stampanoni Bassi et al., 2019; Bernáez Timón et al., 2022) and is driven by network topology (Giannakakis et al., 2023). Moreover, it has been recently demonstrated that biological plasticity mechanisms are highly redundant in the sense that any observed neural connectivity or recorded activity can be achieved with a variety of distinct, unrelated learning rules (Ramesh, 2023). This observed redundancy of learning rules in biological settings complements our results and suggests that the function of plasticity rules cannot be studied independently of the connectivity and topology of the networks they are acting on.

The optimization of functional plasticity in neural networks is a promising research direction both as a means to understand biological learning processes and as a tool for building more autonomous artificial systems. Our results suggest that reward-modulated plasticity is highly adaptable to different environments and can be incorporated into larger systems that solve complex tasks.

Future work

This work studies a simplified toy model of neural network learning in stochastic environments. Future work could be built on this basic framework to examine more complex reward distributions and sources of environmental variability. Moreover, a greater degree of biological realism could be added by studying more plausible network architectures (possibly derived from connectomics data) and more sophisticated plasticity rule parametrizations.

Additionally, our foraging simulations were constrained by limited computational resources and were far from exhaustive. Further experiments can investigate environments with different constraints, food distributions, and multiple seasons as well as the inclusion of plasticity on the motor parts of the artificial organisms.

Acknowledgements

This work was supported by a Sofja Kovalevskaja Award from the Alexander von Humboldt Foundation. EG and SK thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for their support. We acknowledge the support from the BMBF through the Tübingen AI Center (FKZ: 01IS18039A). AL is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 39072764.

References

- Baras, D. and Meir, R. (2007). Reinforcement Learning, Spike-Time-Dependent Plasticity, and the BCM Rule. *Neural Computation*, 19(8):2245–2279.
- Bernáez Timón, L., Ekelmans, P., Konrad, S., Nold, A., and Tchumachenko, T. (2022). Synaptic plasticity controls the emergence of population-wide invariant representations in balanced network models. *Phys. Rev. Res.*, 4:013162.
- Biesialska, M., Biesialska, K., and Costa-jussà, M. R. (2020). Continual lifelong learning in natural language processing: A survey. In *Proceedings of the 28th International Conference on Computational Linguistics*. International Committee on Computational Linguistics.
- Bing, Z., Baumann, I., Jiang, Z., Huang, K., Cai, C., and Knoll, A. (2019). Supervised learning in snn via reward-modulated spike-timing-dependent plasticity for a target reaching vehicle. *Frontiers in Neurorobotics*, 13.
- Burms, J., Caluwaerts, K., and Dambre, J. (2015). Reward-modulated hebbian plasticity as leverage for partially embodied control in compliant robotics. *Frontiers in Neurorobotics*, 9.
- Butz, M., Steenbuck, I., and van Ooyen, A. (2014). Homeostatic structural plasticity increases the efficiency of small-world networks. *Frontiers in Synaptic Neuroscience*, 6.
- Caroni, P., Donato, F., and Muller, D. (2012). Structural plasticity upon learning: regulation and functions. *Nature reviews. Neuroscience*, 13(7):478–490.
- Citri, A. and Malenka, R. (2008). Synaptic plasticity: Multiple forms, functions, and mechanisms. *Neuropsychopharmacology : official publication of the American College of Neuropsychopharmacology*, 33:18–41.
- Confavreux, B., Zenke, F., Agnes, E., Lillicrap, T., and Vogels, T. (2020). A meta-learning approach to (re)discover plasticity rules that carve a desired function into a neural network. In Larochele, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 16398–16408. Curran Associates, Inc.
- Deb, K. (2011). *Multi-objective Optimisation Using Evolutionary Algorithms: An Introduction*, pages 3–34. Springer London, London.
- Dunlap, A. S. and Stephens, D. W. (2009). Components of change in the evolution of learning and unlearned preference. *Proceedings of the Royal Society B: Biological Sciences*, 276(1670):3201–3208.
- Dunlap, A. S. and Stephens, D. W. (2016). Reliability, uncertainty, and costs in the evolution of animal learning. *Current Opinion in Behavioral Sciences*, 12:73–79. Behavioral ecology.
- Ellefsen, K. O. (2014). The evolution of learning under environmental variability. pages 649–656.
- Eskridge, B. E. and Hougen, D. F. (2012). Nurturing promotes the evolution of learning in uncertain environments. In *2012 IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL)*, pages 1–6.
- Fawcett, T. W., Hamblin, S., and Giraldeau, L.-A. (2012). Exposing the behavioral gambit: the evolution of learning and decision rules. *Behavioral Ecology*, 24(1):2–11.
- Feldman, D. E. (2009). Synaptic mechanisms for plasticity in neocortex. *Annual Review of Neuroscience*, 32(1):33–55. PMID: 19400721.
- Florian, R. V. (2007). Reinforcement Learning Through Modulation of Spike-Timing-Dependent Synaptic Plasticity. *Neural Computation*, 19(6):1468–1502.
- Giannakakis, E., Vinogradov, O., Buendia, V., and Levina, A. (2023). Recurrent connectivity structure controls the emergence of co-tuned excitation and inhibition. *bioRxiv*.
- Guttenberg, N. (2019). Evolutionary rates of information gain and decay in fluctuating environments. ALIFE 2019: The 2019 Conference on Artificial Life:365–371.
- Jordan, J., Schmidt, M., Senn, W., and Petrovici, M. A. (2021). Evolving interpretable plasticity for spiking networks. *eLife*, 10:e66273.
- Kerr, B. and Feldman, M. (2003). Carving the cognitive niche: Optimal learning strategies in homogeneous and heterogeneous environments. *Journal of Theoretical Biology*, 220(2):169–188.
- Khajehabdollahi, S., Prosi, J., Giannakakis, E., Martius, G., and Levina, A. (2022). When to Be Critical? Performance and Evolvability in Different Regimes of Neural Ising Agents. *Artificial Life*, 28(4):458–478.
- Lange, R. T. and Sprekeler, H. (2020). Learning not to learn: Nature versus nurture in silico.
- Lee, C. and Lee, A. (2020). Clinical applications of continual learning machine learning. *The Lancet Digital Health*, 2:e279–e281.
- Legenstein, R., Pecevski, D., and Maass, W. (2008). A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback. *PLOS Computational Biology*, 4(10):1–27.
- Magee, J. C. and Grienberger, C. (2020). Synaptic plasticity forms and functions. *Annual Review of Neuroscience*, 43(1):95–117. PMID: 32075520.
- Najarro, E. and Risi, S. (2020). Meta-learning through hebbian plasticity in random networks. In Larochele, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20719–20731. Curran Associates, Inc.
- Nolfi, S. and Parisi, D. (1996). Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior*, 5(1):75–98.
- Open Ended Learning Team, Stooke, A., Mahajan, A., Barros, C., Deck, C., Bauer, J., Sygnowski, J., Trebacz, M., Jaderberg, M., Mathieu, M., McAleese, N., Bradley-Schmieg, N., Wong, N., Porcel, N., Raileanu, R., Hughes-Fitt, S., Dalibard, V., and Czarnecki, W. M. (2021). Open-ended learning leads to generally capable agents.
- Papini, M. R. (2012). *Evolution of Learning*, pages 1188–1192. Springer US, Boston, MA.

- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., and Wermter, S. (2019). Continual lifelong learning with neural networks: A review.
- Pedersen, J. W. and Risi, S. (2021). Evolving and merging hebbian learning rules: Increasing generalization by decreasing the number of rules. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '21*, page 892–900, New York, NY, USA. Association for Computing Machinery.
- Ramesh, P. (2023). *GANs schön kompliziert: Applications of Generative Adversarial Networks*. PhD thesis, University of Tübingen.
- Schmidhuber, J. (1987). Evolutionary principles in self-referential learning. on learning now to learn: The meta-meta-meta...-hook. Diploma thesis, Technische Universität München, Germany.
- Snell-Rood, E. C. (2013). An overview of the evolutionary causes and consequences of behavioural plasticity. *Animal Behaviour*, 85(5):1004–1011. Special Issue: Behavioural Plasticity and Evolution.
- Snell-Rood, E. C. and Steck, M. K. (2019). Behaviour shapes environmental variation and selection on learning and plasticity: review of mechanisms and implications. *Animal Behaviour*, 147:147–156.
- Stampanoni Bassi, M., Iezzi, E., Gilio, L., Centonze, D., and Buttari, F. (2019). Synaptic plasticity shapes brain connectivity: Implications for network topology. *International Journal of Molecular Sciences*, 20(24).
- Thangarasa, V., Miconi, T., and Taylor, G. W. (2020). Enabling continual learning with differentiable hebbian plasticity.
- Thornton, A. and Clutton-Brock, T. (2011). Social learning and the development of individual and group behaviour in mammal societies. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 366:978–87.
- Yaman, A., Iacca, G., Mocanu, D. C., Coler, M., Fletcher, G., and Pechenizkiy, M. (2021). Evolving Plasticity for Autonomous Learning under Changing Environmental Conditions. *Evolutionary Computation*, 29(3):391–414.
- Zenke, F. and Gerstner, W. (2017). Hebbian plasticity requires compensatory processes on multiple timescales. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1715):20160259.

Appendix B Network bottlenecks and task structure control the evolution of interpretable learning rules in a foraging agent

Contains the full text of the publication: Giannakakis, E., Khajehabdollahi, S., and Levina, A. (2024). Network Bottlenecks and Task Structure Control the Evolution of Interpretable Learning Rules in a Foraging Agent. *Artificial Life*, pages 1–18.

Network bottlenecks and task structure control the evolution of interpretable learning rules in a foraging agent

Emmanouil Giannakakis,^{1,2} Sina Khajehabdollahi,¹ and Anna Levina^{1,2}

¹*Department of Computer Sciences, University of Tübingen, Germany*

²*Max Planck Institute for Biological Cybernetics, Tübingen, Germany*

Developing reliable mechanisms for continuous local learning is a central challenge faced by biological and artificial systems. Yet, how the environmental factors and structural constraints on the learning network influence the optimal plasticity mechanisms remains obscure even for simple settings. To elucidate these dependencies, we study meta-learning via evolutionary optimization of simple reward-modulated plasticity rules in embodied agents solving a foraging task. We show that unconstrained meta-learning leads to the emergence of diverse plasticity rules. However, regularization and bottlenecks to the model help reduce this variability, resulting in interpretable rules. Our findings indicate that the meta-learning of plasticity rules is very sensitive to various parameters, with this sensitivity possibly reflected in the learning rules found in biological networks. When included in models, these dependencies can be used to discover potential objective functions and details of biological learning via comparisons with experimental observations.

I. INTRODUCTION

A hallmark of living organisms is their ability to adapt to their environment and assimilate new information to modify their behaviour. It is unclear how the ability to learn first evolved [1], yet their advantages are clear. Natural environments are too complex for all the necessary information to be encoded genetically [2] and more importantly, they keep changing during an organism's lifetime in ways that cannot be anticipated [3, 4]. The capacity to learn is so beneficial in evolutionary terms that most organisms accept high costs, such as increased energy consumption [5] and the need for lengthy nurturing periods [6] in order to maintain an ability to learn throughout their lifetimes.

The most prominent mechanism by which learning occurs in biological organisms is via changes to the strengths of synaptic connections between neurons [7–9]. Synaptic plasticity is observed across a wide variety of organisms and remains active throughout most organisms' lifetimes [10]. Starting from the early days of computational neuroscience [11, 12], many experimental and theoretical studies have focused on identifying the basic principles of synaptic modification and developing simplified models of these processes [13–16].

Using an automated discovery of optimal plasticity rules for solving particular tasks, *meta-learning of plasticity*, can be successfully applied for training artificial neural networks. There the advantage of plasticity lies in its time-continuous nature (compatible with life-long learning) and its local nature [17–19]. Often this meta-learning is implemented using evolutionary computations [20–22]. Several applications following this approach can compete with state-of-the-art machine learning algorithms on various complex tasks [23–25].

The meta-learning of synaptic plasticity rules in artificial networks is a promising approach for investigating the functions of synaptic plasticity in biological networks. Experimental measurements of synaptic plasticity can

correlate the impact that pre- and post-synaptic activity have on the synaptic strength but can rarely identify the objective function that the learning process is optimising for or the exact impact different plasticity parameters have on the learning task [26]. Optimizing learning rules in artificial networks can help uncover the functions of different plasticity forms in biological networks [27–32].

Still, most of the meta-learning studies focus on the inference of the optimal rules for minimising specific objective functions while ignoring other biophysical constraints that may influence the learning process. In particular, while it has been shown that network structure [33] can exert a strong effect on the performance of synaptic plasticity and that multiple compensatory processes heavily interact with classical plasticity mechanisms in biological systems [34], such factors are rarely considered in meta-learning studies of plasticity. Additionally, a recent study [35] has demonstrated that meta-learning can uncover diverse families of plasticity rules leading to identical learning outcomes. While this finding encourages a more holistic understanding of plasticity that focuses on shared learning outcomes, rather than individual rules, it nevertheless raises the need for methods to classify and assess learning rules in different settings.

Here, extending previous work [36], we model embodied agents equipped with plastic sensory networks that learn about the agents environment during a foraging task. We study the effect that different structural features of the network and parameters of the evolutionary process have on the form of evolved functional reward-modulated plasticity rules. We find that the apparent redundancy of evolved learning rules can be significantly reduced by the introduction of simple regularization techniques and information bottlenecks. We further find that small changes in the plastic network's structure (such as changes in the activation function of a neuron or the method of weight normalization) can have a significant effect on the form of the meta-learned plasticity rules. Finally, we separate the plastic sensory networks from the

moving agents and we attempt to re-create the plasticity rules that co-evolved with the agents motor networks. We find that different assumptions about the function of the motor networks as well as the function that the plasticity performs in the embodied system can lead to the emergence of distinct rules in the static sensory networks. We argue that such effects can enable us to infer important facts about the objective functions and learning processes of biological learning systems.

II. METHODS

Our model consists of an environment with food particles of different values and agents that can sense the nearest food to them and decide what to eat. To this end, the agent must move to the target food using its motors and the motor network that controls them. The sensory system of the agents can learn using synaptic plasticity, which is encoded genetically and evolves together with the agent’s motor network. In the following, we describe the parts of the model in more detail.

A. Environment

A number of food particles (usually 100, unless otherwise specified) are randomly placed in a 300×300 , 2D space with periodic boundary conditions. Each food particle consists of various amounts of N ingredients with positive (food) or negative (poison) values. The value of a food particle is a weighted sum of its ingredients. To predict the reward value of a given resource, the agent must learn the values of these ingredients by interacting with the environment.

The values of the N ingredients $W^c = \{W_1^c, \dots, W_N^c\}$, with $W_i^c = -1 + \frac{2(i-1)}{N-1}$, $i = 1, \dots, N$ are equally spaced between -1 and 1. Each food particle is encoded as an N -dimensional vector $X_t = (x_1, \dots, x_N)$ is presented, where the value x_i , $i \in \{1, \dots, N\}$ represents the quantity of the ingredient i . The value of the food particle is $W^c X^T = \sum_i W_i^c \cdot x_i$. We draw x_i independently form a uniform distribution on the $[0, 1]$ interval ($x_i \sim U(0, 1)$). Each time a food particle is eaten, it is re-spawned with the same value somewhere randomly on the grid (following the setup of [37]). This re-spawning algorithm guarantees that even if the agents have already evolved to eat only non-poisonous food, the mean of the food value distribution in the environment remains ≈ 0 .

B. Agent

The agent (Fig. 1a) consists of a motor network, which controls their movement across the 2D space and a plastic sensory network that learns the values of different ingredients via synaptic plasticity and provides the agent

with information about the value (food or poison) of the nearest food (Fig. 1b).

1. Motor network

An agent’s motor network is not plastic but trained via evolution across generations and remains constant throughout an agent’s lifetime. The motor network consists of 4 layers. An input layer with sensors that provide the agents with information such as the distance from the nearest food, the angle between the current velocity and the nearest food direction, its own velocity, and its own energy level (sum of consumed food values). Finally, one of the sensors is connected to the plastic sensory network, which provides the agent with an estimation of the nearest food’s value. Two hidden layers (of 30 and 15 neurons with tanh activation) process these inputs. The network’s outputs are the angular and linear acceleration of the agent.

2. Sensory network

The sensory network consists of N sensory neurons (one for each type of the ingredients, loosely resembling the olfactory system of *drosophila melanogaster* [38, 39]) that are projecting to a single post-synaptic neuron signalling the food value to the motor network. At each time step, an input $X_t = (x_1, \dots, x_N)$, representing the nearest food to the agents, is given. The postsynaptic neuron outputs an assessment of the food X_t value as $y_t = g(W_t X_t^T)$. Throughout the paper, g will be either the identity function, in which case the prediction neuron is linear, or a step-function; however, it could be any other nonlinearity, such as a sigmoid or ReLU. When a food particle is consumed, the agent receives the real value of the input $R_t = W^c X_t^T$ as feedback for learning. Once an agent consumes a food particle, the input X_t , output y_t and feedback R_t are used by the plasticity rule to update the sensory weights $W_{t+1} \leftarrow W_t + \Delta W_t$ where $\Delta W_t = F(X_t, y_t, R_t)$.

C. Plasticity rule parametrization

Reward-modulated plasticity is one of the most promising explanations for biological credit assignment [40]. In our network, the plasticity rule that updates the weights of the sensor network is a reward-modulated rule which is parameterized as a linear combination of the input, the output, and the reward at each time step:

$$\Delta W_t = F(X_t, y_t, R_t) = \eta_p [R_t \cdot \underbrace{(\theta_1 X_t y_t + \theta_2 y_t + \theta_3 X_t + \theta_4)}_{\text{Reward Modulated}} + \underbrace{(\theta_5 X_t y_t + \theta_6 y_t + \theta_7 X_t + \theta_8)}_{\text{Hebbian}}]. \quad (1)$$

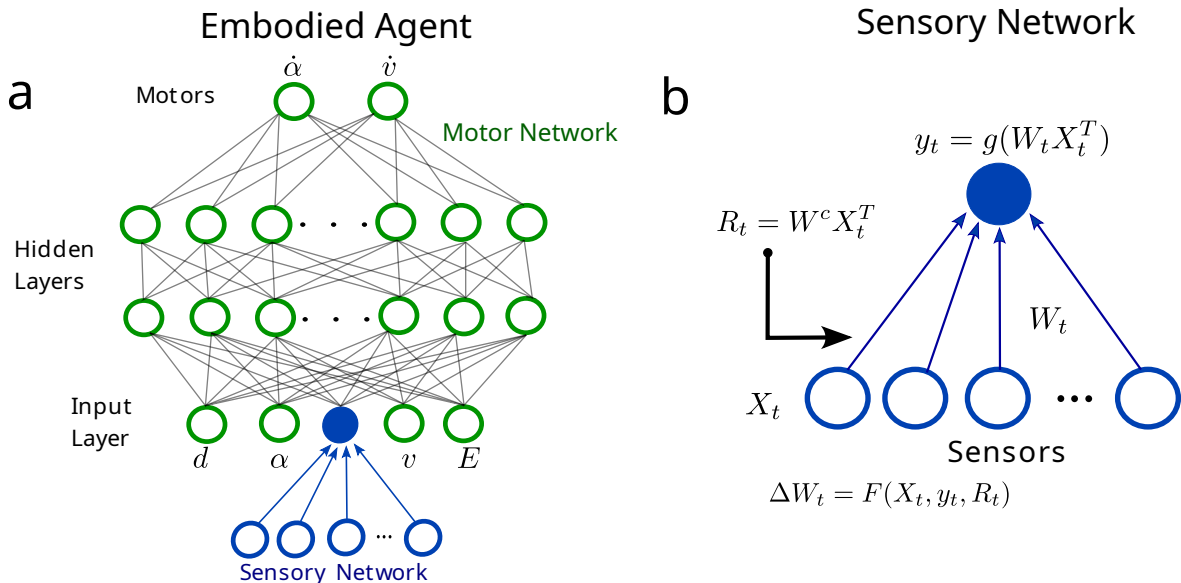


FIG. 1. Structure of the neural network controlling the embodied foraging agent. **a.** A diagram of the network controlling the foraging agent. The sensor layer receives inputs at each time step (the ingredients of the nearest food). The output of that network is given as input to the motor network, along with the distance d and angle α to the nearest food, the current velocity v , and energy E of the agent. These signals are processed through two hidden layers to the final output of motor commands as the linear and angular acceleration of the agent **b.** Details of the sensory network. The sensor layer receives inputs representing the quantity of each ingredient of the nearest food at each time step. The agent outputs its assessment of the food's value y_t , and when a food particle is consumed it receives the true value R_t as feedback; it finally uses this information to update the weight matrix according to the plasticity rule.

Additionally, after each plasticity step, the weights are normalized to maintain a constant sum, an important step for the stabilization of Hebbian-like plasticity rules [34]. Unless specified otherwise, the normalization happens via mean subtraction.

We use a genetic algorithm to optimize the learning rate η_p and amplitudes of different terms $\theta = (\theta_1, \dots, \theta_8)$. The successful plasticity rule after many food presentations must converge to a weight vector that predicts the correct food values (or allows the agent to correctly decide whether to eat a food or avoid it).

To have comparable results, we divide $\theta = (\theta_1, \dots, \theta_8)$ by $\theta_{\max} = \max_k |\theta_k|$. So that $\theta/\theta_{\max} = \theta^{\text{norm}} \in [-1, 1]^8$. We then multiply the learning rate η_p with θ_{\max} to maintain the rule's evolved form unchanged, $\eta_p^{\text{norm}} = \eta_p \cdot \theta_{\max}$. In the following, we always use normalized η_p and θ , omitting norm .

D. Evolutionary algorithm

To evolve the plasticity rule and the moving agents' motor networks, we use a simple genetic algorithm with elitism [41]. Each agent is initialized in a separate instance of the 2D environment. The agents' evolvable parameters (motor weights and plasticity parameters) are

initialized at random (drawn from a Gaussian distribution), then the sensory network is trained by the plasticity rule and after 5000 time steps the agent is evaluated. The energy of an agent at the beginning of its lifetime is initialized to be positive ($E_0 = 2$) but if the agent's energy becomes negative during its lifetime (due to eating too many negative food particles), the agent can no longer move. The fitness of an agent is its energy at the end of its lifetime, given as the sum of the values of the foods it consumed:

$$E = \sum_{t \in T_{\text{eat}}} R_t \quad (2)$$

where T_{eat} are the timesteps where the agents consumed a food particle. After each generation, the best-performing agents (top 10 % of the population size) are selected and copied into the next generation. The remaining 90 % of the generation is repopulated with mutated copies of the best-performing agents. We mutate agents by adding independent Gaussian noise ($\sigma = 0.1$) to its parameters. Unless specified otherwise, we train a population of 100 agents for 250 generations.

E. Curriculum

The task of evolving the motor network weights and plasticity rule from scratch is very hard. To assist the evolutionary process of the moving agents, we developed a curriculum that progressively initializes the agents' sensory networks farther from the ground truth, making correct plasticity rule more and more essential to distinguish eatable from poisonous food particles rapidly. To do so, we add a normally distributed noise ($\mathcal{N}(0, \sigma)$) to the ground truth values with the progressively increasing variance per curriculum $\sigma \in \{0, 1, 2, 3, 4, 5\}$. The populations train within each curriculum for 250 generations before moving to the next curriculum step. All results presented are from populations trained until the final curriculum step.

III. SUMMARY OF PREVIOUS RESULTS

In previous work [36], we studied the effect of different factors (environmental fluctuations, reliability of sensory input, and task complexity) on the evolved learning rate and form of reward-modulated plasticity rules. In particular, we investigated the speed at which agents need to learn about their environment (i.e., learning rate). Here, we extend that study by focusing on the impact of network and task parameters on the resulting plasticity rules.

IV. RESULTS

A. The agents evolve to solve the foraging task

We begin by evolving 20 populations of 100 agents with a linear sensory readout ($g(x) = x$). At each step of the curriculum, the agents evolve a functional motor network and a plasticity rule for the sensory network. The converged sensory weights allow the agent's motor network to adequately interpret the quality of food, enabling it to effectively approach food particles with positive values while avoiding the ones with negative values (Fig. 2a).

We see that after ~ 100 evolutionary steps (Fig. 2b), the agents reach a stable high fitness (> 250). To put this value into context, we train 5 populations of "Oracle" agents without a plastic sensory network, who are given the real values of nearby foods as direct input. These populations of agents reach a maximum fitness of 383 ± 37 at the end of their training. Thus, the plastic agents' fitness is about as good as they could possibly get. This effectively means that the best-performing agents can accurately navigate and rapidly learn the correct ingredient values via their plastic sensory networks and subsequently only consume foods with positive values.

The weights of the best-performing agent's sensory networks subject to synaptic plasticity converge to the val-

ues strongly correlated (Pearson correlation coefficient of 0.92 ± 0.09) but not identical to the correct ingredient values (Fig. 2c). Thus, the evolutionary algorithm produces learning rules that can reliably learn a rough approximation of the real distribution of ingredient values, enabling the motor network to decide which food particles the agent should consume.

B. Redundancy in the plasticity rules and its diminution by the information bottleneck

Once the evolution has converged, we look at the evolved plasticity parameters $\theta = (\theta_1, \dots, \theta_8)$ of the best-performing agents of each population (Fig. 2d). We observe that despite the fact that all the rules seem to be converging to similar sensory weights (Fig. 2c), the plasticity parameters are quite different between different rules and effectively present no visible pattern. Such redundancy of plasticity rules has been observed in previous studies [35] of meta-learning.

We hypothesize that the redundancy of learning rules is due to the superfluous information the sensory network passes on to the motor network. For good task performance, the motor network only needs to know whether to approach or avoid a food particle (depending on whether its value is positive or negative). Since the output of the sensory network is a scalar value, any learning rule which produces output that allows the network to distinguish between positive and negative foods will perform equally well.

To test this hypothesis, we introduce a bottleneck of information propagation between the sensory and motor networks by using a step-function nonlinearity on the output of the sensory network:

$$g(x) = \begin{cases} 1, & \text{if } x \geq 0, \\ -1, & \text{if } x < 0. \end{cases} \quad (3)$$

After convergence, we observe that despite the variability of the resulting learning rules being high, a general pattern begins to emerge in the form of the rules. Specifically, the third parameter θ_3 , is consistently positive and close to 1 (Fig. 2e). This means that the term $R_t \cdot X_t$, which correlates the presence or absence of ingredients with positive/negative rewards becomes prominent in the rule, indicating some form of Hebbian-like learning.

C. The information bottleneck improves performance and generalizability

We compare the performance of scalar and binary sensory readout agents in multiple realisations of the 2D environment. In both cases, we observe a high level of variability in the resulting fitnesses. This variability is in part due to the fact that depending on the sequence of food particles an agent encounters early on during its

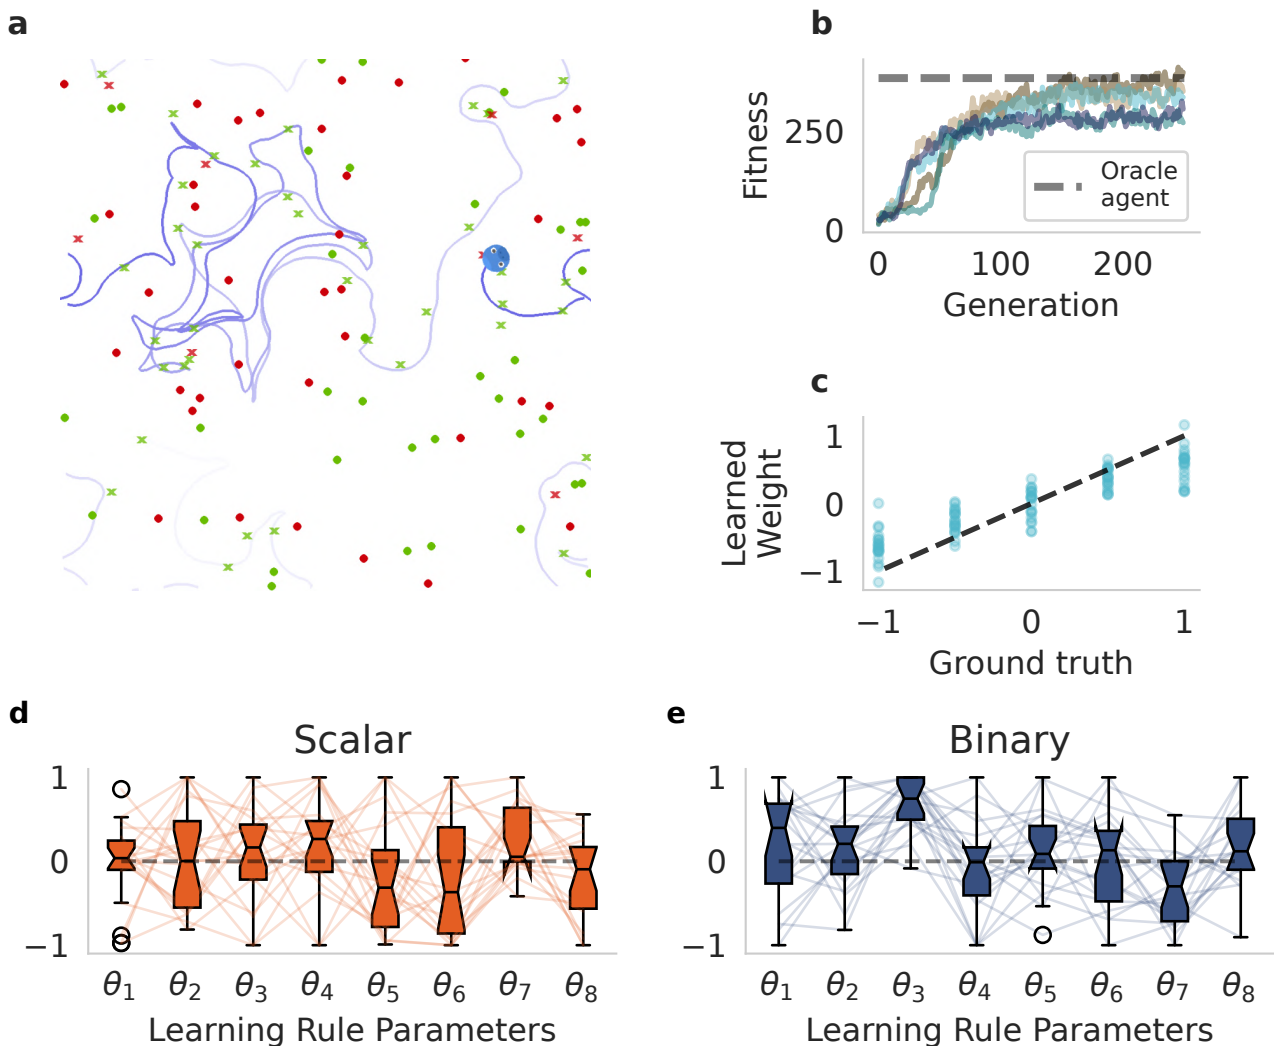


FIG. 2. Diverse learning rules lead to high performance in the foraging task **a**. The trajectory of an agent (grey line) in the 2D environment. A well-trained agent will approach and consume food with positive values (green dots) and avoid negative food (red dots). The \times signs denote the locations of food particles consumed by the agent **b**. The fitness of the best performing agent from 5 evolving populations (Eq. 2) increases over generations of the evolutionary algorithm. The dotted gray line indicates the maximum fitness of the "Oracle" non-plastic agents who are given the correct food values as input **c**. The learned weights of the sensory network (blue dots) correlate strongly with the actual ingredient values (Pearson cc of 0.92 ± 0.09). **d**, **e**. The evolved plasticity rules across 20 runs for agents with scalar/binary readout, respectively.

lifetime, an agent can either learn the ingredient value distribution, which will lead to a relatively high fitness or encounter too many negative foods and remain stagnant during its lifetime (due to the energy constraint), which leads to a fitness of ≈ 0 (Fig. 3b, d). This risk of an agent encountering too many negative foods at the beginning of its lifetime means that even very well-performing agents will fail for a small fraction of environment initializations. Nevertheless, despite the high variability, we observe a significantly higher average fitness among the agents with binary (Fig. 3d, average fitness $\langle E \rangle = 287$), compared to the ones with scalar sensory readouts (Fig. 3b, average

fitness $\langle E \rangle = 137$), suggesting that the information bottleneck helps with task performance.

We then examine how well the evolved rules can generalize between different motor networks. We swap the sensory and motor networks of different evolved agents and test the fitness of the resulting mixed agents (Fig. 3a). We observe a statistically significant reduction ($p_{\text{value}} = 0.00087$) in the average fitness (Fig. 3c) of agents recombined in this way compared to agents with co-evolved sensory and motor networks in the case of scalar sensory readouts (here and later we use paired t-test comparing the mean fitness across 100 environment realizations for

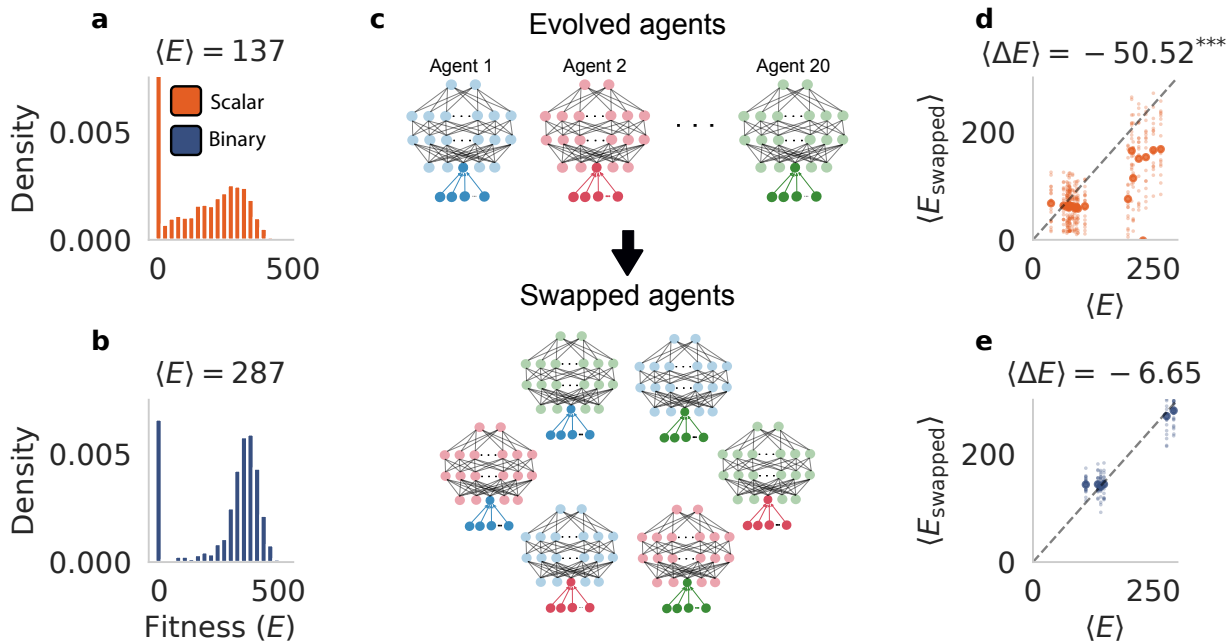


FIG. 3. Agents with binary sensory readouts perform and generalize better than those with scalar readouts. **a, b.** The histogram of fitnesses of the top-ranking agent in each of the 20 runs over 100 independent environment realizations for scalar/binary networks times, respectively. **c.** Schematic of swapping networks: Motor and sensor networks are swapped between the fittest agents from each run. **d, e.** The swapped networks' fitnesses plotted against the mean fitness of the original configurations for the scalar/binary networks, respectively. The mean fitnesses for the swapped agents were significantly different (paired two-sided t -test) for the networks with a scalar readout ($p_{\text{value}} = 0.00087$) but not for the ones with a binary readout ($p_{\text{value}} = 0.1457$).

each tested agent against the mean fitness of the swapped agents with the same motor network - mean across 100 environments for each swapped agenda and 19 different swapped versions - that is approximately Gaussian). This suggests that the motor network evolves to interpret the specific output of a given learning rule, and thus, the differences between the converged sensory weights of different rules (Fig. 2c) make a significant difference in the resulting sensory readout. On the contrary, in the case of binary readouts (Fig. 3f), the average fitness remains the same, i.e. agents perform equally well with learning rules that co-evolved with the motor networks of different agents ($p_{\text{value}} = 0.1457$). Thus, adding the information bottleneck in the sensory network, not only improves performance, but it makes the agents much more generalizable and enables modular evolution

D. Regularization of the plasticity parameters leads to interpretable rules

We now test how the regularization of the plasticity rules decreases the outcome's variability. To this end, we regularize the plasticity parameters $\theta = \{\theta_1, \dots, \theta_8\}$ using an L^1 norm, implemented as a weight decay of these

parameters after each generation: $\theta_i \leftarrow \theta_i - \kappa \cdot \text{sign}(\theta_i)$, where κ is a weight-decay velocity taken in further experiments to be $\kappa = 0.05$. We evolve agents with the scalar and binary versions of the sensory network, and we observe that evolved plasticity rules show minimal rule patterns, different between binary and scalar networks (Fig. 4c, e).

In the case of agents with a scalar sensory readout, the plasticity parameters $\theta = (\theta_1, \dots, \theta_8)$ converge on approximately the same point for most of the evolved agents (Fig. 4c). In particular, $\theta_3 \rightarrow 1$, $\theta_5 \rightarrow -1$, $\theta_i \rightarrow 0$ for all other i , and thus the learning rule converges to:

$$\Delta W_t = \eta_p [\theta_3 X_t R_t + \theta_5 X_t y_t] \approx \eta_p X_t (R_t - y_t). \quad (4)$$

Since by definition $g(x) = x$ in this experiment, thus $y_t = g(W_t X_t^T) = W_t X_t^T$ and $R_t = W^c X_t^T$. For the weight update we get:

$$\Delta W_t = \eta_p X_t (W^c - W_t) X_t^T \quad (5)$$

Thus, $\Delta W_t \rightarrow 0$ as $W \rightarrow W^c$, and consequently, this learning rule will match the agent's weight vector with the vector of ingredient values in the environment. We further observe a slight increase in the scalar agent's average fitness (Fig. 4a), compare first and second violins).

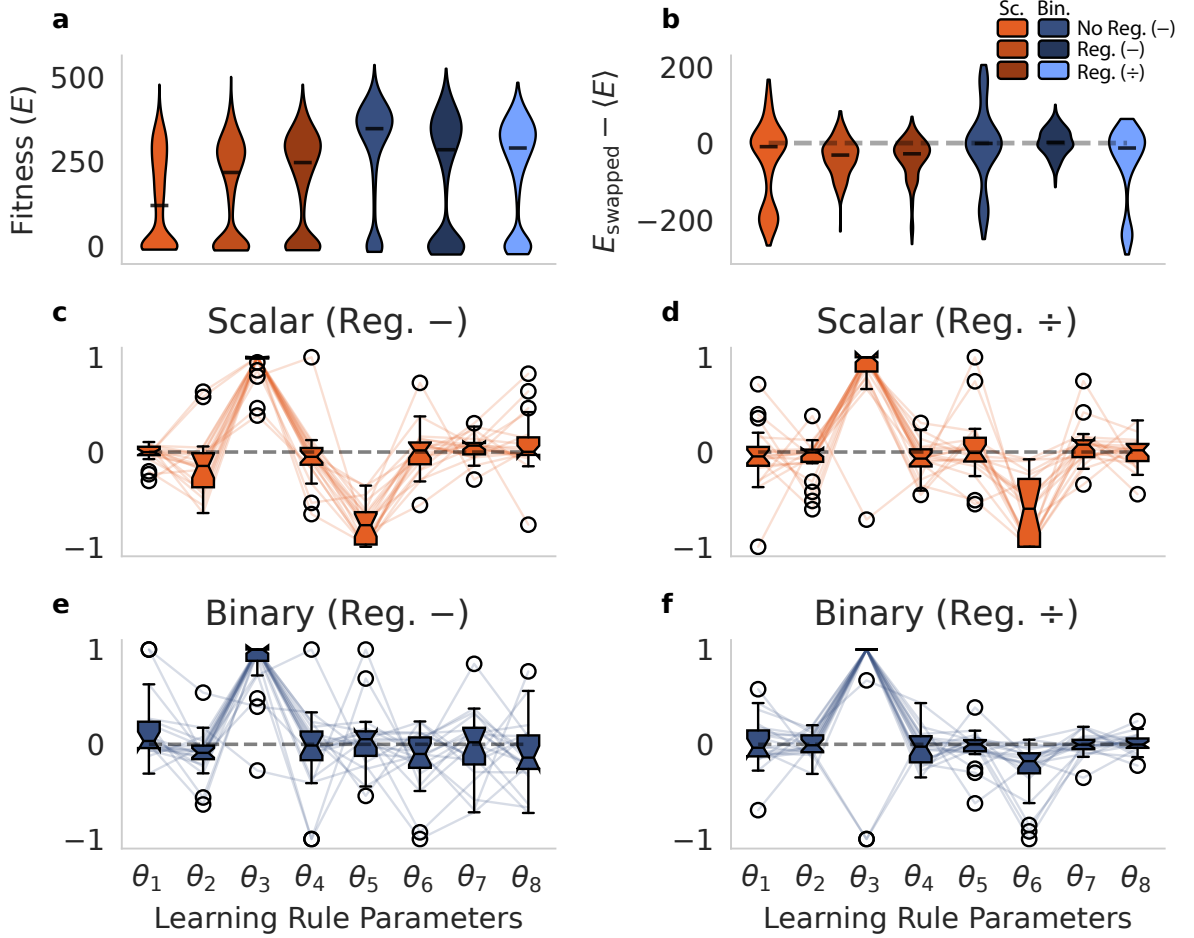


FIG. 4. *Plasticity rules converge with regularization. a. The fitness distribution for different simulations of scalar (Sc) and binary (Bin) sensory readouts, regularization of the plasticity parameters (Reg.) and subtractive (-) vs divisive (\div) weight normalization. b. The difference in fitness between the original and swapped agents for different simulations. c. The regularized rules developing in networks with scalar sensory readout for subtractive normalization. d. Same for divisive normalization. e. The regularized rules developing in networks with binary sensory readout for subtractive normalization. f. Same for divisive normalization.*

Finally, we note that when mixing different scalar agent’s sensory and motor networks, as in the previous section, the same reduction in the resulting fitness we observed for non-normalized rules persists. This suggests that despite the overall pattern being the same, differences between learning rules are still significant enough to require specialization of the motor networks for particular rules.

In the case of a binary readout, we see that all parameters converge to ≈ 0 except for θ_3 that converges to 1 (Fig. 4e). This leads to a learning rule of the form:

$$\Delta W_t = \eta_p \theta_3 X_t \cdot R_t \approx \eta_p R_t \cdot (x_1, \dots, x_n) \quad (6)$$

This rule clearly does not depend on the network’s sensory weights W_t and thus has no fixed points. Still, we see that the rule creates a Hebbian-like learning process: if a $w_i > 0$ and the reward is positive $R_t > 0$, then the absolute value of w_i would increase; the same if $w_i < 0$ and

$R_t < 0$, then the absolute value of w_i would also increase. Thus the rule strengthens the weights that are positively correlated with positive rewards (and also increases the absolute value of the weights negatively correlated with negative rewards) Overall, the learning rule leads to a sensory output y_t that is highly correlated with the sign of the reward R_t but develops weights with diverging absolute values.

For the binary sensory readouts, similar to the experiment in the previous section, swapping motor and sensory networks of different agents does not lead to a significant reduction in fitness, which indicates that binary readout networks maintain their generalizability and modularity under regularization of the plasticity parameters.

E. The method of weight normalization impacts the evolved learning rule

Most studies of Hebbian-like plasticity include a weight normalization mechanism that is necessary for stabilizing the learning dynamics over time [34, 42]. However, a particular form of normalization can be essential for the learning outcome. As we have established in the previous sections, introducing an information bottleneck between the sensory and motor networks leads to superior performance and generalizability. Still, we showed that the resulting learning rule generates divergent weights. We hypothesize that this can be changed by adjusting the weight normalization.

So far, for normalizing the weights, we have used a mean subtraction after each plasticity step [43, 44], which maintains the mean weight at 0:

$$\begin{aligned} W_t &= W_{t-1} + \Delta W_t \\ W_t &\leftarrow W_t - \langle W_t \rangle \end{aligned} \quad (7)$$

However, it is just one option out of a wide range of mechanisms for normalizing synaptic weights that have been proposed in neuroscientific and AI studies [45]. Moreover, some normalization mechanisms have been shown to have a significant effect on the performance of Hebbian learning [46]. To examine how the choice of weight normalization mechanisms may affect the evolution of the learning rule, we repeated the experiments described in the previous section for both network settings (scalar and binary) using a divisive normalization. This mechanism maintains the sum of the absolute values of the weights constant and equal to a given target S_g which is set to $S_g = 3$ for all the experiments:

$$\begin{aligned} W_t &= W_{t-1} + \Delta W_t \\ W_t &\leftarrow W_t \cdot \frac{S_g}{\sum_{i=1}^N |w_i(t)|}, \end{aligned} \quad (8)$$

Using this weight normalization mechanism after each plasticity step, we evolve 20 populations each, for both scalar and binary sensory networks. We see that the networks with a scalar sensory readout converge to a different learning rule (Fig. 4d) than the equivalent networks with subtractive normalization (Fig. 4c). Specifically, we see that instead of $\theta_5 \rightarrow -1$ and $\theta_6 \rightarrow 0$, under the divisive normalization $\theta_5 \rightarrow 0$ and $\theta_6 < 0$, which makes the learning rule to take the form:

$$\Delta W_t = \eta_p [\theta_1 X_t R_t + \theta_6 y_t] \quad (9)$$

where $\theta_1 \approx 1$ and $\theta_6 < 0$. In combination with the divisive normalization, this rule converges to a close approximation of the correct ingredient distribution W^c .

The trained agents maintain a similar fitness to previous experiments (Fig. 4b), and as expected from networks with scalar sensory readouts, the performance declines significantly ($p_{\text{value}} = 0.0003$, paired two-sided t-test)

when we test agents with motor and sensory networks that did not co-evolve.

The networks with a binary sensory readout evolve a learning rule similar to all other binary readout networks (i.e. most parameters converge to the vicinity of 0 except for $\theta_3 \rightarrow 1$). A small but significant difference between the networks that evolved with divisive normalization and those that evolved with subtractive normalization is that the former evolves a parameter $\theta_6 < 0$ (Fig. 4f). The same pattern is observed much more prominently for networks with a scalar readout (Fig. 4d), which suggests that a negative θ_6 parameter is important for networks with divisive normalization.

The divisive normalization successfully constrains the learned weights to be rather small (in contrast to other binary networks whose plasticity converges to very large sensory weights), and this does not seem to affect their fitness that remains relatively high (Fig. 4a). Also, as with the other binary readout networks, swapping motor and sensory networks between different agents does not lead to a significant reduction in performance ($p_{\text{value}} = 0.067$, paired two-sided t-test).

F. Trainable nonlinearity on the sensory readout

We now assess the impact the nonlinearity of the sensory network has on the evolved learning rule. To test this, we allow the steepness of the nonlinearity to evolve by setting:

$$g(x, \alpha) = \frac{2}{1 + e^{-\alpha x}} - 1 \quad (10)$$

and making α an evolvable parameter. For small values of α the function becomes effectively linear, while for large values $\alpha \rightarrow \infty$ it becomes a step function.

We evolve 40 populations of agents with the trainable sigmoid nonlinearity on the sensory network, half with subtractive and half with divisive weight normalization. For both types of normalization, we observe that in most cases, the resulting learning rules look similar to the ones evolved for the networks with a binary readout (Fig. 5a, b). Still, the two cases are visibly different in terms of the steepness of the evolved sigmoid functions. The divisive normalization evolves higher α (and consequently steeper nonlinearities), while the networks with subtractive normalization evolve relatively steep sigmoids (Fig. 5c), but much more diverse than the ones with divisive normalization (Fig. 5d).

In terms of fitness, both network types perform similarly (Fig. 5e), but when testing agents with motor and sensory networks that did not evolve together (Fig. 5f), the networks with subtractive normalization show a significant ($p_{\text{value}} = 0.00098$) difference in performance, while for the ones with divisive normalization, the performance remains largely unchanged ($p_{\text{value}} = 0.267$). We hypothesize this is due to the motor network adjusting to the specific slope of the sigmoid nonlinearity of its own

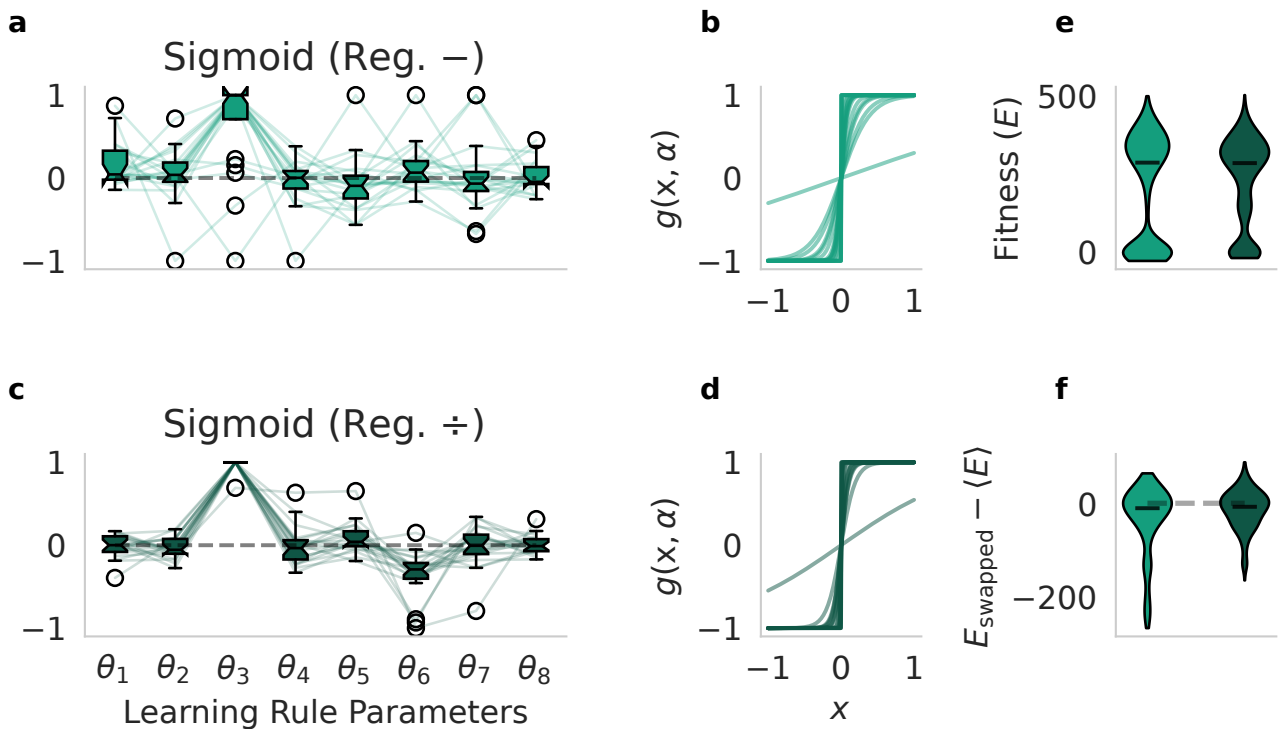


FIG. 5. Trainable readout nonlinearity leads to different rules depending on weight normalization. **a.** The evolved plasticity rules across 20 runs for agents with trainable sigmoid sensory network outputs and subtractive weight normalization. **b.** The sigmoid nonlinearities evolve a range of different slopes for subtractive normalization. **c.** The evolved plasticity rules with divisive weight normalization are similar to the ones developed for the binary networks with fixed nonlinearity, Fig. 4. **d.** The evolved sigmoid nonlinearities are still broad but steeper in networks trained with divisive normalization. **e.** The fitness distribution for subtractive ($-$, light green) vs divisive (\div , dark green) weight normalization. **f.** The fitness difference between the original and swapped agents.

sensory network (Fig. 5b,d), which makes the more diverse nonlinearities of the networks with subtractive normalization generate some fraction of the swapped agents that cannot perform (fitness close to 0) thus also perform significantly worse on average.

G. Static agents: The distribution of presented foods affects the emerging rule in static networks

To better study the convergence of plasticity rules, we remove the motor network and leave only the sensory network subject to the plasticity rule. For the following experiments, we define a new *static agent* whose plasticity rule we optimize to perform a simplified version of the task: to decide whether a presented food particle should be consumed or not. We use regularization (from Section IV D) and divisive weight normalization to avoid diverging weights. A new food particle is presented to the sensory network each of 125 time steps, and if the output of the network is positive $y_t = 1$ (which we take to mean the agent decides to consume the food particle), feedback is given to the agent and the sensory weights are updated

according to the network’s learning rule. The network’s fitness is assessed by the sum of the values of the foods it consumed.

After 15 generations, the plasticity parameters converge, at which point the average trained agent at the end of its lifetime can recognize poisonous foods with $> 99\%$ accuracy. However, the evolved parameters of plasticity rules differ from the ones evolved in the binary-readout moving agents with regularization and divisive normalization (compare Fig. 6a and Fig. 4f). Specifically, in contrast to the rules from the moving agents where $\theta_3 \rightarrow 1$ and $\theta_1 \rightarrow 0$, in the static networks, we observe that on average $0 < \theta_1, \theta_3 < 1$. We particularly note that in the static networks $\theta_1 + \theta_3 \approx 1$. Moreover, in the static networks, we see that $\theta_6 \rightarrow 0$ in contrast to the moving agent’s plasticity rules where on average $\theta_6 < 0$ (Fig. 6a). We hypothesize that this discrepancy is due to the fact that in the static networks, the plasticity step only occurs when $y_t = 1$, which effectively means the plasticity rule becomes:

$$\Delta W_t = \eta_p [\theta_1 X_t R_t y_t + \theta_3 X_t R_t] \stackrel{y_t=1}{=} \eta_p (\theta_1 + \theta_3) X_t R_t \approx \eta_p X_t R_t \quad (11)$$

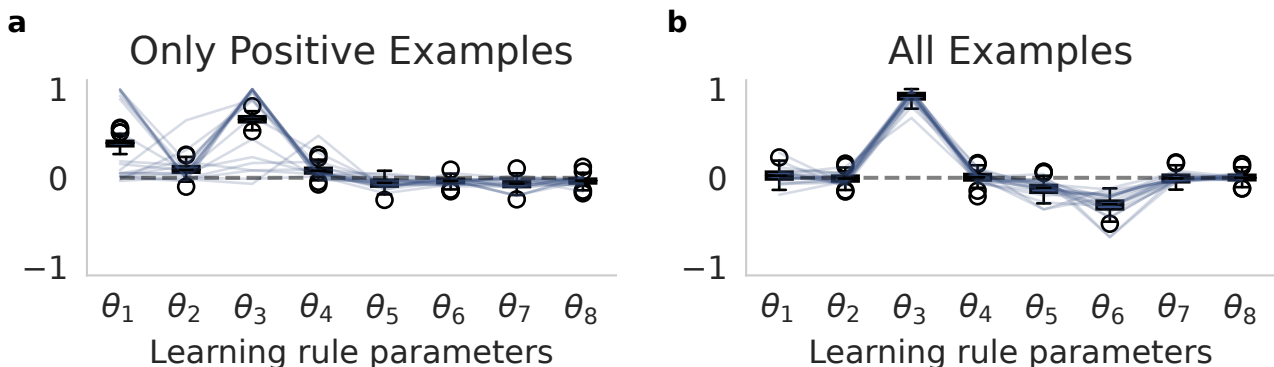


FIG. 6. The distribution of presented food values influences the evolved sensory learning rules. The evolved plasticity rules across 20 runs for static sensory networks with binary outputs, trained with divisive normalization. **a.** The evolved rules when the plasticity step happens only when the network’s readout is positive $y_t = 1$. **b.** The evolved rules when the plasticity step happens after every training step.

which is similar to the plasticity rule that evolved in the moving agent’s sensory networks, albeit without the negative θ_6 term (Fig. 4f).

To test this hypothesis, we repeat the experiment, but now the plasticity step happens at every time step, regardless of the sensory network’s output. We now observe that the resulting learning rule is much more similar to the equivalent rule of the moving agents (specifically, we see that $\theta_1 \rightarrow 0, \theta_3 \rightarrow 1$ and $\theta_1 < 0$ (Fig. 6b).

This suggests that the original discrepancy between the evolved rules of the static and moving networks was due to the assumption of perfect accuracy in the motor network when simulating only the sensory network (only the food particles to whom the sensory network assigns positive values were consumed). In reality, during a training generation, the moving agents tend to consume quite a few food particles to whom the sensory network assigns a negative value ($y_t = -1$) due to imperfections in the motor network and environmental variability (see the example trajectory, Fig. 2a). Thus, the fact that the embodied agents’ motor networks make mistakes leads to the development of a different learning rule from the one that the agent would evolve in a simplified scenario, where the locomotion of the agent is removed, and the food is instead presented to it sequentially.

H. Different objective functions lead to different evolved rules in static networks

We finally repeat the same experiment by testing static agents with a linear readout ($g(x) = x$) while following the same normalization standard as before (L1 regularization, divisive weight normalization). In this experiment, the readout is a scalar, and thus we train the networks to predict a given food particle’s real value as accurately as possible. To do so, we use the mean squared error (MSE) between the actual value of a given food par-

ticle and the sensory network’s prediction as a training loss for the plasticity parameters $L = \sum_t (R_t - y_t)^2$.

We see that the resulting rule (Fig. 7a) is not matching the equivalent rule from the moving networks (Fig. 4d), but rather evolved the classic Hebbian form $\Delta W_t \approx \eta_p X_t (R_t - y_t)$, associated with subtractive normalization in the moving agents (Fig. 4a).

We hypothesize that the difference in the resulting rules is due to the loss we use to train the static network. In the moving agents, the way the output of the sensory network is processed by the motor network is unclear, but due to the nature of the task, we can relatively safely assume that the only information the motor network really uses is whether a presented food has a positive or negative value. Thus, by testing the static network on the accuracy of its prediction of the presented food’s value, we optimized for a different goal than the EA does for the motor networks.

To test this assumption, we retrain the agents with a scalar readout y_t , but we assess its performance using the same loss as in the binary networks (if $y_t > 0$, the presented food is consumed, and the agent’s total fitness is the sum of all consumed foods). Thus, the plasticity rule has access to a scalar value of the readout (unlike the binary networks, where the y_t that the plasticity sees is always 1 or -1) but is optimized to perform a binary version of the task; decide whether a food is to be consumed rather than predict its value. We find that the resulting rule in this case (Fig. 7b) is more similar to the equivalent rule for the moving agents (Fig. 4d), with a strongly negative θ_6 parameter (albeit the static networks also develop a slightly negative θ_5 parameter which is absent for the moving agents’ rules). Nevertheless, the similarity suggests that the evolutionary pressure the moving agent’s plasticity rule is under is better approximated by the decision rather than the prediction task in the static networks.

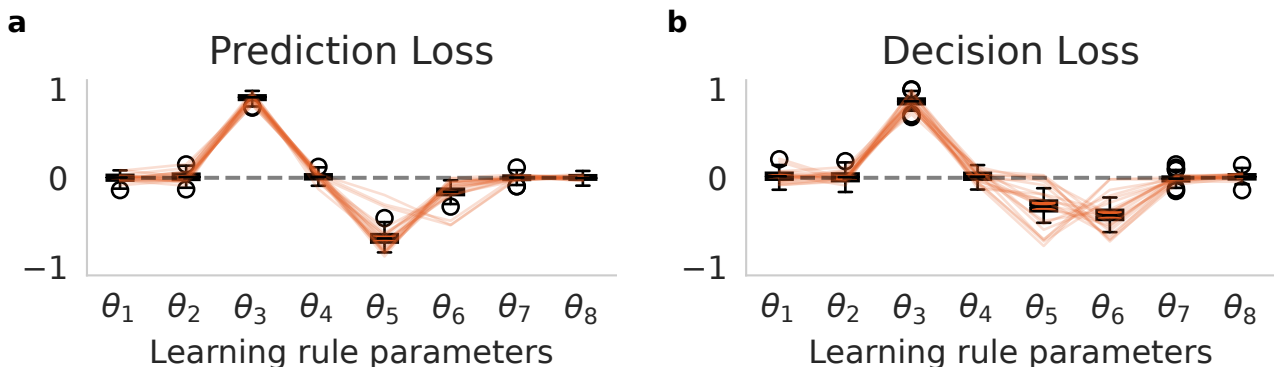


FIG. 7. The objective function of sensory learning determines the evolved learning rules. The evolved plasticity rules across 20 runs for static sensory networks with scalar outputs and divisive normalization. **a.** The evolved rules when the network is optimized to predict the presented food particle’s value **b.** The evolved rules when the network is optimized to decide whether a presented food particle should be consumed or not.

V. DISCUSSION

We study the emergence of plasticity rules in simple embodied agents that learn to perform a foraging task in a 2D environment. We show how a simple evolutionary algorithm can optimize the different parameters of a linear reward-modulated plasticity rule for solving sensory tasks with sufficient accuracy for guiding the foraging behaviour of embodied agents. Reward-modulated plasticity has been extensively studied as a plausible mechanism for credit assignment in the brain [40, 47, 48] and has found several applications in artificial intelligence and robotics tasks [23, 49]. Here, we demonstrate how such rules can be well-tuned to take into account different parameters and produce optimal behaviour in relatively complex systems with multiple interacting parts.

Recent studies on meta-learning plasticity rules have conclusively shown that identical learning outcomes can be reached via a diverse family of distinct learning rules [35]. Our findings suggest that learning rule redundancy can appear even in very simple systems, such as a single-layer feed-forward linear network. We propose a potential cause of this degeneracy in our network; as the relatively complex motor network is allowed to read out and process the outputs from the plastic network, any consistent information coming out of these outputs can be potentially interpreted in a behaviorally useful way. We find that while introducing an information bottleneck between the sensory and motor networks can enforce some rudimentary structure to the emerging rules, the variability remains high. Still, the introduction of a simple regularization scheme in the evolutionary process [30] forces the development of minimal rules that follow interpretable patterns and allow analytical treatment of the learning process.

Traditionally, plasticity rules are formalized in terms of the pre and post-synaptic activity of individual neurons, despite evidence of heavy bidirectional interaction

between synaptic plasticity mechanisms, the structure of individual neurons [50], the topology of neural networks [33, 51] and various compensatory mechanisms [34]. Here, we show that small changes in the neuronal nonlinearity of a single readout unit (linear vs step function) or the weight normalization mechanism (subtractive vs divisive) can strongly affect the evolutionary trajectory of reward-modulated plasticity rules, and lead to convergence on significantly different forms. Additionally, we see that the use of different activation functions and weight normalization schemes can impact the ability of learning rules to generalize across networks. Increased generalizability can potentially enable the mosaic evolution of artificial learning systems, a process that is well documented in biological systems [52, 53] and is associated with increased specialization and complexity [54].

We finally study how abstracting the learning process that takes place in the sensory network away from its interaction with the motor network affects the evolved learning rules. We find that small assumptions about important components in the learning process of the moving agents (such as the distribution of values of the food particles the agent consumes) or the precise objective the plasticity is optimizing for (decision vs prediction) can have a strong effect on the resulting plasticity rules. Following the insights of earlier studies [27, 28, 30, 32], this observation can be used for reverse engineering biologically observed plasticity rules and uncovering their objective function via simulation of simplified (and thus more controlled) learning processes. In particular, comparing the form of a learning rule observed in a complex, embodied system with simplified learning rules, optimized to solve known objective functions, can offer insights into the function of the learning rule in the original system, whose objective function cannot be directly known.

The optimization of functional plasticity in neural networks is a promising research direction both as a means to understand biological learning processes and as a tool

for building more autonomous artificial systems. Our results suggest that reward-modulated plasticity is highly adaptable to different environments and can be incorporated into larger systems that solve complex tasks.

VI. ACKNOWLEDGEMENTS

This work was supported by a Sofja Kovalevskaja Award from the Alexander von Humboldt Foundation.

EG and SK thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for their support. We acknowledge the support from the BMBF through the Tübingen AI Center (FKZ: 01IS18039A). AL is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 39072764.

-
- [1] Mauricio R. Papini. *Evolution of Learning*, pages 1188–1192. Springer US, Boston, MA, 2012.
- [2] Emilie C. Snell-Rood. An overview of the evolutionary causes and consequences of behavioural plasticity. *Animal Behaviour*, 85(5):1004–1011, 2013. Special Issue: Behavioural Plasticity and Evolution.
- [3] Kai Olav Ellefsen. The evolution of learning under environmental variability. pages 649–656, 07 2014.
- [4] Aimee S Dunlap and David W Stephens. Reliability, uncertainty, and costs in the evolution of animal learning. *Current Opinion in Behavioral Sciences*, 12:73–79, 2016. Behavioral ecology.
- [5] Frederic Mery and Tadeusz J. Kawecki. A cost of long-term memory in *Drosophila*. *Science*, 308(5725):1148–1148, 2005.
- [6] Alex Thornton and Tim Clutton-Brock. Social learning and the development of individual and group behaviour in mammal societies. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 366:978–87, 04 2011.
- [7] Ami Citri and Robert Malenka. Synaptic plasticity: Multiple forms, functions, and mechanisms. *Neuropsychopharmacology : official publication of the American College of Neuropsychopharmacology*, 33:18–41, 02 2008.
- [8] Daniel E. Feldman. Synaptic mechanisms for plasticity in neocortex. *Annual Review of Neuroscience*, 32(1):33–55, 2009. PMID: 19400721.
- [9] Wickliffe Abraham, Owen Jones, and David Glanzman. Is plasticity of synapses the mechanism of long-term memory storage? *npj Science of Learning*, 4, 12 2019.
- [10] Jonathan Power and Bradley Schlaggar. Neural plasticity across the lifespan. *Wiley Interdisciplinary Reviews: Developmental Biology*, 6, 11 2016.
- [11] Donald Hebb. The organisation of behaviour, 1948.
- [12] Elie Bienenstock, Leon Cooper, and Paul Munro. Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 2:32–48, 02 1982.
- [13] Henry D. I. Abarbanel, R. Huerta, and M. I. Rabinovich. Dynamical model of long-term synaptic plasticity. *Proceedings of the National Academy of Sciences*, 99(15):10132–10137, 2002.
- [14] Matthias Hennig. Theoretical models of synaptic short term plasticity. *Frontiers in Computational Neuroscience*, 7, 2013.
- [15] K. P. S. D. Kumarapathirana, Don Kulasiri, Sandhya Samarasinghe, and Jingyi Liang. Computational modelling of synaptic plasticity: A review of models, parameter estimation using deep learning, and stochasticity. pages 1–7, 12 2021.
- [16] Jeffrey Magee and Christine Grienberger. Synaptic plasticity forms and functions. *Annual Review of Neuroscience*, 43, 07 2020.
- [17] Y. Bengio, S. Bengio, and J. Cloutier. Learning a synaptic learning rule. In *IJCNN-91-Seattle International Joint Conference on Neural Networks*, volume ii, pages 969 vol.2–, 1991.
- [18] Andrea Soltoggio, Kenneth O. Stanley, and Sebastian Risi. Born to learn: The inspiration, progress, and future of evolved plastic artificial neural networks. *Neural Networks*, 108:48–67, 2018.
- [19] Eseoeghene Ben-Iwhiwhu, Pawel Ladosz, Jeffery Dick, Wen-Hua Chen, Praveen Pilly, and Andrea Soltoggio. Evolving inborn knowledge for fast adaptation in dynamic pomdp problems. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference, GECCO '20*, page 280–288, New York, NY, USA, 2020. Association for Computing Machinery.
- [20] Elias Najarro and Sebastian Risi. Meta-learning through hebbian plasticity in random networks. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20719–20731. Curran Associates, Inc., 2020.
- [21] Joachim Winther Pedersen and Sebastian Risi. Evolving and merging hebbian learning rules: Increasing generalization by decreasing the number of rules. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '21*, page 892–900, New York, NY, USA, 2021. Association for Computing Machinery.
- [22] Anil Yaman, Giovanni Iacca, Decebal Constantin Mocanu, Matt Coler, George Fletcher, and Mykola Pechenizkiy. Evolving Plasticity for Autonomous Learning under Changing Environmental Conditions. *Evolutionary Computation*, 29(3):391–414, 09 2021.
- [23] Jeroen Burms, Ken Caluwaerts, and Joni Dambre. Reward-modulated hebbian plasticity as leverage for partially embodied control in compliant robotics. *Frontiers in Neurorobotics*, 9, 2015.
- [24] Luke Metz, Niru Maheswaranathan, Brian Cheung, and Jascha Sohl-Dickstein. Learning to learn without labels, 2018.
- [25] Jack Lindsey. Learning to learn with feedback and local plasticity. In *Real Neurons & Hidden Units: Future directions at the intersection of neuroscience and artificial*

- intelligence @ NeurIPS 2019*, 2019.
- [26] Blake Richards, Timothy Lillicrap, Philippe Beaudoin, Y. Bengio, Rafal Bogacz, Amelia Christensen, Claudia Clopath, Rui Costa, Archy Berker, Surya Ganguli, Colleen Gillon, Danijar Hafner, Adam Kepecs, Nikolaus Kriegeskorte, Peter Latham, Grace Lindsay, Kenneth Miller, Richard Naud, Christopher Pack, and Konrad Kording. A deep learning framework for neuroscience. *Nature Neuroscience*, 22:1761–1770, 11 2019.
- [27] Basile Confavreux, Friedemann Zenke, Everton Agnes, Timothy Lillicrap, and Tim Vogels. A meta-learning approach to (re)discover plasticity rules that carve a desired function into a neural network. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 16398–16408. Curran Associates, Inc., 2020.
- [28] Jakob Jordan, Maximilian Schmidt, Walter Senn, and Mihai A Petrovici. Evolving interpretable plasticity for spiking networks. *eLife*, 10:e66273, oct 2021.
- [29] Danil Tyulmankov, Guangyu Robert Yang, and L.F. Abbott. Meta-learning synaptic plasticity and memory addressing for continual familiarity detection. *Neuron*, 110(3):544–557.e8, 2022.
- [30] Navid Shervani-Tabar and Robert Rosenbaum. Meta-learning biologically plausible plasticity rules with random feedback pathways. *Nature Communications*, 14, 03 2023.
- [31] Basile Confavreux, Poornima Ramesh, Pedro J. Goncalves, Jakob H. Macke, and Tim P. Vogels. Meta-learning families of plasticity rules in recurrent spiking networks using simulation-based inference. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [32] Yash Mehta, Danil Tyulmankov, Adithya E. Rajagopalan, Glenn C. Turner, James E. Fitzgerald, and Jan Funke. Model-based inference of synaptic plasticity rules. *bioRxiv*, 2023.
- [33] Emmanouil Giannakakis, Oleg Vinogradov, Victor Buen-dia, and Anna Levina. Recurrent connectivity structure controls the emergence of co-tuned excitation and inhibition. *bioRxiv*, 2023.
- [34] Friedemann Zenke and Wulfram Gerstner. Hebbian plasticity requires compensatory processes on multiple timescales. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1715):20160259, 2017.
- [35] Poornima Ramesh, Basile Confavreux, Pedro J. Goncalves, Tim P. Vogels, and Jakob H. Macke. Indistinguishable network dynamics can emerge from unlike plasticity rules. *bioRxiv*, 2023.
- [36] Emmanouil Giannakakis, Sina Khajehabdollahi, and Anna Levina. Environmental variability and network structure determine the optimal plasticity mechanisms in embodied agents. volume ALIFE 2023: Proceedings of the 2023 Artificial Life Conference, 07 2023.
- [37] Sina Khajehabdollahi, Jan Prosi, Emmanouil Giannakakis, Georg Martius, and Anna Levina. When to Be Critical? Performance and Evolvability in Different Regimes of Neural Ising Agents. *Artificial Life*, 28(4):458–478, 11 2022.
- [38] Yoshinori Aso, Daisuke Hattori, Yang Yu, Rebecca M Johnston, Nirmala A Iyer, Teri-TB Ngo, Heather Dionne, LF Abbott, Richard Axel, Hiromu Tanimoto, and Gerald M Rubin. The neuronal architecture of the mushroom body provides a logic for associative learning. *eLife*, 3:e04577, dec 2014.
- [39] Benjamin Fabian and Silke Sachse. Experience-dependent plasticity in the olfactory system of drosophila melanogaster and other insects. *Frontiers in Cellular Neuroscience*, 17, 2023.
- [40] Robert Legenstein, Dejan Pecevski, and Wolfgang Maass. A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback. *PLOS Computational Biology*, 4(10):1–27, 10 2008.
- [41] Kalyanmoy Deb. *Multi-objective Optimisation Using Evolutionary Algorithms: An Introduction*, pages 3–34. Springer London, London, 2011.
- [42] Terry Elliott. An Analysis of Synaptic Normalization in a General Class of Hebbian Models. *Neural Computation*, 15(4):937–963, 04 2003.
- [43] Kenneth D. Miller and David J. C. MacKay. The Role of Constraints in Hebbian Learning. *Neural Computation*, 6(1):100–126, 01 1994.
- [44] Geoffrey J. Goodhill and Harry G. Barrow. The Role of Weight Normalization in Competitive Learning. *Neural Computation*, 6(2):255–269, 03 1994.
- [45] Yang Shen, Julia Wang, and Saket Navlakha. A Correspondence Between Normalization Strategies in Artificial and Biological Neural Networks. *Neural Computation*, 33(12):3179–3203, 11 2021.
- [46] Samuel Eckmann and Julijana Gjorgjieva. Synapse-type-specific competitive hebbian learning forms functional recurrent networks. *bioRxiv*, 2022.
- [47] Răzvan V. Florian. Reinforcement Learning Through Modulation of Spike-Timing-Dependent Synaptic Plasticity. *Neural Computation*, 19(6):1468–1502, 06 2007.
- [48] Dorit Baras and Ron Meir. Reinforcement Learning, Spike-Time-Dependent Plasticity, and the BCM Rule. *Neural Computation*, 19(8):2245–2279, 08 2007.
- [49] Zhenshan Bing, Ivan Baumann, Zhuangyi Jiang, Kai Huang, Caixia Cai, and Alois Knoll. Supervised learning in snn via reward-modulated spike-timing-dependent plasticity for a target reaching vehicle. *Frontiers in Neuro-robotics*, 13, 2019.
- [50] Simon d’Aquin, Andras Szonyi, Mathias Mahn, Sabine Krabbe, Jan Gründemann, and Andreas Lüthi. Compartmentalized dendritic plasticity during associative learning. *Science*, 376(6590):eabf7052, 2022.
- [51] Markus Butz, Ines Steenbuck, and Arjen van Ooyen. Homeostatic structural plasticity increases the efficiency of small-world networks. *Frontiers in Synaptic Neuroscience*, 6, 2014.
- [52] Stephanie Fong, Björn Rogell, Mirjam Amcoff, Alexander Kotrschal, Wouter van der Bijl, Séverine D. Buechel, and Niclas Kolm. Rapid mosaic brain evolution under artificial selection for relative telencephalon size in the guppy (*Poecilia reticulata*). *Science Advances*, 7(46):eabj4314, 2021.
- [53] Erika L Schumacher and Bruce A Carlson. Convergent mosaic brain evolution is associated with the evolution of novel electrosensory systems in teleost fishes. *eLife*, 11:e74159, jun 2022.
- [54] Biagio D’Aniello, Anna Di Cosmo, Anna Scandurra, and Claudia Pinelli. Mosaic and concerted brain evolution: The contribution of microscopic comparative neuroanatomy in lower vertebrates. *Frontiers in Neuroanatomy*, 13, 2019.

Appendix C Dendritic nonlinearities and synapse-type specific input clustering enable the development of input selectivity in diverse settings

Contains the full text of the manuscript: Giannakakis, E., Bird, A., Jedlicka, P., Cuntz, H., and Levina, A. (In Preparation). Dendritic nonlinearities and synapse-type specific input clustering enable the development of input selectivity in diverse settings.

Dendritic nonlinearities and synapse-type specific input clustering enable the development of input selectivity in diverse settings.

Emmanouil Giannakakis,^{1,2} Alex D. Bird,^{3,4} Peter Jedlicka,³ Hermann Cuntz,^{3,4,5} and Anna Levina^{1,2}

¹*Department of Computer Sciences, University of Tübingen, Germany*

²*Max Planck Institute for Biological Cybernetics, Tübingen, Germany*

³*Computer-Based Modelling in the field of 3R Animal Protection, Faculty of Medicine, Justus Liebig University Giessen, Giessen, Germany*

⁴*Ernst Strüngmann Institute (ESI) for Neuroscience in cooperation with Max Planck Society, Frankfurt am Main, Germany*

⁵*Frankfurt Institute for Advanced Studies, Frankfurt am Main, Germany*

Pyramidal neurons are characterized by a complex dendritic structure that enables connectivity across large networks. Computations within dendritic branches augment somatic computations and influence multiple functions, including the shaping of neuron’s orientation selectivity. However, the influences of different aspects of dendritic architectures and specific dendritic wiring on the synaptic plasticity mechanisms via which orientation preference emerges during development remain unclear. Here, we study how dendritic structures can determine the emergence of specific orientation tunings by acting as inductive biases for competitive excitatory and inhibitory Hebbian plasticity. We first show that dendritic nonlinearities can significantly increase the robustness of plasticity to input noise. Moreover, using simulation-based inference, we find that the shape of the dendritic nonlinearity can effectively control the tuning width of the emerging orientation preference. Finally, we show how differences in the synaptic distribution of E/I inputs on the dendritic tree control the stability of learned orientation preference and the sharpness of tuning. Our findings indicate that different features of dendritic structures can jointly regulate the emergence and form of orientation selectivity in pyramidal neurons by influencing the learning dynamics of synaptic plasticity.

I. INTRODUCTION

Pyramidal neurons throughout the cortex are characterized by an intricate morphology [1, 2] and highly nonlinear dynamics across their dendrites [3, 4], which endow them with the ability to process complex inputs [5–7]. The dendritic structure of pyramidal neurons [8] is particularly important in this respect, since their intrinsic dynamics [9, 10] are known to play an important role in driving spiking activity [11, 12], controlling neural information processing [13–16] and contributing to plasticity [17–20]. Moreover, the dendritic trees of pyramidal neurons are known to develop very differently in different cortical areas [21–23] and specific parts of the dendritic tree have been found to receive diverse types of feedforward and feedback input [24, 25], which they subject to a complex nonlinear integration process [26–29].

The shape of dendritic nonlinearities [30] varies significantly, even among the dendrites of the same neuron [16, 31] and is crucial for efficient computations [32, 33] and the regulation of local synaptic plasticity [34]. Additionally, the specificity of inputs on distinct parts of the dendritic tree [35, 36] has been linked to various aspects of information processing, including the tuning of orientation-selective neurons. For example, a study of layer 2/3 pyramidal neurons in the ferret visual cortex [37], showed that clustering of similarly tuned inputs on the dendritic tree was correlated with sharper tuning. A later study on mouse V1 visual cortex [38] also found that similar visual inputs on the dendritic tree tend to cluster and related this finding to orientation selectivity. Finally, a recent study [39] showed how dendritic morphology and the distribution of E and I synapses on the

dendritic tree correlate with orientation tuning in mouse V1 pyramidal neurons.

These findings create a strong case for a link between the distribution of synapses along the dendritic tree and the orientation tuning of the neuron. However, orientation selectivity is commonly thought to emerge via Hebbian synaptic plasticity [40]. Indeed, the idea that different variations of Hebbian learning could produce orientation selectivity has been known for decades [41, 42]. A significant number of theoretical studies have focused on proposing different protocols that account for additional constraints such as robustness and stability [43–45], and the inclusion of selective inhibitory neurons with plastic connections [46–49]. Thus, a question arises about the potential interaction between the learning and dendritic structures. Specifically, whether and how the emergence of orientation selectivity via Hebbian-like synaptic plasticity can be affected by different dendritic structures.

The impact of dendritic dynamics on synaptic plasticity has been established in several studies, showing that dendritic action potentials can induce long-term potentiation [50, 51] and demonstrating the impact of synaptic clustering on dendrites on memory encoding [52–54]. Moreover, theoretical studies have shown how synaptic plasticity rules accounting for dendritic morphology can better capture different experimental observations [55, 56]. Nevertheless, the impact of different active dendritic structures on the learning dynamics of Hebbian rules at spatially clustered synapses has not been explicitly studied before.

Here, we study how the shape of dendritic nonlinearities and the distribution of E/I synapses along the dendritic tree can effectively control the learning dynamics

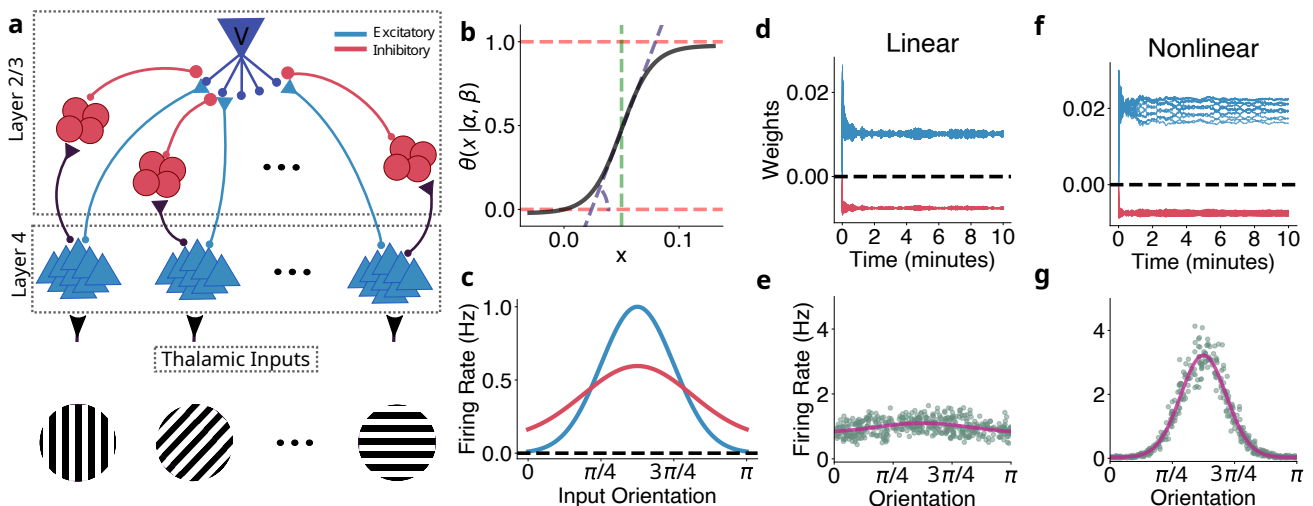


FIG. 1. **A dendritic nonlinearity makes learning orientation selectivity robust to noise.** (a) A diagram of the modelled circuit. Groups of orientation-tuned excitatory and inhibitory neurons target a single postsynaptic cell with plastic connections. The blue and red lines indicate plastic excitatory and inhibitory connections, respectively, while the black arrows indicate fixed excitatory connections. (b) The dendritic nonlinearity is a sigmoid parametrized by its slope (α) and centre (β). (c) The tuning of an excitatory and an inhibitory (scaled) neuron centred at $\pi/2$. Inhibitory tuning is broader. (d) Noisy inputs lead to a collapse of the developing weight trajectories when inputs are summed linearly. (e) This lack of diversity in the learned weights prevents the emergence of input selectivity. (f) The introduction of dendritic nonlinearities increases diversity among the learned weights, which leads to the emergence of (g) robust orientation selectivity (here centered at $\pi/2$).

and orientation tuning developing via interacting E/I Hebbian plasticity. We find that simple dendritic structures can significantly increase the robustness to noise of a previously proposed plasticity protocol [49]. We show that the steepness and threshold of the dendritic nonlinearity and the clustering of similar E and I inputs on the dendritic tree control the sharpness of the learned orientation preference. Our findings highlight the interdependence between neuronal structures and synaptic learning and offer testable predictions that could guide future experimental work on the development of selectivity.

II. RESULTS

We model the development of orientation selectivity in an abstract pyramidal neuron receiving input from multiple orientation selective subpopulations (each tuned to a different orientation) of excitatory and inhibitory neurons, Fig. 1a). The feedforward input connections to the pyramidal neuron are plastic, following classic Hebbian modification with additional synapse-type specific competitive normalization [49] and homeostatic plasticity [57] that maintains a constant post-synaptic firing rate.

A. The presence of dendritic nonlinearities enables the learning of noisy inputs

The plasticity protocol with competitive normalization, proposed by [49], relies on the input covariance structure to create orientation selectivity via interacting excitatory and inhibitory plasticity. However, as has been previously shown [58], the presence of input noise damages the input covariance structure and compromises the emergence of selectivity. We begin by investigating whether the introduction of a rudimentary dendritic structure can result in the emergence of selectivity despite the presence of strong decorrelating noise.

Starting with the rate neuron model from [49], we add a number of dendrites Fig. 1a), which we model as sigmoid functions parametrized by their centre and slope (Fig. 1b). Incoming E/I inputs to the neuron are randomly assigned to a dendrite, and the total input to each dendrite passes through the sigmoid nonlinearity before being added to the somatic voltage. Additionally, to increase biological realism we introduce broader inhibitory tuning (Fig. 1c) in the input population [59] and an inhibitory delay [60], which makes the learning task even more challenging.

The addition of noise in the input population decorrelates the co-tuned inputs, which, in the absence of dendritic nonlinearities, leads to a collapse of the learning dynamics to a single point (Fig. 1d). The lack of diversity in the learned weights, essentially ensuring the same post-synaptic response regardless of the tuning of the presyn-

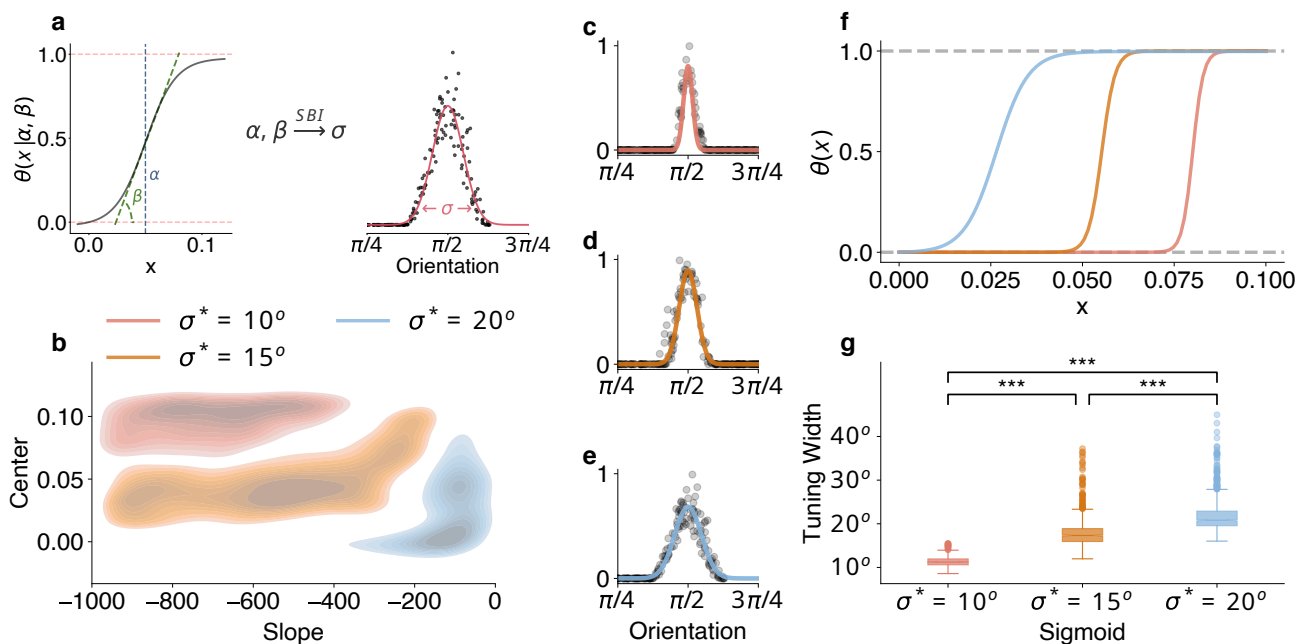


FIG. 2. **The shape of the dendritic nonlinearity controls the sharpness of the emergent tuning.** **a.** Using simulation-based inference, we map the slope and centre of the dendritic nonlinearity to the sharpness of the emerging tuning. **b.** The posterior distribution of the nonlinearity parameters for different postsynaptic tuning widths. **c.** An illustration of the postsynaptic neuron responses for neurons with the three MAP (c-e) dendritic nonlinearities. The lines illustrate response curves with the target tuning widths $\sigma^* = 10^\circ$, $\sigma^* = 15^\circ$ and, $\sigma^* = 20^\circ$ respectively. The neuron responses match the target tuning curves quite well. For illustration purposes, the responses have been scaled to 1 and centred at $\pi/2$. **f.** Sampled nonlinearities for the three different target tuning widths. **g.** The distribution of the learned tuning widths for the MAP nonlinearity of each posterior. Despite the stochastic nature of the data, different nonlinearities lead to statistically significant differences in the resulting tuning width.

naptic inputs, effectively prevents the emergence of input selectivity (Fig. 1e).

The introduction of dendritic nonlinearities greatly ameliorates this effect. In particular, we notice that a nonlinear summation of different inputs before the total input is added to the somatic voltage, significantly increases the diversity of the learned weights (Fig. 1f), which leads to the emergence of stable orientation selectivity in the post-synaptic neuron. Specifically, after a relatively brief period of rapid change, the learned weights settle to fixed values that follow the correlation structure of the input population (connection weights from similarly tuned neurons will have similar values) which leads to the post-synaptic neuron producing consistent and diverse responses to the presentation of different stimuli (Fig. 1g).

This result is relatively robust to the parameters of the sigmoid nonlinearity. Specifically, we observe that almost any sufficiently steep sigmoid will produce orientation selectivity in the face of noisy input. However, we note that steeper nonlinearities with more right-shifted centres (higher thresholds) tend to produce more robust and sharper orientation tuning.

B. The shape of the nonlinearity controls the sharpness of the emergent tuning

Given that the presence of a nonlinearity generally improves robustness to input noise, we turn to the question of what impact the precise shape of the nonlinearity has on tuning and selectivity. We quantify the tuning width of the postsynaptic neuron by the standard deviation of a Gaussian, fitted to the neuron's responses to different stimuli, and study how this tuning width is affected by changes in the shape of the dendritic nonlinearity.

Due to the highly stochastic nature of the learning dynamics and neural responses, we employ a statistical approach, using simulation-based inference [61] to estimate the impact of different sigmoid shapes on the resulting tuning. Specifically, we wish to estimate the posterior $p(\alpha, \beta | \sigma^*)$, i.e. the most likely parameters of the nonlinearity (centre α and slope β) that would lead to a given tuning width σ^* (Fig. 2a). Using a simplified form of filter SBI (fSBI) [62], we begin by sampling sigmoid parameters from a uniform prior and simulating the learning trajectory of a network for each of them. Once an adequate number of samples has been collected, we filter out the ones that do not satisfy a number of metrics. Specifically, we look at the weight development of the collected

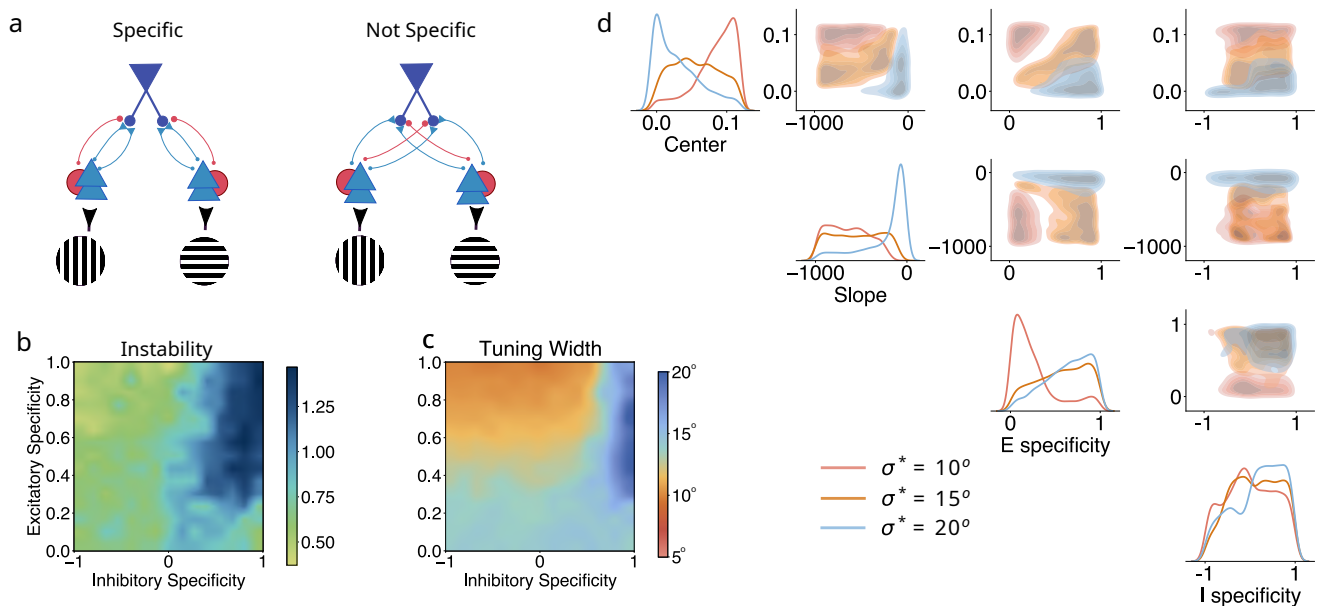


FIG. 3. **Varying levels of E and I dendritic specificity impact the synaptic learning outcomes.** **a.** We define a metric of dendritic specificity, which controls the likelihood that inputs from an input group will project to a specific dendrite as opposed to projecting randomly across the dendritic tree. We evaluate plasticity outcomes for different combinations of E and I specificity. We find that: **b.** While the E specificity barely impacts stability, high I specificity creates instabilities in the learning dynamics while negative I specificity ensure maximally stable weight trajectories. **c.** On the contrary, when it comes to tuning width, the I specificity is largely irrelevant, while the E specificity appears to have a major effect, with high E specificity leading to sharper orientation tuning. **d.** Finally, we perform simulation-based inference for both the dendritic nonlinearity parameters and the E/I specificity. We find that, while I specificity does not influence the tuning sharpness much, the E specificity can interact with the nonlinearity parameters (in particular the centre) to control the emerging tuning width.

samples and filter for high weight stability (i.e no drift in the preferred orientation), strong E/I weight co-tuning (ensuring E and I connectivity profiles match) and high R^2 for the fitted tuning curve on the post-synaptic neuron’s responses. Using the remaining samples, which give us the likelihood $p(\sigma^*|\alpha, \beta)$, we estimate the posterior using a neural posterior estimator.

After this fitting process, we look at the resulting posterior distributions for different target tunings of $\sigma^* = 10^\circ$, $\sigma^* = 15^\circ$ and $\sigma^* = 20^\circ$. We observe that the distributions are quite distinct (Fig. 2b), with sharper target tuning leading to right-shifted centre and steeper slope for the nonlinearity. In order to establish the accuracy of the inference process, we sample from the posteriors and verify the different tuning widths (Fig. 2c-e) of networks simulated with the corresponding dendritic nonlinearities (Fig. 2f). Finally, using the maximum a posteriori (MAP) of each network, we simulate multiple networks and establish that they lead to statistically significant differences between the resulting orientation tunings (Fig. 2g). Our findings, suggest that the shape of the nonlinearity can very reliably direct the learning dynamics towards developing different tuning widths.

C. The distribution of E/I synapses on dendrites affects the learning

Having shown how the presence and shape of the dendritic nonlinearity affect the emergence of orientation tuning, we return to our original question of the impact that E/I synaptic distribution across the active dendritic tree has on the tuning.

In our model, we simulate N dendrites and N input groups (each selective to one orientation), which allows us to make a 1-1 mapping, nominally “assigning” a dendrite to each input group. We define a dendritic specificity metric S_A , $A \in \{E, I\}$, such that $S_A = 1$ if the neurons of type A from each input group target exactly only their assigned dendrite (i.e. each dendrite receives input only from neurons selective to a specific orientation). $S_A = 0$ if inputs target any dendrite with equal probability (this is the case we have studied so far with $S_E = S_I = 0$) and intermediate values indicate the likelihood that neurons from a given group target their corresponding dendrite preferentially compared to target any dendrite randomly (Fig. 3a). Finally, while for excitatory inputs $S_E \in [0, 1]$, we allow negative specificity for inhibitory inputs, with $S_I \in [-1, 1]$. If $S_I < 0$ the specificity indicates that the inhibitory inputs are less likely to target their orientation’s corresponding dendrite as opposed to other

dendrites, with $S_I = -1$ is denoting a situation where inhibitory inputs never target their orientation’s corresponding dendrite.

We simulate networks with a fixed dendritic nonlinearity for different dendritic specificities in the range $S_E \times S_I \in [0, 1] \times [-1, 1]$ and measure the stability of the learned weights as well as the resulting tuning width of the postsynaptic neuron. We first observe that high inhibitory specificity tends to lead to higher instability, while $S_I < 0$ usually leads to more stable weight dynamics (Fig. 3b). Additionally, excitatory specificity has virtually no impact on stability, but controls the resulting orientation tuning, with higher excitatory specificity leading to sharper tuning and more random excitatory targeting of dendrites leading to broader neuron responses (Fig. 3c). This finding is in agreement with the experimental studies linking clustering of similar inputs in the same parts of the dendritic tree with sharper tuning [37, 39].

We finally, we repeat the SBI posterior estimation, fitting both the nonlinearity (slope and centre) and specificity (E and I) parameters (Fig. 3d). For the nonlinearity parameters, we see that the same pattern observed in the original case (Fig. 2b) is maintained (i.e. sharper tuning is associated with steeper and right-shifted sigmoids). Additionally, we see that the excitatory specificity interacts with the sigmoid shape, especially the tuning, in a way that allows for the two parameters to compensate for each other in order for a specific tuning width to emerge. Finally, as expected from the grid search on the specificity parameters (Fig. 3c), we see that the inhibitory specificity has almost no influence on the tuning, a result that agrees with the experimental findings of [39].

III. DISCUSSION

Branched dendritic trees fundamentally allow for connectivity between neurons across large networks and so allow the brain to carry out complex computations at scale [63]. The dendrites themselves, however, also provide a space upon which important functions can be implemented at the level of a single neuron [6, 64]. Here we have shown how nonlinear dendrites enable neurons to robustly learn orientation preferences via competitive Hebbian plasticity, and further how the spatial specificity of connectivity interacts with the shape of local nonlinearities to control the sharpness of the learned neuronal selectivity.

We have demonstrated this effect on an abstract model of a visual sensory neuron, where the computational task is relatively well-defined, but believe that the relationship between synaptic spatial specificity and the parameters of local nonlinearities will apply generally across different neural circuits. The question of how neurons can form and make use of dendritic compartments to enhance their computational power has been investigated in a number of other studies. Two in particular, are highly comple-

mentary to our results here. Legenstein and Maass (2011) used experimental data on the plasticity of dendritic nonlinearities in different cell types [56, 65] and investigated how competition between excitable branches could allow the development of clustered effective connectivity and solve a feature binding problem [20]. There are two major differences between the mechanisms investigated in this study and ours. Firstly, they take an initial case of uniformly distributed excitatory inputs; in the current study we consider how preferential spatial targeting, as observed in some systems [66], might provide an inductive bias that shapes learning. Secondly, they allow for plasticity in the coupling of dendritic activity in each compartment to the soma, effectively the magnitudes of the dendritic nonlinearities. We rather consider the question of how nonlinearities with different thresholds and sensitivities, but homogeneous across the neuron, might lead to the emergence of different orientation tunings. Kirchner and Gjorgjieva (2021) considered a mechanistic model of the emergence of synaptic organisation on developing cortical dendrites [67]. Their study did not focus on dendritic nonlinearities, but found that steeper activation functions lead to spatially smaller and more highly correlated clusters of synapses. Both of these papers imply interesting extensions to our work; by allowing both the dendritic nonlinearity and the spatial specificity to vary dynamically with inputs it might be possible for single neurons to efficiently develop a range of responses to distinct inputs and participate in multiple coding ensembles [68–70].

Overall, our results have demonstrated how two major advantages allowed by spatially extended dendrites, the separation and hierarchical filtering of inputs, interact and enable the robust development of useful computational function.

IV. MATERIALS & METHODS

A. Neuron Model

We model a single layer 2/3 pyramidal neuron with 10 dendrites. Each dendrite sums the incoming E and I and passes the sum through a simple nonlinear function. The input to each dendrite d is given by:

$$I_d = \sum_{i \in S^E(d)} W_i^E s_i^E - \sum_{i \in S^I(d)} W_i^I s_i^I - \lambda \cdot v, \quad (1)$$

where $S^E(d)$ and $S^I(d)$ the indices of the excitatory and inhibitory input neurons targeting the dendrite d respectively, W_i^E and W_i^I the connection strengths of the i th excitatory and the i th inhibitory input neurons, s_i^E and s_i^I the activity of the i th excitatory and the i th inhibitory input neurons, v is the somatic voltage and λ is a small leak constant. The dendritic inputs are summed in the

soma, whose voltage is given by:

$$\tau \frac{dv}{dt} = -v + \sum_{d \in D} \theta(I_d), \quad (2)$$

where τ is the membrane time constant, D is the set of dendrites and $\theta(x)$ is the dendritic nonlinearity, given by:

$$\theta(x) = \sigma(x) - \sigma(0) \quad (3)$$

the subtraction of $\sigma(0)$ ensures that $\theta(0) = 0$, which prevents violations of Dale's law due to shifts in the nonlinearity. $\sigma(x)$ is the sigmoid function:

$$\sigma(x) = \frac{1}{1 + e^{\beta \cdot (x - \alpha)}} \quad (4)$$

The parameters α, β determine the slope and centre of the function respectively and are optimized via evolutionary computation. Finally, the firing rate of the neuron is given as:

$$r(t) = \kappa \cdot [v(t) - v_0]_+^n \quad (5)$$

Following [49] we set $\kappa = 1, v_0 = 0.25$ and $n = 2$ for all our simulations.

B. Input

The input consists of 200 excitatory and 50 inhibitory neurons, divided in 10 groups of equal size (each containing 20 excitatory and 5 inhibitory neurons). Each group is tuned to a specific orientation, equally spaced between 0° and 180° .

Every 200 mS an input with orientation θ_t^{stim} uniformly drawn from the interval $[0^\circ, 180^\circ]$ is presented to the neuron. The response to the stimulus of the input neurons is then given by:

$$FR_i = R_{\max} \cdot \exp \left[\frac{(\theta_t^{\text{stim}} - \theta_i)^2}{2\sigma_i^2} \right] \quad (6)$$

where R_{\max} is the neuron's maximum firing rate, θ_i is the neuron's preferred orientation and σ_i is the neuron's tuning width. The inhibitory neurons follow the same orientation preferences and activation function, but they have a significantly broader tuning of 45° as opposed to 30° and respond to stimuli with a delay of 20 mS.

Finally, the noise of the network is quantified by $n_s \in [0, 1]$ and the input to the postsynaptic neuron from a presynaptic neuron i is given by:

$$\text{Inp}_i = (1 - n_s) \cdot FR_i + n_s \cdot \xi_u \quad (7)$$

where $\xi_u \sim U(0, 2R_{\max})$. This maintains a similar average input strength for different levels of noise.

C. Plasticity

We use a very simple Hebbian [71] learning rule for both E and I synapses

$$\Delta W_k^X = f(X_k, r) = \eta_X \cdot X_k \cdot r, \quad X \in \{E, I\} \quad (8)$$

In order to stabilize the weights and promote the development of input selectivity, we apply synapse type-specific [49] competitive normalization of weights in the post synaptic neuron

$$W_k^X \leftarrow W_k^X \cdot \frac{W_{\text{target}}^X}{\sum_i W_i^X}, \quad X \in \{E, I\} \quad (9)$$

Finally, in order to maintain a stable postsynaptic firing rate across experiments, we scale the E target weight sum W_{target}^E as:

$$W_{\text{target}}^E \leftarrow (1 - \epsilon) \cdot W_{\text{target}}^E - \epsilon \cdot W_{\text{target}}^E \cdot [r - \rho_0], \quad (10)$$

where ρ_0 is a target firing rate and ϵ is a normalization rate.

D. Analysis of learned connectivity and activity

We are trying to achieve the emergence of stable orientation selectivity in the postsynaptic neuron. To measure the extent to which this objective has been achieved in a simulation, we use the following metrics:

1. Weight Instability

The learned weights need to stabilize after a period of learning and remain more or less constant in order to produce consistent orientation selectivity in the postsynaptic neuron.

We measure this by an instability metric, defined as the standard deviation of the weights after a time point t_0 at which the plasticity would be expected to have converged:

$$S = \sum_k \text{std}(W_k^E(t))_{t > t_0} + \sum_k \text{std}(W_k^I(t))_{t > t_0} \quad (11)$$

where $W_k^E(t)$ and $W_k^I(t)$ the vectors of the development over time of the k -th excitatory and inhibitory weight respectively.

2. E/I Weight co-tuning

In addition to have a stable feedforward connectivity, we additionally expect a co-tuning of the E and I connectivity. Specifically, the inhibitory connections should match the excitatory connections coming from groups with similar orientation preference.

We quantify this with a weight balance metric, which is measured as the Pearson correlation coefficient of the average E and I learned weights for each input group:

$$B = \frac{\text{cov}(\bar{W}^E, \bar{W}^I)}{\text{std}(\bar{W}^E) \cdot \text{std}(\bar{W}^I)} \quad (12)$$

where \bar{W}^E and \bar{W}^I are the average learned weights for each E and I input group (i.e. population of neurons with a given orientation selectivity).

3. Orientation Tuning

We measure the postsynaptic neuron’s orientation preference by fitting the average responses to different stimuli to a Gaussian with periodic boundaries, given by:

$$R(x) = A \cdot \exp\left[\frac{(x - \mu)^2}{2\sigma^2}\right] + R_0, \quad (13)$$

where A is the amplitude, μ is the mean indicating the preferred orientation, σ is the standard deviation,

which we use as a proxy for the tuning width and R_0 is a baseline firing rate.

The parameters (A, μ, σ, R_0) are fitted to the neuron’s average responses during the presentation of each stimulus and the R^2 of the fit is used to estimate how well the neuron’s response function resembles a Gaussian.

E. Simulation-based inference

The SBI process for estimating the posteriors described in sections [IIB](#) and [IIC](#) was performed using the [sbi](#) package from [\[61\]](#).

V. ACKNOWLEDGEMENTS

This work was supported by a Sofja Kovalevskaja Award from the Alexander von Humboldt Foundation. EG thanks the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support. We acknowledge the support from the BMBF through the Tübingen AI Center (FKZ: 01IS18039B). AL is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 39072764.

-
- [1] Ruth Benavides-Piccione, Lidia Blazquez-Llorca, Asta Kastanauskaitė, Isabel Fernaud-Espinosa, Silvia Tapia-González, and Javier DeFelipe. Key morphological features of human pyramidal neurons. *Cerebral Cortex*, 34(5):bhae180, 2024.
- [2] Karlijn I Van Aerde and Dirk Feldmeyer. Morphological and physiological characterization of pyramidal neuron subtypes in rat medial prefrontal cortex. *Cerebral cortex*, 25(3):788–805, 2015.
- [3] Giancarlo La Camera, Alexander Rauch, David Thurbon, Hans-R Luscher, Walter Senn, and Stefano Fusi. Multiple time scales of temporal response in pyramidal and fast spiking cortical neurons. *Journal of neurophysiology*, 96(6):3448–3464, 2006.
- [4] Xiao-Jing Wang. Calcium coding and adaptive temporal computation in cortical pyramidal neurons. *Journal of Neurophysiology*, 79(3):1549–1566, 1998.
- [5] Christof Koch and Idan Segev. The role of single neurons in information processing. *Nature neuroscience*, 3(11):1171–1177, 2000.
- [6] Panayiota Poirazi, Terrence Brannon, and Bartlett W Mel. Pyramidal neuron as two-layer neural network. *Neuron*, 37(6):989–999, 2003.
- [7] A Polsky, B Mel, and J Schiller. Computational subunits in thin dendrites of pyramidal cells. *Nature Neuroscience*, 7(6):621–627, 2004.
- [8] Nelson Spruston. Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neuroscience*, 9(3):206–221, 2008.
- [9] Fritjof Helmchen, Karel Svoboda, Winfried Denk, and David W Tank. In vivo dendritic calcium dynamics in deep-layer cortical pyramidal neurons. *Nature neuroscience*, 2(11):989–996, 1999.
- [10] Matthew E Larkum, Jiameng Wu, Sarah A Duverdin, and Albert Gidon. The guide to dendritic spikes of the mammalian cortex in vitro and in vivo. *Neuroscience*, 489:15–33, 2022.
- [11] Joseph Bastian and Jerry Nguyenkim. Dendritic modulation of burst-like firing in sensory neurons. *Journal of neurophysiology*, 85(1):10–22, 2001.
- [12] Leonardo L Gollo, Osame Kinouchi, and Mauro Copelli. Active dendrites enhance neuronal dynamic range. *PLoS computational biology*, 5(6):e1000402, 2009.
- [13] Spencer L Smith, Ikuko T Smith, Tiago Branco, and Michael Häusser. Dendritic spikes enhance stimulus selectivity in cortical neurons in vivo. *Nature*, 503(7474):115–120, 2013.
- [14] Guy Eyal, Huibert D Mansvelder, Christiaan PJ de Kock, and Idan Segev. Dendrites impact the encoding capabilities of the axon. *Journal of Neuroscience*, 34(24):8063–8071, 2014.
- [15] Tiago Branco, Beverley A Clark, and Michael Häusser. Dendritic discrimination of temporal input sequences in cortical neurons. *Science*, 329(5999):1671–1675, 2010.
- [16] Alexandre Payeur, Jean-Claude Béique, and Richard Naud. Classes of dendritic information processing. *Current opinion in neurobiology*, 58:78–85, 2019.
- [17] Nace L Golding, Nathan P Staff, and Nelson Spruston. Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature*, 418(6895):326–331, 2002.
- [18] Björn M Kampa, Johannes J Letzkus, and Greg J Stuart. Dendritic mechanisms controlling spike-timing-dependent synaptic plasticity. *Trends in neurosciences*, 30(9):456–463, 2007.

- [19] Panayiota Poirazi and Bartlett W Mel. Impact of active dendrites and structural plasticity on the memory capacity of neural tissue. *Neuron*, 29(3):779–796, 2001.
- [20] Robert Legenstein and Wolfgang Maass. Branch-specific plasticity enables self-organization of nonlinear computation in single neurons. *Journal of Neuroscience*, 31(30):10787–10802, 2011.
- [21] Ruth Benavides-Piccione, Farid Hamzei-Sichani, Inmaculada Ballesteros-Yanez, Javier DeFelipe, and Rafael Yuste. Dendritic size of pyramidal neurons differs among mouse cortical regions. *Cerebral cortex*, 16(7):990–1001, 2006.
- [22] Serena Bianchi, Cheryl D Stimpson, Amy L Bauernfeind, Steven J Schapiro, Wallace B Baze, Mark J McArthur, Ellen Bronson, William D Hopkins, Katerina Semendeferi, Bob Jacobs, et al. Dendritic morphology of pyramidal neurons in the chimpanzee neocortex: regional specializations and comparison to humans. *Cerebral cortex*, 23(10):2429–2436, 2013.
- [23] Ruth Benavides-Piccione, Inmaculada Ballesteros-Yáñez, Javier DeFelipe, and Rafael Yuste. Cortical area and species differences in dendritic spine morphology. *Journal of neurocytology*, 31:337–346, 2002.
- [24] Edward M Callaway. Feedforward, feedback and inhibitory connections in primate visual cortex. *Neural Networks*, 17(5-6):625–632, 2004.
- [25] Mehmet Fişek, Dustin Herrmann, Alexander Egea-Weiss, Matilda Cloves, Lisa Bauer, Tai-Ying Lee, Lloyd E Russell, and Michael Häusser. Cortico-cortical feedback engages active dendrites in visual cortex. *Nature*, 617(7962):769–776, 2023.
- [26] Bartlett W Mel. Synaptic integration in an excitable dendritic tree. *Journal of neurophysiology*, 70(3):1086–1101, 1993.
- [27] Jeffrey C Magee. Dendritic integration of excitatory synaptic input. *Nature Reviews Neuroscience*, 1(3):181–190, 2000.
- [28] Ning-long Xu, Mark T Harnett, Stephen R Williams, Daniel Huber, Daniel H O’Connor, Karel Svoboda, and Jeffrey C Magee. Nonlinear dendritic integration of sensory and motor input during an active sensing task. *Nature*, 492(7428):247–251, 2012.
- [29] Christoph Schmidt-Hieber, Gabija Toleikyte, Laurence Aitchison, Arnd Roth, Beverley A Clark, Tiago Branco, and Michael Häusser. Active dendritic integration as a mechanism for robust and precise grid cell firing. *Nature neuroscience*, 20(8):1114–1121, 2017.
- [30] Alexandra Tran-Van-Minh, Romain D Cazé, Thérèse Abrahamsson, Laurence Cathala, Boris S Gutkin, and David A DiGregorio. Contribution of sublinear and supralinear dendritic integration to neuronal computations. *Frontiers in cellular neuroscience*, 9:67, 2015.
- [31] Attila Losonczy, Judit K Makara, and Jeffrey C Magee. Compartmentalized dendritic plasticity and input feature storage in neurons. *Nature*, 452(7186):436–441, 2008.
- [32] Balazs B Ujfalussy, Judit K Makara, Tiago Branco, and Mate Lengyel. Dendritic nonlinearities are tuned for efficient spike-based computations in cortical circuits. *Elife*, 4:e10056, 2015.
- [33] Guy Major, Matthew E Larkum, and Jackie Schiller. Active properties of neocortical pyramidal neuron dendrites. *Annual review of neuroscience*, 36(1):1–24, 2013.
- [34] Andreas Frick, Jeffrey Magee, and Daniel Johnston. Ltp is accompanied by an enhanced local excitability of pyramidal neuron dendrites. *Nature neuroscience*, 7(2):126–135, 2004.
- [35] Matthew E Larkum and Thomas Nevian. Synaptic clustering by dendritic signalling mechanisms. *Current opinion in neurobiology*, 18(3):321–331, 2008.
- [36] Anthony Zador, Brenda Claiborne, and Thomas Brown. Nonlinear pattern separation in single hippocampal neurons with active dendritic membrane. *Advances in neural information processing systems*, 4, 1991.
- [37] Daniel E Wilson, David E Whitney, Benjamin Scholl, and David Fitzpatrick. Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex. *Nature neuroscience*, 19(8):1003–1009, 2016.
- [38] M Florencia Iacaruso, Ioana T Gasler, and Sonja B Hofer. Synaptic organization of visual space in primary visual cortex. *Nature*, 547(7664):449–452, 2017.
- [39] Simon Weiler, Drago Guggiana Nilo, Tobias Bonhoeffer, Mark Hübener, Tobias Rose, and Volker Scheuss. Orientation and direction tuning align with dendritic morphology and spatial connectivity in mouse visual cortex. *Current Biology*, 32(8):1743–1753, 2022.
- [40] Larry F Abbott and Sacha B Nelson. Synaptic plasticity: taming the beast. *Nature neuroscience*, 3(11):1178–1183, 2000.
- [41] Chr Von der Malsburg. Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik*, 14(2):85–100, 1973.
- [42] Kenneth D Miller, Joseph B Keller, and Michael P Stryker. Ocular dominance column development: analysis and simulation. *Science*, 245(4918):605–615, 1989.
- [43] Elie L Bienenstock, Leon N Cooper, and Paul W Munro. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, 2(1):32–48, 1982.
- [44] Nathan Intrator and Leon N Cooper. Objective function formulation of the bcm theory of visual cortical plasticity: Statistical connections, stability conditions. *Neural Networks*, 5(1):3–17, 1992.
- [45] C Charles Law and Leon N Cooper. Formation of receptive fields in realistic visual environments according to the bienenstock, cooper, and munro (bcm) theory. *Proceedings of the National Academy of Sciences*, 91(16):7797–7801, 1994.
- [46] Tim P Vogels, Henning Sprekeler, Friedemann Zenke, Claudia Clopath, and Wulfram Gerstner. Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science*, 334(6062):1569–1573, 2011.
- [47] Claudia Clopath, Tim P Vogels, Robert C Froemke, and Henning Sprekeler. Receptive field formation by interacting excitatory and inhibitory synaptic plasticity. *BioRxiv*, page 066589, 2016.
- [48] Everton J Agnes and Tim P Vogels. Co-dependent excitatory and inhibitory plasticity accounts for quick, stable and long-lasting memories in biological networks. *Nature Neuroscience*, 27(5):964–974, 2024.
- [49] Samuel Eckmann, Edward James Young, and Julijana Gjorgjieva. Synapse-type-specific competitive hebbian learning forms functional recurrent networks. *Proceedings of the National Academy of Sciences*, 121(25):e2305326121, 2024.
- [50] Stefan Remy and Nelson Spruston. Dendritic spikes induce single-burst long-term potentiation. *Proceedings of the National Academy of Sciences*, 104(43):17192–17197,

- 2007.
- [51] Jeffrey C Magee and Daniel Johnston. A synaptically controlled, associative signal for hebbian plasticity in hippocampal neurons. *Science*, 275(5297):209–213, 1997.
- [52] Arvind Govindarajan, Raymond J Kelleher, and Susumu Tonegawa. A clustered plasticity model of long-term memory engrams. *Nature Reviews Neuroscience*, 7(7):575–583, 2006.
- [53] George Kastellakis, Denise J Cai, Sara C Mednick, Alcino J Silva, and Panayiota Poirazi. Synaptic clustering within dendrites: an emerging theory of memory formation. *Progress in neurobiology*, 126:19–35, 2015.
- [54] George Kastellakis and Panayiota Poirazi. Synaptic clustering and memory formation. *Frontiers in molecular neuroscience*, 12:300, 2019.
- [55] Christian Ebner, Claudia Clopath, Peter Jedlicka, and Hermann Cuntz. Unifying long-term plasticity rules for excitatory synapses by modeling dendrites of cortical pyramidal neurons. *Cell reports*, 29(13):4295–4307, 2019.
- [56] P Jesper Sjöstrom, Ede A Rancz, Arnd Roth, and Michael Häusser. Dendritic excitability and synaptic plasticity. *Physiological reviews*, 88(2):769–840, 2008.
- [57] Gina G Turrigiano and Sacha B Nelson. Homeostatic plasticity in the developing nervous system. *Nature reviews neuroscience*, 5(2):97–107, 2004.
- [58] Emmanouil Giannakakis, Oleg Vinogradov, Victor Buendía, and Anna Levina. Structural influences on synaptic plasticity: The role of presynaptic connectivity in the emergence of e/i co-tuning. *PLOS Computational Biology*, 20(10):e1012510, 2024.
- [59] Kazuhiro Sohya, Katsuro Kameyama, Yuchio Yanagawa, Kunihiko Obata, and Tadaharu Tsumoto. Gabaergic neurons are less selective to stimulus orientation than excitatory neurons in layer ii/iii of visual cortex, as revealed by in vivo functional ca2+ imaging in transgenic mice. *Journal of Neuroscience*, 27(8):2145–2149, 2007.
- [60] Wen-pei Ma, Bao-hua Liu, Ya-tang Li, Z Josh Huang, Li I Zhang, and Huizhong W Tao. Visual representations by cortical somatostatin inhibitory neurons—selective but with weak and delayed responses. *Journal of Neuroscience*, 30(43):14371–14379, 2010.
- [61] Alvaro Tejero-Cantero, Jan Boelts, Michael Deistler, Jan-Matthis Lueckmann, Conor Durkan, Pedro J. Gonçalves, David S. Greenberg, and Jakob H. Macke. sbi: A toolkit for simulation-based inference. *Journal of Open Source Software*, 5(52):2505, August 2020.
- [62] Basile Confavreux, Poornima Ramesh, Pedro J Gonçalves, Jakob H Macke, and Tim Vogels. Meta-learning families of plasticity rules in recurrent spiking networks using simulation-based inference. *Advances in Neural Information Processing Systems*, 36:13545–13558, 2023.
- [63] D Chklovskii. Synaptic connectivity and neuronal morphology: Two sides of the same coin. *Neuron*, 43(5):609–617, 2004.
- [64] M London and M Häusser. Dendritic computation. *Annu. Rev. Neurosci.*, 28:503–532, 2005.
- [65] Judit K. Makara, Attila Losonczy, Quan Wen, and Jeffrey C. Magee. Experience-dependent compartmentalized dendritic plasticity in rat hippocampal CA1 pyramidal neurons. *Nat Neurosci*, 12(12):1485–1487, December 2009.
- [66] Jianhua Cang and David A. Feldheim. Developmental mechanisms of topographic map formation and alignment. *Annual Review of Neuroscience*, 36(Volume 36, 2013):51–77, July 2013.
- [67] Jan H Kirchner and Julijana Gjorgjieva. Emergence of local and global synaptic organization on cortical dendrites. *Nature Communications*, 12(1):4005, 2021.
- [68] Zsuzsanna Varga, Hongbo Jia, Bert Sakmann, and Arthur Konnerth. Dendritic coding of multiple sensory inputs in single cortical neurons in vivo. *Proceedings of the National Academy of Sciences*, 108(37):15420–15425, September 2011.
- [69] Roman Makarov, Michalis Pagkalos, and Panayiota Poirazi. Dendrites and efficiency: Optimizing performance and resource utilization. *Current Opinion in Neurobiology*, 83:102812, December 2023.
- [70] Ali Choucry, Masanori Nomoto, and Kaoru Inokuchi. Engram mechanisms of memory linking and identity. *Natl Rev Neurosci*, 25(6):375–392, 2024.
- [71] Wulfram Gerstner. Hebbian learning and plasticity. *From neuron to cognition via computational neuroscience*, pages 0–25, 2011.

Appendix D Structural influences on synaptic plasticity: The role of presynaptic connectivity in the emergence of E/I co-tuning

Contains the full text of the publication: Giannakakis, E., Vinogradov, O., Buendía, V., and Levina, A. (2024). Structural influences on synaptic plasticity: The role of presynaptic connectivity in the emergence of e/i co-tuning. *PLOS Computational Biology*, 20(10), 1–22.


RESEARCH ARTICLE

Structural influences on synaptic plasticity: The role of presynaptic connectivity in the emergence of E/I co-tuning

Emmanouil Giannakakis^{1,2}, Oleg Vinogradov^{1,2}, Victor Buendía^{1,2}, Anna Levina^{1,2*}

1 Department of Computer Science, University of Tübingen, Tübingen, Germany, **2** Max Planck Institute for Biological Cybernetics, Tübingen, Germany

* anna.levina@uni-tuebingen.de



OPEN ACCESS

Citation: Giannakakis E, Vinogradov O, Buendía V, Levina A (2024) Structural influences on synaptic plasticity: The role of presynaptic connectivity in the emergence of E/I co-tuning. *PLoS Comput Biol* 20(10): e1012510. <https://doi.org/10.1371/journal.pcbi.1012510>

Editor: Julijana Gjorgjieva, Technical University of Munich: Technische Universität München, GERMANY

Received: December 12, 2023

Accepted: September 25, 2024

Published: October 31, 2024

Copyright: © 2024 Giannakakis et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Code for all our simulations is available at the GitHub repository: <https://github.com/MGiannakakis/NetworkAssembliesandSTDP/>.

Funding: AL, EG, VB, OV received funding from the Alexander von Humboldt Foundation via Sofja Kovalevskaja Award (<https://www.humboldt-foundation.de/en/>). MG received funding from the BMBF (<https://www.bmbf.de/>) through the Tübingen AI Center (FKZ: 01IS18039A). The

Abstract

Cortical neurons are versatile and efficient coding units that develop strong preferences for specific stimulus characteristics. The sharpness of tuning and coding efficiency is hypothesized to be controlled by delicately balanced excitation and inhibition. These observations suggest a need for detailed co-tuning of excitatory and inhibitory populations. Theoretical studies have demonstrated that a combination of plasticity rules can lead to the emergence of excitation/inhibition (E/I) co-tuning in neurons driven by independent, low-noise signals. However, cortical signals are typically noisy and originate from highly recurrent networks, generating correlations in the inputs. This raises questions about the ability of plasticity mechanisms to self-organize co-tuned connectivity in neurons receiving noisy, correlated inputs. Here, we study the emergence of input selectivity and weight co-tuning in a neuron receiving input from a recurrent network via plastic feedforward connections. We demonstrate that while strong noise levels destroy the emergence of co-tuning in the readout neuron, introducing specific structures in the non-plastic pre-synaptic connectivity can re-establish it by generating a favourable correlation structure in the population activity. We further show that structured recurrent connectivity can impact the statistics in fully plastic recurrent networks, driving the formation of co-tuning in neurons that do not receive direct input from other areas. Our findings indicate that the network dynamics created by simple, biologically plausible structural connectivity patterns can enhance the ability of synaptic plasticity to learn input-output relationships in higher brain areas.

Author summary

Many studies of learning in biological neural networks have focused on how plausible plasticity rules shape individual connections between neurons in a recurrent network or in feed-forward projections. However, in the latter case, the presynaptic network properties, such as clustered connectivity, strongly influence population dynamics and, thus, the learning process of the projections. Here, we aim to close this gap by showing how non-plastic network structure can strongly influence the outcomes of synaptic learning. We

funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

show that unstructured recurrent connectivity and the presence of noise can significantly reduce the ability of synaptic plasticity to separate presynaptic input populations and coordinate excitatory and inhibitory connection development, while the introduction of overlapping clustered structures in fixed or plastic recurrent connectivity can boost synaptic learning. Using Bayesian inference, we identify the optimal connectivity structures for a recurrent network and demonstrate the strong effects that relatively simple connectivity patterns can have on the ability of a network to learn via local plasticity.

Introduction

Stimulus selectivity, the ability of neurons to respond differently to distinct stimuli, is one of the primary mechanisms for encoding information in the nervous system. This selectivity can range from simple orientation selectivity in lower sensory areas [1–3] to more complex spatio-temporal pattern selectivity in higher areas [4–6]. Such selectivity is shown to be self-organized under the influence of structured input, enabling, for example, the emergence of visual orientation preference in non-visual sensory areas upon rewiring [7] or changing the whiskers representation in the barrel cortex of rats depending on the level of sensory input [8]. The mechanisms underlying the emergence of input selectivity have been the subject of extensive investigation, both through experimental and computational modelling studies, and still remain under active discussion [9–13].

Despite initially attributing stimulus-selectivity to excitatory neurons and their network structure, we now know that inhibitory neurons are also tuned to stimuli, and the coordination of the E/I currents is a central component of efficient neural computation [14, 15]. In particular, it has been shown that excitatory and inhibitory inputs are often correlated [16], with preferred stimuli eliciting stronger excitatory and, with the small delay, stronger inhibitory responses compared to the non-preferred stimuli [17, 18]. This co-tuning of excitation and inhibition is theorized to be beneficial for a variety of computations such as gain control [19, 20], visual surround suppression [21, 22], novelty detection [23] and optimal spike-timing [15, 24].

Although it is still unclear how E/I co-tuning emerges, the dominant view is that it arises via the interaction of several synaptic plasticity mechanisms [25], a hypothesis that has been reinforced by the findings of multiple theoretical studies over the last decade. First, it has been demonstrated that different inhibitory plasticity rules can match static excitatory connectivity [26–29]. More recently, it was also shown that various combinations of plasticity and diverse normalisation mechanisms allow for the simultaneous development of matching excitatory and inhibitory connectivity in feedforward settings [12, 30–32]. Moreover, a variety of plasticity mechanisms has been associated with the formation of stable assemblies [33–37], the creation of E/I balance [38] and the emergence of tuning selectivity [39] in recurrent networks.

However, so far, most of the theoretical studies of synaptic plasticity have focused on identifying the optimal parameters of individual learning rules and normalization mechanisms for specific tasks and ignore the ways in which these mechanisms act within complex network structures that may influence their function. Specifically, biological networks are characterized by highly non-trivial connectivity structures that are known to display varied degrees of clustering and neuron type-specific connectivity patterns [40]. Such network structures give rise to distinct neural dynamics; for example, clustering in recurrent network structure introduces correlations in the activity of similarly tuned neurons [41] and complex interaction between subpopulations of neurons [42]. These types of dynamics fundamentally alter the statistics of

population activity that most synaptic plasticity mechanisms rely on for modifying synaptic strength.

Here, we investigate how the development of matching E/I input selectivity in a downstream neuron via synaptic plasticity can be driven by the structure of recurrent connectivity in the input network. We combine excitatory and inhibitory plasticity rules [26, 31, 43] in the feedforward connections of spiking network to develop detailed co-tuning of excitatory and inhibitory connectivity, and we demonstrate that the ability of these plasticity mechanisms to create co-tuning is significantly reduced in the presence of noise and (non-plastic) random recurrent connectivity between the input neurons. We further show that the effects of recurrence and noise on the population activity that drives the formation of matching E/I feedforward weights on a downstream neuron can be fully ameliorated by the introduction of synapse-type specific assemblies of neurons, characterized by local excitation and relatively homogeneous inhibition, an often-observed pattern of cortical connectivity [44, 45]. Our findings demonstrate that network structure can, by influencing population dynamics, significantly modulate the capacity of synaptic plasticity to generate input selectivity in downstream neurons. This highlights a synergistic interaction between structural connectivity and learning mechanisms that can enhance the computational capabilities of brain networks.

Results

We begin by reproducing previously reported results, validating the emergence of co-tuning in a plastic feedforward network, and introducing measures to capture weight diversity and co-tuning for subsequent analysis. Next, we show that the introduction of strong noise or a random static recurrent connectivity in the presynaptic networks impairs the development of co-tuning by destroying the correlation structure in the activity of the presynaptic population. We then illustrate how specific structures in the static recurrent connectivity can restore the ability of plastic synapses to generate co-tuning. Using analytical results from a reduced linear neural mass model and Bayesian inference for the full network, we identify the optimal connectivity structures in static networks, demonstrating that optimal connectivity is influenced by network sparsity. Finally, we simulate fully plastic networks, confirming that our key observations hold true.

Co-tuning and its self-organization by synaptic plasticity in a low-noise feedforward setting

We simulate a single postsynaptic read-out unit driven by a population of $N = 1000$ neurons. The pre-synaptic population is divided into M groups G_i , $i \in \{1, \dots, M\}$. Each group is comprised of $n = N/M$ neurons, of which 80% are excitatory and 20% are inhibitory. These neurons are driven by an identical, group-specific Poisson spike train—a shared external input. Additionally, each neuron receives low-intensity independent external noise [26] that prevents unrealistic total synchrony between the input neurons. This setting, depicted in (Fig 1a), leads to correlated firing among neurons of the same input group (Fig 1b), which is a common setting for studying the effect of different plasticity rules [12, 26, 31].

Input selectivity develops when the post-synaptic neuron responds differently to inputs from different groups (by adjusting its firing rate). This happens when the average excitatory feedforward projections are sufficiently diverse between pre-synaptic groups (Fig 1d and 1f), which leads to groups with stronger feedforward connections eliciting stronger post-synaptic responses upon activation. Moreover, connections from neurons with highly correlated firing (i.e., from the same group) should have a similar strength. To quantify this feature of the

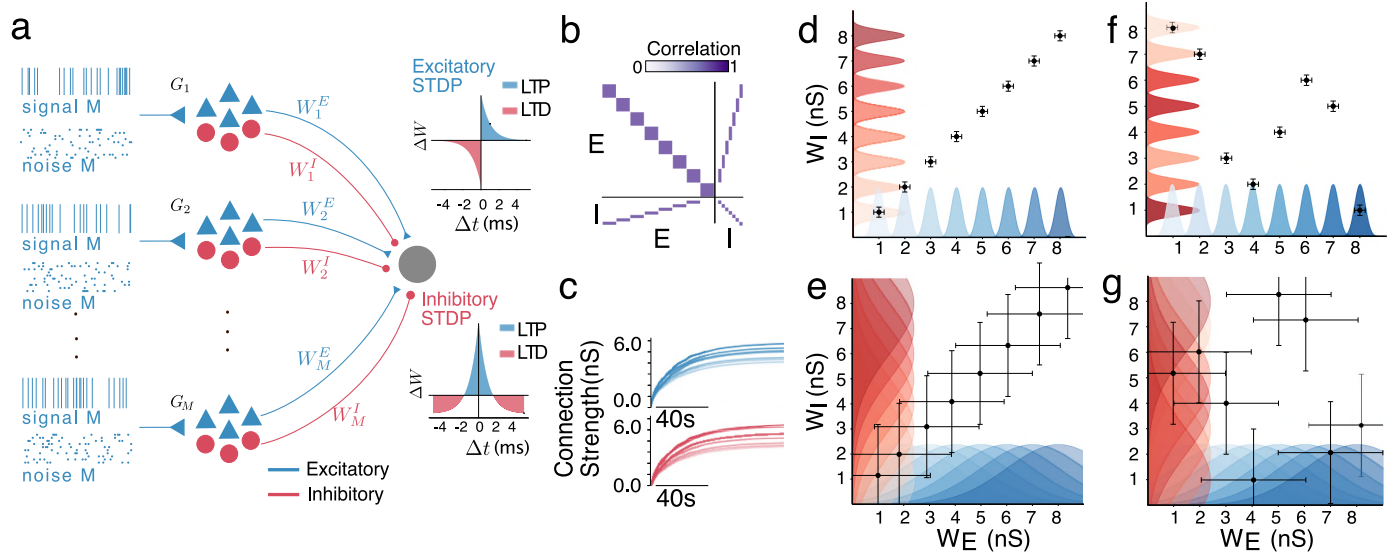


Fig 1. Emergence of co-tuning in a feedforward network. **a.** A diagram of the feedforward network with plastic connections from the different inputs group to the readout neuron. **b.** The correlation matrix of the network’s activity. In the absence of noise, neurons of the same input group are highly correlated. **c.** The development of average E and I weights in an ideal feedforward network with very low noise leads to co-tuned and diverse feedforward connectivity. **d.** An illustration of feedforward connectivity that exhibits both co-tuning and diversity. Different input groups are clearly distinct, and the E and I weights for each group are correlated. The colours of the distribution indicate different groups, and the shading (light to dark) is matched for the corresponding E and I populations; the points and error bars indicate the mean and std of the E and I connectivity of each group. **e.** A co-tuned but not diverse connectivity. E/I weight correlation is maintained, but there is hardly any distinction between groups. **f.** A diverse connectivity without E/I weight co-tuning. While each group is distinct from the others, there is no coordination of the E and I connections from the same group. **g.** In the absence of weights co-tuning and diversity, the feedforward connectivity lacks any discernible pattern.

<https://doi.org/10.1371/journal.pcbi.1012510.g001>

network, we define a diversity metric,

$$D = 1 - \frac{1}{M \cdot \text{Std}(W^E)} \sum_{i=1}^M \text{Std}(W_{G_i}^E), \tag{1}$$

where $W^E \in \mathbb{R}^N$ is the set of excitatory feedforward connection weights, $W_{G_i}^E \in \mathbb{R}^{N/M}$ is the subset of excitatory feedforward connection weights from input group i and $\text{Std}(\cdot)$ denotes the standard deviation. Diversity $D \in [0, 1]$ equals unity when the feedforward connections from the same group are the same but differ across groups; D is close to zero when the feedforward connections from each group follow the same distribution and different groups cannot be meaningfully distinguished (Fig 1e and 1g).

Brain networks are characterized by a balance of excitatory and inhibitory currents [18, 46, 47]. For a neuron to be balanced, the average inhibitory current must be equal to the average excitatory current during a (relatively short—usually a few ms) time window. Depending on the temporal resolution of this canceling out, the balance can be more “loose” or “tight”, with detailed (“tight”) balance associated with efficient coding [14, 15] and the ability to encode for multiple stimuli [19] (Further discussion in S1 Text Section A). In our specific setting, due to the highly correlated firing in each group, detailed balance can be achieved by matching the relative strengths of the excitatory and inhibitory weights from each group.

To quantify the E/I weight co-tuning, which generates the detailed balance in our simplified network, we use the Pearson correlation coefficient between the mean excitatory and

inhibitory weights of each group,

$$CT_W = \frac{\text{Cov}(\langle W_G^E \rangle, \langle W_G^I \rangle)}{\text{Std}(\langle W_G^E \rangle) \cdot \text{Std}(\langle W_G^I \rangle)}, \quad (2)$$

where $\langle W_G^A \rangle = (\langle W_{G_1}^A \rangle, \langle W_{G_2}^A \rangle, \dots, \langle W_{G_m}^A \rangle)$, $A \in \{I, E\}$ and $\langle W_{G_i}^A \rangle$ is the average projection weight from the excitatory ($A = E$) or inhibitory ($A = I$) neurons in group i , and $\text{Cov}(x_1, x_2)$ denotes the covariance of the variables x_1 and x_2 . In networks with strong weight co-tuning CT_W , the strength of incoming E and I currents is highly correlated (Fig 1d).

We verify that high diversity ($D \approx 1$) and weight co-tuning ($CT_W \approx 1$) can organically emerge via a combination of plasticity mechanisms in the feedforward connections whose trajectories are initialized at the same (small) value. Specifically, the excitatory connections follow the triplet Hebbian STDP rule [48], and the inhibitory connections follow a symmetric rule [26]. We additionally use competitive normalization in both the inhibitory and excitatory connections, which amplifies small transient differences between the firing rates of different input groups and leads to the development of input selectivity [31]. Since neurons in different groups are independent, while neurons in the same group share a strong correlation (Fig 1b), the plasticity protocol generates very strongly correlated E/I weights and strong input selectivity, as shown in (Fig 1c). (For more details on the network's firing activity, see S1 Text Section C).

Postsynaptic weights with high diversity can be well separated, while those with co-tuned weights produce a match between E and I incoming connections. (Fig 1d) illustrates four possible situations. Observe that it is possible to have weights that are, e.g., diverse (so they can be distinguished) but not co-tuned (so E-I are not correlated) and vice versa: E-I weights are correlated, but weights cannot be separated.

Noise and recurrent connectivity compromise the ability of STDP to produce E/I co-tuning

Strictly feed-forward networks with relatively low noise levels are unrealistic approximations of complex cortical circuits (which are characterized by noisy inputs and complex recurrent connectivity), and thus their dynamics might deviate significantly from those observed in experiments. Thus, we introduce noise and non-plastic recurrent connectivity in our pre-synaptic network, both ubiquitously present in biological networks [49, 50]. First, we investigate how they individually affect the emerging E/I co-tuning by changing the structure of the correlations between the neurons of the input network. Then, we examine ways in which different connectivity structures can ameliorate these effects.

We vary the level of noise by changing the fraction of input spikes that are specific to each neuron (noise) vs the shared (signal) input (Fig 2a). This allows us to control the signal-to-noise ratio while keeping the average number of incoming spikes the same. As the noise intensity increases, the cross-correlations within each input group decrease, while the cross-correlations between neurons of different input groups remain very low, (Fig 2b). The effect of this in-group decorrelation is an increased variability in the learned projections to the postsynaptic neuron from neurons of the same input group and, thus, a decrease in the resulting diversity (Fig 2). At the same time, this decorrelation has a much weaker effect on the ability of the plasticity to match E and I feedforward weights from the same groups. This is reflected in the slower reduction of the E/I weight co-tuning, which visibly declines only once the noise becomes overwhelmingly stronger than the input (more than 80% incoming spikes are not shared between neurons of the same group, Fig 2c). Raster plots illustrating the dynamics for different values of noise are shown in (Fig 2d).

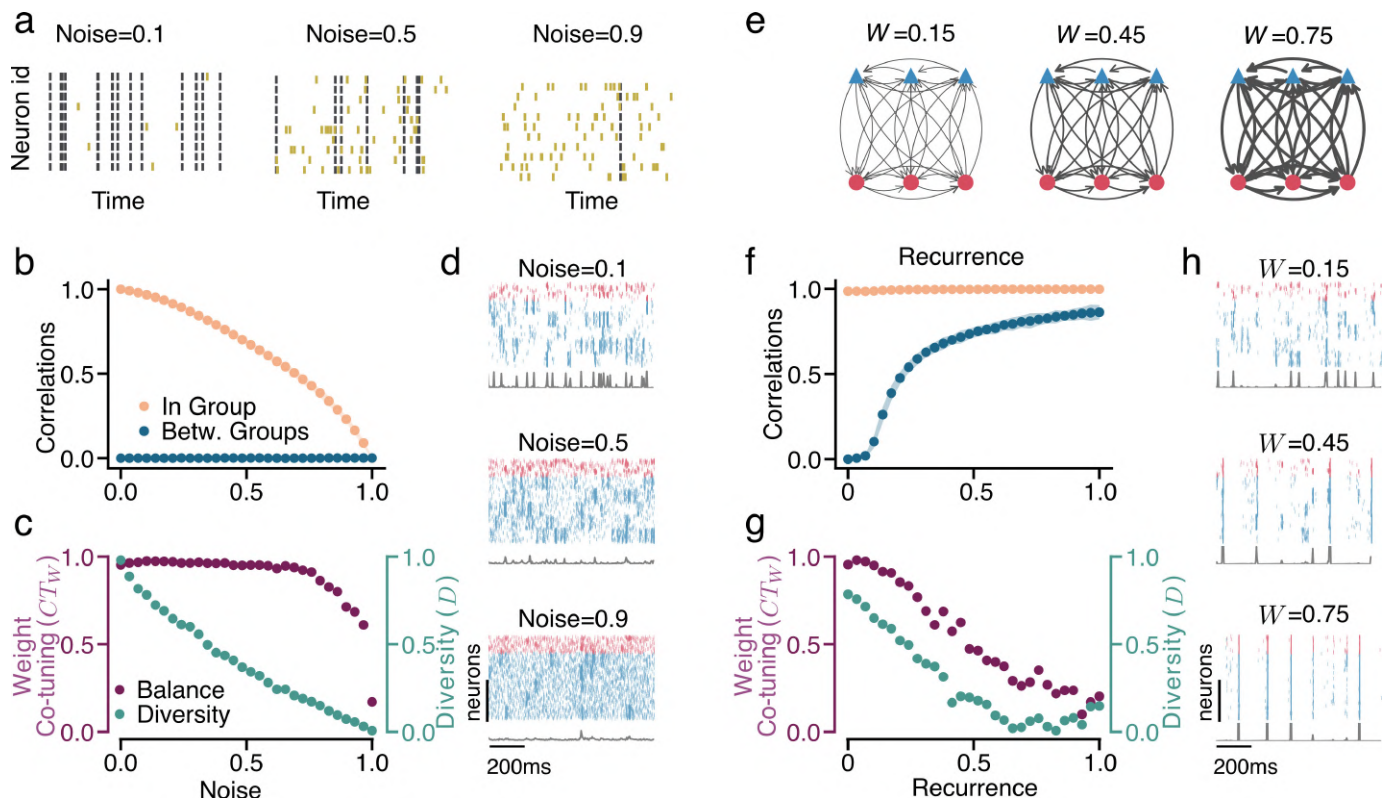


Fig 2. Noise and Recurrence Destroy E/I Co-Tuning. **a.** An illustration of increasing levels of noise in a single input group. In low noise settings, all neurons of the same groups fire at the same time, while as the noise level increases, each neuron fires more and more individual spikes. Joint inputs are shown as black, and individual noise spikes are yellow. **b.** The increase in noise leads to a decrease in in-group correlation (orange), while the between-group correlation (blue) remains low. **c.** An increase in the input noise leads to a reduction in diversity (teal) and, for larger noise intensities, also co-tuning of E/I weights (purple). **d.** Raster plots of input populations' activity (red, inhibitory neurons; blue, excitatory ones; gray, corresponding firing rate). As the noise increases, the spiking in each group becomes more asynchronous. The traces below (gray line) show the spike count over all neurons in 2ms bins. **e.** An illustration of recurrent connectivity (for three input groups). The coupling parameter W controls the mean connection strength. As W increases (indicated by thicker connections in the diagram), more of the input a neuron receives comes from other neurons in the recurrently connected network than from the feedforward input. **f.** An increase in the recurrent coupling strength W leads to an increase in the between-group correlation (blue), while the in-group correlation (orange) remains high. **g.** The decrease in weight co-tuning and diversity with an increase in the coupling strength. **h.** The spiking activity becomes more synchronous across groups as the coupling strength increases. The traces below (gray line) show the spike count over all neurons in 2ms bins. Simulation parameters not indicated in the text can be found in Tables A, B and C in [S1 Text](#).

<https://doi.org/10.1371/journal.pcbi.1012510.g002>

Recurrent connectivity in the pre-synaptic network introduces cross-correlations between the neurons from different input groups, which compromises both diversity and E/I weight co-tuning. To test the extent of this effect, we connect the N presynaptic neurons (creating a non-plastic recurrent network) with connection probability p and use coupling strength W (which denotes the mean synaptic strength) as a control parameter. Initially, we only consider fully-connected networks ($p = 1$). By changing W , we can control the ratio between the input received from the feedforward connections (whose rate and connection strength are fixed) and the other neurons in the network via recurrent connectivity (Fig 2e). The recurrent connectivity increases cross-correlations between groups while maintaining the high correlation within each input group, (Fig 2f). The effect of these cross-correlations is stronger than the effect of the noise since they affect both the diversity and the weight co-tuning, both of which decline as the recurrent connections become stronger (Fig 2g). As with noise, raster plots for different recurrent connectivity strengths are shown in (Fig 2h).

The combination of noise and recurrent connectivity affects both in-group and between-group correlations (Fig 3c and 3d), resulting in a reduction of weight co-tuning and diversity

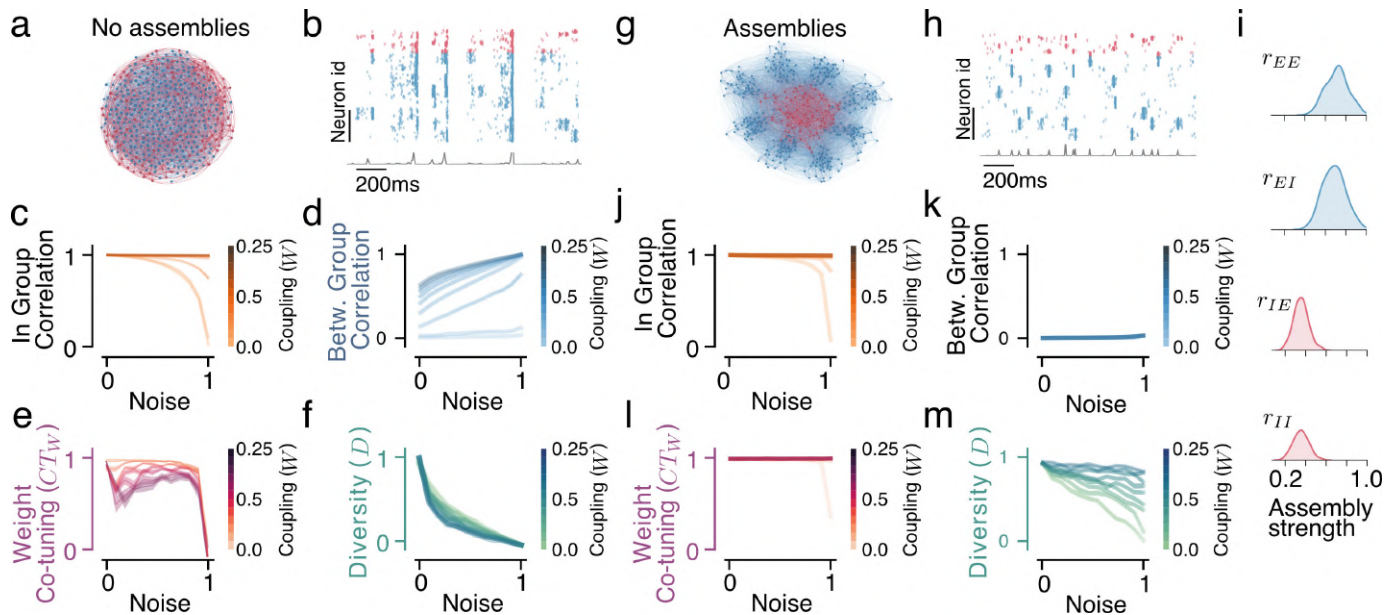


Fig 3. Optimized assemblies of neurons restore the co-tuning in recurrent noisy networks. **a.** Diagram of the network with uniform connectivity **b.** The network activity is characterized by synchronous events across groups. The traces below (gray line) show the spike count over all neurons in 2ms bins. The in-group (**c.**) and between-group (**d.**) correlations for different levels of noise and recurrent connection strength in the uniformly connected network lead to a reduction in the weight co-tuning (**e.**) and weight diversity (**f.**) metrics for different noise and recurrent strength combinations. **g.** Diagram of the network with optimal assembly structure **h.** The network activity becomes more decorrelated across groups. The traces below (gray line) show the spike count over all neurons in 2ms bins. **i.** Approximate posterior distributions of optimal excitatory and inhibitory assemblies strength. The in-group (**j.**) and between-groups (**k.**) correlations for different levels of noise and recurrent connection strength can be almost fully restored by the assembly structure, leading to a restoration of (**l.**) strong weight co-tuning and (**m.**) weight diversity.

<https://doi.org/10.1371/journal.pcbi.1012510.g003>

as the noise and recurrent connection strength increase, (Fig 3e and 3f). The effects of combined noise and unstructured recurrence are not only a sum of their independent effects, but they can lead to novel effects compounding the impact on the emergence of input selectivity. For example, for strong recurrence (Fig 3d), increasing noise levels seem to lead to an increase of between-group correlations (a counter-intuitive effect, given the decorrelating effects of noise). This is due to the absence of synchronous input to neurons of the same group (due to increased noise, i.e., reduced in-group correlation), which makes synchronisation across groups (and consequently the increase of the between-group correlation) due to recurrent input easier. More details on the effects of noise and recurrence on the network firing are discussed in S1 Text Section C and D).

We develop a formal description of the effect of noise and recurrence on the correlation structure in a simplified linear neural mass model. To this end, we consider $M = 8$ mesoscopic units instead of the previously studied M inter-connected groups, represented by continuous rate variables $x_j(t), j = 1, \dots, M$. These units evolve in time, subject to stochastic white noise. The linear approximation is justified for any system at a stationary state with a constant average firing rate, and it serves as a simplified model for a wide range of parameters of the spiking network (for details on the linear model and its relation to the spiking network, (see S1 Text Section F and H).

In this simplified case, it is possible to derive analytical equations for all the relevant in- and between-group covariances, which yield the correlation coefficients. These correlations are the solution to a linear system of equations, which can be obtained exactly using numerical methods. Furthermore, one can find close-form solutions in some simple scenarios. For example, in the case of a completely homogeneous network, where all coupling weights are the same,

correlation coefficients can be written explicitly (see [S1 Text](#) Section F and H). If the coupling strength increases $W \rightarrow +\infty$, all correlations grow to 1 as $1 - \mathcal{O}(1/W^2)$. On the contrary, if we reduce the noise $r \rightarrow 0_+$, the correlations will decrease to $MW/(M-1)^2 + \mathcal{O}(r^{-2})$. Both cases eliminate any possible differentiation between the groups, thus compromising the ability of the plasticity mechanisms to create high diversity $D \approx 1$. Another observation is that in the linear network, increasing noise affects the correlation coefficient quadratically, while coupling increases it linearly. Therefore, since $r < 1$, increasing the coupling has a larger impact on the co-tuning, a consequence that is recovered in the spiking network, consistent with the results shown in [\(Fig 2b\)](#) and [\(Fig 2e\)](#).

Neuronal type-specific assemblies restore the ability of STDP to produce co-tuning

The homogeneous all-to-all connectivity ([Fig 3a and 3b](#)) that we have examined so far is not a realistic assumption and could be particularly detrimental to the self-organization of co-tuning in higher areas. Thus, we examine the impact of different types of inhomogeneous connectivity. In particular, using the idea of functional assemblies (strongly connected neurons that are tuned to the same stimulus [[25](#)]), we study whether stronger recurrent connectivity between neurons of the same input group can introduce the necessary correlation structure in the population activity that will allow the plasticity to produce weight diversity and co-tuning.

We maintain the total recurrent input to a neuron constant (the fraction of input coming from the signal/noise vs. the other neurons in the recurrent network, excluding the input from the feedforward connections) while using the ratio of input coming from neurons of the same vs. other input groups as a control parameter. Since we want to vary this ratio independently for each connection type, we define a metric of assembly strength as:

$$r_{ab} = \frac{C_{in}^{ab}}{C_{in}^{ab} + C_{out}^{ab}} = \frac{W_{in}^{ab} \cdot \bar{r}}{W_{in}^{ab} \cdot \bar{r} + (M-1) \cdot W_{out}^{ab} \cdot \bar{r}} = \frac{W_{in}^{ab}}{W_{in}^{ab} + (M-1) \cdot W_{out}^{ab}}, \tag{3}$$

where C_{in}^{ab} is the total recurrent input a neuron of type b receives from neurons of type a , for $a, b \in \{E, I\}$, of its own input group and C_{out}^{ab} is the total recurrent input the neuron receives from other input groups, W_{in}^{ab} the connection strength between neurons of the same group, W_{out}^{ab} the connection strength between neurons of different groups and \bar{r} is the average firing rate of network neurons (we assume a uniform firing across the network, which can be simplified out of the equation).

We vary assembly strengths for each type of connection r_{EE} , r_{EI} , r_{IE} , and r_{II} while keeping the total recurrent input to a neuron $C_{in}^{ab} + C_{out}^{ab} = \frac{p \cdot N}{M} \cdot (W_{in}^{ab} + (M-1) \cdot W_{out}^{ab}) =: p \cdot N \cdot W$ constant. Here W is the average coupling strength, and p is the recurrent connection probability. As in the network without assemblies, for now, we consider fully connected networks ($p = 1$). Thus, we can vary the fraction of input coming to a neuron from its own input group without changing the total recurrent E, or I input it receives.

Since the structure of the feedforward connectivity (diversity and weight co-tuning) that the plasticity converges to is controlled by the correlation structure of the inputs, we can use the correlations as a proxy that is easier to optimize than the connectivity metrics. Specifically, we want to maximize the in-group and minimize the between-group correlations, and we seek the assembly structure that leads to that objective. In the reduced linear neural mass model, we compute the optimal assembly strengths (see [S1 Text](#) Section H) analytically and find that strong local excitation and dispersed inhibition restore the desired correlation structure in the network's activity (see [S1 Text](#) Section I). We find that for all combinations of noise and sufficiently strong recurrent connectivity, strong excitatory assemblies (high r_{EE} and r_{EI}) and

uniform inhibitory connectivity (low r_{IE} and r_{II}) allow the correlating excitatory currents to remain mostly within the input the input group/assembly and maintain high in-group correlation, while the diffused inhibitory currents reduce correlations between groups (for a more detailed discussion of this mechanism see [S1 Text Section J](#)). Still, since the reduced model does not account for many essential features of the spiking network, like sparsity of connections, in-group interactions between neurons of the same type, and non-stationary dynamical states of the groups, the analytic solution obtained for the linear neural mass model can serve to develop intuition, but the results need to be validated for the recurrent spiking network.

We now study the effect of various assembly strengths on weight co-tuning and diversity. Instead of directly assessing it, we turn again to the impact of assembly strength on the spiking network's activity correlation structure. Thus, we search for combinations r_{EE} , r_{EB} , r_{IE} , r_{II} that lead to the correlation structure (high in-group and low between-group correlations) that is associated with strong E/I weight co-tuning ($CT_W \approx 1$) and maximum weight diversity ($D \approx 1$). To this end, we use sequential Approximate Bayesian Computation (ABC) [[51](#)] to minimize a loss function defined to be zero when the in-group correlations are equal one and all between-group correlations vanish (for details, see [Methods](#)).

This method allows us to find the approximate posterior distribution of network parameters (the four assembly strengths) that minimize the loss. Afterward, we verify whether connectivity parameters sampled from the approximate posterior lead to the emergence of diversity and co-tuning in the post-synaptic neuron.

Networks with optimized assemblies largely regain the ability to develop E/I co-tuning despite the noise and the non-plastic recurrent connectivity. Assembly strengths that are drawn from the approximate posterior result in a correlation structure very similar to the one observed in a feedforward/low noise network ([Fig 3j and 3k](#)), which allows the plasticity to produce a near-optimal structure in the feedforward connections ([Fig 3l and 3m](#)). We find that the optimal assembly structure involves very strong $E \leftarrow E$ and $I \leftarrow E$ assemblies and medium-strength $E \leftarrow I$ and $I \leftarrow I$ ([Fig 3g and 3i](#)). For details on the impact of assemblies on the network firing and the learned connectivity, see [S1 Text Section C](#) and [S1 Text Section D](#).

This connectivity pattern is similar to the optimal pattern of the reduced linear model, albeit with the difference that the reduced model predicted optimal performance for uniform inhibitory weights (i.e., no inhibitory assemblies, $r_{IE} = r_{II} = 0$). This difference can be attributed to the more complex dynamics of the spiking network that require some degree of local inhibition to prevent extreme synchronization (see [S1 Text Section G](#)), which can negatively impact the STDP's ability to produce co-tuning.

This partial specificity of inhibitory recurrent connectivity can be linked to the role of inhibitory tuning in stabilizing network dynamics at the cost of reduced network feature selectivity [[52](#)]. In theory, the optimal connectivity pattern in the network would promote competition between different groups, for which completely uniform inhibitory connectivity would be ideal [[42](#)]. However, the instability in the population activity induced by such connectivity is detrimental to the emergence of E/I weight co-tuning and input selectivity. Therefore, an intermediate level of specificity in inhibitory recurrent connectivity achieves a balance by maximizing between-group competition (and thus the desired correlation structure) while maintaining stable network dynamics (see [S1 Text Section G](#)).

The sparsity of a network's recurrent connectivity shifts the optimal assembly structure

Biological neural networks are usually very sparsely connected [[53–55](#)], and the sparsity of connections is associated with distinct dynamics [[56](#)]. We observed that the impact of noise




Fig 4. Assemblies improve co-tuning and allow for co-tuning in sparse networks. **a.** weight co-tuning (purple) and diversity (teal) in the networks with assemblies compared to non-structured networks (dashed lines), error bars—standard deviation. The noise level is 0.1 for all sparsities, and the coupling is 1.7 (scaled by $1/p$). **b.** Loss for the sparser networks is higher, which results in the overall worse performance; the inset shows the loss for 50 accepted samples at the last ABC step. **c.** Posterior distributions of all assembly strengths change with sparsity. Sparser networks require weaker inhibitory assemblies (more uniform connections) to produce co-tuning.

<https://doi.org/10.1371/journal.pcbi.1012510.g004>

and recurrence on the deterioration of weight co-tuning and diversity in sparse networks without assemblies is qualitatively similar to fully connected networks. Thus, we examined the ability of neuronal assemblies to produce activity that restores weight co-tuning and diversity in sparsely connected recurrent networks that receive noisy input.

The optimal assembly strength values depend on sparsity levels. We use ABC to discover the approximate posterior distribution of assembly strengths for 5 different levels of sparsity, corresponding to the probability of connection $p = 1.0$, $p = 0.75$, $p = 0.5$, $p = 0.25$, and $p = 0.1$ (Fig 4a). We preserve the total input per neuron across different sparsity levels by scaling the coupling strength inversely proportional to p .

As sparsity increases, the ability of assemblies to improve the tuning diminishes. After 21 ABC steps, the overall loss is larger for the sparse networks than for fully connected networks and increases with sparseness, (Fig 4b). Therefore, despite an improvement in the tuning metrics for most sparse networks (compare the dashed and solid lines in (Fig 4a), particularly diversity is strongly affected by sparseness and cannot be recovered by assemblies to the same extent as for the fully connected networks, (Fig 4a). This reduced effectiveness is expected, given the smaller number of connections and the greater variance in the network's connectivity.

The optimal strength of most assemblies is reduced as the connection probability is decreased (Fig 4c). Specifically, we find that all but $E \leftarrow E$ assemblies should be weaker in sparser networks, with the greatest decrease observed in the $I \leftarrow I$ assemblies that completely disappear for very sparse networks. This could be due to a reduced (compared to fully connected networks) need for within-group recurrent inhibition to prevent completely synchronized behaviours.

Structured connectivity promotes the emergence of co-tuning in fully plastic recurrent networks

In the previous sections, we analyzed a fixed recurrent presynaptic network that projected onto a single postsynaptic neuron via plastic synapses. While this setup provides valuable insights into how structural features can shape population activity for input selectivity to emerge via STDP, it does not fully capture the dynamics of biological neural networks. In this section, we extend our model to a fully plastic and fully recurrent network, offering a more realistic approximation of the behaviour observed in actual biological systems.

In addition to making the recurrent connections between input neurons plastic, we treat the readout neuron as part of the network, and thus, we also introduce sparse feedback

connections from the readout neuron to the input neurons (Fig 5a). Since the recurrent connections are now plastic, we cannot implement the assemblies by changing the in-group connection strength. Thus, we use a sparse network (connection probability $p = 0.25$) and control the in-group connection probability. Following the same logic as with the in-group and between-group connection weights (described in Eq 3), we produce in-group and between-group connection probabilities such that the total connection probability (and consequently the total recurrent input) remains constant while the ratio of recurrent input received by a neuron from neurons of its own vs from other groups varies.

We simulate the network with three connectivity structures. A fully random one, where the sparse connections are implemented without any regard for the neurons' input groups; a fully clustered one, where each input group is almost fully connected, and connections between groups are extremely sparse, and finally, a putative optimally clustered network, which is structured according to the connectivity derived in the previous section for the fixed-weights sparse recurrent network via the Bayesian optimization for the networks of sparsity of $p = 0.25$ (Fig 4c). We measure the co-tuning and weight diversity of the readout neuron, which has now been embedded in the network, projecting randomly with plastic connections to the neurons in the presynaptic network with probability p .

The fully random network displays very low co-tuning and diversity; the fully clustered network develops stronger co-tuning and diversity, which are then surpassed by the optimally clustered network (Fig 5b and 5c). Thus, the putatively optimally clustered network outperforms the other two: after the convergence of the plasticity, the population activity becomes much less noisy, and the activity of different groups can be more easily distinguished, which presumably also enables the development of input selectivity in the post-synaptic neuron (Fig 5d, 5e and 5f). The correlation structure of the converged plastic networks verifies this observation, with the putative optimally clustered network having very high in-group and near-zero between-group correlation. This structure is introduced at the beginning of the simulation by

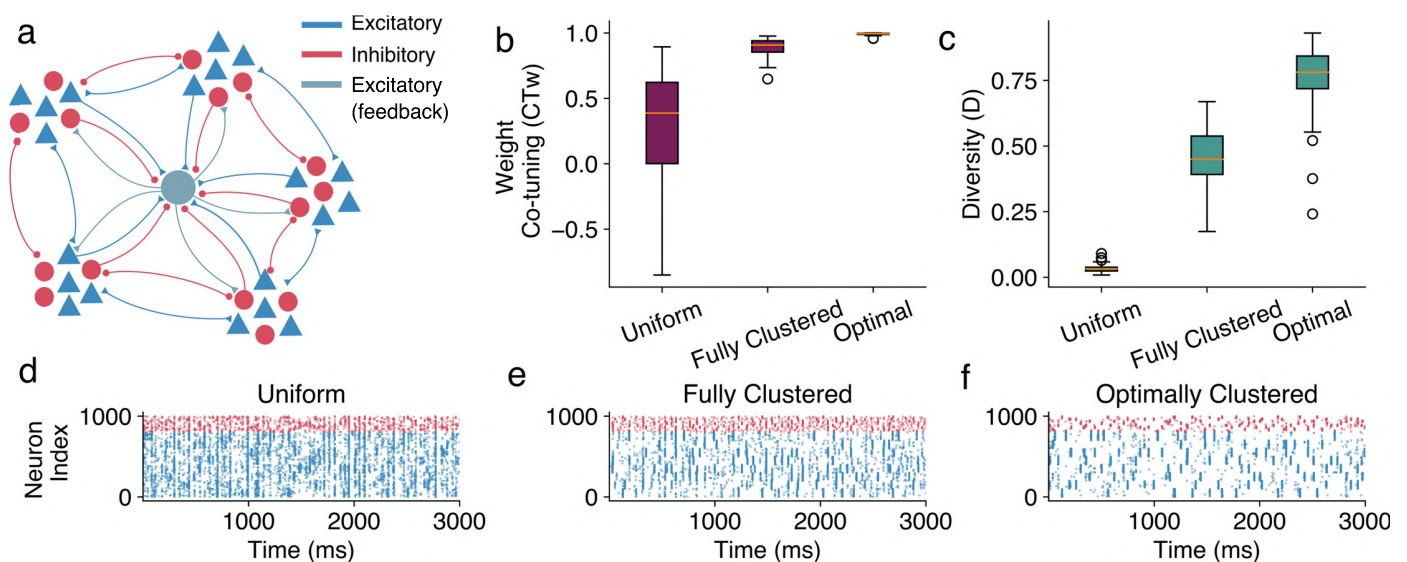


Fig 5. Structured recurrent connectivity can drive synaptic learning even in fully plastic networks. **a.** A diagram of the fully recurrent plastic network. A comparison of the **(b.)** weight co-tuning (CT_w) and **(c.)** diversity (D) metrics for plastic networks with fully random, fully clustered and optimal connectivity structure. Spike trains of the converged networks with **(d.)** fully random, **(e.)** fully clustered, and **(f.)** optimal recurrent connectivity structures. The activity of different input groups is much more correlated, allowing for easier discrimination in the last network.

<https://doi.org/10.1371/journal.pcbi.1012510.g005>

the structural connectivity and is maintained throughout the learning process, driving the development of input selectivity in the postsynaptic neuron.

Directly optimizing the fully plastic network is not computationally feasible because, for simulation-based inference, we would need to simulate the network's activity until all the plastic connections converge in each simulation. However, the connectivity from the static sparse presynaptic network appears to give a good prior for a beneficial connectivity in the fully plastic networks.

Our findings suggest that structured connectivity can drive the statistics of activity even in fully plastic networks and enable the development of specific connectivity in neurons that do not receive direct input from other areas.

Discussion

Synaptic plasticity is theorized to be responsible for the formation of input selectivity across brain hierarchies, including in brain areas that only receive input from highly recurrent networks. Here, we demonstrated how the structure of non-plastic presynaptic recurrent connectivity could hinder or boost the ability of synaptic plasticity mechanisms [12, 26, 27, 31] to generate input selectivity in neurons of higher areas. We find that strong excitatory connectivity among neurons tuned to the same input, combined with broader inhibition, creates population activity with a beneficial statistical structure that enables the formation of co-tuned projections by plasticity, potentially fostering input selectivity.

How different plasticity mechanisms shape neural connectivity, such as the formation of E/I co-tuning [12, 26, 27, 31] in feedforward networks or neural assemblies in recurrent networks [26, 34–36], has been a topic of extensive theoretical research. Nevertheless, the opposite effect—the ways in which fixed connectivity can shape the effects of synaptic plasticity—has only been studied in very specialized cases [57]. This omission partially obscures the two-way interaction between connectivity and synaptic plasticity in biological neural networks. While synaptic plasticity constantly modifies some aspects of neural connectivity, it acts within the many constraints of network structure that are either constant throughout an organism's lifetime or change via structural adaptation mechanisms that act on timescales slower than synaptic plasticity [58, 59]. Effectively, this means that synaptic plasticity relies on population activity originating from networks with highly intricate connectivity structures very different from those of random networks [33, 60]. Given that most synaptic plasticity mechanisms fundamentally depend on the statistics imposed by network activity, it is reasonable to assume that the network structure highly impacts the behaviour of synaptic plasticity.

Cortical connectivity is known to be highly clustered [61], and the clustering has functional as well as spatial determinants. For example, neurons that share common inputs [62] or targets [63] are more likely to form recurrent connections between themselves [64, 65]. Moreover, excitatory cells with similar receptive fields are known to form strong reciprocal connections [66], which determine neural responses. Additionally, cortical networks have been shown to present specific correlation structures early in development [67], suggesting that recurrent cortical connectivity is at least partially structured before sensory inputs are present.

Clustered networks display distinct dynamics, including competition between clusters [42, 68] and slower timescales [41, 69], both of which can be useful for computations. Additionally, there is strong evidence that groups of highly interconnected neurons (neuronal assemblies) share common functions within recurrent networks [37, 70, 71]. Moreover, evidence has accumulated [72, 73] that different neuron types (excitatory and inhibitory subtypes) follow distinct spatial connectivity patterns, which have implications for neural computation [74]. Thus, our findings complement the ongoing research on computational implications of recurrent

neural connectivity in biological networks by suggesting a link between specific, fixed connectivity patterns and local learning in feedforward projections to downstream populations via synaptic plasticity.

While the connectivity pattern we identify in our study is biologically plausible, the extent to which it is realized *in vivo* remains unclear. Further experimental studies in the connectivity patterns of different neuron types are necessary to model different network connectivities and study their dynamics effectively. On a theoretical side, more research is needed to uncover whether the synapse-type specific assembly structures we identify can emerge in plastic networks without any prior structure. Specifically, while there have been several studies investigating the emergence of stable assemblies via STDP [33–35, 75], the protocols by which assemblies of different pre-selected strengths for each connection type could arise remain unclear. Potential mechanisms include structural plasticity and variation in the learning rates of different synaptic types. Additionally, the presence of different regulatory interneurons, which have already been studied in the context of assembly formation [34], could play a role in modulating the relative assembly strengths of different connections.

For our study, we parameterized the network connectivity by adopting a quantitative metric for the strength of different types of neuronal assemblies. This resulted in a low-dimensional parameter space and allowed us to use rejection sampling-based ABC [51] to infer the optimal assembly strengths. One limitation of this technique is that it suffers from the curse of dimensionality and typically requires a large simulation budget [76, 77]. Therefore, extensions of the current work using higher dimensional connectivity parameters or simultaneous inference of the connectivity and neuron parameters will require more efficient simulation-based methods such as neural posterior estimation [78]. Alternatively, direct optimization of each recurrent weight via gradient-based methods [79, 80] may uncover more intricate connectivity patterns that are not limited to the specific network parametrization we chose.

To summarize, we identified how particular presynaptic connectivity structures could be a favourable or detrimental substrate for plasticity to develop co-tuning of excitation and inhibition on neuronal projections. Our study is the first step in illuminating the two-way dependence between the non-plastic structural features of a network’s connectivity and synaptic plasticity, which can motivate further research on this intricate interaction.

Materials and methods

Neuron model

We modelled all neurons of our networks as leaky integrate-and-fire (LIF) neurons with leaky synapses [81]. The evolution of their membrane potential is given by the ODE:

$$C_m \cdot \frac{dV(t)}{dt} = g_{\text{leak}} \cdot (V_{\text{rest}} - V(t)) + g_I(t) \cdot (V_I - V(t)) + g_E(t) \cdot (V_E - V(t)), \quad (4)$$

where V_{rest} is the neuron’s resting potential, V_E , V_I are the excitatory and inhibitory reversal potentials and g_{leak} the leak conductance. Additionally, the excitatory and inhibitory conductances g_E , g_I decay exponentially over time and get boosts upon excitatory or inhibitory pre-synaptic spiking, respectively, as

$$\begin{aligned} \frac{dg_E(t)}{dt} &= -\frac{g_E(t)}{\tau_E} + \bar{g}_E \cdot \sum_j W_j^E \cdot \sum_f \delta(t - t_j^f), \\ \frac{dg_I(t)}{dt} &= -\frac{g_I(t)}{\tau_I} + \bar{g}_I \cdot \sum_j W_j^I \cdot \sum_f \delta(t - t_j^f). \end{aligned} \quad (5)$$

Here t_j^f denotes the time at which the f -th spike of the j -th neuron happened and $\delta(t)$ is the Dirac's delta function. When the membrane potential reaches the spiking threshold V_{th} , a spike is emitted, and the potential is changed to a reset potential V_{reset} . Finally, the neurons have an absolute refractory period (during which no spikes are emitted even when the spiking threshold is reached) between spikes of $\tau_{ref} = 5ms$.

We would like to remark that both $W_i^E, W_i^I > 0$, since the effect of inhibition is encoded on Eq (4). However, for illustration purposes inhibitory weights are currents are shown to be negative. For instance, this happens in (Fig 1b). An alternative neuron model is discussed in S1 Text Section B.

Network input

The external input to each of the 1000 pre-synaptic neurons is the mixture of two Poisson spike trains. The first Poisson spike train is shared with all the other neurons of the same group, while the second Poisson spike train is the individual noise of the neuron,

$$C_{total} = C_{signal} + C_{noise}, \tag{6}$$

where $C_{signal} \sim \text{Poisson}((1 - c) \cdot f_0)$ and $C_{noise} \sim \text{Poisson}(c \cdot f_0)$. Here, f_0 is the total firing rate of the input, and c is the strength of the noise. C_{signal} is the same for all neurons of the same input group, while C_{noise} is individual to each neuron.

Recurrent connectivity

The recurrent connectivity is implemented in two different versions for the fixed and plastic versions.

Non-plastic recurrent connectivity. The non-plastic recurrent connectivity between the input neurons is defined by a coupling strength parameter W , which defines the average synaptic strength and a connection probability p , which defines the sparsity. The connectivity is implemented as follows:

At first, an adjacency matrix A is defined, which implements an Erdős–Rényi connectivity with connection probability p (i.e., a connection between any two neurons is implemented with connectivity p , independently of any other connection). Then, using the coupling strength parameter W , and the given assembly strength for each connection type r_{EE}, r_{EI}, r_{IE} , and r_{II} , we extract the parameters W_{in}^{ab} and W_{out}^{ab} for each connection type ($a, b \in \{E, I\}$) according to Eq 3. The inhibitory weights $W_{in}^{IE}, W_{out}^{IE}, W_{in}^{II}$ and W_{out}^{II} are scaled by a parameter g_s which is set to counterbalance the slower inhibitory synapse dynamics and the smaller number of I neurons. This scaling leads to an approximately balanced network across implementations. Once the connectivity strengths are calculated, for each pre and post-synaptic neuron pair i and j , we set the connection between them as

$$w_{ij} = \begin{cases} 0, & \text{if } A_{ij} = 0 \\ \alpha_{ij} \sim |\mathcal{N}(W^{ij}, c \cdot W^{ij})|, & \text{if } A_{ij} = 1 \end{cases} \tag{7}$$

where W^{ij} is the appropriate connectivity strength ($W_{in}^{ab}, W_{out}^{ab}$ for $a, b \in \{E, I\}$) depending on the neuron type of neurons i and j and whether they belong to same assembly or not. The parameter c , which scales the standard deviation, was normally set to 0.1, but we also examined narrower and broader distributions with similar results.

We finally considered an alternative, log-normal distribution of weights, which increased variability but largely lead to the same results.

Plastic recurrent connectivity. In the plastic recurrent network case, instead of varying the connection strengths for in-group and between-group connections, we use the total connection probability p and given assembly strength for each connection type r_{EE} , r_{EB} , r_{IE} , and r_{IB} , we extract parameters p_{in}^{ab} and p_{out}^{ab} for each connection type ($a, b \in \{E, I\}$) according to Eq 3. These parameters give the probability that a connection of a specific type is implemented in the recurrent network, which creates an inhomogeneous adjacency matrix A that implements the different levels of clustering for each connection type.

Once the adjacency matrix has been defined, we set the initial connection strength for pre and post-synaptic neuron pair i and j as:

$$w_{ij} = \begin{cases} 0, & \text{if } A_{ij} = 0 \\ \alpha_{ij} \sim |\mathcal{N}(W, c \cdot W)|, & \text{if } A_{ij} = 1 \end{cases} \tag{8}$$

where W is the coupling strength parameter, which we scale for inhibitory connections, similarly to the non-plastic network. The resulting connectivity reflects the initial conditions for the plastic recurrent network and the non-zero connections are updated according to the same plasticity protocol that is used to learn the feedforward connectivity in the networks with fixed recurrent connectivity.

Plasticity

Triplet excitatory STDP. The excitatory connections are modified according to a simplified form of the triplet STDP rule [43], which has been shown to generalize the Bienenstock–Cooper–Munro (BCM) rule [9] for higher-order correlations [48]. In our implementation of the triplet rule, the firing rates of the pre-synaptic excitatory neurons and the post-synaptic neuron are approximated by traces with two different timescales (we use the same timescales for the fast and slow traces of the pre and postsynaptic neuron):

$$\frac{dy_k^E(t)}{dt} = -\frac{y_k^E(t)}{\tau_1^{estdp}} + \sum_f \delta(t - t_k^f), \tag{9a}$$

$$\frac{dz_k^E(t)}{dt} = -\frac{z_k^E(t)}{\tau_2^{estdp}} + \sum_f \delta(t - t_k^f), \tag{9b}$$

$$\frac{dx_1(t)}{dt} = -\frac{x_1(t)}{\tau_1^{estdp}} + \sum_f \delta(t - t_x^f), \tag{9c}$$

$$\frac{dx_2(t)}{dt} = -\frac{x_2(t)}{\tau_2^{estdp}} + \sum_f \delta(t - t_x^f), \tag{9d}$$

where $\tau_1^{estdp} < \tau_2^{estdp}$ are the two timescales of the plasticity rule, $y_k^E(t)$, $z_k^E(t)$ and $x_1(t)$, $x_2(t)$ represent the slow and fast traces of the k -th excitatory pre-synaptic and the single post-synaptic neuron respectively while t_k^f and t_x^f are their respective firing times. The function $\delta(x)$ represents a Dirac's delta. The connection weights are updated upon pre and post-synaptic spiking according to

$$\Delta W_k^E = \eta_E \cdot A_{LTP} \cdot x_2(t) \cdot y_k^E(t) \cdot \sum_f R(t - t_x^f) - \eta_E \cdot A_{LTD} \cdot x_1(t) \cdot z_k^E(t) \cdot \sum_f R(t - t_k^f), \tag{10}$$

where η_E is the excitatory learning rate and A_{LTP}, A_{LTD} the amplitudes of long term depression and potentiation respectively. Despite scaling down the LTD amplitude to account for higher presynaptic firing rates, the rule remains slightly LTD dominated in our experiments, a setting that has been observed in experimental studies [82, 83]. The function $R(x)$ is defined as:

$$R(x) = \begin{cases} 1, & x = 0 \\ 0, & x \neq 0 \end{cases} \tag{11}$$

In numerical simulation, if spikes have been produced in the last few timesteps, so in practice $R(t) = 1$ for $t \in [-\delta t, 0]$ for a small δt . Hence, $R(t)$ is a rectangular function. Different parameters and learning rules for the excitatory plasticity are discussed in (S1 Text Section E.i and E.ii).

Inhibitory STDP. We used the learning rule first proposed in [26] for the inhibitory connections. Approximations of the firing rates are kept via a trace for each of the pre-synaptic inhibitory neurons as well as the post-synaptic neuron,

$$\frac{dy_k^I}{dt} = -\frac{y_k^I}{\tau^{istdp}} + \sum_f \delta(t - t_k^f), \tag{12a}$$

$$\frac{dx}{dt} = -\frac{x}{\tau^{istdp}} + \sum_f \delta(t - t_x^f), \tag{12b}$$

where τ^{istdp} is the single timescale of the plasticity rule, y_k^I and x are the traces of the the k_{th} inhibitory pre-synaptic and the single post-synaptic neuron, and t_k^f, t_x^f are their respective spike times. The connection weights are updated upon pre and post-synaptic spiking as

$$\Delta W_k^I = \eta_I \cdot (x(t) - 2\rho_0 \tau^{istdp}) \cdot \sum_f R(t - t_k^f) + \eta_I \cdot y_k^I(t) \cdot \sum_f R(t - t_x^f). \tag{13}$$

Here, η_I is the inhibitory learning rate, and ρ_0 is the target firing rate of the post-synaptic neuron. The rectangular function $R(t)$ is defined in Eq (10).

Weight normalization. Due to the instability of the triplet STDP rule, some normalization mechanism is needed to constrain weight development. We use a modified version of the competitive normalization protocol proposed in [31], which we adapt for spiking neurons.

Specifically, we normalize the k -th connection every time there is a weight update (i.e., upon pre or postsynaptic spiking):

$$W_k^A(t) \leftarrow (1 - \eta_N) \cdot W_k^A(t) + \eta_N \cdot W_k^A(t) \cdot \frac{W_{target}^A}{\sum_{i=1}^{N_A} W_i^A(t)}, \quad A \in \{E, I\}. \tag{14}$$

Where W_{target}^A is the target total weight of each connection type and η_N is the normalization rate. In the recurrent plastic networks, the W_{target}^A for the recurrent neurons is determined by the coupling strength W . The normalization pushes the sum of the excitatory and the sum of the inhibitory feedforward connections weights close to the set target total weights W_{target}^E and W_{target}^I over time. The implications of implementing a regular normalization step (on every time step only when spiking occurs) are discussed in S1 Text Section E.iv.1. Moreover, the implications of using different normalization rates η_N are discussed in S1 Text Section E.iv.2.

An alternative way to stabilize the weights via subtractive normalization of only the excitatory synapses [12, 23, 33] was also considered leading to comparable results (see S1 Text Section E.v).

Approximating the posterior distribution of the model parameters

To estimate the set of parameters that lead to high in-group correlations and low out-group correlations, we used simulation-based inference [76]. The basic idea is to use simulation with known parameters to approximate the full posterior distributions for the model given the required output, i.e., the distribution of parameters and samples from which produce the required correlation structure. We use sequential Approximate Bayesian Computation (ABC) [51] to approximate the posterior distribution. We define a loss function that maximizes in-group correlations and minimizes between-group correlations:

$$\mathcal{L} = -\alpha C_{in}^2 - \beta [(1 - C_{out}^{EE})^2 + (1 - C_{out}^{EI})^2 + (1 - C_{out}^{II})^2] \tag{15}$$

We define a uniform prior $p(\theta)$. A set of parameters $\theta = [r_{ee}, r_{ei}, r_{ie}, r_{ii}]$ is sampled from it and used to run the simulations for 3 seconds. From the simulation results, correlations are computed, which allows us to obtain the loss. We accept a parameter set if the loss is below the error ϵ , and keep sampling until the number of accepted samples is 60. We use the kernel density estimate on the accepted samples to obtain an approximate posterior. Next, we rescale this approximate posterior with the original prior to obtain a proposal distribution that we use as a prior in the next step of the ABC. In each step, we reduce ϵ by setting it to the 75th percentile of the losses for the accepted samples (see [51] for more details). As a rule, we run 20 to 30 steps of the sequential ABC until the loss converges. We run separate fits for networks with different levels of sparsity with connection probabilities $p = 0.1, 0.25, 0.5, 0.75, 1.0$. The fitting was done using a modified version of the simple-abc toolbox <https://github.com/rcmorehead/simpleabc/> for python.

Reduced model

The dynamics of the system can be studied analytically using a simplified, reduced linear model. Here, each pair of variables (x_i, y_i) represents the excitatory and inhibitory mean firing rate of a neuron group. In theory, these variables display complicated non-linear interactions that arise from the microscopic details of the LIF spiking network and synapse dynamics. However, in the stationary state –and away from any critical point– a linearised model can capture the essential features of the correlations between different populations.

Internal noise, modelled as independent Poisson trains to each individual neuron, becomes Gaussian white noise in the large-population limit, characterized by zero mean and variance σ_{int} . Each population is affected by different internal fluctuations. For simplicity, external noise, which is applied as the same train of Poisson spikes to all the neurons inside an input group, will also be approximated as a Gaussian white noise of mean η_0 and variance σ_{ext} .

Therefore, the simplified linear model reads:

$$\dot{x}_i = ax_i + by_i + \frac{1}{M-1} \sum_{j \neq i} (W^{E \leftarrow E} x_j + W^{E \leftarrow I} y_j) + \sigma_{int} \xi_i^x(t) + \sigma_{ext} \eta_i(t) + \eta_0, \tag{16a}$$

$$\dot{y}_i = cx_i + dy_i + \frac{1}{M-1} \sum_{j \neq i} (W^{I \leftarrow E} x_j + W^{I \leftarrow I} y_j) + \sigma_{int} \xi_i^y(t) + \sigma_{ext} \eta_i(t) + \eta_0, \tag{16b}$$

where M is the number of populations, a, b, c, d are parameters controlling in-group recurrent coupling, and $W^{E \leftarrow E}, W^{E \leftarrow I}, W^{I \leftarrow E}, W^{I \leftarrow I}$ are couplings between different clusters. Internal noise for each population is represented by $\xi_i^{x,y}(t)$, while external noise is notated as $\eta_i(t)$.

All noises are uncorrelated, meaning that

$$\langle \xi_i^{cc'} \xi_j^{c'} \rangle = \delta_{cc'} \delta_{ij} \delta(t - t'), \quad (17a)$$

$$\langle \xi_i^{cc'}(t) \eta_j(t') \rangle = 0 \quad \forall i, j, t, t', \quad (17b)$$

$$\langle \eta_i(t) \eta_j(t') \rangle = \delta_{ij} \delta(t - t'), \quad (17c)$$

with $c, c' = \{x, y\}$, and where $\langle \dots \rangle$ represents an ensemble average, i.e., an average over noise realizations. From this model, it is possible to obtain closed equations for Pearson correlation coefficients (see [S1 Text](#) Section H for details). Notice that stochastic differential equations are never complete without an interpretation, and we choose to interpret these in the Itô sense, which will be relevant for computations. Tables of all the parameters used in our simulations are given in [S1 Text](#) Section K Tables A, B and C.

Supporting information

S1 Text. Supplementary information.
(PDF)

Acknowledgments

EG thanks the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support, and OV thanks the International Max Planck Research School for the Mechanisms of Mental Function and Dysfunction and the Joachim Herz Foundation for support. AL is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1—Project number 39072764.

Author Contributions

Conceptualization: Emmanouil Giannakakis, Anna Levina.

Formal analysis: Emmanouil Giannakakis, Victor Buendía.

Funding acquisition: Anna Levina.

Investigation: Emmanouil Giannakakis, Oleg Vinogradov.

Methodology: Emmanouil Giannakakis, Oleg Vinogradov, Victor Buendía.

Project administration: Anna Levina.

Resources: Anna Levina.

Software: Emmanouil Giannakakis, Oleg Vinogradov, Victor Buendía.

Supervision: Anna Levina.

Visualization: Emmanouil Giannakakis, Oleg Vinogradov, Victor Buendía.

Writing – original draft: Emmanouil Giannakakis, Oleg Vinogradov, Victor Buendía, Anna Levina.

Writing – review & editing: Emmanouil Giannakakis, Oleg Vinogradov, Victor Buendía, Anna Levina.

References

1. Priebe NJ. Mechanisms of orientation selectivity in the primary visual cortex. *Annual review of vision science*. 2016; 2(1):85–107. <https://doi.org/10.1146/annurev-vision-111815-114456> PMID: 28532362
2. Osorio D. Directionally selective cells in the locust medulla. *Journal of Comparative Physiology A*. 1986; 159(6):841–847. <https://doi.org/10.1007/BF00603737> PMID: 3806440
3. Ringach DL, Shapley RM, Hawken MJ. Orientation selectivity in macaque V1: diversity and laminar dependence. *Journal of neuroscience*. 2002; 22(13):5639–5651. <https://doi.org/10.1523/JNEUROSCI.22-13-05639.2002> PMID: 12097515
4. Riesenhuber M, Poggio T. Hierarchical models of object recognition in cortex. *Nature neuroscience*. 1999; 2:1019–25. <https://doi.org/10.1038/14819> PMID: 10526343
5. Scholl B, Tan AYY, Corey J, Priebe NJ. Emergence of Orientation Selectivity in the Mammalian Visual Pathway. *Journal of Neuroscience*. 2013; 33(26):10616–10624. <https://doi.org/10.1523/JNEUROSCI.0404-13.2013> PMID: 23804085
6. Rossi LF, Harris KD, Carandini M. Spatial connectivity matches direction selectivity in visual cortex. *Nature*. 2020; 588(7839):648–652. <https://doi.org/10.1038/s41586-020-2894-4> PMID: 33177719
7. Sharma J, Angelucci A, Sur M. Induction of visual orientation modules in auditory cortex. *Nature*. 2000; 404(6780):841–847. <https://doi.org/10.1038/35009043> PMID: 10786784
8. Feldman DE, Brecht M. Map plasticity in somatosensory cortex. *Science*. 2005; 310(5749):810–815. <https://doi.org/10.1126/science.1115807> PMID: 16272113
9. Bienenstock E, Cooper L, Munro P. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*. 1982; 2(1):32–48. <https://doi.org/10.1523/JNEUROSCI.02-01-00032.1982> PMID: 7054394
10. Blais BS, Intrator N, Shouval H, Cooper LN. Receptive Field Formation in Natural Scene Environments: Comparison of Single-Cell Learning Rules. *Neural Computation*. 1998; 10(7):1797–1813. <https://doi.org/10.1162/089976698300017142> PMID: 9744898
11. Hansel D, van Vreeswijk C. The mechanism of orientation selectivity in primary visual cortex without a functional map. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*. 2012; 32(12):4049–4064. <https://doi.org/10.1523/JNEUROSCI.6284-11.2012> PMID: 22442071
12. Clopath C, Vogels TP, Froemke RC, Sprekeler H. Receptive field formation by interacting excitatory and inhibitory synaptic plasticity. *bioRxiv*. 2016;.
13. Brito CSN, Gerstner W. Nonlinear Hebbian Learning as a Unifying Principle in Receptive Field Formation. *PLOS Computational Biology*. 2016; 12(9):1–24. <https://doi.org/10.1371/journal.pcbi.1005070> PMID: 27690349
14. Deneve S, Machens C. Efficient codes and balanced networks. *Nature Neuroscience*. 2016; 19:375–382. <https://doi.org/10.1038/nn.4243> PMID: 26906504
15. Zhou S, Yu Y. Synaptic E-I Balance Underlies Efficient Neural Coding. *Frontiers in Neuroscience*. 2018; 12. <https://doi.org/10.3389/fnins.2018.00046> PMID: 29456491
16. Singer W. Recurrent dynamics in the cerebral cortex: Integration of sensory evidence with stored knowledge. *Proceedings of the National Academy of Sciences*. 2021; 118(33):e2101043118. <https://doi.org/10.1073/pnas.2101043118> PMID: 34362837
17. Isaacson J, Scanziani M. How Inhibition Shapes Cortical Activity. *Neuron*. 2011; 72(2):231–243. <https://doi.org/10.1016/j.neuron.2011.09.027> PMID: 22017986
18. Okun M, Lampl I. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nature Neuroscience*. 2008; 11(5):535–537. <https://doi.org/10.1038/nn.2105> PMID: 18376400
19. Vogels T, Abbott LF. Gating Multiple Signals through Detailed Balance of Excitation and Inhibition in Spiking Networks. *Nature neuroscience*. 2009; 12:483–91. <https://doi.org/10.1038/nn.2276> PMID: 19305402
20. Bhatia A, Moza S, Bhalla US. Precise excitation-inhibition balance controls gain and timing in the hippocampus. *eLife*. 2019; 8:e43415. <https://doi.org/10.7554/eLife.43415> PMID: 31021319
21. Ozeki H, Finn I, Schaffer E, Miller K, Ferster D. Inhibitory Stabilization of the Cortical Network Underlies Visual Surround Suppression. *Neuron*. 2009; 62:578–92. <https://doi.org/10.1016/j.neuron.2009.03.028> PMID: 19477158
22. Rubin D, Van Hooser S, Miller K. The Stabilized Supralinear Network: A Unifying Circuit Motif Underlying Multi-Input Integration in Sensory Cortex. *Neuron*. 2015; 85(2):402–417. <https://doi.org/10.1016/j.neuron.2014.12.026> PMID: 25611511
23. Schulz A, Miehl C, Berry I, Michael J, Gjorgjieva J. The generation of cortical novelty responses through inhibitory plasticity. *eLife*. 2021; 10:e65309. <https://doi.org/10.7554/eLife.65309> PMID: 34647889

24. Wehr M, Zador AM. Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature*. 2003; 426:442–446. <https://doi.org/10.1038/nature02116> PMID: 14647382
25. Wu YK, Miehl C, Gjorgjieva J. Regulation of circuit organization and function through inhibitory synaptic plasticity. *Trends in Neurosciences*. 2022; 45(12):884–898. <https://doi.org/10.1016/j.tins.2022.10.006> PMID: 36404455
26. Vogels TP, Sprekeler H, Zenke F, Clopath C, Gerstner W. Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks. *Science*. 2011; 334(6062):1569–1573. <https://doi.org/10.1126/science.1211095> PMID: 22075724
27. Luz Y, Shamir M. Balancing Feed-Forward Excitation and Inhibition via Hebbian Inhibitory Synaptic Plasticity. *PLOS Computational Biology*. 2012; 8(1):1–12. <https://doi.org/10.1371/journal.pcbi.1002334> PMID: 22291583
28. Hellyer PJ, Jachs B, Clopath C, Leech R. Local inhibitory plasticity tunes macroscopic brain dynamics and allows the emergence of functional brain networks. *NeuroImage*. 2016; 124:85–95. <https://doi.org/10.1016/j.neuroimage.2015.08.069> PMID: 26348562
29. Khajehabdollahi S, Giannakakis E, Prosi J, Levina A. Reservoir computing with self-organizing neural oscillators. *The 2021 Conference on Artificial Life Proceedings*. 2021;.
30. Agnes EJ, Luppi AI, Vogels TP. Complementary Inhibitory Weight Profiles Emerge from Plasticity and Allow Flexible Switching of Receptive Fields. *Journal of Neuroscience*. 2020; 40(50):9634–9649. <https://doi.org/10.1523/JNEUROSCI.0276-20.2020> PMID: 33168622
31. Eckmann S, Young EJ, Gjorgjieva J. Synapse-type-specific competitive Hebbian learning forms functional recurrent networks. *Proceedings of the National Academy of Sciences*. 2024; 121(25):e2305326121. <https://doi.org/10.1073/pnas.2305326121> PMID: 38870059
32. Agnes EJ, Vogels TP. Co-dependent excitatory and inhibitory plasticity accounts for quick, stable and long-lasting memories in biological networks. *Nature Neuroscience*. 2024; 27(5):964–974. <https://doi.org/10.1038/s41593-024-01597-4> PMID: 38509348
33. Litwin-Kumar A, Doiron B. Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nature communications*. 2014; 5:5319. <https://doi.org/10.1038/ncomms6319> PMID: 25395015
34. Lagzi F, Bustos MC, Oswald AM, Doiron B. Assembly formation is stabilized by Parvalbumin neurons and accelerated by Somatostatin neurons. *bioRxiv*. 2021;.
35. Mackwood O, Naumann LB, Sprekeler H. Learning excitatory-inhibitory neuronal assemblies in recurrent networks. *eLife*. 2021; 10:e59715. <https://doi.org/10.7554/eLife.59715> PMID: 33900199
36. Zenke F, Agnes E, Gerstner W. Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nature Communications*. 2015; 6:6922. <https://doi.org/10.1038/ncomms7922> PMID: 25897632
37. Miehl C, Gjorgjieva J. Stability and learning in excitatory synapses by nonlinear inhibitory plasticity. *PLOS Computational Biology*. 2022; 18(12):1–34. <https://doi.org/10.1371/journal.pcbi.1010682> PMID: 36459503
38. Effenberger F, Jost J, Levina A. Self-organization in balanced state networks by STDP and homeostatic plasticity. *PLoS computational biology*. 2015; 11(9):e1004420. <https://doi.org/10.1371/journal.pcbi.1004420> PMID: 26335425
39. Larisch R, Gönner L, Teichmann M, Hamker FH. Sensory coding and contrast invariance emerge from the control of plastic inhibition over emergent selectivity. *PLOS Computational Biology*. 2021; 17(11):1–37. <https://doi.org/10.1371/journal.pcbi.1009566> PMID: 34843455
40. Jiang X, Shen S, Cadwell CR, Berens P, Sinz F, Ecker AS, et al. Principles of connectivity among morphologically defined cell types in adult neocortex. *Science*. 2015; 350(6264):aac9462. <https://doi.org/10.1126/science.aac9462> PMID: 26612957
41. Litwin-Kumar A, Doiron B. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience*. 2012; 15:1498–1505. <https://doi.org/10.1038/nn.3220> PMID: 23001062
42. Lagzi F, Atay F, Rotter S. Bifurcation analysis of the dynamics of interacting subnetworks of a spiking network. *Scientific Reports*. 2019; 9:1–17. <https://doi.org/10.1038/s41598-019-47190-9> PMID: 31388027
43. Pfister JP, Gerstner W. Triplets of Spikes in a Model of Spike Timing-Dependent Plasticity. *Journal of Neuroscience*. 2006; 26(38):9673–9682. <https://doi.org/10.1523/JNEUROSCI.1425-06.2006> PMID: 16988038
44. Wu GK, Arbuckle R, Hua Liu B, Tao HW, Zhang LI. Lateral Sharpening of Cortical Frequency Tuning by Approximately Balanced Inhibition. *Neuron*. 2008; 58(1):132–143. <https://doi.org/10.1016/j.neuron.2008.01.035> PMID: 18400169

45. Sutter ML, Loftus WC. Excitatory and Inhibitory Intensity Tuning in Auditory Cortex: Evidence for Multiple Inhibitory Mechanisms. *Journal of Neurophysiology*. 2003; 90(4):2629–2647. <https://doi.org/10.1152/jn.00722.2002> PMID: 12801894
46. Liu G. Local structural balance and functional interaction of excitatory and inhibitory synapses in hippocampal dendrites. *Nature Neuroscience*. 2004; 7(4):373–379. <https://doi.org/10.1038/nn1206> PMID: 15004561
47. Sukenik N, Vinogradov O, Weinreb E, Segal M, Levina A, Moses E. Neuronal circuits overcome imbalance in excitation and inhibition by adjusting connection numbers. *Proceedings of the National Academy of Sciences*. 2021; 118(12). <https://doi.org/10.1073/pnas.2018459118> PMID: 33723048
48. Gjorgjieva J, Clopath C, Audet J, Pfister JP. A triplet spike-timing-dependent plasticity model generalizes the Bienenstock–Cooper–Munro rule to higher-order spatiotemporal correlations. *Proceedings of the National Academy of Sciences*. 2011; 108(48):19383–19388. <https://doi.org/10.1073/pnas.1105933108> PMID: 22080608
49. Ecker AS, Tolias AS. Is there signal in the noise? *Nature neuroscience*. 2014; 17(6):750–751. <https://doi.org/10.1038/nn.3722> PMID: 24866037
50. Aponte DA, Handy G, Kline AM, Tsukano H, Doiron B, Kato HK. Recurrent network dynamics shape direction selectivity in primary auditory cortex. *Nature Communications*. 2021; 12(1):314. <https://doi.org/10.1038/s41467-020-20590-6> PMID: 33436635
51. Beaumont MA, Cornuet JM, Marin JM, Robert CP. Adaptive approximate Bayesian computation. *Biometrika*. 2009; 96(4):983–990. <https://doi.org/10.1093/biomet/asp052>
52. Lagzi F, Fairhall AL. Emergence of co-tuning in inhibitory neurons as a network phenomenon mediated by randomness, correlations, and homeostatic plasticity. *Science Advances*. 2024; 10(12):eadi4350. <https://doi.org/10.1126/sciadv.adi4350> PMID: 38507489
53. Seeman SC, Campagnola L, Davoudian PA, Hoggarth A, Hage TA, Bosma-Moody A, et al. Sparse recurrent excitatory connectivity in the microcircuit of the adult mouse and human cortex. *eLife*. 2018; 7:e37349. <https://doi.org/10.7554/eLife.37349> PMID: 30256194
54. Wildenberg GA, Rosen MR, Lundell J, Paukner D, Freedman DJ, Kasthuri N. Primate neuronal connections are sparse in cortex as compared to mouse. *Cell Reports*. 2021; 36(11):109709. <https://doi.org/10.1016/j.celrep.2021.109709> PMID: 34525373
55. Barral J, D Reyes A. Synaptic scaling rule preserves excitatory–inhibitory balance and salient neuronal network dynamics. *Nature Neuroscience*. 2016; 19(12):1690–1696. <https://doi.org/10.1038/nn.4415> PMID: 27749827
56. Brunel. Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons. *Journal of Computational Neuroscience*. 2000; 8. PMID: 10809012
57. Giannakakis E, Hutchings F, Papasavvas CA, Han CE, Weber B, Zhang C, et al. Computational modeling of the long-term effects of brain stimulation on the local and global structural connectivity of epileptic patients. *PLOS ONE*. 2020; 15(2):1–21. <https://doi.org/10.1371/journal.pone.0221380> PMID: 32027654
58. Tetzlaff C, Kolodziejcki C, Markelic I, Wörgötter F. Time scales of memory, learning, and plasticity. *Biological Cybernetics*. 2012; 106. <https://doi.org/10.1007/s00422-012-0529-z> PMID: 23160712
59. Zenke F, Gerstner W, Ganguli S. The temporal paradox of Hebbian learning and homeostatic plasticity. *Current Opinion in Neurobiology*. 2017; 43:166–176. <https://doi.org/10.1016/j.conb.2017.03.015> PMID: 28431369
60. Negrón A, Getz MP, Handy G, Doiron B. The mechanics of correlated variability in segregated cortical excitatory subnetworks. *bioRxiv*. 2023; <https://doi.org/10.1101/2023.04.25.538323> PMID: 37162867
61. Hilgetag C, Kaiser M. Clustered Organization of Cortical Connectivity. *Neuroinformatics*. 2004; 2:353–60. <https://doi.org/10.1385/NI:2:3:353> PMID: 15365196
62. Yoshimura Y, Dantzker J, Callaway E. Excitatory cortical neurons form fine-scale functional networks. *Nature*. 2005; 433:868–73. <https://doi.org/10.1038/nature03252> PMID: 15729343
63. Brown S, Hestrin S. Intracortical circuits of pyramidal neurons reflect their long-range axonal targets. *Nature*. 2009; 457:1133–6. <https://doi.org/10.1038/nature07658> PMID: 19151698
64. Rossi LF, Harris KD, Carandini M. Spatial connectivity matches direction selectivity in visual cortex. *Nature*. 2020; 588(7839):648–652. <https://doi.org/10.1038/s41586-020-2894-4> PMID: 33177719
65. Ding Z, Fahey PG, Papadopoulos S, Wang EY, Celii B, Papadopoulos C, et al. Functional connectomics reveals general wiring rule in mouse visual cortex. *bioRxiv*. 2023;.
66. Cossell L, Iacaruso M, Muir D, Houlton R, Sader E, Ko H, et al. Functional organization of excitatory synaptic strength in primary visual cortex. *Nature*. 2015; 518. <https://doi.org/10.1038/nature14182> PMID: 25652823

67. Smith G, Hein B, Whitney D, Fitzpatrick D, Kaschube M. Distributed network interactions and their emergence in developing neocortex. *Nature Neuroscience*. 2018; 21. <https://doi.org/10.1038/s41593-018-0247-5> PMID: 30349107
68. Lagzi F, Rotter S. Dynamics of Competition between Subnetworks of Spiking Neuronal Networks in the Balanced State. *PLOS ONE*. 2015; 10(9):1–28. <https://doi.org/10.1371/journal.pone.0138947> PMID: 26407178
69. Zeraati R, Shi YL, Steinmetz N, Gieselmann M, Thiele A, Moore T, et al. Intrinsic timescales in the visual cortex change with selective attention and reflect spatial connectivity. *Nature Communications*. 2023; 14. <https://doi.org/10.1038/s41467-023-37613-7> PMID: 37012299
70. Badin AS, Fermani F, Greenfield SA. The Features and Functions of Neuronal Assemblies: Possible Dependency on Mechanisms beyond Synaptic Transmission. *Frontiers in Neural Circuits*. 2017; 10. <https://doi.org/10.3389/fncir.2016.00114> PMID: 28119576
71. Umbach GS, Tan RJ, Jacobs J, Pfeiffer BE, Lega BC. Flexibility of functional neuronal assemblies supports human memory. *Nature Communications*. 2021; 13.
72. Hofer S, Ko H, Pichler B, Vogelstein J, Ros H, Zeng H, et al. Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nature neuroscience*. 2011; 14:1045–52. <https://doi.org/10.1038/nn.2876> PMID: 21765421
73. Levy RB, Reyes AD. Spatial Profile of Excitatory and Inhibitory Synaptic Connectivity in Mouse Primary Auditory Cortex. *Journal of Neuroscience*. 2012; 32(16):5609–5619. <https://doi.org/10.1523/JNEUROSCI.5158-11.2012> PMID: 22514322
74. Mongillo G, Rumpel S, Loewenstein Y. Inhibitory connectivity defines the realm of excitatory plasticity. *Nature neuroscience*. 2018; 21(10):1463–1470. <https://doi.org/10.1038/s41593-018-0226-x> PMID: 30224809
75. Zenke F, Hennequin G, Gerstner W. Synaptic Plasticity in Neural Networks Needs Homeostasis with a Fast Rate Detector. *PLOS Computational Biology*. 2013; 9(11):1–14. <https://doi.org/10.1371/journal.pcbi.1003330> PMID: 24244138
76. Cranmer K, Brehmer J, Louppe G. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*. 2020; p. 201912789. <https://doi.org/10.1073/pnas.1912789117> PMID: 32471948
77. Lueckmann JM, Boelts J, Greenberg D, Goncalves P, Macke J. Benchmarking simulation-based inference. In: *International Conference on Artificial Intelligence and Statistics*. PMLR; 2021. p. 343–351.
78. Gonçalves PJ, Lueckmann JM, Deistler M, Nonnenmacher M, Öcal K, Bassetto G, et al. Training deep neural density estimators to identify mechanistic models of neural dynamics. *eLife*. 2020; 9:e56261. <https://doi.org/10.7554/eLife.56261> PMID: 32940606
79. Bellec G, Scherr F, Subramoney A, Hajek E, Salaj D, Legenstein R, et al. A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications*. 2020; 11. <https://doi.org/10.1038/s41467-020-17236-y>
80. Zenke F, Vogels TP. The Remarkable Robustness of Surrogate Gradient Learning for Instilling Complex Function in Spiking Neural Networks. *Neural Computation*. 2021; 33(4):899–925. https://doi.org/10.1162/neco_a_01367 PMID: 33513328
81. Destexhe A, Rudolph M, Paré D. The high-conductance state of neocortical neurons in vivo. *Nature Reviews Neuroscience*. 2003; 4(9):739–751. <https://doi.org/10.1038/nrn1198> PMID: 12951566
82. Debanne D, Gähwiler BH, Thompson SM. Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *The Journal of physiology*. 1998; 507(Pt 1):237. <https://doi.org/10.1111/j.1469-7793.1998.237bu.x> PMID: 9490845
83. Feldman DE. Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron*. 2000; 27(1):45–56. [https://doi.org/10.1016/S0896-6273\(00\)00008-8](https://doi.org/10.1016/S0896-6273(00)00008-8) PMID: 10939330

Supplementary information

Structural influences on synaptic plasticity: the role of presynaptic connectivity in the emergence of E/I co-tuning

Emmanouil Giannakakis, Oleg Vinogradov, Victor Buendía, Anna Levina

Contents

A. E/I co-tuning and input selectivity	1
B. The effects of inhomogeneous connectivity on the population activity are independent of the neuron model	2
C. Firing behaviour of the presynaptic network	2
D. Learned connectivity and the dynamics of the post-synaptic neuron	4
E. The effects of inhomogeneous connectivity are independent of the plasticity protocols' details	5
i. Excitatory plasticity: Parameters of the Triplet rule	5
ii. Excitatory plasticity: Triplet vs Pair rule	7
iii. Inhibitory learning rule target rate	7
iv. Alternative implementations of the competitive Normalization	8
1. Regular vs event-based normalization steps	8
2. "Soft" vs. "Strict" normalization: The impact of the normalization rate	9
v. Subtractive Normalization with modified input	9
F. Convergence of weights to an eigenvector of a modified covariance matrix under plasticity	10
G. Perturbation of the optimal assembly strengths leads to diverse effects on the network's activity	11
H. Reduced model calculations	12
i. Derivation of the equations	12
ii. Solutions for the homogeneous network	14
I. Clustering optimisation	14
J. The inferred connectivity structure encourages competition between assemblies	16
K. Tables of parameters	17
References	18

A. E/I co-tuning and input selectivity

In our study, we use the weight diversity as a proxy for input selectivity. Verify that this mapping is reasonable even for networks with very detailed E/I balance. Specifically, we model 2 networks with fixed, perfectly balanced connectivity. The first of the two has tuned excitation (Fig Aa) and totally flat inhibition (leading to loose balance), while the second has near-perfect co-tuning of its incoming E and I incoming weights, which leads to tight balance (Fig Ab).

We sequentially activate different input groups and record the post-synaptic neuron's response in terms of incoming currents and firing rate. We see that while there is a small difference between the incoming currents in the two cases (Fig Ab and d), the firing rate of the postsynaptic neuron is significantly different between the two networks (Fig Ae).

In particular, we see that compared to the network with flat inhibition, the responses of the network with the co-tuned E/I weights are significantly reduced but still clearly different for different inputs, allowing discrimination. This is consistent with recent theoretical findings [1] about inhibitory tuning reducing input selectivity.

However, the co-tuned network displays a reduced firing rate, which suggests the possibility for more efficient encoding as well as a partial encoding for non-preferred inputs (i.e., a small but detectable response for the input group associated with small input weights) in contrast to the neuron with flat inhibition, which is fully silent for the non-preferred inputs (i.e., those for which it receives stronger inhibition than excitation). This agrees with earlier findings on the benefits of weights co-tuning in terms of efficiency and coding [2].

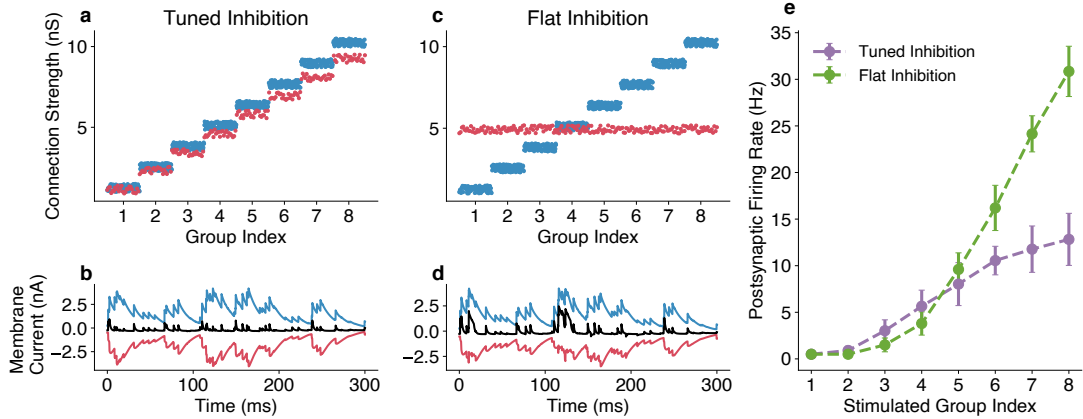


FIG. A. **E/I Weights and input selectivity:** **a.** The feedforward connectivity and **b.** the resulting postsynaptic currents for a network with flat co-tuned E/I feedforward weights feedforward inhibition. **c.** The feedforward connectivity and **d.** the resulting postsynaptic currents for a network with flat feedforward inhibition. **e.** The firing of the postsynaptic neuron upon activation of different input groups for each network. Flat inhibition leads to sharper responses while co-tuning leads to lower FR and a more gradual increase of activity for different input activation.

This suggests that while the detailed co-tuning of feedforward E/I weights might lead to a reduction in the sharpness of the post-synaptic neuron’s tuning, it still enables discrimination between inputs and, indeed, can encode for a broader range of inputs (responses even for non-preferred stimuli) in a more energetically efficient way (lower firing rate).

B. The effects of inhomogeneous connectivity on the population activity are independent of the neuron model

In order to ensure that our results are independent of the exact neuron model we are using, we repeated the bayesian fitting with a simplified network of LIF neurons with current based synapses.

Specifically, the voltage is given by:

$$\frac{dV(t)}{dt} = (V_{\text{rest}} - V(t))/\tau + \sum_j W_j^E \cdot \sum_f \delta(t - t_j^f) - \sum_j W_j^I \cdot \sum_f \delta(t - t_j^f), \quad (1)$$

where V_{rest} is the neuron’s resting potential and τ is the membrane timescale. Here t_j^f denotes the time at which the f -th spike of the j -th neuron happened. When the membrane potential reaches the spiking threshold V_{th} , a spike is emitted, and the potential is changed to a reset potential V_{reset} .

We repeat the fitting with ABC for a network with this model, using the same objective function (maximizing in-group and minimizing between-group correlations) and we find that the optimal assembly strength distribution for all four connection types we identified for the original network remains very similar to the ones we identified for the original model.

C. Firing behaviour of the presynaptic network

In order to better understand the effects that the introduction of noise and non-plastic recurrent connections has on the population activity of the presynaptic network, we calculate a number of metrics that quantify population activity for different connectivity structures.

At first, we look at a feedforward network with relatively low noise (noise intensity is set to 0.15). This setting produces activity that is ideal for the plasticity to produce diverse and co-tuned E/I weights.

As expected, due to the Poisson input and the absence of any recurrent interactions, the networks have a broad ISI distribution, a relatively narrow distribution of the CV of the ISI and a Fano Factor very close to 1 (Fig Ba-c). The activity of the network is highly correlated within groups, but different groups fire independently (Fig Bd).

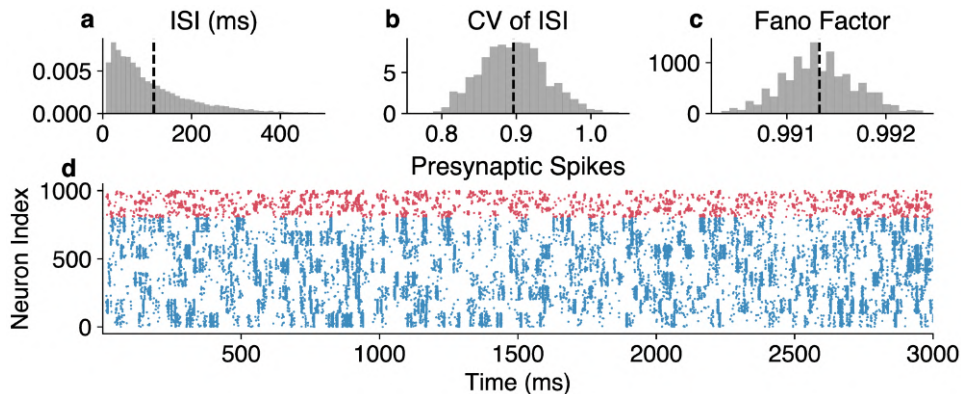


FIG. B. **The firing of a feedforward network:** **a.** The inter-spike interval distribution. **b.** The distribution of the Coefficient of variation ((CVs) of the ISIs) **(c)**The distribution of the Fano Factor is very close to 1, indicating the Poisson input the neurons receive. **(d)** A raster plot visualizing 3 seconds of the network's activity (Here blue indicates excitatory and red inhibitory neurons)

We then examine a case that is particularly detrimental to the emergence of co-tuning. We introduce unstructured (no assemblies) recurrent connectivity ($p = 0.5$ and $W = 2$) and high noise (noise intensity is set to 0.6).

The resulting distribution of the inter-spike intervals becomes somewhat skewed, as does the distribution of the CV of the ISI, while the Fano Factor remains very close to 1 (Fig Ca-c). The network develops occasional large bursts of synchronized activity involving neurons from all groups (Fig Cd).

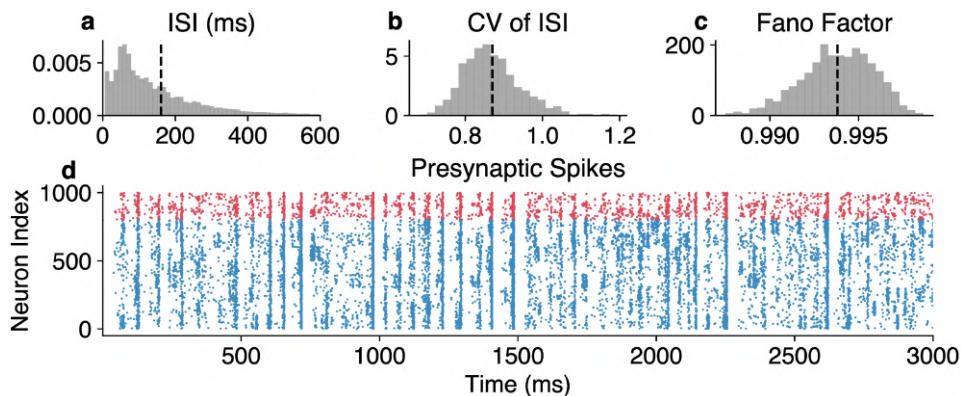


FIG. C. **The firing of a network with strong noise (noise intensity = 0.6) and strong random recurrence ($p = 0.5$ and $W = 2$):** **a.** The inter-spike interval distribution. **b.** The distribution of the Coefficient of variation ((CVs) of the ISIs) **(c)**The distribution of the Fano Factor remains very close to 1. **(d)** A raster plot visualizing 3 seconds of the network's activity (Here blue indicates excitatory and red inhibitory neurons)

Finally, we look at the previous network, albeit with the introduction of optimal assembly structure ($r_{EE} = r_{EI} = 1.0$ and $r_{IE} = r_{II} = 0.4$).

This network is characterized by short synchronized firing periods of each group, during which neurons of other groups are largely silent (Fig Dd). This leads to a bimodal distribution of the mean ISI and a very skewed distribution of the CV of the ISI, which indicates relatively irregular firing (largely due to the combination of bursts and longer periods of silence). Finally, the Fano Factor remains close to 1 (Fig Da-c).




FIG. D. **The firing of a network with near-optimal clustering** ($r_{EE} = r_{EI} = 1.0$ and $r_{IE} = r_{II} = 0.4$): **a.** The inter-spike interval distribution. **b.** The distribution of the Coefficient of variation ((CVs) of the ISIs. The average increases slightly compared to the previous networks. **(c)**The distribution of the Fano Factor remains very close to 1. **(d)** A raster plot visualizing 3 seconds of the network’s activity (Here blue indicates excitatory and red inhibitory neurons)

D. Learned connectivity and the dynamics of the post-synaptic neuron

Since synaptic plasticity depends not only on presynaptic network statistics but also on postsynaptic neuron activity, we examine the learned connectivity of different networks (the ones we examined in the last section) and visualize the activity of the postsynaptic neuron.

At first, the feedforward, low-noise network leads to a very strongly co-tuned E/I connectivity (Fig Ea) and diverse weights between different groups. The incoming currents are highly correlated (Fig Eb) and the post-synaptic neuron fires (Fig Ec) relatively sparsely ($\langle \text{ISI} \rangle = 0.398$, $\text{CV}_{\text{ISI}} = 1.15$, Fano Factor = 0.991).

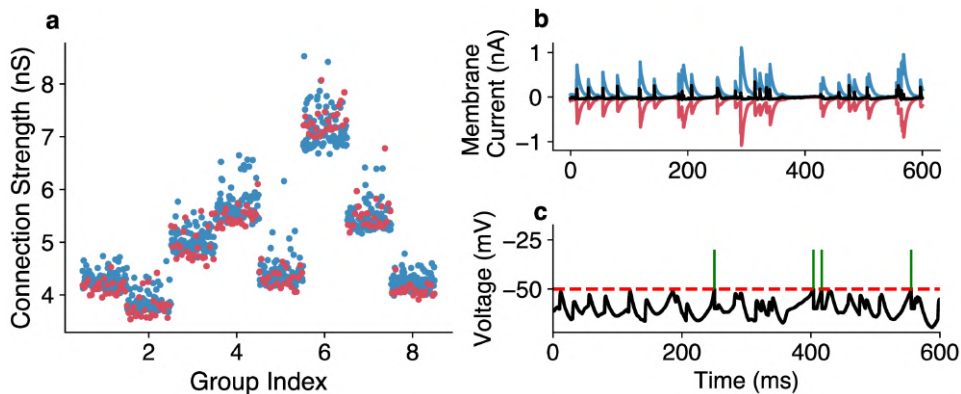


FIG. E. **Connectivity and firing of a postsynaptic neuron receiving input from a feedforward network:** **a.** The learned connectivity is co-tuned and diverse. **b.** The incoming currents are tightly balanced. **c** The voltage trace and spikes of the readout neuron.

On the contrary, the noisy and recurrent network fails to develop any structure in the learned connectivity. The coordinated firing across groups, visualized in (Fig Fa), leads to occasional large current influxes (Fig Fb), but otherwise, the neuron remains relatively tightly balanced. The neuron’s firing (Fig Fc) becomes slightly more frequent ($\langle \text{ISI} \rangle = 0.21$, $\text{CV}_{\text{ISI}} = 1.23$, Fano Factor = 0.991).

Finally, the network with the optimized assembly structure, due to the restored statistics of the input (Fig Ga), develops co-tuned E/I connectivity and relatively diverse weights between groups. The currents arriving to the postsynaptic neurons are tightly balanced (Fig Gb) and the firing of the postsynaptic neuron (Fig Gc) becomes again more sparse and relatively irregular ($\langle \text{ISI} \rangle = 0.61$, $\text{CV}_{\text{ISI}} = 1.14$, Fano Factor = 0.993).

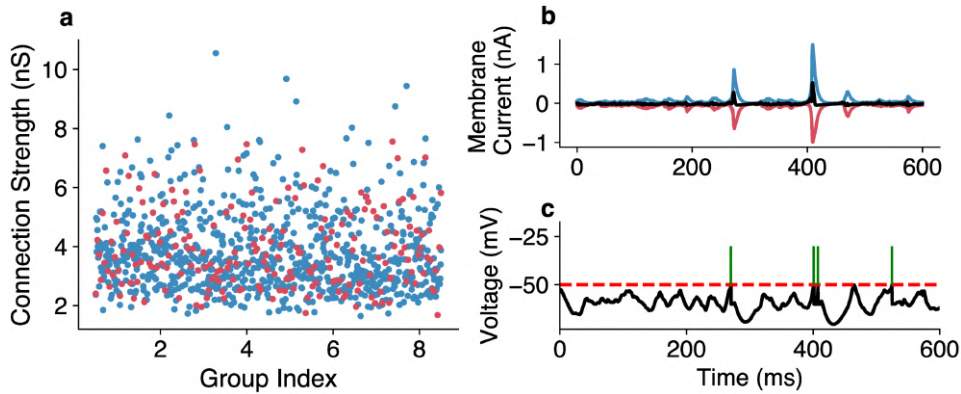


FIG. F. **Connectivity and firing of a postsynaptic neuron receiving input from a noisy, recurrent network:** **a.** The learned connectivity is completely unstructured. **b.** The incoming currents are tightly balanced despite occasional large influxes. **c** The voltage trace and spikes of the readout neuron.

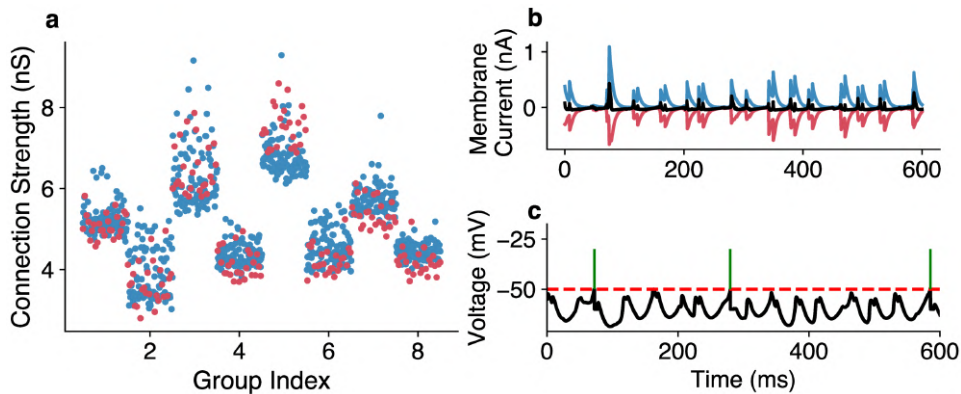


FIG. G. **Connectivity and firing of a postsynaptic neuron receiving input from a network with optimal assembly structure:** **a.** The learned connectivity is restored to be co-tuned and diverse. **b.** The incoming currents are tightly balanced. **c** The voltage trace and spikes of the readout neuron.

E. The effects of inhomogeneous connectivity are independent of the plasticity protocols' details

We examined whether our results are dependent on the particular plasticity protocol we used, and we verified that they hold for alternative plasticity mechanisms. Specifically, we tested whether a variety of different plasticity protocols produces co-tuning in the simple feedforward case, whether the effects of noise and recurrence on the resulting connectivity are consistent, and whether the inferred optimal fixed pre-synaptic recurrent connectivity restores the ability of the plasticity to produce co-tuning.

Starting with the plasticity protocol we used in our original experiments, we visualize the development of excitatory (Fig Ha) and inhibitory weights (Fig Hb), the matching resulting connectivity (Fig Hc) and the incoming E/I currents to the post-synaptic neuron after convergence (Fig Hd) in a setting with optimal connectivity for comparison with the other plasticity protocols simulated with the same connectivity and noise levels.

i. Excitatory plasticity: Parameters of the Triplet rule

In our experiments presented in the main text, we have used a simplified form of the triplet STDP rule for the excitatory synapses, which relies on identical slow and fast timescales for the pre and postsynaptic traces. We first tested how robust our results are to changes on the LTD/LTP ratio (varying the $A_{LTD} \in [0.05, 1.2]$ and $A_{LTP} \in [0.05, 1.0]$) as well as the timescale of the fast and slow traces. We found that qualitatively our results are robust to these changes in the plasticity protocol.

Moreover, in order to ensure that our results did not depend on the particular form of the rule we used, we

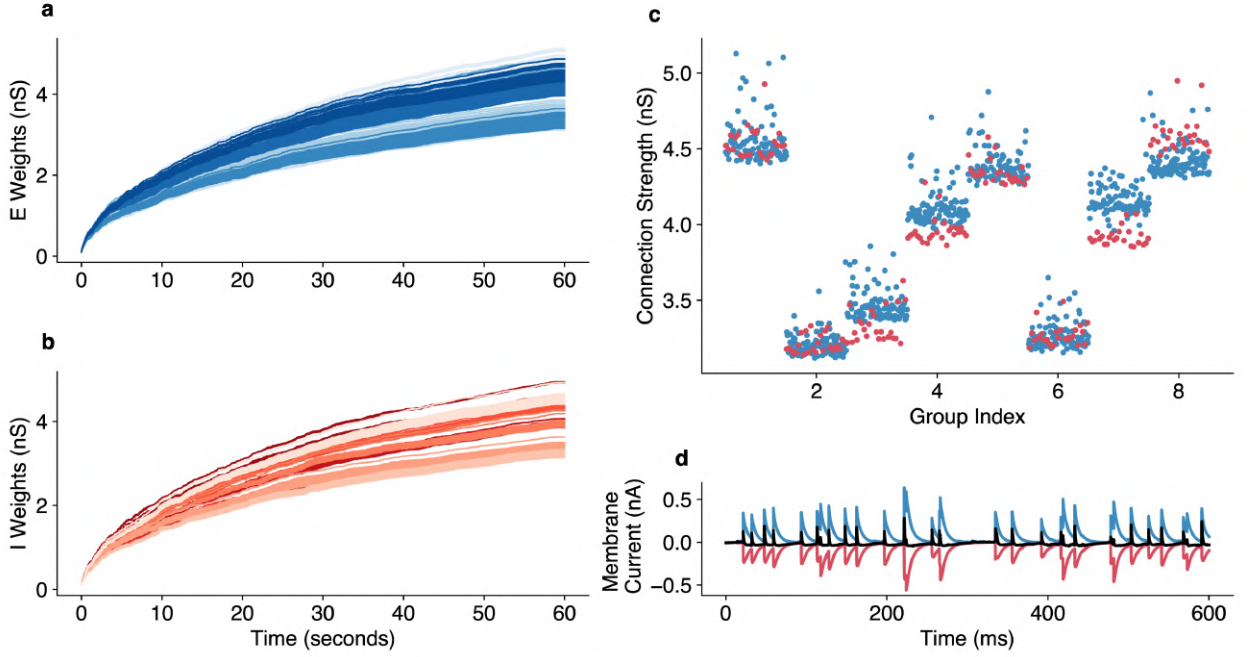


FIG. H. **Weight development under the original protocol.** The development of **a.** excitatory and **b.** inhibitory weights over a minute of simulation time. **c.** The convergence point of excitatory (blue) and inhibitory weights (red). The inhibitory weights are scaled to account for the slower synaptic constant of the inhibitory synapses and the smaller number of inhibitory neurons. **d.** The resulting E/I currents on the postsynaptic neuron are tightly balanced.

replicated our main simulations using the original implementation of the triplet rule from [3], which implemented different timescales for the pre and postsynaptic fast and slow traces.

Specifically, in this implementation, the pre and post-synaptic activity is tracked by the traces:

$$\frac{dy_k^E(t)}{dt} = \frac{-y_k^E(t)}{\tau_{\text{pre}}^{\text{fast}}} + \sum_f \delta(t - t_k^f), \quad (2a)$$

$$\frac{dz_k^E(t)}{dt} = \frac{-z_k^E(t)}{\tau_{\text{pre}}^{\text{slow}}} + \sum_f \delta(t - t_k^f), \quad (2b)$$

$$\frac{dx_1(t)}{dt} = \frac{-x_1(t)}{\tau_{\text{post}}^{\text{fast}}} + \sum_f \delta(t - t_x^f), \quad (2c)$$

$$\frac{dx_2(t)}{dt} = \frac{-x_2(t)}{\tau_{\text{post}}^{\text{slow}}} + \sum_f \delta(t - t_x^f), \quad (2d)$$

where $\tau_{\text{pre}}^{\text{fast}} = 16.8$ ms, $\tau_{\text{pre}}^{\text{slow}} = 101$ ms, $\tau_{\text{post}}^{\text{fast}} = 33.7$ ms and $\tau_{\text{post}}^{\text{slow}} = 125$ ms following the parameters from [3, 4]. Additionally, $y_k^E(t)$, $z_k^E(t)$ and $x_1(t)$, $x_2(t)$ represent the slow and fast traces of the k -th excitatory pre-synaptic and the single post-synaptic neuron respectively while t_k^f and t_x^f are their respective firing times. The function $\delta(x)$ represents a Dirac's delta. The connection weights are updated upon pre and post-synaptic spiking according to

$$\Delta W_k^E = \eta_E \cdot y_k^E(t) \cdot (A_2^+ + A_3^+ \cdot x_2(t)) \cdot \sum_f R(t - t_k^f) - \eta_E \cdot x_1(t) \cdot (A_2^- + A_3^- \cdot z_k^E(t)) \sum_f R(t - t_k^f), \quad (3)$$

where $A_2^- = 0.7$, $A_3^- = 0.023$, $A_2^+ = 7.5 \cdot 10^{-8}$, $A_3^+ = 0.93$ and $\eta_E = 10^{-2}$, following the parameters of [3, 4]. We see that under this protocol (Fig Ia, b), the weights develop very similarly to the one we used for our experiments (Fig Ha, b) under optimal connectivity. Moreover, we see that strong co-tuning emerges (Fig Ic) as well as tightly balanced E/I currents to the postsynaptic neuron as expected (Fig Id).

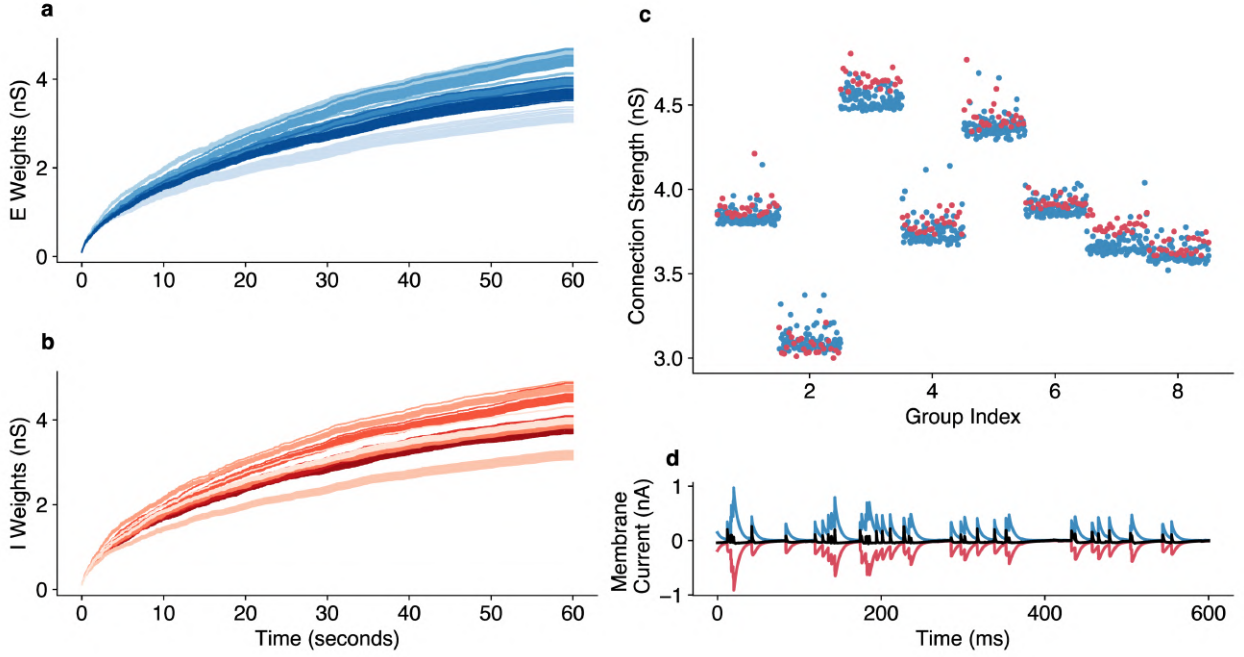


FIG. 1. **Weight development with the original triplet rule parameters.** The development of **a.** excitatory and **b.** inhibitory weights over a minute of simulation time. **c.** The convergence point of excitatory (blue) and inhibitory weights (red). The inhibitory weights are scaled to account for the slower synaptic constant of the inhibitory synapses and the smaller number of inhibitory neurons. **d.** The resulting E/I currents on the postsynaptic neuron are tightly balanced.

ii. Excitatory plasticity: Triplet vs Pair rule

We further examined whether a different Hebbian learning rule in the excitatory synapses has any impact on our results. In particular, we replaced the triplet rule [3] with a classic spike pair Hebbian rule [5], which relies on a single pre and postsynaptic trace:

$$\frac{dy_k^E(t)}{dt} = \frac{-y_k^E(t)}{\tau} + \sum_f \delta(t - t_k^f), \quad (4a)$$

$$\frac{dx(t)}{dt} = \frac{-x(t)}{\tau} + \sum_f \delta(t - t_x^f), \quad (4b)$$

where $y_k^E(t)$ the trace of the k -th presynaptic excitatory neuron, $x(t)$ the trace of the postsynaptic neuron and $\tau = 10$ ms the time constant of the decay. The weight update happens as:

$$\Delta W_k^E = \eta_E \cdot A_{LTP} \cdot y_k^E(t) \cdot \sum_f R(t - t_x^f) - \eta_E \cdot A_{LTD} \cdot x(t) \sum_f R(t - t_k^f), \quad (5)$$

where $A_{LTP} = 1.0$ and $A_{LTD} = 0.2$ same as for the simplified triplet rule. Using the same competitive, synapse-type specific normalization, we found that in this setting, the weight development as well as the impact of noise and recurrence are similar to our original setting.

iii. Inhibitory learning rule target rate

We additionally examined several different target rates ρ_0 (ranging from 0 to 6 Hz) for the inhibitory plasticity rule (setting the rate to $\rho_0 = 0$, i.e., making the inhibitory plasticity into a pure symmetric Hebbian rule, as was used in

the original study of the normalization mechanism by [6]). Besides an expected change in the postsynaptic firing rate after the convergence of the weights, we did not observe any other noticeable changes in our findings regarding the emergence of co-tuning and input selectivity.

iv. Alternative implementations of the competitive Normalization

1. Regular vs event-based normalization steps

For reasons of numerical simulation speed, we implemented the weight normalisation in an asynchronous manner. Specifically, we apply a normalisation step after every weight update occurs (i.e., after each postsynaptic spike all connections get normalised and after each presynaptic spike the corresponding connection gets normalised). Since the presynaptic population has a homogeneous firing rate (i.e., all neurons spike approximately the same number of times) all connections are normalised approximately equally often, which makes this implementation behave similarly to implementing the normalisation step on each time step of the simulation.

To test that this assumption is correct, we simulated a network with a regular normalization update, where all connections are normalized simultaneously every time step ($dt = 0.1$ ms). Besides a modification of the rate of the normalization η_N (necessary to counter the higher frequency of normalization in the regular case), we maintain the exact same parameters as in the original experiment.

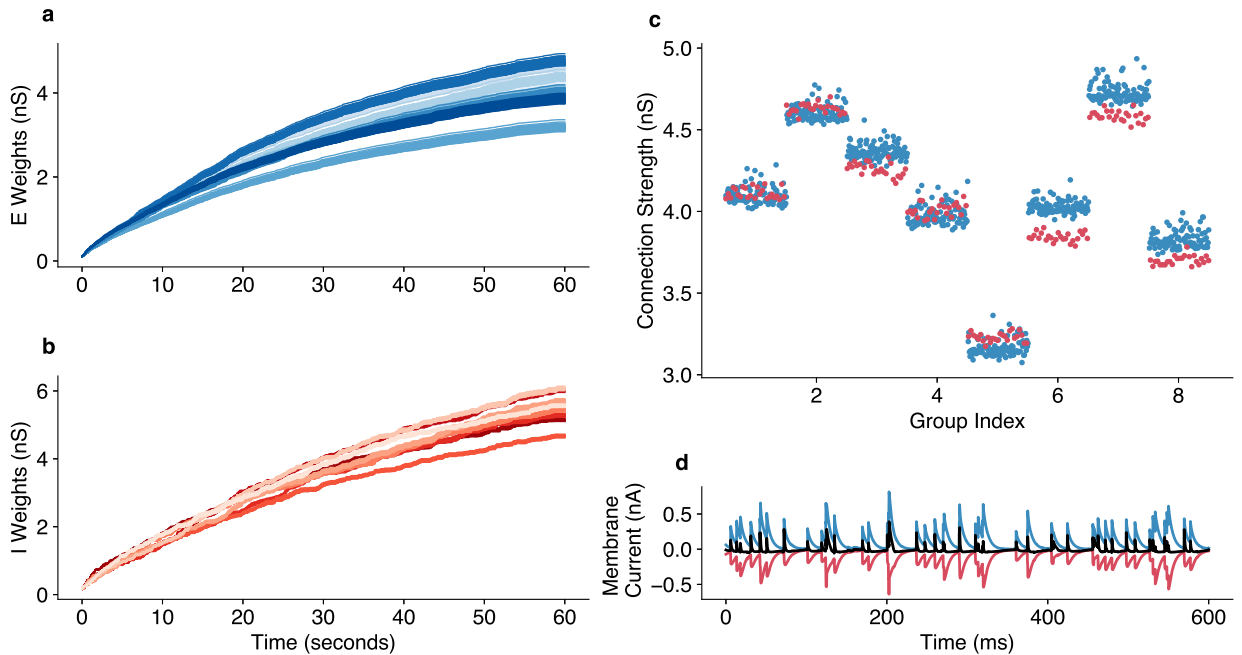


FIG. J. **Weight development with a regular normalization update.** The development of **a.** excitatory and **b.** inhibitory weights over a minute of simulation time. **c.** The convergence point of excitatory (blue) and inhibitory weights (red). The inhibitory weights are scaled to account for the slower synaptic constant of the inhibitory synapses and the smaller number of inhibitory neurons. **d.** The resulting E/I currents on the postsynaptic neuron are tightly balanced.

We observe that in this setting, the weight development happens very similar to the original experiment (Fig Ja, b), producing co-tuning in cases with optimal input connectivity (Fig Jc) and tight E/I current balance (Fig Jd). In summary, we find that the two normalization methods behave similarly for our setting, suggesting that our findings are independent of how the weight normalization is implemented. However, in settings different from ours, where the pre-synaptic rates are largely inhomogeneous, we would expect that the two normalization methods would produce different results (since some weights would be normalized more often than others), which might lead to different weight dynamics. Thus, the numerical convenience of the asynchronous normalization updates cannot be generally used in simulations of other types of networks with potentially very different distributions of firing rates.

2. “Soft” vs. “Strict” normalization: The impact of the normalization rate

As described in the “Methods” section of the main text, we implement our normalization step as follows:

$$W_k^A(t) \leftarrow (1 - \eta_N) \cdot W_k^A(t) + \eta_N \cdot W_k^A(t) \cdot \frac{W_{target}^A}{\sum_{i=1}^{N_A} W_i^A(t)}, \quad A \in \{E, I\}. \quad (6)$$

where $\eta_N = 3 \cdot 10^{-3}$ to match the equivalent excitatory and inhibitory learning rates. This essentially implements a “soft” normalization, where the total sum of the incoming weights to the postsynaptic neuron is not kept constant in each timestep but rather pushed towards maintaining a sum as close as possible to the target weight sum over time.

However, if we treated the normalization strictly as the preservation of synaptic resources over time following [6], we would need to set $\eta_N = 1$, in order to preserve the exact weight sum constant for each time step. This approach would impose a very strict condition, which, in combination with the small learning rates of the E and I plasticity, would prevent weight diversification.

One way of countering this is by massively increasing the E and I learning rates, but this tends to make the learning dynamics unstable. Another way of solving this problem is by implementing the normalisation regularly but not on every time step (a solution that has been used previously in [4]), for example, implementing it every 20 ms. This does lead to stable learning, but it also occasionally promotes winner-take-all connectivity, which is not in itself problematic, but may be undesirable for some types of coding. The emergence of winner-take-all connectivity can, in turn, be potentially countered by faster inhibitory plasticity relative to the excitatory plasticity, but finding the exact parameters can involve a fair amount of fine-tuning. For our experiment, the “soft” which can be biologically justified as a competition for synaptic resources among multiple neurons, happening on a very slow timescale, clearly leads to more plausible and stable results.

v. Subtractive Normalization with modified input

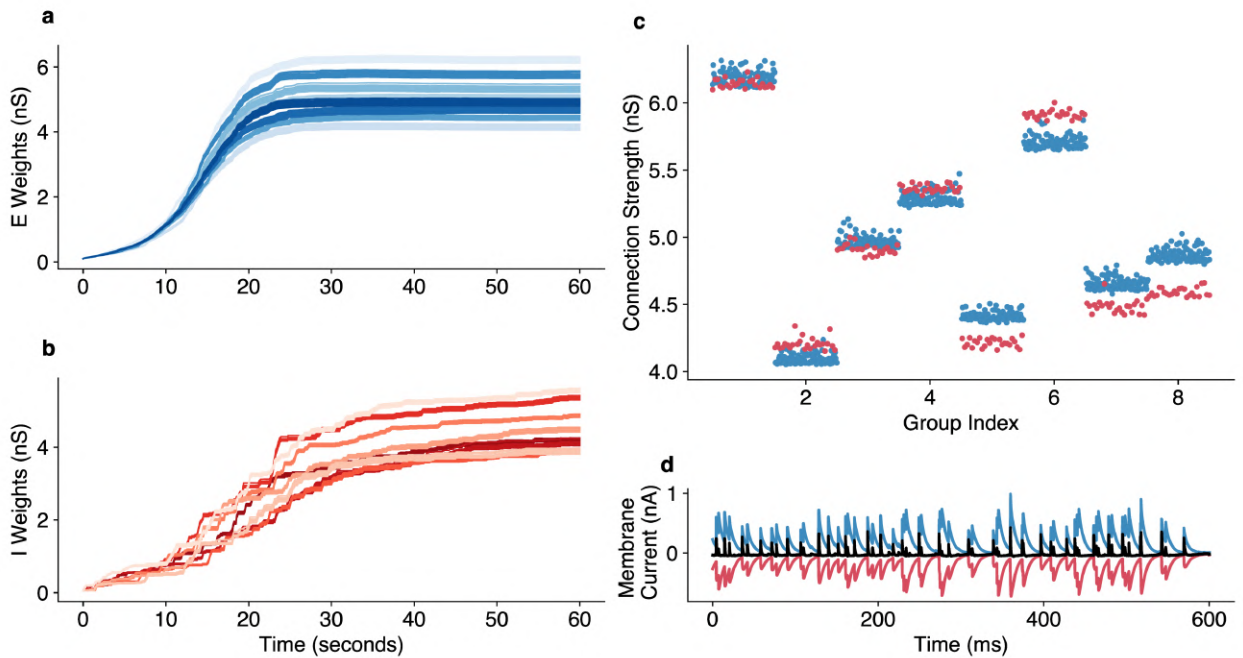


FIG. K. **Alternative normalization protocol.** The development of **a.** excitatory and **b.** inhibitory weights over a minute of simulation time. **c.** The convergence point of excitatory (blue) and inhibitory weights (red). The inhibitory weights are scaled to account for the slower synaptic constant of the inhibitory synapses and the smaller number of inhibitory neurons. **d.** The resulting E/I currents on the postsynaptic neuron are tightly balanced.

Finally, in order to fully demonstrate that our results are independent of the exact normalization protocol, we studied both the triplet [3] and pair [5] excitatory Hebbian rules combined with a subtractive normalization mechanism that

has been previously used in plasticity studies [4, 7]. Following our setting for the multiplicative normalization, we applied this rule also as a “soft” normalization in the excitatory synapses (the inhibitory synapses are not normalized in this setting):

$$W_k^E(t) \leftarrow (1 - \eta_N) \cdot W_k^E(t) + \eta_N \cdot W_k^E(t) \cdot \frac{\sum_{i=1}^{N_E} W_i^E - W^{\text{target}}}{N_E}. \quad (7)$$

In 2016 [8], it was demonstrated that subtractive normalization only on the excitatory connections will lead to all the weights converging on the same point due to the inhibitory plasticity creating a moving threshold. In order to prevent this collapse of the receptive field, enforced inhomogeneity on the firing rates of different groups is needed. We solved this problem by giving the network’s input as pulses of 100 mS during which some of the input groups firing rate quadruples. This enforces inhomogeneous firing rates, which result in the emergence of stable, diverse, and co-tuned feedforward connectivity (Fig Ka - c). We also verified that the post-synaptic neuron is tightly balanced (Fig Kd) and maintains a stable firing rate.

We verified that the co-tuning achieved in this setting, similar to the mechanism presented in the main text, suffers from the introduction of noise and recurrent connectivity. Furthermore, the assembling principles that we derived for the original network seem to have a similarly beneficial effect on this setting, restoring the original covariance structure of the network’s activity and leading to detailed co-tuning between the excitatory and inhibitory feedforward connections.

F. Convergence of weights to an eigenvector of a modified covariance matrix under plasticity

The study that introduced the competitive synapse type-specific normalization [6] analytically predicted that for Hebbian E and I feedforward plasticity, the convergence point of the weights is an eigenvector of the modified covariance matrix (Fig La):

$$\bar{C} = \left\langle \begin{pmatrix} EE^T & -EI^T \\ IE^T & -II^T \end{pmatrix} \right\rangle \quad (8)$$

where E, I are the activities of the excitatory and inhibitory populations, respectively.

We test whether in our modified plasticity protocol (simplified Triplet STDP in E connections, slower normalization), the convergence point can be approximated by an eigenvector of the above matrix.

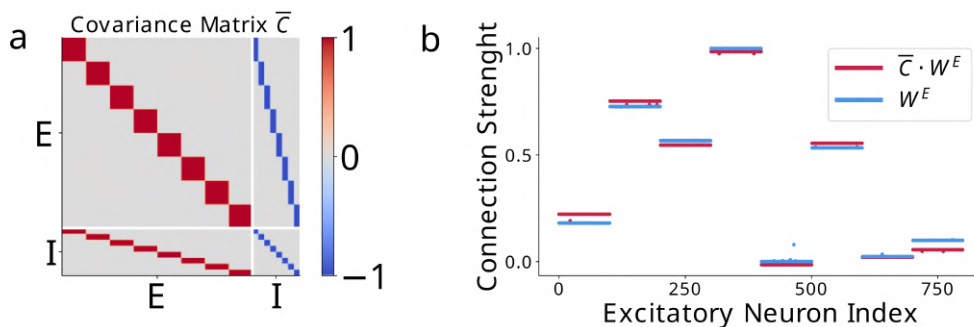


FIG. L. **Weights converge to an eigenvector of the covariance matrix.** **a:** The estimated covariance matrix for a feedforward network. **b:** We verify that the convergence point of the weights is an eigenvector of the covariance matrix.

Specifically, we multiply the converged weight vector with a numerical calculation of the covariance (estimated via binning of the spike trains with a bin size of 1 mS), for different noise and recurrence settings, and we verify that the resulting product is approximately equal to a multiple of the original weight matrix (Fig Lb). This indicates that despite the differences in the learning protocol, the plasticity converge point is largely controlled by the covariance structure of the population activity.

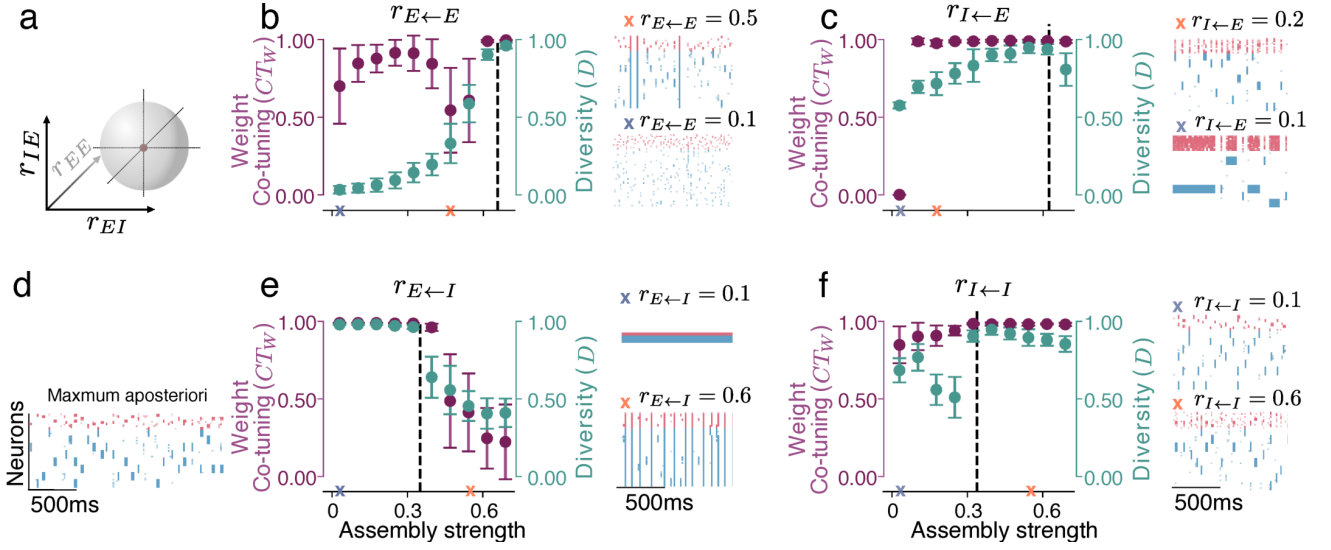


FIG. M. **Changes of the various assemblies strengths differently affect weight co-tuning and diversity.** **a.** We sequentially vary the value of each assembly strength while keeping the rest of the parameters fixed at the maximum a posteriori (MAP) solution (**d**). **b.** A decrease in the $E \leftarrow E$, assembly strength ($R_{E \leftarrow E}$) introduces synchronous burst-like events that jeopardize co-tuning and weight diversity. Further reduction of the $E \leftarrow E$ assembly strength results in sparse, asynchronous spiking, which significantly reduces tuning quality. **c.** Reduction of $I \leftarrow E$ assembly strength first leads to synchronous inhibitory firing across groups, and further reduction leads to persistent activity of the whole inhibitory population combined with bursts of excitatory activity that prevent the development of diversity. **e.** Decreasing the $E \leftarrow I$ assembly strength leads to persistent activation of a single group of neurons (which affects the tuning only marginally). While the increase of the $E \leftarrow I$ assembly strength leads to synchronized behavior of the whole network. **f.** Weakening the $I \leftarrow I$ assemblies decreases the weight diversity by introducing occasional synchronous bursts in the network while strengthening them leads to asynchronous inhibitory activity.

G. Perturbation of the optimal assembly strengths leads to diverse effects on the network's activity

To investigate the relative importance of the different connection types, we perturb various assemblies away from the optimal solutions inferred with ABC and check the resulting changes on the network activity as well as how fast the weight co-tuning and diversity deteriorate.

We find that $E \leftarrow I$ assemblies have a strong impact on the resulting network dynamics and, consequently, the statistics that determine the emergence of co-tuning (Fig Me). Thus, when $E \leftarrow I$ assemblies are too strong, the network exhibits synchronous behaviour that does not allow discrimination between neuron groups and, consequently, the emergence of co-tuned connectivity. In contrast, when $E \leftarrow I$ assemblies are too weak, one group of neurons is constantly active. This eventually leads to perfect discrimination of only a single input.

$E \leftarrow E$ assemblies also have a strong effect on the resulting learned connectivity. When they are reduced, the weight diversity rapidly reduces, while the weight co-tuning metric first decreases and then restores. This behaviour is related to two different bifurcations in the population dynamics. First, synchronous full-network bursts emerge on top of the stimulus-induced firing (Fig Ma). Then, as $E \leftarrow E$ are reduced further, the network enters an asynchronous irregular state (Fig Ma).

Perturbations of $I \leftarrow I$ and $I \leftarrow E$ assemblies have a weaker effect on the weight co-tuning and diversity. Specifically, when $I \leftarrow I$ assemblies are weakened, the network displays synchronous bursts on top of the stimulus-driven activity that only minimally reduces the weight co-tuning and diversity (Fig Md). When $I \leftarrow E$ assemblies are reduced, the weight diversity slowly decreases due to a decrease in the diversity of inhibitory weights. Thus, the activity of the inhibitory population becomes more disorganized (Fig Mc). On the other hand, a network with uniform $I \leftarrow E$ connections (i.e., without $I \leftarrow E$ assemblies), shows a dramatic reduction in the weight co-tuning metric because of very strong activation of a single neuron group accompanied by the simultaneous activity of all inhibitory neurons (Fig Mc).

H. Reduced model calculations

i. Derivation of the equations

In this Appendix, we start from the main text Eqs. 14,

$$\dot{x}_i = ax_i + by_i + \frac{1}{M-1} \sum_{j \neq i} (W^{E \leftarrow E} x_j + W^{E \leftarrow I} y_j) + \sigma_{\text{int}} \xi_i^x(t) + \sigma_{\text{ext}} \eta_i(t) + \eta_0, \quad (9a)$$

$$\dot{y}_i = cx_i + dy_i + \frac{1}{M-1} \sum_{j \neq i} (W^{I \leftarrow E} x_j + W^{I \leftarrow I} y_j) + \sigma_{\text{int}} \xi_i^y(t) + \sigma_{\text{ext}} \eta_i(t) + \eta_0. \quad (9b)$$

where we consider M groups composed by excitatory $x_i(t)$ and inhibitory $y_i(t)$ populations ($i = 1, \dots, n$) coupled linearly (see also Methods). Internal noise of each population is represented by $\xi_i^x(t)$ and $\xi_i^y(t)$ for excitatory and inhibitory populations, respectively. On the other hand, external noise $\eta_i(t)$ is shared among both populations. All the noises have zero mean and correlations

$$\langle \xi_i^c \xi_j^{c'} \rangle = \delta_{cc'} \delta_{ij} \delta(t - t'), \quad (10a)$$

$$\langle \xi_i^c(t) \eta_j(t') \rangle = 0 \quad \forall i, j, t, t', \quad (10b)$$

$$\langle \eta_i(t) \eta_j(t') \rangle = \delta_{ij} \delta(t - t'). \quad (10c)$$

We will derive closed expressions for the correlation coefficients. We would like to remark that $\langle \cdot \rangle$ means an ensemble average over noise realizations. All stochastic equations are to be interpreted in the Itô convention[9].

First of all, we redefine the noise terms, which will prove convenient later to simplify the algebra. For this reason, we define

$$\xi_i^1(t) = \sigma_{\text{int}} \xi_i^x(t) + \sigma_{\text{ext}} \eta_i(t), \quad (11a)$$

$$\xi_i^2(t) = \sigma_{\text{int}} \xi_i^y(t) + \sigma_{\text{ext}} \eta_i(t), \quad (11b)$$

which are Gaussian white noises with zero mean and correlation matrix

$$\langle \xi_i^1(t) \xi_j^1(t') \rangle = \langle \xi_i^2(t) \xi_j^2(t') \rangle = \delta_{ij} \delta(t - t') (\sigma_{\text{int}}^2 + \sigma_{\text{ext}}^2), \quad (12a)$$

$$\langle \xi_i^1(t) \xi_j^2(t') \rangle = \delta_{ij} \delta(t - t') \sigma_{\text{ext}}^2. \quad (12b)$$

To start with, one can obtain the average values for the stationary rates by applying averages to both sides of equations (9) and imposing the stationary state condition, $\langle \dot{x}_i \rangle = \langle \dot{y}_i \rangle = 0$. Once this is done, it is immediate to solve the resulting linear system and check that $\langle x_i^* \rangle, \langle y_i^* \rangle \propto \eta_0$, where the star (*) indicates that these values correspond to the stationary state. Hence, making $\eta_0 = 0$ the mean values vanish. One can demonstrate that correlations do not depend on η_0 , and hence we can make $\eta_0 = 0$ without loss of generality. Conceptually, this means just shifting up or down the baseline of fluctuations of the firing rate, which does not affect the fluctuations themselves.

To compute correlations, we need to evaluate the second-order moments between different populations as $\langle x_i y_j \rangle$ or $\langle x_i^2 \rangle$. A possible way of doing this is starting from the analytical solution of the multidimensional Orstein-Uhlenbeck process [10]. However, this approach will yield a linear system with $N(N+1)/2$ variables to solve for, which are all the elements of the (symmetric) correlation matrix. But all the groups are identical (or *indistinguishable*), so we expect correlations not to depend on the particular population chosen. Therefore, all the equations will be reduced to just 6 covariances: $\langle x_i x_j \rangle, \langle x_i y_j \rangle, \langle y_i y_j \rangle, \langle x_i^2 \rangle, \langle x_i y_i \rangle$ and $\langle y_i^2 \rangle$.

In this context, it is conceptually simpler to obtain equations for the evolution of the second moments and then evaluate them in the stationary state. Here, we report in detail the computation of two of these moments as an example, giving just the final answer for the other four, which is performed in an analogous way.

First, we define $X_{ij} = x_i x_j$, and then look for the time evolution of X_{ij} , i.e., \dot{X}_{ij} . Notice that this is a non-linear change of variables, and thus Itô's lemma is required. The lemma tells us that if we have a change of variables $z = z(x)$, then one has to include the second-order terms in the expansion,

$$dz = \underbrace{\sum_{i=1}^N \partial_{x_i} z dx_i}_{\text{Chain rule}} + \frac{1}{2} \underbrace{\sum_{i=1}^N \partial_{x_i} \partial_{x_j} z dx_i dx_j}_{\text{Itô's lemma}}. \quad (13)$$

The terms dx_i can be obtained as $\dot{x}_i dt$. It is important to remark that in this procedure noise terms are rewritten as the differential of the Wiener processes, i.e., $\eta_i(t)dt = dW_i$. After applying the Itô lemma above, only terms up to order dt should be taken into account. Notice that $dW_i(t) \propto \sqrt{dt}$ [10]. Finally, one just divides again by dt to recover the stochastic differential equation and applies the ensemble average.

For X_{ij} , this reads as

$$\begin{aligned} \frac{d\langle x_i x_j \rangle}{dt} &= \langle \dot{x}_i x_j \rangle + \langle x_i \dot{x}_j \rangle + \frac{1}{2} \langle \dot{x}_i \dot{x}_j \rangle = \\ &= a \langle x_i x_j \rangle + b \langle y_i x_j \rangle + \frac{W^{E \leftarrow E}}{M-1} \sum_{k \neq i} \langle x_k x_j \rangle + \frac{W^{E \leftarrow I}}{M-1} \sum_{k \neq i} \langle y_k x_j \rangle + \\ &+ a \langle x_i x_j \rangle + b \langle x_i y_j \rangle + \frac{W^{E \leftarrow E}}{M-1} \sum_{k \neq j} \langle x_i x_k \rangle + \frac{W^{E \leftarrow I}}{M-1} \sum_{k \neq j} \langle x_i y_k \rangle + \\ &\quad \langle \xi_i^1 x_j \rangle + \langle x_i \xi_j^1 \rangle + \langle \xi_i^1 \xi_j^1 \rangle + \mathcal{O}(dt^2), \end{aligned} \quad (14)$$

where all the averages between the noise and the variable yield 0, due to Itô's prescription. The next step is to simplify the sums involving correlations. As discussed above, since clusters are indistinguishable, all the terms are exactly the same. However, the dummy index k will also take the value of the fixed index, $k = j$, and this has to be taken into account separately since the in-group is different to the between-group one. Then,

$$\sum_{k \neq i} \langle x_k x_j \rangle = (M-2) \langle x_i x_j \rangle + \langle x_i^2 \rangle, \quad (15)$$

allowing us to simplify the equation. At this step we simplify the notation by letting $X_{ij} = \langle x_i x_j \rangle$, $Z_{ij} = \langle x_i y_j \rangle$, $Y_i = \langle y_i^2 \rangle$, etc., leading to

$$\frac{1}{2} \dot{X}_{ij} = \left(a + \frac{M-2}{M-1} W^{E \leftarrow E} \right) X_{ij} + \left(b + \frac{M-2}{M-1} W^{E \leftarrow I} \right) Z_{ij} + \frac{1}{M-1} [W^{E \leftarrow E} X_i + W^{E \leftarrow I} Z_i]. \quad (16)$$

The same procedure can be repeated for all the other correlations, such as

$$\begin{aligned} \frac{d\langle y_i^2 \rangle}{dt} &= \langle 2y_i \dot{y}_i \rangle + \frac{1}{2} \langle 2\dot{y}_i^2 \rangle = \\ &= 2c \langle x_i y_i \rangle + 2d \langle y_i^2 \rangle + \frac{2W^{I \leftarrow E}}{M-1} \sum_{k \neq i} \langle y_i x_k \rangle + \frac{2W^{I \leftarrow I}}{M-1} \sum_{k \neq i} \langle y_i y_k \rangle + 2 \langle y_i \xi_i^2 \rangle + \langle \xi_i^2 \xi_i^2 \rangle = \\ &\quad 2c \langle Z_i \rangle + 2d \langle Y_i \rangle + 2W^{I \leftarrow E} \langle Z_{ij} \rangle + 2W^{I \leftarrow I} \langle Y_{ij} \rangle + \sigma_{\text{int}}^2 + \sigma_{\text{ext}}^2, \end{aligned} \quad (17)$$

where now the correlation between noises yields a non-vanishing value. This operation is repeated with all the remaining terms, in order to find a linear system of differential equations with 6 variables and 6 equations,

$$\frac{1}{2} \dot{X}_i = a X_i + b Z_i + [W^{E \leftarrow E} X_{ij} + W^{E \leftarrow I} Z_{ij}] + \frac{1}{2} (\sigma_{\text{int}}^2 + \sigma_{\text{ext}}^2), \quad (18a)$$

$$\frac{1}{2} \dot{Y}_i = c Z_i + d Y_i + [W^{I \leftarrow E} Z_{ij} + W^{I \leftarrow I} Y_{ij}] + \frac{1}{2} (\sigma_{\text{int}}^2 + \sigma_{\text{ext}}^2), \quad (18b)$$

$$\dot{Z}_i = c X_i + (a+d) Z_i + b Y_i + [W^{I \leftarrow E} X_{ij} + (W^{E \leftarrow E} + W^{I \leftarrow I}) Z_{ij} + W^{E \leftarrow I} Y_{ij}] + \sigma_{\text{ext}}^2, \quad (18c)$$

$$\frac{1}{2} \dot{X}_{ij} = \left(a + \frac{M-2}{M-1} W^{E \leftarrow E} \right) X_{ij} + \left(b + \frac{M-2}{M-1} W^{E \leftarrow I} \right) Z_{ij} + \frac{1}{M-1} [W^{E \leftarrow E} X_i + W^{E \leftarrow I} Z_i], \quad (18d)$$

$$\frac{1}{2} \dot{Y}_{ij} = \left(d + \frac{M-2}{M-1} W^{I \leftarrow I} \right) Y_{ij} + \left(c + \frac{M-2}{M-1} W^{I \leftarrow E} \right) Z_{ij} + \frac{1}{M-1} [W^{I \leftarrow I} Y_i + W^{I \leftarrow E} Z_i], \quad (18e)$$

$$\dot{Z}_{ij} = \left(c + \frac{M-2}{M-1} W^{I \leftarrow E} \right) X_{ij} + \left(b + \frac{M-2}{M-1} W^{E \leftarrow I} \right) Y_{ij} + \left(a+d + \frac{M-2}{M-1} (W^{E \leftarrow E} + W^{I \leftarrow I}) \right) Z_{ij} + \quad (18f)$$

$$+ \frac{1}{M-1} [W^{I \leftarrow E} X_i + W^{E \leftarrow I} Y_i + (W^{E \leftarrow E} + W^{I \leftarrow I}) Z_i]. \quad (18g)$$

This system can be solved in the stationary limit when all the derivatives of the left-hand side are zero. From these, one is able to obtain the Pearson correlation coefficients. Correlation with itself is always unity, thus there are only four coefficients remaining: the correlation between excitation and inhibition inside a group $C_{EI}^{\text{int}} = Z_i^*/\sqrt{X_i^*Y_i^*}$, and all three between-group correlations, $C_{EE}^{\text{ext}} = X_{ij}^*/X_i^*$, $C_{II}^{\text{ext}} = Y_{ij}^*/Y_i^*$, and $C_{EI}^{\text{ext}} = Z_{ij}^*/\sqrt{X_i^*Y_i^*}$.

ii. Solutions for the homogeneous network

In some special cases, it is possible to give a simple solution in closed form for the correlation coefficients. One example is the homogeneous network: when all weights are identical, and an intrinsic decay is added to both the excitatory and inhibitory populations (i.e., with $c = W^{E \leftarrow E} = W^{I \leftarrow E} = +W$, $b = W^{I \leftarrow E} = W^{I \leftarrow I} = -W$ and $a = W - 1$, $d = -W - 1$) the solution reads

$$C_{EI}^{\text{int}} = \frac{r^2 (M - 1)^2 + W^2 (1 - r)^2 M}{\sqrt{M^2 (1 - r)^4 W^4 + (M - 1) (1 - r)^2 W^2 [M(2 - 4(1 - r)r) - (1 - r)^2] + (M - 1)^4 [2(r - 1)r + 1]^2}}, \quad (19a)$$

$$C_{EI}^{\text{ext}} = \frac{W^2 (1 - r)^2 M}{\sqrt{M^2 (1 - r)^4 W^4 + (M - 1) (1 - r)^2 W^2 (M(4(r - 1)r + 2) - (r - 1)^2) + (M - 1)^4 (2(r - 1)r + 1)^2}}, \quad (19b)$$

$$C_{EE}^{\text{ext}} = \frac{(1 - r)^2 W ((W + 1) M - 1)}{M(1 - r)^2 W^2 + (M - 1)(1 - r)^2 W + (1 - M)^2 [1 - 2(1 - r)r]}, \quad (19c)$$

$$C_{II}^{\text{ext}} = \frac{(1 - r)^2 W ((W - 1) M + 1)}{M(1 - r)^2 W^2 - (M - 1)(1 - r)^2 W + (1 - M)^2 [1 - 2(1 - r)r]}, \quad (19d)$$

where we defined r as the signal-to-noise ratio, $\sigma_{\text{int}} = r\sigma$, $\sigma_{\text{ext}} = (1 - r)\sigma$. This analytical solution has some interesting features. First, notice it does not depend on the total amount of noise σ that the system receives, but only on the ratio between external and internal noise. Second, if $W \rightarrow \infty$ all correlations go to 1, making the diversity between groups vanish. Expanding in series around $\epsilon = 1/W = 0$, one can see that all coefficients are $C = 1 - \mathcal{O}(1/W^2)$ for large coupling.

It is also possible to study the limiting values of the noise. $r = 1$ makes all the between-group correlations equal to zero, while coupling determines the in-group value. On the other hand, when $r \ll 1$, one gets

$$C_{EI}^{\text{int}} \simeq C_{EI}^{\text{ext}} \simeq \frac{MW^2}{(M - 1)^2} + \mathcal{O}(r^2), \quad (20a)$$

$$C_{EE}^{\text{ext}} \simeq -C_{II}^{\text{ext}} \simeq \frac{W}{M - 1} + \mathcal{O}(r^2). \quad (20b)$$

meaning that the external correlations grow linearly with the coupling, but quadratically with the signal-to-noise ratio: a small increase in coupling needs to be followed by a larger increase in signal intensity in order to recover the previous tuning. As a result, the coupling has a larger impact on tuning than the signal-to-noise ratio, an effect that can be measured in the full spiking network.

Finally, we see that between-group correlations also tend to zero as the limit $M \rightarrow +\infty$ is taken, since in that case, the input that a module receives from all others becomes just white noise. A finite number of clusters (or finite connectivity among them) is thus required for tuning.

I. Clustering optimisation

Optimization of clustering for a fully connected network can be done by minimizing a loss function that depends on the correlations. A simple possibility is to employ minimum squares,

$$\mathcal{L}^{\text{an}}[C; p] = (1 - C_{EI}^{\text{int}})^2 + (C_{EE}^{\text{ext}})^2 + (C_{EI}^{\text{ext}})^2 + (C_{II}^{\text{ext}})^2. \quad (21)$$

The solution and associated optimal correlations are shown in N. There are several key remarks following from this figure:




FIG. N. **Optimal clustering for the analytical results.** The selected values of the clustering lead to large decorrelations between different groups. Large correlations in-group can be achieved for high values of recurrence and external noise. In this case, highly clustered excitation and homogeneous inhibition are recovered.

1. Even for very large clustering and extremely low signal-to-noise ratio, the clustering is able to provide an in-group correlation close to unity combined with low between-group correlations, thus ensuring co-tuning.
2. Inhibition to excitation never clusters. Inhibitory neurons act over excitatory individuals regardless of their cluster.
3. Excitatory connections are clustered. In particular, excitatory connections always project to inhibitory neurons in their own cluster but not to other ones. Excitatory-to-excitatory connectivity is also strongly clustered, except for large coupling.
4. Inhibition controls excitation for large W . If one keeps highly clustered excitation and increases the coupling, the dynamics of single modules become unstable at a critical value $W_c(r)$. However, the network can remain stable if the excitatory clustering is reduced and the amount of inhibition in the group increases, which can be accomplished by increasing r_{II} .
5. When the signal-to-noise ratio is close to one, clustering becomes mostly irrelevant, since the system is driven by the external input, which allows co-tuning easily.

Notice that the optimization algorithm automatically finds solutions where the equations are well-defined –i.e., where the system reaches a stationary state– thus selecting to increase the inhibitory clustering when W goes over the instability threshold (Fig N).

Therefore, the analytical approach is able to find a good candidate for optimal clustering depending on the network dynamics. Although it cannot be directly applied to the spiking network, which is able to display richer dynamics, it tells us that, as a rule of thumb, excitatory clustering should be as high as possible while avoiding crossing an instability threshold. If this happens, inhibition needs to be increased.

J. The inferred connectivity structure encourages competition between assemblies

The distribution of optimal assembly strengths we identify (Fig Oa), consists in very strong excitatory assemblies and much weaker inhibitory assemblies, leading to strong excitatory connections among neurons of the same input group and more spread-out inhibitory connections that also target neurons from other groups.

In our network, this global inhibition setting promotes competition between different assemblies which generates the desired correlation structure in the population activity. Essentially, when a group receives an external input, the strong excitatory connections within the group (projecting to both E and I in-group neurons) will lead to high activity of all the neurons (both E and I) in the group. However, unlike the E connections, which largely target other neurons inside the group, the I connections (targeting both E and I neurons) are mostly directed toward the neurons of other groups, which means that the high inhibitory activity inside the group would lead to the suppression of the activity (in both E and I neurons) of other groups.




FIG. O. **a** The posterior distribution of assembly strengths. It promotes very localized excitatory and relatively broad inhibitory connectivity. **b** A diagrammatic representation of how the inferred connectivity is tuned for decorrelating the activities of different groups

K. Tables of parameters

We mostly used the neuron model parameters from the original inhibitory STDP paper [11].

Network Model		
Symbol	Description	Value
N	Number of neurons	1000
N_E	Number of E neurons	800
N_I	Number of I neurons	200
M	Number of input groups	8
g_{leak}	Leak conductance	10 nS
V_{rest}	Resting potential	-60 mV
V_{reset}	Reset potential	-60 mV
V_{th}	Spiking threshold	-50 mV
V_E	Excitatory reversal potential	0 mV
V_I	Inhibitory reversal potential	-80 mV
C_m	Membrane capacitance	200 pF
τ_{ref}	Absolute refractory period	5 ms
τ_E	Decay time constant of E conductance	5 ms
τ_I	Decay time constant of I conductance	10 ms
$\overline{g_E}$	E weight scaling constant	1.4 nS
$\overline{g_I}$	I weight scaling constant	3.5 nS

TABLE A. The parameters used in all simulations (unless otherwise specified) for the network and neuron models

Plasticity Rules		
Symbol	Description	Value
τ_1^{estdp}	Slow eSTDP timescale	50 ms
τ_2^{estdp}	Fast eSTDP timescale	10 ms
η_E	eSTDP learning rate	0.0025
A_{LTP}	Long term potentiation amplitude	1.0
A_{LTD}	Long term depression amplitude	0.2
τ^{istdp}	iSTDP timescale	10 ms
η_I	iSTDP learning rate	0.01
ρ_0	iSTDP target firing rate	3 Hz
η_N	Normalization learning rate	0.003
W_{target}^E	Excitatory normalization target	5.0
W_{target}^I	Inhibitory normalization target	5.0

TABLE B. The parameters used in all simulations (unless otherwise specified) for the plasticity rules

ABC Optimization		
Symbol	Description	Value
α	weight of in-group correlation	0.1
β	weight of between-group correlations	0.3

TABLE C. The parameters used for the loss of the ABC

-
- [1] F. Lagzi and A. L. Fairhall, Emergence of co-tuning in inhibitory neurons as a network phenomenon mediated by randomness, correlations, and homeostatic plasticity, *Science Advances* **10**, eadi4350 (2024), <https://www.science.org/doi/pdf/10.1126/sciadv.adi4350>.
 - [2] S. Deneve and C. Machens, Efficient codes and balanced networks, *Nature Neuroscience* **19**, 375 (2016).
 - [3] J. Gjorgjieva, C. Clopath, J. Audet, and J.-P. Pfister, A triplet spike-timing-dependent plasticity model generalizes the bienstock-cooper-munro rule to higher-order spatiotemporal correlations, *Proceedings of the National Academy of Sciences* **108**, 19383 (2011), <https://www.pnas.org/content/108/48/19383.full.pdf>.
 - [4] A. Schulz, C. Miehl, I. Berry, Michael J, and J. Gjorgjieva, The generation of cortical novelty responses through inhibitory plasticity, *eLife* **10**, e65309 (2021).
 - [5] N. Caporale and Y. Dan, Spike timing-dependent plasticity: A hebbian learning rule, *Annual review of neuroscience* **31**, 25 (2008).
 - [6] S. Eckmann, E. J. Young, and J. Gjorgjieva, Synapse-type-specific competitive hebbian learning forms functional recurrent networks, *Proceedings of the National Academy of Sciences* **121**, e2305326121 (2024).
 - [7] A. Litwin-Kumar and B. Doiron, Formation and maintenance of neuronal assemblies through synaptic plasticity, *Nature communications* **5**, 5319 (2014).
 - [8] C. Clopath, T. P. Vogels, R. C. Froemke, and H. Sprekeler, Receptive field formation by interacting excitatory and inhibitory synaptic plasticity, *bioRxiv* 10.1101/066589 (2016), <https://www.biorxiv.org/content/early/2016/07/29/066589.full.pdf>.
 - [9] One could argue that external noise should be interpreted as Stratonovich and internal as Itô. Since both noises are additive, this difference is not so relevant, and we treat both noises as Itô for simplicity.
 - [10] C. Gardiner, *Stochastic Methods: A Handbook for the Natural and Social Sciences*, Springer Series in Synergetics (Springer, 2009).
 - [11] T. P. Vogels, H. Sprekeler, F. Zenke, C. Clopath, and W. Gerstner, Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks, *Science* **334**, 1569 (2011), <https://science.sciencemag.org/content/334/6062/1569.full.pdf>.

Appendix E Modularity in E/I networks controls the dynamical regime and optimizes computational capabilities

Contains the full text of the manuscript: Giannakakis, E., Buendía, V., Vinogradov, O., and Levina, A. (In Preparation). Modularity in E/I networks controls the dynamical regime and optimizes computational capabilities.

Modularity in E/I networks controls the dynamical regime and optimizes computational capabilities

Emmanouil Giannakakis,^{1,2,*} Victor Buendía,^{1,2,3,*} Oleg Vinogradov,^{1,2} and Anna Levina^{1,2}

¹*Department of Computer Sciences, University of Tübingen, Germany*

²*Max Planck Institute for Biological Cybernetics, Tübingen, Germany*

³*Department of Computing Sciences, Bocconi University, 20136 Milano, Italy*

Structural modularity is a ubiquitous feature of biological neural networks and has been associated with rich dynamics that can impact computational performance. However, brain networks are characterized by diverse neuron types with distinct connectivity patterns. Still, the impact of cell-type specific modular structures on network dynamics and computational capacity has not been explicitly investigated before. Here, using methods from statistical mechanics, we demonstrate how different levels of excitatory and inhibitory modularity can regulate the dynamical state of a balanced E/I network. We further showcase how the dynamical state generated by the E/I topology can affect the network’s performance using a reservoir computing framework. Our findings propose analytical methods for the study of E/I networks with inhomogeneous connectivity and demonstrate how, by regulating a network’s dynamical state, cell-type specific network topologies can optimize its performance in complex temporal tasks.

I. INTRODUCTION

Biological neural networks are characterized by a pronounced inhomogeneity in their connectivity, which prominently manifests by the formation of modules or clusters of highly interconnected areas that sparsely connect to different areas [1–3]. Modular and clustered connectivity structures have been shown to significantly impact network dynamics, including controlling population bursting [4] and cascading dynamics [5] as well as increasing the robustness of attractor dynamics [6]. Additionally, the dynamics associated with modular networks have been shown to match the dynamics of cortical populations [7, 8] and impact information processing [9] and task performance [10].

Thus, understanding the impact of the modular topology on the network’s dynamics is essential to grasp how the cortex extract and process information. However, biological neural networks are characterized by a large diversity of neuron types, each of whom follows distinct connectivity patterns [11], including different types of inhibitory neurons, which form inhomogeneous connectivity structures [12, 13]. Despite this, most theoretical studies focusing on the dynamics of clustered and modular brain networks, focus on excitatory neurons alone [7, 14, 15] and the few studies combining excitation and inhibition largely focus on the dynamics of balanced networks under spatial constraints [16, 17]. However, despite promising results on the impact of E/I clusters [18, 19] on network dynamics, the potential impact of neuron-type specific modular structures for population dynamics and information processing has not been studied before.

Since different levels of excitatory clustering with uniform inhibition has already demonstrated to create very rich dynamics characterized by long autocorrelations [14],

the impact of different E and I modular structures on population dynamics could have serious implications on the computational capabilities of neuronal networks. For instance, clustered spiking E/I networks have already demonstrated to be able to learn more efficiently in the context of reinforcement learning [20].

Here we study the dynamics and the computational capabilities of a modular, balanced rate network with cell-type specific connectivity. Using *dynamic mean-field* analysis, we develop a full map of the network’s dynamics, explicitly showing how different levels of E/I modularity influence the network’s dynamical state. Then, we link network dynamics with the network’s capacity to perform a complex time-series prediction task, using an *echo state network* (ESN) to assess the computational capabilities of the system. We map the relation between different dynamical states, as defined by the network’s topology and the ESN’s performance. Finally, we study how the structure of the external input interacts with the network’s topology. We examine how the interplay between E/I structured recurrent connectivity and input specificity ultimately impacts task performance.

Our results provide a mapping between the structure of random E/I network, characterized by arbitrary levels of modularity, and the resulting population dynamics. We find that networks with stronger clustered excitation can prevent the network from entering into a chaotic state, boosting its performance. In general, we observe that larger heterogeneity in the neurons’ average rates helps to boost the network’s performance at the task. Our findings indicate that structured networks obeying Dale’s law benefit from distinct E/I connectivity patterns to optimize their ability to perform complex tasks and suggest a possible link between cortical connectivity modularity and the emergence of population dynamics beneficial for intricate computations.

* These authors contributed equally.




Fig 1. Different levels of E and I modularity control network dynamics and optimizes computational capabilities. We investigate the impact of synapse-type specific recurrent connectivity on the dynamics of a balanced network and evaluate its impact on the network’s computational capabilities. Using dynamic mean-field theory, we map different network topologies to distinct dynamical states and, using a reservoir computing framework, we show how the analytically-derived dynamical state controls the network’s performance in a time series prediction task.

II. RESULTS

We use a network of N rate-neurons split into M modules of N/M neurons. Each module contains the same number of excitatory and inhibitory neurons, which form excitatory and inhibitory subpopulations. Thus, the network has $2M$ populations that we index with Greek letters $\mu \in \{1, \dots, 2M\}$, where odd populations are excitatory and even are inhibitory and populations from the same module have consequent indices. To map the populations to their respective modules, we define an index function

$$I : \{1, \dots, 2M\} \rightarrow \{1, \dots, M\}, I(\mu) = M_\mu,$$

where M_μ is an index of the module containing the population μ . Additionally, we define $J : \{1, \dots, 2M\} \rightarrow \{I, E\}$, $J(\mu) = I$ if μ is inhibitory population and $J(\mu) = E$ otherwise. We vary the network’s modularity (how strongly modules connect within each other compared to connected across modules) while maintaining the total number of connections in the network constant. To do so, we define a modularity metric $r \in [0, 1]$ such that for $r = 0$, the network connectivity is uniform (no modules) and for $r = 1$ the modules are fully disconnected. Specifically, the probability of connection from a neuron j in population ν to a neuron i in population μ is given by:

$$p_{ij}^{\mu\nu} = \begin{cases} \frac{p_0}{M} \cdot [1 + (M-1) \cdot r_{J(\mu)J(\nu)}], & I(\mu) = I(\nu) \\ \frac{p_0}{M} \cdot [1 - r_{J(\mu)J(\nu)}], & \text{otherwise} \end{cases} \quad (1)$$

where the probability p_0 controls the overall sparsity of the network. Here, we study the effects of different modularity levels for the excitatory and inhibitory population by varying the $r_E = r_{IE} = r_{EE}$ and $r_I = r_{II} = r_{EI}$, thus allowing the modularity to depend only on the identity of the presynaptic neuron. This way, we control the fractions of excitatory/inhibitory connections each neuron receives from its module to the total number of incoming excitatory/inhibitory connections (Fig. 1, Fig. 2 a).

A. The level of excitatory vs inhibitory modularity controls the network’s dynamics

We start by analysing the dynamical properties of E/I modular networks in the absence of external inputs. It has been shown that such dynamical properties are correlated with the computational capabilities of the networks. Specifically, in the case of echo-state networks, the echo-state property requires that the network’s largest Lyapunov exponent is below unity to avoid the divergence of the system [21–23]. However, being at the *edge of chaos* positively impacts the network’s performance [22, 24, 25] and hence locating such a transition is of foremost importance.

In a random heterogeneous network, the transition from a stable fixed point to chaotic dynamics is characterised by the spectral radius crossing the stability threshold. At this point, many eigenmodes are destabilized at the same time, leading to chaotic dynamics for individual nodes. In contrast, if a single eigenvalue crosses the threshold a single mode is destabilized, lead-

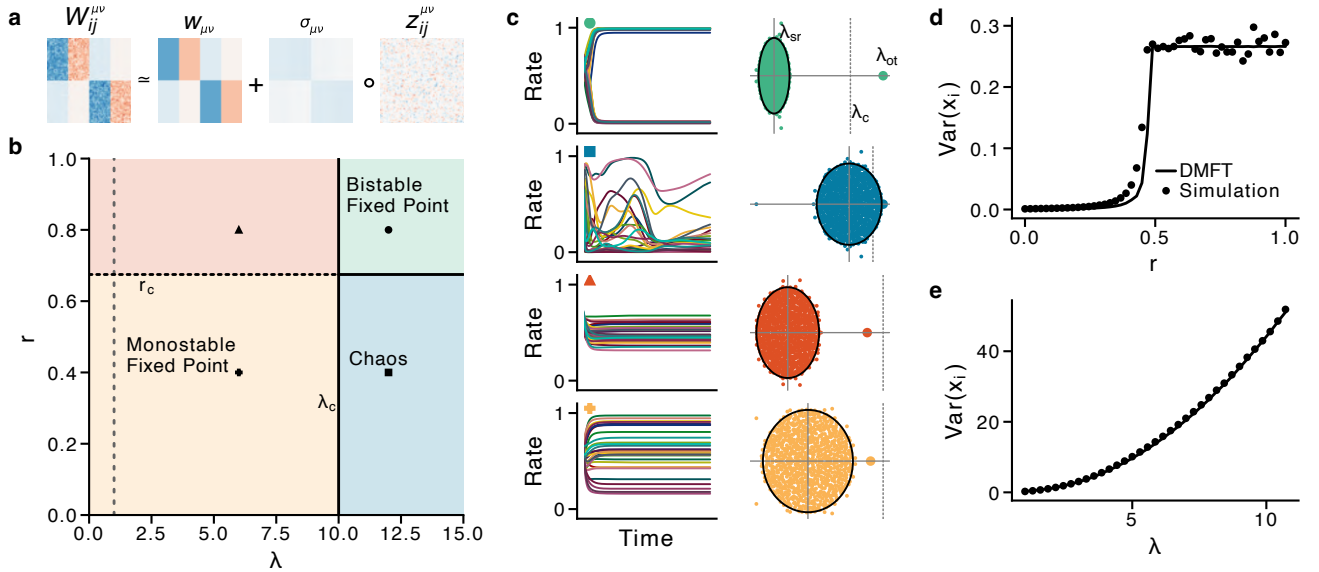


Fig 2. Dynamics of the reservoir network. **a.** Visual decomposition of the $N \times N$ connectivity matrix as mean and variance components. **b.** Phase diagram of the recurrent network in the absence of external input, as predicted by topology and dynamic mean-field theory. Continuous lines represent bifurcations where the stability of fixed points changes. The discontinuous line shows an abrupt change in the variance around the monostable fixed point. The dashed vertical line indicates the section used in panel (d). **c.** Examples trajectories of individual neurons for networks of $N = 3000$ for all 4 different phases of the diagram (left) and the corresponding spectrum of the connectivity matrix. The outlier eigenvalue, as predicted by the reduced matrix $w_{\mu\nu}$, is shown by a larger dot. The dotted grey lines indicate the threshold of instability. **d.** Variance of the distribution of fixed points around their mean as a function of the clustering for a fixed spectral radius. A continuous line indicates the theoretical results from the DMFT, while dots are values obtained averaging 10 networks of size $N = 1000$. **e.** Same as (d), but fixing $r = 0.1$ and changing the spectral radius instead.

ing to a new stable fixed point or limit cycle. Recently, Dahmen *et al.* have highlighted the differences between both transitions and its impact in the correlation structure of neurons [26].

Hence, we first find under which conditions our network's spectral radius can potentially cross an instability; these theoretical results will be the basis to later compare the performance of our reservoirs. The details of the analytical derivation are written step-by-step in the Supplementary I, but we outline the main results and their consequences here.

Firstly, we characterise our recurrent network using a $N \times N$ matrix in which each block is approximated by random Gaussian weights (Fig. 2a). This is a very robust approximation, used successfully in similar setups [27–30]. With it, the network can be written as

$$W_{ij}^{\mu\nu} \simeq \frac{w_{\mu\nu}}{N} + \frac{\sigma_{\mu\nu}}{\sqrt{N}} z_{ij}^{\mu\nu}, \quad (2)$$

The matrix $z_{ij}^{\mu\nu} \sim \mathcal{N}(0, 1)$ is composed by standard Gaussian variables, and $w_{\mu\nu}$, $\sigma_{\mu\nu}$ encode the average and standard deviation of each module.

The spectrum of the full matrix \hat{W} is completely determined by that of the reduced mean and variance matrices, \hat{w} and $\hat{\sigma}$ [29, 31]. In particular, the largest eigenvalue of \hat{w} yields a potential discrete outlier of \hat{W} , while the

square root of the largest eigenvalue of $\hat{\sigma}^2$ returns its spectral radius. We leverage this to study analytically the spectrum of the matrices in a simple, but illustrative case, in which we have only two modules ($M = 2$) with $r_{EE} = r_{IE} = r_E \equiv r$ and $r_{EI} = r_{II} = r_I \equiv 1 - r$. Therefore, when excitatory neurons are very clustered, inhibition is practically random and vice versa. This setting allows us to maintain a constant level of modularity, while varying the type of neurons that are module-specific. In Sup. Info. II, we show that all dynamics regimes present in the system when excitation and inhibition modularity change independently appear also just by moving r . Then, to have an overview of the system's dynamical properties, it suffices to use r as a control parameter for the topology.

In this case, the reduced matrices take a simple form (see Supplementary I) and can be diagonalized exactly. The outlier and spectral radius for \hat{W} are given by

$$\lambda_{\text{ot}} = p(2r - 1)w_0, \quad (3)$$

$$\lambda_{\text{sr}} = \sqrt{\sigma_w^2 + w_0^2(1 - p)}, \quad (4)$$

where w_0 and σ_w^2 are the mean and variance of the weight distribution, respectively. The first surprising feature should be the fact that for small r we can have a large

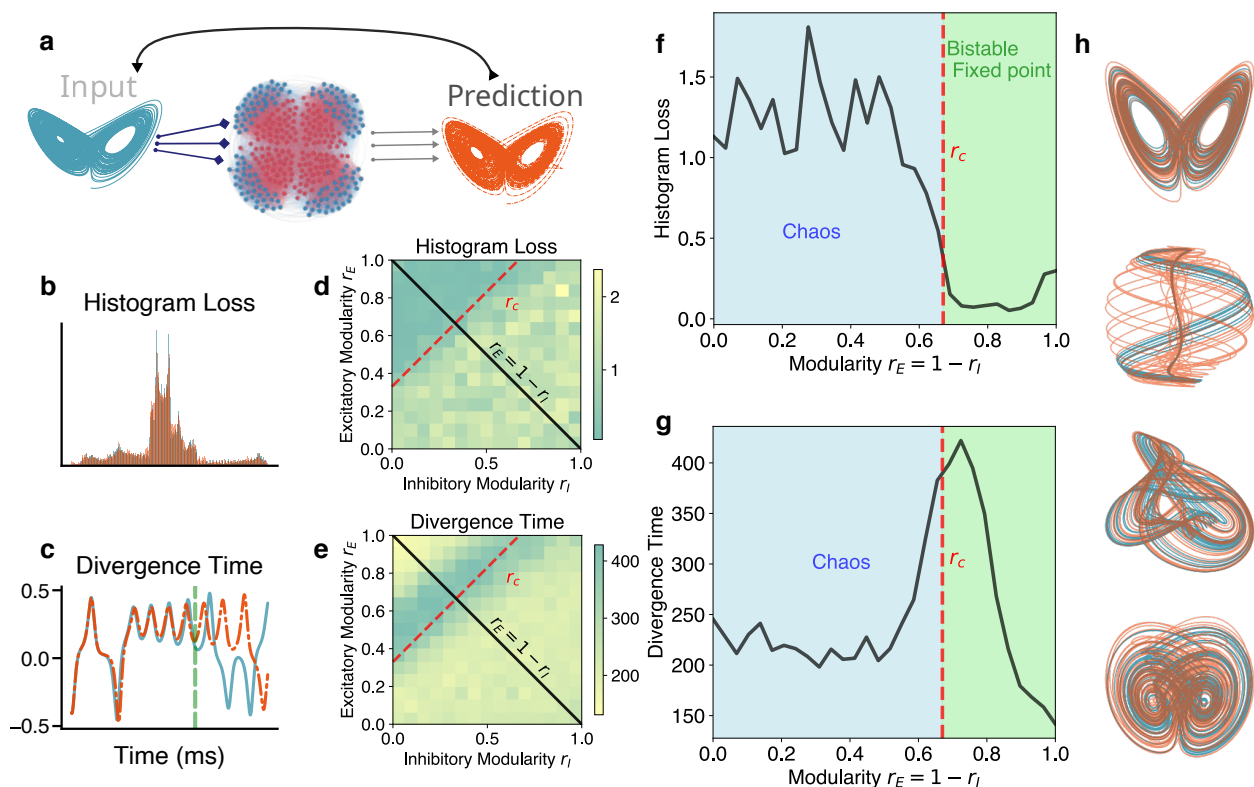


Fig 3. The dynamical state impacts the computational performance of the ESN. **a.** An outline of the Echo State Network (ESN) **b.** The histogram loss evaluates the ESN’s performance by comparing the distribution of points in the predicted and target time series. **c.** The divergence time of the predicted and target time series gives us a second measure of the ESN’s performance. **d.** The histogram and **e.** divergence time losses for different E and I modularity values and a spectral radius of $\lambda = 6.5$ indicate that, as expected from the analytics, the computational capabilities are optimized near the point of the chaotic transition. **f.** The average histogram and **g.** divergence time losses for different levels of E and I modularity (due to the symmetry in E and I modularity, we study the change along the $r_E = 1 - r_I$ line) largely matches the analytically estimated network dynamics. We plot here the performance of the ESN in reconstructing the Lorenz attractor. **h.** Reconstructions of four different chaotic attractors by the reservoir. The blue lines show the input data and the orange lines the ESN’s prediction.

discrete spectral outlier. Even in a balanced network, the topological structure is able to generate a spectral outlier, and make our system function similarly to a excitation dominated network. As discussed above, the network cannot experience a transition to a chaotic regime when such an isolated eigenvalue is dominating. For the chaotic transition to happen, the spectral disk must cross the instability first. This is possible only when the outlier has been merged inside the spectral disk. The first point at when this can happen is when the outlier touches the spectral radius, i.e., at $\lambda_{\text{ot}}(r_c) = \lambda_{\text{sr}}(r_c)$. Solving for r_c , we get

$$r_c = \frac{1}{2} + \frac{1}{2} \sqrt{\frac{\sigma_w^2 + w_0^2(1-p)}{pw_0}}. \quad (5)$$

Having, $r < r_c$ is a necessary condition to have chaotic dynamics. However, it is not sufficient: once the outlier has been merged inside the spectral disk, the disk needs to be large enough in order to overcome the insta-

bility threshold. We would like to remark that r_c holds independently of the size of such spectral radius (see Supplementary A for more details). This means that as soon as the spectral radius of \hat{W} is large enough to produce chaos, crossing $r = r_c$ leads to a chaotic regime (Fig. 2b).

Thus, the immediate next question is: where is the instability threshold? what is the minimal spectral radius needed to generate chaos? We used a dynamical mean-field theory (DMFT) approach [27, 29] to compute the threshold of instability. The DMFT approach allows one to substitute all variables by a set of $2M$ representative nodes, one for each module. Complete derivation details are given in Supplementary A. We obtain that the stability of the fixed point and the transition happen when the largest eigenvalue λ_S of

$$S_{\mu\nu} = a_\mu \sigma_{\mu\nu}^2 \langle [\phi'(y_\nu^*)]^2 \rangle \quad (6)$$

becomes larger than unity. a_μ is the fraction of neurons in a module μ (in our case, all $a_\mu = N_\mu/(2M)$), while y^*

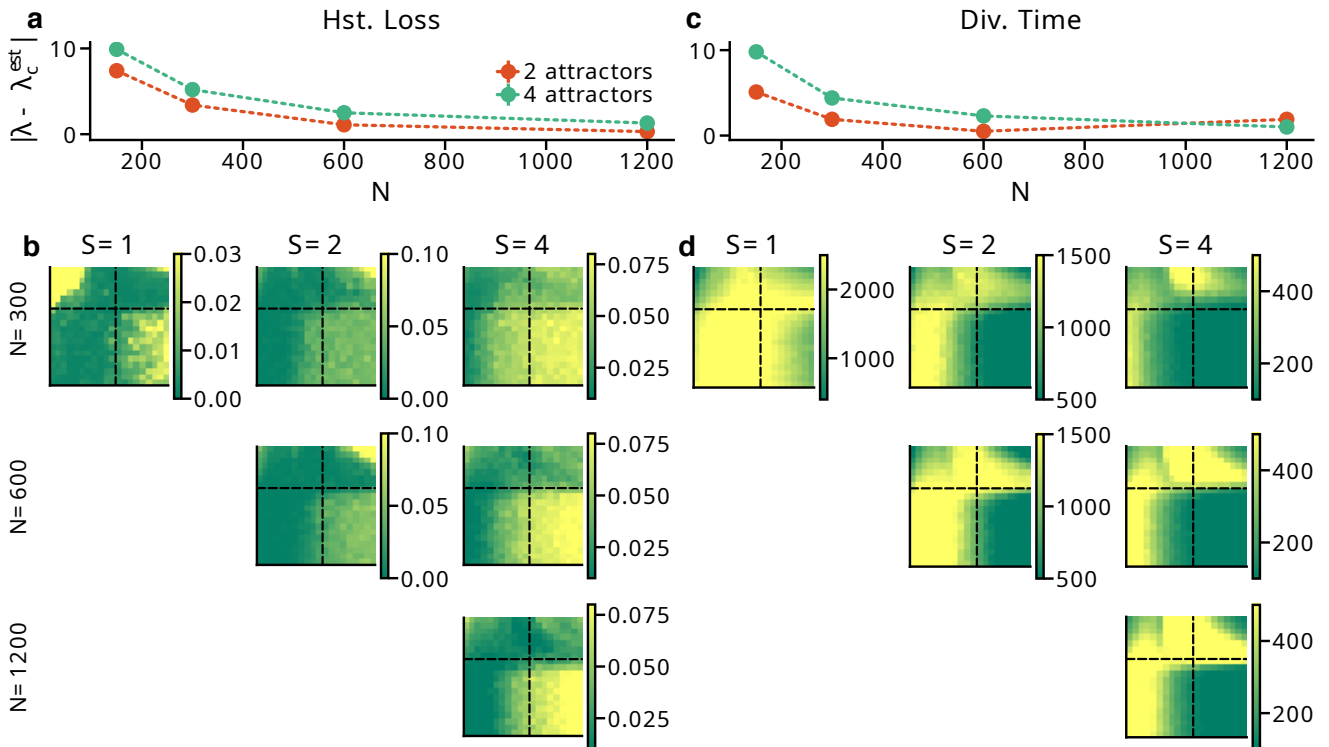


Fig 4. The effect of network size and number of predicted attractors. **a.** Distance from the estimated transition to the critical spectral radius λ_c for increasing system sizes. The network is tasked with the reconstruction of 2 (red) or 4 (green) chaotic attractors at the same time. The distance to criticality shifts with the number of attractors, but decreases with the system size. To estimate the transition point, we find the point where the loss increases abruptly for each value of r below r_c , and then average. **b.** Loss landscape for different system sizes and number of oscillators to reconstruct. **(c),(d)** Same as (a), (b) but using the divergence time instead of the histogram loss.

indicates the value of the fixed point and $\langle[\phi'(y_\nu^*)]^2\rangle$ is an average over the distribution of fixed points. Observe that in the setting of DMFT it is assumed that all nodes in the same module have equivalent behaviour.

In the absence of modularity or structure, the largest eigenvalue λ_S of the matrix \hat{S} is just the spectral radius of the network multiplied by $\langle[\phi'(y_\nu^*)]^2\rangle$; hence, we refer to λ_S as the *generalized spectral radius*, since it generalizes its role in the modular case. In order to determine the location to the transition to chaos it suffices to see when the generalized spectral radius becomes larger than unity (Fig. 2c).

B. Computational capabilities depend on the dynamical state

Having an understanding of how network modularity affects dynamics, we examine how these impact the networks' ability to perform computational tasks. To this end we use the recurrent networks as reservoirs in an Echo State network (ESN) [32] whose readout layer is trained to predict the trajectory of chaotic timeseries (Fig. 3a).

We were evaluating a network's performance based on two metrics: the distance between the distribution of points in the target and predicted time series (histogram loss Fig. 3b) and the first time point where the two time series diverge (divergence time Fig. 3c). Due to the chaotic nature of the target time series, divergence between the data and the learned trajectory will happen. After the divergence, however, the reservoir trajectory might continue to be close to the set of points in the chaotic attractor or not. Thus, the divergence time quantifies how accurately the reservoir can replicate the specific trajectory presented, while the histogram loss measures how well the system learns the attractor set in the phase space. We report both metrics for all our experiments.

First, we test the performance with the Lorenz attractor, a three-dimensional chaotic attractor commonly used for time series prediction tasks. We begin by testing whether the performance is affected for different levels of modularity in the E and I recurrent connections for a fixed spectral radius $\lambda = 10$. We see that as expected from the analysis of the network's dynamics, the distribution of losses is largely symmetric with respect to the $r_E = 1 - r_I$ line (Fig. 3d, e). This suggests that the reser-

voir's performance is primarily determined by the ratio of E to I modularity rather than the exact value of either variable.

When we look at how the losses change as we change the E and I modularity in opposite directions (i.e. we set $r = r_E = 1 - r_I$), we verify that, for a large spectral radius, the loss map closely matches the dynamical states predicted by the DMFT analysis (Fig. 3 f, g). The chaotic state is clearly detrimental to the network's performance, and the area around the chaotic transition leads to superior performance, most clearly for the divergence time (Fig. 3g). Effectively, the modularity can bring networks of a sufficiently large spectral radius to a state that ensures good computational performance. Specifically, the optimal modularity (clustered excitation and relatively uniform inhibition) brings the network near, but not beyond, the chaotic transition point, a result consistent with what is known about reservoir performance [32, 33].

Finally, in order to ensure that our results generalize to other chaotic time series and are not specific to the Lorenz attractor, we repeat our experiment with three other chaotic attractors (Aizawa, Chen and Halvorsen). Despite changes in the average loss for different attractors and small changes in the loss distributions, the overall patterns remain the same, suggesting that our estimation of the network's dynamical state can be effectively used to predict the area in which the ESN's performance is maximized.

C. The simultaneous prediction of multiple chaotic attractors shifts the optimal connectivity structures in small networks

We test whether the reservoir can simultaneously reconstruct multiple independent chaotic time series. Specifically, we present the network with combinations of the four chaotic attractors (Lorenz, Aizawa, Halvorsen and Chen) we tested individually, creating higher dimensional inputs (6, 9 or 12 dimensional for 2, 3 or 4 attractors respectively). While the network generally succeeds in reconstructing time series consisting in combinations of multiple attractors, as would be expected due to the higher complexity of the task, the average performance decreases as the number of simultaneously predicted attractors increases (Fig. 4b, d).

We test the network performance for different combinations of E/I modularity (varying $r = r_E = 1 - r_I$) and spectral radius λ . We note that as we increase the number of presented attractors, the loss map shifts away from the predicted bifurcation diagram, with the chaotic transition appearing to happen for lower spectral radius λ . This shift seems to be explainable as a finite-size effect. To see it, we performed experiments with reservoirs of increasing sizes. Smaller networks do not only perform worse, but also display a shifted transition to criticality, very similar to the shift we observe for networks of fixed size presented with multiple attractors. We verify that

optimality at the edge of chaos is recovered as the network size is gradually increased, as shown in Fig. 4a, c. In particular, we find that as we increase network size, the loss function reservoirs predicting multiple attractors slowly moves towards the recurrent network's phase diagram. The finite-size shift of a critical point is a well-known effect in the theory of phase transitions [34] but here we also observe it for varying input dimensionality.

We would like to remark that the finite-size shifting is inherent to the input properties, and not only because of a shifting of the critical dynamics of the reservoir network alone. One can see this because for a fixed size, transition shifts when the number of neurons considered is increased. To fully describe this phenomenon, one should characterize the dynamics of the recurrent network while driven by chaotic input. This is a theoretical open question which falls outside the scope of our work, and which should be tackled separately.

Finally, it is interesting to note that the difficulty of the task for small networks highlights the best dynamical regimes for the reservoir. The bistable region confers the reservoir the capability to predict the time series for the longest time, while having a good score on the 4 oscillators reconstruction even for $N = 300$. Observe that this region, above r_c , is only available for structured networks. Hence, modularity is able to boost the reservoir's capabilities in comparison to the random counterpart. Moreover, the effect seems to be extended to very large spectral radius when the network size grows, allowing for reservoirs able to optimally solve the task with $\lambda \gg \lambda_c$.

D. Performance of the network is mostly independent to the modularity of input connections

So far, each neuron received randomly selected input from one of the multiple input attractors. We hypothesize that random input may be suboptimal and ask if structured input to specific E and I populations has an impact on the network's performance.

To test this, we simulate a network divided into 4 modules, which is tasked with the simultaneous reconstruction of all 4 chaotic attractors. We assign an attractor to each module and define an input modularity metric r^{In} which controls how specific is the input from a given attractor. This is analogous to our recurrent modularity metrics. Thus, for $r^{In} = 1$, each module receives input only from a single attractor, while for $r^{In} = 0$ the connections are assigned randomly, and a single neuron can receive input from any attractor with equal probability. Since the input can be positive or negative, we study the impact of different excitatory and inhibitory input modularity r_E^{In} and r_I^{In} , similar to the recurrent connections.

To systematically map out the potential optimal solution, we use Bayesian inference. The additional input modularity parameters of the network allows the system to have many potential solutions that optimize the reconstruction task with multiple attractors. Since we lack

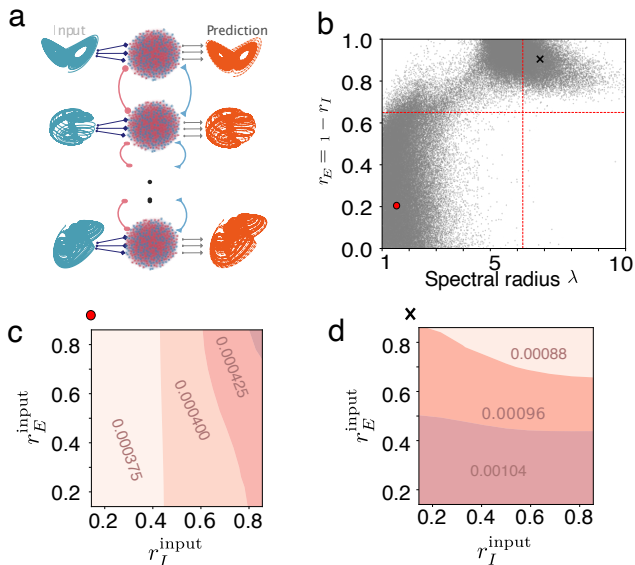


Fig 5. Module-specific input connections have a minor effect on the ESN performance. **a.** An outline of the network with multiple attractors. We study how to optimally distribute feedforward E and I connections from different attractors to different network modules. **b.** The posterior distribution of the network parameters obtained via simulation based inference. **c, d.** The reservoir losses for different levels of E and I modularity in the two indicated points of the posterior

analytics for input-driven networks, we used simulation-based inference (SBI) [35]. Given a set of summary statistics, SBI allows us to obtain a posterior of the parameters which are likely to reproduce the data. In our case, we look for the distributions of the spectral radius, recurrent modularity, and input modularity given the average histogram and divergence loss for all four attractors. We run $3 \cdot 10^5$ simulations with $S = 4$ inputs and recurrent network size $N = 300$ and fit to approximate the posterior distribution of the parameters given the losses.

The distribution of the optimal parameters for the recurrent parameters, (Fig. 5b), shows that there are two optimal solutions. These are perfectly consistent with (Fig. 4) (e.g., compare with the divergence time for $N = 300$ and $S = 4$, our simulation conditions). One solution is found in the area with high spectral radius and high excitatory modularity, inside the bistable region, and another one with low spectral radius and modularity, in the analytically inferred monostable stable. We will refer to these solutions as the bistable solution and monostable solution, respectively. Observe that for the simulation conditions, these solutions do not cover the entire regime of the phase space, but only a part of it. On the other hand, the marginal distribution of the input modularity was close to uniform, indicating that the system is mostly independent of changes in the input modularity (Fig. 5cd).

However, in certain cases the marginal distribution of the parameters can hinder possible dependencies that appear when the other parameters are fixed - i.e. in the conditional distributions [36]. Therefore, we further check the impact of input modularity in these two optimal recurrent structures (Fig. 5b). Then, we simulate the network in both the bistable and monostable solutions to compute their average loss (Fig. 5)c, d).

We found that although the effect of the input modularity was small, the system could still benefit from the structured input in two different ways. For the monostable solution, modularity of both excitatory and inhibitory input connections is detrimental for the dynamics and for the inhibitory inputs modularity has a negative effect. In contrast, in the bistable solution case, the system can benefit from module-specific inhibitory and excitatory connections. The modularity of excitatory connections generally has a positive effect.

These differences explain a deviation from the uniform distribution that we observed in the marginal distributions. However, we should emphasize that the effect is rather small compared to the effect of recurrent modularity.

III. DISCUSSION

In this paper, we have analysed a structured balanced network with different levels E/I modularity, linking its dynamical regimes to its computational capabilities in a reservoir computing setting. We have explicitly shown how the network topology can directly influence its dynamics by shaping its spectra. Moreover, using the network as an echo state network, we can reconstruct different chaotic time series, both independently or simultaneously. We match the different dynamical states with the network's performance, which is optimized at the edge of chaos.

Our findings suggest that a structure of relatively high excitatory modularity combined with relatively low inhibitory modularity represents a particularly important transition point. Networks with weak inhibitory modularity can maintain very high computational performance even at a very large spectral radius. This particular topology has been shown before to generate beneficial dynamics for unsupervised learning [37] and matches connectivity patterns observed in cortical networks [13, 38]. For instance, orientation selectivity in V1 has traditionally regarded excitatory neurons as tuned, and inhibitory neurons as weakly tuned. Since similarly tuned neurons are more likely to connect to each other, this translates into a stronger modularity for exciting with broad inhibition. Thus we believe that our results are consistent with current biological evidence, and could shed light on the computational aspects of E/I networks.

In order to obtain the full phase diagram of system analytically, we have used a dynamical mean-field approach. Our theoretical results are consistent with previous lit-

erature in the topic [27–29]; however, Aljadeff *et al.* [28] do not consider any external input to the network, nor a mean for any of the modules; Mastrogiuseppe and Ostojic [29] study in detail only unstructured E/I dynamics, and Kadmon and Sompolinsky [27] do not explicitly study the consequences of topology. Furthermore, since in our case the sigmoid does not vanish at zero, our system can be seen as driven by external, quenched random noise. Therefore, previous results cannot be directly applied and needed to be extended. We have extended our results to deal with this situation, while developing an approximated method to obtain the location of the transitions in an accurate way. Given the generality of our setup, we believe that the dynamic mean-field presented here will constitute an advance to study related systems in both computational neuroscience and reservoir computing applications.

Finally, we have studied the simultaneous reconstruction of several attractors, as well as the consequences of the input structure. As our network is composed of several modules, the reconstruction might have improved when the input is distributed in a particular way among the modules. The relationship between the input and the network structure is still an open question for the efficiency of reservoir computers, and it is particularly interesting in the context of neuroscience, where input to the cortex is highly structured [39–41].

Our findings suggest that in balanced, E/I recurrent networks, cell-type specific modular connectivity structures may offer precise control of network dynamics, enabling the network to tune its distance from the chaotic state, which optimized computational capabilities in challenging temporal tasks. These results offer insights into a potential link between the structure and function of biological neural networks, and motivate further work in the potential computational benefits of structured connectivity in artificial networks.

IV. METHODS

A. Neurons

At each timestep, the dynamics of the network’s activity $x = (x_1, \dots, x_N)$ are given by the following equations

$$\dot{x}_i = -x_i + \phi \left(\sum_{j=1}^N W_{ij} x_j + s_i(t) \right), \quad (7)$$

where \hat{W} is the recurrent connectivity matrix and $s_i(t)$ is the external input to each node. The nonlinearity ϕ is the sigmoid

$$\phi(x) = \frac{1}{1 + e^{-x}}. \quad (8)$$

Notice that in our setup, $\phi(0) = 1/2 \equiv \phi_0$.

B. Input

Each neuron receives as input one dimension of a D -dimensional time series $\vec{C}(t)$:

$$\vec{s}(t) = \hat{W}_{\text{Input}} \vec{C}(t), \quad (9)$$

where \hat{W}_{Input} is a $N \times D$ matrix, where a single non-zero entry per row with weight $w_{\text{Input}} = 0.5$. In this way, each node receives a single dimension from the time series.

C. Echo state network (ESN)

The simulation of the ESN is divided in two phases:

Stimulus-driven: The time series $\vec{C}(t)$ is given as input for T time steps, leading to the recurrent network activity $\hat{X} = [\vec{x}_1, \dots, \vec{x}_T]$ obtained by stacking network’s state at each time. Using Ridge regression, a linear layer \hat{W}_R is trained to predict the trajectory of the time series. Thus, using the network activity at a given time step, \vec{x}_t the linear layer predicts the next step of the time series $\vec{C}(t+1) \equiv \vec{C}_{t+1}$. The target is constructed by stacking the time series vectors in columns, $\hat{y}^{\text{target}} = [\vec{C}_2, \dots, \vec{C}_{T+1}]$, and regression yields

$$\hat{W}_R = \hat{y}^{\text{target}} \hat{X}^\dagger (\hat{X} \hat{X}^\dagger + \beta \hat{I})^{-1}, \quad (10)$$

where β is the Ridge regression’s normalization constant and \dagger stands for transpose.

Closed-Loop: Once the training is complete, we replace the input from the given time series \vec{C}_{T+1} , with the prediction generated via the trained linear layer $\hat{W}_R \hat{x}_T^\dagger$. If Ridge regression was successful, then $\vec{x}_{T+1}^\dagger = \hat{W}_R \hat{x}_T^\dagger$ and the recurrent network will produce the same activity as if its input had been the actual time series.

We continue in the same way for T' time steps, using the linear layer’s prediction as the next input for the recurrent network, creating a fully closed loop which results in the reconstructed time series $\hat{X}' = [\vec{x}'_{T+1}, \dots, \vec{x}'_{T+T'}]$

D. Evaluation

To evaluate the performance of the ESN, we compare the reconstructed time series \hat{X}' with the actual corresponding time series $\hat{Y} = \{C_{T+1}, \dots, C_{T+T'}\}$. In an ideally functioning ESN, the dynamics of the reservoir should be rich enough to capture the trajectory of the time series and enable a very precise reconstruction, so that $\hat{X}' \approx \hat{Y}$ over the entire testing period. We use two metrics to quantify the ESN performance:

Divergence Time. We estimate the time τ at which the reconstructed trajectory and the original one start to diverge, by obtaining the τ such that $|C_\tau - x'_\tau| > \epsilon$ for a given small threshold ϵ .

Histogram Loss. We compute the histogram of each dimension of the time series separately, both for the original series and the reconstruction. Then, we compute the mean square difference between both histograms for each dimension, and sum among them all.

Each one of the metrics measures a different aspect of the reservoir prediction. On one hand, divergence time captures the quality of the prediction for individual time series. A long divergence time means that the reservoir is able to imitate the behaviour of the chaotic system for long time. However, it does not give information about the statistical properties of the reconstruction after the divergence happened. On the other hand, histogram loss can be very small as long as the reservoir trajectory is predicting the stable basin of the chaotic attractor. Notice that one can have a short divergence time with a small histogram loss, meaning that a reservoir might correctly predict the chaotic basin of attraction without being performant at individual series, but usually both metrics are correlated.

E. Simulation-Based Inference

We use simulation-based inference [42] to find the optimal excitatory and inhibitory input modularity (r_E^{in} and

r_I^{in}) together with the spectral radius (λ) and recurrent modularity (r). We randomly sampled 500,000 parameters from a uniform prior distribution (r , r_E^{in} and r_I^{in} between 0 and 1; λ between 1 and 10). We then fed four attractors to the reservoir simultaneously and computed their divergence time and histogram loss. We then obtain an average loss for all of them as $L^{\text{total}} = \frac{1}{4} \sum_n^{N=4} \frac{L_n^{\text{hist}}}{D_n + 1}$. Next, we fit a normalizing spine flow as a neural posterior estimator to predict the distribution of four parameters given the loss function ($p(r_E^{in}, r_I^{in}, r, \lambda | L^{\text{total}})$). Next, we sampled from this posterior given $L = 0$. For Fig.5 c and d, we additionally refitted the parameters r_E^{in}, r_I^{in} given fixed spectral radius and modularity (400,000 samples).

V. ACKNOWLEDGMENTS

This work was supported by a Sofja Kovalevskaja Award from the Alexander von Humboldt Foundation. EG and SK thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support, and OV thanks the International Max Planck Research School for the Mechanisms of Mental Function and Dysfunction and the Joachim Herz foundation for support. We acknowledge the support from the BMBF through the Tübingen AI Center (FKZ: 01IS18039B). AL is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 39072764.

-
- [1] Claus C Hilgetag and Marcus Kaiser. Clustered organization of cortical connectivity. *Neuroinformatics*, 2:353–360, 2004.
 - [2] Miguel Valencia, MA Pastor, MA Fernández-Seara, J Artieda, J Martinerie, and M Chavez. Complex modular structure of large-scale brain networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 19(2), 2009.
 - [3] Olaf Sporns and Richard F Betzel. Modular brain networks. *Annual review of psychology*, 67(1):613–640, 2016.
 - [4] Wilfredo Blanco, Joel Tabak, and Richard Bertram. Population bursts in a modular neural network as a mechanism for synchronized activity in kndy neurons. *PLOS Computational Biology*, 20(7):e1011820, 2024.
 - [5] Aram Galstyan and Paul Cohen. Cascading dynamics in modular networks. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 75(3):036109, 2007.
 - [6] Neeraj Pradhan, Subinay Dasgupta, and Sitabhra Sinha. Modular organization enhances the robustness of attractor network dynamics. *Europhysics Letters*, 94(3):38004, 2011.
 - [7] Hideaki Yamamoto, Satoshi Moriya, Katsuya Ide, Takeshi Hayakawa, Hisanao Akima, Shigeo Sato, Shigeru Kubota, Takashi Tani, Michio Niwano, Sara Teller, et al. Impact of modular organization on dynamical richness in cortical networks. *Science advances*, 4(11):eaau4914, 2018.
 - [8] Mikael Lundqvist, Martin Rehn, Mikael Djurfeldt, and Anders Lansner. Attractor dynamics in a modular network model of neocortex. *Network: Computation in Neural Systems*, 17(3):253–276, 2006.
 - [9] James C Houk. Information processing in modular circuits linking basal ganglia and cerebral cortex. 1994.
 - [10] Louis Kirsch, Julius Kunze, and David Barber. Modular networks: Learning to decompose neural computation. *Advances in neural information processing systems*, 31, 2018.
 - [11] Casey M Schneider-Mizell, Agnes L Bodor, Derrick Brittain, JoAnn Buchanan, Daniel J Bumbarger, Leila Elabady, Clare Gamlin, Daniel Kapner, Sam Kinn, Gayathri Mahalingam, et al. Cell-type-specific inhibitory circuitry from a connectomic census of mouse visual cortex. *bioRxiv*, 2023.
 - [12] Robert B Levy and Alex D Reyes. Spatial profile of excitatory and inhibitory synaptic connectivity in mouse primary auditory cortex. *Journal of Neuroscience*, 32(16):5609–5619, 2012.
 - [13] Petr Znamenskiy, Mean-Hwan Kim, Dylan R Muir, M Florencia Iacaruso, Sonja B Hofer, and Thomas D Mrsic-Flogel. Functional specificity of recurrent inhibition in visual cortex. *Neuron*, 112(6):991–1000, 2024.
 - [14] Ashok Litwin-Kumar and Brent Doiron. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience*, 15:1498–

- 1505, 09 2012.
- [15] Hesam Setareh, Moritz Deger, Carl C. H. Petersen, and Wulfram Gerstner. Cortical dynamics in presence of assemblies of densely connected weight-hub neurons. *Frontiers in Computational Neuroscience*, 11, 2017.
- [16] Robert Rosenbaum and Brent Doiron. Balanced networks of spiking neurons with spatially dependent recurrent connections. *Phys. Rev. X*, 4:021039, May 2014.
- [17] Robert Rosenbaum, Matthew Smith, Adam Kohn, Jonathan Rubin, and Brent Doiron. The spatial structure of correlated neuronal variability. *Nature neuroscience*, 20, 10 2016.
- [18] Thomas Rost, Moritz Deger, and Martin Nawrot. Winnerless competition in clustered balanced networks: inhibitory assemblies do the trick. *Biological Cybernetics*, 112, 04 2018.
- [19] Vahid Rostami, Thomas Rost, Alexa Riehle, Sacha J. van Albada, and Martin P. Nawrot. Excitatory and inhibitory motor cortical clusters account for balance, variability, and task performance. *bioRxiv*, 2022.
- [20] Philipp Weidel, Renato Duarte, and Abigail Morrison. Unsupervised learning and clustered connectivity enhance reinforcement learning in spiking neural networks. *Frontiers in Computational Neuroscience*, 15, 2021.
- [21] Izzet B. Yildiz, Herbert Jaeger, and Stefan J. Kiebel. Revisiting the echo state property. *Neural Networks*, 35:1–9, November 2012.
- [22] Mantas Lukoševičius. A Practical Guide to Applying Echo State Networks. In Grégoire Montavon, Geneviève B. Orr, and Klaus-Robert Müller, editors, *Neural Networks: Tricks of the Trade: Second Edition*, Lecture Notes in Computer Science, pages 659–686. Springer, Berlin, Heidelberg, 2012.
- [23] Ken Caluwaerts, Francis Wyffels, Sander Dieleman, and Benjamin Schrauwen. The spectral radius remains a valid indicator of the Echo state property for large reservoirs. In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, pages 1–6, August 2013.
- [24] Nils Bertschinger and Thomas Natschläger. Real-time computation at the edge of chaos in recurrent neural networks. *Neural computation*, 16(7):1413–1436, 2004.
- [25] Junjie Jiang and Ying-Cheng Lai. Model-free prediction of spatiotemporal dynamical systems with recurrent neural networks: Role of network spectral radius. *Physical Review Research*, 1(3):033056, October 2019.
- [26] David Dahmen, Sonja Grün, Markus Diesmann, and Moritz Helias. Second type of criticality in the brain uncovers rich multiple-neuron dynamics. *Proceedings of the National Academy of Sciences*, 116(26):13051–13060, June 2019.
- [27] Jonathan Kadmon and Haim Sompolinsky. Transition to Chaos in Random Neuronal Networks. *Physical Review X*, 5(4):041030, November 2015.
- [28] Johnatan Aljadeff, Merav Stern, and Tatyana Sharpee. Transition to Chaos in Random Networks with Cell-Type-Specific Connectivity. *Physical Review Letters*, 114(8):088101, February 2015.
- [29] Francesca Mastrogiuseppe and Srdjan Ostojic. Intrinsic-generated fluctuating activity in excitatory-inhibitory networks. *PLOS Computational Biology*, 13(4):e1005498, April 2017.
- [30] Francesca Mastrogiuseppe and Srdjan Ostojic. Linking Connectivity, Dynamics, and Computations in Low-Rank Recurrent Neural Networks. *Neuron*, 99(3):609–623.e29, August 2018.
- [31] Isabelle D. Harris, Hamish Meffin, Anthony N. Burkitt, and Andre D. H. Peterson. Effect of sparsity on network stability in random neural networks obeying Dale’s law. *Physical Review Research*, 5(4):043132, November 2023.
- [32] Herbert Jaeger. The “echo state” approach to analysing and training recurrent neural networks-with an erratum note. *Bonn, Germany: German National Research Center for Information Technology GMD Technical Report*, 148(34):13, 2001.
- [33] Herbert Jaeger and Harald Haas. Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science (New York, N.Y.)*, 304:78–80, 05 2004.
- [34] J. J. Binney, N. J. Dowrick, A. J. Fisher, and M. E. J. Newman. *The Theory of Critical Phenomena: an introduction to the renormalization group*. Oxford Univ. Press, 2001.
- [35] Alvaro Tejero-Cantero, Jan Boelts, Michael Deistler, Jan-Matthis Lueckmann, Conor Durkan, Pedro J. Gonçalves, David S. Greenberg, and Jakob H. Macke. sbi: A toolkit for simulation-based inference. *Journal of Open Source Software*, 5(52):2505, August 2020.
- [36] Michael Deistler, Jakob H. Macke, and Pedro J. Gonçalves. Energy-efficient network activity from disparate circuit parameters. *Proceedings of the National Academy of Sciences*, 119(44):e2207632119, November 2022. Publisher: Proceedings of the National Academy of Sciences.
- [37] Emmanouil Giannakakis, Oleg Vinogradov, Victor Buendía, and Anna Levina. The topology of e/i recurrent networks regulates the effects of synaptic plasticity. *bioRxiv*, 2023.
- [38] Benjamin Scholl, Daniel E Wilson, Juliane Jaepel, and David Fitzpatrick. Functional logic of layer 2/3 inhibitory connectivity in the ferret visual cortex. *Neuron*, 104(3):451–457, 2019.
- [39] X Jiang, S Shen, CR Cadwell, P Berens, F Sinz, AS Ecker, S Patel, and AS Tolias. Principles of connectivity among morphologically defined cell types in adult neocortex. *science* 350, aac9462, 2015.
- [40] Dirk Feldmeyer, Joachim Lübke, R Angus Silver, and Bert Sakmann. Synaptic connections between layer 4 spiny neurone-layer 2/3 pyramidal cell pairs in juvenile rat barrel cortex: physiology and anatomy of interlaminar signalling within a cortical column. *The Journal of physiology*, 538(3):803–822, 2002.
- [41] Ho Ko, Sonja B Hofer, Bruno Pichler, Katherine A Buchanan, P Jesper Sjöström, and Thomas D Mrsic-Flogel. Functional specificity of local synaptic connections in neocortical networks. *Nature*, 473(7345):87–91, 2011.
- [42] Pedro J Gonçalves, Jan-Matthis Lueckmann, Michael Deistler, Marcel Nonnenmacher, Kaan Öcal, Giacomo Bassetto, Chaitanya Chintaluri, William F Podlaski, Sara A Haddad, Tim P Vogels, David S Greenberg, and Jakob H Macke. Training deep neural density estimators to identify mechanistic models of neural dynamics. 9:e56261. Publisher: eLife Sciences Publications, Ltd.

Supplementary information

Emmanouil Giannakakis*, Victor Buendía*, Oleg
Vinogradov, Sina Khajehabdollahi, Anna Levina

(Dated: December 6, 2024)

Contents

I. Dynamic Mean-Field Theory	2
A. Single population	2
B. Modular Networks	7
C. Numerical implementation	13
II. Dynamical state in r_e, r_i phase space	15
III. Attractors	16
References	17

I. Dynamic Mean-Field Theory

A. Single population

Set up

We start by a simplified model to showcase our methods, and then we will further expand them to include Dale's law and modular topology. Let us assume we have N neurons, each one characterised by a continuous firing rate $x_i > 0$ for $i = 1, 2, \dots, N$. Then, the equations for the dynamical state of the reservoirs read

$$\dot{x}_i = -x_i + \phi \left(\sum_{j=1}^N W_{ij} x_j \right), \quad (1)$$

where $\phi(x) = 1/(1 + \exp(-x))$ is a sigmoidal transfer function and \hat{W} the connectivity matrix. For simplicity we will start assuming that $\hat{W} \sim \mathcal{N}(0, \lambda^2/N)$. The variance is set in order to fix the spectral radius of the matrix, given by $\sqrt{N \text{Var}(W_{ij})} = \lambda$.

Observe that we have selected our transfer function in such a way that $\phi(0) = 1/2$ and $\phi(x) > 0 \quad \forall x$, which ensures Dale's law by keeping all $x > 0$ through the dynamics. The most common choice for nonlinearities in neural networks are those such that $\phi(x \leq 0) = 0$. However, in our case, a capped nonlinearity would cut the input to the reservoir, making

impossible for the echo-state network to learn. Our choice of the nonlinearity allows the reservoir to fulfill the task while respecting Dale's law.

The first step we will pursue is a change of variable, $\vec{y} = \hat{W}\vec{x}$. The dynamics of $y(t)$ are now

$$\dot{y}_i = -y_i + \sum_{j=1}^N W_{ij}\phi(y_j). \quad (2)$$

The difference between the variables x_i and y_i is that of rate activity to a membrane potential, and their dynamics have been shown to be equivalent [1]. In our case, y_i represents the input current to the i -th neuron.

The main issue is finding the fixed points of the system and the location of the critical point. If the sigmoidal function vanishes at the origin, $\phi(0) \equiv \phi_0 = 0$, then the system has a fixed point $y_i^* = 0$ for all i , whose stability is lost at $\lambda_c = 1/\phi'(0)$ (continue reading below). However, since we consider $\phi_0 > 0$ the fixed point will be a disordered state. In order to find its properties, we first perform a static mean-field analysis.

Static mean-field

The fixed points fulfill the stationary equation

$$y_i^* = \sum_{j=1}^N W_{ij}\phi(y_j^*), \quad (3)$$

which is transcendental and cannot be solved analytically.

The key point is to realise that we do not need the exact value of the fixed points, but rather their distribution, as all relevant observables will be averages over the disorder, i.e., averages over the distribution of y^* . Since the variables y_i are the input currents, and these are obtained by summing all the inputs to the i -th neuron, we can approximate their distribution by a Gaussian. We will describe the distributions by its mean, $\langle y^* \rangle$ and variance $\sigma_y^2 = \langle (y^*)^2 \rangle - \langle y^* \rangle^2$.

The moments can be computed self-consistently from equation (3). The main strategy is to assume that the weights and the currents are independent from each other, so $\langle W_{ij}y_j^* \rangle \simeq \langle W_{ij} \rangle \langle y_j^* \rangle$. Thus, taking averages in eq. (3) and using that $\langle W_{ij} \rangle = 0$, one immediately




FIG. 1. **Mean-field for a single population.** **a)** Comparison of the simulated (step histogram) fixed points of the currents y^* with theoretical prediction (continuous line) for a network of $\lambda = 6$, $N = 3000$. **b).** Variance of the trajectory of $y(t)$, which serves as an indicator for a transition to chaos. Critical point is marked with a dashed vertical line. Size $N = 300$ to match the reservoirs. Other parameters: $r = 0.7$

finds $\langle y^* \rangle = 0$, so $\sigma_y^2 = \langle (y^*)^2 \rangle$ directly. The variance of the distribution can be obtained by multiplying eq. (3) by itself and taking averages,

$$\begin{aligned} \langle (y_i^*)^2 \rangle = \sigma_y^2 &= \left\langle \left(\sum_j W_{ij} \phi(y_j^*) \right)^2 \right\rangle = \sum_{j,k} \langle W_{ij} W_{ik} \rangle \langle \phi(y_j^*) \phi(y_k^*) \rangle = \\ &= \sum_{j,k} \frac{\lambda^2}{N} \delta_{jk} \langle \phi(y_j^*) \phi(y_k^*) \rangle = \lambda^2 \langle \phi(y^*)^2 \rangle. \end{aligned} \quad (4)$$

where we used that $\langle W_{ij} W_{ik} \rangle = \frac{\lambda^2}{N} \delta_{jk}$ and that $\langle \phi(y_j^*)^2 \rangle = \langle \phi(y^*)^2 \rangle$ for every j .

We still need to deal with the non-linearity of the right-hand side of the equation. However, observe that we assumed that the y^* are Gaussian, so by definition of average value

$$\sigma_y^2 = \lambda^2 \int_{-\infty}^{+\infty} \frac{dy}{\sqrt{2\pi\sigma_y^2}} \exp\left(-\frac{y^2}{2\sigma_y^2}\right) [\phi(y)]^2. \quad (5)$$

This is an integral equation for σ_y^2 that can be numerically solved by e.g. an iterative method.

An example of the determination of the mean and variance of the currents and the critical point can be found in Fig. 1a.

Dynamic Mean-Field

We now construct a dynamic mean field theory for the Gaussian network. We are interested to see when our fixed point y_i^* loses its stability. The way to do so is linearising the system. Making a change of variable $y_i = y_i^* + \varepsilon_i$, expanding the sigmoid function $\phi(y_i)$ to first order around y_i^* and using the stationary state equation (3) leads to

$$\dot{\varepsilon}_i = -\varepsilon_i + \sum_{j=1}^N W_{ij} \phi'(y_j^*) \varepsilon_j \quad (6)$$

This is equivalent to the computation of the system's Jacobian. We'll employ basic concepts from dynamical mean-field theory (DMFT) to see where the bifurcation happens. The idea behind DMFT is to assume that the input to neuron i is modeled by a Gaussian stochastic process whose correlations match those of the system. In our case it is defined as

$$\eta_i(t) = \sum_{j=1}^N W_{ij} \phi'(y_j^*) \varepsilon_j. \quad (7)$$

so that our system becomes linear,

$$\dot{\varepsilon}_i = -\varepsilon_i + \eta_i(t). \quad (8)$$

The averages over this noise should be taken with respect to time, as well as with respect to the disorder, and we assume both averages can be taken independently of each other. The mean of the noise is zero, since $\langle W_{ij} \rangle = 0$. The autocorrelation function is given by

$$\begin{aligned} \langle \eta_i(t) \eta_i(t + \tau) \rangle &= \sum_{j,k} \langle W_{ij} W_{ik} \rangle \langle \phi'(y_j^*) \phi'(y_k^*) \rangle \langle \varepsilon_j(t) \varepsilon_k(t + \tau) \rangle = \\ &= \lambda^2 \langle (\phi'(y^*))^2 \rangle \cdot \frac{1}{N} \sum_j \langle \varepsilon_j(t) \varepsilon_j(t + \tau) \rangle = \\ & \qquad \qquad \qquad \lambda^2 \langle (\phi'(y^*))^2 \rangle \Delta(\tau). \quad (9) \end{aligned}$$

The autocorrelation of the noise is related to the autocorrelation of the ε_i variables, given by $\Delta(\tau) = N^{-1} \sum_j \langle \varepsilon_j(t) \varepsilon_j(t + \tau) \rangle$. We can write a differential equation for $\Delta(\tau)$ by Fourier transforming eq. (8), then multiplying it by its complex conjugate, taking averages and transforming back to the time domain [2]. After the process is done, we are left with

$$\ddot{\Delta} = \Delta - \langle \eta(t)\eta(t + \tau) \rangle = \left(1 - \lambda^2 \langle (\bar{\phi}'(y^*))^2 \rangle\right) \Delta(\tau). \quad (10)$$

We have found a condition for the destabilization of the autocorrelation function of the ε variables, which will mark the transition to chaos. Formally showing that this transition leads to a chaotic state is more complicated and outside of the scope of this paper, it has been previously demonstrated [2, 3].

Therefore, the critical condition for chaos is given by

$$\lambda_c = \sqrt{\frac{1}{\langle \phi'(y^*)^2 \rangle}}. \quad (11)$$

The difficulty of evaluating the location of the critical point passes by evaluating the average of the derivative of the sigmoidal function at the fixed points. Notice that there are two remarkable cases that are worth to check separately.

The first case is the sigmoidal vanishing at the origin, $\phi(0) = 0$. In this case, the fixed point is $y_i^* = 0$ for all nodes. Here, one sees that $\langle \phi'(y^*)^2 \rangle = \phi'(0)^2$, so the critical point is exactly located at $\lambda_c = 1/\phi'(0)$. This classical result can be obtained even from simple stability analysis. The second one is a tightly balanced network. Here, after generating the Gaussian random weights, these are adjusted so $\sum_j W_{ij} = 0$ exactly. If one decomposes the nonlinearity as $\phi(y) = \bar{\phi}(y) + \phi_0$, it follows immediately that $y_i^* = 0$ is again a fixed point, recovering $\lambda_c = 1/\phi'(0)$.

In our case, the average can be integrated using the mean and variance previously obtained in the static mean field analysis,

$$\langle \phi'(y^*)^2 \rangle = \int_{-\infty}^{+\infty} \frac{dy}{\sqrt{2\pi\sigma_y^2}} \exp\left(-\frac{y^2}{2\sigma_y^2}\right) [\phi'(y)]^2. \quad (12)$$

This is one of the main results of the analytics, but it needs to be extended to account for different interconnected populations. A comparison in the single population case can be seen in Fig. 1a.

B. Modular Networks

Set up

Up to this point, we have considered a homogeneous, Gaussian random matrix \hat{W} with zero mean. We will now expand the results above by considering that weights respect Dale's law, and that they are distributed in M modules. We will assume that each module has an excitatory and inhibitory population, giving in total $2M$ populations. In order to differentiate the modules, we add subscripts to the indices of the matrices and vectors, so that a_i^μ indicates that i -th element of the μ -th module. We will assume that the distribution of weights in each population can be approximated by a Gaussian. Thus, we assume that the connectivity matrix of each module –each block in the connectivity matrix \hat{W} – can be approximated by a Gaussian random matrix. Under this assumption, our network can be written as

$$W_{ij}^{\mu\nu} = \frac{w_{\mu\nu}}{N} + \frac{\sigma_{\mu\nu}}{\sqrt{N}} z_{ij}^{\mu\nu}, \quad (13)$$

where $z_{ij}^{\mu\nu} \sim \mathcal{N}(0, 1)$, and $w_{\mu\nu}, \sigma_{\mu\nu}$ are $2M \times 2M$ matrices which encode the average and variance of modules. The N, \sqrt{N} factors ensure that the network's eigenvalues do not grow with the system size, converging in the thermodynamic limit. Contrary to our homogeneous case, we allow the network to have non-zero mean.

The first step is to find the expression of $w_{\mu\nu}$ and $\sigma_{\mu\nu}$ that better approximate our reservoir's clustered networks. Our network is sparse, with a certain probability p_c to connect two neurons. If two neurons are connected, then its weight is sampled from a uniform distribution, between $[w_0, w_0 + \Delta w]$. The mean and variance of such a uniform distribution are given by $w = w_0 + \Delta w/2$ and $\sigma_w^2 = (\Delta w)^2/12$, respectively. The resulting mean and variance can be found by the law of total variance, giving $w p_c$ for the mean and $p_c (\sigma_w^2 + w^2(1 - p_c))$ for the variance.

The probability of connection p_c also depends on the neurons being on the same module or not. Neurons that belong to the same cluster are connected with probability $p_{\text{in},X} = p(1 + (M - 1)r_X)/M$, while those connecting among different modules are connected with probability $p_{\text{out},X} = p(r_X - 1)/M$, with X referring to the kind of module (excitatory-excitatory, inhibitory-excitatory, etc.).

With all this information, it is possible to write the reduced matrices. Using only two different modularity levels for the E and I population, r_e and r_i , they read

$$w_{\mu\nu} = \begin{cases} w_0 p_{\text{in},E}, & \nu < M, \quad \mu = \nu \pmod{M}, \\ w_0 p_{\text{out},E}, & \nu < M, \quad \mu \neq \nu \pmod{M}, \\ -w_0 p_{\text{in},I}, & \nu > M, \quad \mu = \nu \pmod{M}, \\ -w_0 p_{\text{out},I}, & \nu > M, \quad \mu \neq \nu \pmod{M}, \end{cases} \quad (14)$$

for the means and

$$\sigma_{\mu\nu} = \begin{cases} p_{\text{in},E} (\sigma_w^2 + w_0^2 (1 - p_{\text{in},E})), & \nu < M, \quad \mu = \nu \pmod{M}, \\ p_{\text{out},E} (\sigma_w^2 + w_0^2 (1 - p_{\text{out},E})), & \nu < M, \quad \mu \neq \nu \pmod{M}, \\ p_{\text{in},I} (\sigma_w^2 + w_0^2 (1 - p_{\text{in},I})), & \nu > M, \quad \mu = \nu \pmod{M}, \\ p_{\text{out},I} (\sigma_w^2 + w_0^2 (1 - p_{\text{out},I})), & \nu > M, \quad \mu \neq \nu \pmod{M}, \end{cases} \quad (15)$$

for the variances. For simplicity, from this point on we work on the diagonal $r_e = 1 - r, r_i = r$.

Now that we have characterized our reduced matrices, we can start analysing the spectra of \hat{W} . The dynamical behaviour of our system will be dominated by the largest eigenvalue of the full-network \hat{W} . Fortunately, for large N the largest eigenvalues can be predicted by the spectra of the reduced matrices \hat{w} and $\hat{\sigma}$.

The largest eigenvalue of $w_{\mu\nu}$ corresponds to a discrete outlier of \hat{W} , while the square root of the largest eigenvalue of $\sigma_{\mu\nu}^2$ yields its spectral radius. For two modules, $M = 2$ both can be easily computed, since they are just the eigenvalues of 4×4 matrices,

$$\lambda_{\text{ot}} = p(2r - 1)w_0, \quad (16)$$

$$\lambda_{\text{sr}} = \sqrt{\sigma_w^2 + w_0^2(1 - p)}. \quad (17)$$

Observe that although \hat{W} is balanced in the thermodynamic limit (since $\sum_{\nu} w_{\mu\nu} = 0$) the outlier can get big and positive, which is due to the structure of the network. In contrast, a perfectly balanced network will always have $\lambda_{\text{ot}} = 0$. In our case, if r is small enough, the isolated outlier can become stable, leading to a saddle-node bifurcation, which does not cause any chaotic dynamics. For chaos to happen, the spectral disk must cross the

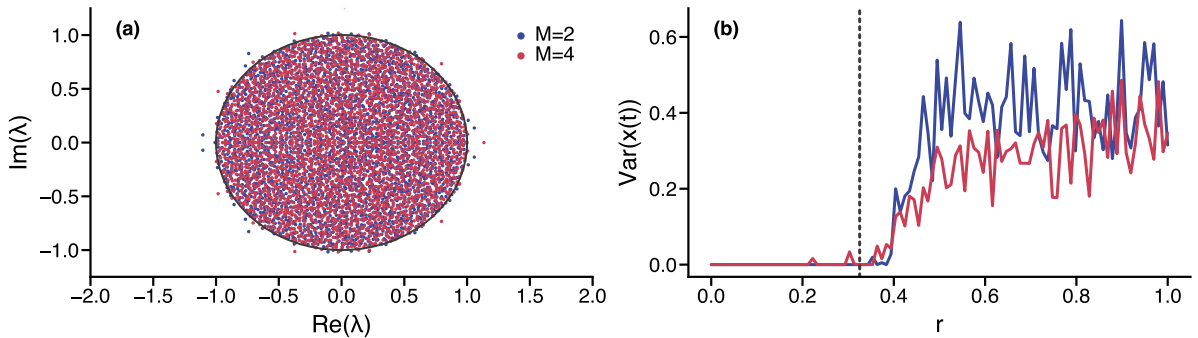


FIG. 2. **Network’s dynamical properties are independent of number of modules** (a) Spectrum two networks with 2 and 4 modules and $N = 3000$, normalized to have unit spectral radius. (b) Variance of the individual rate, which indicates a chaotic transition for fixed spectral radius $\lambda = 10$ and increasing modularity for the two previous networks. The transition to chaos happens at λ_{ot} (eq. 18), which is the value computed for $M = 2$.

instability. Therefore, a necessary condition for chaos is that the outlier λ_{ot} must be inside of the spectral radius, which leads to a critical overlap coefficient,

$$r_c = \frac{1}{2} - \frac{1}{2} \sqrt{\frac{\sigma_w^2 + w_0^2(1-p)}{pw_0}}. \quad (18)$$

It is important to remark that although λ_{ot} becomes negative at $r = 1/2$, when E and I are symmetric, the outlier blends with the spectral disk before, thus rendering the ”asymmetry” in the bifurcation point which benefits a structure where E is clustered and I more homogeneous.

This result is independent of the spectral radius of \hat{W} . One can repeat the above calculation by re-scaling the full matrix to have any desired spectral radius, since it enters linearly in both λ_{ot} and λ_{sr} . Hence, r_c does not depend on the spectral radius: one just needs to impose the spectral radius to be large enough in order to cross the instability, and r_c becomes the ”edge to chaos”. In this case, networks present a very reliable transition to chaos, which boosts the performance of reservoirs. Interestingly, the value r_c obtained for $M = 2$ seems to be a good approximation for larger values of M , such as $M = 4$ used in the main text. A phase diagram as a function of r_c for different numbers of modules can be seen in Fig. 2.

Static Mean Field

In the single population case, we assumed that all neurons received a Gaussian input with the same mean and average. In order to extend the calculation to multiple populations, the key idea is to allow each population to receive an input with different mean and variance. The input received in different populations might be correlated. Hence, we go from a unidimensional Gaussian distribution for the input to a multidimensional one [4, 5].

The dynamics of the network can be written in our new notation as

$$\dot{y}_i^\mu = -y^\mu + \sum_{\nu=1}^{2M} \sum_{j=1}^{N_\nu} W_{ij}^{\mu\nu} \phi(y_j^\nu), \quad (19)$$

where N_μ represents the number of neurons in population μ . The averages of the input distribution at stationary state will be represented by $\langle y_\mu^* \rangle$, and are given by

$$\begin{aligned} \langle y_i^{*\mu} \rangle = \langle y_\mu^* \rangle &= \sum_{\nu} \sum_j \langle W_{ij}^{\mu\nu} \phi(y_j^{*\nu}) \rangle = \sum_{\nu} \sum_j \left[\frac{w_{\mu\nu}}{N} + \frac{\sigma_{\mu\nu}}{\sqrt{N}} \langle z_{ij}^{\mu\nu} \rangle \right] \langle \phi(y_j^{*\nu}) \rangle = \\ &= \sum_{\nu} \frac{w_{\mu\nu}}{N} \sum_j \langle \phi(y_j^\nu) \rangle = \sum_{\nu} a_\nu w_{\mu\nu} \langle \phi(y^{*\nu}) \rangle, \quad (20) \end{aligned}$$

where we have defined $a_\mu = N_\mu/N$ as the fraction of individuals in population μ . Notice that the input to population μ is the sum of the the average rate of the other populations weighted by the average connectivity $w_{\mu\nu}$ connecting both populations, as one would expect.

The computation of the variance now is more complicated. For it, we need to employ an assumption which is common in the physics of disordered neuronal networks, which is that the rates at different sites are uncorrelated, $\langle \phi(y_i^\mu) \phi(y_j^\nu) \rangle \simeq \langle \phi(y_i^\mu) \rangle \langle \phi(y_j^\nu) \rangle$ when $\mu \neq \nu$ and $i \neq j$. Although this is also true for the single population, we did not need to apply it before. Here, it will help to simplify a lot of expressions and find closed equations for the variance. In order to light the notation, we will write $\phi(y_i^{*\mu}) \equiv \phi_i^\mu$. Then, the second moment is

$$\begin{aligned}
\langle y_i^{*\mu} y_j^{*\nu} \rangle &= \sum_{\alpha\beta} \sum_{kl} \langle W_{ik}^{\mu\alpha} W_{jl}^{\nu\beta} \rangle \langle \phi_k^\alpha \phi_l^\beta \rangle = \sum_{\alpha\beta} \sum_{kl} \left[\frac{w_{\mu\alpha} w_{\nu\beta}}{N^2} + \frac{\sigma_{\mu\alpha} \sigma_{\nu\beta}}{N} \langle z_{ik}^{\mu\alpha} z_{jl}^{\nu\beta} \rangle \right] \langle \phi_k^\alpha \phi_l^\beta \rangle = \\
&= \sum_{\alpha\beta} \left[\frac{w_{\mu\alpha} w_{\nu\beta}}{N^2} \sum_{kl} \langle \phi_k^\alpha \phi_l^\beta \rangle + \delta_{ij} \delta_{\mu\nu} \delta_{\alpha\beta} \frac{\sigma_{\mu\alpha} \sigma_{\nu\beta}}{N} \sum_{kl} \delta_{kl} \langle \phi_k^\alpha \phi_l^\beta \rangle \right] = \\
&\qquad \sum_{\alpha\beta} \frac{w_{\mu\alpha} w_{\nu\beta}}{N^2} \sum_{kl} \langle \phi_k^\alpha \phi_l^\beta \rangle + \delta_{ij} \delta_{\mu\nu} \sum_{\alpha} a_{\alpha} \sigma_{\mu\alpha}^2 \langle \phi_{\alpha}^2 \rangle. \quad (21)
\end{aligned}$$

Observe that we have not worked out the first term, as it will simplify when the averages are subtracted to compute the covariance. Performing the computations yields

$$\begin{aligned}
\langle y_i^{*\mu} y_j^{*\nu} \rangle - \langle y_{\mu}^* \rangle \langle y_{\nu}^* \rangle &= \sum_{\alpha\beta} \frac{w_{\mu\alpha} w_{\nu\beta}}{N^2} \sum_{kl} \left[\langle \phi_k^\alpha \phi_l^\beta \rangle - \langle \phi_k^\alpha \rangle \langle \phi_l^\beta \rangle \right] + \delta_{ij} \delta_{\mu\nu} \sum_{\alpha} a_{\alpha} \sigma_{\mu\alpha}^2 \langle \phi_{\alpha}^2 \rangle = \\
&\qquad \sum_{\alpha} a_{\alpha} \left(\frac{w_{\mu\alpha} w_{\nu\alpha}}{N} [\langle \phi_{\alpha}^2 \rangle - \langle \phi_{\alpha} \rangle^2] + \delta_{ij} \delta_{\mu\nu} \sigma_{\mu\alpha}^2 \langle \phi_{\alpha}^2 \rangle \right). \quad (22)
\end{aligned}$$

In order to simplify the first term, one breaks down the summation over α, β in two terms, first for $\alpha = \beta$ and then for $\alpha \neq \beta$. The second term always refers to different neurons, so the covariance for this whole term vanishes. In the case $\alpha = \beta$, we again break down the sum for $k = l$ and $k \neq l$, cancelling the second term due to our approximation. Finally, we can take out the indices i, j , since all the covariances are equal between neurons that belong to the same populations. The result is

$$\langle y_{\mu}^* y_{\nu}^* \rangle - \langle y_{\mu}^* \rangle \langle y_{\nu}^* \rangle = \sum_{\alpha} a_{\alpha} \left(\frac{w_{\mu\alpha} w_{\nu\alpha}}{N} [\langle \phi_{\alpha}^2 \rangle - \langle \phi_{\alpha} \rangle^2] + \delta_{\mu\nu} \sigma_{\mu\alpha}^2 \langle \phi_{\alpha}^2 \rangle \right). \quad (23)$$

Some observations apply now. If we assume we have a single population with no mean, we recover the result from the previous section. We see that the role of spectral radius λ is now played by the reduced matrix $\hat{\sigma}$, and that the averages of different populations contribute only to the covariances, which scale with a term $\mathcal{O}(1/N)$. This means that in the thermodynamic limit, correlations between the input in different populations will be weak compared to the variance, and the first term can be dropped. Hence, we end up with equations only for the variance,

$$\langle (y_{\mu}^*)^2 \rangle - \langle y_{\mu}^* \rangle^2 = \sum_{\alpha} a_{\alpha} \sigma_{\mu\alpha}^2 \langle \phi_{\alpha}^2 \rangle. \quad (24)$$




FIG. 3. **Mean-field for an example of a single module with E and I population network.**

a. Comparison of the simulated (step histogram) fixed points of the currents y^* with theoretical prediction (continuous line) for a network consisting of $M = 2$ Gaussian populations of averages $w_{11} = w_{21} = +1$, $w_{12} = w_{22} = -1$ and $\sigma_{\mu\nu} \sim \mathcal{N}(0, \lambda)$, size $N = 3000$. Each color represents the result for a different population. **b.** Variance of the trajectory of $y(t)$, which serves as an indicator for a transition to chaos. Critical point is marked with a dashed vertical line. Size $N = 300$ to match the reservoirs.

Hence, equations for mean and variance lead to a that this is a set of $2M$ nonlinear integral equations that need to be solved numerically, since the averages $\langle \phi_\mu \rangle$ are integrals over a Gaussian distribution with the estimated statistics.

Once we obtained the probability distribution for the fixed points, we go back again to the dynamic mean field analysis. Fig. 3a shows how the theory agrees with simulations for a simplified case, while Fig. 4 shows results for our modular networks.

Dynamic Mean Field

Again, the idea is study the behaviour of perturbations around the fixed point. Writing $y_i^\mu = y_i^{*\mu} + \varepsilon_j^\mu$ one can Taylor expand the non-linearities around the fixed point, finding the system's Jacobian. As in the case of a single population, the input to the perturbations is interpreted as noise matching the correlations of the system. The main difference with the previous case is that now the noise is a $2M$ -dimensional vector,

$$\eta_\mu(t) = \sum_{\nu=1}^{2M} \sum_{j=1}^{N_\nu} W_{ij}^{\mu\nu} \phi'(y_j^{*\mu}) \varepsilon_j. \quad (25)$$

The average of the noise is now non-zero,

$$\begin{aligned}
\langle \eta_\mu(t) \rangle &= \sum_\nu \sum_j \langle W_{ij}^{\mu\nu} \phi'(y_j^{*\nu}) \varepsilon_j^\nu \rangle = \frac{1}{N} \sum_\nu w_{\mu\nu} \sum_j \langle \phi'(y_j^{*\nu}) \rangle \langle \varepsilon_j^\nu \rangle = \\
&= \sum_\nu w_{\mu\nu} a_\nu \langle \phi'(y_\nu^*) \rangle \langle \varepsilon_\nu \rangle. \quad (26)
\end{aligned}$$

Again, the averages are assumed be all equal to a single value, going out of sum, yielding the number N_ν of neurons in population ν .

Similarly to the static case, the second moment of the noise can be shown to be

$$\begin{aligned}
H_{\mu\nu} &= \langle \eta_\mu(t) \eta_\nu(t + \tau) \rangle = \\
&\sum_{\alpha\beta} \frac{w_{\mu\alpha} w_{\nu\beta}}{N^2} \sum_{kl} \langle \phi'_k{}^\alpha \phi'_l{}^\beta \rangle \langle \varepsilon_k^\alpha(t) \varepsilon_l^\beta(t + \tau) \rangle + \delta_{\mu\nu} \sum_\alpha a_\alpha \sigma_{\mu\alpha}^2 \langle \phi_\alpha'^2 \rangle \langle \varepsilon_\alpha(t) \varepsilon_\alpha(t + \tau) \rangle. \quad (27)
\end{aligned}$$

The first term is identical to that of the static mean field, and it can worked out the same way to find again a contribution $\mathcal{O}(1/N)$ to noise. Then, we can drop this term, so μ and ν are not correlated. Setting $H_\mu \equiv H_{\mu\mu}$ and following the same procedure used for a single population, the differential equation for the covariance Δ_μ can be directly generalized by using H_μ ,

$$\ddot{\vec{\Delta}} = \vec{\Delta} - \vec{H} = (\hat{I} - \hat{S}) \vec{\Delta}, \quad (28)$$

where $S_{\mu\nu} = \sigma_{\mu\nu}^2 \langle \phi'(y_\nu^*)^2 \rangle a_\nu$. Therefore, the destabilization of the autocorrelation happens when the largest eigenvalue of of \hat{S} , λ_S becomes larger than unity, generalizing the previous results.

C. Numerical implementation

The implementation has two main challenges: first, solving the self-consistent equations for the mean and covariances of the fixed points, and then use these in order to find the location of the critical point.

The solution of the fixed point equations can be easily implemented through an iterative algorithm,

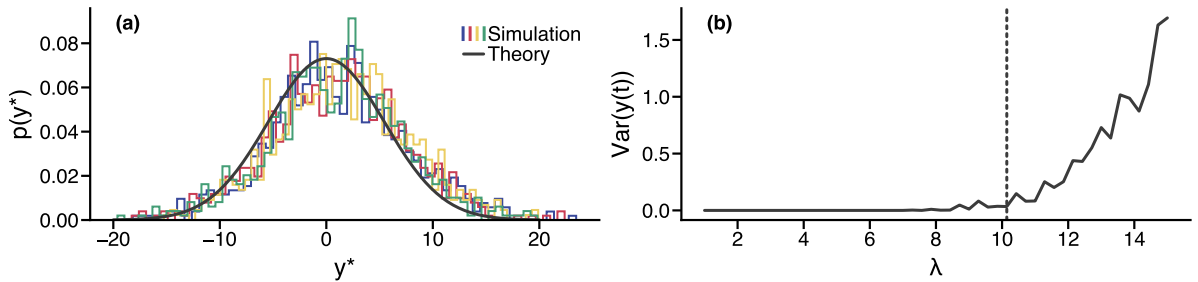


FIG. 4. **Mean-field for a modular network.** **a)** Comparison of the simulated (step histogram) fixed points of the currents y^* with theoretical prediction (continuous line) for a network consisting of $M = 4$ modules, size $N = 3000$. Theory predicts all modules having a similar distribution, and hence only the average distribution is shown. Each color represents the result for a different population. **b.** Variance of the trajectory of $y(t)$, which serves as an indicator for a transition to chaos. Critical point is marked with a dashed vertical line. Size $N = 300$ to match the reservoirs. Other parameters: $r = 0.7$

$$\langle y_\mu^* \rangle(n+1) = f(\langle y_\mu^* \rangle(n), \Sigma_{\mu\nu}^*(n)), \quad (29)$$

$$\Sigma_{\mu\nu}^*(n+1) = g(\langle y_\mu^* \rangle(n), \Sigma_{\mu\nu}^*(n)), \quad (30)$$

since all fixed point equations are already casted in this form. We found this method to be very fast and get accurate solutions for all cases. The dependence of the right hand side on the mean and averages are through averages such as $\langle \phi_\mu^2 \rangle$. For the single population case, we solve the averages by integrating with the trapezoid rule. For multiple populations, doing the integrals would be very costly, so we use a Monte Carlo method: we generate a large number of Gaussian random variables with the multivariate normal of means $\langle y_\mu^* \rangle(n)$ and covariances $\Sigma_{\mu\nu}^*(n)$, apply the nonlinearity that we desire to integrate, and then average the result.

For the determination of the critical point, on single populations it suffices to fix a spectral radius λ , solve the iterative equations to obtain the mean and variance of the network, and compute $F(\lambda) = 1 - \lambda^2 \langle \phi^2 \rangle$. Then, one finds the root $F(\lambda) = 0$, for example by plotting it.

On multiple populations, the situation becomes way more convoluted, as there is no obvious single control parameter. Since the spectral radius of the matrix depends only on the largest eigenvalue of $\sigma_{\mu\nu}$, a clean solution is to scale $\lambda \sigma_{\mu\nu}$, leaving $w_{\mu\nu}$ untouched, as




FIG. 5. **All transitions lie in the $r_e = 1 - r_i$ diagonal.** (a) Average variance of the individual trajectories of neurons in the network for two spectral radius. The large one (left) shows a transition to chaos, displayed by the increase in each trajectory’s variance. The smaller one (right) has no transitions. (b) Variance of the stationary rate values for the network. In both cases, a transition can be seen, symmetric with respect to the desired diagonal.

we do in Fig. 3, which gives very clean solutions. However, for the reservoirs we want to ensure that the largest eigenvalue is set to the selected λ value. Hence, for all our experiments we divide the sampled network by the largest eigenvalue and then multiply by the target λ . In the Gaussian approximation, where all eigenvalues are inside of the spectral disk[6] this procedure sets the spectral radius, but in the original matrix, which respects Dale’s law, a small amount of outliers can fall outside. Thus we would not normalize our theory networks the same way we do with the reservoirs. To alleviate this problem, we can compute the mismatch between these two results. Then, we compute the spectral radius as the square root of the largest eigenvalue λ_{sr} of $\hat{\sigma}^2$ and the average largest eigenvalue of our reservoir network $\langle \lambda_{\text{max}} \rangle$. If no outlier fell outside the spectral disk, these two would coincide. However, the reservoir network’s average largest eigenvalue is slightly larger, giving $\lambda_{\text{sr}} / \langle \lambda_{\text{max}} \rangle \simeq 0.92$. We can multiply our critical point prediction, shown in Fig. 4b, by this ratio in order to produce the prediction for criticality in reservoirs seen in the main text.

II. Dynamical state in r_e, r_i phase space

We have controlled that all the dynamical regimes of interest are in the $r_e = r, r_i = 1 - r$ diagonal. To show this, we run the dynamics of the recurrent network 100 times for each value of r_e and r_i and check that the results are symmetric to the diagonal of interest.

In order to detect different states, we obtain the average variance of individual trajectories, as well as the variance of the stationary firing rates across all neurons for two different spectral radius. One can see (Fig. 5) that for small values of the spectral radius, there's no chaotic state (average variance of trajectories is always low), and a transition from monostable to bistable can be found, hinted by the change in rate's variance. On the other hand, with high spectral the transition to chaos is visible by an increase of single trajectory variance. Variance across system's rates is not so informative, as this measure assumes rates are all stationary –which is false for the chaotic case– but the transition is seen due to increase of spread of rate's mean. Observe that the shape of all transitions is such that they are captured in the diagonal of interest, which justifies working in a single dimension for the modularity.

III. Attractors

We use the following 3-dimensional chaotic attractors as input time series (individually or all of them simultaneously) to our network:

Lorenz:

$$\begin{aligned}\dot{x} &= \sigma(y - x), \\ \dot{y} &= \rho x - y - xz, \\ \dot{z} &= xy - \beta z,\end{aligned}\tag{31}$$

for $\sigma = 10$, $\rho = 28$ and $\beta = 2.667$

Aizawa:

$$\begin{aligned}\dot{x} &= (z - b)x - y, \\ \dot{y} &= x + (z - b)y, \\ \dot{z} &= c + az - z^3/3 - (x^2 + y^2)(1 + ez) + fzx^3,\end{aligned}\tag{32}$$

for $a = 0.95$, $b = 0.7$, $c = 0.6$, $d = 3.5$, $e = 0.25$ and $f = 0.1$

Halvorsen:

$$\begin{aligned}
\dot{x} &= -\alpha x - 4y - 4z - y^2, \\
\dot{y} &= -\alpha y - 4x - 4z - z^2, \\
\dot{z} &= -\alpha z - 4x - 4y - x^2,
\end{aligned} \tag{33}$$

for $\alpha = 1.4$

Chen:

$$\begin{aligned}
\dot{x} &= a(y - x), \\
\dot{y} &= (c - a)x - xz + cy, \\
\dot{z} &= xy - bz,
\end{aligned} \tag{34}$$

for $a = 40$, $b = 3$ and $c = 28$

-
- [1] Kenneth D. Miller and Francesco Fumarola. Mathematical Equivalence of Two Common Forms of Firing Rate Models of Neural Networks. *24(1):25–31*.
 - [2] Moritz Helias and David Dahmen. *Statistical Field Theory for Neural Networks*. Lecture Notes in Physics. Springer Cham, 2020.
 - [3] Haim Sompolinsky and Andrea Crisanti. Chaos in Random Neural Networks. *Physical Review Letters*, 61(3):259–262, 1988.
 - [4] Johnatan Aljadeff, Merav Stern, and Tatyana Sharpee. Transition to Chaos in Random Networks with Cell-Type-Specific Connectivity. *Physical Review Letters*, 114(8):088101, February 2015.
 - [5] Johnatan Aljadeff, David Renfrew, Marina Vegué, and Tatyana O. Sharpee. Low-dimensional dynamics of structured random networks. *Physical Review E*, 93(2):022302, February 2016.
 - [6] Except the single outlier coming from $w_{\mu\nu}$, as discussed above. However, if this outlier is inside or to the left of the disk it plays no role in this discussion.

Appendix F Emergent mechanisms for long timescales depend on training curriculum and affect performance in memory tasks

Contains the full text of the publication: Khajehabdollahi, S., Zeraati, R., Giannakakis, E., Schäfer, T. J., Martius, G., and Levina, A. (2024). Emergent mechanisms for long timescales depend on training curriculum and affect performance in memory tasks. In *The Twelfth International Conference on Learning Representations*.

EMERGENT MECHANISMS FOR LONG TIMESCALES DEPEND ON TRAINING CURRICULUM AND AFFECT PERFORMANCE IN MEMORY TASKS

Sina Khajehabdollahi^{1,2,*}, **Roxana Zeraati**^{1,3,*}, **Emmanouil Giannakakis**^{1,3},
Tim J. Schäfer^{1,3}, **Georg Martius**^{1,2}, **Anna Levina**^{1,3}

¹ University of Tübingen, Germany

² Max Planck Institute for Intelligent Systems, Tübingen, Germany

³ Max Planck Institute for Biological Cybernetics, Tübingen, Germany

* These authors contributed equally to this work.

{firstname.lastname}@uni-tuebingen.de

ABSTRACT

Recurrent neural networks (RNNs) in the brain and *in silico* excel at solving tasks with intricate temporal dependencies. Long timescales required for solving such tasks can arise from properties of individual neurons (single-neuron timescale, τ , e.g., membrane time constant in biological neurons) or recurrent interactions among them (network-mediated timescale, τ_{net}). However, the contribution of each mechanism for optimally solving memory-dependent tasks remains poorly understood. Here, we train RNNs to solve N -parity and N -delayed match-to-sample tasks with increasing memory requirements controlled by N , by simultaneously optimizing recurrent weights and τ s. We find that RNNs develop longer timescales with increasing N , but depending on the learning objective, they use different mechanisms. Two distinct curricula define learning objectives: sequential learning of a single- N (single-head) or simultaneous learning of multiple N s (multi-head). Single-head networks increase their τ with N and can solve large- N tasks, but suffer from catastrophic forgetting. However, multi-head networks, which are explicitly required to hold multiple concurrent memories, keep τ constant and develop longer timescales through recurrent connectivity. We show that the multi-head curriculum increases training speed and stability to perturbations, and allows generalization to tasks beyond the training set. This curriculum also significantly improves training GRUs and LSTMs for large- N tasks. Our results suggest that adapting timescales to task requirements via recurrent interactions allows learning more complex objectives and improves the RNN’s performance.

1 INTRODUCTION

The interaction of living organisms with their environment requires the concurrent processing of signals over a wide range of timescales, from short timescales of coding sensory stimuli (Bathellier et al., 2008; Panzeri et al., 2010; Safavi et al., 2023) to longer timescales of cognitive processes like working memory (Jonides et al., 2008). The diverse timescales of these tasks are reflected in the dynamics of the neural populations performing the corresponding computations in the brain (Murray et al., 2014; Cavanagh et al., 2020; Gao et al., 2020; Zeraati et al., 2022). At the same time, artificial neural networks performing memory-demanding tasks (speech (Graves et al., 2013), handwriting (Graves, 2013), sketch (Ha & Eck, 2018), language (Bowman et al., 2015), time series prediction (Chung et al., 2014; Torres et al., 2021), music composition (Boulanger-Lewandowski et al., 2012)) need to process the temporal dependency of sequential data over variable timescales. Recurrent neural networks (RNNs) (Elman, 1990; Hochreiter & Schmidhuber, 1997; Lipton et al., 2015; Yu et al., 2019) have been introduced as a tool that can learn such temporal dependencies using back-propagation through time.

In biological networks, diverse neural timescales emerge via a variety of interacting mechanisms. Timescales of individual neurons in the absence of recurrent interactions are determined by cellu-

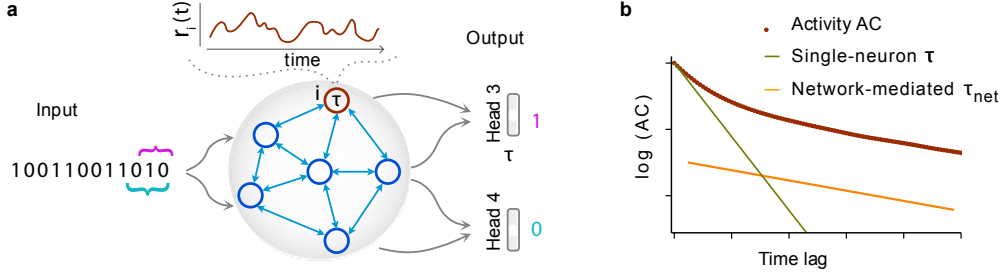


Figure 1: Schematics of network structure and timescales. **a.** An outline of the network. A binary sequence is given as input to a leaky RNN, with each neuron’s intrinsic timescale being a trainable parameter τ . The illustration shows the N -parity task with readout heads for $N = 3$ and $N = 4$. **b.** An illustration of the manifestation of different timescales (single-neuron and network-mediated) on the autocorrelation (AC) of a network neuron (see also Fig. S2).

lar and synaptic processes (e.g., membrane time constant) that vary across brain areas and neuron types (Gjorgjieva et al., 2016; Duarte et al., 2017). However, recurrent interactions also shape neural dynamics introducing network-mediated timescales. The strength (Ostojic, 2014; Chaudhuri et al., 2015; van Meegen & van Albada, 2021) and topology (Litwin-Kumar & Doiron, 2012; Chaudhuri et al., 2014; Zeraati et al., 2023; Shi et al., 2023) of recurrent connections give rise to network-mediated timescales that can be much longer than single-neuron timescales.

Heterogeneous and tunable single-neuron timescales have been proposed as a mechanism to adapt the timescale of RNN dynamics to task requirements and improve their performance (Perez-Nieves et al., 2021; Tallec & Ollivier, 2018; Quax et al., 2020; Yin et al., 2020; Fang et al., 2021; Smith et al., 2023b; Jain et al., 2020). In these studies, the time constants of individual neurons are trained together with network connectivity. For tasks with long temporal dependencies, the distribution of trained timescales becomes heterogeneous according to the task’s memory requirements (Perez-Nieves et al., 2021). Explicit training of single-neuron timescales improves network performance in benchmark RNN tasks in rate (Tallec & Ollivier, 2018; Quax et al., 2020) and spiking (Yin et al., 2020; Fang et al., 2021; Perez-Nieves et al., 2021) networks and leads to greater robustness (Perez-Nieves et al., 2021) and adaptability to novel stimuli (Smith et al., 2023b). While these studies propose the adaptability of single-neuron timescales as a mechanism for solving time-dependent tasks, the exact contribution of single-neuron and network-mediated timescales in solving tasks is unknown.

Here, we study how single-neuron and network-mediated timescales shape the dynamics and performance of RNNs trained on long-memory tasks. We show that the contribution of each mechanism in solving such tasks largely depends on the learning objective defined by the curriculum. Challenging common beliefs in the field, we identify settings where trainable single-neuron timescales offer no advantage in solving temporal tasks. Instead, adapting RNNs’ timescales using network-mediated mechanisms improves training speed, stability and generalizability.

2 MODEL

We approximate the effect of the membrane timescale of biological neurons by equipping each RNN-neuron with a trainable leak parameter τ , defining the single-neuron timescale (Fig. 1a). The activity of each neuron evolves over discrete time steps t governed by:

$$r_i(t) = \left[\left(1 - \frac{\Delta t}{\tau_i} \right) \cdot r_i(t - \Delta t) + \frac{\Delta t}{\tau_i} \cdot \left(\sum_{j \neq i} W_{ij}^R \cdot r_j(t - \Delta t) + W_i^I \cdot S(t) + b^R + b^I \right) \right]_{\alpha} \quad (1)$$

where $[\cdot]_{\alpha}$ is the leaky ReLU function with negative slope α , given by:

$$[x]_{\alpha} = \begin{cases} x, & x \geq 0 \\ \alpha \cdot x, & x < 0. \end{cases} \quad (2)$$

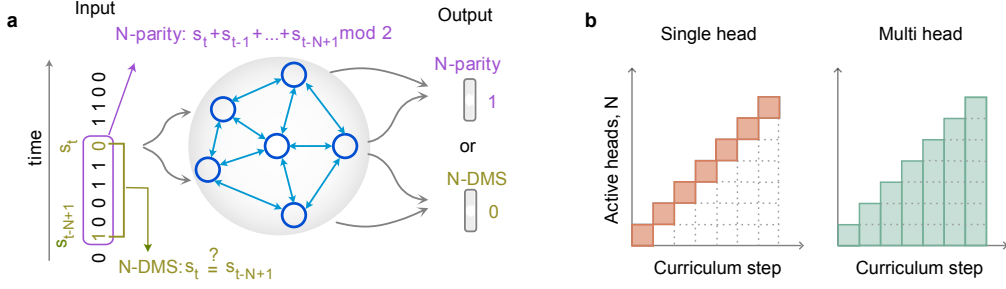


Figure 2: Schematic description of the tasks and curricula **a**. An outline of the network and tasks. In both tasks, the network receives a binary input sequence, one bit at each time step. **b**. In the single-head curriculum, only one read-out head is trained at each curriculum step, while in the multi-head curriculum, a new read-out head is added at each step without removing the older heads.

For all networks, we use $\alpha = 0.1$. We obtain similar results using a different type or location of nonlinearity (Appendix A). W^R, b^R , and W^I, b^I are the recurrent and input weights and biases, respectively, S is the binary input given to the network at each time step, and $\tau_i \geq 1$ is the trainable timescale of the neuron. Unless otherwise stated, the time step is $\Delta t = 1$ (other Δt discussed in Appendix B). Each RNN has 500 neurons. When $\tau = 1$, a neuron becomes memory-less (in isolation) as its current state does not directly depend on its past activity, i.e., memory can only be stored at the network level via interactions. In contrast, for $\tau > 1$, the neuron’s activity depends on its past activity, and the dependency increases with τ . In the limit of $\tau \rightarrow \infty$, the neuron’s activity is constant, and the input has no effect.

The dynamics of each neuron can be characterized by two distinct timescales: (i) single neuron timescale τ , (ii) network mediated timescale τ_{net} . τ gives the intrinsic timescale of a neuron in the absence of any network interaction, while τ_{net} is shaped by the combination of τ and the learned connectivity and represents the effective timescale of the neuron’s activity within the network. τ_{net} is generally a function of τ and recurrent weights: $\tau_{\text{net}} = f(\tau, W^R)$ (Ostojic, 2014; Chaudhuri et al., 2014; Shi et al., 2023) and $\tau_{\text{net}} \geq \tau$ (Shi et al., 2023). For networks with linear dynamics, τ_{net} can be directly estimated from the eigenvalues of the connectivity matrix normalized by τ (Chaudhuri et al., 2014). For nearest-neighbor connectivity or mean-field dynamics, it is also possible to derive τ_{net} analytically for nonlinear networks. However, a general analytical solution does not exist. Instead, τ_{net} can be effectively estimated from the decay rate of the autocorrelation (AC) function. The AC is defined as the correlation coefficient between the time series and its copy, shifted by time t' , called the time lag. For the neuron’s activity, it can be computed as

$$\text{AC}_i(t') = \frac{1}{\hat{\sigma}_i^2(T-t')} \sum_{t=0}^{T-t'} (r_i(t) - \hat{\mu}_i)(r_i(t-t') - \hat{\mu}_i), \quad (3)$$

where $\hat{\mu}_i$ and $\hat{\sigma}_i^2$ are the sample mean and variance of $r_i(t)$. To estimate τ_{net} , we drive the network by uncorrelated binary inputs sampled from a Bernoulli distribution.

The AC of a neuron’s activity, defined by Eq. 1, can be approximated by two distinct timescales which appear as two slopes in logarithmic-linear coordinates (Fig. 1b) (Shi et al., 2023). The steep initial slope indicates τ , and the shallower slope indicates τ_{net} . In the same way, we characterize the timescale of collective network dynamics by computing population activity (summed activity of all neurons within a network) timescale τ_{pop} , which reflects the timescale of network dynamics as a whole. To avoid AC bias in our estimates (Zeraati et al., 2022), we use long simulations ($T = 10^5$ time steps). We simulate each network for 10 trials (i.e. 10 distinct realizations of inputs) and compute the average AC of each neuron across trials. To estimate τ_{net} , we fit the average AC with a single- ($\tau_{\text{net}} = \tau$) and with a double-exponential ($\tau_{\text{net}} > \tau$) decay function using the nonlinear least-squares method. Then, we use the Akaike Information Criterion (AIC) (Akaike, 1974) to select the best-fitting model. For most neurons (above 95%), Bayesian information criterion (BIC) selects the same model (Fig. S1) and previous work (Pasula, 2023) indicates that AIC provides similar results as the sum of three information criteria (AIC, BIC and Hannan-Quinn information criteria).

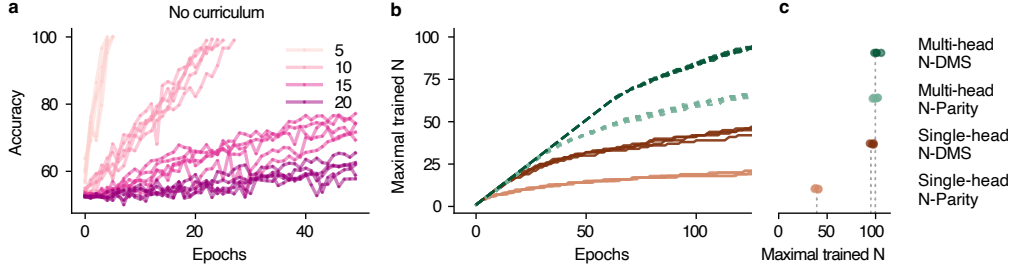


Figure 3: Training performance depends on the curriculum. **a.** Accuracy of training the networks (N -parity task) without a curriculum increases slowly, especially when $N > 10$. For each N , 5 models are independently trained for 50 epochs or until reaching $> 98\%$ accuracy. **b.** Multi-head (dashed) trained networks are solving larger N s than single-head (solid) within the same training time (colors in c). **c.** The maximum trained N for each task/curriculum at the end of training (1000 epochs or solving $N = 101$, whichever comes first). Gray lines - mean value across 4 networks.

When the double-exponential model is selected, the slowest of two timescales indicates τ_{net} . For most fits, we obtain a large coefficient of determination, confirming a good quality of fit (Fig. S2).

3 SETUP

3.1 TASKS

In both tasks (Fig. 2a), a binary sequence S is given as the input, one bit at each time step. We train the networks on sequences with lengths uniformly chosen from the interval $L \in \{N + 2, 4N\}$.

N -delayed match-to-sample (N -DMS): The network outputs 1 or 0 to indicate whether the digit presented at current time t matches the digit presented at time $t - N + 1$. To update the output at every time step, the network needs to store the values and order of the last N digits in memory.

N -parity: The network outputs the binary sum (XOR) of the last N digits. N -parity has a similar working memory component as N -DMS, but requires additional computations (binary sum).

3.2 TRAINING

We train single-neuron timescales $\tau = \{\tau_1, \dots, \tau_{500}\}$, W^R, b^R, W^I, b^I , and a linear readout layer via back-propagation through time using a stochastic gradient descent optimizer with Nesterov momentum and a cross-entropy loss. Each RNN is trained on a single Nvidia GeForce 2080ti for 1000 epochs, 3 days, or until the $N = 101$ task is solved, whichever comes first. RNNs are trained without any regularization. Including L2 regularization achieves comparable performance.

Single-head: Starting with $N = 2$, we train the network to reach an accuracy of 98%. We then use the trained network parameters to initialize the next network that we train for $N + 1$ (Fig. 2b).

Multi-head: As with the single-head networks, we begin with a network solving a task for $N = 2$, but once a threshold accuracy of 98% is reached, a new readout head is added for solving the same task for $N + 1$, preserving the original readout heads. At each curriculum step, all readout heads are trained simultaneously (the loss is the sum of all readout heads' losses) so that the network does not forget how to solve the task for smaller N s (Fig. 2b).

4 RESULTS

4.1 PERFORMANCE UNDER DIFFERENT CURRICULA

Necessity of curriculum: Our objective is to learn the largest possible N in each task. First, we test whether a good performance can be achieved for high N in either task without any curricula. We find that for both tasks (see Appendix E for N -DMS results), networks struggle to reach high accuracy for $N > 10$ (Fig. 3a, Fig. S13). However, using either curriculum significantly boosts the network's capacity to learn tasks with larger N (Fig. 3b).




Figure 4: Importance of single-neuron timescales for different curricula. **a.** The maximum N solved in the N -bit parity task after 1000 epochs (reaching an accuracy of 98%). X-label indicates training constraints: $\tau = 1, 2$ or 3 - fixed τ with only weights being trained, “Trained” allows training of τ . In the single-head curriculum, models rely on training τ , whereas in the multi-head curriculum, τ fixed $\tau = 1$ is as good as training τ . Horizontal bars - mean. **b.** The mean and standard deviation (STD) of the trained τ s increase with N in single-head networks. In contrast, in multi-head networks, the mean τ decreases towards 1, and the STD remains largely constant. The mean and STD are computed across neurons within each network (up to the maximum N shared between all trained networks). Shading - variability across 4 trained networks.

Comparison of curricula: Task performance differs significantly between curricula. The single-head networks can reliably reach $N \approx 35$ for the N -parity and $N \approx 90$ for the N -DMS task. The difference in performance between the two tasks is expected because the N -DMS task is much easier than the N -parity task. However, the multi-head networks can reliably reach $N \geq 100$ for both tasks and require fewer training epochs to reach 98% accuracy for each N (Fig. 3b, c). Networks that are trained using an intermediate curriculum between the two extremes of single- and multi-head (i.e., solving simultaneously $H < N$ tasks for $N, N-1, \dots, N-H+1$ with gradually increasing N s) exhibit an intermediate performance (Appendix F). Overall, the multi-head networks outperform the single-head ones in terms of performance (maximum N reached) and the required training time. Moreover, the multi-head curriculum significantly improves the training of other recurrent architectures such as GRU and LSTM to perform large- N tasks (Appendix G), suggesting that the multi-objective curriculum can generally improve learning long-memory tasks.

The superior performance of the multi-head networks may be counter-intuitive since they solve the task for every $N \leq m$ at the m -th step of the curriculum, whereas the single-head networks only solve it for exactly $N = m$. However, we find that the single-head networks suffer from catastrophic forgetting: a network trained for larger N cannot perform the task for smaller N s it was trained for, even after retraining the readout weights (Appendix H). These results suggest that explicit prevention of catastrophic forgetting by learning auxiliary tasks (i.e. tasks with N smaller than objective, $N < m$) facilitates learning large N s. Interestingly, training directly on a multi- N task without using an explicit curriculum results in the emergence of the multi-head curriculum: networks learn to first solve small- N tasks and then large- N tasks (Appendix I), supporting the use of the multi-head step-wise strategy.

Necessity of training τ : We examine the impact of training single-neuron timescales on training performance. We compare the training performance of networks with a fixed $\tau \in \{1, 2, 3\}$ shared across all neurons versus networks with trainable timescales. In the single-head curriculum, the training performance with fixed τ is significantly worse than when we train τ (Fig. 4a). On the other hand, in multi-head networks, training performance is the same for fixed $\tau = 1$ and trainable τ cases, but steeply declines for fixed $\tau \geq 2$ (Fig. 4b). See (Fig. S10) for N -DMS results. These results indicate that single-head networks rely on τ for solving the task, whereas multi-head networks only use τ to track the timescale of updating the input and rely on other mechanisms to hold the memory.

4.2 MECHANISMS UNDERLYING LONG TIMESCALES

To uncover the mechanisms underlying the difference between the two curricula, we study how τ , τ_{net} , τ_{pop} and recurrent weights change with increasing task difficulty N . One can expect that as the timescale of the task (mediated by N) increases, neurons would develop longer timescales to integrate the relevant information. Such long timescales can arise either by directly modulating τ for each neuron or through recurrent interactions between neurons reflected in τ_{net} and τ_{pop} .

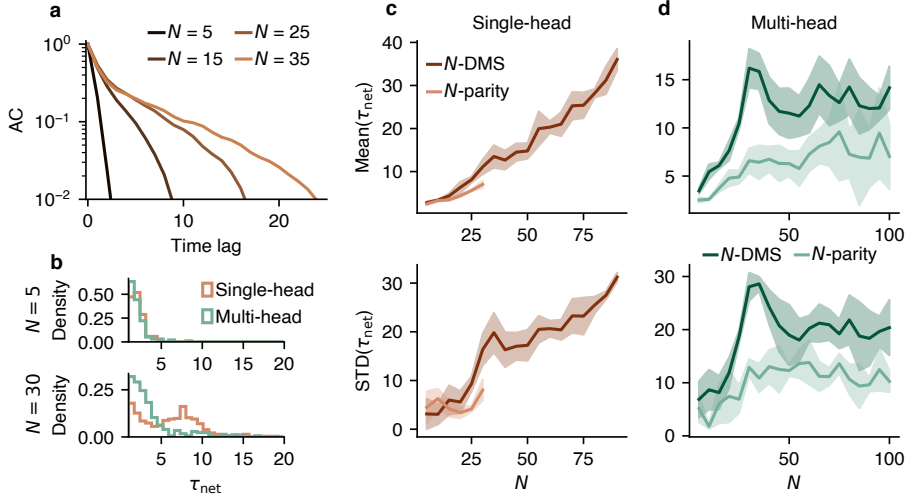


Figure 5: The emergence of network-mediated timescales depends on the curriculum. **a.** Example average ACs of all the neurons within a single-head network, N -parity task. The ACs of individual neurons’ activity decay slower with increasing N . **b.** Distributions of the network-mediated timescales τ_{net} for single and multi-head networks solving N -parity task for $N = 5$ and $N = 30$. The distribution becomes broader for higher N . **c, d.** The mean and STD of the network-mediated timescale τ_{net} increase with N in both tasks. The mean and STD are computed across neurons within each network. Shades - variability across 4 trained networks.

Dependence of τ on N : The two curricula adjust their τ to N in distinct ways: single-head networks increase their τ with N , but multi-head networks prefer $\tau \rightarrow 1$ (Fig. 4b). For the single-head curriculum, the mean and variance of τ increase with N , suggesting that not only τ s become longer as the memory requirement grows, but they also become more heterogeneous. We obtain similar results for networks trained without curriculum (Fig. S3). On the contrary, in multi-head networks, the average τ decreases with N , approaching $\tau = 1$. The trend of $\tau \rightarrow 1$ is consistent with the fact that multi-head networks with fixed $\tau = 1$ performed as well as networks with trained τ .

Dependence of τ_{net} and τ_{pop} on N : Network-mediated timescales τ_{net} and τ_{pop} generally increase with N in both curricula (Fig. 5, Fig. S4). τ_{net} reflects the contribution of τ and recurrent weights in dynamics of individual neurons. In single-head networks, the mean and variance of τ_{net} follow a similar trend as τ (Fig. 5c), suggesting that changes in τ_{net} can arise from changes in τ . The mean and variance of τ_{net} in multi-head networks increase with N up to some intermediate N , but the pace of increase reduces gradually and saturates for very large N s (Fig. 5d, top). Given the small τ in multi-head networks, long τ_{net} can only arise from recurrent interactions between neurons. τ_{pop} is the timescale of collective network dynamics (sum of all neurons’ activations) and arises from interactions between neurons within the whole network. τ_{pop} exhibits a clear increase with N for both tasks and curricula with comparable values (Fig. S4). These results indicate that in both curricula, collective network dynamics become slower with increasing N , but due to differences in τ , the two curricula employ distinct mechanisms to achieve this.

Dependence of connectivity on N : Multi-head networks have, on average, almost the same total incoming positive and negative weights (with a slight tendency towards larger total negative weights as N increases), leading to relatively balanced dynamics (Fig. 6a). On the other hand, single-head networks have a stronger bias towards more inhibition (negative weights) as N increases. The strong negative weights in single-head networks are required to create stable dynamics in the presence of long single-neuron timescales (Appendix J).

Dimensionality of dynamics: The dimensionality of population activity (measured as the number of principal components that explain 90% of the variance) increases almost linearly with N in the N -parity task but sub-linearly in the N -DMS task, using both curricula, reflecting distinct computational requirements for each task (Fig. 6b, S20, Appendix K). Since computations should be

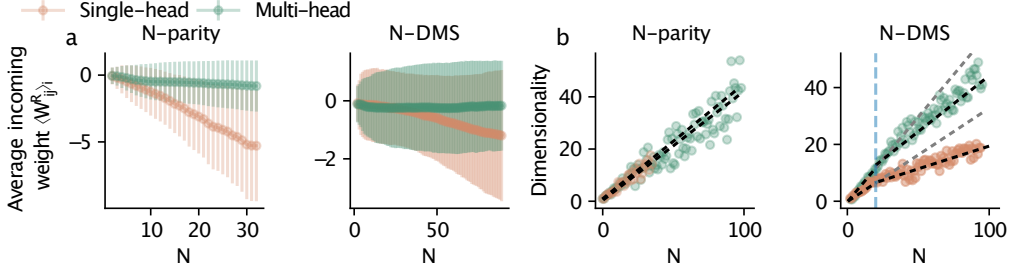


Figure 6: Learned recurrent connectivity and dimensionality of population activity. **a.** The average incoming weight to a neuron remains close to zero for multi-head networks but becomes strongly negative as N increases in single-head networks (example RNN). Error bars - \pm STD. **b.** The dimensionality of the activity increases linearly with N in N -parity and sub-linearly in N -DMS task. Dashed lines - linear fit computed for all N s (N -parity task) and independently for $N \in [0, 20]$ (gray line - extension for visual guidance) and $N \in [20, 100]$ (N -DMS task). Blue line - $N = 20$.

performed at every time step, the increase in dimensionality may be required to map different input patterns (that grow with N) to the same outputs.

Our findings suggest that both curricula give rise to networks with slower and higher dimensional collective dynamics with increasing N , but via distinct mechanisms. In single-head networks, single-neuron properties τ play an important role in creating slow dynamics, which are then stabilized by stronger inhibition in the network. However, in multi-head networks, the slow dynamics should arise from recurrent network interactions. The significant difference in performance of the two curricula suggests that the second mechanism is more effective in solving the task.

4.3 IMPACT OF DIFFERENT CURRICULA ON NETWORKS ROBUSTNESS

To compare the robustness and retraining capability between the two curricula, we investigate changes in network accuracy resulting from ablations, perturbations, and retraining networks on unseen N . We measure the effects on network performance using a relative accuracy metric with respect to the originally trained network, defined as $\text{acc}_{\text{rel}} := (\text{acc} - 0.5) / (\text{acc}_{\text{base}} - 0.5)$, where acc_{base} represents the accuracy of the network before any perturbations or ablations, acc is the measured accuracy after the intervention, and 0.5 is a chance level used for the normalization. If $\text{acc}_{\text{rel}} = 1$, the intervention did not change the accuracy; when $\text{acc}_{\text{rel}} \approx 0$, the intervention reduced the accuracy to chance level. All accuracies are evaluated on the maximal trained N .

Ablation: To examine the relative impact of neurons with different trained τ on network performance, we ablate individual neurons based on their τ and measure the performance without retraining (Appendix L). Specifically, we compare the effect of ablating the 20 longest (4% of the network) and 20 shortest timescale neurons from the network (Fig. 7a,b). For small N and both curricula, ablating individual neurons has only minimal effect (less than 1%) on accuracy. However, for larger N , we observe a considerable difference in the importance of neurons. Single-head networks rely strongly on long-timescale neurons (Fig. 7a), such that ablating them reduces the performance much more than for short-timescale neurons. In contrast, multi-head networks exhibit greater robustness against ablation, and their accuracy is more affected when short-timescale neurons (i.e., neurons with $\tau = 1$) are ablated (Fig. 7b). Note that in single-head networks, the average of the 20 longest single-neuron timescales is 2.7 times longer than in multi-head networks.

Perturbation: We perturb W^R and τ with strength ε as (Wu et al., 2020)

$$\tilde{W}^R = W^R + \varepsilon \frac{\xi_W}{\|\xi_W\|} \|W^R\|, \quad \tilde{\tau} = \tau + \varepsilon \left| \frac{\xi_\tau}{\|\xi_\tau\|} \right| |\tau|, \quad \xi_W \sim \mathcal{N}(0, \mathbb{I}^{n \times n}), \quad \xi_\tau \sim \mathcal{N}(0, \mathbb{I}^n). \quad (4)$$

$\|\cdot\|$ represents Frobenius norm and $|\cdot|$ the absolute value. τ is perturbed positively to avoid $\tau < 1$. Multi-head networks are more robust to perturbations (Fig. 7c,d). The robustness to changes in W^R is noteworthy since these networks rely on connectivity to mediate long timescales.

Retraining: We evaluate the performance of networks that solve $N = 16$, when retrained without curriculum (without training on intermediate N s) for an arbitrary higher N , after 20 epochs. Multi-




Figure 7: Multi-head networks are more robust to ablation and perturbation, and better retrainable. (a, b) Ablating long-timescale neurons largely decreases the performance of single-head networks (a), while multi-head networks (b) are more affected by the ablation of short-timescale neurons ($N = 30$). (c, d) Multi-head networks are more robust against perturbations of recurrent connectivity W^R and τ than single-head networks (note the log-scale x-axis, $N = 30$). (e) Multi-head networks retrain faster: They achieve higher relative accuracy when retrained for 20 epochs without a curriculum for a higher N . Networks trained to solve $N = 16$ are retrained for 20 epochs to solve higher N without a curriculum to compare re-trainability. Bars - mean; dots and lines - 4 networks for each curriculum; error bars and shades - \pm STD.

head networks show a superior retraining ability compared to single-head networks for at least 10 N beyond $N = 16$ (Fig. 7e). This finding suggests that multi-head networks are better at learning the underlying task and adjust faster to a new, larger N even when skipping intermediate N s.

5 RELATION TO NEUROSCIENCE

Continuous-time setting: So far, we described the tasks and network dynamics in discrete time with time-step $\Delta t = 1$. However, to make the connection to more realistic settings, such as neuroscience, we need to describe the dynamics in continuous time. For this purpose, we consider that each input digit is presented to the network for a certain time T . In neuroscience experiments (e.g., DMS task), T is often set to 250-500 ms (Meyer et al., 2011; Qi et al., 2011; Kim & Sejnowski, 2021).

First, we show that for a fixed T , the discretization time step Δt does not affect the networks’ performance and dynamics. We train the networks with each curriculum using a variety of Δt (Appendix B). In these networks, the performance on the test data is comparable for different Δt , even for values much smaller than what was included during the training, mimicking continuous-time dynamics (Fig. S5 c-d). Moreover, similar to our previous results, we find that single-head networks increase their τ with N , while multi-head networks try to reach $\tau \rightarrow T$ (Fig. S5 a,b).

Next, we test how changes in input presentation time T affect the dynamics. We set $T = k\Delta t$, where $\Delta t = 1$. We find that for all k , single-head networks increase τ with N , whereas multi-head networks’ τ tends to k , thus matching the timescales of the input changes (Appendix C).

Timescales and learning in the biological neural networks: Our findings suggest that networks that are required to solve tasks with larger memory requirements should develop longer timescales. This largely agrees with findings in the brain: higher cortical areas that are involved in cognitive processes with larger memory requirements (e.g., working memory, evidence accumulation) have longer timescales than sensory areas (Murray et al., 2014). Moreover, we show that developing longer network timescales via changes in network connectivity is a superior solution (in terms of performance and stability) than using longer single-neuron timescales. This is consistent with findings that neural timescales in primate visual cortex adapt to task demands via recurrent network interactions rather than biophysical time constants (Zeraati et al., 2023). Moreover, this result aligns with the learning strategy of biological neural networks, which primarily relies on changes in synaptic strengths rather than modifying the biophysical time constants of individual neurons. While such changes do happen in biology (via protein turnover (Sun & Schuman, 2022), calcium currents (Tiganj et al., 2015), and other mechanisms), synaptic strength modification is overwhelmingly the mechanism by which biological networks learn.

6 RELATED WORK

Previous works independently investigated the role of neuronal and network-mediated timescales in solving memory tasks and proposed inconsistent solutions. Studies focusing on neuronal aspect suggested heterogeneous and adaptable neuronal properties (e.g., membrane time constant) as an optimal mechanism (Perez-Nieves et al., 2021; Mahto et al., 2021; Smith et al., 2023a; Quax et al., 2020). At the same time, other studies presented that network-mediated mechanisms like balanced dynamics (Lim & Goldman, 2013), strong inhibition (Kim & Sejnowski, 2021) or homeostatic plasticity (Cramer et al., 2020; 2023) can create timescales required for memory tasks. For a single neuron modeled with multiple memory units, long timescales were shown to be instrumental in solving memory tasks (Spieler et al., 2023). Here, we explicitly compare these mechanisms and show that while both can be useful for learning long-memory tasks, applying network-mediated mechanisms leads to faster training and more robust solutions.

We find that the difference between mechanisms is revealed mainly in the context of distinct learning objectives defined by curricula. This is an important distinction with previous work, since the role of timescales has been often studied when RNNs solve a single task (e.g., single-head DMS), without considering learning dynamics or the potential for catastrophic forgetting. We relate the mechanisms of task-dependent timescale with the learning dynamics of RNNs across curricula. The use of curricula in our study is inspired by previous work suggesting curriculum learning as a fitness landscape-smoothing mechanism that can enable the gradual learning of highly complex tasks (Elman, 1993; Bengio et al., 2009; Krueger & Dayan, 2009) and be used to uncover distinct learning mechanisms (Kepple et al., 2022; Dekker et al., 2022). Here, we extend these findings by demonstrating how different curricula can push networks towards adopting different strategies to develop slow collective dynamics required for solving long-memory tasks.

7 DISCUSSION

We find that to solve long-memory tasks, RNNs develop high-dimensional activity with slow timescales via two distinct combinations of connectivity and single-neuron timescales. While single-head networks crucially rely on the long single-neuron timescales to perform the task, multi-head networks prefer a constant single-neuron timescale and solve the task relying only on the long timescales emerging from recurrent interactions. We show that developing long timescales via recurrent interactions instead of single-neuron properties is optimal for learning memory tasks and leads to more stable and robust solutions, which can be a beneficial strategy for brain computations.

Our findings suggest that training networks on *sets* of related memory tasks instead of a single task improves performance and robustness. By progressively shaping the loss function with a curriculum to include performance evaluations on sub-tasks that are known to correlate with the desired task, we can smooth the loss landscape of our network to allow training for difficult tasks that were previously unsolvable. In this way, choosing an appropriate curriculum can act as a powerful regularization.

Limitations: Our study considers two relatively simple tasks with explicitly controllable memory requirements. In follow-up studies, it would be important to test our observations in more sophisticated tasks and investigate whether our results apply to other architectures and optimizers. Additionally, our approach is suitable only for a set of tasks with controllably increasing memory requirements, where the different versions of the same task can be simultaneously performed on the same data (multi-head training). This is a relatively strong constraint, and future research expanding our findings could focus on generalizing the multi-head curriculum for training more realistic tasks. Time series reconstruction is a potential task that can be used to uncover generative dynamical systems from data (Durstewitz et al., 2023). We proposed a potential experiment in Appendix D.

Our current model is a crude approximation of biological neural networks, and more plausible architectures (spiking models, distinct neuron types) could be studied. Finally, biological neural networks can produce long timescales via various other mechanisms we did not consider here (short-term plasticity (Hu et al., 2021), adaptation (Salaj et al., 2021; Beiran & Ostojic, 2019), synaptic delays, etc). A follow-up study could investigate whether our findings extend to more plausible networks incorporating such additional mechanisms.

ACKNOWLEDGMENTS

This work was supported by a Sofja Kovalevskaja Award from the Alexander von Humboldt Foundation, endowed by the Federal Ministry of Education and Research (SK, RZ, EG, AL), the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC number 2064/1 - Project number 390727645 (RZ, EG), and Else Kröner Medical Scientist Kolleg “ClinbrAI: Artificial Intelligence for Clinical Brain Research” (TJS). We acknowledge the support from the BMBF through the Tübingen AI Center (FKZ: 01IS18039B), International Max Planck Research School for the Mechanisms of Mental Function and Dysfunction (IMPRS-MMFD), and International Max Planck Research School for Intelligent Systems (IMPRS-IS). We thank Victor Buendía for valuable discussions.

REFERENCES

- H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974. doi: 10.1109/TAC.1974.1100705. URL <https://ieeexplore.ieee.org/document/1100705>.
- Brice Bathellier, Derek L. Buhl, Riccardo Accolla, and Alan Carleton. Dynamic ensemble odor coding in the mammalian olfactory bulb: Sensory information at different timescales. *Neuron*, 57(4):586–598, 2008. ISSN 0896-6273. doi: <https://doi.org/10.1016/j.neuron.2008.02.011>. URL <https://www.sciencedirect.com/science/article/pii/S0896627308001347>.
- Manuel Beiran and Srdjan Ostojic. Contrasting the effects of adaptation and synaptic filtering on the timescales of dynamics in recurrent networks. *PLOS Computational Biology*, 15(3):e1006893, March 2019. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1006893. URL <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1006893>.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML ’09*, pp. 41–48, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605585161. doi: 10.1145/1553374.1553380. URL <https://doi.org/10.1145/1553374.1553380>.
- Nicolas Boulanger-Lewandowski, Yoshua Bengio, and Pascal Vincent. Modeling temporal dependencies in high-dimensional sequences: Application to polyphonic music generation and transcription. *arXiv preprint arXiv:1206.6392*, 2012. URL <https://arxiv.org/abs/1206.6392>.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. *arXiv preprint arXiv:1511.06349*, 2015. URL <https://arxiv.org/abs/1511.06349>.
- Sean E. Cavanagh, Laurence T. Hunt, and Steven W. Kennerley. A Diversity of Intrinsic Timescales Underlie Neural Computations. *Frontiers in Neural Circuits*, 14, 2020. ISSN 1662-5110. doi: 10.3389/fncir.2020.615626. URL https://www.frontiersin.org/articles/10.3389/fncir.2020.615626/full?field=&id=615626&journalName=Frontiers_in_Neural_Circuits. Publisher: Frontiers.
- Rishidev Chaudhuri, Alberto Bernacchia, and Xiao-Jing Wang. A diversity of localized timescales in network activity. *eLife*, 3, January 2014. ISSN 2050-084X. doi: 10.7554/eLife.01239. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3895880/>.
- Rishidev Chaudhuri, Kenneth Knoblauch, Marie-Alice Gariel, Henry Kennedy, and Xiao-Jing Wang. A Large-Scale Circuit Mechanism for Hierarchical Dynamical Processing in the Primate Cortex. *Neuron*, 88(2):419–431, October 2015. ISSN 0896-6273. doi: 10.1016/j.neuron.2015.09.008. URL <http://www.sciencedirect.com/science/article/pii/S0896627315007655>.

- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014. URL <https://arxiv.org/abs/1412.3555>.
- Benjamin Cramer, David Stöckel, Markus Kreft, Michael Wibrall, Johannes Schemmel, Karlheinz Meier, and Viola Priesemann. Control of criticality and computation in spiking neuromorphic networks with plasticity. *Nature communications*, 11(1):2853, 2020. URL <https://www.nature.com/articles/s41467-020-16548-3>.
- Benjamin Cramer, Markus Kreft, Sebastian Billaudelle, Vitali Karasenko, Aron Leibfried, Eric Müller, Philipp Spilger, Johannes Weis, Johannes Schemmel, Miguel A Muñoz, et al. Auto-correlations from emergent bistability in homeostatic spiking neural networks on neuromorphic hardware. *Physical Review Research*, 5(3):033035, 2023. URL <https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.5.033035>.
- Ronald B Dekker, Fabian Otto, and Christopher Summerfield. Curriculum learning for human compositional generalization. *Proceedings of the National Academy of Sciences*, 119(41):e2205582119, 2022. URL <https://www.pnas.org/doi/10.1073/pnas.2205582119>.
- Renato Duarte, Alexander Seeholzer, Karl Zilles, and Abigail Morrison. Synaptic patterning and the timescales of cortical dynamics. *Current Opinion in Neurobiology*, 43:156–165, April 2017. ISSN 0959-4388. doi: 10.1016/j.conb.2017.02.007. URL <https://www.sciencedirect.com/science/article/pii/S0959438817300545>.
- Daniel Durstewitz, Georgia Koppe, and Max Ingo Thurm. Reconstructing computational system dynamics from neural data with recurrent neural networks. *Nature Reviews Neuroscience*, pp. 1–18, 2023. URL <https://www.nature.com/articles/s41583-023-00740-7>.
- Jeffrey L Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990. URL https://onlinelibrary.wiley.com/doi/abs/10.1207/s15516709cog1402_1.
- Jeffrey L Elman. Learning and development in neural networks: The importance of starting small. *Cognition*, 48(1):71–99, 1993. URL <https://www.sciencedirect.com/science/article/pii/0010027793900584>.
- Wei Fang, Zhaofei Yu, Yanqi Chen, Timothee Masquelier, Tiejun Huang, and Yonghong Tian. Incorporating Learnable Membrane Time Constant to Enhance Learning of Spiking Neural Networks, August 2021. URL <http://arxiv.org/abs/2007.05785>. arXiv:2007.05785 [cs].
- Richard Gao, Ruud L van den Brink, Thomas Pfeffer, and Bradley Voytek. Neuronal timescales are functionally dynamic and shaped by cortical microarchitecture. *eLife*, 9:e61277, November 2020. ISSN 2050-084X. doi: 10.7554/eLife.61277. URL <https://doi.org/10.7554/eLife.61277>. Publisher: eLife Sciences Publications, Ltd.
- Julijana Gjorgjieva, Guillaume Drion, and Eve Marder. Computational implications of biophysical diversity and multiple timescales in neurons and synapses for circuit performance. *Current Opinion in Neurobiology*, 37:44–52, April 2016. ISSN 0959-4388. doi: 10.1016/j.conb.2015.12.008. URL <https://www.sciencedirect.com/science/article/pii/S0959438815001865>.
- Alex Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013. URL <https://arxiv.org/abs/1308.0850>.
- Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, pp. 6645–6649. Ieee, 2013. URL <https://arxiv.org/abs/1303.5778>.
- David Ha and Douglas Eck. A neural representation of sketch drawings. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=Hy6GHpkCW>.

- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997. URL <https://direct.mit.edu/neco/article-abstract/9/8/1735/6109/Long-Short-Term-Memory?redirectedFrom=fulltext>.
- Brian Hu, Marina E Garrett, Peter A Groblewski, Douglas R Ollerenshaw, Jiaqi Shang, Kate Roll, Sahar Manavi, Christof Koch, Shawn R Olsen, and Stefan Mihalas. Adaptation supports short-term memory in a visual change detection task. *PLoS computational biology*, 17(9):e1009246, 2021. URL <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1009246>.
- Shailee Jain, Vy Vo, Shivangi Mahto, Amanda LeBel, Javier S Turek, and Alexander Huth. Interpretable multi-timescale models for predicting fMRI responses to continuous natural speech. *Adv. Neural Inf. Process. Syst.*, 33:13738–13749, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/9e9a30b74c49d07d8150c8c83b1ccf07-Abstract.html>.
- John Jonides, Richard L. Lewis, Derek Evan Nee, Cindy A. Lustig, Marc G. Berman, and Katherine Sledge Moore. The Mind and Brain of Short-Term Memory. *Annual Review of Psychology*, 59(1):193–224, 2008. doi: 10.1146/annurev.psych.59.103006.093615. URL <https://doi.org/10.1146/annurev.psych.59.103006.093615>.
- Daniel R. Kepple, Rainer Engelken, and Kanaka Rajan. Curriculum learning as a tool to uncover learning principles in the brain. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=TpJMvo0_pu-.
- Robert Kim and Terrence J. Sejnowski. Strong inhibitory signaling underlies stable temporal dynamics and working memory in spiking neural networks. *Nature Neuroscience*, 24(1):129–139, January 2021. ISSN 1546-1726. doi: 10.1038/s41593-020-00753-w. URL <https://www.nature.com/articles/s41593-020-00753-w>.
- Kai A. Krueger and Peter Dayan. Flexible shaping: How learning in small steps helps. *Cognition*, 110(3):380–394, 2009. ISSN 0010-0277. doi: <https://doi.org/10.1016/j.cognition.2008.11.014>. URL <https://www.sciencedirect.com/science/article/pii/S0010027708002850>.
- Sukbin Lim and Mark Goldman. Balanced cortical microcircuitry for maintaining information in working memory. *Nature neuroscience*, 16, 2013. doi: 10.1038/nn.3492. URL <https://www.nature.com/articles/nn.3492>.
- Zachary C Lipton, John Berkowitz, and Charles Elkan. A critical review of recurrent neural networks for sequence learning. *arXiv preprint arXiv:1506.00019*, 2015. URL <https://arxiv.org/abs/1506.00019>.
- Ashok Litwin-Kumar and Brent Doiron. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience*, 15(11):1498–1505, November 2012. ISSN 1546-1726. doi: 10.1038/nn.3220. URL <https://www.nature.com/articles/nn.3220>. Number: 11 Publisher: Nature Publishing Group.
- Shivangi Mahto, Vy Ai Vo, Javier S. Turek, and Alexander Huth. Multi-timescale representation learning in {Istm} language models. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=9ITXiTrAoT>.
- Travis Meyer, Xue-Lian Qi, Terrence R Stanford, and Christos Constantinidis. Stimulus selectivity in dorsal and ventral prefrontal cortex after training in working memory tasks. *Journal of neuroscience*, 31(17):6266–6276, 2011. URL <https://www.jneurosci.org/content/31/17/6266>.
- John D. Murray, Alberto Bernacchia, David J. Freedman, Ranulfo Romo, Jonathan D. Wallis, Xinying Cai, Camillo Padoa-Schioppa, Tatiana Pasternak, Hyojung Seo, Daeyeol Lee, and Xiao-Jing Wang. A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience*, 17(12):1661–1663, December 2014. ISSN 1546-1726. doi: 10.1038/nn.3862. URL <https://www.nature.com/articles/nn.3862>.

- Srdjan Ostojic. Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons. *Nature Neuroscience*, 17(4):594–600, April 2014. ISSN 1546-1726. doi: 10.1038/nn.3658. URL <https://www.nature.com/articles/nn.3658>.
- Stefano Panzeri, Nicolas Brunel, Nikos K. Logothetis, and Christoph Kayser. Sensory neural codes using multiplexed temporal scales. *Trends in Neurosciences*, 33(3):111–120, March 2010. ISSN 0166-2236. doi: 10.1016/j.tins.2009.12.001. URL <http://www.sciencedirect.com/science/article/pii/S0166223609002008>.
- Pranay Pasula. Real world time series benchmark datasets with distribution shifts: Global crude oil price and volatility. *arXiv preprint arXiv:2308.10846*, 2023. URL <https://arxiv.org/abs/2308.10846>.
- Nicolas Perez-Nieves, Vincent C. H. Leung, Pier Luigi Dragotti, and Dan F. M. Goodman. Neural heterogeneity promotes robust learning. *Nature Communications*, 12(1):5791, October 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-26022-3. URL <https://www.nature.com/articles/s41467-021-26022-3>.
- Xue-Lian Qi, Travis Meyer, Terrence R Stanford, and Christos Constantinidis. Changes in prefrontal neuronal activity after learning to perform a spatial working memory task. *Cerebral cortex*, 21(12):2722–2732, 2011. URL <https://academic.oup.com/cercor/article/21/12/2722/295413>.
- Silvan C. Quax, Michele D’Asaro, and Marcel A. J. van Gerven. Adaptive time scales in recurrent neural networks. *Scientific Reports*, 10(1):11360, July 2020. ISSN 2045-2322. doi: 10.1038/s41598-020-68169-x. URL <https://www.nature.com/articles/s41598-020-68169-x>. Number: 1 Publisher: Nature Publishing Group.
- Shervin Safavi, Matthew Chalk, Nikos Logothetis, and Anna Levina. Signatures of criticality in efficient coding networks. *bioRxiv*, pp. 2023–02, 2023. URL <https://www.biorxiv.org/content/10.1101/2023.02.14.528465v1>.
- Darjan Salaj, Anand Subramoney, Ceca Krausnikovic, Guillaume Bellec, Robert Legenstein, and Wolfgang Maass. Spike frequency adaptation supports network computations on temporally dispersed information. *Elife*, 10:e65459, 2021. URL <https://elifesciences.org/articles/65459>.
- Yan-Liang Shi, Roxana Zeraati, Anna Levina, and Tatiana A Engel. Spatial and temporal correlations in neural networks with structured connectivity. *Physical Review Research*, 5(1):013005, 2023. URL <https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.5.013005>.
- Jimmy T.H. Smith, Andrew Warrington, and Scott Linderman. Simplified state space layers for sequence modeling. In *The Eleventh International Conference on Learning Representations*, 2023a. URL <https://openreview.net/forum?id=Ai8Hw3AXqks>.
- Jimmy T.H. Smith, Andrew Warrington, and Scott Linderman. Simplified state space layers for sequence modeling. In *The Eleventh International Conference on Learning Representations*, 2023b. URL <https://openreview.net/forum?id=Ai8Hw3AXqks>.
- Aaron Spieler, Nasim Rahaman, Georg Martius, Bernhard Schölkopf, and Anna Levina. The elm neuron: an efficient and expressive cortical neuron model can solve long-horizon tasks. *arXiv preprint arXiv:2306.16922*, 2023. URL <https://arxiv.org/abs/2306.16922>.
- Chao Sun and Erin M. Schuman. Logistics of neuronal protein turnover: Numbers and mechanisms. *Molecular and Cellular Neuroscience*, 123:103793, 2022. ISSN 1044-7431. doi: <https://doi.org/10.1016/j.mcn.2022.103793>. URL <https://www.sciencedirect.com/science/article/pii/S1044743122000999>.
- Corentin Tallec and Yann Ollivier. Can recurrent neural networks warp time? *arXiv preprint arXiv:1804.11188*, 2018. URL <https://arxiv.org/abs/1804.11188>.

- Zoran Tiganj, Michael E. Hasselmo, and Marc W. Howard. A simple biophysically plausible model for long time constants in single neurons. *Hippocampus*, 25(1):27–37, 2015. doi: <https://doi.org/10.1002/hipo.22347>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/hipo.22347>.
- José F Torres, Dalil Hadjout, Abderrazak Sebaa, Francisco Martínez-Álvarez, and Alicia Troncoso. Deep learning for time series forecasting: a survey. *Big Data*, 9(1):3–21, 2021. URL <https://www.liebertpub.com/doi/10.1089/big.2020.0159>.
- Alexander van Meegen and Sacha J. van Albada. Microscopic theory of intrinsic timescales in spiking neural networks. *Physical Review Research*, 3(4):043077, October 2021. doi: [10.1103/PhysRevResearch.3.043077](https://doi.org/10.1103/PhysRevResearch.3.043077). URL <https://link.aps.org/doi/10.1103/PhysRevResearch.3.043077>.
- Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. *Advances in Neural Information Processing Systems*, 33:2958–2969, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/1ef91c212e30e14bf125e9374262401f-Abstract.html>.
- Bojian Yin, Federico Corradi, and Sander M. Bohté. Effective and Efficient Computation with Multiple-timescale Spiking Recurrent Neural Networks. In *International Conference on Neuromorphic Systems 2020, ICONS 2020*, pp. 1–8, New York, NY, USA, July 2020. Association for Computing Machinery. ISBN 978-1-4503-8851-1. doi: [10.1145/3407197.3407225](https://doi.org/10.1145/3407197.3407225). URL <https://dl.acm.org/doi/10.1145/3407197.3407225>.
- Yong Yu, Xiaosheng Si, Changhua Hu, and Jianxun Zhang. A review of recurrent neural networks: Lstm cells and network architectures. *Neural computation*, 31(7):1235–1270, 2019. URL <https://ieeexplore.ieee.org/document/8737887>.
- Roxana Zeraati, Tatiana A. Engel, and Anna Levina. A flexible Bayesian framework for unbiased estimation of timescales. *Nature Computational Science*, 2(3):193–204, March 2022. ISSN 2662-8457. doi: [10.1038/s43588-022-00214-3](https://doi.org/10.1038/s43588-022-00214-3). URL <https://www.nature.com/articles/s43588-022-00214-3>.
- Roxana Zeraati, Yan-Liang Shi, Nicholas A Steinmetz, Marc A Gieselmann, Alexander Thiele, Tirin Moore, Anna Levina, and Tatiana A Engel. Intrinsic timescales in the visual cortex change with selective attention and reflect spatial connectivity. *Nature Communications*, 14(1):1858, 2023. URL <https://www.nature.com/articles/s41467-023-37613-7>.

APPENDIX

A DIFFERENT TYPES AND LOCATIONS OF NONLINEARITY

In order to verify that our results are robust with respect to the type of nonlinearity used in the network, we train RNNs using two of the most commonly used nonlinearities: ReLU and Tanh. We find that in both cases, the training performance is similar to leaky ReLU, and the development of single-neuron and network-mediated timescales follow the same trajectory as N increases (Fig. S7).

In some implementations of leaky-RNN, the neural self-interaction is linear and located outside of the nonlinearity (cf. equ.1)

$$r_i(t) = \left(1 - \frac{\Delta t}{\tau_i}\right) \cdot r_i(t - \Delta t) + \left[\frac{\Delta t}{\tau_i} \cdot \left(\sum_{i \neq j} W_{ij}^R \cdot r_j(t - \Delta t) + W_i^I \cdot S(t) + b^R + b^I\right)\right]_{\alpha}. \quad (5)$$

with the explicit time discretization Δt . The input is presented for time duration $T = k\Delta t$ with input-update time steps k . In the main text, we chose $k = 1$ and $\Delta t = 1$. We discuss $k > 1$ in Appendix C and $\Delta t < 1$ in Appendix B.

We verify that training RNNs with this implementation gives similar training dynamics and trajectories of τ and τ_{net} with increasing N (Fig. S8), for both curricula. Furthermore, we find that, for large N , ablating neurons with long τ in single-head networks and neurons with short τ in multi-head networks reduces the performance significantly, compatible with the findings in the main text (Fig. S9, cf. Fig. 7). We also verify that the performance of the model depends on the initialization of τ and its trainability in the same way regardless of the location of the nonlinearity (Fig. S11, cf. Fig. 4).

B CHANGING TIME DISCRETIZATION

In computational neuroscience, the single neuron dynamics are typically captured by the differential equations that need to be discretized for running numerical simulations and training networks. However, the discretization can be important for stability and internal representation of the model and the task. In the main text, we used Eq. 5 with $\Delta t = T = 1$. For simplicity of notation, we take in the rest of this section $T = 1$. We train networks with different values of Δt (a different Δt for each training batch), so they can perform the same task independent time discretization. We take $\Delta t = 1/n$ with $n \in \mathbb{N}$ and train the network while keeping the duration of each stimulus presentation in units of time constant (which means that with larger n , it would be presented for more time steps). The flexible framework for time discretization allows us to train with multiple Δt simultaneously. Then, we test whether the network can solve the same task but with Δt not included in their training set.

We find that in networks trained with multiple Δt , the single-neuron timescales τ follow a similar trajectory as the results in the main text, independent of Δt (Fig. S5a,b compared to Fig. 4b). Multi-head networks adjust their τ to converge to $n\Delta t = 1$, while single-head networks increase their individual neuron timescale. Moreover, the networks can generalize (without retraining) the task to smaller Δt than what was included in their training set, in single- and multi-head networks. Interestingly, the performance decreases slowly when Δt becomes smaller than the training set, but abruptly when it becomes larger (Fig. S5c,d). The performance is best when training with multiple Δt , but qualitatively, the result is similar for a single, small enough Δt (Fig. S5e).

C CHANGING THE DURATION OF THE INPUT PRESENTATION

In our tasks, the input contains two timescales. First is the duration of presentation of each input digit $T = k \cdot T_{\text{min}}$, with T_{min} a minimal considered duration of stimulus presentation measured in milliseconds. Second is the timescale of the task’s memory N . In the main text, we consider the situation of $k = 1$, but in general, k acts as a time-rescaling parameter and defines one unit of time for the task performance. Here, we train the RNNs with different values of $k \in \{2, 3, 5, 10\}$ and check the trajectories of changing τ with N depending on k . We find that similar to the case

with $k = 1$, single-head networks trained with $k > 1$ increase their τ with N , while multi-head networks try to keep τ close to k (Fig. S6). Moreover, tasks with $k > 1$ are generally more difficult to solve since the input needs to be tracked over $N \cdot k$ time steps. Hence, as k grows, RNNs would reach smaller N within the same number of training epochs. The changes in values of τ after rescaling with k might be due to nonlinear interactions in the network arising from the combination of different N and k .

D PROPOSED ADDITIONAL TASK: TEMPORAL PATTERN GENERATION

For future research, the task variety can be extended to include the temporal pattern generation, which is a continuous-time task that is often used to evaluate RNNs (Durstewitz et al., 2023). The classic variation of the task involved an RNN receiving either random noise or no input and having to produce a target time series as output (usually a sum of sine waves with different frequencies).

A variation of the task we could consider for testing our model is the following:

Single-head: On the first step of the curriculum, we train the network to produce a single sine wave with frequency f_1 , setting the target sequence to be $y_{N=1} = \sin(2\pi \cdot f_1 \cdot t)$.

Then, for each curriculum step, we complexify the target sequence by setting the new target as:

$$y_{N=m} = \sum_{i=1}^m \sin(2\pi \cdot f_i \cdot t), \quad (6)$$

for $f_1 > f_2 > \dots > f_m$. In this way, as the newly added frequencies decrease, a need arises for the network to develop longer timescales.

Multi-head: Unlike the single-head network where the RNN needs to produce only one target time series $y_{N=m}$ at the m -th step of the curriculum, in the multi-head curriculum, the network produces m output time series $Y = \{y_{N=1}, \dots, y_{N=m}\}$.

E EFFECTS OF TRAINING WITHOUT A CURRICULUM ON THE N -DMS TASK

We investigate the negative effects of not using a curriculum during training for the N -DMS task to extend our results from Fig. 3a. We show in Fig. S13 that similar to the N -Parity task, networks rapidly lose the ability to solve the N -DMS task as N increases when training without a curriculum. Interestingly, the two tasks differ in the way they fail to be solved despite using identical optimizers. In all of our results, the N -DMS task tends to be easier to solve for larger N . However, despite the relative success these networks have with the N -DMS task, their drop-off in training these networks is much steeper when comparing the curves from Fig. 3a and Fig. S13. In Fig. S13, tasks $N < 15$ get solved in only 1 or 2 epochs, however between $15 < N < 20$ the networks rapidly slow down in their ability to train until completely failing for $N > 20$ even when given longer training time. We can infer from these results that different tasks have varying degrees to which they benefit from a particular curriculum.

F INTERMEDIATE CURRICULA: MULTI-HEAD WITH A SLIDING WINDOW

The two curricula discussed in the main text (single-head and multi-head) represent two extreme cases. In the single-head curriculum, at each step of the curriculum, RNNs are trained to solve a new N without requiring to remember the solution to the previous N s. On the other hand, in the multi-head curriculum, RNNs need to remember the solution to all the previous N s in addition to the new N . Here, we test the behavior of curricula that lie in between the two extreme cases.

The intermediate curricula involve the simultaneous training of multiple heads, similar to the multi-head curriculum, but instead of adding new heads at each curriculum step, we train a fixed number of heads and only shift the N s, which they are trained for according to a sliding window. We consider the number of heads to be 10, and start the training for $N \in [2, \dots, 11]$. In the next steps of the curriculum, we use the already trained network to initialize another network which we train for $N + w$ (e.g., $N \in [2 + w, \dots, 11 + w]$), where $w \in \{1, 3, 5\}$ indicates the size of the sliding

window. For each w , we train 4 different networks (i.e., 4 different initialization). For the following analyses, we trained the networks on the N -parity task.

We find that networks trained with the multi-head-sliding curriculum generally demonstrate an in-between behavior compared with the extreme curricula, but the results also depend on the size of the sliding window. Within 1000 training epochs, the maximal N these networks can solve (with $> 98\%$ accuracy) is in between the maximal N of single- and multi-head curricula, depending on the sliding window. Networks with a larger sliding window can solve a higher maximal N , indicating that a large sliding window not only does not slow down the training but also provides a more efficient curriculum to learn higher N s (Fig. S12a). Moreover, in multi-head-sliding networks, single-neuron (τ) and network-mediated (τ_{net}) timescales have values in between single-head and multi-head curricula (Fig. S12b). However, both τ and τ_{net} grow with N similar to single-head networks, with the pace of growth reducing for larger sliding windows.

Similar to the main text (Fig. 7c,d,e), we perform the perturbation and retraining analysis on multi-head-sliding networks trained with $w = 5$. The relative accuracy after perturbation of recurrent weights W^R and timescales τ for these networks lies between the two extremes (Fig. S14a, b). However, the retraining analysis suggests that multi-head-sliding networks can be retrained better for higher new N s (Fig. S14c,d). If the network is originally retrained for a small N (e.g., $N = 16$), the retraining relative accuracy is similar between multi-head and sliding networks but is larger than single-head networks. For networks trained for larger N s (e.g., $N = 31$), sliding networks exhibit a superior retraining ability compared to the other two curricula. These results suggest that the curriculum with the sliding window helps multi-head networks to better adjust to new N s.

G SINGLE- AND MULTI-HEAD CURRICULA FOR TRAINING GRU AND LSTM

The results presented in the main text were generated using a modified version of a vanilla RNN (leaky-RNN) with an explicit definition of the timescale parameter τ . To test whether the difficulties in training for long memory tasks without curriculum would carry over to recurrent networks that were specifically designed for long memory tasks, we train two other architectures, an LSTM (long short-term memory) and a GRU (gated recurrent unit) on the N -parity task for increasing N , with and without a curriculum. Both the GRU and LSTM have similar network sizes to the RNN with 500 neurons, though they differ in their activation functions (the RNN used a single leakyReLU whereas the GRU/LSTMs have both sigmoids and tanhs for different gates). Furthermore, in contrast with the RNNs, an Adam optimizer is used with learning rate $lr = 10^{-3}$ and the input signals to the models take values $\in \{-1, 1\}$ (to have a zero-mean input signal).

We find that for both architectures, training the networks without a curriculum is extremely slow for large N and relatively unstable for small N and probably requires strict hyper-parameter tuning (Fig. S15a). Without additional hyper-parameter tuning, introducing the multi-head curriculum speeds up the training significantly, and both architectures can easily learn the N -parity task with large N similar to the leaky-RNN (Fig. S15b). Moreover, similar to RNNs, the multi-head curriculum has a higher training speed than the single-head curriculum (Fig. S16). Our results indicate that GRUs and LSTMs are subject to similar training dynamics as RNNs used in the main text and the multi-head curriculum is an optimal curriculum regardless of the RNN architecture. The advantage of using the leaky-RNN architecture is that its parameters are easier to interpret, and it allows us to study better the mechanisms underlying each curriculum by explicitly studying the role of timescales.

H BACKWARD AND FORWARD RETRAINING OF NETWORKS

To understand how trained models develop their ability to create longer timescales throughout the curriculum as well as their backward compatibility and robustness to catastrophic forgetting, we measure the retrainability of models trained on a task with memory N on a different task with memory N^* . We freeze all parameters of a trained network except the final readout layer weights which are retrained on an N^* task. Specifically, we load models trained for $N \in [2, \dots, 19]$ and retrain them on a new $N^* \in [2, \dots, N+2]$, independently for each N^* , for a maximum of 10 epochs or until its accuracy was above 98%. Note that we retrain both single- and multi-head networks as single-head.

We find that the multi-head networks exhibit near-perfect backward compatibility as well as better forward compatibility than the single-head models (Fig. S17), while single-head networks suffer from catastrophic forgetting. For the multi-head networks, the backward compatibility is enforced through the loss function (as is the case in the multi-head curriculum) hence, the necessary representations for $N^* < N$ persist. However, the multi-head curriculum also has positive implications for forward compatibility, which is evident in the off-diagonal entries of the accuracy where $N^* > N$ (to the right of the dotted line) when compared to the single-head values.

I EMERGENCE OF CURRICULUM DURING MULTI-HEAD TRAINING

In the multi-head curriculum, the difficulty of the task increases gradually; a new head with a larger N is added at each step of the curriculum. In the main text, we discussed that networks trained with such a curriculum generally train well up to large N s. Here we ask whether this optimal curriculum can emerge by itself if we train a network with multiple heads, but without any predefined curricula. For this analysis, we train RNNs with 19 ($N \in [2, \dots, 20]$) and 39 heads ($N \in [2, \dots, 40]$) to solve all the available N s simultaneously.

We find that despite the absence of an explicit curriculum, these networks learn the task by generating an internal multi-head curriculum. While all the heads contribute equally to the loss, heads with a smaller N reach the higher accuracy faster (Fig. S18a). However, the speed of training strongly depends on the total number of heads in each network. For the same N , the network with 19 heads reaches the 98% accuracy faster than the network with 39 heads (Fig. S18b), but both networks have a slower training speed when compared to the multi-head curriculum. These results suggest that the multi-head curriculum is an optimal curriculum that can arise naturally during multi-head training and can increase the training speed when applied explicitly.

J ROLE OF STRONG INHIBITORY CONNECTIVITY IN SINGLE-HEAD NETWORKS

The main difference between single- and multi-head networks in terms of connectivity is the stronger inhibitory (negative) connectivity for large N in the single-head networks compared with the relatively balanced connectivity in multi-head networks (Fig. 6a). We hypothesized that larger inhibition in single-head networks is required to keep the dynamics stable in the presence of slow single-neuron timescales τ . To test this hypothesis, we perturb only the inhibitory connections in networks trained with both curricula as:

$$W_{ij}^R = W_{ij} + c \cdot W_{ij}, \quad \forall W_{ij}^R < 0, \quad (7)$$

for a given amount of $c \in [-0.1, 0.1]$. We observe that by reducing the amount of inhibition in single-head networks, the network activity explodes even before reaching the balanced point, i.e., the point when the average incoming weight of neurons becomes 0 (Fig. S19a). On the contrary, multi-head networks are significantly more robust to such perturbations and their activity remains within a reasonable range for a broad range of inhibitory scaling (Fig. S19b).

This difference is most likely attributed to the difference in single-neuron timescales τ between single- and multi-head networks. The single-head networks have a larger average τ compared to the multi-head networks whose average $\tau \approx 1$ for large N . Longer τ leads to neurons with self-sustaining activity, and thus, a stronger inhibition might be required to prevent the runaway activation. Such a relationship can be observed when comparing the average τ and inhibitory strength across networks: for single-head networks as τ grows, the average weight becomes more negative (inhibitory)(Fig. S19c), but such correlation does not exist in multi-head networks (Fig. S19d).

K DEPENDENCE OF DIMENSIONALITY OF POPULATION ACTIVITY ON N

We measure the dimensionality as the number of principal components that explain 90% of the population activity variance. The dimensionality increases with N for both tasks and curricula, but the increase follows a linear relation with N for N -parity task but a sub-linear relation for the N -DMS task (Fig. 6b).

To demonstrate this difference, we fit two separate lines for the data up to $N = 20$ and from $N = 20$ up to the largest N . We observe that for the N -parity task, the slope of two lines largely overlaps, indicating a linear relation. However, for the N -DMS task, the second line clearly has a smaller slope than the first one, indicating a sub-linear growth with N (Fig. S20).

L ABLATION DETAILS

To test whether neurons with fast or slow timescales (τ) are necessary for computations in the trained RNNs we perform the ablation analysis. For this analysis, we compute the relative accuracy of the model (Eq. 4 in the main text) after removing a single neuron. We ablate neuron i by setting all incoming and outgoing associated weights to zero

$$\begin{aligned} W_{ij}^R, W_{ji}^R &= 0 & \forall j \\ W_i^O, W_i^I &= 0 \end{aligned} \quad (8)$$

Here W_{ij} refers to recurrent weights, W_i^O to input weights and W_i^I to readout weights. To measure the relative accuracy, we simulate the RNN forward using random binary inputs for 1000 time steps after 100 time steps of a burn-in period (to reach the stationary state). Then, we evaluate the accuracy of the network at each time step. We repeat this procedure over 10 trials and compute the average and standard deviation of the relative accuracies across trials.

M SIGNIFICANCE OF THE RESPONSES TO PERTURBATIONS OF WEIGHTS AND RETRAINING

We investigate the significance of differences between single- and multi-head networks presented in Fig. 7 using a t-test (two-sided, unpaired). Perturbations are computed 10 times for 4 networks per group with results being pooled across networks. Retraining accuracy is computed once per network. Table 1, 2, and 3 indicates with stars the significance levels corresponding to p-values below $5e-2$, $1e-2$, $1e-3$, $1e-4$, and $1e-5$.

Weight Perturbation Strength	p-value	Significance
1.0e-02	8.8e-03	**
2.2e-02	4.5e-01	n/s
4.6e-02	2.9e-01	n/s
1.0e-01	5.1e-15	*****
2.2e-01	2.2e-39	*****
4.6e-01	8.3e-47	*****
1.0e+00	8.6e-04	***
2.2e+00	7.6e-01	n/s
4.6e+00	6.5e-01	n/s
1.0e+01	9.0e-01	n/s

Table 1: Significance of the weights' perturbation for different perturbation sizes Fig. 7c. Two-sided and unpaired t-test, stars indicate p-values below $5e-2$, $1e-2$, $1e-3$, $1e-4$, and $1e-5$.

Perturbation of τ	p-value	Significance
1.0e-03	4.3e-01	n/s
2.2e-03	8.9e-01	n/s
4.6e-03	5.4e-01	n/s
1.0e-02	1.2e-02	*
2.2e-02	2.3e-14	*****
4.6e-02	9.1e-29	*****
1.0e-01	1.3e-50	*****
2.2e-01	4.1e-35	*****
4.6e-01	4.8e-01	n/s
1.0e+00	6.0e-01	n/s
2.2e+00	1.3e-01	n/s

Table 2: Significance of the τ 's perturbation for different perturbation sizes Fig. 7d. Two-sided and unpaired t-test, stars indicate p-values below $5e-2$, $1e-2$, $1e-3$, $1e-4$, and $1e-5$.

retraining for N	p-value	Significance
17	1.7e-03	**
18	9.7e-04	***
19	1.2e-07	*****
20	3.2e-05	****
21	1.7e-05	****
22	1.6e-05	****
23	7.4e-06	*****
24	6.4e-05	****
25	3.1e-04	***
26	1.7e-03	**
27	1.7e-02	*
28	2.1e-03	**
29	7.5e-03	**
30	1.6e-01	n/s
31	1.5e-01	n/s
32	6.5e-01	n/s

Table 3: Significance of the retraining differences between single and multi-head, Fig. 7e. Two-sided and unpaired t-test, stars indicate p-values below $5e-2$, $1e-2$, $1e-3$, $1e-4$, and $1e-5$.

N CODE AND DATA AVAILABILITY

Codes for training and evaluating the RNNs and reproducing the experiments (e.g., measuring timescales, performing ablations, etc.) together with example trained networks are available on GitHub at https://github.com/LevinaLab/rnn_timescale_public (more details in README).

O SUPPLEMENTARY FIGURES

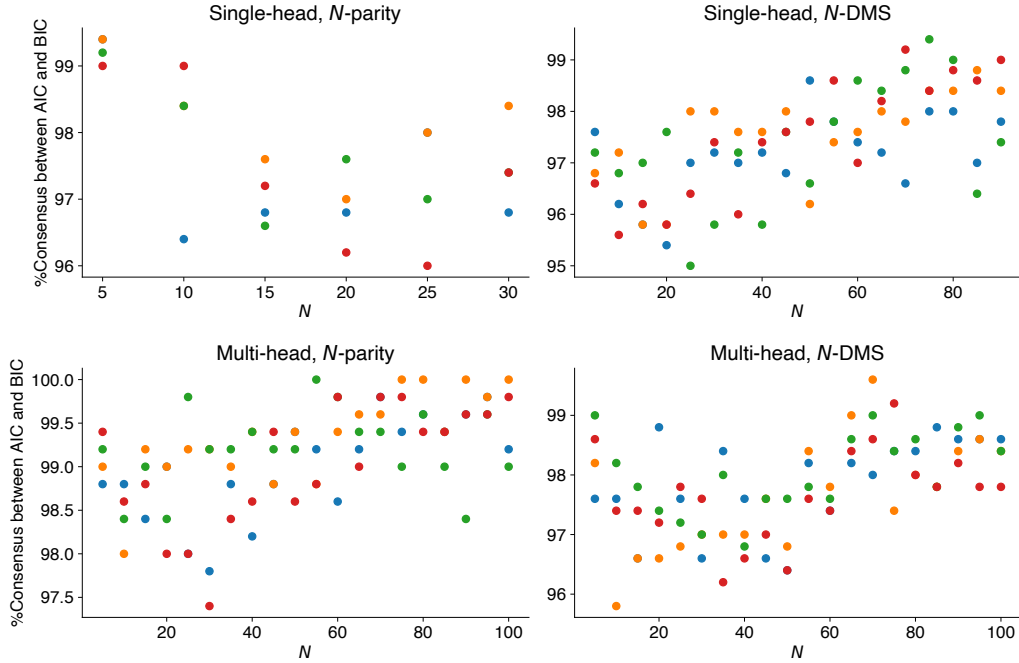


Figure S1: AIC and BIC choose the same model for most neurons. We fitted AC of each neuron’s activity with single- and double-exponential functions and used AIC or BIC to select the best-fitting models. The results show that for above 95% of neurons, the two criteria select the same model. The colors of the dots indicate different networks (4 networks for each task, curriculum and N).




Figure S2: ACs of neurons are well captured with single- or double-exponential fits. Note that the y-axis is in logarithmic coordinates, meaning that deviations between the fit and data AC are much smaller in the AC tail compared to initial time lags. **a, b.** Fitting double and single exponential functions to the AC of example (a) single-timescale ($\tau_{\text{net}} = \tau$) and (b) double timescale ($\tau_{\text{net}} > \tau$) neurons. **c.** Values of coefficient of determination R^2 estimated for all selected fits using AIC are close to 1, indicating a good fit.




Figure S3: Networks trained without curriculum have similar single-neuron (a) and network-mediated (b) timescales to networks trained with the single-head curriculum in the range of N that the no-curriculum-trained networks can learn.




Figure S4: Dependence of population activity timescales τ_{pop} on N . For both tasks and curriculum, the timescale of population activity fluctuations increases with N , indicating a general trend toward slower collective dynamics for tasks with larger memory requirements. Shade - \pm STD across 4 networks.

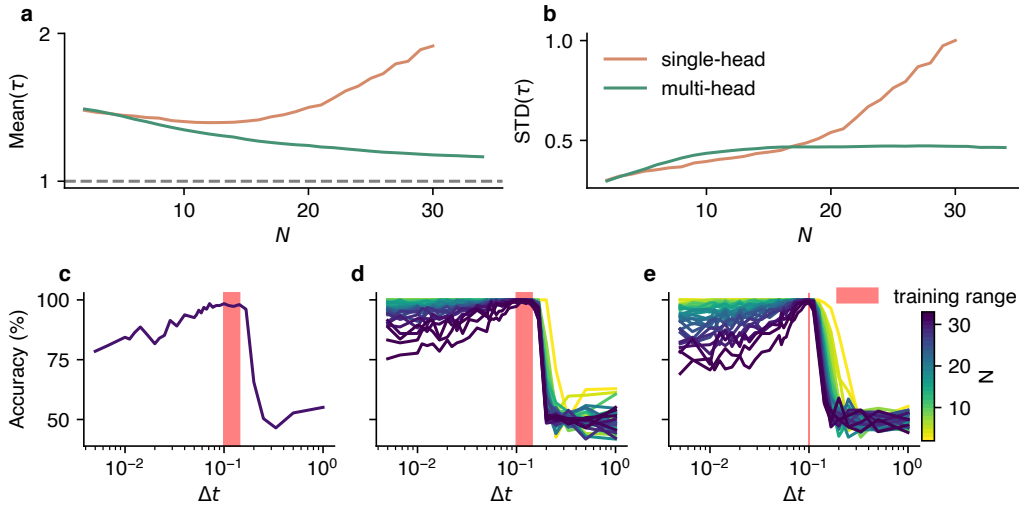


Figure S5: Impact of discretization time-step Δt on training performance. We train networks with $\Delta t = \{\frac{1}{10}, \frac{1}{9}, \frac{1}{8}, \frac{1}{7}\}$, while presenting each input digit for the duration of $T = 1$. (a) Similar to discrete-time networks ($\Delta t = 1$), the mean of single-neuron timescales τ increases with N for single-head networks and decreases towards $T = 1$ for multi-head networks. (b) The standard deviation of τ indicates heterogeneous τ s for single-head networks but constrained values for multi-head networks. (c,d) Single-head (c) and multi-head (d) networks can solve the task above the chance level for Δt smaller than their training regime (indicated by the red rectangle) when trained with multiple Δt . (e) The networks are slightly more inaccurate when trained on only a single $\Delta t = \frac{1}{10}$. Lines and the color bar indicate different N .




Figure S6: Changing the duration of input presentation. Each input digit is presented to RNN for a duration of $T = k\Delta t$, $\Delta = 1$. Single-neuron timescales (τ s) normalized by k remain roughly constant in multi-head networks (i.e. $\tau \rightarrow k\Delta t$), but increase with N in single-head networks (cf. Fig. 4b).




Figure S7: The development of timescales follows similar trajectories when the self-interaction is inside (nonlinear τ) or outside (linear τ) the nonlinearity (leaky-ReLU). Top: network-mediated timescales, bottom: single-neuron timescales. Shades - \pm STD.

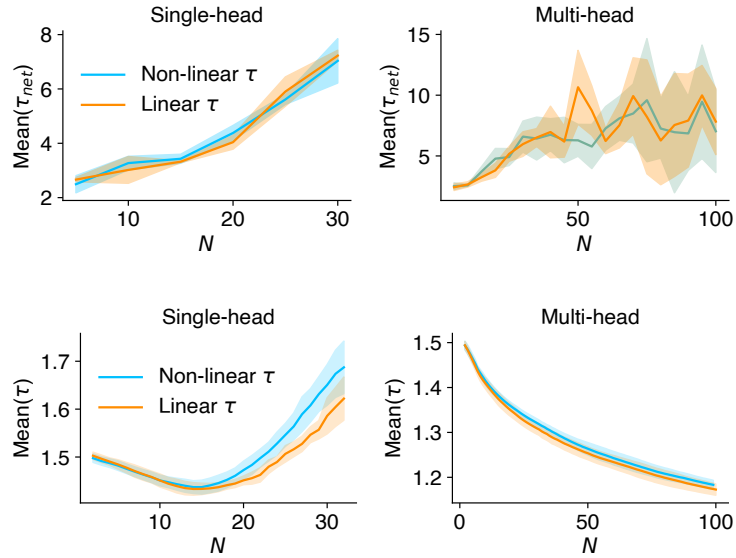


Figure S8: The development of timescales in networks with different nonlinearities follows similar trajectories. Top: network-mediated timescales, bottom: single-neuron timescales. Shades - \pm STD.

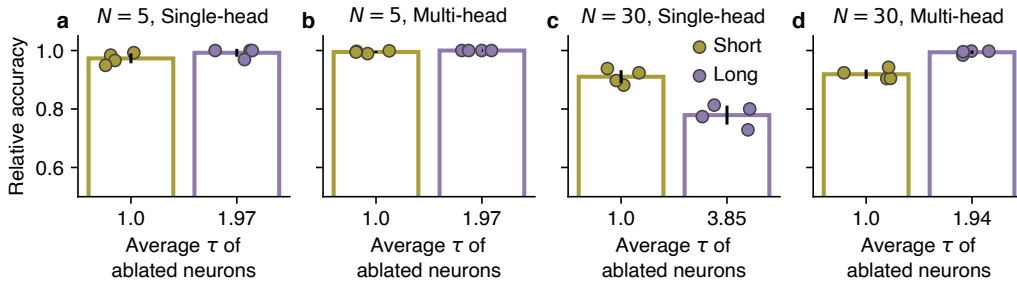


Figure S9: Impact of ablating neurons with distinct timescales on RNNs' performance when neural self-interactions are linear (cf. Fig. 7a,b). **a, b.** Ablating the longest and shortest timescale neurons has minimal effect on network performance when N is small for both curricula. **c, d.** For higher N , ablating long timescale neurons largely decreases the performance of single-head networks, while multi-head networks are more affected by the ablation of short-timescale neurons. Bars - mean, error bars - STD, dots - 4 individual networks.




Figure S10: The maximum N solved in the N -DMS task after 1000 epochs (reaching an accuracy of 98%). Similar to the N -parity task (cf. Fig. 4a), models trained with a single-head curriculum rely more on training τ than the multi-head curriculum networks, which prefer to have a small τ and are more agnostic to it being trainable. Horizontal bars - mean.




Figure S11: The maximum N solved in the N -parity task (for models with the self-interaction outside the nonlinearity) after 1000 epochs (reaching an accuracy of 98%). In the single-head curriculum, models rely on training τ , whereas in the multi-head curriculum, having τ s fixed at 1 value is as good as training them. Results are consistent with the models where the self-interaction is inside the nonlinearity (cf. Fig. 4), Horizontal bars - mean.

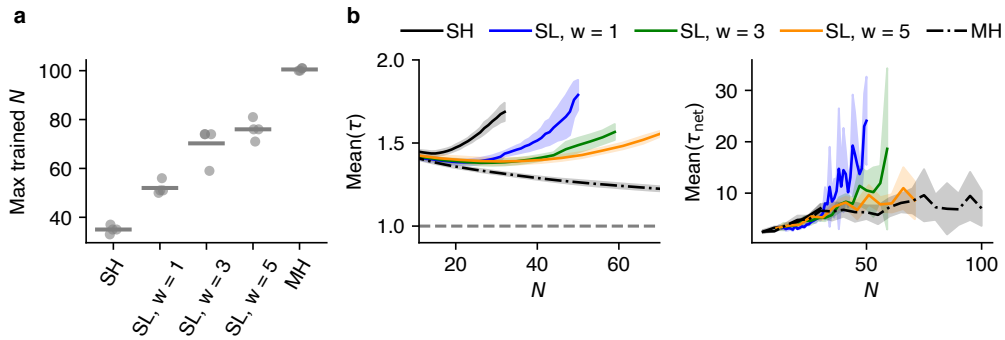


Figure S12: The behavior of networks trained with multi-head-sliding curriculum depends on the size of the sliding window and lies in between extreme curricula. **a.** The maximal trained N (with $> 98\%$ accuracy, within 1000 training epochs) for multi-head-sliding (SL) lies between single-head (SH) and multi-head (MH) networks and increases with the size of sliding window (w). Dots indicate individual networks (4 networks) and the horizontal bars indicate the mean value. **b.** Single-neuron (τ) and network-mediated (τ_{net}) timescales increase with N , but the pace of change reduces as the sliding window grows. Shadings indicate \pm std computed across 4 trained networks.




Figure S13: Training without a curriculum on the N -DMS task. For each N , 4 models are independently trained for 50 epochs or until reaching $> 98\%$ accuracy. Similar to the N -Parity task (cf. Fig. 3a), the ability to solve the task decreases as we increase N . From $N > 20$, we see the network is no longer capable of finding a solution to the task even with longer training time.

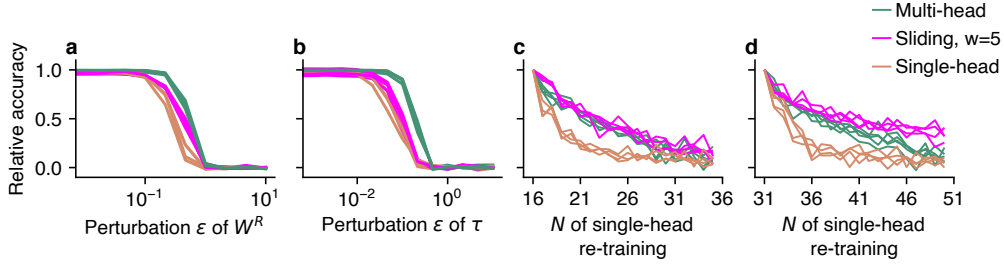


Figure S14: Robustness of networks trained with multi-head-sliding curriculum. **a, b.** Multi-head-sliding networks are more robust than single-head networks but less robust than the multi-head networks against perturbations of recurrent connectivity (a) and trained timescale τ (b). Each line indicates one trained network (4 networks for each curriculum). Shades indicate \pm std computed across 10 trials. **c, d.** retraining of networks trained with different curricula as a single-head network on new N s (for 20 epochs). Multi-head-sliding networks achieve higher relative accuracy when retrained for a higher N in comparison to single-head networks. If originally trained for small N s (c, $N = 16$), they have similar retraining accuracy to multi-head networks, but for larger N s (d, $N = 31$) their accuracy suppresses the multi-head networks. Each line indicates the relative accuracy for one network (4 networks for each curriculum).




Figure S15: Comparing the impact of curriculum on different recurrent architectures. Two different architectures, GRU and LSTM, are trained on the N -Parity task with and without a curriculum. We observe that the GRU and LSTM both exhibit instability when training without a curriculum (**a**), but are comparable to the RNNs with the multi-head curriculum (**b**).




Figure S16: Comparison of single- and multi-head curricula for training LSTMs on N -parity task. Networks trained with the multi-head curriculum can reach a higher N faster than networks trained with the single-head curriculum.




Figure S17: Single-head (a) and multi-head (b) networks loaded for $N \in [2, \dots, 19]$ have new readout heads retrained on new tasks with $N^* \in [2, \dots, N + 2]$. The heat map of the loss and accuracy of these retrained networks (after a maximum of 10 epochs or reaching an accuracy of 98%+) shows the robustness of the multi-head networks to catastrophic forgetting, as well as an improvement towards forward compatibility in the $N^* > N$ region. The dotted line indicates the diagonal.




Figure S18: Emergence of curriculum during multi-head training. **a.** In the absence of an explicit curriculum, multi-head networks solve smaller N s before solving the large N s. The color bar indicates the range of N s. **b.** The speed of training reduces with the increasing number of heads. The network with 19 heads needs fewer epochs to solve the same N (i.e. reaching 98% accuracy) than the network with 39 heads.

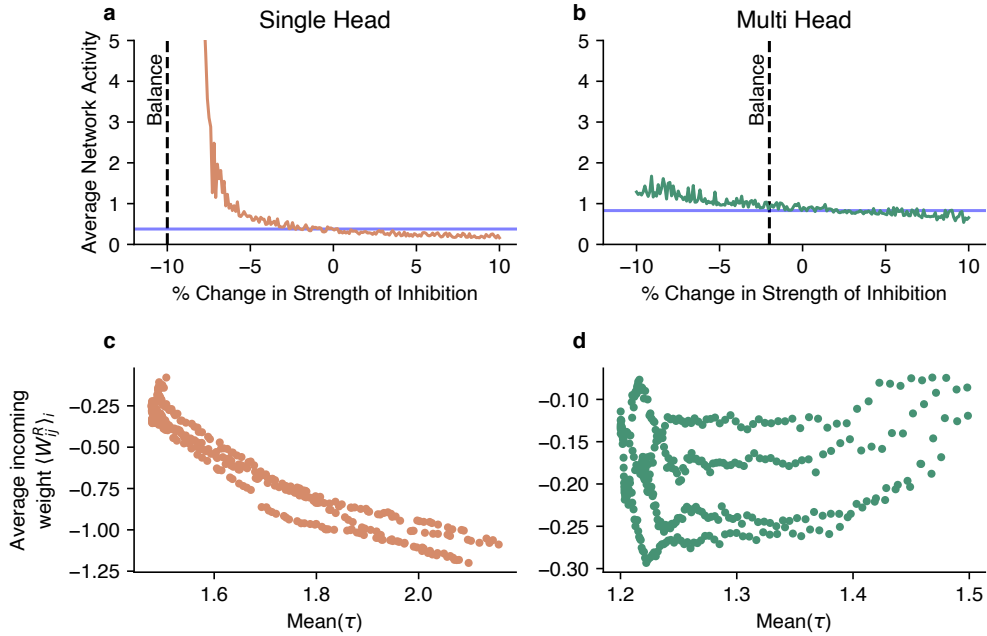


Figure S19: Networks trained with the single-head curriculum require strong inhibitory connectivity. We perturb the inhibitory (negative) connections of a single (a) and a multi-head network (b). We see that the activity of the single-head network explodes as we approach the balanced point (the point where the average of incoming weights becomes 0, indicated by the horizontal blue line). On the contrary, the multi-head network is quite robust and produces activity within a normal range even after the balanced point. (c) In the single-head networks, we observe a negative correlation between the average τ and the average strength of incoming weights for each neuron (i.e. higher τ is correlated with more negative average weight). This relationship is not present for multi-head networks (d) that are largely balanced and maintain small τ .

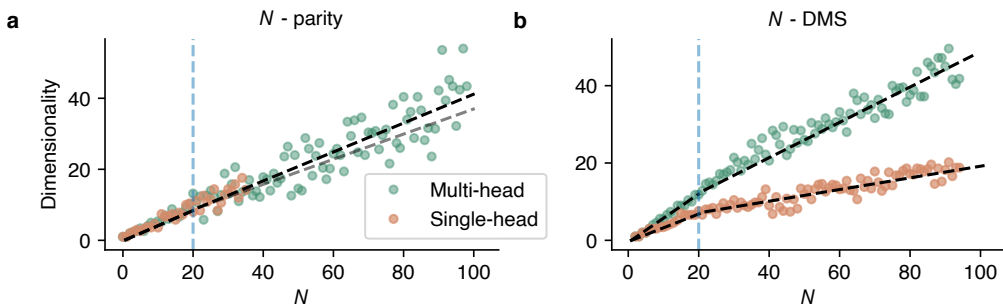


Figure S20: Dimensionality of activity increases approximately linearly with N for the N -parity task and sub-linearly for the N -DMS task. We separately fit the data points for $N \in [0, \dots, 20]$ and $N \in [20, \dots, 100]$ ($N \in [20, \dots, 30]$ for the single-head N -parity network) and we observe that in the N -parity task, the two lines largely coincide, while in the N -DMS case there is a clear change in the slope of the line, suggesting a sub-linear increase with N .

Appendix G Modular Growth of Hierarchical Networks: Efficient, General, and Robust Curriculum Learning

Contains the full text of the publication: Hamidi, M., Khajehabdollahi, S., Giannakakis, E., Schäfer, T. J., Levina, A., and Wu, C. M. (2024). Modular growth of hierarchical networks: Efficient, general, and robust curriculum learning. In ALIFE 2024: Proceedings of the 2024 Artificial Life Conference. MIT Press.

Modular Growth of Hierarchical Networks: Efficient, General, and Robust Curriculum Learning

Mani Hamidi^{1,2,*}, Sina Khajehabdollahi¹, Emmanouil Giannakakis^{1,2}, Tim J. Schäfer^{1,2},
Anna Levina^{1,2}, Charley M. Wu^{1,2}

¹Department of Computer Science, University of Tübingen, Tübingen, Germany

²Max Planck Institute for Biological Cybernetics, Tübingen, Germany

*mani.hamidi@uni-tuebingen.de

Abstract

Structural modularity is a pervasive feature of biological neural networks, which have been linked to several functional and computational advantages. Yet, the use of modular architectures in artificial neural networks has been relatively limited despite early successes. Here, we explore the performance and functional dynamics of a modular network trained on a memory task via an iterative growth curriculum. We find that for a given classical, non-modular recurrent neural network (RNN), an equivalent modular network will perform better across multiple metrics, including training time, generalizability, and robustness to some perturbations. We further examine how different aspects of a modular network's connectivity contribute to its computational capability. We then demonstrate that the inductive bias introduced by the modular topology is strong enough for the network to perform well even when the connectivity within modules is fixed and only the connections between modules are trained. Our findings suggest that gradual modular growth of RNNs could provide advantages for learning increasingly complex tasks on evolutionary timescales, and help build more scalable and compressible artificial networks.

Introduction

From bacteria (Andrews, 1998), to brains (Sporns and Betzel, 2016), to man-made artifacts (Lake et al., 2015), many things are composed of modular components, specialized for different purposes and capable of being recombined in distinct configurations to solve new problems. In cognitive science, modularity (Fodor, 1983) has played an important role in understanding intelligent behavior as composition of modular, symbolic representations (Rubino et al., 2023; Zhou et al., 2024) while evolutionary accounts have explored how selection pressure towards reducing connection costs may favor modular solutions (Clune et al., 2013). By constraining the search space (Happel and Murre, 1994) and favoring more computationally efficient solutions (Yuan et al., 2023), modular architectures (Amer and Maul, 2019) offer a complementary and more biologically plausible account of intelligence (Cosmides and Tooby, 1997), at a time when current trends in deep learning have pursued scale at all costs (Shen et al., 2023).

Taking inspiration from the modular duplication of entire body parts during gene duplication events (Garcia-Fernández, 2005), we explore the functional utility of modular growth in adapting a recurrent network to a memory task of gradually increasing complexity. We find that coordination of modular growth together with a learning curriculum facilitates a surprising array of advantages, both in terms of performance and costs.

Specifically, we focus on the domain of multi-timescale signal processing (Bathellier et al., 2008; Panzeri et al., 2010; Safavi et al., 2023). In tasks such as speech recognition (Graves et al., 2013), time-series prediction (Chung et al., 2014; Torres et al., 2021), and navigation (Morcos and Harvey, 2016), both biological and artificial agents commonly need to represent and remember relatively long timescales in the underlying network dynamics. In artificial neural networks, such dynamics can arise not only via the training of recurrent connectivity, but also by explicitly training the timescales of individual neurons (Perez-Nieves et al., 2021; Tallec and Ollivier, 2018; Quax et al., 2020; Yin et al., 2020; Fang et al., 2021; Smith et al., 2023). Trainable timescales have been linked with improved performance for rate- (Tallec and Ollivier, 2018; Quax et al., 2020) and spiking networks (Yin et al., 2020; Fang et al., 2021; Perez-Nieves et al., 2021). However, recent work has suggested that greater reliance on connectivity, rather than trainable timescales, is associated with superior performance and robustness of RNNs in similar memory tasks (Khajehabdollahi et al., 2024). Thus, these current debates suggest that the interaction between trainable and connectivity-based mechanisms, together with their dependence on architectural and training decisions remains relatively poorly understood, even in simple networks performing rather trivial tasks.

One such architectural distinction between artificial and biological networks is the presence of modular topologies (Litwin-Kumar and Doiron, 2012; Chaudhuri et al., 2014; Zeraati et al., 2023; Shi et al., 2023), which organize various neuronal types into higher level clusters with different timescales (Greengard, 2001) and specific local connectiv-

ity patterns (Schaub et al., 2015). In particular, hierarchical structures in the cortex are associated with a gradual increase in intrinsic timescales (Murray et al., 2014; Manea et al., 2022) that has been linked to the computational requirements of the tasks performed by different cortical regions (Murray et al., 2014).

Although artificial hierarchical networks have been used for temporal tasks in the past (El Hiji and Bengio, 1995), they only explored shallow and static hierarchies that lacked curriculum-based growing mechanisms. Thus, to our knowledge, the role of hierarchical structure in accommodating the emergence of longer timescales has yet to be investigated.

Here, we address this gap by comparing standard RNNs (non-modular) to growing modular networks, using memory tasks with increasing complexity (Fig. 1). Building on the success of different curriculum-based training strategies that support the learning of longer timescales (Khajehabdoollahi et al., 2024), we additionally introduce module-level duplicative growth of the network at every step along the curriculum. This modular strategy is loosely reminiscent of growth over evolutionary timescales that gradually accommodate adaptation to more complex tasks and environments (Lui et al., 2011). Our networks achieve superior performance and robustness, in line with other studies that use neural growth models to train RNNs in a different domain (Najarro et al., 2023).

Our findings suggest that the well-calibrated growth of structured networks can significantly reduce the number of trainable parameters and training steps required to solve complex temporal tasks.

Methods

We train non-modular and modular networks on the N -parity task at increasing complexity levels N and compare their performance, robustness, and learning dynamics. Models are trained using back-propagation-through-time with a stochastic gradient descent optimizer, a cross-entropy loss function, and curriculum learning¹.

Task

An N -parity task is a memory task that is commonly used to assess the capabilities of recurrent network architectures (Stork and Allen, 1992; Hohil et al., 1999). The task requires the accurate retention of a sequence of binary digits to perform a modulo 2 summation over the last N digits in the sequence. Therefore N captures the complexity of the task. The input is a random binary sequence S with length L chosen uniformly from the interval $\{N+2, 4N\}$, with one bit provided to the network at each time step. The network must output the binary sum (XOR) of the last N digits. To update the output at every time step, the network must learn

¹Code is provided at <https://github.com/manihamidi/growing-rnn>

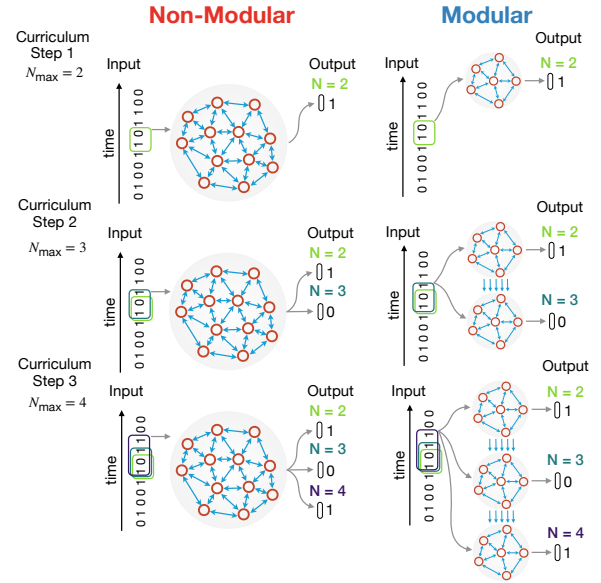


Figure 1: Network Structures. *Non-modular* networks consist of a fixed number of neurons ($M = 20, 54, 91, 128$), with connections that are retrained at each curriculum step. After a given accuracy is reached for a task of length N , a new readout head (consisting of 2 neurons) for $N + 1$ is added and the network is retrained for all previous N . In contrast, the *modular* network adds a much smaller RNN module ($M_m = 5, 10, 15, 20$) for every readout head that is added in the curriculum. Thus, each readout head is attached to a separate RNN module, rather than one large reservoir as in non-modular networks.

to store some representation of the values and order of the last N digits in memory.

Although simple, this task provides a foundation for testing the representation learning and memory capacities of artificial neural networks, while simultaneously allowing us to control the difficulty of the task 1 bit at a time using N . Furthermore, multiple N -parity tasks can be computed for a single sequence, allowing the possibility of training concurrent tasks/readout layers on the same inputs. This concurrent training is used to encourage more universal features that can be shared between the different tasks and to prevent catastrophic forgetting.

The success criterion for completing task N is therefore defined as jointly satisfying i) $> 98\%$ accuracy on task N and ii) maintaining an average of $> 98\%$ accuracy on all previous tasks. Accuracy is reported as an average over multiple tests, with random chance yielding 50%. We report network performance using N_{solved} which describes the largest task that it was able to solve after 60 epochs.

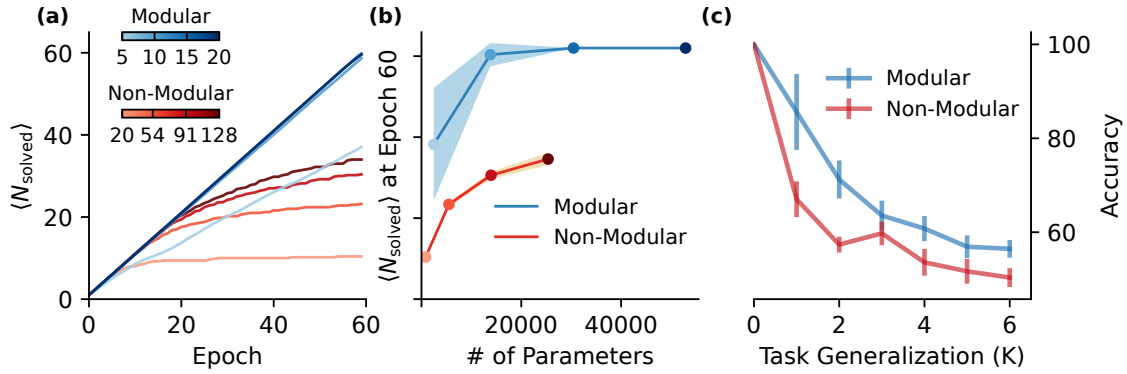


Figure 2: Performance comparison of modular and non-modular architectures with different numbers of neurons (color bars) allocated to their recurrent processing unit. For modular networks, the number of neurons is reported per module. For non-modular networks, the number indicates the total number of neurons in the single recurrent core. These numbers were chosen to allow for a comparable total number of learnable parameters for a given task difficulty N . (a) Learning curves show that modular architectures with sufficient (> 5) neurons can solve a new task at every epoch while the non-modular networks plateau in their learning ability. (b) Pareto frontier of performance after 60 training epochs, shows that the modular architecture always achieves better performance for the same number of parameters. Note that the saturation of performance at $N_{\text{solved}} = 60$ for the modular architecture is due to the choice of training epochs. (c) Generalization performance of networks trained on a task difficulty of N and then tested on tasks of $N + K$. Accuracy is measured as the percentage of correct trials and error bars indicate the standard deviation. Results are averages over 4 non-modular and 3 modular networks, and errorbars show STD.

Networks

We use simple rate neurons with a trainable time constant τ whose activity is defined by:

$$r_i(t) = \left(1 - \frac{\Delta t}{\tau_i}\right) \cdot r_i(t - \Delta t) + \frac{\Delta t}{\tau_i} \cdot [C_t]_\alpha \quad (1)$$

C_t is the input at each timestep t and the non-linearity $[\cdot]_\alpha$ is the leaky ReLU function with negative slope α , given by:

$$[x]_\alpha = \begin{cases} x, & x \geq 0 \\ \alpha \cdot x, & x < 0. \end{cases} \quad (2)$$

During learning, we train both the network connectivity (recurrent and feedforward) as well as the individual neuron timescales τ_i . The trainable timescale τ_i has a minimum value of 1 indicating a neuron that reacts only to the shortest timescales (i.e., the current input C_t), whereas larger values integrate information from longer timescales in the past.

Non-modular networks. In non-modular networks (Jaeger, 2002), the input at each timestep t for neuron i is:

$$C_t = \sum_{j \neq i}^M W_{ij}^R \cdot r_j(t - \Delta t) + W_i^I \cdot S(t) + b_i. \quad (3)$$

Here, W^R is the recurrent connectivity, W^I is the feedforward input connections, $S(t)$ is the input signal, b_i is the neuron bias, and time discretization $\Delta t = 1$.

The non-modular networks follow the multi-head curriculum in (Khajehabdollahi et al., 2024): At the first curriculum

step, a single linear readout head is trained to solve the task for $N = 2$. Upon successfully completing task N , a new readout head is added and is trained to solve the $N + 1$ -parity task. Previous readout heads continue to be trained after every step in the curriculum. Thus, at the m th step of the curriculum, the network has m readout heads solving the task for $N = 2, \dots, m - 1$.

Modular networks. In modular networks, each neuron receives input from the other neurons of the same module, and for modules $m > 1$, neurons also receive input from a feed-forward connection from the previous module $m - 1$:

$$C_t^m = \sum_{j \neq i}^{M_m} W_{ij}^R \cdot r_j^m(t - \Delta t) + \sum_k^{M_m} W_{ik}^{FF} \cdot r_k^{m-1}(t - \Delta t) + W_i^I \cdot S(t) + b_i \quad (4)$$

where W^{FF} is the feed-forward connectivity from the previous module and $r_k^{m-1}(t - \Delta t)$ is the activity of neuron k in module $m - 1$ from the previous timestep. M_m indicates the number of neurons in each module, which is fixed to be one of $M_m \in \{5, 10, 15, 20\}$ in our experiments.

For training the modular networks, we follow a growing curriculum as follows: Starting with a single module with a small population (5, 10, 15, or 20 neurons), we train a linear readout for $N = 2$. Then, at each curriculum step, we add a new module and readout head for each new N .

The connectivity of the new module is a copy of the corresponding (feedforward and recurrent) connectivity for the $N - 1$ module and receives feedforward input from it. Thus each module is specialized to solving the task for a single N . At each curriculum step, we freeze all network connections for all modules except the last and train only the recurrent and feedforward connections (input from sequence and input from previous module) of the last module.

Timescale Estimation

Apart from the trainable parameters τ_i that define the timescale of each neuron explicitly, we also compute the *effective* timescale of a neuron, which is influenced by its intrinsic trained τ_i and the extrinsic modulation of its activity through its connections with the rest of the network. We determine this network-mediated effective timescale for individual neurons by calculating the lagged autocorrelations (AC) of their activity during task performance, following previously developed methods (Khajehabdollahi et al., 2024). The AC of each neuron is then modeled with an exponential function featuring one or two timescales, selecting the best-fit model based on the Akaike Information Criterion (Akaike, 1974). To avoid estimation bias, we utilize long time series (10^5 time steps) of activity (Zeraati et al., 2022).

Robustness

We evaluate robustness by measuring the accuracy of the network after perturbing one of its three trained parameter groups, W^R , W^{FF} or τ . We define the magnitude of the perturbation ε_W as a function of the magnitude of the weights (Wu et al., 2020):

$$\widetilde{W} = W + \varepsilon_W \frac{\xi_W}{\|\xi_W\|} \|W\| \quad (5)$$

where $\xi_W \sim \mathcal{N}(0, \mathbb{I}^{n \times n})$ and $\|\cdot\|$ represents the Frobenius norm. This normalization allows for comparable amounts of perturbation across networks of different types and sizes. We also only perturb τ_i in positive direction to avoid $\tau < 1$.

Results

Our findings are divided into two main sections. In the first part, we demonstrate that modular networks following an iterative growing curriculum outperform equivalent non-modular RNNs in terms of task performance, training speed, generalizability, and robustness to perturbation of their learned connectivity. In the second part, using perturbation and ablation techniques, we investigate how different aspects of modular networks (feedforward vs. recurrent connectivity and trainable vs. effective single-neuron timescales) contribute to their computational capabilities.

Modular vs. Non-Modular Networks

Performance and Generalization We first compare a set of modular networks with different module sizes ($M_m =$




Figure 3: Weight perturbations degrade performance in both modular and non-modular networks of different sizes (colours). (a) Modular networks are more robust to perturbation of connections (modular networks: feedforward and recurrent weights). (b) Non-modular networks are more robust to the perturbation of single-neuron trained timescales.

5, 10, 15, 20) against a set of non-modular networks ($M = 20, 54, 91, 128$) with an equivalent number of trainable parameters (at $N = 30$). Performance is calculated as the average task difficulty solved, N_{solved} , both as a function of training time (Fig. 2a) and the number of parameters (Fig. 2b).

Overall, we find that all modular networks reach a higher task difficulty (N_{solved}) than the corresponding non-modular networks (Fig. 2a). Except for the smallest modular networks ($M_m = 5$), all modular networks maintain a steady, linear progression through the curriculum, solving every task after a single epoch of training. As an extreme test of the limits of the modular network, we were able to reach a maximum of $N = 200$ in one network before ending the simulation, with even greater capabilities being theoretically possible given sufficient training time. The non-modular networks, on the other hand, plateau much earlier in their progression through the curriculum. Thus, for a fixed number of trainable parameters, our results suggest there is a modular architecture that is capable of solving more complex tasks with the same amount of training.

We also explore the impact of a limited budget for learnable parameters as an alternative evaluation of the cost-benefit trade-off between two architectures. Figure 2b shows the Pareto frontiers of the architectures after 60 epochs, indicating that a modular architecture can solve a task with fewer parameters than a non-modular network.

Finally, we test generalization performance on networks of size $M_m = 15$ and $M = 91$, by training them to $N_{solved} = 10$ and then test their ability to solve $N = N_{solved} + K$, for $K = 1, 2, \dots$. For modular networks, the $N = 10$ module is now tested for solving $N = 10 + K$, while the non-modular network is tested on a new read-out head for $N = 10 + K$. We only allow 10 epochs of additional training on the new task. The results are shown in Figure 2c, where although performance decays with K for

both types of networks, modular networks consistently reach a higher accuracy compared to non-modular networks, thus demonstrating superior generalization.

Robustness to Perturbations. Next, we tested the robustness of both network architectures to perturbations of connectivity weights (Fig. 3a) and trained timescale parameters (Fig. 3b). The perturbations were proportional to the magnitude of the weights of the layer being targeted (Eq. 5). For different perturbation levels, we measure the network performance via the number of tasks that the network is still able to solve with accuracy > 0.9 .

Overall, Figure 3 shows the degradation of performance as a function of the size of the perturbation on the connections (both recurrent W_{ij}^R and feedforward W_{ij}^{FF} in the case of the modular networks) and the timescale parameters (τ_i) of the network. Modular networks not only have better initial performance but are also more robust against connectivity perturbations (Fig. 3a). This is visible in the later inflection point, reflecting how larger perturbations are required to degrade the performance of modular networks. However, modular networks are more sensitive to the perturbation of time-scale parameters (Fig. 3b). A special case is modular networks with $M_m = 5$ that appear to be more sensitive to both perturbations than their larger counterparts. For all other cases, the number of neurons does not significantly impact the robustness of either network type. So the inflection point of degrading performance is an inherent feature of network architecture. In sum, modular networks are significantly more robust against connectivity perturbations, but are more sensitive to perturbation of time-scale parameters.

Functional Analysis of Modular Networks

Trained vs. Effective Timescales. To better understand the increased sensitivity of modular networks to timescale perturbations, we examine the mechanism by which long timescales emerge in the networks. We do so by examining both the *trained timescale* parameterized by τ_i^m (Fig. 4a) and also the *effective timescale* inferred from the activity of each neuron (Fig. 4b). In the modular case, the average trained and effective timescales of neurons are reported separately per module solving $N = N_{solved}$. In the non-modular case, the average is taken over all neurons in the entire network at the time when it has just solved the $N = N_{solved}$ task.

First, we find that neurons in modular networks maintain stable trained timescales, $\langle \tau_i \rangle_m$, throughout the curriculum, but non-modular networks decay to the fastest possible rate of one (Fig. 4a). In both cases, these trained timescale values seem too small to account for the long timescales that are required for the large N tasks they can solve. This motivated us to examine the effective timescales, which are a product of the connectivity structure of the whole network.

Figure 4b shows the increase in the average effective

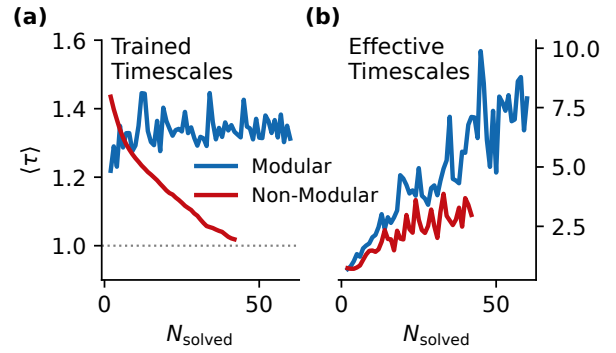


Figure 4: Change in trained and effective timescales for different N_{solved} . **(a)** The trained timescale of the modular network stays the same across modules, while the timescales of the non-modular network converge to 1 as the N_{solved} increases. **(b)** The effective timescales of both networks increase steadily with N_{solved} . A modular network with a module size of 15 neurons and a non-modular network with an equal number of trainable parameters were used. Results are averages over 4 networks.

timescales of neurons in both modular and non-modular networks for increasingly complex tasks N . This observation is consistent with recent observations (Khajehabdollahi et al., 2024) suggesting a circuit-level implementation of memory by learning appropriate connection weights (W) rather than learning slow timescales (τ_i) at the single neuron level. Notably, neurons in the modular networks appear to harbor memory of longer timescales compared to the non-modular counterparts at the same stage in the curriculum.

Recurrent vs Feedforward Connections. We now turn to how the topological constraints imposed by the modular architecture suggest different qualitative roles for feedforward and recurrent connections. For instance, the feedforward connections serve as the only bottleneck through which the recurrent computations of each module are reused by subsequent modules (Fig. 1), making them potentially more vulnerable to noise, whereas the recurrent connections may be more robust. Furthermore, the dependence of long effective neural timescales on appropriate connectivity patterns, encourages us to examine the distinct roles that feedforward (W^{FF}) and recurrent (W^R) weights play in performance.

To do so, we performed separate perturbation analyses targeting either feedforward W^{FF} or recurrent connections W^R . We evaluate performance by measuring the accuracy of each module's prediction for its corresponding task N .

Figure 5a-b shows that the feedforward connections are more sensitive to perturbations relative to the recurrent weights. Furthermore, modules corresponding to more difficult tasks (indicated by darker colors; Fig. 5a-b), suffer more because they are affected by the cumulative effect of all the




Figure 5: Feedforward connections are more sensitive than recurrent connections in modular networks. Here, we use $M_m = 15$ but achieve qualitatively similar results for other network sizes. **(a)** Perturbing feedforward weights affect downstream modules more strongly than earlier modules. **(b)** The recurrent weights exhibit a similar qualitative pattern but are quantitatively more robust against the same levels of perturbations. **(c)** Variability of feedforward versus recurrent connection weights across modules, where variability is inversely related to the degree of conservation (i.e., the amount of shared weights from one module to the next). Accuracy is averaged over 5 networks, 10 perturbations, and 1000 continuous evaluations. Error bars show SEM.

perturbations to preceding modules. Note that the effect of a perturbation is always downstream, due to the sequentially connected architecture of modules, such that a single perturbation equally affects all downstream modules.

To understand the basis for the sensitivity of W^{FF} , we hypothesize that W^{FF} are subject to more fine-tuning via back-propagation, in order to facilitate specialization to each new task. On the other hand, we expect W^R to be more conserved, and subject to less modification from one task to the next. We define the change in the connection weight between nodes i and j across two consecutive modules as $\Delta \bar{W}_{ij}$, where we normalized each weight by its mean and standard deviation to correct for systematic differences across module sizes (LeCun et al., 2012). The variance of the normalized weight change, $\sigma_{\Delta \bar{W}_{ij}}^2$, is the change in the normalized weight, $\Delta \bar{W}_{ij} = \bar{W}_{ij}^m - \bar{W}_{ij}^{m-1}$ between corresponding neurons, i, j , in consecutive modules, $m, m-1$. This serves as an inverse proxy for the degree to which the weights are conserved during training, where lower variance corresponds to greater conservation of weights.

Figure 5c shows how recurrent weights are more conserved than feedforward weights (lower variance). This supports the hypothesis that sensitivity of feedforward connections to perturbations is due to a more general sensitivity to modification, both from noise and from back-propagated error during training.

Weight Freezing. The low variability of the recurrent weights motivated us to test the degree to which a network can continue to learn despite a complete freezing of either its recurrent or the feedforward weights, after passing the very first step in the curriculum. A freezing of weights corresponds to a reduction of computational costs, by converting learnable parameters of a model into fixed biases. In evolu-

tionary biology, this is known as the Baldwin effect (Gruau and Whitley, 1993), where plastic behaviors become fixated in response to environmental stability. In the case of frozen recurrent weights, we trained the first module for $N = 2$ and then duplicated the learned recurrent weights on all subsequent modules, training only the feedforward connections between modules. Conversely, in the case of frozen feedforward weights, we train the feedforward connections between the $N = 2$ and $N = 3$ modules and then duplicate this trained connectivity between all subsequent modules, only training the internal recurrent connections of each module.

Figure 6a shows that networks with frozen feedforward weights fail catastrophically to solve the task for any $N \geq 3$. However, equivalent networks with frozen recurrent connectivity (resembling deep reservoirs; Gallicchio and Micheli, 2021), perform much better, comfortably reaching an $N_{solved} \approx 15$ before stagnating (Fig. 6a) and even outperforming non-modular networks that taper at $N \approx 10$. Thus, while both network types underperform the original network for which both recurrent and feedforward connections are trained, we see that, in line with our previous results, the learnability of the feedforward connections appears to be much more important for task performance.

Weight Duplication. The weight freezing experiments involved the reuse (or duplication) of weights from the previous module followed by a complete blockage of their training. Here we consider duplicative versus random initialization of the weights, and maintain the plasticity of weights during training. Theoretically, duplication can be beneficial by amortizing the costs of learning through the reuse of already-trained weights, or it could be detrimental by initializing the network in a state unsuitable for the new task. Empirically, previous work using weight-agnostic neural networks (Gaier and Ha, 2019), compositional pattern-

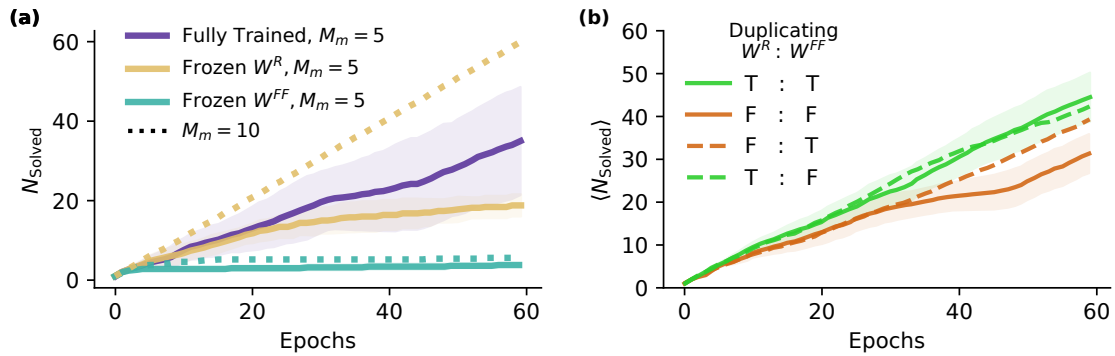


Figure 6: The impact of weight freezing and duplication on feedforward and recurrent weights. **(a)** Weight freezing. Solid lines are a modular network with $M_m = 5$, selected for the strongest differences. Here, we can see that freezing recurrent connections (orange) after solving the first task ($N = 2$) still needs to relatively good performance, whereas freezing feedforward connections (teal) severely impairs learning. In networks of larger module size ($M_m = 10$; dotted lines), the difference between frozen recurrent weights and a fully trained network disappears (both dotted orange), while frozen feedforward weights remain impaired. **(b)** Weight duplication. New modules are initialized with either a duplicated copy (T) of their weights as they appeared in the previous module, or a random initialization of weights (F). Duplication of recurrent weights confers a slight advantage (green), while either duplicative or randomly initialized feedforward weights both confer a slight impairment to performance. Results are averages over 5 networks and shaded regions indicate SEM.

producing networks (Stanley, 2007), and hypernetworks (Ha et al., 2016), have demonstrated the utility of similar “weight-sharing” schemes that support high-performing neural networks with a fraction of the parameter count. Motivated by these past successes, here we explore whether the duplication of recurrent or feedforward components can accelerate the training of the network, or if such methods are detrimental by trapping the behavior of the network in a local optima.

Figure 6b shows the results of these experiments, where weights were either duplicated (T for true) or not (F for false); in case of the latter, they were initialized randomly following a uniform distribution. The lines show the average performance of 6 networks in each of the 4 different conditions, with and without duplication of feedforward W^{FF} and/or recurrent W^R weights. These results show that the duplication of recurrent weights W^R appears to account for an unambiguous advantage (Fig. 6b; green curves), compared to the duplication of feedforward weights, which do not confer any additional benefits. Indeed, duplicating feedforward weights confers little to no improvement to when neither weights are duplicated (Fig. 6b, brown curves). We also experimented with noisy duplication, but given the weak overall effect, we only show the results of exact duplication, without noise. Furthermore, different duplication strategies mattered less at larger module sizes, so we focused only on networks with five nodes per module $M_m = 5$.

These results are consistent with our previous analysis showing that feedforward connections must diverge from their past configurations to specialize for their respective task. Therefore, as long as the feedforward connections re-

main plastic after a growth event, they are molded strongly by the error signal from their new task, making their initial state irrelevant. In contrast, we have already shown how keeping the same or similar recurrent weights across multiple tasks is relatively adaptive. Thus, duplication simply initializes the weights near their last functional state, eliminating the need to retrain them to that stage at every epoch, and allows for minimal fine-tuning to improve their adaptiveness over longer timescales.

Discussion

In sum, we find that modular networks, trained via a growing curriculum to solve a memory task of adjustable difficulty, outperform standard recurrent neural networks (RNNs) of fixed size in terms of accuracy, training time, generalization, and robustness to perturbations of the learned connectivity.

Biological neural networks are highly structured (Hagmann et al., 2008; Sporns and Zwi, 2004; Sporns and Betzel, 2016), with multiple neuron types each following distinct connectivity patterns (Peng et al., 2021; Liu et al., 2023) that form complex circuits repeating across brain regions (Douglas and Martin, 2004; Shepherd, 2011). The repetition of similar complex structures in the brain potentially enables the development of very complex neural circuits from limited genetic information (Rakic, 2009; Geschwind and Rakic, 2013; Stanley, 2007), and may also play a role in more efficient information processing (Bassett and Bullmore, 2006; Sporns, 2013). Moreover, inhomogeneous network topology has been shown to generate non-trivial dynamics (Litwin-Kumar and Doiron, 2012) and act as an inductive bias for local learning (Giannakakis et al., 2023).

Here, our findings add to this literature by suggesting relatively simple topological modularity can be utilized to align network dynamics with task requirements, thus increasing performance and reducing training time and cost.

Our modular architecture allows us to study the roles played by different aspects of the network's structure. In particular, we focus on the distinction between the recurrent weights within each module and the feedforward weights connecting adjacent modules. Our analysis indicates that the training and task specialization of feedforward weights between modules are vastly more important than the recurrent connections within each module. This suggests that the function of the recurrent connectivity is largely limited to creating appropriate and generic dynamical units (and can thus be effectively reused through duplication), while the feedforward connections control the flow of information that is necessary for solving the specifics of a given task.

While our network included trainable single-neuron timescales, which have been utilized in RNNs solving temporal tasks (Tallec and Ollivier, 2018; Quax et al., 2020; Khajehabdollahi et al., 2024), we find that their importance is minimal in terms of network performance. The distribution of trained timescales barely changes across modules, suggesting that the network fully relies on its connectivity (particularly the hierarchical topology) for developing the long timescales necessary for solving temporal tasks. This augments previous findings suggesting that network topology can generate long timescales with relatively random connections (Khajehabdollahi et al., 2024).

This understanding of the relative importance of different network components, allows us to test whether training even fewer parameters can lead to comparable results. We find that training the networks as deep reservoirs (Gallicchio and Micheli, 2021) leads to reasonably good performance despite reducing the trainable parameters by half. This suggests that functional sub-networks can be used as building blocks in larger systems with little or no training, which seems like a promising avenue for future research.

The importance of initializing neural networks in an appropriate dynamical regime has been widely explored (Zierenberg et al., 2020; Khajehabdollahi et al., 2022), and our findings here suggest that hierarchical modular structures could be used to generate networks with beneficial dynamics for learning temporal tasks. Recycling functional structures and incorporating them in complex networks of interacting sub-units is a widespread characteristic of biological networks (Felleman and Van Essen, 1991) and has shown promise in artificial settings (Happel and Murre, 1994; Sharkey, 1996; Amer and Maul, 2019). Our findings suggest that this approach is not only more efficient in terms of the number of trainable parameters, but it can also boost performance as well as the ability to generalize.

Limitations & Future Work. A basic premise of our work was that the power and cost of a neural network are proportional to the number of its trainable parameters. The number of parameters, therefore, provided a basis for comparing neural networks of two different architectures. Our results indicate that one can maintain or even exceed performance while being more frugal in resource expenditure on connections. While this result is relevant both for artificial (Amer and Maul, 2019) and biological networks (Clune et al., 2013), where these connections carry non-trivial costs, it does overlook the relative cost of producing and maintaining neurons in addition to their connections. Our modular architecture incurs a heavy cost in terms of individual neural processing units, scaling linearly with task complexity in proportion to each module's size. Indeed, the ability of modular networks to learn longer effective timescales points to the importance of large but sparse networks in solving a memory task. Future work could address the trade-off between these two resources and their impact on performance.

Finally, even though we found a functional role for duplicative growth in artificial networks, the observed benefits were modest. In evolutionary contexts, relaxed selection (Deacon, 2010) provides a mechanism by which duplication can accommodate enhanced adaptation by fostering synergistic interactions between duplicated modules. In our current setup, each duplicated module is required to solve a new task, subjecting it to heavy selection pressure that prevents it from exploring the entirety of the weight space. Relaxing this constraint by introducing redundancy in modules that are assigned the same task, could amplify the positive effects of duplication in future research.

In conclusion, modular growth offers a promising avenue for training lighter networks with more scalable properties. In the era of large models with increasingly prohibitive costs, there is great interest in developing new tools to compress models with fewer parameters (Cheung et al., 2019) and fine-tune them to perform novel tasks (Hu et al., 2021). Our work, like other evolutionary-inspired approaches (Akiba et al., 2024), can offer new directions towards this end.

Acknowledgements

MH, EG, TS and SK thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for their support. MH, EG, and CMW are supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A. MH and CMW are funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC2064/1—390727645. TS is supported by Else Kröner Medical Scientist Kolleg “ClinbrAI: Artificial Intelligence for Clinical Brain Research”.

References

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.
- Akiba, T., Shing, M., Tang, Y., Sun, Q., and Ha, D. (2024). Evolutionary optimization of model merging recipes.
- Amer, M. and Maul, T. (2019). A review of modularization techniques in artificial neural networks. *Artificial Intelligence Review*, 52:527–561.
- Andrews, J. H. (1998). Bacteria as modular organisms. *Annual review of microbiology*, 52(1):105–126.
- Bassett, D. S. and Bullmore, E. (2006). Small-world brain networks. *The neuroscientist*, 12(6):512–523.
- Bathellier, B., Buhl, D. L., Accolla, R., and Carleton, A. (2008). Dynamic ensemble odor coding in the mammalian olfactory bulb: Sensory information at different timescales. *Neuron*, 57(4):586–598.
- Chaudhuri, R., Bernacchia, A., and Wang, X.-J. (2014). A diversity of localized timescales in network activity. *eLife*, 3.
- Cheung, B., Terekhov, A., Chen, Y., Agrawal, P., and Olshausen, B. (2019). Superposition of many models into one.
- Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Clune, J., Mouret, J.-B., and Lipson, H. (2013). The evolutionary origins of modularity. *Proceedings of the Royal Society B: Biological sciences*, 280(1755):20122863.
- Cosmides, L. and Tooby, J. (1997). The modular nature of human intelligence. *The origin and evolution of intelligence*, pages 71–101.
- Deacon, T. W. (2010). A role for relaxed selection in the evolution of the language capacity. *Proc. Natl. Acad. Sci. U. S. A.*, 107 Suppl 2(Suppl 2):9000–9006.
- Douglas, R. J. and Martin, K. A. (2004). Neuronal circuits of the neocortex. *Annu. Rev. Neurosci.*, 27:419–451.
- El Hahi, S. and Bengio, Y. (1995). Hierarchical recurrent neural networks for long-term dependencies. In *Proceedings of the 8th International Conference on Neural Information Processing Systems*, NIPS’95, page 493–499, Cambridge, MA, USA. MIT Press.
- Fang, W., Yu, Z., Chen, Y., Masquelier, T., Huang, T., and Tian, Y. (2021). Incorporating Learnable Membrane Time Constant to Enhance Learning of Spiking Neural Networks. *arXiv:2007.05785 [cs]*.
- Felleman, D. J. and Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral cortex (New York, NY: 1991)*, 1(1):1–47.
- Fodor, J. A. (1983). *The modularity of mind*. MIT press.
- Gaier, A. and Ha, D. (2019). Weight agnostic neural networks. *Advances in neural information processing systems*, 32.
- Gallicchio, C. and Micheli, A. (2021). *Deep Reservoir Computing*, pages 77–95. Springer Singapore, Singapore.
- Garcia-Fernández, J. (2005). The genesis and evolution of homeobox gene clusters. *Nat. Rev. Genet.*, 6(12):881–892.
- Geschwind, D. H. and Rakic, P. (2013). Cortical evolution: judge the brain by its cover. *Neuron*, 80(3):633–647.
- Giannakakis, E., Vinogradov, O., Buendia, V., and Levina, A. (2023). Recurrent connectivity structure controls the emergence of co-tuned excitation and inhibition. *bioRxiv*.
- Graves, A., Mohamed, A.-r., and Hinton, G. (2013). Speech recognition with deep recurrent neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 6645–6649. Ieee.
- Greengard, P. (2001). The Neurobiology of Slow Synaptic Transmission. *Science*, 294(5544):1024–1030. Publisher: American Association for the Advancement of Science.
- Gruau, F. and Whitley, D. (1993). Adding learning to the cellular development of neural networks: Evolution and the baldwin effect. *Evolutionary computation*, 1(3):213–233.
- Ha, D., Dai, A., and Le, Q. V. (2016). Hypernetworks. *arXiv preprint arXiv:1609.09106*.
- Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J., and Sporns, O. (2008). Mapping the structural core of human cerebral cortex. *PLoS biology*, 6(7):e159.
- Happel, B. L. and Murre, J. M. (1994). Design and evolution of modular neural network architectures. *Neural networks*, 7(6-7):985–1004.
- Hohil, M. E., Liu, D., and Smith, S. H. (1999). Solving the n-bit parity problem using neural networks. *Neural Networks*, 12(9):1321–1323.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. (2021). LoRA: Low-Rank adaptation of large language models.
- Jaeger, H. (2002). Tutorial on training recurrent neural networks, covering bppt, rtrl, ekf and the echo state network approach.
- Khajehabdollahi, S., Prosi, J., Giannakakis, E., Martius, G., and Levina, A. (2022). When to Be Critical? Performance and Evolvability in Different Regimes of Neural Ising Agents. *Artificial Life*, pages 1–21.
- Khajehabdollahi, S., Zeraati, R., Giannakakis, E., Schäfer, T. J., Martius, G., and Levina, A. (2024). Emergent mechanisms for long timescales depend on training curriculum and affect performance in memory tasks. In *The Twelfth International Conference on Learning Representations*.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338.
- LeCun, Y. A., Bottou, L., Orr, G. B., and Müller, K.-R. (2012). Efficient BackProp. In Montavon, G., Orr, G. B., and Müller, K.-R., editors, *Neural Networks: Tricks of the Trade: Second Edition*, pages 9–48. Springer Berlin Heidelberg, Berlin, Heidelberg.

- Litwin-Kumar, A. and Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience*, 15(11):1498–1505. Number: 11 Publisher: Nature Publishing Group.
- Liu, L., Yun, Z., Manubens-Gil, L., Chen, H., Xiong, F., Dong, H., Zeng, H., Hawrylycz, M., Ascoli, G. A., and Peng, H. (2023). Neuronal connectivity as a determinant of cell types and subtypes. *Research Square*.
- Lui, J., Hansen, D., and Kriegstein, A. (2011). Development and evolution of the human neocortex. *Cell*, 146:18–36.
- Manea, A. M., Zilverstand, A., Ugurbil, K., Heilbronner, S. R., and Zimmermann, J. (2022). Intrinsic timescales as an organizational principle of neural processing across the whole rhesus macaque brain. *eLife*, 11:e75540.
- Morcos, A. and Harvey, C. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. *Nature Neuroscience*, 19.
- Murray, J., Bernacchia, A., Freedman, D., Romo, R., Wallis, J., Cai, X., Padoa Schioppa, C., Pasternak, T., Seo, H., Lee, D., and Wang, X.-J. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature neuroscience*, 17.
- Najarro, E., Sudhakaran, S., and Risi, S. (2023). Towards Self-Assembling Artificial Neural Networks through Neural Developmental Programs. volume ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference, page 80.
- Panzeri, S., Brunel, N., Logothetis, N. K., and Kayser, C. (2010). Sensory neural codes using multiplexed temporal scales. *Trends in Neurosciences*, 33(3):111–120.
- Peng, H., Xie, P., Liu, L., Kuang, X., Wang, Y., Qu, L., Gong, H., Jiang, S., Li, A., Ruan, Z., et al. (2021). Morphological diversity of single neurons in molecularly defined cell types. *Nature*, 598(7879):174–181.
- Perez-Nieves, N., Leung, V. C. H., Dragotti, P. L., and Goodman, D. F. M. (2021). Neural heterogeneity promotes robust learning. *Nature Communications*, 12(1):5791.
- Quax, S. C., D’Asaro, M., and van Gerven, M. A. J. (2020). Adaptive time scales in recurrent neural networks. *Scientific Reports*, 10(1):11360. Number: 1 Publisher: Nature Publishing Group.
- Rakic, P. (2009). Evolution of the neocortex: a perspective from developmental biology. *Nature Reviews Neuroscience*, 10(10):724–735.
- Rubino, V., Hamidi, M., Dayan, P., and Wu, C. M. (2023). Compositionality under time pressure.
- Safavi, S., Chalk, M., Logothetis, N., and Levina, A. (2023). Signatures of criticality in efficient coding networks. *bioRxiv*, pages 2023–02.
- Schaub, M. T., Billeh, Y. N., Anastassiou, C. A., Koch, C., and Barahona, M. (2015). Emergence of Slow-Switching Assemblies in Structured Neuronal Networks. *PLOS Computational Biology*, 11(7):e1004196. Publisher: Public Library of Science.
- Sharkey, A. J. C. (1996). On combining artificial neural nets. *Connection science*, 8(3-4):299–314.
- Shen, L., Sun, Y., Yu, Z., Ding, L., Tian, X., and Tao, D. (2023). On efficient training of large-scale deep learning models: A literature review.
- Shepherd, G. M. (2011). The microcircuit concept applied to cortical evolution: from three-layer to six-layer cortex. *Frontiers in neuroanatomy*, 5:30.
- Shi, Y.-L., Zeraati, R., Levina, A., and Engel, T. A. (2023). Spatial and temporal correlations in neural networks with structured connectivity. *Physical Review Research*, 5(1):013005.
- Smith, J. T., Warrington, A., and Linderman, S. (2023). Simplified state space layers for sequence modeling. In *The Eleventh International Conference on Learning Representations*.
- Sporns, O. (2013). Network attributes for segregation and integration in the human brain. *Current opinion in neurobiology*, 23.
- Sporns, O. and Betzel, R. F. (2016). Modular brain networks. *Annual review of psychology*, 67:613–640.
- Sporns, O. and Zwi, J. D. (2004). The small world of the cerebral cortex. *Neuroinformatics*, 2:145–162.
- Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8:131–162.
- Stork, D. G. and Allen, J. D. (1992). How to solve the n-bit parity problem with two hidden units. *Neural networks*, 5(6):923–926.
- Tallec, C. and Ollivier, Y. (2018). Can recurrent neural networks warp time? *arXiv preprint arXiv:1804.11188*.
- Torres, J. F., Hadjout, D., Sebaa, A., Martínez-Álvarez, F., and Troncoso, A. (2021). Deep learning for time series forecasting: a survey. *Big Data*, 9(1):3–21.
- Wu, D., Xia, S.-T., and Wang, Y. (2020). Adversarial weight perturbation helps robust generalization. *Advances in Neural Information Processing Systems*, 33:2958–2969.
- Yin, B., Corradi, F., and Bohtë, S. M. (2020). Effective and Efficient Computation with Multiple-timescale Spiking Recurrent Neural Networks. In *International Conference on Neuro-morphic Systems 2020, ICONS 2020*, pages 1–8, New York, NY, USA. Association for Computing Machinery.
- Yuan, X., Savarese, P., and Maire, M. (2023). Accelerated training via incrementally growing neural networks using variance transfer and learning rate adaptation. *Advances in Neural Information Processing Systems*, 36.
- Zeraati, R., Engel, T. A., and Levina, A. (2022). A flexible Bayesian framework for unbiased estimation of timescales. *Nature Computational Science*, 2(3):193–204.
- Zeraati, R., Shi, Y.-L., Steinmetz, N. A., Gieselmann, M. A., Thiele, A., Moore, T., Levina, A., and Engel, T. A. (2023). Intrinsic timescales in the visual cortex change with selective attention and reflect spatial connectivity. *Nature Communications*, 14(1):1858.

Zhou, H., Nagy, D. G., and Wu, C. M. (2024). Harmonizing program induction with rate-distortion theory. *arXiv preprint arXiv:2405.05294*.

Zierenberg, J., Wilting, J., Priesemann, V., and Levina, A. (2020). Tailored ensembles of neural networks optimize sensitivity to stimulus statistics. *Phys. Rev. Res.*, 2:013115.