Discovering plasticity rules for learning and resilience in neural circuits

David G Bell

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Adrienne Fairhall, Chair

Paul Wiggins

Marcel den Nijs

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

Discovering plasticity rules for learning and resilience in neural circuits

David G Bell

Chair of the Supervisory Committee:
Professor Adrienne Fairhall
Department of Neurobiology and Biophysics

While modern supervised and reinforcement learning algorithms can train neural networks to solve a wide range of tasks, the brain often operates in data-sparse regimes where such extensive supervision is unavailable. This thesis argues that the brain succeeds in these settings by leveraging inductive biases about the tasks it is likely to encounter. These biases are embedded in initial connectivity, cell-type structure, and critically, in synaptic plasticity rules. Here, we investigate how unsupervised synaptic plasticity can shape neural circuits prior to extensive behavioral experience.

In the first part of this thesis, we study plasticity in the zebra finch song system. In collaboration with researchers at the California Institute of Technology, we examine the restoration of singing behavior following viral perturbation of nucleus HVC, a premotor region essential for song production. Adult male zebra finches transiently lose song after viral manipulation but recover within approximately two weeks. Strikingly, birds prevented from practicing during early recovery subsequently require less practice to regain song, suggesting that recovery is partially unsupervised. We model this process using several unsupervised plasticity mechanisms, including spike timing-dependent and homeostatic plasticity. While standard homeostatic rules restore regular spiking activity in a network model of HVC, they fail to reproduce experimentally observed synaptic reorganization. We therefore propose a local population-level homeostatic rule that recruits previously silent neurons, accounting for both activity recovery and synaptic changes.

In the second chapter, we employ meta-learning, a technique by which biologically plausible learning rules are learned via a supervised procedure, to discover biologically plausible plasticity rules that organize robust sequential dynamics in HVC-like networks. In this framework, candidate unsupervised plasticity rules are optimized by a supervised outer loop to maximize a task objective. Starting from disordered connectivity, the learned rules reliably self-organize networks into sequence-generating circuits resembling those observed in vivo. Analysis of resulting rules reveals that plasticity on recurrent excitatory synapses generalizes Oja's rule, replacing the classical Hebbian term with a spike timing-dependent component. We further show that learned plasticity rules can compensate for continual synaptic turnover and that learned inhibitory plasticity enhances the precision and robustness of sequential dynamics.

In the final chapter, we apply meta-learning to the self-organization of neural integrators—circuits that generate long timescales via carefully tuned structure to maintain representations of sensory inputs. Such integrators underlie functions including head direction coding and oculomotor control. We hypothesize that unsupervised plasticity can shape these circuits from weak structural priors. Using meta-learning, we identify plasticity rules that reliably organize integration dynamics without requiring previously hypothesized anti-Hebbian mechanisms. Instead, the learned rules rely heavily on three-factor plasticity. In a simplified model, we demonstrate how such three-factor mechanisms can tune integrator circuitry and stabilize persistent dynamics.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

So many people made the completion of this thesis possible. Foremost, I would thank my adviser, Adrienne Fairhall, whose scientific sensibility, enthusiasm for gathering our community, and prose are unmatched. I would also like to thank Alison Duffy, who served in many ways as a second mentor, and whose patience and endless enthusiasm were instrumental in every project I completed. I have deeply appreciated and benefited from the companionship of my lab mates: Patrick Zhang, Ilse Dippenaar, Faeze Aminmansoor, Scott Sterrett, Rich Pang, as well as folks not in the lab, but who make the Computational Neuroscience Center what it is, including Jessica Huszar, Po-Chen Kuo, Srinidhi Naidu, Leenoy Meshulam, and Eric Shea-Brown.

I also have a host of collaborators to thank. At the Allen Institute, Cindy Poo, Bruno Cruz, Tiffany Ona, and Lukas Braun. I am particularly thankful to have met Laura Driscoll, whose kindness and excitement about science are infectious and who greatly helped shape the narrative of this thesis. I would would like to thank the members of the Lois lab at CalTech, including Carlos Lois, Bo Wang, and Zsofia Török.

I would like to thank the members of my committee, who graciously made time for me: Jason Detwiler, Paul Wiggins, Marcel den Nijs, Eli Shlizerman, as well as Eric and Adrienne. My PhD cohort (too many to name!), but you all were so kind and excited, and it meant so much to me that we didn't just take classes together.

Finally, I'd like to thank my friends and family. You kept me sane and moving forward and gave these years their joy.

DEDICATION

To my family and friends.

Chapter 1

INTRODUCTION

In understanding our remarkable capacity to move, communicate, think, and problem-solve, evolution provides essential context. It is extremely probable that the brain did not arise as a general-purpose computing device, but as an organ shaped by natural selection to support survival and reproduction under specific ecological pressures. In this view, we expect the rules by which the brain updates itself to be molded by these pressures. This evolutionary origin imposes a set of constraints that can guide our understanding of neural computation and learning. This thesis is heavily informed by three such constraints: (1) that neural circuits may be optimized not for general problem solving, but for a restricted set of behaviorally relevant tasks; (2) that sample efficiency and rapid acquisition of reasonable performance may be prioritized over eventual peak performance; and (3) that robustness, reliability, and reparability may be favored over fragile optimality. Together, these constraints suggest a view of biological learning as an extensive bootstrapping process, in which strong architectural priors and local learning rules enable fast, dependable performance across a modest number of critical domains, rather than maximal flexibility or asymptotic optimality.

1.1 What is a biological learning rule?

Changes in the strength of connections between neurons are widely thought to underlie learning in animals [1]. In vertebrates, these connections are primarily chemical synapses. Depolarization of a neuron's membrane potential propagates along its axon and triggers the release of neurotransmitters at synaptic terminals, where they bind to receptors on the dendrites of postsynaptic neurons. This synaptic transmission modulates postsynaptic membrane potentials, altering the likelihood of subsequent depolarization and producing cascading effects across neural circuits.

Early work in *Aplysia*, a sea slug, demonstrated that habituation and dishabituation of a

gill-withdrawal reflex could be attributed to changes in the efficacy of excitatory synapses, providing the first link between a learned behavior and the strength of chemical synapses [2, 3]. Subsequent studies in vertebrates connected the repetitive activation of excitatory pathways in the hippocampus to long-lasting changes in synaptic strength, suggesting a mechanism by which neural activity induced by experience could be consolidated into the synaptic machinery [4]. Together, these results suggest that learning and memory are mediated, at least in part, by activity-dependent modifications of synaptic strength.

Despite extensive progress in identifying the mechanisms of synaptic plasticity in specific experimental paradigms, there is currently no universal theory governing the evolution of synaptic efficacy. Synaptic change has been linked to a wide range of local biophysical processes and properties, including the activity of pre- and postsynaptic neurons, the relative timing of this activity [5], synaptic size [6], and neuromodulator-dependent changes in ionic concentrations and conductances [7]. How these diverse signals are integrated to guide learning in general circuits, beyond well-characterized special cases, remains an open question.

1.2 Overview of neural plasticity rules

Below, we review what is currently known about the mechanisms by which brains modify synaptic strength in response to neural activity. These mechanisms span a diverse range of timescales, and may depend on the activity of a single neuron or the activity of a population of presynaptic cells.

1.2.1 Long-term plasticity

Perhaps the most famous form of plasticity, long-term plasticity refers to activity-dependent changes in synaptic efficacy that last up to hours or days. Its discovery in the hippocampus, a brain area implicated in learning and memory, by Bliss and colleagues provided some of the first insight into the mechanisms by which memories might be formed and maintained [4, 8]. It has since been shown that long-term potentiation requires activity in both the pre- and postsynaptic neurons, making it a compelling potential anatomical substrate for associative learning [9, 10]: it was found that weak synapses could be potentiated if their pre- and postsynaptic neurons were co-activated, while suppression of postsynaptic response by

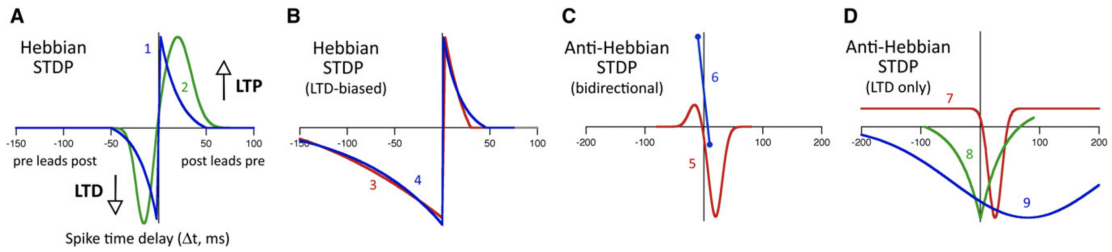

Figure 1.1: **Types of STDP (A)** Classic STDP, in which post leading pre leads to potentiation and pre leading post leads to depression. **(B)** Classic STDP, but with a large window of depression. **(C)** Anti-Hebbian STDP, in which the sign of the typical kernel is reversed. **(D)** All inhibitory STDP in which both time lags lead to depression. Figure reproduced from Feldman.

hyperpolarizing or voltage clamping the postsynaptic cell eliminated LTP [11, 12], confirming the predictions of Hebb roughly 30 years prior [1]. Most synapses capable of undergoing long-term potentiation (LTP) are also capable of undergoing depression (LTD), suggesting flexible, bidirectional modification of synapses in response to learning signals [13]. High and low frequency presynaptic firing rates are thought to mediate these two regimes, respectively [14]. These changes were later found to depend on activation of NMDA (N-methyl-D-aspartate) receptors; pharmacological blockade of these receptors was shown to impair spatial learning and memory [15, 16].

1.2.2 Spiking timing-dependent plasticity (STDP)

An incredible theoretical insight by Wulfram Gerstner and colleagues first led to the prediction that long-term potentiation might exhibit dependence on the precise timing and ordering of pre- and postsynaptic spikes [17]. It was later experimentally confirmed that the sign and magnitude of LTP and LTD could depend on the ordering and time lag between pre- and post-synaptic spikes as well as the frequency of these events [5, 18]. In its canonical form, STDP describes potentiation that occurs when presynaptic activity leads postsynaptic activity, and

depression that occurs when it follows. The magnitude of the synaptic changes generally correlates with the proximity in time of the two spikes [5]. Plasticity typically requires many (60-100) pairings of pre- and postsynaptic spikes [14]. In electric fish, it was discovered that pre leading post induced depression and post leading pre induced potentiation, corresponding to a sign change of the typical kernel [19], which has hence been termed anti-Hebbian spike timing dependent plasticity. An important variant of anti-Hebbian STDP induces depression at both positive and negative time lags of pre and post [20]. Across a range of cell types and organisms, forms of STDP vary widely, suggesting that the exact form used by a brain circuit may be specific to the computation it performs [14]. Further generalizations of STDP exist, such as triplet plasticity, which generalizes pair-based rules and considers triplets of spikes [21].

1.2.3 Short-term plasticity (STP)

Short-term plasticity (STP) is a considerably faster form of plasticity that transiently modifies the efficacy of synapses on timescales from milliseconds to minutes. STP arises from the mechanics of molecular transmission between the two neurons, and can both enhance and reduce the effect of future stimulation of the postsynaptic cell [22]. Depression can result from the depletion of readily available neurotransmitter-carrying vesicles, feedback inhibition of presynaptic receptors, or receptor desensitization at the postsynaptic target [23]. Facilitation may result from residual calcium, wherein calcium fails to fully clear between spikes and increases the release probability upon future spikes [24], the existence of calcium sensors, which bind Ca^{2+} and make additional spikes more probably, [25], and the saturation of Ca^{2+} buffers [26]. These changes can be observed after paired pulse stimulation with short inter-pulse duration (20 ms), and also after longer trains of stimuli. The effects of STP are typically most noticeable in synapses that exhibit a high probability of neurotransmitter release prior to stimulation [13]. It was shown in populations of pyramidal neurons that STP could affect the code by which cells communicate by modifying the postsynaptic response to presynaptic firing rates and the temporal coherence of inputs depending on the level of neurotransmitter depletion [27]. Beyond shaping ongoing neural responses, STP has also

been theorized to be an important part of working memory [28, 29].

1.2.4 Homeostatic mechanisms

Another prominent class of plasticity rules describes mechanisms that serve to maintain a target level of activity for an individual neuron, i.e. a form of homeostatic regulation. One well studied plasticity rule by which homeostatic regulation of firing rate occurs is synaptic scaling, wherein neurons maintain a target firing rate by scaling up or down their incoming synapses, typically on timescales of hours to days [30]. This form of plasticity was first discovered in cultures of neocortical pyramidal neurons, which were found to restore a baseline level of activity in response to chronic activity suppression or blockade of GABA-mediated inhibition [6]. Synaptic scaling rescales synapses multiplicatively (hence its name), which is theorized to preserve the relative efficacy of distribution of inputs a postsynaptic neuron receives, while normalizing the firing rate of the cell. The rules for scaling depend on synapse type: inhibitory synapses tend to strengthen in response to increased postsynaptic activity [30]. Adaptation of the intrinsic excitability of neurons has also been shown to maintain a target firing rate [31, 32]. In cultures of pyramidal neurons, this was shown to occur through adaptation of voltage-dependent conductances [33]. Theoretical work has proposed that these two distinct forms of homeostatic plasticity may regulate different aspects of recovery from perturbation [34].

1.2.5 Three factor plasticity

While most early work in plasticity focused on the modulation of synapses by the activity of pre- and postsynaptic neurons, an emerging literature has begun to describe plasticity in terms of three factor interactions. So-called three factor plasticity rules describe extensions to two-factor plasticity rules in which the pairwise rule is modulated or gated by a third factor [35–37]. A wide range of neuromodulators, including dopamine, adrenaline, noradrenaline, acetylcholine, and serotonin have been shown to modify effective plasticity rules by, in particular circumstances, modifying the STDP window, affecting the learning kernel shape (e.g. converting an asymmetric kernel to a symmetric one or changing the sign of synaptic

modification), and gating synaptic change [7]. GABA, the neurotransmitter that typically mediates inhibition, has also been shown to affect plasticity of excitatory synapses in particular cases [38].

1.2.6 Behavioral timescale synaptic plasticity

In a recent landmark paper, it was shown that the formation of hippocampal CA1 place fields is mediated by a novel form of plasticity that does not depend on postsynaptic activity, but rather pairing between presynaptic activity and calcium plateau potentials in the postsynaptic neuron [39]. Potentiation via this plasticity is extremely rapid, requiring 1-2 pairings *in vivo*, and 5 when the plateau potential is artificially induced. Further, the timescale of the pairing window was seconds long, in stark contrast to the typical ∼20 millisecond timescale of LTP, leading to the term "behavioral timescale" plasticity. The discovery of this mechanism has challenged a host of historical assumptions about the typical timescales of learning rules and the rates at which they permanently modify synapses.

1.3 Learning rules in the context of living organisms

How might we expect biological learning rules to function given the pressures imposed by evolution? Here, we discuss a number of heuristics, constraints, and challenges that inform the nature of learning rules in the brain.

1.3.1 Learning rules can take advantage of the specific structure of brain circuits

In many cases, biological learning rules operate within circuits designed for specific, behaviorally relevant tasks. This is particularly true of invertebrates, like *Aplysia*, *C. elegans*, and the fruit fly *Drosophila*. In these organisms, neurons reliably take particular roles and connect to one another via largely stereotyped wiring diagrams. A compelling theory is that the structure of brain circuits does not simply provide a promising starting point for generalized learning rules, but that biological learning rules are adapted to function within specific circuits to achieve particular aims. Thus, if the structure of a circuit is arranged to solve a particular task, the learning rule operating within it can leverage the circuit's

structure to accelerate learning.

An example of the synergy between the structure of a neural circuit and a learning rule acting upon it can be found within the head direction system of *Drosophila*. This circuit, located within the central complex, tracks the orientation of the insect's head direction relative to salient landmarks [40, 41]. Glutamatergic neurons within the ellipsoid body, known as E-PG cells, connect to one another in a ring-like topology and fire persistently to represent particular angles of this orientation [42]. The structure of this circuit appears to be arranged for the specific task of tracking a continuous periodic variable. It has been shown that the persistent activity representation of the insect's heading integrates self-motion cues, allowing the fly to navigate in darkness. In addition to exciting their neighbors, E-PG neurons receive retinotopically organized inhibitory visual inputs via a separate cell type, known as ring neurons. The direct mapping of visual input to this circuit aligns the heading representation naturally with input from the eyes. Learned depression of the connections between ring and E-PG neurons can be used to correct estimates of head direction based on incoming visual cues. A recently discovered plasticity rule has been shown to facilitate this depression based on activity in ring neurons and coincident release of a neuromodulator, octopamine, leveraging the topology of the circuit and conveniently distributed visual input to facilitate rapid corrections to heading via visual cues [43].

Although less common, vertebrates also possess circuits that appear to be dedicated to specific tasks. The zebra finch brain contains a set of nuclei (non-laminar collections of tens of thousands of neurons) dedicated to song learning and production. These nuclei are arranged in two pathways; one is believed to be an "experimental" pathway by which song is learned, while the other consolidates successful changes to song discovered by the first. These two pathways converge in premotor area RA, which sits upstream of the bird's vocal organ. A specialized learning rule was discovered in which inputs from one pathway to this area modulate plasticity on inputs from the other, intriguingly with opposite sign [44]. This plasticity rule may act to mediate the balance in drive from the experimental pathway to the more stereotyped one as learned song is consolidated.

1.3.2 Biological learning may favor early suboptimal performance over eventual optimal performance

While learning via trial and error can eventually lead to optimality, evolutionary pressures may rather have crafted neural circuits and paired learning rules that exhibit extreme sample efficiency. As highlighted in a recent perspective, a host of animals rapidly become competent at survival-essential behaviors following birth, suggesting that sufficient early performance is often an evolutionary bottleneck [45]. In extreme cases, behaviors may require no or nearly no learning, as in the weaving of a web by a spider [46] or the flight of mice in response to looming stimuli. Such bottlenecks may also explain the lack of optimality of animals in certain domains: if the circuitry required to perform some behavior optimally does not merit its energetic cost, or worse, impedes survival at early stages, learning rules may reflect simpler, more heuristic strategies.

1.3.3 Robustness and reliability

Neurons and synapses are a dynamic and therefore challenging substrate to support intelligence. It has been estimated that synapses between pyramidal neurons in mouse hippocampal area CA1 have a mean lifetime of two weeks [47]. Assuming independence, this result implies complete turnover of the synaptic population in roughly 2-3 times that timescale. In areas of cortex, the longevity of synapses is estimated to be longer: an estimated 73% of dendritic spines were stable for at least 8 days in mouse somatosensory cortex [48]. In addition to synaptic turnover, cell death, whether natural or induced by trauma or disease, further contributes to the shifting nature of the brain's structure. In extreme cases, brain areas demonstrate remarkable preservation of function in the face of perturbations, including transection, viral infection, and ablation of inputs. A number of examples will be covered in Section 2.4. Taken together, these findings imply that the learning algorithms employed by the brain are able to compensate for these complications to enable behavior.

1.3.4 Locality

Perhaps the clearest constraint on the nature of synaptic learning rules is the fact that information used by plasticity to modify the synapse must be local to the synapse. Despite an extremely high degree of connectivity between cortical neurons (cortical neurons receive roughly 10,000 synapses in some estimates [49]), an individual cell may still receive no direct information about the state of neurons to which it is not directly connected or spatially proximal. Thus, the brain must coordinate its neurons on the basis of partial information about its global state. This constraint has proven particularly problematic in making correspondences between learning in the brain and in artificial networks, which are typically trained using backpropagation [50], the supervised learning algorithm that has enabled the modern age of artificial intelligence, which requires neurons to have access to nonlocal information.

1.4 The search for learning rules via meta-learning

Given the large space of potential learning rules and the many factors upon which they might depend, a compelling approach is to start not from the learning rule, but rather from the desired function of a circuit, and to ask which rules are consistent with that function [51]. An emerging approach, termed "meta-learning," has been to express the desired function of a circuit via a loss function, and then to tune the parameters of a local plasticity rule via a supervised optimization procedure to minimize the loss. The earliest known approach of this nature was reported in Bengio et al.. Meta-learning typically unfolds via an inner and outer loop. In the inner loop, the dynamics of a network and a parameterized local plasticity rule are allowed to evolve over time. The outer loop then evaluates the resulting dynamics or network structure via the loss function, and modifies the parameters of the learning rule in order to minimize the loss in future iterations. Typical choices for the supervised optimization procedure include evolutionary strategies [51] and gradient-based methods [53].

Meta-learning presents a number of advantages over classic computational approaches to discovering plasticity rules. First, it alleviates the process of tuning plasticity rule parameters by hand; this aspect is particularly useful when tuning multiple plasticity rules within the

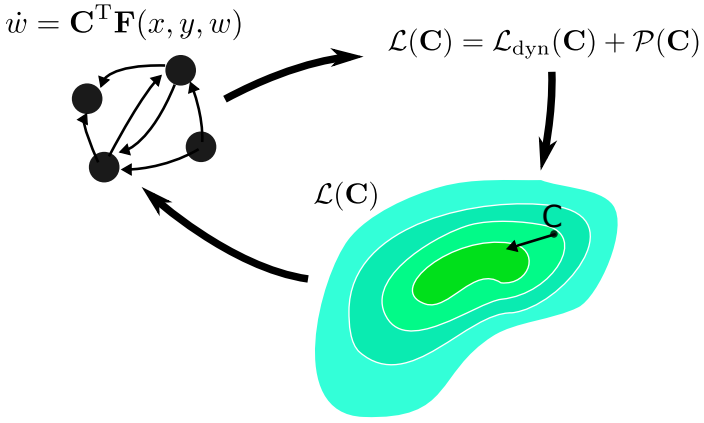

Figure 1.2: **Schematic of meta-learning** Meta-learning proceeds via an inner and outer loop. In the inner loop, synapses evolve in simulation along with network dynamics according to a rule parameterized by a set of parameters, C, which weight a set of basis functions, F. Once run forward, the results of the simulation are evaluated with respect to a loss function, $\mathcal{L}(C)$, which is used to tune C in the space of possible parameterizations. An additional penalty, $\mathcal{P}(C)$ can be added to encourage sparsity in the basis.,

same network, in which case the interaction between distinct rules can have unexpected consequences [54]. In addition, the loss function may be defined as average or median loss over a number of networks, imposing the requirement that a discovered rule perform well in a number of network contexts. This approach may also allow researchers avoid biases about the types of plasticity that enable a particular function. Automated discovery may lead to the discovery of novel plasticity mechanisms, and reveal in certain contexts that many rules may lead to similar function [55].

While meta-learning is often effective in finding plausible plasticity rules, interpreting its results can be challenging. Even relatively low-dimensional bases can quickly become difficult to interpret, and this can be exacerbated if the basis functions are themselves parameterized combinations of many variables [54]. Another issue is that of long-term stability of the joint plasticity rule and network dynamics: if dynamics are only unrolled for a finite period, meta-learning does not guarantee stability of longer timescales. Some authors solve this by imposing penalties for weights that do not converge over the course of the simulation period, but this may rule out plasticity rules that rely on continual synaptic change [54].

To date, research in meta-learning plasticity has largely had two foci. The first is the search for a biologically plausible replacement for backpropagation [52, 56–59]. Here, the aim has been to discover a single rule that could act as a replacement for typical error-propagation methods, but which does not violate the idea that the synapse should update its strength based on locally available information. These studies primarily focus on feedforward network settings in which error signals are randomly projected back to neurons in the network. In all work we are aware of, the performance of backpropagation has typically stood as an upper bound on discovered plasticity rules on tasks considered difficult by modern machine learning standards.

The second area of focus is the discovery of unsupervised plasticity mechanisms. Here, potential rules are not functions of supervised targets or error, and instead evolve weights purely from local variables, such as pre- and postsynaptic activity, synapse size, inputs from distinct cell types, and membrane voltage. Work in this setting has focused on recovering known plasticity rules [51], discovering novel timing-based kernels that organize cortical-like population spiking dynamics [58, 60], continual familiarity detection [61], representation

learning [62], and embodiment [63].

While not the focus of this thesis, a complementary line of work uses meta-learning to search for architectural priors capable of enabling rapid learning and efficient generalization. Research in this space has considered compressible structural priors [64, 65], and has shown that biologically inspired network architectures can lead to rapid learning on naturalistic tasks [66].

1.5 Summary of contributions

The bulk of this thesis explores the application of automated procedures to discover local, non-feedback-mediated plasticity rules that organize sequential and integration dynamics. In Chapter 2, I describe my work modeling the outcome of an experimental manipulation of a sequence generating network in the bird brain. In this work, I implemented a network model in which I explored learning rules that are capable of restoring the sequence and that satisfied constraints from experimental measurements of synaptic strength. We proposed that a novel population level learning rule can explain the experimental data. The complexities of proposing and testing many variants of rules motivated my interest in a more automated approach. Thus, in Chapter 3, I take a meta-learning approach to discovered plasticity rules that lead to the self-organization and maintenance of sequential dynamics in networks and excitatory and inhibitory neurons. In Chapter 4, I apply this meta-learning strategy to a more challenging problem of broad interest in neuroscience: the self-organization of line attractor dynamics. For a small circuit model, I derive a novel set of rules that are capable of solving this problem. I introduce a novel approach to visualize and characterize the resulting high-dimensional learning rules. Finally, I propose reduced sets of rules and test their similarity to the discovered rule sets.

Chapter 2

RESILIENCE AND RECOVERY THROUGH PLASTICITY AND NETWORK ARCHITECTURE IN THE ZEBRA FINCH

Some of the contents of this chapter are published in Nature Neuroscience as Wang et al..

2.1 Introduction

Brains must reliably perform computations necessary for survival and reproduction. They face a host of challenges in doing so, however, including the effects of disease, cell and synapse death, aging, and trauma. Here, we investigate the mechanisms behind the remarkable reliability of brain circuits by studying the perturbation and recovery of the dedicated circuit that underlies singing behavior in the zebra finch. In this chapter, we discuss several aspects of this reliability, including the plasticity rules that may serve to repair the circuit and the dynamics of the circuit itself, which display little sensitivity to change in parameters.

The bulk of this chapter describes modeling and theory results that accompanied a set of experiments in which the motor system that governs singing behavior in the songbird was disrupted by a set of cell type-specific viral perturbations. Here, select cell types within HVC, a premotor nucleus essential to singing behavior in adult male zebra finches, were disrupted, resulting in the eventual silencing of song. However, after roughly ten days, animals largely recovered their ability to sing. We investigate a number of plasticity mechanisms and their potential to effect recovery of typical dynamics within a model of HVC, including homeostatic plasticity, spike-timing dependent plasticity, and a hypothetical population-level homeostatic rule. Under the hypothesis that these rules remain continually active in the absence of perturbations, we describe aspects of plasticity rules that are necessary for stable sequential activity.

A second set of results focuses on the dynamics of HVC, which experimental evidence suggests provides a timing signal for downstream pre-motor areas. We first develop a simple model of the dynamics within HVC based on the hypothesis that its structure is essentially

a feedforward excitatory network with winner-take-all inhibition [68]. This feedforward, sequential connectivity can support a propagating mode of excitation which serves as a clock for the downstream motor areas. We show within our model that the speed of this propagating mode through the network varies minimally as the strength of excitatory and inhibitory synapses are changed, while the number of spikes in the propagating mode and the average frequency of these spikes can change more freely.

2.2 Preliminaries on the zebra finch

Adult male zebra finches sing a highly stereotyped song composed of syllables, the smallest unit of song, which are roughly 100 ms long and strung together into *motifs* with intervening silences. These motifs are repeated, resulting in *bouts*, which are often 1-2 seconds long. Mature song production is not innate, but rather is learned from a tutor when the bird is a juvenile. Song acquisition can generally be divided into three stages: subsong, plastic song, and crystallized song. During the subsong phase (beginning 25-30 days post-hatch), the bird's vocalizations are similar to babbling observed in infants [69]. During the plastic song phase (beginning 50 days post-hatch), the juvenile's song begins to take the shape of the tutor's, but both the acoustic features of individual syllables and their orderings are variable. In the crystallized phase of song (beginning 90-120 days post-hatch), the ordering of syllables within the motif and acoustic features solidify into the song the bird will sing for the remainder of its life [70].

A juvenile's eventual acquisition of song is predicated on exposure to tutor song. Auditory exposure to mature song is sufficient to trigger the process of imitation. This exposure must occur during a first critical period (25-65 days post-hatch), resulting in the bird acquiring a memorized template of song. Acquisition of the template is distinct from the process of vocal experimentation: zebra finches exposed to tutor song but isolated from a tutor by the time they begin singing nevertheless learn to imitate the tutor song [71, 72]. Birds denied exposure to tutor song during this critical period will never develop song [73, 74].

The process of song imitation involves a feedback loop of experimentation and self-evaluation. Auditory interventions that prevent birds from accurately evaluating their song, such as deafening or playing noises that mask the bird's own vocalizations, result in malformed

song [75, 76]. Reinforcement learning is thought to underlie this iterative learning process [77–79]. A number of studies have correlated dopamine neuron activity during song with classic reward prediction error signals [80, 81]. Population dopamine activity is predictive of a song's evolution in acoustic feature space, and is consistent with a single vocal target [80].

2.2.1 The neural substrate of song acquisition and production in the zebra finch

Twin neural pathways underlie the acquisition and production of song in the zebra finch. The first, known as the Anterior Forebrain Pathway (AFP), is thought to be responsible for the acquisition and maintenance of song. The AFP begins in HVC (formerly, the High Vocal Center, now simply referred to as a proper noun) and continues through the basal ganglia nucleus Area X, to the Dorsal Lateral Nucleus of the Medial Thalamus (DLM), to the Lateral Magnocellular Nucleus of the Anterior Nidopallium (LMAN). LMAN then projects to the Robust Nucleus of the Archistriatum. The Descending Motor Pathway (DMP), forms a complementary path, beginning in HVC, continuing to RA, which projects to the dorsal medial nucleus (DM) and the tracheosyringeal part of the hypoglossal nucleus (nXIIts), which innervates the syrinx, the bird's vocal organ. The DMP is thought to be responsible for mature song production.

A series of studies have disambiguated the roles of these two pathways. Lesions to RA and HVC have demonstrated their necessity for singing behavior in birds with crystallized song, with bilateral lesions of HVC resulting in permanent ablation of song [82, 83]. The same lesions in the female have no effect on their calls, distinguishing these nuclei as underlying learned song as opposed to innate vocalizations. Lesions to Area X have no effect on crystallized song, but abolish song learning [84, 85]. Further, X lesions paired with deafening results in the eventual deterioration of the acoustic features that distinguish song, suggesting that the AFP is important for both song learning and maintenance.

2.2.2 HVC as a sequence-generating circuit

Recordings from HVC demonstrate that excitatory neurons that project to RA (HVC(RA) neurons) exhibit sparse, song-locked activity, firing in 1-3 bursts of 4-6 spikes per motif

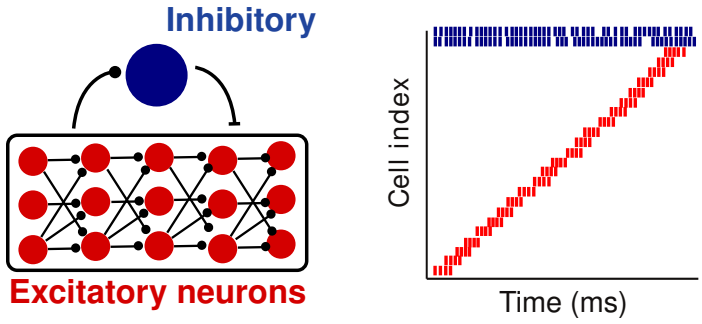

Figure 2.1: **Putative structure and dynamics of HVC** (Left) Putative structure of HVC. The red excitatory neurons represent the HVC(RA) cells, which connect to one another in feedforward manner. These neurons are mediated by recurrent inhibition implemented by inhibitory interneurons. (Right) Putative dynamics of HVC during song. Excitatory neurons fire in one burst while interneurons fire tonically.

[86]. Population HVC(RA) activity is consistent with a uniform distribution of firing times throughout song, leading to the belief that HVC acts as a timing signal for singing behavior [87]. Cooling of HVC but not RA results in a slowing of song tempo consistent with HVC acting as the locus of timing for song. The sparse sequential dynamics of HVC(RA) neurons are hypothesized to provide a timing scaffold for song that is independent of the spectral properties of song. Such a representation of time has been shown in certain models to simplify the credit assignment process by breaking down song learning into a moment by moment problem of driving the correct set of motor neurons [35, 88, 89]. RA also demonstrates song-locked firing, but its constituent neurons do not fire as sparsely as HVC(RA) neurons [90–92], and a major portion of its excitatory input comes from HVC [83].

The sparse song-locked dynamics of HVC(RA) neurons led to the hypothesis that HVC is fundamentally a sequence generating circuit [93]. Models have suggested these dynamics might arise via excitatory projections between HVC(RA) neurons, which form a feedforward network capable of generating propagating modes of excitation. Support for this hypothesis includes experiments that revealed that the depolarizations of HVC(RA) membrane potentials occur shortly before burst onset, suggesting these neurons receive most of their excitatory

input at once [68], in contrast to predictions made by "ramping" models in which HVC(RA) neurons integrated more temporally distributed input. Further, cooling HVC results in a slowing of song tempo consistent with increased synaptic transition time through a chain-like network whereas cooling RA does not have the same effect on timing [94]. Finally, a chain model is consistent with connectivity statistics measured by slice electrophysiology: HVC(RA) $\rightarrow$ HVC(RA) coupling is sparse but existent, while connectivity between HVC(RA) and interneurons in either direction is common (75% in each direction) [95]. Further, approximately 50% of HVC(RA) neuron pairs inhibit one another disynaptically, consistent with a model of widespread global inhibition [95]. While some models suggest interneuron connectivity with the HVC(RA) population is purely random [68], others have suggested this inhibition is structured, leading to enhanced precision in the firing times of HVC(RA) neurons [95, 96]. However, it is clear that excitatory cells carry considerable timing information in this circuit since altering the timing of gaps in inhibition on HVC(RA) neurons does not reliably alter their firing time [95].

2.2.3 Models of HVC chain dynamics

The earliest known model of chain dynamics is the "synfire chain" popularized by Abeles, who showed coincident arrival of inputs onto a postsynaptic target could produce a stable propagation in a chain-like network of excitatory neurons. In this model, dispersion of activity caused by membrane or synaptic transmission dynamics competes with the amplification of inputs via a nonlinear gain to create a steadily propagating mode [98]. Li and Greenside demonstrated this model could explain the propagation of a 2-6 spike burst of excitation within a 1-D chain. Models of sequence generation within HVC typically additionally include recurrent inhibition implemented through the interneuron population [68, 100, 101]. The inclusion of these inhibitory cells significantly alleviates the problem of fine tuning feedforward excitatory connections in these models, as blanket inhibition serves to regulate the growth of the propagating mode and enforces winner-take-all dynamics in which the network can only propagate one packet of excitation at a time [102]. Chang and Jin derived conditions under which a packet of spikes can propagate in a simplified conductance-coupled model of a 1D

inhibition-dominated network, where 1D here refers to the fact that neuron i is only coupled in an excitatory fashion to neuron $i+1$. Their model has the following dynamics:

$$\tau \frac{dV_j}{dt} = E_R + I - V_j(t) + I_s, \tag{2.1}$$

where V_j is the membrane potential of neuron j, I is a constant external input, and I_s is a synaptic current given by

$$I_s = \sum_0^\infty \left(-G_E^{j,\,s_n} V_j(t) + G_I \left(E_I - V_j(t) \right) \right) \tau \delta(t - t_n), \tag{2.2}$$

where $G_I > 0$ and $G_E > 0$ are the inhibitory and excitatory gains, $E_I < 0$ is the reversal potential the inhibitory inputs (E_E is assumed zero), τ is a time constant, and s_n is the index of the neuron spiking at time t_n. The neuron is assumed to spike when V_j reaches Θ, after which the neuron is reset to $V_j = V_r$.

Two conditions must be met: the first is that inhibition must be sufficiently strong, which can be expressed as:

$$\frac{G_I}{G_I + G_E} E_I \leq \Theta, \tag{2.3}$$

where G_I and G_E are the inhibitory and excitatory gains, E_I is the reversal potential of the inhibitory inputs (E_E is assumed zero), and Θ is the threshold voltage. If this condition is not satisfied, the number of spikes propagating through the network will grow indefinitely. The second condition is that propagating excitation must elicit at least one spike in the subsequent neuron while the inhibition is in a steady state condition, which is predicated on G_E being sufficiently large given G_i (exact conditions can be found in Chang and Jin).

A number of authors have also suggested that nonlinear spiking dynamics in excitatory neurons in the sequence, which leads to bursting, might contribute to the robustness of spiking behavior [68, 103]. This mechanism can decrease the jitter in timing of a neuron's activity by increasing the stereotypy of excitation a neuron delivers to postsynaptic targets when it fires. It was shown experimentally that HVC(RA) bursts occur on top of a long-lasting calcium depolarization, although the same neurons do not show a burst response to intracellular current injection [68].

2.2.4 Plasticity in feedforward sequential networks and models of HVC

Sequences of activity have been shown to organize under a number of supervised and unsupervised mechanisms [100, 101, 104–108]. Here, we review the most common plasticity mechanisms evoked to explain the formation and maintenance of chain-like networks that generate sequences of activity, including models of HVC. For unsupervised rules, spike timing-dependent plasticity (STDP) is the most common mechanism evoked. STDP in combination with other bounding mechanisms has been shown to organize sequences under external inputs with no average temporal ordering [100, 101, 104, 106] and also inputs with some average ordering [105, 108].

Spike timing-dependent plasticity

Most unsupervised rules that lead to sequential activity leverage some form of spiking-dependent plasticity, by which neurons that fire in sequence lead to potentiation of the synapse that connects them. This mechanism is a form of Hebbian learning [1] in which the precise timing of inputs matters. It was first discovered in cultured hippocampal neurons [5], and connected to sequential propagation of activity and sequence learning by Abbott and Blum, who showed that the synaptic modification rule,

$$\Delta w_{ij} = \int_0^t H(t') f_i(s(t+t')) f_j(s(t))\, dt\, dt', \tag{2.4}$$

could lead to entrainment of sequential input into sequences of activity. Here, $f_i(s(t))$ is the firing rate of neuron i when the sensory input is equal to $s(t)$ and $H(t')$ is the STDP kernel, which is typically chosen to be antisymmetric and positive for positive values of t'. Under this rule, the synapse $j \to i$ will grow if $f_i(s(t))$ tends to peak after $f_j(s(t))$. Resulting networks have a feedforward structure that reflects the ordering of sensory inputs.

Bounding STDP

In isolation, such STDP will permit synapses to grow in unbounded fashion [108], which leads to compression of the learned sequence and eventual collapse in the dynamics. To avoid this, STDP must either be switched off after some finite amount of time or be counterbalanced

with another mechanism. The literature has explored a wide array of mechanisms that could counterbalance STDP, including forms of firing rate homeostasis, heterosynaptic plasticity, bounds on the size of individual synapses, and synaptic pruning. Here, we discuss the relative advantages and disadvantages to these mechanisms.

The first mechanism by which STDP might be bounded is firing rate homeostasis. [6] first noted that neocortical neurons could scale their afferent synapses in an activity-dependent manner. In culture, chronic activity blockade resulted in an increase in miniature excitatory postsynaptic currents (mEPSCs), a measure of spontaneous neurotransmitter release thought to correlate with the strength of synapses. The scaling of the amplitude of these currents suggested the compensatory mechanism was postsynaptic, but forms of homeostatic plasticity have also been discovered at the neuromuscular junction that operate via presynaptic changes [109]. These mechanisms are thought to adjust all synapses afferent to a neuron in order to move it toward a target firing rate. In some systems this adjustment occurs in a multiplicative manner.

This simplest form of homeostatic plasticity may be written as

$$\Delta w_{ij}(t) = \alpha\, w_{ij}(x_{\text{target}} - x_i(t)), \tag{2.5}$$

where w_{ij} is the synapse size, x_{target} is the target firing rate of the neuron, α is a scaling factor, and synaptic modification occurs to all synapses w_{ij} in order to modify the postsynaptic rate $x_i(t)$. This rule, however, is generally insufficient to bound STDP if it is linear in pre- and postsynaptic firing rate, e.g.

$$\Delta w_{ij} = \int_0^t H(t')x_i(t+t')x_j(t)\,dt'. \tag{2.6}$$

To solve this, firing rate homeostasis can be made second order in the postsynaptic firing as in Oja's rule [110] and Pereira and Brunel. If the STDP rule itself saturates as a function of pre- or postsynaptic activity, then a homeostatic rule like Eq. 2.6 can be sufficient to bound synaptic growth.

STDP can also be bounded by heterosynaptic plasticity or multiplicative scaling of synapses. These rules typically normalize the total synaptic strength onto a postsynaptic

neuron in an L1 or L2 fashion,

$$w_{ij} \leftarrow w_{ij} / \left(\sum_k w_{ik}^p \right)^{1/p}, \tag{2.7}$$

where $p \in \{1, 2\}$ for an L1 or L2 norm, respectively. E proved that under slow weight update and stationary input, synaptic normalization via an L2 norm is equivalent to adding a term proportional to $-w_{ij}x_i^2$ to an additive learning rule. If input is not regular and stationary, synaptic normalization differs from Oja's second order term in that it does not require activity to take effect.

A number of other mechanisms have been used in models to bound STDP, including bounding of single synapses, in which a hard bound is installed to prevent growth in synaptic strength beyond a chosen value [101, 106]. This mechanism typically is paired with synaptic scaling, a bound on total weight growth [101], or synaptic pruning, as in Tupikov and Jin, where a limit to the number of outgoing or incoming synapses from a neuron is imposed.

Models of sequence organization not based in STDP

Error-based plasticity rules are also capable of generating sequences of activity. Backpropagation and FORCE learning [111] have both demonstrated the capacity to learn target sequences of activity [89, 112]. The change in synaptic connectivity required to induce sequential activity can be very small compared to the typical regime induced by STDP-like rules. Sequential, but periodic, activity has also been shown to arise in networks with extremely sparse connectivity regulated solely by synaptic scaling [113].

2.3 Unsupervised restoration of a complex learned after large-scale neuronal perturbation

We now describe the main results of this chapter. All experiments were performed by Carlos Lois's lab at the California Institute of Technology. Figures are reproduced, with minor modifications, from Wang et al..

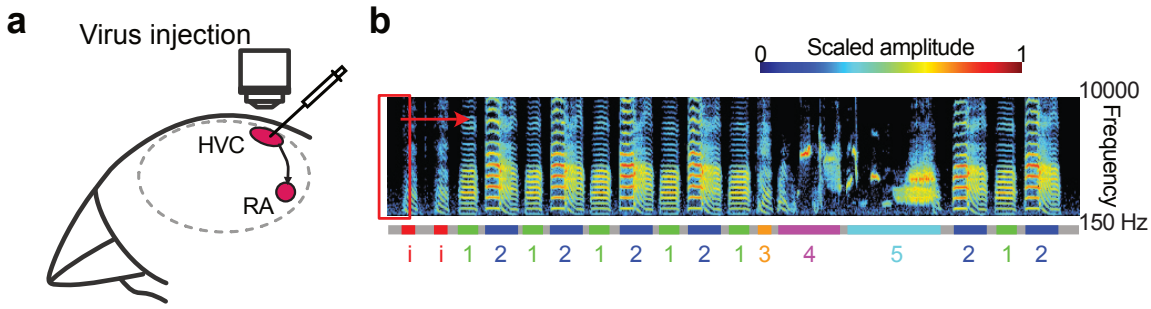

Figure 2.2: **Schematic of viral perturbation and spectrogram of unperturbed song** **(a)** Schematic drawing illustrating the visual guided virus delivery into HVC (see Methods Wang et al.). **(b)** Spectrogram of a motif from an unperturbed animal. The syllables are indicated by different colors and numbers in the spectrogram.

2.3.1 Experimental results

Twin viral perturbations of HVC were performed in adult male zebra finches via a bilaterally injected lentiviral vector (Fig. 2.2a). The vector has previously been shown to selectively infect HVC(RA) projection neurons [114, 115]. The first set of perturbations virally expressed a bacterial voltage-gated Na^+ channel known as NaChBac [116], which is known to be active at membrane voltage levels at which vertebrate Na+ channels are inactive and induces long-lasting depolarizations (>1 s) in neurons once activated [117–119]. This is substantially longer than the typical depolarization timescale in vertebrates, which is often on the order of milliseconds. In the second set of perturbations, the same lentivirus was used to deliver tetanus toxin (TeNT), an enzyme which cleaves synaptobrevin, a protein required for neurotransmitter release. While the toxin leaves electrical activity of the neuron unaltered, it prevents communication with downstream targets. Paired whole-cell patch clamp recordings were used to confirm that TeNT perturbation prevented HVC projection neurons from driving their interneuron counterparts (see Wang et al.).

Both NaChBac and TeNT viral perturbations induced song degradation in similar manners

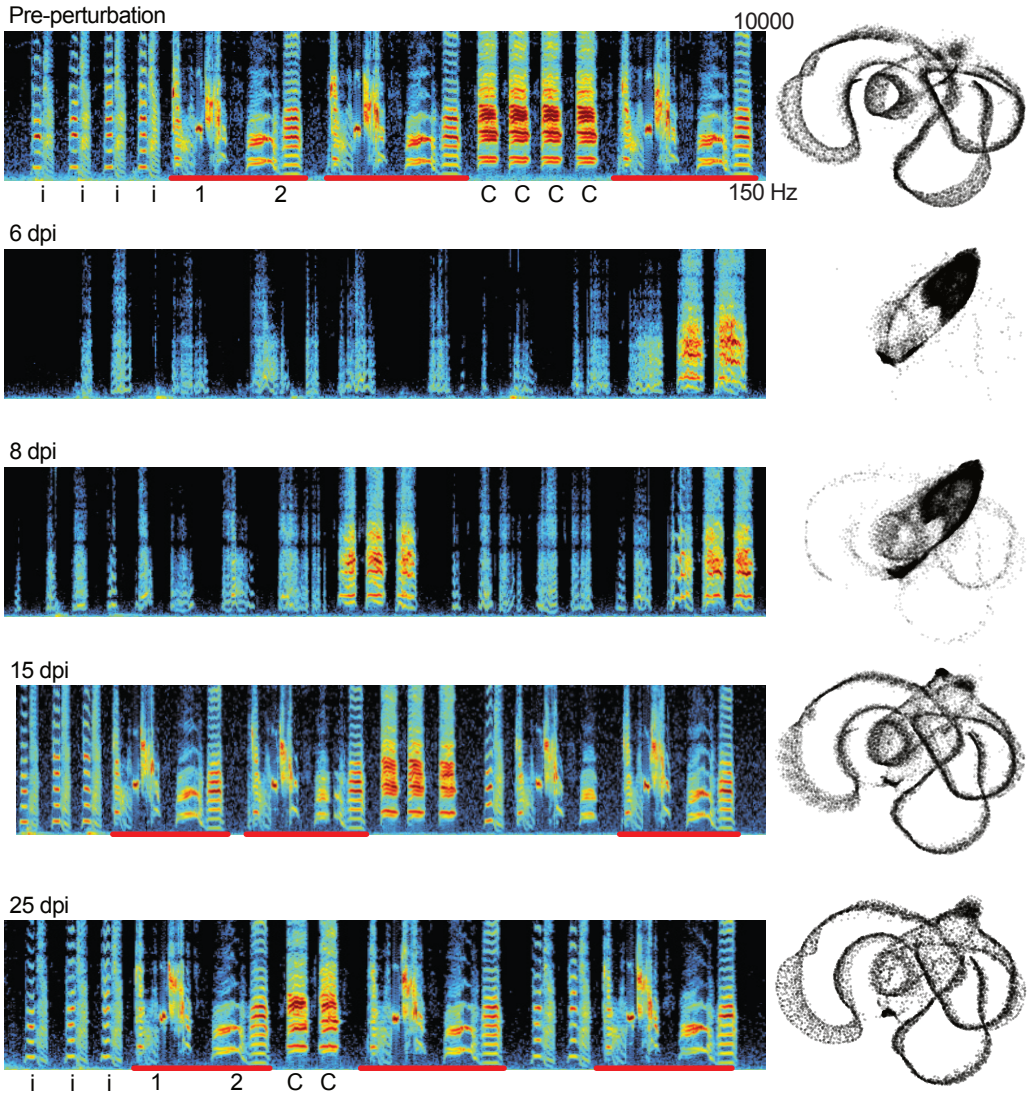

Figure 2.3: **Song degradation and recovery after selective large-scale perturbation of excitatory neurons** Example spectrograms and UMAP visualizations of the song of a bird injected with LV-NaChBac.

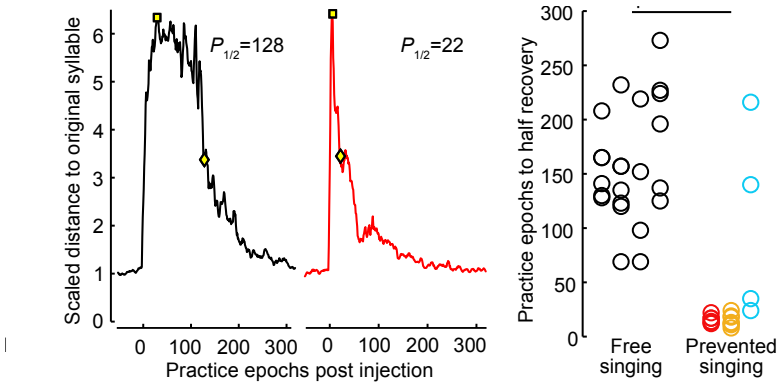

Figure 2.4: **Syllable recovery vs number of practice epochs with or without song prevention** (Left) Example plots of syllable recovery vs number of practice epochs with or without prevention. (Right) Group data of the number of practice epochs to reach half recovery of each syllable, each column is one bird.

(Fig. 2.3). NaChBac-infected animals showed no changes to their songs for the first 2-3 days post infection (dpi), after which the temporal and acoustic aspects of song fully degraded. TeNT-infected animals showed signs of degradation at 1 dpi. Following both perturbations, songs consisted of shorter, weaker, and noisier syllables, where syllables were defined as any continuous song element. Remarkably, by 20-25 dpi, the song features of song were fully recovered. TeNT-infected neurons were still muted at >25 dpi, even after the song had recovered.

To assess whether recovery of song depended on practice, TeNT-infected animals were prevented from singing during a two-week period following infection, after which they were allowed to sing freely. After singing tens of renditions in a single day, animals produced song that was as recovered as animals that had sung thousands of songs since perturbation. An analysis of recovery as characterized by acoustic features of song (see Wang et al., Methods) revealed individual syllables recovered with less practice when prevented from singing (Fig. 2.4). This result suggested that offline mechanisms may be key to the recovery of regular HVC activity.

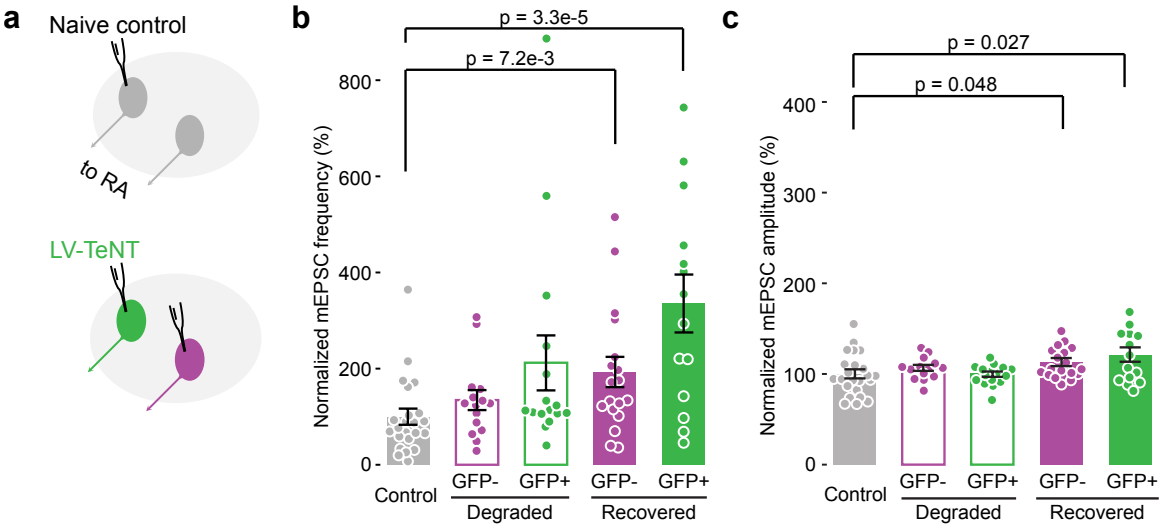

Figure 2.5: **Measurements of mEPSCs of HVC(RA) neurons reveal increased frequency of spontaneous events (a)** Schematic illustrating whole-cell recordings from HVC(RA) neurons in birds injected with LV-TeNT or naive controls; **(b-c)** Group data of mEPSCs. "Degraded" indicates recordings at 5 dpi when the song was degraded. "Recovered" indicates recordings at 25 dpi, after the song had fully recovered.

To understand the mechanisms by which HVC recovered from TeNT infection, whole-cell recordings of unmanipulated HVC(RA) neurons were performed. These neurons showed no change in their intrinsic excitability or inhibitory synaptic inputs. However, the frequency of mEPSC events increased substantially over the contralateral control, suggesting excitatory inputs to these neurons increased in number, presynaptic release probability, or strength (Fig. 2.5). These effects were only observed in ipsilateral hemisphere upon unilateral injection, suggesting the mechanism of recovery was local.

2.3.2 Modeling results

Using several models, we explored how synaptic plasticity mechanisms could contribute to unsupervised recovery of network activity after perturbation and to account for the observed synaptic changes. Song recovery in animals that were prevented from singing suggests that plasticity initially occurs largely locally within HVC rather than through behavioral feedback from practice. We sought to determine which mechanisms are consistent with restoration of the sequence as well as with our physiological findings about synaptic strength changes. We modeled HVC as an excitatory-inhibitory (E-I) network, with HVC(RA) neurons connected to each other in a feedforward, polychronous chain and recurrently connected to interneurons (Fig. 2.6a, top) [68, 100, 120]. We then inactivated varying fractions of the HVC(RA) population (Fig. 2.6a, bottom) to mimic the silencing of neurons caused by TeNT expression. To explore mechanisms that may enable recovery, we first implemented spike-timing dependent plasticity (STDP) between excitatory neurons (E→E), which is integral to many models of sequence self-organization [100, 101, 105, 121], and a form of downward firing rate homeostasis, which scaled down an HVC(RA) neuron's afferents if its activity exceeded a threshold (see Methods, Eq. 1). STDP alone is unlikely to rebuild the feedforward sequence once broken, as this learning rule requires activity in both pre- and postsynaptic neurons to strengthen synapses. It is possible that the addition of noise can drive random activity that may allow rebuilding of the sequence (see Fig 2.13). We found that STDP can indeed restore sequential activity with the help of independent random inputs across the network, but that this restoration only occurs for weak perturbations (Fig 2.9a).

Further, this mechanism does not produce the increase in excitatory connectivity relative to the level before perturbation in HVC(RA) neurons observed by our electrophysiological measurements (Fig 2.9b, Fig 2.10a-b).

As STDP alone was insufficient to restore the sequential dynamics in our model of HVC, we next considered whether cell-autonomous homeostatic mechanisms may enable recovery of the network, as was previously proposed to account for recovery from lesioning in Nif, a nucleus upstream of HVC [122]. While homeostasis may act either on intrinsic excitability or synaptic inputs, our recordings showed that unmanipulated HVC(RA) neurons did not show changes in intrinsic excitability post-perturbation (see Wang et al., Extended Data Fig. 7), but rather in their synaptic inputs (Fig 2.5b). Thus, we added a cell-autonomous homeostatic rule into our model, based on scaling excitatory synaptic inputs to individual excitatory neurons to maintain their firing rates (Fig 2.6b). The implementation of this rule is consistent with reports that have found activity-dependent synaptic homeostasis in other circuits with sequential dynamics [123]. We found that adding this form of cell-autonomous synaptic homeostasis in excitatory neurons reliably restored sequential network activity, and that individual cells fired at close to their original times in the sequence (Fig. 2.7a and Fig. 2.8b). In addition, implementing STDP between excitatory and inhibitory neurons (E→I) enabled the rebound of input to inhibitory neurons observed in our experiments (Fig. 2.8a, blue line). However, in models employing STDP alone or both STDP and cell-autonomous synaptic homeostasis, the recovered HVC outputs remained reduced in proportion to the percentage of neurons inactivated, weakening accordingly the drive to downstream regions (Fig. 2.9b, d). Additionally, these models did not reproduce the relative increase in excitatory synaptic input to unmanipulated excitatory neurons as revealed by our experiments (Fig 2.8a, Fig. 2.9b, d).

A potential mechanism by which HVC might restore the strength of its output to downstream targets is the recruitment of neurons that initially do not participate in the sequential dynamics, defined here as "silent" neurons (Fig. 2.6c). Multiple experiments have indicated that a large fraction of HVC(RA) neurons do not fire during song, suggesting a possible redundancy [86, 93, 124]. We hypothesized that the presence of such silent neurons–assumed to be HVC(RA) neurons connected within the network, but whose inputs

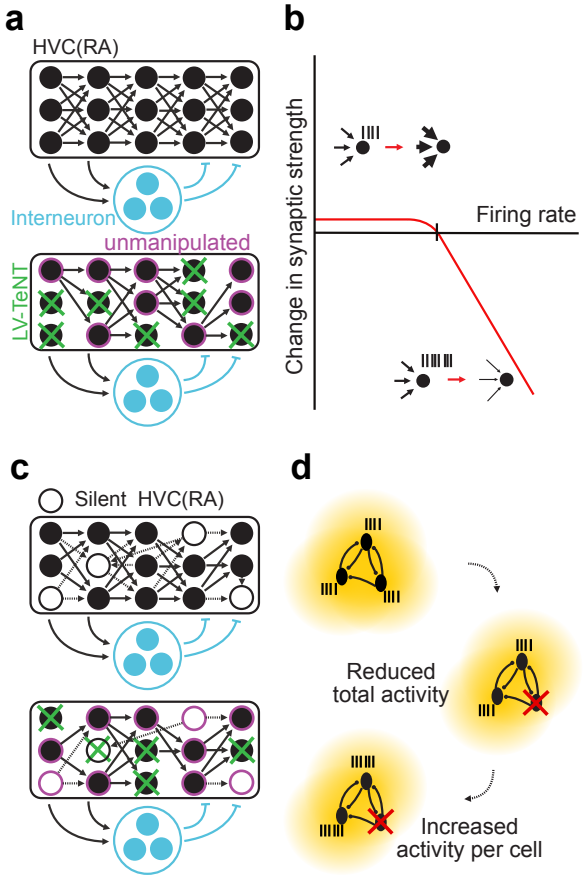

Figure 2.6: **Schematics of sequence network manipulations and plasticity rules (a)** Schematic diagram of the neuronal organization in the model. **(b)** Schematic illustrating the single-cell homeostatic plasticity rule implemented. **(c)** Schematic diagram illustrating that initially inactive neurons were recruited into the network. **(d)** Schematic diagram illustrating that a population-level homeostatic plasticity, governing the summed firing activity of all neurons, was also implemented in the model.

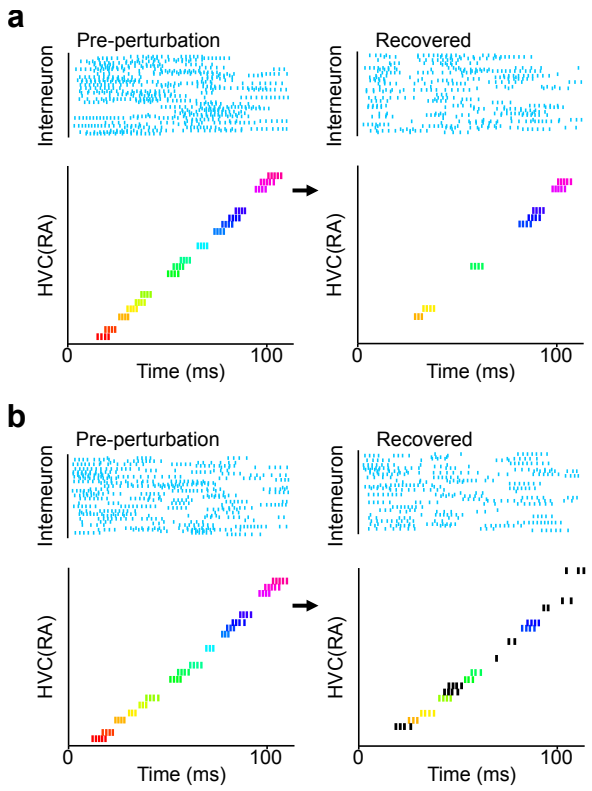

Figure 2.7: **Rasters of activity in pre-perturbation and recovered sequence networks** **(a)** Spike raster plots showing the sequential dynamics generated by HVC neurons before perturbation and after recovery (with only single-cell homeostatic plasticity implemented). **(b)** Spike raster plots showing the sequential dynamics generated by HVC cells before perturbation and after recovery.

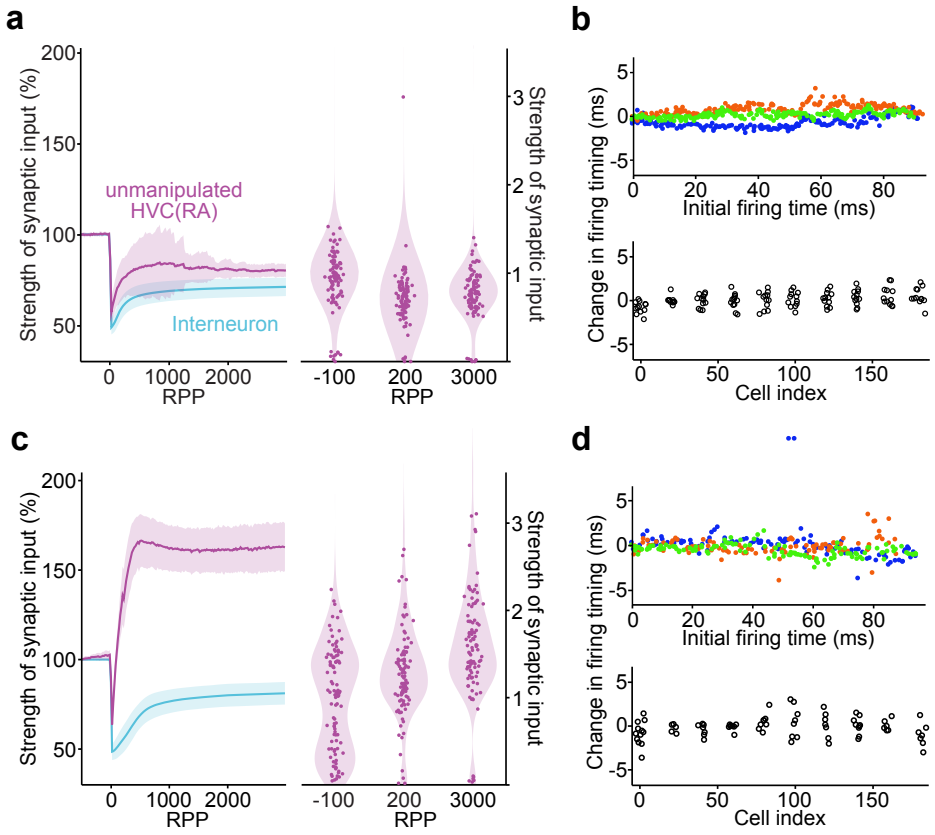

Figure 2.8: **Comparison of synaptic strengths and timing jitter between networks recovering via synaptic scaling and a population homeostatic rule (a)** (Left) Plot of the normalized total excitatory synaptic input per neuron against renditions. RPP, rendition(s) post perturbation. (Right) Scatter plot of the normalized total excitatory synaptic input received by each HVC(RA) neuron at three time points; **(b)** Plots of the normalized total firing activity of all functioning neurons; **(c)** (Left) Plot of the normalized total excitatory synaptic input per neuron against number of renditions. Dashed lines are adapted from panel g for comparison. (Right) Scatter plot of the normalized total excitatory synaptic input received by each HVC(RA) neuron at three time points; **(d)** Plots of the normalized total firing activity of all functioning neurons.

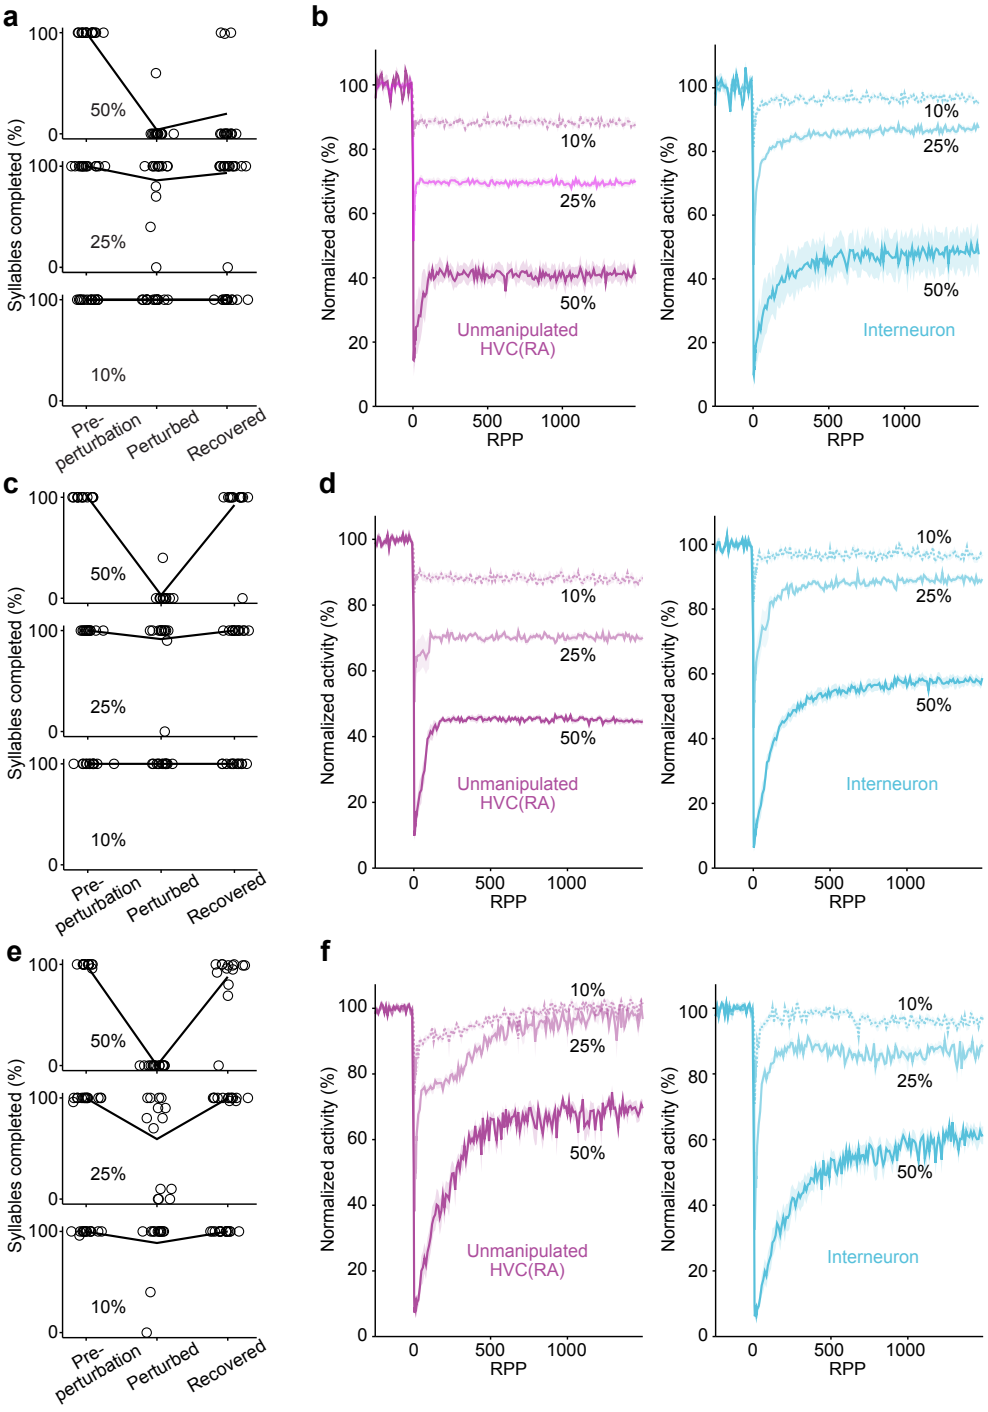

Figure 2.9: **Comparison of HVC network recovery under STDP, synaptic scaling, and a population homeostatic rule**

Figure 2.9: **Comparison of HVC network recovery under STDP, synaptic scaling, and a population homeostatic rule (a-b)** are generated from networks with STDP and a downward, single-cell, firing rate homeostasis rule implemented; **(c-d)** with STDP and 2-sided, single-cell, firing rate homeostasis; **(e-f)** with silent cells and population-level homeostasis added to rules in (a-b); **(a,c,e)** Percentage of modeled syllables that successfully completed following >80% of activations in three time windows (before perturbation, perturbed and recovered) in response to different degrees of perturbation (10, 25, and 50%). Each point represents a network. Lines represent averages over all networks (n = 15); **(b,d,f)** The normalized total firing activity of all functional HVC(RA) (left) or interneurons (right) plotted against renditions. Solid curves show the average of multiple networks (n = 15) and shading represents s.e.m.. RPP, renditions post perturbation.

are subthreshold during song–might provide additional resilience by allowing the sequential dynamics to be partially carried by newly recruited HVC neurons when active constituents of the network fail. While such shifts in participation could be due to the loss of inhibition onto silent cells following the loss of excitatory neurons, this picture of recovery is inconsistent with our experiments in that it does not require the observed increase in the excitatory inputs onto HVC(RA) neurons. We therefore hypothesized that silent neurons may be recruited into the sequence through a form of homeostasis that is sensitive to the activity of the population, for which there is an emerging body of support [125–127]. We modeled one potential mechanism by which such population homeostasis can be achieved: synaptic scaling based on the activity-dependent release of a secreted factor (Fig. 2.6d) such as BDNF and TNFα, both of which have been shown to regulate local network activity in a non-cell-autonomous manner [128–131]. The recruitment of previously silent neurons by population homeostasis enabled the most complete recovery. The total activity of the network recovered to a greater degree than models without a form of population homeostasis (Fig 2.9e-f), and the dynamics of the sequence in terms of numbers of participating neurons and their temporal resolution most closely resembles the state before perturbation (Fig. 2.7b,

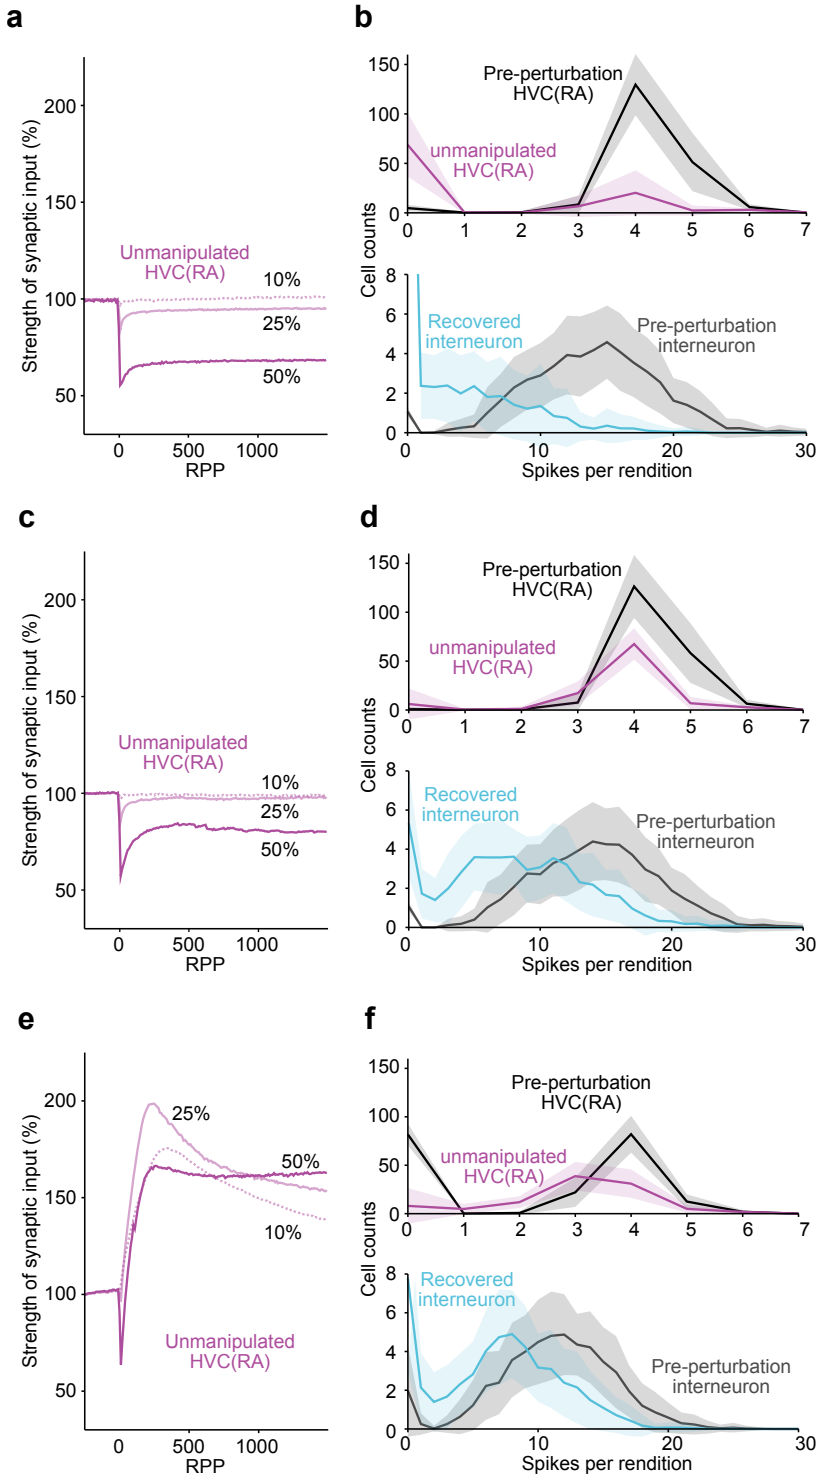

Figure 2.10: **Comparison of evolution of bulk synaptic strengths under STDP, synaptic scaling, and a population homeostatic rule**

Figure 2.10: **Comparison of evolution of bulk synaptic strengths under STDP, synaptic scaling, and a population homeostatic rule (a-b)** are generated from networks with STDP and a downward, single-cell, firing rate homeostasis rule implemented; **(c-d)** are generated from networks with STDP and 2-sided, single-cell, firing rate homeostasis; **(e-f)** are generated from networks with silent cells and population-level homeostasis added to rules in (a-b). **(a,c,e)** Normalized total excitatory synaptic input into each unmanipulated HVC(RA) neuron after perturbations. Note that with population-level plasticity and silent neurons, an increase in E→E weight over its pre-perturbation value was introduced; **(b,d,f)** Distributions of single-cell firing rates per rendition for HVC(RA) (top) and interneurons (bottom) before and following 50% perturbation. Pre-perturbation distributions in black or gray. Shading represents standard error.

Fig. 2.8d). Furthermore, consistent with our experiments, E→E synaptic inputs increased by 100% (Fig. 2.8c). This increase was primarily accounted for by strengthened inputs to silent neurons (Fig. 2.12). Finally, the recruitment of silent neurons (which may be either mature cells already present or new neurons produced during adult neurogenesis [132]) may initially add noise to song production, consistent with the observation of an increase in the song's variability, peaking at 10 dpi, approximately 4-6 days after the acoustic features of the songs were maximally degraded. Despite the addition of new neurons into the sequence dynamics, the timing of the originally active neurons was preserved after recovery, and timing jitter did not increase beyond that displayed by our model without silent cells (Fig. 2.8d). This is consistent with the reemergence of timing of within-syllable acoustic structure immediately after song prevention. The participation of new cells in the sequence could also contribute to noise in the acoustic structure and account for the deviations in acoustic structure observed immediately after song prevention, before the bird is able to use auditory feedback to tune the new HVC-RA synapses.

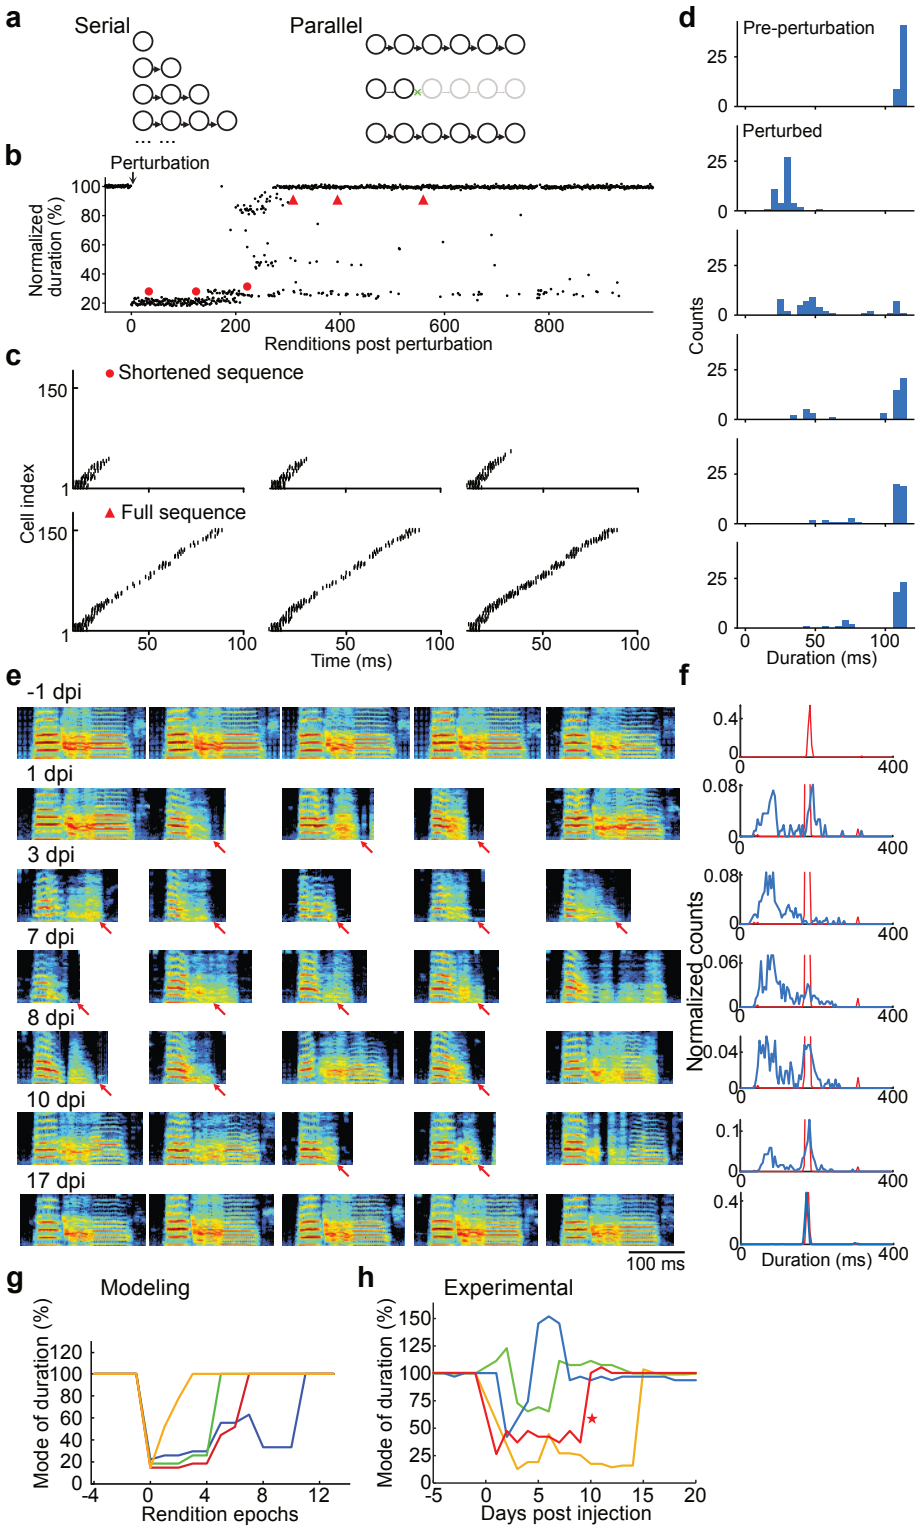

Figure 2.11: **Saltatory recovery of syllable duration**

Figure 2.11: **Saltatory recovery of syllable duration**

(a) Schematic diagrams showing two types of potential circuit recovery mechanisms. (Left) Serial – the recovery of sequential firing requires building of feedforward synaptic chains step by step from the breaking point, such that the duration should regrow continuously. (Right) Parallel – all links in a sequence are repaired simultaneously, so that the recovery of the full sequence can be abrupt, when the broken links are fixed; **(b)** Plot of the normalized duration of modeled sequential firing against number of renditions. Note that the recovery of the full syllable duration was not continuous, but instead it occurred by a sudden leap from a shortened state; **(c)** Example raster plots showing the sequential spiking of modeled HVC(RA) neurons, picked from multiple time points indicated by red dots (shortened sequence) or triangles (full sequence) in panel b; **(d)** Probability density distributions of the durations of the modeled syllables at different times (before perturbation, during the degraded period, during the recovery period, and after full recovery), ordered chronologically. Note the bimodal distribution of the duration of a single syllable during the recovery phase; **(e)** Example spectrograms picked from the nearest k neighbors of one syllable at multiple days before and post injection of LV-TeNT. Red arrows mark the shortened/truncated syllables found between 1 to 10 dpi; **(f)** Probability density distributions of the duration of the k neighboring syllables, ordered so that each row of panels in e and f is from the same day. Note the bimodal distribution of durations for a single syllable found during the period of song degradation and recovery; **(g)** Plots of the mode of the duration of the modeled sequences against rendition epochs, showing that the recovery of sequence duration occurs in a saltatory, rather than in a continuous, manner; **(h)** Plots of the mode duration of actual song syllables, which shows a saltatory recovery similar to that predicted by the model. The red curve marked with a star is made from the syllable shown in panels e and f.

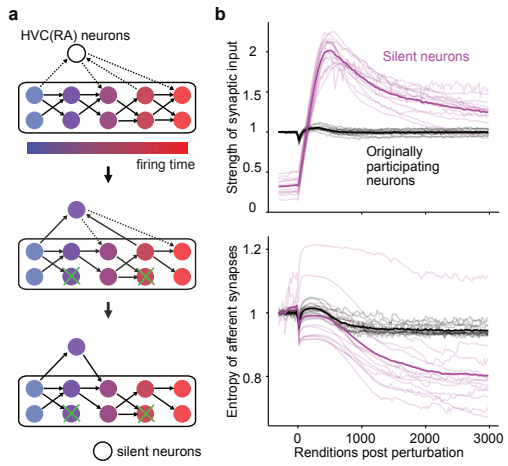

Figure 2.12: **Strengthening and refinement of inputs to silent neurons (a)** Schematic showing perturbation and recovery in a model with both silent neurons and population-level homeostasis. Perturbing neurons lowers total network activity. In response to lowered activity, the population-level homeostatic rule drives an increase in excitatory synaptic strength, which strengthens inputs to all neurons and potentially drives silent neurons to fire. Once active, STDP further refines and selectively strengthens the connections between originally participating neurons and silent neurons. **(b)** (Top) The normalized total excitatory synaptic input received by individual unperturbed neurons plotted against renditions. Color indicates neurons that were active or silent pre-perturbation. Silent neurons undergo a much larger increase in synaptic strength because of their less-structured, weaker synaptic connections. The eventual decrease in synaptic strength is due to the interplay between STDP and the downward firing rate homeostasis. However, note that total average synaptic strength remains higher than pre-perturbation values. (Bottom) The normalized entropy of afferent synapses of unperturbed neurons plotted against renditions. Each light-colored curve represents a single network realization. Solid curves represent the average of each group. Decreased entropy in the originally silent population is due to the strengthening of local feedforward synapses and the weakening of unstructured recurrent synapses via STDP and downward firing rate homeostasis.

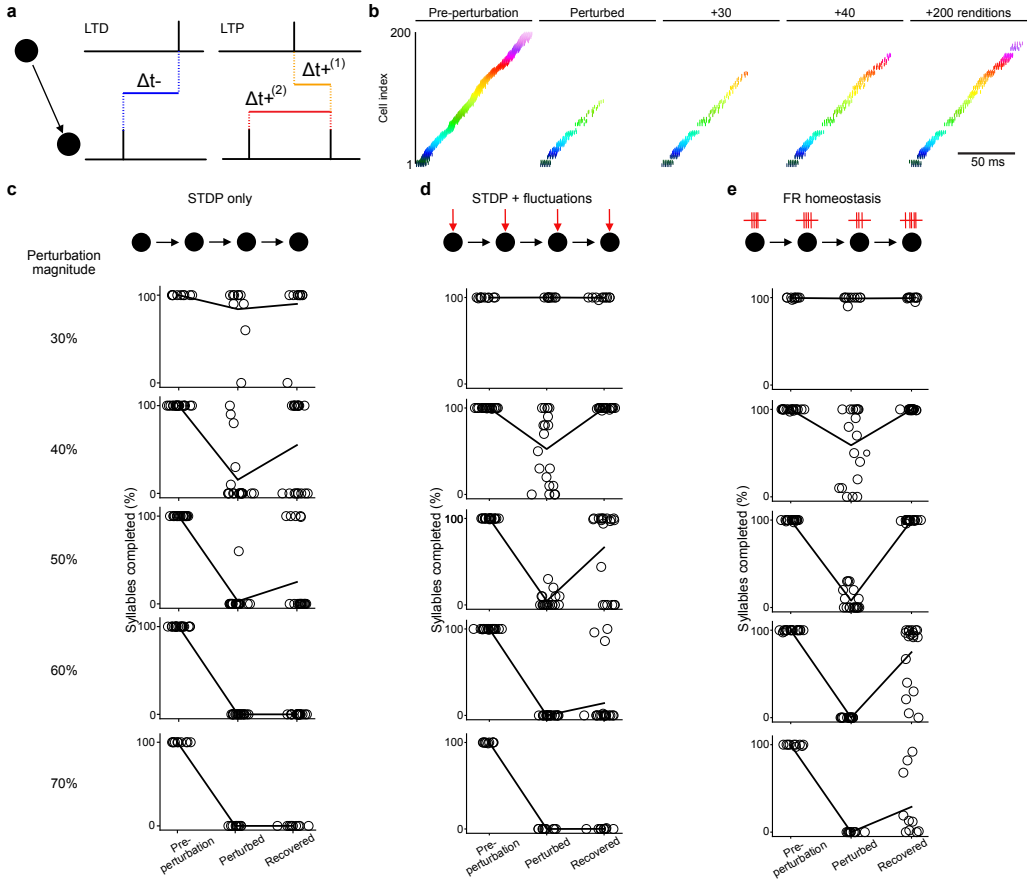

Figure 2.13: **Recovery of modeled sequences with STDP (a)** Schematic showing triplet STDP implemented both in E→E and E→I synapses (see Methods). Potentiation is mediated by triplets of spikes; depression is dictated by a pairwise rule. See Methods for details; **(b)** Spike raster plots showing the sequential dynamics generated by HVC neurons before and after perturbation with only STDP and downward firing rate homeostasis implemented (see Eq. 2 in Methods). Colors represent the firing timing of each cell before perturbation. Note that the sequence regenerates serially and all cells recapture their original firing timing; **(c-e)** Percentage of syllables completed prior to (pre-perturbation), just after (perturbed), and 3000 renditions after (recovered) perturbation for networks with (c) STDP alone, (d) STDP plus Poisson inputs, (e) or firing rate homeostatic plasticity for 30% to 70% perturbations in increments of 10%. Note that networks with firing rate homeostatic mechanisms possessed larger recoverable regimes.

Saltatory recovery of syllable duration

We can further use our model to explore predictions for the recovery process. The core of the model is that the sequence is reliably restored through a parallel homeostatic process, in contrast to the serial regrowth one would expect in a timing-dependent process as demonstrated in previous modeling of HVC sequence organization [100] (Fig. 2.11a and Fig. 2.13a-b). A prediction of the model is that the sequence should recover in an abrupt fashion: all links are repaired simultaneously, so that by the time any broken link is restored, all other links tend to have recovered (Fig. 2.11b-d). STDP alone can also lead to abrupt recoveries when HVC neurons receive parallel noisy input, but as noted above, for experimentally constrained noise levels, these recoveries typically occur for weak perturbations (Fig. 2.13c,d) and are less common than abrupt recoveries under a homeostatic mechanism. Thus, we find that homeostasis is a generally more robust mechanism of recovery. In our experiments, we found many examples of such saltatory recovery in the song, whereby syllables show a bimodal temporal distribution (Fig. 2.11e-h and see Wang et al. Supplementary Fig. 7), with syllables rapidly alternating between two durations from rendition to rendition until eventually they permanently recover their original duration. Syllables with these bimodal temporal distributions are only present in renditions that include song segments that bear acoustic resemblance to the original syllables. This suggests that a "skeleton" of redundant feedforward connections persists, allowing the sequence to reemerge when all links are repaired. Thus, the skeleton can act as a scaffold upon which the homeostatic plasticity can rebuild the dynamics in an unsupervised manner. As this redundancy can exist very 'locally' in the chain (i.e. only between adjacent links), it is still consistent with an extremely sparse network [86, 93], and may be a key contributor to resilience in HVC.

2.3.3 Methods

Leaky Integrate and Fire Neurons

The membrane potentials V_j of neurons in all networks were modeled using leaky integrate-and-fire dynamics:

$$C_m \frac{dV_j}{\mathrm{dt}} = g^{(l)}\left(E^{(l)} - V_j\right) + g^{(i)}\left(E^{(i)} - V_j\right) + g^{(e)}\left(E^{(e)} - V_j\right) + I_{\text{ext}} \tag{2.8}$$

where $C_m = 1\ F/cm^2$ is the membrane capacitance, $g^{(l)}$, $E^{(l)}$, $g^{(e)}$, $E^{(e)}$, $g^{(i)}$, $E^{(i)}$, the leak, excitatory, and inhibitory conductances and reversal potentials, and $I_{\text{ext}} \sim\ (0, \sigma = 0.1\, nA)$, a white noise current input. For excitatory neurons, $g^{(l)} = 0.25\ mS/cm^2$ and $E^{(l)} = -70\ mV$; for inhibitory neurons, $g^{(l)} = 0.4\ mS/cm^2$ and $E^{(l)} = -53\ mV$. For all neurons, $E^{(e)} = 0$ and $E^{(i)} = -90\ mV$. When V_j reaches threshold, $V_{\text{th}} = -43\ mV$, the neuron spikes, and the voltage is reset to V_{re} ($-65\ mV$ for excitatory, $-53\ mV$ for inhibitory) after a refractory period $t_r = 1\ ms$. The excitatory and inhibitory conductances are driven by incoming spike trains, represented by delta functions at times ${t_k}^{(s)}$, where s indexes all spike times for an upstream neuron k, filtered with a time constant, $\tau = 4\ ms$:

$$\tau \frac{d{g_j}^{(x)}}{dt} = -{g_j}^{(x)} + w_{kj}\delta\left(t - {t_k}^{(s)}\right). \tag{2.9}$$

In all modeling, we assume that neurons are connected with a certain synaptic strength, which does not differentiate between the number and strength of individual contact sites between a pair of neurons.

Network Architecture

We explored two different network architectures and the resulting dynamics. In the first network, we assumed all cells participate in the sequential dynamics. In the second network, we assumed that only a portion of the excitatory cells initially participate in the dynamics [68, 86]. To implement this, we assumed each neuron is silent with probability $p_s = 0.4$.

HVC Network

Following recent work, we modeled HVC as a feedforward, polychronous network [106, 120]. The network is composed of 200 excitatory (E) and 50 inhibitory (I) neurons. In order to define the approximate difference in firing times between any neuron pair, we assign a coordinate c_i to each E cell, where c_i is the i^{th} sum of random variables uniformly distributed over [0, 1]. Pairs of active excitatory cells are initially connected with a fixed weight

$W = 4e{-}5/\left(1 - \mathrm{p_s}\right)$ if the difference in their indices c is greater than zero and less than $c^* = 10$, ensuring feedforward propagation. Silent cells are assumed to be connected within the network but with much lower probability, with no sensitivity to sequence order, and with cell-by-cell heterogeneity. Specifically, if one of a pair of neurons is silent, the cells are connected with probability 0.8 with a weight w_{ij} drawn from an exponential distribution with a cell-specific mean $a_i W/2$, where $a_i$is uniformly distributed on $[0,\ 1]$. To implement recurrent inhibition observed in HVC [133], inhibitory neurons receive connections from excitatory cells with probability $p_{e,i} = 0.075$ (0.125 for networks with silent cells) and weight $w_{e,i} = 3.5\text{e-}4$ ($3.5e{-}5$ if the presynaptic cell is silent) and vice versa with probability $p_{i,e} = 0.5$ and weight $w_{i,e} = 4e{-}5$.

Following observations that axonal delays between HVC(RA) projectors are relatively long (1 - 7.5 ms) [120] and that HVC(RA) projectors typically synapse onto inhibitory interneurons close to their soma and other excitatory cells far from their soma [134], we implemented axonal delays in our model that reflected longer E-E axonal delays and relatively shorter E-I and I-E delays. The delays between pairs of active E cells were

$$d_{ij} = 6\frac{|c_i - c_j|}{c^*}\,\mathrm{ms} \tag{2.10}$$

if $|c_i - c_j| \ < c^*$.

When one of the pair was inactive or $|c_i - c_j| \ \geq \ c^*$, the delay was chosen randomly from a uniform distribution on $[0,\ 6]$ ms. This leads to E→E axonal delays that were 3 ms on average and roughly uniformly distributed. E→I and I→E delays were uniformly set to 0.5 ms. We found that a comparatively fast inhibitory pathway stabilized sequence dynamics by enabling inhibition to respond rapidly to changes in excitation.

Plasticity Rules

We then allowed networks to evolve under both firing rate homeostasis and spiking timing-dependent plasticity, which are evidenced by the synaptic reorganization among HVC neurons shown by our patch-clamp results [123, 135–137]. All synapses subject to plasticity were given a lower bound $w_{\min} = 1e{-}8$. After each trial, synaptic strength w_{ij} evolved according

to

$$\Delta w_{\mathrm{ij}} = \beta_f {\Delta_{\mathrm{ij}}}^f + \beta_{\mathrm{pop}} {\Delta_{\mathrm{ij}}}^{\mathrm{pop}} + \beta_{\mathrm{STDP}} {\Delta_{\mathrm{ij}}}^{\mathrm{STDP}} \tag{2.11}$$

where weights are updated due to single-cell firing rate homeostasis (Δ^f), local activity homeostasis (Δ^{pop}), and STDP (Δ^{STDP}), respectively, and β_f=0.025, $\beta_{\mathrm{pop}} = 0.01$, and $\beta_{\mathrm{STDP}} = 1.5e{-}4$. β_f was chosen such that firing rate homeostasis could bound potentiation due to STDP. β_{pop} was chosen to be small so that recruitment of formerly silent neurons would occur slowly. All E→E synaptic strengths had upper bound ${w_{e,e}}^{\max} = 1e{-}3$, all E→I ${w_{e,i}}^{\max} = 2w_{e,i}$. For simplicity, we did not introduce plasticity in the inhibitory synaptic inputs to excitatory neurons. We made this simplification because in our electrophysiological data showing the TeNT perturbation's impact to HVC's synaptic structure, the average strength of these synapses did not change.

Firing rate homeostasis:

Single-cell firing rate homeostasis moves a neuron's firing rate toward a set point ${r_j}^{(0)}$ according to:

$${\Delta_{\mathrm{ij}}}^f = w_{\mathrm{ij}} \left(a - \Theta \left(r_j - {r_j}^{(0)} \right) \left(r_j - {r_j}^{(0)} \right)^2 \right) \tag{2.12}$$

where r_j is the average firing rate of neuron j in that trial (refer to Fig. 5), Θ is the Heaviside function, and $a = 0.1$ is a small constant that mediates synaptic scaling when ${r_j}^{(0)} > r_j$. In our model with STDP and downward firing rate homeostasis and our model with local population homeostasis, we set $a = 0$.

Local population activity homeostasis:

Taking inspiration from literature that has shown homeostasis may operate on a network scale, we included in our final model a form of homeostasis that permits individual neurons to monitor and respond to the activity of their neighbors. We implemented here one potential mechanism by which such local population activity homeostasis might be achieved, based on the TNFα pathway [131]. We assume each E neuron secretes a chemical factor which diffuses locally in space, the concentration of which is proportional to the neuron's own activity level. All E neurons are assumed to monitor the local concentration of this factor and adjust their incoming excitatory synapses in order to maintain a target concentration. To implement this

form of local population homeostasis, each excitatory neuron was first assigned a location by uniformly sampling the space within a unit sphere. The local factor concentration that an excitatory neuron senses is then

$$m_j = \frac{1}{\sqrt{2\pi\sigma^2}} \sum_i r_i \exp\left[-\frac{1}{2}\frac{(x_j - x_i)^2}{\sigma^2}\right], \tag{2.13}$$

where x_j is the vector representing the location of neuron j, and σ is a parameter controlling the spatial extent of each neuron's diffuse release ($\sigma^2 = 0.03$). The corresponding update to local population homeostasis is

$$\Delta_{\mathrm{ij}}{}^{\mathrm{pop}} = w_i j \frac{e^{\gamma\left({m_j}^{(0)} - m_j\right)} - 1}{e^{\gamma\left({m_j}^{(0)} - m_j\right)} + 1}, \tag{2.14}$$

where ${m_j}^{(0)}$ is the local concentration setpoint of neuron j and $\gamma = 0.1$ dictates the strength of the local population homeostasis near the setpoint. Local concentration setpoints were chosen by computing the average local concentration over 100 activations of the network prior to the perturbation.

Hebbian plasticity:

We found that an antisymmetric pairwise STDP rule with a reasonable time constant (20 ms) did not maintain the relative firing times of a sequence when the spike trains of successively firing neurons overlapped. A burst timing dependent plasticity (BTDP) rule [106] that low pass filtered spike trains before applying pairwise STDP was able to maintain relative firing times, but could lead to contraction of the sequence if noise enabled neurons to spike earlier than their typical firing times. In contrast, we found that a triplet STDP rule well maintained the firing times of neurons in a feedforward, excitatory network. The minimal triplet rule introduced by Pfister and Gerstner depends on triplets of spikes (post-pre-post) for potentiation and pairwise interactions for depression [21]. The update for the triplet rule for excitatory neurons was

$$\Delta_{\mathrm{ij}}^{\mathrm{STDP}} = A_+ \left(w^{\max} - w_{\mathrm{ij}}\right) \; {k_{\mathrm{ij}}}^{+,\,1}\left({t_j}^{(l)}\right) \; {k_j}^{+,\,2}\left({t_j}^{(l)} - \epsilon\right) - A_- w_{\mathrm{ij}} \; {k_{\mathrm{ij}}}^{-}\left({t_i}^{(m)}\right) \tag{2.15}$$

where $w^{\max}$ differs for E→E and E→I synapse, m and l index the spike times of the pre- and postsynaptic neuron, respectively, ϵ is a small positive constant, and the k variables keep

track of relative spike timing and implement the eligibility window:

$$\tau_1{}^{+}\frac{dk_{\mathrm{ij}}{}^{+,\,1}(t)}{\mathrm{dt}} = -k_{\mathrm{ij}}{}^{+,\,1}(t) + \delta\left(t - t_i{}^{(m)}\right) \tag{2.16}$$

$$\tau_2{}^{+}\frac{dk_j{}^{+,\,2}(t)}{\mathrm{dt}} = -k_j{}^{+,\,2}(t) + \delta\left(t - t_j{}^{(l)}\right) \tag{2.17}$$

$$\tau^{-}\frac{dk_{\mathrm{ij}}{}^{-}(t)}{\mathrm{dt}} = -k_{\mathrm{ij}}{}^{-}(t) + \delta\left(t - t_j{}^{(l)}\right) \tag{2.18}$$

The constants regulating the relative strength of potentiation and depression, $A_+ = 5$ and $A_- = 2$ (25 and 0, respectively, for E→I synapses), and the timescales of STDP, $\tau^- = 33.7\ ms$, $\tau_1{}^+ = 16.8\ ms$, and $\tau_2{}^+ = 40\ ms$, were inspired by the parameters given by Pfister and Gerstner in their triplet STDP model of neurons in the visual cortex. We found sequences were most stable when when potentiation due to triplets was implemented in "nearest spike" spike fashion, i.e. $k_{\mathrm{ij}}{}^{+,\,1}(t)$ and $k_j{}^{+,\,2}(t)$ were bounded by $[0,\ 1]$. To stabilize E→I STDP, each inhibitory neuron was assigned a total excitatory synaptic input bound determined by the neuron's total excitatory synaptic input at the beginning of the simulation. When the bound was exceeded, E→I weight w_{ij} was rescaled as

$$w_{\mathrm{ij}} \to \frac{\sum_k w_{\mathrm{kj}}{}^{(0)}}{\sum_k w_{\mathrm{kj}}} w_{\mathrm{ij}}, \tag{2.19}$$

where $w_{\mathrm{kj}}{}^{(0)}$ represents the size of the synapse prior to the first activation of the network. No plasticity takes place on I→E connections.

Activation of Networks

At the beginning of each trial, networks were active for 10 ms, after which each of the first 10 neurons of the network was driven by an independent burst (4 spikes, 660 Hz), with onset times drawn from a Gaussian (mean 10 ms, STD 1 ms). Input weights were chosen to produce reasonable spiking behavior in the first layers of the network. Networks were simulated for an additional 115 ms following the stimulus. The time step for all simulations was 0.1 ms.

Fluctuating Input

For models in which fluctuating input was supplied to projector neurons, each neuron received independent poisson input ($\lambda = 50\ Hz$) with input weight $w_p = 1.6e{-}4$. The input rate and weight were constrained by experimental studies in which project cell activity increased a factor <3 when $GABA_A$ antagonist gabazine was locally infused within HVC. We mimicked this procedure *in silico* by removing inhibition from an individual projector cell and observing the change in activity during activation of the network.

Network Initialization Procedure

To initialize networks, an activation of each network was first simulated once without plasticity. Neurons that fired during this trial were assumed to be active. Active neurons were assigned a uniform firing rate setpoint ($r^{(0)} = 3$) and were subject to cell-autonomous firing rate homoeostasis as given in Eq. (1) and STDP for 3000 additional activations of the network. For networks with population homeostasis, $\{m_i\}$, the local secreted factor concentrations were computed for activations 1100-1200, and then averaged to set $\left\{m_i{}^{(0)}\right\}$; from this point, population homeostasis was permitted to act on the network.

Simulated Tetanus Toxin Perturbation of Neurons

Tetanus toxin perturbation of the network was simulated by randomly selecting a cell with probability p_T for perturbation and removing all its outgoing connections over 5 consecutive trials. p_T was chosen such that $1 - (1 - p_T)^5$ equaled the perturbation percentage reported. Perturbed cells were assumed to not contribute to the population activity level. Networks were then allowed to evolve for 3000 renditions according to the plasticity rules described above. We classified networks as 'recovered' if for >80% of renditions 2900 to 3000 (i) each group of 20 contiguous neurons (cell indices 0-19, 20-39, etc.) had at least one active neurons, (ii) the average spike time of cells 0-19 preceded that of 20-39, and so on, and (iii) network activity persisted for >10 ms after the onset of activation (typical network sequence duration was 100 ms initially).

Characterization of graded and saltatory recoveries

To classify recoveries as saltatory or graded, we computed the variance in activation length within a running time window (20 activations), and found the maximum of this trajectory. We reasoned that graded recoveries should produce narrow distributions of activation lengths within short windows of time, whereas saltatory recoveries should contain a short period in which a wide range of activation lengths could be produced. This analysis is presented in Supplementary Fig. 6.

2.4 A sequence model of HVC has tempo largely invariant of connection strengths in the network

How do the parameters of a sequence model affect its timing? While there is some evidence to suggest that sequences within HVC may be initiated by input from other nuclei [138], activity within HVC is thought to dictate song tempo on an intra-syllabic, if not broader, scale. This hypothesis establishes a link between song tempo and the propagation speed of the excitatory mode in a sequence model of HVC. Song tempo is incredibly consistent, responding only minimally to perturbations to HVC, including transection [139], ablation of inputs [122], and cooling [94].

Here, we attempt to calculate the propagation speed of a packet of excitation as it travels through an HVC-like network in order to predict how perturbations to the circuit might affect song tempo. We numerically compute the propagation speed and demonstrate it only changes minimally in response to the strength of feedforward excitation and recurrent inhibition, while the number of spikes in the propagating mode changes in compensatory fashion. In this model, the strength of feedforward excitation is offset by depression of the membrane voltage of HVC(RA) cells that grows as a function of the number of propagating spikes. The result is that the propagating packet largely travels at the same speed through the network, though the size of the packet may vary as a function of strengths of couplings chosen.

2.4.1 Background

Prior work has shown that feedforward excitatory sequences with recurrent inhibition possess a unique propagating stable mode [99, 102]. This mode is an attractor in the sense that the traveling packet of excitation will eventually stabilize, reaching a discrete time-translationally invariant state. If neurons are ordered according to their time of firing by index i, this invariance can be expressed as

$$r_i(t) = r_{i+1}(t + \Delta t), \tag{2.20}$$

where $r_i(t)$ is the firing rate of neuron i at time t. This invariance equation is wave-like in the limit of an infinitely large network in which neurons are indexed by a continuous variable [140].

We now focus on a simplified setting in which the feedforward structure of HVC is a reduced one-by-one structure, i.e. excitatory neuron i only drives neuron $i + 1$, which has been studied before in [99, 102]. We further assume that each neuron is connected to a single inhibitory neuron, which in turn recurrently inhibits every unit in the sequence. Excitatory neurons connect to one another in sequence with coupling strength W_E and to the inhibitory neuron with strength W_{EI}. The inhibitory neuron in turn connects back to all excitatory neurons with strength W_{IE}, which is assumed to be negative.

We assume each neuron obeys the typical current-coupled leaky integrate and fire dynamics given by

$$\tau_m \dot{V}(t) = -V(t) + W_E I_E(t) + W_{IE} I_I(t), \tag{2.21}$$

where $V(t)$ is the membrane voltage of the neuron, τ_m is the membrane time constant, and I_E and I_I denote the incoming excitatory and inhibitory currents to the cell, and W_E and W_{IE} denote the strength of the coupling. We assume $W_E \geq 0$ and $W_{IE} \leq 0$. This equation may be rewritten using the Green's function, giving

$$V(t) = \frac{1}{\tau_m} \int_0^t e^{-(t-t')/\tau_m} \left(W_E I_E(t') + W_{IE} I_I(t') \right) dt', \tag{2.22}$$

When $V(t)$ hits a threshold voltage V_{th}, $V(t)$ is reset to 0, and a spike is emitted. The neuron then enters a refractory period of duration t_r during which input current is ignored.

Assumption 1: A spike in a presynaptic neuron evokes a current block in a postsynaptic neuron given by

$$I_\alpha(t) = \Theta(t)\,\Theta(\tau_\alpha - t), \tag{2.23}$$

where Θ denotes the Heaviside function, and τ_α dictates the duration of the current response.
Assumption 2: The current evoked by a set of N spikes arriving at a constant inter-spike interval, Δt_d, can be approximated as a constant current of size,

$$I_N^{(\Delta t_d)}(t) = \sum_{n=0}^{N-1} I_\alpha(t - n\Delta t_d) \approx \tilde{I}\,\Theta(t)\,\Theta(t^{**} - t)\,, \tag{2.24}$$

where

$$t^{**} = (N-1)\Delta t_d + \tau_\alpha, \tag{2.25}$$

and

$$\tilde{I} = \frac{N\tau_\alpha}{t^{**}}. \tag{2.26}$$

The total charge evoked by the set of spikes is $N\tau_\alpha$. This approximation is valid so long as $\Delta t_d \lesssim \tau_\alpha$. The expression $I_N^{(\Delta t_d)}(t)$ has appropriate limits; for example,

$$\lim_{\Delta t_d \to 0} I_N^{(\Delta t_d)}(t) = N\,\Theta(t)\,\Theta(\tau_\alpha - t)\,. \tag{2.27}$$

Calculating the membrane voltage for a constant current

For an arbitrary constant excitatory current of magnitude I_0 and duration t_I, we have from Eq. (2.22) that

$$V(t) = \begin{cases} 0, & t < 0, \\ W_E I_0 \left(1 - e^{-t/\tau_m}\right), & 0 \le t < t_I, \\ W_E I_0 \left(1 - e^{-t_I/\tau_m}\right) e^{-(t-t_I)/\tau_m}, & t \ge t_I. \end{cases} \tag{2.28}$$

Calculating the evoked number of spikes and their average ISI when inhibition is constant

We now consider the response of a postsynaptic neuron driven by a presynaptic burst of size N and inter-spike interval Δt_d when there is a constant inhibitory current I_I^0. We first

compute the time to the first spike in the postsynaptic neuron, t^*. Assuming the postsynaptic neuron has been driven by $I_I(t) = I_I^0$ for a long time, Eq. (2.22) becomes

$$V(t) = W_{IE} I_I^0 + \frac{W_E}{\tau_m} \int_0^t e^{-(t-t')/\tau_m} I_E(t')\, dt'. \tag{2.29}$$

We set

$$I_E(t) = \tilde{I}\, \Theta(t)\, \Theta(t^{**} - t)\,, \tag{2.30}$$

and solve for t^* such that $V(t^*) = V_{\text{th}}$:

$$V_{\text{th}} = W_{IE} I_I^0 + \frac{W_E}{\tau_m} \int_0^{t^*} e^{-(t^*-t')/\tau_m} \tilde{I}\, \Theta\big(t^{**} - t'\big)\, dt' \tag{2.31}$$

$$= W_{IE} I_I^0 + W_E \tilde{I} \left(1 - e^{-t^*/\tau_m}\right). \tag{2.32}$$

If $t^* > t^{**}$, no spikes are evoked. If $t^* \leq t^{**}$,

$$t^* = -\tau_m \log\left(1 - \frac{V_{\text{th}} - W_{IE} I_I^0}{W_E \tilde{I}}\right). \tag{2.33}$$

After the first spike, the membrane is reset to zero. The time to the next spike $\tilde{t}$ satisfies

$$V_{\text{th}} = \left(W_E \tilde{I} + W_{IE} I_I^0\right) \left(1 - e^{-\tilde{t}/\tau_m}\right), \tag{2.34}$$

since the membrane is set the reset voltage after spiking, which here we choose to be zero. Thus, the time to the next spike is shorter than t^* if $W_{IE} I_I^0 < 0$. We can rearrange Eq. 2.34 to find

$$\tilde{t} = -\tau_m \log\left(1 - \frac{V_{\text{th}}}{W_E \tilde{I} + W_{IE} I_I^0}\right), \tag{2.35}$$

Including a refractory period t_r, the average ISI is

$$\Delta t_d' = \tilde{t} + t_r. \tag{2.36}$$

The total number of evoked spikes is

$$N' = \begin{cases} 1 + \left\lfloor \dfrac{t^{**} - t^*}{\Delta t_d'} \right\rfloor, & t^* \leq t^{**}, \\ 0, & t^* > t^{**}, \end{cases} \tag{2.37}$$

since every spike after the first requires the same amount of time under the input current, $\Delta t_d'$, to induce.

Movement in phase space of the dynamics

With equations Eqs. 2.37 and 2.36 in hand, which define N' and $\Delta t_d'$ in terms of N and Δt_d, we can now the compute the evolution of a spike train of N spikes and inter-spike interval Δt_d as it propagated through the 1D network.

The phase flow is defined by the mapping

$$(N, \Delta t_d) \to \left(N', \Delta t_d'\right), \tag{2.38}$$

where the lefthand side describes the activity of a neuron in the sequence and the righthand side describes the activity of the neuron that follows it.

Recurrent inhibition

This reduced description still lacks a critical component of the HVC model: recurrent inhibition. We assume the dynamics of the inhibitory neuron are linear and fast. We found rapid inhibitory time constants to be essential to the generation of stable sequential activity in the modeling work performed in Wang et al.. Under these assumptions, the dynamics of the inhibitory current to all excitatory neurons is

$$I_I = W_{EI} \sum_i I_E(t). \tag{2.39}$$

Assuming steady-state dynamics, a reasonable approximation for the inhibitory current is

$$I_I = W_{EI} \frac{N \tau_\alpha}{t^*}. \tag{2.40}$$

To see this, we consider that each neuron outputs a current for duration t^{**}, and, at steady state, neurons activate with a delay t^*. The average number of neurons active at any moment in time is t^{**}/t^*, where each neuron contributes current $\tilde{I}$, giving

$$\begin{aligned} I_I &= W_{EI} \frac{t^{**}}{t^*} \tilde{I} \\ &= W_{EI} \frac{N \tau_\alpha}{t^*}. \end{aligned} \tag{2.41}$$

We can make the assumption that the inhibitory current evolves quickly, and rewrite the phase-space mapping as

$$(N, \Delta t_d, I_I) \to \left(N', \Delta t_d', I_I'\right). \tag{2.42}$$

We must also replace the constant inhibitory current in the equations above with the new dynamical current. The final recursive equations read

$$N' = 1 + \left\lfloor \frac{(N-1)\Delta t_d + \tau_\alpha - t^*}{\Delta t_d'} \right\rfloor, \tag{2.43}$$

$$\Delta t_d' = -\tau_m \log\left(1 - \frac{V_{\text{th}}}{W_E \tilde{I} + W_{IE} I_I}\right) + t_r, \tag{2.44}$$

$$I_I' = W_{EI} \frac{N\tau_\alpha}{t^*}, \tag{2.45}$$

$$t^* < (N-1)\Delta t_d + \tau_\alpha, \tag{2.46}$$

where

$$t^* = -\tau_m \log\left(1 - \frac{V_{\text{th}} - W_{IE} I_I}{W_E \tilde{I}}\right). \tag{2.47}$$

Since the inhibitory neuron is fully linearized, the product $W_{IE}W_{EI}$ is redefined as W_I, the strength of the recurrent inhibition. With this redefinition, at steady state, we have

$$W_{EI} I_I = W_I \frac{N\tau_\alpha}{t^*}. \tag{2.48}$$

Stationary solutions

Stationary solutions satisfy

$$(N, \Delta t_d, I_I) \to (N, \Delta t_d, I_I). \tag{2.49}$$

To find the dynamical attractor of the system, Eqs. 2.43-2.45 can be run forward recursively and checked for convergence while ensuring compliance to Eq. 2.46. Propagation speed ρ was defined as the inverse delay between successive neurons, $\rho = 1/t^*$.

2.4.2 Results

We ran Eqs. 2.43-2.45 for 1000 steps starting from a range of initial conditions $N \in [1, 10]$ and $\Delta t_d \in [1, 5]$ ms (see Table2.2 for all model parameters). The initial inhibitory current was always assumed to be zero. Trajectories typically traversed a path through large N and I_I before settling down to lower values of both N and I_I that were stable (Fig 2.14, bottom). Concurrently, the inter-spike interval of the propagating mode approached the refractory

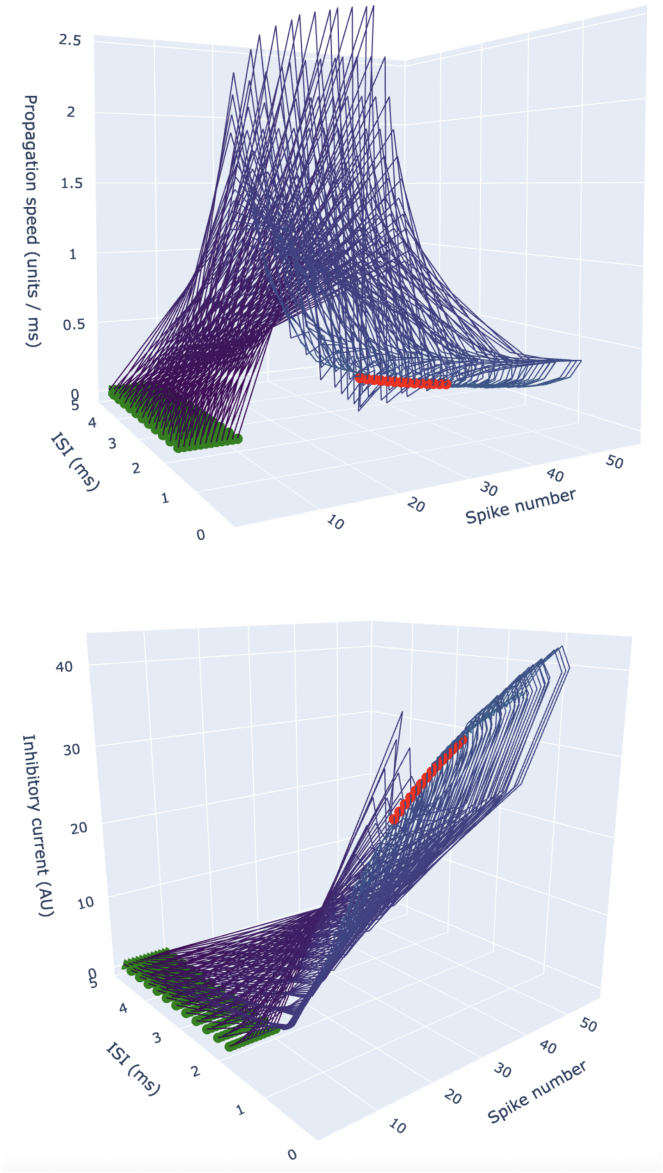

Figure 2.14: **Trajectories for reduced model of sequence dynamics (Top)** Evolution of trajectories that begin with an initial spike number and frequency given by green points. Trajectories evolve along purple lines in the space of propagation speed, interspike interval (ISI), and spike number. Most trajectories converge at red dot, while trajectories that begin with a single spike do not reach the steady state condition. **(Bottom)** Equivalent plots shown in the space of spike number, ISI, and inhibitory current that is broadcast to all excitatory cells. Trajectories that begin with a high initial spike count or small interspike interval take a path to steady state condition (red dot) that passes though states of high inhibitory current, which then settles down.

period of the excitatory neurons. We found that network activity converged to one of several nearby fixed points distinguished by discrete values of N (Fig 2.14, red dots). We found that N varied inversely with the propagation speed, ρ, suggesting that bursts with fewer spikes propagate more rapidly in these networks (Fig 2.14, top).

To determine how the dynamical attractor varied as a function of coupling strengths, we repeated the procedure above for $W_E \in [0, 20]$ and $W_I \in [0, -1]$. We screened for stable modes by ensuring N was nonzero through all 1000 steps and confirming that the number of propagating spikes converged. We found stable modes existed for a large part of parameter space, except when W_I grew large enough given the value of W_E, indicating that sufficiently strong recurrent inhibition abolishes propagating modes (Fig. 2.15d). Also, when $W_I = 0$, a single value of W_E was stable, corresponding to a fine tuning that allowed a burst of spikes to neither grow nor decay, as previously shown in Li and Greenside.

Parameter	**Description**	**Value**
τ_m	Membrane time constant	5 ms
V_{th}	Baseline membrane potential	0.5
t_r	Refractory period	0.5 ms
τ_α	Synaptic time constant	3 ms

Table 2.1: Analytical parameters and default values.

When $W_I < 0$, the propagation speed was largely invariant to changes in W_E and W_I (Fig. 2.15a). An exception occurred at small values of both feedforward excitation and recurrent inhibition, which produced modestly increased propagation speeds. The ratio of the fastest and slowest modes was approximately $\sim$1.8, with the majority of parameter space producing slower modes (Fig. 2.17, left).

In contrast, N showed large sensitivity to changes in both parameters, decreasing with stronger recurrent inhibition and increasing with stronger feedforward excitation (Fig. 2.15b). Propagating modes consistent with $\sim$4-6 spikes, as reported in zebra finch HVC [93], were found at the edge of parameter space, just below values of W_I large enough to suppress

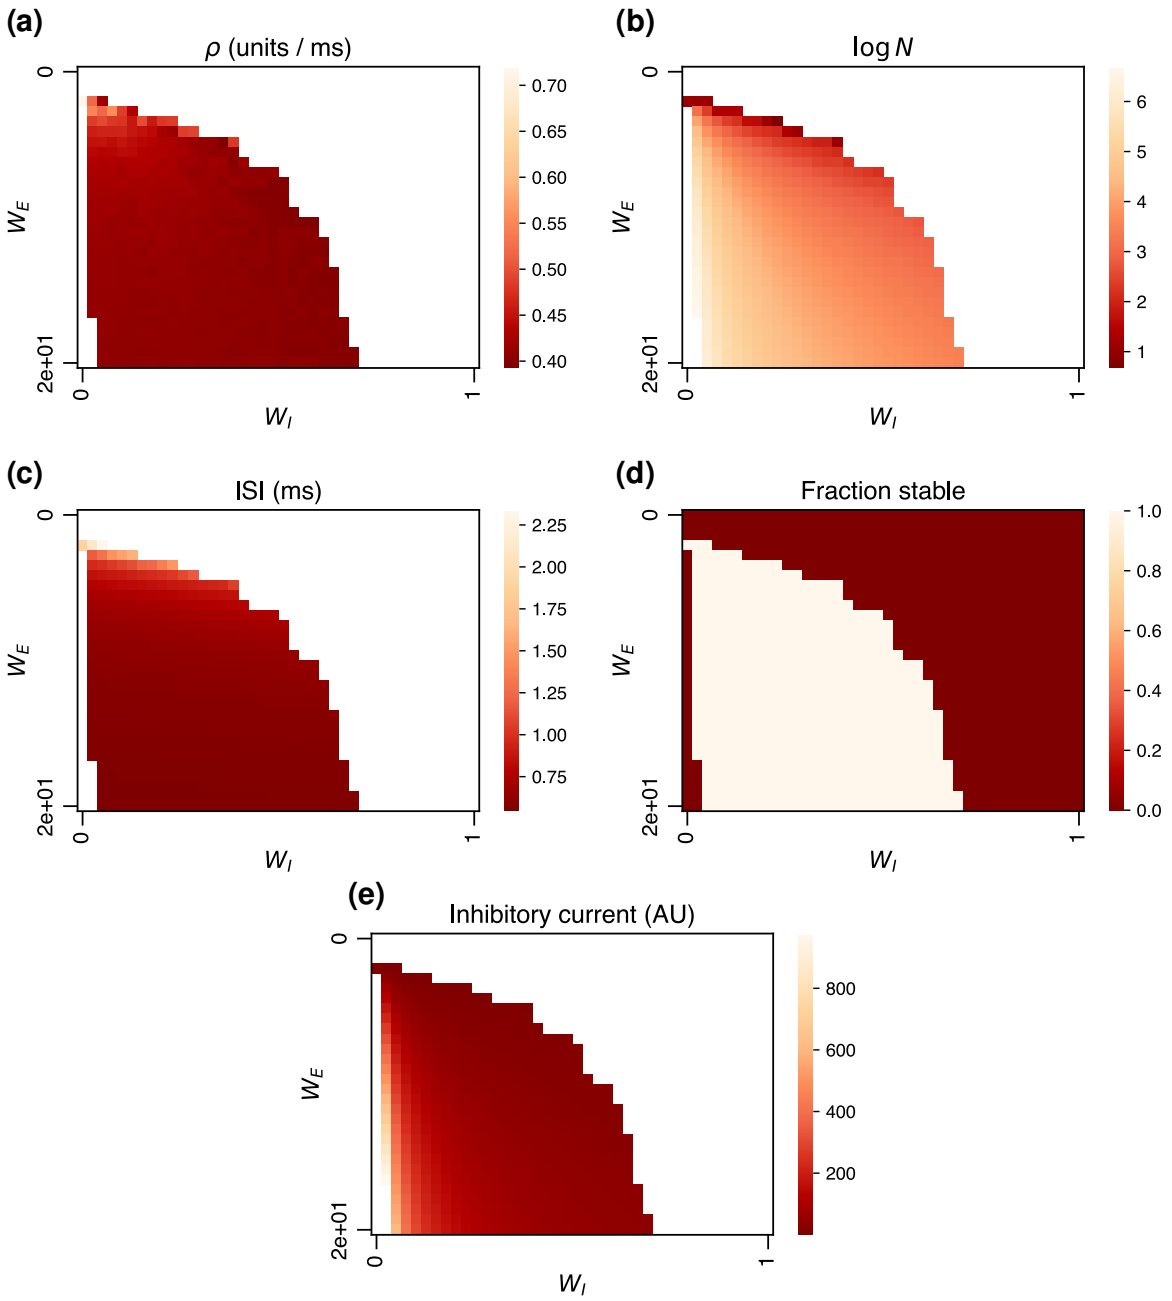

Figure 2.15: **Properties of stable propagating stable of a one-by-one chain with recurrent inhibition** Properties graphed as a function of strength of feedforward excitation and excitatory input to single global interneuron. Only points where the fraction of stable trajectories exceeded 80% shown. **(a)** Propagation speed in units of neurons per millisecond.**(b)** The interspike interval (inversely proportional to frequency) of excitatory cells. **(c)** Number of spikes fired by each neuron in stable mode. **(d)** Region of parameter space with stable modes. **(e)** Global inhibitory current.

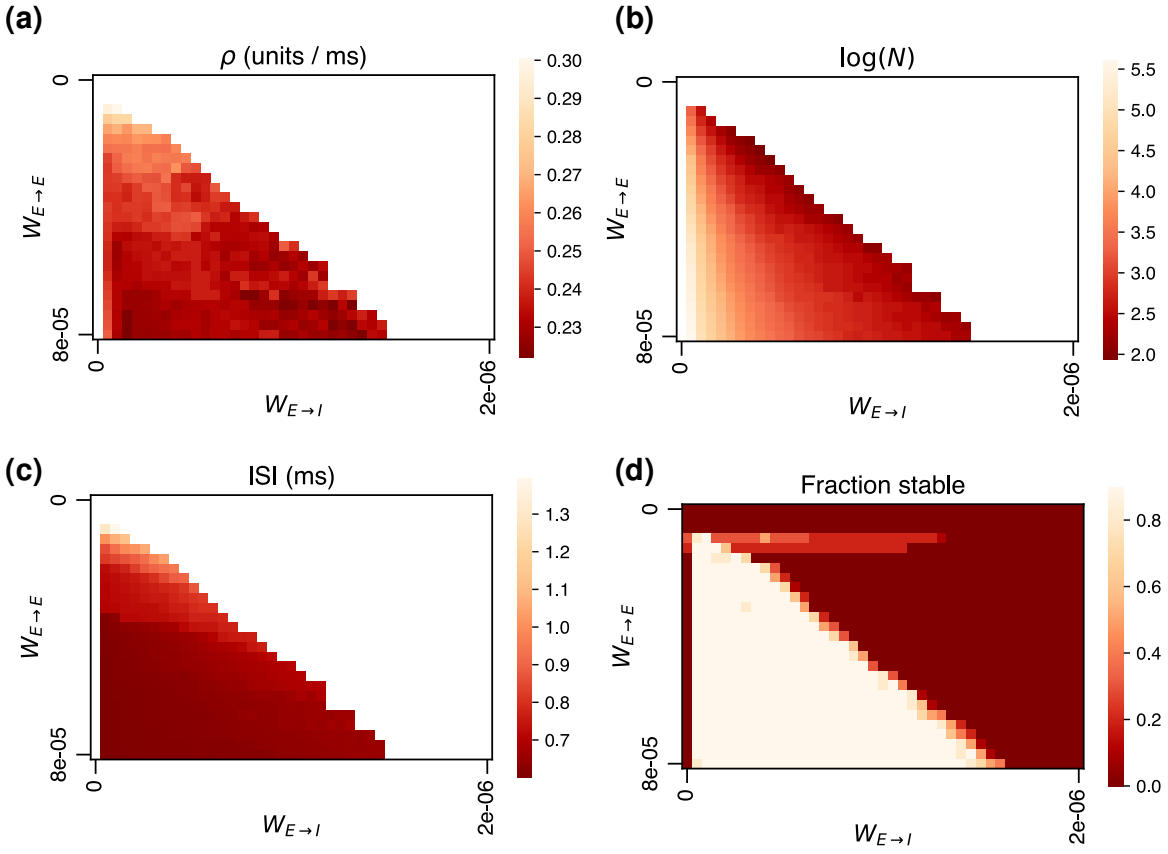

Figure 2.16: **Properties of stable propagating stable of a one-by-one chain with recurrent inhibition** Properties graphed as a function of strength of feedforward excitation and excitatory input to single global interneuron. Only points where the fraction of stable trajectories exceeded 80% shown. **(a)** Propagation speed in units of neurons per millisecond.**(b)** The interspike interval (inversely proportional to frequency) of excitatory cells. **(c)** Number of spikes fired by each neuron in stable mode. **(d)** The fraction of stable trajectories for the condition (out of 10 simulations per condition).

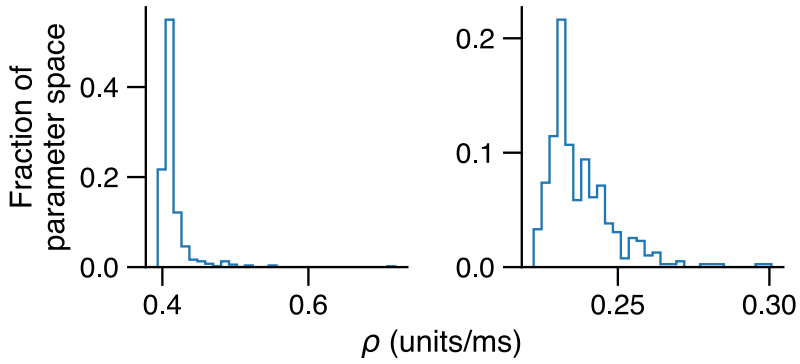

Figure 2.17: **Fraction of parameter space represented by each value of ρ (Left)** Results from numerical calculations. Distribution of ρ is strongly peaked at a low value. **(Right)** Results from simulations of current-coupled leaky integrate and fire network. Shape of distribution is a qualitative match for predictions of theory.

propagating modes. Weaker values of W_I produced modes with orders of magnitude more spikes, suggesting this parameter regime is unrealistic for HVC. In this region of parameter space, large inhibitory currents are produced by the large number of spikes in the excitatory population (Fig. 2.15e), which would require an extremely active interneuron population. We additionally found typical inter-spike interval of propagating excitation was largely invariant to changes in coupling strengths, except at values of weak excitation and inhibition, where it slowed down modestly (Fig. 2.15c).

We checked our numerical solutions against simulations of a 1D chain of current-coupled excitatory neurons with recurrent inhibition, aiming for qualitative agreement. Specifically, we simulated a feedforward network of 250 excitatory neurons, coupled in a one-by-one fashion with strength $W_{E\to E}$, and connected to a population of 50 inhibitory neurons with strength $w_{ij} \sim \mathcal{N}(W_{E\to I}, 0.1 \times W_{E\to I})$. Each inhibitory neuron connected to every excitatory neuron with strength $W_{I\to E}/N_{\text{inh}}$, where N_{inh} is the number of inhibitory neurons. The dynamics of the membrane potential of each neuron, $V_i(t)$, evolved according to

$$C_m \dot{V}_i(t) = g_l(E_l - V_i(t)) + I_i(t), \tag{2.50}$$

where C_m is the membrane capacitance, g_l is the leak conductance, and E_l is the leak potential, and the input current, $I_i(t)$, is given by

$$I_i(t) = \sum_{t_j^{(n)}} \int_0^\infty w_{ij}\, I_\alpha(t - t_j^{(n)})\, \delta(t - t_j^{(n)}) dt, \tag{2.51}$$

where $t_j^{(n)}$ is the n^{th} spike of presynaptic neuron j, w_{ij} is the synapse $j \to i$, and $I_\alpha(t)$ is a kernel that converts spikes into current pulses according to

$$I_\alpha(t) = \Theta\left(t/\tau_\alpha\right) \frac{t}{\tau_\alpha} e^{-t/\tau_\alpha}. \tag{2.52}$$

When $V_i(t)$ reached the threshold voltage, V_{th}, a spike was emitted and the neuron was clamped to potential V_{reset} for a refractory period, t_r.

Parameter	**Description**	**Value**
C_m	Membrane capacitance	1 μF
g_l	Leak conductance	1e-4 S
E_l	Leak reversal potential	-70 mV
V_{th}	Threshold membrane potential	-50 mV
V_{reset}	Reset voltage	-70 mV
$t_r^{(E)}$	Refractory period (excitatory)	0.5 ms
$t_r^{(I)}$	Refractory period (inhibitory)	0 ms
$W_{I \to E}$	Coupling strength $I \to E$	-6e-7
τ_α	Synaptic time constant	1.6 ms

Table 2.2: Spiking network model parameters and default values.

To empirically determine the propagating modes within these networks, the activity of first excitatory neuron in the feedforward sequence was clamped to a high frequency (1500 Hz, 2 ms long) burst of Poisson-distributed spikes (3 spikes on average). The dynamics of the network were then simulated for 600 ms. We assessed convergence of the dynamics by ensuring that (1) standard deviation in the number of spikes fired by neurons 90-100 was less

than 1 and either (2) the last excitatory neuron fired at some point or (3) the simulation ended while an excitatory neuron was still active. If these criteria were met, the spikes per neuron was determined by computing the mean of neurons 50-100. The propagation speed was determined by computing the time lag between the first spikes of neurons 50 and 100 and dividing by the number of neurons in this segment. We computed these values over a wide range of parameters, $W_{E\to E} \in [0, 8e{-}5]$ and $W_{E\to I} \in [0, 2e{-}6]$, with $W_{I\to E}$ held fixed.

Spiking network simulations qualitatively agreed with the results of our model (Fig 2.16). The propagation speed was the largest at small values of $W_{E\to E}$ and $W_{E\to I}$, but all across parameter space, the distribution of propagation speeds was strongly peaked at a slow value (Fig 2.17, right). N was strongly correlated with $W_{E\to E}$ and inversely correlated with $W_{E\to I}$. The ISI of the propagating burst was largest at small values of $W_{E\to E}$ and $W_{E\to I}$.

The specific predictions made by this model are though changes to the circuit that reduce the strength of solely HVA(RA)→HVA(RA) monosynaptic connections or the strength of recurrent inhibition should be reflected in the size of the propagating burst, but not its speed nor its frequency. Joint perturbations to both coupling strengths, however, may have a minimal impact on both N and ρ, as the dynamics are less sensitive to changes in W_E and W_I proportionally toward zero. This idea is consistent with reports that bilateral HVC transection only results a moderate ($\sim$5%) slowing of song [139].

Chapter 3

SEARCHING FOR RULES: MODELING SEQUENCE SELF-ORGANIZATION AND RECOVERY VIA PLASTICITY

The contents of this chapter are published in Neural Information Processing Systems as Bell et al..

Intrinsic dynamics within the brain can accelerate learning by providing a prior scaffolding for dynamics aligned with task objectives. Such intrinsic dynamics should self-organize and self-sustain in the face of fluctuating inputs and biological noise, including synaptic turnover and cell death. An example of such dynamics is the formation of sequences, a ubiquitous motif in neural activity. The sequence-generating circuit in zebra finch HVC provides a reliable timing scaffold for motor output in song and demonstrates a remarkable capacity for unsupervised recovery following perturbation. Inspired by HVC, we seek a local plasticity rule capable of organizing and maintaining sequence-generating dynamics despite continual network perturbations. We adopt a meta-learning approach introduced by Confavreux et al, which parameterizes a learning rule using basis functions constructed from pre- and postsynaptic activity and synapse size, with tunable time constants. Candidate rules are simulated within initially random networks, and their fitness is evaluated according to a loss function that measures the fidelity with which the resulting dynamics encode time. We use this approach to introduce biological noise, forcing meta-learning to find robust solutions. We first show that, in the absence of perturbation, meta-learning identifies a temporally asymmetric generalization of Oja's rule that reliably organizes sparse sequential activity. When synaptic turnover is introduced, the learned rule incorporates an additional form of homeostasis, better maintaining sequential dynamics relative to other previously proposed rules. Additionally, inspired by recent findings demonstrating plasticity in synapses from inhibitory interneurons in HVC, we explore the role of inhibitory plasticity in sequence-generating circuits. We find that learned plasticity adjusts both excitation and inhibition

in response to manipulations, outperforming rules applied only to excitatory connections. We demonstrate how plasticity acting on both excitatory and inhibitory synapses can better shape excitatory cell dynamics to scaffold timing representations.

3.1 Introduction and related work

How computational structures are organized and maintained within the brain is a central question within neuroscience. While feedback is clearly essential for learning, self-organization of neural circuits can unfold without feedback, e.g. during development. Brains have evolved specific cell types with nonrandom spatial organization, plasticity rules, and connectivity that likely introduce a strong set of inductive biases on the information processing they perform. How might organization of useful circuit dynamics be established and maintained throughout life without the need for feedback? Recent work suggests self-organized computations, once established, can accelerate learning and improve performance when experience is limited: Nicola and Clopath [89] demonstrated that a stable high dimensional time signal could improve a network's performance on sequential motor tasks (Fig. 3.1a). Najarro and Risi [63] learned Hebbian plasticity that orchestrated spontaneous walking behavior in quadruped agents; similar work has shown architectural priors increase the sample efficiency and generalization of RL approaches to locomotion [66, 142]. Additionally, in RL settings, supplying agents with a time input permits them to adopt time-dependent policies [143].

The ability of computational primitives, such as timing representations, to self-organize is challenged by the shifting structure of neural circuits. Synaptic loss, synaptogenesis, cell death, and neurogenesis pose challenges for all learning algorithms, but particularly for self-organization which must be based solely on local information rather than global task performance.

Here, we aim to find plasticity rules that self-organize and maintain one useful computational primitive: sparse, sequential activity. Such activity is widely seen in many areas of the brain including hippocampus [144], cortex [145], and basal ganglia [146]. In the songbird zebra finch, area HVC (used as a proper noun), a cortical-like region, displays sequential activity representing time [93], reducing the problem of motor learning to driving the correct motor neuron at the correct moment [88]. Extensive literature has explored how such

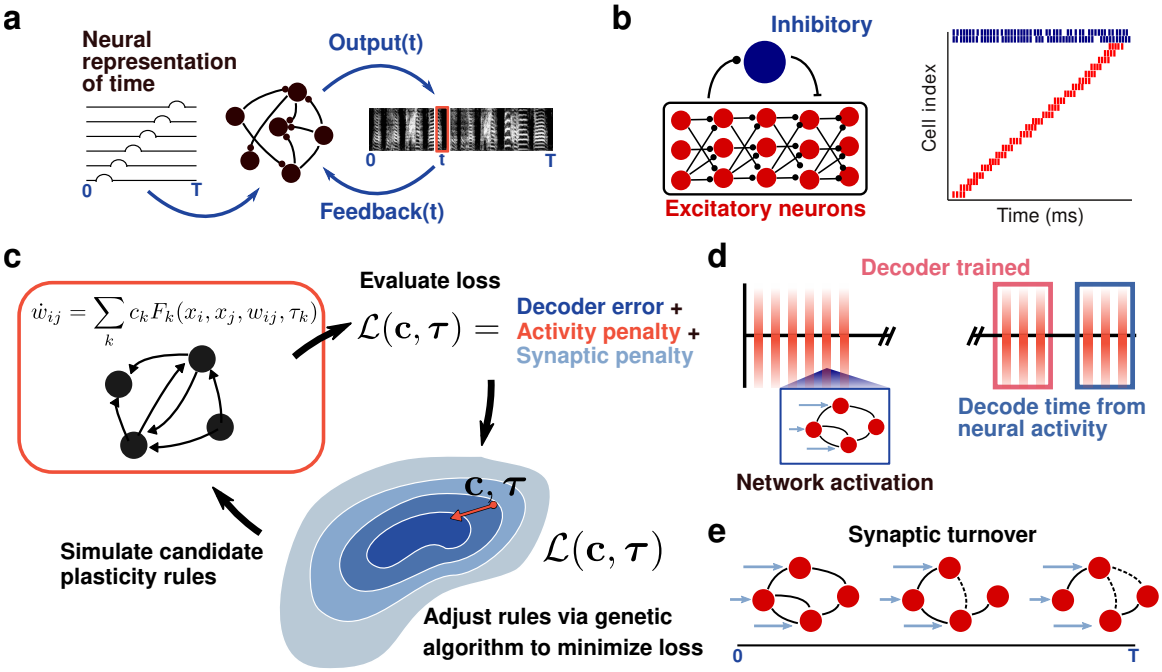

Figure 3.1: **Meta-learning approach to discovering plasticity rules that organize sequences (a)** In zebra finch song learning, a neural representation of time (left) in HVC simplifies the sequential motor learning task of producing the correct spectral output. **(b)** Putative network structure of zebra finch HVC: a feedforward, excitatory network with recurrent inhibition (left). HVC excitatory neurons (red) fire sparsely in time while interneurons (blue) fire tonically (right). **(c)** Strategy for learning plasticity underlying sequence organization: candidate plasticity rules, parameterized by a set of coefficients and time constants, are simulated. A loss function is evaluated on the resulting dynamics, and new candidate rules are generated. **(d)** Test procedure for representation of time. Networks are activated 400 times (red bars). From the final 50 activations, six are chosen to train a decoder and six to test the representation by decoding time from neural activity. **(e)** Discovery of robust plasticity rules is encouraged by introducing synaptic turnover, the stochastic addition and removal of synapses, into simulations.

sequence-generating circuits could emerge in the absence of feedback [100, 101, 104–107, 147], but has largely focused on either how these structures organize or how they self-maintain, using guessed plasticity rules, and neglecting the effects of ongoing synaptic noise. Further, previous work on sequence organization within HVC has focused on plasticity between excitatory (E) neurons. Recent experimental findings show unsupervised recovery of HVC dynamics is accompanied by changes in both E→E and also inhibitory-to-excitatory (I→E) synaptic strength [67].

Here, instead of imposing a guessed rule, we ask which self-supervised plasticity rules can organize and maintain sequential dynamics within a network. We employ meta-learning, a supervised method to learn learning rules [52, 56, 57, 61, 148], in order to discover rules that self-organize a sequence. Our approach stems from a rich history of learning local plasticity rules, including rules that extract representations from data [62], enhance artificial agent performance on familiarity and navigation tasks [53, 63], and explore biologically-plausible replacements or complements to backpropagation [52, 57, 149]. In this study, we parameterize the space of plasticity rules with a basis of activity- and synapse size-dependent terms. The set of coefficients weighting these terms and associated time constants are adjusted to minimize a loss function. Inspired by the HVC context, we pose the loss in terms of the accuracy with which the time since an initial network input can be decoded from the circuit dynamics. We first consider E→E rules alone, and then add I→E and E→I plasticity. We then introduce perturbations to the circuit and investigate which learned plasticity rules promote circuit stability. We find that meta-learned rules for self-organized sequence generation and maintenance contain distinct forms of spike-timing dependent plasticity, homeostasis, and network bounds, which outperform previously proposed sequence-organizing rules in the presence of noise and that plasticity on reciprocal connectivity to inhibitory neurons confers additional stability on the network dynamics. Our main contribution is the exploration of unsupervised and unrewarded plasticity via meta-learning that organizes and maintains a specific and biologically relevant computational motif, a sequence.

3.1.1 Background on zebra finch physiology

In the zebra finch, nucleus HVC contains excitatory neurons that fire sparsely (typically in one burst of spikes) during song and are purportedly arranged in a feedforward structure (Fig. 3.1b) [68]. A subset of these cells, known as $\text{HVC}_{(\text{RA})}$ neurons, project to downstream nucleus RA (robust nucleus of the archistriatum), which in turn projects to vocal neurons of the syrinx and to the brainstem, which regulates respiration [93]. HVC receives excitatory projections from nucleus Uva, which controls the onset of song syllables [150] and provides input for the duration of song [151]. $\text{HVC}_{(\text{RA})}$ neurons inhibit each other disynaptically via a population of inhibitory interneurons [95]. Remarkably, singing behavior can persist when the nucleus is transected [139], demonstrating its resilience.

3.2 Results

3.2.1 Learning biologically plausible plasticity on E→E synapses that organizes a sequence-generating circuit

We first learned a plasticity rule on E→E synapses that organizes a randomly connected network into a sequence-generating circuit in the absence of any perturbation. We initially constrained plasticity to E→E synapses to compare with previously proposed rules, which have largely only considered E→E plasticity. We did not include rules that imposed hard bounds on the size of individual synapses or on the collective strength of all synapses onto a neuron, as we aimed to learn plasticity rules that could permit flexible rescaling of connections in response to perturbations, as has been observed experimentally [67].

We use a network of 25 E and 8 I threshold-linear neurons. Each neuron fired according to $x_j(t) = [V_j(t) - b]^+$, where $V_j(t)$ evolves via $\tau_m \dot{V}_j(t) = -V_j(t) + \sum_i w_{ij} x_i(t)$. Here, w_{ij} is the weight of the synapse $i \to j$, τ_m is the membrane time constant, and b is the bias. Initial connectivity (Fig. 3.2d) was random, but contained no I→I connectivity as is the case in HVC [95, 133] (see Supp. Sec. 3.5 for all network model details).

Meta-learning procedure

We adopt a meta-learning approach pioneered by Bengio et al. [52] and extended by Confavreux et al [51]. We parameterize a set of plasticity rules with coefficients c_k and time constants τ_k, such that individual synapses, w_{ij}, evolve according to

$$\dot{w}_{ij}(t) = \Theta(|w_{ij}(t)|) \sum_k c_k F_k(x_i(t), x_j(t), w_{ij}(t), \tau_k) \tag{3.1}$$

where x_i and x_j are pre- and postsynaptic activities, respectively, Θ is the Heaviside function, and F_k is the k^{th} term in the plasticity rule (see Fig. 3.3a or Supp. Sec. 3.6 for all terms). The basis includes terms that filter pre- and postsynaptic activities with decaying exponentials (denoted e.g. $\tilde{x}_i$) as these terms convey information about the durations of activations and relative ordering, not just their instantaneous rates. Synapses evolving under Eq. 3.1 were bounded so that they obeyed Dale's law.

We next define a loss function that evaluates the quality of the sequential dynamics organized under a chosen $\{c_k, \tau_k\}$ and attempt to minimize it using an evolutionary strategy, Covariance Matrix Adaptation (CMA-ES) [152]. We use CMA-ES to sample from the space of possible $\{c_k, \tau_k\}$ and evaluate the loss at each point by simulating 10 randomly initialized networks under the given rule and evaluating the resultant dynamics at the end of the simulation. Each simulation is divided into 400 activations of 110 ms. At $t = 10$ ms of each activation, a single fixed neuron is driven by a strong kick of excitation [106]. Following this, all other neurons in the network receive Poisson distributed input for a period of 65 ms. A fraction of this Poisson input is held fixed from trial to trial, mirroring input to HVC from the nucleus Uva, which likely does not provide a fully stochastic signal to the downstream area, [150]. The total loss for a given rule is simply the sum of losses across each of the 10 networks.

Loss function

To learn sequences, we define a loss function based on three principles: (1) elapsed time since initial network activation should be readily decoded from network activity, (2) total network activity should be sparse, and (3) total synaptic change in the network should be minimized.

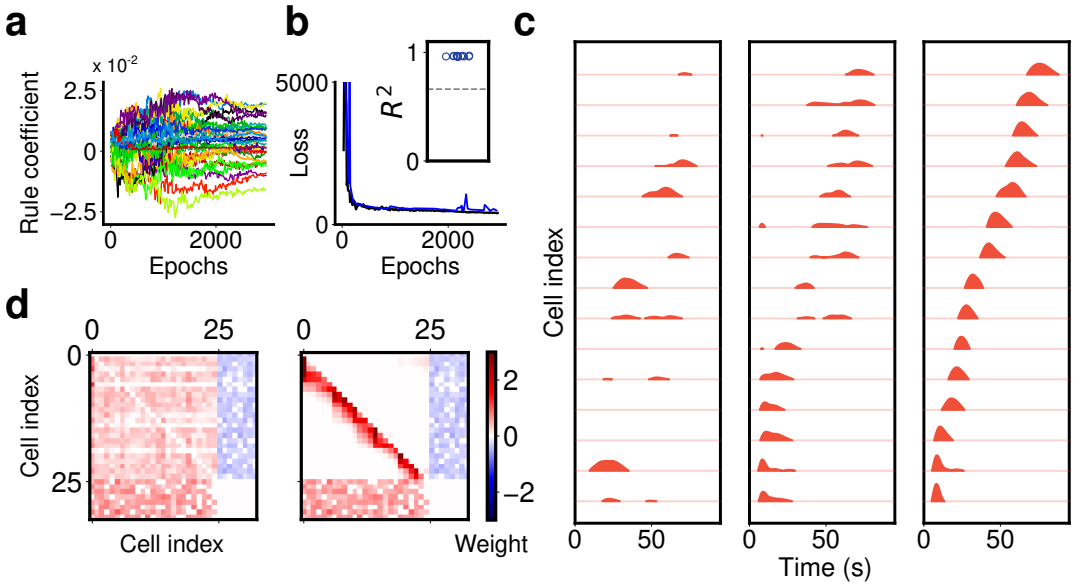

Figure 3.2: **Meta-learning discovers unsupervised local plasticity rules that organize sequential activity (a)** Evolution of coefficients and time constants during meta-learning. **(b)** Training loss (black) and test loss (blue) during meta-learning. Inset: median decoding accuracy for 14 learned rules across 100 networks (blue points) compared to no plasticity (dashed gray line). **(c)** Meta-learned plasticity rules generate sequence dynamics. Network activity during the first activation (left), 200^{th} (middle), and 400^{th} (right), sorted by ordering of mean firing time of the final activation (425^{th}). Note: plasticity rule is held fixed during simulation. **(d)** Weight matrices at activation 1 and 400, sorted based on mean firing time of the final activation. Connectivity between E cells organizes into a feedforward structure.

The latter two principles impose the assumption that effective plasticity does not require excess network activity or synaptic change to stabilize network function. Our loss is

$$\mathcal{L}(\mathbf{c}, \boldsymbol{\tau}) = \mathcal{L}_{\text{dec}}(\mathbf{c}, \boldsymbol{\tau}) + \lambda_{\text{a}} \sum_i \int_0^T x_i(t) dt + \lambda_{\text{s}} \mathcal{P}(\mathbf{c}, \boldsymbol{\tau}), \tag{3.2}$$

where

$$\mathcal{P}(\mathbf{c}, \boldsymbol{\tau}) = \sum_{i,j,k} \int_0^T |c_k F_k(x_i(t), x_j(t), w_{ij}(t), \tau_k)| \Theta(|w_{ij}(t)|) dt. \tag{3.3}$$

$\mathcal{P}(\mathbf{c}, \boldsymbol{\tau})$ penalizes the all synaptic changes due to each component of $\mathbf{F}$. This L_1-like penalty on $\mathbf{F}$ penalizes each term not by the size of the term's coefficient, c_k, but by the quantity of synaptic change it evokes. While penalizing the magnitude of c_k is standard [51, 57], we take this approach to compare different terms on a common scale, as each component of $\mathbf{F}$ has differing dependence on x_i, x_j, and w_{ij}. λ_{a} and λ_{s} are positive constants weighting the activity and synaptic change penalties.

To determine the decoder loss $\mathcal{L}_{\text{dec}}(\mathbf{c}, \boldsymbol{\tau})$, the activity of networks was sampled at 500 time points of six activations of the network. From these, a linear decoder was constructed and used to decode the activity at 200 time points of six subsequent activations. Note that since the initial kick of excitation was time-locked to beginning of each activation, constructing a decoder to read out time elapsed since the beginning of the activation was equivalent to decoding time elapsed since the initial excitatory kick was presented.

Learned E→E plasticity induces sequences using temporally asymmetric Hebbian learning and a peak postsynaptic activity bound

Meta-learning reliably found plasticity rules that organized random E synaptic connectivity into feedforward structures that generated sequences when activated (Fig. 3.2d-e). The organized structures were not grouped into links, as in a synfire chain, but were better described by a kernel in which the strength of a synapse between two cells depended on the lag between their mean firing times (Fig. 3.2d, right). The set of rules discovered by meta-learning are not sparse in the space of $\{c_k, \tau_k\}$ (Fig. 3.3a). To test whether learned plasticity rules truly required all terms with nonzero coefficients, we compared 'dropout' variants of the discovered rule, in which one coefficient within $\{c_k\}$ was set to zero, to its

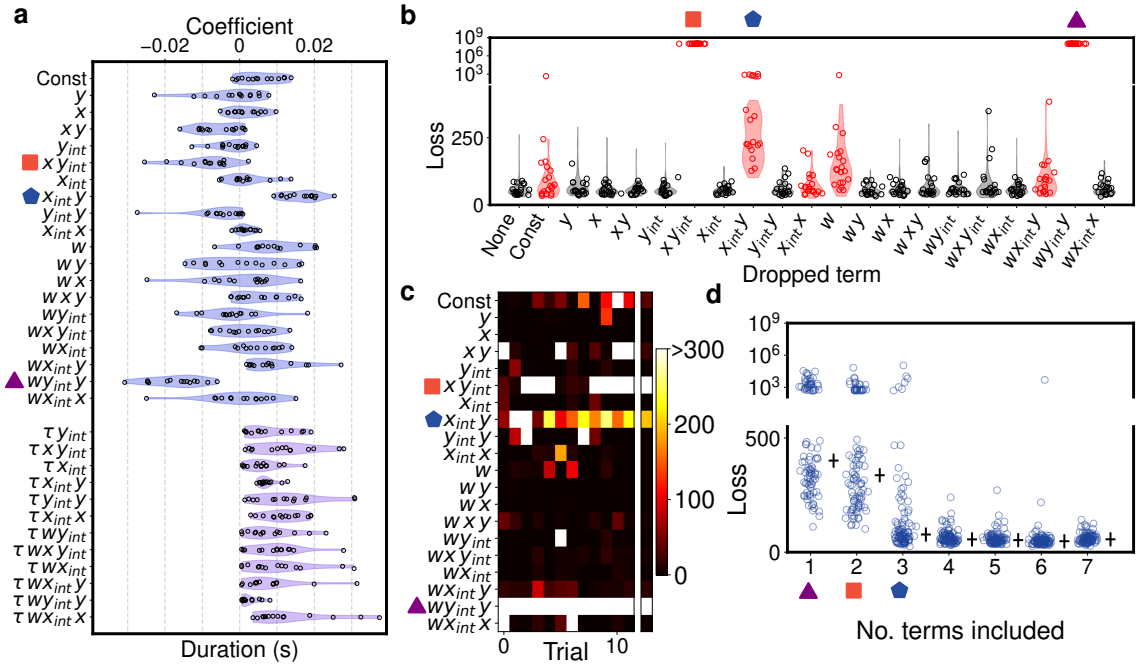

Figure 3.3: **Perturbing learned plasticity rules reveals dependence on temporally asymmetric Hebbian learning and a postsynaptic activity bound (a)** Distributions of coefficients (blue) and time constants (purple) of basis terms across N=14 training instances (individual instances in black). The basis set consists of functions of pre and post-synaptic neural activity and synaptic weights. x = activity of pre-synaptic neuron; y = activity of post-synaptic neuron; w = weight of synapse; $x_{int} = \tilde{x}$ (x filtered with τ). **(b)** Test loss when individual coefficients are set to zero. **(c)** Difference in median loss between full learned solutions run on 100 test networks and solutions with one term dropped. Separate column denotes median loss for terms across all trials. **(d)** Progressive refitting of model in order of impact on median loss with and without term. Symbols indicate the term added at each refitting. Losses for 100 test networks shown (blue); medians shown as black crosses.

unaltered form. We computed the loss of these variants on test networks to determine if a term's absence impacted the loss (Fig. 3.3b). Computing the change in the median loss of networks organized by the learned solution and a dropout variant across 14 discovered rules revealed reliable trends in the importance of various terms (Fig. 3.3c). We refit these terms progressively in order of the impact on loss across all training runs and found a sharp elbow at 3 terms: a temporally asymmetric Hebbian learning term, $\tilde{x}_i x_j$ (blue pentagon in Fig. 3.3), its complement $x_i \tilde{x}_j$ (orange square), and a term second order in the postsynaptic activity multiplied by the synapse size, $w_{ij}\tilde{x}_j x_j$ (purple triangle). The first term was consistently learned with a positive coefficient while the latter two were always almost negative (Fig. 3.3a), rendering the effective learning rule

$$\dot{w}_{ij} = c_0 \tilde{x}_i x_j - c_1 x_i \tilde{x}_j - c_2 w_{ij} \tilde{x}_j x_j. \tag{3.4}$$

where all $\{c_i\}$ are non-negative and time constants are different for each term. The 3 most important terms can be interpreted as a temporally asymmetric generalization of Oja's rule in that the Hebbian learning term, $x_i x_j$, is replaced by the first two terms in Eq. 3.4, which depend on the relative timing of x_i and x_j. The time constant of the third term was on average short (∼1 ms), making this term akin to $w_{ij} x_j^2$, the normalizing term of Oja's rule.

3.2.2 Biological noise alters the learned plasticity rules

We next asked whether ongoing disruptions to network structure alter which plasticity rules are meta-learned. To explore this, we introduced synaptic turnover to the simulation phase of the meta-learning loop. Synaptic turnover is a stochastic process by which existing synapses disappear and new, small synapses emerge (Fig. 3.1e). Prior to each network activation, all connections were updated according to

$$w_{ij} \leftarrow \begin{cases} 0 & |w_{ij}| > 0 \quad \text{and} \quad x_{\text{ST}} < p_{\text{ST}} \\ w_{ij} & |w_{ij}| > 0 \quad \text{and} \quad x_{\text{ST}} \geq p_{\text{ST}} \\ \epsilon & |w_{ij}| = 0 \quad \text{and} \quad x_{\text{ST}} < p_{\text{ST}} \\ 0 & |w_{ij}| = 0 \quad \text{and} \quad x_{\text{ST}} \geq p_{\text{ST}}, \end{cases} \tag{3.5}$$

where $x_{ST} \sim U[0,1]$, p_{ST} is the probability of single synapse turnover per activation, and ϵ is a small positive (negative) constant if the presynaptic cell is excitatory (inhibitory). Since plasticity is unable to act on connections of size 0 (see Eq. 3.1), synaptic turnover determines the set of synapses available to the plasticity rule. To ensure learned rules were robust to a spectrum of rates of synaptic turnover, only half the networks used to evaluate the batch loss underwent this process.

We found that meta-learned rules were able to organize persistent representations of time despite synaptic turnover, with performance near that of rules learned on unperturbed networks. Our term-sensitivity analysis (Fig 3.4a) showed that solutions again heavily depended on temporally asymmetric Hebbian learning, i.e. $\tilde{x}_i x_j$ (light blue hexagon), and the bound on postsynaptic activity, $w_{ij}\tilde{x}_j x_j$ (purple triangle); however, we frequently found dependence on two additional terms: one that constantly strengthened all synapses (dark blue pentagon) and an activity bound independent of synapse size (orange square). Refitting the plasticity rule in order of impact on loss demonstrated that these 4 terms recapitulated most of the success of the learned solutions. The effective rule may be written as

$$\dot{w}_{ij} = c_0 + c_1 \tilde{x}_i x_j - c_2 \tilde{x}_j x_j - c_3 w_{ij} \tilde{x}_j x_j, \tag{3.6}$$

where again all $\{c_i\}$ are non-negative. Synapses for which postsynaptic activity, x_j, remains chronically small will be potentiated by c_0; however, adding this potentiating term also necessitates an activity bound that does not scale with synapse size, such as the term with prefactor c_2 in Eq. 3.6. To understand this impact of this term, consider a synapse between two neurons whose typical firing times are far apart, i.e. $\tilde{x}_i x_j$ is nearly zero, the fixed point under this rule is

$$\langle w_{ij} \rangle = \max\left(\frac{c_0}{c_3 \langle \tilde{x}_j x_j \rangle} - \frac{c_2}{c_3}, 0 \right) \tag{3.7}$$

if $\langle \tilde{x}_j x_j \rangle > 0$, where $\langle \cdot \rangle$ denotes the time average. Thus, a large enough choice of c_2 prevents every synapse in the network from growing, enforcing sparsity and decreasing the risk of a neuron changing its firing time upon loss of its original inputs. We additionally note that Eq. 3.6 does not contain $-x_i \tilde{x}_j$, which appeared in the reduced rule learned in the absence of synaptic turnover (Eq. 3.4). This may be because the roles of $-x_i \tilde{x}_j$ and $-\tilde{x}_j x_j$ are partially redundant: both terms can suppress synapses that run counter to the sequential

dynamics in the network. When constant potentiation of all synapses occurs, the term in $\tilde{x}_j x_j$ is preferable as it offsets constant potentiation of all synapses. In the unperturbed context, where constant synaptic growth is unnecessary, the term in $\tilde{x}_j x_j$ is problematic in that it can set all afferent synapses to a driven neuron to zero (whereas $-w_{ij}\tilde{x}_j x_j$ cannot). Thus, the term in $x_i \tilde{x}_j$ becomes preferable in the unperturbed context.

Comparison to existing models of sequence formation

Do these discovered learning rules more robustly encode timing than previously proposed rules when the circuit is disrupted with biologically relevant noise? We hypothesized this would be true given that discovered rules do not impose hard bounds on the size of single synapses, total synaptic strength onto a neuron, nor number of synapses, as other models of sequence formation have [100, 101, 104, 106]. The absence of these constraints permits compensatory rescaling of synapses in response to disruptions. We compared meta-learned rules trained with and without synaptic turnover to a previously proposed sequence learning rule that used multiplicative asymmetric Hebbian learning, a single synapse bound, and bounds on the total strength of synapses onto and out of individual neurons [101]. Each rule was applied to 100 test networks for 400 activations without disruption. Following this, a decoder was constructed to read out time from the neural activity, and then synapses in the networks were turned over during an additional 150 activations, after which the loss was evaluated using the constructed decoder (Fig. 3.4c). We compared the best versions of each rule when synapses were turned over (with probability $p_{ST} = 0.00072$) at each activation, equivalent to a 90% probability of individual synapse survival during the disruption period (Fig. 3.4d, light purple points), and when they were not turned over (Fig. 3.4d, grey points; see Supp. Sec. 3.8). We used the Kruskal-Wallis H test to test for equality of medians. We found that meta-learned rules trained with and without synaptic turnover outperformed the rule based based on rigid synapse constraints ($p = 2.5\,\mathrm{x}\,10^{-8}$, Cohen's $d = -0.45$ and $p = 7.2\,\mathrm{x}\,10^{-7}$, Cohen's $d = -0.44$, respectively), while in the absence of perturbation, medians were not distinct after 4-fold Bonferroni correction ($p = 0.016$, Cohen's $d = -0.29$ and $p = 0.029$, Cohen's $d = -0.30$, respectively). When we studied the connectivity structure

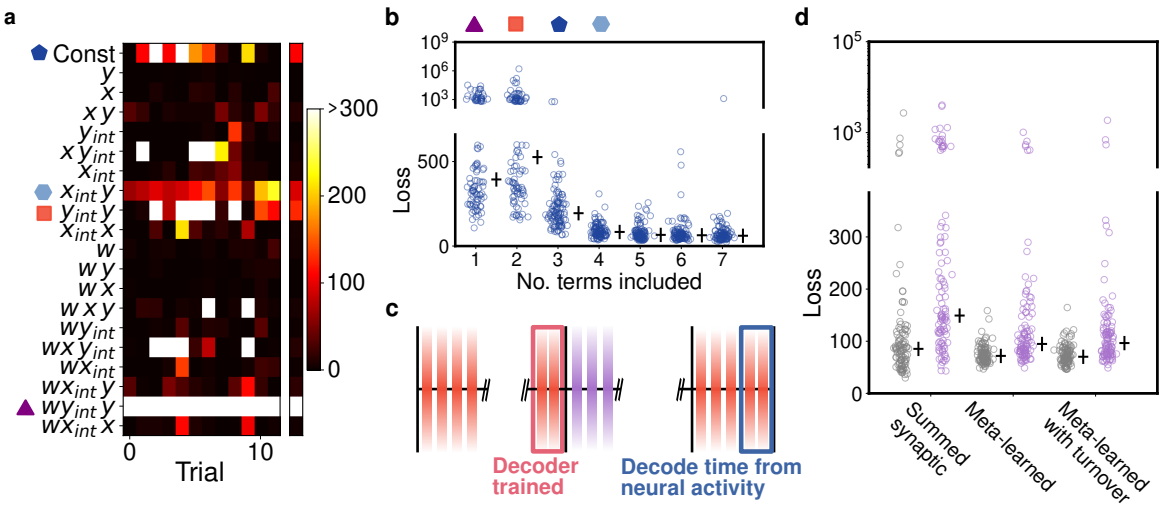

Figure 3.4: **Discovered rules organized dense feedforward structures (a)** Median impact of loss when coefficient of given plasticity is set to zero for a given learned solution. **(b)** Refitting terms in order of median impact on loss shows sequence generation in synaptic turnover context is well captured by 4 terms. **(c)** Task structure for comparison between different rules. Rules are given 400 activations to organize dynamics. At the end, the decoder is constructed. Networks then undergo synaptic perturbation for 150 activations. Finally, the decoder attempts to decode time from the resulting dynamics to determine the loss. **(d)** Comparison between multiplicative Hebbian learning and summed synaptic bound, meta-learned without synaptic turnover, and meta-learned with synaptic turnover on time encoding task illustrated in (c). Trials shown in grey do not include synaptic turnover; light purple include synaptic turnover. Black crosses indicate median values for each condition.

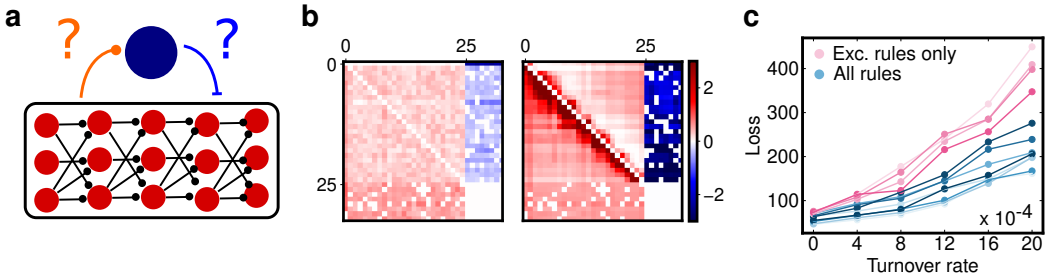

Figure 3.5: **Learning plasticity on all synapses** **(a)** Schematic of synapses upon which we newly allow plasticity. **(b)** Weight matrices on activations 1 and 400 of a network evolving under plasticity rules learned on all 3 sets of synapses. I synapses increase. **(c)** Comparison of performance of rules learned only on E→E synapses (red) (N=5) versus all sets of synapses (blue) (N=8) when training includes synaptic turnover across varying rates of synaptic turnover.

of networks organized by meta-learned rules, we found that the discovered rule generated denser feedforward connectivity in comparison to other plasticity rules with alternative forms of Hebbian learning and heterosynaptic competition (see Supp. Sec. 3.9).

3.2.3 Including inhibitory plasticity

Meta-learning allows the exploration of multiple plasticity rules operating on distinct sets of synapses within the same circuit, as might arise if there are multiple cell types [153, 154]. In particular, I→E plasticity has been the focus of much recent work [155–160]. The interaction of many plasticity rules is challenging to analyze theoretically, but meta-learning allows exploration of these interactions [54, 60]. Recent evidence suggests there may be multiple forms of plasticity within sequence-generating circuits: Wang et al. [67] increased the intrinsic excitability of *in vivo* $\text{HVC}_{(\text{RA})}$ neurons and found that the strength of both E→I and I→E connections could dynamically shift in response. Targeted cells received increased total inhibitory synaptic strength and decreased excitatory strength. To understand the role of this additional plasticity, we meta-learned plasticity on all sets of synapses within the circuit

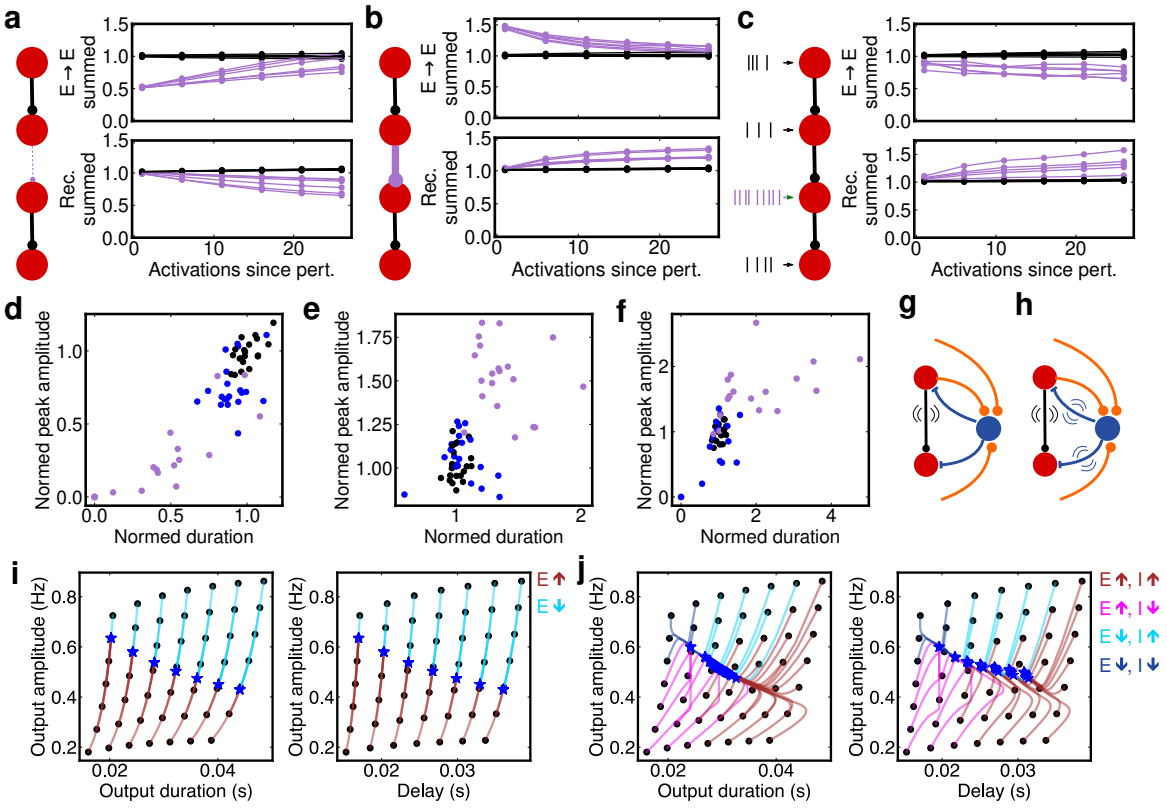

Figure 3.6: **Network perturbations reveal homeostatic compensation in E→E and I→E synapses**

Figure 3.6: **Network perturbations reveal homeostatic compensation in E→E and I→E synapses (a)** Response of the magnitude of summed E→E weights $|W_i^{\text{exc}}| = |\sum_k w_{k,i}^{E\to E}|$ (top) and summed recurrent weight $|W_i^{\text{rec}}| = |\sum_k w_{i,k}^{\text{E}\to\text{I}} w_{k,i}^{\text{I}\to\text{E}}|$ (bottom) to the imposed scaling down of E→ E weights to a single cell. Values for neuron with perturbed inputs shown in light purple; all others black. Each line represents average values for a single learned rule over N=20 networks. **(b)** Same as (a), but for imposed upscaling of E→ E weights to one cell (light purple). **(c)** Same as (a), but frequency of stochastic inputs to one cell is greatly increased (light purple) relative to all other cells (black). **(d-f)** Postsynaptic response of affected neurons in 20 networks plotted in the space of normalized peak amplitude and duration for manipulations (a-c). Black dots represent pre-perturbation responses; light purple, immediately after perturbation; blue, 25 activations after perturbation. **(g)** Schematic of feedforward motif maintaining homeostasis via its E input. **(h)** Schematic of feedforward motif with homeostasis on E→E and I→E synapses. **(i)** Phase flow in space of output firing envelope duration and amplitude (left) for a single neuron. Black points are first responses of the neuron for various inputs; blue stars represent responses after 200 activations. Note: initial and final output durations and delays are roughly the same. Colors of trajectories indicate how the neuron's input weight changes to satisfy the plasticity rule: see legend. **(j)** If the forms of homeostasis on E→E and I→E maintain different aspects of the postsynaptic response, an attractor forms in the duration and peak amplitude of the postsynaptic response (left). The response delay is also constrained (right).

(Fig. 3.5a) both with and without synaptic turnover. Specifically, we attempted to learn three independent plasticity rules that operated on three distinct groups of synapses (E→E, E→I, and I→E; see Section 3.10).

E→I and I→E synaptic plasticity improves decoding of time, particularly in the presence of synaptic turnover

Meta-learning uncovered triplets of plasticity rules that successfully organized initially random networks into sequence generators (Fig. 3.5b). When trained with turnover in E→E synapses, solutions that acted upon E→I and I→E synapses in addition to E→E outperformed solutions that only acted upon E→E connections, particularly when the rate of synaptic turnover was high (Fig. 3.5c), suggesting this additional plasticity played an important role in maintaining network dynamics through perturbation.

To investigate how rules acting on all synapses generated improved time representations, we repeated the dropout analysis. We found that E→E plasticity within these learned triples was largely similar to the rules previously learned on E→E synapses alone (Eqs. 3.4 and 3.6): solutions were consistently sensitive to the removal of $\tilde{x}_i x_j$ and $w_{ij}\tilde{x}_j x_j$, which appeared consistently with positive and negative coefficients, respectively. Further, dependence on these terms persisted when we trained networks with turnover on E→E synapses or I→E synapses (Supp. Fig. 6). As expected, we also found that solutions depended heavily upon terms that acted upon E→I and I→E synapses, but this plasticity was more difficult to interpret due to increased trial to trial variability in the discovered rules. We found, however, that the E→I plasticity rule consistently depended on the second order presynaptic term $\tilde{x}_i x_i$, which always appeared with a positive coefficient, suggesting that E cells project to inhibitory counterparts with a strength that increases with the E cell level of activity. An implication of dependence on this term is that the strength of an E neuron's recurrent inhibition, defined as $|W_i^{\text{rec}}| = |\sum_k w_{i,k}^{\text{E}\to\text{I}} w_{k,i}^{\text{I}\to\text{E}}|$, where i is the index of the E cell and k indexes the I cells to which it projects, might depend on its level of activity. Thus, ablation of excitatory inputs to an E cell might cause its recurrent, disynaptic inhibition to lower in a manner that homeostatically restores its firing.

Recurrent inhibition of E neurons is homeostatic in networks with learned plasticity

As the role of E→I and I→E plasticity was not completely clarified by our perturbations so far, we next investigated how this plasticity adjusted synapses coupled to an individual E cell when its typical input was manipulated. Noting the dependence of the E→I plasticity rule on $\tilde{x}_i x_i$, we hypothesized that a targeted neuron's recurrent inhibition and excitatory afferents might be adjusted in concert to restore its typical firing pattern. We used discovered learning rules to organize sequences and then performed three varieties of *in silico* manipulations of individual E cells within these networks. In the first, we scaled down the excitatory afferents to the targeted E cell by 50% (Fig. 3.6a, diagram). In the second, we scaled up the same connections by 50% (Fig. 3.6b). In the last, we increased the rate of the targeted cell's input Poisson process by a factor of 10 (Fig. 3.6c). This final manipulation mirrored the viral insertion of NaChBac to $\mathrm{HVC_{RA}}$ neurons in Wang et al., which causes these cells to become hyper-excitable [67, 161, 162]. Wang et al. found that manipulated cells recruited additional inhibition and weakened excitatory afferents.

We found that scaling down a targeted E cell's excitatory afferents resulted in a rescaling of those excitatory weights and an accompanying decrease in $|W_i^{\mathrm{rec}}|$ (Fig. 3.6a). We further found that these synaptic changes restore the initial firing pattern of the E cell: in Fig 3.6d, we plot the initial responses of targeted E cells across 20 self-organized sequences in the space of duration and peak amplitude of response (black points). Scaling down the excitatory afferents initially causes both the peak amplitude and duration of response to decrease (Fig. 3.6d; light purple points), but these responses largely recover after ∼25 activations of the network (Fig. 3.6d; blue points). When we instead strengthened E afferents to the targeted cell, we found these connections weakened and recurrent inhibition strengthened in compensatory fashion (Fig. 3.6b). In these networks, we found that targeted E cell responses that were initially lengthened with increased peak amplitude (Fig. 3.6e; light purple points) were reduced back to their pre-perturbation values of duration and amplitude (Fig. 3.6e; black points show initial responses; blue points, responses after 25 activations). Finally, rendering the targeted E cell hyper-excitable caused it to scale down its excitatory afferents and increase its recurrent inhibition in a manner that restored its typical firing pattern. In

summary, plasticity rules on excitatory afferents and recurrent inhibition operate in tandem to maintain the firing pattern of the neuron (see Supp. Sec. 3.11 for additional details).

Two forms of homeostasis create an attractor in postsynaptic response and timing

How might modulation of recurrent inhibition contribute differently to the modulation of excitation in preserving network dynamics? We hypothesized that distinct plasticity on different sets of synapses might confer a robust representation of time if these rules governed distinct aspects of the desired network function. For instance, within our HVC-like circuit, peak firing of E neurons might be controlled by E→E plasticity while total activity might be controlled by I→E plasticity. Since excitatory and inhibitory inputs to E cells differ in their timescales (E inputs are transient while I inputs are relatively tonic), we reasoned plasticity on both sets of synapses should increase the control of the postsynaptic response. While it is already known that multiple plasticity mechanisms can sharpen the responses of neurons to stimuli [155], prior work has not addressed whether similar plasticity might be leveraged to preserve the timing of such responses, which is crucial to timing representations. To explore this, we constructed a simplified model of a single neuron responding to a broad range of excitatory inputs, characterized by varying peak amplitudes and durations, and a tonic inhibitory input. We compared the responses of the neuron when both the excitatory and inhibitory input synapses (Fig. 3.6h) evolved under plasticity rules to responses produced when only the excitatory synapse evolved under a plasticity rule (Fig. 3.6g) and found the two rule model was able to better constrain the duration, peak amplitude, and timing of the postsynaptic response the E neuron (Fig. 3.6i-j; see Supp. Sec. 3.12 for full description of reduced model).

3.3 Discussion and limitations

Meta-learning plasticity rules via stochastic optimization is a promising technique, but suffers a number of limitations. One, the optimization process becomes expensive as the size of the rule basis, the number of neurons in the network, and amount of simulation time required grows. Further, CMA-ES may require many epochs to converge on good solutions. Training E→E plasticity (plasticity on all synapses) across 10 networks in batch typically required 24

(72) hours of compute on 30 Cascade Lake or Ice Lake Intel CPU cores to yield reasonable solutions. Two, meta-learning tends to generate different solutions based on the seed; due to the expensive nature of each trial, we did not carry out enough trials to claim full knowledge of the solution space. Three, the plasticity rules learned were quite dense in our choice of basis, limiting interpretation, and we ultimately employed perturbations to better understand the critical terms. Four, the choice of basis limits the space of discoverable rules; for instance, we did not include feedback-modulated plasticity in this study. Five, the initial connectivity of the circuit likely has a strong bearing on the sort of plasticity that successfully can leverage it [163, 164].

Prior to any experience, intrinsic, self-organized dynamics within the brain can serve as powerful priors that can accelerate and shape successful, feedback driven learning. In this work, we study how one such computational primitive could emerge by adapting a meta-learning procedure to learn the learning rules that self-organize and maintain robust representations of time in neural dynamics. Meta-learning discovers a temporally asymmetric (STDP-like) generalization of Oja's rule that organizes and maintains sparse, sequential activity out of initially random connectivity, which outperforms other models of sequence generation in the presence of synaptic turnover by permitting flexible rescaling of inputs to restore dynamics. Additionally we found that plasticity rules learned on all sets of synapses outperform plasticity rules applied only to E connections. Through a toy model, we show how plasticity on all synapses could confer extra timing stability if the plasticity in distinct sets of synapses act on different moments of an E neuron's activity.

In this work, we have developed a paradigm to understand how computational primitives might self-organize within neural circuits and selected sequences as our test example. Future work could study the emergence of other canonical forms of neural dynamics that have been widely identified in brain activity and serve as fundamental components of computation, such as line attractors or limit cycles, or how plasticity rules generating such components interact with rules requiring feedback.

3.4 Acknowledgements and disclosure of funding

We would like to thank Carlos Lois, Zsofia Török, Bo Wang, Patrick Zhang, and Leenoy Meshulam for useful discussions. This work was supported by the Simons Collaboration for the Global Brain and an NIH BRAIN grant (5R01NS104925). This work was facilitated through the use of advanced computational, storage, and networking infrastructure provided by the Hyak supercomputer system and funded by the STF at the University of Washington.

3.5 Network simulations

3.5.1 Network initialization

For all network simulations, network were initialized with excitatory units connected all-to-all with weights drawn from $U[0.2\, w^{(0)}_{E\to E}, w^{(0)}_{E\to E}]$. E→I and I→E weights were nonzero with probability 0.8 with values drawn from $\mathcal{N}(w^{(0)}_{E\to I}, 0.3\, w^{(0)}_{E\to I})$ and $\mathcal{N}(w^{(0)}_{I\to E}, 0.3\, |w^{(0)}_{I\to E}|)$, respectively. Any weights violating Dale's law were set to zero.

3.5.2 Single neuron dynamics

All neurons were modeled as rate-based units with threshold linear activations, with continuous firing rates evolving according to

$$x_j(t) = [V_j(t) - b]^+ , \tag{3.8}$$

where $V_j(t)$ evolved via

$$\tau_m \dot{V}_j(t) = -V_j(t) + \sum_i w_{ij} x_i(t) \tag{3.9}$$

Here, w_{ij} is the weight of the synapse $i \to j$, τ_m is the membrane time constant, and b is the bias. Values of τ_m differed across cell types (E and I), but were consistent within cell type. Eqs. 3.8 and 3.9 were solved numerically with timestep size $\delta t = 0.1$ ms.

Table 3.1: Environment parameters

Parameter	Description	Value
$w^{(0)}_{E\to E}$	Maximum initial E→E weight	0.4
$w^{(0)}_{E\to I}$	Mean initial E→I weight	0.5
$w^{(0)}_{I\to E}$	Mean initial I→E weight	-0.3
$\tau_m^{(E)}$	Excitatory cell membrane time constant	10 ms
$\tau_m^{(I)}$	Inhibitory cell membrane time constant	0.1 ms
b	Excitatory cell resting potential	0.1

3.5.3 Network input

We adopt the assumptions made in Jun and Jin [100] and Tupikov and Jin [106] and drive a single neuron with a strong kick of excitation at the start of each trial. Following this, all other neurons in the network received two filtered Poisson process inputs, a "frozen" (stereotyped activation to activation) input and a stochastic input. These processes had rates $\lambda_{\text{frozen}} = \lambda^{(0)} p_{\text{fix}}$ and $\lambda_{\text{stoch}} = \lambda^{(0)}(1 - p_{\text{fix}})$, respectively, where $p_{\text{fix}} = 0.75$ and $\lambda^{(0)} = 80$ Hz. Arrivals from both processes were then smoothed into an input current via an alpha function kernel such that arrivals $\{t_i\}$ became input current

$$I(t) = \sum_{\{t_i\}} ae \frac{t - t_i}{\tau_{\text{input}}} e^{-(t-t_i)/\tau_{\text{input}}} \Theta(t - t_i), \tag{3.10}$$

where $\tau_{\text{input}} = 3$ ms and a was 0.09 and 0.02 for excitatory and inhibitory neurons, respectively. Input current to the initially driven neuron was modeled in a similar fashion: a block of 10 contiguous arrivals was transformed into a current via Eq. 3.10.

3.6 Meta-learning basis

To discover plasticity rules, we leverage a meta-learning approach pioneered by Bengio et al.[52] and more recently extended by Confavreux et al[51]. Specifically, we parameterize a set of plasticity rules with coefficients c_k and time constants τ_k, such that individual synapses

evolve according to

$$\dot{w}_i = \sum_k c_k F_k(x_i, y_i, w_i, \tau_k) \tag{3.11}$$

where x_i and y_i are the pre and postsynaptic activity of synapse i, respectively, w_i is the synapse size, and F_k is the k^{th} term in the plasticity rule. The terms of $\mathbf{F}$ are

$$\mathbf{F}(x_i, y_i, w_i) = \left\{ \begin{array}{c} 1 \\ y_i \\ x_i \\ x_i\, y_i \\ \tilde{y}_i^{(4)} \\ x_i\, \tilde{y}_i^{(5)} \\ \tilde{x}_i^{(6)} \\ \tilde{x}_i^{(7)}\, y_i \\ \tilde{y}_i^{(8)}\, y_i \\ \tilde{x}_i^{(9)}\, x_i \\ w_i \\ w_i y_i \\ w_i x_i \\ w_i x_i\, y_i \\ w_i \tilde{y}_i^{(14)} \\ w_i x_i\, \tilde{y}_i^{(15)} \\ w_i \tilde{x}_i^{(16)} \\ w_i \tilde{x}_i^{(17)}\, y_i \\ w_i \tilde{y}_i^{(18)}\, y_i \\ w_i \tilde{x}_i^{(19)}\, x_i \end{array} \right\}, \tag{3.12}$$

where

$$\tilde{x}_i^{(k)}(t) = \int_0^t e^{-(t-t')/\tau_k} x_i(t') dt', \tag{3.13}$$

$$\tilde{y}_i^{(k)}(t) = \int_0^t e^{-(t-t')/\tau_k} y_i(t')dt', \tag{3.14}$$

and τ_k is a learned time constant specific to element k of $\mathbf{F}$. Note, we rescaled terms that were higher order in pre- or postsynaptic activity by a factor 0.1 per order in order to limit the number exploding simulations.

3.7 Optimization via CMA-ES

To optimize rule coefficients and time constants, we used Covariance Matrix Adaptation (CMA-ES) [152], a technique that iteratively samples from a high dimensional Gaussian placed on the space of plasticity rule parameters. The mean of the Gaussian was initialized to 0 for all rule coefficient parameters and 5 ms for all time constants. The covariance matrix was initially diagonal with the standard deviation set to 0.003 and 3 ms for rule coefficients and time constants, respectively. We found that initializing the covariance matrix with small values relative to the search space led to fewer simulations in which weight growth was unbounded. The search space for rule coefficients and time constants was bounded to $[-10, 10]$ and $[0.05, 40]$ ms, respectively. The population size for each epoch of training was typically 14.

Each evaluation of a candidate learning rules was performed using the aggregate loss from 10 randomly initialized networks (see Supp. Section 3.5.1) to encourage CMA-ES to find plasticity rules that gracefully generalized to unobserved initial network structures and inputs. Importantly, the seeds of these simulation were held fixed from trial to trial; in a separate set of experiments, we attempted to draw a fresh set of networks and inputs for evaluation of the loss function, but found this drastically slowed learning.

3.8 Comparison to existing model of sequence formulation

We compared discovered rules for sequence generation to a prior plasticity rule introduced by Fiete et al. [101]. The rule relies on spike timing dependent plasticity and bounds on both individual synapses as well as bounds on the total synaptic strength onto and out of a neuron. While the rule was originally formulated for both binary neurons and spiking networks, we

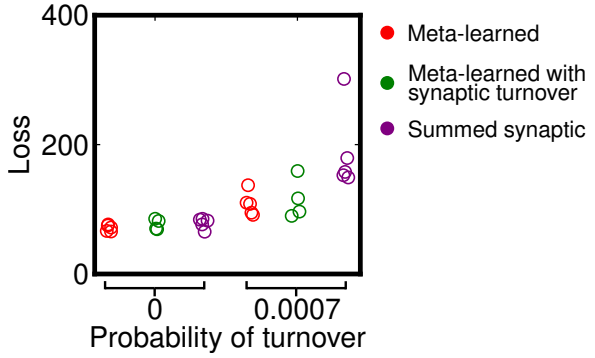

Figure 3.7: Five best versions of meta-learned rules trained with and without synaptic turnover compared to summed synaptic learning rule. Points represent median loss for each condition.

adapt it here for continuous firing rate units. Under this rule, synapses evolve according to

$$\dot{w}_{ij} = \alpha w_{ij}(\tilde{x}_i x_j - x_i \tilde{x}_j) - \alpha\beta \left(\left[\sum_k w_{kj} - W_{\max} \right]^+ + \left[\sum_{k'} w_{ik'} - W_{\max} \right]^+ \right), \qquad (3.15)$$

where all w_{ij} are additionally bounded by some $w_{\max} < W_{\max}$. The first term in this rule implements a multiplicative (in that potentiation and depression are proportional to the synapse size) spike timing-dependence plasticity kernel (or the continuous firing rate analog). The second term bounds the sum of all synapses onto a neuron and out of a neuron to $W_{\max}$. Fiete et al. [101] demonstrated this rule generates wide synfire chains if groups of N neurons receive correlated noise inputs and $w_{\max} = W_{\max}/N$.

To compare Eq. 3.15 to the discovered rule, we first optimized $(\alpha, \beta, w_{\max}, W_{\max})$ to minimize the decoder error (first term of Eq. 2 in the main text) in a setup identical to that described in Sec. 3.5.1. Once optimized, we used discovered rules trained with and without synaptic turnover and Eq. 3.15 to organize sequence generating networks by activating networks 400 times; 6 activations from the end of this period were used to train a linear decoder without regularization that attempted to decode time from the neural dynamics at 500 time points throughout each activation. Following this, networks underwent synaptic

turnover for an additional 150 activations; synaptic turnover then ceased, and networks were given an additional 50 activations, after which the trained decoder's performance was evaluated on 200 time points from an additional 6 activations of the network. Results are shown in Supp. Fig. 3.7 for the best 5 learned rules of each class.

3.9 Learned rule generates dense feedforward connectivity relative to alternatives

Temporally asymmetric Hebbian learning and a mechanism that induces competition between synapses that drive the same postsynaptic target is a common recipe for sequence organization [101, 104, 106, 107]. Meta-learning selected a particular form of these two components both when synaptic turnover was included in training and when it was not. This led us to study the connectivity of sequences organized under different possible choices of these components (asymmetric Hebbian learning and synaptic competition) and compare them to those of the meta-learned rule. We fit the following learning rules using CMA-ES. In general, all coefficients and time constants were free parameters tuned by the model to produce the best encoding of time. The rules were

$$\dot{w}_{ij} = \tilde{x}_i x_j - x_i \tilde{x}_j - \beta \left(\left[\sum_k w_{kj} - W_{\max} \right]^+ + \left[\sum_{k'} w_{ik'} - W_{\max} \right]^+ \right), \tag{3.16}$$

$$\dot{w}_{ij} = w_{ij}(\tilde{x}_i x_j - x_i \tilde{x}_j) - \beta \left(\left[\sum_k w_{kj} - W_{\max} \right]^+ + \left[\sum_{k'} w_{ik'} - W_{\max} \right]^+ \right), \tag{3.17}$$

$$\dot{w}_{ij} = \tilde{x}_i x_j - x_i \tilde{x}_j - \tilde{x}_j x_j, \tag{3.18}$$

$$\dot{w}_{ij} = w_{ij}(\tilde{x}_i x_j - x_i \tilde{x}_j) - w_{ij} \tilde{x}_j x_j, \tag{3.19}$$

and

$$\dot{w}_{ij} = \tilde{x}_i x_j - x_i \tilde{x}_j - w_{ij} \tilde{x}_j x_j. \tag{3.20}$$

We fit Eqs. 3.16 and 3.17 both with and without a single synapse bound. The results are shown Fig. 3.8. We found that Eqs. 3.16 and 3.17 only grew dense feedforward structures

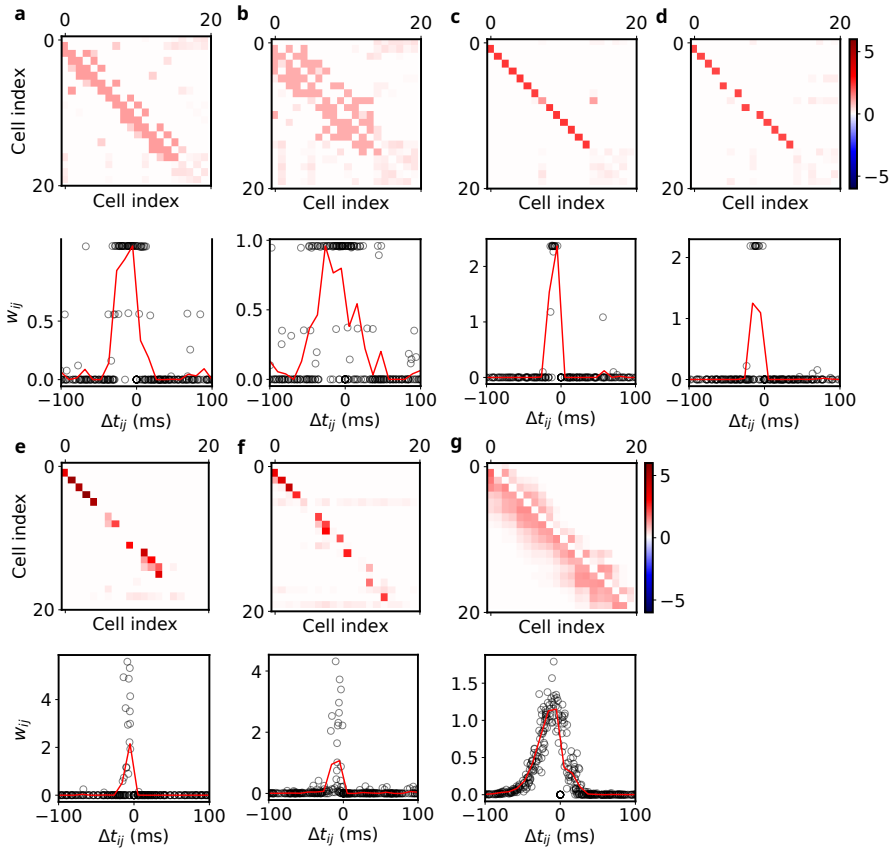

Figure 3.8: **Discovered rule organizes dense feedforward structures**

Typical connectivity for sequences grown with different selections of Hebbian learning and heterosynaptic bound. **(a)** Connectivity matrix organized by additive Hebbian learning, summed synaptic bound, and a single synapse bound (top). Weight as a function of delay in firing time for all synapses in the network (bottom). Individual instances shown in black; smoothed averages shown in red. Note: most nonzero synapses are at single synapse bound. **(b)** Same as (a), but with multiplicative Hebbian learning. **(c)** Same as (a), but with single synapse bound removed; sequences become 1D, i.e. each neuron receives exactly one input and projects to exactly one downstream neuron. **(d)** Same as (c), but with single synapse bound removed. Sequences become 1D. **(e)** Connectivity organized by equation 10. **(f)** Connectivity organized when both Hebbian learning and firing rate bound are multiplicative. **(g)** Connectivity organized by the discovered rule, Eq. (6). Connectivity is dense and weight size depends on the relative time delay between firing times.

when a single synapse bound was imposed (Fig. 3.8, a-b). Removal of this bound generated 1D chains (Fig. 3.8, c-d). Eqs. 3.18 and 3.19 were generally less successful in organizing sequential activity, but grew structures where neurons received sparse synapses.

To explain these connectivity patterns, we considered the fixed points of individual synapses within them operating under the relevant rules. We first consider the effective learned rule of the meta-learned plasticity, Eq. 3.20. For w_{ij} to be fixed, we require $\langle \dot{w}_{ij} \rangle = 0$. Assuming synaptic evolution is slow, we find

$$\langle w_{ij} \rangle = \frac{\langle \tilde{x}_i x_j \rangle - \langle x_i \tilde{x}_j \rangle}{\langle \tilde{x}_j x_j \rangle}. \tag{3.21}$$

The synapse size is directly proportional to average potentiation under the Hebbian learning, as in Oja's rule, leading to connectivity shown in Fig. 3.8g.

We now consider what occurs if the postsynaptic bound does not scale with the synapse size, i.e. Eq. 3.18. Here, the potentiation under Hebbian learning must be equal on average to the postsynaptic activity bound, which is identical for all synapses that drive j. Thus the competition between synapses is winner-take-all, unless multiple synapses receive identical potentiation. We speculate that there are in general few presynaptic firing times relative to the postsynaptic firing times that receive equal potentiation under the first two terms of Eq. 3.18, thus sequences formed by this rule tend to be sparse (Fig 3.8c). The same argument applies to Eq. 3.16 when the single synapse bound is not imposed and for Eq. 3.19: assuming the synapse is nonzero, 3.19 may be rewritten as

$$\frac{d \log(w_{ij})}{dt} = \tilde{x}_i x_j - x_i \tilde{x}_j - \tilde{x}_j x_j, \tag{3.22}$$

which, after setting the right side to zero, is identical to the preceding case. In summary, that the last term in Eq. 3.20 is of greater order in the synapse size than the first two terms is important for generating dense connectivity without the use of single synapse bounds.

3.10 Meta-learning plasticity on all synapses

For completeness, we include the results of term-sensitivity analyses performed on solutions learned on unperturbed networks, networks with E→E turnover, and I→E turnover (Fig. 3.12). Note that E→E plasticity was largely similar to rules trained only on E→E synapses.

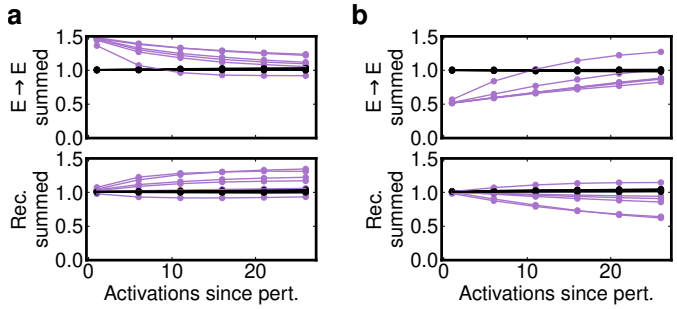

Figure 3.9: **Synaptic changes under rules learned with E→E turnover**
(a) Response of the magnitude of summed E→E weights $|W_i^{\text{exc}}| = |\sum_k w_{k,i}^{E\to E}|$ (top) and summed recurrent weight $|W_i^{\text{rec}}| = |\sum_k w_{i,k}^{\text{E}\to\text{I}} w_{k,i}^{\text{I}\to\text{E}}|$ (bottom) to the imposed scaling up of E→ E weights to a single cell. Values for neuron with perturbed inputs shown in light purple; all others black. Each line represents average values for a single learned rule over N=20 networks. **(b)** Same as (a), for imposed scaling down of E→ E weights to a single cell.

Formally, synapses of projection group G evolved according to

$$\dot{w}_{ij}^{(G)} = \Theta(|w_{ij}^{(G)}|) \sum_k c_k^{(G)} F_k^{(G)}(x_i, x_j, w_{ij}^{(G)}, \tau_k^{(G)}), \tag{3.23}$$

where $F_k^{(G)}$ is the k^{th} term of the plasticity rule acting on group $G \in$ {E→E, E→I, I→E}. i (j) indexes the presynaptic (postsynaptic) neurons of G.

3.11 Single neuron perturbation experiments in networks with plasticity on all synapses

We performed three manipulations on 20 test networks under twelve different learned plasticity rules: six learned under E→E synaptic turnover and six under I→E synaptic turnover. In each manipulation, only the afferent synapses or inputs to a single neuron were modified. The targeted cell was always excitatory and randomly chosen. The three manipulations were:

1. Scaling down the excitatory afferents to an E cell by 50%.

2. Scaling up the excitatory afferents to an E cell by 50%.

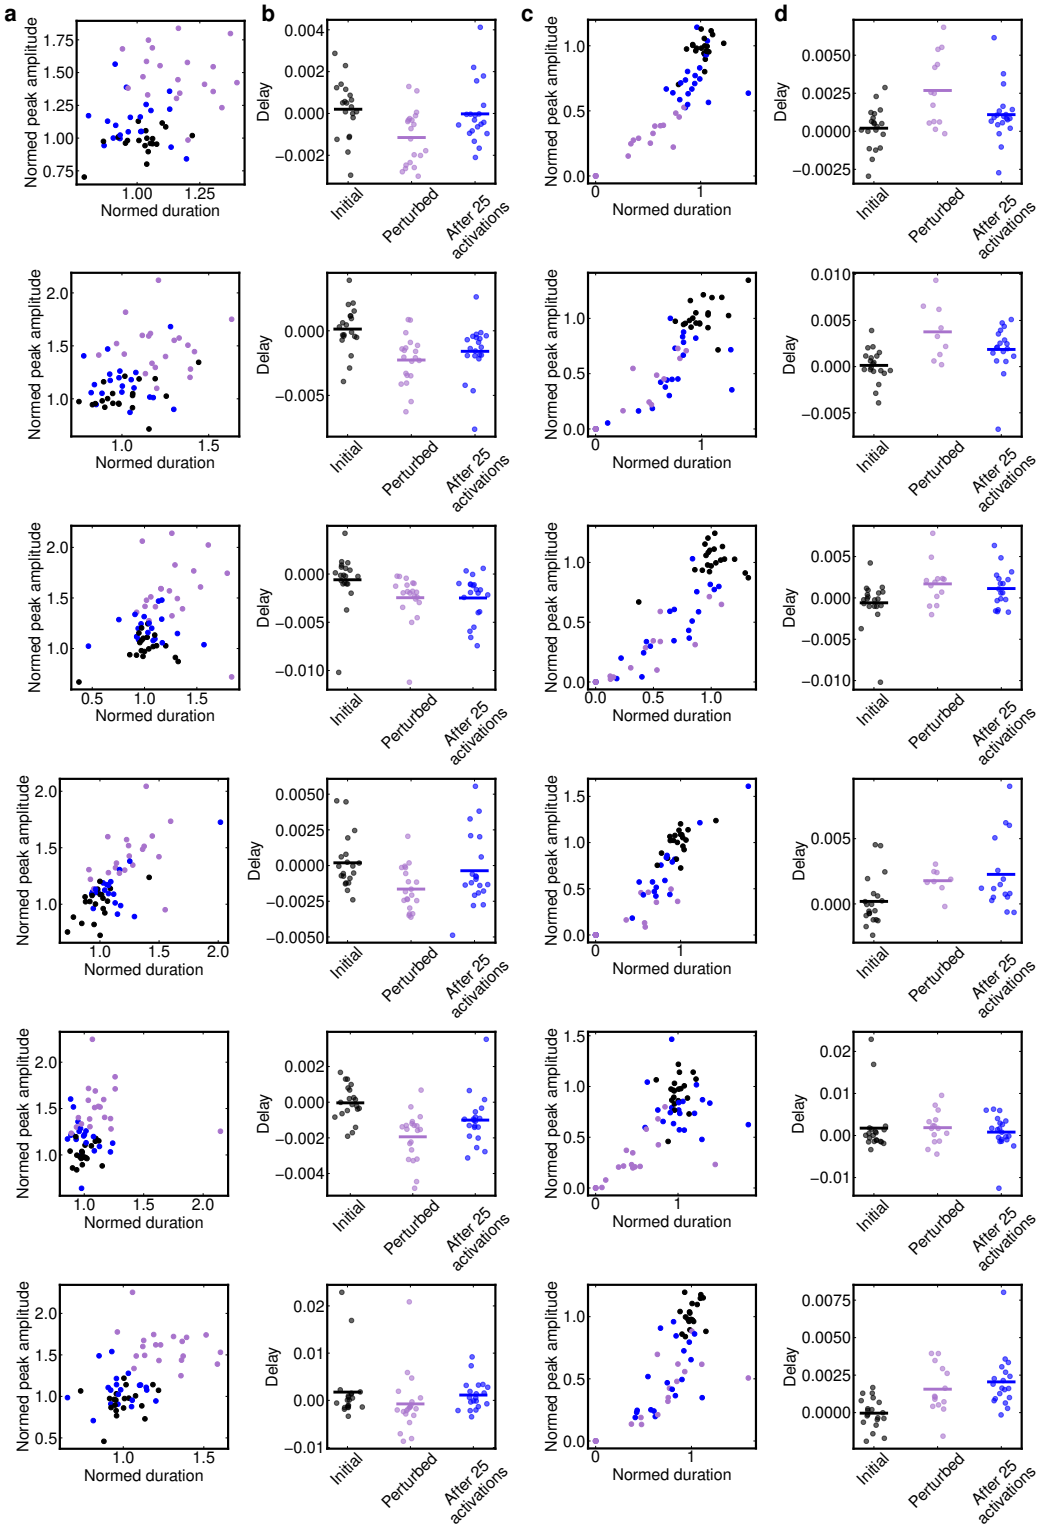

Figure 3.10: **Network response to single cell manipulations under all rules learned with E→E turnover**

Figure 3.10: **Network response to single cell manipulations under all rules learned with E→E turnover (a)** Evolution of E cell postsynaptic firing patterns in the space of normalized peak amplitude and duration for six learned plasticity rules (each row represents one learned rule) when a targeted E cell's afferent synapses are scaled up by 50%. Black points represent pre-perturbation values, light purple represent immediately post-perturbation, and blue represent 25 activations after perturbation. **(b)** Analysis of the timing of targeted E cell firing under the same rules and perturbations as (a). Note delays are relative to a pre-perturbation activation not pictured. Scaling up the excitatory afferents of the E cell initially shifts the firing time forward, but plasticity restores the original firing time of the E cell. **(c)** Same as (a), but with the excitatory afferents to a targeted E cell scaled down by 50%. Plasticity rules restore the initial peak amplitude and duration of targeted E cell responses. **(d)** Same as (b) when E cell afferents are scaled down by 50%. Perturbations initially delay firing of targeted cells, but four of six learned rules partially restore original firing time of targeted cells.

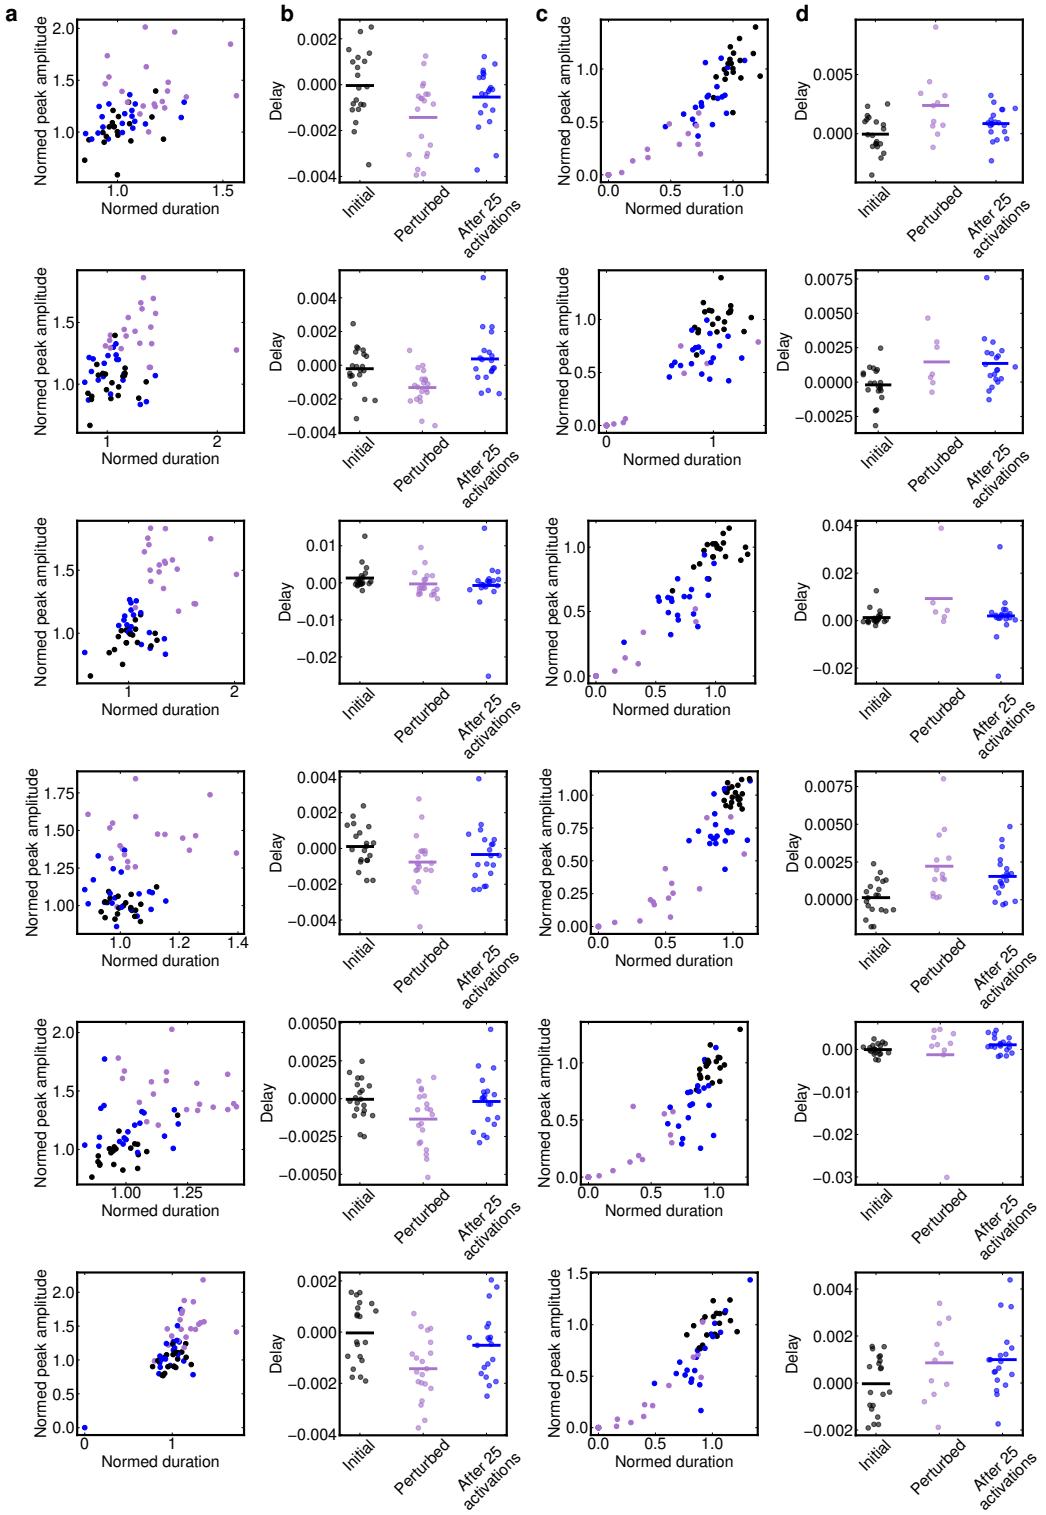

Figure 3.11: **Network response to single cell manipulations under all rules learned with I→E turnover** Identical to Supp. Fig. 3.10 (see caption) but for rules learned under I→E turnover.

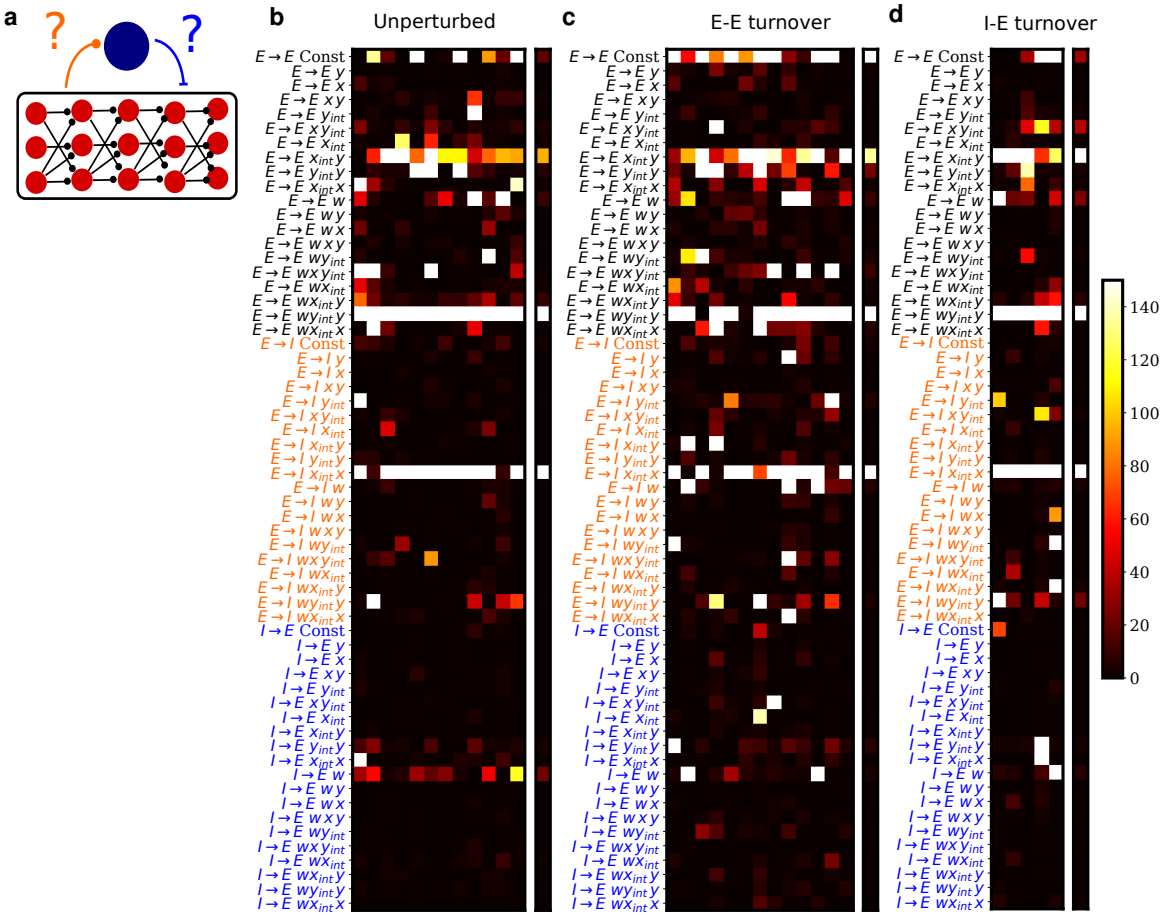

Figure 3.12: **Learning plasticity on all synapses**

(a) Schematic of synapses upon which additional plasticity will be learned. In previous simulations, these connections were held static. **(b)** Difference in medians between distribution of losses of full solutions run on 100 test networks and distribution of losses when one term is dropped on the same test set for solutions learned in the absence of synaptic turnover (left), **(c)** with turnover on E→E synapses, and **(d)** turnover on I→E synapses. Color of label indicates the synapses upon which each term acts. Isolated columns denote median loss for each term and type of perturbation. Note, temporally asymmetric Hebbian learning and heterosynaptic activity bound on E→E synapses remain amongst the most important terms by this measure across all types of perturbation. A presynaptic term on E→I synapses is also consistently important.

3. Scaling up the stochastic input to a single E cell by a factor 10 ($\lambda^{*}_{\text{stoch}} = 10\,\lambda_{\text{stoch}}$).

In each case, plasticity rules were applied to initially random networks (see Section 3.5) for 400 network activations, after which the target neuron was manipulated and the synaptic change effected by the plasticity rules was measured at six subsequent time points, each spaced five activations apart.

We characterized the activity and synaptic connectivity of targeted cells. In particular, we recorded the sum of all excitatory weights onto the target cell as well as the strength of the recurrent inhibition, computed as

$$|W_i^{\text{rec}}| = |\sum_k w_{i,k}^{\text{E}\to\text{I}} w_{k,i}^{\text{I}\to\text{E}}|. \tag{3.24}$$

In Fig. 6 and Supp. Fig. 3.9, we report these connectivity values normalized by pre-perturbation network averages across all neurons.

We additionally recorded changes in activity that stemmed from these manipulations and changes plasticity made to the network. We computed the peak amplitudes, durations, and timing of target cell firing. We defined the duration as $\int x_i(t)dt/x_i^{(\text{peak})}$ and the timing as $\int t x_i(t)dt/\int x_i(t)dt$. In Fig. 6d-f, Supp. Figs. 3.10a, c and 3.11a, c, we report peak amplitudes and durations normalized by pre-perturbation averages of the targeted cell. We report firing times relative to a pre-perturbation average (Supp. Figs. 3.10b, d and 3.11b, d).

We found that rules learned under different synpatic turnover settings and within the same setting largely converged in terms of function. Nearly all sets of rules increased (decreased) excitatory afferents onto the targeted cell and decreased (increased) recurrent inhibition in response to imposed scaling down (up) of excitatory afferents to the targeted cell (Fig 6a-c, Supp. Fig. 3.9). Plasticity rules acted to restore the peak amplitude, duration, and timing of an E cell's pre-perturbation activations (Supp. Figs. 3.10 and 3.11).

3.12 Reduced model of dual homeostatic control of E cell responses

We posit that the mechanism behind the improved robustness of the timing representation in networks with excitatory and inhibitory plasticity is compound homeostatic regulation of E cell postsynaptic response on different sets of synapses. To investigate this, we construct a

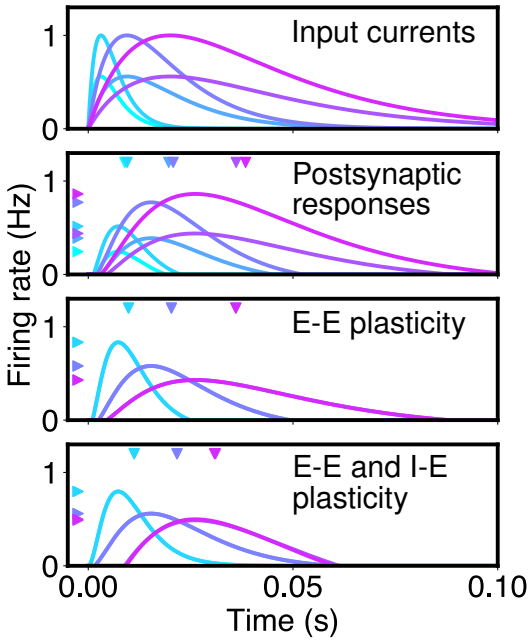

Figure 3.13: **Two plasticity mechanisms can correct timing**

Two plasticity mechanisms can better restore timing and postsynaptic firing pattern under a variety of inputs. Top panel: inputs to reduced neuron model described in Supp. Sec. 3.12. Second panel: initial firing rate responses of neuron to different inputs (colors matched to input). Third panel: firing rate responses of neuron evolving under E plasticity alone (see Supp. Sec. 3.12). Downward-pointing arrows indicate mean firing time of response and leftward-point arrows indicated peak firing rate. Bottom panel: firing rate responses of neuron evolving under E plasticity and I plasticity. Note mean firing times of responses and peak rates are more tightly clustered relative to those of responses of neuron with E plasticity alone.

model of a single neuron responding to a spectrum of excitatory inputs and investigate the stability of its postsynaptic response under different plasticity rules. Specifically, we compute the mean time relative to input onset and the magnitude of response under three different plasticity schemes:

1. Homeostatic conservation of squared magnitude of response on excitatory afferent

2. Homeostatic conservation of squared magnitude on excitatory afferent; conservation of magnitude on inhibitory afferent

We model the neuron as threshold linear unit, whose response is

$$r_i(t) = \left[w_i \int_0^t e^{-(t-t')/\tau_{\text{in}}} r_{\text{in}}(t')dt' - b_i \right]^+ \tag{3.25}$$

Here, w_i is the feedforward excitatory weight, b_i is the bias after i inputs, and τ_{in} is the membrane time constant of the neuron. We assume that inhibition to the neuron is relatively tonic and is incorporated into b_i. We allow w_i and b_i to adjust after each input according to

$$w_{i+1} = w_i + \gamma_w \left(A_w - \int_0^\infty r_i^2(t)dt \right) \tag{3.26}$$

and

$$b_{i+1} = b_i + \gamma_b \left(\int_0^\infty r_i(t)dt - A_b \right), \tag{3.27}$$

where A_w, γ_w, A_b and γ_b are non-negative homeostatic set points and rates for the feedforward weight and bias term, respectively. We drove the neuron with an alpha function,

$$r_{\text{in}}(t) = e a_{\text{in}}(t/\tau_{\text{in}}) e^{-t/\tau_{\text{in}}}, \tag{3.28}$$

starting from $t = 0$ and tracked the evolution of the postsynaptic response in the phase space of peak response amplitude and response duration. We found a suitable choice of model parameters created a stable fixed point in this space. We further computed the delay between the average firing time of the response, $\int t r_i(t)dt / \int r_i(t)dt$, and the onset of input after 200 activations and found that plasticity tended to constrain the range of possible delays (Fig. 6j, right). We compared these dynamics to those of a neuron with fixed bias, and found it did not conserve the duration nor delay of the postsynaptic response (Fig. 6j).

In Supp. Fig. 3.13, we demonstrate how plasticity on I→E synapses in tandem with E→E plasticity improves timing robustness. A key aspect of this model is that it assumes inhibition is relatively constant while excitation is transient. Thus, I→E plasticity permits an E cell to adjust its resting potential in a manner that shortens or lengthens its postsynaptic response. Inputs that are too short (long) relative to the homeostatic setpoint are lengthened by decreasing (increasing) inhibition. E→E plasticity can then rescale the amplitude of the response to match the desired setpoint.

Chapter 4

SEARCHING FOR RULES II: SELF-ORGANIZATION OF INTEGRATION CIRCUITRY

Integration—the accumulation of information over time—is a fundamental computation underlying navigation, decision-making, and memory. The dominant circuit model, the continuous attractor neural network (CANN), explains integration through finely tuned, symmetric connectivity that supports marginally stable states. However, data from systems such as the Drosophila central complex suggest that real connectivity is heterogeneous, implying that symmetry may only exist in a coarse-grained sense, if at all. While continuous attractor networks are known to be compatible with heterogeneity if trained in supervised fashion, we hypothesize that unsupervised local plasticity rules may shape integration circuitry, requiring less trial and error to achieve effective integration. To this end, we take a meta-learning approach to identify such rules in small networks of excitatory and inhibitory neurons. Briefly, the learned rule is parameterized as the Taylor expansion of an arbitrary function that depends on synapse size as well as eligibility traces of pre- and post-synaptic activity and input from other cell types. In networks of size N=5, we discover rules that form bump attractors given only weak spatial priors and inputs to be integrated. In a simpler two neuron system, we find rules that generate line attractor dynamics. Surprisingly, rules that shape N=5 networks do not consistently implement differential plasticity, a form of anti-Hebbian learning hypothesized to shape attractor networks. Instead, we find learned plasticity consistently uses input-based potentiation to shape network structure, and additionally exhibits sensitivity to the sequencing of input, pre-, and postsynaptic activity. We hypothesize this form of plasticity can be used to tune attracting states in a network and demonstrate this idea in a simplified model. Together, these results suggest that unsupervised local plasticity rules may provide a biologically plausible mechanism for the self-organization and robustness of integration circuits.

4.1 Introduction

A host of animal behaviors depend fundamentally on the brain's ability to integrate and retain incoming sensory cues. For example, *Drosophila* integrate self-motion and visual cues to maintain a persistent activity representation of heading [40–42]; zebrafish integrate optic flow to align themselves with current flow via the optomotor response [165, 166]; non-human primates maintain persistent activity representations of accumulated evidence in perceptual decision making tasks [167]; and grid cells in the entorhinal cortex of mammals perform path integration computations [168, 169]. A predominant circuit model proposed to underlie these capabilities is the continuous attractor neural network (CANN) [169–175]. These networks can be constructed as a combination of symmetric short-range excitation and long-range inhibition. This structure supports a continuum of marginally stable activity states, typically far exceeding the number of neurons in the network (Fig. 4.1b).

A challenge to the biological plausibility of this architecture is that it requires fine-tuning. Perturbation to symmetric connectivity causes the underlying dynamical system of a CANN to bifurcate into a series of fixed points, limiting the capacity of the network to represent many unique stimulus values in unique activity states of the network (Fig. 4.1b-c, e) [173, 176]. The requirement for fine-tuning is problematic because real connectomes are not symmetric, implying that effective symmetry would need to be "hidden" in otherwise heterogeneous connectivity [176–178]. In models, establishing integration dynamics in networks that are not initialized with tuned connectivity typically requires supervised learning [178–180] (Fig 4.1f, second from left). While effective at establishing integration dynamics, the view that integration circuit is established entirely through trial and error learning is challenged by recent evidence suggesting some circuits may be intrinsically capable of integration: zebrafish can accurately align themselves with noisy optic flow even when reared in darkness or strobe lit conditions, and rapidly gain this ability after brain-wide sodium blockade [166]. The sparsity of experience required to perform integration-dependent behaviors suggests extensive error-driven learning may not be the organizing principle in the underlying neural circuitry.

Here, we ask whether integration and retention of incoming inputs can occur through self-organization, i.e. intrinsic dynamics and plasticity that does not depend on error-based

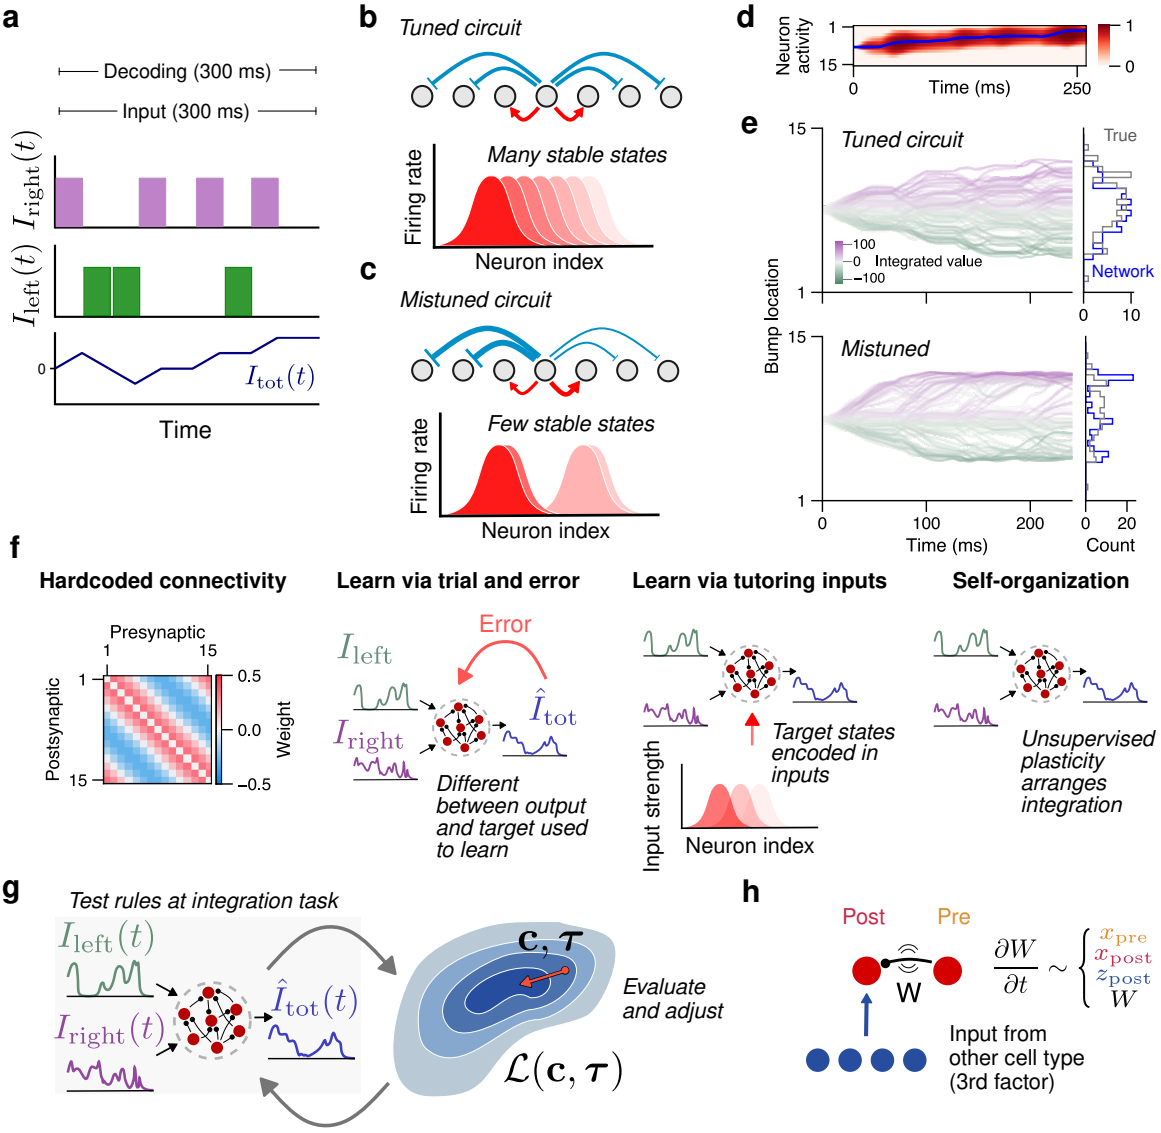

Figure 4.1: **Approach to learning rules that organize precisely tuned connectivity required for integration**

Figure 4.1: **Approach to learning rules that organize precisely tuned connected required for integration (a)** Illustration of integration task. Given a series of right and left cues, the goal is to output the integrated difference in the two signals (blue line). **(b)** Schematic of a tuned continuous attractor, containing short-range excitation and long-range inhibition, which supports many stable states (bottom). **(c)** Mistuned continuous attractor (top) supports comparatively fewer states (bottom). **(d)** A continuous attractor integrates incoming sensory cues and its dynamics tracks the summed difference in right and left cues (blue line). **(e)** (Top) Many trajectories through the same attractor, colored by the integrated value of the incoming cues. Histogram shows the distribution of targets and final states is approximately normal. (Bottom) Mistuning creates a set of fixed points that warp the integration dynamics of the network. **(f)** Schematics for hypotheses for how integrators are formed in the brain. Connectivity could be predetermined (far left), learned via error feedback to the network (second from left), tutored into the network by inputs that specify target network states (second from right), or fully self-organized (far right). **(g)** Schematic of meta-learning approach to discovering plasticity rules. **(h)** Schematic of factors that can influence the evolution of a synapse, W: activity of the pre- and postsynaptic neuron as well as net input from distinct cell types.

feedback or supervision. We take path integration as our canonical example, whereby a time-varying cue must be integrated to keep track of total displacement. To discover plasticity rules that can organize integration circuitry without feedback, we use evolutionary algorithms to learn local synaptic update rules within a basis of possible terms that depend on activity, synapse size, timing, and gating factors. We study a model based on the central complex of the fly, in which network cell types and structure that reflect the ellipsoid body and protocerebral bridge. To our knowledge, we find the first rules capable of organizing integration without supervision and that remain continually active. We then use analysis of a simplified network example to interpret the learned rule sets.

4.2 Related work

A large literature addresses the fine-tuning problem in continuous attractor networks, whereby deviations from symmetry in connectivity induce drift along the marginally stable manifold [173, 176, 178, 181]. Although theory shows that continuous attractors can arise even with non-symmetric connectivity [176, 177], and that slightly untuned networks can still support long memory timescales [178], most mechanistic explanations for the emergence of continuous attractors in neural circuits invoke feedback-based learning rules [178, 182].

An alternative architectural approach replaces a continuous manifold with a sequence of adjacent fixed points, often termed quasi-continuous attractor networks [29, 180, 183, 184]. While these models eliminate representational drift, they typically exhibit reduced sensitivity to weak inputs and require large network sizes, potentially limiting their applicability to integration circuits composed of relatively few neurons [181]. Synaptic facilitation has also been shown to suppress drift by reinforcing co-active neuronal populations [29, 185]. Finally, control-theoretic approaches based on negative derivative feedback—implemented via time-offset signal subtraction—can mitigate memory decay and confer robustness to common circuit perturbations [186, 187].

Many authors have investigated the role of synaptic plasticity in organizing continuous attractor networks that perform integration. Broadly, these approaches fall into two classes. The first relies on tutoring or supervisory inputs [29, 179, 180], in which desired network states or trajectories are externally imposed and used by synaptic plasticity to generate internal

error signals (Fig 4.1f, second from right). It remains unclear whether such approaches ultimately produce truly continuous attractors or instead yield quasi-continuous manifolds composed of discretized states.

The second class focuses on fully unsupervised plasticity mechanisms, including homeostatic processes [188] and so-called differential plasticity [189, 190]. However, these rules are typically studied in networks whose connectivity is already close to that of a continuous attractor, raising the question of whether plasticity alone can give rise to such structures from more generic initial conditions.

Given the large space of possible learning rules, an emerging approach is to use a supervised objective to tune a local plasticity rule, sometimes referred to as meta-learning. In these approaches, the tuning process often proceeds by simulating the local rule in a network, evaluating its task performance, and then modifying the parameters of the local rule via the supervised learning algorithm. Studies in meta-learning have been divided between searches for biologically plausible replacements for backpropagation [52, 56, 57, 148] and automated discovery of unsupervised plasticity capable of fulfilling particular functions [51, 53, 61, 62, 141, 191].

4.3 Results

To better understand how the continuous attractors could be self-organized within finite-sized networks, we sought to meta-learn plasticity rules that could shape networks into integration circuitry purely from the dynamics induced by input. Our approach was to simulate networks of firing rate neurons that evolved under a parametrized plasticity rule as they received two streams of input, and then iteratively tune the parameters of the rule to maximize a supervised objective, path integration of the input.

4.3.1 Parametrization of the learning rule

We parametrized a learning rule as a linear combination of basis terms which depended on factors like the activity of pre- and postsynaptic neurons, the synapse size, and gross input from distinct cell types. This last factor allows for heterosynaptic forms of plasticity,

which have increasingly gained interest in experimental and theoretical literature [44, 192–194]. Some terms in the basis also included eligibility traces of these factors, which were incorporated to capture timing information. For example, rules that depend on the relative ordering of activity of neurons are not implementable without these trace. In all cases, these traces were formed by convolving the factor with a decaying exponential. The weight update due to a particular learning rule is

$$\dot{w}_{ij} = \sum_k c_k F_k(x_i, x_j, w_{ij}, z_i, \{\tau_k\}), \tag{4.1}$$

where x_j and x_i are the pre- and postsynaptic firing rate, $\{\tau_k\}$ are the time constants of the eligibility traces, if applicable, and z_i is the total input current from the other cell type to the postsynaptic neuron (see Section 4.5.3). While a more expressive and general nonlinear combination of factors influencing synaptic strength can be generated by parameterizing the learning rule differently, e.g. by a feedforward network as in [191], we opted for this approach as the solution can be more interpretable [141].

4.3.2 Initial network architecture

We assume networks are not initially entirely unstructured, but instead begin with a structural prior. We modeled this prior on the connectivity of the head direction circuit found in *Drosophila*, which is remarkably stereotyped [195]. Our networks contained two cell types, which refer to as cell type 1 and 2. Cells of the first type initially randomly interconnect via excitatory synapses and disynaptically inhibit one another via a single inhibitory neuron. Two additional groups of excitatory cells (cell type 2) of equal size to cell type 1 population receive uniform "right" and "left" current inputs. We assumed connectivity between cell type 1 and 2 reflected the shifting structure found in projections between E-PG neurons and the protocerebral bridge (Fig. 4.2).

4.3.3 Integration task

We first studied a case in which networks received inputs for a 300 ms period. While the time scale of working memory is often an order of magnitude or two larger, we first sought to test

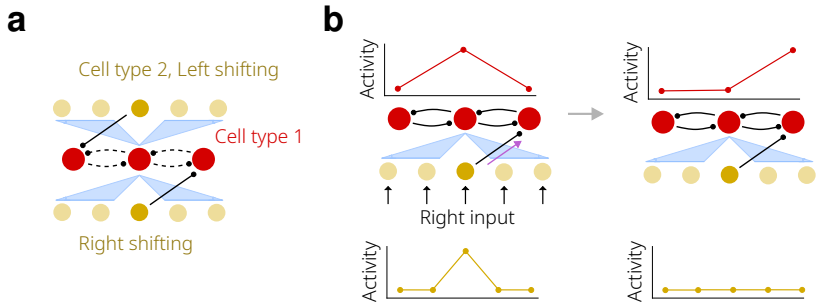

Figure 4.2: **Sketch of structural prior in networks (a)** Schematic of structural prior for cell type 1↔2 connectivity. Each cell type 1 neuron disinhibits two cell type 2 neurons, one which projects laterally to the right, the other to the left. Blue triangles denote blanket inhibition to all other cell type 2 neurons. Connections denoted with dashed lines are not initialized with a structural prior. **(b)** Function of structural prior. When reciprocal connections between cell type 1 neurons are well tuned, a bump of excitation will combine with a global right (left) input to cell type 2 neurons to produce a shifting signal that moves the bump right (left).

the hypothesis that integration could be self-organized from input cues alone. A challenge in this setting is that inputs the network receives may be asymmetric, which has prevented past network models from learning stable attractor dynamics [29].

Our procedure is as follows: after a 20 ms delay, networks received a string of input cues for the remaining time divided into 10 ms blocks. During each block, either a right or left input, implemented as a smoothed Poisson process, was presented with probability $p_{\text{in}}/2$ or no input was provided with probability $1 - p_{\text{in}}$.

4.3.4 Loss function and optimization of the learning rule

The full loss function depended only on the coefficients and time constants of the learning rule, $(\boldsymbol{c}, \boldsymbol{\tau})$. Explicitly,

$$\mathcal{L}(\mathbf{c}, \boldsymbol{\tau}) = \mathcal{L}_{\text{dec}}(\mathbf{c}, \boldsymbol{\tau}) + \lambda_{\text{a}} \sum_i \int_0^T x_i(t)dt + \lambda_{\text{s}} \mathcal{P}(\mathbf{c}, \boldsymbol{\tau}), \tag{4.2}$$

where the first term is the decoder loss, $\lambda_{\text{dec}}(1 - R^2)$, the second term is an activity penalty, and the third term is a synaptic penalty that discourages synaptic change,

$$\mathcal{P}(\mathbf{c}, \boldsymbol{\tau}) = \sum_{i,j,k} \int_0^T |c_k F_k(x_i(t), x_j(t), w_{ij}(t), \tau_k)| \Theta(|w_{ij}(t)|) dt. \tag{4.3}$$

We optimized rules by adjusting $(\boldsymbol{c}, \boldsymbol{\tau})$ via covariance matrix adaptation (CMA-ES) [152]. To evaluate the loss function, we batch evaluated a selection of $(\boldsymbol{c}, \boldsymbol{\tau})$ on 10 networks, and varied the value of p_{in} in even intervals from 0.25 to 1 when training on each network. Optimizations typically required tens of thousands of meta-iterations to converge.

4.3.5 Learned unsupervised plasticity organizes integration in finite networks

Meta-learning discovered sets of plasticity rules that shaped central complex-like networks into integrators. When run on new network initializations, these rules typically modified connectivity between cell type 1 neurons to be largely symmetric and strongest near the diagonal (Fig 4.3b). Resultant weight matrices were reminiscent of the cosine tuning used to implement ring attractors, but did not contain a full rotational symmetry due to the nature of the task. These connections support a "bump" of activity that passes between the

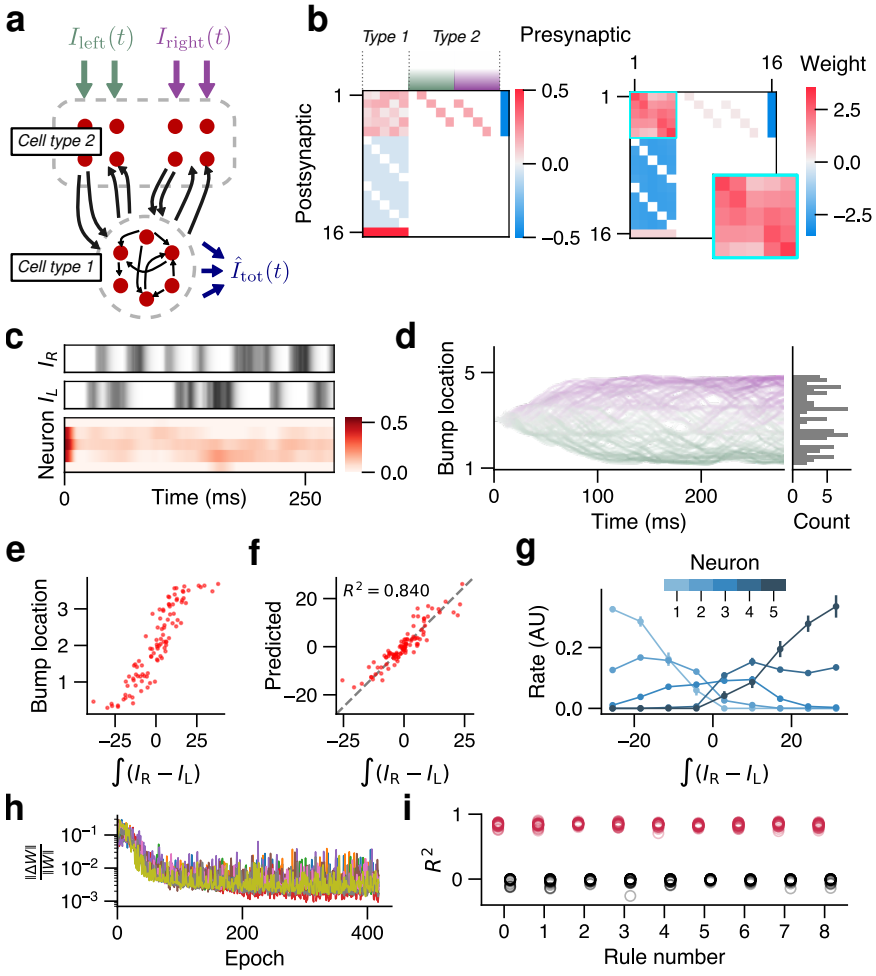

Figure 4.3: **Meta-learned rules organize integration dynamics in central complex-like networks**

Figure 4.3: **Meta-learned rules organize integration dynamics in central complex-like networks (a)** Schematic of central complex-like network on top of which plasticity rules are learned. Two distinct groups of neurons in cell type, neurons 6-10 and 11-15, receive global right and left inputs, respectively. Cell type 2 neurons send excitatory projections to cell type 1, neurons 1-5, **(b)**, and cell type 1 neurons reciprocally inhibit cell type 2. Connectivity between cell type 1 and 2 is assumed to have a scaffold that outlines the desired manifold. Connections between cell type 1 neurons are drawn from a uniform distribution (**(b)**, left), but converge to a diagonal structure (right) that permits integration. **(c)** Tracking the "bump location", i.e. the average location of activity in the network ($\frac{1}{\sum_k r_k}\sum_i i\, r_i$), reveals neural trajectories separate to encode integrated value. **(d)** Bump location reflects current integrated value. **(e)** Tuning curves for each of the 5 neurons in cell type 1. **(g)** Representation of distinct values of $\int (I_{\mathrm{R}} - I_{\mathrm{L}})$ in terms of the firing rate of different neurons. **(h)** Trial over trial relative synaptic change averaged across all synapses between units in cell type 1. **(i)** Comparison of R^2 between networks organized by plasticity rules (red) versus networks with no plasticity (black) across learned rules. Each circle represents a network (N=30).

neurons of cell type 1 in response to inputs to the network (Fig 4.3c). The bump location strongly correlates with the integrated difference in input cues (Fig 4.3d-e), and the true displacement of the path integration can be decoded from the activity of cell type 1 neurons to high accuracy (R^2=0.85±0.02 for the best rule discovered, R^2=-0.02±0.02 without plasticity, n=30) (Fig 4.3f). Individual cell type 1 neurons in the network display graded tuning to particular values of the displacement (Fig 4.3g). To determine how the discovered plasticity rules drive synaptic strengths, we measured the average change in synaptic strength across trials, $\langle|\Delta w_{ij}|/|w_{ij}|\rangle_{ij}$, and found that synapses in shaped networks only changed about 1% epoch over epoch by the time decoding occurred, indicating connectivity largely converged.

4.3.6 Memory in self-organized integrating circuits is based on long lasting dynamics

While the learned networks above integrate inputs, they do not necessarily retain the results over time. Thus, we next added a requirement that they retain the final value stably over time by appending a 700 ms "hold" period to the end of the integration period during which the integrated value could still be decoded from network activity (Fig. 4.4a) .

Meta-learning discovered plasticity rules that reliably solved this task . While a number of learning rules generated networks whose dynamics slowly decayed to a set of fixed points across the hold period, a subset shaped networks in which the dynamics of cell type 1 neurons appeared to freeze during the hold period. In Fig 4.4b, we show the dynamics of a network shaped by the most successful rule discovered. At the end of the hold period, the network dynamics displayed graded tuning as a function of the displacement of inputs (Fig. 4.4c-d). Across networks organized by this rule, we linearized the dynamics about activity states at the start of hold period. During this period cell type 2 neurons were generally quiescent, thus these linearizations yielded six pairs of eigenvectors and eigenvalues, corresponding to five cell type 1 neurons and the single global inhibitory neuron. Fig. 4.4e shows an example eigenspectrum for a linearization. For activity states within networks organized by the most successful rule, we found the real parts of all eigenvalues were generally negative, while typically one eigenvalue was close to zero, indicating a near marginal stability. Fig. 4.4f displays a histogram of the largest real eigenvalue across 100 activity states from 10

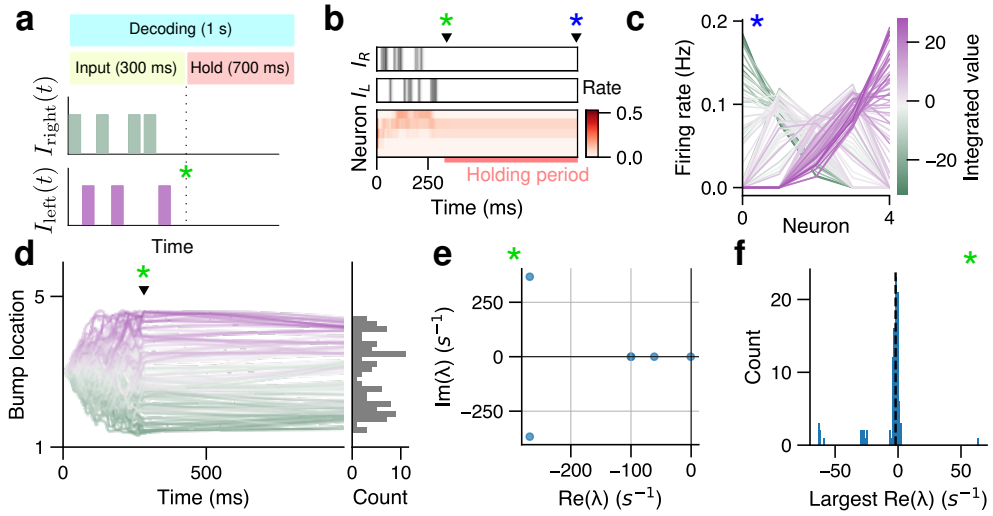

Figure 4.4: **Plasticity rules trained on integration task with a long holding period maintain long-lasting representations of integrated inputs (a)** In altered version of integration task, inputs are presented for an initial 300 ms period, followed by a long holding period, from which integrated inputs are still decoded. **(b)** Plasticity rules recurrent cell type 1 connectivity into diagonal, symmetric structure. **(c)** Network activity states at the end of the hold period reflect continuous encoding of integrated inputs. **(d)** Example of network activity over a single epoch of integration. Top two plots show right and left input currents to network. Bottom plot depicts cell type 1 activity. **(e)** Plot of bump peak across time for a single integration epoch.

networks. The median largest real eigenvalue for this rule was -1.95 s^{-1}, indicating activity states typically persisted for 500 ms or longer. A small fraction of these states had leading eigenvalues that indicated rapid collapse, and, in one case, rapid expansion.

4.3.7 Probing learned plasticity rules reveals mechanisms of self-organization

Having established that learned plasticity rules could organize integrators, we next attempted to probe the learned rules to reverse engineer their function. Direct interpretation of meta-learned plasticity rules is often difficult, particularly if the basis of terms is large [196]. This

difficulty arises from the fact that coefficients of terms in the basis do not necessarily have the same units and their contribution to dynamics greatly depends on the dynamics themselves.

Plasticity "knockouts" reveal rules are not sparse in the basis of possible terms

We first performed a set of plasticity "knockouts" in which the coefficients of individual terms within learned plasticity rules were set to zero in order to assess the dependence of the rule on the term [141]. We chose to focus on the learned plasticity rules that regulate cell type 1↔1 connections, as these weights displayed the most structural change over the course of self-organization. In Fig. 4.5, we show the results of these knockouts. The most consistent result was dependence on high order firing rate terms, like $w_{ij}x_i^4$ and $w_{ij}x_j^4$, which appear regularly with a negative coefficient and presumably act to bound the activity of individual neurons. Removal of these terms resulted in the complete collapse of the plasticity rule's performance. Reliance on other terms was inconsistent, but removal of many resulted in modest decreases in performance: independent removal of 11±3 and 15±3 terms resulted in a decrease of 0.1 in R^2 for the regular and long hold-trained plasticity rules, respectively, indicating plasticity rules were not sparse.

Integrator networks have two different dynamical regimes: one in which they receive input and one in which they do not. In a regime with no input, plasticity should work to support static dynamics. However, in a regime in which inputs arrive, the network dynamics should be allowed to change to match integration trajectories without impedance from local plasticity. We therefore reasoned that three-factor plasticity might be useful in shaping these network since it provides a way to change the effective plasticity rule during these two different regimes. We found that removal of three-factor terms (in which cell type 2 activity modulates cell type 1→1 synapses) only inconsistently affected rule performance (Fig. 4.5a-b, final eight rows). Despite this, we found that simultaneous ablation of all eight three-factor terms greatly reduced rule performance across all learned rules (Fig. 4.5c-d), suggesting combinations of three-factor plasticity play an important role in organizing integrator circuits and also that single term ablation may not fully capture the importance of individual terms to the full rule.

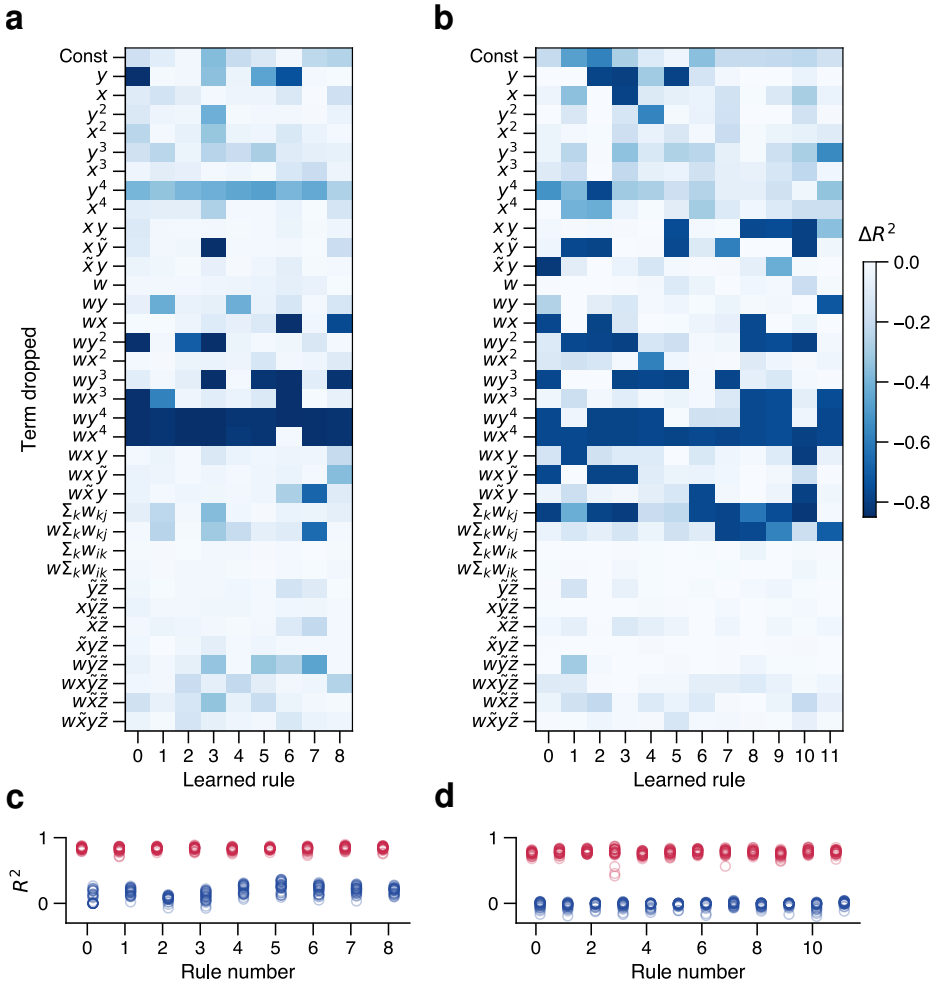

Figure 4.5: **Plasticity "knockouts"** Results of zeroing out individual coefficients of plasticity rules (each column represents one learned rule) in terms of change in R^2 relative to the unperturbed rule. Each row represents median change in R^2 after zeroing out of one term across the set of rules. **(a)** Perturbations for plasticity rules trained on the short (300 ms) integration task **(b)** Same as (a) for the long hold variation. **(c)** Comparison of R^2 for N=30 networks when all three factor terms are zeroed out (blue) versus left in (red) across 9 discovered rules for 300 ms integration task. **(d)** Same as (c) for long hold variation across 12 discovered rules.

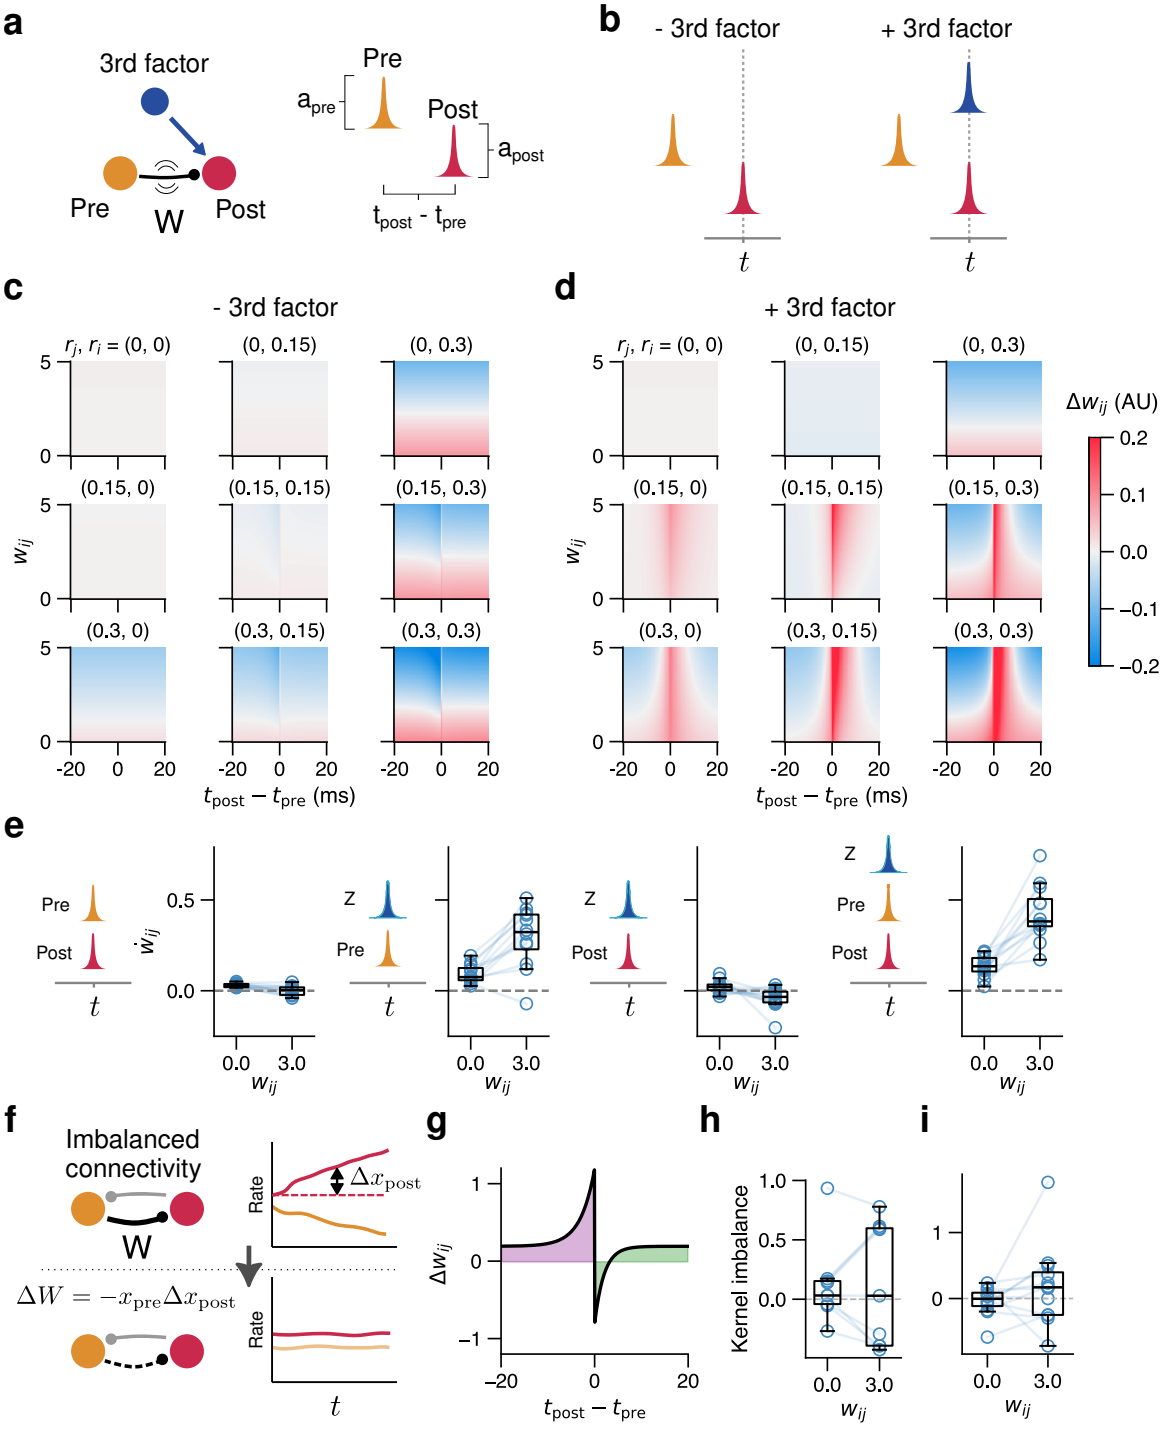

Figure 4.6: **Probing learned plasticity rules with delta function inputs**

Figure 4.6: **Probing learned plasticity rules with delta function inputs (a)** Evolution of a synapse (W) is computed as a function of pre- and postsynaptic activity and third factor input from cell type 2 (top). When the third factor is zero, weight change is dictated by amplitudes of pre and post synaptic timing and the time lag between them (bottom). **(b)** Schematic of relative activity timing for (d). Third factor is assumed to co-occur with postsynaptic activity. **(b)** Plots of weight change induced when the third factor is zero. Each plot shows the weight change as a function of synapse size (W) and relative pre and post lag for a given set of pre and postsynaptic activity amplitudes. **(d)** Induced change when third factor co-occurs with postsynaptic activity at a fixed amplitude. **(e)** Weight change induced by pairing and triplet activity in the pre- and postsynaptic neurons and third factor input across varying synapse size. For large synapse sizes, coincident presynaptic activity and third factor and coincident pre, post, third factor elicits potentiation while coincident pre and post and coincident post and third factor elicit depression. **(f)** Cartoon of calculation of kernel asymmetry. Area under the curve relative to zero from the negative post - pre time lag (purple) is subtracted from area for positive time lags (green). **(g)** Kernel imbalance across meta-learned plasticity rules (each line is one rule) across values of synapse size.

Computing weight change for special inputs

To gain a deeper intuition of how learning rules organized integrators, we next computed the change in a type 1→1 weight, w_{ij}, induced by its pre- and postsynaptic activity (x_j, x_i) and the current input of the cell type 2 population to the postsynaptic neuron, which we denote z_i (Fig 4.6a). One can compute the effective kernels for delta function inputs:

$$\begin{aligned} x_j(t) &= r_j \delta(t - t_{\text{pre}}), \\ x_i(t) &= r_i \delta(t - t_{\text{post}}), \\ z_i(t) &= z_i \delta(t - t_{\text{z}}). \end{aligned} \tag{4.4}$$

We then compute the resulting weight change as the integral over the period during which the neurons and inputs are nonzero. While learning rules are often expressed simply as a function of the relative timing of pre- and post-spiking, here the kernel is in general a function of the three input amplitudes, the three time lags derived from t_{pre}, t_{post}, and t_{z}, and the synapse size, w_{ij}; see Section 4.5.7 for details. This flexibility reflects literature that notes synaptic modification can depend on the synapse size [5, 6].

Presynaptic activity and third factor mediate potentiation; postsynaptic activity and third factor mediate depression

In Fig 4.6b and d, we show the weight change induced by these inputs. In b, we assume no third factor input ($z_i = 0$), and in d, we assume third factor input occurs concurrently with postsynaptic activity and with a fixed amplitude. We found that the magnitude of weight change typically scaled with r_j and r_i, and tended toward depression when w_{ij} grew large. When $z_i > 0$, we found that range of w_{ij} for which concurrent activity in the pre- and postsynaptic neurons resulted in potentiation was extended.

In Fig 4.6e, we show weight change resulting from simultaneous activity and third factor input. Across learning rules, we found that paired pre- and postsynaptic activity, pre and third factor, post and third factor, and all three factors simultaneously resulted in weak potentiation when the synapse was small ($w_{ij} = 0$). When the synapse was instead roughly equal to the largest size synapses reached during self-organization, paired post and third factor led to weak depression, and paired pre- and third factor and the simultaneous presence

of all three factors led to potentiation. In all cases, we tested whether the population mean was greater or less than zero (one-sample t-test; see Table 4.1). We note that this form of potentiation based on third factor and presynaptic activity is similar to a previously proposed rule, however the third factor in that case the third factor is a target signal and not an input to be integrated [179].

To confirm intuition generated from these numerical calculations matched the manner in which rules changed synapses in simulations, we computed the correlation of weight updates with the pre- and postsynaptic activity and third factor input. If the numerical calculations are correct, we might expect weight change in self-organized networks to strongly correlate with third-factor inputs that co-occur with presynaptic activity or simultaneous pre- and postsynaptic activity. We chose to study the rule shown in Fig. 4.6b, d. When we computed these correlations over the last 10 epochs of networks organize by this rule, we found weight change was most correlated with third-factor input and presynaptic activity that occurred shortly before the weight change (see Table 4.3 and Fig. 4.11). Weight change was also strongly correlated with the presence of all three factors occurring shortly before the weight change. Weight changes were significantly less correlated with individual factors as well as other possible products, like simultaneous pre- and postsynaptic activity.

Differential plasticity is not consistently a feature of learned plasticity rules

Previous work has explored differential plasticity as a mechanism for tuning neural integrator. This plasticity typically has the form

$$\dot{w}_{ij} = -\dot{x}_i x_j. \tag{4.5}$$

In the absence of inputs, differential plasticity can act to suppress spontaneous drift of activity in a neural network by weakening connections that lead to the propagation of activity between neurons, and has been shown to tune attractor networks in prior work [189, 190]. Such a rule can arise from an anti-Hebbian learning if positive time lags between i and j lead to depression and negative time lags lead to potentiation (see Xie and Seung). A key prediction about differential plasticity is that it should be inactive when inputs to be

Pairing	w	p (mean > 0)	p (mean < 0)
Pre, post	0	$\mathbf{1.1 \times 10^{-6}}$	1.0
Pre, post	3	0.46	0.54
Pre, third	0	$\mathbf{2.7 \times 10^{-5}}$	1.0
Pre, third	3	$\mathbf{3.4 \times 10^{-5}}$	1.0
Post, third	0	0.027	0.97
Post, third	3	0.98	**0.017**
Pre, post, third	0	$\mathbf{3.4 \times 10^{-6}}$	1.0
Pre, post, third	3	$\mathbf{6.0 \times 10^{-7}}$	1.0

Table 4.1: Summary of one-sample t-tests. P values below 0.025 are shown in bold. Significance threshold chosen based on Bonferroni correction.

integrated are present: because differential plasticity opposes changes in firing rate, it must be disabled when a network receives inputs that are intended to change its state.

To assess whether our meta-learned rules implement differential plasticity, we computed the synaptic change induced by varying time lags between pre- and postsynaptic activity of cell type 1 neurons. We then computed the total asymmetry of the kernel by calculating the area under the curve (relative to zero) when the time lag was positive and subtracting the area for which the time lag was negative within a given window (Fig. 4.6e). We tested whether the population mean of kernel asymmetry was nonzero for different values of w_{ij} and across different window sizes ($\tau_w = \{10, 20, 30\}$ ms) and found in all cases it was not (one sample t-test against zero; τ_w=10 ms, w=0, p=0.60, w=3, p=0.25; τ_w=20 ms, w=0, p=0.59, w=3, p=0.31; τ_w=30 ms, w=0, p=0.55, w=3, p=0.32). These results suggest differential plasticity of the kind proposed in Xie and Seung was not consistently the mechanism that organized neural integrators in networks.

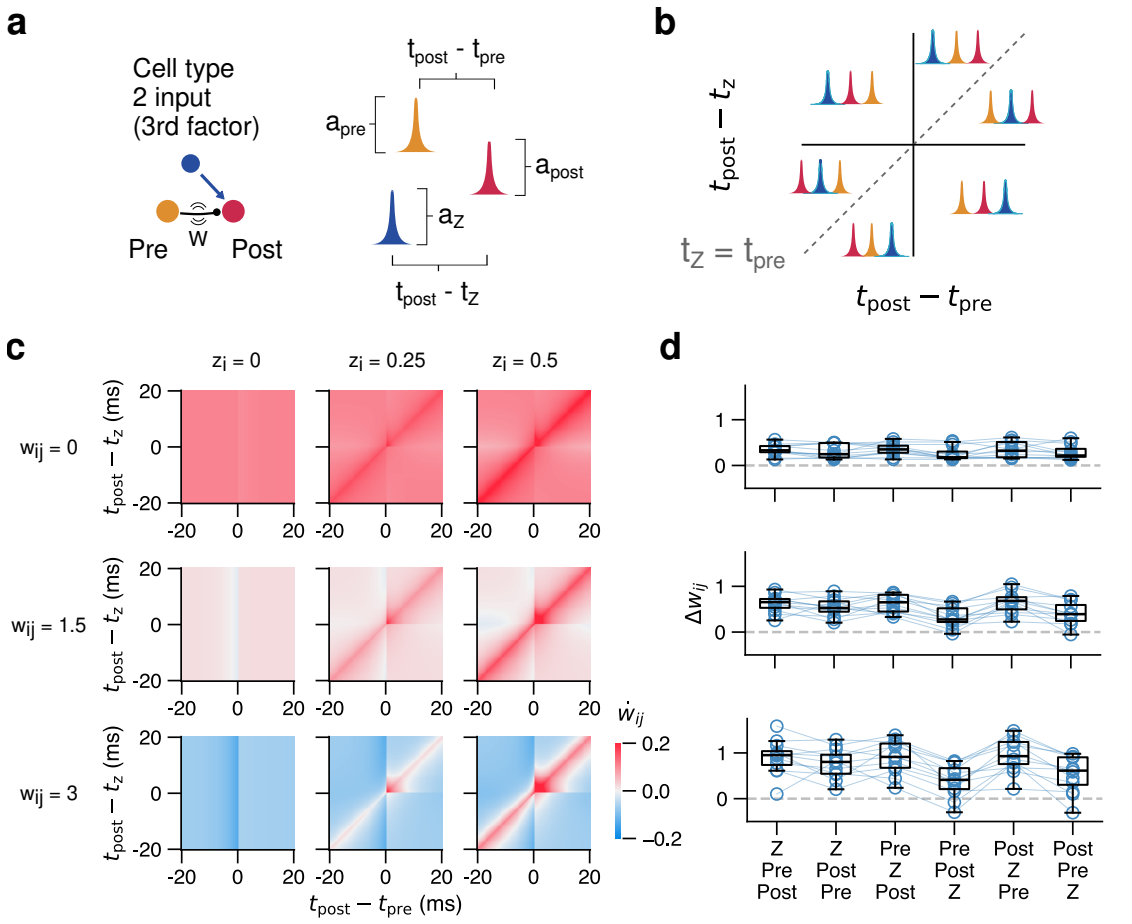

Figure 4.7: **Probing 1→1 plasticity rules based on timing of pre- and postsynaptic activity and input from cell type 2 points toward timing-dependent 3 factor potentiation (a)** Schematic of network activity in terms of amplitudes and relevant time lags. **(b)** Visualizations of the orderings in which activities occur on plots in (c). **(c)** Rate of synaptic change based on the timing of pre and postsynaptic activity of cell type 1 and input from cell type 2 across different values of the synapse size and the magnitude of input from cell type 2. **(d)** Rate of synaptic change based on the relative timings of pre and postsynaptic activity of cell type 1 and input from cell type 2 across all rules. From top to bottom, $w_{ij} = 0, 1.5, 3$.

A timing-based three factor rule may help tune attracting states

If differential plasticity is not consistently the mechanism by which attracting states are organized, what other mechanism might learned plasticity use? To answer this, we examined the weight change enacted by learned rules as we varied the timing of pre-, postsynaptic activity and third factor input. We fixed the amplitudes of these signals and varied the time lag between them (Fig. 4.7a). These three signals can be arranged into six possible orderings (Fig. 4.7b). In Fig. 4.7c, we plot the potentiation induced as a function of these time lags, oriented relative to the timing of the postsynaptic activity (single plot), and additionally varied the strength of the synapse and the amplitude of the third factor input (across plots). Consistent patterns across learning rules included a transition from potentiation to depression as synapse size increased and potentiation corresponding to nearly simultaneous presynaptic activity and third factor input (a red diagonal band in plots with $z_i \neq 0$). Rules additionally displayed dependence three factor terms (pre, post, third factor). We tested whether plasticity rules preferentially potentiated or depressed particular orderings of these factors and found, across rules, that sequences terminating in the third factor were either potentiated less than other sequences or depressed (4.7d). An interpretation of this rule is that plasticity rewards transition of activity between pre- and postsynaptic neurons that follow or co-occur with third factor input and penalizes transitions that precede third factor input.

4.3.8 Input-dependent rules can self-organize a two neuron network into a line attractor

We tested the hypothesis that components of rules discovered via meta-learning were sufficient to organize integration in a simple circuit. (Fig 4.8a). Specifically, we wondered if a plasticity rule based on three discovered terms (1) potentiation based on coincident third factor and presynaptic activity, (2) depression based on coincident third factor and postsynaptic activity, and (3) a term sensitive to the timing of all three factors might be sufficient to tune attractor dynamics (Fig 4.8b). We wrote this as

$$\dot{w}_{ij} = \alpha_1 z_i x_j - \alpha_2 z_i x_i + \alpha_3 z_i x_j \dot{x}_i, \tag{4.6}$$

where all terms in α are positive coefficients. The first two terms represent two-factor interactions between the third factor and pre- or postsynaptic activity, and the third term depends on the presynaptic activity, third factor, and derivative of the postsynaptic activity. Asymmetric timing-sensitive rules can be approximated by rules that include a derivative if the timescale of the rules is generally shorter than the timescale of fluctuations in firing rates (see Section 4.7.1).

We chose to study a reduced model in which the cell type 1 population consists of two threshold linear neurons that were recurrently connected to themselves and each other. Accordingly, the cell type 2 population consisted of four neurons, two for each group of shifting neurons. All other details were the same as in the meta-learned case. We initialized recurrent cell type 1 connectivity randomly, and allowed the cross synapses (w_{ij}, $i \neq j$) to evolve according to the plasticity rule over the course of 40,000 epochs.

Under the assumption that tonic input and recurrent inhibition to each neuron is equal, the condition for a line attractor in which neurons trade off activity are well known (see Section 4.6). This condition is

$$w_{ij} = w_{jj} - 1. \tag{4.7}$$

That is, cross synapse must be exactly one less than the recurrent synapse of the same presynaptic cell. We found randomly initialized networks converged to this condition under the plasticity rule (Fig 4.8c, right). By the last epoch, the activity of the network reflected the displacement of the input currents, whereas untrained networks did not reflect this coding 4.8d). We computed the correlation between the bump location within the network $x_2/(x_1 + x_2)$, and the displacement, and found they correlated extremely well. In 4.8e, we plot R^2, which is the square of this correlation, for plastic and fixed networks. Plastic networks were significantly better at encoding the displacement of inputs in their dynamics than untrained networks (plastic: $R^2 = 0.93 \pm 0.03$, fixed: $R^2 = 0.06 \pm 0.06$, N=5).

4.4 Discussion

We have demonstrated in a minimal network that unsupervised plasticity and a structural prior is sufficient to organize integration dynamics. Meta-learned plasticity rules shaped

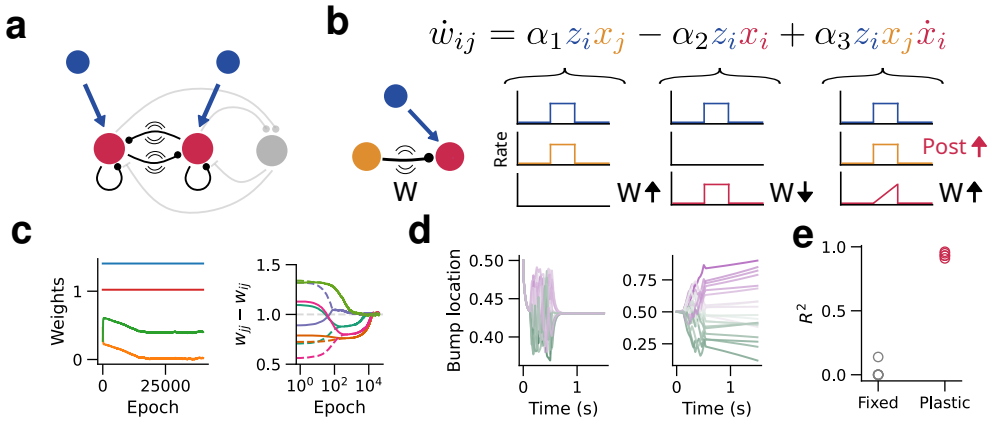

Figure 4.8: **Input-dependent plasticity rules organize a two-dimensional network into a line attractor (a)** Schematic of simplified threshold linear network. Two cell type 1 neurons (red) receive projections from cell type 2 neurons (blue) and mutually inhibit from via an inhibitory cell. **(b)** Effective plasticity rule. **(c)** Weight evolution under plasticity rule. Reciprocal synapses (orange and green) evolve over time (left). Plotting $w_{jj} - w_{ij}$ denotes proximity to the tuned condition ($w_{jj} - w_{ij} = 1$). Solid lines denote $w_{11} - w_{21}$ and dashed lines denote $w_{22} - w_{12}$. Color indicates network (N=5). **(d)** Bump location in network, $x_2/(x_1 + x_2)$, as a function of time for the final epoch. An untrained network rapidly converges to a fixed point (left). A plastic network shows integration dynamics (right). **(e)** Comparison of R^2 between fixed (randomly initialized) and plastic networks after 40,000 integration epochs.

networks with prior connectivity reflecting the connectivity in *Drosophila* central complex into bump attractor networks that encode inputs in persistent activity. The meta-plasticity procedure further succeeds in finding rules capable of building networks with long timescale retention of the integrated inputs.

By probing meta-learning plasticity, we discovered patterns across rules that make predictions about plasticity in integration circuits. We find that three-factor plasticity is an important element of self-organizing integration circuitry. Our rules also suggest that activity in the input cell type (type 2) and presynaptic activity in the representation cell type should lead to potentiation, while input cell activity paired with postsynaptic activity should lead to depression. We find that classical differential plasticity, in which anti-Hebbian learning can stabilize attractor networks, is not a necessary component of learned plasticity. Three-factor plasticity within learned rules also demonstrates sensitivity to the ordering of inputs, suggesting a potential mechanism by which integration dynamics could be organized within networks. While we do not prove this aspect of plasticity is essential to its ability to organize integrators, we demonstrate in a reduced, two-neuron network that such plasticity can organize integrators. This form of plasticity suggests, in opposition to differential plasticity, that updates to synapses occur while networks receive inputs rather than when they do not.

4.5 Methods

4.5.1 Rate neuron model

All simulations feature threshold linear firing rate neurons, whose dynamics evolved according to

$$\tau_m \frac{d}{dt} s_i(t) = -s_i(t) + \sum_j w_{ij} x_i(t) + u_i(t), \tag{4.8}$$

$$x_i(t) = \sigma(s_i), \tag{4.9}$$

where $x_i(t)$ is the firing rate for neuron i, w_{ij} is the synapse from $j \to i$, $u_i(t)$ is the input to neuron i, s_i is the activation of neuron i, τ_m is the membrane time constant, and σ is a firing rate nonlinearity, which depends on the cell type, γ. For all excitatory neurons,

the nonlinearity is a rectified tanh function, $\sigma(s) = \tanh([s - \beta^{(\gamma)}]^+)$, where $\beta^{(\gamma)}$ is a cell type-specific threshold. For the single inhibitory neuron, $\sigma(s) = [s]^+$.

4.5.2 Drosophila-inspired structural prior

To mimic the *Drosophila* central complex, N=16 neuron networks were divided into 4 groups. Cell type 1 contained 5 neurons. Cell type 2 contained two groups, "right" and "left" shift neurons, which differ in their projections to cell type 1 and the inputs they receive. Finally, one inhibitory neuron served to impose blanket inhibition on cell type 1 neurons. Connections between cell type 1 neurons were chosen according to a uniform distribution

$$w_{ij}^{(1)\to(1)} = U[0, w_0^{(1)\to(1)}] \tag{4.10}$$

For the set of right-shifting and left-shifting neurons, respectively,

$$\begin{aligned} w_{ij}^{(2)\to(1)} &= w_0^{(2)\to(1)}\delta(i-j+1) \\ w_{ij}^{(2)\to(1)} &= w_0^{(2)\to(1)}\delta(i-j-1) \end{aligned} \tag{4.11}$$

And projections from cell type 1 to 2 were

$$w_{ij}^{(1)\to(2)} = w_0^{(1)\to(2)}(1-\delta(i-j)) \tag{4.12}$$

All cell type 1 neurons were coupled to the inhibitory neuron with weight $w^{E\to I}$ and received reciprocal connections from the inhibitory neurons with strength $w^{I\to E}$.

4.5.3 Meta-learning rules and basis

We meta-learn a unique plasticity for each set of projections: 1→1, 1→2, 2→1. For each projection type, the rule takes the form

$$\dot{w}_{ij} = \sum_k c_k F_k(x_i, x_j, w_{ij}, z_i, \tau_k), \tag{4.13}$$

where F_k is the set of basis functions composed of two factor, three factor, and weight bound-like terms. The combined two factor and weight-bound terms and three factor terms

Parameter	Description	Value
τ_m	Membrane time constant for cell type 1 and 2	10 ms
$\tau_m^{(I)}$	Membrane time constant inhibitory cell	1 ms
$\beta^{(1)}$	Cell type 1 firing rate nonlinearity threshold	0.1
$\beta^{(2)}$	Cell type 2 firing rate nonlinearity threshold	0
$w_0^{(1)\to(1)}$	Max recurrent weight (type 1 $\to$ type 1)	$0.9/N^{(1)}$
$w_0^{(2)\to(1)}$	Shift weight (type 2 $\to$ type 1)	$0.9/N^{(1)}$
$w_0^{(1)\to(2)}$	Feedback weight (type 1 $\to$ type 2)	$-0.3/N^{(1)}$
$w^{E\to I}$	Excitatory $\to$ inhibitory weight	$2.5/N^{(1)}$
$w^{I\to E}$	Inhibitory $\to$ excitatory weight	$-4.5/N^{(1)}$

Table 4.2: Parameters of the central complex network model.

are

$$F^{(\text{pair})} = \begin{bmatrix} x_i & w_{ij}x_i \\ x_j & w_{ij}x_j \\ x_i^2 & w_{ij}x_i^2 \\ x_j^2 & w_{ij}x_j^2 \\ x_i^3 & w_{ij}x_i^3 \\ x_j^3 & w_{ij}x_j^3 \\ x_i^4 & w_{ij}x_i^4 \\ x_j^4 & w_{ij}x_j^4 \\ x_i x_j & w_{ij}x_i x_j \\ x_j\tilde{x}_i & w_{ij}x_j\tilde{x}_i \\ \tilde{x}_j x_i & w_{ij}\tilde{x}_j x_i \\ \sum_k w_{kj} & w_{ij}\sum_k w_{kj} \\ \sum_k w_{ik} & w_{ij}\sum_k w_{ik} \end{bmatrix} \qquad F^{(\text{3-factor})} = \begin{bmatrix} \tilde{x}_i\tilde{z}_i \\ \tilde{x}_i\tilde{z}_i x_j \\ \tilde{x}_j\tilde{z}_i \\ \tilde{x}_j\tilde{z}_i x_i \\ w_{ij}\tilde{x}_i\tilde{z}_i \\ w_{ij}\tilde{x}_i\tilde{z}_i x_j \\ w_{ij}\tilde{x}_j\tilde{z}_i \\ w_{ij}\tilde{x}_j\tilde{z}_i x_i \end{bmatrix} \tag{4.14}$$

All tildes denote convolutions with decaying exponentials, i.e.

$$\tilde{x}(t) = \int_0^t e^{-(t-t')/\tau_k}\, x(t')\, dt', \tag{4.15}$$

where the time constant, τ_k, is specific to the basis term k. Note that all three factor terms have two time constants. z_i is the bulk current input from the other cell type. For example, for the 1→1 rule,

$$z_i = \sum_l w_{il}^{(2)\to(1)} x_l^{(2)}, \tag{4.16}$$

where l indexes all cell type 2 neurons. As cell type 2 neurons do not interconnect, we assume there are no three factor terms that operate on the 1→2 projections.

4.5.4 Integration task

The integration task was divided into 420 equivalent epochs during which networks received a stream of right and left current inputs to integrate. At the end of each epoch, networks dynamics were reset but the weight matrix evolved by plasticity rules was preserved between epochs.

At the start of each epoch ($t = 10$ ms) a subset of cell type 1 neurons (cells 2-4) received a constant pulse of activity (duration = 10 ms, amplitude = 0.05) to break the symmetry of the quiescent state. At $t = 20$ ms, networks received inputs for a 280 ms period. Inputs were divided into 10 ms blocks. During each block, a right or left input, implemented as a smoothed Poisson process, was presented with probability $p_{\text{in}}/2$ or no input was provided with probability $1 - p_{\text{in}}$. To generate stochastic input during the block, we filtered a Poisson process ($\lambda = 2000$ Hz) with an alpha kernel,

$$\alpha(t - t') = \frac{e(t - t')}{\tau_\alpha} e^{-(t-t')/\tau_\alpha} \Theta(t - t'), \tag{4.17}$$

where $\tau_\alpha = 3$ ms. If the input was a right input, it was presented globally to cell type 2 neurons 6-10, otherwise it was presented to neurons 11-15. For the long hold trials, we extended the simulation past the initial 300 ms of input with 700 ms of quiescence.

Performance on the integration task was determined by the accuracy with which the true displacement could be decoded from the activity of cell type 1 neurons using a Lasso

regression ($\alpha = 0.05$). To construct the decoder, we randomly sampled the activity of the network at $400T$ time points (where T is the length of the epoch in seconds) between the onset of input and the end of the epoch. To evaluate performance, we sampled the same number of time points for epochs 320-420. Performance is captured by the R^2 of the decoder.

4.5.5 Meta-learning procedure

We optimized the loss function,

$$\mathcal{L}(\mathbf{c}, \boldsymbol{\tau}) = \mathcal{L}_{\text{dec}}(\mathbf{c}, \boldsymbol{\tau}) + \lambda_{\text{a}} \sum_i \int_0^T x_i(t) dt + \lambda_{\text{s}} \mathcal{P}(\mathbf{c}, \boldsymbol{\tau}). \tag{4.18}$$

The three terms in the loss correspond to a decoder loss, $\lambda_{\text{dec}}(1 - R^2)$, an activity penalty, which discourages solutions in which neurons are highly active, and a synaptic change penalty,

$$\mathcal{P}(\mathbf{c}, \boldsymbol{\tau}) = \sum_{i,j,k} \int_0^T |c_k F_k(x_i(t), x_j(t), w_{ij}(t), \tau_k)| \Theta(|w_{ij}(t)|) dt. \tag{4.19}$$

which computes the path length of the synaptic change per term in the basis.

To compute the loss of a plasticity rule, we initialized $N_{\text{batch}} = 10$ networks from a fixed seed. The n^{th} network was driven with input corresponding to $p_{\text{in}} = 0.25 + 0.75(n - 1)/(N_{\text{batch}} - 1)$. The loss of the rule was the average of the loss computed according to Eq. 4.18.

We optimize this loss via Covariance Matrix Adaptation (CMA) [152]. We choose 0 as the initial value for all coefficients and 0.005 (5 ms) as the starting value for all time constants, $\sigma_0 = 10^{-3}$. The population size per generation was typically fixed to 30. We additionally rescaled a number of terms to encourage CMA to explore reasonable parts of parameter space.

4.5.6 Linearizations of dynamics

To determine the stability of fixed points in organized networks, we compute the eigenvalues of the Jacobian of activations, which are identical to the eigenvalues of the Jacobian of the firing rates. The Jacobian is

$$J_s = T^{-1}(WD - I), \tag{4.20}$$

where $T = \mathrm{diag}(\tau_\mathrm{m})$, W is the weight matrix, and

$$D_{ij} = \left.\frac{\delta\sigma(s_i)}{\delta s_j}\right|_{s^*} \tag{4.21}$$

We performed this linearization only on the active subset of the network when there was no input. Thus, cell type 2 neurons were always excluded. The resulting linearizations describe a six-dimensional space, corresponding to five cell type 1 neurons and one inhibitory cell).

4.5.7 Constructing empirical synaptic update rules

To probe learned synaptic update rules, we compute the synaptic update resulting from pairs (or in some cases, triplets) of delta function-like activations. These updates will depend on the amplitudes of the activations and their relative time lags. We first consider the pairwise rules active between cell type 2. We take

$$\begin{aligned} x_i &= a_i\delta(t) \\ x_j &= a_j\delta(t-\Delta t). \end{aligned} \tag{4.22}$$

Terms in only pre or postsynaptic activity:

$$\begin{aligned} \Delta w_{ij} &= \int_0^\infty \dot{w}_{ij}\,dt \\ &= \int_0^\infty w_{ij}^\alpha x_i\,dt \\ &= \int_0^\infty w_{ij}^\alpha a_i\delta(t)\,dt \\ &= w_{ij}^\alpha a_i \end{aligned} \tag{4.23}$$

where $\alpha \in \{0,\ 1\}$.

For rules with pre and postsynaptic activity and a filter on on either the pre or post, we may also consider the time lag between the delta functions, Δt_{ij}. For instance, for

$\dot{w}_{ij}(t) = w_{ij}(t)\tilde{x}_i(t)x_j(t)$:

$$\begin{aligned}
\Delta w_{ij} &= \int_0^\infty \dot{w}_{ij}dt \\
&= \int_0^\infty w_{ij}\tilde{x}_i x_j dt \\
&= \int_0^\infty w_{ij}\left(\int_0^t e^{-(t-t')/\tau}x_i(t')dt'\right)x_j(t)dt \\
&= \int_0^\infty w_{ij}\left(\int_0^t e^{-(t-t')/\tau}a_i\delta(t')dt'\right)a_j\delta(t-\Delta t_{ij})dt \\
&= \begin{cases} a_i a_j w_{ij} e^{-\Delta t_{ij}/\tau} & \Delta t_{ij} \geq 0 \\ 0 & \Delta t < 0 \end{cases}
\end{aligned} \tag{4.24}$$

Finally, some terms include filters on two distinct favors. For example, $\dot{w}_{ij}(t) = \tilde{x}_i z_j$, where $z_j = \sum_k w_{kj}x_k$, where k indexes the neurons of a particular cell type. We assume

$$z_k = a_z\delta(t - \Delta t_{iz}). \tag{4.25}$$

Then, if $\Delta t_{iz} \geq 0$,

$$\begin{aligned}
\Delta w_{ij} &= \int_0^\infty \tilde{x}_i(t)\tilde{z_j}(t)dt \\
&= \int_0^\infty \left(\int_0^t e^{-(t-t')/\tau_i}a_i\delta(t'-\Delta t_{iz})dt'\right)\left(\int_0^t e^{-(t-s)/\tau_z}a_z\delta(s)ds\right)dt \\
&= a_i a_z \int_{\Delta t_{iz}}^\infty e^{-(t+\Delta t_{iz})/\tau_i}e^{-t/\tau_z}dt \\
&= a_i a_z \left(\frac{1}{\tau_i}+\frac{1}{\tau_z}\right)^{-1} e^{-\Delta t_{iz}/\tau_z}.
\end{aligned} \tag{4.26}$$

When $\Delta t_{iz} < 0$,

$$\Delta w_{ij} = a_i a_z \left(\frac{1}{\tau_i}+\frac{1}{\tau_z}\right)^{-1} e^{\Delta t_{iz}/\tau_i}. \tag{4.27}$$

4.6 Conditions for line-attracting dynamics in N=2 networks of threshold linear neurons

We consider a 2 neuron network in hopes of building intuition for self-organizing integration applicable to larger networks. The rationale for studying integration in larger networks is threefold. First, circuits thought to be responsible for integration generally consist of many,

interconnected neurons. Second, integration circuits that have many units are more robust to noise: unit noise orthogonal to the integration manifold is fully suppressed, while noise along the manifold is dampened by the central limit theorem. Consider an integration circuit realized as a single, self-connected neuron. All unit noise is aligned with the integration manifold, which is generally undesirable. Finally, integration circuits build of many neurons may have the opportunity to bifurcate to a series of fixed points when mistuned as opposed to a single fixed point, allow the circuit to maintain some integration capacity in the event of damage to the network.

Here, we consider of reduced system of two threshold linear units. The evolution of firing rates is given by

$$\tau \frac{d\mathbf{x}}{dt} = W\,[\mathbf{x}]^+ - \mathbf{x} + \mathbf{u}, \tag{4.28}$$

where

$$\mathbf{x} = \begin{bmatrix} x_1 \\ \\ x_2 \end{bmatrix}, \quad W = \begin{bmatrix} w_{11} - w_1^I & w_{12} - w_1^I \\ \\ w_{21} - w_2^I & w_{22} - w_2^I \end{bmatrix}, \quad \mathbf{u} = \begin{bmatrix} u_1 \\ \\ u_2 \end{bmatrix}. \tag{4.29}$$

We can search for generalized conditions for line attracting dynamics by setting $\frac{d\mathbf{x}}{dt} = 0$, yielding

$$\begin{aligned} (w_{11} - w_1^I - 1)x_1 + (w_{12} - w_1^I)x_2 + u_1 = 0 \\ (w_{22} - w_2^I - 1)x_2 + (w_{21} - w_2^I)x_1 + u_2 = 0 \end{aligned} \tag{4.30}$$

From these equations, we first determine

$$x_2 = -\frac{u_2 + (w_{21} - w_2^I)x_1}{w_{22} - w_2^I - 1} \tag{4.31}$$

and then that

$$\left((w_{11} - w_1^I - 1) - \frac{(w_{12} - w_1^I)(w_{21} - w_2^I)}{w_{22} - w_2^I - 1}\right) x_1 - \frac{(w_{12} - w_1^I)u_2}{w_{22} - w_2^I - 1} + u_1 = 0. \tag{4.32}$$

If this is to hold for any $x_i \geq 0$, then

$$(w_{11} - w_1^I - 1)(w_{22} - w_2^I - 1) = (w_{12} - w_1^I)(w_{21} - w_2^I) \tag{4.33}$$

$$(w_{12} - w_1^I)u_2 = (w_{22} - w_2^I - 1)u_1 \tag{4.34}$$

And by symmetry, we also get

$$(w_{21} - w_2^I)u_1 = (w_{11} - w_1^I - 1)u_2. \tag{4.35}$$

Note that if Eqs. 4.34 and 4.37 hold, then Eq. 4.33 holds automatically. We then find a stable line of activity states corresponding to

$$x_2 = -\frac{u_1 + \frac{u_1}{u_2}(w_{21} - w_2^I)x_1}{w_{12} - w_1^I}, \tag{4.36}$$

where the choice of w_{12} and w_{21} fixes w_{11} and w_{22} through (13) and (14). Note that if $u_1 = u_2$ and $w_1^I = w_2^I$, these constraints simplify to

$$\begin{aligned} w_{21} &= w_{11} - 1 \\ w_{12} &= w_{22} - 1. \end{aligned} \tag{4.37}$$

4.7 Connection between filtered activity and derivative-like terms

Consider the filtration of activity $x(t)$ via a decaying exponential with time constant τ.

$$\tilde{x}(t) = \int_0^t e^{-(t-s)/\tau} x(s)\, ds. \tag{4.38}$$

First, we define $v = (t - s)/\tau$. Substituting $s = t - \tau v$ and $dv = -ds/\tau$ and changing the integration bounds, we have:

$$\tilde{x}(t) = \int_t^0 e^{-v} x(t - \tau v)\, (-\tau dv) = \tau \int_0^t e^{-v} x(t - \tau v)\, dv \tag{4.39}$$

For small τ, $x(t - \tau v)$ will only contribute if v is small, thus we can substitute:

$$x(t-v) = x(t) - \tau v x'(t) \; + \; ... \tag{4.40}$$

This yields:

$$\tilde{x}(t) = \tau \int_0^t e^{-v} \left(x(t) - \tau v x'(t) \right) \, dv \tag{4.41}$$

Assuming $\tau \ll t$, we can use the identity:

$$\int_0^\infty v^n e^{-v} = n! \tag{4.42}$$

Thus,

$$\tilde{x}(t) = \tau x(t) - \tau^2 \dot{x}(t) \tag{4.43}$$

4.7.1 Approximations of three-point plasticity rules

To investigate the consequences of the effective filter imposed on input, pre-, and postsynaptic activity imposed by Fig. 4.7, we repeated the analysis performed in Xie and Seung, extended to include correlations between all three factors.

Explicitly, w_{ij} evolves according to a three-factor correlation function

$$\Delta w_{ij} = \int_{-\tau}^{\tau} \int_{-\tau}^{\tau} du \, dv \, C_{ij}(u,v) f(u,v) \tag{4.44}$$

where the three factor correlation is given by

$$C_{ij}(u,v) = \int_0^T dt \; z_i(t) \, x_i(t+u) \, x_j(t+v), \tag{4.45}$$

where u, v denote time lags between the post- and presynaptic activity and the input, respectively.

We now assume that x_i and x_j vary slowly compared to the timescale of the filter $f(u,v)$. Under this assumption,

$$\begin{aligned} x_i(t+u) &\approx x_i(t) + u\dot{x}_i(t), \\ x_j(t+v) &\approx x_j(t) + v\dot{x}_j(t), \end{aligned} \tag{4.46}$$

and

$$C_{ij}(u,v) = \int_0^T dt\, z_i(t)\, \left[x_i(t) + u\dot{x}_i(t)\right]\, \left[x_j(t) + v\dot{x}_j(t)\right]. \tag{4.47}$$

We drop the second order term and plug into Eq 4.44, yielding

$$\Delta w_{ij} = \int_0^T dt\, \left[\, \beta_{00} z_i x_j x_i + \beta_{10} z_i x_j \dot{x}_i + \beta_{01} z_i x_i \dot{x}_j\right], \tag{4.48}$$

where

$$\beta_{mn} = \iint du\, dv\, f(u,v)\, u^m\, v^n. \tag{4.49}$$

Example of effective derivative rule resulting from three-point kernel

As an example, we study the three-point function

$$f(u,v) = e^{-u/\tau_1}\, e^{-(u-v)/\tau_2}\, \Theta(u)\Theta(u-v) - \frac{1}{4} e^{-|u|/\tau_1} e^{-|v|/\tau_2}. \tag{4.50}$$

This kernel potentiates a synapse if postsynaptic activity that follows both presynaptic activity and third factor and depresses if activity occurs in a different order. We can find the first order coefficients by evaluating a series of integrals that reduce to gamma functions, e.g.

$$\begin{aligned} \int_0^\infty u e^{-u/\tau}\, du &= \tau^2 \int_0^\infty x e^{-x}\, dx \\ &= \tau^2 \Gamma(2). \end{aligned} \tag{4.51}$$

We then find

$$\begin{aligned} \beta_{00} &= 0, \\ \beta_{10} &= \tau_1^2 \tau_2, \\ \beta_{01} &= \tau_1 \tau_2 (\tau_1 - \tau_2). \end{aligned} \tag{4.52}$$

Under the conditions that $\tau_1 \approx \tau_2$, then $\beta_{01} \approx 0$, and we have

$$\Delta w_{ij} \approx \tau_1^2 \tau_2 \int_0^T dt\, z_i x_j \dot{x}_i. \tag{4.53}$$

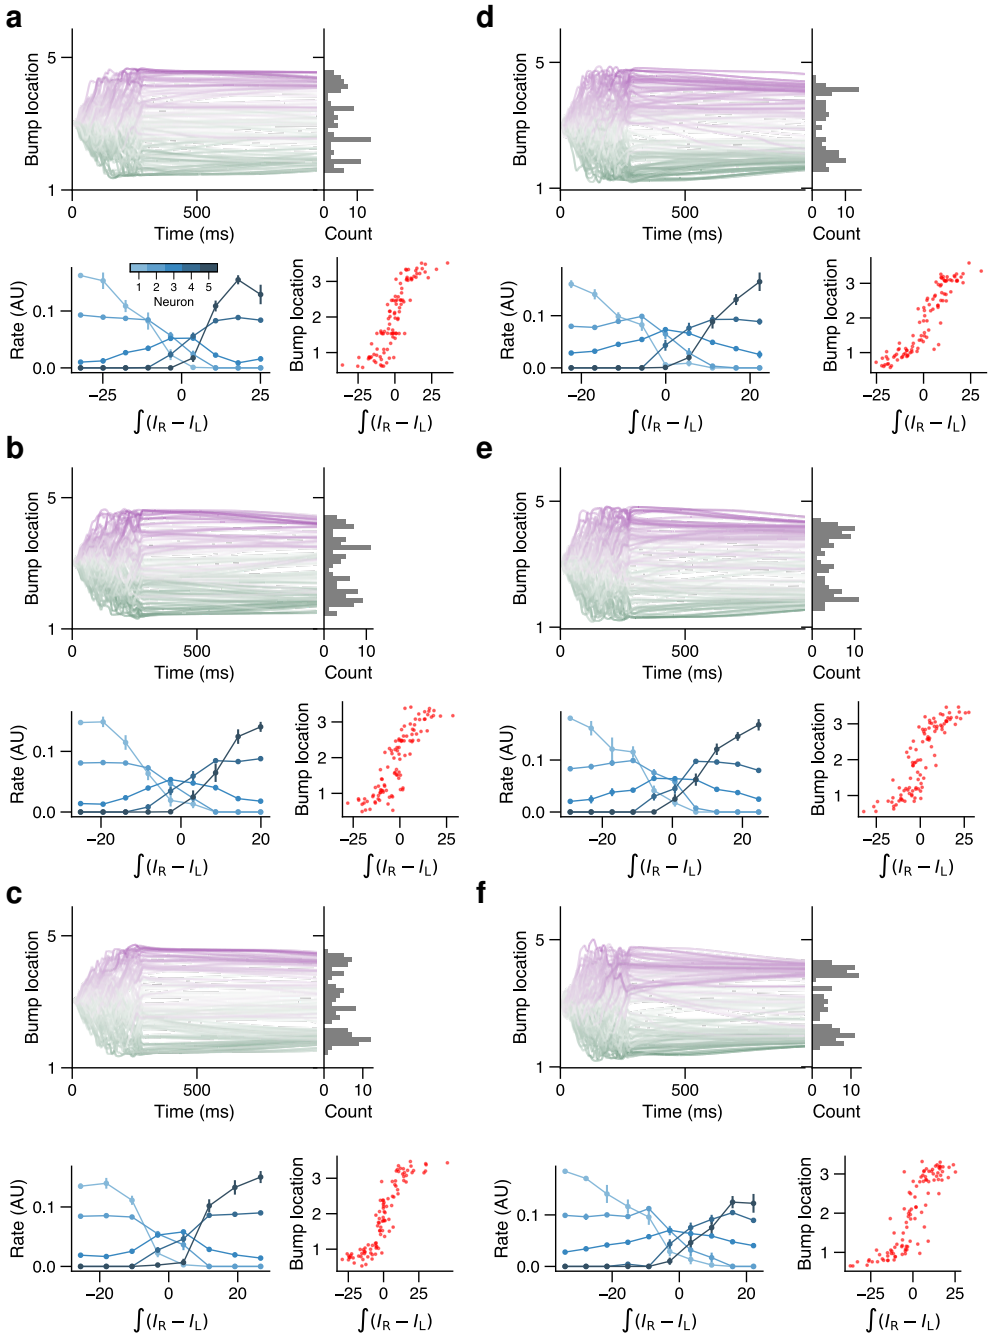

Figure 4.9: **Dynamics of several networks organized by two different learned plasticity rules trained on the long hold integration task**

(a-c) For three distinct network initializations organized by the same learned rule: (Top) Bump trajectories colored by integrated value (see Fig. 4.3). (Bottom left) Tuning curves for type 1 neurons. (Bottom right) Bump location plotted against integrated value for 100 trials. **(d-f)** Same as (a-c), but for a different learned rule.

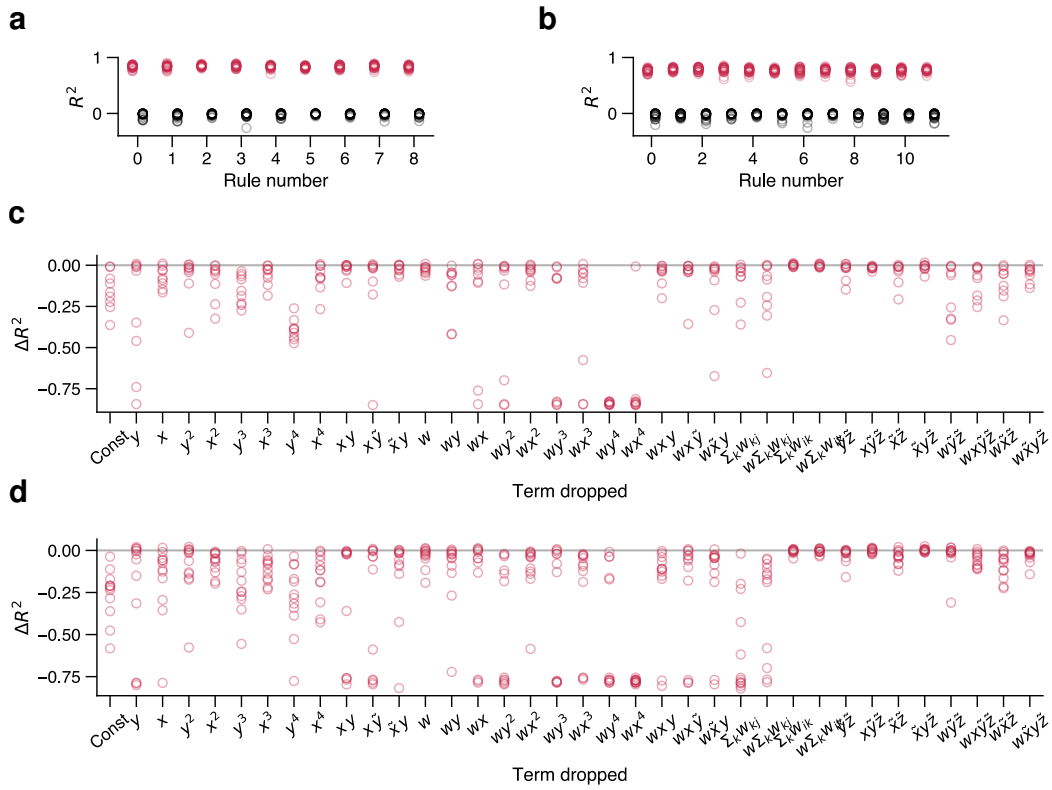

Figure 4.10: **Performance of trained rules and effect of term dropouts**

(a) Performance of all rules trained on short integration task (red) for 30 random network initializations. Black circles indicate the same initialization with no plasticity. **(b)** Same as (a), but for the long hold integration task. **(c)** Median change in R^2 that results from removing a term across 9 rules trained on short integration task. Only plasticity affecting type 1 recurrent synapses was removed. **(d)** Same as (c), but for the long hold integration task.

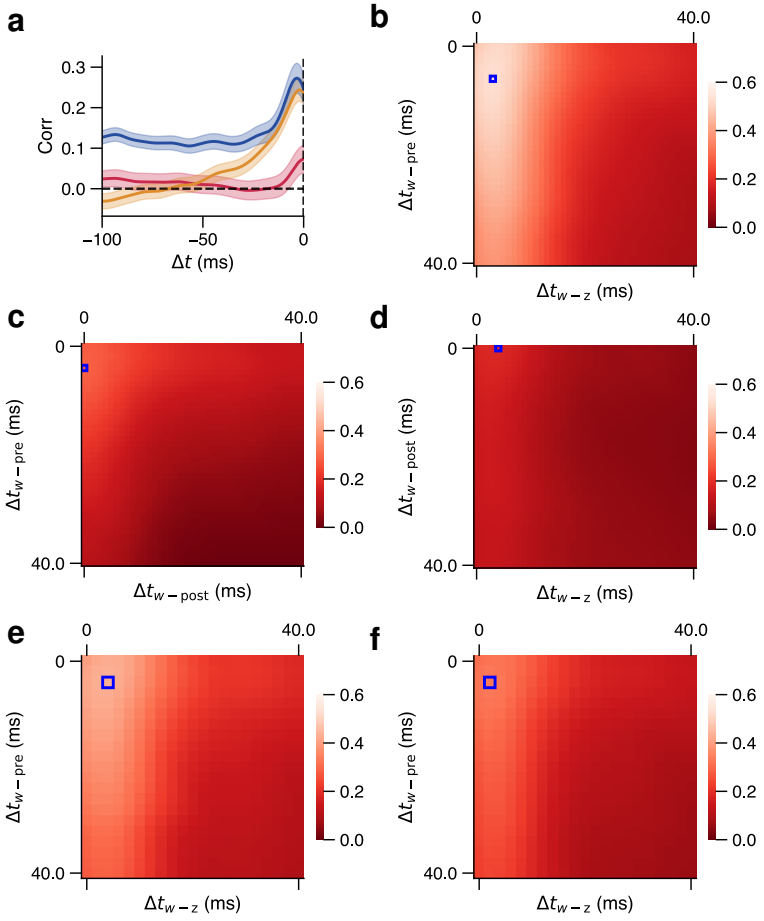

Figure 4.11: **Correlations between weight change and network activity over the final 10 epochs**

(a) Averaged correlations between weight updates and postsynaptic activity (red), presynaptic activity (gold), and third-factor input (blue). Each line represents the average over all type 1↔1 weights over the final 10 epochs of testing (epochs 410-420). **(b)** Averaged correlations between weight updates and the product of presynaptic activity and third factor input at two separate lags relative to the weight update, $\langle \Delta w_{ij}(t),\, x_j(t+\Delta t_{\mathrm{w-pre}}) z_i(t+\Delta t_{\mathrm{w-z}}) \rangle$. **(c)** Same as (b), but for product of post- and presynaptic activity. **(d)** Same as (b), but for product of postsynaptic activity and third factor input. **(e-f)** Averaged correlations between weight updates and the product of presynaptic activity, postsynaptic activity, and third factor input at three separate lags relative to the weight update. In (e), the post-weight update time lag is fixed to 0. In (f), -40 ms. See Table 4.3 for peak correlation values.

Correlation	Peak value (mean ± SEM)	$\Delta t_{\text{w-pre}}$ (ms)	$\Delta t_{\text{w-post}}$ (ms)	$\Delta t_{\text{w-z}}$ (ms)
$\langle x_i, \Delta w_{ij} \rangle$	0.07 ± 0.03	–	-0.1	–
$\langle x_j, \Delta w_{ij} \rangle$	0.24 ± 0.03	-2.0	–	–
$\langle z_i, \Delta w_{ij} \rangle$	0.27 ± 0.04	–	–	-3.4
$\langle x_i z_i, \Delta w_{ij} \rangle$	0.19 ± 0.04	–	0	-4
$\langle x_j z_i, \Delta w_{ij} \rangle$	0.52 ± 0.03	-6	–	-3
$\langle x_j x_i, \Delta w_{ij} \rangle$	0.29 ± 0.02	-4	0	–
$\langle x_j x_i z_i, \Delta w_{ij} \rangle$	0.43 ± 0.03	-4	0	-4

Table 4.3: Peak correlation values and time lags for plots shown in Fig. 4.11.

Chapter 5

DISCUSSION

This dissertation has explored the extent to which unsupervised plasticity rules can shape functional and resilient circuits in the absence of explicit feedback. Across projects, internal circuit dynamics were leveraged by neurons to modify synaptic strengths in ways that promoted desired network function. Together, these studies suggest that rich and robust computation can emerge from plasticity rules that are local, unsupervised, and sensitive to the structure of network activity.

5.1 Designing plasticity rules that explain recovery of singing behavior after viral perturbation of zebra finch HVC

In the first project, I examined whether standard and novel hand-crafted plasticity rules could account for recovery in a spiking network model of HVC following viral perturbation. While conventional homeostatic mechanisms could restore activity propagation in a feedforward network, they failed to reproduce the synaptic reorganization observed experimentally via mEPSPs and mIPSPs (miniature excitatory and inhibitory postsynaptic potentials). These measurements suggested that excitatory synaptic strength or number onto HVC(RA) neurons approximately doubled relative to baseline.

To account for this increase, we proposed a population-level homeostatic rule in which individual neurons sense the activity of nearby cells via secreted factors. This mechanism explained the large-scale increase in synaptic strength through recruitment of previously inactive neurons, thereby redistributing sequential structure across a broader population.

A primary challenge in this work was constructing an appropriate model of HVC. We focused on HVC(RA)-projecting neurons and interneurons; however, another major excitatory population exists in the nucleus: HVC(X) neurons, which project to basal ganglia and exhibit dynamics distinct from HVC(RA) neurons. These neurons are also known to interconnect with interneurons and HVC(RA) cells [133]. Although omitted here for tractability, their

inclusion may influence recovery dynamics and plasticity requirements.

We also incorporated long synaptic delays between HVC(RA) neurons, a relatively recent finding inferred largely through indirect experimental approaches [120]. Model architecture proved critical to recovery. Successful restoration of sequential activity required sufficient network width, particularly when the connection probability between adjacent HVC(RA) neurons was substantially below one. Additionally, as in many excitatory–inhibitory circuit models, relatively fast inhibition was essential to constrain runaway excitation. When excitation and inhibition operated on comparable timescales, the propagating mode fluctuated in amplitude and could become unstable.

Another substantial challenge was establishing the initial stability of candidate plasticity rules. Although the stability of rate-based models can sometimes be analyzed analytically, spiking networks present additional complexity. A striking example emerged when we found that pairwise spike-timing-dependent plasticity (STDP) failed to maintain stable sequential dynamics when neurons fired in bursts. Previous studies had addressed this issue by counting a burst as a single spike. However, in the context of recovery, this approximation proved problematic: because STDP enforces sequential ordering, variability in spike timing following perturbation could reorder neurons within the chain. A solution emerged in the triplet-based plasticity model introduced by Pfister and Gerstner. This formulation reinforced spike ordering while prioritizing bursts over isolated spikes, thereby stabilizing sequential dynamics under realistic firing patterns.

The introduction of a population-level homeostatic mechanism represents one of the more exploratory aspects of this work. While there is evidence that homeostatic plasticity does not operate purely at the level of individual neurons, direct experimental support for population-level homeostasis remains limited. Nevertheless, this modeling work motivates experimental investigation into such mechanisms. Population-level homeostasis may be particularly useful in circuits where neurons serve largely interchangeable roles, facilitating functional replacement or reorganization following perturbation.

A notable feature of this proposed mechanism is its tethering to physical space. Plasticity rules based solely on pre- and postsynaptic activity abstract away from the spatial arrangement of neurons. In contrast, a rule mediated by diffusible factors necessarily depends on

spatial proximity, linking anatomical organization to learning dynamics. We are particularly interested in future work exploring connections between population-level mechanisms and representational drift.

5.2 Meta-learning plasticity rules that organize and maintain sequential dynamics

In the second project, I employed meta-learning to discover plasticity rules capable of self-organizing and maintaining sequential dynamics. With this tool, we found plasticity rules that accomplished this by generalizing Oja's rule, replacing Hebbian learning with a form of spike timing-dependent plasticity. To extract interpretable rules from the meta-learned parameterizations, we developed a pruning procedure in which individual terms were systematically removed and the remaining terms refit in order of importance. This approach revealed compact rules that retained performance while clarifying their functional components. We additionally found we could train rules that were robust to structural alterations; plasticity rules accomplished this by flexibly adapting the strength of connections. Finally, learning additional plasticity on E→I and I→E synapses led to more robust sequential dynamics.

The motivation for this project was largely the sense that space of plasticity rules explored in the first project was incomplete. Confavreux et al. provided inspiration for how one might search for plasticity more generally, without exhaustive search through a catalog of hand-picked rules.

An initial challenge was selection of the loss function: in many cases, meta-learning relieves the researcher of hand-tuning a plasticity rule and replaces this task with hand-tuning a loss function. Initial efforts in this project experimented with the idea of using a supervised template of neural activity as the objective of the loss function. We found, in many cases, that this led to erratic network behavior, as learned plasticity attempted to fit exactly the template provided. In particular, if the target activity of the network cannot be naturally generated under a reasonable set of synaptic weights, which can be difficult to determine *a priori*, plasticity rules may become immensely complex. One can always wire up a network that generates the intended dynamics, but then one has prescribed how the network should

solve a given task.

We found a much more natural way to frame a loss function was in terms of a decoder, which attempted to match a supervised target by reading out from network activity. In this project, we chose to decode elapsed time since an initial stimulus was presented to the network. This objective well captures the goal of a time-keeping network like HVC, while remaining agnostic to how it should be reached. The use of a linear decoder likely forces the network to have a straightforward representation of the task, and also can be trained from a relatively small number of examples.

Regularization of network activity further prevented degenerate solutions, such as slow global ramping across the entire population. Although originally conceived as a sparsity penalty on plasticity-rule terms, regularization also encouraged convergence of synaptic weights and stabilized learning dynamics.

Meta-learning in this case recovered a number of sensible aspects of the plasticity that might shape HVC. First, the bounding term—the same as Oja's—is second order in the postsynaptic firing rate and first order in the size of the synapse. This term presents a natural bound for STDP or any form of Hebbian learning which goes as the product of the pre- and postsynaptic activity [107]. This bound additionally prevents winner-take-all dynamics in terms of the size of synapses that target a postsynaptic cell (see Section 3.9). This allows the sequential structure of the network to be distributed amongst the many neurons in the network, rendering it more robust to perturbation. Finally, the bound is mediated by activity, naturally enabling homeostasis. Loss of afferent connections to a postsynaptic neuron can thus be compensated for by the strengthening of the remaining connections.

Despite its apparent utility, we should remain skeptical of meta-learning and the synaptic plasticity rules it produces. Discovered rules may be extremely tailored to the settings in which they are discovered. We attempted to mitigate this by training rules in batch on many different random initializations. Still, it remains possible that rules may fail if transferred to new settings. A variable we largely did not randomize in this work was the size of the different neuronal populations in a network, as well as the total size of the network. While we found that varying this did not significantly impact trained rule performance in this case, it remains possible that network size may have a bigger impact on other tasks. Future work

may benefit from varying this parameter during the training phase of meta-learning.

5.3 Meta-learning plasticity that organizes neural integrators

In the final project, I moved beyond sequential dynamics and applied meta-learning to a problem of general interest in the field: the organization of neural integrators. Starting from networks whose structural priors loosely resemble the central complex found in Drosophila, we searched for plasticity rules that could organize networks capable of integrating a series of incoming sensory cues without supervised feedback. We found that unsupervised plasticity was capable of performing path integration and maintaining integrated displacement on the timescale of approximately one second.

The discovered plasticity rules acting on excitatory-to-excitatory synapses did not consistently leverage the form of anti-Hebbian learning theorized to suppress drift in activity. Instead, across learned rules, three-factor plasticity displayed consistent patterns, including potentiation as a function of presynaptic activity and third-factor input, and depression as a function of postsynaptic activity and third-factor input. We additionally demonstrated that plasticity was sensitive to the temporal order of activity. We hypothesized that forms of three-factor plasticity may be sufficient to organize neural integrators and used one such rule to demonstrate robust self-organization of a neural integrator in a simplified circuit.

A novel aspect of this work is that it explores the organization of neural integrators in the few-neuron limit. The discrepancy between continuous and discrete attractors becomes less salient in the thermodynamic limit, since larger networks have many more unique states by virtue of their size. This is not the case in small networks, where establishing a continuum of stable states relies entirely on fine-tuning of the network [181]. Here, we attempted to explain how such tuning could emerge via unsupervised plasticity. Plasticity rules accomplished this surprisingly well given the relatively short period dedicated to network self-organization: 280 trials during which the network received inputs to be integrated. Future work might explore how performance depends on the number of trials available for self-organization before testing.

An initial motivation for this project was to explore the coexistence of heterogeneity in connectivity and attractor dynamics. While it is known that heterogeneous neural circuits

can support continuous attractors [177, 178], existing explanations for how these circuits are organized typically invoke supervised learning. In mice, such heterogeneity is associated with substantial variability in the tuning curves of head direction neurons [177]. We hope to extend the meta-learning approach introduced here to cases in which the circuit is explicitly constrained to be heterogeneous, for example by varying the strength of global inhibition across cells. The discovery of such plasticity mechanisms may provide insight into how representations of integrated variables emerge in circuits early in development.

BIBLIOGRAPHY

[1] Donald O. Hebb. *The Organization of Behavior*. Wiley, 1949.

[2] V. Castellucci, H. Pinsker, I. Kupfermann, and E. R. Kandel. Neuronal mechanisms of habituation and dishabituation of the gill-withdrawal reflex in Aplysia. *Science*, 167 (3926):1745–1748, March 1970. ISSN 0036-8075. doi: 10.1126/science.167.3926.1745.

[3] M. Brunelli, V. Castellucci, and E. R. Kandel. Synaptic facilitation and behavioral sensitization in Aplysia: possible role of serotonin and cyclic AMP. *Science*, 194(4270): 1178–1181, December 1976. ISSN 0036-8075. doi: 10.1126/science.186870.

[4] T. V. Bliss and T. Lomo. 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, July 1973. ISSN 0022-3751. doi: 10.1113/jphysiol.1973.sp010273.

[5] Guo-qiang Bi and Mu-ming Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing. *Journal of Neuroscience*, 18:10464–10472, 1998.

[6] Gina G. Turrigiano, Kim R. Leslie, Niraj S. Desai, Lara C. Rutherford, and Sacha B. Nelson. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391:892–896, 1998.

[7] Łukasz Kuśmierz, Takuya Isomura, and Taro Toyoizumi. Learning with three factors: modulating Hebbian plasticity with errors. *Current Opinion in Neurobiology*, 46: 170–177, October 2017. ISSN 0959-4388. doi: 10.1016/j.conb.2017.08.020. URL `https://www.sciencedirect.com/science/article/pii/S0959438817300612`.

[8] Roger A. Nicoll. A Brief History of Long-Term Potentiation. *Neuron*, 93(2):281–290, January 2017. ISSN 1097-4199. doi: 10.1016/j.neuron.2016.12.015.

[9] William B. Levy and Oswald Steward. Synapses as associative memory elements in the hippocampal formation. *Brain Research*, 175(2):233–245, October 1979. ISSN 0006-8993. doi: 10.1016/0006-8993(79)91003-5. URL https://www.sciencedirect.com/science/article/pii/0006899379910035.

[10] B. L. McNaughton, R. M. Douglas, and G. V. Goddard. Synaptic enhancement in fascia dentata: Cooperativity among coactive afferents. *Brain Research*, 157(2):277–293, November 1978. ISSN 0006-8993. doi: 10.1016/0006-8993(78)90030-6. URL https://www.sciencedirect.com/science/article/pii/0006899378900306.

[11] R. Malinow and J. P. Miller. Postsynaptic hyperpolarization during conditioning reversibly blocks induction of long-term potentiation. *Nature*, 320(6062):529–530, April 1986. ISSN 0028-0836. doi: 10.1038/320529a0.

[12] S. R. Kelso and T. H. Brown. Differential conditioning of associative synaptic enhancement in hippocampal brain slices. *Science*, 232(4746):85–87, April 1986. ISSN 0036-8075. doi: 10.1126/science.3952501.

[13] Ami Citri and Robert C. Malenka. Synaptic Plasticity: Multiple Forms, Functions, and Mechanisms. *Neuropsychopharmacology*, 33(1):18–41, January 2008. ISSN 1740-634X. doi: 10.1038/sj.npp.1301559. URL https://www.nature.com/articles/1301559.

[14] Daniel E. Feldman. The Spike-Timing Dependence of Plasticity. *Neuron*, 75(4):556–571, August 2012. ISSN 0896-6273. doi: 10.1016/j.neuron.2012.08.001. URL https://www.sciencedirect.com/science/article/pii/S0896627312007039.

[15] R. G. Morris, E. Anderson, G. S. Lynch, and M. Baudry. Selective impairment of learning and blockade of long-term potentiation by an N-methyl-D-aspartate receptor antagonist, AP5. *Nature*, 319(6056):774–776, March 1986. ISSN 0028-0836. doi: 10.1038/319774a0.

[16] R. G. Morris. Synaptic plasticity and learning: selective impairment of learning rats and blockade of long-term potentiation in vivo by the N-methyl-D- aspartate receptor antagonist AP5. *Journal of Neuroscience*, 9(9):3040–3057, September 1989. ISSN

0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.09-09-03040.1989. URL https://www.jneurosci.org/content/9/9/3040.

[17] W. Gerstner, R. Kempter, J. L. van Hemmen, and H. Wagner. A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383(6595):76–81, September 1996. ISSN 0028-0836. doi: 10.1038/383076a0.

[18] H. Markram, J. Lübke, M. Frotscher, and B. Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275(5297):213–215, January 1997. ISSN 0036-8075. doi: 10.1126/science.275.5297.213.

[19] Curtis C. Bell, Angel Caputi, and Kirsty Grant. Physiology and Plasticity of Morphologically Identified Cells in the Mormyrid Electrosensory Lobe. *Journal of Neuroscience*, 17(16):6409–6423, August 1997. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.17-16-06409.1997. URL https://www.jneurosci.org/content/17/16/6409.

[20] Jiang-teng Lu, Cheng-yu Li, Jian-Ping Zhao, Mu-ming Poo, and Xiao-hui Zhang. Spike-Timing-Dependent Plasticity of Neocortical Excitatory Synapses on Inhibitory Interneurons Depends on Target Cell Type. *Journal of Neuroscience*, 27(36):9711–9720, September 2007. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.2513-07.2007. URL https://www.jneurosci.org/content/27/36/9711.

[21] Jean-Pascal Pfister and Wulfram Gerstner. Triplets of Spikes in a Model of Spike Timing-Dependent Plasticity. *Journal of Neuroscience*, 26(38):9673–9682, September 2006. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.1425-06.2006. URL https://www.jneurosci.org/content/26/38/9673.

[22] B. Katz and R. Miledi. A study of synaptic transmission in the absence of nerve impulses. *The Journal of Physiology*, 192(2):407–436, September 1967. ISSN 0022-3751. doi: 10.1113/jphysiol.1967.sp008307. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC1365564/.

[23] Robert S. Zucker and Wade G. Regehr. Short-Term Synaptic Plasticity. *Annual*

Review of Physiology, 2002. doi: 10.1146/annurev.physiol.64.092501.114547. URL https://sci-hub.ru/10.1146/annurev.physiol.64.092501.114547.

[24] B. Katz and R. Miledi. The role of calcium in neuromuscular facilitation. *The Journal of Physiology*, 195(2):481–492, 1968. ISSN 1469-7793. doi: 10.1113/jphysiol.1968.sp008469. URL https://onlinelibrary.wiley.com/doi/abs/10.1113/jphysiol.1968.sp008469. _eprint: https://physoc.onlinelibrary.wiley.com/doi/pdf/10.1113/jphysiol.1968.sp008469.

[25] Martin Geppert, Yukiko Goda, Robert E. Hammer, Cai Li, Thomas W. Rosahl, Charles F. Stevens, and Thomas C. Südhof. Synaptotagmin I: A major Ca2+ sensor for transmitter release at a central synapse. *Cell*, 79(4):717–727, November 1994. ISSN 00928674. doi: 10.1016/0092-8674(94)90556-8. URL https://linkinghub.elsevier.com/retrieve/pii/0092867494905568.

[26] A. Rozov, N. Burnashev, B. Sakmann, and E. Neher. Transmitter release modulation by intracellular Ca2+ buffers in facilitating and depressing nerve terminals of pyramidal cells in layer 2/3 of the rat neocortex indicates a target cell-specific difference in presynaptic calcium dynamics. *The Journal of Physiology*, 531(Pt 3):807–826, March 2001. ISSN 0022-3751. doi: 10.1111/j.1469-7793.2001.0807h.x.

[27] Misha V. Tsodyks and Henry Markram. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences*, 94(2):719–723, January 1997. doi: 10.1073/pnas.94.2.719. URL https://www.pnas.org/doi/abs/10.1073/pnas.94.2.719.

[28] Gianluigi Mongillo, Omri Barak, and Misha Tsodyks. Synaptic Theory of Working Memory. *Science*, 319(5869):1543–1546, March 2008. doi: 10.1126/science.1150769. URL https://www.science.org/doi/10.1126/science.1150769. Publisher: American Association for the Advancement of Science.

[29] S.M. Stringer, T.P. Trappenberg, E.T. Rolls, and I.E.T.d. Araujo. Self-organizing continuous attractor networks and path integration: one-dimensional models of head

direction cells. *Network: Computation in Neural Systems*, 13(2):217–242, January 2002. ISSN 0954-898X, 1361-6536. doi: 10.1080/net.13.2.217.242. URL `https://www.tandfonline.com/doi/full/10.1080/net.13.2.217.242`.

[30] Gina Turrigiano. Too Many Cooks? Intrinsic and Synaptic Homeostatic Mechanisms in Cortical Circuit Refinement. *Annual Review of Neuroscience*, 2011. doi: 10.1146/annurev-neuro-060909-153238. URL `https://sci-hub.ru/10.1146/annurev-neuro-060909-153238`.

[31] Eve Marder and Jean-Marc Goaillard. Variability, compensation and homeostasis in neuron and network function. *Nature Reviews. Neuroscience*, 7(7):563–574, July 2006. ISSN 1471-003X. doi: 10.1038/nrn1949.

[32] Wei Zhang and David J. Linden. The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nature Reviews. Neuroscience*, 4(11):885–900, November 2003. ISSN 1471-003X. doi: 10.1038/nrn1248.

[33] N. S. Desai, L. C. Rutherford, and G. G. Turrigiano. Plasticity in the intrinsic excitability of cortical pyramidal neurons. *Nature Neuroscience*, 2(6):515–520, June 1999. ISSN 1097-6256. doi: 10.1038/9165.

[34] Yue Kris Wu, Keith B. Hengen, Gina G. Turrigiano, and Julijana Gjorgjieva. Homeostatic mechanisms regulate distinct aspects of cortical circuit dynamics. *Proceedings of the National Academy of Sciences*, 117(39):24514–24525, September 2020. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1918368117. URL `https://pnas.org/doi/full/10.1073/pnas.1918368117`.

[35] Michael A. Farries and Adrienne L. Fairhall. Reinforcement learning with modulated spike timing–dependent synaptic plasticity. *Journal of neurophysiology*, 98(6):3648–3665, 2007. doi: 10.1152/jn.00364.2007.

[36] Nicolas Frémaux, Henning Sprekeler, and Wulfram Gerstner. Functional requirements for reward-modulated spike-timing-dependent plasticity. *Journal of Neuroscience*, 30

(40):13326–13337, 2010. URL https://www.jneurosci.org/content/30/40/13326?utm_source=TrendMD&utm_medium=cpc&utm_campaign=JNeurosci_TrendMD_0.

[37] Nicolas Frémaux and Wulfram Gerstner. Neuromodulated Spike-Timing-Dependent Plasticity, and Theory of Three-Factor Learning Rules. *Frontiers in Neural Circuits*, 9: 85, 2015. ISSN 1662-5110. doi: 10.3389/fncir.2015.00085.

[38] Vincent Paille, Elodie Fino, Kai Du, Teresa Morera-Herreras, Sylvie Perez, Jeanette Hellgren Kotaleski, and Laurent Venance. GABAergic circuits control spike-timing-dependent plasticity. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 33(22):9353–9363, May 2013. ISSN 1529-2401. doi: 10.1523/JNEUROSCI.5796-12.2013.

[39] Katie C. Bittner, Aaron D. Milstein, Christine Grienberger, Sandro Romani, and Jeffrey C. Magee. Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science*, 357(6355):1033–1036, September 2017. doi: 10.1126/science.aan3846. URL https://www.science.org/doi/10.1126/science.aan3846.

[40] Johannes D. Seelig and Vivek Jayaraman. Neural dynamics for landmark orientation and angular path integration. *Nature*, 521(7551):186–191, May 2015. ISSN 0028-0836. doi: 10.1038/nature14446. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC4704792/.

[41] Daniel Turner-Evans, Stephanie Wegener, Hervé Rouault, Romain Franconville, Tanya Wolff, Johannes D Seelig, Shaul Druckmann, and Vivek Jayaraman. Angular velocity integration in a fly heading circuit. *eLife*, 6:e23496, May 2017. ISSN 2050-084X. doi: 10.7554/eLife.23496. URL https://doi.org/10.7554/eLife.23496.

[42] Daniel B. Turner-Evans, Kristopher T. Jensen, Saba Ali, Tyler Paterson, Arlo Sheridan, Robert P. Ray, Tanya Wolff, Scott Lauritzen, Gerald M. Rubin, Davi Bock, and Vivek Jayaraman. The neuroanatomical ultrastructure and function of a biological ring attractor. *Neuron*, 108(1):145–163.e10, October 2020. ISSN 0896-6273. doi:

10.1016/j.neuron.2020.08.006. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8356802/.

[43] Mark H. Plitt, Daniel B. Turner-Evans, Jessica C. Co, Aryanna Layden, Mark Eddison, Robert P. Ray, Vivek Jayaraman, and Yvette E. Fisher. Octopamine instructs head direction plasticity. *bioRxiv*, pages 2025–12, 2025. URL https://www.biorxiv.org/content/10.64898/2025.12.11.693783.abstract.

[44] W. Hamish Mehaffey and Allison J. Doupe. Naturalistic stimulation drives opposing heterosynaptic plasticity at two inputs to songbird cortex. *Nature Neuroscience*, 18(9):1272–1280, September 2015. ISSN 1546-1726. doi: 10.1038/nn.4078. URL https://www.nature.com/articles/nn.4078.

[45] Anthony M. Zador. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications*, 10(1):3770, August 2019. ISSN 2041-1723. doi: 10.1038/s41467-019-11786-6.

[46] James P. Porter. The Habits, Instincts, and Mental Powers of Spiders, Genera, Argiope and Epeira. *The American Journal of Psychology*, 1906. doi: 10.2307/1412250. URL https://sci-hub.ru/10.2307/1412250.

[47] Alessio Attardo, James E. Fitzgerald, and Mark J. Schnitzer. Impermanence of dendritic spines in live adult CA1 hippocampus. *Nature*, 523(7562):592–596, July 2015. ISSN 1476-4687. doi: 10.1038/nature14467.

[48] Anthony J. G. D. Holtmaat, Joshua T. Trachtenberg, Linda Wilbrecht, Gordon M. Shepherd, Xiaoqun Zhang, Graham W. Knott, and Karel Svoboda. Transient and persistent dendritic spines in the neocortex in vivo. *Neuron*, 45(2):279–291, January 2005. ISSN 0896-6273. doi: 10.1016/j.neuron.2005.01.003.

[49] Abbott, Varela, Sen, and Nelson. Synaptic Depression and Cortical Gain Control. *Science*, 1997. doi: 10.1126/science.275.5297.221. URL https://sci-hub.ru/10.1126/science.275.5297.221.

[50] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, October 1986. ISSN 1476-4687. doi: 10.1038/323533a0. URL https://www.nature.com/articles/323533a0.

[51] Basile Confavreux, Friedemann Zenke, Everton J. Agnes, Timothy Lillicrap, and Tim P. Vogels. A meta-learning approach to (re)discover plasticity rules that carve a desired function into a neural network. In *34th Conference on Neural Information Processing Systems*, Vancouver, Canada, 2020.

[52] Y. Bengio, S. Bengio, and J. Cloutier. Learning a synaptic learning rule. ii:969 vol.2–, 1991. doi: 10.1109/IJCNN.1991.155621.

[53] Thomas Miconi, Jeff Clune, and Kenneth O. Stanley. Differentiable plasticity: training plastic neural networks with backpropagation, July 2018. URL http://arxiv.org/abs/1804.02464. arXiv:1804.02464.

[54] Basile Confavreux, Everton J. Agnes, Friedemann Zenke, Henning Sprekeler, and Tim P. Vogels. Balancing complexity, performance and plausibility to meta learn plasticity rules in recurrent spiking networks, June 2024. URL https://www.biorxiv.org/content/10.1101/2024.06.17.599260v1. Pages: 2024.06.17.599260 Section: New Results.

[55] Poornima Ramesh, Basile Confavreux, Pedro J. Gonçalves, Tim P. Vogels, and Jakob H. Macke. Indistinguishable network dynamics can emerge from unalike plasticity rules. *eLife*, 13, March 2024. doi: 10.7554/eLife.94411.1. URL https://elifesciences.org/reviewed-preprints/94411.

[56] Jack Lindsey and Ashok Litwin-Kumar. Learning to Learn with Feedback and Local Plasticity. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 21213–21223. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/f291e10ec3263bd7724556d62e70e25d-Paper.pdf.

[57] N. Shervani-Tabar and R. Rosenbaum. Meta-learning biologically plausible plasticity

rules with random feedback pathways. *Nat Commun*, 14(1):1805, 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-37562-1.

[58] Jakob Jordan, Maximilian Schmidt, Walter Senn, and Mihai A Petrovici. Evolving interpretable plasticity for spiking networks. *eLife*, 10:e66273, October 2021. ISSN 2050-084X. doi: 10.7554/eLife.66273. URL `https://doi.org/10.7554/eLife.66273`.

[59] Thomas Miconi. Learning to acquire novel cognitive tasks with evolution, plasticity and meta-meta-learning. In *Proceedings of the 40th International Conference on Machine Learning*, pages 24756–24774. PMLR, July 2023. URL `https://proceedings.mlr.press/v202/miconi23a.html`.

[60] Basile Confavreux, Poornima Ramesh, Pedro J Gonçalves, Jakob H Macke, and Tim P Vogels. Meta-learning families of plasticity rules in recurrent spiking networks using simulation-based inference. .

[61] D. Tyulmankov, G. R. Yang, and L. F. Abbott. Meta-learning synaptic plasticity and memory addressing for continual familiarity detection. *Neuron*, 110(3):544–557.e8, 2022. ISSN 0896-6273 (Print) 0896-6273. doi: 10.1016/j.neuron.2021.11.009.

[62] Luke Metz, Niru Maheswaranathan, Brian Cheung, and Jascha Sohl-Dickstein. Meta-Learning Update Rules for Unsupervised Representation Learning, February 2019. URL `http://arxiv.org/abs/1804.00222`. arXiv:1804.00222.

[63] Elias Najarro and Sebastian Risi. Meta-Learning through Hebbian Plasticity in Random Networks, April 2022. URL `http://arxiv.org/abs/2007.02686`. arXiv:2007.02686.

[64] Sergey Shuvaev, Divyansha Lachi, Alexei Koulakov, and Anthony Zador. Encoding innate ability through a genomic bottleneck. *Proceedings of the National Academy of Sciences*, 121(38):e2409160121, September 2024. doi: 10.1073/pnas.2409160121. URL `https://www.pnas.org/doi/abs/10.1073/pnas.2409160121`. Publisher: Proceedings of the National Academy of Sciences.

[65] Qiufeng Wang, Xu Yang, Shuxia Lin, Jing Wang, and Xin Geng. Learngene: Inheriting

Condensed Knowledge from the Ancestry Model to Descendant Models, June 2023. URL http://arxiv.org/abs/2305.02279. arXiv:2305.02279 [cs].

[66] Nikhil X. Bhattasali, Venkatesh Pattabiraman, Lerrel Pinto, and Grace W. Lindsay. Neural Circuit Architectural Priors for Quadruped Locomotion, October 2024. URL http://arxiv.org/abs/2410.07174. arXiv:2410.07174.

[67] B. Wang, Z. Torok, A. Duffy, D. G. Bell, S. Wongso, T. A. F. Velho, A. L. Fairhall, and C. Lois. Unsupervised restoration of a complex learned behavior after large-scale neuronal perturbation. *Nat Neurosci*, 2024. ISSN 1097-6256. doi: 10.1038/s41593-024-01630-6.

[68] Michael A. Long, Dezhe Z. Jin, and Michale S. Fee. Support for a synaptic chain model of neuronal sequence generation. *Nature*, 468(7322):394–399, November 2010. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature09514. URL http://www.nature.com/articles/nature09514.

[69] Allison J. Doupe and Patricia K. Kuhl. Birdsong and human speech: common themes and mechanisms. *Annual Review of Neuroscience*, 22:567–631, 1999.

[70] Klaus Immelmann. Song development in the zebra finch and other estrildid finches. *Bird Vocalizations*, pages 61–74, 1969.

[71] Philippe Adret. Operant conditioning, song learning and imprinting to taped song in the zebra finch. *Animal Behaviour*, 46(1):149–159, 1993.

[72] Ofer Tchernichovski, Partha P. Mitra, Thomas Lints, and Fernando Nottebohm. Dynamics of the vocal imitation process: how a zebra finch learns its song. *Science*, 291 (5513):2564–2569, 2001.

[73] Philip H. Price. Developmental determinants of structure in zebra finch song. *Journal of Comparative and Physiological Psychology*, 93(2):260–277, 1979.

[74] H. Williams, K. Kilander, and M. L. Sotanski. Song learning in zebra finches: the role of social factors. *Animal Behaviour*, 46:215–230, 1993.

[75] Masakazu Konishi. The role of auditory feedback in the control of vocalization in the white-crowned sparrow. *Zeitschrift für Tierpsychologie*, 22:770–783, 1965.

[76] Yoko Funabiki and Masakazu Konishi. Long memory in song learning by zebra finches. *Journal of Neuroscience*, 23(17):6928–6935, 2003.

[77] Ila R. Fiete, Michale S. Fee, and H. Sebastian Seung. Gradient estimation through dynamic perturbation of neural activity. *Physical Review Letters*, 98(4):048104, 2007.

[78] Kenji Doya and Terrence J. Sejnowski. A novel reinforcement model of birdsong vocalization learning. *Advances in Neural Information Processing Systems*, 7, 1995.

[79] Michale S. Fee and Jesse H. Goldberg. A hypothesis for basal ganglia-dependent reinforcement learning in the songbird. *Neuroscience*, 198:152–170, 2011.

[80] Andrew Duffy et al. Dopamine neuron activity predicts song learning in zebra finches. *Nature*, 605:292–297, 2022.

[81] Vikram Gadagkar et al. Dopamine neurons encode performance error in singing birds. *Science*, 354(6317):1278–1282, 2016.

[82] H. B. Simpson and D. S. Vicario. Brain pathways for learned and unlearned vocalizations differ in zebra finches. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 10(5):1541–1556, May 1990. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.10-05-01541.1990.

[83] F. Nottebohm, T. M. Stokes, and C. M. Leonard. Central control of song in the canary, Serinus canarius. *The Journal of Comparative Neurology*, 165(4):457–486, February 1976. ISSN 0021-9967. doi: 10.1002/cne.901650405.

[84] S. W. Bottjer, E. A. Miesner, and A. P. Arnold. Forebrain lesions disrupt development but not maintenance of song in passerine birds. *Science*, 224(4651):901–903, May 1984. ISSN 0036-8075. doi: 10.1126/science.6719123.

[85] Farida Sohrabji, Ernest J. Nordeen, and Kathy W. Nordeen. Selective impairment of song learning following lesions of a forebrain nucleus in the juvenile zebra finch.

Behavioral and Neural Biology, 53(1):51–63, January 1990. ISSN 0163-1047. doi: 10.1016/0163-1047(90)90797-A. URL https://www.sciencedirect.com/science/article/pii/016310479090797A.

[86] Alexei A. Kozhevnikov and Michale S. Fee. Precise timing of neuronal firing in a songbird motor circuit. *Nature*, 450:1096–1100, 2007.

[87] Galen F. Lynch, Tatsuo S. Okubo, Alexander Hanuschkin, Richard H. R. Hahnloser, and Michale S. Fee. Rhythmic Continuous-Time Coding in the Songbird Analog of Vocal Motor Cortex. *Neuron*, 90(4):877–892, May 2016. ISSN 0896-6273. doi: 10.1016/j.neuron.2016.04.021. URL https://www.cell.com/neuron/abstract/S0896-6273(16)30108-8.

[88] Ila R. Fiete, Richard H.R. Hahnloser, Michale S. Fee, and H. Sebastian Seung. Temporal sparseness of the premotor drive is important for rapid learning in a neural network model of birdsong. *Journal of neurophysiology*, 92:2274–2282, 2004. doi: https://doi.org/10.1152/jn.01133.2003.

[89] W Nicola and C Clopath. Supervised learning in spiking neural networks with force training. *Nat Commun*, 8(1):2208, 2017. ISSN 2041-1723. doi: 10.1038/s41467-017-01827-3.

[90] A. C. Yu and D. Margoliash. Temporal patterns of neural activity in the songbird motor nucleus ra. *Journal of Neurophysiology*, 76:1735–1747, 1996.

[91] Zhiyong Chi and Daniel Margoliash. Temporal precision and temporal drift in brain and behavior of zebra finch song. *Neuron*, 32:899–910, 2001.

[92] Anthony Leonardo and Michale S. Fee. Ensemble coding of vocal control in birdsong. *Journal of Neuroscience*, 25:652–661, 2005.

[93] R. H. Hahnloser, A. A. Kozhevnikov, and M. S. Fee. An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature*, 419(6902):65–70, 2002. ISSN 0028-0836 (Print) 0028-0836. doi: 10.1038/nature00974.

[94] Michael A. Long and Michale S. Fee. Using temperature to analyse temporal dynamics in the songbird motor pathway. *Nature*, 456:189–194, 2008.

[95] G. Kosche, D. Vallentin, and M. A. Long. Interplay of inhibition and excitation shapes a premotor neural sequence. *J Neurosci*, 35(3):1217–27, 2015. ISSN 0270-6474 (Print) 0270-6474. doi: 10.1523/jneurosci.4346-14.2015.

[96] J. Cannon, N. Kopell, T. Gardner, and J. Markowitz. Neural sequence generation using spatiotemporal patterns of inhibition. *PLoS Comput Biol*, 11(11):e1004581, 2015. ISSN 1553-734X (Print) 1553-734x. doi: 10.1371/journal.pcbi.1004581.

[97] Moshe Abeles. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, 1991.

[98] Markus Diesmann, Marc-Oliver Gewaltig, and Ad Aertsen. Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402:529–533, 1999.

[99] MengRu Li and Henry S. Greenside. Stable propagation of a burst through a one-dimensional homogeneous excitatory chain model of songbird nucleus hvc. *Physical review. E, Statistical, nonlinear, and soft matter physics*, 74 1 Pt 1:011918, 2006. URL `https://api.semanticscholar.org/CorpusID:14083629`.

[100] J Jun and D Jin. Development of Neural Circuitry for Precise Temporal Sequences through Spontaneous Activity, Axon Remodeling, and Synaptic Plasticity. *PLOS ONE*, 2(8):e723, August 2007. ISSN 1932-6203. doi: 10.1371/journal.pone.0000723. URL `https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0000723`. Publisher: Public Library of Science.

[101] IR Fiete, W Senn, CZ Wang, and RH Hahnloser. Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron*, 65:563–76, 2010. doi: 10.1016/j.neuron.2010.02.003.

[102] Wonil Chang and Dezhe Z. Jin. Spike propagation in driven chain networks with dominant global inhibition. *Physical Review E*, 79(5):051917, May 2009. ISSN 1539-

3755, 1550-2376. doi: 10.1103/PhysRevE.79.051917. URL https://link.aps.org/doi/10.1103/PhysRevE.79.051917.

[103] Dezhe Jin. Fast population bursting and its computational significance. *Neural Computation*, 19:201–226, 2007.

[104] P. Zheng and J. Triesch. Robust development of synfire chains from multiple plasticity mechanisms. *Front Comput Neurosci*, 8:66, 2014. ISSN 1662-5188 (Print) 1662-5188. doi: 10.3389/fncom.2014.00066.

[105] Murray JM and Escola GS. Learning multiple variable-speed sequences in striatum via cortical tutoring. *Elife*, 6, 2017. doi: 10.7554/eLife.26084.

[106] Yevhen Tupikov and Dezhe Z Jin. Addition of new neurons and the emergence of a local neural circuit for precise timing. page 43, 2021.

[107] U. Pereira and N. Brunel. Unsupervised learning of persistent and sequential activity. *Front Comput Neurosci*, 13:97, 2019. ISSN 1662-5188 (Print) 1662-5188. doi: 10.3389/fncom.2019.00097.

[108] L. F. Abbott and K. I. Blum. Functional significance of long-term potentiation for sequence learning and prediction. *Cerebral Cortex*, 6:406–416, 1996.

[109] Gina G. Turrigiano. Homeostatic plasticity in neuronal networks. *Nature Reviews Neuroscience*, 5:97–107, 2004.

[110] Oja E. A simplified neuron model as a principal component analyzer. *J Math Biol*, 15 (3):267–73, 1982. doi: 10.1007/BF00275687.

[111] David Sussillo and L. F. Abbott. Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63:544–557, 2009.

[112] Kanaka Rajan, Christopher D. Harvey, and David W. Tank. Recurrent network models of sequence generation and memory. *Neuron*, 90:128–142, 2016.

[113] Dean V. Buonomano. Temporal information processing in the nervous system. *Nature Reviews Neuroscience*, 5:1–12, 2004.

[114] William A. Liberti et al. Unstable neurons underlie a stable learned behavior. *Nature Neuroscience*, 19:1665–1671, 2016.

[115] Jeffrey E. Markowitz et al. Mesoscopic patterns of neural activity support songbird vocal learning. *Nature*, 528:551–556, 2015.

[116] D. Ren, B. Navarro, H. Xu, L. Yue, Q. Shi, and D. E. Clapham. A prokaryotic voltage-gated sodium channel. *Science*, 294(5550):2372–2375, December 2001. ISSN 0036-8075. doi: 10.1126/science.1065635.

[117] Shuyin Sim, Salome Antolin, Chia-Wei Lin, Yingxi Lin, and Carlos Lois. Increased cell-intrinsic excitability induces synaptic changes in new neurons in the adult dentate gyrus that require Npas4. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 33(18):7928–7940, May 2013. ISSN 1529-2401. doi: 10.1523/JNEUROSCI.1571-12.2013.

[118] Chia-Wei Lin, Shuyin Sim, Alice Ainsworth, Masayoshi Okada, Wolfgang Kelsch, and Carlos Lois. Genetically increased cell-intrinsic excitability enhances neuronal integration into adult brain circuits. *Neuron*, 65(1):32–39, January 2010. ISSN 1097-4199. doi: 10.1016/j.neuron.2009.12.001.

[119] Mingshan Xue, Bassam V. Atallah, and Massimo Scanziani. Equalizing excitation-inhibition ratios across visual cortical neurons. *Nature*, 511(7511):596–600, July 2014. ISSN 1476-4687. doi: 10.1038/nature13321.

[120] Robert Egger, Yevhen Tupikov, Margot Elmaleh, Kalman A. Katlowitz, Sam E. Benezra, Michel A. Picardo, Felix Moll, Jörgen Kornfeld, Dezhe Z. Jin, and Michael A. Long. Local Axonal Conduction Shapes the Spatiotemporal Properties of Neural Sequences. *Cell*, 183(2):537–548.e12, October 2020. ISSN 1097-4172. doi: 10.1016/j.cell.2020.09.019.

[121] A. Veliz-Cuba, H. Z. Shouval, K. Josić, and Z. P. Kilpatrick. Networks that learn the precise timing of event sequences. *J Comput Neurosci*, 39(3):235–54, 2015. ISSN 0929-5313. doi: 10.1007/s10827-015-0574-4.

[122] Timothy M. Otchy, Steffen B. E. Wolff, Juliana Y. Rhee, Cengiz Pehlevan, Risa Kawai, Alexandre Kempf, Sharon M. H. Gobes, and Bence P. Ölveczky. Acute off-target effects of neural circuit manipulations. *Nature*, 528(7582):358–363, December 2015. ISSN 1476-4687. doi: 10.1038/nature16442. URL `https://www.nature.com/articles/nature16442`.

[123] Jimok Kim and Richard W. Tsien. Synapse-specific adaptations to inactivity in hippocampal circuits achieve homeostatic gain control while dampening network reverberation. *Neuron*, 58(6):925–937, June 2008. ISSN 0896-6273. doi: 10.1016/j.neuron.2008.05.009. URL `https://pmc.ncbi.nlm.nih.gov/articles/PMC2561251/`.

[124] E. D. Jarvis, C. Scharff, M. R. Grossman, J. A. Ramos, and F. Nottebohm. For whom the bird sings: context-dependent gene expression. *Neuron*, 21(4):775–788, October 1998. ISSN 0896-6273. doi: 10.1016/s0896-6273(00)80594-2.

[125] Hajime Hirase, Xavier Leinekugel, András Czurkó, Jozsef Csicsvari, and György Buzsáki. Firing rates of hippocampal neurons are preserved during subsequent sleep episodes and modified by novel awake experience. *Proceedings of the National Academy of Sciences*, 98(16):9386–9390, July 2001. doi: 10.1073/pnas.161274398. URL `https://www.pnas.org/doi/10.1073/pnas.161274398`.

[126] Edden Slomowitz, Boaz Styr, Irena Vertkin, Hila Milshtein-Parush, Israel Nelken, Michael Slutsky, and Inna Slutsky. Interplay between population firing stability and single neuron dynamics in hippocampal networks. *eLife*, 4:e04378, January 2015. ISSN 2050-084X. doi: 10.7554/eLife.04378. URL `https://doi.org/10.7554/eLife.04378`.

[127] Stéphanie Trouche, Pavel V. Perestenko, Gido M. van de Ven, Claire T. Bratley, Colin G. McNamara, Natalia Campo-Urriza, S. Lucas Black, Leon G. Reijmers, and David Dupret. Recoding a cocaine-place memory engram to a neutral engram in the hippocampus. *Nature Neuroscience*, 19(4):564–567, April 2016. ISSN 1546-1726. doi: 10.1038/nn.4250.

[128] Eric C. Beattie, David Stellwagen, Wade Morishita, Jacqueline C. Bresnahan,

Byeong Keun Ha, Mark Von Zastrow, Michael S. Beattie, and Robert C. Malenka. Control of synaptic strength by glial TNFalpha. *Science*, 295(5563):2282–2285, March 2002. ISSN 1095-9203. doi: 10.1126/science.1067859.

[129] A. F. Schinder and M. Poo. The neurotrophin hypothesis for synaptic plasticity. *Trends in Neurosciences*, 23(12):639–645, December 2000. ISSN 0166-2236. doi: 10.1016/s0166-2236(00)01672-6.

[130] Ye Wang, Wing-Yu Fu, Kit Cheung, Kwok-Wang Hung, Congping Chen, Hongyan Geng, Wing-Ho Yung, Jianan Y. Qu, Amy K. Y. Fu, and Nancy Y. Ip. Astrocyte-secreted IL-33 mediates homeostatic synaptic plasticity in the adult hippocampus. *Proceedings of the National Academy of Sciences of the United States of America*, 118 (1):e2020810118, January 2021. ISSN 1091-6490. doi: 10.1073/pnas.2020810118.

[131] David Stellwagen and Robert C. Malenka. Synaptic scaling mediated by glial tnf-alpha. *Nature*, 440:1054–1059, 2006.

[132] Constance Scharff, John R Kirn, Matthew Grossman, Jeffrey D Macklis, and Fernando Nottebohm. Targeted Neuronal Death Affects Neuronal Replacement and Vocal Behavior in Adult Songbirds. *Neuron*, 25(2):481–492, February 2000. ISSN 0896-6273. doi: 10.1016/S0896-6273(00)80910-1. URL `https://www.sciencedirect.com/science/article/pii/S0896627300809101`.

[133] Richard Mooney and Jonathan F. Prather. The HVC Microcircuit: The Synaptic Basis for Interactions between Song Motor and Vocal Plasticity Pathways. *Journal of Neuroscience*, 25(8):1952–1964, 2005. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.3726-04.2005. URL `https://www.jneurosci.org/content/25/8/1952`. Publisher: Society for Neuroscience _eprint: https://www.jneurosci.org/content/25/8/1952.full.pdf.

[134] Jörgen Kornfeld, Sam E. Benezra, Rajeevan T. Narayanan, Fabian Svara, Robert Egger, Marcel Oberlaender, Winfried Denk, and Michael A. Long. EM connectomics reveals axonal target variation in a sequence-generating network. *eLife*, 6:e24364, March 2017. ISSN 2050-084X. doi: 10.7554/eLife.24364.

[135] Tara C. Thiagarajan, Erika S. Piedras-Renteria, and Richard W. Tsien. alpha- and betaCaMKII. Inverse regulation by neuronal activity and opposing effects on synaptic strength. *Neuron*, 36(6):1103–1114, December 2002. ISSN 0896-6273. doi: 10.1016/s0896-6273(02)01049-8.

[136] Tara C. Thiagarajan, Maria Lindskog, and Richard W. Tsien. Adaptation to synaptic inactivity in hippocampal neurons. *Neuron*, 47(5):725–737, September 2005. ISSN 0896-6273. doi: 10.1016/j.neuron.2005.06.037.

[137] Edward B. Han and Charles F. Stevens. Development regulates a switch between post- and presynaptic strengthening in response to activity deprivation. *Proceedings of the National Academy of Sciences of the United States of America*, 106(26):10817–10822, June 2009. ISSN 1091-6490. doi: 10.1073/pnas.0903603106.

[138] Felix W. Moll, Devorah Kranz, Ariadna Corredera Asensio, Margot Elmaleh, Lyn A. Ackert-Smith, and Michael A. Long. Thalamus drives vocal onsets in the zebra finch courtship song. *Nature*, 616(7955):132–136, April 2023. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-023-05818-x. URL `https://www.nature.com/articles/s41586-023-05818-x`.

[139] Barish Poole, Jeffrey E. Markowitz, and Timothy J. Gardner. The Song Must Go On: Resilience of the Songbird Vocal Motor Pathway. *PLoS ONE*, 7(6):e38173, June 2012. doi: 10.1371/journal.pone.0038173. URL `https://pmc.ncbi.nlm.nih.gov/articles/PMC3387175/`.

[140] Davide Spalla, Isabel Maria Cornacchia, and Alessandro Treves. Continuous attractors for dynamic memories. *eLife*, 10:e69499, September 2021. ISSN 2050-084X. doi: 10.7554/eLife.69499. URL `https://doi.org/10.7554/eLife.69499`. Publisher: eLife Sciences Publications, Ltd.

[141] David Bell, Alison Duffy, and Adrienne Fairhall. Discovering plasticity rules that organize and maintain neural circuits. *bioRxiv: The Preprint Server for Biology*, page 2024.11.18.623688, November 2024. ISSN 2692-8205. doi: 10.1101/2024.11.18.623688.

[142] Nikhil X Bhattasali, Anthony M Zador, and Tatiana A Engel. Neural Circuit Architectural Priors for Embodied Control.

[143] Jane X Wang, Zeb Kurth-Nelson, Dhruva Tirumala, Hubert Soyer, Joel Z Leibo, Remi Munos, Charles Blundell, Dharshan Kumaran, and Matt Botvinick. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*, 2016.

[144] Pastalkova E, Itskov V, Amarasingham A, and Buzsáki G. Internally generated cell assembly sequences in the rat hippocampus. *Science*, 321(5894):1322–7, sep 2008. doi: 10.1126/science.1159775.

[145] Luczak A, Barthó P, Marguet SL, Buzsáki G, and Harris KD. Sequential structure of neocortical spontaneous activity in vivo. *Proc Natl Acad Sci*, 104(1), 2007. doi: 10.1073/pnas.0605643104.

[146] Gustavo B.M Mello, Sofia Soares, and Joseph J Paton. A Scalable Population Code for Time in the Striatum. *Current Biology*, 25(9):1113–1122, may 2015. ISSN 0960-9822. doi: 10.1016/j.cub.2015.02.036. URL `https://doi.org/10.1016/j.cub.2015.02.036`. Publisher: Elsevier.

[147] Felix Weissenberger, Florian Meier, Johannes Lengler, Hafsteinn Einarsson, and Angelika Steger. Long synfire chains emerge by spike-timing dependent plasticity modulated by population activity. *International Journal of Neural Systems*, 27(08): 1750044, 2017. doi: 10.1142/S0129065717500447. URL `https://doi.org/10.1142/S0129065717500447`. PMID: 28982282.

[148] Samy Bengio, Yoshua Bengio, Jocelyn Cloutier, and Jan Gecsei. On the optimization of a synaptic learning rule. In *Optimality in Biological and Artificial Networks?*, pages 265–287. Routledge, 2013.

[149] Hector Garcia Rodriguez, Qinghai Guo, and Timoleon Moraitis. Short-Term Plasticity Neurons Learning to Learn and Forget, June 2022. URL `http://arxiv.org/abs/2206.14048`. arXiv:2206.14048.

[150] F. W. Moll, D. Kranz, A. Corredera Asensio, et al. Thalamus drives vocal onsets in the zebra finch courtship song. *Nature*, 616:132–136, 2023. doi: 10.1038/s41586-023-05818-x. URL https://doi.org/10.1038/s41586-023-05818-x.

[151] H. H. Danish, D. Aronov, and M. S. Fee. Rhythmic syllable-related activity in a songbird motor thalamic nucleus necessary for learned vocalizations. *PLoS ONE*, 12 (6):e0169568, 2017. doi: 10.1371/journal.pone.0169568.

[152] Anne Auger and Nikolaus Hansen. A restart cma evolution strategy with increasing population size. volume 2, pages 1769–1776, 01 2005. doi: 10.1109/CEC.2005.1554902.

[153] Manuel A. Castro-Alamancos and Barry W. Connors. Distinct forms of short-term plasticity at excitatory synapses of hippocampus and neocortex. *Proceedings of the National Academy of Sciences*, 94(8):4161–4166, April 1997. doi: 10.1073/pnas.94.8.4161. URL https://www.pnas.org/doi/10.1073/pnas.94.8.4161. Publisher: Proceedings of the National Academy of Sciences.

[154] Rylan S. Larsen and P. Jesper Sjöström. Synapse-type-specific plasticity in local circuits. *Current opinion in neurobiology*, 35:127, August 2015. doi: 10.1016/j.conb.2015.08.001. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC5280068/.

[155] 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(6062):1569–73, 2011. ISSN 0036-8075. doi: 10.1126/science.1211095.

[156] Ashok Litwin-Kumar and Brent Doiron. Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nature communications*, 5(1):5319, 2014.

[157] Yotam Luz and Maoz Shamir. Balancing Feed-Forward Excitation and Inhibition via Hebbian Inhibitory Synaptic Plasticity. *PLOS Computational Biology*, 8(1):1–12, January 2012. doi: 10.1371/journal.pcbi.1002334. URL https://doi.org/10.1371/journal.pcbi.1002334. Publisher: Public Library of Science.

[158] Waitzmann F, Wu YK, and Gjorgjieva J. Top-down modulation in canonical cortical

circuits with short-term plasticity. *Proc Natl Acad Sci*, 121(16), 2024. doi: 10.1073/pnas.2311040121.

[159] Ziyi Gong and Nicolas Brunel. Inhibitory Plasticity Enhances Sequence Storage Capacity and Retrieval Robustness, April 2024. URL `https://www.biorxiv.org/content/10.1101/2024.04.08.588573v1`. Pages: 2024.04.08.588573 Section: New Results.

[160] Fereshteh Lagzi and Adrienne L. Fairhall. Emergence of co-tuning in inhibitory neurons as a network phenomenon mediated by randomness, correlations, and homeostatic plasticity. *Science Advances*, 10(12):eadi4350, March 2024. ISSN 2375-2548. doi: 10.1126/sciadv.adi4350. URL `https://www.science.org/doi/10.1126/sciadv.adi4350`.

[161] Shuyin Sim, Salome Antolin, Chia-Wei Lin, Yingxi Lin, and Carlos Lois. Increased cell-intrinsic excitability induces synaptic changes in new neurons in the adult dentate gyrus that require npas4. *Journal of Neuroscience*, 33(18):7928–7940, 2013. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.1571-12.2013. URL `https://www.jneurosci.org/content/33/18/7928`.

[162] C. W. Lin, S. Sim, A. Ainsworth, M. Okada, W. Kelsch, and C. Lois. Genetically increased cell-intrinsic excitability enhances neuronal integration into adult brain circuits. *Neuron*, 65(1):32–39, jan 2010. doi: 10.1016/j.neuron.2009.12.001.

[163] Ashok Litwin-Kumar, Kameron Decker Harris, Richard Axel, Haim Sompolinsky, and L. F. Abbott. Optimal Degrees of Synaptic Connectivity. *Neuron*, 93(5):1153–1164.e7, March 2017. ISSN 1097-4199. doi: 10.1016/j.neuron.2017.01.030.

[164] Kaushik J. Lakshminarasimhan, Marjorie Xie, Jeremy D. Cohen, Britton A. Sauerbrei, Adam W. Hantman, Ashok Litwin-Kumar, and Sean Escola. Specific connectivity optimizes learning in thalamocortical loops. *Cell Reports*, 43(4):114059, April 2024. ISSN 2211-1247. doi: 10.1016/j.celrep.2024.114059.

[165] Armin Bahl and Florian Engert. Neural circuits for evidence accumulation and decision

making in larval zebrafish. *Nature Neuroscience*, 23(1):94–102, 2020. ISSN 1546-1726. doi: 10.1038/s41593-019-0534-9.

[166] Dániel L. Barabási, Gregor F. P. Schuhknecht, and Florian Engert. Functional neuronal circuits emerge in the absence of developmental activity. *Nature Communications*, 15 (1):364, January 2024. ISSN 2041-1723. doi: 10.1038/s41467-023-44681-2. URL https://www.nature.com/articles/s41467-023-44681-2. Publisher: Nature Publishing Group.

[167] M. N. Shadlen and W. T. Newsome. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86(4): 1916–1936, October 2001. ISSN 0022-3077. doi: 10.1152/jn.2001.86.4.1916.

[168] Torkel Hafting, Marianne Fyhn, Sturla Molden, May-Britt Moser, and Edvard I. Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806, August 2005. ISSN 1476-4687. doi: 10.1038/nature03721.

[169] Yoram Burak and Ila R. Fiete. Accurate Path Integration in Continuous Attractor Network Models of Grid Cells. *PLOS Computational Biology*, 5(2):e1000291, February 2009. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1000291. URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1000291.

[170] Shun-ichi Amari. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27(2):77–87, June 1977. ISSN 1432-0770. doi: 10.1007/BF00337259. URL https://doi.org/10.1007/BF00337259.

[171] William Skaggs, James Knierim, Hemant Kudrimoti, and Bruce McNaughton. A Model of the Neural Basis of the Rat' s Sense of Direction. In *Advances in Neural Information Processing Systems*, volume 7. MIT Press, 1994. URL https://proceedings.neurips.cc/paper_files/paper/1994/hash/024d7f84fff11dd7e8d9c510137a2381-Abstract.html.

[172] R Ben-Yishai, R L Bar-Or, and H Sompolinsky. Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences*, 92(9):3844–3848,

April 1995. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.92.9.3844. URL https://pnas.org/doi/full/10.1073/pnas.92.9.3844.

[173] H. S. Seung. How the brain keeps the eyes still. *Proceedings of the National Academy of Sciences of the United States of America*, 93(23):13339–13344, November 1996. ISSN 0027-8424. doi: 10.1073/pnas.93.23.13339.

[174] K. Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *Journal of Neuroscience*, 16(6):2112–2126, March 1996. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.16-06-02112.1996. URL https://www.jneurosci.org/content/16/6/2112.

[175] A. Compte, N. Brunel, P. S. Goldman-Rakic, and X. J. Wang. Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex (New York, N.Y.: 1991)*, 10(9):910–923, September 2000. ISSN 1047-3211. doi: 10.1093/cercor/10.9.910.

[176] Ábel Ságodi, Guillermo Martín-Sánchez, Piotr Sokół, and Il Memming Park. Back to the Continuous Attractor.

[177] David G. Clark, L. F. Abbott, and Haim Sompolinsky. Symmetries and Continuous Attractors in Disordered Neural Circuits, January 2025. URL https://www.biorxiv.org/content/10.1101/2025.01.26.634933v1. Pages: 2025.01.26.634933 Section: New Results.

[178] Ran Darshan and Alexander Rivkind. Learning to represent continuous variables in heterogeneous neural networks. page 51, 2022.

[179] R. H. R. Hahnloser. Emergence of neural integration in the head-direction system by visual supervision. *Neuroscience*, 120(3):877–891, 2003. ISSN 0306-4522. doi: 10.1016/s0306-4522(03)00201-x.

[180] Pantelis Vafidis, David Owald, Tiziano D'Albis, and Richard Kempter. Learning accurate path integration in a ring attractor model of the head direction system.

Technical report, bioRxiv, March 2021. URL https://www.biorxiv.org/content/10.1101/2021.03.12.435035v1. Section: New Results Type: article.

[181] Marcella Noorman, Brad K. Hulse, Vivek Jayaraman, Sandro Romani, and Ann M. Hermundstad. Maintaining and updating accurate internal representations of continuous variables with a handful of neurons. *Nature Neuroscience*, 27(11):2207–2217, November 2024. ISSN 1546-1726. doi: 10.1038/s41593-024-01766-5. URL https://www.nature.com/articles/s41593-024-01766-5. Publisher: Nature Publishing Group.

[182] Valerio Mante, David Sussillo, Krishna V. Shenoy, and William T. Newsome. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503 (7474):78–84, November 2013. ISSN 1476-4687. doi: 10.1038/nature12742. URL https://www.nature.com/articles/nature12742.

[183] M. S. Goldman, A. Compte, and X. J. Wang. Neural Integrator Models. In Larry R. Squire, editor, *Encyclopedia of Neuroscience*, pages 165–178. Academic Press, Oxford, January 2009. ISBN 978-0-08-045046-9. URL https://www.sciencedirect.com/science/article/pii/B9780080450469014340.

[184] S. M. Stringer, E. T. Rolls, and T. P Trappenberg. Self-organizing continuous attractor network models of hippocampal spatial view cells. *Neurobiology of Learning and Memory*, 83(1):79–92, January 2005. ISSN 1074-7427. doi: 10.1016/j.nlm.2004.08.003. URL https://www.sciencedirect.com/science/article/pii/S1074742704000978.

[185] Alexander Seeholzer, Moritz Deger, and Wulfram Gerstner. Stability of working memory in continuous attractor networks under the control of short-term plasticity. *PLOS Computational Biology*, 15(4):e1006928, April 2019. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1006928. URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1006928. Publisher: Public Library of Science.

[186] Sukbin Lim and Mark S. Goldman. Balanced cortical microcircuitry for maintaining information in working memory. *Nature Neuroscience*, 16(9):1306–1314, September 2013.

ISSN 1546-1726. doi: 10.1038/nn.3492. URL https://www.nature.com/articles/nn.3492. Publisher: Nature Publishing Group.

[187] Sukbin Lim and Mark S. Goldman. Balanced cortical microcircuitry for spatial working memory based on corrective feedback control. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 34(20):6790–6806, May 2014. ISSN 1529-2401. doi: 10.1523/JNEUROSCI.4602-13.2014.

[188] Alfonso Renart, Pengcheng Song, and Xiao-Jing Wang. Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks. *Neuron*, 38 (3):473–485, May 2003. ISSN 0896-6273. doi: 10.1016/s0896-6273(03)00255-1.

[189] Xiaohui Xie and H. Sebastian Seung. Spike-based Learning Rules and Stabilization of Persistent Neural Activity. In *Advances in Neural Information Processing Systems*, volume 12. MIT Press, 1999. URL https://proceedings.neurips.cc/paper_files/paper/1999/hash/c0560792e4a3c79e62f76cbf9fb277dd-Abstract.html.

[190] Jintao Gu and Sukbin Lim. Unsupervised learning for robust working memory. *PLOS Computational Biology*, 18(5):e1009083, May 2022. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1009083. URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1009083. Publisher: Public Library of Science.

[191] Basile Confavreux, Zoe P M Harrington, Maciej Kania, Poornima Ramesh, Anastasia N Krouglova, Panos A Bozelos, Jakob H Macke, Andrew M Saxe, Pedro J Goncalves, and Tim P Vogels. Memory by a thousand rules: Automated discovery of multi-type plasticity rules reveals variety & degeneracy at the heart of learning. .

[192] Lang Wang and Arianna Maffei. Inhibitory Plasticity Dictates the Sign of Plasticity at Excitatory Synapses. *The Journal of Neuroscience*, 34(4):1083–1093, January 2014. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.4711-13.2014. URL https://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.4711-13.2014.

[193] Jonathan Mapelli, Daniela Gandolfi, Antonietta Vilella, Michele Zoli, and Albertino Bigiani. Heterosynaptic GABAergic plasticity bidirectionally driven by the activity

of pre- and postsynaptic NMDA receptors. *Proceedings of the National Academy of Sciences*, 113(35):9898–9903, August 2016. doi: 10.1073/pnas.1601194113. URL https://www.pnas.org/doi/abs/10.1073/pnas.1601194113.

[194] 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, May 2024. ISSN 1546-1726. doi: 10.1038/s41593-024-01597-4. URL https://www.nature.com/articles/s41593-024-01597-4.

[195] Philipp Schlegel, Yijie Yin, Alexander S. Bates, Sven Dorkenwald, Katharina Eichler, Paul Brooks, Daniel S. Han, Marina Gkantia, Marcia dos Santos, Eva J. Munnelly, Griffin Badalamente, Laia Serratosa Capdevila, Varun A. Sane, Alexandra M. C. Fragniere, Ladann Kiassat, Markus W. Pleijzier, Tomke Stürner, Imaan F. M. Tamimi, Christopher R. Dunne, Irene Salgarella, Alexandre Javier, Siqi Fang, Eric Perlman, Tom Kazimiers, Sridhar R. Jagannathan, Arie Matsliah, Amy R. Sterling, Szi-chieh Yu, Claire E. McKellar, Marta Costa, H. Sebastian Seung, Mala Murthy, Volker Hartenstein, Davi D. Bock, and Gregory S. X. E. Jefferis. Whole-brain annotation and multi-connectome cell typing of Drosophila. *Nature*, 634(8032):139–152, October 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07686-5. URL https://www.nature.com/articles/s41586-024-07686-5.

[196] Basile Confavreux, Everton J. Agnes, Friedemann Zenke, Henning Sprekeler, and Tim P. Vogels. Balancing complexity, performance and plausibility to meta learn plasticity rules in recurrent spiking networks. *PLOS Computational Biology*, 21(4):e1012910, April 2025. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1012910. URL https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1012910.