

EMERGENT NEURAL COMPUTATION FROM THE INTERACTION OF DIFFERENT FORMS OF PLASTICITY

EDITED BY: Cristina Savin, Matthieu Gilson and Friedemann Zenke
PUBLISHED IN: Frontiers in Computational Neuroscience



frontiers Research Topics



frontiers

Frontiers Copyright Statement

© Copyright 2007-2016 Frontiers Media SA. All rights reserved.

All content included on this site, such as text, graphics, logos, button icons, images, video/audio clips, downloads, data compilations and software, is the property of or is licensed to Frontiers Media SA ("Frontiers") or its licensees and/or subcontractors. The copyright in the text of individual articles is the property of their respective authors, subject to a license granted to Frontiers.

The compilation of articles constituting this e-book, wherever published, as well as the compilation of all other content on this site, is the exclusive property of Frontiers. For the conditions for downloading and copying of e-books from Frontiers' website, please see the Terms for Website Use. If purchasing Frontiers e-books from other websites or sources, the conditions of the website concerned apply.

Images and graphics not forming part of user-contributed materials may not be downloaded or copied without permission.

Individual articles may be downloaded and reproduced in accordance with the principles of the CC-BY licence subject to any copyright or other notices. They may not be re-sold as an e-book.

As author or other contributor you grant a CC-BY licence to others to reproduce your articles, including any graphics and third-party materials supplied by you, in accordance with the Conditions for Website Use and subject to any copyright notices which you include in connection with your articles and materials.

All copyright, and all rights therein, are protected by national and international copyright laws.

The above represents a summary only. For the full conditions see the Conditions for Authors and the Conditions for Website Use.

ISSN 1664-8714

ISBN 978-2-88919-788-0

DOI 10.3389/978-2-88919-788-0

About Frontiers

Frontiers is more than just an open-access publisher of scholarly articles: it is a pioneering approach to the world of academia, radically improving the way scholarly research is managed. The grand vision of Frontiers is a world where all people have an equal opportunity to seek, share and generate knowledge. Frontiers provides immediate and permanent online open access to all its publications, but this alone is not enough to realize our grand goals.

Frontiers Journal Series

The Frontiers Journal Series is a multi-tier and interdisciplinary set of open-access, online journals, promising a paradigm shift from the current review, selection and dissemination processes in academic publishing. All Frontiers journals are driven by researchers for researchers; therefore, they constitute a service to the scholarly community. At the same time, the Frontiers Journal Series operates on a revolutionary invention, the tiered publishing system, initially addressing specific communities of scholars, and gradually climbing up to broader public understanding, thus serving the interests of the lay society, too.

Dedication to Quality

Each Frontiers article is a landmark of the highest quality, thanks to genuinely collaborative interactions between authors and review editors, who include some of the world's best academicians. Research must be certified by peers before entering a stream of knowledge that may eventually reach the public - and shape society; therefore, Frontiers only applies the most rigorous and unbiased reviews.

Frontiers revolutionizes research publishing by freely delivering the most outstanding research, evaluated with no bias from both the academic and social point of view.

By applying the most advanced information technologies, Frontiers is catapulting scholarly publishing into a new generation.

What are Frontiers Research Topics?

Frontiers Research Topics are very popular trademarks of the Frontiers Journals Series: they are collections of at least ten articles, all centered on a particular subject. With their unique mix of varied contributions from Original Research to Review Articles, Frontiers Research Topics unify the most influential researchers, the latest key findings and historical advances in a hot research area! Find out more on how to host your own Frontiers Research Topic or contribute to one as an author by contacting the Frontiers Editorial Office: researchtopics@frontiersin.org

EMERGENT NEURAL COMPUTATION FROM THE INTERACTION OF DIFFERENT FORMS OF PLASTICITY

Topic Editors:

Cristina Savin, IST Austria, Austria

Matthieu Gilson, Universitat Pompeu Fabra, Spain

Friedemann Zenke, Stanford University, USA

From the propagation of neural activity through synapses, to the integration of signals in the dendritic arbor, and the processes determining action potential generation, virtually all aspects of neural processing are plastic. This plasticity underlies the remarkable versatility and robustness of cortical circuits: it enables the brain to learn regularities in its sensory inputs, to remember the past, and to recover function after injury.

While much of the research into learning and memory has focused on forms of Hebbian plasticity at excitatory synapses (LTD/LTP, STDP), several other plasticity mechanisms have been characterized experimentally, including the plasticity of inhibitory circuits (Kullmann, 2012), synaptic scaling (Turrigiano, 2011) and intrinsic plasticity (Zhang and Linden, 2003). However, our current understanding of the computational roles of these plasticity mechanisms remains rudimentary at best. While traditionally they are assumed to serve a homeostatic purpose, counterbalancing the destabilizing effects of Hebbian learning, recent work suggests that they can have a profound impact on circuit function (Savin 2010, Vogels 2011, Keck 2012). Hence, theoretical investigation into the functional implications of these mechanisms may shed new light on the computational principles at work in neural circuits.

This Research Topic of Frontiers in Computational Neuroscience aims to bring together recent advances in theoretical modeling of different plasticity mechanisms and of their contributions to circuit function. Topics of interest include the computational roles of plasticity of inhibitory circuitry, metaplasticity, synaptic scaling, intrinsic plasticity, plasticity within the dendritic arbor and in particular studies on the interplay between homeostatic and Hebbian plasticity, and their joint contribution to network function.

Citation: Savin, C., Gilson, M., Zenke, F., eds. (2016). Emergent Neural Computation from the Interaction of Different Forms of Plasticity. Lausanne: Frontiers Media.
doi: 10.3389/978-2-88919-788-0

Table of Contents

- 04 Editorial: Emergent Neural Computation from the Interaction of Different Forms of Plasticity**
Matthieu Gilson, Cristina Savin and Friedemann Zenke
- 06 Models of Metaplasticity: A Review of Concepts**
Pierre Yger and Matthieu Gilson
- 20 Addendum: Models of Metaplasticity: A Review of Concepts**
Pierre Yger and Matthieu Gilson
- 21 Homeostatic role of heterosynaptic plasticity: models and experiments**
Marina Chistiakova, Nicholas M. Bannon, Jen-Yung Chen, Maxim Bazhenov and Maxim Volgushev
- 43 Enhanced polychronization in a spiking network with metaplasticity**
Mira Guise, Alistair Knott and Lubica Benuskova
- 55 Learning structure of sensory inputs with synaptic plasticity leads to interference**
Joseph Chrol-Cannon and Yaochu Jin
- 68 Excitatory and inhibitory STDP jointly tune feedforward neural circuits to selectively propagate correlated spiking activity**
Florence I. Kleberg, Tomoki Fukai and Matthieu Gilson
- 83 Robust development of synfire chains from multiple plasticity mechanisms**
Pengsheng Zheng and Jochen Triesch
- 93 Dynamic stability of sequential stimulus representations in adapting neuronal networks**
Renato C. F. Duarte and Abigail Morrison
- 113 Unsupervised discrimination of patterns in spiking neural networks with excitatory and inhibitory synaptic plasticity**
Narayan Srinivasa and Youngkwan Cho
- 136 Learning and stabilization of winner-take-all dynamics through interacting excitatory and inhibitory plasticity**
Jonathan Binas, Ueli Rutishauser, Giacomo Indiveri and Michael Pfeiffer
- 149 Adaptation of short-term plasticity parameters via error-driven learning may explain the correlation between activity-dependent synaptic properties, connectivity motifs and target specificity**
Umberto Esposito, Michele Giugliano and Eleni Vasilaki
- 167 RM-SORN: a reward-modulated self-organizing recurrent neural network**
Witali Aswolinskiy and Gordon Pipa
- 182 Emergence of task-dependent representations in working memory circuits**
Cristina Savin and Jochen Triesch



Editorial: Emergent Neural Computation from the Interaction of Different Forms of Plasticity

Matthieu Gilson^{1†}, Cristina Savin^{2*†} and Friedemann Zenke^{3†}

¹ Computational Neuroscience Group, Department of Technology and Information of Communication, Universitat Pompeu Fabra, Barcelona, Spain, ² Institute of Science and Technology Austria, Klosterneuburg, Austria, ³ Neural Dynamics and Computation Lab, Department of Applied Physics, Stanford University, Stanford, CA, USA

Keywords: homeostatic plasticity, inhibitory plasticity, learning, neural network, neuromodulation

More than 60 years later, Hebb's prophecy "neurons that fire together wire together" (Hebb, 1949; Shatz, 1992) prevails as one of the cornerstones of modern neuroscience. Nonetheless, it is becoming increasingly evident that there is more to neural plasticity than the strengthening of synapses between co-active neurons. Experiments have revealed a plethora of synaptic and cellular plasticity mechanisms acting simultaneously in neural circuits. How such diverse forms of plasticity collectively give rise to neural computation remains poorly understood. The present Research Topic approaches this question by bringing together recent advances in the modeling of different forms of synaptic and neuronal plasticity. Taken together, these studies argue that the concerted interaction of diverse forms of plasticity is critical for circuit formation and function.

A first insight from this Research Topic underscores the importance of the time scale of homeostatic plasticity to avoid runaway dynamics of Hebbian plasticity. While known homeostatic processes act slowly, on the timescale of hours to days, existing theoretical models invariably use *fast homeostasis*. Yger and Gilson (2015) review a body of theoretical work arguing that rapid forms of homeostatic control are in fact critical for stable learning and thus should also exist in biological circuits. Following a similar line of thought, Chistiakova et al. (2015) review experimental and theoretical literature which suggests that the role of rapid homeostasis could be filled by heterosynaptic plasticity. Alternatively, other mechanisms can achieve a similar stabilizing effect, as long as they are fast, for instance the rapid homeostatic sliding threshold in Guise et al. (2015). These findings raise questions concerning the purpose of slow homeostasis and metaplasticity. Since non-modulated plasticity leads to "interference" between memories when confronted with rich environmental stimuli (Chrol-Cannon and Jin, 2015), it is tempting to hypothesize that certain slow homeostatic mechanisms may correct for this (Yger and Gilson, 2015).

The second development reflected in this Research Topic concerns the interactions between *excitatory and inhibitory* (E/I) plasticity. Multiple studies independently stress the importance of such interactions for shaping circuit selectivity and decorrelating network activity during learning. Kleberg et al. (2014) demonstrate how spike-timing-dependent plasticity at excitatory (eSTDP) and inhibitory (iSTDP) synapses drives the formation of selective signaling pathways in feed-forward networks. Together they ensure excitatory-inhibitory balance and sharpen neuronal responses to salient inputs. Moreover, by systematically exploring different iSTDP windows, the authors show that anti-symmetric plasticity, in which pre-post spike pairs lead to potentiation of an inhibitory synapse, are most efficient at establishing pathway-specific balance. Zheng and Triesch (2014) confirm the relevance of e/iSTDP for propagating information in a recurrent network. Their model also highlights the importance of other forms of plasticity, in particular *intrinsic plasticity* and *structural plasticity* for robust synfire-chain learning.

Beyond information propagation, Duarte and Morrison (2014) show that E/I plasticity allows recurrent neural networks to form internal representations of the external world and to perform non-linear computations with them. They find that the decorrelating action

OPEN ACCESS

Edited and reviewed by:

Si Wu,
Beijing Normal University, China

*Correspondence:

Cristina Savin
cristina.savin@ist.ac.at

[†]These authors have contributed
equally to this work.

Received: 02 November 2015

Accepted: 13 November 2015

Published: 30 November 2015

Citation:

Gilson M, Savin C and Zenke F (2015)
Editorial: Emergent Neural
Computation from the Interaction of
Different Forms of Plasticity.
Front. Comput. Neurosci. 9:145.
doi: 10.3389/fncom.2015.00145

of inhibitory plasticity pushes the network away from states with poor discriminability. These results are corroborated by Srinivasa and Cho (2014), who show that such representations can be efficiently picked up by downstream layers. Networks shaped by both e- and iSTDP learn to discriminate between neural activity patterns in a self-organized fashion, whereas networks with only one form of plasticity perform worse. Binas et al. (2014) show that the interplay of E/I plasticity in recurrent neural networks can form robust winner-take-all (WTA) circuits, important for solving a range of behaviorally relevant tasks (e.g., categorization or decision making). Using a novel mean-field theory for network dynamics and plasticity, the authors characterize parameter regions in which stable WTA circuits emerge autonomously through the interaction of E/I plasticity.

While most work presented here focuses on long-term plasticity, Esposito et al. (2015), study the interactions between Hebbian and *short-term plasticity* (STP) at excitatory synapses. The authors postulate a form of *metaplasticity* that adjusts the properties of STP to minimize circuit error. This model provides a normative interpretation for experimentally observed variability in STP properties across neural circuits and its close link to network connectivity motifs. While detailed error computation as assumed here is biologically implausible, reward-related information could be provided by neuromodulators (in particular, dopamine), which are known to regulate circuit dynamics and plasticity.

The functional importance of *neuromodulation* is explored in two papers. First, Aswolinskiy and Pipa (2015) systematically compare reward-dependent vs. supervised and unsupervised learning across a broad range of tasks. They find that, when

combined with suitable homeostatic plasticity mechanisms, reward-dependent synaptic plasticity can yield a performance similar to abstract supervised learning. Second, Savin and Triesch (2014) use a similar circuit model to study how reward-dependent learning shapes random recurrent networks into working memory circuits. They show that the interaction between dopamine-modulated STDP and homeostatic plasticity is sufficient to explain a broad range of experimental findings regarding the coding properties of neurons in prefrontal circuits. More generally, these results enforce the idea that reward-dependent learning is critical for shifting the limited neural resources toward the computations that matter most in terms of behavioral outcomes.

Taken together, the contributions to this Research Topic suggest that circuit-level function emerges from the complex, but well-orchestrated interplay of different forms of neural plasticity. To learn how neuronal circuits self-organize and how computation emerges in the brain it is therefore vital to focus on interacting forms of plasticity. This sets the scene for exciting future research in both theoretical and experimental neuroscience.

FUNDING

CS acknowledges funding from the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme (FP7/2007-2013) under REA grant agreement n° [291734]. MG acknowledges funding from the Human Brain Project (PR10513 :EC-7PM-HBP). FZ was supported by the SNSF (Swiss National Science Foundation) Mobility fellowship P2ELP3_161836.

REFERENCES

- Aswolinskiy, W., and Pipa, G. (2015). RM-SORN: a reward-modulated self-organizing recurrent neural network. *Front. Comput. Neurosci.* 9:36. doi: 10.3389/fncom.2015.00036
- Binas, J., Rutishauser, U., Indiveri, G., and Pfeiffer, M. (2014). Learning and stabilization of winner-take-all dynamics through interacting excitatory and inhibitory plasticity. *Front. Comput. Neurosci.* 8:68. doi: 10.3389/fncom.2014.00068
- Chistiakova, M., Bannon, N. M., Chen, J.-Y., Bazhenov, M., and Volgushev, M. (2015). Homeostatic role of heterosynaptic plasticity: models and experiments. *Front. Comput. Neurosci.* 9:89. doi: 10.3389/fncom.2015.00089
- Chrol-Cannon, J., and Jin, Y. (2015). Learning structure of sensory inputs with synaptic plasticity leads to interference. *Front. Comput. Neurosci.* 9:103. doi: 10.3389/fncom.2015.00103
- Duarte, R. C. F., and Morrison, A. (2014). Dynamic stability of sequential stimulus representations in adapting neuronal networks. *Front. Comput. Neurosci.* 8:124. doi: 10.3389/fncom.2014.00124
- Esposito, U., Giugliano, M., and Vasilaki, E. (2015). Adaptation of short-term plasticity parameters via error-driven learning may explain the correlation between activity-dependent synaptic properties, connectivity motifs and target specificity. *Front. Comput. Neurosci.* 8:175. doi: 10.3389/fncom.2014.00175
- Guise, M., Knott, A., and Benuskova, L. (2015). Enhanced polychronization in a spiking network with metaplasticity. *Front. Comput. Neurosci.* 9:9. doi: 10.3389/fncom.2015.00009
- Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory*. New York, NY: Wiley & Sons.
- Kleberg, F. I., Fukai, T., and Gilson, M. (2014). Excitatory and inhibitory STDP jointly tune feedforward neural circuits to selectively propagate correlated spiking activity. *Front. Comput. Neurosci.* 8:53. doi: 10.3389/fncom.2014.00053
- Savin, C., and Triesch, J. (2014). Emergence of task-dependent representations in working memory circuits. *Front. Comput. Neurosci.* 8:57. doi: 10.3389/fncom.2014.00057
- Shatz, C. J. (1992). The developing brain. *Sci. Am.* 267, 60–67.
- Srinivasa, N., and Cho, Y. (2014). Unsupervised discrimination of patterns in spiking neural networks with excitatory and inhibitory synaptic plasticity. *Front. Comput. Neurosci.* 8:159. doi: 10.3389/fncom.2014.00159
- Yger, P., and Gilson, M. (2015). Models of metaplasticity: a review of concepts. *Front. Comput. Neurosci.* 9:138. doi: 10.3389/fncom.2015.00138
- Zheng, P., and Triesch, J. (2014). Robust development of synfire chains from multiple plasticity mechanisms. *Front. Comput. Neurosci.* 8:66. doi: 10.3389/fncom.2014.00066

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2015 Gilson, Savin and Zenke. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.



Models of Metaplasticity: A Review of Concepts

Pierre Yger^{1,2*} and Matthieu Gilson³

¹ Sorbonne Université, UPMC Univ Paris06 UMR5968, Paris, France, ² Institut de la Vision, INSERM, U968, Centre National de la Recherche Scientifique, UMR7210, Paris, France, ³ Computational Neurosciences Group, Departament de Tecnologies de la Informació i les Comunicacions, Universitat Pompeu Fabra, Barcelona, Spain

Part of hippocampal and cortical plasticity is characterized by synaptic modifications that depend on the joint activity of the pre- and post-synaptic neurons. To which extent those changes are determined by the exact timing and the average firing rates is still a matter of debate; this may vary from brain area to brain area, as well as across neuron types. However, it has been robustly observed both *in vitro* and *in vivo* that plasticity itself slowly adapts as a function of the dynamical context, a phenomena commonly referred to as metaplasticity. An alternative concept considers the regulation of groups of synapses with an objective at the neuronal level, for example, maintaining a given average firing rate. In that case, the change in the strength of a particular synapse of the group (e.g., due to Hebbian learning) affects others' strengths, which has been coined as heterosynaptic plasticity. Classically, Hebbian synaptic plasticity is paired in neuron network models with such mechanisms in order to stabilize the activity and/or the weight structure. Here, we present an oriented review that brings together various concepts from heterosynaptic plasticity to metaplasticity, and show how they interact with Hebbian-type learning. We focus on approaches that are nowadays used to incorporate those mechanisms to state-of-the-art models of spiking plasticity inspired by experimental observations in the hippocampus and cortex. Making the point that metaplasticity is an ubiquitous mechanism acting on top of classical Hebbian learning and promoting the stability of neural function over multiple timescales, we stress the need for incorporating it as a key element in the framework of plasticity models. Bridging theoretical and experimental results suggests a more functional role for metaplasticity mechanisms than simply stabilizing neural activity.

Keywords: synaptic plasticity, metaplasticity, Hebbian learning, homeostasis, STDP

OPEN ACCESS

Edited by:

Friedemann Zenke,
Stanford University, USA

Reviewed by:

Harel Z. Shouval,
University of Texas Medical School at
Houston, USA
Christian Tetzlaff,
Max Planck Institute for Dynamics and
Self-Organization, Germany

*Correspondence:

Pierre Yger
pierre.yger@inserm.fr

Received: 02 July 2015

Accepted: 27 October 2015

Published: 10 November 2015

Citation:

Yger P and Gilson M (2015) Models of
Metaplasticity: A Review of Concepts.
Front. Comput. Neurosci. 9:138.
doi: 10.3389/fncom.2015.00138

1. INTRODUCTION

The brain is made of billions of neurons able to efficiently process the huge flow of information impinging continuously on sensory modalities, extracting relevant data, and producing appropriately timed responses. Even during development (Corlew et al., 2007; Wang et al., 2012) or when lesioned (Young et al., 2007; Beck and Yaari, 2008), the brain has the striking capability to adapt in order to maintain the stability of neural functions. Importantly, this slow adaptation, acting at a timescale of hours or days (Turrigiano and Nelson, 2000; Davis, 2006) is performed in conjunction with fast changes often observed in the so called Hebbian learning (Hebb, 1949). Understanding the mechanisms leading to the dynamical organization of neuronal network via the fine interactions of those two competing processes is therefore a crucial step toward analyzing the stability of the computations performed by cerebral activity.

Following the seminal idea that neurons firing together should wire together (Hebb, 1949), numerous experimental studies have been conducted to unravel part of the links between plasticity and neuronal activity. Nowadays, this so-called Hebbian form of plasticity in the brain has been characterized experimentally in many areas, involving multiple but still misunderstood molecular pathways (see Abbott and Nelson, 2000; Caporale and Dan, 2008, for reviews). While it is commonly assumed that NMDA receptors are the primary actors in long-term potentiation, or LTP (Feldman, 2012), the biochemical pathways for long-term depression (LTD) seem to differ in cortex and in hippocampus (Wang et al., 2005; Bender et al., 2006; Nevian and Sakmann, 2006). In controlled *in vitro* experiments, it has also been shown that LTP and LTD depend on the precise timing of pre- and post-synaptic spikes (Markram et al., 1997; Bi and Poo, 1998), leading to the concept of timing-LTP/LTD or spike-timing-dependent plasticity (STDP).

By acting independently at each synapse without spatial or temporal crosstalk among synapses, Hebbian learning is a form of homosynaptic plasticity that is intrinsically unstable. In point of fact, provided synapses are reinforced when both the pre- and post-synaptic neurons are active, nothing prevents the synapses from strengthening themselves boundlessly, which causes the post-synaptic activity to explode (Rochester et al., 1956; von der Malsburg, 1973; Miller, 1996). While this instability can be avoided by artificially imposing hard boundaries onto the synaptic weights, several learning models came with intrinsic mechanisms regulating the synaptic efficacies (Bienenstock et al., 1982; Oja, 1982) in order to solve this issue in a less fine-tuned manner.

The present paper reviews such mechanisms that aim to tame the positive feedback provided by Hebbian plasticity. In the biology, some homeostatic mechanisms can be viewed as independent from the Hebbian learning that they counterbalance. For example, the sum of synaptic strengths may be up or down regulated to maintain the average post-synaptic firing rate; see Vitteira and Goda (2013) for a review of the biophysics of such mechanisms. In contrast, other processes directly modulate the learning rule itself as a function of the dynamical context, which is referred to as metaplasticity. This concept is the plasticity of the synaptic plasticity itself (Abraham and Bear, 1996; Abraham, 2008), and it is tightly related to the notion of homeostasis (O'Leary and Wyllie, 2011). To ensure the overall stability of the neuronal system, a key role for metaplasticity is to regulate the synaptic update rules in terms of the past history of the activity at the whole neuronal level. Many experiments have demonstrated metaplasticity using distinct protocols (Abraham, 2008). Quite often, it also involves some form of heterosynaptic plasticity, in the sense that the local changes affecting a particular synapse onto a post-synaptic neuron influence the plasticity for neighboring synapses.

The study of the dynamical implications of the interaction between homeostatic mechanisms and Hebbian plasticity requires the integration of experimental data in model studies (Marder and Goaillard, 2006). From a modeler's point of view, interactions between Hebbian learning and its regulating counterpart, either by homeostatic mechanisms or

by metaplasticity, is problematic. The principal reason being that those two distinct forms of plasticity do not act on similar timescales. Following experimental results, it is commonly assumed that synaptic changes triggered by Hebbian plasticity protocols are rather fast (Bliss and Lomo, 1973; Sjöström et al., 2001, 2003; Wang et al., 2005), occurring in the timescale of minutes or faster, while metaplasticity or homeostatic changes are much slower (Abraham and Bear, 1996), in the order of days. The present paper provides a theoretical framework to analyze the interaction between Hebbian and homeostatic plasticities at different timescales. In this way it gives an overarching view of different methods used in the literature to solve the above-mentioned instability issue of Hebbian plasticity. Maintaining the stability only being one of the requirements for proper behavior, we will discuss how homeostatic constraints can also be used to adjust the function implemented by the neural circuits.

2. THE APPARENT ANTAGONISM BETWEEN HEBBIAN AND HOMEOSTATIC PLASTICITY

2.1. Two Divergent Goals

As it has already been observed (Turrigiano and Nelson, 2000; Watt and Desai, 2010; Vitteira and Goda, 2013), Hebbian and homeostatic plasticities are two apparently opposing processes, which compete at the synaptic level to fulfill different goals. Hebbian learning promotes strong or synchronous firing among neurons, which is hypothesized to be a building block for memory storage (Nabavi et al., 2014). In contrast, homeostatic processes counterbalance such intense spiking activity to maintain the global stability in neuronal networks (Turrigiano and Nelson, 2000; Turrigiano, 2008; Pozo and Goda, 2010). Several types of homeostatic processes have been observed at the neuronal level in many brain areas, such as synaptic scaling (Turrigiano et al., 1998) and intrinsic plasticity (Zhang and Linden, 2003).

It has been long known that Hebbian plasticity *alone* is intrinsically unstable (Rochester et al., 1956; von der Malsburg, 1973; Miller, 1996). The entrainment between synapses often force all to grow boundlessly or to a maximal set value; in other cases, they may all become silent. To circumvent these issues of traditional rate-based Hebbian learning, weight normalization can be introduced to prevent the runaway of synapses (Oja, 1982; Miller, 1996). In the context of spiking activity, STDP has been termed “temporally Hebbian” when it promotes synchronous firing. Weight-dependent STDP update rules, which induces more LTD than LTP for strong synapses, provide a fixed point in the learning dynamics (van Rossum et al., 2000; Gütiğ et al., 2003). Although this ensures some stability, it may dramatically change the weight distribution from being bimodal to being unimodal. In the case of a narrow unimodal weight distribution, competition induced by STDP among synapses is weakened between pathways with distinct characteristics (e.g., rate, correlation), which is not functionally interesting. For weight-dependent STDP, this trade-off compromise is only fulfilled in a given parameter range. In recurrent networks

especially, the synaptic specialization by competition may be severely impaired without fine tuning (Morrison et al., 2007; Gilson and Fukai, 2011).

2.2. Two Different Timescales

Most of the plasticity protocols performed *in vitro* are based on either input stimulation at a high/low frequency leading to LTP/LTD (Bliss and Lomo, 1973) or STDP-type pairings of pre-post spikes (Markram et al., 1997; Bi and Poo, 1998; Sjöström et al., 2001; Froemke and Dan, 2002; Wang et al., 2005). The typical protocol used in cortical or hippocampal slices to elicit STDP *in vitro* using spike pairs is represented in **Figure 1A**: a spike is triggered at the pre-synaptic neuron and another at the post-synaptic neuron with time difference $\delta t = t_{\text{pre}} - t_{\text{post}}$. This pairing is repeated approximately 60 times with frequency $f_{\text{pairing}} = 1$ Hz in order to see a robust change in the weight: it has been shown that after an induction phase, the total weight change evolves non-linearly up to a saturation plateau, at around 60–100 pairings (Froemke et al., 2006), which corresponds to the number of protocol repetitions in most studies.

For modelers, this STDP protocol leads to the simplified view of the time-difference window in **Figure 1B**, where a single pre spike followed by a post spike will trigger LTP, whereas post followed by pre causes LTD. This is clearly an over-simplification of a much more complex phenomenon. Just to mention some limitations of this simplified view, it has been shown that if the frequency f_{pairing} of the pairing is changed, the typical STDP curve with LTP for $\delta t < 0$ and LTD for $\delta t > 0$ is dramatically modified (Sjöström et al., 2001). Depression is only visible for low

frequency pairings, when pairings are performed with $\delta t < 0$ and $f_{\text{pairing}} < 20$ Hz. For $f_{\text{pairing}} > 20$ Hz, however, synapses undergo LTP irrespective of the sign for δt . Moreover, several *in vitro* studies on cortical pyramidal neurons showed that the canonical shape of the STDP curve for such pre-post pairings strongly depends of the position of the synapse along the dendritic tree (Froemke et al., 2005; Letzkus et al., 2006; Kampa et al., 2007), as well as the post-synaptic voltage (Artola et al., 1990). Those experimental findings led to the refinements of initial STDP models based on the curve, in order to incorporate the observed effects for triplets of spikes, spike bursts, clamping the post-synaptic membrane potential and so on (Pfister and Gerstner, 2006; Clopath et al., 2010; El Boustani et al., 2012; Graupner and Brunel, 2012; Yger and Harris, 2013).

Despite those efforts, there is a point that is almost never considered: STDP changes are not instantaneous. In most experiments, when plasticity protocols are performed, the resulting weight is recorded up to 30 min later. The curve in **Figure 1B** corresponds to the corresponding weight change divided by the number of pairings. In models of classical (van Rossum et al., 2000; Song and Abbott, 2001) and weight-dependent (van Rossum et al., 2000; Gütig et al., 2003; Morrison et al., 2007; Gilson and Fukai, 2011) STDP, its final value is the results of additive instantaneous and independent weight updates following each pairing. In fact, even elaborate models consider the linear summation of weight updates, even when contributions are restricted to neighboring spikes (Burkitt et al., 2004). Only a few attempts have been done to change this property that is convenient for theory, such as probabilistic

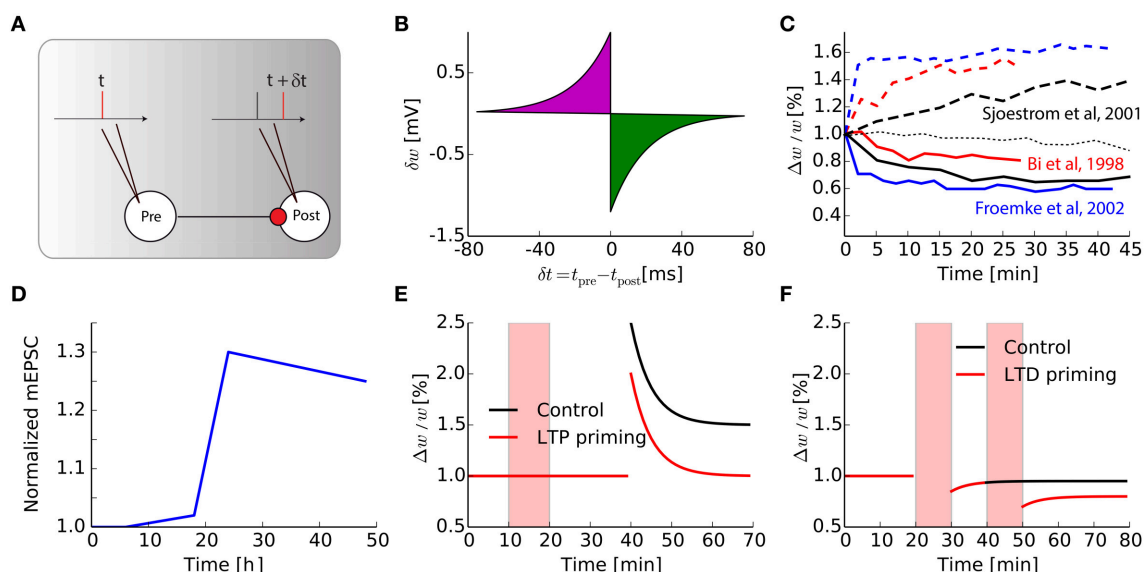


FIGURE 1 | Intrinsic timescale of Hebbian learning. (A) The classical STDP pairing protocols widely used in the literature. **(B)** Synaptic modification for one pair of pre- and post-synaptic spikes, as a function of their relative timing. **(C)** Evolution as a function of time of a single synaptic weight, after an STDP protocol, for various papers taken from the literature, both for LTP or LTD protocols [dash-dotted thin black line is the null-line for Sjöström et al. (2001)]. **(D)** Adapted from Keck et al. (2013), Normalized mEPSC amplitude in a layer 5 cell in the mice visual cortex following a lesion in the retina. **(E)** Adapted from Huang et al. (1992), Prior synaptic activity triggered during the red shaded area (LTP priming, red curve) reduces LTP in CA1 hippocampus compared to control without pre-activation (black curve). **(F)** Adapted from Mockett et al. (2002), Low frequency stimulation (LFS, red shaded areas) influences non-linearly the amount of LTD in CA1 hippocampus: black curve, (control with only one LFS), red curve (two consecutive LFS).

models of STDP (Appleby and Elliott, 2005). By re-examining the weight traces found in the STDP literature (Bi and Poo, 1998; Sjöström et al., 2001; Froemke and Dan, 2002; Froemke et al., 2006) and reproduced in **Figure 1C**, it can be seen that the weights actually evolve continuously *in vitro*. Therefore, plasticity should better be seen as a phenomenon that is triggered by a stimulation event and evolves toward a new equilibrium with a time constant $\tau_{\text{Hebb}} \simeq 10$ min.

Now considering that Hebbian plasticity induces such a transient synaptic change, the question arises about its interaction with homeostatic plasticity. Those processes, either intrinsic or synaptic, are assumed to be much slower. For example synaptic scaling, one of the numerous mechanisms of homeostasis, takes place *in vitro* with a time constant τ_{homeo} of the order of a day (Turrigiano and Nelson, 2000), and *in vivo* during the 2–3 days after an abrupt change, as observed for neurons in the visual cortex following visual deprivation (Hengen et al., 2013; Keck et al., 2013). **Figure 1D**, adapted from Keck et al. (2013), shows the amplitude of miniature EPSC in V1 neurons after a bilateral lesion in the adult retina: after an initial period of about a day, amplitudes are scaled up to compensate for the reduced inputs. Together, these results stress the fact that Hebbian and homeostatic processes have distinct timescales. Understanding the biological mechanisms responsible for those changes at the molecular level is necessary to gain a better insight on the interaction between them, especially *in vivo* where synapses are constantly bombarded by spikes.

2.3. Primings as an Evidence for Metaplasticity

Although on a first approximation it may appear that $\tau_{\text{homeo}} \gg \tau_{\text{Hebb}}$, several experiments show that those two timescales may be more interleaved. In hippocampal slices, it has been shown in so-called priming experiments that the activation of a synapse before its reactivation modulates the plasticity triggered later at that particular synapse (see **Figures 1E,F**). In **Figure 1E** that is adapted from Huang et al. (1992), weak tetanic priming stimulations can reduce the amount of LTP obtained during a strong subsequent tetanic stimulation; note that the effects last more than 1 h. On the contrary, the LTD pathway seems to be facilitated when the synapse is preactivated a few hours before the plasticity protocol (Christie and Abraham, 1992; Wang, 1998; Mockett et al., 2002). This is illustrated in **Figure 1F**, adapted from Mockett et al. (2002), where the effect lasts at least 2 h. Those primings experiments suggest the existence of long-lasting regulation mechanisms, acting over large time constants, which counteracts the effect of Hebbian learning. This modulation of the Hebbian plasticity by preactivation of the synaptic pathway is a direct application of the so-called metaplasticity (Abraham and Bear, 1996), i.e., the plasticity of the learning rules themselves.

3. MATHEMATICAL FORMALISM

To formally study the interactions between Hebbian and homeostatic plasticity, we use the following mathematical formalism. We consider a Poisson neuron (Kempter et al., 1999) with N synapses indexed by i , corresponding to the input firing

rates r_i ; for STDP examples, we also define the input cross-covariances c_{ij} between neurons i and j . The equations for the output firing rate r_{post} and pre-post covariances $c_{i-\text{post}}$ between synapse i and the post-synaptic spike train in a feedforward scenario are given by

$$\begin{aligned} r_{\text{post}} &= \sum_{1 \leq j \leq N} w_j r_j \\ c_{i-\text{post}} &= \sum_{1 \leq j \leq N} w_j c_{ij} \end{aligned} \quad (1)$$

In order to compare several learning rules in the context of metaplasticity, we consider the following general equations for the evolution of a given weight w_i and a modulation parameter θ :

$$\begin{aligned} \dot{w}_i &= \frac{1}{\tau_{\text{Hebb}}} \Phi(w_i, r_i, r_{\text{post}}, c_{i-\text{post}}, \theta) \\ \dot{\theta} &= \frac{1}{\tau_{\text{homeo}}} [\Psi(r_{\text{post}}) - \theta] \end{aligned} \quad (2)$$

The motivation for these expressions is to model the two timescales explicitly, as previously done for the BCM rule (Bienenstock et al., 1982) and for an extension of the triplet STDP rule (Zenke et al., 2013): τ_{Hebb} and τ_{homeo} are the two time constants at which both Hebbian and homeostatic changes are propagated onto the synapses. The Hebbian plasticity update is embodied in Φ , which also depends on r_{pre} , $c_{\text{pre-post}}$, etc. The parameter θ is global for all synapses of a neuron and interacts or modulates the corresponding weight updates. Typically, it is used to implement a homeostatic mechanism, as we will see for several models of synaptic plasticity that are commonly used in the literature. The present framework could be extended to incorporate other non-linearities in the firing mechanism (e.g., LIF neuron), adaptation or intrinsic plasticity.

3.1. Stability Analysis for the Mean-field Dynamical System

Ignoring correlations and inhomogeneities across synapses, we focus on the analysis of the mean weight $\bar{w} = \sum_j w_j / N > 0$ with mean input rate r_{pre} . The rate Equation (1) simply becomes

$$r_{\text{post}} = \bar{w} r_{\text{pre}} \quad (3)$$

This allows for an easy comparison of the weight dynamics based on polynomial expressions in \bar{w} . Other neuron models usually give more complex mapping between input and output rate/correlations, but the common trend is that they are monotonically increasing function of the weight \bar{w} . This property is the cause for the instability of Hebbian learning, as it increases \bar{w} all the more as r_{post} is large. Therefore, we will review through the example of the Poisson neuron how stabilization mechanisms interact with the Hebbian component.

In order to examine the stability of the mean-field dynamical system (Equation 2) where w_i is replaced by \bar{w} , we consider its Jacobian matrix.

$$\begin{pmatrix} \frac{1}{\tau_{\text{Hebb}}} \left[\frac{\partial \Phi}{\partial \bar{w}} + \frac{\partial \Phi}{\partial r_{\text{post}}} r_{\text{pre}} \right] & \frac{1}{\tau_{\text{Hebb}}} \frac{\partial \Phi}{\partial \theta} \\ \frac{1}{\tau_{\text{homeo}}} \frac{\partial \Psi}{\partial r_{\text{post}}} r_{\text{pre}} & \frac{1}{\tau_{\text{homeo}}} \frac{\partial \Psi}{\partial \theta} \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \quad (4)$$

For the top-left term in the Jacobian, we have used the following equality for the feedforward architecture corresponding to Equation (3): $\frac{\partial r_{\text{post}}}{\partial w} = r_{\text{pre}}$. The eigenvalues of the Jacobian matrix are given by

$$x_{\pm} = \frac{1}{2}(T \pm \sqrt{T^2 - 4D}) \quad (5)$$

where $T = a + d$ is the trace and $D = ad - bc$ is the determinant. To ensure stability for this 2-dimensional dynamical system, these eigenvalues must be real negative. This requires that the following relationships are satisfied.

$$\begin{aligned} T &< 0 \\ 0 &< 4D < T^2 \end{aligned} \quad (6)$$

If, however, the discriminant is positive with the trace still negative ($T < 0$ and $T^2 < 4D$) the system exhibits damped oscillations related to the imaginary eigenvalues. With purely imaginary eigenvalues, we may obtain a limit cycle. Finally, when $D < 0$ or $T > 0$, at least one eigenvalue is positive and can lead to an explosion of the mean weight.

3.2. Competition between Input Pathways

Following Kempster et al. (1999); Gütig et al. (2003); Gilson et al. (2009), we can use and rewrite Equation (1) to study the competition between learning weights.

$$r_{\text{post}} = w_1 r_1 + w_2 r_2 \quad (7)$$

Again ignoring correlations, we obtain the following 3-dimensional learning system.

$$\begin{aligned} \dot{w}_1 &= \frac{1}{\tau_{\text{Hebb}}} \Phi(w_1, r_1, r_{\text{post}}, \theta) \\ \dot{w}_2 &= \frac{1}{\tau_{\text{Hebb}}} \Phi(w_2, r_2, r_{\text{post}}, \theta) \\ \dot{\theta} &= \frac{1}{\tau_{\text{homeo}}} [\Psi(r_{\text{post}}) - \theta] \end{aligned} \quad (8)$$

Considering the equilibrium for the mean weight $\bar{w} = (w_1 + w_2)/2$ to be satisfied, the competition between the two input pathways can be studied for what is called “symmetry breaking” namely the divergence of w_1 and w_2 . This relates to the following differential equation for the weight difference $\Delta w = w_1 - w_2$, which quantifies the tendency for splitting

$$\begin{aligned} \Delta \dot{w} &= \frac{1}{\tau_{\text{Hebb}}} [\Phi(w_1, r_1, r_{\text{post}}, \theta) - \Phi(w_2, r_2, r_{\text{post}}, \theta)] \quad (9) \\ &\simeq \frac{1}{\tau_{\text{Hebb}}} \frac{\partial \Phi}{\partial w} (\bar{w}, \bar{r}_{\text{pre}}, r_{\text{post}}, \theta) \Delta w \\ &\quad + \frac{1}{\tau_{\text{Hebb}}} \frac{\partial \Phi}{\partial r_{\text{pre}}} (\bar{w}, \bar{r}_{\text{pre}}, r_{\text{post}}, \theta) \Delta r \end{aligned}$$

where $\bar{r}_{\text{pre}} = (r_1 + r_2)/2$ and $\Delta r = r_1 - r_2$ is assumed to be small here. The larger positive $\frac{\partial \Phi}{\partial w} (\bar{w}, \bar{r}_{\text{pre}}, r_{\text{post}}, \theta)$ is, the more strongly the weights w_1 and w_2 will move apart from each other.

3.3. Conditions for Joint Stability and Competition for Hebbian Learning with Synaptic Scaling

In general, the equations for stability and competition may turn out to be quite complex, even for the mean-field dynamical system. The ambition here is to describe the general trends for the influence of τ_{Hebb} and τ_{homeo} on the behavior of the dynamical learning system. To illustrate this, we examine the “simple” case of an arbitrary Hebbian-type learning rule with additional synaptic scaling. Inspired by experimental results (Turrigiano and Nelson, 2000) and used in previous studies (van Rossum et al., 2000; Yger and Harris, 2013; Zenke et al., 2013), synaptic scaling is used as a homeostatic mechanism that increases or decreases homogeneously the synaptic weights in order to reach a given firing rate r_{target} . In our generic formulation in Equation (2), this is equivalent to including an additive scaling term Γ in the expression of Φ in addition to the Hebbian contribution H , while θ tracks the post-synaptic firing rate with a timescale τ_{homeo} .

$$\begin{aligned} \Phi(\bar{w}, r_{\text{pre}}, r_{\text{post}}, \theta) &= H(\bar{w}, r_{\text{pre}}, r_{\text{post}}) + \Gamma(\bar{w}, \theta) \quad (10) \\ \Gamma(\bar{w}, \theta) &= \alpha \bar{w} (r_{\text{target}} - \theta) \\ \Psi(r_{\text{post}}) &= r_{\text{post}} \end{aligned}$$

For simplicity, we rewrite the Hebbian contribution using Equation (3) in terms of \bar{w} only: $\tilde{H}(\bar{w}) := H(\bar{w}, r_{\text{pre}}, r_{\text{post}})$. This yields the following expression for the Jacobian in Equation (4):

$$\begin{pmatrix} \frac{\tilde{H}'(\bar{w}) + \alpha(r_{\text{target}} - \theta)}{\tau_{\text{Hebb}}} & \frac{-\alpha \bar{w}}{\tau_{\text{homeo}}} \\ \frac{\tau_{\text{Hebb}}}{\tau_{\text{homeo}}} & \frac{-1}{\tau_{\text{homeo}}} \end{pmatrix} \quad (11)$$

The equilibrium corresponds to the fixed point(s) where $\dot{\bar{w}} = 0$ and $\dot{\theta} = 0$, which implies that $r_{\text{post}} = \theta$ and $\tilde{H}(\bar{w}) = -\alpha \bar{w} (r_{\text{target}} - \theta)$. The trace and determinant of the Jacobian matrix are given by

$$\begin{aligned} T &= \frac{\tilde{H}'(\bar{w}) - \tilde{H}(\bar{w})/\bar{w}}{\tau_{\text{Hebb}}} - \frac{1}{\tau_{\text{homeo}}} \quad (12) \\ D &= \frac{-[\tilde{H}'(\bar{w}) - \tilde{H}(\bar{w})/\bar{w}] + \alpha r_{\text{post}}}{\tau_{\text{Hebb}} \tau_{\text{homeo}}} \end{aligned}$$

As explained above, stability is ensured when the necessary conditions $T < 0$ and $0 < 4D < T^2$ in Equation (6) are met. These three conditions read.

$$\tilde{H}'(\bar{w}) - \frac{\tilde{H}(\bar{w})}{\bar{w}} < \frac{\tau_{\text{Hebb}}}{\tau_{\text{homeo}}} \quad (13)$$

$$\tilde{H}'(\bar{w}) - \frac{\tilde{H}(\bar{w})}{\bar{w}} < \alpha r_{\text{post}} \quad (14)$$

$$\alpha r_{\text{post}} < \frac{1}{4} \left\{ \left[\tilde{H}'(\bar{w}) - \frac{\tilde{H}(\bar{w})}{\bar{w}} \right] \sqrt{\frac{\tau_{\text{homeo}}}{\tau_{\text{Hebb}}}} + \sqrt{\frac{\tau_{\text{Hebb}}}{\tau_{\text{homeo}}}} \right\}^2 \quad (15)$$

The term $\tilde{H}'(\bar{w}) - \tilde{H}(\bar{w})/\bar{w}$ corresponds to the sub/super-linearity of the effective weight update \tilde{H} at the equilibrium \bar{w} , including

the effects of the neuron model. For the simplest Hebbian rule $H(\bar{w}, r_{\text{pre}}, r_{\text{post}}) = r_{\text{pre}} r_{\text{post}}$, $\tilde{H}(\bar{w}) = r_{\text{pre}}^2 \bar{w}$ is linear and we always have $\tilde{H}(\bar{w}) - \tilde{H}(\bar{w})/\bar{w} = 0$. This implies that the first two conditions Equations (13) and (14) are always true, while the third condition Equation (15) reduces to $\alpha r_{\text{post}} < \tau_{\text{Hebb}}/4\tau_{\text{homeo}}$. For the synaptic scaling mechanism, α should be chosen sufficiently large in order to keep the output rate r_{post} close to its target r_{target} . It follows that the third condition may be violated depending on the details of the parameters, in particular when $\tau_{\text{Hebb}} \ll \tau_{\text{homeo}}$. This corresponds to non-real eigenvalues, synonymous with oscillatory dynamics in the weights.

As a second example related to the BCM rule and triplet-STDP as will be detailed later, when \tilde{H} is a quadratic polynomial in \bar{w} with positive second-order coefficient, we have $\tilde{H}'(\bar{w}) - \tilde{H}(\bar{w})/\bar{w} > 0$ for large weights. According to Equation (9), a large positive value for $\tilde{H}'(\bar{w})$ implies strong competition as desired. However, the condition for the negativity of the trace in Equation (13) implies that τ_{homeo} should not be much larger than τ_{Hebb} , as shown previously (Zenke et al., 2013). Then, assuming Hebbian learning to be relatively fast, Equations (14) and (15) define a limited range for the choice of α , out of which divergence or oscillations may occur. As a conclusion, those stability and competition conditions oppose each other and make the fine tuning of the parameters necessary.

4. THE FAMILY OF STDP LEARNING RULES

4.1. Need for Regulation with Classical STDP

As a first example of learning rules, we consider the family of STDP rules to illustrate the interplay between Hebbian learning and synaptic scaling. We show that they fall into the mathematical framework developed in Section 3. To start, without any additional homeostatic regulation based on θ , we recall that the convergence of the weight depends on the fixed points of Φ only. The original version of STDP simply describes the effect for pairs of input-output spikes using the well-known temporal window in **Figure 1B**, which determines the weight update as a function of spike-time difference. All contributions are then summed over time to obtain the total weight update. The net effect denoted by H here can be decomposed into two terms, for the neuronal firing rates and covariances, respectively (Kempster et al., 1999; Gilson et al., 2009). In our framework based on the Poisson neuron (see Section 3), this gives the following differential equation for the mean weight \bar{w}

$$\begin{aligned} \dot{\bar{w}} &= \frac{1}{\tau_{\text{Hebb}}} H(\bar{w}, r_{\text{pre}}, r_{\text{post}}, c_{\text{pre-post}}) \\ &= \frac{1}{\tau_{\text{Hebb}}} (A r_{\text{pre}} r_{\text{post}} + B c_{\text{pre-post}}) \\ &= \frac{1}{\tau_{\text{Hebb}}} (A r_{\text{pre}}^2 + B c_{\text{pre}}) \bar{w} \end{aligned} \quad (16)$$

where the typical area under the curve $A < 0$ corresponds to more LTD than LTP for the rate contribution, while $B > 0$ describes LTP due to the temporal interaction for correlated

inputs. The last line is obtained using Equation (3), where the mean weight update can be expressed as a linear function of the weight from a macroscopic point of view. We obtain a first-order polynomial similar to that for classical Hebbian learning, where the coefficient depends on the input correlation. Two behaviors can occur for this system: for sufficiently strong input correlations c_{pre} , the factor for \bar{w} becomes positive and the fixed point unstable, so positive weights are potentiated in a Hebbian fashion and diverge; otherwise weights are depressed and converge to the fixed point $\bar{w} = 0$. For a pool of synapses, competition is ensured provided $\frac{\partial H}{\partial \bar{w}} = (A r_{\text{pre}}^2 + B c_{\text{pre}})/\tau_{\text{Hebb}} > 0$, which occurs for sufficiently strong input correlation here. In that case, the diverging learning dynamics can result in a bimodal distribution when a positive upper bound is set (Kempster et al., 1999; Song and Abbott, 2001).

To change the fixed-point structure and enforce stability, one can add a penalty term on the weight update based on the current value of the weight (Oja, 1982). A usual example found in the literature uses a polynomial in \bar{w} , which leads to the following expression for Φ .

$$\dot{\bar{w}} = \frac{1}{\tau_{\text{Hebb}}} (A r_{\text{pre}}^2 + B c_{\text{pre}}) \bar{w} - \alpha \bar{w}^n \quad (17)$$

The key point here is that $\dot{\bar{w}}$ is a first-order polynomial in w for classical STDP, so $n \geq 2$ stabilizes the system (Tetzlaff et al., 2011). Synaptic scaling maintains the synaptic competition while preventing weights from taking too high values, at the cost of not being able to control the post-synaptic firing rate, and without having any relationship to the real homeostatic timescale. Although that previous work studied in depth the interaction of synaptic scaling with more complex Hebbian learning rule, the temporal dynamics when the two processes are not acting on the same timescale is still poorly understood.

4.2. Synaptic Scaling Mechanism Targeting a Fixed Firing Rate Requires Fine Tuning

In order to target a fixed firing rate, weight normalization as previously defined is not sufficient. One must add a constraint enforcing the post-synaptic neuron to scale all its input weights such that, on average, a desired firing rate is maintained. Following previous studies (van Rossum et al., 2000; Yger and Harris, 2013; Zenke et al., 2013), it can be implemented by the term Γ as in Equation (10), which depends on the difference between a running estimate of the post-synaptic firing rate and a desired firing rate r_{target} . The expression for Φ with the STDP contribution H and Ψ then read

$$\begin{aligned} \Phi(\bar{w}, r_{\text{post}}, c_{\text{pre-post}}, \theta) &= H(\bar{w}, r_{\text{post}}, c_{\text{pre-post}}) \\ &\quad + \alpha \bar{w} (r_{\text{target}} - \theta) \\ \Psi(r_{\text{post}}) &= r_{\text{post}} \end{aligned} \quad (18)$$

The constant α defines the strength of the homeostasis on the mean weight \bar{w} , while τ_{homeo} determines the timescale of the smoothing of the r_{post} estimate tracked by θ .

The analysis in Section 3 states that τ_{Hebb} , τ_{homeo} and α must be chosen so as to avoid instability and trivial solutions where all

weights become silent. As it has been shown for other learning rules (Cooper et al., 2004; Zenke et al., 2013), the running estimate θ of the post-synaptic firing rate have to be rather fast, otherwise the system is subject to strong oscillations. To illustrate the problem, suppose we have a neurons targeting $r_{\text{target}} = 1$ Hz, with $r_{\text{pre}} = 0.9$ Hz, $c_{\text{pre}} = 0.1$, $A = -0.1$ and $B = 1$ (see Section 4.1). The value of α is varied between 0.01, 0.1, and 1. As we can see on **Figure 2A**, the convergence to the fixed point can be pretty fast if $\tau_{\text{Hebb}} = \tau_{\text{homeo}}$, and if α is strong enough to counterbalanced the Hebbian force that depresses synapses here; see panels with $\alpha \geq 0.1$, insets show the trajectory in the phase space (\bar{w}, θ) as function of time. However, we can see on **Figure 2B** that when $\tau_{\text{homeo}} \gg \tau_{\text{Hebb}}$, as it is found *in vivo* (Keck et al., 2013), strong oscillations emerge for strong value $\alpha = 1$. There is a fine tuning required between those two competing forces. To circumvent the problem, the use of a Proportional-Integral (PI) controller was incorporated in some study (van Rossum et al., 2000; Yger and Harris, 2013), but even when it prevents some oscillations from occurring, it does not abolish the requirement that τ_{Hebb} and τ_{homeo} should not be order of magnitudes apart.

4.3. Similar Stability Issues Occur for Weight-dependent and Triplet STDP

The analysis and the observations performed previously can be extended to several STDP-like learning rules. For example, a simple version of the weight-dependent STPD learning rule (van Rossum et al., 2000; Morrison et al., 2007; Gilson and Fukai, 2011) with linearly increasing LTD as a function of the weight and constant LTP gives

$$\begin{aligned} H(\bar{w}, r_{\text{post}}, c_{\text{pre-post}}) &= (A_+ + A_- \bar{w}) r_{\text{pre}} r_{\text{post}} \\ &\quad + B c_{\text{pre-post}} \\ &= (A_+ r_{\text{pre}}^2 + B c_{\text{pre}}) \bar{w} + A_- r_{\text{pre}}^2 \bar{w}^2 \end{aligned} \quad (19)$$

Again Equation (3) was used to obtain the second-order polynomial in \bar{w} . In **Figure 3B** that depicts the convergence of the system in a similar fashion to **Figure 2** with typical values for the parameters ($A_+ = 0.1$, $A_- = -0.3$, $B = 1$, $c_{\text{pre}} = 0.1$), the convergence is achieved if homeostatic coupling is weak ($\alpha = 0.1$). However, large oscillations arise for strong coupling ($\alpha = 1$) and when the ratio between the homeostatic and Hebbian timescales is large.

Likewise, the triplet STDP model (Pfister and Gerstner, 2006) corresponds to

$$\begin{aligned} H(\bar{w}, r_{\text{post}}, c_{\text{pre-post}}) &= (A_+ r_{\text{post}} + A_-) r_{\text{pre}} r_{\text{post}} \\ &\quad + B c_{\text{pre-post}} \\ &= (A_- r_{\text{pre}}^2 + B c_{\text{pre}}) \bar{w} + A_+ r_{\text{pre}}^3 \bar{w}^2 \end{aligned} \quad (20)$$

where $A_+ > 0$, $A_- < 0$ for the LTP and LTD rate contributions, respectively, as well as $B > 0$ for the correlation contribution. Again, for standard values of the parameters $A_+ = 0.05$, $A_- = -0.2$, $B = 1$, $c_{\text{pre}} = 0.1$ (Pfister and Gerstner, 2006), we see in **Figure 3B** the same qualitative behavior as with weight-dependent STDP.

The similarity can be explained by the fact that both Equations (19) and (20) are quadratic polynomials in \bar{w} . The difference between the two rules lies in the signs of the coefficients. Nevertheless, we have for the scaling term $\Gamma(\bar{w}, \theta) = \alpha \bar{w}(r_{\text{target}} - \theta) \simeq \alpha r_{\text{target}} \bar{w} - \alpha r_{\text{pre}} \bar{w}^2$, where we have used $\theta \simeq r_{\text{post}} = r_{\text{pre}} \bar{w}$. This means that, when Γ overpowers the STDP contribution to enforce stability with a large α , the coefficient for \bar{w} in Φ is negative in both cases. It ensures stability, but generates similar oscillations for large τ_{homeo} . The intuitive explanation is that large values for τ_{homeo} cause the gradient to have a strong horizontal component in the phase space (\bar{w}, θ) of **Figure 3**, which often implies oscillations around the fixed point.

4.4. Trade-off between Stability and Competition

While we analyzed the dynamical behavior of the learning rules for the mean weight to assess their implications for stability, we now examine the situation for two inputs in order to study how competition can be affected by this interaction between homeostatic and Hebbian learning. This yields an extra differential equation as explained in 3.2. **Figure 4** illustrates the evolution of the weights w_1 and w_2 , as well as θ , for the three learning rules previously mentioned combined with synaptic scaling, and show how competition can take place. We consider two input pathways with the same input rates $r_{1/2}$, but different levels of correlation: $c_{\text{pre}}^1 = 0.1$ and $c_{\text{pre}}^2 = 0.05$. The homeostatic mechanism targets the fixed firing rate $r_{\text{target}} = 2$ Hz. As shown in **Figures 4A,C**, strong competition is observed for both classical STDP and triplet STDP, leading to $w_2 = 0$ for the pathway with weaker correlation $c_2 < c_1$ (Kempster et al., 1999; van Rossum et al., 2000; Song and Abbott, 2001). For weight-dependent STDP, the competition is much weaker in **Figure 4B**. Nevertheless, in all cases, increasing the ratio $\tau_{\text{homeo}}/\tau_{\text{Hebb}}$ introduces oscillations of the weights during competition, exactly as previously observed for the mean weight \bar{w} . We also see that an increased strength for the homeostatic force ($\alpha = 0.5$ in the bottom row of **Figure 4**) does not solve the stability issue when $\tau_{\text{homeo}} \gg \tau_{\text{Hebb}}$, but causes larger fluctuations.

5. METAPLASTIC LEARNING RULES

The previous section showed the common trend for STDP learning rules paired with synaptic scaling targeting a desired firing rate: a large time constant to estimate the post-synaptic firing rate gives rise to instability or potentially large oscillations in the weights. Now we examine a second category of stabilizing mechanisms, where the homeostatic mechanism is implemented directly in the metaplastic learning rule; see Yeung et al. (2004) for an example for calcium-based regulation. Metaplasticity is often used to enforce a homeostatic behavior on the neural system and we will stick to this function here. Without loss of generality, we ignore correlations in the learning rules and focus on rate-based rules.

5.1. The Bienenstock-Cooper-Munro (BCM) Learning Rule

In order to extend rules based on correlations of rates (Oja, 1982) and approach the problem of synaptic competition via weight normalization, Bienenstock et al. (1982) designed a model of synaptic plasticity that was able to reproduce phenomenologically several observations made *in vivo*. Their so-called BCM rule is a physical theory of learning in the visual cortex; see Cooper and Bear (2012) for a review. The mechanism

consists in an efficient way to balance and regulate the amount of plasticity according to past activity by means of a heterosynaptic process.

Practically, a sliding threshold determines the boundary between LTP above and LTD below, and evolves according to the square of the postsynaptic firing rate (Bienenstock et al., 1982). In our formalism, this can be taken care of by a temporal tracking of r_{post}^2 using θ as the threshold variable with $\tau_{\text{homeo}} \gg \tau_{\text{Hebb}}$, such that $\theta \simeq \langle r_{\text{post}}^2 \rangle$ with the angular brackets indicating the

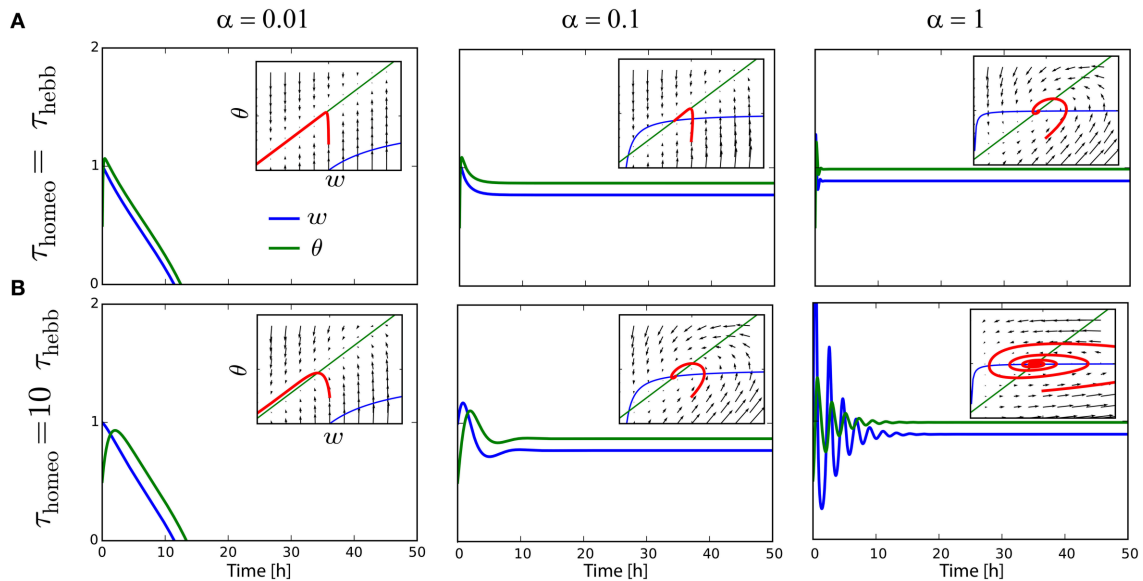


FIGURE 2 | Interplay between Hebbian and homeostatic timescales for pairwise STDP. (A) Evolution of the weight w and the running estimate θ of the post-synaptic firing rate as function of time, for $\tau_{\text{Hebb}} = \tau_{\text{homeo}} = 10$ min and various gain α for the heterosynaptic scaling. Insets show the trajectory in the phase space (w, θ) . **(B)** Same as **(A)** with a slower homeostatic scaling: $\tau_{\text{homeo}} = 10$ $\tau_{\text{Hebb}} = 100$ min.

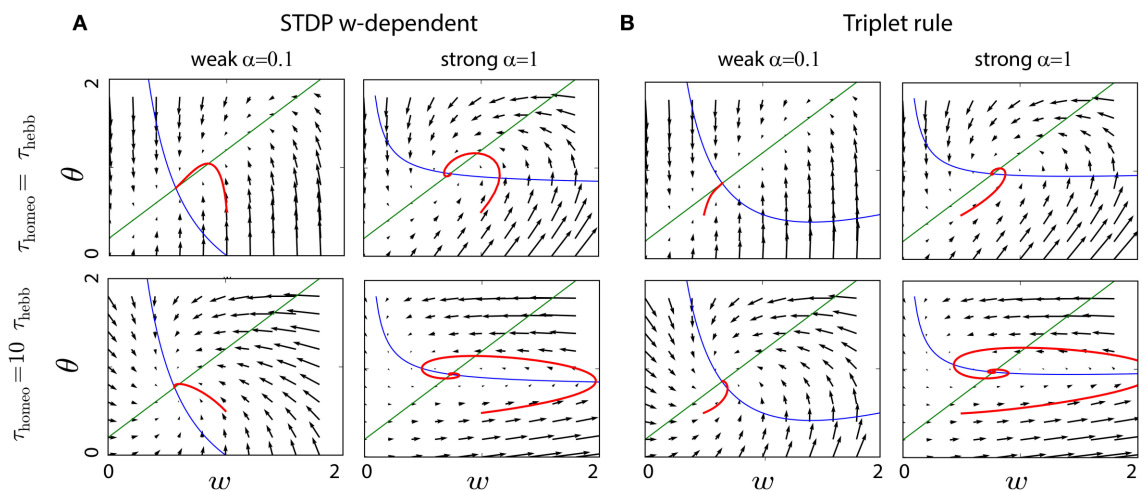


FIGURE 3 | Interplay between Hebbian and homeostatic timescales for different learning rules and homeostatic forces. (A) Pairwise STDP with weight dependent modification in Equation (19). Left column: convergence in the phase space (w, θ) for a fast homeostatic force ($\tau_{\text{Hebb}} = \tau_{\text{homeo}} = 10$ min, upper row), or for a slow homeostatic force ($\tau_{\text{homeo}} = 10\tau_{\text{Hebb}}$, lower row). Right column is the same, but with a stronger drive $\alpha = 1$ for the homeostatic force. **(B)** Same as **(A)** for the triplet learning rule (Pfister and Gerstner, 2006), see Equation (20).

average over the randomness. The expression for Φ is a second-order polynomial in r_{post} (Bienenstock et al., 1982), which finally gives

$$\begin{aligned}\Phi(\bar{w}, r_{\text{post}}, \theta) &= r_{\text{pre}} r_{\text{post}} (r_{\text{post}} - \theta) \\ \Psi(r_{\text{post}}) &= r_{\text{post}}^2\end{aligned}$$

Here Φ has a similar form to that for the triplet rule in Equation (20), but the boundary between potentiation and depression is now given by θ .

It is known that the BCM formalism can be subject to strong oscillations, when the timescales for the two differential equations are too far apart (Cooper et al., 2004; Toyozumi et al., 2014). In **Figure 5A**, even when $\tau_{\text{Hebb}} = \tau_{\text{homeo}}$, weight oscillations are present. Moreover, for a slightly larger ratio $\tau_{\text{Hebb}}/\tau_{\text{homeo}}$, the oscillations can destroy the convergence of the system when the weights hit the lower bound 0, as illustrated in **Figure 5B**.

5.2. Modulation of STDP Depending on the Post-synaptic Firing Rate

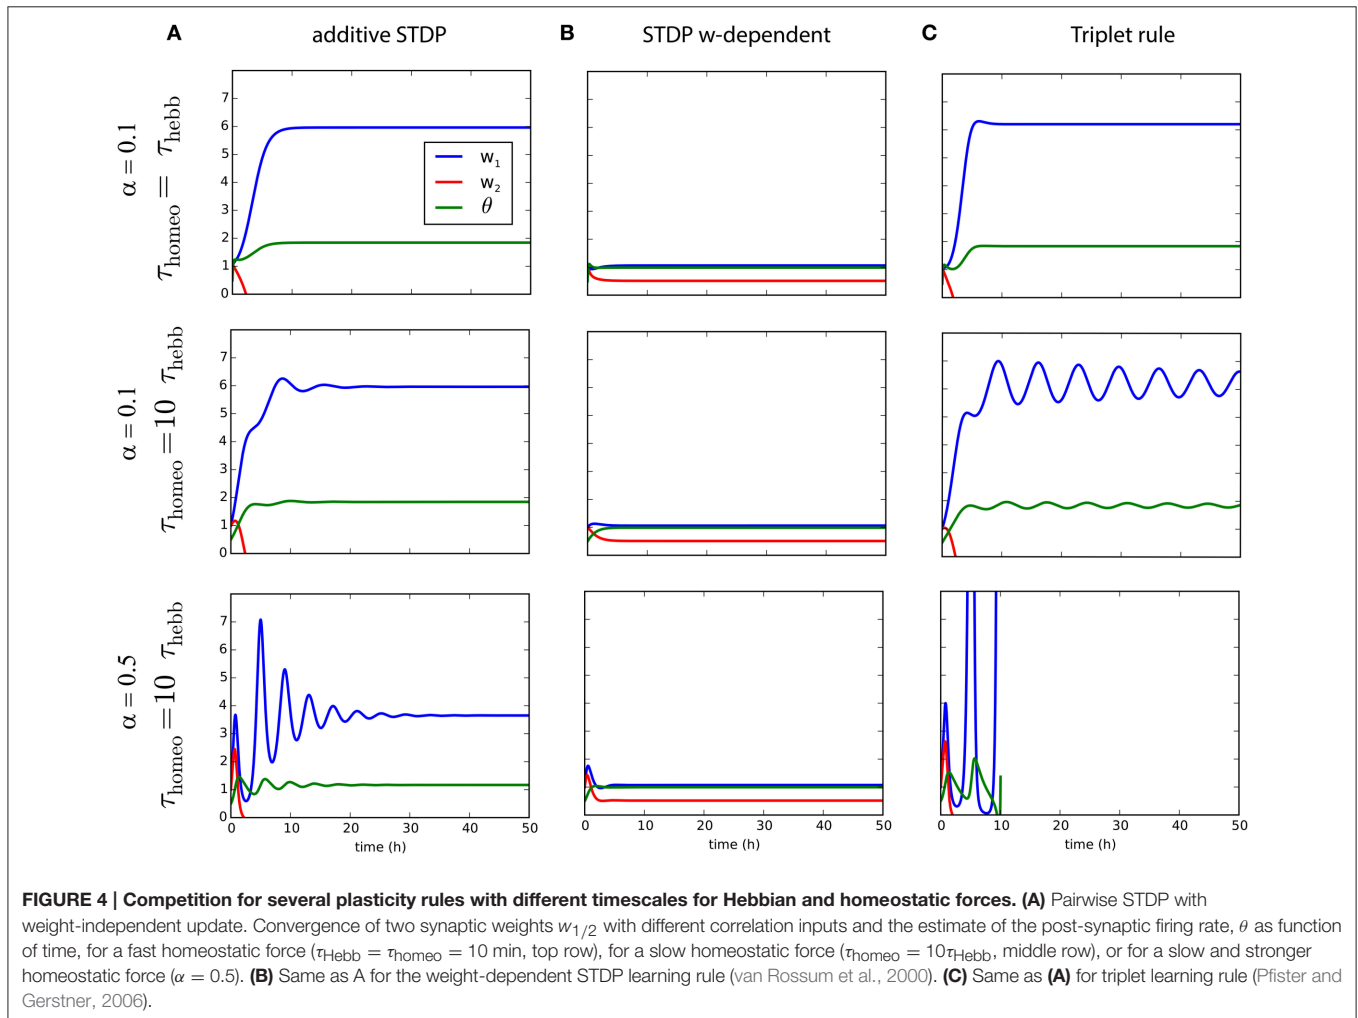
In order to stabilize the triplet STDP rule (Pfister and Gerstner, 2006) in recurrent networks, further studies (Clopath et al., 2010;

Zenke et al., 2013) scaled the amount of LTD in terms of a smoothed average of the firing of the post-synaptic neuron. This modulation of LTD actually brings the triplet STDP rule closer to the BCM rule, by implementing a regulation of the threshold between effective LTP and LTD. In our formalism, the rule used by Zenke et al. (2013) can be implemented for rates as

$$\begin{aligned}\Phi(\bar{w}, r_{\text{post}}, \theta) &= r_{\text{pre}} r_{\text{post}} (A_+ r_{\text{post}} + A_- \theta^2 / r_{\text{target}}) \\ \Psi(r_{\text{post}}) &= r_{\text{post}}\end{aligned}$$

Note that the difference here compared to BCM is that θ tracks r_{post} and not r_{post}^2 , and the limit between depression and potentiation is related to θ^2 . As in Equation (20), we have $A_+ > 0$ and $A_- < 0$.

Figure 6 compares the evolution for this metaplastic triplet STDP rule with classical STDP combined with synaptic scaling: we can clearly see that the resulting dynamics is strongly affected by the ratio between Hebbian and homeostatic time constants in both cases. The trajectories of w and θ in the same phase space as before show several types of instability, from weight (and rate) explosion for slow tracking with large τ_{homeo} (Zenke et al., 2013) to oscillations when $\tau_{\text{homeo}} = \tau_{\text{Hebb}}$. As before, slow tracking yields a gradient with a strong horizontal component,



hence oscillations. The limit cycle in the top panel of **Figure 6B** only happens for some limited range of the parameters, but this illustrates the severe instability issues even for the simple dynamical system considered here.

5.3. Non-linearly Gated STDP Rules

Another direction of research (Senn et al., 2001; El Boustani et al., 2012) introduced non-linearity in the effect of the Hebbian term, by turning it on and off depending on the past pre- and post-synaptic activity of the neuron. Taking a simplified version with a similar mechanism for both LTP and LTD, we consider

$$\begin{aligned}\Phi(\bar{w}, r_{\text{post}}, \theta) &= \|r_{\text{pre}}r_{\text{post}} \\ &\quad - f_+(\theta)\|_+ - \|r_{\text{pre}}r_{\text{post}} - f_-(\theta)\|_+ \quad (21) \\ \Psi(r_{\text{post}}) &= r_{\text{pre}}r_{\text{post}}\end{aligned}$$

where $\|x\|_+$ is a non linear function equal to x if $x > 0$ and 0 otherwise. Now Ψ is such that θ embodies a smoothed average of the pre-post correlations with the time constant τ_{homeo} . When instantaneous correlations are higher than thresholds $f_{\pm}(\theta)$, for potentiation or depression respectively, plasticity effectively occurs. In the general case, f_{\pm} could be any non-linear functions, and do not even need to rely on the same timescales (El Boustani et al., 2012). The simulations in **Figure 6C** correspond the simple case where $f_{\pm}(\theta) = a_{\pm} = \pm 0.4$ are constant. The problem with those non-linearities is that it becomes hard to perform an mathematical analysis of the equilibrium. As with other rules, we observe the same effect

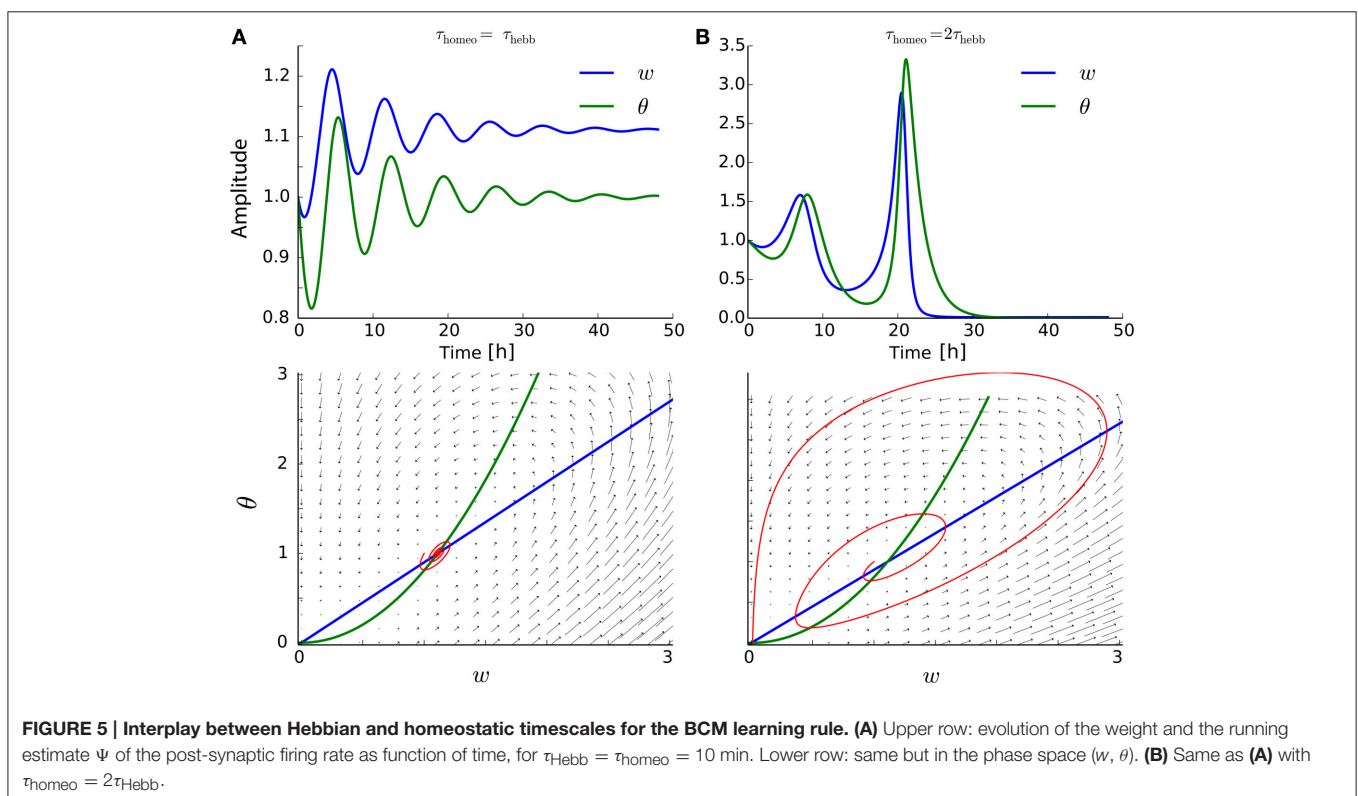
of a large τ_{homeo} on the gradient and the same qualitative conclusion that slow tracking implies the slow convergence of the system.

5.4. Toward More Complex Models

The stability problem arises because, at the equilibrium state, the Hebbian and homeostatic mechanisms compete to balance each other, but they do not act on the same timescale. As pointed out recently (Toyoizumi et al., 2014), a solution can be found when considering that both do not interact linearly, i.e., summing their effects at the synapses, but rather work in a multiplicative manner to determine the synaptic weight. To be more precise, the model developed by Toyoizumi et al. (2014) can be integrated within our framework modulo a slightly more generic formulation for the equation in θ . The model states that $w = \rho H$, where those two quantities are governed by the following system of differential equations.

$$\begin{aligned}\dot{\rho} &= \frac{1}{\tau_{\text{hebb}}} [(\rho_{\text{max}} - \rho) \|r_{\text{pre}}r_{\text{post}} - A_+\|_+ \\ &\quad - (\rho - \rho_{\text{min}}) \|A_- - r_{\text{pre}}r_{\text{post}}\|_+] \\ \dot{H} &= \frac{1}{\tau_{\text{homeo}}} H(1 - r_{\text{post}})\end{aligned}$$

Even if the lower and upper weight bounds ρ_{min} and ρ_{max} depend on H , the model can be written in a generalized version of Equation (2), using $\dot{w} = H\dot{\rho} + \dot{H}\rho$ with simply $\theta = H$. The final expression resembles non-linearly gated plasticity with an additional synaptic scaling, but involves further refinements compared to Equation (21).



6. DISCUSSION AND PERSPECTIVES

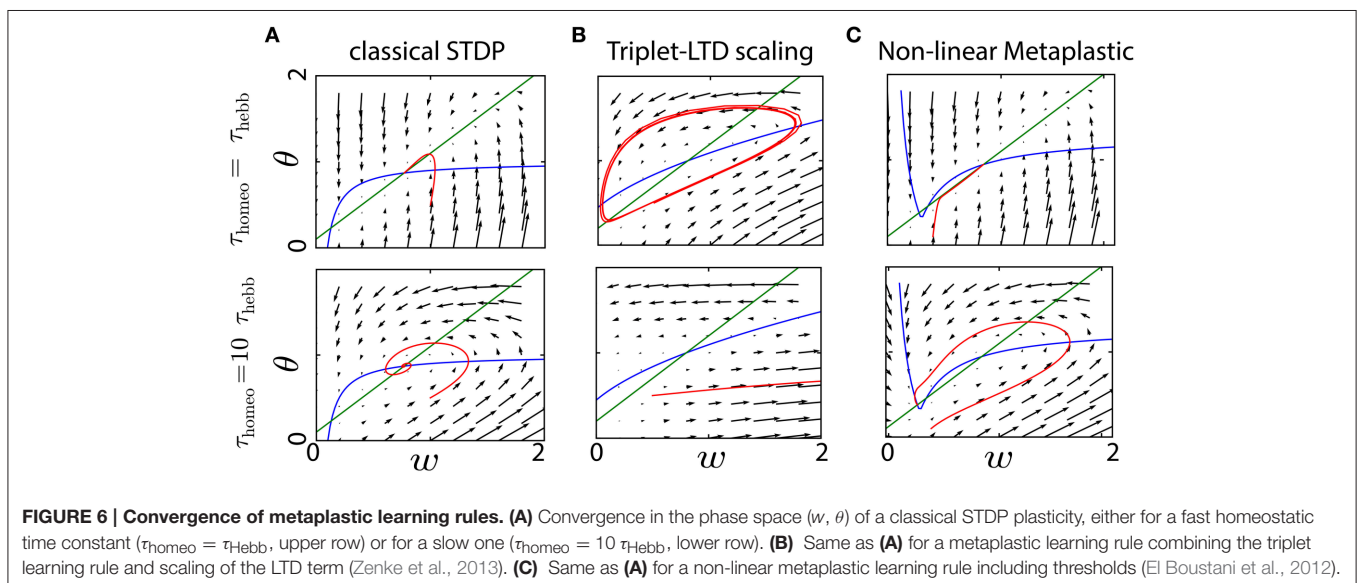
In this paper, we have reviewed various homeostatic mechanisms that are used in recent state-of-the-art plasticity models to regulate Hebbian-type learning. We have focused on two main categories of models: (1) homeostatic synaptic scaling as an independent process that competes with the Hebbian force via an additive term, and (2) metaplastic rules, for which the Hebbian contribution is modulated in an homeostatic fashion. In both cases, the regulation is performed via an estimate of the neural activity (often the post-synaptic firing rate r_{post}) smoothed with a timescale τ_{homeo} , whereas the Hebbian update corresponds to another timescale τ_{Hebb} . We have shown for most models that, when $\tau_{\text{homeo}} \gg \tau_{\text{Hebb}}$, undesired behaviors such as oscillations in the synaptic weights occur, in particular in the case where the homeostatic force is strong. Moreover, competition and stability correspond to conflicting constraints on the parameters, which requires fine-tuning. There is thus a trade-off between the strength of the homeostatic regulation that must compete with the Hebbian drive without perturbing the convergence to a fixed point for the weights. Stability in the weights at a macroscopic level is necessary to ensure stability of the neural functions; note that we have not considered noise in the dynamics of individual weights here, but rather their mean for given pathways.

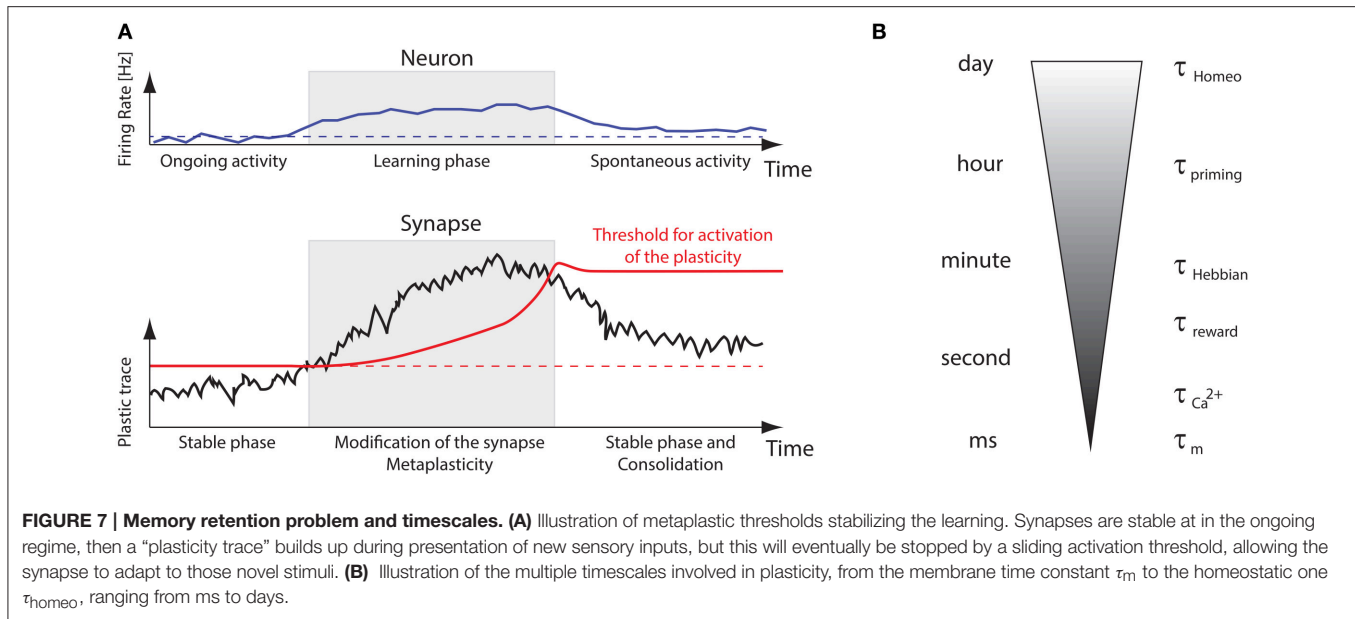
This constraint on the timescales τ_{Hebb} and τ_{homeo} is problematic in regards of available experimental data, as many of them point to slow homeostatic processes (Turrigiano et al., 1998) in comparison with Hebbian processes for which typically $\tau_{\text{Hebb}} \simeq 10$ min. Other models not considered here exhibit similar behavior, for example a homeostatic regulation obtained via intrinsic plasticity (see Zheng et al., 2013) for an example based on spike-threshold adaptation. As a conclusion, the control of the firing rate of the post-synaptic neuron should be taken care of by a mechanisms at a fast timescale, say few minutes at the maximum. Conversely, we point out that homeostatic mechanisms operating on a much slower timescale should be

related to other functions than maintaining the neural activity in a given range.

This claim is supported by several experimental and theoretical findings. Spiking activity of neurons *in vivo* is known to be sparse and highly irregular. Most V1 neurons display Poissonian or supra Poisson spike-count variability in response to low dimensional stimuli such as bars and gratings (Dean, 1981). Even *in vitro*, they fire as Poisson sources, irregularly, with a coefficient of variation for their inter-spike intervals close to 1 (Nawrot et al., 2008). The origin of this irregular activity observed in the sub-threshold voltage and/or in spiking activity is linked to synaptic activity (Paré et al., 1998; Destexhe and Paré, 1999), and because it has been observed experimentally that excitatory and inhibitory conductances are closely balanced (Froemke et al., 2007; Okun and Lampl, 2008), such a fine balance has to be maintained by the system (Renart et al., 2010). Therefore, there is a crucial need for compensatory mechanisms that may interfere or act in concert with Hebbian learning to not only keep the neuron's firing rate within a certain range, but also guarantee this balance (Vogels et al., 2011), or the irregularity of the spiking discharge (Pozzorini et al., 2013). Weight normalization has also been studied in depth in the context of emergence of ocular dominance in order to adjust the competition between synaptic pathways, switching from winner-take-all to winner-share-all behaviors for example (Miller, 1996).

We should discuss several limitations of our study related to the proposed mathematical framework. We have focused on very simple and canonical models of synaptic plasticity, ignoring the fine morphological structure of the neurons. It was shown that the shape of the temporal learning window represented in **Figure 1B** depends on the synaptic position on the dendritic tree (Letzkus et al., 2006; Kampa et al., 2007). More importantly, homeostatic regulation or plasticity thresholds exhibit variability and affect predominantly neighboring synapses *in vivo* (Harvey and Svoboda, 2007). Therefore, we only address the temporal crosstalk between Hebbian and homeostatic plasticity at the





largest scale and the question of defining the spatial extent for heterosynaptic mechanisms remains open. Nevertheless, we expect our conclusions to hold locally for groups of synapses that can be isolated and experience homogeneous processes.

In the general dynamical system in Equations (2) considered here, the timescales are explicitly defined via τ_{Hebb} and τ_{homeo} . In more complex dynamical systems involving noise and attractors, implicit time constants can emerge in a population of synapses (Tetzlaff et al., 2013). Usually, they are slow time constants though, and cannot be used for fast control of the rate, but rather to implement long-lasting memory patterns in the synaptic weights. Another limitation of our conclusions is that we only consider a feedforward model. To extend those to networks with plastic recurrent connections, the mathematical formalism should be modified to account for the case of synapses with the same pre- and post-synaptic firing rates $r_{post} = r_{pre} = r_{rec}$, and likewise the correlations $c_{pre-post} = c_{rec}$. Those quantities follow the consistency equations.

$$r_{rec} = \frac{r_{in}}{1 - w} \quad \text{and} \quad c_{rec} = \frac{c_{in}}{(1 - w)^2}$$

A similar analysis can be done to predict the behavior of learning rules and compare them. The difference compared to the feedforward case is that rates and correlations contributions to the weight updates are not of the same order. This implies that oscillations or other instability effects induced by spike synchrony are more likely to be amplified in recurrent networks than those due to firing rates. It remains that stability can be studied similarly via the Jacobian matrix. Note also that noise in firing and learning dynamics, as well as heterogeneity in neuron and network parameters, may help to prevent “pathological” weight trajectories such as limit cycles, as they smooth the dynamical landscape and degenerate too stereotypical situations.

Beyond those technical details, the puzzling question with plasticity is how synapses can store relevant information while neurons are constantly bombarded by spiking activity *in vivo*. This ongoing input stimulation is quite often considered to be noise in models, which impairs stability of dynamical systems over long time-scales. Although this issue has been addressed theoretically for various models (Clopath et al., 2008; Billings and van Rossum, 2009; Gilson and Fukai, 2011; Tetzlaff et al., 2013; Zenke et al., 2013), it suggests that additional timescales are necessary to properly combine short-term and long-term properties such that the system learns fast and slowly forgets. **Figure 7A** recapitulates several timescales involved in learning and memory. In essence, for the neural system to retain memories, synaptic plasticity should only be turned on by metaplasticity when “new” incoming stimuli impinge neurons. Once this novelty has been learnt, metaplasticity should stop synaptic changes. Then a selection process should trim all newly formed memories to keep only appropriate ones (Frey and Morris, 1997). This is illustrated in **Figure 7B**, where several interleaved timescales interact to bridge all mechanisms, from the effective membrane time constant τ_m (order of ms) that interacts with STDP to the homeostatic time constants τ_{homeo} , which can range from hours to days (Turrigiano et al., 1998; Turrigiano and Nelson, 2004). Calcium signals can act as activity buffers at a timescale $\tau_{Ca^{2+}}$ (Artola et al., 1990; Shouval et al., 2002; Yeung et al., 2004; Graupner and Brunel, 2012), whereas reward signals or neuromodulation would affect plasticity at a larger timescale τ_{reward} (Izhikevich, 2007), comparable to the one observed for Hebbian changes ($\tau_{Hebbian}$). There is also evidence for a control of intrinsic excitability, synaptic scaling at the post-synaptic density, adaptation of the pre-synaptic neurotransmitter release (Davis, 2006). As most models incorporate only a few of those at a time, we stress the need for a better understanding of the complex interactions that may arise when bringing together those experimentally observed mechanisms.

REFERENCES

- Abbott, L. F., and Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nat. Neurosci.* 3, 1178–1183. doi: 10.1038/81453
- Abraham, W. C. (2008). Metaplasticity: tuning synapses and networks for plasticity. *Nat. Rev. Neurosci.* 9, 387. doi: 10.1038/nrn2356
- Abraham, W. C., and Bear, M. F. (1996). Metaplasticity: the plasticity of synaptic plasticity. *Trends Neurosci.* 19, 126–130.
- Appleby, P. A., and Elliott, T. (2005). Synaptic and temporal ensemble interpretation of spike-timing-dependent plasticity. *Neural Comput.* 17, 2316–2336. doi: 10.1162/0899766054796879
- Artola, A., Bröcher, S., and Singer, W. (1990). Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature* 347, 69–72.
- Beck, H., and Yaari, Y. (2008). Plasticity of intrinsic neuronal properties in CNS disorders. *Nat. Rev. Neurosci.* 9, 357–369. doi: 10.1038/nrn2371
- Bender, V. A., Bender, K. J., Brasier, D. J., and Feldman, D. E. (2006). Two coincidence detectors for spike timing-dependent plasticity in somatosensory cortex. *J. Neurosci.* 26, 4166–4177. doi: 10.1523/JNEUROSCI.0176-06.2006
- Bi, G. Q., and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Bienenstock, E. L., Cooper, L. N., and Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J. Neurosci.* 2, 32–48.
- Billings, G., and van Rossum, M. C. W. (2009). Memory retention and spike-timing-dependent plasticity. *J. Neurophysiol.* 101, 2775–2788. doi: 10.1152/jn.91007.2008
- Bliss, T. V. P., and Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232, 331–356.
- Burkitt, A. N., Meffin, H., and Grayden, D. B. (2004). Spike-timing-dependent plasticity: the relationship to rate-based learning for models with weight dynamics determined by a stable fixed point. *Neural Comput.* 16, 885–940. doi: 10.1162/089976604773135041
- Caporale, N., and Dan, Y. (2008). Spike timing-dependent plasticity: a Hebbian learning rule. *Annu. Rev. Neurosci.* 31, 25–46. doi: 10.1146/annurev.neuro.31.060407.125639
- Christie, B. R., and Abraham, W. C. (1992). Priming of associative long-term depression in the dentate gyrus by theta frequency synaptic activity. *Neuron* 9, 79–84.
- Clopath, C., Büsing, L., Vasilaki, E., and Gerstner, W. (2010). Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nat. Neurosci.* 13, 344–352. doi: 10.1038/nn.2479
- Clopath, C., Ziegler, L., Vasilaki, E., Büsing, L., and Gerstner, W. (2008). Tag-trigger-consolidation: a model of early and late long-term-potential and depression. *PLoS Comput. Biol.* 4:e1000248. doi: 10.1371/journal.pcbi.1000248
- Cooper, L. N., and Bear, M. F. (2012). The BCM theory of synapse modification at 30: interaction of theory with experiment. *Nat. Rev. Neurosci.* 13, 798–810. doi: 10.1038/nrn3353
- Cooper, L. N., Intrator, N., Blais, B. S., and Shouval, H. Z. (2004). *Theory of Cortical Plasticity*. New Jersey, NJ: World Scientific.
- Corlew, R., Wang, Y., Ghermazien, H., Erisir, A., and Philpot, B. D. (2007). Developmental switch in the contribution of presynaptic and postsynaptic NMDA receptors to long-term depression. *J. Neurosci.* 27, 9835–9845. doi: 10.1523/jneurosci.5494-06.2007
- Davis, G. W. (2006). Homeostatic control of neural activity: from phenomenology to molecular design. *Annu. Rev. Neurosci.* 29, 307–323. doi: 10.1146/annurev.neuro.28.061604.135751
- Dean, A. F. (1981). The relationship between response amplitude and contrast for cat striate cortical neurones. *J. Physiol.* 318, 413–427.
- Destexhe, A., and Paré, D. (1999). Impact of network activity on the integrative properties of neocortical pyramidal neurons *in vivo*. *J. Neurophysiol.* 81, 1531–1547.
- El Boustani, S., Yger, P., Frégnac, Y., and Destexhe, A. (2012). Stable learning in stochastic network states. *J. Neurosci.* 32, 194–214. doi: 10.1523/JNEUROSCI.2496-11.2012
- Feldman, D. E. (2012). The spike-timing dependence of plasticity. *Neuron* 75, 556–571. doi: 10.1016/j.neuron.2012.08.001
- Frey, U., and Morris, R. G. (1997). Synaptic tagging and long-term potentiation. *Nature* 385, 533–536.
- Froemke, R. C., Merzenich, M. M., and Schreiner, C. E. (2007). A synaptic memory trace for cortical receptive field plasticity. *Nature* 450, 425–429. doi: 10.1038/nature06289
- Froemke, R. C., Poo, M.-M., and Dan, Y. (2005). Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature* 434, 221–225. doi: 10.1038/nature03366
- Froemke, R. C., Tsay, I. A., Raad, M., Long, J. D., and Dan, Y. (2006). Contribution of individual spikes in burst-induced long-term synaptic modification. *J. Neurophysiol.* 95, 1620–1629. doi: 10.1152/jn.00910.2005
- Froemke, R. C. R., and Dan, Y. (2002). Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature* 416, 433–438. doi: 10.1038/416433a
- Gilson, M., Burkitt, A. N., Grayden, D. B., Thomas, D. A., and van Hemmen, J. L. (2009). Emergence of network structure due to spike-timing-dependent plasticity in recurrent neuronal networks. I. Input selectivity–strengthening correlated input pathways. *Biol. Cybern.* 101, 81–102. doi: 10.1007/s00422-009-0319-4
- Gilson, M., and Fukai, T. (2011). Stability versus neuronal specialization for STDP: long-tail weight distributions solve the dilemma. *PLoS ONE* 6:e25339. doi: 10.1371/journal.pone.0025339
- Graupner, M., and Brunel, N. (2012). Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. *Proc. Natl. Acad. Sci. U.S.A.* 109, 21551–21551. doi: 10.1073/pnas.1109359109
- Gütig, R., Aharonov, R., Rotter, S., and Sompolinsky, H. (2003). Learning input correlations through nonlinear temporally asymmetric Hebbian plasticity. *J. Neurosci.* 23, 3697–3714.
- Harvey, C. D., and Svoboda, K. (2007). Locally dynamic synaptic learning rules in pyramidal neuron dendrites. *Nature* 450, 1195–1200. doi: 10.1038/nature06416
- Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory*, Vol. 44. New York, NY: Science Education.
- Hengen, K. B., Lambo, M. E., Van Hooser, S. D., Katz, D. B., and Turrigiano, G. G. (2013). Firing rate homeostasis in visual cortex of freely behaving rodents. *Neuron* 80, 335–342. doi: 10.1016/j.neuron.2013.08.038
- Huang, Y.-Y. Y., Colino, A., Selig, D. K. D., and Malenka, R. C. R. (1992). The influence of prior synaptic activity on the induction of long-term potentiation. *Science* 255, 730–733.
- Izhikevich, E. M. (2007). Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cereb. Cortex* 17, 2443–2452. doi: 10.1126/science.1346729
- Kampa, B. M., Letzkus, J. J., and Stuart, G. J. (2007). Dendritic mechanisms controlling spike-timing-dependent synaptic plasticity. *Trends Neurosci.* 30, 456–463. doi: 10.1016/j.tins.2007.06.010
- Keck, T., Keller, G. B., Jacobsen, R. I., Eysel, U. T., Bonhoeffer, T., and Hübener, M. (2013). Synaptic scaling and homeostatic plasticity in the mouse visual cortex *in vivo*. *Neuron* 80, 327–334. doi: 10.1016/j.neuron.2013.08.018
- Kempter, R., Gerstner, W., Van Hemmen, J. L., and Hemmen, J. V. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E* 59, 4498–4514.
- Letzkus, J. J., Kampa, B. M., and Stuart, G. J. (2006). Learning rules for spike timing-dependent plasticity depend on dendritic synapse location. *J. Neurosci.* 26, 10420–10429. doi: 10.1523/JNEUROSCI.2650-06.2006
- Marder, E., and Goaillard, J.-M. (2006). Variability, compensation and homeostasis in neuron and network function. *Nat. Rev. Neurosci.* 7, 563–574. doi: 10.1038/nrn1949
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–215.
- Miller, K. (1996). Synaptic economics: competition and cooperation in synaptic plasticity. *Neuron* 17, 371–374.
- Mockett, B., Coussens, C., and Abraham, W. C. (2002). NMDA receptor-mediated metaplasticity during the induction of long-term depression by low-frequency stimulation. *Eur. J. Neurosci.* 15, 1819–1826. doi: 10.1046/j.1460-9568.2002.02008.x

- Morrison, A., Aertsen, A., and Diesmann, M. (2007). Spike-timing-dependent plasticity in balanced random networks. *Neural Comput.* 19, 1437–1467. doi: 10.1162/neco.2007.19.6.1437
- Nabavi, S., Fox, R., Proulx, C. D., Lin, J. Y., Tsien, R. Y., and Malinow, R. (2014). Engineering a memory with LTD and LTP. *Nature* 511, 348–352. doi: 10.1038/nature13294
- Nawrot, M. P., Boucsein, C., Rodriguez Molina, V., Riehle, A., Aertsen, A., and Rotter, S. (2008). Measurement of variability dynamics in cortical spike trains. *J. Neurosci. Methods* 169, 374–390. doi: 10.1016/j.jneumeth.2007.10.013
- Nevian, T., and Sakmann, B. (2006). Spine Ca²⁺ signaling in spike-timing-dependent plasticity. *J. Neurosci.* 26, 11001–11013. doi: 10.1523/JNEUROSCI.1749-06.2006
- Oja, E. (1982). Simplified neuron model as a principal component analyzer. *J. Math. Biol.* 15, 267–273.
- Okun, M., and Lampl, I. (2008). Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat. Neurosci.* 11, 535–537. doi: 10.1038/nn.2105
- O’Leary, T., and Wyllie, D. J. (2011). Neuronal homeostasis: time for a change? *J. Physiol.* 589, 4811–4826. doi: 10.1113/jphysiol.2011.210179
- Paré, D., Shink, E., Gaudreau, H., Destexhe, A., and Lang, E. J. (1998). Impact of spontaneous synaptic activity on the resting properties of cat neocortical pyramidal neurons *in vivo*. *J. Neurophysiol.* 79, 1450–60.
- Pfister, J.-P., and Gerstner, W. (2006). Triplets of spikes in a model of spike timing-dependent plasticity. *J. Neurosci.* 26, 9673–9682. doi: 10.1523/JNEUROSCI.1425-06.2006
- Pozo, K., and Goda, Y. (2010). Unraveling mechanisms of homeostatic synaptic plasticity. *Neuron*, 66, 337–351. doi: 10.1016/j.neuron.2010.04.028
- Pozzorini, C., Naud, R., Mensi, S., and Gerstner, W. (2013). Temporal whitening by power-law adaptation in neocortical neurons. *Nat. Neurosci.* 16, 942–948. doi: 10.1038/nn.3431
- Renart, A., de la Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., and Harris, K. D. (2010). The asynchronous state in cortical circuits. *Science* 327, 587–590. doi: 10.1126/science.1179850
- Rochester, N., Holland, J., Haibt, L., and Duda, W. (1956). Tests on a cell assembly theory of the action of the brain, using a large digital computer. *IRE Trans. Inf. Theory* 2, 80–93. doi: 10.1109/TIT.1956.1056810
- Senn, W., Markram, H., and Tsodyks, M. (2001). An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Comput.* 13, 35–67. doi: 10.1162/089976601300014628
- Shouval, H. Z., Bear, M. F., and Cooper, L. N. (2002). A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. *Proc. Natl. Acad. Sci. U.S.A.* 99, 10831–10836. doi: 10.1073/pnas.152343099
- Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32, 1149–1164. doi: 10.1016/S0896-6273(01)00542-6
- Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2003). Neocortical LTD via coincident activation of presynaptic NMDA and cannabinoid receptors. *Neuron* 39, 641–654. doi: 10.1016/S0896-6273(03)00476-8
- Song, S., and Abbott, L. F. (2001). Cortical development and remapping through spike timing-dependent plasticity. *Neuron* 32, 339–350. doi: 10.1016/S0896-6273(01)00451-2
- Tetzlaff, C., Kolodziejewski, C., Timme, M., Tsodyks, M., and Wörgötter, F. (2013). Synaptic scaling enables dynamically distinct short- and long-term memory formation. *PLoS Comput. Biol.* 9:e1003307. doi: 10.1371/journal.pcbi.1003307
- Tetzlaff, C., Kolodziejewski, C., Timme, M., and Wörgötter, F. (2011). Synaptic scaling in combination with many generic plasticity mechanisms stabilizes circuit connectivity. *Front. Comput. Neurosci.* 5:47. doi: 10.3389/fncom.2011.00047
- Toyozumi, T., Kaneko, M., Stryker, M. P., and Miller, K. D. (2014). Modeling the dynamic interaction of Hebbian and homeostatic plasticity. *Neuron* 84, 497–510. doi: 10.1016/j.neuron.2014.09.036
- Turrigiano, G. G. (2008). The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell* 135, 422–435. doi: 10.1016/j.cell.2008.10.008
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., and Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391, 892–896.
- Turrigiano, G. G., and Nelson, S. B. (2000). Hebb and homeostasis in neuronal plasticity. *Curr. Opin. Neurobiol.* 10, 358–364. doi: 10.1016/S0959-4388(00)00091-X
- Turrigiano, G. G., and Nelson, S. B. (2004). Homeostatic plasticity in the developing nervous system. *Nat. Rev. Neurosci.* 5, 97–107. doi: 10.1038/nrn1327
- van Rossum, M. C., Bi, G. Q., and Turrigiano, G. G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neurosci.* 20, 8812–8821.
- Vitureira, N., and Goda, Y. (2013). The interplay between Hebbian and homeostatic synaptic plasticity. *J. Cell Biol.* 203, 175–186. doi: 10.1083/jcb.201306030
- Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* 334, 1569–1573. doi: 10.1126/science.1211095
- von der Malsburg, C. (1973). Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik* 14, 85–100.
- Wang, H.-X., Gerkin, R. C., Nauen, D. W., and Bi, G.-Q. (2005). Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nat. Neurosci.* 8, 187–193. doi: 10.1038/nn1387
- Wang, L., Fontanini, A., and Maffei, A. (2012). Experience-dependent switch in sign and mechanisms for plasticity in layer 4 of primary visual cortex. *J. Neurosci.* 32, 10562–10573. doi: 10.1523/JNEUROSCI.0622-12.2012
- Wang, X. J. (1998). Calcium coding and adaptive temporal computation in cortical pyramidal neurons. *J. Neurophysiol.* 79, 1549–1566.
- Watt, A. J., and Desai, N. S. (2010). Homeostatic plasticity and STDP: keeping a neuron’s cool in a fluctuating world. *Front. Synaptic Neurosci.* 2:5. doi: 10.3389/fnsyn.2010.00005
- Yeung, L. C., Shouval, H. Z., Blais, B. S., and Cooper, L. N. (2004). Synaptic homeostasis and input selectivity follow from a calcium-dependent plasticity model. *Proc. Natl. Acad. Sci. U.S.A.* 101, 14943–14948. doi: 10.1073/pnas.0405551101
- Yger, P., and Harris, K. D. (2013). The Convallis rule for unsupervised learning in cortical networks. *PLoS Comput. Biol.* 9:e1003272. doi: 10.1371/journal.pcbi.1003272
- Young, J. M., Waleszczyk, W. J., Wang, C., Calford, M. B., Dreher, B., and Obermayer, K. (2007). Cortical reorganization consistent with spike timing- but not correlation-dependent plasticity. *Nat. Neurosci.* 10, 887–895. doi: 10.1038/nn1913
- Zenke, F., Hennequin, G., and Gerstner, W. (2013). Synaptic plasticity in neural networks needs homeostasis with a fast rate detector. *PLoS Comput. Biol.* 9:e1003330. doi: 10.1371/journal.pcbi.1003330
- Zhang, W., and Linden, D. J. (2003). The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nat. Rev. Neurosci.* 4, 885–900. doi: 10.1038/nrn1248
- Zheng, P., Dimitrakakis, C., and Triesch, J. (2013). Network self-organization explains the statistics and dynamics of synaptic connection strengths in cortex. *PLoS Comput. Biol.* 9:e1002848. doi: 10.1371/journal.pcbi.1002848

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2015 Yger and Gilson. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.



Addendum: Models of Metaplasticity: A Review of Concepts

Pierre Yger^{1,2*} and Matthieu Gilson³

¹ Sorbonne Université, Université Pierre et Marie Curie Univ Paris06 UMR5968, Paris, France, ² Institut de la Vision, Institut National de la Santé et de la Recherche Médicale, U968, Centre National de la Recherche Scientifique, UMR7210, Paris, France, ³ Computational Neuroscience Group, DTIC, Universitat Pompeu Fabra, Barcelona, Spain

Keywords: synaptic plasticity, metaplasticity, Hebbian learning, homeostasis, STDP

An addendum on

Models of Metaplasticity: A Review of Concepts

by Yger, P., and Gilson, M. (2015). *Front. Comput. Neurosci.* 9:138. doi: 10.3389/fncom.2015.00138

OPEN ACCESS

Edited and reviewed by:

Friedemann Zenke,
Stanford University, USA

*Correspondence:

Pierre Yger
pierre.yger@inserm.fr

Received: 26 August 2015

Accepted: 18 December 2015

Published: 28 January 2016

Citation:

Yger P and Gilson M (2016)
Addendum: Models of Metaplasticity:
A Review of Concepts.
Front. Comput. Neurosci. 10:4.
doi: 10.3389/fncom.2016.00004

AUTHOR CONTRIBUTIONS

All authors listed, have made substantial, direct and intellectual contribution to the work, and approved it for publication.

ACKNOWLEDGMENTS

MG acknowledges funding from FP7 FET ICT Flagship Human Brain Project (604102).

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2016 Yger and Gilson. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Homeostatic role of heterosynaptic plasticity: models and experiments

Marina Chistiakova¹, Nicholas M. Bannon¹, Jen-Yung Chen², Maxim Bazhenov² and Maxim Volgushev^{1*}

¹ Department of Psychology, University of Connecticut, Storrs, CT, USA, ² Department of Cell Biology and Neuroscience, University of California, Riverside, Riverside, CA, USA

OPEN ACCESS

Edited by:

Friedemann Zenke,
Stanford University, USA

Reviewed by:

Paul Miller,
Brandeis University, USA
Rui Ponte Costa,
University of Oxford, UK

*Correspondence:

Maxim Volgushev,
Department of Psychology, University
of Connecticut, 406 Babbidge Road
Unit 1020, Storrs, CT 06269-1020,
USA
maxim.volgushev@uconn.edu

Received: 19 January 2015

Accepted: 25 June 2015

Published: 13 July 2015

Citation:

Chistiakova M, Bannon NM, Chen J-Y,
Bazhenov M and Volgushev M (2015)
Homeostatic role of heterosynaptic
plasticity: models and experiments.
Front. Comput. Neurosci. 9:89.
doi: 10.3389/fncom.2015.00089

Homosynaptic Hebbian-type plasticity provides a cellular mechanism of learning and refinement of connectivity during development in a variety of biological systems. In this review we argue that a complimentary form of plasticity—heterosynaptic plasticity—represents a necessary cellular component for homeostatic regulation of synaptic weights and neuronal activity. The required properties of a homeostatic mechanism which acutely constrains the runaway dynamics imposed by Hebbian associative plasticity have been well-articulated by theoretical and modeling studies. Such mechanism(s) should robustly support the stability of operation of neuronal networks and synaptic competition, include changes at non-active synapses, and operate on a similar time scale to Hebbian-type plasticity. The experimentally observed properties of heterosynaptic plasticity have introduced it as a strong candidate to fulfill this homeostatic role. Subsequent modeling studies which incorporate heterosynaptic plasticity into model neurons with Hebbian synapses (utilizing an STDP learning rule) have confirmed its ability to robustly provide stability and competition. In contrast, properties of homeostatic synaptic scaling, which is triggered by extreme and long lasting (hours and days) changes of neuronal activity, do not fit two crucial requirements for a hypothetical homeostatic mechanism needed to provide stability of operation in the face of on-going synaptic changes driven by Hebbian-type learning rules. Both the trigger and the time scale of homeostatic synaptic scaling are fundamentally different from those of the Hebbian-type plasticity. We conclude that heterosynaptic plasticity, which is triggered by the same episodes of strong postsynaptic activity and operates on the same time scale as Hebbian-type associative plasticity, is ideally suited to serve a homeostatic role during on-going synaptic plasticity.

Keywords: synaptic plasticity, homosynaptic plasticity, STDP, Hebbian plasticity, homeostasis, heterosynaptic plasticity, runaway dynamics

Introduction: Three Forms of Synaptic Plasticity

Normal operation of the brain requires the maintenance of balances across various neuronal and synaptic features, and keeping key factors within their operating ranges. This is achieved by a variety of homeostatic mechanisms operating at multiple nested levels, such as maintenance of excitation/inhibition balance and total activity at the network level, or homeostasis of synaptic weights at the single-cell level (van Vreeswijk and Sompolinsky, 1996; Haydon, 2001; Burrone and Murthy, 2003; Turrigiano, 2011; Davis, 2013). Synaptic weights are subject to changes caused

by a variety of mechanisms mediating multiple forms of plasticity (Bliss and Collingridge, 1993; Abbott and Nelson, 2000; Malenka and Bear, 2004; Sjostrom et al., 2007; Chistiakova and Volgushev, 2009; Feldman, 2009). In the present review we will consider the relation between different forms of plasticity and synaptic homeostasis. We ask: which forms of plasticity bring synaptic weights out of balance and which forms may serve to restore their balance?

The multitude of forms of synaptic plasticity can be segregated into three broad types, distinguished by the differential activity patterns required for their induction, distinct functions served in learning systems, and diverse computational roles. The first, *homosynaptic plasticity*, requires presynaptic activation of the synapse for the induction. By definition, it occurs only at synapses that were directly involved in activation of a cell during the induction, for example during afferent tetanization or pairing procedure (**Figure 1A**). This form of plasticity is also called input specific, and, if the induction follows Hebbian-type learning rules, associative (Bliss and Collingridge, 1993). For induction of associative plasticity correlated activity of pre and postsynaptic neurons is crucial. Associative plasticity underlies a multitude of phenomena in the nervous system, ranging from refinement of connectivity during development (“neurons that fire together wire together”) to extraction of causal relations between events in the environment in Pavlovian conditioning and other types of associative learning as well as motor learning. The second form, *heterosynaptic plasticity*, can be induced at synapses that were not active during the induction of homosynaptic plasticity. Thus, it is not limited to active synapses, but may also change non-active synapses after episodes of strong postsynaptic activity. As the majority of synapses onto a cell are not presynaptically activated during a typical induction protocol, heterosynaptic plasticity typically affects a larger population of synapses than homosynaptic plasticity does. Heterosynaptic plasticity can be induced at unstimulated synapses by typical induction protocols, such as afferent tetanization or a pairing procedure (**Figure 1A**), but also by purely postsynaptic protocols such as intracellular tetanization: bursts of spikes evoked by depolarizing pulses (**Figure 1B**). Homosynaptic and heterosynaptic plasticity have complementary computational properties, making both forms necessary for normal operation of neural systems with plastic synapses (Chistiakova and Volgushev, 2009; Chen et al., 2013; Chistiakova et al., 2014).

The third form of plasticity considered in this review is *homeostatic synaptic scaling*, which is induced by prolonged (hours/days) and dramatic changes of activity, and leads to compensatory scaling of synaptic weights (Turrigiano et al., 1998, recently reviewed in Burrone and Murthy, 2003; Turrigiano and Nelson, 2004; Rich and Wenner, 2007; Rabinowitch and Segev, 2008; Watt and Desai, 2010; Turrigiano, 2012). Prolonged silencing of neurons by tetrodotoxin (TTX) leads to up-scaling of synaptic weights, while prolonged elevation of activity leads to down-scaling of the synapses (**Figure 1C**). Homeostatic synaptic scaling is triggered by deviation of firing rate from a target level (Turrigiano et al., 1998; van Rossum et al., 2000) and can be considered as a mechanism of firing rate homeostasis. Computationally, homeostatic synaptic scaling may have the

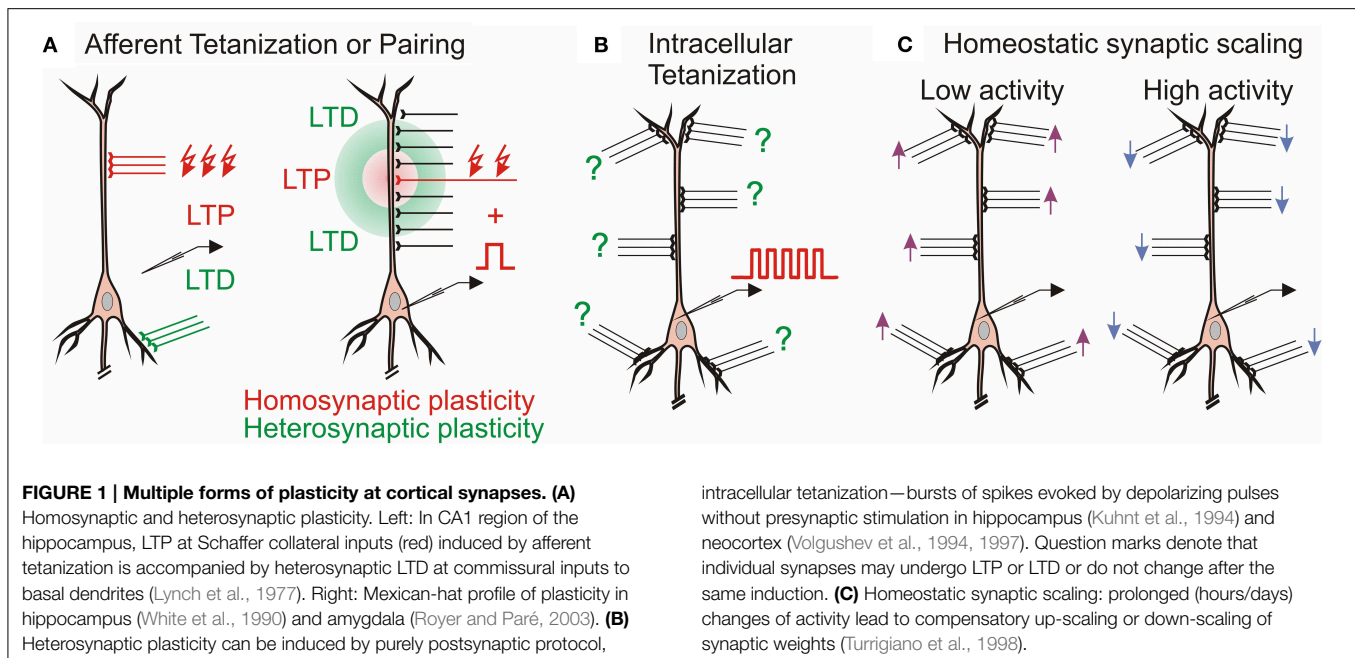
effect of a “delayed” normalization if changes of postsynaptic firing were due to changes at a subset of inputs, or scale all synapses proportionally if postsynaptic firing rates were changed because of uniform silencing (or uniform up-regulation of activity) at all inputs. Because homeostatic synaptic scaling is triggered by overall activity level, irrespective of which specific synapses contributed to the activation of a neuron, but changes the weights of all synapses of a cell proportionally, it may include changes of both homosynaptic (those which were active) and heterosynaptic (those which were not active) inputs.

In this review we will consider: (i) Why do learning systems need forms of plasticity, additional to homosynaptic plastic changes governed by associative, Hebbian-type learning rules? (ii) What are the requirements for these additional forms of plasticity, outlined in theoretical and modeling studies? (iii) Which biological candidate mechanisms express properties that fulfill these requirements?

Homosynaptic Plasticity: Why It Cannot Work Alone

Synaptic plasticity induced according to Hebbian-type rules mediates the formation and refinement of connectivity patterns in the nervous system during development, and underlies various types of associative learning throughout life. However, for at least two reasons, associative synaptic plasticity governed by Hebbian-type rules cannot be the only type of plasticity in a learning system. First, Hebbian-type learning rules impose an intrinsic positive feedback on synaptic weight changes, thus making the system prone to runaway dynamics of synaptic weights and activity. Second, they introduce only a weak degree of competition between synapses. Competition is indispensable for building sensory representations during development, and is instrumental in learning that involves differentiation.

The propensity for runaway dynamics in a system with Hebbian-type synapses originates from the positive feedback on synaptic weight changes intrinsic to this type of plasticity rule. Potentiation, by making synapses stronger, increases the chance that these synapses will contribute to the firing of a neuron, and will be potentiated further. Similarly, depression of synapses decreases the chance that these synapses will contribute to the firing of a neuron, and thus decreases their chances for subsequent potentiation, but increases the probability that they undergo further depression. Indeed, neuron models with Hebbian-type rules for synaptic plasticity expresses runaway dynamics of synaptic weights. In a computational model of a single neuron with symmetrical windows for potentiation and depression within STDP rules, and receiving inputs from neurons with Poisson-distributed spikes intervals, synaptic weights get potentiated to the maximal value (**Figure 2**). In the same model but with STDP rules biased toward depression, synapses tend to be depressed, with the weights of many of them declining to zero (**Figure 3**). As a result, with an unchanged level of presynaptic activity the runaway of synaptic weights causes runaway dynamics of postsynaptic spiking. Runaway potentiation of synaptic weights to maximal values leads to



overexcitability of neurons and excessive postsynaptic activity. In the **Figure 2** model, firing rate of the postsynaptic model neuron increased from 1.8 Hz during the first 10 s, to 6.3 Hz during the last 10 s of simulation. Runaway depression leads to a decrease of postsynaptic activity, and eventually to the silencing of neurons. In the example in **Figure 3**, the postsynaptic neuron becomes essentially silent by the end of simulation despite a three-fold increase of the firing rate of presynaptic neurons. Both runaway potentiation and runaway depression scenarios lead to the disturbance of input-output relations of neurons (Chen et al., 2013), thus compromising computational abilities of neuronal networks (Skorheim et al., 2014). Moreover, runaway potentiation of synaptic weights and associated runaway activity are energetically unsustainable and may lead to pathology. To counteract runaway dynamics of synaptic weights and activity imposed by Hebbian-type plasticity rules, additional plasticity mechanisms are necessary. Such mechanisms would keep synaptic weights and neuronal networks with plastic synapses within their operational range, and thus maintain homeostasis of synaptic weights and neuronal activity despite unbalancing perturbations introduced by associative plasticity.

The idea of competition, which maintains that while undergoing changes, synapses might compete for shared limited resources such as available energy, molecules or plasticity factors, is well-grounded from the biological perspective (e.g., Miller, 1996; Peters et al., 2004). Synaptic competition is absolutely necessary for the development of sensory representations (Wiesel and Hubel, 1963; Aitkin et al., 1970; Merzenich et al., 1975; Thompson et al., 1983; Feldman, 2009), and is instrumental in a broad class of learning tasks that involve discrimination (Skorheim et al., 2014). However, Hebbian-type learning rules introduce only a weak, if any, degree of competition (Miller, 1996), restricted to the synapses receiving distinct input patterns (Zhang et al., 2011; see below for further discussion). Therefore,

to support intrinsic competition between synapses, a mechanism for plastic changes outside of the Hebbian learning rule is required.

To summarize, homosynaptic plasticity, while being a major driving force for synaptic changes mediating associative learning, imposes positive feedback on synaptic changes which creates energetically and computationally unstable runaway dynamics. It also does not provide the required degree of synaptic competition known to be necessary for many types of learning. Modeling studies show the need for additional mechanisms which could constrain and balance Hebbian plasticity. Biological neuronal systems do possess such mechanisms, as evidenced by stable operation and diverse dynamics of synaptic changes over a broad range of conditions. It has not been clear, however, which of the multitude of proposed physiological mechanisms is/are able to serve these roles.

Modeling Perspective

The Need for Homeostasis

The need for maintaining homeostasis of key variables of neuronal operation, such as synaptic weights, the amount of synaptic drive received by a neuron and neuronal activity, as well as for strong synaptic competition in learning systems with Hebbian-type rules has been appreciated since early theoretical studies (von der Malsburg, 1973; Miller and MacKay, 1994; Miller, 1996). To some extent, this can be achieved by careful adjustment of STDP rules. With an appropriate negative bias between the windows for potentiation and depression, STDP can indeed lead to stabilization of the neurons' mean firing rate (Song et al., 2000; Kempter et al., 2001; Gütiğ et al., 2003; Babadi and Abbott, 2010). For certain ranges of presynaptic firing rates, fine-tuned STDP can also support synaptic competition, by driving synaptic weights to the maximal value or to zero (e.g.,

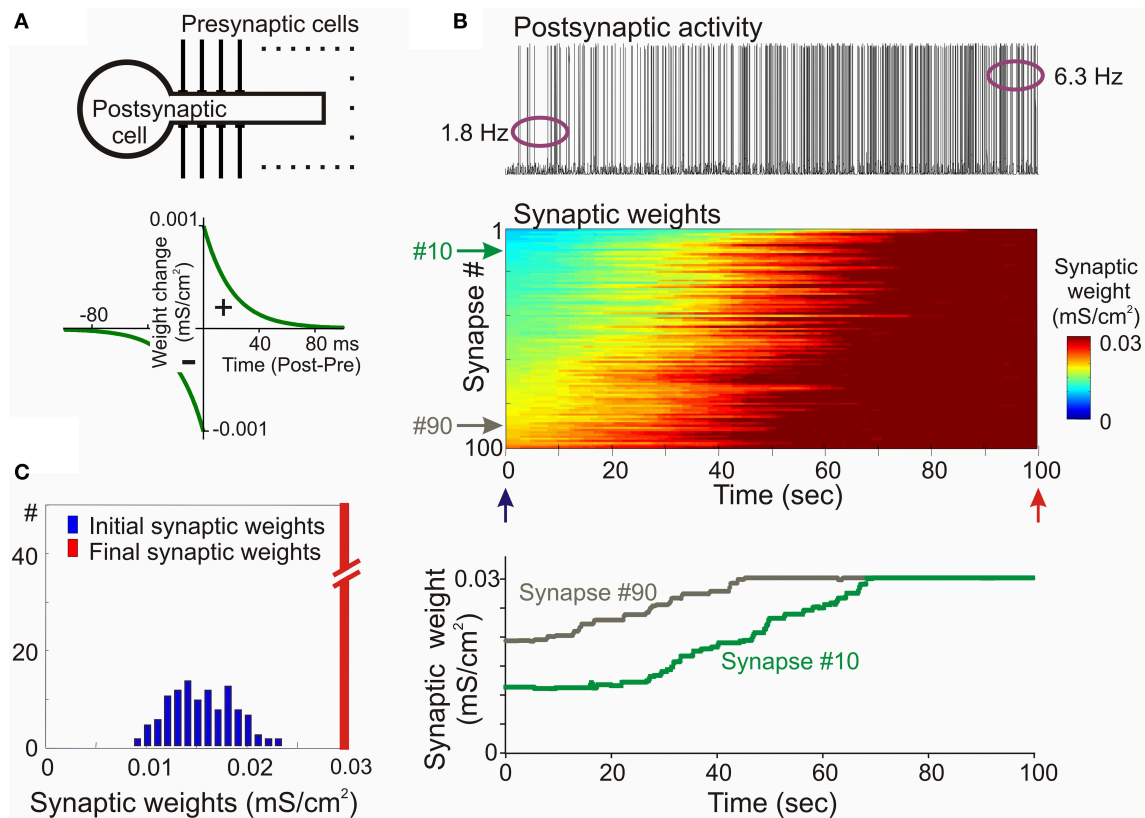


FIGURE 2 | Synaptic activity produced by weakly correlated inputs leads to runaway dynamics of synaptic weights in a model with symmetrical STDP mechanism. (A) A scheme of a model neuron and STDP learning rule. The model neuron consisted of axosomatic and dendritic compartments, receiving 100 synaptic inputs from 100 presynaptic neurons. Each presynaptic neuron fired action potentials at ~ 1 Hz, with Poisson distributed interspike intervals. In simulations shown in this figure, firing of input neurons was mildly correlated (averaged cross-correlation 0.348 ± 0.05). The STDP learning rule had symmetrical potentiation and depression windows ($\tau^+ = \tau^- = 20$ ms; $a^+ = a^- = 10^{-3}$ mS/cm²). **(B)** Membrane potential trace of the model neuron (top); changes of the weights

of 100 synapses (middle), color coded, with synapses sorted by their synaptic weights at the beginning of experiment; and changes of the weights of synapses #10 and #90 (bottom). In this model with symmetrical STDP learning rule synaptic inputs expressed runaway dynamics, and all inputs were potentiated to the maximum by the end of the simulation. **(C)** Distributions of synaptic weights at the beginning (blue, at 20 ms) and at the end (red, at 100 s) of simulation experiment shown in **(B)**. Note runaway dynamics of synaptic weights leading to their saturation at the extreme value (0.03 mS/cm²; red bar length is out of scale) and associated increase of the firing rate of postsynaptic neuron from 1.8 to 6.3 Hz. (Modified, with permission from Chen et al., 2013).

Song et al., 2000; van Rossum et al., 2000; Gütiğ et al., 2003; Morrison et al., 2008). In all of these scenarios, the steady state distribution of synaptic weights depended on the fine-tuning of model parameters, such as the firing rate, the overall balance between excitatory and inhibitory inputs, and the fine details of plasticity rules (e.g., temporal shift or temporal jitter of plasticity windows, or weight-dependence of synaptic changes Song et al., 2000; Gütiğ et al., 2003; Morrison et al., 2008; Babadi and Abbott, 2010).

A problem with this solution is that it works only for certain combinations of internal features (STDP rules) and external events, input activity pattern and correlations. For STDP rules, it imposes strict requirements on the relative strength (amplitude and duration) of potentiation and depression windows which are compatible with the stable operation of a neuron (see below, **Figure 7A** and related text). Experimental evidence, however, has demonstrated an enormous heterogeneity in the

width and magnitude of STDP windows in different synapses, cells, developmental stages, and conditions of neuromodulation (Nishiyama et al., 2000; Sjöström et al., 2001; Zhou et al., 2005; Haas et al., 2006; Feldman, 2009). Although depression-biased STDP rules had been reported for some synapses, a general requirement for a negative integral of STDP rules is not compatible with the experimentally observed variety of learning rules. Moreover, STDP rules adjusted to maintain stability in a neuron subject to a certain pattern of external drive may still lead to runaway dynamics when activity changes, for example if the level of correlation of activity at a subset of the inputs were increased (Gütiğ et al., 2003).

Positive feedback on synaptic changes can be counteracted by weight-dependence of plastic changes—a mechanism suggested theoretically (Oja, 1982) and confirmed experimentally (Bi and Poo, 1998; van Rossum et al., 2000; Hardingham et al., 2007). Weight dependence of plastic changes dictates that weaker

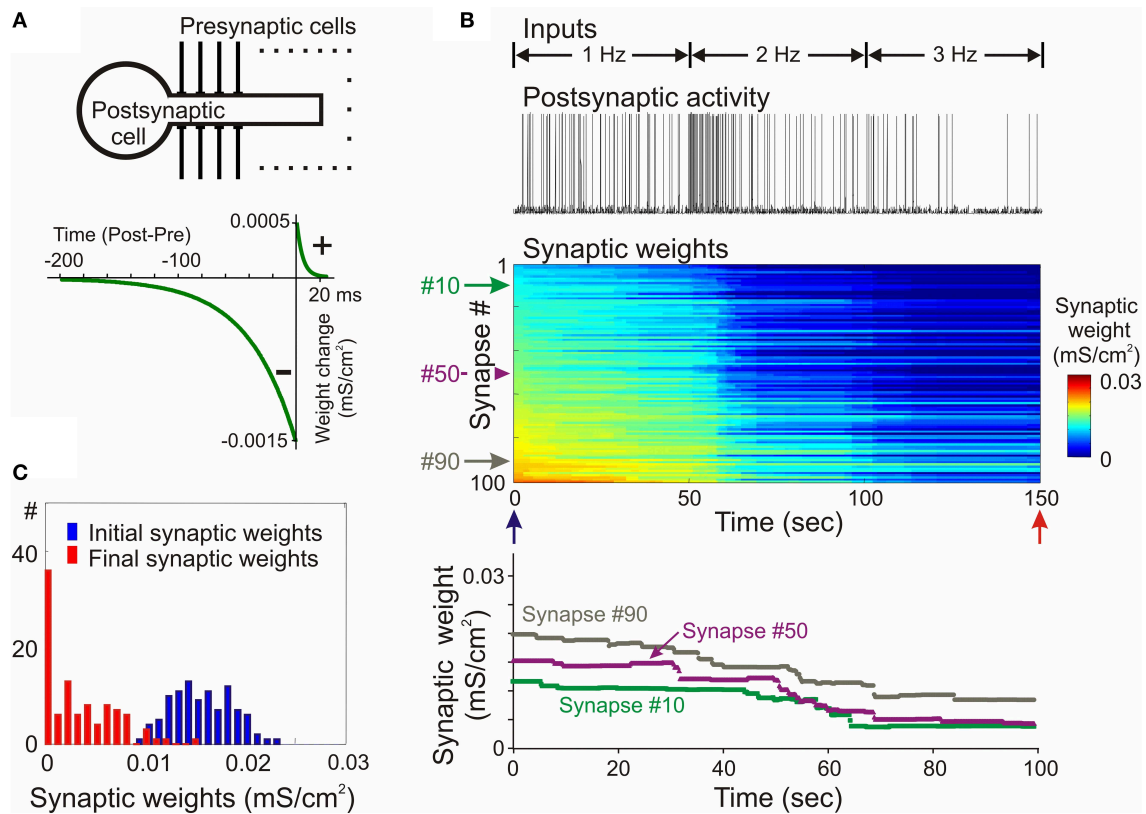


FIGURE 3 | Synaptic activity produced by weakly correlated inputs leads to runaway dynamics of synaptic weights in a model with negatively biased STDP mechanism. (A) A scheme of the model neuron and STDP learning rule. The STDP learning rule with negative bias ($\tau^+ = 5$ ms, $a^+ = 0.5 \times 10^{-3}$ mS/cm², $\tau^- = 40$ ms, $a^- = 1.5 \times 10^{-3}$ mS/cm²). **(B)** Membrane potential trace of the model neuron (top) receiving input from presynaptic neurons firing at an average rate of 1 Hz during first 50 s of simulation, 2 Hz during 50–100 s and 3 Hz during 100–150 s, as indicated; changes of the weights of 100 synapses (middle), color coded, with synapses sorted

by their synaptic weights at the beginning of experiment; and changes of the weights of synapses #10, #50, and #90 (bottom). In this model with negatively biased STDP learning rule synaptic inputs expressed runaway dynamics toward the minimum value. **(C)** Distributions of synaptic weights at the beginning (blue, at 20 ms) and at the end (red, at 150 s) of simulation experiment shown in **(B)**. Note runaway dynamics of synaptic weights leading to saturation at zero of about 40% of synapses, and associated dramatic decrease of postsynaptic firing rate despite a 3-fold increase of presynaptic firing. (Modified, with permission from Chen et al., 2013).

synapses can potentiate more, while stronger synapses express less potentiation, and in the limit do not change or even depress (Hardingham et al., 2007). Depending on the details of implementation, weight dependence imposes bounds on synaptic weights without preventing their saturation at extreme values or grouping around the values over which an amplitude increase turns into a decrease. Results of theoretical analyses and computer simulations show that stability of the activity level of model neurons can indeed be achieved by implementing weight-dependence of plastic changes (van Rossum et al., 2000; Gütiç et al., 2003). Weight-dependence is often used in combination with depression-biased STDP rules, which further improves the stability of the system (Kempster et al., 2001; Gütiç et al., 2003; Morrison et al., 2008; Gilson and Fukai, 2011).

Normalization as Mechanism of Stability

A simple and robust method of stabilization of the synaptic drive of neurons and of neuronal activity is normalization.

Normalization has been commonly used in modeling and theoretical analyses of neurons and networks with plastic synapses since early studies (von der Malsburg, 1973; Oja, 1982). The concept of normalization is that following an induction of plasticity in a subset of synapses on a cell, the weights of all synapses on that cell are readjusted so that their sum (or squared sum) remains constant. Normalization has typically been implemented via multiplicative or subtractive methods. In multiplicative normalization (von der Malsburg, 1973; Oja, 1982), each weight is multiplied by an amount necessary to maintain the overall net weight (van Ooyen, 2001). In the subtractive method, all synaptic weights are changed by a fixed amount regardless of their weight: a decrease to compensate for the effect of homosynaptic potentiation or an increase to compensate for the effect of homosynaptic depression (Miller and MacKay, 1994). Normalization, either multiplicative or subtractive, robustly prevents runaway dynamics of activity, because it maintains synaptic drive of a cell at a certain level.

It does not, however, prevent runaway dynamics of individual synaptic weights. Indeed training of models with normalization typically leads to a bimodal distribution of synaptic weights, with the weights of “winner” synapses bunched around the maximal value and weights of other synapses gathered around zero (e.g., Song et al., 2000; van Rossum et al., 2000; Gütiç et al., 2003; Morrison et al., 2008).

Both methods of normalization introduce synaptic competition (see below). The use of normalization for achieving both stability of the activity level and synaptic competition has been further elaborated in later studies (e.g., Kempter et al., 2001; Elliott and Shadbolt, 2002; Wu and Yamaguchi, 2006; Finelli et al., 2008).

Note that because normalization affects all synapses of the neuron irrespective of their recent activity, it postulates the existence of heterosynaptic plasticity—changes of synapses which were not activated during plasticity induction.

Homeostatic Mechanism: The Trigger and the Timescale

In most modeling studies homeostasis of total synaptic weight and activity is achieved by normalization, implemented directly into the learning rules which update synaptic weights in each iteration. This introduces to the models, as an automatic consequence of their design, two important features: first, the trigger for plasticity also inevitably triggers the normalization, and second, the normalization operates on the same time scale as the plasticity rule. These two features may represent additional characteristics required from a homeostatic mechanism in systems with Hebbian-type learning rules.

The trigger for the homeostatic mechanism which protects synaptic weights and neuronal activity from runaway in the face of Hebbian plasticity is not unequivocally identified. The operation of neurons and neuronal networks is regulated by a multitude of homeostatic mechanisms operating at different levels, with a number of parameters that are tracked and a number of signals that could trigger homeostatic mechanisms. At the single-neuron level, biologically plausible candidates that regulate plasticity and constrain total synaptic weights could range from changes of the level of activity-dependent calcium influx (e.g., Yeung et al., 2004), to competition for limited available energy, intracellular resources or neurotrophic factors (e.g., Frey and Morris, 1997, 1998; Elliott and Shadbolt, 2002; Fonseca et al., 2004). These processes are initiated by the same signals which trigger synaptic plasticity (e.g., rise of intracellular calcium), and therefore by design accompany synaptic plasticity and operate on a similar time scale. At the network level, changes of the firing rate of a neuron or a neuronal population can be used as a trigger for homeostatic synaptic scaling (e.g., van Rossum et al., 2000; Zenke et al., 2013). These possibilities are not mutually exclusive, and some are inter-related, e.g., changes of firing rate inevitably lead to changes of calcium influx in the dendrites via voltage-gated calcium channels activated by backpropagating spikes (Spruston et al., 1995; Golding et al., 2002; Waters et al., 2003; Sjöström and Häusser, 2006). Because of their relation to global variables such as concentration of calcium in neurons or firing rate of neuronal populations, mechanisms of homeostatic regulation of synaptic weights might be embedded

into a multi-level system of neuronal homeostasis and thus could be triggered by signals from several different levels.

The time scale on which mechanism(s) of synaptic homeostasis should operate is better understood. Results from the studies in which normalization was implemented as a separate mechanism with an individual time constant converge at the conclusion that to effectively counteract runaway dynamics imposed by associative plasticity this time constant should be short. In one of the first models which implemented homeostatic scaling of synaptic weights regulated by changes of the postsynaptic firing rate, a relatively short time constant of 100 s was used (van Rossum et al., 2000). This time scale is comparable to the time scale of synaptic changes induced by STDP or other Hebbian-type rules. Although in the model of van Rossum et al. (2000) the homeostatic scaling was primarily used to introduce synaptic competition, the same mechanism by design stabilizes the activity level. In a theoretical analysis of one-trial sequence learning of place-fields in the hippocampus, Wu and Yamaguchi (2006) concluded that for learning processes that occur within minutes, the physiological mechanism that constrains synaptic weights must also operate rapidly. The relation between the timescale of a homeostatic mechanism and its ability to maintain the stability of a system with Hebbian-type plastic synapses has been directly addressed in a recent study by Zenke et al. (2013). The authors found that to achieve robust stability of the system, Hebbian-type plastic synapses must be complemented by a homeostatic mechanism operating on a time scale of seconds to minutes, which is comparable to the time scale of plasticity itself (Zenke et al., 2013). This conclusion is compatible with a wealth of evidence from computational studies. In fact, to the best of our knowledge, in all computer models which used normalization or mechanisms inspired by homeostatic synaptic scaling for the purpose of preventing runaway dynamics and/or introducing synaptic competition, these mechanisms were implemented on a fast time scale, the same or comparable to the time scale of the Hebbian plasticity.

The Need for Competition

Theoretical analyses have demonstrated that the establishment and refinement of sensory representations during development (Wiesel and Hubel, 1963; Aitkin et al., 1970; Merzenich et al., 1975; Thompson et al., 1983; Feldman, 2009) requires strong competition between projecting fibers. An increase of the weights of synapses formed by some fibers takes place at the expense of synapses formed by other fibers, which must decrease their weights (von der Malsburg, 1973; Miller and MacKay, 1994; van Ooyen, 2001). In the mature brain, synaptic competition is instrumental, for example, in learning tasks that involve discrimination (Skorheim et al., 2014). Competition between synapses, for example for limited available resources, could also be one of the natural ways in which the magnitude of possible weight change is restricted, thus preventing excessive increases of synaptic strengths. Early theoretical studies had already suggested that competition may involve a broad range of physiologically-restricted resources, such as receptor molecules, surface area, energy resources or plasticity factors (von der Malsburg, 1973; Oja, 1982).

Homosynaptic plasticity can change synaptic weights in both directions, and thus can introduce a certain degree of competition. However, this competition is restricted to synapses which are subject to specific patterns of presynaptic activation, and is determined by presynaptic activity patterns which may be independent of each other. For example, inputs expressing high frequency activity will be potentiated, and those consistently active at low frequencies will be depressed. In the framework of STDP, synapses from presynaptic neurons that are repeatedly active shortly before the postsynaptic spikes and thus fall into potentiation window will be strengthened, while those repeatedly active shortly after postsynaptic spikes, during the depression window, will be weakened. Such scenarios, though possible in theory, impose strict requirements on the patterns of input activity and their relationship to the details of plasticity rules for potentiation and depression. Moreover, in both cases, synaptic changes are restricted to active synapses. Because bidirectional changes rely on specific patterns of presynaptic activity, this mechanism represents competition between external activation patterns, but has no relation to biologically-plausible forms of cell-intrinsic competition, such as competition between synapses of the same cell for limited resources (e.g., Frey and Morris, 1997, 1998; Elliott and Shadbolt, 2002; Fonseca et al., 2004). This latter point is important because models in which competition results from implementing an underlying physiological mechanism have stronger explanatory and predictive power than those with competition imposed as a mathematical convenience (van Ooyen, 2001; Elliott and Shadbolt, 2002). Therefore, ideally, competition should be a consequence of the learning rule and not require explicit additional rules (Gerstner and Kistler, 2002).

Normalization introduces competition that is not restricted to activated synapses, because any change of synaptic weights induced by associative plasticity at a group of active synapses is accompanied by an opposite-direction change of the weights of all other synapses. Multiplicative and subtractive normalization introduce a different degree of competition between synapses, and thus may lead to different final distributions of synaptic weights and distinct functional connectivity (Miller and MacKay, 1994; Miller, 1996). For example, in a model of development of the visual cortex, multiplicative normalization leads to the development of receptive fields with a “graded” distribution of the inputs, such that most of the inputs that expressed correlated activity during the training period were represented. In contrast, with subtractive normalization the final receptive fields were restricted to a subset of inputs which expressed the strongest correlation while other inputs to a cell, including those weakly correlated, were eliminated (Miller and MacKay, 1994).

To summarize, theoretical analysis clearly shows the necessity of mechanism(s) that (i) counteract positive feedback imposed by Hebbian-type rules on synaptic weights and neuronal activity and prevent their runaway dynamics, and (ii) introduce synaptic competition. It has also identified the following features of candidate mechanisms. First, such mechanisms should be able to robustly support the stability of network operation under a broad range of conditions, such as experimentally observed variability in the details of plasticity rules, and diverse patterns of activity. Second, it should prevent runaway dynamics of synaptic

weights, but also support synaptic competition. It is not clear if both roles are related and served by a single mechanism, such as the normalization, or mediated by diverse mechanisms. Third, candidate mechanisms should include heterosynaptic plasticity, because changes of synaptic weights at non-active synapses seem to be necessary for achieving stable dynamics of model neurons and networks with plastic synapses. Finally, it should operate on a timescale comparable to the timescale of Hebbian-type synaptic plasticity. This latter requirement would be automatically fulfilled if homeostatic mechanism(s) in question and Hebbian-type plasticity share the trigger and are mediated by overlapping intracellular machinery.

Biological Candidate: Homeostatic Synaptic Scaling

Experimental Phenomena

A form of plasticity which has received much attention as a potential mechanism of stabilization of neuronal activity is the phenomenon of homeostatic synaptic scaling (Turrigiano et al., 1998; recently reviewed in Burrone and Murthy, 2003; Turrigiano and Nelson, 2004; Rich and Wenner, 2007; Watt and Desai, 2010; Turrigiano, 2012). Homeostatic synaptic scaling is defined as compensatory up- or down- scaling of synaptic weights triggered by prolonged dramatic changes of neuronal activity, whereby synaptic weights adjust to counteract changes of activity. Synaptic weights scale up after hours and days of activity blockade by tetrodotoxin (TTX; Turrigiano et al., 1998) or hyperpolarization of neurons caused by expression of inwardly rectifying potassium channels (Burrone et al., 2002), and scale down after prolonged increases of activity by blockade of inhibition (Turrigiano et al., 1998). This scaling is multiplicative and thus maintains the relative strength of existing synaptic weights. Because homeostatic synaptic scaling is triggered by changes of neuronal firing and acts to neutralize these changes, it can be considered a mechanism of firing rate homeostasis. Originally discovered in dissociated neuron cultures, scaling has been also demonstrated in the whole brain after hours of deafferentation (Becker et al., 2013; Keck et al., 2013; Vlachos et al., 2013). Homeostatic synaptic scaling in the neocortex is developmentally regulated, such that layer 4 neurons express a transient ability to scale, and subsequently neurons from layers 2/3 demonstrate the phenomena which persist through adulthood (Turrigiano, 2012).

Although originally homeostatic scaling was described as a global cell-wide process, later studies raised the possibility that scaling could be quasi-local (at the spatial resolution of dendritic branches) which could serve for localized normalization of weights in a computationally useful way (Burrone et al., 2002; Branco et al., 2008; Rabinowitch and Segev, 2008). Realistically homeostatic scaling is not a unitary phenomenon, but involves a multitude of synaptic mechanisms such as scaling of synaptic weight at the postsynaptic site (Turrigiano et al., 1998) and changes of release probability (Bacci et al., 2001; Murthy et al., 2001; Thiagarajan et al., 2005). Homeostatic regulation of synaptic drive may be also achieved with nonsynaptic

mechanisms such as changes of the intrinsic excitability of neurons (Zhang and Linden, 2003; Karmarkar and Buonomano, 2006). Possible triggers, in addition to originally suggested changes of the postsynaptic firing rate, include changes in transmitter release or activation of postsynaptic receptors, and changes of calcium influx (Burrone and Murthy, 2003; Hou et al., 2008; Fong et al., 2015). A common feature of the events which activate homeostatic plasticity is their long duration, hours and days, and extreme magnitude, such as complete blockade of activity or elimination of sensory input after peripheral lesions, or substantial increase of activity induced by a blockade of inhibition. The activated mechanisms are aimed to counteract the effects of these changes on neuronal activity and push the firing rate back to a set point.

Homeostatic synaptic scaling, operating alongside other plasticity mechanisms, has an established role in the compensatory plastic changes observed in the visual cortex in response to dramatic distortions of normal development, such as ocular dominance plasticity in experimental paradigms of monocular deprivation (Desai et al., 2002; Watt and Desai, 2010; Vituriera et al., 2012; Keck et al., 2013; Lambo and Turrigiano, 2013). It may also be involved in maintaining activity level during normal development, especially in periods of synaptogenesis and pruning (Desai et al., 2002; Turrigiano, 2012). It has been proposed to contribute to the maintainance of normal patterns of sleep oscillations after thalamic lesions (Lemieux et al., 2014). However, while synaptic scaling may be suited to adjust long term and drastic alterations of activity, and supplement other homeostatic mechanisms operating during development, two features make it a poor candidate for serving the acute constraining role necessary to combat runaway dynamics imposed by Hebbian-type plasticity rules: the time scale and the trigger for induction.

Timescale of Homeostatic Synaptic Scaling and Runaway Dynamics

Homeostatic synaptic scaling operates on the timescale of hours and days. The “rapid” scaling takes place after 4 h of complete silencing of cultured neurons with TTX (Ibata et al., 2008). This timescale is compatible with that of developmental processes, such as the formation of sensory representations in norm and pathology (Wiesel and Hubel, 1963; Aitkin et al., 1970; Merzenich et al., 1975; Thompson et al., 1983; Feldman, 2009), but it is at least two orders of magnitude slower than the timescale of associative plasticity, which changes synaptic weights within minutes or even tens of seconds. Runaway dynamics of synaptic weights and activity can be induced by Hebbian-type learning rules within seconds or minutes (e.g., Chen et al., 2013; Zenke et al., 2013). Mechanisms of homeostatic synaptic scaling will be engaged and start affecting synaptic weights after a system has been in a runaway state for hours, which prevents its normal operation. Because of this fundamental difference of the timescales, homeostatic synaptic scaling cannot mediate synaptic competition or normalize synaptic weights during on-going associative synaptic plasticity. For the same reason, it is also not suitable for counteracting runaway dynamics induced by associative plasticity. This inconsistency between time scales of

slow homeostatic scaling and fast associative learning has been pointed out by Wu and Yamaguchi (2006) who concluded that synaptic scaling does not seem to work for fast learning. A recent theoretical study confirmed this conclusion, demonstrating that for achieving robust stability of a system with Hebbian-type plastic synapses, the mechanism that maintains homeostasis and prevents runaway dynamics must operate on a time scale comparable to the plasticity itself (Zenke et al., 2013).

Realism of Experimental Paradigm: Trigger for Homeostatic Synaptic Scaling

One further concern regarding possible involvement of homeostatic synaptic scaling in balancing synaptic changes induced by Hebbian-type plasticity is the severity of changes that are required to trigger the scaling. Typically, scaling-up is induced by a complete silencing of activity for many hours by TTX application (Turrigiano et al., 1998; Turrigiano, 2012). Such dramatic and global changes of activity are neither likely, nor compatible with normal operation of the brain. A recent study demonstrated that 6 h after complete bilateral retinal lesions, activity in the visual cortex is reduced to ~60% of pre-lesion level (Keck et al., 2013). The authors did see evidence for homeostatic synaptic scaling after the lesions, but noted that homeostatic scaling alone could not explain the observed recovery of activity in the deprived cortex (Keck et al., 2013). Note that even this more “modest” intervention represents extreme pathology. In contrast, Hebbian-type synaptic plasticity can be induced by far more subtle events (e.g., see Chistiakova and Volgushev, 2009 for number of spikes in typical plasticity-induction protocols), and activity changes that may result from associative synaptic plasticity might be also far less dramatic.

In theory, the requirement for a dramatic and prolonged change of activity for triggering the homeostatic synaptic scaling and its very long time scale could be related issues: if the rate of change is very low, it will take a long time until a change becomes detectable. Indeed, in experiments on cultured neurons, the amplitude of miniature EPSCs increased progressively during TTX application (e.g., Turrigiano et al., 1998; Ibata et al., 2008). Averaged rates of change in the amplitudes of miniature EPSCs calculated from these experiments were at ~2% increase per hour during complete blockade of spiking with TTX, and ~0.6% decrease per hour during bicuculline-induced increase of activity. Note that thousands to tens of thousands of “extra” spikes were generated during hours of elevated activity, or were missing during hours of reduced activity. Existing experimental evidence neither provides proof, nor allows us to exclude the possibility that homeostatic synaptic scaling can indeed be triggered by physiological-range changes of activity. It is clear, however, that even if it is induced by physiological changes of activity, these changes are both too small and too slow to be able to counteract tendency for runaway dynamics induced by Hebbian-type learning.

Computational Properties of Homeostatic Synaptic Scaling

Depending on the pattern of changes of the input, homeostatic synaptic scaling may have diverse effects on the operation of

neuronal networks: the per design homeostatic effect on activity, a normalizing effect on synaptic weights and normalization-related competition, but also a destabilizing effect on synaptic weights and neuronal activity.

Normalization of synaptic weights by the mechanism of homeostatic synaptic scaling can be understood as following. Consider a simplistic situation in which postsynaptic firing of a neuron is proportional to its total synaptic drive. Potentiation of a portion of synapses would lead to a firing rate increase. To counteract this increase, a mechanism of firing rate homeostasis would scale down all synapses to restore the target firing rate, and therefore also the total synaptic drive, thus performing multiplicative normalization of synaptic weights. This effect might be considered a “delayed normalization,” as the feedback from a change in activity to synaptic scaling operates via the slow loop of firing rate homeostasis. The time scale of this delayed normalization is determined by the time scale of homeostatic scaling. To attain in computer models normalization-derived properties such as synaptic competition or prevention of runaway dynamics of activity, the homeostatic scaling is often implemented on a relatively short time scale of seconds or minutes. However, these models might not reflect computational properties of experimentally observed homeostatic synaptic scaling, because reported timescale is at least two orders of magnitude longer (hours and days). In fact, we are not aware of a computational study in which homeostatic synaptic scaling with experimentally-observed features, specifically the requirement of at least 4 h of altered activity level to produce observable synaptic changes, was shown to be effective in preventing runaway dynamics or supporting synaptic competition during on-going learning.

One further factor that limits the relevance of homeostatic synaptic scaling for maintaining the normal operation of neuronal networks is that this kind of plasticity is triggered by lasting and drastic changes of overall activity, such as a persistent increase or decrease of the firing rate. In biological neuronal networks, the level of activity is subject to energy constraints and is tightly controlled by diverse fast-scale mechanisms operating at both the network and at the cellular levels. At the network level, activity is controlled by inhibition, including recurrent inhibition, which by design limits both the magnitude and the duration of episodes of elevated activity. Strong recurrent inhibition mitigates changes of activity level even when external input changes dramatically, allowing neuronal networks operate in a regime of dynamically balanced excitation and inhibition (Wehr and Zador, 2003; Okun and Lampl, 2008; Ozeki et al., 2009; Dorn et al., 2010; Sun et al., 2010). For example, complete binocular retinal lesions which resulted in an extreme change of the afferent input to the visual cortex, led to only a ~40% reduction of activity in visual cortex of the mouse (Keck et al., 2013). Moreover, inhibitory plasticity can adjust the strength of inhibition and maintain the excitatory/inhibitory balance in neurons and neuronal networks (Vogels et al., 2011; Luz and Shamir, 2012).

At the level of synapses, mechanisms regulating the input strength on a fast time scale include short-term plasticity, vesicle recycling and fast retrograde signaling. Episodes of strong

presynaptic activity lead to depletion of the ready-to-release pool of vesicles, thus limiting release during the following seconds and minutes, or setting a new, lower, steady state of release (Abbott et al., 1997; Tsodyks and Markram, 1997; Varela et al., 1997, 1999; Markram et al., 1998; Sussillo et al., 2007; Costa et al., 2013). Episodes of strong postsynaptic firing and depolarization lead to activation of retrograde signaling that reduces transmitter release (Pitler and Alger, 1992; Wilson and Nicoll, 2001; Freund et al., 2003; Hashimoto et al., 2007). Strong pre and postsynaptic activity is associated with the release of adenosine and cyclic adenosine-phosphates from neurons and glial cells and thus the elevation of extracellular adenosine levels in a local area where active synapses and neurons are located (Pascual et al., 2005; Wall and Dale, 2008; Halassa et al., 2009; Lovatt et al., 2012). Because adenosine has a suppressive effect on synaptic transmission in the neocortex and hippocampus (e.g., Dunwiddie and Haas, 1985; Scanziani et al., 1992; Thompson et al., 1992; Kerr et al., 2013; Bannon et al., 2014; Zhang et al., 2015), this would hinder further buildup of activity at this location. These are just few examples of mechanisms which operate on a substantially faster scale than homeostatic synaptic scaling and which might effectively combat excessive lasting changes of activity and restore activity level long before homeostatic synaptic scaling is activated.

The long time scale of homeostatic synaptic scaling is compatible with the time scale of developmental processes, or compensatory processes during recovery from injury. Homeostatic synaptic scaling may play a role in maintenance of overall activity level during normal development, especially in periods of synaptogenesis and pruning (Desai et al., 2002; Turrigiano, 2012). It may also play a role in pathological conditions, when the mechanisms which maintain the activity level during normal operation, are impaired or overloaded and cannot cope with drastic changes of activity caused by pathology. Indeed, evidence for homeostatic synaptic scaling has been reported in the visual cortex after binocular retinal lesions (Keck et al., 2013) and in the dentate gyrus after denervation (Becker et al., 2013; Vlachos et al., 2013). However, in pathological conditions the “homeostatic” synaptic scaling may also have a destabilizing effect on synaptic weights and neuronal activity. For example, if activity is reduced temporarily (e.g., because of a reversible injury to a peripheral sensory apparatus), the up-scaling of synaptic weights could lead to over-excitability of neurons when input firing recovers. Thus, because synaptic up-scaling follows activity changes with a delay of several hours, it may lead to an over-shoot of activity when the input is recovering. Indeed, network models of pathological conditions, lesions or deafferentation, show that homeostatic synaptic scaling helps to recover normal activity patterns after small and moderate deafferentation, but leads to post-traumatic seizures if the degree of deafferentation is above a certain threshold (about 80%; Houweling et al., 2005; Fröhlich et al., 2008; Volman et al., 2011a,b). One of the contributing mechanisms here may be the formation of new silent synapses during prolonged silencing of activity, which leads to enhancement of LTP induction (Arendt et al., 2013). If a similar process takes place after cortical damage, potentiation of these new synapses after partial recovery of the

activity will further amplify the increase of the overall synaptic drive, which might facilitate the development of seizures or other pathological activity patterns.

To summarize, homeostatic synaptic scaling represents a set of mechanisms which are triggered by extreme and long lasting (hours and days) changes of neuronal activity, and serve to counteract firing rate changes by up- or down-scaling synaptic weights. These mechanisms operate on time scales which are orders of magnitude longer than the time scale at which associative plasticity is induced. Therefore, they would not be engaged or expressed until runaway dynamics had created an unstable and saturated network, which generates dramatically altered activity for hours. Thus, experimentally observed properties of homeostatic synaptic scaling do not fit two crucial requirements for a hypothetical mechanism which maintains stability of operation and provides synaptic competition in systems with Hebbian-type learning rules. Both the trigger and the time scale of synaptic scaling are fundamentally different from those of the Hebbian-type plasticity.

Biological Candidate: Heterosynaptic Plasticity

Experimental Phenomena

Heterosynaptic plasticity refers to changes at synapses which were not active during the induction of plasticity (**Figure 1**). Heterosynaptic LTD accompanying the induction of LTP was first described in the hippocampus shortly after the phenomenon of LTP was discovered (Lynch et al., 1977). In CA1 pyramidal neurons, induction of LTP of Schaffer collateral-commissural synapses at apical dendrites was accompanied by LTD at inputs to basal dendrites made by commissural fibers that were not stimulated during the induction (**Figure 1A**, left). Vice versa, induction of LTP at the basal dendrites was accompanied by LTD at the apical dendrites. Heterosynaptic LTD accompanying the induction of homosynaptic LTP clearly has potential for both balancing plastic changes and supporting synaptic competition.

Spatial distribution of LTP and LTD studied in structures with a regular organization of their inputs, such as the hippocampus or amygdala, revealed a bi-phasic Mexican-hat type profile (White et al., 1990; Royer and Paré, 2003). Induction of LTP at a set of synapses was accompanied by a weaker heterosynaptic LTP at nearby inputs, and heterosynaptic LTD at more distant inputs (**Figure 1A**, right). A symmetrical profile of heterosynaptic changes was observed around the site of LTD induction: weaker LTD at close distances and LTP at more distant inputs (Royer and Paré, 2003). Because the amount of potentiation and depression in these profiles was balanced, this type of heterosynaptic plasticity can provide a powerful local mechanism of both normalization of synaptic weights and synaptic competition.

In the CA1 region of the hippocampus, pairing of one input to a pyramidal neuron led to potentiation not only of that stimulated synapse, but also of synapses formed by nearby fibers on that neuron, and even on nearby neurons (Bonhoeffer et al., 1989; Kossel et al., 1990; Schuman and Madison, 1994; Engert and Bonhoeffer, 1997).

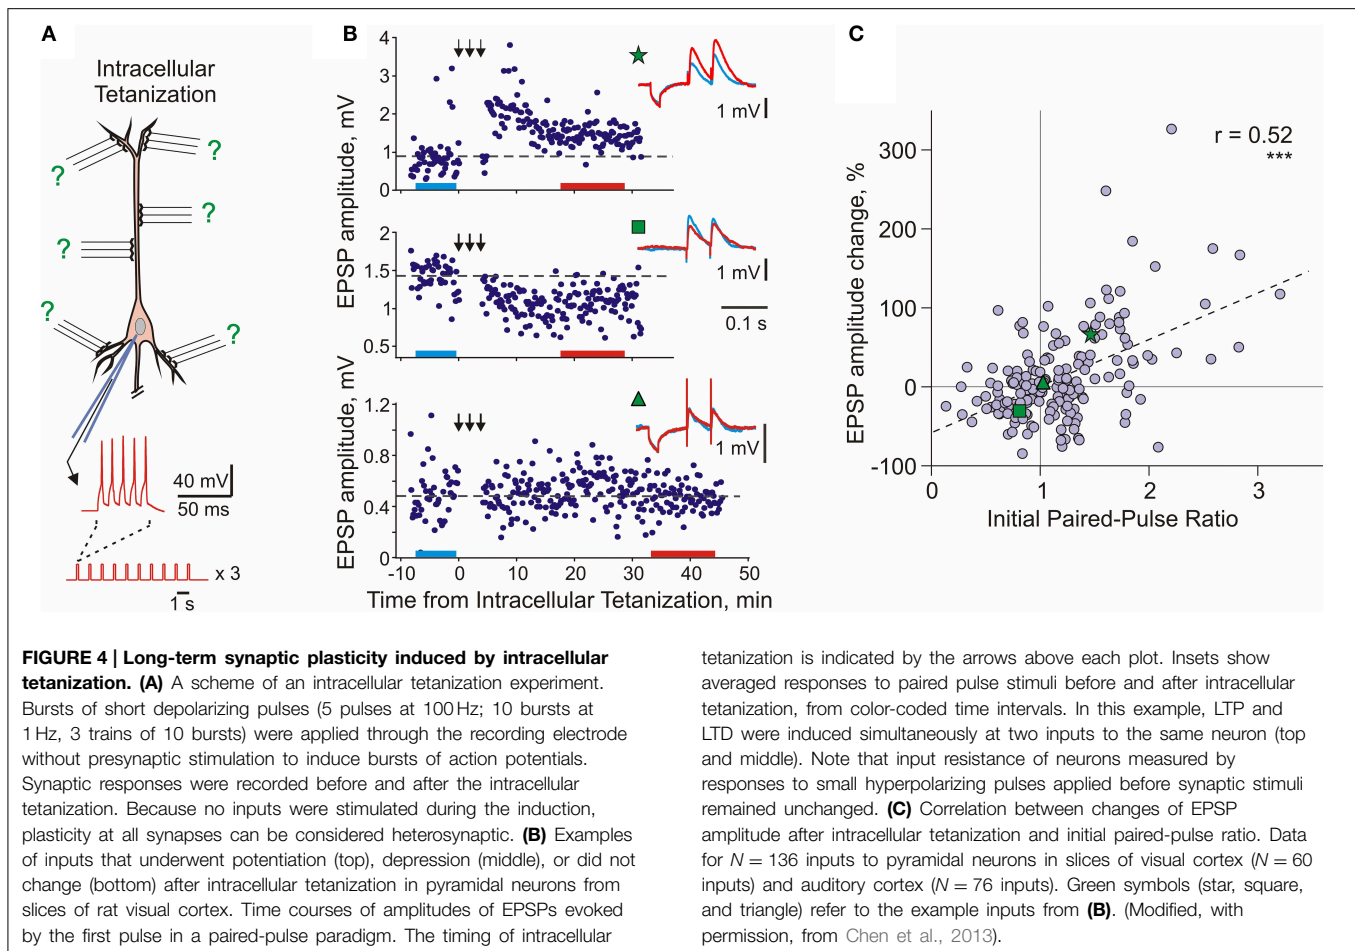
This evidence for heterosynaptic plasticity indicates that presynaptic activation of the synapse is not a strict requirement for plasticity induction. Indeed, long term plasticity can be induced by purely postsynaptic protocols. In the hippocampus and neocortex, photolysis of caged Ca^{2+} (Neveu and Zucker, 1996a,b; Yang et al., 1999) or postsynaptic spiking (Kuhnt et al., 1994; Volgushev et al., 1994, 2000; Cummings et al., 1996; Chistiakova and Volgushev, 2009; Lee et al., 2012) is sufficient to induce plasticity.

Trigger for Heterosynaptic Plasticity

Heterosynaptic changes are triggered by acute rises of intracellular Ca^{2+} concentration (Yang et al., 1999; Balaban et al., 2004; Lee et al., 2012), thus sharing the trigger with Hebbian-type plasticity (Lisman, 1989; Artola and Singer, 1993; Cummings et al., 1996). The required rises of $[\text{Ca}^{2+}]$ can be produced by bursts of action potentials backpropagating throughout the dendritic tree (Spruston et al., 1995; Staubli and Ji, 1996; Larkum et al., 1999; Golding et al., 2002; Waters et al., 2003; Lisman and Spruston, 2005; Sjöström and Häusser, 2006; Remy and Spruston, 2007). Chelation of intracellular calcium impairs induction of heterosynaptic plasticity (Lee et al., 2012). In addition to the shared calcium dependence, intracellular mechanisms of homosynaptic and heterosynaptic plasticity overlap, as indicated by at least partial occlusion between homo- and hetero-synaptic plastic changes (Kuhnt et al., 1994; Cummings et al., 1996; Neveu and Zucker, 1996a,b; Volgushev et al., 1999; Yang et al., 1999). Thus, heterosynaptic plasticity is induced by the same protocols, occurs at the same timescale, and shares mechanisms with Hebbian-type plasticity (see below, **Figure 5** and related text for further discussion).

Properties of Heterosynaptic Plasticity

Heterosynaptic, long-term plastic changes can be induced in hippocampal and neocortical neurons by intracellular tetanization—bursts of spikes evoked by short depolarizing pulses applied through the recording electrode (**Figure 4A**; Kuhnt et al., 1994; Volgushev et al., 1994, 1997, 1999, 2000; Chistiakova and Volgushev, 2009; Lee et al., 2012). The rationale behind the intracellular tetanization protocol as a tool to study heterosynaptic plasticity is the following. Each neuron in the neocortex receives thousands of synaptic inputs, but activation of only a fraction of these inputs, few dozens to hundreds, is necessary to evoke spikes. Repetitive activation of a fraction of inputs and repetitive firing of the postsynaptic cell can, under certain conditions, induce synaptic plasticity. During the induction, all synapses but for those of the activated fraction will experience postsynaptic activity without activation of their presynaptic fibers. This situation, postsynaptic activity without presynaptic activation, is mimicked by the intracellular tetanization (**Figure 4A**). Because none of the synaptic inputs was stimulated during the intracellular tetanization, any changes of synaptic transmission after the intracellular tetanization can be considered heterosynaptic. It is important to note that postsynaptic activity during intracellular tetanization (~150 spikes) is both compatible with activity patterns observed *in vivo*, for example during visual stimulation (e.g., Volgushev et al.,



2003), and comparable to postsynaptic activity during typical plasticity-induction protocols (see Chistiakova and Volgushev, 2009 for comparison of number of spikes in plasticity-induction protocols).

Following intracellular tetanization, amplitudes of synaptic responses could increase, decrease or not change (Figure 4B). The amplitude changes occurred fast, on the same time scale as homosynaptic changes. Moreover, intracellular tetanization could simultaneously induce LTP and LTD in two independent inputs onto one cell (Figure 4B top and middle). The direction of plastic change of a synaptic input was correlated with the initial paired-pulse ratio, a measure which is inversely related to release probability (Figure 4C, Volgushev et al., 1997, 2000; Lee et al., 2012; Chen et al., 2013). Inputs which initially had a low release probability (high initial paired-pulse ratio) were typically potentiated. Inputs that had a high release probability (low initial paired-pulse ratio) were typically depressed or did not change. Thus, the direction of heterosynaptic changes depends on initial properties of a synapse, and is determined at each synapse individually. Weight-dependence is one further similar feature of heterosynaptic and homosynaptic plasticity: it has been also reported for LTP and LTD induced by afferent tetanization or by a pairing procedure in the hippocampus and neocortex

(van Rossum et al., 2000; Sjöström et al., 2001; Hardingham et al., 2007).

The weight-dependence of heterosynaptic plasticity might reflect history-dependent predispositions of synaptic inputs to undergo potentiation or depression (Volgushev et al., 1997, 2000; Chistiakova and Volgushev, 2009). Weak synaptic inputs with low release probability, such as those which underwent depression in the past, are less susceptible to further depression yet have a stronger predisposition for potentiation. Strong synapses with a high release probability, such as those recently potentiated, have a higher predisposition for depression. The notion of the predisposition of synapses for plastic changes is closely related to the ideas of a sliding threshold between depression and potentiation in the BCM rule (Bienenstock et al., 1982; Yeung et al., 2004) and metaplasticity – history-dependent changes of the ability of synapses to undergo potentiation or depression (Abraham and Bear, 1996; Clem et al., 2008).

Thus, heterosynaptic plasticity induced by strong postsynaptic activity has properties which make it an ideal candidate for counteracting runaway dynamics of synaptic weights and mediating synaptic competition. Heterosynaptic plasticity, while not requiring presynaptic activity at the synapse for the induction, has the same trigger (rise of intracellular calcium),

partially overlapping mechanisms of expression, and operates on the same time scale as homosynaptic plasticity. Moreover, heterosynaptic changes can be induced by the same protocols which are typically used to induce homosynaptic plasticity.

Heterosynaptic Plasticity in Published Studies: Meta-Analysis

This latter conclusion stays in apparent contradiction to the wealth of publications reporting that amplitude of responses in non-activated or control inputs did not change, and, more generally, to the notion of input specificity of homosynaptic plasticity. We suggest that this contradiction could be due to the fact that heterosynaptic changes are bidirectional but balanced. To test this conjecture, we re-analyzed results from eight papers on STDP of excitatory inputs to layer 2/3 or layer 5 pyramidal neurons in slices from somatosensory, visual or auditory areas of rat neocortex (Feldman, 2000; Sjöström et al., 2001; Birtoli and Ulrich, 2004; Watt et al., 2004; Letzkus et al., 2006; Nevian and Sakmann, 2006; Hardingham et al., 2007; our data from Chistiakova et al., 2014). In all of these papers clear cases of homosynaptic LTP or LTD are presented. **Figure 5** illustrates the results of 36 experimental series from these papers, as the averaged change of response amplitude (diamond symbol) after the pairing or control procedure and the range covered by ± 2 SD. This range includes 95% of normally distributed values, however because number of measurements contributing to each experimental series was not high, typically between $N = 4$ and $N = 20$, the actual measured values did not necessarily cover the whole ± 2 SD range. **Figure 5** illustrates several important points. First, most “No change or unpaired” groups (**Figure 5**, blue) and especially “AP bursts only” groups (**Figure 5**, blue-gray-pink bars) have high variance, with the ranges of response amplitude changes overlapping substantially with the ranges of homosynaptic changes after LTP and LTD protocols. This implies that “No change” and “AP bursts only” groups must have contained individual cases of potentiation and depression, which were heterosynaptic LTP and LTD in experiments in which no presynaptic stimulation was applied. Moreover, because the averages were not significantly different from zero potentiation and depression were balanced. Second, although on average LTP protocols increased and LTD protocols decreased response amplitude (**Figure 5**, magenta and green), the effects were highly variable, often including changes in the opposite directions. Assuming continuous distributions of EPSP amplitude changes over the mean ± 2 SD range, the LTP and LTD protocols might have induced plastic changes of both signs. In 8 out of 11 LTP groups the value of mean -2 SD is well below zero, suggesting that some of the inputs were depressed. In 5 out of 9 LTD groups mean $+2$ SD reaches well above zero, suggesting that some inputs were potentiated. In most LTP and LTD groups (17 out of 20) there should have been inputs which did not change (**Figure 5**). This suggests that factors other than timing, such as synaptic predispositions for plasticity, might have contributed to the final effect of the plasticity-induction protocol on response amplitude. This conjecture is supported by the results of Hardingham et al. (2007), who found that the same protocol could induce either potentiation or depression.

The direction of the EPSP amplitude change was correlated with the release probability of the synapse before the plasticity induction. This finding is corroborated by our results (**Figure 4C**, modified from Chen et al., 2013). Finally, the range of EPSP amplitude changes in unpaired inputs was typically smaller than the range of amplitude changes induced by spike burst-only protocols (**Figure 5**, d,e,h). This may reflect competition of plastic synapses for limited resources. In this scenario, pairing may facilitate access to resources for homosynaptic plasticity at paired inputs via a mechanism of synaptic tagging (Frey and Morris, 1997, 1998; Fonseca et al., 2004) or a similar process, thus leaving fewer resources available for heterosynaptic changes at unpaired inputs. Spikes-only protocols leave more resources available for heterosynaptic changes, and thus induce heterosynaptic plasticity of a larger amplitude. Note that this latter point (larger variance after spike-burst only protocols as compared to changes in un-paired inputs) is suggested by our meta-analysis (Chistiakova et al., 2014), though limited number of studies precludes statistical analysis. This question needs to be tested in future work.

The above analysis demonstrates that both potentiation and depression might have been induced in individual unpaired inputs and also by spike bursts-only protocols (**Figure 5**). However, when averaged across the inputs, EPSP amplitude changes were not significant because of the balanced nature of heterosynaptic plasticity. It is also important to note that in papers specifically aimed at investigating heterosynaptic plasticity, it was readily induced by regular pairing (Nishiyama et al., 2000; Huang et al., 2008; Arami et al., 2013), afferent tetanization (Cummings et al., 1996; Staubli and Ji, 1996; Chevaleyre and Castillo, 2003; Royer and Paré, 2003; Bauer and LeDoux, 2004; Pascual et al., 2005; Nugent et al., 2007; Wöhrle et al., 2007), or purely postsynaptic protocols (e.g., Pockett et al., 1990; Christofi et al., 1993; Volgushev et al., 1994, 1997, 1999, 2000; Cummings et al., 1996; Lee et al., 2012). This analysis substantiates our conclusion that induction of homosynaptic plasticity by a typical pairing procedure used in STDP studies is accompanied by induction of heterosynaptic plasticity in unpaired inputs.

To summarize, heterosynaptic plasticity induced by intracellular tetanization expresses properties that are well suited for serving as a robust mechanism of normalization of synaptic weights: (i) it depresses strong and potentiates weak synapses thus preventing runaway dynamics of synaptic weights, (ii) it is induced at non-active synapses by the same protocols which induce homosynaptic plasticity, providing for explicit competition and (iii) it operates on the same time scale as homosynaptic plasticity.

Modeling Heterosynaptic Plasticity

Heterosynaptic Plasticity Robustly Prevents Runaway Dynamics

To test the hypothesis that heterosynaptic plasticity can prevent runaway dynamics of synaptic weights and activity (Volgushev et al., 2000; Chistiakova and Volgushev, 2009) we used a

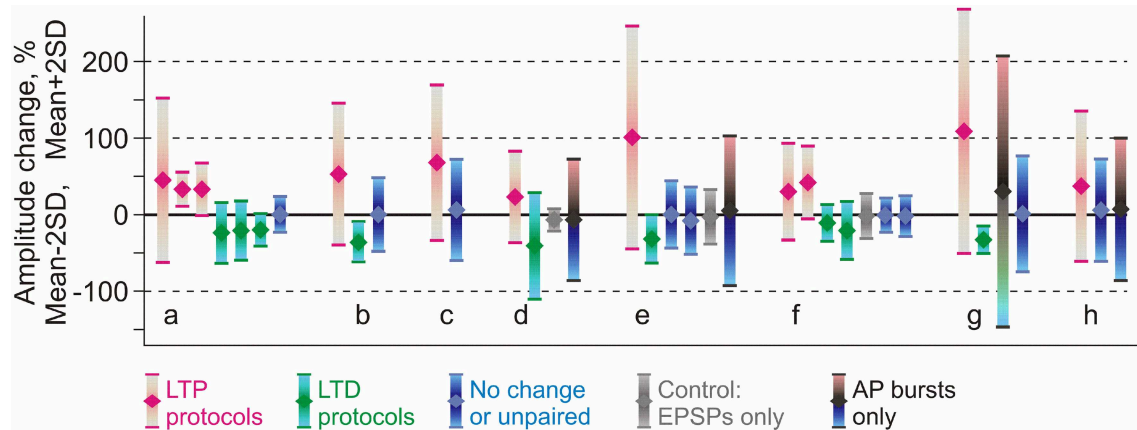


FIGURE 5 | Comparison of reported changes of response amplitude at inputs that were active during the induction (homosynaptic, input-specific) and those not active during the induction (heterosynaptic). The plot shows results of 36 experimental series (bars) from eight papers (groups of bars) on pairing-induced long-term plasticity (STDP), in which the mean amplitude changes were reported together with the SD (or SEM) and number of observations. Each bar shows an average (diamond symbol) change of EPSP amplitude after pairing procedure ± 2 SD. This range includes 95% of normally distributed values. Magenta: changes after LTP protocols (post after pre). Green: changes after LTD protocols (pre after post). Blue: range of EPSP amplitudes after protocols that did not lead to significant changes of the averaged response (such as interval between pre and post spikes outside

plasticity windows). Gray: range of EPSP amplitudes after only presynaptic stimulation without postsynaptic spikes. Black, bars from cyan to pink (in d,e,h): range of EPSP amplitudes after bursts of postsynaptic spikes only, without presynaptic stimulation. Data for excitatory inputs to L2/3 or L5 pyramidal neurons from somatosensory, visual or auditory cortex, from the following papers: Feldman (2000) (a); Sjöström et al. (2001) (b); Watt et al. (2004) (c); Birtoli and Ulrich (2004) (d); Nevian and Sakmann (2006) (e); Letzkus et al. (2006) (f); Hardingham et al. (2007) (g); Chistiakova et al. (2014) (h). Results from Hardingham et al. (2007) (g) present the LTP and LTD data selected by the direction of the change. The third bar in this group shows LTP and LTD data pooled together. Details of experimental protocols can be found in original papers. (Modified, with permission, from Chistiakova et al., 2014).

neuron model with synaptic weight changes governed by STDP rules and heterosynaptic plasticity with experimentally observed properties (Chen et al., 2013). The model neuron received inputs from 100 simulated presynaptic neurons, firing action potentials with Poisson distributed interspike intervals. Activity of presynaptic neurons was mildly correlated, with an averaged cross-correlation between pairs of spike trains of 0.35 ± 0.05 . Averaged presynaptic firing at 1 Hz led to the firing of the postsynaptic model neuron at ~ 1.8 Hz. Synaptic weight changes were governed either by STDP rules (STDP-only models) or by STDP rules complemented with heterosynaptic plasticity (STDP + heterosynaptic plasticity models). Heterosynaptic plasticity was implemented according to experimental data (Volgushev et al., 2000; Chistiakova and Volgushev, 2009; Lee et al., 2012; Chen et al., 2013). It was triggered by increases of intracellular calcium concentration above a threshold level, and affected all synapses in a weight-dependent manner: the probability of synaptic change, its direction, and its magnitude depended on the initial weight. These dependences were implemented using the Equations (1) and (2).

$$P = 3000 \times (W_{\text{syn}} - W_{\text{max}}/2)^2 + 0.1 \quad (1)$$

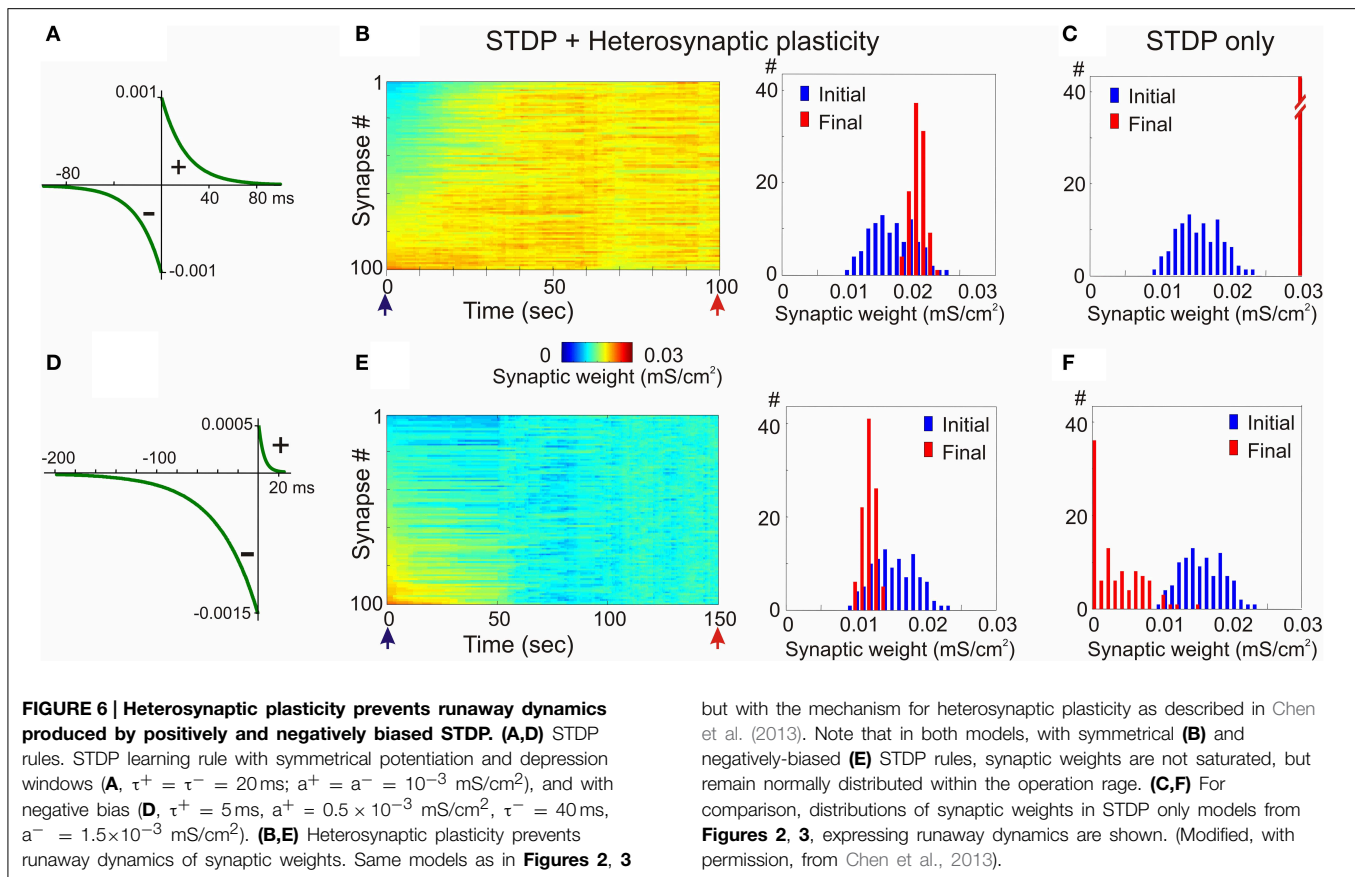
where P is the probability of the synaptic change, W_{syn} is the current synaptic strength and $W_{\text{max}} = 0.03$ mS/cm² is the maximal synaptic strength. According to Equation (1), P is equal to 0.1 for synapses with intermediate strength, and P equals to ~ 0.775 for synapses with maximal or minimal strength. The

change of synaptic weight dW_{syn} was calculated according to following equation:

$$dW_{\text{syn}} = (((1/(1 + \exp((W_{\text{syn}} - (0.5 \times W_{\text{max}})) \times 100))) - 0.5) + \sigma \times 0.02) \times 0.0001 \quad (2)$$

In this equation, dW_{syn} indicates the change of synaptic strength and σ is a random variable drawn from Gaussian distribution with zero mean and standard deviation of 3. A detailed description of the model, implementation of plasticity rules and discussion of parameters can be found in the original publication (Chen et al., 2013).

In the first example, STDP with symmetrical windows for potentiation and depression was used (Figures 6A–C; same STDP rules as in Figure 2). In the model with the STDP-only learning rule, synaptic weights expressed clear runaway dynamics and were saturated at the maximal value by the end of a 100s simulation (Figures 2A,C, 6C). This led to a profound increase of the postsynaptic firing rate despite unchanged presynaptic firing (1 Hz throughout the simulation). The postsynaptic firing rate increased from ~ 1.8 Hz during the first 10 s of simulation, to ~ 6.3 Hz during the last 10 s of simulation. Implementing in the model heterosynaptic plasticity with experimentally observed properties in addition to STDP-rules effectively prevented runaway dynamics of synaptic weights and activity. Synaptic weights in this model increased slightly, but did not saturate. The new stable distribution of synaptic weights



but with the mechanism for heterosynaptic plasticity as described in Chen et al. (2013). Note that in both models, with symmetrical (**B**) and negatively-biased (**E**) STDP rules, synaptic weights are not saturated, but remain normally distributed within the operation range. (**C,F**) For comparison, distributions of synaptic weights in STDP only models from **Figures 2, 3**, expressing runaway dynamics are shown. (Modified, with permission, from Chen et al., 2013).

around the new mean value was completely located within the operational range of synapses (**Figure 6B**). Firing rate of the postsynaptic neuron slightly increased from ~ 1.8 to ~ 2.6 Hz.

In the second example model, the STDP-rule was strongly biased toward depression (**Figure 6D**, same STDP rule as in **Figure 3**). In the STDP-only model, synaptic weights expressed clear runaway dynamics toward the minimum value (**Figures 3A,C, 6F**). The decreased synaptic weights were not producing sufficient depolarization to maintain spiking of the postsynaptic neuron, therefore averaged firing rate of presynaptic neurons was increased to 2 Hz and then to 3 Hz (**Figure 3A**). Even with the three-fold increase of presynaptic firing rate, the postsynaptic neuron became silent. A portion of synaptic weights was saturated at the minimum value (**Figures 3C, 6F**). Heterosynaptic plasticity effectively prevented the runaway dynamics of synaptic weights toward zero and silencing of the cell (**Figure 6E**). These results demonstrate that heterosynaptic plasticity can prevent runaway of synaptic weights to either extreme.

The stabilizing effect of heterosynaptic plasticity on synaptic weights and activity is long-lasting and robust. Heterosynaptic plasticity was able to keep synaptic weights and activity levels within an operational range for models with different calcium thresholds for plasticity induction, models subject to different patterns of presynaptic activity, and over a broad range of parameters of STDP learning rules. This latter point is illustrated

in **Figure 7**. To explore how changing the parameters of STDP rules affects the stability of operation of the model neuron, we systematically varied the amplitude and the time constant of the potentiation window of STDP, while keeping the depression window constant. The set of tested STDP rules covered a range from those strongly dominated by depression, to those with balanced potentiation and depression windows, as well as those dominated by potentiation (see insets in **Figure 7**). As an indicator of runaway dynamics we used the deviation from normality (D'Agostino-Pearson's K^2 -test) of synaptic weight distribution after 100 s of simulation. Note that the shape of the final distribution of synaptic weights is determined by a multitude of factors operating at each synapse. These factors include, but are not restricted to, initial synaptic weight, the set of plasticity mechanisms operating at a synapse and the specifics of these mechanisms, and the pattern of presynaptic activity experienced by a synapse (see below, **Figure 8**). Experimentally measured distributions of synaptic weights, usually asymmetrical and close to log-normal (e.g., Song et al., 2005), might reflect the broad variety of unique combinations of these factors at individual synapses. Because in each simulation presented here these factors were similar for all synapses, synaptic weights converged to the same value. Convergence of synaptic weights to a value within the operation range resulted in a normal distribution, while convergence to one of the extremes resulted in a distribution which deviated from normality. Therefore, we

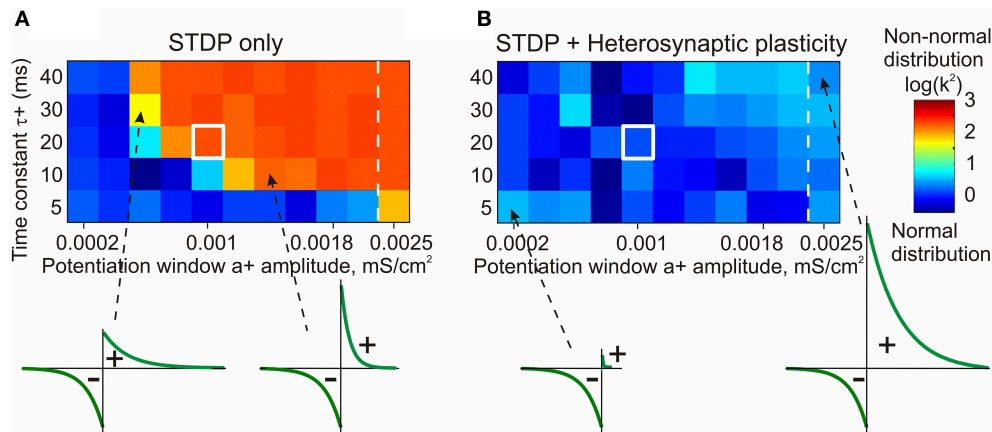


FIGURE 7 | STDP only model fails to prevent runaway dynamics of synaptic weights in broad range of STDP rules (A).

Heterosynaptic plasticity makes a broad range of STDP parameters compatible with stable operation of neurons (B). Each box in the grids shows the D'Agostino-Pearson's K^2 -test for normality of synaptic weight distribution after 100s of simulations with different STDP potentiation windows, with a^+ and τ^+ as indicated on the X and Y axes. White square indicates symmetrical STDP learning rule. Synaptic weight distributions with high K^2 -test values (>50) indicating deviation

from normality, typically contain most of the weights saturated at maximal or minimal values. Note that in simulations with the STDP only model (A), only few STDP rules, with strong bias toward depression, did not lead to runaway dynamics. Most STDP rules, including examples shown in the bottom, led to runaway dynamics of synaptic weights. In contrast, the model with STDP and heterosynaptic plasticity (B) did not express runaway dynamics over the whole range of tested STDP rules, including those extremely unbalanced (insets, bottom). (Modified, with permission, from Chen et al., 2013).

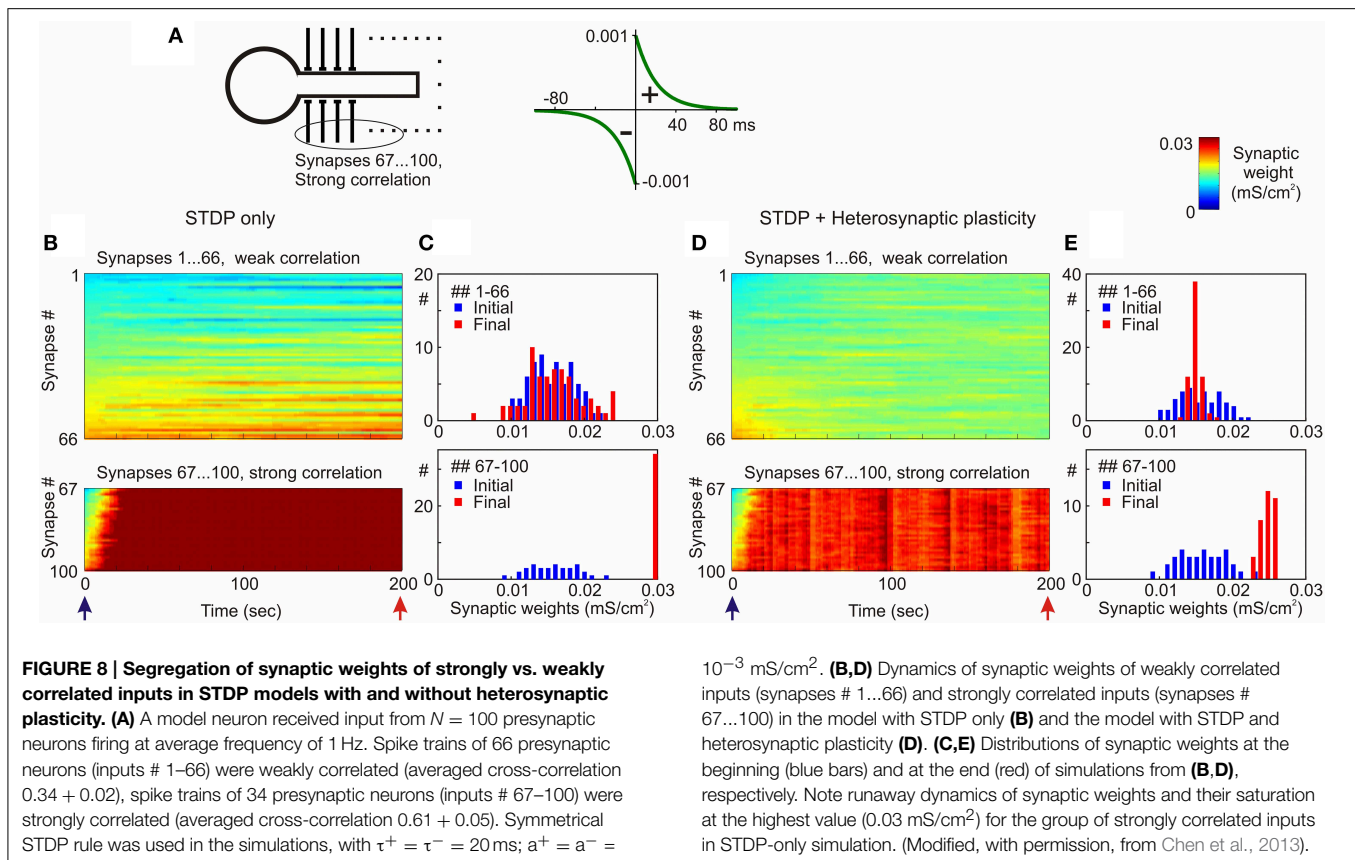
used a test of normality and deviation from normality of the final distribution of synaptic weights as an indicator of runaway dynamics.

The STDP-only model expressed non-saturating behavior only with a limited sub-set of tested STDP rules, in which the window for potentiation was smaller than the depression window (Figure 7A, blue area of the matrix). As soon as the potentiation window of the STDP rule was $\sim 75\%$ of the depression window or stronger, synaptic weights and postsynaptic firing invariably expressed runaway dynamics (Figure 7A, orange/red area of the matrix). In fact, the range of STDP windows compatible with stable dynamics of synaptic weights was somewhat overestimated in these experiments. In cases with a strongly dominating window for depression, the synaptic drive was reduced below the level necessary to evoke postsynaptic spiking before synaptic weights were saturated. After the postsynaptic spiking had ceased, synaptic weights did not change any more per STDP design. Addition of heterosynaptic plasticity to the model robustly prevented runaway dynamics over the whole range of tested STDP parameters, from almost exclusively depressing STDP rules, to those strongly dominated by potentiation (Figure 7B). Joint action of STDP and heterosynaptic plasticity led to the stable distribution of synaptic weights within their operational range. The new equilibrium point of the synaptic weight distribution depended on the relative strength of potentiation and depression windows: in models equipped with a stronger potentiation window of STDP the final distributions of synaptic weights were shifted toward higher values. Notably, heterosynaptic plasticity, by preventing runaway dynamics of synaptic weights, also kept averaged firing rate around the operating point (Chen et al., 2013).

Thus, heterosynaptic plasticity with experimentally observed properties provides a robust stabilizing mechanism, which makes possible the stable operation of neurons expressing a broad range of STDP parameters. This is an important feature because experimental evidence indeed shows wide variations of STDP windows for potentiation and depression and of their relative strength in neurons and synaptic connections of different types (Abbott and Nelson, 2000; Nishiyama et al., 2000; Sjöström et al., 2001; Froemke et al., 2005; Zhou et al., 2005; Haas et al., 2006; Caporale and Dan, 2008; Feldman, 2009).

Heterosynaptic Plasticity Permits Segregation of Inputs and Supports Competition

Despite its strong stabilizing effect, heterosynaptic plasticity does not prevent segregation of the weights of synapses which have diverse properties or are subject to diverse input patterns, and supports synaptic competition. In the examples illustrated in Figure 8, inputs to the model neuron were segregated into two groups, one with a high and the other with a low correlation of presynaptic firing. In the STDP-only model, inputs that were strongly correlated were rapidly potentiated and saturated at the maximum value, while the weights of weakly correlated inputs changed little (Figures 8B,C). In the model with both forms of plasticity, STDP and heterosynaptic, inputs from both groups remained unsaturated. The groups of weakly and strongly correlated inputs formed two clearly separate distributions, both within the operational range of synaptic weights (Figures 8D,E; Chen et al., 2013). Similarly, segregation of synaptic weights was observed when the two groups of inputs differed by their average firing frequency rather than by their correlation. Examples from Figures 6, 8 illustrate that in the model with both forms



of plasticity, STDP and heterosynaptic, the location of the steady-state distribution of a group of synaptic weights depends on the balance of several factors, such as the specifics of plasticity rules at these synapses (Figure 6), the level of correlation of presynaptic firing (Figure 8) or firing frequency.

The origin of synaptic competition arising from heterosynaptic plasticity can be understood as following. Heterosynaptic plasticity triggered by episodes of strong postsynaptic activity pushes all synapses, including those not recently activated, toward an equilibrium point. Because heterosynaptic plasticity is triggered by the same episodes of activity which induce homosynaptic changes, the induction of homosynaptic potentiation or depression does not simply push activated synapses toward the maximum or minimum, but also pushes all non-active synapses toward a separate equilibrium point. The existence of two different target weights for active vs. inactive inputs creates a contrast of forces which drive weight changes at active vs. non-active synapses. This facilitates the segregation of weights of differentially active synapses by plasticity-inducing episodes of postsynaptic activity. By driving synaptic weights toward an equilibrium point within the operating range, heterosynaptic plasticity also prevents their saturation, and supports ongoing differentiation of the weights of synapses which experience different activity. For example, if an initially large number of synapses were active in synchrony and were potentiated, but later on only a portion

of them remained consistently active, the background level of competition provided by heterosynaptic plasticity would be able to suppress the remaining synapses, thus allowing for selection of only the relevant group—a process that may mediate the differentiation stage of learning. In this scenario, synapses compete for maintaining their weights at increased or decreased values set by homosynaptic plasticity, but will be driven to the heterosynaptic equilibrium point if other synapses, but not themselves, are active.

Thus, heterosynaptic plasticity facilitates segregation and competition between groups of synaptic inputs exhibiting diverse properties, such as the frequency or correlation of presynaptic firing, or details of plasticity rules. Moreover, it helps to preserve the ability of a neuron with plastic synapses for further learning: unsaturated synapses have a higher potential for further changes than those potentiated to the maximum or depressed to zero by STDP-only learning rules.

To summarize, heterosynaptic plasticity with experimentally observed properties is a strong candidate mechanism for counteracting the runaway dynamics which is imposed on synaptic weights and activity by the positive feedback of Hebbian-type learning rules. It robustly prevents runaway dynamics over a broad range of activity patterns and details of Hebbian-plasticity rules, such as the balance of STDP windows for potentiation and depression. Heterosynaptic plasticity does not prevent segregation of synaptic weights, and can support synaptic

competition. Moreover, it shares the trigger, has overlapping mechanisms and operates on the same time scale as Hebbian-type plasticity. This combination of features makes heterosynaptic plasticity an ideal candidate mechanism of homeostatic control of synaptic weight changes.

Biological Candidates: Other Mechanisms Counteracting Runaway

Several further mechanisms may contribute to counteracting the tendency for runaway of synaptic weights and activity. One is saturation of plasticity: in a series of potentiation-inducing tetanizations, the magnitude of the response increase after each subsequent tetanization is diminished until the ability for further potentiation is eventually lost altogether (Colino et al., 1992; Huang et al., 1992). Another mechanism is weight-dependence of plasticity, whereby the magnitude of potentiation is smaller at strong synapses which likely were already potentiated, than at weak synapses, which did not experience prior potentiation or were previously depressed (van Rossum et al., 2000; Sjöström et al., 2001; Hardingham et al., 2007). One further mechanism is a sliding calcium threshold for potentiation and depression, whereby depending on the history of recent activity and synaptic changes, the thresholds for potentiation and depression or intracellular calcium homeostasis change (Bienenstock et al., 1982; Yeung et al., 2004). These notions and mechanisms contribute to the concept of metaplasticity—history-dependent changes of the ability of synapses to undergo further plastic changes (Abraham and Bear, 1996; Clem et al., 2008). These mechanisms are inherent to Hebbian-type plasticity rules, and thus are ideally suited to shape the ability of synapses to change. By imposing negative feedback on homosynaptic plastic changes, they clearly can limit the runaway tendency, and thus decrease the instability of a system with plastic synapses. A drawback of these mechanisms, as of any mechanisms governing homosynaptic plasticity, is that they require presynaptic activation and cannot affect inactive synapses. This requirement limits the ability of these mechanisms to serve as regulators of global, cell-wide synaptic homeostasis.

A family of non-synaptic mechanisms regulating intrinsic excitability of neurons is not restricted to activated synapses and thus does not have this limitation. These mechanisms can change the excitability of an activated dendritic branch or a whole neuron, and thus affect all respective synapses (Bliss and Lomo, 1973; Daoudal et al., 2002; Zhang and Linden, 2003; Frick et al., 2004; Karmarkar and Buonomano, 2006; Fink and O'Dell, 2009; Sehgal et al., 2013). Excitability changes may counteract synaptic changes, thus having a homeostatic effect (Zhang and Linden, 2003; Karmarkar and Buonomano, 2006), or enhance and amplify synaptic changes, thus having an anti-homeostatic effect (Frick et al., 2004; Fink and O'Dell, 2009; see Sehgal et al., 2013 for recent review).

Several mechanisms may counteract the development of runaway activity, even in cases in which runaway potentiation or depression of individual synaptic weights had not been prevented. Short-term plasticity determines transient changes

in transmitter release occurring on the temporal scale from milliseconds to several seconds (Zucker and Regehr, 2002), and thus is an important factor shaping synaptic responses to sequences of presynaptic spikes (Abbott et al., 1997; Markram et al., 1998; Abbott and Regehr, 2004; Richardson et al., 2005; Sussillo et al., 2007; Costa et al., 2013). Long-term plasticity is partially expressed presynaptically and thus alters short-term plasticity (e.g., Bekkers and Stevens, 1990; Markram and Tsodyks, 1996; Schulz, 1997; Volgushev et al., 1997), affects synaptic responses to sequences of spikes, and the amplitude of steady-state responses to repetitive presynaptic spikes (Markram and Tsodyks, 1996). Because depletion of synaptic vesicles and short-term depression of transmitter release are proportional to release probability, increase of release probability in association with LTP would lead to a stronger short-term depression, while a decrease of release probability associated with LTD would lead to a weaker short-term depression. As a result, changes of steady-state synaptic responses (and thus of synaptic drive of the postsynaptic neuron) resulting from sequences of presynaptic spikes will be less pronounced than the potentiation or depression of responses to single spikes. The magnitude of this attenuating effect will strongly depend on the relative contribution of presynaptic mechanisms to LTP/LTD expression, time constants of short-term facilitation and depression in a synapse, and on presynaptic firing rate. For example, let us assume that 50% of LTP or LTD magnitude is expressed presynaptically as an increase or a decrease of the release probability. In connections with strong facilitation or strong depression (see Costa et al., 2013, Table 1), the magnitude of LTP or LTD of steady-state responses at ~4 Hz will then be ~15% less than LTP or LTD measured with single-pulses (calculated according to Equations (5)–(7) in Costa et al., 2013). For connections with less pronounced facilitation and depression (depression, facilitation and facilitation-depression in Table 1 in Costa et al., 2013), and lower presynaptic firing rates, the effect will be weaker, between 1 and 7%.

Further mechanisms limiting the level of postsynaptic activity include negative feedback on transmitter release via fast retrograde signaling (e.g., Pitler and Alger, 1992; Wilson and Nicoll, 2001; Freund et al., 2003; Hashimoto et al., 2007), or activity-dependent changes of the extracellular level of adenosine (Pascual et al., 2005; Wall and Dale, 2008; Halassa et al., 2009; Lovatt et al., 2012), which has suppressive effect on synaptic transmission in the neocortex and hippocampus (e.g., Dunwiddie and Haas, 1985; Scanziani et al., 1992; Thompson et al., 1992; Kerr et al., 2013; Bannon et al., 2014; Zhang et al., 2015).

Finally, tight control of activity in neurons and neuronal networks is achieved by inhibition, including recurrent inhibition. Strong recurrent inhibition by design limits changes of activity level even during dramatic changes of external input allowing neuronal networks to operate in a regime of dynamically balanced excitation and inhibition (Wehr and Zador, 2003; Okun and Lampl, 2008; Ozeki et al., 2009; Dorn et al., 2010; Sun et al., 2010). Inhibitory plasticity can adjust the strength of inhibition and maintain the excitatory/inhibitory balance in neurons and neuronal networks (Vogels et al., 2011; Luz and Shamir, 2012).

These mechanisms, together with heterosynaptic plasticity, might contribute to a multi-level system of homeostatic control of synaptic changes and neuronal activity.

Summary and Conclusions

A long history of research supports the hypothesis that homosynaptic plasticity provides a powerful cellular mechanism of learning in a variety of biological systems. In this review we argued that a complimentary form of plasticity—heterosynaptic plasticity—represents a necessary cellular component for homeostatic regulation of synaptic weights and neuronal activity. The necessary properties of a homeostatic mechanism which acutely constrains the runaway dynamics imposed by Hebbian associative plasticity have been well-articulated by theoretical and modeling studies. The experimentally observed properties of heterosynaptic plasticity have introduced it as a strong candidate to fulfill this homeostatic role, and subsequent modeling studies which incorporate heterosynaptic plasticity into model neurons with Hebbian-type learning synapses have confirmed its ability to robustly provide stability and competition. In contrast, properties of homeostatic synaptic scaling, which is triggered by extreme and long lasting (hours and days) changes of neuronal activity, do not fit two crucial requirements for a hypothetical homeostatic mechanism needed to provide stability of operation in the face of on-going associative synaptic changes. Both the trigger and the time scale of homeostatic synaptic scaling are fundamentally different from those of the Hebbian-type plasticity.

Heterosynaptic plasticity, which operates on the same time scale and is triggered by similar activity episodes as homosynaptic plasticity, introduces a normalizing driving force that counterbalances a tendency for runaway dynamics of synaptic weights imposed by homosynaptic plasticity. As a result the system maintains synapses within an operational range, preserving the dynamic range for their changes. This allows it to modify synapses in response to a new experience—new learning. Segregation of synaptic weights and competition between synapses are achieved by the differential driving forces

for the weight changes at active (homosynaptic) and inactive (heterosynaptic) synapses. At strongly activated homosynaptic sites, the associative driving force may be dominant, leading to net potentiation or depression of these sub-populations of synapses. Concurrently, the stabilizing effect of heterosynaptic plasticity dominates at the vast number of synapses which are inactive at that moment. As a consequence, every spike, or burst of spikes, becomes a homeostatic signal to the cell. Because homosynaptic and heterosynaptic changes are induced by the same activity patterns, and take place on the same time scale, the weight of a synapse is determined by the balance of homosynaptic LTP, homosynaptic LTD, and the normalizing force of heterosynaptic plasticity. This allows networks to update the relative strength of inputs while keeping synapses within their operational range, preserving their abilities for further adjustments, and maintaining the activity of neurons and networks in a stationary regime. Importantly, heterosynaptic plasticity allows robust homeostasis of synaptic weights and activity over a wide range of parameters: details of STDP rules, Ca^{2+} thresholds, and frequencies and correlations of presynaptic activity. Therefore, heterosynaptic plasticity expresses all of the desired features of an intrinsic homeostatic mechanism for stabilizing synaptic weight dynamics after learning.

A state of the neural system, e.g., controlled by different neuromodulators, may influence the relative balance of homo- and hetero-synaptic plasticity promoting either associative changes or synaptic homeostasis. Thus, hetero- and homo-synaptic forms of plasticity interact, with their balance depending on the state of the network, and therefore have to be studied in combination as integrative components of the whole plasticity system.

Acknowledgments

We are grateful to James Chrobak and Ian Stevenson for comments and improvement of English. Supported by NIH grant R01 MH087631. MV was partially supported by Humboldt Research Award from the Alexander von Humboldt-Foundation.

References

- Abbott, L. F., and Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nat. Neurosci.* 3, 1178–1183. doi: 10.1038/81453
- Abbott, L. F., and Regehr, W. G. (2004). Synaptic computation. *Nature* 431, 796–803. doi: 10.1038/nature03010
- Abbott, L. F., Varela, J. A., Sen, K., and Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science* 275, 220–224. doi: 10.1126/science.275.5297.221
- Abraham, W. C., and Bear, M. F. (1996). Metaplasticity: the plasticity of synaptic plasticity. *Trends Neurosci.* 19, 126–130. doi: 10.1016/S0166-2236(96)80018-X
- Aitkin, L. M., Anderson, D. J., and Brugge, J. F. (1970). Tonotopic organization and discharge characteristics of single neurons in nuclei of the lateral lemniscus of the cat. *J. Neurophysiol.* 33, 421–440.
- Arami, M. K., Sohya, K., Sarihi, A., Jiang, B., and Yanagawa, Y. (2013). Reciprocal homosynaptic and heterosynaptic long-term plasticity of corticogeniculate projection neurons in layer VI of the mouse visual cortex. *J. Neurosci.* 33, 7787–7798. doi: 10.1523/JNEUROSCI.5350-12.2013
- Arendt, K. L., Sarti, F., and Chen, L. (2013). Chronic inactivation of a neural circuit enhances LTP by inducing silent synapse formation. *J. Neurosci.* 33, 2087–2096. doi: 10.1523/JNEUROSCI.3880-12.2013
- Artola, A., and Singer, W. (1993). Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends Neurosci.* 16, 480–487. doi: 10.1016/0166-2236(93)90081-V
- Babadi, B., and Abbott, L. F. (2010). Intrinsic stability of temporally shifted spike-timing dependent plasticity. *PLoS Comput. Biol.* 6:e1000961. doi: 10.1371/journal.pcbi.1000961
- Bacci, A., Coco, S., Pravettoni, E., Schenk, U., Armano, S., Frassoni, C., et al. (2001). Chronic blockade of glutamate receptors enhances presynaptic release and downregulates the interaction between synaptophysin-synaptobrevin-vesicle-associated membrane protein 2. *J. Neurosci.* 21, 6588–6596.
- Balaban, P., Chistiakova, M., Malyshev, A., and Volgushev, M. (2004). Dependence of calcium influx in neocortical cells on temporal structure of depolarization,

- number of spikes, and blockade of NMDA receptors. *J. Neurosci. Res.* 76, 481–487. doi: 10.1002/jnr.20104
- Bannon, N. M., Zhang, P., Ilin, V., Chistiakova, M., and Volgushev, M. (2014). Modulation of synaptic transmission by adenosine in layer 2/3 of the rat visual cortex *in vitro*. *Neuroscience* 260, 171–184. doi: 10.1016/j.neuroscience.2013.12.018
- Bauer, E. P., and LeDoux, J. E. (2004). Heterosynaptic long-term potentiation of inhibitory interneurons in the lateral amygdala. *J. Neurosci.* 24, 9507–9512. doi: 10.1523/JNEUROSCI.3567-04.2004
- Becker, D., Zahn, N., Deller, T., and Vlachos, A. V. (2013). Tumor necrosis factor alpha maintains denervation-induced homeostatic synaptic plasticity of mouse dentate granule cells. *Front. Cell. Neurosci.* 7:257. doi: 10.3389/fncel.2013.00257
- Bekkers, J. M., and Stevens, C. F. (1990). Presynaptic mechanism for long-term potentiation in the hippocampus. *Nature* 346, 724–729. doi: 10.1038/346724a0
- Bi, G. Q., and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Bienenstock, E. L., Cooper, L. N., and Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J. Neurosci.* 2, 32–48.
- Birtoli, B., and Ulrich, D. (2004). Firing mode-dependent synaptic plasticity in rat neocortical pyramidal neurons. *J. Neurosci.* 24, 4935–4940. doi: 10.1523/JNEUROSCI.0795-04.2004
- Bliss, T. V., and Collingridge, G. L. (1993). A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361, 31–39. doi: 10.1038/361031a0
- Bliss, T. V., and Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232, 331–356. doi: 10.1113/jphysiol.1973.sp010273
- Bonhoeffer, T., Staiger, V., and Aersten, A. (1989). Synaptic plasticity in rat hippocampal slice cultures: local “Hebbian” conjunction of pre- and postsynaptic stimulation leads to distributed synaptic enhancement. *Proc. Natl. Acad. Sci. U.S.A.* 86, 8113–8117. doi: 10.1073/pnas.86.20.8113
- Branco, T., Staras, K., Darcy, K. J., and Goda, Y. (2008). Local dendritic activity sets release probability at hippocampal synapses. *Neuron* 59, 475–485. doi: 10.1016/j.neuron.2008.07.006
- Burrone, J., and Murthy, V. N. (2003). Synaptic gain control and homeostasis. *Curr. Opin. Neurobiol.* 13, 560–567. doi: 10.1016/j.conb.2003.09.007
- Burrone, J., O’Byrne, M., and Murthy, V. N. (2002). Multiple forms of synaptic plasticity triggered by selective suppression of activity in individual neurons. *Nature* 420, 414–418. doi: 10.1038/nature01242
- Caporale, N., and Dan, Y. (2008). Spike timing-dependent plasticity: a Hebbian learning rule. *Annu. Rev. Neurosci.* 31, 25–46. doi: 10.1146/annurev.neuro.31.060407.125639
- Chen, J. Y., Lonjers, P., Lee, C., Chistiakova, M., Volgushev, M., and Bazhenov, M. (2013). Heterosynaptic plasticity prevents runaway synaptic dynamics. *J. Neurosci.* 33, 15915–15929. doi: 10.1523/JNEUROSCI.5088-12.2013
- Chevalyere, V., and Castillo, P. E. (2003). Heterosynaptic LTD of hippocampal GABAergic synapses: a novel role of endocannabinoids in regulating excitability. *Neuron* 38, 461–472. doi: 10.1016/S0896-6273(03)00235-6
- Chistiakova, M., Bannon, N. M., Bazhenov, M., and Volgushev, M. (2014). Heterosynaptic plasticity: multiple mechanisms and multiple roles. *Neuroscientist* 20, 483–498. doi: 10.1177/1073858414529829
- Chistiakova, M., and Volgushev, M. (2009). Heterosynaptic plasticity in the neocortex. *Exp. Brain Res.* 199, 377–390. doi: 10.1007/s00221-009-1859-5
- Christof, G., Nowicky, A. V., Bolsover, S. R., and Bindman, L. Y. (1993). The postsynaptic induction of nonassociative long-term depression of excitatory synaptic transmission in rat hippocampal slices. *J. Neurophysiol.* 69, 219–229.
- Clem, R. L., Celikel, T., and Barth, A. L. (2008). Ongoing *in vivo* experience triggers synaptic metaplasticity in the neocortex. *Science* 319, 101–104. doi: 10.1126/science.1143808
- Colino, A., Huang, Y. Y., and Malenka, R. C. (1992). Characterization of the integration time for the stabilization of long-term potentiation in area CA1 of the hippocampus. *J. Neurosci.* 12, 180–187.
- Costa, R. P., Sjöström, P. J., and van Rossum, M. C. (2013). Probabilistic inference of short-term synaptic plasticity in neocortical microcircuits. *Front. Comput. Neurosci.* 7:75. doi: 10.3389/fncom.2013.00075
- Cummings, J. A., Mulkey, R. M., Nicoll, R. A., and Malenka, R. C. (1996). Ca^{2+} signaling requirements for long-term depression in the hippocampus. *Neuron* 16, 825–833. doi: 10.1016/S0896-6273(00)80102-6
- Daoudal, G., Hanada, Y., and Debanne, D. (2002). Bidirectional plasticity of excitatory postsynaptic potential (EPSP)-spike coupling in CA1 hippocampal pyramidal neurons. *Proc. Natl. Acad. Sci. U.S.A.* 99, 14512–14517. doi: 10.1073/pnas.222546399
- Davis, G. W. (2013). Homeostatic signaling and the stabilization of neural function. *Neuron* 80, 718–728. doi: 10.1016/j.neuron.2013.09.044
- Desai, N. S., Cudmore, R. H., Nelson, S. B., and Turrigiano, G. G. (2002). Critical periods for experience-dependent synaptic scaling in visual cortex. *Nat. Neurosci.* 5, 783–789. doi: 10.1038/nn878
- Dorn, A. L., Yuan, K., Barker, A. J., Schreiner, C. E., and Froemke, R. C. (2010). Developmental sensory experience balances cortical excitation and inhibition. *Nature* 465, 932–936. doi: 10.1038/nature09119
- Dunwiddie, T. V., and Haas, H. L. (1985). Adenosine increases synaptic facilitation in the *in vitro* rat hippocampus: evidence for a presynaptic site of action. *J. Physiol.* 369, 365–377. doi: 10.1113/jphysiol.1985.sp015907
- Elliott, T., and Shadbolt, N. R. (2002). Multiplicative synaptic normalization and a nonlinear Hebb rule underlie a neurotrophic model of competitive synaptic plasticity. *Neural Comput.* 14, 1311–1322. doi: 10.1162/089976602753712954
- Engert, F., and Bonhoeffer, T. (1997). Synapse specificity of long-term potentiation breaks down at short distances. *Nature* 388, 279–284. doi: 10.1038/41815
- Feldman, D. E. (2000). Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron* 27, 45–56. doi: 10.1016/S0896-6273(00)00008-8
- Feldman, D. E. (2009). Synaptic mechanisms for plasticity in neocortex. *Annu. Rev. Neurosci.* 32, 33–55. doi: 10.1146/annurev.neuro.051508.135516
- Finelli, L. A., Haney, S., Bazhenov, M., Stopfer, M., and Sejnowski, T. J. (2008). Synaptic learning rules and sparse coding in a model sensory system. *PLoS Comput. Biol.* 4:e1000062. doi: 10.1371/journal.pcbi.1000062
- Fink, A. E., and O’Dell, T. J. (2009). Short trains of theta frequency stimulation enhance CA1 pyramidal neuron excitability in the absence of synaptic potentiation. *J. Neurosci.* 29, 11203–11214. doi: 10.1523/JNEUROSCI.1450-09.2009
- Fong, M. F., Newman, J. P., Potter, S. M., and Wenner, P. (2015). Upward synaptic scaling is dependent on neurotransmission rather than spiking. *Nat. Commun.* 6, 6339. doi: 10.1038/ncomms7339
- Fonseca, R., Nägerl, U. V., Morris, G. M., and Bonhoeffer, T. (2004). Competing for memory: hippocampal LTP under regimes of reduced protein synthesis. *Neuron* 44, 1011–1020. doi: 10.1016/j.neuron.2004.10.033
- Freund, T. F., Katona, I., and Piomelli, D. (2003). Role of endogenous cannabinoids in synaptic signaling. *Physiol. Rev.* 83, 1017–1066. doi: 10.1152/physrev.00004.2003
- Frey, U., and Morris, R. G. (1998). Synaptic tagging: implications for late maintenance of hippocampal long-term potentiation. *Trends Neurosci.* 21, 181–188. doi: 10.1016/S0166-2236(97)01189-2
- Frey, U., and Morris, R. G. M. (1997). Synaptic tagging and long-term potentiation. *Nature* 385, 533–536. doi: 10.1038/385533a0
- Frick, A., Magee, J., and Johnston, D. (2004). LTP is accompanied by an enhanced local excitability of pyramidal neuron dendrites. *Nat. Neurosci.* 7, 126–135. doi: 10.1038/nn1178
- Froemke, R. C., Poo, M., and Dan, Y. (2005). Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature* 434, 221–225. doi: 10.1038/nature03366
- Fröhlich, F., Bazhenov, M., and Sejnowski, T. J. (2008). Pathological effect of homeostatic synaptic scaling on network dynamics in diseases of the cortex. *J. Neurosci.* 28, 1709–1720. doi: 10.1523/JNEUROSCI.4263-07.2008
- Gerstner, W., and Kistler, W. M. (2002). Mathematical formulations of Hebbian learning. *Biol. Cybern.* 87, 404–415. doi: 10.1007/s00422-002-0353-y
- Gilson, M., and Fukai, T. (2011). Stability versus neuronal specialization for STDP: long-tail weight distributions solve the dilemma. *PLoS ONE* 6:e25339. doi: 10.1371/journal.pone.0025339
- Golding, N. L., Staff, N. P., and Spruston, N. (2002). Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature* 418, 326–331. doi: 10.1038/nature00854

- Gütig, R., Aharonov, R., Rotter, S., and Sompolsinsky, H. (2003). Learning input correlations through nonlinear temporally asymmetric Hebbian plasticity. *J. Neurosci.* 23, 3697–3714.
- Haas, J. S., Nowotny, T., and Abarbanel, H. D. I. (2006). Spike-timing-dependent plasticity of inhibitory synapses in entorhinal cortex. *J. Neurophysiol.* 96, 3305–3313. doi: 10.1152/jn.00551.2006
- Halassa, M. M., Florian, C., Fellin, T., Munoz, J. R., Lee, S. Y., Abel, T., et al. (2009). Astrocytic modulation of sleep homeostasis and cognitive consequences of sleep loss. *Neuron* 61, 213–219. doi: 10.1016/j.neuron.2008.11.024
- Hardingham, N. R., Hardingham, G. E., Fox, K. D., and Jack, J. B. (2007). Presynaptic efficacy directs normalization of synaptic strength in layer 2/3 rat neocortex after paired activity. *J. Neurophysiol.* 97, 2965–2975. doi: 10.1152/jn.01352.2006
- Hashimoto, Y., Ohno-Shosaku, T., and Kano, M. (2007). Endocannabinoids and synaptic function in the CNS. *Neuroscientist* 13, 127–137. doi: 10.1177/1073858406296716
- Haydon, P. G. (2001). GLIA: listening and talking to the synapse. *Nat. Rev. Neurosci.* 2, 185–193. doi: 10.1038/35058528
- Hou, Q., Zhang, D., Jarzyl, L., Hugarir, R. L., and Man, H. Y. (2008). Homeostatic regulation of AMPA receptor expression at single hippocampal synapses. *Proc. Natl. Acad. Sci. U.S.A.* 105, 775–780. doi: 10.1073/pnas.0706447105
- Houweling, A. R., Bazhenov, M., Timofeev, I., Steriade, M., and Sejnowski, T. J. (2005). Homeostatic synaptic plasticity can explain post-traumatic epileptogenesis in chronically isolated neocortex. *Cereb. Cortex* 15, 834–845. doi: 10.1093/cercor/bhh184
- Huang, Y., Yasuda, H., Sarihi, A., and Tsumoto, T. (2008). Roles of endocannabinoids in heterosynaptic long-term depression of excitatory synaptic transmission in visual cortex of young mice. *J. Neurosci.* 28, 7074–7083. doi: 10.1523/JNEUROSCI.0899-08.2008
- Huang, Y. Y., Colino, A., Selig, D. K., and Malenka, R. C. (1992). The influence of prior synaptic activity on the induction of long-term potentiation. *Science* 255, 730–733. doi: 10.1126/science.1346729
- Ibata, K., Sun, Q., and Turrigiano, G. G. (2008). Rapid synaptic scaling induced by changes in postsynaptic firing. *Neuron* 57, 819–826. doi: 10.1016/j.neuron.2008.02.031
- Karmarkar, U. R., and Buonomano, D. V. (2006). Different forms of homeostatic plasticity are engaged with distinct temporal profiles. *Eur. J. Neurosci.* 23, 1575–1584. doi: 10.1111/j.1460-9568.2006.04692.x
- Keck, T., Keller, G. B., Jacobsen, R. I., Eysel, U. T., Bonhoeffer, T., and Hübner, M. (2013). Synaptic scaling and homeostatic plasticity in the mouse visual cortex *in vivo*. *Neuron* 80, 227–234. doi: 10.1016/j.neuron.2013.08.018
- Kempter, R., Gerstner, W., and van Hemmen, J. L. (2001). Intrinsic stabilization of output rates by spike-based Hebbian learning. *Neural Comput.* 13, 2709–2741. doi: 10.1162/089976601317098501
- Kerr, M. I., Wall, M. J., and Richardson, M. J. E. (2013). Adenosine A1 receptor activation mediates the developmental shift at layer 5 pyramidal cell synapses and is a determinant of mature synaptic strength. *J. Physiol.* 591, 3371–3380. doi: 10.1113/jphysiol.2012.244392
- Kossel, A., Bonhoeffer, T., and Bolz, J. (1990). Non-Hebbian synapses in rat visual cortex. *Neuroreport* 1, 115–118. doi: 10.1097/00001756-199010000-00008
- Kuhnt, U., Kleschevnikov, A. M., and Voronin, L. L. (1994). Long-term enhancement of synaptic transmission in the hippocampus after tetanization of single neurones by short intracellular current pulses. *Neurosci. Res. Comm.* 14, 115–123.
- Lambo, M. E., and Turrigiano, G. G. (2013). Synaptic and intrinsic homeostatic mechanisms cooperate to increase L2/3 pyramidal neuron excitability during a late phase of critical period plasticity. *J. Neurosci.* 33, 8810–8819. doi: 10.1523/JNEUROSCI.4502-12.2013
- Larkum, M. E., Kaiser, K. M. M., and Sakmann, B. (1999). Calcium electrogenesis in distal apical dendrites of layer 5 pyramidal cells at a critical frequency of back-propagating action potentials. *Proc. Natl. Acad. Sci. U.S.A.* 96, 14600–14604. doi: 10.1073/pnas.96.25.14600
- Lee, C. M., Stoelzel, C., Chistiakova, M., and Volgushev, M. (2012). Heterosynaptic plasticity induced by intracellular tetanization in layer 2/3 pyramidal neurons in rat auditory cortex. *J. Physiol.* 590, 2253–2271. doi: 10.1113/jphysiol.2012.228247
- Lemieux, M., Chen, J. Y., Lonjers, P., Bazhenov, M., and Timofeev, I. (2014). The impact of cortical deafferentation on the neocortical slow oscillation. *J. Neurosci.* 34, 5689–5703. doi: 10.1523/JNEUROSCI.1156-13.2014
- Letzkus, J. J., Kampa, B. M., and Stuart, G. J. (2006). Learning rules for spike timing-dependent plasticity depend on dendritic synapse location. *J. Neurosci.* 26, 10420–10429. doi: 10.1523/JNEUROSCI.2650-06.2006
- Lisman, J., and Spruston, N. (2005). Postsynaptic depolarization requirements for LTP and LTD: a critique of spike timing-dependent plasticity. *Nat. Neurosci.* 8, 839–841. doi: 10.1038/nn0705-839
- Lisman, J. E. (1989). A mechanism for the Hebb and the anti-Hebb processes underlying learning and memory. *Proc. Natl. Acad. Sci. U.S.A.* 86, 9574–9578. doi: 10.1073/pnas.86.23.9574
- Lovatt, D., Xu, Q., Liu, W., Takano, T., Smith, N. A., Schnerrmann, J., et al. (2012). Neuronal adenosine release, and not astrocytic ATP release, mediates feedback inhibition of excitatory activity. *Proc. Natl. Acad. Sci. U.S.A.* 109, 6265–6270. doi: 10.1073/pnas.1120997109
- Luz, Y., and Shamir, M. (2012). Balancing feed-forward excitation and inhibition via Hebbian inhibitory synaptic plasticity. *PLoS Comput. Biol.* 8:e1002334. doi: 10.1371/journal.pcbi.1002334
- Lynch, G. S., Dunwiddie, T., and Gribkoff, V. (1977). Heterosynaptic depression: a postsynaptic correlate of long-term potentiation. *Nature* 266, 737–739. doi: 10.1038/266737a0
- Malenka, R. C., and Bear, M. F. (2004). LTP and LTD: an embarrassment of riches. *Neuron* 44, 5–21. doi: 10.1016/j.neuron.2004.09.012
- Markram, H., and Tsodyks, M. (1996). Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature* 382, 807–810. doi: 10.1038/382807a0
- Markram, H., Wang, Y., and Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *Proc. Natl. Acad. Sci. U.S.A.* 95, 5323–5328. doi: 10.1073/pnas.95.9.5323
- Merzenich, M. M., Knight, P. L., and Roth, G. L. (1975). Representation of cochlea within primary auditory cortex in the cat. *J. Neurophysiol.* 38, 231–249.
- Miller, K. D. (1996). Synaptic economics: competition and cooperation in synaptic plasticity. *Neuron* 17, 371–374. doi: 10.1016/S0896-6273(00)80169-5
- Miller, K. D., and MacKay, D. J. C. (1994). The role of constraints in Hebbian learning. *Neural Comput.* 6, 100–126. doi: 10.1162/neco.1994.6.1.100
- Morrison, A., Aertsen, A., and Diesmann, M. (2008). Spike-timing-dependent plasticity in balanced random networks. *Neural Comput.* 19, 1437–1467. doi: 10.1162/neco.2007.19.6.1437
- Murthy, V. N., Schikorski, T., Stevens, C. F., and Zhu, Y. (2001). Inactivity produces increase in neurotransmitter release and synapse size. *Neuron* 18, 599–612. doi: 10.1016/S0896-6273(00)80301-3
- Neveu, D., and Zucker, R. S. (1996a). Postsynaptic level of $[Ca^{2+}]_i$ needed to trigger LTD and LTP. *Neuron* 16, 619–629. doi: 10.1016/S0896-6273(00)80081-1
- Neveu, D., and Zucker, R. S. (1996b). Long-lasting potentiation and depression without presynaptic activity. *J. Neurophysiol.* 75, 2157–2160.
- Nevian, T., and Sakmann, B. (2006). Spine Ca^{2+} signaling in spike-timing-dependent plasticity. *J. Neurosci.* 26, 11001–11013. doi: 10.1523/JNEUROSCI.1749-06.2006
- Nishiyama, M., Hong, K., Mikoshiba, K., Poo, M., and Kato, K. (2000). Calcium stores regulate the polarity and input specificity of synaptic modification. *Nature* 408, 584–588. doi: 10.1038/35046067
- Nugent, F. S., Penick, E. C., and Kauer, J. A. (2007). Opioids block long-term potentiation of inhibitory synapses. *Nature* 446, 1086–1090. doi: 10.1038/nature05726
- Oja, E. (1982). A simplified neuron model as a principal component analyzer. *J. Math. Biol.* 15, 267–273. doi: 10.1007/BF00275687
- Okun, M., and Lampl, I. (2008). Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat. Neurosci.* 11, 535–537. doi: 10.1038/nn.2105
- Ozeki, H., Finn, I. M., Schaffer, E. S., Miller, K. D., and Ferster, D. (2009). Inhibitory stabilization of the cortical network underlies visual surround suppression. *Neuron* 62, 578–592. doi: 10.1016/j.neuron.2009.03.028
- Pascual, O., Casper, K. B., Kubera, C., Zhang, J., Revilla-Sanchez, R., Sul, J. Y., et al. (2005). Astrocytic purinergic signaling coordinates synaptic networks. *Science* 310, 113–116. doi: 10.1126/science.1116916
- Peters, A., Schweiger, U., Pellerin, L., Hubold, C., Oltmanns, K. M., Conrad, M., et al. (2004). The selfish brain: competition for energy resources. *Neurosci. Biobehav. Rev.* 28, 143–180. doi: 10.1016/j.neubiorev.2004.03.002

- Pitler, T. A., and Alger, B. E. (1992). Postsynaptic spike firing reduces synaptic GABAA responses in hippocampal pyramidal cells. *J. Neurosci.* 12, 4122–4132.
- Pockett, S., Brookes, N. H., and Bindman, L. J. (1990). Long-term depression at synapses in slices of rat hippocampus can be induced by bursts of postsynaptic activity. *Exp. Brain Res.* 80, 196–200. doi: 10.1007/BF00228861
- Rabinowitch, I., and Segev, I. (2008). The interplay between homeostatic synaptic plasticity and functional dendritic compartments. *J. Neurophysiol.* 96, 276–283. doi: 10.1152/jn.00074.2006
- Remy, S., and Spruston, N. (2007). Dendritic spikes induce single-burst long-term potentiation. *Proc. Natl. Acad. Sci. U.S.A.* 104, 17192–17197. doi: 10.1073/pnas.0707919104
- Rich, M. M., and Wenner, P. (2007). Sensing and expressing homeostatic synaptic plasticity. *Trends Neurosci.* 30, 119–125. doi: 10.1016/j.tins.2007.01.004
- Richardson, M. J. E., Melamed, O., Silberberg, G., Gerstner, W., and Markram, H. (2005). Short-term synaptic plasticity orchestrates the response of pyramidal cells and interneurons to population bursts. *J. Comput. Neurosci.* 18, 323–331. doi: 10.1007/s10827-005-0434-8
- Royer, S., and Paré, D. (2003). Conservation of total synaptic weight through balanced synaptic depression and potentiation. *Nature* 422, 518–522. doi: 10.1038/nature01530
- Scanziani, M., Capogna, M., Gähwiler, B. H., and Thompson, S. M. (1992). Presynaptic inhibition of miniature excitatory synaptic currents by baclofen and adenosine in the hippocampus. *Neuron* 9, 919–927. doi: 10.1016/0896-6273(92)90244-8
- Schulz, P. E. (1997). Long-term potentiation involves increases in the probability of neurotransmitter release. *Proc. Natl. Acad. Sci. U.S.A.* 94, 5888–5893. doi: 10.1073/pnas.94.11.5888
- Schuman, E. M., and Madison, D. V. (1994). Locally distributed synaptic potentiation in the hippocampus. *Science* 263, 532–536. doi: 10.1126/science.8290963
- Sehgal, M., Song, C., Ehlers, V. L., and Moyer, J. R. Jr. (2013). Learning to learn - intrinsic plasticity as a metaplasticity mechanism for memory formation. *Neurobiol. Learn. Mem.* 105, 186–199. doi: 10.1016/j.nlm.2013.07.008
- Sjöström, P. J., and Häusser, M. (2006). A cooperative switch determines the sign of synaptic plasticity in distal dendrites of neocortical pyramidal neurons. *Neuron* 51, 227–238. doi: 10.1016/j.neuron.2006.06.017
- Sjostrom, P. J., Rancz, E. A., Roth, A., and Häusser, M. (2007). Dendritic excitability and synaptic plasticity. *Physiol. Rev.* 88, 769–840. doi: 10.1152/physrev.00016.2007
- Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2001). Rate, timing and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32, 1149–1164. doi: 10.1016/S0896-6273(01)00542-6
- Skorheim, S., Lonjers, P., and Bazhenov, M. (2014). A spiking network model of decision making employing rewarded STDP. *PLoS ONE* 9:e90821. doi: 10.1371/journal.pone.0090821
- Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.* 3, 919–926. doi: 10.1038/78829
- Song, S., Sjostrom, P. J., Reigl, M., Nelson, S., and Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol.* 3:e68. doi: 10.1371/journal.pbio.0030068
- Spruston, N., Schiller, Y., Stuart, G., and Sakmann, B. (1995). Activity-dependent action potential invasion and calcium influx into hippocampal CA1 dendrites. *Science* 268, 297–300. doi: 10.1126/science.7716524
- Staubli, U. V., and Ji, Z. X. (1996). The induction of homo- vs. heterosynaptic LTD in area CA1 of hippocampal slices from adult rats. *Brain Res.* 714, 169–176. doi: 10.1016/0006-8993(95)01523-X
- Sun, Y. J., Wu, G. K., Liu, B. H., Li, P., Zhou, M., Xiao, Z., et al. (2010). Fine-tuning of pre-balanced excitation and inhibition during auditory cortical development. *Nature* 465, 927–931. doi: 10.1038/nature09079
- Sussillo, D., Toyozumi, T., and Maass, W. (2007). Self-tuning of neural circuits through short-term synaptic plasticity. *J. Neurophysiol.* 97, 4079–4095. doi: 10.1152/jn.01357.2006
- Thiagarajan, T. C., Lindskog, M., and Tsien, R. W. (2005). Adaptation to synaptic inactivity in hippocampal neurons. *Neuron* 47, 725–737. doi: 10.1016/j.neuron.2005.06.037
- Thompson, I. D., Kossut, M., and Blakemore, C. (1983). Development of orientation columns in cat striate cortex revealed by 2-deoxyglucose autoradiography. *Nature* 301, 712–715. doi: 10.1038/301712a0
- Thompson, S. M., Haas, H. L., and Gähwiler, B. H. (1992). Comparison of the actions of adenosine at pre- and postsynaptic receptors in the rat hippocampus *in vitro*. *J. Physiol.* 451, 347–363. doi: 10.1113/jphysiol.1992.sp019168
- Tsodyks, M. V., and Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci. U.S.A.* 94, 719–723. doi: 10.1073/pnas.94.2.719
- Turrigiano, G. G. (2011). Too many cooks? Intrinsic and synaptic homeostatic mechanisms in cortical circuit refinement. *Annu. Rev. Neurosci.* 34, 89–103. doi: 10.1146/annurev-neuro-060909-153238
- Turrigiano, G. G. (2012). Homeostatic synaptic plasticity: local and global mechanisms for stabilizing neuronal function. *Cold Spring Harb. Perspect. Biol.* 4:a005736. doi: 10.1101/cshperspect.a005736
- Turrigiano, G. G., and Nelson, S. B. (2004). Homeostatic plasticity in the developing nervous system. *Nat. Rev. Neurosci.* 5, 97–107. doi: 10.1038/nrn1327
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., and Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391, 892–896. doi: 10.1038/36103
- van Ooyen, A. (2001). Competition in the development of nerve connections: a review of models. *Network* 12, R1–R47. doi: 10.1080/net.12.1.1.47
- van Rossum, M. C., Bi, G. Q., and Turrigiano, G. G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neurosci.* 20, 8812–8821.
- van Vreeswijk, C., and Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274, 1724–1726. doi: 10.1126/science.274.5293.1724
- Varela, J., Sen, K., Gibson, J., Fost, J., Abbott, L. F., and Nelson, S. (1997). A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex. *J. Neurosci.* 17, 7926–7940.
- Varela, J. A., Song, S., Turrigiano, G. G., and Nelson, S. B. (1999). Differential depression at excitatory and inhibitory synapses in visual cortex. *J. Neurosci.* 19, 4293–4304.
- Vitureira, N., Letellier, M., and Goda, Y. (2012). Homeostatic synaptic plasticity: from single synapses to neural circuits. *Curr. Opin. Neurobiol.* 22, 516–521. doi: 10.1016/j.conb.2011.09.006
- Vlachos, A., Helias, M., Becker, D., Diesmann, M., and Deller, T. (2013). NMDA-receptor inhibition increases spine stability of denervated mouse dentate granule cells and accelerates spine density recovery following entorhinal denervation *in vitro*. *Neurobiol. Dis.* 59, 267–276. doi: 10.1016/j.nbd.2013.07.018
- Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* 334, 1569–1573. doi: 10.1126/science.1211095
- Volgushev, M., Chistiakova, M., Balaban, P., and Eysel, U. T. (2000). Retrograde signaling with nitric oxide at neocortical synapses. *Eur. J. Neurosci.* 12, 4255–4267. doi: 10.1046/j.0953-816X.2000.01322.x
- Volgushev, M., Mittmann, T., Chistiakova, M., Balaban, P., and Eysel, U. T. (1999). Interaction between intracellular tetanization and pairing-induced long-term synaptic plasticity in the rat visual cortex. *Neuroscience* 93, 1227–1232. doi: 10.1016/S0306-4522(99)00265-1
- Volgushev, M., Pernberg, J., and Eysel, U. T. (2003). Gamma-frequency fluctuations of the membrane potential and response selectivity in visual cortical neurons. *Eur. J. Neurosci.* 17, 1768–1776. doi: 10.1046/j.1460-9568.2003.02609.x
- Volgushev, M., Voronin, L. L., Chistiakova, M., and Singer, W. (1994). Induction of LTP and LTD in visual cortex neurons by intracellular tetanization. *Neuroreport* 5, 2069–2072. doi: 10.1097/00001756-199410270-00020
- Volgushev, M., Voronin, L. L., Chistiakova, M., and Singer, W. (1997). Relations between long-term synaptic modifications and paired-pulse interactions in the rat neocortex. *Eur. J. Neurosci.* 9, 1656–1665. doi: 10.1111/j.1460-9568.1997.tb01523.x
- Volman, V., Sejnowski, T. J., and Bazhenov, M. (2011a). Topological basis of epileptogenesis in a model of severe cortical trauma. *J. Neurophysiol.* 106, 1933–1942. doi: 10.1152/jn.00458.2011
- Volman, V., Bazhenov, M., and Sejnowski, T. J. (2011b). Pattern of trauma determines the threshold for epileptic activity in a model of cortical

- deafferentation. *Proc. Natl. Acad. Sci. U.S.A.* 108, 15402–15407. doi: 10.1073/pnas.1112066108
- von der Malsburg, C. (1973). Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik* 14, 85–100. doi: 10.1007/BF00288907
- Wall, M., and Dale, N. (2008). Activity-dependent release of adenosine: a critical re-evaluation of mechanism. *Curr. Neuropharmacol.* 6, 329–337. doi: 10.2174/157015908787386087
- Waters, J., Larkum, M., Sakmann, B., and Helmchen, F. (2003). Supralinear Ca²⁺ influx into dendritic tufts of layer 2/3 neocortical pyramidal neurons *in vitro* and *in vivo*. *J. Neurosci.* 23, 8558–8567.
- Watt, A. J., and Desai, N. S. (2010). Homeostatic plasticity and STDP: keeping a neuron's cool in a fluctuating world. *Front. Synaptic Neurosci.* 2:5. doi: 10.3389/fnsyn.2010.00005
- Watt, A. J., Sjöström, P. J., Häusser, M., Nelson, S. B., and Turrigiano, G. G. (2004). A proportional but slower NMDA potentiation follows AMPA potentiation in LTP. *Nat. Neurosci.* 7, 518–524. doi: 10.1038/nn1220
- Wehr, M., and Zador, A. M. (2003). Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* 426, 442–446. doi: 10.1038/nature02116
- White, G., Levy, W. B., and Steward, O. (1990). Spatial overlap between populations of synapses determines the extent of their associative interaction during the induction of long-term potentiation and depression. *J. Neurophysiol.* 64, 1186–1198.
- Wiesel, T. N., and Hubel, D. H. (1963). Effects of visual deprivation on morphology and physiology of cells in the cats lateral geniculate body. *J. Neurophysiol.* 26, 978–993.
- Wilson, R. I., and Nicoll, R. A. (2001). Endogenous cannabinoids mediate retrograde signalling at hippocampal synapses. *Nature* 410, 588–592. doi: 10.1038/35069076
- Wöhrl, R., von Haebler, D., and Heinemann, U. (2007). Low-frequency stimulation of the direct cortical input to area CA1 induces homosynaptic LTD and heterosynaptic LTP in the rat hippocampal-entorhinal cortex slice preparation. *Eur. J. Neurosci.* 25, 251–258. doi: 10.1111/j.1460-9568.2006.05274.x
- Wu, Z., and Yamaguchi, Y. (2006). Conserving total synaptic weight ensures one-trial sequence learning of place fields in the hippocampus. *Neural Netw.* 19, 547–563. doi: 10.1016/j.neunet.2005.06.048
- Yang, S. N., Tang, Y. G., and Zucker, R. S. (1999). Selective induction of LTP and LTD by postsynaptic [Ca²⁺]_i elevation. *J. Neurophysiol.* 81, 781–787.
- Yeung, L. C., Shouval, H. Z., Blais, B. S., and Cooper, L. N. (2004). Synaptic homeostasis and input selectivity follow from a calcium-dependent plasticity model. *Proc. Natl. Acad. Sci. U.S.A.* 101, 14943–14948. doi: 10.1073/pnas.0405555101
- Zenke, F., Hennequin, G., and Gerstner, W. (2013). Synaptic plasticity in neural networks needs homeostasis with a fast rate detector. *PLoS Comput. Biol.* 9:e1003330. doi: 10.1371/journal.pcbi.1003330
- Zhang, J., Ackman, J. B., Xu, H. P., and Crair, M. C. (2011). Visual map development depends on the temporal pattern of binocular activity in mice. *Nat. Neurosci.* 15, 298–307. doi: 10.1038/nn.3007
- Zhang, P., Bannon, N. M., Ilin, V., Volgushev, M., and Chistiakova, M. (2015). Adenosine effects on inhibitory synaptic transmission and excitation-inhibition balance in the rat neocortex. *J. Physiol.* 593, 825–841. doi: 10.1113/jphysiol.2014.279901
- Zhang, W., and Linden, D. J. (2003). The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nat. Rev. Neurosci.* 4, 885–900. doi: 10.1038/nrn1248
- Zhou, Y. D., Acker, C. D., Netoff, T. I., Sen, K., and White, J. A. (2005). Increasing Ca²⁺ transients by broadening postsynaptic action potentials enhances timing-dependent synaptic depression. *Proc. Natl. Acad. Sci. U.S.A.* 102, 19121–19125. doi: 10.1073/pnas.0509856103
- Zucker, R. S., and Regehr, W. G. (2002). Short-term synaptic plasticity. *Annu. Rev. Physiol.* 64, 355–405. doi: 10.1146/annurev.physiol.64.092501.114547

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2015 Chistiakova, Bannon, Chen, Bazhenov and Volgushev. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.



Enhanced polychronization in a spiking network with metaplasticity

Mira Guise, Alistair Knott* and Lubica Benuskova

Department of Computer Science, University of Otago, Dunedin, New Zealand

Edited by:

Cristina Savin, Institute of Science and Technology Austria, Austria

Reviewed by:

Matthieu Gilson, Universitat

Pompeu Fabra, Spain

Daniel Carl Miner, Frankfurt Institute for Advanced Studies, Germany

*Correspondence:

Alistair Knott, Department of Computer Science, University of Otago, Owheo Building, 133 Union Street East, Dunedin 9016, New Zealand
e-mail: alik@cs.otago.ac.nz

Computational models of metaplasticity have usually focused on the modeling of single synapses (Shouval et al., 2002). In this paper we study the effect of metaplasticity on network behavior. Our guiding assumption is that the primary purpose of metaplasticity is to *regulate synaptic plasticity*, by increasing it when input is low and decreasing it when input is high. For our experiments we adopt a model of metaplasticity that demonstrably has this effect for a single synapse; our primary interest is in how metaplasticity thus defined affects network-level phenomena. We focus on a network-level phenomenon called *polychronicity*, that has a potential role in representation and memory. A network with polychronicity has the ability to produce non-synchronous but precisely timed sequences of neural firing events that can arise from strongly connected groups of neurons called polychronous neural groups (Izhikevich et al., 2004). Polychronous groups (PNGs) develop readily when spiking networks are exposed to repeated spatio-temporal stimuli under the influence of spike-timing-dependent plasticity (STDP), but are sensitive to changes in synaptic weight distribution. We use a technique we have recently developed called Response Fingerprinting to show that PNGs formed in the presence of metaplasticity are significantly larger than those with no metaplasticity. A potential mechanism for this enhancement is proposed that links an inherent property of integrator type neurons called spike latency to an increase in the tolerance of PNG neurons to jitter in their inputs.

Keywords: metaplasticity, STDP, spiking network, polychronous neural group, memory, spike latency, synaptic weight, synaptic drive

1. INTRODUCTION

The term *metaplasticity* describes the ability of neurons to modulate their overall levels of synaptic plasticity as a function of recent inputs. Models of LTP/LTD induction that include a metaplastic mechanism have been around for some time, with the Bienenstock, Cooper and Munro (BCM) learning rule and its sliding modification threshold being a significant early influence (Bienenstock et al., 1982). The BCM rule provides a rate-dependent model of the tipping point between LTD and LTP induction based on instantaneous neural firing rates. More recently, Izhikevich and Desai (2003) have combined the BCM rule with spike-timing-dependent plasticity (STDP), a learning rule that takes the precise spike timing of pre- and post-synaptic neurons into account. The addition of a BCM sliding modification threshold to an STDP learning rule has also been used to explain experimental data showing hetero-synaptic LTD of untetanized inputs in a model of a hippocampal dentate granule cell (Benuskova and Abraham, 2007). The precise mechanism behind metaplasticity is still an open question despite receiving much recent attention (for a review see Abraham, 2008). Intensive research into the cellular processes behind metaplasticity has uncovered multiple mechanisms that both cooperate and compete, with the balance between the various mechanisms varying between different brain regions.

Models based on BCM define a *modification threshold* for LTP/LTD induction that is dynamically altered as a function of previous post-synaptic spike activity. When spiking activity increases, the modification threshold is also increased and it therefore becomes harder to induce subsequent LTP and easier to induce LTD. A decrease in spiking activity produces the opposite effect, with LTP induction becoming easier and LTD induction becoming more difficult. The BCM learning rule defines a single modification threshold, while later versions have defined separate thresholds for LTP and LTD (Ngezahayo et al., 2000). In one example of a single threshold model (taken from Benuskova and Abraham, 2007), the relationship between the modification threshold and synaptic change is defined as follows:

$$\begin{aligned}A_{LTP}(t) &= A_{LTP}(0)(1/\theta_M(t)) \\ A_{LTD}(t) &= A_{LTD}(0)(\theta_M(t))\end{aligned}\quad (1)$$

where $\theta_M(t)$ is the current value of the modification threshold, $A_{LTP}(0)$ and $A_{LTD}(0)$ are the baseline amplitudes of synaptic change, and $A_{LTP}(t)$ and $A_{LTD}(t)$ are the new amplitudes. Typically, these models assume that the metaplastic modification threshold is determined primarily by the post-synaptic firing rate (e.g., Beňušková et al., 2001), although this assumption is still

open to debate (Hulme et al., 2012). Shouval et al. (2002) suggest that the modification threshold is more directly set by the levels of intracellular Ca^{2+} while (Izhikevich and Desai, 2003) suggest that synaptic size might also be an influence.

In a synapse whose learning is governed by a spike-timing-dependent plasticity (STDP) rule, the direction and magnitude of neural plasticity is determined not only by factors that govern the level of synaptic input, but also by the precise timing of pre-synaptic and post-synaptic spikes. Changes in synaptic plasticity cannot therefore be predicted from either the post-synaptic firing rate or the total synaptic input alone (Izhikevich and Desai, 2003). In this scenario, the conditions under which the modification threshold should be modified relate to the consistency of these timings in the recent history of the synapse. We use the term *synaptic drive* to describe these conditions: strongly correlated spike trains with pre- before post-synaptic spike timings are defined as producing a positive drive on synaptic plasticity, whilst post- before pre-synaptic spike trains with identical firing rates are defined as having a negative synaptic drive.

In the current paper we define a model of metaplasticity that is determined by the direction and magnitude of synaptic drive, and also by the size of each of the synaptic connections onto the cell (Delorme et al., 2001; Guetig et al., 2003). The model is therefore both spike-timing-dependent and reactive to synaptic weight extremes i.e., it resists synaptic pruning and opposes synaptic weights that grow too large. We have chosen to define the metaplastic modification threshold in this model as a cell-level property that integrates the changes in plasticity that are occurring at each synapse. The choice of a cell-level property, rather than defining a modification threshold at each synapse, allows the metaplasticity model to be integrated into a larger model of network behavior and is supported by a recent finding that metaplastic effects can be seen in non-primed dendritic compartments (Hulme et al., 2012). Previous computational models of metaplasticity have typically focused on the modeling of single synapses, although reports on the effect of metaplasticity at network-level have recently started to appear (Clopath et al., 2010; Zenke et al., 2013). Like any computational model of the synapse, the model of metaplasticity we use in our experiments is motivated by a mixture of mechanistic and computational considerations. Some components in the model aim to account for specific empirically identified biological mechanisms in the synapse. Other components are included to implement a particular theoretical claim about the *function* of metaplasticity—namely that it serves to regulate synaptic plasticity, by increasing it when input is low, and decreasing it when it is high (see e.g., Hulme et al., 2012; Zenke et al., 2013). In both cases, the model we use is heavily based on existing models of synaptic plasticity, though it also includes novel mechanistic and novel functional components.

The primary focus of our study is the impact of metaplasticity on an empirically observed property of spiking neural networks called *polychronicity*. Polychronous neural groups (or PNGs) are connected groups of neurons that can be activated together to produce polychronization, a non-synchronous but precisely timed sequence of neural firing events (Izhikevich et al., 2004). These stimulus-specific firing signatures form reproducible patterns that are observable in the firing data generated by the

network. Polychronization requires that the connection weights between PNG neurons be *adapted* to support a sequential chain of neural firing (Martinez and Paugam-Moisy, 2009). With an STDP learning rule, this adaptation occurs readily when spiking neural networks are exposed to repeated spatio-temporal stimuli. The STDP rule combined with repeated stimulation potentiates intra-group connection weights and prunes non-contributing connections, leading to the preferential selection of stimulus-dependent polychronous groups (Izhikevich et al., 2004).

Given that polychronous groups evolve via selective enhancement of the connections between PNG neurons, it is often assumed that the stability of adapted PNGs over an extended period requires that these same connections be maintained. However, polychronicity requires only that the combined input to PNG neurons be sufficient to produce firing within a precise temporal window. Theoretically therefore, PNG neurons can remain stable within the group even if the weight value on some of their afferent connections wanders randomly, so long as other input connections evolve their weights to compensate. This proposed independence of PNG stability from the weight values of specific synapses leaves the weights free to support other aspects of the network dynamics such as competing or co-activating PNGs, the maintenance of the balance between excitation and inhibition (van Vreeswijk and Sompolinsky, 1996; Vogels et al., 2005), and *mixture states* of synchronization and desynchronization in the network firing activity (Lubenov and Siapas, 2008).

Evidence for polychronicity in biological networks has been technically difficult to establish, although precise spatio-temporal firing patterns observed in rat and monkey cortical neurons provide some supporting evidence (Villa et al., 1999; Shmiel et al., 2006). However, in simulated networks the process of isolating structural PNGs or detecting PNG activation is straightforward (e.g., Izhikevich et al., 2004; Martinez and Paugam-Moisy, 2009). We use a recently developed technique called Response Fingerprinting to test whether polychronicity both persists and is stable within our model metaplastic regime (Guise et al., 2014).

A modified STDP rule that includes a metaplastic mechanism is likely to have a significant effect on PNG formation and may also be more biologically plausible than existing STDP rules. Lazar et al. (2007) report improvements in both network performance and stability using a combination of intrinsic plasticity with STDP to produce a reduction in synaptic saturation. A metaplastic modification to the STDP rule has the potential to maintain synaptic weights more centrally in the range and may therefore produce a similar performance advantage. However, the formation of polychronous groups has a significant effect on synaptic weights, resulting in a characteristic bimodal weight distribution that opposes this predicted centralizing effect. Polychronizing pathways are very dependent on strong connections that support convergent input to PNG neurons, and therefore any network mechanism that affects the synaptic weight distribution is predicted to have a significant effect on PNG formation. Given the opposing effects of PNG formation and metaplasticity on synaptic weight distributions, it is not clear whether PNG formation will be supported in networks with the new metaplastic mechanism, and if it is supported, what the effect will be on PNG size.

2. METHODS

2.1. METAPLASTICITY MODEL

2.1.1. Methodological preliminaries

As mentioned in the introduction, the model of metaplasticity we implemented in our experiments was designed to accommodate a mixture of mechanistic and computational considerations. The computational considerations are uppermost: we assume that the key purpose of the metaplastic mechanism we are modeling is to *regulate synaptic plasticity*, by limiting the range of weights within any given synapse, forcing weights away from both their upper and lower extremes. Accordingly, a key design goal for our computational model is that it produces this effect. At the same time, we want the model to make as much reference as possible to empirically identified mechanisms in the synapse, so that the regulatory effect can be linked as much as possible to physiological processes. On both counts our model draws heavily on existing models of metaplasticity. A useful point of reference is the model of metaplasticity of Zenke et al. (2013). Like our model, this model implements an assumption that the main purpose of metaplasticity is to regulate synaptic plasticity. However, our model incorporates a slightly different set of mechanistic components to achieve this effect. We will draw attention to these differences as the model is introduced.

One difference to mention straight away is that our model defines a metaplastic modification threshold that is a neuron-level property: the threshold value for a given neuron is computed from a weighted average of the threshold values of its afferent synapses. The model is therefore best described in two sections: a synapse-level model that weights the size of each synapse according to the current direction and magnitude of synaptic change (synaptic drive); and a neuron-level model that is computed as a weighted sum of the individual synaptic values.

2.1.2. Synapse-level model

The metaplasticity model at the level of each individual synapse is defined by a weighting function that computes a weighted value for each synapse. The weighting function takes values representing the current synaptic weight and synaptic drive as arguments, and returns a weighted value representing the resistance to synaptic weight change. The synaptic drive is dependent on both the level of synaptic input and the precise timing of that input relative to a back-propagating dendritic signal. In our simulated network, we approximate the synaptic drive with a synaptic derivative, an instantaneous measure of the direction and magnitude of change at each synapse that is an explicit value from the original network simulation code (Izhikevich, 2006b).

The desired weighting function needs to exert little influence when synaptic weights are within bounds, but must step in with increasing resistance as the weights approach the upper or lower weight limits. One possibility, that we use throughout this paper, is as follows:

$$f(d_i, w_i) = re^{p(map(d_i))(w_i - min)} - re^{p(10 - map(d_i))(max - w_i)} \quad (2)$$

where:

p = precision

r = resistance

min = minimum soft limit

max = maximum soft limit

d_i = derivative of synapse i , $d_i = \lim_{\Delta t \rightarrow 0} \frac{\Delta w_i}{\Delta t}$

w_i = weight of synapse i

$$map(x) = \begin{cases} 0 & \text{if } 0.5(x + 10) < 0 \\ 10 & \text{if } 0.5(x + 10) > 10 \\ 0.5(x + 10) & \text{if } 0 \leq 0.5(x + 10) \leq 10 \end{cases}$$

The *map* function maps the normal range of synaptic derivative values (-10 to $+10$) into the range $0-10$, and clips values outside of this normal range. The precision (p) and resistance (r) parameters control the curvature and amplitude of the function. The weight limits are determined by the parameters *min* and *max* that specify uncapped *soft* limits rather than capped hard limits on synaptic weights. **Figure 1A** shows the full picture for both parameters (d_i and w_i) over the weight range $0-10$ and limiting the range of the synaptic derivative to ± 10 . Most combinations of d_i and w_i generate a weighting term that is close to zero, and the resulting surface is therefore largely flat for these combinations. However, the surface exponentially rises and falls in opposing corners, producing maximum resistance to increases in synaptic weight when the synaptic drive is positive and the weight is already large, and maximum resistance to decreases in synaptic weight when the synaptic drive is negative and the weight is already small. However, the function generates little resistance to large weights if the synaptic drive is negative, or to small weights if the synaptic drive is positive, providing no impediment to migration of synaptic weights away from the weight limits.

2.1.3. Neuron-level model

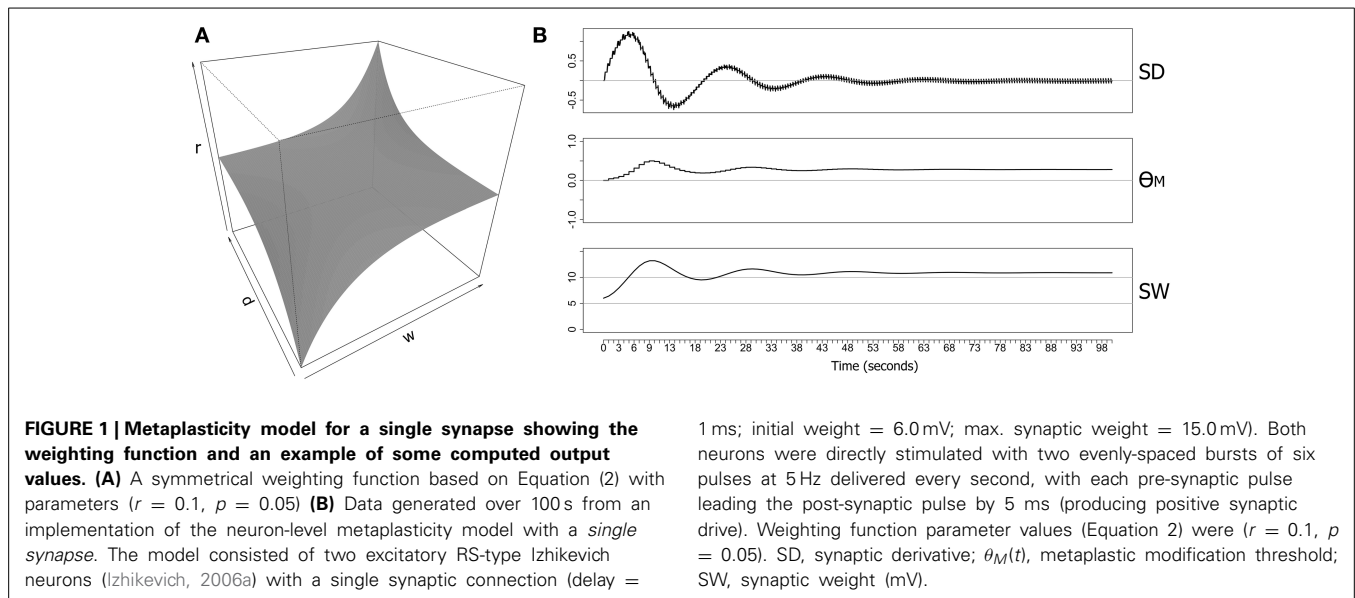
The *modification threshold* (θ_M) in this model is a neuron-level property that determines the ease of subsequent synaptic change. Unlike previous models, the modification threshold is dependent on synaptic drive and therefore only indirectly dependent on the post-synaptic spike rate or the total input. When synaptic drive is strongly positive, θ_M increases, and subsequent LTP induction becomes harder (and LTD induction becomes easier). A negative synaptic drive produces the opposite effect by causing $\theta_M(t)$ to decrease.

The modification threshold is computed as a range-limited average of the weighting function output for each of the afferent synaptic connections. The weighting function outputs represent a *weighted synaptic derivative* for each synapse, and the average therefore represents the integrated synaptic drive across all synaptic inputs. The weighting mechanism assumes that a global metaplastic signal interacts with the local conditions (particularly the synaptic size) at each synapse.

Given a weighting function $f(d_i, w_i)$, the modification threshold is computed as follows:

$$\theta_M(t) = \tanh \left(I \frac{\sum_{i=1}^n f(d_i, w_i)}{n} \right) \quad (3)$$

where:



$I = \text{inertia}$

$n = \text{number of synapses}$

$f(d_i, w_i) = \text{a weighting function}$

The hyperbolic tangent limits the range of the modification threshold to ± 1 , while the *inertia* parameter controls the rate of change of θ_M i.e., the sensitivity of the modification threshold to small changes in the average weighted synaptic derivative. Here we are assuming a biophysical process that maps a wide input range into a narrower response range such as in the proposed power-law relationships between stimulus strength and perceived intensity (MacKay, 1963).

Equations (2) and (3) are novel elements in a model of metaplasticity, in that they assume a generalized postsynaptic activity function rather than the ‘spike counter’ assumed by existing models (see e.g., Benuskova and Abraham, 2007; Zenke et al., 2013). However, this departure is justified in the light of recent evidence that synaptic plasticity can be homeostatically regulated by the cell-wide history of synaptic activity through a calcium-dependent but action potential-independent mechanism (Hulme et al., 2012).

Given the modification threshold, the amplitudes of synaptic change in LTP and LTD can now be calculated as follows:

$$A_{LTP}(t) = A_{LTP}(0) - (A_{LTP}(0)\theta_M(t)) \quad (4)$$

$$A_{LTD}(t) = A_{LTD}(0) + (A_{LTD}(0)\theta_M(t)) \quad (5)$$

The two equations in (4) are symmetrical by design: if $\theta_M(t)$ is positive then the LTP amplitude (A_{LTP}) decreases, and the LTD amplitude (A_{LTD}) increases by the same proportion. Unlike the equation described by (Benuskova and Abraham, 2007) (see Equation 1), the LTP and LTD amplitudes in Equation (4) are modified in direct proportion to the modification threshold (θ_M) and the current baseline amplitudes. In contrast, each of the amplitudes of synaptic change in the equation of Benuskova and

Abraham (2007) differ in their relationship to the modification threshold: A_{LTP} is inversely proportional to θ_M , while A_{LTD} is directly proportional. If spike activity is low but consistent, the equation of Benuskova and Abraham (2007) has the potential to create a dramatic imbalance between A_{LTP} and A_{LTD} that allows synaptic weight to increase without limit. Clopath et al. (2010) and Zenke et al. (2013) introduced models in which only the LTD amplitude is metaplastically modified and the LTP amplitude stays constant. However, we do not have any neurobiological evidence why only LTD would be subject to homeostatic control and LTP not, therefore we assume that magnitudes of both are metaplastically modified.

Metaplasticity models based on post-synaptic spike rate restrict the modification threshold to positive values. However, in the current model (based on synaptic drive) and in models based on post-synaptic membrane potential, both positive and negative values are allowed (Ngezhahayo et al., 2000). We limit the range of the modification threshold to ± 1 (above) so that both the LTP and LTD amplitudes have the range 0–2 times the baseline amplitudes, and have default values of $A_{LTP}(0)$ and $A_{LTD}(0)$ respectively. Although there is no explicit limit on synaptic growth in Equation (4), the symmetry between the equations for LTP and LTD limits the degree of imbalance between the two.

It is worth emphasizing that because the metaplastic modification threshold is calculated from an average of the values returned by the weighting function, the resistance to synaptic weights that near the limits also applies *only on average*. Therefore, individual synapses are allowed to grow without limit so long as the average across all synaptic inputs is within the allowed weight range. Synaptic pruning can still therefore occur, even if the value for weight limit resistance in Equation (2) is large. Likewise, while individual synapses are allowed to grow large, they will increasingly dominate the weighted average as they grow, providing an implicit limit to their growth.

2.2. NETWORK SIMULATIONS

2.2.1. Networks

For network simulations we use a spiking neural network platform we have developed called Spinula (Guise et al., 2013) that is based on the reference implementation from Izhikevich (2006b). Twenty different networks were generated for these experiments, with each network composed of 1000 Izhikevich neurons (800 excitatory RS type and 200 inhibitory FS type). Within each network each simulated neuron was connected to 100 randomly selected post-synaptic neurons with the restriction that inhibitory neurons were connected to excitatory neurons only. Inhibitory connections were assigned a 1 ms conduction delay while excitatory connections were randomly assigned a delay in the range 1–20 ms. Connection weights were initialized to the values +3.0 mV (for excitatory weights) and −2.0 mV (for inhibitory weights). Each network was then matured for 2 h by exposure to a 1 Hz random input under the influence of a spike-timing-dependent plasticity (STDP) rule. The STDP rule was temporally asymmetric and with parameters as in (Izhikevich, 2006a) i.e., $A_+ = 0.1$ and $A_- = 0.12$. Random input was generated by an independent Poisson process on each neuron.

2.2.2. Training

Networks were trained on a 5 Hz stimulus with a 1 Hz random input for 180 s (internal simulation time). Guise et al. (2014) have previously reported that PNG size reaches a plateau within around 2 min with this training protocol. Each stimulus was composed of forty firing events arranged in an ascending pattern (see **Figure 2** for an example of the *Ascending* pattern, and Guise et al., 2014, for further details). Metaplasticity-related parameters were $r = 0.1$; $p = 0.5$; inertia = 0.2; maximum synaptic weight = 10.0 mV (hard limit). Following training, synaptic weight distributions were generated from the saved synaptic weights for each network. The number of neurons participating in PNG activation was assessed by generating a Response Fingerprint for each network.

2.2.3. Response fingerprinting

The effect of the metaplasticity model on large networks of 100,000 synapses was examined using a technique we have recently developed called *Response Fingerprinting* (Guise et al., 2014, for implementation details see Guise et al., 2013a). A Response Fingerprint is a probabilistic representation of PNG activation that describes the spatio-temporal pattern of firing within a network in response to an input stimulus. It consists of a set of time windows within which specified neurons are likely to fire with empirically determined probabilities; information can be combined across time windows using Bayesian techniques to derive an aggregate estimate of the likelihood of the stimulus. The effect of metaplasticity on the ability of a network to polychronise was assessed by comparing the Response Fingerprints generated by the network with and without metaplasticity enabled. Response Fingerprints were generated by profiling the firing event data in the presence of a 1 Hz random background and identifying peaks in the histograms using a final consistency threshold of 0.75, a measure of the consistency of spiking within each peak region.

2.2.4. Connection activation

The presentation at fixed intervals of a known stimulus (one on which a network has been trained) produces a regular pattern of firing reflecting the activation of a PNG. The network connections can be partitioned into those that are regularly activated by the stimulus and those that are not, allowing an examination of the differential effect of metaplasticity on connections that participate in PNG activation vs. non-participating connections. The partitioning procedure involves attempting a *fit* for each of the 100,000 connections in each network to firing data generated from the network in response to the stimulus: for each connection and each pair of firing events in the firing data, we label the connection as *active* if $\text{connection length} \leq \text{time difference} \leq \text{connection length} + \text{jitter}$, otherwise the connection is labeled *non-active* i.e., if the time difference between firing events is longer

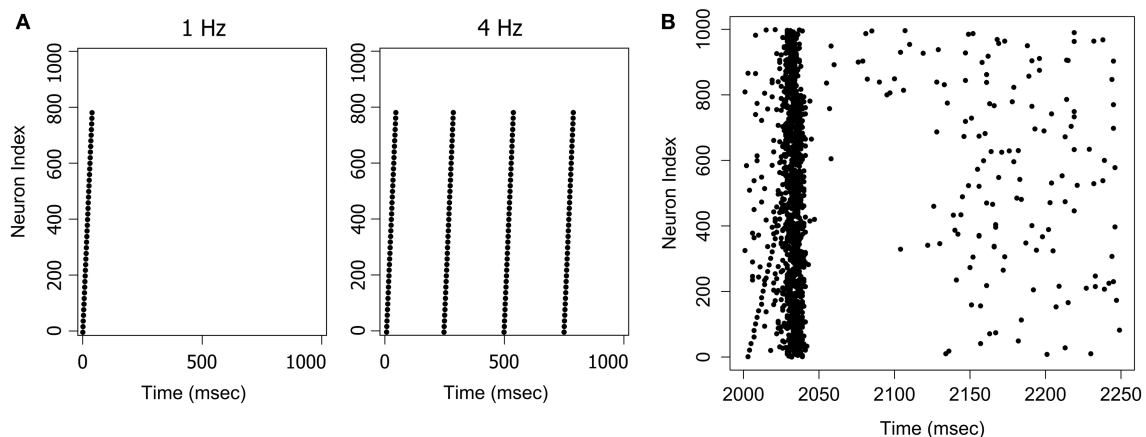


FIGURE 2 | Examples of stimuli and the stimulus response. (A) The *Ascending* pattern as a 1 Hz or 4 Hz stimulus. **(B)** The response to a 4 Hz stimulus. A network trained on the *Ascending* pattern was repeatedly presented with the same pattern at 4 Hz with a 1 Hz random background. The figure shows a randomly selected response frame between $t = 2000$ and

$t = 2250$ ms. The input pattern can also be seen as an ascending sequence of firing events in the first 40 ms of the frame. The network responds to this pattern with an avalanche-like burst of activity, that builds as the signal arrives, and then terminates quite suddenly when activity in the pool of inhibitory neurons reaches a critical threshold.

than the connection length by *some small amount* then the pre-synaptic spike was probably a contributor to the post-synaptic firing event and can be considered to be a part of the PNG activation. The allowed variation or *jitter* is typically set to 2 ms.

2.2.5. Input space response

The Input Space Response (ISR) of a neuron is produced by varying the firing times of each of the afferent neurons over a defined range and recording which of the resulting spatio-temporal patterns produces consistent firing of the post-synaptic neuron. This *input space* of potential firing patterns has the same dimension as the number of inputs to the target neuron, and the *active input space* is the subset of the input space that produces firing of the target. For example, neuron 4 in **Figure 8** has three input neurons. If the firing times of 1, 2 and 3 are systematically varied over the range 1–20 ms (keeping connection delays fixed) then the combination of all inputs produces a $20 \times 20 \times 20$ cube of spatio-temporal patterns, only some of which produce firing of neuron 4. With just three inputs the cube may conveniently be flattened to two dimensions by taking the difference between each pair of firing times i.e., $(t_1 - t_2)$ and $(t_2 - t_3)$, where (t_1, t_2, t_3) are the firing times of neurons 1, 2, and 3. Only two difference pairs are required as the remaining difference (t_1, t_3) is constrained by the other two. This 2D projection has the additional benefit of removing redundancies, as many of the patterns in the original cube are just shifted versions of the same spatio-temporal pattern.

3. RESULTS

The intention of the metaplasticity model was to force the synaptic weight values away from the extremes and toward the middle of the weight range. However, in networks with many afferent connections we might expect this effect to be diluted by the large number of synaptic inputs onto each neuron. Nevertheless, the metaplasticity model attempts to maintain a central weight for each synapse *on average*, and might therefore be expected to increase the number of non-saturated weights in the network. Significantly, this predicted effect opposes the bimodal weight distributions observed during PNG formation. It is unclear which of these effects will be stronger, the PNG-formation effect that moves synaptic weights toward the limits through STDP, or the metaplasticity effect that moves weights toward the center of the range.

3.1. OVERALL EFFECTS

3.1.1. Weight distributions

The results on twenty large networks of 100,000 synapses each are shown in **Figure 3**. Metaplasticity was found to have a significant effect on the distribution of excitatory synaptic weights in each network (inhibitory weights are non-plastic and are therefore not shown in **Figure 3**). For convenience, each of the eighty thousand excitatory connections was categorized into just one of three weight groups as follows: synaptic connections with zero weight (*pruned synapses*); connections with the maximum

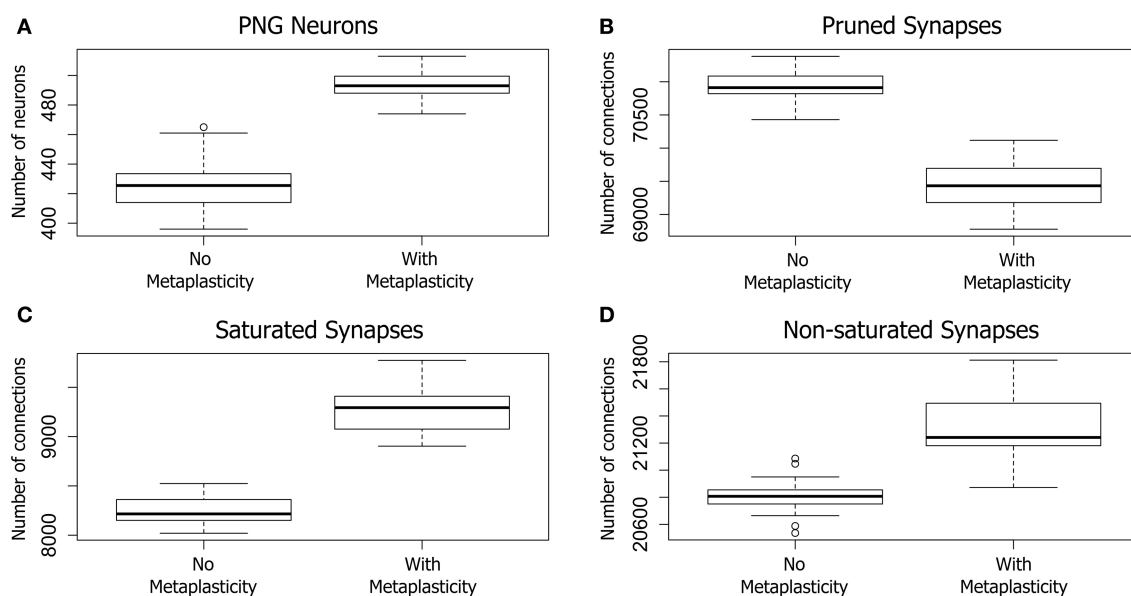


FIGURE 3 | The effect of metaplasticity on twenty large networks showing the change in PNG size (A), or changes in synaptic weight distributions (B–D) with metaplasticity either enabled or disabled.

The boxes in each box-and-whisker plot show the location of the middle 50% of the data, while whiskers show either the maximum (minimum) value or 1.5 times the interquartile range (IQR). Outliers that are outside 1.5 times the IQR are shown as circles (Crawley, 2012). A. Change in the average number of PNG neurons (PNG size). (B–D) Change in average synaptic weight distributions. Each of the 80,000 excitatory connections in each network was assigned to one of the following categories: Pruned

(synaptic weight of zero), Saturated (maximum synaptic weight), Other (non-zero and non-saturated synaptic weight). Data: PNG Size = (means: with metaplasticity = 493, no metaplasticity = 426) (paired t -test: $t = 15.0106$, $p < 0.001$ (2-tailed), d.f. = 19). Pruned = (means: with metaplasticity ≈ 69400 , no metaplasticity ≈ 71000) (paired t -test: $t = 18.0874$, $p < 0.001$ (2-tailed), d.f. = 19). Saturated = (means: with metaplasticity ≈ 9300 , no metaplasticity ≈ 8300) (paired t -test: $t = 20.4666$, $p < 0.001$ (2-tailed), d.f. = 19). Non-saturated = (means: with metaplasticity ≈ 21300 , no metaplasticity ≈ 20800) (paired t -test: $t = 8.2596$, $p < 0.001$ (2-tailed), d.f. = 19).

synaptic weight (*saturated synapses*); and the remaining connections that were neither pruned or saturated (*non-saturated connections*). The overall effect of metaplasticity on these networks was a shift in the weight distribution toward larger weights when metaplasticity is enabled.

Pruned synapses were particularly affected (see **Figure 3B**). The number of pruned synapses dropped significantly when metaplasticity was enabled relative to the number with metaplasticity disabled, producing an increase in the number of *effective connections* (i.e., those with non-zero weight). On average, around 1500 additional connections were added to the network when metaplasticity was enabled and these were distributed between both saturated and non-saturated connections, affecting the counts for these weight groups. The number of saturated connections was therefore significantly increased with approximately 1000 additional connections becoming saturated when metaplasticity was enabled (see **Figure 3C**). There was also a significant increase in the number of non-saturated connections (**Figure 3D**): around 500 additional non-saturated connections were observed with metaplasticity enabled, relative to a network with no metaplasticity as originally predicted from the single synapse model.

3.1.2. PNG Size

Of particular relevance to the focus of this paper, metaplasticity also produced a significant increase in the *average PNG size* across networks (see **Figure 3A**): the number of PNG neurons was significantly higher with metaplasticity enabled than with metaplasticity disabled. There was also a significant increase in the excitatory firing rate measured at the end of the training period

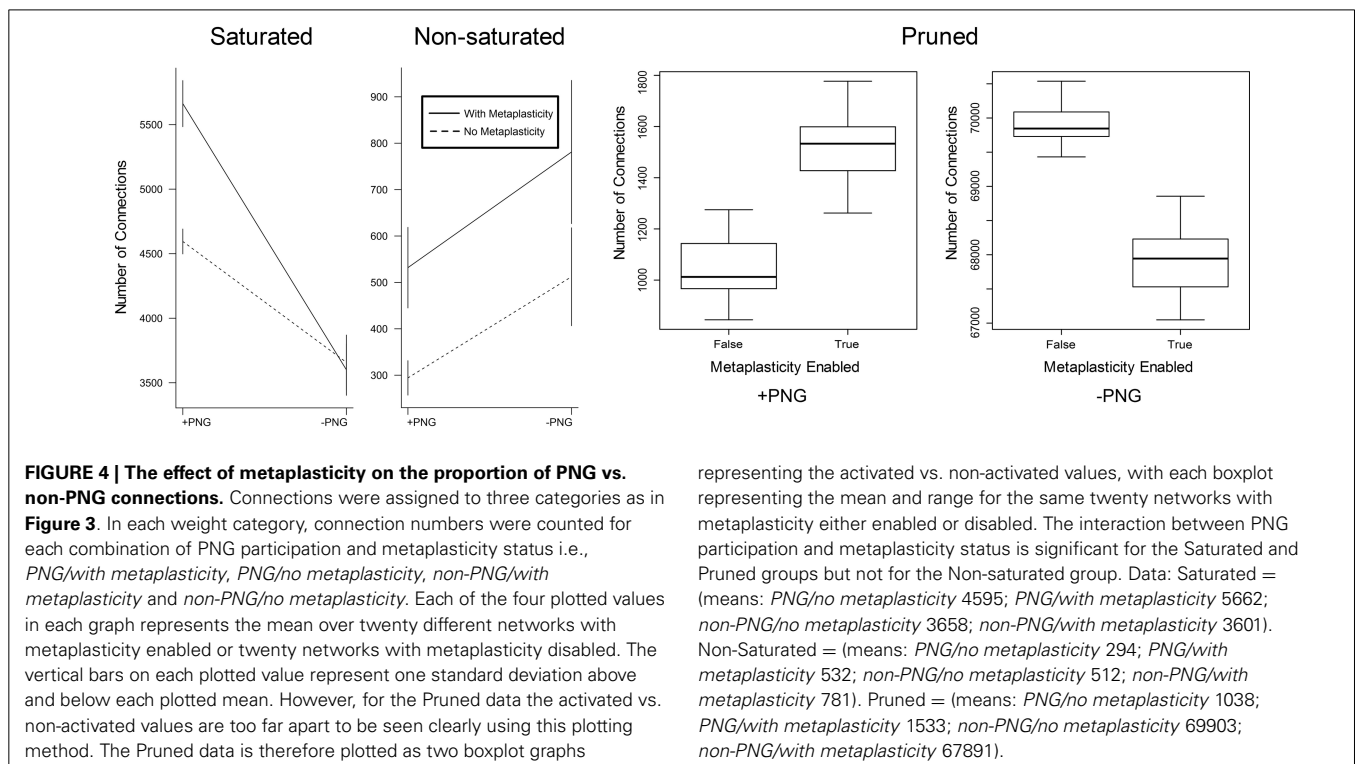
when networks were trained with metaplasticity enabled [$t = 17.7123$, $p < 0.001$ (2-tailed), d.f. = 19; mean (enabled) = 6.0; mean (disabled) = 5.1] (results not shown).

3.2. EFFECTS ON PNG CONNECTIONS

Given the observed increase in PNG size when metaplasticity is enabled, it is worth considering the differential effect of metaplasticity on the weight distributions of connections that do or do not participate in PNG activation. This entails detecting those connections that are regularly activated by the stimulus, allowing the network connections to be partitioned into *PNG connections* (i.e., those that participate in PNG activation) and *non-PNG connections* (i.e., those that do not).

3.2.1. Weight distributions

The effect of metaplasticity on the proportion of PNG vs. non-PNG connections in each of the weight groups of **Figure 3** can be seen in **Figure 4**. There is a significant interaction between the metaplasticity status of the networks and the PNG participation of the connections for some but not all of these weight groups. Saturated weights increase for PNG connections when metaplasticity is enabled, but not for non-PNG connections. For the non-saturated weight group both PNG and non-PNG connections increase in numbers when metaplasticity is enabled, but with no significant interaction for this weight group. Particularly notable is that, despite the overall decrease in pruned weights observable in **Figure 4**, the number of pruned weights in the PNG group actually increases when metaplasticity is enabled. These effects of metaplasticity are small, but given the strongly recurrent structure of these networks, they might still have important consequences on the network dynamics.



3.2.2. PNG Size

The results in Section 3.1.2 show a significant increase in the number of neurons involved in PNG activation when metaplasticity is enabled. A technique for partitioning connections allows an alternative view of PNG activation size in terms of the number of participating connections. **Figure 5** shows the effect of metaplasticity on PNG connection counts for each of the 20 independent networks in **Figure 3**. Unsurprisingly, given the previously observed increase in the number of PNG neurons, enabling metaplasticity produces a significant increase in the total number of PNG connections in each network. Interestingly, most of this increase comes from additional excitatory connections that are recruited into the PNG activation when metaplasticity is enabled.

3.3. EFFECTS OF VARIATION IN THE METAPLASTICITY PARAMETERS

All of the effects reported above used the same values for the resistance (r) and precision (p) metaplasticity parameters ($r = 0.1$ and $p = 0.5$). In this section we briefly discuss some experiments with alternative parameter values. **Figure 6** shows the effect of a random selection of alternative values on the PNG size distributions. The size distribution with metaplasticity disabled is shown on the left for comparison. These results allow a few preliminary observations. Firstly, there is a strong interaction between the two metaplasticity parameters: for instance setting r to very small values has the same effect as disabling metaplasticity, regardless of the value of p . Secondly, metaplasticity certainly has a positive effect on PNG size over a considerable range of values of r and p , so the effects we observed are not due to fortuitous or carefully tweaked settings of these parameter values.

3.4. A ROLE FOR SPIKE LATENCY?


A particularly interesting direction of research has been the interaction of metaplasticity with a rarely studied phenomenon called *spike latency* that is an intrinsic property of the integrator type neurons employed as excitatory cells in this study (Izhikevich, 2007). Spike latency is the delay in spike generation that occurs when a neuron is stimulated at near threshold levels. A simple

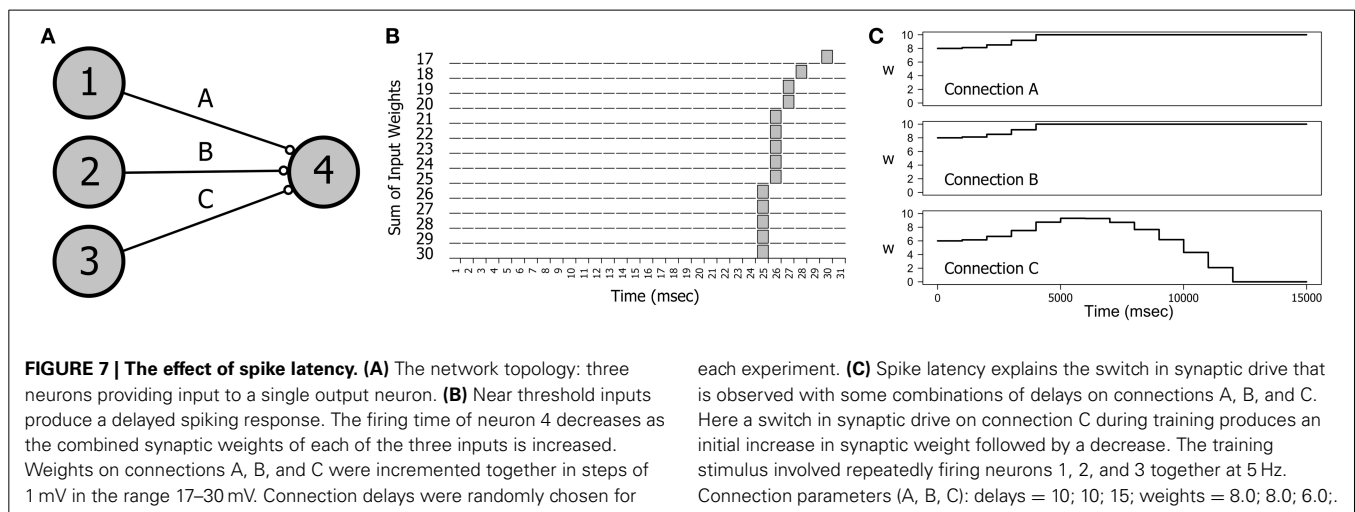
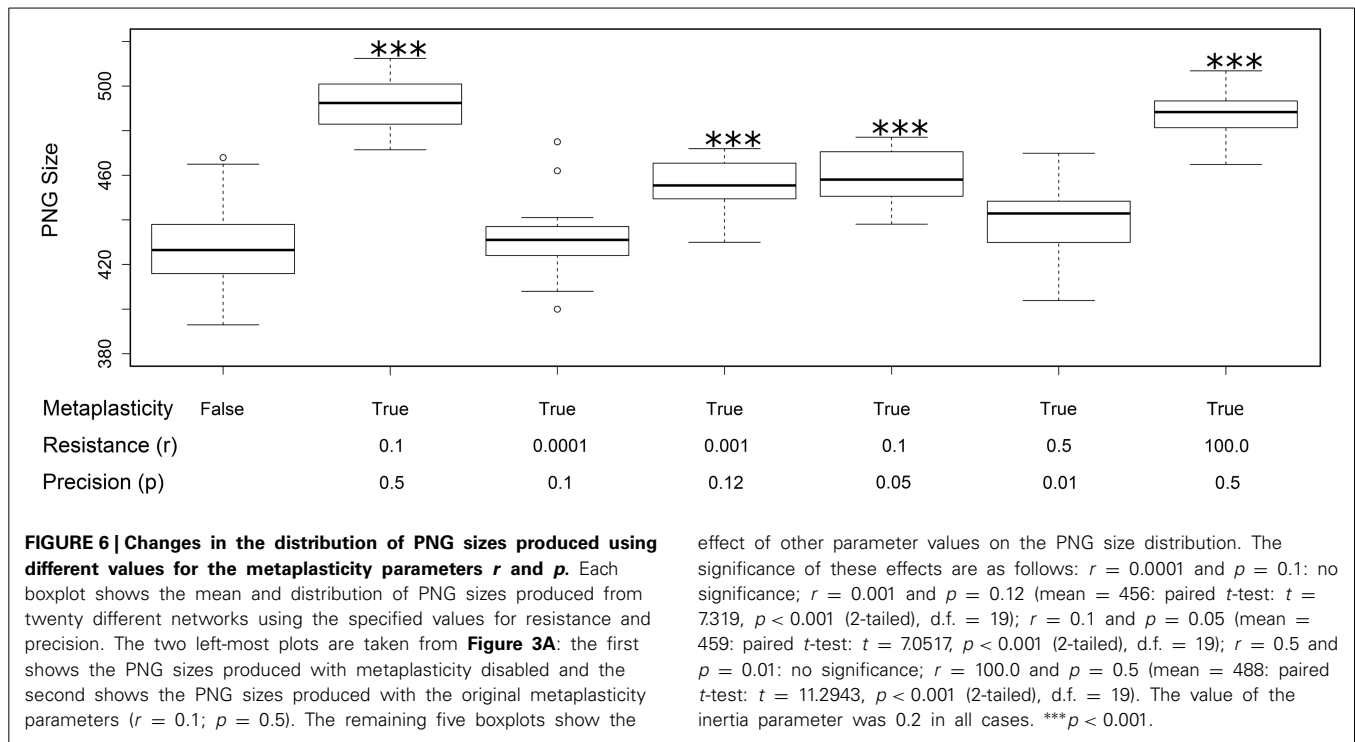
demonstration can be seen in **Figures 7A,B**. **Figure 7A** shows a small network of four neurons in which neurons 1, 2, and 3 provide input to neuron 4. In **Figure 7B** we see the effect of varying levels of stimulation on the firing time of neuron 4 as the connection weights are incremented together in fixed-sized steps. If the input level is barely superthreshold (at 17 mV), neuron 4 spikes at around 30 ms (including connection delays). However, as the input level is increased the firing time of neuron 4 migrates backwards until all three connection weights are saturated.

Spike latency can explain some unusual results in the dynamics of connection weights. If a network such as the one in **Figure 7A** is repeatedly stimulated with a firing pattern that is congruent with the connection delays then the interaction of STDP with the convergent impulses arriving on neuron 4 produces a strong positive synaptic drive that causes the weights on all three connections to increase to saturation and stay there. However, small changes in the network parameters can produce the effect demonstrated in **Figure 7C** in which the weight of Connection C first increases and then decreases. This effect was engineered by decreasing the initial weight on C and making the Connection C delay just a little longer than the delays on A and B. In the first 5 s of training the spike arrival time on C occurs *before* spiking of neuron 4, as is also true of Connections A and B. However, as the combined connection weights increase causing the firing time of neuron 4 to migrate backwards, the spike arrival time on C occurs *after* neuron 4 firing, producing synaptic depression on C.

We hypothesize that spike latency is involved in the underlying mechanism that supports the stability of polychronization, and hence in the ability of PNGs to extend. In large networks with recurrent connections, the effect of recurrent input and other factors such as random firing influence the firing probabilities of PNG neurons in response to subsequent activating stimuli, resulting in complex and unpredictable dynamics. Nevertheless, PNGs are able to exist and even extend, despite this input variability that threatens their stability. The neurons in a polychronous group are exposed to a wide range of spatio-temporal input patterns that we refer to as an *input space*. Individual PNG neurons fire in response to only some of these input patterns, and this subset of the input space we term the *active input space*. Input patterns of particular significance in the active input space are those that result from polychronization in neighboring PNG neurons. However, even these polychronising input patterns can occur with considerable jitter in impulse arrival times due to the complex dynamics of the network. It therefore seems to us that a mechanism that expands the size of the active input space (i.e., the range of patterns that produce neural firing) will increase the firing probability of each PNG neuron in response to the current wave of polychronization. Expansion of the active input space should therefore increase the stability of polychronization, resulting in extended polychronization and an increase in PNG size.

Spike latency allows the precise firing timing to be a function of the level of afferent input, potentially allowing an increase in the range of inputs that produce firing. Our current hypothesis is that spike latency allows increased flexibility in the precise timing of neural firing, producing an expansion of the input space for each PNG neuron. To see how this might work, consider the network shown in **Figure 8A**. This potential polychronous group





is composed of six neurons and is derived from the four neuron network of **Figure 7**. Varying the firing times of the initial three input neurons (1, 2, and 3) produces a wide range of spiking patterns on neuron 4 that together define the input space. With three inputs, this input space is a three dimensional cube that includes the subset of patterns that produce firing of neuron 4.

Figure 8B shows the input space of neuron 4. For convenience, the three-dimensional input space has been flattened to two dimensions by taking the difference between each pair of firing times for the three input neurons. Nevertheless, the figure represents the entire input space i.e., all possible spatio-temporal patterns onto neuron 4 that can be generated if each input neuron is allowed to independently vary its firing time in the range

0–20 ms. Each circle in **Figure 8B** represents a pattern, with filled circles denoting those patterns that produce firing of neuron 4 (the active input space). The larger the proportion of the available input space that is consumed by the active input space, the more flexible the neuron is to jitter in its spatio-temporal inputs.

We can also examine the active input space of neuron 6 relative to these same three input neurons (1, 2, and 3) as shown in **Figure 9**. The active input space for neuron 6 determines the firing probability of neuron 6 under variable conditions, and hence determines the ability of the potential PNG in **Figure 9A** to extend beyond neuron 4. The left column of **Figure 9** (Non-optimized) shows changes in the input space as the phase is shifted through each of four different delays on the 4–6 connection. Importantly,

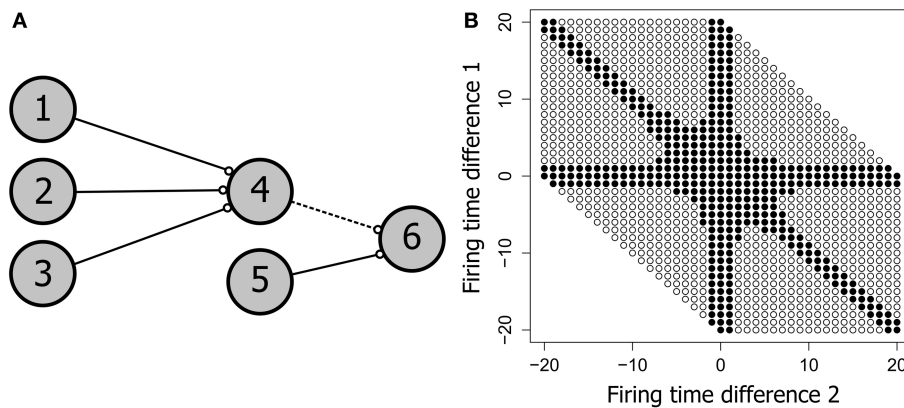


FIGURE 8 | The neural input space. (A) A potential polychronous group consisting of six neurons. **(B)** The input space of neuron 4, showing the patterns that produced firing of neuron 4 as filled circles and patterns that failed to produce firing as unfilled circles. Neurons 1, 2, 3, and 4 have the same topology as in **Figure 7**. The two axes represent the firing time differences ($t_1 - t_2$ and $t_2 - t_3$), where (t_1, t_2, t_3) are the firing times of neurons 1, 2, and 3. All connection weights are initially set to the maximum

synaptic weight of 10 mV. Neuron 4 therefore requires at least two convergent inputs to reach firing threshold. The ability of each firing time combination to produce firing of neuron 4 is tested over 400 trials in fixed-sized 250 ms test frames. Consistent firing requires that neuron 4 fire in at least 100 of the 400 test frames. When weights are saturated, just two of the three inputs are required for firing of neuron 4, producing the *arms* in the input space diagram.

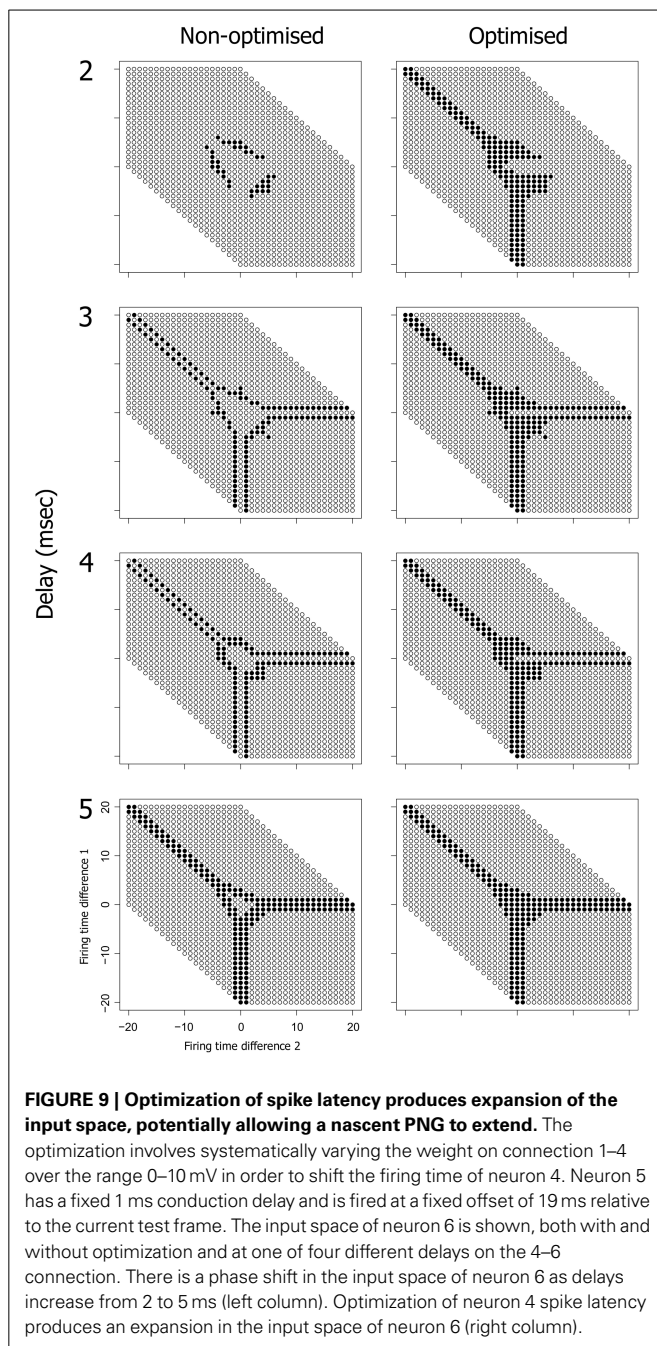
the active input space for many of these phases can be expanded by shifting the firing time of neuron 4 to a time that is congruent with the 4–6 and 5–6 connection delays [right column of **Figure 9** (Optimized)]. These shifts in neuron 4 firing time are produced by changes in the connection weights of the three input neurons (1, 2, and 3) and the effects of spike latency. For now we have performed this optimization for each pattern in the input space by trialing each of ten weight steps on the input neuron connections and selecting weights that produce firing of neuron 6 (if any). An important research question, and one that has yet to be resolved, is whether the interaction of any of the known biologically plausible mechanisms such as metaplasticity and STDP can produce stability enhancement of polychronous groups through a mechanism that optimizes the active input space of each PNG neuron.

4. DISCUSSION

The BCM model famously introduced the idea of a sliding modification threshold in which the tipping point for LTP/LTD induction is determined by the average of recent spiking activity in the post-synaptic cell. Many subsequent models of metaplasticity have followed the BCM model in defining a spike-activity-dependent modification threshold, although these models are typically independent of synaptic size and are not able to prevent synapses from becoming arbitrarily large. In the current study we employ a model of metaplasticity in which the modification threshold ($\theta_M(t)$) is not spike-activity-dependent but is instead set more directly from the current synaptic weight and the level of a spike-timing-dependent variable, the synaptic drive. An implementation of this model was shown to have a significant effect on the size of polychronous groups in large recurrent networks. An understanding of the underlying mechanism for this enhancement is likely to shed light on the principles of PNG formation and perhaps therefore also on the processes of memory formation and storage.

A spike-timing-dependent learning rule appears to be a significant contributor to synaptic plasticity in many parts of the brain. As shown by (Izhikevich and Desai, 2003), BCM-like behavior can be reproduced with an STDP learning rule using uncorrelated or weakly correlated firing of pre- and post-synaptic cells, provided that the spike interaction model conforms to a variant of nearest-neighbor. Biologically realistic spiking patterns are likely to have both weakly correlated and strongly correlated components, with polychronous firing patterns providing an important example of the latter. With strongly correlated spiking patterns the direction and magnitude of synaptic plasticity is no longer determined by the post-synaptic spike rate alone: spike trains with pre- before post-synaptic spike timings produce an upwards or positive drive on synaptic plasticity, whilst post- before pre-synaptic spike trains with identical firing rates produce the opposite effect, a downwards or negative synaptic drive.

As discussed in Section 2.1.2, our metaplastic mechanism was found to maintain the weight of a single synapse within predefined limits without reaching maximum capped values. However, when translated into a large scale network composed of one hundred thousand synapses this moderating influence was considerably diluted and capping at the global weight limits was no longer achieved. Nevertheless, our modeling results show that metaplasticity has a small but significant effect on the distribution of synaptic weights in the network, producing an overall shift toward larger weights. Networks with metaplasticity show a decrease in the number of pruned synapses, and an increase in the number of saturated and non-saturated synapses (**Figure 3**). This trend toward stronger weights is particularly noticeable within the PNG connection group where there is a significant preference for saturated weights relative to the non-PNG connection group. However, there is also a significant increase in pruned synapses within this group, in contrast to the overall trend observed in **Figure 3**.



Other effects of metaplasticity include an increase in the excitatory firing rate, and an increase in the number of PNG connections. Much of the increase in the size of the PNG connection group is due to an increase in participating excitatory connections, although both excitatory and inhibitory connections show an increase in participation with metaplasticity (Figure 5). Perhaps the most interesting finding from the current study was the sensitivity of PNG size to these small metaplasticity-induced changes in network parameters. Small changes in weight distributions produced a 16% increase in PNG size, suggesting that factors that alter the network connectivity have a strong influence

on the stability of neural circuits based on polychronization. A more refined version of the current metaplastic model with carefully tuned parameters might therefore substantially influence the efficiency of polychronization.

Together these results suggest that neurons that participate in polychronization prefer a smaller number of stronger afferent connections relative to non-participating neurons. A high level account of the effects of metaplasticity on PNG size might therefore be constructed by observing the overall match between the effects of metaplasticity on the synaptic weight distribution (i.e., more saturated and non-saturated weights), and the preference of PNG connections for saturated weights. However, a deeper explanation is required that describes the underlying mechanism whilst accounting for the pruning of PNG connection weights. To this end we have explored a number of avenues such as the effect of metaplasticity on the temporal firing precision, or on the evolution of synaptic weights over time. The effect of spike latency on the active input space of PNG neurons has been a particularly interesting research direction. Initial results show that spike latency allows increased flexibility in the precise timing of neural firing, and that expansion of the active input space for each neuron can be achieved by optimization of spike latency. We speculate that a mechanism that optimizes the active input space of each PNG neuron might produce stability enhancement of polychronous groups through the interaction of metaplasticity with a biologically plausible learning rule such as STDP.

REFERENCES

- Abraham, W. C. (2008). Metaplasticity: tuning synapses and networks for plasticity. *Nat. Rev. Neurosci.* 9, 387–387. doi: 10.1038/nrn2356
- Beňušková, L., Rema, V., Armstrong-James, M., and Ebner, F. F. (2001). Theory for normal and impaired experience-dependent plasticity in neocortex of adult rats. *Proc. Natl. Acad. Sci. U.S.A.* 98, 2797–2802. doi: 10.1073/pnas.051346398
- Benuskova, L., and Abraham, W. C. (2007). STDP rule endowed with the BCM sliding threshold accounts for hippocampal heterosynaptic plasticity. *J. Comput. Neurosci.* 22, 129–133. doi: 10.1007/s10827-006-0002-x
- Bienenstock, E. L., Cooper, L. N., and Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J. Neurosci.* 2, 32–48.
- Clopath, C., Büsing, L., Vasilaki, E., and Gerstner, W. (2010). Connectivity reflects coding: a model of voltage-based stdp with homeostasis. *Nat. Neurosci.* 13, 344–352.
- Crawley, M. J. (2012). *The R book*. New York, NY: John Wiley & Sons.
- Delorme, A., Perrinet, L., and Thorpe, S. J. (2001). Networks of integrate-and-fire neurons using rank order coding b: spike timing dependent plasticity and emergence of orientation selectivity. *Neurocomputing* 38, 539–545. doi: 10.1016/S0925-2312(01)00403-9
- Guetig, R., Aharonov, R., Rotter, S., and Sompolsinsky, H. (2003). Learning input correlations through nonlinear temporally asymmetric hebbian plasticity. *J. Neurosci.* 23, 3697–3714. Available online at: <http://www.jneurosci.org/content/23/9/3697.full>
- Guise, M., Knott, A., and Benuskova, L. (2013). *Spinula: Software for Simulation and Analysis of Spiking Network Models (OUCS-2013-03)*. Technical report, Department of Computer Science, University of Otago, Dunedin.
- Guise, M., Knott, A., and Benuskova, L. (2013a). *Response Fingerprinting: A Probabilistic Method for Evaluating the Network Response to Stimuli (OUCS-2013-10)*. Technical report, Department of Computer Science, University of Otago, Dunedin.
- Guise, M., Knott, A., and Benuskova, L. (2014). A Bayesian model of polychronicity. *Neural Comput.* 26, 2052–2073. doi: 10.1162/NECO.a.00620
- Hulme, S. R., Jones, O. D., Ireland, D. R., and Abraham, W. C. (2012). Calcium-dependent but action potential-independent bcm-like metaplasticity in the hippocampus. *J. Neurosci.* 32, 6785–6794. doi: 10.1523/JNEUROSCI.0634-12.2012

- Izhikevich, E. M., and Desai, N. S. (2003). Relating stdp to bcm. *Neural Comput.* 15, 1511–1523. doi: 10.1162/089976603321891783
- Izhikevich, E. M., Gally, J. A., and Edelman, G. M. (2004). Spike-timing dynamics of neuronal groups. *Cereb. Cortex* 14, 933–944. doi: 10.1093/cercor/bhh053
- Izhikevich, E. M. (2006a). Polychronization: computation with spikes. *Neural Comput.* 18, 245–282. doi: 10.1162/089976606775093882
- Izhikevich, E. M. (2006b). *Reference Software Implementation for the Izhikevich Model: Minimal Spiking Network that can Polychronize*.
- Izhikevich, E. M. (2007). *Dynamical Systems in Neuroscience*. Cambridge, MA: MIT press.
- Lazar, A., Pipa, G., and Triesch, J. (2007). Fading memory and time series prediction in recurrent networks with different forms of plasticity. *Neural Netw.* 20, 312–322. doi: 10.1016/j.neunet.2007.04.020
- Lubenov, E. V., and Siapas, A. G. (2008). Decoupling through synchrony in neuronal circuits with propagation delays. *Neuron* 58, 118–131. doi: 10.1016/j.neuron.2008.01.036
- MacKay, D. (1963). Psychophysics of perceived intensity: a theoretical basis for fechner's and stevens' laws. *Science* 139, 1213–1216. doi: 10.1126/science.139.3560.1213-a
- Martinez, R., and Paugam-Moisy, H. (2009). "Algorithms for structural and dynamical polychronous groups detection," in *Artificial Neural Networks – ICANN 2009, Vol. 5769 of Lecture Notes in Computer Science*, eds C. Alippi, M. Polycarpou, C. Panayiotou, and G. Ellinas (Berlin; Heidelberg: Springer Berlin Heidelberg), 75–84.
- Ngezahayo, A., Schachner, M., and Artola, A. (2000). Synaptic activity modulates the induction of bidirectional synaptic changes in adult mouse hippocampus. *J. Neurosci.* 20, 2451–2458. Available online at: <http://www.jneurosci.org/content/20/7/2451.full>
- Shmiel, T., Drori, R., Shmiel, O., Ben-Shaul, Y., Nadasdy, Z., Shemesh, M., et al. (2006). Temporally precise cortical firing patterns are associated with distinct action segments. *J. Neurophysiol.* 96, 2645–2652. doi: 10.1152/jn.00798.2005
- Shouval, H. Z., Castellani, G. C., Blais, B. S., Yeung, L. C., and Cooper, L. N. (2002). Converging evidence for a simplified biophysical model of synaptic plasticity. *Biol. Cybern.* 87, 383–391. doi: 10.1007/s00422-002-0362-x
- van Vreeswijk, C. and Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274, 1724–1726. doi: 10.1126/science.274.5293.1724
- Villa, A. E., Tetko, I. V., Hyland, B., and Najem, A. (1999). Spatiotemporal activity patterns of rat cortical neurons predict responses in a conditioned task. *Proc. Natl. Acad. Sci. U.S.A.* 96, 1106–1111. doi: 10.1073/pnas.96.3.1106
- Vogels, T. P., Rajan, K., and Abbott, L. (2005). Neural network dynamics. *Annu. Rev. Neurosci.* 28, 357–376. doi: 10.1146/annurev.neuro.28.061604.135637
- Zenke, F., Hennequin, G., and Gerstner, W. (2013). Synaptic plasticity in neural networks needs homeostasis with a fast rate detector. *PLoS Comput. Biol.* 9:e1003330. doi: 10.1371/journal.pcbi.1003330

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 31 March 2014; accepted: 15 January 2015; published online: 05 February 2015.

Citation: Guise M, Knott A and Benuskova L (2015) Enhanced polychronization in a spiking network with metaplasticity. *Front. Comput. Neurosci.* 9:9. doi: 10.3389/fncom.2015.00009

This article was submitted to the journal *Frontiers in Computational Neuroscience*. Copyright © 2015 Guise, Knott and Benuskova. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Learning structure of sensory inputs with synaptic plasticity leads to interference

Joseph Chrol-Cannon and Yaochu Jin *

Department of Computer Science, Faculty of Engineering and Physical Sciences, University of Surrey, Guildford, UK

OPEN ACCESS

Edited by:

Matthieu Gilson,
Universitat Pompeu Fabra, Spain

Reviewed by:

Pierre Yger,
Institut de la Vision, France
Naoki Hiratani,
The University of Tokyo, Japan

*Correspondence:

Yaochu Jin,
Department of Computer Science,
Faculty of Engineering and Physical
Sciences, University of Surrey,
Guildford, Surrey GU2 7XH, UK
yaochu.jin@surrey.ac.uk

Received: 29 January 2015

Accepted: 20 July 2015

Published: 05 August 2015

Citation:

Chrol-Cannon J and Jin Y (2015)
Learning structure of sensory inputs
with synaptic plasticity leads to
interference.
Front. Comput. Neurosci. 9:103.
doi: 10.3389/fncom.2015.00103

Synaptic plasticity is often explored as a form of unsupervised adaptation in cortical microcircuits to learn the structure of complex sensory inputs and thereby improve performance of classification and prediction. The question of whether the specific structure of the input patterns is encoded in the structure of neural networks has been largely neglected. Existing studies that have analyzed input-specific structural adaptation have used simplified, synthetic inputs in contrast to complex and noisy patterns found in real-world sensory data. In this work, input-specific structural changes are analyzed for three empirically derived models of plasticity applied to three temporal sensory classification tasks that include complex, real-world visual and auditory data. Two forms of spike-timing dependent plasticity (STDP) and the Bienenstock-Cooper-Munro (BCM) plasticity rule are used to adapt the recurrent network structure during the training process before performance is tested on the pattern recognition tasks. It is shown that synaptic adaptation is highly sensitive to specific classes of input pattern. However, plasticity does not improve the performance on sensory pattern recognition tasks, partly due to synaptic interference between consecutively presented input samples. The changes in synaptic strength produced by one stimulus are reversed by the presentation of another, thus largely preventing input-specific synaptic changes from being retained in the structure of the network. To solve the problem of interference, we suggest that models of plasticity be extended to restrict neural activity and synaptic modification to a subset of the neural circuit, which is increasingly found to be the case in experimental neuroscience.

Keywords: synaptic plasticity, spiking neural networks, recurrent neural networks, inference, pattern recognition

1. Introduction

Recurrent neural networks consisting of biologically based spiking neuron models have only recently been applied to real-world learning tasks under a framework called reservoir computing (Maass et al., 2002; Buonomano and Maass, 2009). The models of this framework use a recurrently connected set of neurons driven by an input signal to create a non-linear, high-dimensional temporal transformation of the input that is used by single layer perceptrons (Rosenblatt, 1958) to produce desired outputs. This restricts the training algorithms to a linear regression task, while still allowing the potential to work on temporal data in a non-linear fashion.

Given an initially generated static connectivity, reservoir computing is based on the principle of random projections of the input signal in which the network structure is completely independent

of the input patterns. In these models, the only features learned by the trainable parameters of the perceptron readout are the correlations between the randomly projected features and the desired output signal.

We believe that learning in neural networks should go further than supervised training based on error from the output. All synapses should adapt to be able to encode the structure of the input signal and ideally, should not rely on the presence of a desired output signal from which to calculate an error with the actual output. The neural activity generated by the input signal should provide enough information for synapses to adapt and encode properties of the signal in the network structure. By applying unsupervised adaptation to the synapses in the form of biologically derived plasticity rules (Bienenstock et al., 1982; Bi and Poo, 1998; Wittenberg and Wang, 2006) it is hoped to provide the means for the recurrently connected neurons of the network to learn a structure that generates more effective features than a completely random projection that is not specific to the input data.

On a conceptual level, unsupervised learning is important in the understanding of how synaptic adaptation occurs because it is still unknown what the sources of supervised signals are in the brain, if any exist. From early work on synaptic self-organization (Hebb, 1949), the principle of learning has rested on correlations in neural activity becoming associated together and forming assemblies that activate simultaneously. These structures are thought to encode invariances in the sensory input that are key in developing the ability to recognize previously encountered patterns.

In this work we will explore the impact of applying several biologically derived plasticity mechanisms on three temporal sensory discrimination tasks. Two forms of spike-timing dependent plasticity (STDP) (Bi and Poo, 1998; Wittenberg and Wang, 2006) will be tested, along with the Bienenstock-Cooper-Munro (BCM) rule (Bienenstock et al., 1982). The sensory tasks will include real-world speech and video data of human motion. Synaptic plasticity will be applied in an unsupervised pre-training phase, before the supervised regression of the perceptron readout occurs. We will compare the impact that plasticity has on the performance in these tasks and also analyze the specific structural adaptation of the weight matrices between each of the classes of input sample in each task. A method will be introduced to evaluate the extent to which the synaptic changes encode class-specific features in the network structure.

Interference between different samples is a well-established phenomenon in sequentially trained learning models (McCloskey and Cohen, 1989; Ratcliff, 1990; French, 1999). When presented to a learning model, an input pattern will cause specific changes to be made in the models parameters—in the case of neural networks, the synapses. However, during this encoding process, existing structure in the synaptic values is interfered with. In this way, consecutive input patterns disrupt previously learned features, sometimes completely. This effect is known as forgetting. It is of direct concern to neural networks trained on sensory recognition tasks that consist of spatio-temporal patterns projected through a common neural processing pathway. We will quantify the level of interference

between the synaptic parameters for each tested plasticity model being applied to each type of sensory data.

Existing studies report that adapting neural circuits with plasticity improves their performance on pattern recognition tasks (Yin et al., 2012; Xue et al., 2013) but there is no analysis of how the adaptation of synaptic parameters leads to this result. On the other hand, work that does detailed analysis on the structural adaptation of the network does so using synthetic input patterns that are already linearly separable (Toutounji and Pipa, 2014) or Poisson inputs projecting to single and recurrently connected neurons (Gilson et al., 2010). For a review of work applying plasticity models to improve the general properties of neural networks, the reader is referred to Chrol-Cannon and Jin (2014a).

The experiments undertaken in this work will be performed on a typical reservoir computing model with its recurrent connections adapted with plasticity. Two main angles of analysis are made; we determine the strength of input specific synaptic adaptation and the extent to which consecutive inputs interfere within the synapses. Both of these are achieved through analysis of the change in weight matrix in response to each pattern.

2. Results

2.1. Training Recurrent Networks with Plasticity

Our training and analysis is performed on a typical liquid state machine (LSM) model (Maass et al., 2002) that is trained to correctly classify temporal input patterns of sensory signals. Details of the models and simulations can be found in the Section 4. Here we present an overview of the experimental procedure.

An LSM consists of recurrently connected spiking neurons in which transient activity of the neurons is driven by time-series input sequentially exciting their membrane potential. In order for an output to be produced from the network and used to train a supervised readout, a snapshot must be taken of the transient activity which we call the state vector. This vector is weighted and summed to produce an output, the weights of which are trained with linear regression.

In our experiments we adapt the recurrent connections with synaptic plasticity before taking the state vectors used for pattern recognition. We intend to change the synaptic weights from their initial random structure, to values that are adapted to the general statistics of the input signals. After this pre-training process, we take the state vectors for each sample in the data set and use it to train a set of readouts to recognize labeled patterns in the data. Performance of pattern recognition is only a small aspect of our analysis of synaptic adaptation through plasticity. The analysis methodology described in the next subsection requires the information of how each sample of input causes unique adaptation of the synapses. Therefore, for convenience, when collecting the liquid state vectors of a given sample from the neural activity, we also compute the synaptic change during the presentation of that sample and store the weight matrix adaptation.

Figure 1 illustrates the three step process just described, delineated into; a pre-training phase of synaptic plasticity, a collection of the liquid state vectors and weight adaptation

matrices, and a supervised training phase of linear readouts for pattern recognition.

2.2. Description of Sensory Inputs

Complex sensory signals are projected through a common set of nerve fibers to cortical regions that must learn to distinguish between them based on differences in their spatial-temporal features.

Three sensory recognition tasks are selected, among which two of them consist of real audio and video signals of human speech and motion. For all tasks, the neural network output is trained to respond uniquely to each of the different types of input sample and therefore be able to perform effective recognition between them. Also, sample specific synaptic adaptations are analyzed to determine if unique structure is learned within the network due to synaptic plasticity.

The auditory task is to distinguish between nine different speakers based on short utterances of the vowel /ae/. Each of the 640 samples consists of a frequency “spectrogram” that plots frequency intensity over a sequence of audio time frames. **Figure 2** plots an example sample from each of the nine speakers.

The visual task is to distinguish between six types of human behavior; boxing, clapping, waving, walking, running and jogging. The 2391 samples are video sequences of many different subjects performing those six motions. There is a


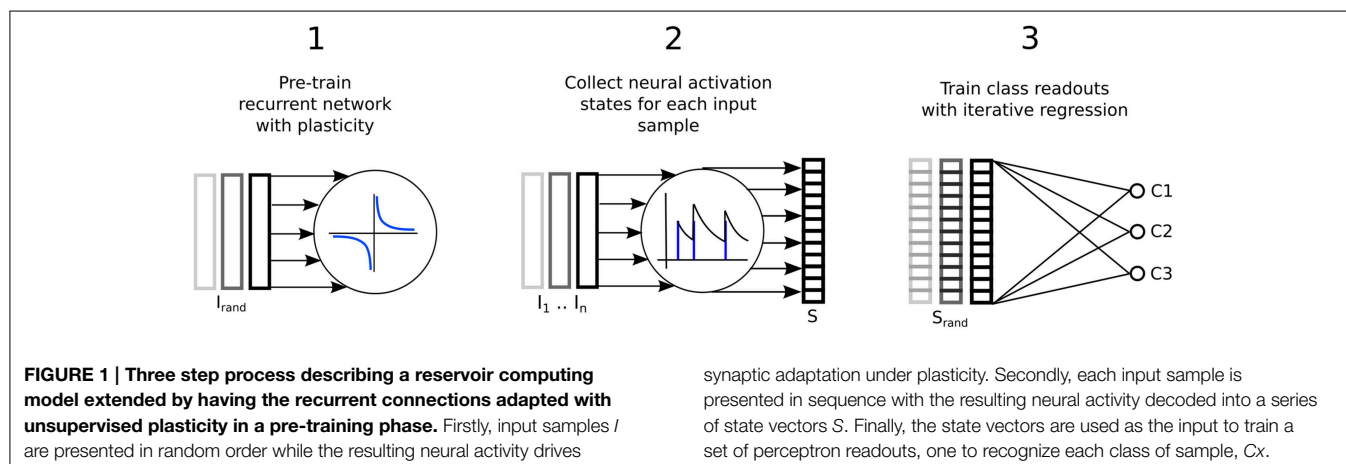
simple pre-processing stage that converts the video data into a sparse representation before being used as input to the neural network. Extracted still frames and processed features are plotted in **Figure 3** for one subject performing each of the six behaviors.

A synthetic data set is generated to model a low spatial dimension but very high frequency temporal structure, in contrast to the previous two sensory tasks. Three functions generate time-varying single dimensional signals that the network learns to distinguish between. A complete description and method for generating the data is described in Jaeger (2007) (**Figure 4**) illustrates part of this signal.

The auditory and visual tasks are described in Kudo et al. (1999) and Schuldt et al. (2004), respectively, with data availability also provided.

2.3. Analysis of Synaptic Adaptation

Synaptic weight adaptation matrices form the basis of the analysis in this work. **Figure 5** depicts the process of these matrices being collected and used for analysis of class-specific synaptic plasticity. Firstly, synaptic plasticity is applied to the network to adapt a baseline weight matrix that reflects the general statistics of the input patterns in the data set. Secondly, each the weight adaptation matrix is collected for each sample and these are grouped by class and also into two sets based on the training and






FIGURE 3 | Human motion samples for the six types of behavior in the KTH visual discrimination task. This illustration consists of different behaviors from a single person, while the whole data set contains 25 persons. **Top row:** Still frames from example video samples; boxing,

clapping, waving, walking, running and jogging. **Bottom row:** Features extracted corresponding to the samples above, according to Equations (13) and (14). Features are the raw time-series activity used as input to the neural network.




FIGURE 4 | Plot of 500 of the 50,000 data samples generated according to Jaeger's tri-function system recognition time-series task (Jaeger, 2007).

testing data division. Finally, the Euclidean distance is calculated between each weight matrix, with the average distance between each set plotted in a type of “confusion matrix” in which a low distance indicates high similarity between the adaptation of synaptic parameters.

In the confusion matrix just described, if the diagonal values are lower than the others it means that synaptic plasticity is sensitive to the structural differences in input samples that are labeled as different classes. The stronger the diagonal trend, the more sensitive plasticity is to features of the input. It means that plasticity learns to distinguish class labels, such as different speakers or human actions, without ever being exposed to the labels themselves *a priori*.

The weight adaptation matrices are also used to estimate the amount of interference between different input samples within the synaptic parameters. This is described further later in the Results Section.

2.4. Learning Input-Specific Adaptations using Plasticity

We wish to test the hypothesis that synaptic plasticity is encoding a distinct structure for input samples of different labels. For the speech task, these labels consist of different speakers and for the video recognition task the labels consist of different human behaviors.

The data sets are divided evenly into two. Each subset is used to train a recurrently connected network for 10,000 iterations, selecting a sample at random on each iteration. The changes to the weight matrix due to plasticity are recorded for each sample presentation. This is then used to create a class-specific average weight change for each of the class labels in both of the sample subsets. Finally, we calculate the Euclidean distance between each class in one set and each class in the other according to the following formula:

$$\text{Dist}(C_{lab}^X, C_{lab}^Y) = \sum_{i=1}^N |\Delta W_i(C_{lab}^X) - \Delta W_i(C_{lab}^Y)| \quad (1)$$

Where C_{lab} denote class labels, X and Y distinguish the separated sets of samples, ΔW is the change in weight matrix for a presented sample, N is the number of synapses, and i the synapse index.

This effectively produces a confusion matrix of similarity in the synaptic weight change for different classes of input. Having lower values on the descending diagonal means that there is structural adaptation that is specific to the class of that column compared with the similarity between structural adaptations of two different classes.

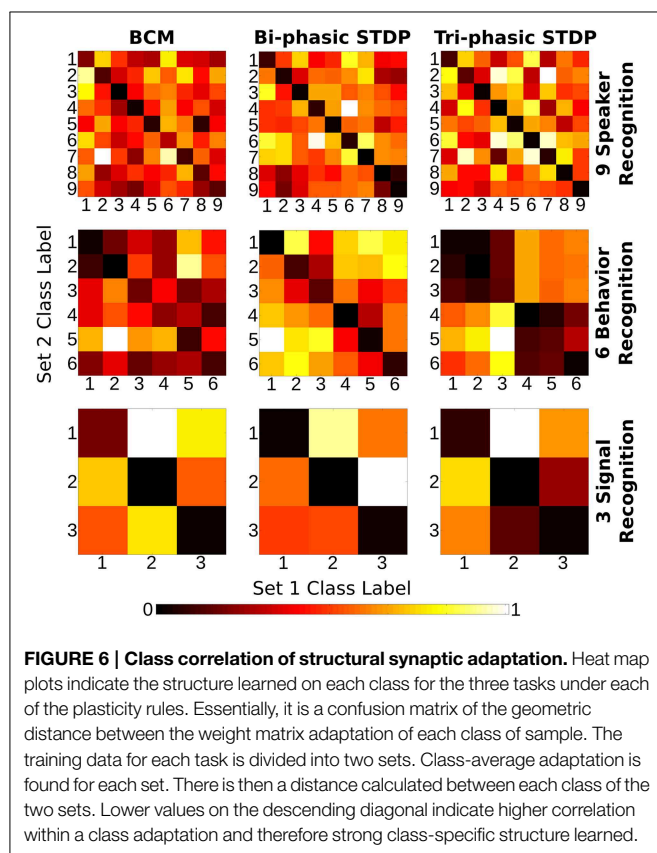



Figure 6 shows the “weight change confusion matrices” described above, for each plasticity model applied to all sensory tasks (nine experiments in total). All of the experiments show at least some stronger similarity in the descending diagonals and most are stark in this manner. It is certainly a strong enough pattern to show that through the many iterations of training, each of the plasticity models have become sensitive to the particular structure of the sensory input signals so that each different class of sample will give rise to changes in synaptic strength that

TABLE 1 | Classification error rates.

	Static	BCM	STDP	TP-STDP
Tri-func	0.153	0.157	0.204	0.138
KTH	0.283	0.3	0.333	0.383
Vowels	0.089	0.086	0.092	0.086

Values averaged over 10 trials with random seed based on system clock. SD did not exceed 0.03 for all values. Bold values indicate lowest error rate.

are distinct from other classes compared with the similarity to themselves. We re-iterate that the class labels were not used in any way in the plasticity models themselves and so the differences in the weight change arise from the input signals alone.

There are a few exceptions to the strong diagonal patterns in **Figure 6**. This means that some classes are not effectively distinguished from each other; speakers 8/9 with bi-phasic STDP, behaviors 1/2 with BCM, behaviors 1/2/3, and 4/5/6 with tri-phasic STDP. The latter confusion corresponds to the behaviors of boxing/clapping/waving and walking/running/jogging. From the similarity of those input features shown in the lower panes of **Figure 3**, it is evident why this confusion might occur.

2.5. Classification Performance with Plasticity

Perhaps the ultimate goal of neural network methods when applied to sensory tasks is the ability to accurately distinguish different types of input sample by their patterns. We compare the error rates achieved by our neural network on the three sensory tasks, with and without the different forms of plasticity used in this work. **Table 1** lists the error rates achieved for each of the learning tasks with the different plasticity rules active in a pre-training phase in addition to a static network with fixed internal synapses.

From the error rates in **Table 1** it is evident that pre-training the network with synaptic plasticity can make insignificant improvements in lowering the error rate. However, the results here indicate that it can have a greater negative impact than a positive one. In the KTH human behavior data set, all three plasticity models increase the error rate by between 1.7 and 10%.

Conversely, the best improvement was found on the tri-function signal recognition task with tri-phasic STDP at only 1.5%.

It is clear from the network output that pre-training with synaptic plasticity is not a suitable method for this class of model. This does not contradict the result that plastic synapses are learning useful, input-specific structure. However, it does suggest that the structure being learned is not effectively utilized in the generation of a network output. We next investigate interference between synaptic changes to determine if the structural learning is retained in the network or if interference is a barrier for effective application of synaptic plasticity.

2.6. Synaptic Interference

When a model adapts incrementally to sequentially presented input, existing patterns that have been learned by the model parameters are prone to be overwritten by learning new patterns. This is known as interference. The work that has studied this effect (McCloskey and Cohen, 1989; Ratcliff, 1990; French, 1999), test the ability to recognize previously presented input after the model has been trained on new ones in order to estimate how much learning has been undone. When new training leaves the model unable to recognize old patterns, it is said there has been catastrophic interference and forgetting.

We introduce a method of measuring interference directly in synaptic parameters instead of the model output. Our measure is described in detail in the Section 4. I^{total} directly quantifies all synaptic changes that are overwritten.

The interference for each of our experiments is listed in **Table 2**. In all but one of the experiments the interference level is between 82 and 96%. Most of the learned structure for each class of input is forgotten as consecutive samples overwrite each other's previous changes. Bi-phasic STDP applied to speaker recognition has the lowest level of interference at 58%.

To further explore interference and visualize the impact of plasticity, synaptic changes will be analyzed directly. **Figure 7** is an illustrative example in which a reduced network size of 35 neurons is used to improve visual clarity of the plotted patterns. It is an example for the speaker recognition task with BCM plasticity with similar figures for the other experiments given in Supplementary Figures 1–8. It shows the adaptation of the synaptic weight matrix produced by each speaker in the voice recognition task. This is plotted against the activity level for each neuron, **S**, and the readout weights, **R**, that are trained to generate an output that is sensitive to that given speaker. Each of these sub plots is the average response taken over all sample presentations from that speaker. This makes a whole chain of effect visible: from

the synaptic change of an internal network connection, to the average neuron state for a given speaker, to the selective weights of the readout for that speaker. For all to be working well in a cohesive system, we expect that a positive weight change should correspond with a neuron activation unique to the class which would in turn improve the recognition ability of the readout to identify that class.

The sections of the class weight matrix highlighted in green in **Figure 7**, highlight an example where synaptic interference is occurring between different types of pattern. Directly opposing features in the weight matrix adaptations show the samples negating each other's changes. However, the same features are also most distinctively class specific.

Any synapse can only change in two directions: positively or negatively, which means that a single synapse can only adapt to distinguish between two mutually exclusive kinds of input pattern. If n synapses are considered in combination, then the number of input patterns that can be discriminated becomes 2^n in ideal theoretical conditions. **Figure 7** illustrates this principle in practice with regards to the nine speaker recognition tasks. The adapted synapses labeled (a) can clearly distinguish speaker {#1} from speakers {#2, #3} but cannot distinguish {#2} from {#3}. Similarly, the adapted synapses labeled (b) can distinguish speakers {#1, #6, #8} from speakers {#3, #4, #9} but cannot distinguish speakers within either of those sets. However, if the synapses (a–d) are considered in combination, then all speakers can be distinguished by synaptic plasticity changes alone.

Figure 7 also shows the weight changes are not correlated with the neural activity or readout weights. For plasticity to improve the accuracy of sensory discrimination, it would be expected that synapses would strengthen for class specific neural activity and weaken for common neural activity. This is not the case in our results.

3. Discussion

3.1. Evolution of Synaptic Weights

Our main conclusions are drawn from the observation that the synaptic plasticity models tested become sensitive to specific class labels during a competitive process of synaptic interference between input patterns. For our conclusions to be generally applicable to recurrent neural circuits and liquid state machines in particular, we must demonstrate that synaptic weights reach some stability during pre-training and that the neural activity dynamics are working in a balanced regime.

Figure 8 shows a series of plots taken at 1, 100, and 1000 input iterations that show the evolving distributions of synaptic weights and inter-spike intervals (ISI) for each of the experiments performed in this work.

In general, the plots show that between the first and 100th pattern, the synaptic weights are adapted significantly by plasticity, with a corresponding—but more subtle—change in the distribution of ISIs. While there is also some level of change in weights between the 100th and 1000th iteration, the level is far smaller, which indicates that the synapses are converging on a common structure. However, it is important to note that for simulations even up to 10,000 iterations there is always

TABLE 2 | Synaptic interference.

	BCM	STDP	TP-STDP
Tri-func	0.82	0.8	0.88
KTH	0.92	0.93	0.96
Vowels	0.96	0.58	0.9

Values averaged over 10 trials with random seed based on system clock. SD did not exceed 0.07 for all values. Bold values indicate lowest interference.

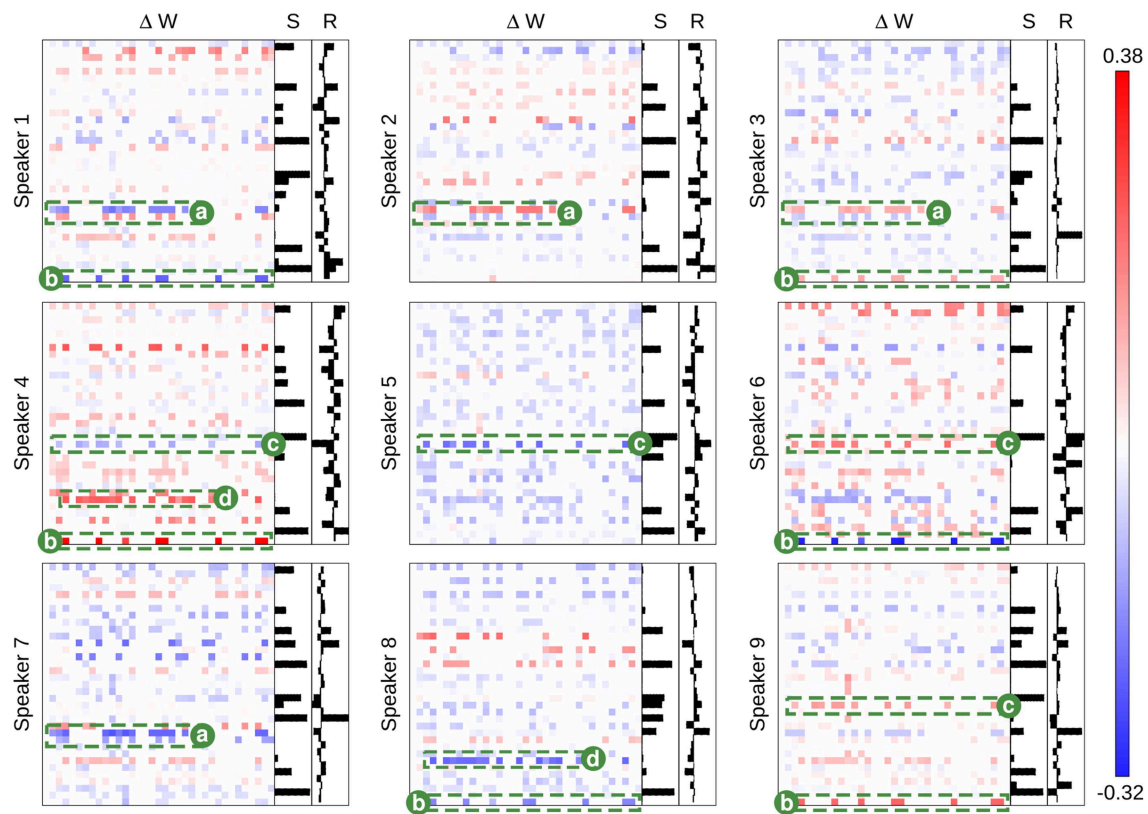


FIGURE 7 | The class-specific synaptic adaptation for the 9 class speaker recognition task under BCM plasticity. The main heat maps in each subplot show the adaptation of the weight matrix (synapses) after the presentation of voice input data from each speaker. Blue values show a reduction in synaptic strength and red values show an increase. Each $N \times N$ weight matrix has *pre-neurons* on the x-axis and *post-neurons* on the y-axis.

The bar-chart, **S**, shows the average neuron activation for each class. The bar-chart, **R**, shows the learned readout weights. Labeled synapses **a–d** indicate key structural changes that are selective between different speakers. Each label alone can distinguish between two sets of speaker. Taken all together, the labeled synapses adapt specifically to each speaker in a unique pattern, learning a distinct network structure for each one.

some low level of synaptic change. The plasticity models tested never stabilize to a point in which there is no further synaptic adaptation, even when we repeatedly present a single input sample.

Each of the plasticity models drives the synaptic weights to a different kind of distribution. STDP creates a bi-modal distribution that drives most weights to the extremes: 0 and 10, with a few that are in a state of change leading up to each boundary. It leads to a structure with more full strength synapses than zeroed. TP-STDP and BCM plasticity leads to sparser connectivity that drives most weights to zero. In particular, TP-STDP only maintains a small number of weak connections due to the narrow window of potentiation being surrounded by depressive regions that suppress most connections. BCM includes an implicit target level of post-synaptic activity that encourages some synapses to take larger values but doesn't drive them to their maximum.

The distribution of ISIs give an indication of the dynamics of the neural activity. The plots in **Figure 8** show that a balance between completely sparse and saturated activity is maintained during the simulation. The shape of the ISI

distributions tend to stabilize between 100 and 1000 sample presentations.

The above observations provide some evidence that the results presented in this article are not simply an artifact of a particular choice of model parameters but are observed for a normally functioning liquid state machine.

3.2. Unsupervised Plasticity Learns Label Specific Structure

Both STDP and BCM models adapt the synapses of a network in distinctive patterns according to which type of sample is being presented to the network. We can conclude that presenting a training signal with the sample label is not required for plasticity to learn specific information for complex sensory inputs from different sources. This result holds for the speech, visual and benchmark pattern recognition tasks. To achieve this feat, we hypothesize that plasticity drives the synaptic parameters to a structure that represents an average between all input samples. Once converged, any further input stimulus will drive the synaptic parameters in a unique direction away from this average structure. On balance, scrambled

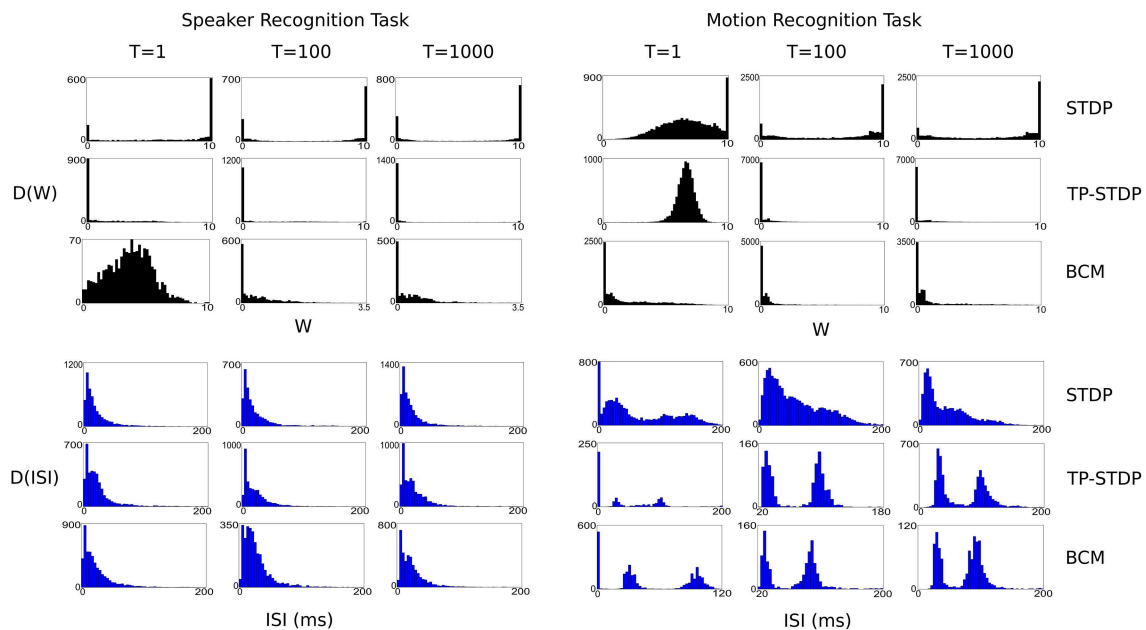


FIGURE 8 | Plots of evolving synaptic weight (black, top) and ISI (blue, bottom) distributions given for each recognition task and each plasticity model. The plots are snapshots of the parameter distributions after 1, 100, and 1000 input samples have been presented during pre-training.

presentation of random inputs keeps the network in this sensitive state.

3.3. Uniformly Applied Plasticity Leads to Synaptic Interference

We show synaptic plasticity spends most of its action counteracting previous changes and overwriting learned patterns. The same patterns of synaptic adaptation that distinctly characterizes each class of input are the same ones that reverse adaptations made by other inputs.

Plasticity is applied uniformly to all synapses. All neurons in a recurrent network produce activity when given input stimulus. Combined, these factors mean that any input sample will cause the same synapses to change. This leads to synaptic competition, interference and ultimately, forgetting.

3.4. Local Plasticity Required to Overcome Interference

To overcome the problem of interference, the mechanisms of plasticity need to be restricted to adapt only a subset of the synapses for any given input stimulus. There is much existing research that supports this conclusion and a number of possible mechanisms that can restrict the locality of plasticity.

It has been shown *in vivo* (using fMRI and neurological experiment) that synaptic plasticity learns highly specific adaptations early in the visual perceptual pathway (Karni and Sagi, 1991; Schwartz et al., 2002). Simulated models of sensory systems have demonstrated that sparsity of activity is essential for sensitivity to input-specific features (Finelli et al., 2008; Barranca et al., 2014). In fact, in a single-layer, non-recurrent structure, STDP is shown to promote sparsity in

a model olfactory system (Finelli et al., 2008). Conversely, in recurrent networks, STDP alone is unable to learn input specific structure because it “over-associates” (Bourjaily and Miller, 2011). Strengthened inhibition was used to overcome this problem and combined with reinforcement learning to produce selectivity in the output (Bourjaily and Miller, 2011). By promoting sparsity, the lack of activity in most of the network will prevent activity-dependent models of plasticity in adapting those connections.

Reward modulated plasticity has also been widely explored in simulated (Gavornik et al., 2009; Darshan et al., 2014) and biological experiment (Li et al., 2013; Lepousez et al., 2014). Input-specific synaptic changes are shown to be strongest in the presence of a reward signal (Gavornik et al., 2009; Lepousez et al., 2014). Lasting memories (synaptic changes not subject to interference), are also seen to rely on a process of re-consolidation consisting of fear conditioning (Li et al., 2013). A reinforcement signal based on either reward or fear conditioning can be effectively used to restrict synaptic changes in a task dependent context such as sensory pattern recognition.

Another way to restrict synaptic changes in a task dependent way is to rely on a back-propagated error signal that has well-established use in artificial neural networks. This might be achieved in a biologically plausible way through axonal propagation (Kempster et al., 2001) or top-down cortical projections sending signals backwards through the sensory pathways (Schäfer et al., 2007). Top-down neural function in general is thought to be essential in determining structure in neural networks (Sharpee, 2014), providing a context for any adaptations. A molecular mechanism for the retro-axon signals required for back-propagation has been proposed (Harris,

2008). However, in general these retro-axon signals are known to be important for neural development but may be too slowly acting to learn sensory input.

3.5. Learning Input Structure Does Not Necessarily Improve Performance

Structural adaptation with plasticity in the pre-training phase, while specific, may not be utilized by the output produced by the network readout. This could be due to the following reasons. Firstly, there is a disparity in the neural code. The output from a recurrent spiking network model is currently decoded as a rate code. In contrast, synaptic plasticity updates structure in a way that depends on the precise temporal activity of neural spikes. Secondly, information content is reduced. While creating associations between co-activating neurons, Hebbian forms of plasticity may also increase correlations and reduce information and separation. These can determine the computational capacity of a recurrent network model (Chrol-Cannon and Jin, 2014b). Both discrepancies could be barriers for the effective application of plasticity to improve pattern recognition. Therefore, new frameworks of neural processing should be based directly on the adapting synapses. This will lead to functional models of neural computing that are not merely improved by synaptic plasticity, but that rely on it as an integral element.

This finding contrasts with some existing work that shows pre-training with plasticity including STDP (Xue et al., 2013) and BCM (Yin et al., 2012) can improve performance in a recurrent spiking network. To address this discrepancy we note that pre-training might improve the general computational properties of recurrent networks without learning input-specific structure. Furthermore, if this is the case, the likelihood of plasticity leading to an improvement will largely depend on how well-tuned the initial parameters of the network are before the pre-training phase begins.

4. Materials and Methods

4.1. Simulation Procedure

The three step procedure depicted in **Figure 1** for training an LSM with plasticity is now described below in pseudocode. Where relevant, some of the expressions within the pseudocode refer to equations that can be found in subsequent subsections where the models for neurons, connectivity, plasticity and pre-processing of inputs can also be found.

Firstly, the following section of pseudocode illustrates the pre-training process in which the recurrent synaptic connections are adapted with plasticity. Input samples are selected at random (scrambled) for a total number of *preTrainIterations* which is 10,000. For a single input sample, each of the time-series frames is presented to the network in sequence by setting the input current of the connected neurons to $W_{in}[x][c] \cdot S[f][x] \cdot inputScale$. The *inputScale* is 20, which is based on the neuron membrane model selected. The neural activity of the network is then simulated for *frameDuration* which is 30 ms. Plasticity is calculated and updated in between each frame of input in a sample. Neural activity is reset for the next input sample.

```
// pre-train recurrent neurons with
                                plasticity
for each iteration I in preTrainIterations
  select random sample S from trainingSamples
  for each frame f in S
    for each attribute x in f
      for each connection c in Cin
        c.input( $W_{in}[x][c] \cdot S[f][x] \cdot$ 
                                inputScale)
    for each timestep t in frameDuration
      neurons.simulateActivity()
      // Equations 2, 3, 4
      synapses.applyPlasticity()
      // Equations 8, 9, 10
      neurons.resetActivity()
```

Secondly, we collect the reservoir states for each sample. The simulation procedure is essentially the same as in pre-training but iterates once for each sample in the dataset. Activity feature vectors are stored in *S.fv* and weight matrix adaptation in *S.dw*.

```
// collect neural activation state vectors
baseWeights.value ← synapses.value
for each sample S in trainingSamples
  for each frame f in S
    for each attribute x in f
      for each connection c in Cin
        c.input( $W_{in}[x][c] \cdot S[f][x] \cdot$ 
                                inputScale)
    for each timestep t in frameDuration
      neurons.simulateActivity()
      // Equations 2, 3, 4
      synapses.applyPlasticity()
      // Equations 8, 9, 10
      S.fv ← neurons.filteredSpikes()
      // Equation 5
      S.dw ← synapses.value - baseWeights.value
      neurons.resetActivity()
      synapses.value ← baseWeights.value
```

Finally, for determining the pattern recognition performance of the LSM, we train a set of readouts using least mean squares regression. There is one readout to predict the presence of each possible class of input. For a total of *readoutTrainingIterations* that is set to 100,000, a randomly selected samples state vector *fv* will be used to adapt the readout weights. The desired signal will be set to 1 for the readout matching the sample class and 0 for the others. For predicting class labels on the training and testing data, the readout with the maximum value for a given *fv* is selected to predict the class (winner takes all).

```
// train readouts with linear regression
for each iteration I in readoutTrainIterations
  select random feature vector fv from
                                trainingSamples.fv
  for each class readout R in nClass
    if R.classLabel = fv.classLabel
```

```

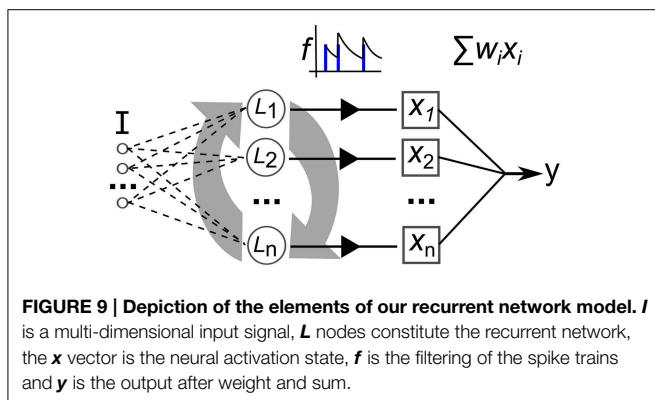
// boost readout for matching class
R.output ← R.lms(fv, 1)
// Equations 6, 7
else
// suppress other readouts
R.output ← R.lms(fv, 0)
// Equations 6, 7
prediction P ← max(R.output)
if P.classLabel ≠ fv.classLabel
errorSum ← errorSum + 1
errorCumulative ← errorSum ÷ I

```

4.2. Recurrent Network

The neural network model used in this work is illustrated in **Figure 9**. Recurrently connected neurons, indicated by L are stimulated by current I that is the sum total of injected current from the input signal, I_{inj} and stimulating current from the pre-synapses, I_{rec} . The total current I perturbs the membrane potential that is modeled with a simple model that matches neuron spiking patterns observed in biology (Izhikevich, 2003). This method for modeling the spiking activity of a neuron is shown to reproduce most naturally occurring patterns of activity (Izhikevich, 2004). The real-valued inputs are normalized between 0 and 1, which are multiplied by a scaling factor of 20 before being injected as current into L . Input connections number $0.2 \cdot \text{network size}$, projected randomly to the network nodes. Weights are uniformly initialized at random between 0 and 1. The video data set used in this work consists of significantly higher dimension inputs—768 features—than the other data sets. Therefore, in this case each feature only projects to one neuron, initially selected at random (a neuron can have connections from multiple inputs). Also the synaptic weights are scaled by 0.25.

The network activity dynamics are simulated for 30 ms for each frame of data in a time-series input sample. This value is chosen as it roughly approximates the actual millisecond delay between digital audio and video data frames. Then, the resulting spike trains produced by each of the neurons are passed through a low-pass filter, f , to produce a real valued vector used to train a linear readout with the iterative, stochastic gradient descent method (each described in the next section).



In our experiments the network consists of 35 or 135 spiking neurons (weight matrix plots consist of 35, performance trials consist of 135) with the ratio of excitatory to inhibitory as 4:1. Neurons are connected with static synapses i.e., the delta impulse (step) function. Connectivity is formed by having $N^2 \cdot C$ synapses that each have source and target neurons drawn according to uniform random distribution, where N is the number of neurons and C is 0.1, the probability of a connection between any two neurons. Weights are drawn from two Gaussian distributions; $\mathcal{N}(6, 0.5)$ for excitatory and $\mathcal{N}(-5, 0.5)$ for inhibitory. When plasticity adapts the reservoir weights, w_{max} is clamped at 10 and w_{min} at -10 . All parameters for excitatory and inhibitory neuron membranes are taken from Izhikevich (2003). The equations for the membrane model are as follows:

$$v' = 0.04v^2 + 5v + 140 - u + I \quad (2)$$

$$u' = a(bv - u) \quad (3)$$

With the spike firing condition:

$$\text{if } v > 30\text{mV} \text{ then } \begin{cases} v \leftarrow c \\ u \leftarrow u + d \end{cases} \quad (4)$$

Parameters for the above equations are; $a = 0.2$, $b = 0.2$, $c = -65$, $d = 8$ for excitatory neurons and; $a = 0.1$, $b = 0.2$, $c = -65$, $d = 2$ for inhibitory neurons.

4.3. Trained Readout

To generate a real-valued output from the discrete spiking activity, the spike train from each neuron is convolved with a decaying exponential according to Equation (5). The vector of values produced is then weighted with the readout weight matrix and summed to produce a single output value, shown in Equation (6).

$$x_i = f(S(t)) = \max \left(\sum_{t=1}^T \exp \left(\frac{-S(t)}{\tau} \right) \right) \quad (5)$$

$$y = \sum_{i=1}^n x_i \cdot w_i \quad (6)$$

The state vector for a neuron is denoted by x_i , the filter function is $f()$ and the spike train is $S(t)$. The maximum number of time-steps in $S(t)$ is T , in this case 50. The decay constant τ is 6 ms.

The maximum value is taken from the low-pass filtered values in Equation (5) in order to detect the highest level of burst activity in the given neuron. We take this approach under the assumption that burst activity is more representative of spiking neural computation than a sum total of the firing rate.

These output weights are updated according to the iterative, stochastic gradient descent method: Least Mean Squares, given in Equation (7).

$$w_i \leftarrow w_i + \mu(y_d - y_o)x_i. \quad (7)$$

Here, y_d is the desired output, y_o is the actual output, x_i is the input taken from a neuron's filtered state, and μ is a small learning rate of 0.005. The weight from x_i to the output is w_i . For the classification tasks of pattern recognition, y_d takes the values of 0 or 1 depending if the class corresponding to the readout is the label of the current input sample.

4.4. Synaptic Plasticity Models

Three synaptic plasticity mechanisms are employed in this study, each of them based on the Hebbian postulate (Hebb, 1949) of "neurons that fire together, wire together." Each mechanism is outlined as follows:

4.4.1. BCM Plasticity

The BCM rule (Bienenstock et al., 1982) is a rate based Hebbian rule that also regulates the post-neuron firing rate to a desired level. It works on a temporal average of pre- and post-synaptic activity. The BCM rule is given in Equation (8). The regulating parameter is the dynamic threshold θ_M , which changes based on the post-synaptic activity y in the following function: $\theta_M = E[y]$, where $E[\cdot]$ denotes a temporal average. In our case, $E[\cdot]$ is calculated as an exponential moving average of the post-synaptic neurons membrane potential. The exponential decay coefficient used for this is 0.935. As the membrane potential is model-dependant, we normalize it between 0..1 in real-time by continuously updating max and min variables of previous values. There is also a uniform decay parameter ϵw set as 0.0001 that slowly reduces connection strength and so provides a means for weight decay, irrespective of the level of activity or correlation between pre-synaptic inputs and post synaptic potential. A plot of the BCM weight change is presented in Supplementary Figure 9.

$$\Delta w = y(y - \theta_M)x - \epsilon w \quad (8)$$

4.4.2. Bi-phasic STDP

The STDP rule depends on the temporal correlation between pre- and post-synaptic spikes. The synaptic weight change is computed based on the delay between the firing times of the pre- and post- neuron. This is described in a fixed "learning window" in which the y-axis is the level of weight change and the x-axis is the time delay between a pre- and post-synaptic spike occurrence. The bi-phasic STDP rule consists of two decaying exponential curves (Song et al., 2000), a positive one to potentiate in-order spikes, and a negative one to depress out-of-order spikes. This rule was derived from experimental work carried out on populations of neurons *in vitro* (Markram et al., 1997; Bi and Poo, 1998). Bi-phasic STDP is given in Equation (9).

$$\Delta w(\Delta t) = \begin{cases} A_+ \cdot \exp\left(\frac{-\Delta t}{\tau_+}\right) & \text{if } t > 0 \\ -A_- \cdot \exp\left(\frac{\Delta t}{\tau_-}\right) & \text{if } t \leq 0 \end{cases} \quad (9)$$

A_+ and A_- are the learning rates for the potentiation and depression, respectively. Δt is the delay of the post-synaptic spike occurring after the transmission of the pre-synaptic spike. τ_+ and τ_- control the rates of the exponential decrease in plasticity across the learning window. For our experiments the learning window is symmetric with $A_+ = A_- = 0.15$ and $\tau_+ = \tau_- = 20$ ms.

4.4.3. Tri-phasic STDP

A tri-phasic STDP learning window consists of a narrow potentiating region for closely correlated activity but depressing regions on either side: for recently uncorrelated activity, and for correlated but late activity. This learning window has been observed *in vitro*, most notably in the hippocampi, between areas CA3 and CA1 (Wittenberg and Wang, 2006). The tri-phasic STDP is given in Equation (10).

$$\Delta w(\Delta t) = A_+ \exp\left(\frac{-(\Delta t - 15)^2}{200}\right) - A_- \exp\left(\frac{-(\Delta t - 15)^2}{2000}\right) \quad (10)$$

The learning rates are set as $A_+ = 0.25$ and $A_- = 0.1$. Both STDP learning windows are plotted in Supplementary Figure 10.

4.5. Synaptic Interference Measure

We wish to quantify interference directly between synaptic adaptations of plasticity. Our formulation of synaptic interference is based on the synaptic changes from sequentially presented samples. Synaptic adaptation for a given class of sample is called ΔW_t and average adaptation for all others are ΔW_o . Interference must be calculated individually for each class of sample, I_t^{class} , and averaged together to get the overall interference, I^{total} . The equations are as follows:

$$I_t^{\text{class}} = \frac{1}{N} \sum_{i=1}^N [\Delta W_{ti} \cdot \Delta W_{oi} < 0][|\Delta W_{ti}| < |\Delta W_{oi}| \cdot C_n] \quad (11)$$

$$I^{\text{total}} = \sum_{t=1}^{C_n} \frac{I_t^{\text{class}}}{C_n} \quad (12)$$

Where I is interference, N is the number of synapses, C_n is the number of competing sample classes and ΔW is a vector of synaptic changes. Subscript i denotes the parameter index, subscript t denotes samples of a given class "this" and subscript o denotes samples of all "other" classes.

In Equation (11), the first set of Iverson brackets returns 1 if synaptic adaptation of a given class is of a different sign than that of the average adaptation of other class samples. The second set of Iverson Brackets returns 1 only if the magnitude of the synaptic adaptation of a class is less than the average weight adaptation of other classes multiplied by the total number. This leads to us taking a conservative measure of synaptic interference where we will only flag interference within a synapse for a class of pattern if the weight change is in a different direction to the average as well as being lower in magnitude than the total weight adaptation of other inputs.

4.6. Synthetic Signal Data

A synthetic benchmark task is taken from a study performed with Echo State Networks (Jaeger, 2007), a similar type of network model to the one we employ, but using continuous rate-based neurons instead. The task is to predict which of three signal generating functions is currently active in producing a time-varying input signal. To generate a sample of the signal at a given time step, one of the three following function types is

used; (1) A sine function of a randomly selected period, (2) A chaotic iterated tent map, (3) A randomly chosen constant. The generator is given some low probability, 0.05, of switching to another function at each time-step. The full method of generating the data is described in Jaeger (2007). A short window of the generated signal is plotted in **Figure 4**.

4.7. Speaker Recognition Data

A speaker recognition task is a classification problem dealing with mapping time-series audio input data to target speaker labels. We use a data set taken from Kudo et al. (1999) which consists of utterances of nine male Japanese speakers pronouncing the vowel /ae/. The task is to correctly discriminate each speaker based on the speech samples. Each sample is comprised of a sequence of 12 feature audio frames. The features of each frame are the LPC cepstrum coefficients. The sample sequence ranges between 7 and 29 frames. The dataset is divided into training and testing sets of 270 and 370 samples each, respectively. Note that unlike the benchmark data used in this report, the samples are not in a consecutive time-series, yet each sample consists of a time-series sequence of audio frames.

4.8. Pre-processing of the Human Motion Data

A visual task is selected to test high dimensional spatial-temporal input data. The KTH data set (Schuldt et al., 2004) consists of 2391 video files of people performing one of six actions; boxing, clapping, waving, walking and jogging. There are 25 different subjects and the samples cover a range of conditions that are described in more detail in Schuldt et al. (2004). Each video sample is taken at 25 frames per second and down sampled to a resolution of 160×120 pixels. We process the raw video sequences according to a formula shown in the following equations:

$$M(t) = \|\Delta(I_1, I_2), \dots, \Delta(I_{N-1}, I_N)\| \quad (13)$$

$$M(t, i) = \begin{cases} 1 & \text{if } M(t, i) \geq 0.2 \cdot \max(M(\cdot)) \\ 0 & \text{else} \end{cases} \quad (14)$$

The final input matrix M is indexed by time-frames, t and spatial samples i . Column vectors I_n are individual frames, re-shaped into one dimension. Each sample contains up to a total of N frames. In plain language, this process essentially further down samples by a factor of 0.2 and calculates the difference between pixels in consecutive frames, which are then used as the new input features. Each frame is then re-shaped into a single dimensional column vector then appended together to form an input matrix in which each column is used as the neural network input at consecutive time steps. **Figure 3** shows frames extracted from an example of each type on motion along with the corresponding processed features.

Author Contributions

Conception and design of the work was by YJ and JC. Experiments were performed by JC. Analysis and interpretation were undertaken by JC and YJ. Manuscript was written by JC and revised by YJ.

Acknowledgments

JC was supported by Engineering and Physical Sciences Research Council's Doctoral Training Grant through University of Surrey.

Supplementary Material

The Supplementary Material for this article can be found online at: <http://journal.frontiersin.org/article/10.3389/fncom.2015.00103>

References

- Barranca, V. J., Kovacic, G., Zhou, D., and Cai, D. (2014). Sparsity and compressed coding in sensory systems. *PLoS Comput. Biol.* 10:e1003793. doi: 10.1371/journal.pcbi.1003793
- Bi, G. Q., and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Bienenstock, E. L., Cooper, L. N., and Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J. Neurosci.* 2, 32–48.
- Bourjaily, M. A., and Miller, P. (2011). Synaptic plasticity and connectivity requirements to produce stimulus-pair specific responses in recurrent networks of spiking neurons. *PLoS Comput. Biol.* 7:e1001091. doi: 10.1371/journal.pcbi.1001091
- Buonomano, D. V., and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nat. Rev. Neurosci.* 10, 113–25. doi: 10.1038/nrn2558
- Chrol-Cannon, J., and Jin, Y. (2014a). Computational modeling of neural plasticity for self-organization of neural networks. *Biosystems* 125, 43–54. doi: 10.1016/j.biosystems.2014.04.003
- Chrol-Cannon, J., and Jin, Y. (2014b). On the correlation between reservoir metrics and performance for time series classification under the influence of synaptic plasticity. *PLoS ONE* 9:e101792. doi: 10.1371/journal.pone.0101792
- Darshan, R., Leblois, A., and Hansel, D. (2014). Interference and shaping in sensorimotor adaptations with rewards. *PLoS Comput. Biol.* 10:e1003377. doi: 10.1371/journal.pcbi.1003377
- Finelli, L. A., Haney, S., Bazhenov, M., Stopfer, M., and Sejnowski, T. J. (2008). Synaptic learning rules and sparse coding in a model sensory system. *PLoS Comput. Biol.* 4:e1000062. doi: 10.1371/journal.pcbi.1000062
- French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends Cogn. Sci.* 3, 128–135. doi: 10.1016/S1364-6613(99)01294-2
- Gavornik, J. P., Shuler, M. G. H., Loewenstein, Y., Bear, M. F., and Shouval, H. Z. (2009). Learning reward timing in cortex through reward dependent expression of synaptic plasticity. *Proc. Natl. Acad. Sci. U.S.A.* 106, 6826–6831. doi: 10.1073/pnas.0901835106
- Gilson, M., Burkitt, A., and Van Hemmen, L. J. (2010). Sdp in recurrent neuronal networks. *Front. Comput. Neurosci.* 4:23. doi: 10.3389/fncom.2010.00023
- Harris, K. D. (2008). Stability of the fittest: organizing learning through retroaxonal signals. *Trends Neurosci.* 31, 130–136. doi: 10.1016/j.tins.2007.12.002
- Hebb, D. (1949). *The Organization of Behavior*. New York, NY: Wiley.
- Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Trans. Neural Netw.* 14, 1569–1572. doi: 10.1109/TNN.2003.820440
- Izhikevich, E. M. (2004). Which model to use for cortical spiking neurons? *IEEE Trans. Neural Netw.* 15, 1063–1070. doi: 10.1109/TNN.2004.832719

- Jaeger, H. (2007). *Discovering Multiscale Dynamical Features with Hierarchical Echo State Networks*. Technical report, Jacobs University Bremen.
- Karni, A., and Sagi, D. (1991). Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. *Proc. Natl. Acad. Sci. U.S.A.* 88, 4966–4970. doi: 10.1073/pnas.88.11.4966
- Kempter, R., Leibold, C., Wagner, H., and van Hemmen, J. L. (2001). Formation of temporal-feature maps by axonal propagation of synaptic learning. *Proc. Natl. Acad. Sci. U.S.A.* 98, 4166–4171. doi: 10.1073/pnas.061369698
- Kudo, M., Toyama, J., and Shimbo, M. (1999). Multidimensional curve classification using passing-through regions. *Pattern Recognit. Lett.* 20, 1103–1111. doi: 10.1016/S0167-8655(99)00077-X
- Lepousez, G., Nissant, A., Bryant, A. K., Gheusi, G., Greer, C. A., and Lledo, P.-M. (2014). Olfactory learning promotes input-specific synaptic plasticity in adult-born neurons. *Proc. Natl. Acad. Sci. U.S.A.* 111, 13984–13989. doi: 10.1073/pnas.1404991111
- Li, Y., Meloni, E. G., Carlezon, W. A. Jr., Milad, M. R., Pitman, R. K., Nader, K., et al. (2013). Learning and reconsolidation implicate different synaptic mechanisms. *Proc. Natl. Acad. Sci. U.S.A.* 110, 4798–4803. doi: 10.1073/pnas.1217878110
- Maass, W., Natschläger, T., and Markram, H. (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* 14, 2531–2560. doi: 10.1162/089976602760407955
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic apss and epsps. *Science* 275, 213–215. doi: 10.1126/science.275.5297.213
- McCloskey, M., and Cohen, N. (1989). Catastrophic interference in connectionist networks: the sequential learning problem. *Psychol. Learn. Motiv.* 24, 109–164. doi: 10.1016/S0079-7421(08)60536-8
- Ratcliff, R. (1990). Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychol. Rev.* 97, 285–308. doi: 10.1037/0033-295X.97.2.285
- Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol. Rev.* 65, 386–408. doi: 10.1037/h0042519
- Schäfer, R., Vasilaki, E., and Senn, W. (2007). Perceptual learning via modification of cortical top-down signals. *PLoS Comput. Biol.* 3:e165. doi: 10.1371/journal.pcbi.0030165
- Schuldt, C., Laptev, I., and Caputo, B. (2004). “Recognizing human actions: a local SVM approach,” in *ICPR Proceedings of the 17th International Conference on Pattern Recognition*, Vol. 3 (Cambridge, UK: IEEE), 32–36. doi: 10.1109/ICPR.2004.1334462
- Schwartz, S., Maquet, P., and Frith, C. (2002). Neural correlates of perceptual learning: a functional mri study of visual texture discrimination. *Proc. Natl. Acad. Sci. U.S.A.* 99, 17137–17142. doi: 10.1073/pnas.242414599
- Sharpee, T. O. (2014). Function determines structure in complex neural networks. *Proc. Natl. Acad. Sci. U.S.A.* 111, 8327–8328. doi: 10.1073/pnas.1407198111
- Song, S., Miller, K. D., and Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.* 3, 919–926. doi: 10.1038/78829
- Toutounji, H., and Pipa, G. (2014). Spatiotemporal computations of an excitable and plastic brain: Neuronal plasticity leads to noise-robust and noise-constructive computations. *PLoS Comput. Biol.* 10:e1003512. doi: 10.1371/journal.pcbi.1003512
- Wittenberg, G. M., and Wang, S. S. (2006). Malleability of spike-timing-dependent plasticity at the ca3-ca1 synapse. *J. Neurosci.* 26, 6610–6617. doi: 10.1523/JNEUROSCI.5388-05.2006
- Xue, F., Hou, Z., and Li, X. (2013). Computational capability of liquid state machines with spike-timing-dependent plasticity. *Neurocomputing* 122, 324–329. doi: 10.1016/j.neucom.2013.06.019
- Yin, J., Meng, Y., and Jin, Y. (2012). A developmental approach to structural self-organization in reservoir computing. *IEEE Trans. Auton. Ment. Dev.* 4, 273–289. doi: 10.1109/TAMD.2012.2182765

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2015 Chrol-Cannon and Jin. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.



Excitatory and inhibitory STDP jointly tune feedforward neural circuits to selectively propagate correlated spiking activity

Florence I. Kleberg*, Tomoki Fukai and Matthieu Gilson

Lab for Neural Circuit Theory, RIKEN Brain Science Institute, Wako, Japan

Edited by:

Cristina Savin, IST Austria, Austria

Reviewed by:

Maoz Shamir, Boston University, USA

Paul Miller, Brandeis University, USA

Ashok Litwin-Kumar, University of

Pittsburgh, USA

***Correspondence:**

Florence I. Kleberg, Lab for Neural Circuit Theory, RIKEN Brain Science Institute, 2-1 Hirosawa, Wako, Saitama, 351-0198, Japan
e-mail: kleberg@brain.riken.jp

Spike-timing-dependent plasticity (STDP) has been well established between excitatory neurons and several computational functions have been proposed in various neural systems. Despite some recent efforts, however, there is a significant lack of functional understanding of inhibitory STDP (iSTDP) and its interplay with excitatory STDP (eSTDP). Here, we demonstrate by analytical and numerical methods that iSTDP contributes crucially to the balance of excitatory and inhibitory weights for the selection of a specific signaling pathway among other pathways in a feedforward circuit. This pathway selection is based on the high sensitivity of STDP to correlations in spike times, which complements a recent proposal for the role of iSTDP in firing-rate based selection. Our model predicts that asymmetric anti-Hebbian iSTDP exceeds asymmetric Hebbian iSTDP for supporting pathway-specific balance, which we show is useful for propagating transient neuronal responses. Furthermore, we demonstrate how STDPs at excitatory–excitatory, excitatory–inhibitory, and inhibitory–excitatory synapses cooperate to improve the pathway selection. We propose that iSTDP is crucial for shaping the network structure that achieves efficient processing of synchronous spikes.

Keywords: STDP, spike-timing, plasticity, inhibition, disynaptic, correlation, excitation–inhibition balance

1. INTRODUCTION

Activity-dependent plasticity of synaptic connections between neurons is crucial for cortical circuit development and memory (Böhme et al., 1993; Hensch et al., 1998). Spike-timing-dependent plasticity (STDP) describes the change in synaptic weights where long-term potentiation (LTP) and long-term depression (LTD) depend on the precise timing of presynaptic and postsynaptic action potentials. STDP has been observed for excitatory glutamatergic synapses in a great diversity of brain structures, such as the hippocampus (Magee and Johnston, 1997; Bi and Poo, 1998; Debanne et al., 1998), the cerebellum of the electric fish (Bell et al., 1997), the neocortex (Markram et al., 1997; Sjöström et al., 2001), and the optic nerve in *Xenopus* (Zhang et al., 1998). An extensive body of theoretical work has uncovered many interesting properties of excitatory STDP (eSTDP): it can select input pathways based on their spike-time correlation (Kempster et al., 1999; Song et al., 2000; Gjorgjieva et al., 2011), it can generate a stable distribution of weights (van Rossum et al., 2000; Gütig et al., 2003; Gilson and Fukai, 2011), it can perform selection of phase-locking in population firing (Gerstner et al., 1996; Senn and Buchs, 2003), it favors the emergence of functional neuronal assemblies (Izhikevich et al., 2004; Clopath et al., 2010), it stabilizes slow oscillations in recurrent networks (Kang et al., 2008) and it allows for rewiring of connections in the developing visual cortex (Song and Abbott, 2001; Senn and Buchs, 2003; Young et al., 2007).

There is also evidence for STDP at inhibitory GABAergic synapses, or iSTDP, (Woodin et al., 2003; Haas et al., 2006;

Kodangattil et al., 2013). However, our understanding of the mechanistic implications of iSTDP remains limited, in spite of the key role of inhibition in signal processing in the cortex (van Vreeswijk and Sompolinsky, 1996; Anderson et al., 2000; Wehr and Zador, 2003; Haider et al., 2006, 2013; Maffei et al., 2006; Rudolph et al., 2007). Considering the abundance of inhibitory interneurons, e.g., in the cortex (Markram et al., 2005), remarkably few types have been tested for the plastic properties of their synapses. Theoretical knowledge of the dynamics and functional implications of iSTDP is also rudimentary, although interest in this direction has increased recently. Extending a simple homeostatic control of firing rate, iSTDP can generate a balance between excitatory and inhibitory inputs onto a neuron (Vogels et al., 2011). In addition, unspecific, but sufficiently strong inhibition developed by iSTDP can enhance competition between excitatory synapses subject to eSTDP (Luz and Shamir, 2012). Interestingly, neither experimental nor theoretical approaches provide a consensus for the shape of the iSTDP learning window, in contrast to eSTDP for which the temporally Hebbian nature (LTP for pre-post pairing, LTD for post-pre pairing) is observed and addressed in the vast majority of cases.

The present paper aims to compare the effect of distinct iSTDP window shapes on the structure of synaptic weights, and endeavors to clarify the role of iSTDP in tuning neuronal responses. Previous studies have shown the precise timing of spikes to convey an important part of information about stimuli in sensory pathways (Riehle et al., 1997; Jackson et al., 2003;

Maldonado et al., 2008; Kilavik et al., 2009; Putrino et al., 2010). Moreover, neurons are sensitive to precise timings of spikes (Rossant et al., 2011). In this context of neural temporal coding, we examined the transmission of temporally correlated spikes in a feedforward circuit equipped with eSTDP and iSTDP. Such neural architectures with joint feedforward excitation and inhibition have been found in various brain structures (Buzsaki, 1984; Davis et al., 1996). We incorporated in our model an important property of the afferent inputs observed experimentally in many feedforward neural architectures: inhibition is delayed compared to excitation with a short time lag (Pouille and Scanziani, 2001; Wilent and Contreras, 2004; Gabernet et al., 2005; Silberberg and Markram, 2007; Tan et al., 2008; Stokes and Isaacson, 2010), which allows for precise temporal gating (Kremkow et al., 2010). We found that iSTDP with specifically anti-Hebbian properties enforces a balanced structure in the synaptic weights, which supports efficient processing of near-coincident spikes.

2. RESULTS

We examined the joint development of excitatory and inhibitory synapses subject to STDP in a feedforward circuit. We consider two circuit architectures. First, for a single neuron with direct excitatory and inhibitory inputs, we examine how the shape of the iSTDP window affects the evolution of synaptic weights. Since there is no current consensus about a single type of iSTDP (Woodin et al., 2003; Haas et al., 2006; Kodangattil et al., 2013), this comparison allows us to link the shape of iSTDP learning windows to their functional implications. Second, we examine the recruitment of interneurons in a more realistic circuit with monosynaptic excitation and disynaptic inhibition. In both cases, we focus on how the emerging weight structure tunes the propagation of spike volleys in the circuit.

2.1. THEORETICAL PREDICTION OF WEIGHT SPECIALIZATION DEPENDING ON iSTDP WINDOW

In order to study the weight dynamics for different iSTDP learning windows, we consider a simplified feedforward circuit (SFC) that consists of a single postsynaptic Poisson neuron excited by excitatory and inhibitory spike trains (Figure 1A). Following experimental observations, excitatory and inhibitory inputs have correlated spiking activity (Okun and Lampl, 2008). In addition, inhibition arrives with a delay d (Okun and Lampl, 2008; Atallah and Scanziani, 2009). The inhibitory delay mimics a disynaptic pathway, as compared to monosynaptic excitation (Figure 1B). The temporal correlations between spikes trains in Figure 1B are governed by the time constant τ_{in} (Figure 1C). All Synapses are plastic.

Excitatory weights are modified by a temporally Hebbian eSTDP rule (Gilson and Fukai, 2011), corresponding to the blue learning window in Figure 1D: a presynaptic spike preceding a postsynaptic spike leads to potentiation. The eSTDP update includes log-STDP weight dependence, which produces a long-tailed distribution of weights (Gilson and Fukai, 2011). For every pair of a pre- and a postsynaptic spikes, the weight w_e is modified by a quantity that depends

on the current value of w_e and the spike-time difference $\Delta t = t_{pre} - t_{post}$:

$$\Delta w_e = W_e(\Delta t, w_e) = \begin{cases} \eta_e \exp\left(\frac{\Delta t}{\tau_{LTP}^e}\right) a_{LTP} \exp\left(-\frac{C_{LTP} w_e}{w_0}\right) & \text{for } \Delta t < 0, \\ \eta_e \exp\left(-\frac{\Delta t}{\tau_{LTD}^e}\right) a_{LTD} \exp\left(\frac{\log(1 + w_e/w_0 C_{LTD})}{\log(1 + C_{LTD})}\right) & \text{for } \Delta t > 0. \end{cases} \quad (1)$$

The time constants $\tau_{LTP}^e = 17$ ms $\tau_{LTD}^e = 34$ ms and coefficients $a_{LTP} = 1$ and $a_{LTD} = -0.5$ determine the shape of the eSTDP window. η_e is the learning rate. The log-style weight dependence scales the LTD curve and ensures a stable fixed point at $w_0 = 0.065$ for uncorrelated inputs; $C_{LTD} = 5$ enforces sufficiently strong competition between the incoming weights onto a given neuron. An exhaustive eSTDP parameter list is given in Table 1.

For inhibitory synapses, we test three types of additive iSTDP windows, shown in orange in Figure 1E:

- Hebbian (Haas et al., 2006; Luz and Shamir, 2012; Kodangattil et al., 2013), with pre-post LTP;
- Anti-Hebbian, with post-pre LTP;
- Symmetric (Vogels et al., 2011), with which LTP occurs for pre-post and post-pre spike pairings.

For every spike pair the inhibitory weight is updated with

$$\Delta w_i = W_i(\Delta t) = \begin{cases} \eta_i p \exp\left(-\frac{\Delta t}{\tau_{post}^i}\right) & \text{for } \Delta t > 0, \\ \eta_i q \exp\left(\frac{\Delta t}{\tau_{pre}^i}\right) & \text{for } \Delta t < 0. \end{cases} \quad (2)$$

The right and left sides of the iSTDP window can be either LTP or LTD depending on the sign of p and q , respectively. Table 2 lists the values of p and q for the three windows employed in the theoretical model and in the simulations. For all iSTDP window types, total LTP exceeds total LTD; for anti-Hebbian, Hebbian, and symmetric (corrected), the difference LTP—LTD is set equal. Additionally, $\tau_{pre}^i = \tau_{post}^i$ for all iSTDP.

To stabilize iSTDP, each inhibitory weight is decreased by a small amount α for every presynaptic spike (Vogels et al., 2011), independently of the iSTDP contribution.

$$W_i \rightarrow W_i - \eta_i \alpha \quad (3)$$

Our aim is to show the emergence of balance between excitation and inhibition through eSTDP and iSTDP. By balance we mean the simultaneous increase of excitatory and inhibitory weights, or weight balance, as opposed to increase of excitatory weights without increase of inhibitory weights. Balance known as the cancelation of currents onto a neuron (e.g., Vogels et al., 2011) can but need not follow from weight balance. Unless otherwise stated, balance in this study means weight balance. Using our analysis based on the Poisson neuron model (Materials and Methods), we evaluate the *expected change* in mean synaptic strengths for both sets of weights. The weight update is determined by the interplay of the iSTDP window, the input spike-time

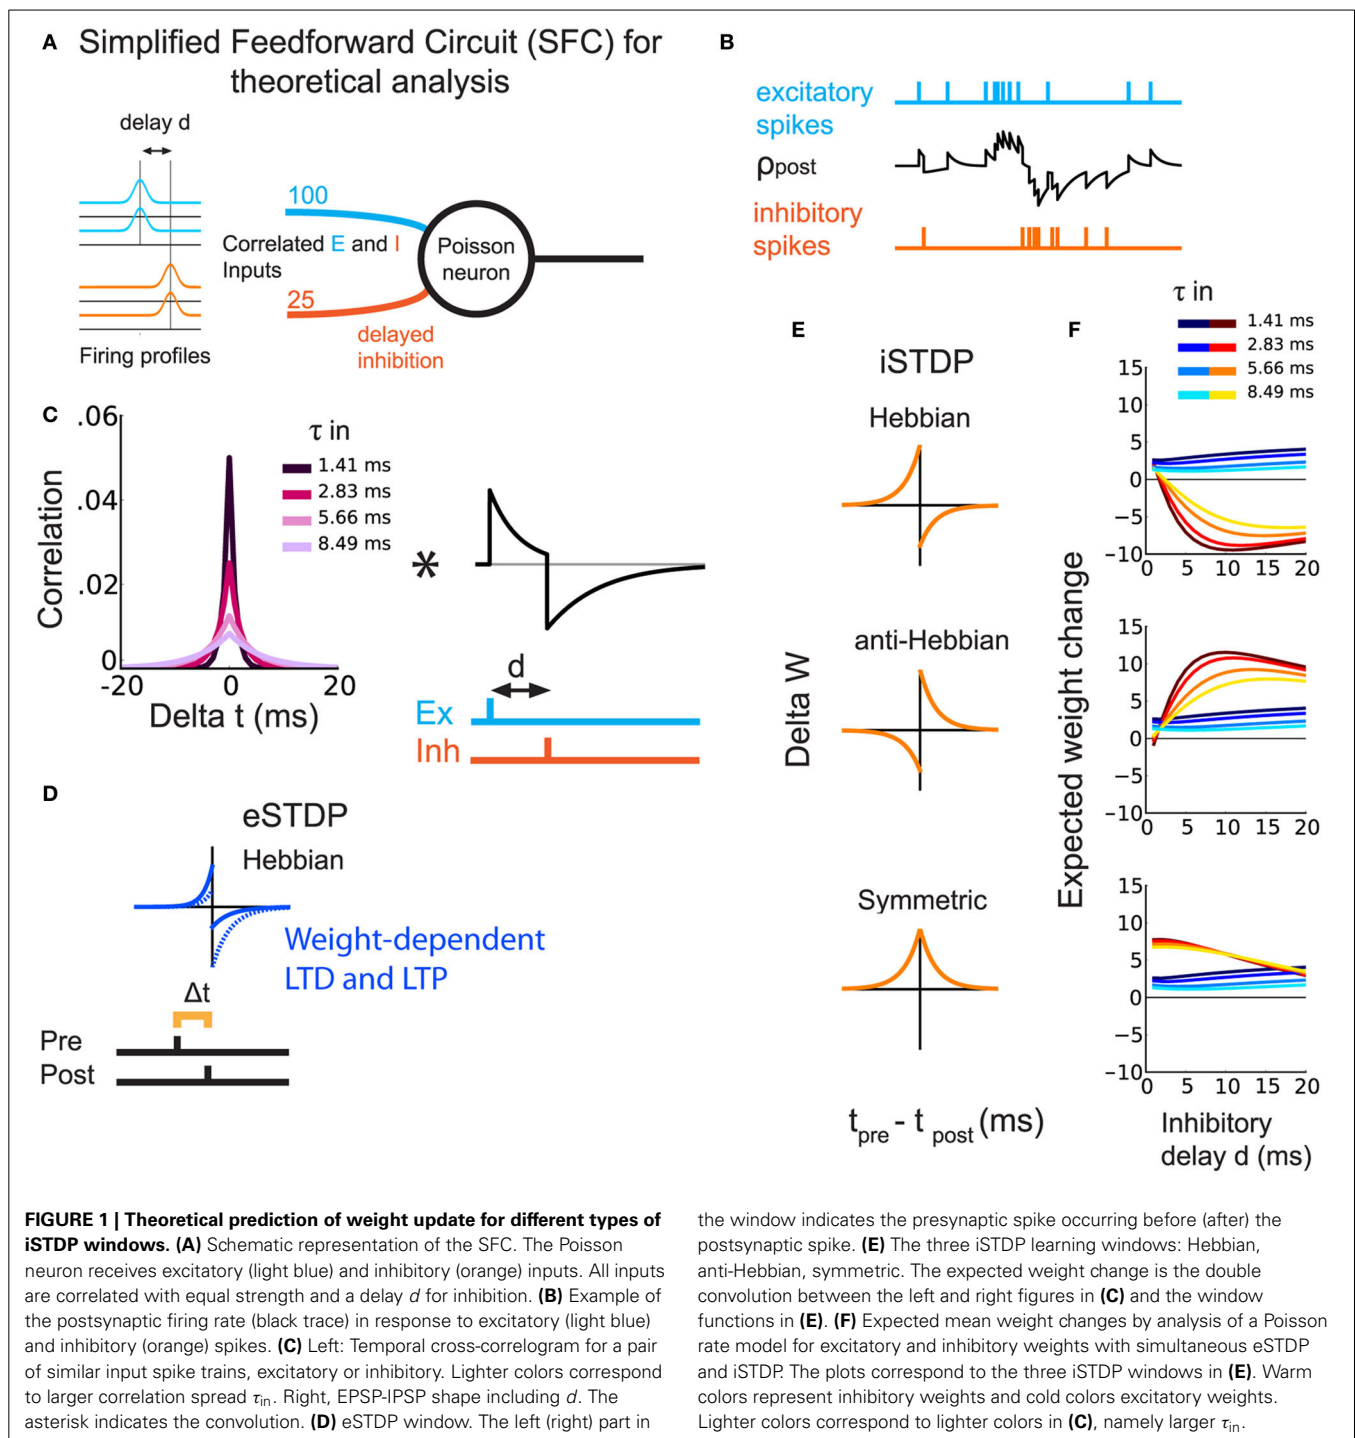


FIGURE 1 | Theoretical prediction of weight update for different types of iSTDP windows. (A) Schematic representation of the SFC. The Poisson neuron receives excitatory (light blue) and inhibitory (orange) inputs. All inputs are correlated with equal strength and a delay d for inhibition. **(B)** Example of the postsynaptic firing rate (black trace) in response to excitatory (light blue) and inhibitory (orange) spikes. **(C)** Left: Temporal cross-correlogram for a pair of similar input spike trains, excitatory or inhibitory. Lighter colors correspond to larger correlation spread τ_{in} . Right, EPSP-IPSP shape including d . The asterisk indicates the convolution. **(D)** eSTDP window. The left (right) part in

the window indicates the presynaptic spike occurring before (after) the postsynaptic spike. **(E)** The three iSTDP learning windows: Hebbian, anti-Hebbian, symmetric. The expected weight change is the double convolution between the left and right figures in **(C)** and the window functions in **(E)**. **(F)** Expected mean weight changes by analysis of a Poisson rate model for excitatory and inhibitory weights with simultaneous eSTDP and iSTDP. The plots correspond to the three iSTDP windows in **(E)**. Warm colors represent inhibitory weights and cold colors excitatory weights. Lighter colors correspond to lighter colors in **(C)**, namely larger τ_{in} .

cross-correlograms and the postsynaptic response (EPSPs+IPSPs; **Figure 1C**). All expected weight changes in **Figure 1F** (and in subsequent sections, weights themselves) are shown after division by the excitatory equilibrium weight w_0 , given in **Table 1**. The influence of the inhibitory delay d on the expected change in excitatory weights is weak—slight increase when d becomes larger—and does not depend much on the iSTDP learning window (**Figure 1F**, curves in cold colors). However, d affects

the evolution of inhibitory weights, as shown by the curves in hot colors in **Figure 1F**. For Hebbian iSTDP, inhibitory weights decrease with a stronger effect for larger delays (≥ 5 ms). Conversely, anti-Hebbian iSTDP causes weights to increase. Symmetric iSTDP leads to a potentiation that weakens for large delays.

In all cases, larger values for the input correlation width τ_{in} decrease the effect of both eSTDP and iSTDP (curves in lighter

Table 1 | General STDP parameters.

STDP parameters: SFC	Description	Value
W_e, W_i	eSTDP or iSTDP window	
$\tau_{pre}^e, \tau_{post}^e$	eSTDP window time constants	34 ms, 17 ms
$\tau_{pre}^i, \tau_{post}^i$	iSTDP window time constants	30 ms
–	Excitatory start-up weights	random $U[0, w_0 \times 3]$
–	Inhibitory start-up weights	0
w_0	Equilibrium weight for eSTDP	0.065
η_e	Excitatory learning rate	$w_0 \times 0.39$
η_i	Inhibitory learning rate	0.075
α	inhibitory presynaptic single-spike contribution	0.2
C_{LTD}	LTD scaling for eSTDP	5
C_{LTP}	LTP scaling for eSTDP	50
STDP PARAMETERS: FFC		
w_0	Equilibrium weight for eSTDP onto output neuron	0.08
η_e	Excitatory learning rate	$w_0 \times 0.78$
η_i	Inhibitory learning rate	0.02
–	Excitatory start-up weights	random $U[0, 1]$
–	Excitatory-to-Inhibitory start-up weights	random $U[0, 1]$
–	Inhibitory start-up weights	1
$\tau_{pre}^i, \tau_{post}^i$	iSTDP window time constants	20 ms
w_0^i	Equilibrium weight for eSTDP onto interneurons	$w_0 \times 2$
–	Excitatory start-up weights onto interneurons	random $U[0, 1]$

Table listing all eSTDP and iSTDP parameters, with the exception of the iSTDP window parameters (Table 2). The notation “random $U[a, b]$ ” denotes a random distribution between a and b .

Table 2 | iSTDP window parameters.

iSTDP type	Hebbian	anti-Hebbian	Symmetric	Symmetric (equal LTP/LTD)
p	–1	1.5	1.5	0.25
q	1.5	–1	1.5	0.25

Table listing the different iSTDP windows. p and q indicate the amplitude of the right and left side of the iSTDP window, respectively.

colors in Figure 1F). In fact, Hebbian and anti-Hebbian iSTDP curves exhibit a delay for which the weight change is maximal. That “best” delay increases when τ_{in} is large. Symmetric iSTDP is less affected by τ_{in} .

In summary, the simultaneous strengthening of correlated excitatory and inhibitory inputs (i.e., the emergence of balance) should occur when iSTDP has an anti-Hebbian LTP component in this simple circuit (anti-Hebbian and symmetric iSTDP), and

when inhibitory input spikes arrive after postsynaptic spikes with a sufficiently large d (axonal delay in a feedforward inhibitory circuit).

2.2. EMERGENCE OF A DETAILED BALANCE BETWEEN EXCITATORY AND INHIBITORY WEIGHTS

Next, we verify our theoretical predictions for the SFC with simulations using a LIF neuron (Materials and Methods: details of the simulated SFC, Equations 8, 9). In contrast to Figure 1, the SFC in Figure 2A includes a distractor pathway with random, uncorrelated inputs (Figure 2A, dark blue and red lines) besides the correlated inputs (light blue and orange lines).

A typical example of synaptic weight evolution with anti-Hebbian iSTDP is shown in Figure 2B1. The weights from random inputs remain weak (dark blue and red traces), whereas the weights from the correlated excitatory inputs (light blue traces) and inhibitory inputs (orange traces) are strengthened, indicating the development of within-pathway balance, or *detailed balance* (Vogels and Abbott, 2009). In detailed balance, the excitatory and inhibitory inputs to strong weights on a postsynaptic neuron have correlated spike times (or spike rates, Vogels and Abbott, 2009). In contrast, when excitation is balanced with inhibition from a different signal pathway, excitation and inhibition are not necessarily correlated, which we may call *global balance*. Both types of balance will be evaluated in the sections below. The histograms of the final weight distributions of this example show the development of weight structure for excitation and inhibition. Excitatory weights exhibit a long-tail distribution that follows from the log-type weight dependence used for eSTDP (Gilson and Fukai, 2011) (Figure 2B2: top). The distribution of inhibitory weights has a long tail as well (Figure 2B2: bottom), but looks more bimodal for smaller τ_{in} due to increased competition between the weights (not shown).

As in our analytical approach, we compared the effect of different iSTDP windows on the inhibitory weights of an inhibitory–excitatory pathway with correlated spike-times. The comparison of the Hebbian, anti-Hebbian, and symmetric iSTDP windows agrees with the theoretical predictions of expected drift in weights in Figure 1F. Correlated inhibitory weights increase with both symmetric (Figure 2D, black curves) and anti-Hebbian iSTDP (magenta curves). Their final equilibrium value depends on d (see also Figure 2C): short delays are preferred only by symmetric iSTDP (Figure 2D, black curves). Anti-Hebbian iSTDP leads to a larger increase in inhibitory weights than symmetric iSTDP for larger delays d , and small τ_{in} . We also test an additional version of symmetric iSTDP: apart from the window with the same maximal amplitude as Hebbian and anti-Hebbian iSTDP (Figure 1E, bottom; Figure 2D, black curves), we apply a symmetric rule with the same amount of LTP-LTD area (equalized symmetric iSTDP; gray curves). Since this equalized window only leads to very small changes, we conclude that the amplitude of LTP is the crucial factor, not the overall LTP/LTD ratio. Lastly, inhibitory weights vanish to zero with Hebbian iSTDP (red curves). These findings confirm that the neuron first becomes driven by the correlated excitatory inputs through eSTDP; then, when excitatory inputs dictate postsynaptic firing times, correlated inhibitory inputs follow up through iSTDP. The increase

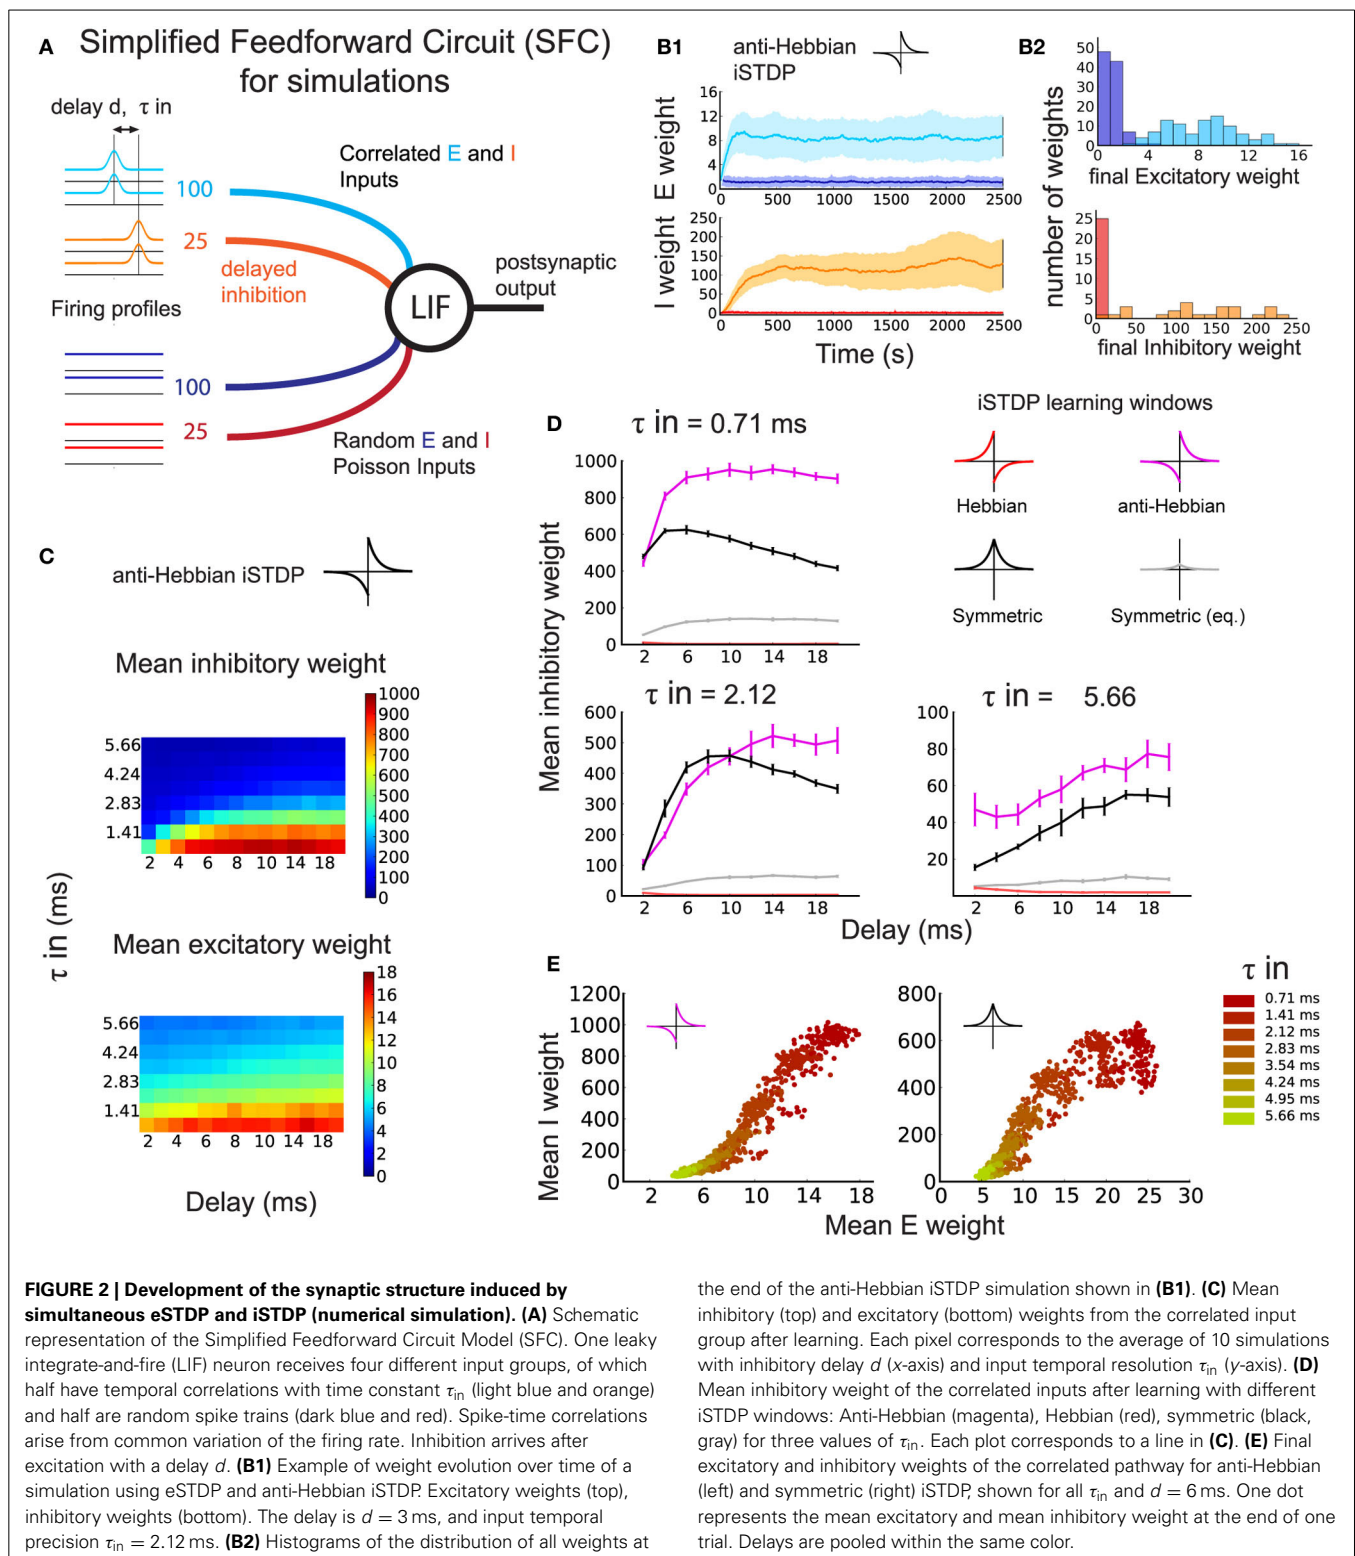


FIGURE 2 | Development of the synaptic structure induced by simultaneous eSTDP and iSTDP (numerical simulation). (A) Schematic representation of the Simplified Feedforward Circuit Model (SFC). One leaky integrate-and-fire (LIF) neuron receives four different input groups, of which half have temporal correlations with time constant τ_{in} (light blue and orange) and half are random spike trains (dark blue and red). Spike-time correlations arise from common variation of the firing rate. Inhibition arrives after excitation with a delay d . (B1) Example of weight evolution over time of a simulation using eSTDP and anti-Hebbian iSTDP. Excitatory weights (top), inhibitory weights (bottom). The delay is $d = 3$ ms, and input temporal precision $\tau_{in} = 2.12$ ms. (B2) Histograms of the distribution of all weights at

the end of the anti-Hebbian iSTDP simulation shown in (B1). (C) Mean inhibitory (top) and excitatory (bottom) weights from the correlated input group after learning. Each pixel corresponds to the average of 10 simulations with inhibitory delay d (x-axis) and input temporal resolution τ_{in} (y-axis). (D) Mean inhibitory weight of the correlated inputs after learning with different iSTDP windows: Anti-Hebbian (magenta), Hebbian (red), symmetric (black, gray) for three values of τ_{in} . Each plot corresponds to a line in (C). (E) Final excitatory and inhibitory weights of the correlated pathway for anti-Hebbian (left) and symmetric (right) iSTDP, shown for all τ_{in} and $d = 6$ ms. One dot represents the mean excitatory and mean inhibitory weight at the end of one trial. Delays are pooled within the same color.

in inhibitory weight is determined by d , the timing of inhibitory spikes to the excitatory spikes (and therefore to the output spikes), together with the shape of the iSTDP window. Note that we set α such that weights from background inputs remain weak.

The potentiation of excitatory and inhibitory weights with both anti-Hebbian and symmetric iSTDP exhibits a balance between correlated excitation and inhibition, as illustrated in **Figure 2E**. Stronger excitatory weights induced by eSTDP are

counterbalanced by stronger inhibition due to iSTDP. This phenomenon depends on the input correlation precision τ_{in} (smaller values in darker color), but not significantly on the delay d . The matching is not linear and depends on the learning window.

In summary, we find that simulations with the LIF neuron confirm the theoretical results with the Poisson neuron. Detailed balance in the weights from the correlated pathway can arise if the iSTDP window is anti-Hebbian or symmetric, but not if it is Hebbian.

2.3. SHARPENING THE NEURONAL RESPONSE IN THE SFC

While detailed balance between excitatory and inhibitory weights can arise through anti-Hebbian or symmetric iSTDP (Figure 2), anti-Hebbian iSTDP may increase inhibition to the point where it will dominate over excitation for the postsynaptic neuron, as shown in Figure 3A. This follows partly because we use additive iSTDP, which strongly potentiates inhibitory weights (with our choice of parameters). Whether inhibition dominates (Rudolph et al., 2007) or not, the detailed balance weight specialization underlies the tuning of the SFC function in propagating spike volleys.

As τ_{in} governs the temporal width of input spike volleys, we evaluate τ_{out} for the postsynaptic response (Figure 3B). To do so, we detect volleys whose coincident spikes exceed a threshold as “events.” We then build a peristimulus time histogram (PSTH) of the postsynaptic spikes with respect to the input events (Figure 3B: right). An example for the simulation in Figure 3B with anti-Hebbian iSTDP is shown in Figures 3C1,C2 for two values of d .

The PSTH obtained from the simulations with specific inhibition (leading to *detailed balance*; red curve) is compared to two control conditions:

- the inhibitory weights are swapped to obtain unspecific inhibition (leading to *global balance*; green curve)
- the inhibitory inputs are omitted (black curve).

The three conditions are characterized by different mean firing rates. For $\tau_{in} = 2.12$ in the SFC with anti-Hebbian iSTDP, $r_{out} = 2.21$ sp/s for specific inhibition, 2.23 sp/s for unspecific inhibition, and 38.94 sp/s for excitation only.

The difference in response width τ_{out} between specific inhibition and the unspecific inhibition control represents the particular contribution of detailed balance on output response sharpening. Likewise, removing the inhibitory inputs from the circuit and taking the difference with the unspecific inhibition control should reveal the response sharpening due to the general presence of inhibition (global balance). To evaluate the relative change of spiking probability induced by the input stimuli, we normalize the PSTH with respect to the mean postsynaptic firing rate (see Materials and Methods for details). This gives a signal/noise ratio (SNR) for the output spikes following an event in this detection task. As can be seen in Figure 3C (right), both the specific inhibition (red) and unspecific inhibition (green) enhance the SNR. The postsynaptic response is even sharper with the specific inhibition circuit for small delays (Figure 3C1,

detailed balance, red curve). This occurs when inhibition is timed with excitation (arrow). For larger d as in Figure 3C2, 8 ms, this sharpening vanishes, as inhibition cannot arrive sufficiently early right after excitation. In that case, the performance is closer to that with unspecific inhibition (global balance, green curves). This sharpening is efficient for all $\tau_{in} = 0.71$ –5.66 ms, in the range of the delay d , as illustrated in Figure 3D. Note that very small delays d prevent a proper weight structure from developing with anti-Hebbian iSTDP, thus the sharpening of the response fails (Figure 3D: top: red curve). The principle can be explained by the presence of inhibition lowering the output firing rate, which increases the SNR of the neuron’s response (Figure 3E: top). Additionally, precisely timed inhibition coming right after excitation further sharpens the response and improves the SNR (Figure 3E: bottom). Figure 3F summarizes the performance of the sharpening by the emerged detailed weight balance, as compared to the global balance with unstructured inhibition or in the absence of inhibition.

In our model, symmetric iSTDP performed similarly to anti-Hebbian iSTDP (Figure 3G). A small difference is that global inhibition contributed slightly more to the sharpening of the response than the inhibition in detailed balance. The performance is a bit better for anti-Hebbian iSTDP because the weights grow stronger. Lastly, because the weight structure does not develop with Hebbian iSTDP, no significant difference in τ_{out} is observed (Figure 3H: left). Actually, the weights from the random input group are not weakened to zero in the Hebbian case (Supplementary Figure 1), so the unspecific inhibition control condition, in which weights between the two pathways are swapped, leads to a slightly better performance (Figure 3H: right).

We conclude that detailed balance, as achieved by anti-Hebbian and symmetric iSTDP, in combination with Hebbian eSTDP, can lead to the temporal restriction of a postsynaptic response to correlated input spikes. Brief delays in inhibition prove most beneficial for this sharpening, though the exact optimal delay is dependent on the input correlation precision τ_{in} .

2.4. RECRUITMENT OF DISYNAPTIC INHIBITORY PATHWAY WITH DELAY SELECTION

Finally, we explicitly model inhibitory interneurons in our circuit in order to examine how they are recruited in a more realistic architecture. In our FFC model in Figure 4A, two correlated input pathways (dark and light blue) compete against each other. The inhibitory inputs contain heterogeneous axonal delays. Here we focus on anti-Hebbian iSTDP, which proved efficient in developing feedforward inhibition in the previous sections. All excitatory synapses are subject to eSTDP as in Figure 2. In the example simulation in Figure 4B, the excitatory weights onto the output neuron (top) specialize to the dark blue group (“winning group”) at the expense of the light blue group (“losing group”). Note that in general, each group has 50% chance of winning because we use sufficiently competitive eSTDP (Gilson and Fukai, 2011). The inputs onto the interneurons specialize in a similar fashion, as shown for two different examples in Figure 4B (middle and bottom).

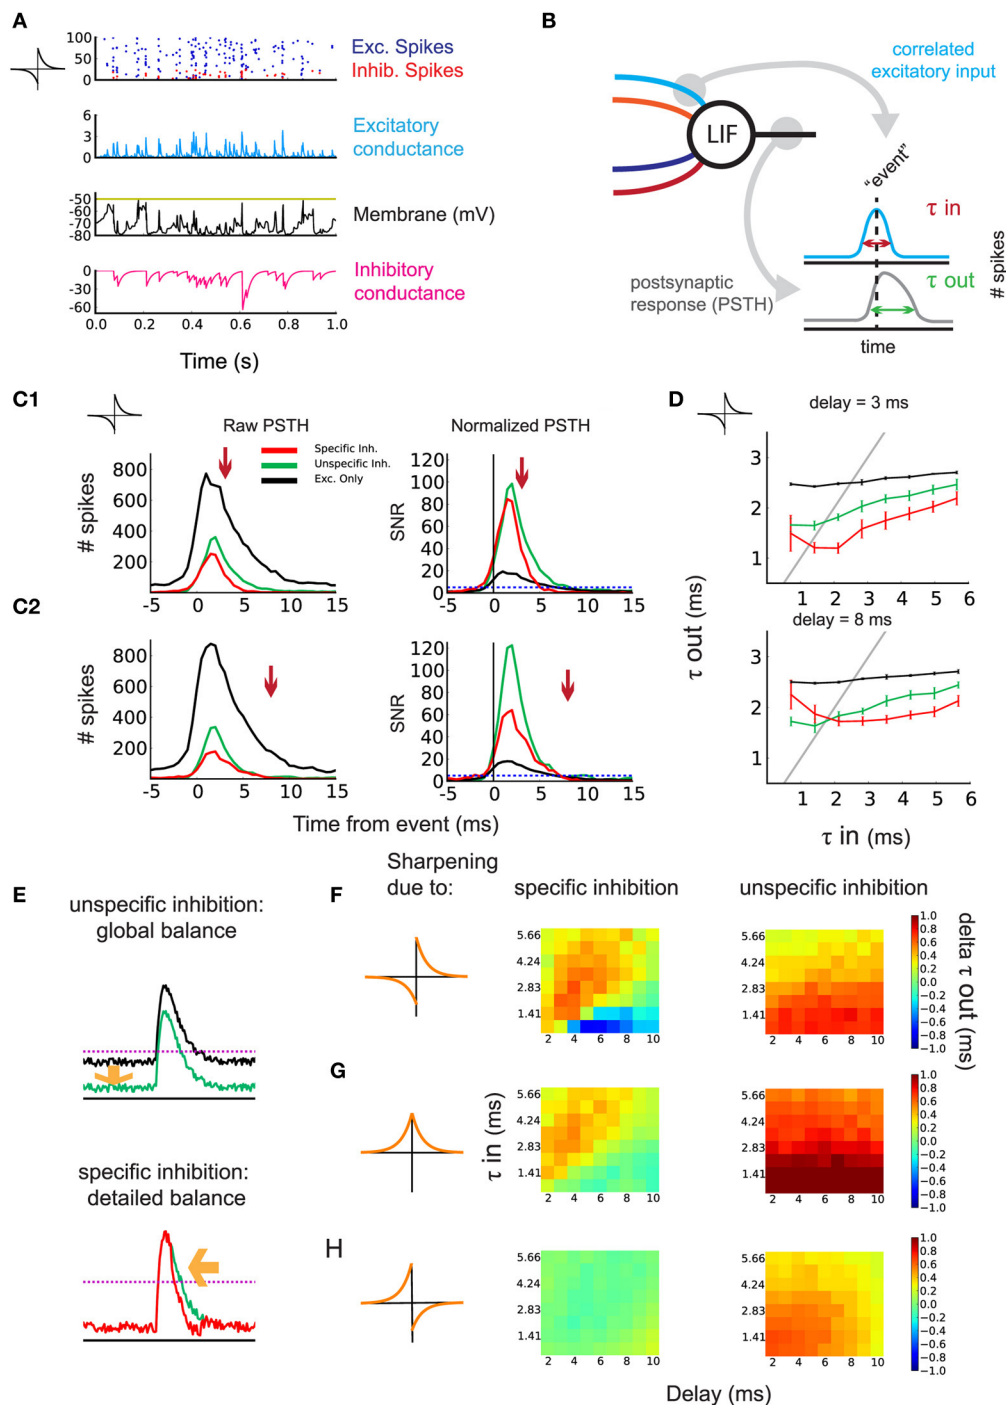


FIGURE 3 | Sharpening of the postsynaptic response by timed inhibition. **(A)** Neuronal state after learning. One second activity after training for the simulation in **Figure 2B**: raster plot of the input excitatory/inhibitory spikes (blue/red), excitatory/inhibitory conductance (blue/magenta), and voltage (black). **(B)** Schematic indicating the construction of the PSTH. Events are detected using the correlated excitatory inputs (blue). Then, postsynaptic spikes that occur in a given window around the event are counted (gray). Response efficiency is evaluated by the temporal width of the PSTH τ_{out} . **(C1)** Effect of inhibition on the response of the postsynaptic neuron to correlated events for $\tau_{in} = 2.12$ ms and $d = 3$ ms. Left: example of raw PSTH for postsynaptic spike count. Comparison of detailed balance (red) with the control of global balance (green) and no inhibition (black). The arrow indicates incoming specific inhibition. Right: signal/noise ratio (SNR) obtained by normalizing the PSTHs. **(C2)** Same as in **(C1)** but with $d = 8$ ms. **(D)** Response sharpening for different values of τ_{in} and τ_{out} . The gray unit line represents instances where the output width and the input width are equal. Top: $d = 3$ ms. Bottom: $d = 8$ ms. Legend corresponds to **(C)**. **(E)** schematic indicating the effect of detailed balance and global balance on the response shape. **(F)** anti-Hebbian iSTDP learning window and the contribution of detailed and global balance to the sharpening of the response. The difference in τ_{out} is shown for varying d (x-axis) and the input τ_{in} (y-axis). Left: difference in τ_{out} between detailed balance and global balance. Warm colors indicate the response is sharper through detailed balance compared to global balance. Right: difference in τ_{out} between global balance and no inhibition. **(G)** Same as in **(D)** but for symmetric iSTDP. **(H)** Same as in **(D)** but for Hebbian iSTDP, where no detailed balance emerged.

signal/noise ratio (SNR) obtained by normalizing the PSTHs. **(C2)** Same as in **(C1)** but with $d = 8$ ms. **(D)** Response sharpening for different values of τ_{in} and τ_{out} . The gray unit line represents instances where the output width and the input width are equal. Top: $d = 3$ ms. Bottom: $d = 8$ ms. Legend corresponds to **(C)**. **(E)** schematic indicating the effect of detailed balance and global balance on the response shape. **(F)** anti-Hebbian iSTDP learning window and the contribution of detailed and global balance to the sharpening of the response. The difference in τ_{out} is shown for varying d (x-axis) and the input τ_{in} (y-axis). Left: difference in τ_{out} between detailed balance and global balance. Warm colors indicate the response is sharper through detailed balance compared to global balance. Right: difference in τ_{out} between global balance and no inhibition. **(G)** Same as in **(D)** but for symmetric iSTDP. **(H)** Same as in **(D)** but for Hebbian iSTDP, where no detailed balance emerged.

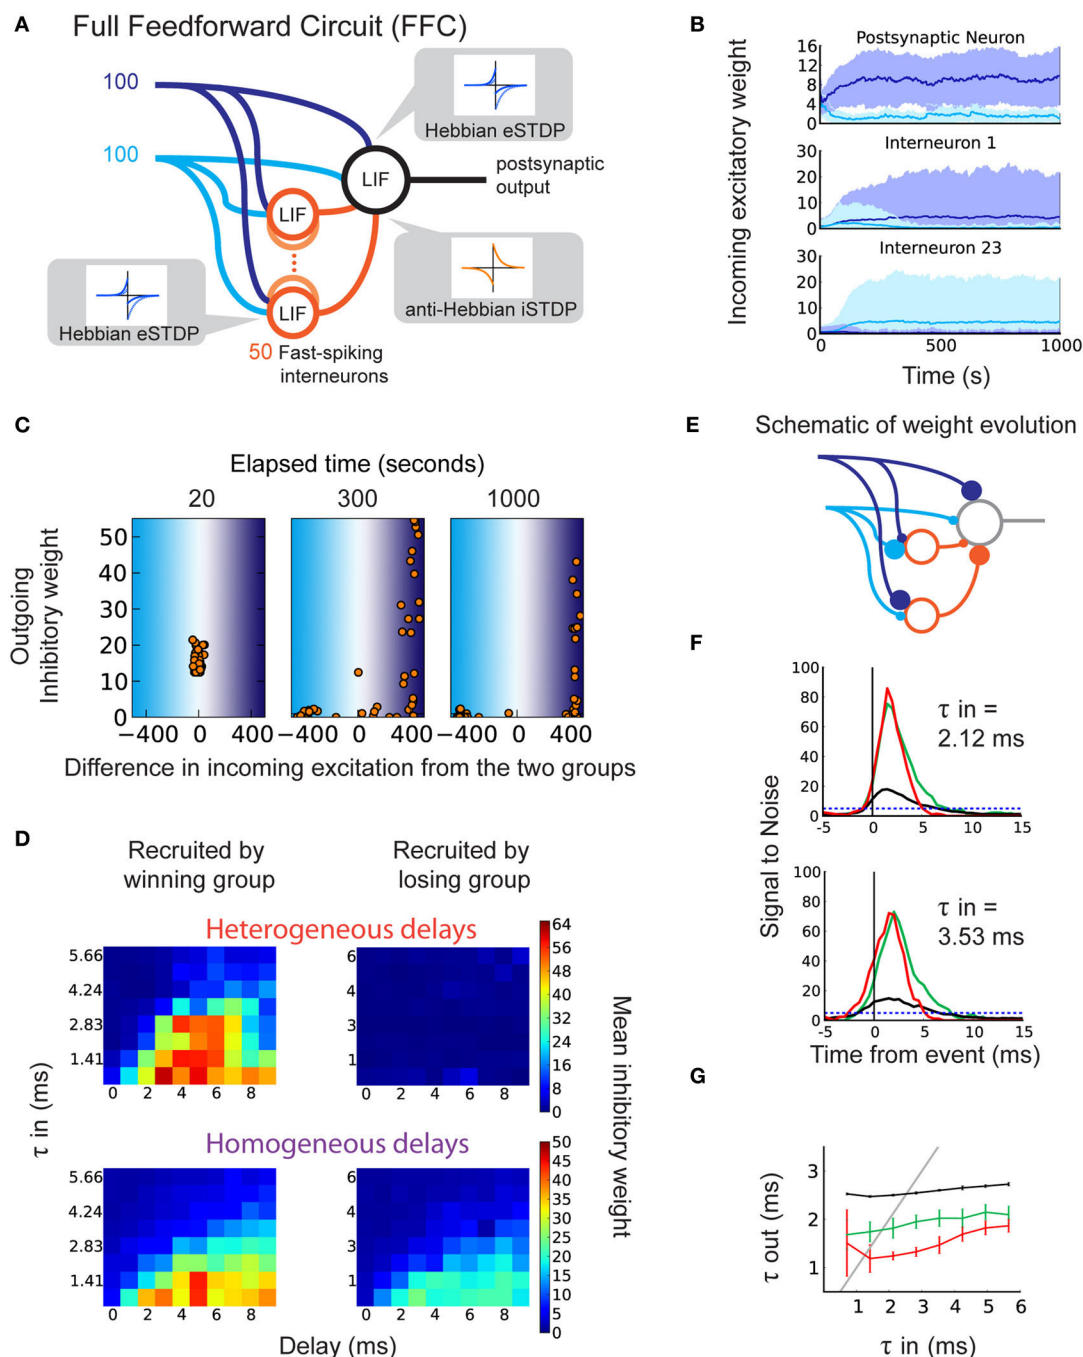


FIGURE 4 | Selection of delays in a disinaptic pathway by iSTDP. (A)

Schematic representation of the full feedforward circuit model (FFC). The postsynaptic neuron receives excitatory input from two correlated groups that compete. Inhibition onto this neuron is provided via 50 fast-spiking interneurons (orange circles). Each interneuron receives inputs from both excitatory groups. The interneurons have axonal delays between 0 and 9 ms. All synapses are plastic, with learning windows shown by the insets. **(B)** The evolution of excitatory weights from the two input groups onto the output neuron (top) and onto two of the 50 interneurons (middle, bottom) in one example trial. Weights from group 1 (dark blue inputs) increase beyond those from group 2 (light blue inputs). Here, group 1 is the “winning group.” Interneuron 1 (23) receives more input from the dark (light) blue group. **(C)** Example of weight evolution onto interneurons and the subsequent change in inhibitory synaptic weights during the simulation (after 20, 300, 1000 s). Each

dot represents one of the 50 interneurons. The x-axis indicates the difference in total weights between the two input groups onto the interneuron. The right (left) part corresponds to interneurons specializing to the dark (light) blue input group. The y-axis indicates the weight of the inhibitory synapse onto the postsynaptic neuron. **(D)** Inhibitory weights after learning depend on the axonal delays of interneurons (x-axis) and specialization of their input weights, in the FFC with heterogeneous delays (top; each horizontal line represents an average over 10 simulations), and in the FFC with homogeneous delays (bottom; each square represents 10 simulations). **(E)** Schematic of the recruitment of interneurons and the consequence on the inhibitory weights, leading to detailed balance. **(F)** SNR of the response to correlated events in the FFC with heterogeneous delays. Top: for $\tau_{in} = 2.12$ ms. Bottom: for $\tau_{in} = 3.54$ ms. **(G)** Relationship between τ_{in} and τ_{out} for the FFC model: specific inhibition (red), unspecific inhibition (green), and without inhibitory interneurons (black).

After the specialization of excitatory synapses, inhibitory synapses start to become potentiated. We find that the structure in the inhibitory synapses develops only for interneurons that specialize to the same group as the output neuron (**Figure 4C**: right part of last panel; **Figure 4D**, top). This is a consequence of the correlation between the spike trains fired by the interneurons that specialize to the winning group, and the output neuron spike train. Conversely, interneurons that specialize to the losing group do not match their spike times to postsynaptic spikes, and their weights remain weak (**Figure 4C**: left part of last panel). Inhibitory and excitatory inputs onto the output neuron become correlated, making detailed balance possible (**Figure 4C**, last panel; **Figures 4D,E**). The use of homogeneous delays in the FFC still achieves detailed balance (**Figure 4D**, bottom) though the difference in inhibitory weight is smaller. This is because there is no competition between winner-recruited neurons of different delays (**Figure 4D** bottom, left), and loser-recruited neurons can more easily adjust their firing times to postsynaptic firing if they receive a small amount of input from the winning group (**Figure 4D** bottom, right).

Importantly, LTP in the inhibitory weights depends on the axonal delay of their interneurons, in a similar manner as the SFC (**Figure 4D**: left; **Figure 2C**: top). For broader input spike volleys with larger τ_{in} , short delays are *not* selected by anti-Hebbian iSTDP. This ensures that inhibition will not cut off the output response before sufficiently many inputs are integrated. Similarly, late-arriving inhibition does not affect the sharpening of the response, therefore there is no need for its weight to be increased for this function. The adequately timed inhibition that follows excitation results in a sharper response to correlated events in the FFC (**Figures 4F,G**). The comparison with unspecific inhibition (global balance) for which inhibitory weights are swapped with interneurons specialized to the losing group confirms that precise timing between excitation and inhibition is important for the response sharpening (red curves versus green curves in **Figures 4F,G**). As in the SFC, the response to τ_{in} in the range of 1–5 ms benefits most from the detailed balance (**Figure 4G**).

To test the robustness of the FFC against noise, we modified the FFC by adding random uncorrelated inputs into the interneurons and the output neuron, and decreased the number of inputs from the correlated pathways (Noisy Full Feedforward Circuit, Noisy FFC; **Supplementary Figure 2A**). Detailed balance emerged as in the FFC, inhibitory synapses showed delay-dependent potentiation (**Supplementary Figure 2B**), and the response from the output neuron was sharpened (**Supplementary Figures 2C,D**, red curve). Detailed balance and the sharpening role of inhibition are therefore robust against noise.

We conclude that in the more realistic FFC and Noisy FFC, eSTDP determines the specialization of both the output neurons and the interneurons, and anti-Hebbian iSTDP selects the interneurons with intermediate delays, which leads to sharpening of the response.

3. DISCUSSION

This study showed how eSTDP and iSTDP can jointly structure synapses in feedforward neural circuits to control downstream firing. We found that the temporally anti-Hebbian (post-pre LTP)

component of iSTDP is crucial to achieve a balance between excitatory and inhibitory weights given correlated inputs, and assuming an inhibitory delay in the order of a few milliseconds. Moreover, interneurons can be recruited by Hebbian eSTDP in a self-organized fashion to develop inhibition through iSTDP onto output neurons. By selecting adequate delays in this disinaptic inhibition scheme, iSTDP sharpens the output firing response, enhancing the propagation of spike volleys.

3.1. INPUT TIMING AND TYPES OF iSTDP LEARNING WINDOW

We investigated how the interplay between eSTDP and iSTDP shapes the excitatory and inhibitory weight distributions. In our model, correlations in inhibition follow correlations in excitation by a delay of up to 10 ms (**Figures 1, 2, and 4**), which agrees with experimental observations at the order of a few milliseconds in the auditory (Wehr and Zador, 2003) and somatosensory cortices (Gabernet et al., 2005). For such input signals, we found that both anti-Hebbian and symmetric iSTDP windows generate a detailed balance between excitatory and inhibitory weights (see SFC in **Figure 2**). In contrast, Hebbian iSTDP leads to the weakening of all synapses: Due to the inhibitory delay and the timescale of the input correlations, a large portion of the inhibitory spikes fall into the LTD part of the window.

There is, to our knowledge, currently no experimental evidence of this kind of anti-Hebbian iSTDP. Some studies show evidence of anti-Hebbian STDP in excitatory synapses in the electric fish (Han et al., 2000; Harvey-Girard et al., 2010), in the dorsal cochlear nucleus (Tzounopoulos et al., 2004), and in corticostriatal synapses (Fino et al., 2009). Anti-Hebbian STDP has also been the subject of theoretical studies (e.g., Roberts and Bell, 2000; Rumsey and Abbott, 2004, 2006; Carnell, 2009), but again only in the context of excitatory synapses. Our study is the first one to show a functional role for the anti-Hebbian LTP in iSTDP. Anti-Hebbian LTP is also part of the symmetric iSTDP learning rule that is the subject of a recent theoretical iSTDP study by Vogels et al. (2011), showing the exact balancing of excitation and inhibition. In our model, the output neuron is dominated by strong inhibition after learning, meaning that the balance between excitatory and inhibitory weights leads to a different firing regime than in their results (Vogels et al., 2011). This follows because of our choice of inputs, which induces strong LTP via iSTDP.

Vogels et al. (2011) showed that symmetric iSTDP can lead inhibitory feedforward connections to detailed balance with fixed excitation, by letting inhibition adapt to the firing rate of each input pathway. We propose that iSTDP can ensure pathway-specific balance between excitation and inhibition, even if firing rates are constant and excitation is growing simultaneously with eSTDP. Since symmetric iSTDP contains an anti-Hebbian element (namely, post-pre LTP), detailed balance will follow as long as there is a positive delay in the inhibitory input (e.g., Pouille and Scanziani, 2001; Wehr and Zador, 2003). Our theoretical results show that the expected increase in weights does not depend on input firing rate. If, however, firing rates are unequal between input groups, we still expect our current results to hold, as long as spike pairing-based effects dominate those coming from the rate differences.

Our findings are in contrast with Hebbian iSTDP, which has been found experimentally in the entorhinal cortex (Haas et al., 2006) and in the ventral tegmental area (Kodangattil et al., 2013). If sufficient inhibition from other sources is present, synapses corresponding to uncorrelated inputs may be potentiated by Hebbian iSTDP, leading to a “reversed detailed balance”; a scenario in which inhibitory inputs from all but one pathway make up for the excitatory input from the remaining pathway. Although Hebbian iSTDP does not directly support detailed balance in the weights in our model, Hebbian iSTDP may subserve alternative functions in neural circuit processes. Recent theoretical work has shown that Hebbian iSTDP leads to decorrelation of inhibition with respect to excitation, which results in global balance and increased sensitivity to excitatory correlations (Luz and Shamir, 2012). This follows because of the increased sensitivity to input fluctuations when the neuron acts as a coincidence detector, in contrast to the integrator regime (Hong et al., 2012). Another study showed that Hebbian iSTDP also decorrelates spike patterns through lateral connections (Savin et al., 2010). Though these studies indicate that Hebbian iSTDP plays a part in creating global balance, it does not lead to the detailed balance in our feedforward circuit. Alternatively, detailed balance by Hebbian iSTDP may arise if inhibitory delays are negative, for instance when somatic inhibitory inputs precede the excitatory dendritic spike. Inhibitory weight increase will, however, be strongly bounded by the fact that an early inhibitory spike may prevent a postsynaptic spike otherwise caused by late excitation, preventing weight increase.

Another form of inhibitory plasticity, slightly different from iSTDP considered here, is voltage-dependent iLTP (Maffei et al., 2006), which leads to a potentiation in inhibitory synapses when a presynaptic spike precedes a postsynaptic depolarization either without spikes (Maffei et al., 2006), or accompanied by low-frequency spiking (Wang and Maffei, 2014). Modeling approaches have shown that when iLTP is complemented by a homeostatic form of LTD, it is capable of creating sparseness in activation that supports stimulus-pair specificity in recipient neurons (Bourjaily and Miller, 2011a,b). iLTP contains a competitive effect for inhibitory synapses, meaning that the weakest synapses will not manage to decrease post-synaptic firing, therefore missing out on LTP. If the postsynaptic spiking rate is low, as in our study, we expect the inhibitory weight evolutions with iLTP to behave similarly to Hebbian iSTDP without the LTD part. This would not lead to detailed balance, because of the brief delay in inhibition, but global balance might ensue when implemented in a large network.

In view of the large diversity of inhibitory interneurons (Markram et al., 2005), explaining the possible roles of iSTDP in different circuits and interneurons is an important open question that requires further work.

3.2. RECRUITMENT OF DISYNAPTIC INHIBITORY PATHWAY IN FEEDFORWARD NETWORK

In our Full Feedforward Circuit model (FFC model), the excitation–inhibition structure in synaptic weights arises from the recruitment of interneurons: specialization due to eSTDP, followed by the strengthening of inhibition onto output neurons induced by iSTDP. Hebbian eSTDP provides a sufficient degree

of temporal correlation between the selected excitatory and inhibitory pathways onto the output neuron. This correlation is essential for anti-Hebbian iSTDP to select weights from adequate interneurons, whose firing is correlated with the output neuron.

One could also imagine other combinations of eSTDP–iSTDP for the interneurons in the FFC model. For example, if the eSTDP onto the interneurons is anti-Hebbian, excitation and inhibition onto the output neuron become *anti*-correlated. We expect that Hebbian iSTDP for the inhibitory synapses from the interneurons would be an interesting choice in this case, to further reinforce the anticorrelation between excitation and inhibition onto the output neuron.

3.3. CONTROL OF CORRELATED FIRING ACTIVITY

In the feedforward circuit, iSTDP enables the neuron to select inhibition with an adequate delay (Figure 4D), which temporally controls the propagation of the volley of correlated spikes without arriving too early to stop it entirely (Figure 3D: top; Figures 3E,G). Moreover, the selected suitable delays depend on the input temporal precision (τ_{in}): for temporally broader spike volleys, larger delays are recruited (Figure 4D). In this sense, the output firing is sharpened only after sufficiently many inputs have been integrated, in agreement with experimental findings (Gabernet et al., 2005).

It is worth noting that delayed inhibition compared to excitation arises naturally because of the disynaptic pathway (axonal delays of interneurons). For the inputs, we considered sharp correlations at the scale of a few milliseconds, in line with the timescale of input correlations for which neurons in a balanced state are most sensitive, shown both experimentally *in vitro* and theoretically by Rossant et al. (2011). Propagation of spike volleys in networks also requires such fine temporal resolution (Diesmann et al., 1999). Our results suggest that interneurons can control the temporal spread of such spike volleys by adapting an inhibitory cutoff. This function of iSTDP is complementary to the homeostatic stabilization (Pouille et al., 2009) enforced by iSTDP to control the average firing of neurons, as was demonstrated recently (Vogels et al., 2011). In addition to restraining the firing rate, iSTDP can control the temporal output by creating a detailed balance in the synaptic weights, in which precisely timed inhibition limits the output spikes to a narrow temporal window. Thus, our finding is in accordance with previous studies that show that inhibition limits the time for summation and integration of EPSCs (Pouille and Scanziani, 2001; Gabernet et al., 2005). The presence of inhibition improves frequency tuning in excitatory neurons in auditory cortex (Wu et al., 2008). We showed that sharpening of the response only takes place if the inhibitory delay is sufficiently brief. Such short delays can limit the range of intensity tuning in auditory neurons by reducing the EPSP amplitude, controlling the response integration window (Wu et al., 2006). Quick but delayed inhibition after excitation therefore allows only inputs from high intensities to generate spikes in the downstream neuron. The millisecond-range sharpening of the response by inhibition, such as in our model, may therefore be useful for tuning control of a neuron.

For certain delays and τ_{in} , the well-timed inhibition may hyperpolarize the neuron so strongly that the responses exhibit a rebound, after inhibition vanishes (Supplementary Figure 3).

This is because the strong hyperpolarization brings the membrane potential far from the excitatory reversal potential, temporarily boosting subsequent excitatory inputs. We only found rebound responses for anti-Hebbian-based iSTDP, for which inhibitory weights grew strongest. Mechanisms to regulate inhibitory strength within a medium range could prevent this phenomenon, such as an *ad hoc* upper bound on inhibitory weights or weight-dependent iSTDP.

Finally, we also showed that response sharpening was robust to noise in the circuit, even when the correlated inputs were decreased, meaning that our results can be extended to more realistic circuit contexts with larger input numbers. iSTDP and its resulting structure in weights may therefore be useful for the propagation of transient activities in larger circuits, such as a cortical column.

4. MATERIALS AND METHODS

Here we provide details about our analysis to predict the weight changes induced by simultaneously occurring eSTDP and iSTDP in the first section of Results. Then we describe the two neural circuit architectures used in this study, namely the SFC and FFC. Finally, we explain how the PSTHs of the postsynaptic and output neurons are calculated.

4.1. THEORETICAL ANALYSIS OF THE WEIGHT EVOLUTION IN THE SIMPLIFIED FEEDFORWARD CIRCUIT (SFC)

In our theoretical model, a postsynaptic Poisson neuron post receives both excitatory and inhibitory inputs (Figure 1A). All inputs share the same source of correlation, and inhibition is delayed by d compared to excitation.

The firing rate ρ_{post} evolves over time according to the presynaptic inputs:

$$\rho_{\text{post}}(t) = \sum_k w_k^e [\epsilon_e * S_k^e](t) - \sum_m w_m^i [\epsilon_i * S_m^i](t). \quad (4)$$

The k th excitatory input spike train S_k^e is modeled as a time series of Dirac functions: $S_k^e(t) = \sum_s \delta(t - t_s^k)$; likewise, S_m^i is the m th inhibitory spike train. Though ρ_{post} may take on negative values in theory, we assume it is positive on average, and do not consider the case of no postsynaptic spiking. The EPSPs and IPSPs are summed together to obtain ρ_{post} ; $*$ denotes the convolution of functions. For each EPSP at synapse k , the time course of the postsynaptic response for a single spike is described by the normalized kernel functions ϵ_e rescaled by the weight w_k^e . For IPSP at synapse m , the same holds with ϵ_i and w_m^i . In Figure 1, we use a simple exponential decay that is identical for all excitatory synapses with decay time $\tau_e = 3$ ms; likewise $\tau_i = 5$ ms for all inhibitory synapses.

In order to evaluate the expected weight change, we calculate the pre-post spike-time correlations for excitatory and inhibitory inputs. We consider the situation when pre-post correlations are dominated by the effect of input correlations. Actually, we use spike-time covariances defined as (Gilson et al., 2010)

$$C_{k,l}^{e,e}(t, \Delta t) = \text{Cov}[S_k^e, S_l^e](t, \Delta t) := \langle S_l^e(t) S_k^e(t + \Delta t) \rangle - \langle S_l^e(t) \rangle \langle S_k^e(t + \Delta t) \rangle. \quad (5)$$

The angular brackets $\langle \dots \rangle$ denote the ensemble average over the randomness from the stochastic process. Considering spike trains with constant average firing rates and fixed pair-wise correlations, we can omit the dependence on t in Equation (5). For the configuration described in Figure 1A, excitatory inputs are homogeneously correlated between them, as well as inhibitory inputs. However, the correlation between an excitatory and an inhibitory inputs involves the delay d . Denoting by $C_0(\Delta t)$ the homogeneous covariance corresponding to Figure 1C, we have

$$\begin{aligned} C_{k,l}^{e,e}(\Delta t) &= C_0(\Delta t), \\ C_{m,n}^{i,i}(\Delta t) &= C_0(\Delta t), \\ C_{k,m}^{e,i}(\Delta t) &= C_0(\Delta t - d). \end{aligned} \quad (6)$$

All covariances are defined in a similar manner to Equation (5). For the k th excitatory input, the covariance $\text{Cov}_{k,\text{post}}^e$ is given by the input covariance on which the postsynaptic response (EPSPs-IPSPs) operates:

$$\begin{aligned} C_{k,\text{post}}^e(\Delta t) &:= \text{Cov}[S_k^e, S_{\text{post}}](t, \Delta t), S_{\text{post}} \propto \rho_{\text{post}} \\ &= \text{Cov} \left[S_k^e, \left(\sum_l w_l^e \epsilon_e * S_l^e - \sum_n w_n^i \epsilon_i * S_n^i \right) \right](\Delta t) \\ &= \sum_l w_l^e [C_{k,l}^{e,e} * \epsilon_e](\Delta t) - \sum_n w_n^i [C_{k,n}^{e,i} * \epsilon_i](\Delta t) \\ &= \sum_l w_l^e [C_0 * \epsilon_e](\Delta t) - \sum_n w_n^i [C_0 * \epsilon_i](\Delta t - d). \end{aligned} \quad (7)$$

The subsequent STDP weight update is given by the integral value of the learning window $W_e(u)$ with the pre-post covariance $C_{k,\text{post}}^e(-u)$, which yields:

$$\begin{aligned} \Delta w_k^e &= [C_{k,\text{post}}^e * W_e](0) \\ &= \sum_l w_l^e [C_0 * \epsilon_e * W_e](0) - \sum_n w_n^i [C_0 * \epsilon_i * W_e](d) \end{aligned} \quad (8)$$

Similarly, the pre-post covariance and the expected change for the m th inhibitory weight is given by:

$$\begin{aligned} C_{m,\text{post}}^i(\Delta t) &= \sum_l w_l^e [C_0 * \epsilon_e](\Delta t + d) - \sum_n w_n^i [C_0 * \epsilon_i](\Delta t), \\ \Delta w_m^i &= \sum_l w_l^e [C_0 * \epsilon_e * W_i](-d) \\ &\quad - \sum_n w_n^i [C_0 * \epsilon_i * W_i](0) \end{aligned} \quad (9)$$

These formulas are used to generate Figure 1F.

4.2. DETAILS OF THE SIMULATED SFC

In Figures 2, 3, the SFC consists of a single postsynaptic neuron that receives a total of 200 excitatory and 50 inhibitory inputs. Half of each set of inputs consist of weakly correlated spike trains, whereas the remainder consists of random Poisson spike trains (Figure 2A).

We use a function in Brian Simulator to generate correlated spike trains, which is based on the first method in Brette (2008). The principle of this function is that a doubly stochastic process (or Cox Process) with an average (spike) rate r , underlies a group of inhomogeneous Poisson processes which have rates that fluctuate around r . Final spike trains are derived from these inhomogeneous Poisson processes, and will appear to be homogeneous, correlated spike trains with stationary rate r . These correlated spike trains do not have Poisson statistics, because their autocovariance is modulated by their correlation. In order to have exponential cross-correlation functions (CCF) between these spike trains, the function employs the Ornstein–Uhlenbeck process. The time-constant of the exponential CCF is a parameter called τ_c in Brette (2008) and in Brian Simulator. We focus on the standard deviation of the latencies in input spike volleys (representing input stimuli), τ_{in} , where $\tau_{in} = \tau_c \sqrt{2}$. We apply correlation strength $c = 0.1$ and CCF standard deviation τ_{in} in the range of 0.71–5.66 ms. Correlated inhibition is delayed by d ms. All inputs have the same firing rate $r_{in} = 5$ sp/s.

The postsynaptic neuron is a conductance-based leaky integrate-and-fire (LIF) model. Its membrane potential V obeys:

$$\tau_m \frac{dV}{dt} = E_{leak} - V + g_e(E_e - V) + g_i(E_i - V) \quad (10)$$

With synaptic conductances g_e and g_i , that decay exponentially with conductance trace parameters τ_e and τ_i :

$$\tau_e \frac{dg_e}{dt} = -g_e, \quad \tau_i \frac{dg_i}{dt} = -g_i \quad (11)$$

For every excitatory spike from synapse k , g_e is increased by w_k^e , and for every inhibitory spike from synapse m , g_i by w_m^i . Intrinsic time constants of the neuron are not considered. All simulations are run with BRIAN, a python-based neural simulator (Goodman and Brette, 2008). The simulations last 2500 s each. For plots with error bars and color maps, 10 trials are repeated for the same simulation protocol with each set of values for d and τ_{in} . All SFC variables and parameters are listed in Table 3.

4.3. DETAILS OF THE SIMULATED FFC

In Figure 4, the FFC model incorporates 50 inhibitory interneurons which receive the same excitatory inputs as the output neuron, and project inhibitory connections onto the latter. In contrast to the SFC, two groups of correlated inputs compete against each other (dark blue and light blue lines in Figure 4A). The postsynaptic neuron receives 100 inputs from each group, and each interneuron receives 10 excitatory synapses from each group. The inputs are chosen so that the first interneuron receives excitatory input from spike trains 1–10 from the dark blue group, the second interneuron receives input from spike trains 2–11 from the dark blue group, and so on. The same procedure is performed for inputs from the light blue group. The 50 interneurons only differ from the output neuron by a shorter membrane time constant $\tau_m^i = 5$ ms. The interneurons are not connected to one another and there is no external inhibition source. Each interneuron makes a single inhibitory synapse onto the output neuron,

Table 3 | SFC and FFC variables and parameters.

Theoretical SFC variables	Description	Value
ρ_{post}	Postsynaptic neuron firing rate	
S_k^e, S_l^e	Presynaptic excitatory spiketrain k or l	
S_m^i, S_n^i	Presynaptic inhibitory spiketrain m or n	
ε_e	Synaptic conductance decay function for excitation	
ε_i	Synaptic conductance decay function for inhibition	
w_k^e, w_l^e	Synaptic weight of excitatory input k or l	
w_m^i, w_n^i	Synaptic weight of inhibitory input m or n	
u	Time between two spikes	
τ_e	Excitatory conductance decay constant	3 ms
τ_i	Inhibitory conductance decay constant	5 ms
d	Inhibitory delay	2–20 ms
τ_c	Time constant for input spike correlogram	0.5–4.0 ms
τ_{in}	STD of the latency of the input spike correlogram	$\tau_c \times \sqrt{2}$
c	Correlation	0.1
NUMERICAL EXTRA VARS: SFC and FFC		
V	Postsynaptic neuron membrane voltage	
g_e	Excitatory synaptic conductance	
g_i	Inhibitory synaptic conductance	
NUMERICAL EXTRA PARAMS: SFC and FFC		
r_{in}	Input firing rate	5 Hz
τ_m	Membrane time constant	20 ms
τ_i	Inhibitory conductance decay constant	20 ms
E_{leak}	Leak Potential	–70 mV
E_e	Excitatory reversal potential	0 mV
E_i	Inhibitory reversal potential	–80 mV
–	LIF spike threshold	–50 mV

Table listing all SFC and FFC variables and parameters relating to inputs and neuronal properties.

and axonal delays d are heterogeneous, ranging from 0 to 9 ms (five interneurons for each d).

All synapses are plastic. The excitatory synapses onto both the output neuron and the inhibitory interneurons are subject to the same Hebbian eSTDP learning window. The list of parameters that vary from the SFC is shown in Table 2. The iSTDP window time constants are lower ($\tau_{pre}^i = \tau_{post}^i = 20$ ms), excitatory and inhibitory learning are slowed down ($\eta_e = 0.0624$, $\eta_i = 0.02$), the eSTDP equilibrium value is higher ($w_0 = 0.08$), and the value of the start-up weights is changed (a random number between 0 and 1 for excitation, 1 for inhibition). Total simulation time is 1000 s.

To test the robustness against noise, we modify the FFC by adding 400 excitatory random inputs onto the interneurons and output neuron (Noisy FFC in Supplementary Figure 2A, green inputs) and decrease the size of the correlated groups to 50 (dark and light blue inputs). The number of interneurons is increased to 120. Other parameters are unchanged (see Table 3 for all FFC parameters).

4.4. ANALYSIS OF THE TEMPORAL ACUITY OF THE POSTSYNAPTIC RESPONSE

We evaluate how the iSTDP learning rule, via the resulting weight distributions, shapes the postsynaptic response to correlated input activity. To do so we run the SFC simulation with fixed weights for 300 s.

Volleys of input spikes (“events”) are detected by binning the spike times of the 100 correlated inputs in the SFC in bins with width 0.5 ms, and counting the spikes in a sliding window of duration τ_{in} . When the spike count exceeds a threshold, the time of the event is set to the center of the sliding window. Events in neighboring windows are discarded. The window spike count threshold is determined for each τ_{in} such that the average number of events per second is as close as possible to r_{in} without exceeding it.

To evaluate the temporal acuity of the spikes fired in response to such events, we count the postsynaptic spikes in bins of 0.5 ms. This yields a peri-stimulus time histogram (PSTHs) around the time of the events, as shown in **Figure 3B**. We then evaluate the temporal acuity of the response of the output neuron to input stimuli by computing the sharpness of the PSTH.

Not only the latencies of spikes following the event, but also the excess of spikes compared to the baseline output firing rate contributes to the temporal acuity of the response. We obtain the average firing rate during the entire 300 s simulation, F_0 . The number of spikes in each bin of the PSTH is then divided by F_0 , yielding the “normalized” PSTH as a deviation from average activity. This deviation is also the signal to noise ratio (SNR: **Figures 3C1,C2**). The temporal acuity of the response input events is then evaluated through the standard deviation of the normalized PSTH, τ_{out} . τ_{out} is computed over the time window 0, +10 ms inside the PSTH (0 is the time of the event).

To study how the emerged inhibitory weight structure affects τ_{out} , we compare the outcome of simulations to two controls:

- “specific inhibition”: the weights are as in the numerical simulation of the SFC;
- “unspecific inhibition control”: the inhibitory weights are swapped between the two input groups, leading to equivalent total inhibition, but abolishing the relation between the temporal structure of the spike trains and the strength of the weights that depended on them through iSTDP;
- “excitation only control”: the inhibitory inputs are removed completely, leaving only the excitatory inputs.

In the unspecific inhibition control for the SFC, we aim to destroy the detailed weight structure that emerges, but preserve strong feedforward inhibition. After swapping the inhibitory weights between the correlated and random inputs, the weight strengths are adjusted down to obtain a postsynaptic firing rate similar to the specific inhibition configuration. Weight corrections are not performed for trials with mean weight smaller than 1. Excitatory weights are unchanged in all three conditions. In the FFC, the same three scenarios are applied. For specific inhibition, all weights are as obtained from the simulation. To obtain

the unspecific inhibition condition, excitatory weights onto the interneurons are swapped between the winner and loser input pathways. The result of this manipulation is that an interneuron receiving strong inputs from the dark blue group and weak inputs from the light blue group, changes to receiving weak inputs from the dark blue group and strong inputs from the light blue group. This procedure leads to qualitatively the same control as in the SFC. Inhibitory weights are not adjusted further as in the SFC. For the excitation only control, the interneurons are omitted.

AUTHOR CONTRIBUTIONS

In this study, Matthieu Gilson and Florence I. Kleberg designed the experiments, Florence I. Kleberg performed the experiments and analyzed the data, and Florence I. Kleberg, Matthieu Gilson, and Tomoki Fukai wrote the paper.

FUNDING

The research was funded by RIKEN.

ACKNOWLEDGMENTS

The authors would like to thank Hideaki Shimazaki for helpful discussions, and Tom Sharp, Anthony DeConstanzo and Pierre Yger for their help writing the article.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <http://www.frontiersin.org/journal/10.3389/fncom.2014.00053/abstract>

Supplementary Figure 1 | Equilibrium weights of the random inhibitory inputs in the SFC. Mean inhibitory weight for Hebbian iSTDP (red curve), anti-Hebbian iSTDP (magenta curve), symmetric iSTDP (black curve), and symmetric with equal total LTP (gray curve). Mean final weights are shown for three τ_{in} .

Supplementary Figure 2 | Robustness of weight structure development and response sharpening in the presence of noise. (A) Noisy full feedforward circuit model with eSTDP and iSTDP. There are 120 interneurons instead of 50. The correlated input groups have decreased to 50 inputs each, and an additional 400 random inputs project onto the output neuron. Each interneuron also receives 60 random inputs. For eSTDP, $w_0 = 0.037$. Other parameters are as in the FFC. **(B)** Delay-dependent inhibitory weight strengthening of interneurons recruited by the winning group (top) and absence of inhibitory weight increase for interneurons recruited by the losing group (bottom). **(C)** Effect of inhibition on the response of the postsynaptic neuron in response to correlated events for $\tau_{in} = 2.12$ ms (left) and 3.54 ms (right). Comparison of the signal/noise ratio (SNR) between specific inhibition (red), the control of unspecific inhibition (green) and excitation only (black). **(D)** τ_{in} and τ_{out} results for the Noisy FFC with specific inhibition (red), unspecific inhibition (green), and with only excitation (black) for various τ_{in} .

Supplementary Figure 3 | Rebound of the neuronal spike probability after the arrival of strong inhibition in the SFC. The rebound in spiking probability is visible for specific inhibition (red curves). In the specific inhibition case, a rebound response is observed (orange arrow). The red arrow indicates the moment inhibition kicks in, in the specific inhibition case. The rebound response is shown for $\tau_{in} = 0.71$ ms, $d = 5$, and 6 ms.

REFERENCES

- Anderson, J., Carandini, M., and Ferster, D. (2000). Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *J. Physiol.* 484, 909–926.
- Atallah, B., and Scanziani, M. (2009). Instantaneous modulation of gamma oscillation frequency by balancing excitation with inhibition. *Neuron* 62, 566–577. doi: 10.1016/j.neuron.2009.04.027
- Bell, C., Han, V., Sugawara, Y., and Grant, K. (1997). Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* 387, 170–170.
- Bi, G., and Poo, M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Böhme, G., Bon, C., Lemaire, M., Reibaud, M., Piot, O., Stutzmann, J., et al. (1993). Altered synaptic plasticity and memory formation in nitric oxide synthase inhibitor-treated rats. *Proc. Natl. Acad. Sci. U.S.A.* 90, 9191–9194. doi: 10.1073/pnas.90.19.9191
- Bourjaily, M. A., and Miller, P. (2011a). Excitatory, inhibitory, and structural plasticity produce correlated connectivity in random networks trained to solve paired-stimulus tasks. *Front. Comput. Neurosci.* 5:37. doi: 10.3389/fncom.2011.00037
- Bourjaily, M. A., and Miller, P. (2011b). Synaptic plasticity and connectivity requirements to produce stimulus-pair specific responses in recurrent networks of spiking neurons. *PLoS Comput. Biol.* 7:e1001091. doi: 10.1371/journal.pcbi.1001091
- Brette, R. (2008). Generation of correlated spike trains. *Neural Comput.* 21, 188–215. doi: 10.1162/neco.2009.12.07-657
- Buzsaki, G. (1984). Feed-forward inhibition in the hippocampal formation. *Prog. Neurobiol.* 22, 131–153. doi: 10.1016/0301-0082(84)90023-6
- Carnell, A. (2009). An analysis of the use of Hebbian and Anti-Hebbian spike time dependent plasticity learning functions within the context of recurrent spiking neural networks. *Neurocomputing* 72, 685–692. doi: 10.1016/j.neucom.2008.07.012
- Clopath, C., Büsing, L., Vasilaki, E., and Gerstner, W. (2010). Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nat. Neurosci.* 13, 344–352. doi: 10.1038/nn.2479
- Davis, K., Miller, R., and Young, E. (1996). Effects of somatosensory and parallel-fiber stimulation on neurons in dorsal cochlear nucleus. *J. Neurophysiol.* 76, 3012–3024.
- Debanne, D., Beat, H., and Thompson, S. (1998). Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *J. Physiol.* 507, 237–248. doi: 10.1111/j.1469-7793.1998.237bu.x
- Diesmann, M., Gewaltig, M., and Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402, 529–533. doi: 10.1038/990101
- Fino, E., Paille, V., Deniau, J., and Venance, L. (2009). Asymmetric spike-timing dependent plasticity of striatal nitric oxide-synthase interneurons. *Neuroscience* 160, 744–754. doi: 10.1016/j.neuroscience.2009.03.015
- Gabernet, L., Jadhav, S., Feldman, D., Carandini, M., and Scanziani, M. (2005). Somatosensory integration controlled by dynamic thalamocortical feed-forward inhibition. *Neuron* 48, 315–327. doi: 10.1016/j.neuron.2005.09.022
- Gerstner, W., Kempter, R., van Hemmen, J., and Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature* 383, 76–68. doi: 10.1038/383076a0
- Gilson, M., Burkitt, A., and van Hemmen, L. (2010). STDP in recurrent neuronal networks. *Front. Comput. Neurosci.* 4:23. doi: 10.3389/fncom.2010.00023
- Gilson, M., and Fukai, T. (2011). Stability versus neuronal specialization for STDP: long-tail weight distributions solve the dilemma. *PLoS ONE* 6:e25339. doi: 10.1371/journal.pone.0025339
- Gjorgjieva, J., Clopath, C., Audet, J., and Pfister, J. (2011). A triplet spike-timing-dependent plasticity model generalizes the Bienenstock-Cooper-Munro rule to higher-order spatiotemporal correlations. *Proc. Natl. Acad. Sci. U.S.A.* 108, 19383–19388. doi: 10.1073/pnas.1105933108
- Goodman, D., and Brette, R. (2008). Brian: a simulator for spiking neural networks in python. *Front. Neuroinform.* 2:5. doi: 10.3389/neuro.11.005.2008
- Gütig, R., Aharonov, R., Rotter, S., and Sompolinsky, H. (2003). Learning input correlations through nonlinear temporally asymmetric Hebbian plasticity. *J. Neurosci.* 23, 3697–714.
- Haas, J., Nowotny, T., Abarbanel, H., Zavala, B., and Landisman, C. (2006). Spike-timing-dependent plasticity of inhibitory synapses in the entorhinal cortex. *J. Neurophysiol.* 96, 3305–3313. doi: 10.1152/jn.00551.2006
- Haider, B., Duque, A., Hasenstaub, A., and McCormick, D. (2006). Neocortical network activity *in vivo* is generated through a dynamic balance of excitation and inhibition. *J. Neurosci.* 26, 4535–4545. doi: 10.1523/JNEUROSCI.5297-05.2006
- Haider, B., Häusser, M., and Carandini, M. (2013). Inhibition dominates sensory responses in the awake cortex. *Nature* 493, 97–100. doi: 10.1038/nature11665
- Han, V., Grant, K., and Bell, C. (2000). Reversible associative depression and nonassociative potentiation at a parallel fiber synapse. *Neuron* 27, 611–622. doi: 10.1016/S0896-6273(00)00070-2
- Harvey-Girard, E., Lewis, J., and Maler, L. (2010). Burst-induced anti-Hebbian depression acts through short-term synaptic dynamics to cancel redundant sensory signals. *J. Neurosci.* 30, 6152–6169. doi: 10.1523/JNEUROSCI.0303-10.2010
- Hensch, T., Fagioli, M., Mataga, N., Stryker, M., Baekkeskov, S., and Kash, S. (1998). Local GABA circuit control of experience-dependent plasticity in developing visual cortex. *Science* 282, 1504–1508. doi: 10.1126/science.282.5393.1504
- Hong, S., Ratté, S., Prescott, S., and De, Schutter, E. (2012). Single neuron firing properties impact correlation-based population coding. *J. Neurosci.* 32, 1413–1428. doi: 10.1523/JNEUROSCI.3735-11.2012
- Izhikevich, E., Gally, J., and Edelman, G. (2004). Spike-timing dynamics of neuronal groups. *Cereb. Cortex* 14, 933–944. doi: 10.1093/cercor/bhh053
- Jackson, A., Gee, V., Baker, S., and Lemon, R. (2003). Synchrony between neurons with similar muscle fields in monkey motor cortex. *Neuron* 38, 115–125. doi: 10.1016/S0896-6273(03)00162-4
- Kang, S., Kitano, K., and Fukai, T. (2008). Structure of spontaneous UP and DOWN transitions self-organizing in a cortical network model. *PLoS Comput. Biol.* 4:e1000022. doi: 10.1371/journal.pcbi.1000022
- Kempter, R., Gerstner, W., and van Hemmen, J. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E* 59, 4498–4514. doi: 10.1103/PhysRevE.59.4498
- Kilavik, B., Confais, J., Gruen, S., and Riehle, A. (2009). Long-term modifications in motor cortical dynamics induced by intensive practice. *J. Neurosci.* 29, 12653–12663. doi: 10.1523/JNEUROSCI.1554-09.2009
- Kodangattil, J. N., Dacher, M., Authement, M. E., and Nugent, F. S. (2013). Spike timing-dependent plasticity at GABAergic synapses in the ventral tegmental area. *J. Physiol.* 591, 4699–4710. doi: 10.1113/jphysiol.2013.257873
- Kremkow, J., Aertsen, A., and Kumar, A. (2010). Gating of signal propagation in spiking neural networks by balanced and correlated excitation and inhibition. *J. Neurosci.* 30, 15760–15768. doi: 10.1523/JNEUROSCI.3874-10.2010
- Luz, Y., and Shamir, M. (2012). Balancing feed-forward excitation and inhibition via Hebbian inhibitory synaptic plasticity. *PLoS Comput. Biol.* 8:e1002334. doi: 10.1371/journal.pcbi.1002334
- Maffei, A., Nataraj, K., Nelson, S., and Turrigiano, G. (2006). Potentiation of cortical inhibition by visual deprivation. *Nature* 443, 81–84. doi: 10.1038/nature05079
- Magee, J., and Johnston, D. (1997). A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science* 275, 209–213. doi: 10.1126/science.275.5297.209
- Maldonado, P., Babul, C., Singer, W., Rodriguez, E., Berger, D., and Grün, S. (2008). Synchronization of neuronal responses in primary visual cortex of monkeys viewing natural images. *J. Neurophysiol.* 100, 1523–1532. doi: 10.1152/jn.00076.2008
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–215. doi: 10.1126/science.275.5297.213
- Markram, H., Toledo-rodriguez, M., Wang, Y., and Gupta, A. (2005). Interneurons of the neocortical inhibitory system. *Nat. Rev. Neurosci.* 5, 793–807. doi: 10.1038/nrn1519
- Okun, M., and Lampl, I. (2008). Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat. Neurosci.* 11, 535–537. doi: 10.1038/nn.2105
- Pouille, F., Marin-Burgin, A., Adesnik, H., Atallah, B., and Scanziani, M. (2009). Input normalization by global feedforward inhibition expands cortical dynamic range. *Nat. Neurosci.* 12, 1577–1585. doi: 10.1038/nn.2441
- Pouille, F., and Scanziani, M. (2001). Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science* 293, 1159–1163. doi: 10.1126/science.1060342

- Putrino, D., Brown, E., Mastaglia, F., and Ghosh, S. (2010). Differential involvement of excitatory and inhibitory neurons of cat motor cortex in coincident spike activity related to behavioral context. *J. Neurosci.* 30, 8048–8056. doi: 10.1523/JNEUROSCI.0770-10.2010
- Riehle, A., Gruen, S., Diesmann, M., and Aertsen, A. (1997). Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278, 1950–1953. doi: 10.1126/science.278.5345.1950
- Roberts, P., and Bell, C. (2000). Computational consequences of temporally asymmetric learning rules: II. Sensory image cancellation. *J. Comput. Neurosci.* 9, 67–83. doi: 10.1023/A:1008938428112
- Rossant, C., Leijon, S., Magnusson, A., and Brette, R. (2011). Sensitivity of noisy neurons to coincident inputs. *J. Neurosci.* 31, 17193–1206. doi: 10.1523/JNEUROSCI.2482-11.2011
- Rudolph, M., Pospischil, M., Timofeev, I., and Destexhe, A. (2007). Inhibition determines membrane potential dynamics and controls action potential generation in awake and sleeping cat cortex. *J. Neurosci.* 27, 5280–5290. doi: 10.1523/JNEUROSCI.4652-06.2007
- Rumsey, C., and Abbott, L. (2004). Equalization of synaptic efficacy by activity- and timing-dependent synaptic plasticity. *J. Neurophysiol.* 91, 2273–2280. doi: 10.1152/jn.00900.2003
- Rumsey, C., and Abbott, L. (2006). Synaptic democracy in active dendrites. *J. Neurophysiol.* 96, 2307–2318. doi: 10.1152/jn.00149.2006
- Savin, C., Joshi, P., and Triesch, J. (2010). Independent component analysis in spiking neurons. *PLoS Comput. Biol.* 6:e1000757. doi: 10.1371/journal.pcbi.1000757
- Senn, W., and Buehl, N. (2003). Spike-based synaptic plasticity and the emergence of direction selective simple cells: mathematical analysis. *J. Comput. Neurosci.* 14, 119–138. doi: 10.1023/A:1021935100586
- Silberberg, G., and Markram, H. (2007). Disynaptic inhibition between neocortical pyramidal cells mediated by Martinotti cells. *Neuron* 53, 735–746. doi: 10.1016/j.neuron.2007.02.012
- Sjöström, P., Turrigiano, G., and Nelson, S. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32, 1149–1164. doi: 10.1016/S0896-6273(01)00542-6
- Song, S., and Abbott, L. (2001). Cortical development and remapping through spike timing-dependent plasticity. *Neuron* 32, 339–350. doi: 10.1016/S0896-6273(01)00451-2
- Song, S., Miller, K., and Abbott, L. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.* 3, 919–926. doi: 10.1038/78829
- Stokes, C., and Isaacson, J. (2010). From dendrite to soma: dynamic routing of inhibition by complementary interneuron microcircuits in olfactory cortex. *Neuron* 67, 452–465. doi: 10.1016/j.neuron.2010.06.029
- Tan, Z., Hu, H., Huang, Z., and Agmon, A. (2008). Robust but delayed thalamocortical activation of dendritic-targeting inhibitory interneurons. *Proc. Natl. Acad. Sci. U.S.A.* 105, 2187–2192. doi: 10.1073/pnas.0710628105
- Tzounopoulos, T., Kim, Y., Oertel, D., and Trussell, L. (2004). Cell-specific, spike timing-dependent plasticities in the dorsal cochlear nucleus. *Nat. Neurosci.* 7, 719–725. doi: 10.1038/nn1272
- van Rossum, M., Bi, G., and Turrigiano, G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neurosci.* 20, 8812–8821.
- van Vreeswijk, C., and Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274, 1724–1726. doi: 10.1126/science.274.5293.1724
- Vogels, T., and Abbott, L. (2009). Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nat. Neurosci.* 12, 483–491. doi: 10.1038/nn.2276
- Vogels, T., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* 334, 1569–1573. doi: 10.1126/science.1211095
- Wang, L., and Maffei, A. (2014). Inhibitory plasticity dictates the sign of plasticity at excitatory synapses. *J. Neurosci.* 34, 1083–1093. doi: 10.1523/JNEUROSCI.4711-13.2014
- Wehr, M., and Zador, A. (2003). Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* 426, 442–446. doi: 10.1038/nature02116
- Wilent, W., and Contreras, D. (2004). Synaptic responses to whisker deflections in rat barrel cortex as a function of cortical layer and stimulus intensity. *J. Neurosci.* 24, 3985–3998. doi: 10.1523/JNEUROSCI.5782-03.2004
- Woodin, M., Ganguly, K., and Poo, M. (2003). Coincident pre- and post-synaptic activity modifies GABAergic synapses by postsynaptic changes in Cl⁻ transporter activity. *Neuron* 39, 807–820. doi: 10.1016/S0896-6273(03)00507-5
- Wu, G., Arbuckle, R., Liu, B., Tao, H., and Zhang, L. (2008). Lateral sharpening of cortical frequency tuning by approximately balanced inhibition. *Neuron* 58, 132–143. doi: 10.1016/j.neuron.2008.01.035
- Wu, G., Li, P., Tao, H., and Zhang, L. (2006). Nonmonotonic synaptic excitation and imbalanced inhibition underlying cortical intensity tuning. *Neuron* 52, 705–715. doi: 10.1016/j.neuron.2006.10.009
- Young, J., Waleszczyk, W., Wang, C., Calford, M., Dreher, B., and Obermayer, K. (2007). Cortical reorganization consistent with spike timing-but not correlation-dependent plasticity. *Nat. Neurosci.* 10, 887–895. doi: 10.1038/nn0807-1073c
- Zhang, L., Tao, H., Holt, C., Harris, W., and Poo, M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature* 395, 37–44. doi: 10.1038/25665

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 25 December 2013; accepted: 10 April 2014; published online: 07 May 2014.
Citation: Kleberg FI, Fukai T and Gilson M (2014) Excitatory and inhibitory STDP jointly tune feedforward neural circuits to selectively propagate correlated spiking activity. *Front. Comput. Neurosci.* 8:53. doi: 10.3389/fncom.2014.00053
This article was submitted to the journal *Frontiers in Computational Neuroscience*.
Copyright © 2014 Kleberg, Fukai and Gilson. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.



Robust development of synfire chains from multiple plasticity mechanisms

Pengsheng Zheng and Jochen Triesch *

Frankfurt Institute for Advanced Studies, Frankfurt am Main, Germany

Edited by:

Matthieu Gilson, *Universitat Pompeu Fabra, Spain*

Reviewed by:

Paul Miller, *Brandeis University, USA*

Arvind Kumar, *University of Freiburg, Germany*

Matthieu Gilson, *Universitat Pompeu Fabra, Spain*

*Correspondence:

Jochen Triesch, *Frankfurt Institute for Advanced Studies, Ruth-Moufang-Str. 1, 60438 Frankfurt am Main, Germany*
e-mail: triesch@fias.uni-frankfurt.de

Biological neural networks are shaped by a large number of plasticity mechanisms operating at different time scales. How these mechanisms work together to sculpt such networks into effective information processing circuits is still poorly understood. Here we study the spontaneous development of synfire chains in a self-organizing recurrent neural network (SORN) model that combines a number of different plasticity mechanisms including spike-timing-dependent plasticity, structural plasticity, as well as homeostatic forms of plasticity. We find that the network develops an abundance of feed-forward motifs giving rise to synfire chains. The chains develop into ring-like structures, which we refer to as “synfire rings.” These rings emerge spontaneously in the SORN network and allow for stable propagation of activity on a fast time scale. A single network can contain multiple non-overlapping rings suppressing each other. On a slower time scale activity switches from one synfire ring to another maintaining firing rate homeostasis. Overall, our results show how the interaction of multiple plasticity mechanisms might give rise to the robust formation of synfire chains in biological neural networks.

Keywords: synfire chain, recurrent neural network, network self-organization, spike-timing-dependent plasticity, homeostatic plasticity, network motif

1. INTRODUCTION

Precise repetitions of neural activity patterns may serve as an infrastructure for numerous neural functions including sensory processing, motor control, and cognition. Synfire chains have been proposed as a fundamental network structure of the nervous system, which can guarantee a fixed level of network activity while allowing to learn and reproduce complicated spatio-temporal firing patterns (Abeles, 1982). Precise neural firing patterns have been found in many brain areas such as the songbird premotor nucleus (Hahnloser et al., 2002) and motor cortex of behaving monkeys (Prut et al., 1998; Shmiel et al., 2006). Studies on isolated neocortical microcircuits have revealed that spontaneous activity, mediated by a combination of intrinsic and circuit mechanisms, can be temporally precise in the absence of sensory stimulation (Mao et al., 2001; Luczak et al., 2007).

There is great interest in understanding how cortical circuits could acquire and maintain synfire-chain-like structures to give rise to relevant computations. Spike timing-dependent plasticity (STDP) has been proposed as a relevant mechanism in previous studies. Hertz and Prugel-Bennett (1996) tried to develop a synfire chain in a random network by introducing a Hebbian learning rule with one-step delay and n -winner-take-all dynamics. Successful learning required that the same training stimulus was shown to the system repeatedly. These stimuli, represented as sequences of activation patterns, determined the network dynamics which in turn determined the network connectivity due to STDP and other learning rules. The external stimuli were crucial for the synfire chain formation, because these stimuli generally drove the firing sequence of groups of neurons. Along similar lines, Levy et al. (2001) studied networks in the distributed

synchrony activity mode whose dynamics depended on an STDP learning rule and external input. Doursat and Bienenstock (2006) proposed an approach in which a set of seed neurons, a variant of spatiotemporal input, was also found essential for the growth of synfire chains. Similarly, Jun and Jin (2007) investigated an approach that also adopted suprathreshold external input. Hosaka et al. (2008) found that STDP provides a substrate for igniting synfire chains by spatiotemporal input patterns. Clopath et al. (2010) proposed a model of voltage-based STDP with homeostasis behaving similar to a triplet STDP (Pfister and Gerstner, 2006), which could develop variable connectivity patterns. Bourjaily and Miller (2011) studied the incorporation of structural plasticity with a rate-dependent (triplet) form STDP (Pfister and Gerstner, 2006) and the effect on motifs and distribution of synaptic strengths. Kunkel et al. (2011) suggested that biologically motivated plasticity mechanisms in the balanced random network model might lead to the development of feed-forward structures. Other recent approaches employed both different variants of STDP rules and spatiotemporal patterns of stimulation (Iglesias and Villa, 2008; Fiete et al., 2010; Waddington et al., 2012).

Overall, these previous works seem to suggest that the development of synfire chains requires either fine-tuning of model parameters, strong topological constraints on network connectivity, or guidance from strong spatiotemporally patterned training inputs. Here, we show that these limitations can be overcome in a network which combines STDP with additional plasticity mechanisms. We show that synfire chains form spontaneously from randomly initialized self-organizing recurrent networks (SORNs) in the absence of any structured external inputs.

Previous work has shown that SORNs with binary units can learn interesting representations of temporal sequences of sensory inputs (Lazar et al., 2009). Furthermore, we have shown that SORNs reproduce experimental data on the statistics and fluctuations of synaptic connection strengths in cortex and hippocampus, offering a plausible explanation for the experimentally observed approximately log-normal distribution of synaptic efficacies (Zheng et al., 2013). The networks self-organize their structure through a combination of STDP, homeostatic synaptic scaling, structural plasticity, and intrinsic plasticity of neuronal excitability. During network development, the topology adapts as STDP eliminates synaptic connections while structural plasticity adds new ones at a low rate. Meanwhile, the other plasticity mechanisms ensure that the network dynamics remains in a healthy regime.

Here we study the formation of synfire chains in such networks. The networks are initialized with a sparse random connectivity structure and go through dramatic changes in topology with a strong tendency to develop feed-forward motifs. These motifs eventually dominate sub-graph patterns as the network enters into a stable phase where connectivity stays roughly constant. Beyond a simple single feed-forward synfire chain structure, we find multiple ring-shaped chains within one network. The sizes of coactive pools of neurons are influenced by network parameters such as the average firing rate of the excitatory neurons. These results hold true over a wide range of parameters as long as the network operates in a “healthy regime,” supporting the view that synfire chains might be a robust consequence of network self-organization driven by multiple plasticity mechanisms. Overall, our model suggests that the combined action of multiple forms of neuronal plasticity may play an important role in shaping and maintaining cortical circuits and their dynamics, and stereotyped connectivity patterns could arise from the interplay of different plasticity mechanisms at the circuit level.

2. MATERIALS AND METHODS

The network model is identical to the one used by Zheng et al. (2013). It is composed of N^E excitatory and $N^I = 0.2 \times N^E$ inhibitory threshold neurons connected through weighted synaptic connections. Generally, W_{ij} is the connection strength from neuron j to neuron i . W^{EI} denotes inhibitory to excitatory connections, while W^{EE} and W^{IE} denote excitatory-to-excitatory and excitatory-to-inhibitory connections, respectively. The W^{EE} and W^{EI} are initialized as sparse random matrices with connection probabilities of 0.1 and 0.2, respectively.

Connections between inhibitory neurons and self-connections of excitatory neurons are not allowed. The W^{IE} connections are all-to-all and remain fixed at their random initial values which are drawn from a uniform distribution and are then normalized such that the sum of connections entering a neuron is one.

The binary vectors $x(t) \in \{0, 1\}^{N^E}$ and $y(t) \in \{0, 1\}^{N^I}$ denote the activity of the excitatory and inhibitory neurons at time step t , respectively. The network state at time step $t + 1$ is given by

$$x_i(t + 1) = \Theta \left(\sum_{j=1}^{N^E} W_{ij}^{EE}(t) x_j(t) - \sum_{k=1}^{N^I} W_{ik}^{EI}(t) y_k(t) - T_i^E(t) + \xi_{E_i}(t) \right), \quad (1)$$

$$y_i(t + 1) = \Theta \left(\sum_{j=1}^{N^E} W_{ij}^{IE} x_j(t) - T_i^I + \xi_{I_i}(t) \right). \quad (2)$$

The T^E and T^I represent threshold values for the excitatory and inhibitory neurons, respectively. They are initially drawn from a uniform distribution in the interval $[0, T_{\max}^E]$ and $[0, T_{\max}^I]$. $\Theta(\cdot)$ is the Heaviside step function. ξ_{E_i} and ξ_{I_i} are white Gaussian noise processes with $\mu_\xi = 0$ and $\sigma_\xi^2 \in [0.01, 0.05]$. Here one time step corresponds roughly to the duration of an STDP “window.”

The set of W^{EE} synapses adapts via a simplified causal STDP rule, as reported experimentally (Markram et al., 1997; Bi and Poo, 1998),

$$\Delta W_{ij}^{EE}(t) = \eta_{\text{STDP}} (x_i(t) x_j(t - 1) - x_i(t - 1) x_j(t)). \quad (3)$$

η_{STDP} is the learning rate. Note that synaptic weights are eliminated if they would become negative due to this rule. To compensate for the loss of synapses, a structural plasticity mechanism adds new synaptic connections between excitatory cells at a small rate. Specifically, with probability $p_c = 0.2$ a new connection (strength set to 0.001) is added between a randomly chosen pair of unconnected excitatory cells. This models the constant generation of new synaptic contacts observed in cortex and hippocampus (Johansen-Berg, 2007; Yasumatsu et al., 2008).

The incoming excitatory connections to an excitatory neuron are normalized at each time step such that their sum stays constant (Bourne and Harris, 2011). This is achieved by scaling the synapses multiplicatively (Turrigiano et al., 1998; Abbott and Nelson, 2000):

$$W_{ij}^{EE}(t) \leftarrow W_{ij}^{EE}(t) / \sum_j W_{ij}^{EE}(t). \quad (4)$$

A homeostatic (intrinsic) plasticity rule maintains a constant average firing rate in every excitatory neuron,

$$T_i^E(t + 1) = T_i^E(t) + \eta_{\text{IP}} (x_i(t) - H_i^{\text{IP}}), \quad (5)$$

where η_{IP} is the adaption rate and the target firing rates H_i^{IP} of individual neurons are drawn from a uniform distribution in $[\mu_{\text{IP}} - \sigma_{\text{HIP}}, \mu_{\text{IP}} + \sigma_{\text{HIP}}]$. In terms of firing rate homeostasis, there are very fast refractory mechanisms which prevent very high firing rates, and there is somewhat slower spike rate adaptation and very slow intrinsic plasticity as seen in some experiments (Desai et al., 1999; Zhang and Linden, 2003). We chose a simple homeostatic regulation of firing rate for our model that can operate relatively fast depending on the choice of the learning rate.

An inhibitory spike-timing dependent plasticity (iSTDP) rule adjusts the weights from inhibitory to excitatory neurons that balances the amount of excitatory and inhibitory drive that the excitatory neurons receive as reported in recent studies (Haas et al., 2006; Vogels et al., 2011, 2013),

$$\Delta W_{ij}^{EI}(t) = -\eta_{\text{inhib}} y_j(t - 1) (1 - x_i(t) (1 + 1/\mu_{\text{iSTDP}})), \quad (6)$$

where η_{inhib} is the adaption rate, and μ_{iSTDP} is set to 0.1 for all the simulations.

Unless otherwise specified, the simulations are conducted using the following parameters. $\eta_{IP} = 0.01$, $T_{\max}^E = 1$, $T_{\max}^I = 0.5$, $\mu_{IP} = 0.1$, $\sigma_{HIP} = 0$, $\eta_{inhib} = 0.001$, $\sigma_{\xi}^2 = 0.01$. Parameter η_{STDP} decreases monotonically as network size N^E increases, and $\eta_{STDP} = 0.004, 0.002$ and 0.001 for $N^E \in [200, 400]$, $[600, 800]$ and $[1000, 1200]$ respectively.

3. RESULTS

3.1. FEED-FORWARD MOTIFS DOMINATE SUBGRAPH PATTERNS

We simulate 10 networks, and initial weights of each network are randomly selected from uniform, Gaussian, delta (all weights identical), or exponential distributions. After weight initialization, each such network is examined on 10 different sets of network evolution parameters, such as neuron number, learning rates, neuron firing rates, etc. The network connectivity changes due to the action of the different plasticity mechanisms. As observed in Zheng et al. (2013), the network goes through different phases characterized by the number of excitatory-to-excitatory connections present in the network. Eventually, it enters a stable regime where connectivity stays roughly constant. For such stabilized networks we use the Fanmod software (Wernicke, 2005) and its computation of a p -value to analyze network motifs involving 3 and 4 neurons. Here the p -value of a motif is defined as the number of random networks in which it occurred more often than in the original network, divided by the total number of random networks. Therefore, p -values range from 0 to 1, and the smaller the p -value, the more significant is the abundance of the motif. The frequency of a motif occurring in 100 simulated SORN networks is compared to the mean frequency of the motif occurring in 1000 random networks with identical connection probability. We found the network motifs are organized into two distinct groups with p -value = 0 and p -value = 1. **Figure 1** shows the group of motifs always with p -value = 0, all of which reveal a feed-forward structure consistent with a synfire-chain topology.

3.2. EVOLUTION OF NETWORK CONNECTIVITY

The abundance of feed-forward network motifs among groups of 3 and 4 neurons during the stable phase of network evolution already suggests that the network may be forming synfire-chain like structures. To investigate this, we studied the evolution of the network's activity patterns and connectivity during its self-organization. **Figure 2** shows an example. In **Figure 2A** we plot the activity of the first 50 neurons during short 500 time step intervals taken at five different time points of the network's evolution. Excitatory neurons are sorted in all recorded networks according to their activity correlations in the last recorded network (in the stable phase). Thus neurons that are highly correlated during the stable phase are plotted in neighboring rows. While the network initially exhibits quite irregular activity, it spontaneously forms highly structured activity patterns as it develops (also see Figure S1 in the supplementary material, which shows example cross-correlograms of different pairs of neurons). In the particular case shown here, the network forms two subsets of neurons which alternate in exhibiting phases of high firing rates.




Figure 2B shows the evolution of firing correlations among all excitatory neurons in the network. The network forms 8 distinct pools of neurons, with neurons of each pool exhibiting highly synchronized firing. The excitatory weight matrix shown in **Figure 2C** reveals that the network develops two independent circular synfire chains, which we will refer to as synfire-rings. The layers of synfire rings are identified automatically by applying a threshold to the neurons' activity correlations. Due to noise and the interaction of multiple forms of plasticity, a neuron's activity maintains a certain degree of randomness, which leads to positive but non-uniform correlations in each layer. As a result there are some neurons with relatively weaker correlation in each layer in most cases.

In the given example, the first synfire-ring comprises 3 smaller pools of neurons (total of 43 neurons), the second synfire ring comprises 5 larger pools of neurons (total of 157 neurons). The two synfire-rings correspond to two transiently stable activity patterns. As shown in **Figure 2**, activities of the first 43 and remaining 7 neurons, which belong to different rings, are roughly complementary. If one synfire ring becomes active, it tends to activate the inhibitory neurons and thereby suppress activity in the other synfire ring. After a while, however, the intrinsic plasticity mechanism will increase the firing thresholds of neurons belonging to the active synfire-ring and decrease the firing thresholds of the inactive synfire-ring. Over time, this destabilizes the active synfire-ring and eventually leads to the suppressed synfire-ring taking over. The strong competition between

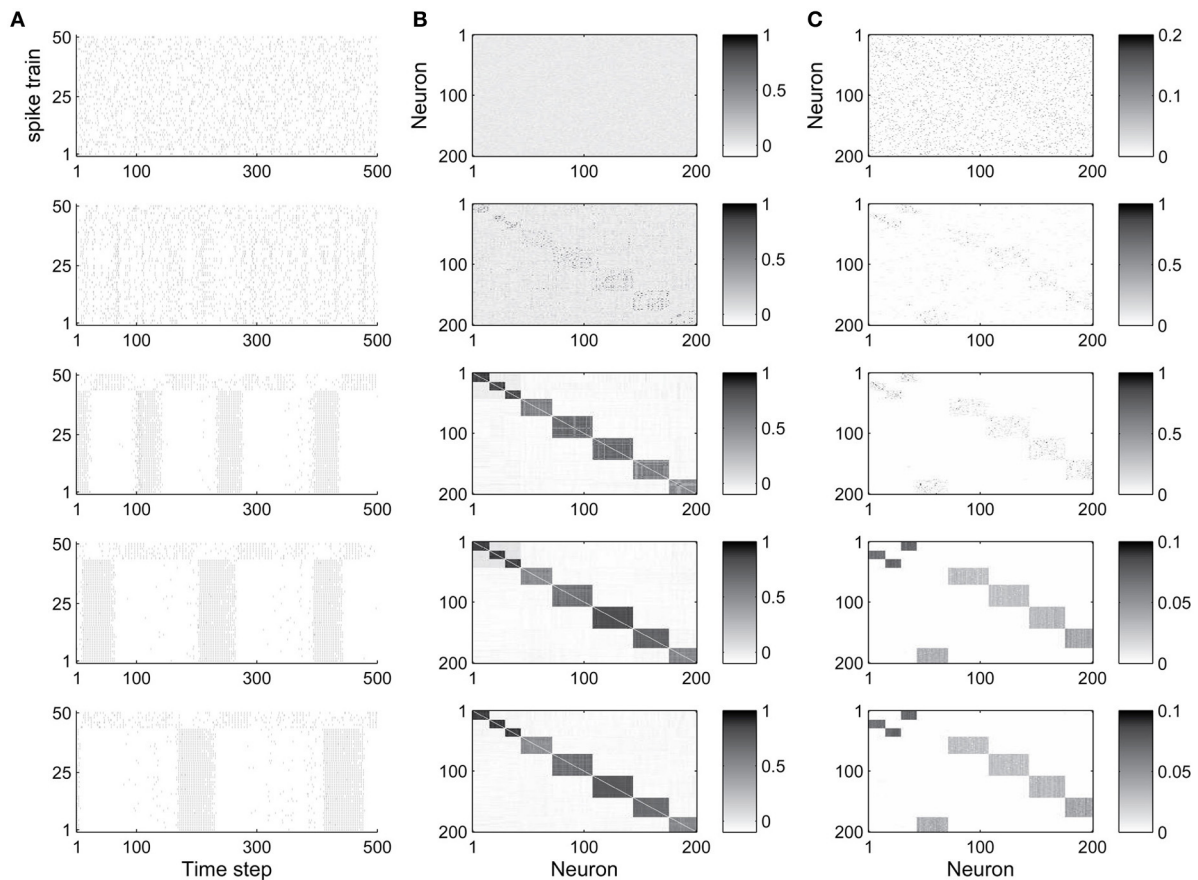


FIGURE 2 | Evolution of network dynamics and connectivity. (A) Spike trains of a set of 50 neurons during different phases of network development. **(B)** Activity correlations between all excitatory units. **(C)** Excitatory connectivity. Gray value represents excitatory synaptic

strength as illustrated in the scale bar. From top to bottom, the five rows show data starting from the 1st (initial phase), 20,000th (early phase), 500,000th (intermediate phase), 1,000,000th (late phase), and 4,000,000th (final phase) time step, respectively.

the synfire rings is due to the widespread inhibition with each inhibitory unit receiving input from all excitatory cells in the network and projecting randomly to one fifth of the excitatory cells (compare Methods).

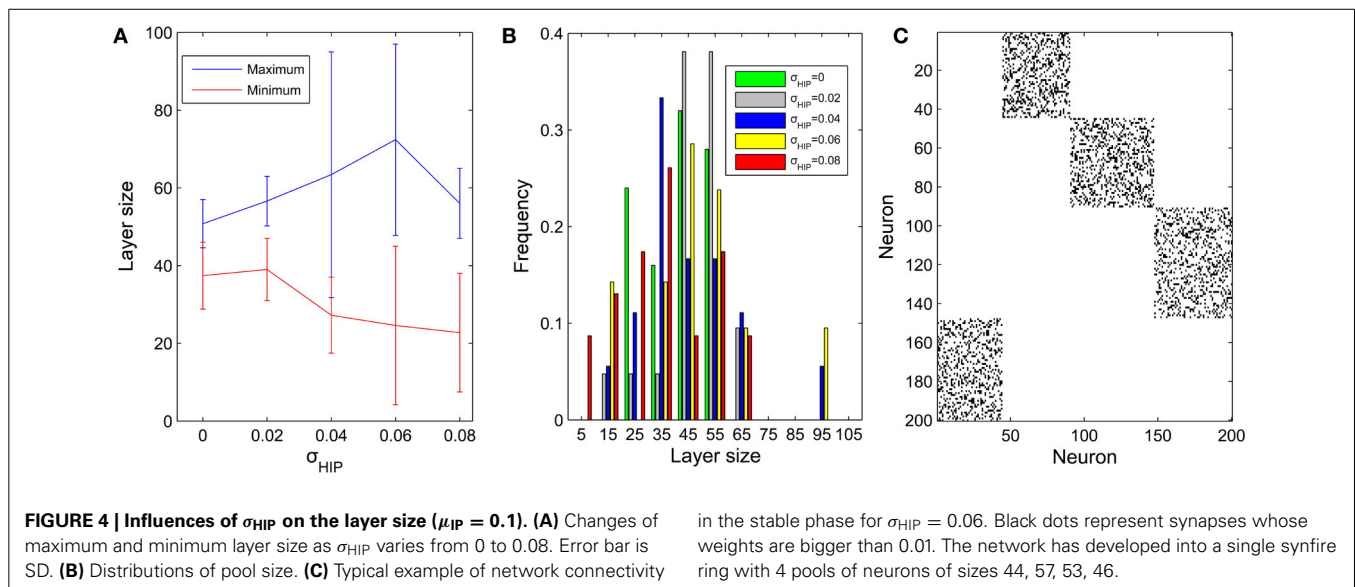
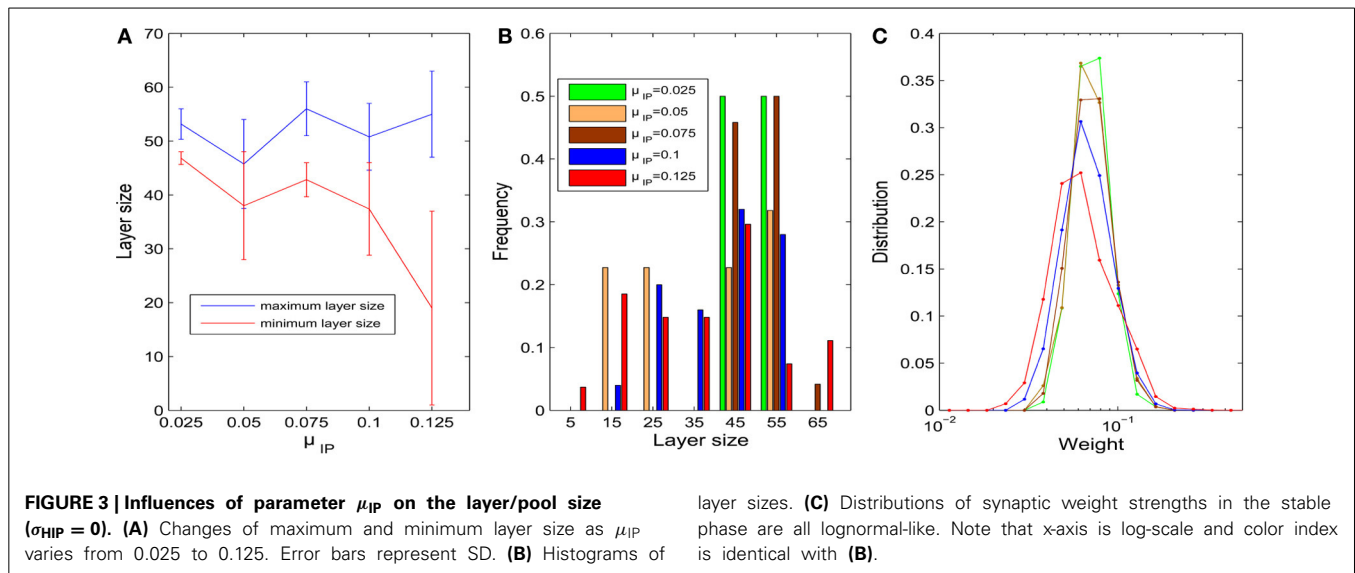
3.3. INFLUENCE OF TARGET FIRING RATE ON SIZES OF NEURONAL POOLS

We next investigate how the sizes of neuronal pools and their connectivity depend on the target firing rates of the neurons in a 200 excitatory neuron network with fixed initial connectivity. The parameter H_i^{IP} sets the target firing rate for the i -th excitatory neuron. These target firing rates are drawn from a uniform distribution in $[\mu_{\text{IP}} - \sigma_{\text{HIP}}, \mu_{\text{IP}} + \sigma_{\text{HIP}}]$.

We first fix $\sigma_{\text{HIP}} = 0$ and study the influence of the target firing rate μ_{IP} . As the target firing rate of the neurons increases, the variability of the sizes of neuronal pools increases. **Figure 3A** plots the average maximum and minimum pool sizes as a function of μ_{IP} . For large μ_{IP} , the maximum layer size tends to get bigger and the minimum layer size tends to be smaller. In addition, the variability of the maximum and especially the minimum layer sizes tends to be largest for the biggest μ_{IP} . **Figure 3B** compares the

histograms of pool sizes for different μ_{IP} . The distribution is very narrow for small μ_{IP} (green bars corresponding to $\mu_{\text{IP}} = 0.025$) and very broad for large μ_{IP} (red bars corresponding to $\mu_{\text{IP}} = 0.125$). In all cases, the final distribution of synaptic strength is lognormal-like which means some weights are way stronger than others. This is shown in **Figure 3C**, which plots this distribution for different μ_{IP} .

We next fix $\mu_{\text{IP}} = 0.1$ and study the effect of the interval size σ_{HIP} of the target firing rates. In a similar way, the diversity of pool sizes grows as σ_{HIP} increases. This holds true for $\sigma_{\text{HIP}} \leq 0.06$ as shown in **Figure 4A**. However, as σ_{HIP} increases more and more neurons are close to silent. The minimum target firing rate of some excitatory neurons is as small as ~ 0.02 when σ_{HIP} reaches 0.08. These neurons barely fire during the network evolution and barely contribute to structuring the network. Therefore, the effective network size is reduced as σ_{HIP} is increased. This may explain why the variability in pool sizes shrinks when σ_{HIP} grows to 0.08. **Figure 4B** compares the distribution of pool sizes for different values of σ_{HIP} . The greatest spread of the distribution is obtained for an intermediate value of $\sigma_{\text{HIP}} = 0.06$. **Figure 4C** shows an example of an excitatory weight matrix in the stable regime for



$\sigma_{HIP} = 0.06$. The network has developed a single synfire ring with 4 pools of neurons whose sizes range from 44 to 57.

3.4. INFLUENCE OF NETWORK SIZE ON SYNFIRED RING STRUCTURE

We next study how the number of synfire rings and the number of neuronal pools or layers depends on the overall network size. To this end, we simulate 40 networks with 200–800 excitatory neurons. As a first measure of network structure we define the number of layers present in the network. **Figure 5A** plots this number as a function of network size (red curve). Not surprisingly, the number of neuronal pools increases as the network gets bigger. As a second index of network structure we measure the fraction of networks of a given size that develop multiple synfire rings. As shown in **Figure 5A** (blue curve) this fraction increases with network size. For networks of 800 neurons it already reaches a value of 0.4 and the increase with network size seems to be faster

than linear for the range of sizes considered. **Figure 5B** shows a typical example of the excitatory weight matrix in a network with 800 neurons and $\mu_{IP} = 0.1$, $\sigma_{HIP} = 0$. This network has developed 4 synfire rings of different sizes. Note that the second one from the top is very small. In **Figure 5C** it is easier to identify it. The sizes of the pools are fairly consistent within a single synfire ring (mean SD is 4.5) but can vary widely across synfire rings (SD is 26.6). The biggest ring in **Figure 5B** has 12 pools, so if activity runs around in this circle, each neuron is activated only every 12th time step, which is less than intrinsic plasticity wants (compare Methods). As shown in **Figure 5C**, the biggest ring is roughly active all the time, and unlike the synfire rings in **Figure 2**, the network could start multiple rings simultaneously. It is worth noting that we achieve the synfire ring structure under a wide range of parameters, excitatory neuron number being one of them. We also run a few simulations with 1000 and 1200 excitatory

neurons, which also develop synfire rings (see Figures S2, S3 in supplementary material).

3.5. MECHANISMS OF SYNFIRES RING FORMATION

With all forms of plasticity present, the network will develop synfire rings spontaneously and robustly over a large range of parameters as long as the network operates in a healthy regime. The results are fully in line with our previous work since we use same network as Zheng et al. (2013), where we discuss in detail the necessity of the different plasticity mechanisms for the networks behavior. So how do these (circular) feed-forward structures come about?

The formation of synfire-chains can be understood as a process of network self-organization driven largely by the STDP rule. **Figure 6** illustrates the process. Consider as an example a strong feed-forward chain from a unit a to a unit b and on to a unit c . According to this structure, there is a high probability that a , b , and c fire in three successive time steps. Standard STDP rules, including the one we are using here, will strengthen the connections in the feed-forward direction and weaken the reverse connections such as the red synapse in **Figure 6A**. This is because of the nature of the STDP rule, which potentiates “causal” firing

patterns (pre before post) and depresses “acausal” firing patterns (post before pre). As shown in **Figure 7A**, the fraction of bidirectional connections plummets during the first stage of network evolution. Thus, a first relevant mechanism in synfire ring formation is the *removal of reciprocal connections*.

A second mechanism in synfire ring formation is the *establishment of parallel pathways*. Consider two units b_1 and b_2 which also happen to be strongly innervated by a (see **Figure 6B**). Because of this, they will tend to be synchronously active with unit b and their activity will be reliably followed by activation of unit c . Because of this correlation structure (b_1 and b_2 likely being active in the time step before c) the weights from b_1 and b_2 onto c , if present, will have a strong tendency to get potentiated. Thus, STDP will potentiate the “missing” connections from b_1 and b_2 onto c establishing additional parallel pathways connecting a and c . In order for STDP to be able to strengthen these connections, they have to either be present from the beginning or become added by the structural plasticity. With this mechanism operating not just at the level of b but at all levels of the network, synfire-chains will develop (see **Figure 6C**). Due to the homeostatic activity regulation, at each time step a certain fraction of neurons in the network will tend to be active. This implicitly

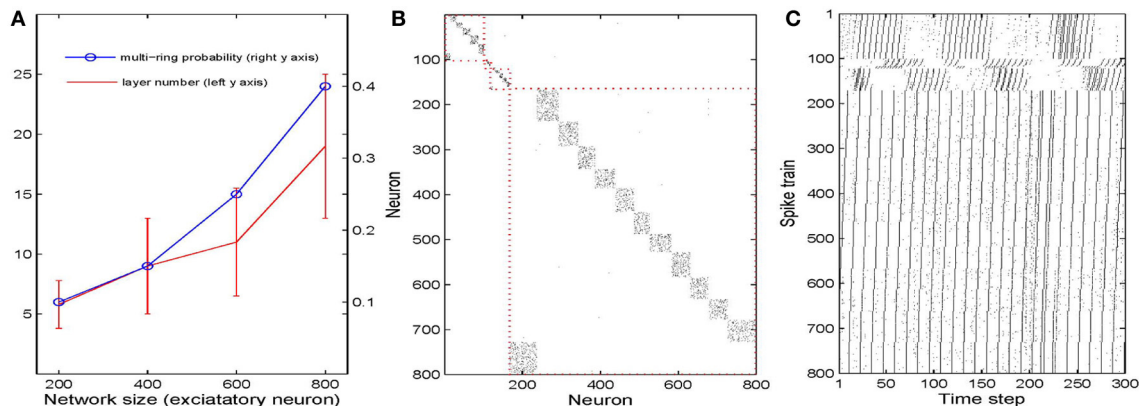


FIGURE 5 | Influences of network size. (A) Changes of multi-ring probability and layer number as network size varies from 200 to 800. Error bars are SD. **(B)** Typical example of network connectivity with four synfire rings in the

stable phase of a 800 excitatory neuron network. Black dots represent synapses whose weights are bigger than 0.01. **(C)** Spike trains of the neurons in **(B)**.




FIGURE 6 | Layered synfire chain structure formation. (A) Removal of reciprocal connection. **(B)** Establishment of parallel pathways. **(C)** Formed synfire chain. Red arrows represent spurious synapses

that are inconsistent with the developing synfire chain structure. Black arrows represent synapses that conform to the synfire route.

regulates the range of layer sizes and limits the breadth of the growing chain. Due to the synaptic scaling, every neuron in a layer receives a certain amount of synaptic input. Moreover, since the network has only a finite number of units and each unit tries to maintain a certain average activity level such that activity cannot die out, it is inevitable that such a chain eventually terminates or connects back to itself thereby forming a synfire ring. A ring-like structure has a competitive advantage against a terminating chain during the formative stage of network development, because a synfire ring will reactivate itself while a terminating chain cannot.

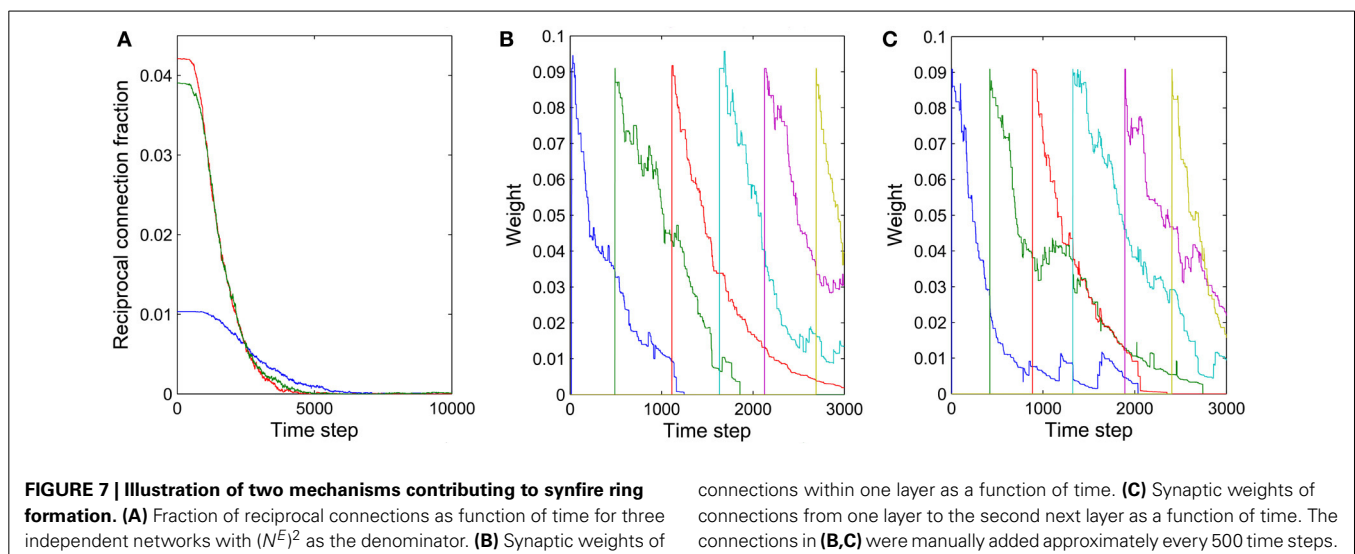
STDP alone can not depress existing synapses that are incompatible with the emerging synfire ring structure. For example, connections within one layer of neurons or connections jumping ahead beyond the immediate next layer (compare red synapses in **Figure 6B**) remain unaltered under perfect synfire chain activity. However, the synaptic scaling mechanism gradually depresses these connections to very small values as the other weights on the synfire route are potentiated. Thus, another relevant mechanism in synfire ring formation is the *competition among synaptic weights onto the same target neuron*.

The synaptic scaling mechanism we use does not remove any such “spurious” synapses, however. This is achieved by STDP. Due to intrinsic membrane noise of the neurons and fluctuations of intrinsic excitability and inhibitory drive, the network’s activity always maintains a random component—even in its stable phase (compare **Figure 2A**). As a consequence, neuron *a* in **Figure 6** could fire right after *c*, which would lead to the removal of a sufficiently depressed spurious connection from *a* to *c*. Such events occur only rarely but they suffice to eliminate such spurious connections if they have already been depressed by synaptic scaling. Thus, a final mechanism in synfire ring formation is the *removal of spurious connections due to random activity fluctuations*. To illustrate this effect, we manually added new connections with rather strong weights of value 0.1 within one layer and between one layer and the layer two steps ahead. These manually added new connections are even stronger than ~70% of the existing connections. **Figures 7B,C** show the fate of

these manually inserted connections that are inconsistent with the dominant synfire-ring structure: within a few thousand time steps their weights decrease to zero as a result of competition among synapses and STDP driven by random activity fluctuations.

The precise outcome of the overall network self-organization depends on the initial conditions (initial network structure and connection weights) and the random activity fluctuations. Feed-forward connections between synfire layers go through strong competition during network evolution as a result of synaptic scaling. Connections that start out strong or are added early have an advantage in this competition and are less prone to removal due to random activity fluctuations. Synapses added in later phases of the network’s evolution are more fragile, which contributes to the stability of already formed synfire rings. **Figure 8** shows examples of weight changes of new synapses that have been added through structural plasticity during the network’s stable phase. In **Figure 8A** we plot synaptic connections that are off any existing synfire ring structure. These connections are removed comparatively quickly. **Figure 8B** illustrates the fate of newly added synaptic connections that are congruent with an existing synfire ring, i.e., they connect a neuron from one pool to a neuron in the successor pool. Interestingly, even these connections tend to be removed eventually. Due to synaptic scaling, they have to compete with many other connections along the synfire ring, which limits their growth and makes them prone to elimination due to random activity fluctuations. It should be clarified that **Figures 7B,C** (unlike **Figure 7A**) and **Figure 8** all study the network in its stable phase. That is the synfire chain is already formed, which is analogous to a prewired synfire chain. In all of these cases, the synfire chain is indeed restored after our perturbation. **Figures 7B,C, 8B** are similar, but they are different in terms of new synapse weight and position definition.

It should be mentioned that the network becomes rather rigid only after synfire chains have been formed. This is important to maintain the stability of synfire chains. At the beginning of network evolution, however, newer synapses are freely competing. During the development from a randomly initialized network






FIGURE 8 | Weights of synaptic connections that have been added by structural plasticity during the stable phase as a function of time. Colors represent different synapses. **(A)** Ten synapses added on

the synfire route. **(B)** Ten synapses that are not on the synfire route. Note that in both cases some of the synapses are eliminated immediately after birth.

to synfire chains, many new synapses are added and stabilized. Besides that, newer synapses could also survive in synfire chains when the network is driven by appropriate strong structured external input (not shown).

As mentioned above, every plasticity mechanism is important for the development of synfire structure. In simulations, we did not observe the formation of synfire rings in networks without synaptic normalization, structural plasticity or STDP of the excitatory connections. Intrinsic plasticity and inhibitory STDP both try to maintain a low average firing rate of excitatory cells, and the formation of synfire structure relies on the presence of both of them. If we switch off one of them, the network suffers from big activity fluctuations from time to time, which usually stop the formation of synfire structure or lead to abnormal network structures exhibiting both extra-large and single-neuron layers (Figure S4 in the supplementary material).

4. DISCUSSION

Since their introduction by Abeles (1982), synfire-chains have been the subject of intense experimental and theoretical investigation. Here we have studied the spontaneous formation of synfire-chains in self-organizing recurrent neural networks (SORNs) shaped by multiple plasticity mechanisms. These networks have been shown to learn effective representations of time-varying inputs (Lazar et al., 2009) and to reproduce data on the statistics and fluctuations of synaptic connections strength in cortex and hippocampus (Zheng et al., 2013). There is also some empirical evidence for their ability to approximate Bayesian inference (Lazar et al., 2011). Despite their simplicity in terms of using binary threshold units operating in discrete time, they have been a useful tool for studying the interaction of different forms of plasticity at the network level. In the present study, we have combined simple spike-timing-dependent plasticity (STDP) rules for

excitatory-to-excitatory and inhibitory-to-excitatory connections with a synaptic normalization and firing rate homeostasis of excitatory units. Furthermore, a structural plasticity rule created new excitatory-to-excitatory connections at a low rate.

The initial connection probability of excitatory to excitatory connections is set to 0.1, which falls in the biologically plausible range. In simulations, we couldn't decrease this probability further, otherwise the network decomposed into unconnected smaller ones. In some cases, the structural plasticity may reconnect these pieces, but not all the time. Generally, it is hard to draw any conclusion from the simulations of such unhealthily initialized networks. In the other extreme, we can increase initial connection probability all the way to a fully connected network and the network will still develop synfire rings.

We found that the STDP mechanism prunes bidirectional connections between pairs of excitatory neurons, which is consistent with previous modeling work (Abbott and Nelson, 2000). It is interesting to note that there maybe layer-specific differences in cortex in terms of the abundance of such bidirectional synaptic connections with layer 5 showing many bidirectional connections in one study (Song et al., 2005), but layer 4-2/3 showing very few (Feldmeyer et al., 2002; Lefort et al., 2009). The pruning of bidirectional connections goes hand in hand with the emergence of feed-forward chains among pools of neurons. This formation of synfire-chains represents a phenomenon of network-self-organization. Partial feed-forward structures between pools of neurons have a tendency to become amplified due to STDP while the homeostatic plasticity mechanisms induce competition among the developing feed-forward structures. These feed-forward chains assume a ring-shaped topology, which we refer to as synfire-rings. We observed that the number of synfire rings, their lengths, and the sizes of their pools are influenced by the distribution of firing rates. The development of "orderly" synfire dynamics in these networks is consistent with previous results

indicating a reduction of chaotic behavior in these networks (Eser et al., 2014).

Previous modeling studies on the formation of synfire-chains have used more realistic model neurons and synapses, but have omitted some of the plasticity mechanisms incorporated into the present model. As was shown previously (Zheng et al., 2013), these mechanisms may be essential for explaining critical aspects of cortical wiring such as the log-normal distribution of excitatory-to-excitatory synaptic efficacies and the pattern of fluctuations of synaptic efficacies. It is clear that a long-tailed, highly skewed distribution of synaptic efficacies may strongly affect synfire dynamics, since the simultaneous activation of only few extremely strong synapses may suffice to elicit an action potential in the postsynaptic neuron. In the present study, lognormal-like statistics of excitatory synaptic connections develop robustly in the network (see **Figure 3C**). To our knowledge, no previous study has investigated synfire dynamics with lognormally distributed excitatory-to-excitatory efficacies. Our model does not only demonstrate synfire dynamics with a biologically realistic distribution of excitatory-to-excitatory synaptic efficacies, it also shows how this distribution and synfire dynamics emerge from fundamental plasticity mechanisms in the absence of any structured input to the network.

Overall, we conclude that the combination of a number of generic plasticity mechanisms is sufficient for the robust formation of synfire chains with synaptic connection statistics matching biological data. Many aspects of our model could be made more realistic. For instance, it will be important to go beyond networks of binary threshold units operating in discrete time steps. We would like to test if similar results can be obtained in more realistic networks of spiking neurons operating in continuous time. Another limitation is that we have assumed identical one time step conduction delays of all synaptic connections. Izhikevich (2006) however found that conduction delays were important for time-locked but not synchronous spiking activity, and managed to generate many more synfire “braids” than the number of neurons in the network. The consideration of heterogeneous conduction delays in a more realistic version of our model is an interesting topic for future work.

AUTHOR CONTRIBUTIONS

Conceived and designed the experiments: Pengsheng Zheng, Jochen Triesch. Performed the experiments: Pengsheng Zheng. Analyzed the data and plotted the results: Pengsheng Zheng. Wrote the paper: Jochen Triesch, Pengsheng Zheng.

ACKNOWLEDGMENTS

FUNDING

This work was supported by the European Commission 7th Framework Programme (FP7/20072013), Challenge 2—Cognitive Systems, Interaction, Robotics, grant agreement No. ICT-IP-231722, project “IM-CLeVeR-Intrinsically Motivated Cumulative Learning Versatile Robots,” by the LOEWE-Program Neuronal Coordination Research Focus Frankfurt (NeFF, <http://www.neff-ffm.de/en/>). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <http://www.frontiersin.org/journal/10.3389/fncom.2014.00066/abstract>

REFERENCES

- Abbott, L. F., and Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nat. Neurosci.* 3, 1178–1183. doi: 10.1038/81453
- Abeles, M. (1982). *Local Cortical Circuits: an Electrophysiological Study*. Berlin; New York: Springer-Verlag. doi: 10.1007/978-3-642-81708-3
- Bi, G. Q., and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Bourjaily, M. A., and Miller, P. (2011). Excitatory, inhibitory, and structural plasticity produce correlated connectivity in random networks trained to solve paired-stimulus tasks. *Front. Comput. Neurosci.* 5:37. doi: 10.3389/fncom.2011.00037
- Bourne, J. N., and Harris, K. M. (2011). Coordination of size and number of excitatory and inhibitory synapses results in a balanced structural plasticity along mature hippocampal CA1 dendrites during LTP. *Hippocampus* 21, 354–373. doi: 10.1002/hipo.20768
- Clopath, C., Busing, L., Vasilaki, E., and Gerstner, W. (2010). Connectivity reflects coding: a model of voltage-based stdp with homeostasis. *Nat. Neurosci.* 13, 344–352. doi: 10.1038/nn.2479
- Desai, N. S., Rutherford, L. C., and Turrigiano, G. G. (1999). Plasticity in the intrinsic excitability of cortical pyramidal neurons. *Nat. Neurosci.* 2, 515–520. doi: 10.1038/9165
- Doursat, R., and Bienenstock, E. (2006). “The self-organized growth of synfire patterns,” in *10th International Conference on Cognitive and Neural Systems (ICCN 2006)* (Massachusetts, MA: Boston University).
- Eser, J., Zheng, P., and Triesch, J. (2014). Nonlinear dynamics analysis of a self-organizing recurrent neural network: chaos waning. *PLoS ONE* 9:e86962. doi: 10.1371/journal.pone.0086962
- Feldmeyer, D., Lubke, J., Silver, R. A., and Sakmann, B. (2002). Synaptic connections between layer 4 spiny neurone-layer 2/3 pyramidal cell pairs in juvenile rat barrel cortex: physiology and anatomy of interlaminar signalling within a cortical column. *J. Physiol. Lond.* 538, 803–822. doi: 10.1111/jphysiol.2001.012959
- Fiete, I. R., Senn, W., Wang, C. Z. H., and Hahnloser, R. H. R. (2010). Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron* 65, 563–576. doi: 10.1016/j.neuron.2010.02.003
- Haas, J. S., Nowotny, T., and Abarbanel, H. D. I. (2006). Spike-timing-dependent plasticity of inhibitory synapses in the entorhinal cortex. *J. Neurophysiol.* 96, 3305–3313. doi: 10.1152/jn.00551.2006
- Hahnloser, R. H. R., Kozhevnikov, A. A., and Fee, M. S. (2002). An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* 419, 65–70. doi: 10.1038/nature00974
- Hertz, J., and Prugel-Bennett, A. (1996). Learning synfire chains: turning noise into signal. *Int. J. Neural Syst.* 7, 445–450. doi: 10.1142/S0129065796000427
- Hosaka, R., Araki, O., and Ikeguchi, T. (2008). STDP provides the substrate for igniting synfire chains by spatiotemporal input patterns. *Neural Comput.* 20, 415–435. doi: 10.1162/neco.2007.11-05-043
- Iglesias, J., and Villa, A. E. P. (2008). Emergence of preferred firing sequences in large spiking neural networks during simulated neuronal development. *Int. J. Neural Syst.* 18, 267–277. doi: 10.1142/S0129065708001580
- Izhikevich, E. M. (2006). Polychronization: computation with spikes. *Neural Comput.* 18, 245–282. doi: 10.1162/089976606775093882
- Johansen-Berg, H. (2007). Structural plasticity: rewiring the brain. *Curr. Biol.* 17, R141–R144. doi: 10.1016/j.cub.2006.12.022
- Jun, J., and Jin, D. (2007). Development of neural circuitry for precise temporal sequences through spontaneous activity, axon remodeling, and synaptic plasticity. *PLoS ONE* 2:e723. doi: 10.1371/journal.pone.0000723
- Kunkel, S., Diesmann, M., and Morrison, A. (2011). Limits to the development of feed-forward structures in large recurrent neuronal networks. *Front. Comput. Neurosci.* 4:160. doi: 10.3389/fncom.2010.00160
- Lazar, A., Pipa, G., and Triesch, J. (2009). SORN: a self-organizing recurrent neural network. *Front. Comput. Neurosci.* 3:23. doi: 10.3389/fncom.2009.10.023.2009

- Lazar, A., Pipa, G., and Triesch, J. (2011). "Emerging bayesian priors in a self-organizing recurrent network," in *Artificial Neural Networks and Machine Learning - Iccnn 2011, Pt II*, Vol. 6792 (Espoo), 127–134.
- Lefort, S., Tamm, C., Sarria, J. C. F., and Petersen, C. C. H. (2009). The excitatory neuronal network of the c2 barrel column in mouse primary somatosensory cortex. *Neuron* 61, 301–316. doi: 10.1016/j.neuron.2008.12.020
- Levy, N., Horn, D., Meilijson, I., and Ruppin, E. (2001). Distributed synchrony in a cell assembly of spiking neurons. *Neural Netw.* 14, 815–824. doi: 10.1016/S0893-6080(01)00044-2
- Luczak, A., Bartho, P., Marguet, S. L., Buzsaki, G., and Harris, K. D. (2007). Sequential structure of neocortical spontaneous activity *in vivo*. *Proc. Natl. Acad. Sci. U.S.A.* 104, 347–352. doi: 10.1073/pnas.0605643104
- Mao, B. Q., Hamzei-Sichani, F., Aronov, D., Froemke, R. C., and Yuste, R. (2001). Dynamics of spontaneous activity in neocortical slices. *Neuron* 32, 883–898. doi: 10.1016/S0896-6273(01)00518-9
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic apss and epsps. *Science* 275, 213–215. doi: 10.1126/science.275.5297.213
- Pfister, J.-P., and Gerstner, W. (2006). Triplets of spikes in a model of spike timing-dependent plasticity. *J. Neurosci.* 26, 9673–9682. doi: 10.1523/JNEUROSCI.1425-06.2006
- Prut, Y., Vaadia, E., Bergman, H., Haalman, I., Slovin, H., and Abeles, M. (1998). Spatiotemporal structure of cortical activity: Properties and behavioral relevance. *J. Neurophysiol.* 79, 2857–2874.
- Shmuel, T., Drori, R., Shmuel, O., Ben-Shaul, Y., Nadasdy, Z., Shemesh, M., et al. (2006). Temporally precise cortical firing patterns are associated with distinct action segments. *J. Neurophysiol.* 96, 2645–2652. doi: 10.1152/jn.00798.2005
- Song, S., Sjostrom, P. J., Reigl, M., Nelson, S., and Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol.* 3:e350. doi: 10.1371/journal.pbio.0030350
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., and Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391, 892–896. doi: 10.1038/36103
- Vogels, T. P., Froemke, R. C., Doyon, N., Gilson, M., Haas, J. S., Liu, R., et al. (2013). Inhibitory synaptic plasticity: spike timing-dependence and putative network function. *Front. Neural Circuits* 7:119. doi: 10.3389/fncir.2013.00119
- Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* 334, 1569–1573. doi: 10.1126/science.1211095
- Waddington, A., Appleby, P. A., De Kamps, M., and Cohen, N. (2012). Triphasic spike-timing-dependent plasticity organizes networks to produce robust sequences of neural activity. *Front. Comput. Neurosci.* 6:88. doi: 10.3389/fncom.2012.00088
- Wernicke, S. (2005). A faster algorithm for detecting network motifs. *Algorithms Bionform. Proc.* 3692, 165–177. doi: 10.1007/11557067-14
- Yasumatsu, N., Matsuzaki, M., Miyazaki, T., Noguchi, J., and Kasai, H. (2008). Principles of long-term dynamics of dendritic spines. *J. Neurosci.* 28, 13592–13608. doi: 10.1523/JNEUROSCI.0603-08.2008
- Zhang, W., and Linden, D. J. (2003). The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nat. Rev. Neurosci.* 4, 885–900. doi: 10.1038/nrn1248
- Zheng, P., Dimitrakakis, C., and Triesch, J. (2013). Network self-organization explains the statistics and dynamics of synaptic connection strengths in cortex. *PLoS Comput. Biol.* 9:e1002848. doi: 10.1371/journal.pcbi.1002848

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 15 February 2014; accepted: 02 June 2014; published online: 30 June 2014.
Citation: Zheng P and Triesch J (2014) Robust development of synfire chains from multiple plasticity mechanisms. *Front. Comput. Neurosci.* 8:66. doi: 10.3389/fncom.2014.00066
This article was submitted to the journal *Frontiers in Computational Neuroscience*.
Copyright © 2014 Zheng and Triesch. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.



Dynamic stability of sequential stimulus representations in adapting neuronal networks

Renato C. F. Duarte^{1,2,3,4*} and Abigail Morrison^{1,2,3,5}

¹ Institute of Neuroscience and Medicine (INM-6) and Institute for Advanced Simulation (IAS-6), Jülich Research Center and JARA, Jülich, Germany

² Bernstein Center Freiburg, Albert-Ludwig University of Freiburg, Freiburg im Breisgau, Germany

³ Faculty of Biology, Albert-Ludwig University of Freiburg, Freiburg im Breisgau, Germany

⁴ School of Informatics, Institute of Adaptive and Neural Computation, University of Edinburgh, Edinburgh, UK

⁵ Faculty of Psychology, Institute of Cognitive Neuroscience, Ruhr-University Bochum, Bochum, Germany

Edited by:

Friedemann Zenke, École
Polytechnique Fédérale de
Lausanne, Switzerland

Reviewed by:

Pierre Yger, École Normale
Supérieure, France
Tim P. Vogels, Oxford University, UK

*Correspondence:

Renato C. F. Duarte, Institute of
Neuroscience and Medicine (INM-6)
and Institute for Advanced
Simulation (IAS-6), Jülich Research
Center and JARA, Building 15.22,
52425 Jülich, Germany
e-mail: r.duarte@fz-juelich.de

The ability to acquire and maintain appropriate representations of time-varying, sequential stimulus events is a fundamental feature of neocortical circuits and a necessary first step toward more specialized information processing. The dynamical properties of such representations depend on the current state of the circuit, which is determined primarily by the ongoing, internally generated activity, setting the ground state from which input-specific transformations emerge. Here, we begin by demonstrating that timing-dependent synaptic plasticity mechanisms have an important role to play in the active maintenance of an ongoing dynamics characterized by asynchronous and irregular firing, closely resembling cortical activity *in vivo*. Incoming stimuli, acting as perturbations of the local balance of excitation and inhibition, require fast adaptive responses to prevent the development of unstable activity regimes, such as those characterized by a high degree of population-wide synchrony. We establish a link between such pathological network activity, which is circumvented by the action of plasticity, and a reduced computational capacity. Additionally, we demonstrate that the action of plasticity shapes and stabilizes the transient network states exhibited in the presence of sequentially presented stimulus events, allowing the development of adequate and discernible stimulus representations. The main feature responsible for the increased discriminability of stimulus-driven population responses in plastic networks is shown to be the decorrelating action of inhibitory plasticity and the consequent maintenance of the asynchronous irregular dynamic regime both for ongoing activity and stimulus-driven responses, whereas excitatory plasticity is shown to play only a marginal role.

Keywords: stimulus representation, synaptic plasticity, excitation/inhibition balance, transient dynamics, asynchronous irregular states, online computation

1. INTRODUCTION

As we navigate the world, we are continuously exposed to dynamic and highly complex streams of multimodal sensory information, which we tend to perceive as a series of discrete and coherently bounded sub-sequences (Schapiro et al., 2013). While these *perceptual events* (Zacks and Tversky, 2001; Zacks et al., 2007) are unfolding, active representations are maintained and ought to be sufficiently discernible by the activity of the processing networks, its attributes being encoded by the distributed responses of specifically tuned neuronal populations that are transiently associated into coherent ensembles (von der Malsburg et al., 2010; Singer, 2013).

Neocortical circuits must therefore self-organize to dynamically adopt relevant representations in a stimulus- and state-dependent manner, while maintaining the necessary sensitivity to allow global shifts in representational space when sudden event transitions occur. The primary source of such sensitivity in stereotypically sparse recurrent networks, such as those encountered in the neocortex, lies in the balance of excitation and inhibition

(Tsodyks and Sejnowski, 1995; van Vreeswijk and Sompolinsky, 1998; Haider et al., 2006). This endows the network with a stable, ongoing background activity, characterized by low-frequency, asynchronous and irregular firing patterns, under stationary conditions (Gerstein and Mandelbrot, 1964; Softky and Koch, 1993; Destexhe et al., 2003; Stiefel et al., 2013). Such a dynamic state provides the substrate for complex responses to external stimuli to develop as a transient spatiotemporal succession of network states (Mazor and Laurent, 2005; Rabinovich et al., 2008; Buonomano and Maass, 2009). External stimuli act as perturbations of the stable ongoing activity, causing transient and variable disruptions of local E/I balance, which necessarily influence the resulting network states. Furthermore, most real world stimulus events occur sequentially and not in isolation. Consequently, the quality of the dynamic representations and characteristics of the resulting *neural trajectories* is highly related to the circuit's ability to adaptively remodel and refine its functional connectivity in an experience-dependent manner so as to counteract the targeted disruptions and acquire the relevant structure of the input stimuli.

Although there is a great variety of biophysical mechanisms involving activity-dependent modifications of various components at different spatial and temporal scales, it is widely acknowledged that the synapse is the primary locus of functional adaptation in the cortex (Abbott and Nelson, 2000), with synaptic modifications providing the basis of learning and memory and allowing purposeful computations to take place. While constituting a diverse set, comprising operations over variable dynamic ranges and involving a multitude of possible functional roles, cortical synapses can be broadly categorized based on the nature of source and target neurons they connect and the effect they exert (excitatory or inhibitory). Understanding and exploring the possible adaptation mechanisms involved in each of these sub-classes and how they interact is important to understand the nature of neural computation. It is reasonable to assume that the required flexibility to support highly complex cognitive computations, relies on the effects of the combined action of multiple, synergistic, plasticity mechanisms.

There is a large body of experimental evidence and theoretical investigations concerning adaptation at excitatory synapses. It has long been experimentally observed that, in cortical pyramidal neurons, the magnitude and direction of change in the strength of a synapse is dependent on the relative timing of pre- and post-synaptic spikes, when they occur within a critical coincidence time window (Gustafsson et al., 1987; Markram et al., 1997; Bi and Poo, 1998). The functional implications of this observation for cortical processing and unsupervised, experience-dependent adaptation have since been the subject of intense investigation and have proven capable to account for several important computational features of cortical processing (for reviews, see e.g., Dan and Poo, 2004, 2006; Sjöström and Gerstner, 2010; Markram et al., 2012). However, despite this progress, attempts to endow recurrent networks with the ability to learn the underlying structure of their inputs using excitatory spike timing dependent plasticity have been largely unsuccessful (Kunkel et al., 2011).

In contrast, research on inhibitory synaptic plasticity is still sparse and its computational role somewhat speculative. Given the ubiquity of inhibition in the cortex (~20% of all cortical neurons are inhibitory, see Braitenberg and Schüz, 1998) and its undeniable role in shaping and stabilizing network dynamics and neuronal excitability, the possible functional implications of dynamic inhibition are of great interest, particularly when interacting with other forms of plasticity (see Kullmann et al., 2012; Vogels et al., 2013 for an overview). Progress in this endeavor is hindered by the complexity and diversity of inhibitory neurons, making it technically challenging to obtain reproducible experimental results and difficult to reconcile the available data. Nevertheless, recent evidence shows that, in cortical networks, GABAergic synapses targeting excitatory neurons are also sensitive to temporally coincident pre- and post-synaptic spiking (Holmgren and Zilberter, 2001; Woodin et al., 2003). To capture this phenomenon, Vogels et al. (2011) studied the computational effects of a simplified, symmetric inhibitory STDP rule in the establishment and robust maintenance of detailed balance between excitation and inhibition, both in a feedforward and in a recurrent configuration, showing that it allows the emergence of stimulus selectivity and memory. Apart from

these self-organized computational roles of inhibitory plasticity, the mechanism implemented by Vogels et al. (2011) has the interesting property of serving as a homeostatic mechanism. It maintains the post-synaptic firing rate under control by dynamically stabilizing the amount of inhibitory and excitatory drive that the excitatory neurons receive, which is particularly relevant in a situation where the excitatory drive is also dynamic, given the possible interdependence between excitatory and inhibitory synaptic plasticity (Wang and Maffei, 2014).

In this work, we consider the combined influence of timing-dependent synaptic plasticity rules operating on different synapse types and analyse its impact on the stability and diversity of global network dynamics as well as their computational implications for online processing of time-varying input streams. Following the general framework of reservoir computing (Lukoševičius and Jaeger, 2009), we explore the properties of information processing based on robust transient dynamics and analyse the influence of plasticity on the development and maintenance of dynamic stimulus representations, while maintaining a stable global dynamics. For that purpose, we implement numerical simulations of biologically realistic networks of leaky integrate-and-fire neurons which incorporate excitatory and inhibitory plasticity, combining well characterized phenomenological models of synaptic plasticity that take into account relevant physiological observations (Morrison et al., 2008; Vogels et al., 2011).

We begin by demonstrating that the balancing effects of these synaptic plasticity rules actively maintain an asynchronous irregular pattern of ongoing, background activity throughout the network, over a much broader range of parameters, compared with networks whose synapses are fixed and static. Furthermore, we establish a relation between dynamical states characterized by a regular, synchronous population firing pattern (which are mostly abolished by the action of plasticity) and a decreased capacity to process generic time-varying input streams, reinforcing the claim that the modulatory actions of plasticity have an important impact on computational performance.

Subsequently, we assess the features of population responses to sequentially occurring stimulus events, modeled as sudden spike bursts across variable numbers of afferent neurons providing a “wake-up” call (Sherman, 2001a) to stimulus-specific sub-populations via targeted, brief disruptions of balance. This stimulus is chosen primarily to be simple but disruptive rather than to model a particular cortical input, however in the following we refer to this type of stimulus as thalamic due to a similarity to the thalamic burst mode of firing (Ramcharan et al., 2000; Sherman, 2001b; Bruno and Sakmann, 2006). We show that the properties of ongoing, dynamic stimulus representations are naturally bound to the stimulus features and the strength of the disruption, but also to the characteristics of the ongoing network activity, that sets the dimensionality of the embedding space over which dynamic representations can unfold. By improving the stability of this ongoing activity and the robustness and reproducibility of response transients, plasticity is shown to benefit the quality of the representations, necessary for subsequent processing by downstream cortical regions.

In the final section of the results, we attempt to disentangle the roles played by the two analyzed plasticity rules in the

development of adequate stimulus representations, concluding that the quality of such representations is largely dependent on the decorrelating actions of inhibitory STDP, which results in the maintenance of AI-type activity across the network. The role played by excitatory STDP only provides a marginal advantage compared to static networks, which is an unexpected result leading us to draw some tentative conclusions and opening up a new set of questions to be addressed in future studies.

2. MATERIALS AND METHODS

In this section, we describe the equations used to model neuronal and synaptic dynamics, the characteristics of the input-dependent tasks, as well as the methods used for numerical simulations and data analysis. A summarized, tabular description of all the models and model parameters used throughout this study is available in the Supplementary Materials.

2.1. NETWORK

2.1.1. Neuron and synapse models

The networks we analyse are composed of $N = 10000$ leaky integrate-and-fire neurons (of which $N^E = 0.8N$ are excitatory and $N^I = 0.2N$ inhibitory) with fixed voltage threshold (Tuckwell, 1988) and conductance-based synapses (Koch, 2004), which capture a broad range of intrinsic properties shared by cortical neurons.

Synaptic interactions between the neurons are modeled as transient conductance fluctuations, so the sub-threshold membrane potential V_i of the i -th neuron ($i = 1, \dots, N_\alpha$) belonging to population α is given by:

$$C_m \frac{dV_i(t)}{dt} = g_{\text{leak}} (V_{\text{rest}} - V_i(t)) + I_i^{\alpha E}(t) + I_i^{\alpha I}(t) + I_i^{\alpha X}(t) \quad (1)$$

where $I_i^{\alpha Y}$ is the sum of all synaptic currents generated by the pre-synaptic neurons of neuron i in population Y . The total synaptic input current onto neuron i is thus the sum of the individual contributions of excitatory (glutamatergic, AMPA-type) synapses (E), inhibitory (GABAergic) synapses (I), and pre-synaptic sources from outside the network (X). The latter models cortical background activity and is assumed, for simplicity, to be non-selective and stochastic, with fixed rate ν_X . When applicable, some neurons belonging to discrete sub-populations receive additional, patterned external stimulation (see Section 2.4). The synaptic current induced in a post-synaptic neuron i in population α when a pre-synaptic neuron j in population β fires is given by:

$$I_{ij}(t) = g_{ij}(t)(V_\beta - V_i(t)) \quad (2)$$

where V_β is the equilibrium/reversal potential of the corresponding synapse. The time course of the synaptic conductance $g_{ij}(t)$ is modeled as an instantaneous rise triggered by each pre-synaptic spike, followed by an exponential decay:

$$\frac{dg_{ij}(t)}{dt} = -\frac{g_{ij}(t)}{\tau_\beta} + \bar{g}^\beta w_{ij}(t) \sum_{t_j} \delta(t - t_j - d) \quad (3)$$

where $\delta(\cdot)$ is the Dirac delta function, t_j are the spike times of the pre-synaptic neuron and d refers to the conduction delay, which

is set to be constant and equal for all synapses, with the value of 1.5 ms.

The peak amplitude of the conductance transient, which determines the “strength” of the synapse, is the product of a constant scaling factor \bar{g}^β , whose value depends on the synapse type and is used to set the scale of the synaptic conductance, and a dimensionless variable w_{ij} , assumed to be dynamic in synapses subjected to activity-dependent adaptation (see Section 2.2) and whose initial value is drawn from a Gaussian distribution with mean μ^β and standard deviation σ^β , which we set to 1 and 0.25, respectively, leading to narrowly distributed initial peak conductances centered around \bar{g} for every synapse type. All synaptic events originating from outside the network are assumed to be excitatory, with the same reversal potential, peak conductance and time constant as recurrent excitatory synapses, i.e., $V_X = V_E$, $\bar{g}^X = \bar{g}^E$ and $\tau_X = \tau_E$.

Following Kumar et al. (2008b), we quantify the effective balance between excitation and inhibition as the approximate ratio of total charges induced at rest:

$$g = \frac{\langle g^{\alpha I} \rangle \tau_I |V_{\text{rest}} - V_I|}{\langle g^{\alpha E} \rangle \tau_E |V_{\text{rest}} - V_E|} \quad (4)$$

with $\langle g^{\alpha I} \rangle = \mu^{I\bar{g}^I}$ and $\langle g^{\alpha E} \rangle = \mu^{E\bar{g}^E}$. Under these conditions, and with all other synaptic parameters fixed and set as described below in Section 2.3.1, we determine the initial value of g to be 0.29γ where γ is the ratio of absolute peak conductances.

The parameter values of the neurons are homogeneous across neuron types and are kept fixed throughout. They were chosen for their biological plausibility and consistency with the experimental literature and previous modeling work (e.g., Compte et al., 2000; Meffin et al., 2004; Kumar et al., 2008b; Vogels et al., 2011; Yger et al., 2011). A complete description of the parameters and their values can be found in the Supplementary Materials.

2.1.2. Network architecture

All the network neurons are laid on the integer points of a 2-dimensional 100×100 regular grid lattice, with periodic boundary conditions and are sparsely and randomly connected. The probability of connection between a target neuron in population α and a source neuron in population β is set to 0.1 for $\alpha\beta \in \{EE, EI, IE, II\}$, such that, on average, each neuron in the network receives a total of $K^E = 0.1 \cdot N^E$ excitatory and $K^I = 0.1 \cdot N^I$ inhibitory, randomly chosen, synaptic inputs from the local network, along with K^X synaptic inputs from outside the network. It is generally assumed that the number of background synapses from external cortical sources, comprising patchy long-range input from the same cortical area as well as input from distant cortical areas (Braitenberg and Schüz, 1998; Kumar et al., 2008a; Kremkow et al., 2010) lies in the same range as the number of local, recurrent excitatory connections, so we set $K^X = K^E$ (Brunel, 2000).

This network structure is relevant mostly for the purpose of visualization when patterned stimuli are delivered to specific, spatially clustered neuronal populations, given that no additional spatial constraints are imposed on the connectivity structure. Furthermore, in networks shaped by plasticity, the connectivity

structure remains unaltered, i.e., plasticity modifies the strength of existing connections only and does not create new synapses or destroy existing ones.

2.2. SYNAPTIC PLASTICITY

In the following, we assume that synapses targeting inhibitory neurons (II and IE) are static, whereas synapses targeting excitatory neurons (EI and EE) are sensitive to pre- and post-synaptic spike times. Although there is increasing evidence for the existence of timing-dependent adaptation mechanisms in synapses within inhibitory populations (Lamsa et al., 2010) and from excitatory to inhibitory neurons (Lu et al., 2007) (II, IE, respectively), their precise mechanisms are highly dependent on the target neuron type, which constitutes an added source of heterogeneity and complexity. Furthermore, in most of our analysis, we assume that the most relevant activity is that which can be propagated to downstream regions, conveying the relevant information for additional processing. For that reason, we focus our attention on the dynamics of the excitatory population, given the known locality of inhibitory connections.

In all experiments on networks incorporating plasticity, synapses are continually plastic. For convenience, notation brevity and consistency, plasticity modifications are not applied directly to the synaptic “strength” (i.e., to the peak amplitude of the conductance transient), but instead, to the dimensionless variable w_{ij} which is subsequently rescaled by a constant factor (see Equation 3).

Both types of plasticity used can be expressed in terms of a synaptic trace variable defined for each neuron that is incremented by 1 at each spike and decreases exponentially in between spikes:

$$\frac{dx_i(t)}{dt} = -\frac{x_i}{\tau_p} + \sum_{t_j} \delta(t - t_j) \quad (5)$$

2.2.1. Excitatory STDP

We term the learning rule applied to the recurrent synapses among the excitatory population excitatory spike-timing dependent plasticity (eSTDP) and adopt the formalism proposed in van Rossum et al. (2000):

$$\Delta w^{EE} = \begin{cases} \lambda \exp(-|\Delta t|/\tau_p), & \text{if } \Delta t > 0 \\ \alpha_{ep} \lambda w^{EE} \exp(-|\Delta t|/\tau_p), & \text{if } \Delta t \leq 0 \end{cases} \quad (6)$$

where $|\Delta t| = t_i^f - t_j^f$ is the absolute difference between a specific pair of spikes of pre-synaptic neuron j and post-synaptic neuron i and τ_p is the time window for potentiation and depression, which we set to be equal ($\tau_p = 20$ ms, following experimental data obtained by Bi and Poo, 1998). The parameter λ sets the magnitude of individual modifications (i.e., the learning rate) and α_{ep} determines the asymmetry between the amount of potentiating and depressing changes. The update rule can be re-written in a differential form that depends on the synaptic trace variables:

$$\frac{dw_{ij}}{dt} = \alpha_{ep} \lambda w^{EE} x_i(t) \delta(t - t_j^f) + \lambda x_j(t) \delta(t - t_i^f) \quad (7)$$

with all propagation delays considered to be dendritic, i.e., spike times are taken at the synapse and no autapses are allowed. The pre- and post-synaptic spikes are paired in an all-to-all scheme (see Morrison et al., 2008).

This “hybrid” learning rule is additive for potentiation and multiplicative for depression, thus incorporating some of the most relevant experimental observations (Bi and Poo, 1998). Importantly, it gives rise to unimodal weight distributions similar to those observed experimentally in the presence of uncorrelated input, while retaining the ability to develop multimodal distributions depending on the input correlation structure.

2.2.2. Inhibitory STDP

We apply the inhibitory spike-timing dependent plasticity (iSTDP) proposed in Vogels et al. (2011) to the weights of synapses between inhibitory and excitatory neurons. The general premise is that pre- and post-synaptic firing that occurs within the relevant coincidence time window (τ_p) should always lead to synaptic potentiation, regardless of the temporal ordering of the spike pair, whereas isolated pre-synaptic spikes lead to synaptic depression.

This rule can be given in terms of the synaptic trace variables (Equation 5) as:

$$\frac{dw_{ij}}{dt} = \eta (x_i(t) - \alpha_{ip}) \delta(t - t_j^f) + \eta x_j(t) \delta(t - t_i^f) \quad (8)$$

where η is a constant learning rate. The parameter α_{ip} sets the amount of synaptic depression upon a pre-synaptic spike and has the value $\alpha_{ip} = 2\rho_0\tau_p$, where ρ_0 is a constant which serves the homeostatic purpose of stabilizing the post-synaptic neuron's firing rate (for further details, see Vogels et al., 2011).

2.3. CONSTRAINING MODEL PARAMETERS

2.3.1. Initial synaptic strengths

Synaptic strengths are adjusted such that the ongoing, background network activity, prior to any patterned input stimulation, mimics the statistics of cortical background activity: inhibition dominated (i.e., $\bar{g}^I/\bar{g}^E > 1$), low rate (1–20 spikes/s), irregular single neuron firing ($CV_{ISI} \simeq 1$) and asynchronous population activity (low average pairwise correlations) (Brunel, 2000; Destexhe et al., 2001; Meffin et al., 2004) (see Section 3.1 and Figure 2).

For that purpose, we started by setting the constant background firing to a low rate of $\nu_X = 5$ spikes/s and tuned \bar{g}^E and \bar{g}^I to obtain the desired activity statistics. This resulted in $\bar{g}^E = 1.8$ nS for excitatory synapses, which leads to an EPSP amplitude of around 1.46 mV at rest and $\bar{g}^I = \gamma \bar{g}^E = 21.6$ nS for inhibitory synapses, leading to IPSP amplitudes of around −1.14 mV at rest, where $\gamma = 12$ determines the absolute strength of inhibition relative to excitation. This parameter combination results in self-consistent firing rates, whereby each neuron fires with a rate approximately equal to the population rate and to the external rate ($\nu_i = \nu_{net} = \nu_X$). It also provides a reasonable match to experimental data: the ratio of IPSP to EPSP amplitudes at rest is $\simeq 0.78$ which lies within the range measured in the cortex (Matsumura et al., 1996). Additionally, the mean coefficient of variation of the single neurons' interspike intervals (CV_{ISI})

is around 1.1 and their mean membrane potential approximately -62 mV (Destexhe et al., 2003).

2.3.2. Plasticity

We take a similar approach to tune the parameters of the plasticity rules. We wish to obtain similar population dynamics in networks whose synapses are subjected to adaptation and in networks whose synapses are fixed, when driven by uncorrelated background input. To constrain the parameter ranges, we assume the plasticity rules are independent and determine their parameters individually.

In order to allow the effects of the two plasticity mechanisms to balance, we set the rate of synaptic modifications to be the same in each case, i.e., $\lambda = \eta = 0.01$. The only free parameter of iSTDP left to tune is the target firing rate, ρ_0 . As mentioned by Vogels et al. (2011), it is quite convenient to control the firing rate with a single parameter, since we can tune the network to a desired operating point by simply setting the desired rate. For consistency, we set $\rho_0 = 5$ spikes/s, to match the average population firing rate displayed by a static network.

To determine the remaining free parameter of the eSTDP rule, α_{ep} (the asymmetry parameter) we look at the equilibrium dynamics of the weight update. We wish to obtain an equilibrium weight distribution whose mean is equal to the mean static weight, i.e., 1.8 nS and an equilibrium firing rate close to the population rate in the static case. Note that, when iSTDP is present, a specific firing rate can be readily achieved as described above. In the absence of iSTDP, the desired features are achieved for a value $\alpha_{ep} \simeq 0.92$ (data not shown).

2.4. INPUT STIMULI

A sequence of stimuli was delivered to specific sub-populations (Figure 1C), in order to perturb the stable background activity. We call these input stimuli “thalamic,” to differentiate them from the unspecific, background stimulation, referred in Section 2.1.1, however an accurate modeling of thalamic activity is not aimed at in this study. Each stimulus was assigned an arbitrary abstract label and subsequently converted into a set of spike trains, according to the following process: a stimulus sequence ($u = \sigma_1, \sigma_2, \dots, \sigma_T$) of a predefined length (T) was built by randomly drawing σ_i from a set of k different stimuli $S_j \in S$, with $1 \leq j \leq k$. Each successive σ_i was then converted into a k -dimensional binary vector \hat{u} , where $\hat{u}_n[j] = 1$ if $u_n = S_j$, for $n = 1, \dots, T$. This binary representation was also used as the target output to train the readout units in a classification task (see Section 2.5.3.1). From the resulting input streams (\hat{u}_n), k independent signals were generated, according to:

$$s_k(t) = \frac{1}{\sigma_u} (\hat{u}_n[k] \times \delta(t - n\Delta)) * g \quad (9)$$

where σ_u determines the signal’s peak amplitude and Δ corresponds to the period of the input sequence, i.e., the duration of each stimulus presentation plus the inter-stimulus interval, assuming regularity of input unit length. The function g is a bi-exponential kernel:

$$g(s) = \exp(-s/\tau_r) - \exp(-s/\tau_d) \quad (10)$$

with rise time $\tau_r = 50$ ms and decay time $\tau_d = 150$ ms (see Figure 1A for a schematic depiction of this input generation process). These independent signals were used to determine the time-dependent firing rates of inhomogeneous Poisson processes, in order to generate N_{aff} input spike trains for each signal s_k . The peak amplitude of the signal thus corresponds to the peak firing rate of a spike burst. Finally, a constant value of 2 spikes/s background activity along with a small amount of Gaussian white noise was added to the signals $s_k(t)$. The resulting input structure is depicted in Figure 1B.

For the experiments in Section 3.2, we use $k = 3$, randomly drawn, independent stimulus classes. In total, $T = 3300$ stimulus samples (comprising 1100 samples of each stimulus class) are presented to the networks. The first 300 samples and corresponding network responses are considered to represent an initial “entraining” period and so are discarded from the analysis. The duration of each stimulus presentation is fixed and set to 200 ms, followed by a 100 ms-long inter-stimulus interval, resulting in a total analyzed simulation time of 900 s. A similar protocol is used in the experiments in Section 3.3, the difference being that the value of k is varied. In such cases we use $1100k$ stimulus samples and discard the first $100k$ samples.

2.5. DATA ANALYSIS

2.5.1. Global network dynamics

In order to properly characterize and compare the dynamic network states in different conditions, we need to adequately quantify the population activity. This analysis focuses on 3 main properties, the average population firing rate, the degree of synchrony and the degree of irregularity of network states.

2.5.1.1. Irregularity. The degree of irregularity of population spiking activity is determined by the coefficient of variation of the interspike intervals (ISI) of each neuron’s spike trains, averaged across all neurons in the population:

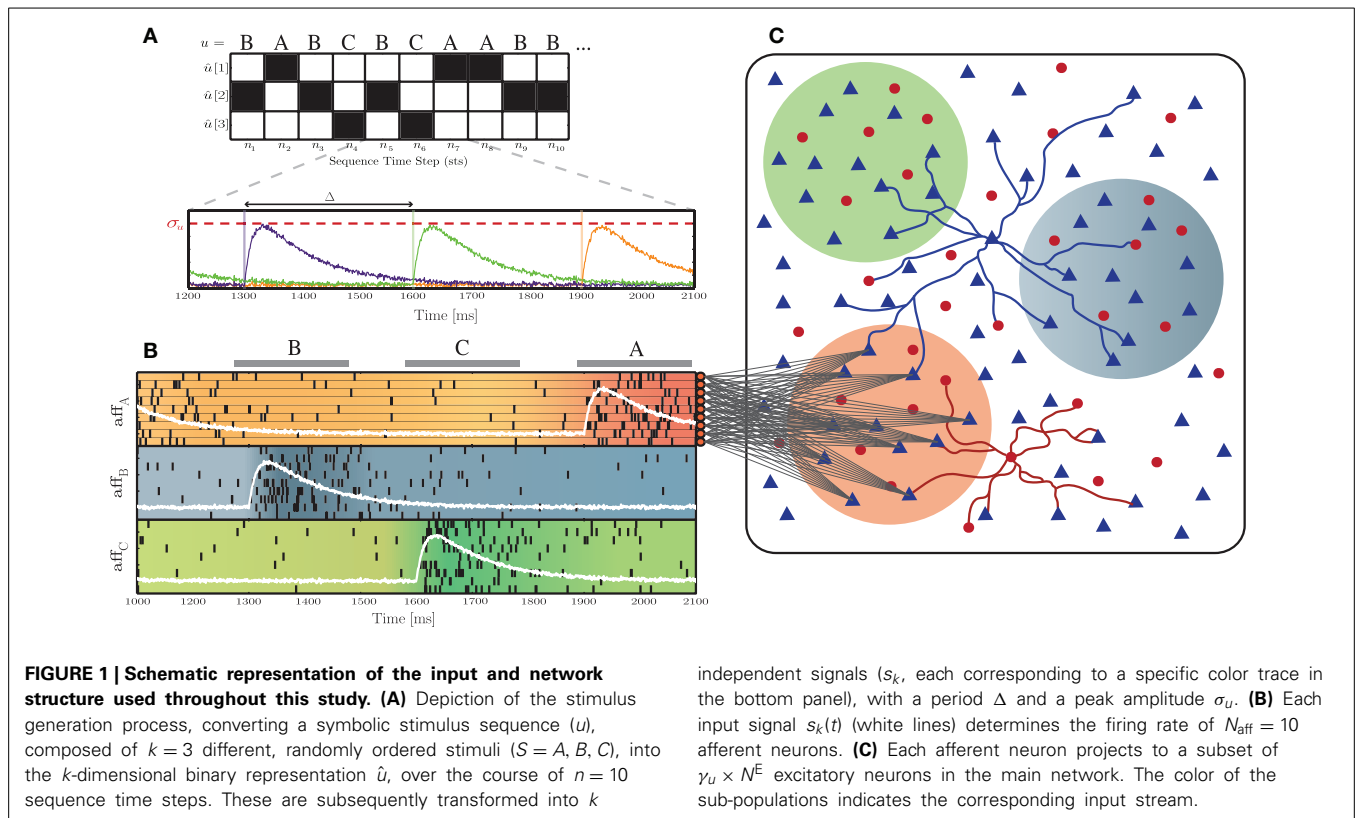
$$CV_{\text{ISI}} = \left\langle \frac{\sigma_i^{\text{ISI}}}{\mu_i^{\text{ISI}}} \right\rangle \quad (11)$$

where $\langle \cdot \rangle$ denotes the average over all neurons, μ_i^{ISI} and σ_i^{ISI} denote the mean and standard deviation of the ISIs of neuron i . The CV_{ISI} provides a good measure of spike train variability over time scales on the order of the mean ISI. An irregularly spiking neuronal population will have CV_{ISI} close to 1 (a value of exactly 1 corresponds to Poissonian firing). CV_{ISI} values close to 0 indicate a regular spiking pattern, whereas values much larger than 1 indicate a bursting firing profile.

2.5.1.2. Synchrony. The degree of synchrony is quantified by the average pairwise correlation coefficient over 500 randomly sampled, disjoint, neuronal pairs (see, e.g., Kumar et al., 2008b):

$$CC_{ij} = \left\langle \frac{\text{Cov}(C_i, C_j)}{\sqrt{\text{Var}(C_i)\text{Var}(C_j)}} \right\rangle \quad (12)$$

where $\langle \cdot \rangle$ denotes the average over all pairs, C_i and C_j represent the spike counts of neurons i and j , computed by



counting the spikes occurring within successive time bins of width 1 ms.

2.5.1.3. AI score. To allow a simpler visualization of the regions encompassing AI-type activity (see Section 3.1), we introduce an additional metric that summarizes the main statistical descriptors and provides a graded measure of the *AI-ness* of the spiking activity, depicted in **Figures 3E,I**. This metric relies on somewhat arbitrary criteria, based on general assumptions, which we establish as the percentage of neurons in the population that fire with a rate ≤ 20 spikes/s and whose $CV_{\text{ISI}} \in [0.8, 1.5]$, in conditions where the average pairwise $CC \leq 0.05$. Objectively, all the parameter combinations that lie within this range reflect AI-type activity, but may correspond to different sub-types of this regime (Ostojic, 2014), with different computational properties. The highest AI scores, in our analysis, reflect the “classical” states, of homogeneous average firing rates.

2.5.2. Global computational power

We adopt the methods introduced in Maass et al. (2005), Legenstein and Maass (2007) to evaluate the generic computational power of neuronal microcircuits, regardless of the precise nature of the circuit. The general premise underlying this approach is that sufficiently different input streams should cause different internal states and hence lead to different, linearly separable outputs.

For this purpose, we can use a much simpler stimulus than previously described in Section 2.4. Consider the microcircuit C , generated with a specific set of parameters, and stimulated

by one of a set of $T = 500$ different input streams (fixed spike patterns), each composed of 4 independent Poisson spike trains, at a rate of 20 spikes/s and a duration of 200 ms. The temporal evolution of the system in response to this input pattern is analyzed and stored in a response matrix, obtained by convolving each neuron’s spike train with an exponential kernel, with 30 ms time constant and temporal resolution equal to the simulation time step, i.e., 0.1 ms (Section 2.6), in order to capture the effect of each spike on the membrane properties of a readout neuron receiving it. This response matrix is then sampled at time point $t_0 = 200$ ms, resulting in the N -dimensional vector $x_u(t_0)$ which contains all the neurons’ responses to input pattern u , or the circuit state (Maass et al., 2002). The procedure is then repeated for all the spike templates, leading to the formation of the state matrix $X \in \mathbb{R}^{N \times T}$. The rank r of the matrix X , calculated by singular value decomposition, corresponds to the number of linearly independent columns of X , i.e., the number of inputs that are mapped into linearly independent circuit states, thus providing a quantitative descriptor of computational performance, or *kernel quality*. If $r \leq k$, a linear readout should be able to separate r classes of inputs (Maass et al., 2005; Legenstein and Maass, 2007).

2.5.3. Stimulus representation

To assess the quality of the input-state mappings, in relation to the underlying dynamical states that the networks achieve under different conditions, we use the following metrics:

2.5.3.1. Readout classification. The network responses to a stimulus sequence of length T are assembled in a state matrix $X \in$

$\mathbb{R}^{N \times T}$, as described in Section 2.5.2, where each column represents the network state in response to one stimulus. These T stimulus-response pairs are subsequently split into train and test samples, with $T_{\text{tr}} = 0.8T$ and $T_{\text{te}} = 0.2T$. A set of k linear readout units are then trained to classify which pattern was presented to the network, using the binary signal \hat{u} as the supervisory signal. We wish to map each N -dimensional state pattern ($X(t)$) to the corresponding input that triggered it (\hat{u}_t), by minimizing the quadratic error $E(\hat{U}, W^{\text{out}}X)$. The synaptic weights from the main network to the readout units (W^{out}) are obtained by ridge regression, i.e., by solving:

$$W^{\text{out}} = \hat{U}X^T(XX^T + \alpha^2\mathbb{I})^{-1} \quad (13)$$

where \hat{U} is a $k \times T_{\text{tr}}$ matrix combining all the binary target patterns, X is the $N \times T_{\text{tr}}$ state matrix, \mathbb{I} is the identity matrix and α is a regularization factor, given by the least squares norm and optimized by 5-fold cross-validation on the training data. In principle, any linear regression method would be applicable, but we chose ridge regularization because of the penalty imposed on the size of the coefficients. It is desirable that the average vector norm of W^{out} be kept small so that the output accurately reflects features of the state space, instead of relying on disproportionate amplification of certain dimensions.

The obtained synaptic weights W^{out} are then used to classify the state responses to the test sequence. Average classification performance is obtained by applying winner-takes-all on the readout output $y(t)$ to determine the label assignments and subsequently quantifying the fraction of correctly classified patterns. To obtain a more fine-grained measurement, we also quantify the performance in classifying each of the individual stimulus patterns using the raw readout output and correlating it with the target binary values, using point-biserial correlation coefficient, which is a suitable statistic to estimate the relationship between a dichotomous variable (target values) and a continuous variable (readout output).

2.5.3.2. Dimensionality reduction and visualization.

Throughout this study, we apply various different methods of dimensionality reduction, which we briefly outline below. It is worth noting that, while for most depictions we chose one particular method, the only criteria that justified this choice was adequate visualization. In every case, several different methods were applied and these results were only included and further discussed if they were consistent across different methods.

These methods are applied to visualize the underlying spatial arrangement of the network states in response to each stimulus pattern (finding structure in the state matrix) and the unfolding trajectory of network states within the time course of single responses to a stimulus, both of which assuming the desirable condition that the input-driven stimulus responses lie within distinct sub-spaces in the N -dimensional state space. In the first case, the low-pass filtered population responses to each stimulus are sampled at time t_0 after stimulus onset (typically $t_0 = 200$ ms, i.e., at the end of each stimulus presentation) and collected in the state matrix $X \in \mathbb{R}^{N \times T}$, following the description in the previous section (Section 2.5.3.1, Readout Classification).

The methods of dimensionality reduction are then applied to X (N features and T samples). We tested several different algorithms for this purpose, namely principal component analysis (PCA), linear discriminant analysis (LDA), spectral embedding and isomap embedding. The first two are based on finding low-dimensional embeddings of the data points, using linear projections of the variables that best explain the original, high-dimensional data, either by identifying the directions in feature space that capture most variance in the data (the top eigenvectors of the data covariance matrix or principal components) or by identifying attributes that account for the most variance between labeled classes (LDA is a supervised method). Spectral and isomap embedding methods, on the other hand, are based on non-linear projections of the data seeking low-dimensional representations that maintain the relative distances (Euclidean distances, in our case) between data points. Spectral embedding (also known as Laplacian eigenmaps) constructs a weighted graph representing the data, using an adjacency matrix based on the pairwise distances between data points. The embedding is subsequently obtained by partial eigenvalue (spectral) decomposition of the graph Laplacian (see e.g., Ng et al., 2001; Belkin and Niyogi, 2003 for a more thorough explanation). Isomap embedding also relies on partial eigenvalue decomposition, but applied to a matrix representing the shortest path lengths between a data point and its nearest neighbors (Tenenbaum et al., 2000).

In the second case, we want to reduce the dimensionality of the data along the time course of single responses to a stimulus. For that purpose, we analyse the full response matrix R , using the low-pass filtered population responses, with $R \in \mathbb{R}^{N \times D}$, where D corresponds to the duration of the stimulus presentation divided by the response time resolution (0.1 ms). We apply techniques that have been previously used to analyse neural data, with a similar goal in mind, namely principal component analysis (PCA) and locally linear embedding (LLE) (see Churchland et al., 2007 for an overview).

2.5.3.3. Input/output correlations. To determine to which degree the activity of each input-specific sub-population (see Sections 2.1.2 and 2.4) becomes specialized to a particular input pattern, we compute firing rate histograms of the output of each population $r_\alpha(t)$ over many sequential trials with each of the input patterns and determine the time-averaged firing rate of these responses \bar{r}_α . These histograms are then correlated with the input signals (s_k) to obtain the correlation coefficient of signal k with population α :

$$C_k^\alpha = \frac{\langle (s_k(t) - \bar{s}_k)(r_\alpha(t) - \bar{r}_\alpha) \rangle}{\sigma_{s_k} \sigma_{r_\alpha}} \quad (14)$$

This procedure allows us to determine the specialization of each population and the impact that each input signal has on each population.

2.6. NUMERICAL SIMULATIONS

All simulations were performed using the NEST simulating environment (Gewaltig and Diesmann, 2007) with an integration

resolution of 0.1 ms. Due to the large memory and computing demands, simulations were carried out on large, parallel computing clusters, using the parallelized kernel of NEST (Morrison et al., 2005). All subsequent calculations and data analysis were performed in Python, using the NumPy and SciPy libraries, as well as the Scikit-learn toolbox (Pedregosa et al., 2011).

3. RESULTS

3.1. IMPACT OF PLASTICITY ON NETWORK DYNAMICS

Numerous *in vivo* recordings in awake, behaving animals have revealed the prevalence of highly irregular and seemingly noisy firing patterns in neocortical circuits (Softky and Koch, 1993; Stiefel et al., 2013). Sub-threshold fluctuations of the neurons' membrane potentials lead to irregularly timed, low frequency spiking (on the order of 1–20 spikes/s) (Gerstein and Mandelbrot, 1964; Destexhe et al., 2003), whereas at the population level, activity is characterized by a low degree of synchrony, with small pairwise correlations between spike trains (Abeles, 1991; Vaadia and Aertsen, 1992; Shadlen and Newsome, 1994). Collectively, these characteristic features of neural activity are generally termed “Asynchronous Irregular” (AI) states and are assumed to constitute the “ground state” of ongoing cortical activity.

The mechanisms underlying this activity regime have been the subject of intense investigation and are known to rely mainly on the balance of excitation and inhibition. Analytical studies of random recurrent networks of IF neurons with static current-based synapses have shown that these systems can display a rich set of behaviors, depending on the intensity of external stimulation and the relative strength of excitation and inhibition (van Vreeswijk and Sompolinsky, 1998; Brunel, 2000; Kumar et al., 2008b). In the cortex, this relationship is highly dynamic and the balance required needs to be actively maintained and tuned to allow the network to operate in suitable regimes. Thus, the activity regimes exhibited by networks with plastic synapses is of particular interest. Recently, Vogels et al. (2011) investigated the transition of a network with iSTDP from non-AI to AI regimes in dependence on the learning rate of the inhibitory plasticity and the strength of excitatory synapses. In this section, we explore the impact of dynamic excitatory and inhibitory synapses on the ongoing activity, by systematically varying the same control parameters investigated in earlier studies on static networks, namely the external input rate ν_X and the inhibitory-excitatory balance g , which can be set via the ratio of absolute peak conductances γ as described in Section 2.1.1. In plastic networks g evolves throughout the simulation as a result of synaptic changes (see Supplementary Materials), so for ease of comparison with the static case we consider the network activity as a function of its initial value. **Figure 2** shows the behavior of an example network as described in Section 2.1, with parameters set according to Section 2.3.1, leading to AI-type activity (**Figure 2A**). As desired and akin to its biological counterpart, the single neuron's spiking activity is highly irregular, with membrane potential hovering slightly below threshold (**Figures 2B,F**). Furthermore, the excitatory and inhibitory synaptic currents impinging onto this neuron are closely balanced (**Figure 2C**). This activity pattern is consistently conserved across the population, leading

to the distributions observed in **Figures 2D,E,G** (these statistical descriptors were computed as described in Section 2.5.1). The synapses are subjected to modifications due to the ongoing activity and the stochastic input, leading to more narrowly distributed synaptic weights compared to the initial condition (**Figure 2H**). However, as the network maintains its AI-type activity during the evolution of synaptic strengths, the distributions of the statistical measures shown in **Figures 2B–F** remain essentially indistinguishable throughout the simulation period.

Our analysis in this section focuses on the most important statistical descriptors of population activity and on the generic computation capacity of the networks. In the first case, the measures are computed as described in Section 2.5.1 over a period of 20 s following a long initial equilibration phase. The network receives no specific input during this analysis, just the external input rate described above. Here, results are obtained from a single realization of each network configuration due to the computational intensity of the parameter scan. In the second case, the network receives independent Poissonian spike trains (as described in Section 2.5.2). Results are averaged over 10 realizations of each network configuration.

The main results of our analysis over a broad parameter range are depicted in **Figure 3** for static (A–E) and plastic (F–J) networks. The presence of plasticity strongly influences network activity. In accordance with the results presented in Kumar et al. (2008b), the static networks exhibited the asynchronous irregular (AI), fast and slow synchronous regular (SR_F , SR_S) and synchronous irregular (SI) regimes, but no asynchronous regular (AR) regimes were observed. In the plastic networks, only SI and AI regimes are observed, indicating that plasticity abolishes regular spiking activity except for a small region, where the external stimulus is weakest ($\nu_X = 0.5$ spikes/s). Static networks with very weak inhibition ($g < 1$) have very high average firing rates, whereas plastic networks have low firing rates for almost all configurations. These results demonstrate that the presence of balanced plasticity makes the existence of the low rate AI dynamical state much more robust in comparison to static networks. The smooth profiles of the measures indicates that a single realization of the network configuration is sufficient to capture them.

We additionally measured the generic computation capacity of these networks, i.e., their ability to separate similar time-varying input streams in the form of fixed spike templates (see Section 2.5.2). Our results reveal that all regimes of the static network have a high generic computation capacity except SR_F . This is demonstrated by the low rank in **Figure 3E** for network configurations in the SR regime identified in **Figure 3D**. In this regime, the dominant excitation and consequent excessive firing hinders a proper stimulus separation. For all other regimes, the rank is maximal, indicating that all the columns of the state matrix are linearly separable, allowing a fine discrimination of input stimuli. As plastic networks abolish the pathological SR regimes, every configuration of parameters leads to maximally separable circuit states (indicated by maximal ranks in **Figure 3J**), thus the presence of plasticity also increases the robustness of generic computation capacity in comparison to static networks.

Based on these results, we were able to select a suitable network configuration for our investigation of the capacity of static

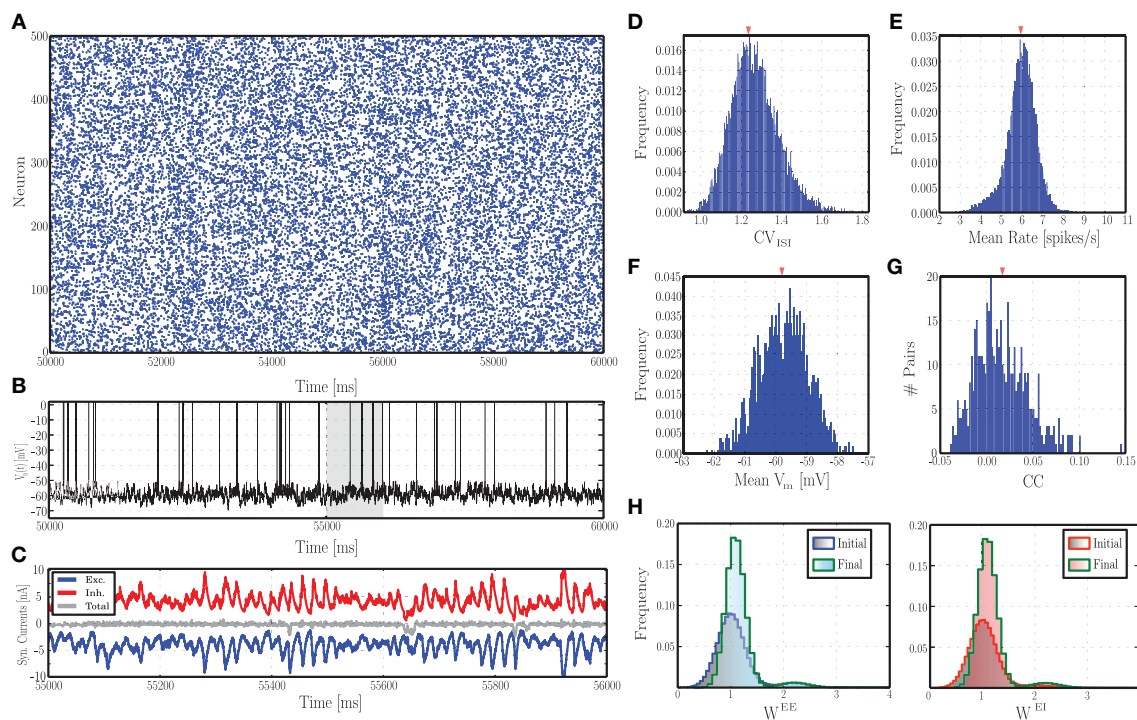


FIGURE 2 | Characteristics of ongoing activity in a network with excitatory and inhibitory STDP. (A) Raster plot depicting the spiking activity recorded in a subset of 500 randomly chosen excitatory neurons, for a period of 10 s, following an initial equilibration phase of 50 s. **(B)** Example of a randomly chosen single neuron's membrane potential during the same time interval. **(C)** Total excitatory (blue) and inhibitory (red) synaptic currents into the neuron whose membrane dynamics is shown above, during a time period of 2 s (highlighted in gray in **B**). The gray line in the middle corresponds to the total current. **(D–G)**

Distributions of the most important descriptors of population activity, namely coefficients of variation of the inter-spike intervals **(D)**, average firing rates **(E)**, mean membrane potentials **(F)**, and pair-wise correlation coefficients **(G)**, computed over 500 pairs. These distributions were obtained from the activity of the entire population (not just the neurons depicted in **A**), recorded over a period of 20 s. **(H)** Initial and final synaptic weight distributions for excitatory (left) and inhibitory (right) weights. Note that these distributions refer to the dimensionless variable w and not the actual synaptic conductances (see Equation 3).

and plastic networks to extract information from structured input (described in Section 2.4), which comprises the main focus of our study. The selected configuration (marked with a star in all panels of **Figure 3**) produces activity with a high AI-score for both types of network. The parameters are $v_X = 5$ spikes/s and $g \simeq 0.29\gamma$, which for $\gamma = 12$ leads to $g \simeq 3.479$ (see also Section 2.3.1).

3.2. STIMULUS DISCRIMINATION

The ongoing network dynamics, when perturbed by an external stimulus pattern, performs a non-linear temporal expansion of its input, projecting it in a high-dimensional state-space as a complex, transient activity pattern (Rabinovich et al., 2008; Lukoševicius and Jaeger, 2009; Maass, 2010). In the following, we investigate whether balanced plasticity allows the network to counteract the effects of stimulation on the local E/I balance and develop stable stimulus representations, making the trajectories of network states more robust and easier to decode while maintaining suitable ongoing population activity.

3.2.1. Effective discrimination with different input features

To better understand the dynamics underlying stimulus representation, we first analyse the absolute difference between static and plastic networks in terms of the performance obtained by readout

neurons trained to classify the responses (**Figure 4A**). To do so, we use input sequences as described in Section 2.4, composed of $k = 3$ randomly ordered and sequentially presented stimuli.

The results show that plastic networks are not invariably better sources of classification information than static networks. When the peak rate of the input burst signals is low ($\sigma_u = 10$ spikes/s), the main differentiating factor is the number of afferent neurons that synapse onto each input population. Both static and plastic networks perform much better in the presence of a stronger input ($N_{\text{aff}} \geq 100$) and when these input neurons connect to a larger sub-population. All other input parameters lead to insufficient discrimination, which is reflected in a readout classification performance at chance level for both network types (see **Figure 4A** and Supplementary Materials).

Increasing the input burst rate allows static networks to outperform plastic ones in conditions where the number of afferent neurons is high ($N_{\text{aff}} = 500$). Conversely, in conditions where the number of afferents is very low ($N_{\text{aff}} = 5$), the input is not strong enough to create a discernible response and both networks perform at a level barely above chance. This performance improves slightly as the number of receiving neurons increases. For intermediate values of afferent neurons, both networks display significantly discriminative responses, with the difference

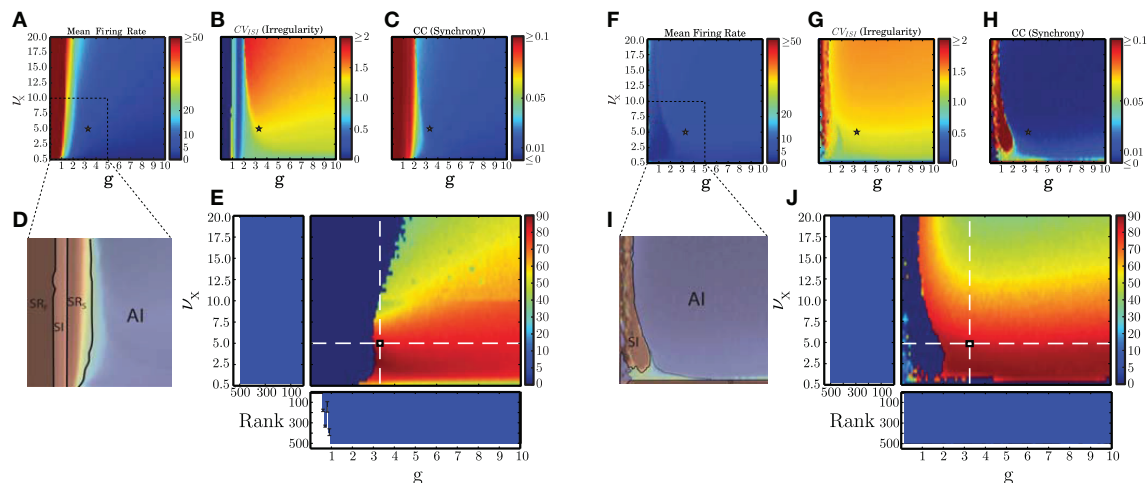


FIGURE 3 | Dynamical network states in static and plastic networks. (A–C) Main properties of static network dynamics as functions of the control parameters ν_x (rate of external Poissonian drive) and g (effective excitation/inhibition balance): average firing rates, irregularity, and synchrony. (D) Schematic depiction of the different network states observed in static networks. This figure was obtained by overlaying (B,C) in the depicted range, which corresponds to the region where the most significant state transitions occur. (E) AI-score expressed as the percentage of neurons in the population that fire with a rate ≤ 20

spikes/s and whose $CV_{ISI} \in [0.8, 1.5]$, in conditions where the average pairwise $CC \leq 0.05$ (see Section 2.5.1). The histograms show the average results of the kernel quality analysis (Section 2.5.2, where Rank refers to the number of linearly separable columns of the state matrix in response to 500 different stimulus templates) along the two main axes (highlighted by the white dashed lines) over 10 analyses per condition. Note that the parameter combinations marked in (A–C) with a small star correspond to the point where these two main axes intersect. (F–J) As in (A–E) but for plastic networks.

favoring mainly plastic networks, particularly if the size of the stimulated population is large ($\gamma_u \geq 0.1$).

In the following sections, we carry out further analysis to uncover the reasons why in some cases plasticity increases the network performance and in other cases decreases it. In order to do this, we isolate three input conditions which lead to different comparative performances of the plastic and static networks. In one configuration, marked by a gray star in **Figure 4A** and examined in Section 3.2.2, plastic and static networks performed flawlessly. In the configuration marked by a white star, there is a clear and significant advantage of having plastic synapses. This is examined Section 3.2.3. Finally, in the configuration marked by a black star, plastic synapses confer a significant disadvantage, which we analyse in Section 3.2.4.

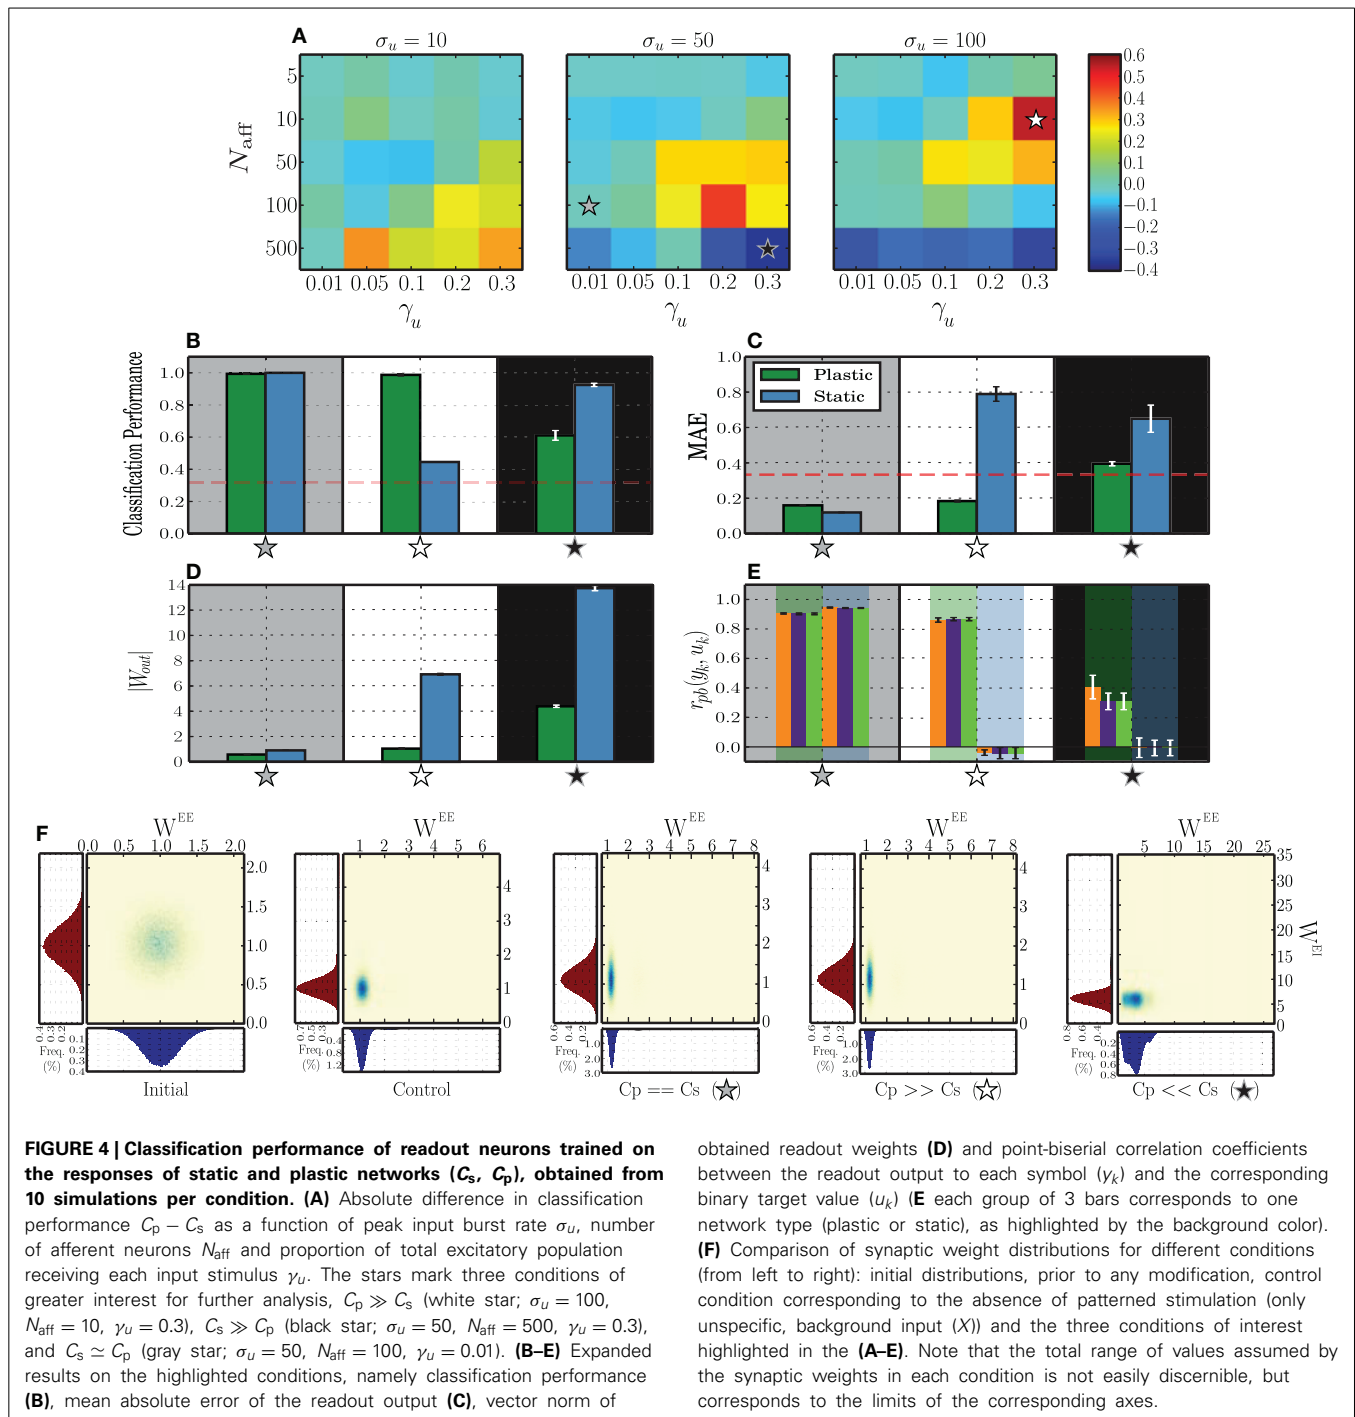
3.2.2. Specialized population responses

We start by analysing the condition where both network types exhibited a high capacity to discriminate the stimulus patterns (configuration marked with a gray star in **Figure 4A**). Each input signal consists of a relatively large number of afferent neurons ($N_{\text{aff}} = 100$), whose peak rate is at an intermediate value (50 spikes/s) and whose target population is very concise, consisting of only 80 excitatory neurons ($0.01 \times N^E$). The stimulus representations developed by both network types are highly specific, allowing the readout to classify with near perfect accuracy and low error (**Figures 4B,C**). Furthermore, the solutions found by the regression algorithm are highly stable and accurately reflect the population activity (low $|W^{\text{out}}|$, see **Figure 4D**) and each readout output y_k is highly correlated with its corresponding target \hat{u}_k (**Figure 4E**).

A closer analysis of the network activity under these conditions provides a straightforward justification for the high discriminability of the responses. As can be seen in **Figure 5A**, upon receiving each stimulus pattern, the responsive sub-populations exhibit a clearly discernible activity that stands out from the background population, with a firing rate 30–40 spikes/s higher than that of the background, unstimulated neurons. This is less obvious in plastic networks, because the inhibitory plasticity rapidly counteracts the disruption of balance in the stimulated neurons, bringing their activity back to the background level within the time-course of a single stimulus presentation (**Figure 5D**). Due to this effect, the plastic network maintains low rates and an AI-score of 86%, whereas the static network decreases to 69% as a result of increased synchrony (data not shown).

In both networks, the strongly localized activity leads to highly specialized network responses whereby each sub-population's firing rates are highly correlated with that of their respective stimulus (**Figures 5B,E**). However, the correlation values are much lower in networks subjected to plasticity and so is the degree to which the population responses are specialized in relation to the background. The slightly degraded discriminability in plastic networks can also be seen by comparing the clustering of the circuit states in response to each pattern. Plastic network states cluster in well defined but less separated regions of state space than static network states (**Figures 5C,F**).

In summary, under input conditions where the stimulus has intermediate strength and the stimulated populations are very small, networks can easily produce a specialized response leading to accurate classification. The main effect of plasticity lies in its ability to maintain globally low average firing rates



(approximately half of those displayed in the corresponding static case) and to ensure the stability and maintenance of the AI state.

3.2.3. Plasticity stabilizes neural trajectories

Several of the conditions depicted in **Figure 4A** resulted in a significant performance advantage for networks incorporating activity-dependent adaptation. To better elucidate the mechanisms underlying such advantage, we focus on the condition where the difference is most evident (highlighted with a white

star in **Figure 4**) and analyse the dynamics of an individual network's responses to each stimulus pattern as they evolve along specific paths through the network's state space. It is worth noting that under the present input conditions (i.e., $\sigma_u = 100$, $N_{\text{aff}} = 10$, $\gamma_u = 0.3$), the responses are not discernible on the basis of a localized increase in firing rate among the stimulated neurons, which is reflected in the low degree of specialization of the population responses (**Figures 6F,L**). Hence, to understand the reasons underlying the performance difference, we must analyse

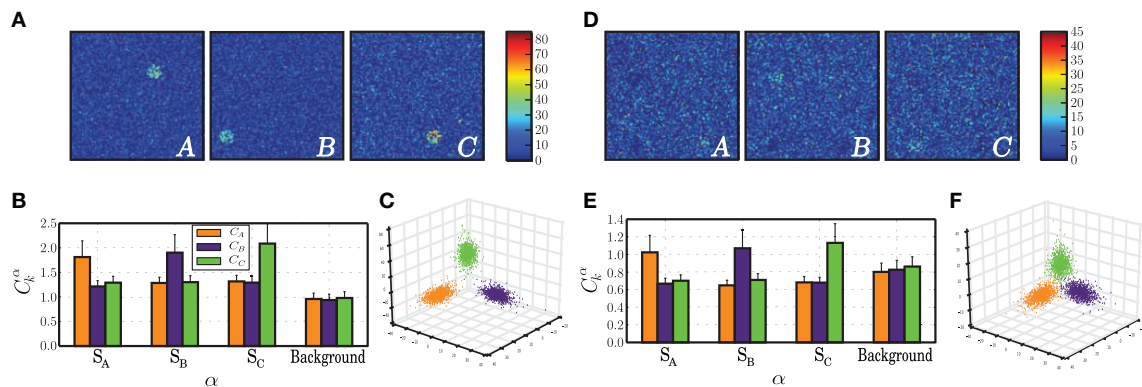


FIGURE 5 | Characteristics of population activity that allow a stable discrimination based on specialized population responses. The results depicted in this figure correspond to the condition $C_s \simeq C_p$ ($\sigma_u = 50$, $N_{\text{aff}} = 100$, $\gamma_u = 0.01$, gray star in **Figure 4**), obtained at the final stages of the stimulation period, i.e., after presenting the network with $\simeq 1000$ samples of each stimulus. **(A)** Snapshot of population activity represented as 2-dimensional firing rate maps, in response to each stimulus pattern (A, B, C) for a static network. The color of each point in the grid represents the

average firing rate of the neuron located in the corresponding position, averaged over an interval spanning the stimulus duration plus the inter-stimulus interval. **(B)** Correlations between the firing rates recorded in each sub-population $\alpha = S_A, S_B, S_C$, Background of a static network in response to each signal k , with each of the input signals s_k , averaged over 100 stimulus-response pairs. **(C)** Projection of network state vectors $x_{u,k}(t)$ in response to each stimulus pattern $u_k(t)$ onto the space spanned by the first 3 PCs. **(D–F)** as in **(A–C)** but for a plastic network.

the high-dimensional response dynamics to each stimulus, in the different network conditions.

Fully dissecting and understanding the dynamics of such high-dimensional dynamical systems is widely recognized as an extremely difficult, if not impossible task. We therefore resort to reduced-dimension descriptions and average measures that attempt to capture the essential phenomena as functions of a few variables expressing the most meaningful relations in the data. This has the advantage of allowing us to visualize the data and thus make some inferences and hypotheses about the underlying dynamics, but the disadvantage of providing a limited scope and ability to test the generality of these hypotheses in relation to the original state space.

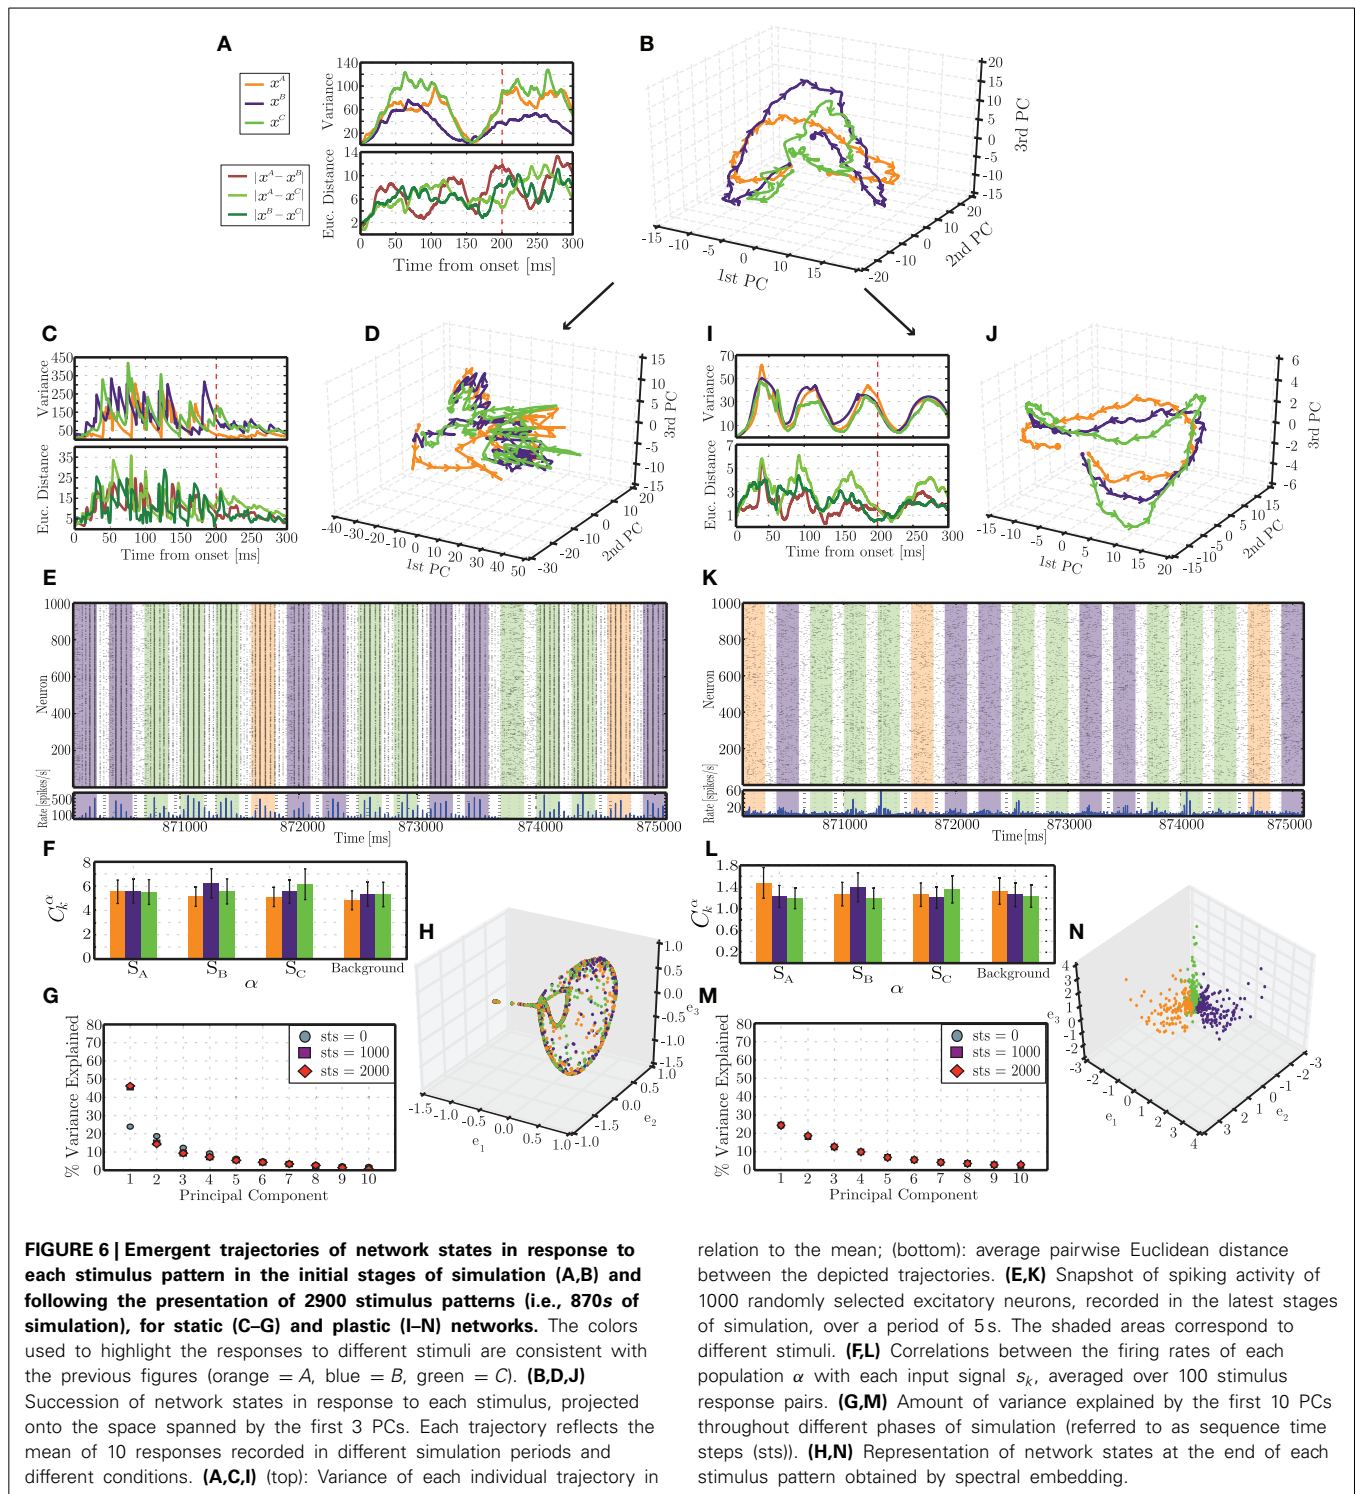
The results depicted in **Figure 6** reflect the activity of individual networks recorded at different points in time during the stimulation period. We will refer to these points as sequence time steps or *sts*. For each analyzed time point, the spiking activity in response to an individual stimulus is first low-pass filtered to create a response matrix containing the circuit states throughout the entire length of the response, until the onset of the subsequent stimulus. The dimensionality of each of these response matrices is then reduced by principal component analysis and their projections in the space spanned by the first three PCs analyzed. The procedure is repeated until 10 responses to each individual stimulus are obtained. We calculate the mean and variance of these responses to determine the stereotypy or variability of the transient activity patterns developed in response to the different stimuli starting from different network conditions.

In the first few sequence time steps (starting from $sts = 0$), the network responses already show a certain degree of stereotypy and the trajectories progress through distinct, albeit overlapping, regions of state space (**Figure 6B**). The average pairwise distances between trajectories show no specific pattern other than an increasing trend (**Figure 6A**, bottom). A striking feature, which

we will come back to later, is the existence of a clear pattern in the variance of the trajectories. These initial results are remarkably similar among the different conditions (static and plastic networks) as well as among different random network instantiations, reflecting the initially similar embedding state space, obtained by tuning the ongoing, background activity dynamics. The trajectories are not exactly the same but tend to occupy similar regions of space and display a very similar pattern of variances.

After being presented with a long sequence of stimuli, the response patterns differ dramatically between the static and the plastic conditions. These results are depicted in the bottom part of **Figure 6** (C–H, for static networks and I–N for plastic networks), and were obtained from $sts = 2900$. The trajectories of network states observed in static networks are now highly variable (with a variance about 4 times larger than in the initial steps; **Figure 6C**, top) and the different stimulus responses clump together, hampering an adequate discrimination (**Figure 6D**). In contrast, the trajectories observed in plastic networks have become more stereotypical, with a maximum variance approximately half of that verified in the initial condition (**Figure 6I**, top) and the responses become more “organized,” consistently unfolding throughout specific paths (**Figure 6J**).

Furthermore, the dimensionality of the response dynamics is also significantly different which has an obvious impact on their linear separability. The dimensionality of the state-space can be inferred by the amount of total variance explained by successive dimensions obtained by PCA. **Figure 6G** shows that, in static networks, after $sts = 1000$, the first PC accounts for just under 50% of the variance, compared to $\sim 25\%$ at $sts = 0$, which stands in clear contrast with the dynamics of plastic networks, where the percentage of explained variance remains invariant along the full stimulation time (**Figure 6M**). The low dimensionality and low separability of the static network’s responses is further demonstrated in the result displayed in **Figure 6H**, which was obtained



by spectral decomposition of the matrix of Euclidean distances between the state vectors (network states at $t_0 = 200$ ms, see the dimensionality reduction section of Section 2.5.3). This figure depicts the existence of a low-dimensional manifold, where all the states in response to the different input patterns lie.

Conversely, networks that have been shaped by plasticity learn to explore the state space much more effectively, partly by virtue

of the maintenance of the AI-type dynamics (**Figure 6K**), which supplies the network with a higher dimensional space over which to develop its responses (**Figures 6M,N**), in contrast with the static network where the activity tends to become more synchronized (**Figure 6E**), thereby increasing the redundancy of the individual neuron's responses resulting in a consequent reduction in dimensionality.

The variance of the analyzed responses also shows, at this stage, a clear periodic pattern where, at relatively constant intervals ($\simeq 70$ ms), the 10 individual trajectories converge (see **Figure 6I**). In the initial state ($sts = 0$), the variances already showed a similar pattern, with a convergence point at $t \simeq 150$ ms after stimulus onset (**Figure 6A**). This is an interesting and somewhat unexpected result. Since these points do not reflect any meaningful property of the stimulus, they must reflect properties of the response dynamics that the network develops. We hypothesize that these points represent spontaneously generated saddle nodes that stabilize the dynamics along the unfolding trajectories, improving robustness and reproducibility. As the different trajectories approach these regions of state-space they are attracted to these points (hence the observed reduction of variance) and, after leaving these regions, the trajectories are repelled and allowed to diverge until the next saddle node captures them.

An additional question that arises from these results is whether the increased discriminability of the population responses in the plastic networks can be accounted for by the macroscopic dynamical state of the network (i.e., the ability to maintain a stable AI activity pattern both for ongoing and stimulus-driven activity) or whether the fine details of the learned synaptic connections are strictly necessary. To address this question, we perform two simple experiments, described in the following.

As discussed in Section 3.1, the effect of plasticity modifies the effective balance g , leading to a final value that is much larger than the initial one (see Supplementary Materials). Therefore, as a result of learning, the network will be strongly dominated by inhibition and placed in a dynamic regime where ongoing activity is more strongly of the AI type (see **Figure 3**). To determine whether the macroscopic dynamical state is sufficient to account for the network performance, we investigate a static network initialized with $g \simeq 12$ (similar to the final value of g obtained in the plastic network). The readout classification performance of this strongly inhibitory network shows a considerable improvement over the more weakly inhibitory network considered in this section ($C_s \simeq 0.9667$ as compared to $C_s \simeq 0.445$), thus reducing the performance difference from 0.54 to 0.02.

This result seems to support the first hypothesis, i.e., that the increased discriminability is due solely to the network's dynamical state. However, a second experiment suggests that this view is too simplistic. We analyse the classification performance obtained if the learned synaptic weights are randomly shuffled, losing any relevant structure. To do so, 3000 stimulus samples are presented to the plastic network, after which its synaptic weights are frozen and plasticity disabled. Subsequently, the recorded weights are randomly shuffled among the existing synapses and the network is exposed to a new sequence of 3000 stimulus samples. The responses to this second set of stimuli is recorded and used to train and test the readout's classification performance (following the same procedure described in Section 2.5.3). In this situation, the classification performance drops to chance level ($C_p \simeq 0.33148$).

Based on these results, we can conclude that the macroscopic dynamical state of the network is critical to achieve a high stimulus discrimination and consequent readout performance. For that reason, a random network can achieve very high performance if its initial state is placed in a strong AI regime. However, if the network connectivity is not random, but pre-structured

by Hebbian learning in response to the training data, the fine details of connectivity that arise from the learning process play a key role in the maintenance of adequate stimulus representations; randomly re-organizing this connectivity structure results in a drop in performance to chance level. So, in this situation, the results do not rely exclusively on the global E/I balance (which is maintained after shuffling), but also require the conservation of the pre-learned weight structure. These phenomena are obviously not independent as the learned connectivity structure emerges to counteract the disruption of balance and to stabilize the activity in the AI regime. Randomly shuffling the synaptic weights may result, for example, in a decreased inhibition toward certain stimulated neurons, that consequently fire excessively and thus destabilize the global network dynamics. Indeed, the activity in the shuffled condition displays a higher amount of synchronous population activity (data not shown).

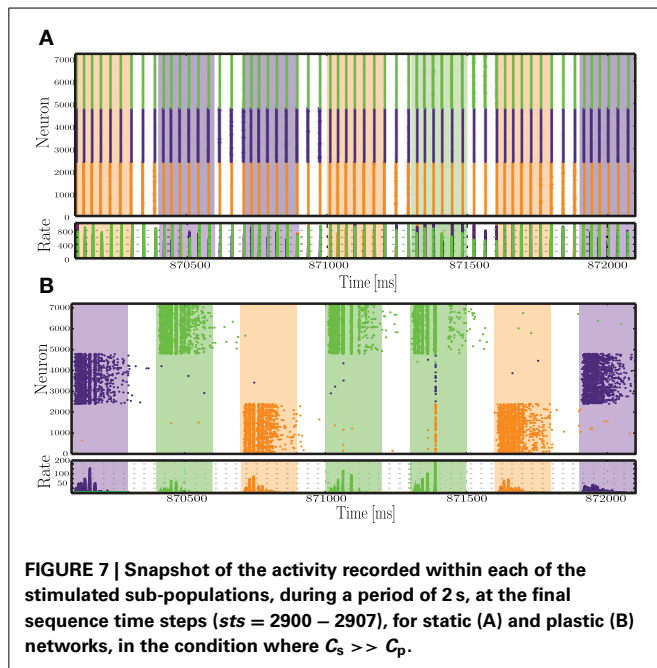
3.2.4. Strong stimulation hinders representation

In some cases, the presence of plasticity reduced the network's ability to represent the input into distinct activity patterns, e.g., the configuration marked with a black star in **Figure 4A**. The conditions that allow this to occur are characterized by intermediate or high peak firing rate and high number of afferents, i.e., very strong input. However, note that although the classification performance is higher for the static network (**Figure 4B**), all other metrics show the reverse effect. The absolute error of the readout output is higher for static networks (**Figure 4C**) and the solutions found by the regression algorithm for the output weights are quite unstable, relying heavily on some state variables in detriment of others (**Figure 4D**). This means that only a certain fraction of the population effectively communicates the relevant information to the readout. Even then, the output does not provide a good match to the target binary values, a result that is further reinforced by the point-biserial correlation between the readout output and the target output, which is close to or below 0 (**Figure 4E**).

Examining the network activity in these conditions provides an idea to the mechanisms underlying these results (**Figure 7**). In plastic networks, the input is too strong and causes the inhibitory synapses to become excessively strong to counteract the equally excessive excitatory drive. The result is an almost completely silenced excitatory population, where the only sparse spiking activity appears as short-lived bursts in immediate response to each input. On the other hand, static networks also develop an unfavorable dynamic state, where most of the activity is punctuated by synchronous, population-wide bursts. The stimulated neurons are briefly and slightly decoupled from the burst, which allows some separation of the responses. The readout algorithm captures mostly the activity of the input populations and heavily amplifies the weights from these neurons. This then leads to an output sequence that, despite a correct label assignment (i.e., the largest output values at each time step are assigned to the correct symbol), consists of disproportionately large values, which justify the large absolute error and the low correlations (**Figures 4C–E**).

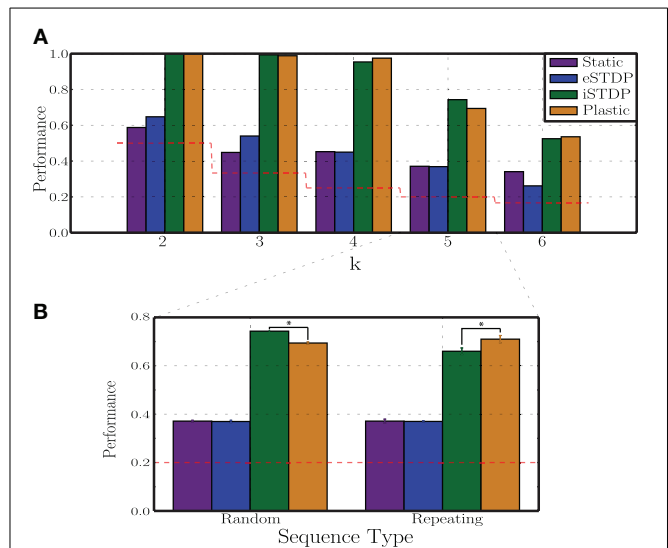
3.3. DIFFERENTIAL EFFECTS OF PLASTICITY

The results presented so far show that the action of plasticity modulates the network's ongoing activity, endowing it with the ability to maintain Asynchronous Irregular states over a



broader range of parameters and abolishing pathological states of Synchronous Regular activity (which we demonstrate to result in an impaired computational performance), when driven by a constant, stochastic and unspecific input (Section 3.1). In addition, as discussed in Section 3.2, if sequentially structured and topographically mapped input patterns are introduced, their interaction with the ongoing activity and the manner in which they disrupt local E/I balance (namely the strength and the spread of the disruption) determines the ability of networks operating in inhibition-dominated regimes to adopt adequate representations, i.e., to utilize bounded stimulus-specific sub-spaces. Different combinations of input features are shown to be able to cause discernible (linearly separable) population responses, regardless of the presence of adaptation. However, the characteristics of the adopted responses demonstrate that the action of plasticity is strictly necessary to maintain a suitable, “healthy” population activity by avoiding the pathological Synchronous Regular regimes toward which static networks are driven in the presence of strong stimulation (Section 3.2.4).

In the previous experiments we did not consider how the two different types of plasticity interact. We now turn our attention to disentangling the roles of the different plasticity mechanisms under study to determine whether the improvements observed in the development of stimulus representations are the product of a combined, synergistic action of these mechanisms, whether one of them plays a dominant role or whether they even counteract each other's effects. As we demonstrate in **Figure 4F** (see also Supplementary Materials), the steady-state weight distributions in the different conditions (disregarding the pathological states observed in the condition $C_s \gg C_p$ Section 3.2.4) do not differ noticeably from those developed in a control condition when no patterned stimuli are delivered, and consequently are not informative about these differential effects of eSTDP and iSTDP. We therefore adopt the configuration of input parameters that leads



to the greatest performance of plastic networks with respect to static networks ($C_p \gg C_s$, condition marked with a white star in **Figure 4**: $\sigma_u = 100$, $N_{\text{aff}} = 10$, $\gamma_u = 0.3$), and assess their performance in situations where the network dynamics is shaped by neither (Static), one (iSTDP/eSTDP) or both (Plastic) of the plasticity mechanisms. We systematically vary the task difficulty by building stimulus sequences with an increasing number of stimuli (k) thus requiring a matching number of discernible network responses in order to be discriminable.

The results of this analysis are depicted in **Figure 8A**. The most striking result is the clear dominance of iSTDP, which is solely responsible for most of the observed performance improvement in relation to the static condition. Working alone, eSTDP is only marginally advantageous in the less demanding task conditions ($k = 2, 3$) and, as the task difficulty increases, its actions result in no net improvements. In the extreme case, the presence of eSTDP can even undermine the network's representational abilities and decrease the overall performance to a level barely above chance (when $k = 6$). On the other hand, iSTDP alone accounts for the majority of the observed results; the addition of eSTDP can even decrease the readout performance ($k = 5$). In most cases, plastic networks with both mechanisms active perform as well as iSTDP alone or worse.

These results demonstrate that the main feature responsible for the increased discriminability of stimulus-driven population responses in plastic networks is the decorrelating action of iSTDP and the consequent maintenance of the AI dynamic regime both for ongoing activity and stimulus-driven responses. These findings contradict our initial hypothesis that a synchronous burst of activity impinging on particular groups of neurons would

bind the neurons belonging to these clusters by introducing correlations in their driven activities, and by so doing aid the discrimination of network states. However, no such clusters formed (see Supplementary Materials), due to the decorrelating action of iSTDP which hinders the efficacy of correlation-driven eSTDP.

The Hebbian nature of eSTDP is well suited to uncover causal relations in the input structure, which prompts the question of whether it would have a more beneficial effect on discriminability if the input contained clear causal relations rather than randomly drawn stimuli (see Section 2.4). We therefore performed an additional test in which the sequences were composed of $k = 5$ stimuli arranged in a fixed pattern, repeated throughout the entire stimulation period. The results of this analysis (shown in **Figure 8B**) demonstrate that, while eSTDP alone does not perform any better on the repetitive pattern (when compared with the random pattern), plastic networks with both eSTDP and iSTDP active perform somewhat better than iSTDP alone. So, even in the presence of a clear causal structure in the input data, the effect of eSTDP seems to be negligible which is an intriguing result. However, it should be noted that, while these results raise some questions, they are not fine-grained enough to allow us to draw broader conclusions. We note, for example, that the inter-stimulus interval (100 ms) used in these experiments is much longer than the relevant time window for eSTDP modifications (20 ms), thus diminishing its ability to learn the causal structure of the sequential input events. Further studies are strictly necessary to clarify some of these results and address the issues they raise.

4. DISCUSSION

A primary function of neocortical circuits lies in their ability to dynamically adopt and sustain reliable representations of sequentially occurring perceptual events in a self-organized and experience-dependent manner (Brosch and Schreiner, 2000; Zacks et al., 2007; Rabinovich et al., 2008; Buonomano and Maass, 2009). They need to maintain the necessary flexibility to adequately respond to sudden transitions that may require a global shift in representational space, while retaining a certain amount of contextual information. These characteristics are necessary for any further processing to occur (such as the dynamic evaluation of sequential dependencies present in the input), however, they entail an apparent contradiction between sensitivity and robustness. It seems probable that this is resolved via functional remodeling and adaptation, involving modifications at different spatial and temporal scales mediated by a combination of different synaptic and intrinsic mechanisms.

In this study, we have explored the relations between several important organizational principles of functional neurodynamics, involving distributed processing in inhibition dominated, sparsely coupled recurrent networks, whose rich ongoing dynamics supports the emergence of stimulus-specific spatiotemporal activity patterns. We have shown that the action of dynamic excitatory and inhibitory synapses, modulated by spike timing-dependent mechanisms, has a significant impact on the robustness and active maintenance of an ongoing activity state characterized by irregular firing that is asynchronous across the network. In this *asynchronous irregular* regime, the network activity is considered to most closely resemble cortical spiking activity

in vivo (Vaadia and Aertsen, 1992; Softky and Koch, 1993; Shadlen and Newsome, 1994; Brunel, 2000; Destexhe, 2009; Ostojic, 2014).

We have additionally established an objective relation between the dynamic states of ongoing activity (characterized by varying degrees of synchrony and regularity) and generic online processing capacity, demonstrating that pathological network-wide synchronization observed in the synchronous regular regime hinders the ability to properly map spatiotemporal input streams into discernible activity states, a process necessary for online computation on time-varying inputs. By abolishing such dynamic regimes (see **Figure 3**), balanced plasticity increases the robustness of generic computational capacity, thus expanding the efficacy of these circuits as information processing devices.

The sequential interaction of spatiotemporal input patterns with the ongoing network activity modifies the dynamics of the stimulated neurons and, via waves of recurrent interactions, also that of the global network on which they are embedded. These modifications are highly heterogeneous, depending on the nature and characteristics of the input stimulus. Experimental evidence shows that increased thalamic input is related to a higher degree of asynchronous activity in sensory cortices (Cohen and Maunsell, 2009; Poulet et al., 2012; Tan et al., 2014), which emphasizes the relevance of AI-type activity both as the ground state (Shadlen and Newsome, 1994, 1998; Vogels et al., 2005) and the active state of cortical activity, even though these two states may be characterized by different statistical features (Ostojic, 2014). In order to ascertain how certain features of the stimulus influence the network responses and modify the observed dynamics, we have driven the networks with specific input stimulus “events,” characterized by spike bursts of different amplitudes (as depicted in **Figure 1**), mimicking the thalamic burst mode of firing (Ramcharan et al., 2000; Sherman, 2001b; Bruno and Sakmann, 2006). These events impinge on a variable number of afferent neurons and target topographically arranged (Thivierge and Marcus, 2007; Silver and Kastner, 2009) subsets of excitatory neurons, thus momentarily disrupting the local E/I balance. The objective was to assess the quality and characteristics of dynamic stimulus representations developed by networks whose synapses are endowed with plasticity, enabling them to counteract the local disturbances, with networks whose synapses are fixed and static, in relation to the strength and spatial distribution of the stimuli.

Our results demonstrate that, in input conditions where the stimulus has intermediate strength but the stimulated populations are very small and spatially concise, the main effect of plasticity lies in its ability to maintain globally low average firing rates (approximately half of those displayed in the corresponding static case) and to ensure the stability and maintenance of the AI state (**Figure 5**). On the other hand, if the input is too strong, comprising the activity of a large number of afferent fibers, activity becomes highly pathological, even if plasticity is present. Whereas plastic networks become largely silent due to excessive inhibition that emerges to counteract the equally excessive excitatory drive, static networks become highly synchronized and fire in short population bursts (see **Figure 7**).

However, strong and highly focussed stimuli are probably not representative of typical cortical input. We also considered

scenarios where the stimulus was weaker and the receiving populations were large and distributed enough to avoid a strong localized response. In these situations, plasticity is shown to be generally beneficial (although not universally so), by allowing the network to efficiently explore a higher dimensional state space, achieved via the maintenance of AI-type activity (**Figures 6K–N**). The reduction in the dimensionality of the dynamic state observed in static networks, on the other hand, is a signature of an increasingly constrained and redundant dynamical space, which is detrimental to an adequate stimulus representation (**Figures 6D–H**). Plasticity is also shown to improve robustness and stereotypy of the successions of network states developed in response to each stimulus pattern (**Figures 6I,J**). Such transient, but trial-to-trial reproducible sequences of neural activity have been demonstrated experimentally in several sensory systems (e.g., Brosch and Schreiner, 2000; Mazor and Laurent, 2005; Broome et al., 2006; Rabinovich et al., 2008) and play a critical role in neural computation.

The pattern observed in the response trajectories demonstrated the existence of regions of negligible variance along the system's response trajectories. We hypothesize these regions to represent saddle nodes, i.e., metastable states, whose temporal order and location is determined by the network's self-organized functional connectivity. This hypothesis is consistent with known principles of neurodynamics (Rabinovich et al., 2006), namely the formation of stable heteroclinic sequences (Rabinovich et al., 2008; Rabinovich and Varona, 2011). The transformation of incoming stimuli into the spatiotemporal activity of a neuronal ensemble is represented as a heteroclinic sequence made up of many saddle nodes, and heteroclinic orbits connecting them, and whose specific architecture is stimulus-dependent and reproducible. Plasticity increases the number of such points along each stimulus representation—from **Figure 6A** (top) to **Figure 6I** (top), the number of low variance points along the response grows from one to three. This finding is interesting, as it suggests that activity-dependent self-organization adjusts the network dynamics in a manner that improves the resilience and reproducibility of each response to a specific stimulus, while maintaining an adequate underlying dynamics that keeps the network sensitive to external modulations.

Obviously, caution is warranted in relation to this interpretation of the data—since we are discussing the dynamics observed in a low-dimensional projection space, no definitive or absolute conclusions may be drawn regarding the original state space. However, the reproducibility of this pattern of results over a range of different random network instantiations and different initial conditions provides some support for the hypothesis. Furthermore, the formation of this stable periodic pattern of variance is only visible after a long training period. Analysing intermediate time points shows a gradual transition, where the number and frequency of low-variance regions varies among stimulus responses (data not shown). Nevertheless, further analysis is necessary to validate this argument. It would be interesting to obtain a low-dimensional formulation of the network dynamics under these conditions and carefully explore it to gain a better insight into the underlying mechanisms. This could be done, for instance, by eigenfunction expansion, which could provide a

reasonable approximate low-dimensional dynamical system that would allow careful analytic treatment.

Additional expansions of the current work could involve the use of different input stimuli, combinations of stimuli or the inclusion of temporal dependencies between sequence elements. Most of the input-dependent results we have analyzed (with the exception of Section 3.3), although involving stimulus sequences, are based solely on stimulus discrimination and representation given that the stimuli are randomly ordered. Under the theory of stable heteroclinic sequences, we would expect that plasticity would allow the network to develop sequence representations, where each element would be dynamically represented by its own saddle node, and full sequence memory would be encoded by a transient motion in state space along the paths specified by these metastable states. It would also be interesting to investigate whether the capacity of the networks to maintain an AI regime when perturbed would allow them to perform balanced amplification of specific activity states, as has recently been demonstrated for networks incorporating optimally tuned inhibition (Hennequin et al., 2014).

In summary, the most relevant conclusion to draw from the current results is that the quality of dynamical representations adopted in response to sequential stimulus patterns is very much dependent on the maintenance of the AI-type activity, which not only provides a stable high-dimensional embedding manifold (in the form of ongoing activity) from which stimulus-specific responses arise, but also shapes the stability and robustness of those responses, that must evolve through bounded trajectories through the network's state-space. In the present study, the ability to maintain these regimes in the face of variable disruptions was achieved by the decorrelating action of iSTDP, which accounts for the results displayed in **Figure 8A**. As the precise difference between spike times is not strictly required for this, it is reasonable to assume that simpler synaptic or intrinsic mechanisms may be equally capable of stabilizing the network's dynamics allowing it to support equally rich dynamical stimulus representations.

These results also raise important and intriguing questions. Given the limited role played by eSTDP in our study, what is its true functional relevance? Some argue that either the functional relevance of eSTDP in the adult cortex has been overstated (Lisman and Spruston, 2005, 2010) or that it must rely on more complex intracellular mechanisms, that are not fully captured by the current formulations (Shouval et al., 2010; Shulz and Jacob, 2010), which are largely based on *in vitro* recordings. We are not in a position to provide definitive answers to address this question, but given that properly representing stimulus events as they unfold over time is a necessary first step toward more complex computations, our demonstration that this ability does not require eSTDP, but relies on the homeostatic process of regulating ongoing activity by active decorrelation, provides some interesting material to this debate and opens up a new set of questions.

Which mechanisms account for the brain's ability to represent stimulus events occurring over variable time scales (most of which much longer than those relevant for STDP modifications) and discover causal relations between them? These are fundamental steps in most cognitive processes, and must rely on some

degree of lasting functional modifications. These processes are likely to rely on a complex interplay of various sub-processes, of which eSTDP and iSTDP may be an integral part of. The pressing need to address this type of questions, spanning multiple spatiotemporal descriptive scales reinforces the relevance of studies involving a synergistic combination of multiple adaptation mechanisms.

FUNDING

Partially funded by the Erasmus Mundus Joint Doctoral Program EuroSPIN, BMBF Grant 01GQ0420 to BCCN Freiburg, the Junior Professor Program of Baden-Württemberg, the Helmholtz Alliance on Systems Biology (Germany), the Initiative and Networking Fund of the Helmholtz Association and the Helmholtz Portfolio theme Supercomputing and Modeling for the Human Brain.

ACKNOWLEDGMENTS

We warmly thank Peggy Seriès for constructive discussions during Renato C. F. Duarte's stay in Edinburgh and for her support of this study as well as Susanne Kunkel for her help with the implementation of plasticity models in NEST. We would like to acknowledge the use of the computing resources provided by bwGrid (<http://www.bw-grid.de>), member of the German D-Grid initiative, funded by the Ministry for Education and Research (Bundesministerium für Bildung und Forschung) and the Ministry for Science, Research and Arts Baden-Württemberg (Ministerium für Wissenschaft, Forschung und Kunst Baden-Württemberg). All simulations were carried out with NEST (<http://www.nest-initiative.org>).

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <http://www.frontiersin.org/journal/10.3389/fncom.2014.00124/abstract>

REFERENCES

- Abbott, L., and Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nat. Neurosci.* 3, 1178–1183. doi: 10.1038/81453
- Abeles, M. (1991). *Corticonics: Neural Circuits of the Cerebral Cortex*, 1st Edn. Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511574566
- Belkin, M., and Niyogi, P. (2003). Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.* 15, 1373–1396. doi: 10.1162/089976603321780317
- Bi, G.-Q., and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Braitenberg, V., and Schüz, A. (1998). *Cortex: Statistics and Geometry of Neuronal Connectivity*. Berlin: Springer.
- Broome, B. M., Jayaraman, V., and Laurent, G. (2006). Encoding and decoding of overlapping odor sequences. *Neuron* 51, 467–482. doi: 10.1016/j.neuron.2006.07.018
- Brosch, M., and Schreiner, C. E. (2000). Sequence sensitivity of neurons in cat primary auditory cortex. *Cereb. Cortex* 10, 1155–1167. doi: 10.1093/cercor/10.12.1155
- Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comput. Neurosci.* 8, 183–208. doi: 10.1023/A:1008925309027
- Bruno, R. M., and Sakmann, B. (2006). Cortex is driven by weak but synchronously active thalamocortical synapses. *Science* 312, 1622–1627. doi: 10.1126/science.1124593
- Buonomano, D. V., and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nat. Rev. Neurosci.* 10, 113–125. doi: 10.1038/nrn2558
- Churchland, M., Yu, B., Sahani, M., and Shenoy, K. (2007). Techniques for extracting single-trial activity patterns from large-scale neural recordings. *Curr. Opin. Neurobiol.* 17, 609–618. doi: 10.1016/j.conb.2007.11.001
- Cohen, M. R., and Maunsell, J. H. R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* 12, 1594–1600. doi: 10.1038/nn.2439
- Compte, A., Brunel, N., Goldman-Rakic, P. S., and Wang, X.-J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex* 10, 910–923. doi: 10.1093/cercor/10.9.910
- Dan, Y., and Poo, M. (2006). Spike timing-dependent plasticity: from synapse to perception. *Physiol. Rev.* 86, 1033–1048. doi: 10.1152/physrev.00030.2005
- Dan, Y., and Poo, M.-M. (2004). Spike timing-dependent plasticity of neural circuits. *Neuron* 44, 23–30. doi: 10.1016/j.neuron.2004.09.007
- Destexhe, A. (2009). Self-sustained asynchronous irregular states and Up-Down states in thalamic, cortical and thalamocortical networks of nonlinear integrate-and-fire neurons. *J. Comput. Neurosci.* 27, 493–506. doi: 10.1007/s10827-009-0164-4
- Destexhe, A., Rudolph, M., Fellous, J.-M., and Sejnowski, T. (2001). Fluctuating synaptic conductances recreate *in vivo*-like activity in neocortical neurons. *Neuroscience* 107, 13–24. doi: 10.1016/S0306-4522(01)00344-X
- Destexhe, A., Rudolph, M., and Paré, D. (2003). The high-conductance state of neocortical neurons *in vivo*. *Nat. Rev. Neurosci.* 4, 739–751. doi: 10.1038/nrn1198
- Gerstein, G. L., and Mandelbrot, B. (1964). Random walk models for the spike activity of a single neuron. *Biophys. J.* 4, 41–68. doi: 10.1016/S0006-3495(64)86768-0
- Gewaltig, M.-O., and Diesmann, M. (2007). NEST (NEural simulation tool). *Scholarpedia* 2:1430. doi: 10.4249/scholarpedia.1430
- Gustafsson, B., Wigström, H., Abraham, W. C., and Huang, Y. Y. (1987). Long-term potentiation in the hippocampus using depolarizing current pulses as the conditioning stimulus to single volley synaptic potentials. *J. Neurosci.* 7, 774–780.
- Haider, B., Duque, A., Hasenstaub, A. R., and McCormick, D. A. (2006). Neocortical network activity *in vivo* is generated through a dynamic balance of excitation and inhibition. *J. Neurosci.* 26, 4535–4545. doi: 10.1523/JNEUROSCI.5297-05.2006
- Hennequin, G., Vogels, T. P., and Gerstner, W. (2014). Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* 82, 1394–1406. doi: 10.1016/j.neuron.2014.04.045
- Holmgren, C. D., and Zilberter, Y. (2001). Coincident spiking activity induces long-term changes in inhibition of neocortical pyramidal cells. *J. Neurosci.* 21, 8270–8277.
- Koch, C. (2004). *Biophysics of Computation: Information Processing in Single Neurons Computational Neuroscience*. Oxford University Press.
- Kremkow, J., Aertsen, A., and Kumar, A. (2010). Gating of signal propagation in spiking neural networks by balanced and correlated excitation and inhibition. *J. Neurosci.* 30, 15760–15768. doi: 10.1523/JNEUROSCI.3874-10.2010
- Kullmann, D. M., Moreau, A. W., Bakiri, Y., and Nicholson, E. (2012). Plasticity of inhibition. *Neuron* 75, 951–962. doi: 10.1016/j.neuron.2012.07.030
- Kumar, A., Rotter, S., and Aertsen, A. (2008a). Conditions for propagating synchronous spiking and asynchronous firing rates in a cortical network model. *J. Neurosci.* 28, 5268–5280. doi: 10.1523/JNEUROSCI.2542-07.2008
- Kumar, A., Schrader, S., Aertsen, A., and Rotter, S. (2008b). The high-conductance state of cortical networks. *Neural Comput.* 20, 1–43. doi: 10.1162/neco.2008.20.1.1
- Kunkel, S., Diesmann, M., and Morrison, A. (2011). Limits to the development of feed-forward structures in large recurrent neuronal networks. *Front. Comput. Neurosci.* 4:160. doi: 10.3389/fncom.2010.00160
- Lamsa, K. P., Kullmann, D. M., and Woodin, M. A. (2010). Spike-timing dependent plasticity in inhibitory circuits. *Front. Syn. Neurosci.* 2:8. doi: 10.3389/fnsyn.2010.00008
- Legenstein, R., and Maass, W. (2007). Edge of chaos and prediction of computational performance for neural circuit models. *Neural Netw.* 20, 323–334. doi: 10.1016/j.neunet.2007.04.017

- Lisman, J., and Spruston, N. (2005). Postsynaptic depolarization requirements for LTP and LTD: a critique of spike timing-dependent plasticity. *Nat. Neurosci.* 8, 839–841. doi: 10.1038/nn0705-839
- Lisman, J., and Spruston, N. (2010). Questions about STDP as a general model of synaptic plasticity. *Front. Syn. Neurosci.* 2:140. doi: 10.3389/fnsyn.2010.00140
- Lu, J.-T., Li, C.-Y., Zhao, J.-P., Poo, M.-M., and Zhang, X.-H. (2007). Spike-timing-dependent plasticity of neocortical excitatory synapses on inhibitory interneurons depends on target cell type. *J. Neurosci.* 27, 9711–9720. doi: 10.1523/JNEUROSCI.2513-07.2007
- Lukoševicius, M., and Jaeger, H. (2009). Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.* 3, 127–149. doi: 10.1016/j.cosrev.2009.03.005
- Maass, W. (2010). “Liquid state machines: motivation, theory, and applications,” in *Computability in Context: Computation and Logic in the Real World*, eds B. Cooper and A. Sorbi (London: Imperial College Press), 275–296.
- Maass, W., Legenstein, R., and Bertschinger, N. (2005). Methods for estimating the computational power and generalization capability of neural microcircuits. *Theor. Comput. Sci.* 17, 865–872.
- Maass, W., Natschlag, T., and Markram, H. (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* 14, 1–27. doi: 10.1162/089976602760407955
- Markram, H., Gerstner, W., and Sjöström, P. J. (2012). Spike-timing-dependent plasticity: a comprehensive overview. *Front. Syn. Neurosci.* 4:2. doi: 10.3389/fnsyn.2012.00002
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–215. doi: 10.1126/science.275.5297.213
- Matsumura, M., Chen, D., Sawaguchi, T., Kubota, K., and Fetz, E. E. (1996). Synaptic interactions between primate precentral cortex neurons revealed by spike-triggered averaging of intracellular membrane potentials *in vivo*. *J. Neurosci.* 16, 7757–7767.
- Mazor, O., and Laurent, G. (2005). Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* 48, 661–673. doi: 10.1016/j.neuron.2005.09.032
- Meffin, H., Burkitt, A. N., and Grayden, D. B. (2004). An analytical model for the “large, fluctuating synaptic conductance state” typical of neocortical neurons *in vivo*. *J. Comput. Neurosci.* 16, 159–175. doi: 10.1023/B:JCNS.0000014108.03012.81
- Morrison, A., Diesmann, M., and Gerstner, W. (2008). Phenomenological models of synaptic plasticity based on spike timing. *Biol. Cybern.* 98, 459–478. doi: 10.1007/s00422-008-0233-1
- Morrison, A., Mehring, C., Geisel, T., Aertsen, A., and Diesmann, M. (2005). Advancing the boundaries of high-connectivity network simulation with distributed computing. *Neural Comput.* 17, 1776–1801. doi: 10.1162/0899766054026648
- Ng, A., Jordan, M., and Weiss, Y. (2001). “On spectral clustering: analysis and an algorithm,” in *Proceedings of Advances in Neural Information Processing Systems* (Cambridge), 849–856.
- Ostojic, S. (2014). Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons. *Nat. Neurosci.* 17, 594–600. doi: 10.1038/nn.3658
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: machine learning in python. *J. Mach. Learn. Res.* 12, 2825–2830.
- Poulet, J. F. A., Fernandez, L. M. J., Crochet, S., and Petersen, C. C. H. (2012). Thalamic control of cortical states. *Nat. Neurosci.* 15, 370–372. doi: 10.1038/nn.3035
- Rabinovich, M., Huerta, R., and Laurent, G. (2008). Transient dynamics for neural processing. *Science* 321:48. doi: 10.1126/science.1155564
- Rabinovich, M., Varona, P., Selverston, A., and Abarbanel, H. (2006). Dynamical principles in neuroscience. *Rev. Mod. Phys.* 78, 1213–1265. doi: 10.1103/RevModPhys.78.1213
- Rabinovich, M. I., and Varona, P. (2011). Robust transient dynamics and brain functions. *Front. Comput. Neurosci.* 5:24. doi: 10.3389/fncom.2011.00024
- Ramcharan, E. J., Gnadt, J. W., and Sherman, S. M. (2000). Burst and tonic firing in thalamic cells of unanesthetized, behaving monkeys. *Vis. Neurosci.* 17, 55–62. doi: 10.1017/S0952523800171056
- Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., and Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nat. Neurosci.* 16, 486–492. doi: 10.1038/nn.3331
- Shadlen, M., and Newsome, W. (1994). Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* 4, 569–579. doi: 10.1016/0959-4388(94)90059-0
- Shadlen, M. N., and Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* 18, 3870–3896.
- Sherman, S. M. (2001a). A wake-up call from the thalamus. *Nat. Neurosci.* 4, 344–346. doi: 10.1038/85973
- Sherman, S. M. (2001b). Tonic and burst firing: dual modes of thalamocortical relay. *Trends Neurosci.* 24, 122–126. doi: 10.1016/S0166-2236(00)01714-8
- Shouval, H. Z., Wang, S. S.-H., and Wittenberg, G. M. (2010). Spike timing dependent plasticity: a consequence of more fundamental learning rules. *Front. Comput. Neurosci.* 4:19. doi: 10.3389/fncom.2010.00019
- Shulz, D. E., and Jacob, V. (2010). Spike-timing-dependent plasticity in the intact brain: counteracting spurious spike coincidences. *Front. Syn. Neurosci.* 2:137. doi: 10.3389/fnsyn.2010.00137
- Silver, M. A., and Kastner, S. (2009). Topographic maps in human frontal and parietal cortex. *Trends Cogn. Sci.* 13, 488–495. doi: 10.1016/j.tics.2009.08.005
- Singer, W. (2013). Cortical dynamics revisited. *Trends Cogn. Sci.* 17, 616–626. doi: 10.1016/j.tics.2013.09.006
- Sjöström, J., and Gerstner, W. (2010). Spike-timing dependent plasticity. *Scholarpedia* 5:1362. doi: 10.4249/scholarpedia.1362
- Softky, W. R., and Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.* 13, 334–350.
- Stiefel, K. M., Englitz, B., and Sejnowski, T. J. (2013). Origin of intrinsic irregular firing in cortical interneurons. *Proc. Natl. Acad. Sci. U.S.A.* 110, 7886–7891. doi: 10.1073/pnas.1305219110
- Tan, A. Y. Y., Chen, Y., Scholl, B., Seidemann, E., and Priebe, N. J. (2014). Sensory stimulation shifts visual cortex from synchronous to asynchronous states. *Nature* 509, 226–229. doi: 10.1038/nature13159
- Tenenbaum, J. B., de Silva, V., and Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science* 290, 2319–2323. doi: 10.1126/science.290.5500.2319
- Thivierge, J.-P., and Marcus, G. F. (2007). The topographic brain: from neural connectivity to cognition. *Trends Neurosci.* 30, 251–259. doi: 10.1016/j.tins.2007.04.004
- Tsodyks, M., and Sejnowski, T. (1995). Rapid state switching in balanced cortical network models. *Netw. Comput. Neural Syst.* 6, 111–124. doi: 10.1088/0954-898X/6/2/001
- Tuckwell, H. C. (1988). *Introduction to Theoretical Neurobiology: Volume 1, Linear Cable Theory and Dendritic Structure*. Cambridge: Cambridge University Press.
- Vaadia, E., and Aertsen, A. (1992). “Coding and computation in the cortex: single neuron activity and cooperative phenomena,” in *Information Processing in the Cortex: Experiments and Theory*, eds A. Aertsen and V. Braitenberg (Berlin: Springer-Verlag), 81–121.
- van Rossum, M. C., Bi, G.-Q., and Turrigiano, G. G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neurosci.* 20, 8812–8821.
- van Vreeswijk, C., and Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Comput.* 10, 1321–1371. doi: 10.1162/089976698300017214
- Vogels, T. P., Froemke, R. C., Doyon, N., Gilson, M., Haas, J. S., Liu, R., et al. (2013). Inhibitory synaptic plasticity: spike timing-dependence and putative network function. *Front. Neural Circuits* 7:119. doi: 10.3389/fncir.2013.00119
- Vogels, T. P., Rajan, K., and Abbott, L. F. (2005). Neural network dynamics. *Ann. Rev. Neurosci.* 28, 357–376. doi: 10.1146/annurev.neuro.28.061604.135637
- Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* 334, 1569–1573. doi: 10.1126/science.1211095
- von der Malsburg, C., Phillips, W., and Singer, W. (2010). *Dynamic Coordination in the Brain: From Neurons to Mind*. Cambridge: MIT Press. doi: 10.7551/mitpress/9780262014717.001.0001
- Wang, L., and Maffei, A. (2014). Inhibitory plasticity dictates the sign of plasticity at excitatory synapses. *J. Neurosci.* 34, 1083–1093. doi: 10.1523/JNEUROSCI.4711-13.2014
- Woodin, M. A., Ganguly, K., and Poo, M.-M. (2003). Coincident pre- and post-synaptic activity modifies GABAergic synapses by postsynaptic changes in Cl⁻ transporter activity. *Neuron* 39, 807–820. doi: 10.1016/S0896-6273(03)00507-5

- Yger, P., El Boustani, S., Destexhe, A., and Frégnac, Y. (2011). Topologically invariant macroscopic statistics in balanced networks of conductance-based integrate-and-fire neurons. *J. Comput. Neurosci.* 31, 229–245. doi: 10.1007/s10827-010-0310-z
- Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., and Reynolds, J. R. (2007). Event perception: a mind-brain perspective. *Psychol. Bull.* 133, 273–293. doi: 10.1037/0033-2909.133.2.273
- Zacks, J. M., and Tversky, B. (2001). Event structure in perception and conception. *Psychol. Bull.* 127, 3–21. doi: 10.1037/0033-2909.127.1.3

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 07 May 2014; accepted: 16 September 2014; published online: 22 October 2014.

Citation: Duarte RCF and Morrison A (2014) Dynamic stability of sequential stimulus representations in adapting neuronal networks. *Front. Comput. Neurosci.* 8:124. doi: 10.3389/fncom.2014.00124

This article was submitted to the journal *Frontiers in Computational Neuroscience*.

Copyright © 2014 Duarte and Morrison. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY).

The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.



Unsupervised discrimination of patterns in spiking neural networks with excitatory and inhibitory synaptic plasticity

Narayan Srinivasa* and Youngkwan Cho

Center for Neural and Emergent Systems, Information and Systems Sciences Department, HRL Laboratories LLC, Malibu, CA, USA

Edited by:

Friedemann Zenke, École Polytechnique Fédérale de Lausanne, Switzerland

Reviewed by:

Andrey Olypher, Georgia Gwinnett College, USA
Guillaume Hennequin, The University of Cambridge, UK
Tilo Schwalger, École Polytechnique Fédérale de Lausanne, Switzerland

*Correspondence:

Narayan Srinivasa, Center for Neural and Emergent Systems, Information and Systems Sciences Department, HRL Laboratories LLC, Malibu Canyon Road, Malibu, CA 90265, USA
e-mail: nsrinivasa@hrl.com

A spiking neural network model is described for learning to discriminate among spatial patterns in an unsupervised manner. The network anatomy consists of source neurons that are activated by external inputs, a reservoir that resembles a generic cortical layer with an excitatory-inhibitory (EI) network and a sink layer of neurons for readout. Synaptic plasticity in the form of STDP is imposed on all the excitatory and inhibitory synapses at all times. While long-term excitatory STDP enables sparse and efficient learning of the salient features in inputs, inhibitory STDP enables this learning to be stable by establishing a balance between excitatory and inhibitory currents at each neuron in the network. The synaptic weights between source and reservoir neurons form a basis set for the input patterns. The neural trajectories generated in the reservoir due to input stimulation and lateral connections between reservoir neurons can be readout by the sink layer neurons. This activity is used for adaptation of synapses between reservoir and sink layer neurons. A new measure called the discriminability index (DI) is introduced to compute if the network can discriminate between old patterns already presented in an initial training session. The DI is also used to compute if the network adapts to new patterns without losing its ability to discriminate among old patterns. The final outcome is that the network is able to correctly discriminate between all patterns—both old and new. This result holds as long as inhibitory synapses employ STDP to continuously enable current balance in the network. The results suggest a possible direction for future investigation into how spiking neural networks could address the stability-plasticity question despite having continuous synaptic plasticity.

Keywords: spiking, STDP, learning, pattern discrimination, stability-plasticity dilemma, reservoir computing, balanced networks, basis sets

INTRODUCTION

A hallmark of biological systems is their ability to learn new knowledge while also exhibiting stability in order to prevent the forgetting of previous knowledge in a dynamically changing world. The nervous system solves this challenging problem in an unsupervised fashion and this problem has been referred to as the *stability-plasticity dilemma* (Grossberg, 1980, 2012).

This problem is further compounded in its complexity by the fact biological systems are open thermodynamic systems where energy and matter constantly flow through them (Katchalsky and Kedem, 1962; Swenson and Turvey, 1991; Kello, 2013). This flow produces variations within the nervous system where action potentials are always generated by neurons such that synaptic strengths are constantly being modulated (Freeman, 2001) to adapt to a changing world, and network structures never stop changing (Pascual-Leone et al., 2005) and all these changes can happen at a variety of spatial and temporal scales.

In a well-known set of experiments by Freeman and Schneider (1982), rabbits were surgically implanted with a rectangular array of electrodes in the olfactory bulb. In one such experiment to test serial conditioning, odor stimuli in the form of sawdust, acetyl acetate, butyric acid and finally sawdust were presented serially

to the rabbits. The neural activity in the bulb electrodes changed with each new odorant. On returning to the first odorant, the sawdust, neural activity was very different from those recorded on the first exposure. However, the rabbits exhibited repeatable behaviors such as avoiding odors that were undesirable while approaching toward other odors that were desirable. How is that the neural activity (or internal representations) in the brain can be so variable and yet the animal can produce stable and repeatable behaviors?

Neural models based on the adaptive resonance theory (Grossberg, 2012) attempt to answer these questions by using firing rate code combined with Hebbian plasticity models. Rate coding is based on the assumption that information is coded coarsely in the number of spikes occurring in a given window of time. The recently proposed *reservoir-computing model* (Maass et al., 2002; Buonomano and Maass, 2009; Maass, 2010) predicts that temporal integration of incoming information and generic non-linear mixing of this information within a liquid or recurrent network of excitatory and inhibitory neurons are primary computational functions of a cortical microcircuit. The state of the network at any given time can be represented by a point in high-dimensional space where each dimension corresponds

to the activity level of a neuron. A temporal sequence of these points forms a neural trajectory. The advantage of computing with neural trajectories is that temporal information is implicitly encoded in them and can be read out by downstream neurons. This approach to computing has received experimental evidence (Hahnloser et al., 2002; Nikolic et al., 2009; Crowe et al., 2010; Long et al., 2010; Bernacchia et al., 2011; Klampfl et al., 2012). These models are based on firing rates of neurons (Jaeger and Haas, 2004; Sussillo and Abbott, 2009; Laje and Buonomano, 2013).

There is mounting evidence for temporal coding in the brain (Rieke et al., 1996; Victor and Purpura, 1996; Van Rullen et al., 2005; Dan and Poo, 2006; Tiesinga et al., 2008) where information is coded in the precise timing of individual spikes from individual neurons. The adaptive resonance models also do not consider spike frequency dependent short-term plasticity (Tsodyks and Markram, 1997; Tsodyks et al., 1998) and spike timing dependent long-term plasticity of excitatory and inhibitory synapses (Markram et al., 1997, 2011; Bi and Poo, 1998; Woodin et al., 2003; Vogels et al., 2011, 2013). Spiking versions of reservoir computing models have shown learning of spatiotemporal patterns (Maass et al., 2002; Maass, 2010) but the reservoir is not plastic in these implementations.

A spiking neural network with spike-driven synaptic dynamics compatible with STDP and short-term synaptic plasticity and with supervisory signals was shown to learn and correctly classify a large number of overlapping patterns (Brader et al., 2007). This network did not consider inhibitory synaptic plasticity dynamics and required plasticity to be turned off after learning. In a previous model, the authors showed that a similar supervisory signal driven spiking neural network learns spatiomotor transformations (Srinivasa and Cho, 2012). It was shown recently that incorporation of synaptic plasticity in the excitatory synapse and network motifs within a spiking reservoir can result in the emergence of long-term memory in the form of sequences of network states (Klampfl and Maass, 2013). However, this model does not have synaptic plasticity in both excitatory and inhibitory synapses. It also did not address the relation of their network to the unsupervised discrimination of patterns.

A spiking neural model with a reservoir type architecture is presented that is composed of a source layer with neurons that are activated by external inputs, a reservoir that resembles a generic cortical layer with an excitatory-inhibitory (EI) network and a sink layer of neurons for readouts. Synaptic plasticity in the form of STDP is imposed on all the excitatory and inhibitory synapses at all times. Using a novel discrimination measure called pattern discriminability index (*DI*), the spiking network is shown to be capable of discriminating between spatial patterns of spiking inputs in an unsupervised manner (i.e., without any explicit supervisory signals or labels) despite continuous synaptic plasticity.

The *DI* can be viewed a generalization of the average Hamming distance (Garcia-Sanchez and Huerta, 2004; Olypher et al., 2012) between neuronal patterns based on relative firing rate distributions. It also has close links to information theoretic measures (Borst and Theunissen, 1999) because it quantifies the amount

of information the output neurons carry about the input patterns presented to the system during training.

MATERIALS AND METHODS

MODEL ARCHITECTURE

The spiking network model proposed in this paper consists of three layers as shown in **Figure 1A**. The source layer contains excitatory neurons that are stimulated by sources external to the network and projected to reservoir neurons. These projections were random and relatively sparse for the sake of simplicity. The reservoir neurons were either excitatory or inhibitory, received projections from source neurons and other reservoir neurons, and projected to other reservoir neurons and neurons in the third layer called the sink layer. The sink neurons received projections from the reservoir neurons but did not project back to the network. The sink neurons were composed of both excitatory and inhibitory neurons.

In this paper, the source layer (layer 1 in **Figure 1A**) contains $K = 900$ neurons (converted from a 30×30 2-D array into a linear array), the reservoir (layer 2 in **Figure 1A**) contains $N = 200$ excitatory and 50 inhibitory neurons (in a 4:1 ratio between excitatory and inhibitory neurons) and $M = 8$ excitatory neurons sink layer (layer 3 in **Figure 1A**) that are recurrently connected to inhibitory neurons in the sink layer. There are four types of synapses depending on the pre- and post-synaptic neuron type at each synapse: $E \rightarrow E$, $E \rightarrow I$, $I \rightarrow E$, and $I \rightarrow I$. The first two types of synapses are excitatory in nature and obey E-STDP rule while the last two types of synapses are inhibitory in nature and obey the I-STDP rule for plasticity. The connectivity between the layers in the network is set randomly with probability c_{ij}^{AB} where the superscripts A and B reflect excitatory (E) or inhibitory (I) type of neuron while subscripts i and j correspond to the sender and receiver layers (**Figure 1A**). All synapses are plastic throughout all simulations and synaptic connections are set randomly. The spiking model simulations were performed using the HRLSim (Minkovich et al., 2014) that is a multiple graphical processing unit (GPU) based spiking simulator in C++.

NEURON MODEL

The leaky integrate and fire neuron (Vogels et al., 2005) is used to model neuronal dynamics with a single compartment and no somatic, dendritic or axonal specialization. In response to multiple input currents coming from excitatory and inhibitory presynaptic neurons in the sets Pre_{ex} and Pre_{inh} , respectively, the membrane potential V for post-synaptic neuron i is determined by:

$$\tau_m \frac{dV_i}{dt} = (V_{rest} - V_i) + (E_{ex} - V_i) \sum_{j \in Pre_{ex}} g_{ex, ij} + (E_{inh} - V_i) \sum_{j \in Pre_{inh}} g_{inh, ij} \quad (1)$$

When V reaches a threshold voltage V_T , the neuron fires a spike (**Figure 1B**), and V is reset to V_{reset} . The output information is encoded into the timing of these spikes. This basic model provides

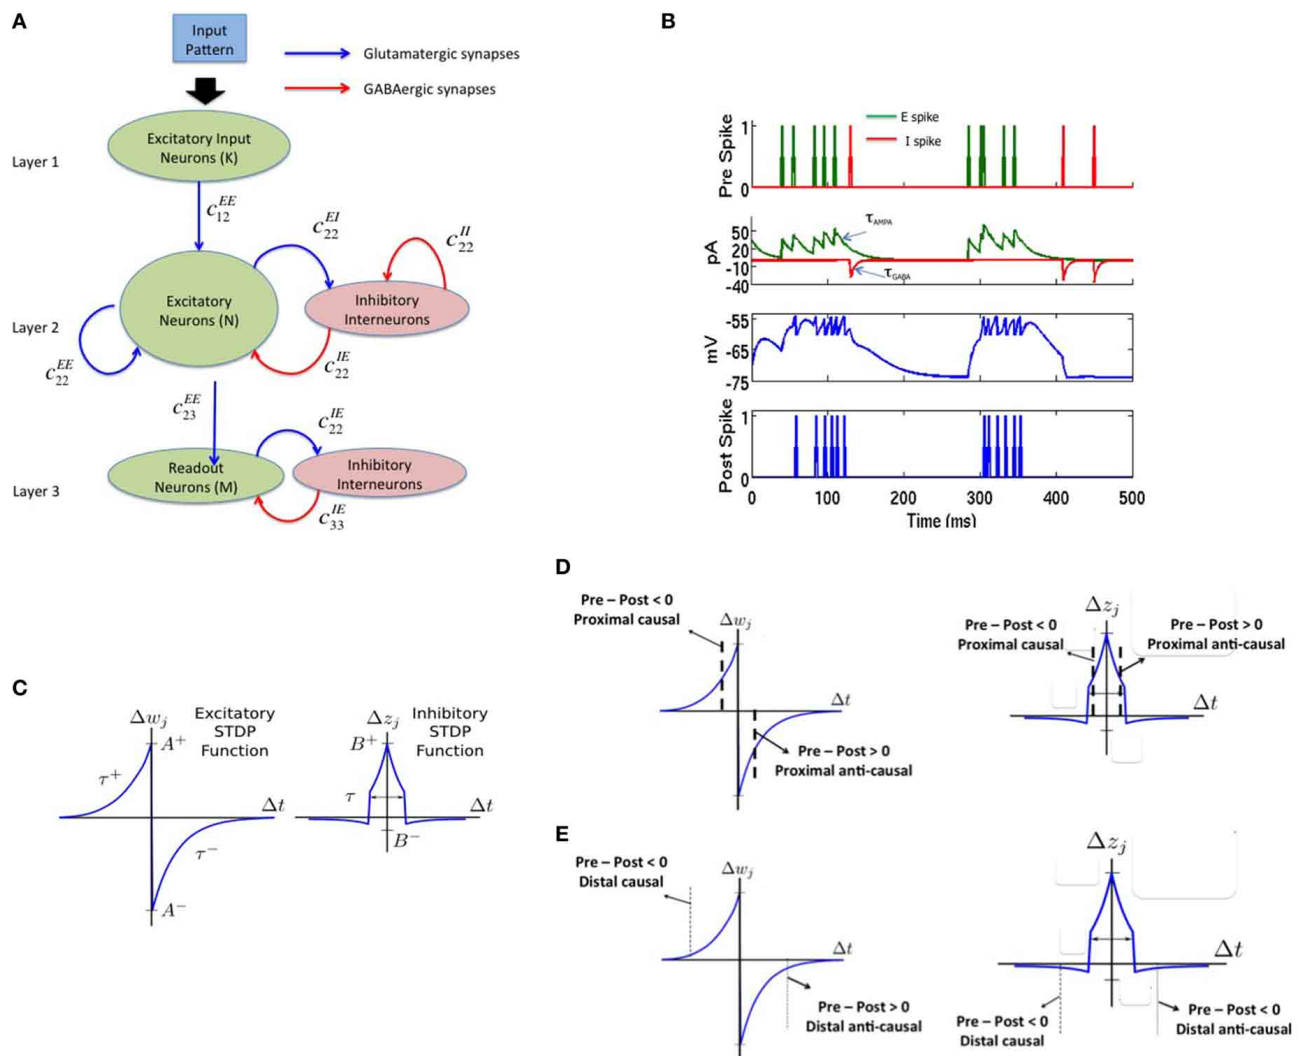


FIGURE 1 | The complete model with (A) a three layered network architecture with a source (layer 1), reservoir (layer 2), and sink (layer 3) neurons. The source neurons receive inputs patterns in spike-encoded form. These spikes are then projected to the excitatory neurons in the reservoir layer that are recurrently connected to other neurons in the excitatory population. The excitatory population of neurons is also connected to an inhibitory population of neurons reciprocally. The inhibitory neurons are recurrently connected to neurons within its population. The connectivity between the various layers in the network are set as: $c_{12}^{EE} = 20\%$, $c_{22}^{EE} = 40\%$, $c_{22}^{EI} = 40\%$, $c_{22}^{II} = 50\%$, $c_{23}^{EE} = 30\%$, $c_{33}^{EI} = 100\%$, and $c_{33}^{II} = 100\%$ for all simulations. Here, $c_{22}^{EI} = 40\%$ means that the connectivity between the E and I neurons in layer 2 is randomly connected at 40% of full connectivity between the two neuron populations. **(B)** The four subplots summarizes the leaky integrate and fire process in a typical neuron in our model. The first subplot shows input spikes from E (green) and I (red) pre-synaptic neurons. The second subplot shows the conversion of these spikes into currents that also includes the AMPA (green) and GABA (red) kinetics. The third subplot shows the integration of membrane voltage trace of the post-synaptic neuron based on the sum of the currents; and the last subplot shows the spikes generated by the post-synaptic neuron when the membrane voltage exceeds V_T . **(C)** The E-STDP is an asymmetric function of the timing difference ($\Delta t = t_{pre} - t_{post}$) between the pre- and post-synaptic spikes at neuron j and the corresponding change in synaptic conductance

Δw_j for $E \rightarrow E$ and $E \rightarrow I$ synapses. The four parameters (A^+ , A^- , τ^+ , τ^-) control the shape of the function and thus the amount of potentiation and depression. The I-STDP is a symmetric function of the timing difference Δt between the pre- and post-synaptic spikes at neuron j and the corresponding change in synaptic conductance Δz_j for $I \rightarrow E$ and $I \rightarrow I$ synapses. The three parameters (B^+ , B^- , τ) control the shape of the function and thus the amount of potentiation and depression. **(D)** Inhibitory STDP interacts with excitatory STDP to favor balance among causal synaptic currents. Presynaptic and post-synaptic spikes can be proximal causal or proximal anti-causal to varying degrees. The dashed lines reflect an example of timing difference for which proximal causality is assumed. **(E)** Presynaptic and post-synaptic spikes can be distal causal or distal anti-causal to varying degrees. The dashed lines reflect an example of timing difference for which distal causality is assumed. The E-STDP and I-STDP combine to form four different interacting regimes: *Balance regime* that occurs for the proximal causal case, where excitatory and inhibitory conductance both increase; *Accelerated Potentiation regime* that occurs for the distal causal case, where excitatory conductance increases albeit by small amounts, while inhibitory conductance decreases by small amounts; *Decelerated Depression regime* that occurs in distal anti-causal case, where excitatory and inhibitory conductance both decrease by small amounts; and *Quiescent regime* that occurs in proximal anti-causal case, where excitatory conductance is strongly decreased and inhibitory conductance is strongly increased.

several control variables for the membrane voltage including conductances g_{ex} (excitatory) and g_{inh} (inhibitory), membrane time constant τ_m , the constant reversal potential for excitatory (E_{ex}), and inhibitory (E_{inh}) synaptic currents, and a fixed voltage threshold for firing V_T at which the neuron fires a spike. Synaptic inputs to the neuron are modeled as conductance changes where a set of excitatory or inhibitory presynaptic spike times, S_{ex} or S_{inh} , respectively, gives conductance dynamics:

$$\frac{dg_{ex}}{dt} = -\frac{g_{ex}}{\tau_{AMPA}} + w \sum_{s \in S_{ex}} \delta(t - s) \quad (2)$$

$$\frac{dg_{inh}}{dt} = -\frac{g_{inh}}{\tau_{GABA}} + z \sum_{s \in S_{inh}} \delta(t - s) \quad (3)$$

Here the time constants τ_{AMPA} and τ_{GABA} approximate the average decay of AMPA and GABA currents respectively (Figure 1B). The value of the excitatory and inhibitory synaptic conductance w and z is controlled by STDP (Figure 1C). In all simulations, $\tau_m = 20$ ms, $V_T = -54$ mV, $V_{rest} = -74$ mV, $V_{reset} = -60$ mV, $E_{ex} = 0$ mV, $E_{inh} = -80$ mV, $\tau_{AMPA} = 40$ ms, and $\tau_{GABA} = 50$ ms. All simulations used Euler integration with a time step of 1 ms (Srinivasa and Jiang, 2013).

EXCITATORY STDP

The E-STDP function modulates the excitatory synaptic weight w based on the timing difference ($t_{pre} - t_{post}$), or Δt , between the spike times of pre- and post-synaptic neuron (Figure 1C). The control parameters $\tau^+ = 20$ ms and $\tau^- = 20$ ms determine the temporal window over which STDP is active. The change in synaptic weight is computed using the additive STDP rule as:

$$w = w_{old} + \Delta w \quad (4)$$

$$\text{where } \Delta w = \begin{cases} A^+ \exp\left(\frac{\Delta t}{\tau^+}\right), & \Delta t < 0 \\ -A^- \exp\left(\frac{-\Delta t}{\tau^-}\right), & \Delta t \geq 0 \end{cases} \quad (5)$$

If $w_{new} > g_{max}^E$, then $w_{new} = g_{max}^E$. On the other hand if $w_{new} < 0$, then $w_{new} = 0$. The factors (A^+ , A^-) correspond to the max synaptic change possible for potentiation and depression respectively at any given time step. The E-STDP parameters are set as: $A^+ = 0.005$ nS and $g_{max}^E = 0.3$ nS. The factor $\beta = |A^- \tau^-| / |A^+ \tau^+|$ which controls the relative amounts of depression to potentiation during learning is set 1.05 that represents a slight bias toward depression (Song et al., 2000). The initial excitatory synaptic weight w was set by picking values randomly in the interval (0, 0.1 nS) for synapses in layers 1 and 2 and was set between (0, 0.2 nS) for synapses between layers 2 and 3.

INHIBITORY STDP

The I-STDP function modulates the inhibitory synaptic weight z (Vogels et al., 2011, 2013; Srinivasa and Jiang, 2013) based on the timing difference Δt between the spike times of corresponding

pre- and post-synaptic neurons (Figure 1C). The synaptic weight is computed as:

$$z = z_{old} + \Delta z \quad (6)$$

The change Δz is governed by the following equations:

$$\Delta z = \begin{cases} B^+ * \exp\left(\frac{-|\Delta t|}{\tau}\right), & \text{if } |\Delta t| \leq \tau \\ -B^- * \exp\left(\frac{-|\Delta t|}{\tau}\right), & \text{if } |\Delta t| > \tau \end{cases} \quad (7)$$

If $z_{new} < 0$ then $z_{new} = 0$. On the other hand, if $z_{new} > g_{max}^I$ then $z_{new} = g_{max}^I$. The I-STDP parameters are set as $B^+ = 0.0015$ nS and $B^- = 0.0003$ nS, $g_{max}^I = 0.2$ nS and $\tau = 10$ ms. The initial inhibitory synaptic weight set by picking values randomly in the interval (0, 0.1 nS) for all synapses.

INTERPLAY BETWEEN E-STDP AND I-STDP FOR BALANCED CURRENTS

Excitatory and inhibitory long-term plasticity are both important, as it is the interplay between these two effects that results in a network with a balance between excitatory and inhibitory currents at each neuron in the reservoir layer. The networks with such a current balance are referred to as *balanced networks* (Vogels et al., 2011; Srinivasa and Jiang, 2013). Networks without inhibitory STDP fail to reach this state for any of a large set of possible network parameters. Figure 1D shows a schematic description of how these two STDP functions combine to create a balanced network.

The inhibitory STDP function is symmetrical supporting an increase in synaptic conductance, i.e., synaptic inhibition, for closely timed pre- and post-synaptic spikes regardless of their order. In contrast, the excitatory STDP function is anti-symmetric and biased toward depressing action. Together, for each of these two STDP functions along the $\Delta t = t_{pre} - t_{post}$ timeline, there are four qualitative regions: proximal causal and anti-causal (Figure 1D), for those spikes that occur relatively close together, and distal causal and anti-causal (Figure 1E), for those that occur farther apart.

INPUT IMAGE ENCODING AND NOISE INJECTION

Each 2-D input image pattern is first converted into a 1-D vector (Figure 2A). The 1-D input image vectors are then converted into spike sequences by an encoding process as follows. The neurons in the input layer are modeled using a Poisson process and each neuron receives an input from one pixel in the image. If a pixel is black in the input image, the neuron is assigned a mean firing rate of $f = 90$ Hz and if it is white, 10% of the source layer neurons with white pixels are assigned a mean firing rate of $f = 10$ Hz to simulate noise in the image. The spike encoding process is generated based on Poisson statistics. Assuming a sampling rate of dt and for a mean firing rate of f Hz for a given pixel, f spikes are generated every $1/dt$ samples. Thus, the probability of spiking at each time step for a given pixel firing at f Hz is $f * dt$. Spike trains are generated for each pixel based on its probability of spiking at each source layer neuron. An example result of this encoding process for input patterns (Figure 2A) is shown in Figure 2B. In all simulations, $dt = 1$ ms as mentioned earlier.

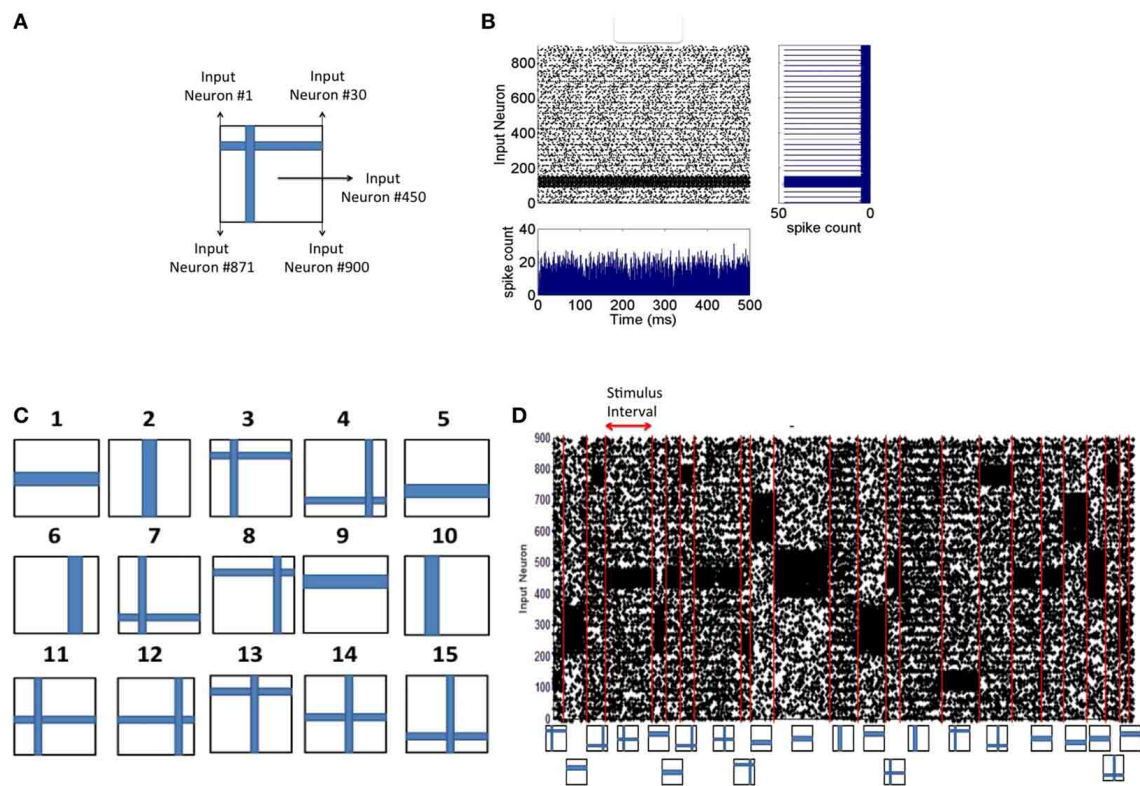


FIGURE 2 | Spike input encoding at the source layer of the network.

(A) An input pattern in the form of 2-D array of pixels is presented to the source neurons of the network that are linearly arranged. See Materials and Methods section for details of spike encoding. (B) The spikes generated during the presentation of the example in (A) for 500 ms duration is shown here. The firing rate at any given time (the bottom subplot) is approximately the same showing no undue bias introduced in the spikes that are fed to the source neurons. The linear arrangement of the source neurons results in the spike frequency plot (shown on the right) for the given input pattern

computed for a duration of 500 ms. (C) Input patterns of the training set consists of $P = 15$ "flag" patterns where each pattern is a binary image array of 30×30 pixels. (D) An example sequence of input patterns after spike encoding at the source neurons is shown here. The duration of the presentation of each pattern varies and is chosen from an exponential distribution with a mean of 30 ms. Red lines demarcate the shift from one training pattern to another. The training patterns are selected from the set of 15 in a random order. The plot shows a total duration of 1800 ms with 10% of the source neurons injected with noise throughout the sequence.

INPUT PATTERN PRESENTATION DURING TRAINING AND TESTING

The training process consists of presenting each input pattern in the training set (Figure 2C) in a random order for a duration drawn from an exponential distribution with a mean of 30 ms (Figure 2D). The network is tested for discriminability at regular intervals (every 10 s) during which synaptic plasticity in the network is turned off. Each input pattern is presented during the testing process in a fixed sequence for d seconds each and the discriminability index is then computed based on the generated firing rate codes (as described below). The process of estimating d is also provided below.

FIRING RATE CODE FOR READOUT NEURONS

The firing rate code for the readout neurons are evaluated only during the testing phase during which each input pattern from the training set is presented to the network for a duration of d seconds for a total duration of d^*P seconds for P patterns. Each pixel in the input image stimulates one neuron in the source layer (Figure 3A). The source neurons are modeled as Poisson spike sources as described above. For each test pattern p , the firing

rates f_i^p of sink neuron i in layer 3 (Figure 3B) can be computed as the total number of spikes emitted during a duration of d seconds. The maximum firing rate f_{max}^p is then estimated from the firing rates of all sink neurons for that test pattern p . The firing rate vector S^p of length M for pattern p is composed of components S_i^p for each sink layer neuron i can be computed as:

$$S_i^p = \begin{cases} 2, & \text{if } 0.9 \leq \frac{f_i^p}{f_{max}^p} < 1.0 \\ 1, & \text{if } 0.4 \leq \frac{f_i^p}{f_{max}^p} < 0.9 \\ 0, & \text{if } \frac{f_i^p}{f_{max}^p} < 0.4 \end{cases} \quad (8)$$

The vector S^p is referred to as the firing rate code and in this example it is a *tertiary* firing rate code (i.e., $C = 3$) because each sink neuron can have three possible states for a given input pattern p . An example of this tertiary code for two different input patterns is shown in Figure 3C. It is possible to use other coding levels such as binary ($C = 2$) or quarternary ($C = 4$) codes. In this paper,

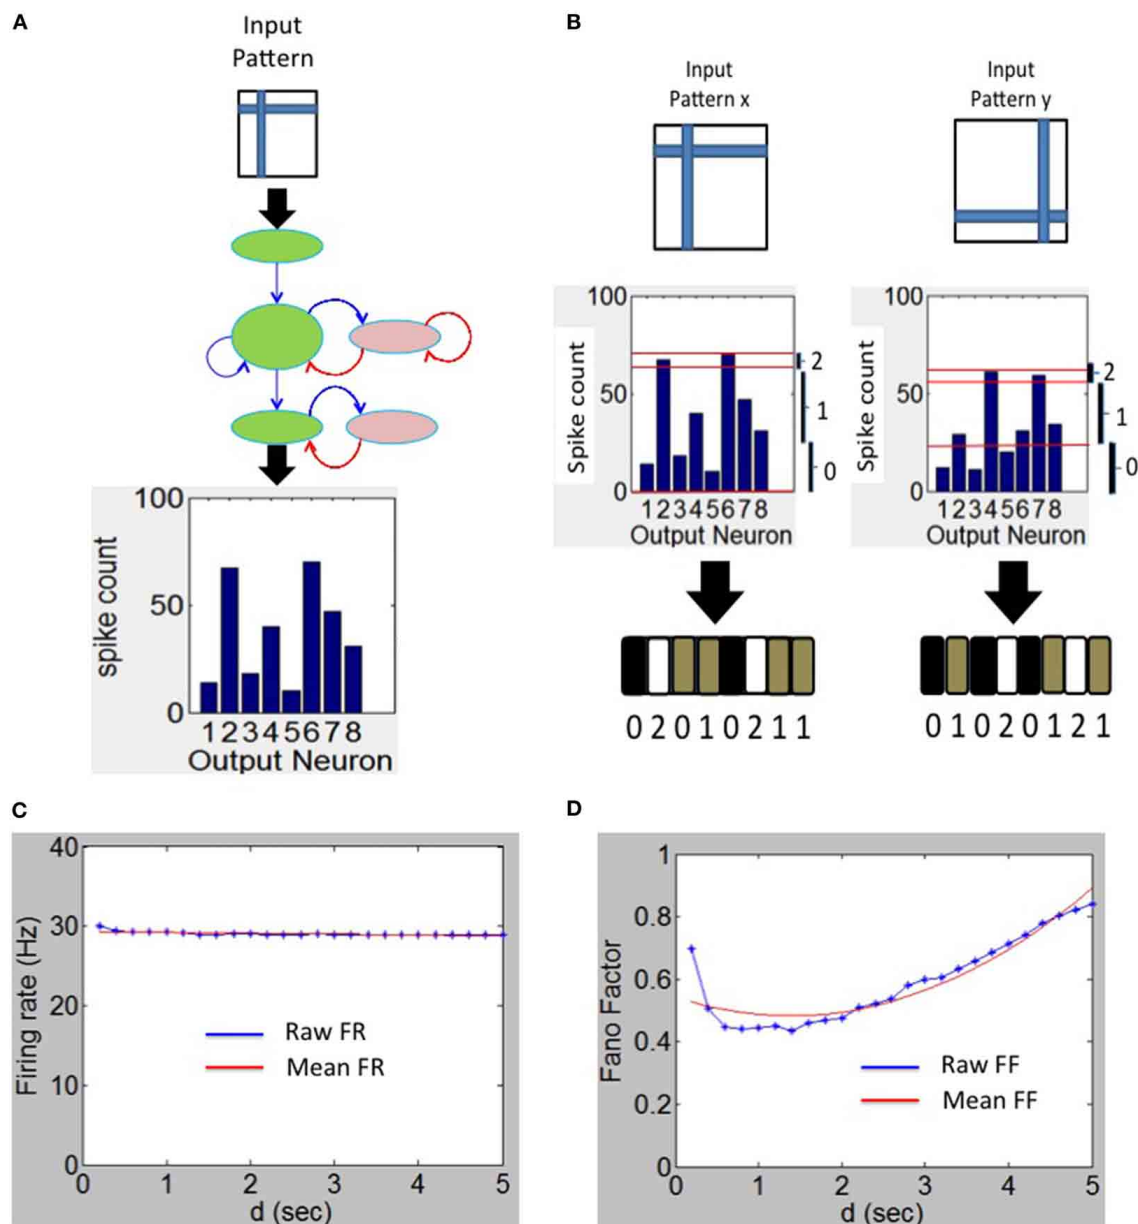


FIGURE 3 | Spike output decoding at the output layer of the network. (A) The output spike histogram is shown for neurons at the sink layer (assumed to have 8 neurons in this example) for a given input pattern presented at the source neurons and (B) this histogram is converted to firing rate code S^p with three possible states for each neuron (black—0

gray—1 and white—2). The firing rate code for two different input patterns results in two different codes as shown here. (C) The mean firing rate of the output neurons during the experiment to determine the duration d using the Fano Factor. (D) The variation of the Fano Factor as a function of the duration d .

$C = 3$ is used as it offers the highest discriminability as explained in the Results section.

ESTIMATING d FOR TESTING

The firing rate code S (above) and the estimation of discriminability index (as explained below) depends upon the duration d of each test pattern presentation. To estimate an appropriate duration d , the Fano factor (Churchland et al., 2010; Eden and Kramer, 2010) was computed from the spikes generated by the

readout neurons by assessing the relationship between variability of the spike counts and duration d .

The Fano factor (FF) is defined as the ratio of sample variance to sample mean of spike counts observed in a time window and the quality of the estimator strongly depends on the length of the window. The FF measures the noise-to-signal ratio and therefore characterizes the neural variability over trials. For example, for a Poisson process, the variance equals the mean spike count for any length of the time window. If the FF has a minimum at some value

of d , this can be an optimal value for d since the firing rate code would be robust at that value (see e.g., Ratnam and Nelson, 2000; Chacron et al., 2001).

The FF was computed for various durations d as follows. The spikes for each test pattern presented for duration of a selected d was first collected for each of the M readout neurons separately. This was repeated for 100 trials to collect a set of $M \times 100$ spike count values. The mean and variance in the spike count was then computed from these values. The ratio of the computed variance to mean gives FF for the selected d and for the selected test pattern. This process was repeated for all remaining $P-1$ test patterns and the resulting average FF was used as the FF for a given duration d . Since the average firing rate of the sink layer neurons was steady between 28 and 30 Hz (Figure 3C), the mean-matching procedure for FF (Churchland et al., 2010) was not used. To estimate the appropriate duration d , the FF was plotted as a function of duration d (Figure 3D) for $M = 8$ and $P = 15$. The minimum FF is ~ 0.48 and occurs at $d = 1.4$ s. We set $d = 1.4$ s in all our simulations.

DISCRIMINABILITY INDEX COMPUTATION

During the learning process, as input patterns are presented, a firing rate code S^p can be computed at the sink layer for each pattern p presented to the source layer as described above. The ternary firing rate code changes as the network is presented with more inputs. This implies that the ternary code cannot be directly used for reliably separating one pattern from another. However, after a few pattern presentations, the ability of the network to discriminate between the patterns becomes stable and reliable.

To verify this, a discriminability index (DI) was computed as follows. At regular intervals (once every 10 s) the network was stopped to probe the state of the network. During this process, the synaptic weights are frozen and each pattern is presented J times for duration of $d = 1.4$ s each. For a given pattern p , the firing rate code S^p was computed for each of the J presentations of p . A *prototype firing rate code* was selected for a given pattern p as the code that is most repeating among the J codes generated. If there are no repeats, one of the J codes as the prototype was selected at random. This process is repeated for each pattern to identify a prototype firing rate code for each input pattern. Using the prototype firing rate codes, the *inter-pattern distance* (D_{inter}) was computed. D_{inter} is defined as the average pair-wise distance between prototype readout codes computed from all possible unique pairs of prototype readout codes generated by the network for a given test set. To calculate D_{inter} , the distance d_i^{pq} between a pair of S codes for two input patterns p and q and for each sink neuron i was computed as:

$$d_i^{pq} = \begin{cases} 0 & \text{if } S_i^p = S_i^q \\ 1 & \text{if } S_i^p \neq S_i^q \end{cases} \quad (9)$$

The distance $D_{inter,i}$ was then computed by using d_i^{pq} for every pair of input patterns p and q for each sink neuron i across all test patterns P as:

$$D_{inter,i} = \frac{\sum_{k=1}^{P-1} \sum_{j=k+1}^P \sum_{i=1}^M d_i^{kj}}{P*(P-1)/2} \quad (10)$$

The maximum value of $D_{inter,i}$ for a readout code can be estimated as follows. Assuming a ternary readout code at the sink layer (i.e., $C = 3$ and that P is odd, the maximum pairwise distance between the readout code at each sink layer neuron i is obtained when the readout is equiprobable with “0” for one third of P input patterns, “1” for another third of inputs and with “2” for the remaining third. The theoretical maximum value of the numerator Equation (10) can be computed as $P^*P/3$ and thus $D_{inter,i}^{max}$ can be computed as $2P/(3*(P-1))$. If P is even, $D_{inter,i}^{max}$ can be similarly computed as $2(P+1)/3*P$. Similarly, for a binary code (e.g., $C = 2$) $D_{inter,i}^{max}$ can be computed to be $(P+1)/(2*P)$ when P is odd and $2P/(3*(P-1))$ when P is even. Thus, $D_{inter,i}^{max}$ can be computed for the general case when C is even as:

$$D_{inter,i}^{max} = \begin{cases} \frac{P(C-1)}{C(P-1)} & \text{if } P \text{ is even} \\ \frac{(P+1)(C-1)}{CP} & \text{if } P \text{ is odd} \end{cases} \quad (11)$$

Similarly $D_{inter,i}^{max}$ for the general case when C is odd can be expressed as:

$$D_{inter,i}^{max} = \begin{cases} \frac{P(C-1)}{C(P-1)} & \text{if } P \text{ is odd} \\ \frac{(P+1)(C-1)}{CP} & \text{if } P \text{ is even} \end{cases} \quad (12)$$

The expression for the inter-pattern distance D_{inter} can be written in terms of $D_{inter,i}$ as:

$$D_{inter} = \sum_{i=1}^M D_{inter,i} \quad (13)$$

By substituting $D_{inter,i}^{max}$ from Equations (11) or (12) (depending upon whether C are even or odd respectively) into Equation (13), the theoretical maximum value of D_{inter} can be computed. For example, if C is even, D_{inter} will be $PM(C-1)/(C(P-1))$ if P is odd and $(P+1)M(C-1)/CP$ if P is even. Thus, if $M = 8$, $C = 2$ and $P = 15$, the theoretical maximum for D_{inter} will be 4.28. It should also be noted that the theoretical maximum for D_{inter} grows linearly with M . The theoretical maximum for D_{inter} will serve as the upper bound for a given set of parameters during learning. This is because there is noise in the network that prevents an equiprobable distribution of readout codes by the network.

An *intra-pattern distance* (D_{intra}) was also computed by presenting the same pattern J times for d seconds each. D_{intra} is defined as the average pair-wise distance between readout codes same as Equation (10) computed from all possible unique pairs of readout codes generated by the network for the same input pattern. This distance provides a measure of an average variation in the response of readout neurons for the same input pattern. This variation can be caused due to noise in the inputs. It should be noted that $J = 10$ in all our simulations.

The discriminability index (DI) is then defined as a product of two measures. The first is called *separability*, ε , that measures the degree of separation of readout codes for a given test set. This measure can be computed as:

$$\varepsilon = 1 - \frac{D_{intra}}{D_{inter}} \quad (14)$$

This measure is akin to computing the Fischer metric (McLachlan, 2004). A small D_{intra} relative to D_{inter} implies that the network can separate the inputs well. Separability is independent of M .

The second measure is called the *uniqueness*, γ , that is defined as the number of unique readout codes produced by the network relative to maximum possible number of unique readout code. This can be expressed as:

$$\gamma = \frac{\#S}{P} \quad (15)$$

where $\#S$ refers to the total number of unique readout codes for a given test set of size P . Uniqueness is dependent on M since high dimensional readout codes generate more unique codes (Kanerva, 1988). The discriminability index (DI) is then computed as:

$$DI = \varepsilon * \gamma \quad (16)$$

High values of DI correspond to readout codes that have a low D_{intra} combined with high D_{inter} or high separability as well as a high uniqueness. The maximum value of DI is 1.0 and its minimum value is typically zero unless $D_{intra} > D_{inter}$. DI is dependent on M since uniqueness is dependent on M (see Appendix for an example calculation of DI).

SYNAPTIC DISTANCE COMPUTATION

In order to analyze for stability of learned codes, the *synaptic distance* was computed to track the synaptic changes between layers of the network for excitatory synapses. Since the E-STDP plasticity rule used in this paper is of the additive type, the resulting distribution of synapses after learning is bimodal in nature (Song et al., 2000). This bimodal distribution is due to competition that occurs among synapses at each neuron. The synapses that cause the post-synaptic neuron to fire more frequently will potentiate to the maximum synaptic weight while the other uncorrelated synapses will depress to a zero. To calculate the synaptic distance, the synaptic weights w_{ij} are converted into a binary weight W_{ij} where $W_{ij} = 1$ if (i.e., $w_{ij} > 0.7 * g_{max}^E$) and $W_{ij} = 0$ otherwise. The synaptic distance $\phi_{kl}(t_1, t_2)$ between excitatory synapses from layer k to layer l at time t_1 with the same synapses at time t_2 can be expressed as:

$$\phi_{kl}(t_1, t_2) = \frac{\sum_{j=1}^{\#l} \sum_{i=1}^{\#k} |W_{ij}(t_2) - W_{ij}(t_1)|}{\#k * \#l} \quad (17)$$

where $\#k$ and $\#l$ correspond to the number of neurons in layer k and l respectively and binary $W_{ij}(t)$ corresponds to the i th synapse in the k th layer that is connected to the j th synapse in layer l at time t . For example, $\#k = K$ for layer 1 and $\#l = N$ for layer 2 in the network. The synaptic distance is the total Hamming distance between the binary weights at two different time steps (t_1, t_2) where $t_2 > t_1$.

In addition to computing the synaptic distance, a shuffled synaptic distance was computed as a control to compare the synaptic weight changes during learning to those that could arise from chance. This distance $\phi_{kl}^{shuffled}(t_1, t_2)$ between excitatory synapses from layer k to layer l at time t_1 with the same

synapses at time t_2 can be expressed as:

$$\phi_{kl}^{shuffled}(t_1, t_2) = \frac{\sum_{k=1}^{\#shuffles} \sum_{j=1}^{\#l} \sum_{i=1}^{\#k} |W_{ij}^{shuffled}(t_2) - W_{ij}(t_1)|}{\#shuffles * \#k * \#l} \quad (18)$$

where $\#shuffles$ is the total number of shuffles that $W_{ij}^{shuffled}(t_2)$ undergoes at time t_2 . In all simulations, $\#shuffles = 10$. By combining the above two measures, a relative synaptic distance measure $\phi_{kl}^{rel}(t_1, t_2)$ can be expressed as:

$$\phi_{kl}^{rel}(t_1, t_2) = \frac{|\phi(t_1, t_2) - \phi^{shuffled}(t_1, t_2)|}{\phi^{shuffled}(t_1, t_2)} \quad (19)$$

If $\phi_{kl}^{rel}(t_1, t_2)$ is closer to 1.0, then $\phi_{kl}(t_1, t_2) \ll \phi_{kl}^{shuffled}(t_1, t_2)$ and that implies that the distance between the synaptic weights at time t_1 and t_2 is very small compared to chance. This implies that the learning has stabilized in the network.

RESULTS

An initial training set was constructed composed of $P = 15$ flag patterns (Figure 2C). The patterns are presented in random order for a duration selected from an exponential distribution with a mean of 30 ms. Each pattern generates a Poisson spike train (Figure 2D) at the source neurons (see Materials and Methods). These spikes generated by each pattern in the input layer are transmitted to the E reservoir neurons in the middle layer for further processing.

BALANCE OF EXCITATION AND INHIBITION DURING LEARNING OF RECEPTIVE FIELDS

As the patterns are presented, STDP in both excitatory (w) and inhibitory (z) synapses helps to achieve a good balance of excitation and inhibition currents in the network (Figure 4A). The synaptic weights w strengthens and creates an imbalance in synaptic currents due to inputs from the source neurons. At the same time the synaptic weights z gets rapidly potentiated due to I-STDP where inhibition increases irrespective of the order of occurrence of pre- and post-synaptic spikes for small timing differences between pre- and post-synaptic spikes. This results in a rapid compensatory increase in inhibitory currents into the neurons effectively preventing the neurons from exceeding V_T more often (Vogels et al., 2011; Srinivasa and Jiang, 2013). Thus, the network is provided with very brief windows of opportunity to learn. These small windows of opportunity correspond to the transients in the balanced current (i.e., brief excursions of net currents above zero as shown by the top plot of Figure 4A). The excitation due to the input pattern is sufficient to overcome inhibition momentarily before inhibition and its modulation via inhibitory plasticity compensates for any discrepancy in the current balance.

As the training patterns are presented under the balanced current regime, synapses between the source layer neurons and each E neurons of the reservoir in layer 2 collectively form receptive fields (Song et al., 2000; Srinivasa and Jiang, 2013). This

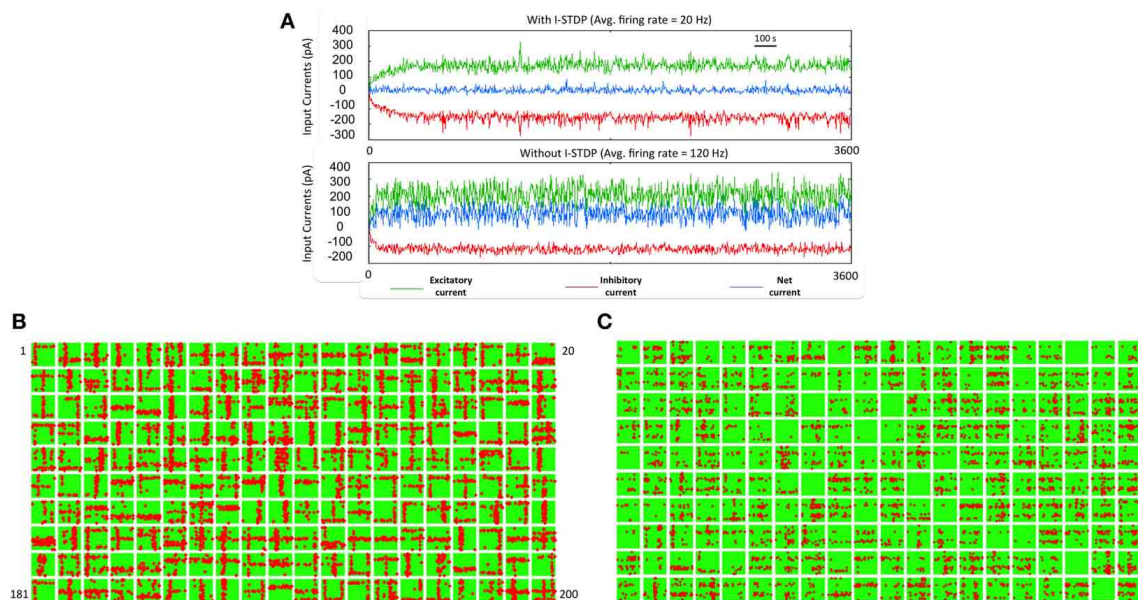


FIGURE 4 | Robust learning of receptive fields requires a balanced

current regime. (A) The sub-plot on top shows the total excitatory (green), and inhibitory (red) currents in the reservoir with I-STDP turned on. The net current (blue) is close to zero with several minor transients from zero. The average firing rate of the network is 20 Hz. If the I-STDP is turned off and with fixed inhibitory synaptic conductance, as shown in the second sub-plot, the total excitatory current is much higher than the total inhibitory current resulting in an unbalanced network with a very high average firing rate of 120 Hz. It may be possible to find a suitable synaptic conductance to achieve a balanced network but the use of I-STDP enables a self-organized process for achieving the balance. **(B)** The synapses between source neurons and the *E* neurons of the reservoir form a

bimodal distribution where most of the synapses are weak with a few strong synapses. This process makes the connectivity between source and reservoir sparse compared its initial connectivity of 40%. Each box shows the synapses in a 30×30 image format from the source neurons to each of the 200 *E* neurons in the reservoir. The red dots within each box correspond to synapses that are greater than $0.7 \cdot g_{max}$ while the rest are shown as green dots. Each box is a receptive field of an *E* neuron in the reservoir and the set of all boxes forms a basis set for the training set. **(C)** The receptive fields with I-STDP turned off are not well-defined as with I-STDP being on. Many receptive fields are not formed at all while many others are have features that do not reflect any structure found in the input patterns.

is achieved by adjusting the strength of these synapses via E-STDP. The red dots within a box (in **Figure 4B**) represent strong synapses between source neurons and an *E* neuron in the reservoir after 1 h of training. This process of synaptic strengthening is incremental and occurs using aggregates of input samples.

When the excitatory synapses alone obeyed the STDP rule and the inhibitory synapses were fixed (i.e., $z = \text{const}$) the excitatory and inhibitory currents are not balanced anymore (**Figure 4A**). There were also many *E* neurons that had no strong synapses. The strong synapses that emerge within each box (**Figure 4B**) appear to have a vertical or horizontal stripe (or both) and resemble the features of the input patterns in the training set. The learning of the receptive fields is also influenced by recurrent connections (i.e., $E \rightarrow E$, and $I \rightarrow I$) as well as mutual connections between the *E* and *I* populations (i.e., $E \rightarrow I$, $I \rightarrow E$) within the reservoir. All excitatory and inhibitory synapses within the reservoir are modified by E-STDP and I-STDP respectively. In order to assess the effect of turning off I-STDP, the receptive fields were analyzed after 1 h of training. Here the inhibitory weights z were randomly initialized between 0 and 1 but fixed throughout the simulations. The receptive fields did not have large variations in connectivity compared to the case where I-STDP was on **Figure 4C**.

The connection strengths for synapses in the reservoir after the presentation of the training set for duration of 1 h shows

that the synapses between any inhibitory pre-synaptic neuron and either *E* or *I* post-synaptic neuron are mostly strong (**Figure 5A**) and the synaptic strengths are distributed in a unimodal fashion (**Figure 5C**). However, the $E \rightarrow E$ synapses between layer 1 and layer 2 are sparse with few strong synapses ($\sim 5\%$ of all synapses). This discrepancy is primarily because E-STDP is anti-symmetrical while I-STDP is symmetrical. In other words, E-STDP is order dependent while I-STDP is not. Thus, I-STDP can potentiate synapses for both proximal causal and anti-causal spikes (**Figure 1D**). However, E-STDP will only potentiate the synapses if the pre-synaptic spike is causal to the post-synaptic spike. This implies that probability of potentiation is much higher for inhibition compared to excitation thus resulting in a bimodal distribution of synaptic strengths with many strong inhibitory synapses.

The synapses between the *E* neurons in the reservoir and the *E* neurons in the sink layer are also modified due to E-STDP (**Figure 5B**). The *E* neurons in the sink layer serve as *readout* neurons (Buonamano and Maass, 2009; Buzsáki, 2010). The distribution of synaptic strengths is unimodal (**Figure 5C**) unlike other $E \rightarrow E$ synapses described above. This is because of the following reason. Initially the strength of the synapses between the reservoir and the readout neurons is small (i.e., between 0 and 0.2). The connectivity between them is also sparse (i.e.,

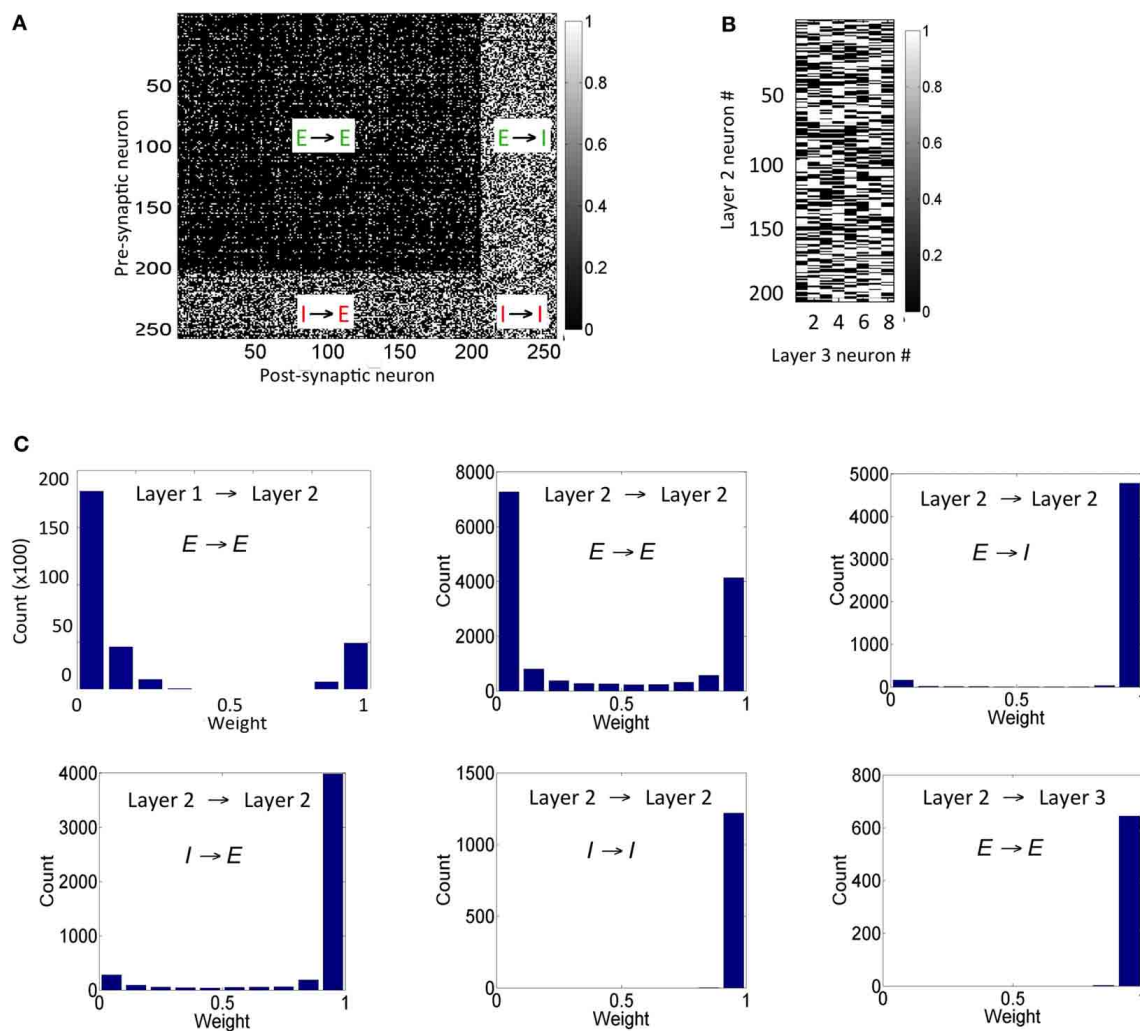


FIGURE 5 | Synaptic states for the network after 1 h of training are shown here. (A) The synaptic weights within the reservoir are shown here. **(B)** The synapses between layer 2 and layer 3 are shown here. **(C)** The distribution of synapses between the E neurons in layer 1 and layer 2 as well as between E neurons within layer 2 is bimodal in nature. The distribution of

synapses between E and I neurons with other I neurons within layer 2 is unimodal. The distribution for synapses between E neurons in layer 2 and E neurons in the layer 3 is also unimodal (see text for details). It should be noted that the weights in these plots are normalized with respect to g_{max}^E and g_{max}^I as appropriate.

$c_{23}^{EE} = 40\%$). Any given input at the source neurons causes a sequence of spiking activity in the reservoir neurons that is signaled to the readout neurons. This sequence of spiking activity among the reservoir neurons is hereinafter referred to as a *neural trajectory*.

Since the readout neurons are driven to fire by the reservoir neurons and by no other means, the temporal causality for their spiking is always from the reservoir neuron to the readout neuron. This results in the strengthening of synapses allocated to a readout unit (due to E-STDP). As the readout neurons fire in response to neural trajectories in the reservoir, all the synapses from the reservoir neurons to the readout neurons strengthen to its max value resulting in a unimodal distribution (Figure 5C). This is unlike the interaction between the source and reservoir neurons or within reservoir neurons where the spiking activity is driven by both feed-forward and lateral connections.

LEARNING TO DISCRIMINATE PATTERNS

In order to assess the pattern discrimination capability, the training set was presented to the network for a total of 3600 s (see Materials and Methods). The firing rate of the readout neurons was monitored after every 10 s of training to test the network's ability to discriminate the input patterns. At these time intervals, plasticity was turned off and each input pattern was presented in a random order for 5 s and the DI metric Equation (16) was computed based on the activity of the readout neurons. The ternary code from the readout codes was plotted for each input pattern at regular intervals during the course of training (Figure 6A). The readout code initially looks alike for all patterns since the network has not really been exposed to all the patterns. As the network is exposed to more training data, the readout codes begin to show more variations. However, when the readout codes for an input pattern at two different times are compared, they appear to

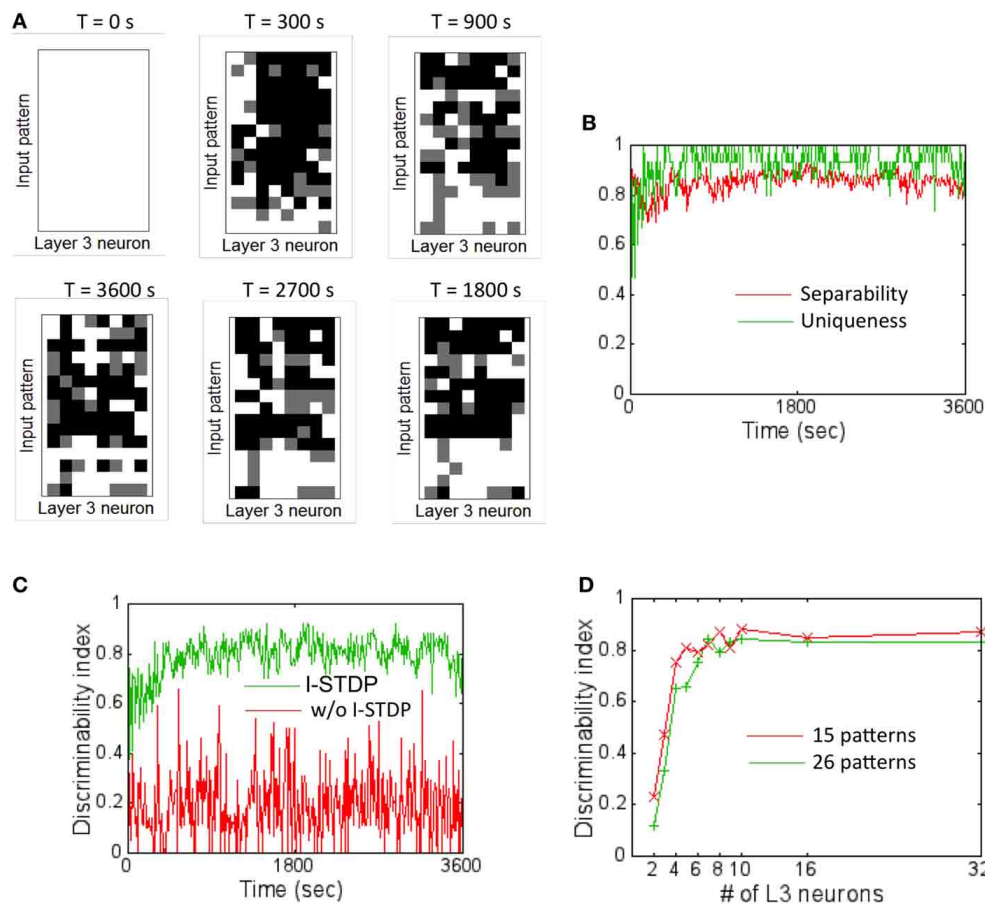


FIGURE 6 | Firing rate codes and DI computation. (A) The firing rate codes were extracted from the spiking activity of the readout neurons (assuming 8 readout neurons). These codes were compared between input patterns at regular intervals during the training process. The inter-pattern readout codes are constantly changing (compare top row of readout code corresponding to input pattern #1) suggesting that the discrimination is very poor if readout codes are compared in an absolute fashion. (B) A plot of separability and uniqueness generated during the first hour of training with 15 input patterns is shown here for the case with I-STDP turned on. The two metrics are consistently high implying that the readout codes are highly separable as well

as very unique. (C) A plot of DI for the case when I-STDP is turned on (green) compared to when it is turned off (red). This DI is computed by multiplying the separability with uniqueness (see text for details). (D) The DI is plotted against various number of readout neurons in the sink layer to assess the minimum number of readout neurons required to achieve high discriminability for two different sizes of training sets. While it is possible to achieve a $DI \sim 0.8$ for the 15 pattern case with only 5 readout neurons, higher number of patterns require more readout neurons. Since the total number of input patterns to be tested is 26 in this paper, a total of 8 readout neurons (or $M = 8$) was assumed for all simulations.

change constantly throughout the duration of the training period. This implies that a static template based discrimination algorithm would not be appropriate here.

The separability and uniqueness were tracked during the first hour of training (Figure 6B). Since the receptive fields form early due to E-STDP in the balanced regime created by the regulatory actions of I-STDP, good separability Equation (14) and uniqueness Equation (15) occur early during the training process. The DI Equation (16) thus rapidly rises to ~ 0.8 (Figure 6C) implying very good discriminability. The DI is however highly unstable and averages to ~ 0.2 when I-STDP is turned off. This is because the average firing rate of the network reaches 120 Hz (Figure 4A). This high firing rate results in poorly formed receptive fields (Figure 5B). This in turn results in very unstable separability and low uniqueness (not shown). Thus, the discriminability is poor and unstable (Figure 6C) when I-STDP is turned off.

The network size is potentially large enough to learn to discriminate many more patterns than used in the training set. However, the number of readout neurons was limited to the minimum required for obtaining good discriminability. To determine the minimum number of readout neurons for the chosen size of training set, the DI was averaged across 10 trials (Figure 6D) for two different sizes of training set: one with 15 patterns and the other with 26 patterns. While for the 16 pattern case, the DI rises to a value close to 0.8 with just four readout neurons, the network requires eight readout neurons to produce an average DI of ~ 0.8 .

When DI is computed using codes other than the ternary code (see Materials and Methods), it was worse (Figure 7A). This was unexpected since the number of possible states for the readout neuron should grow (i.e., 3^8 states for ternary code vs. 5^8 for quinary code) with number of states in the code. The main reason for this unexpected result is because D_{intra} also grew at a faster rate

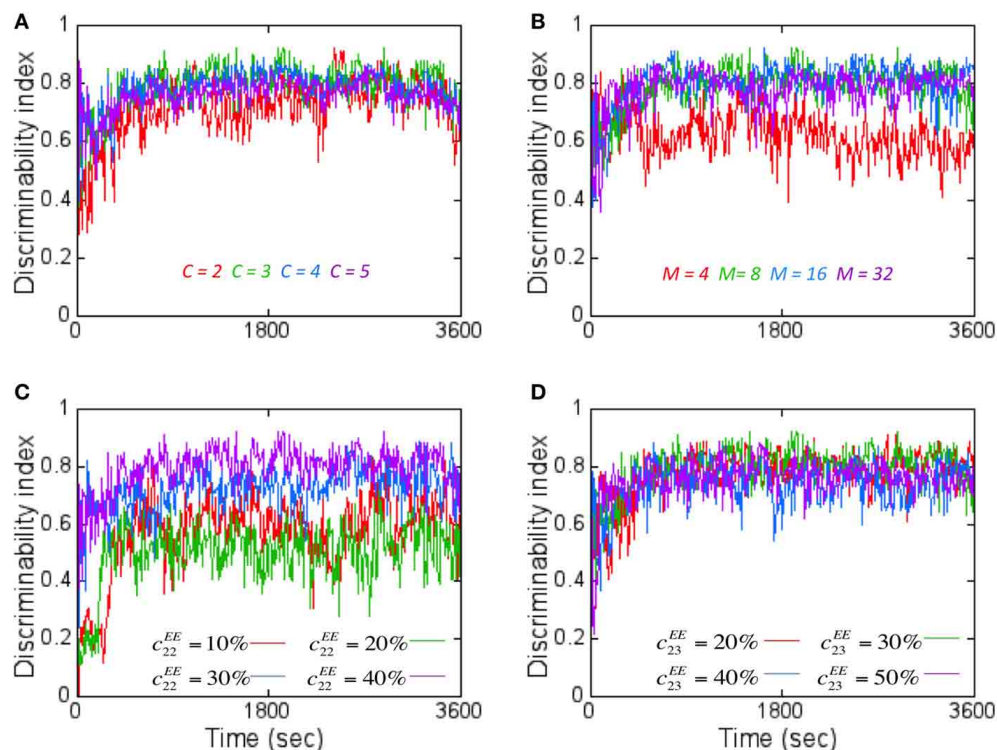


FIGURE 7 | The effect of various network parameters on DI is shown here. (A) The coding level C for the readout code was varied to assess its effects on DI assuming $M = 8$. A coding level $C = 2$ (red) corresponds to just two states (On or Off) for the readout neurons while tertiary (green), quaternary (blue), and quinary (purple) codes correspond to 3–5 states respectively. $C = 3$ (green) produced the best average DI score. **(B)** Increasing the number of readout neurons M affects the DI . The worst average DI was for $M = 4$ while the best average DI was

close for $M = 8, 16$, and 32 . $M = 8$ was chosen for reasons mentioned in **Figure 6**. Here a ternary code ($C = 3$) is assumed. **(C)** The effect of increasing the lateral connectivity within the reservoir neurons does improve the DI metric but the degree of improvement starts to diminish with connectivity beyond 40%. Here $M = 8$ and $C = 3$ is assumed. **(D)** The effect of increasing the connectivity between E neurons from layer 2 to layer 3 results in only a marginal improvement in DI values. Here $M = 8$ and $C = 3$ is assumed.

with higher code values compared to D_{inter} that grew at a slower rate with higher code. Thus, the separability reduces with higher code thus reducing DI .

When the number of readout neurons M was increased, the DI does change (**Figure 7B**). The effect of connectivity on DI was studied by tracking $E \rightarrow E$ connections within the reservoir in layer 2 as well as between layer 2 and layer 3. Adding more connections within the reservoir (i.e., c_{22}^{EE}) improved DI (**Figure 7C**). Similarly, on average, the DI improved when the number of connections between layer 2 and layer 3 (i.e., c_{23}^{EE}) is increased (**Figure 7D**). All simulations primarily used $C = 3$, $M = 8$ and with $c_{22}^{EE} = 40\%$ and $c_{23}^{EE} = 30\%$.

IMPORTANCE OF LEARNED CONNECTIVITY AND FIRING RATE CODE FOR DISCRIMINATION

In order to assess the effect of learned connectivity due to STDP on the discrimination ability of the network, a set of control experiments were performed. During each testing step (at a sampling interval of 10 s) the connectivity between three layers of the network was shuffled in six different ways while maintaining both the synaptic strengths and the total number of synaptic connections intact compared to the network with learned connectivity (or the original network). In the first case, shuffling was

performed between connections between the source layer and the E neurons of the reservoir only. This was accomplished by randomly assigning pre-synaptic neurons from the source layer to post-synaptic E neurons in the reservoir that were different from the learned connections. In the second case, the learned connections between neurons within the reservoir (irrespective of whether they were E or I neurons) were randomly shuffled. In the third case, only the connections from I neurons in reservoir to all other connections (irrespective of whether they were E or I neurons) were randomly shuffled. In the fourth case, only the connections from E neurons in reservoir to all other connections (irrespective of whether they were E or I neurons) were randomly shuffled. In the fifth case the connections between the E neurons in the reservoir and the E neurons in the output layer were randomly shuffled. In the final case, shuffling was performed between all the layers—that is a mixture of shuffling performed for first through fifth cases.

For each of these cases, the DI was computed (after every 10 s for a period of 1 h) by averaging the score after random shuffling 10 times in each case. The DI was worse for the first case compared to the original network (**Figure 8A**). In this case, swapping the connections between the E neurons in the source layer and the E neurons in the reservoir results in disturbing the

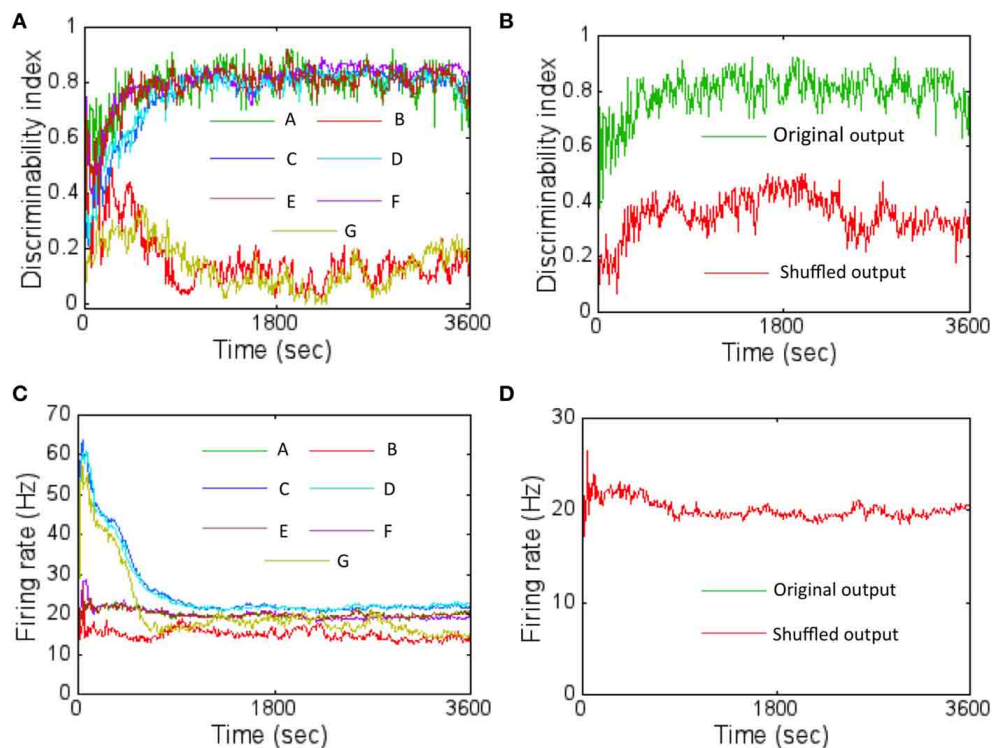


FIGURE 8 | The DI was computed with shuffled network connectivity to assess the importance of the learned connectivity for pattern discrimination. (A) The results from the six control experiments (B: layer 1 → 2 synapses shuffled shown in red; C: layer 2 (*E* or *I*) to layer 2 (*E* or *I*) shown in blue; D—layer 2 (*I*) to layer 2 (*E* or *I*) shown in cyan; E—layer 2 (*E*) to layer 2 (*E* or *I*) shown in brown; F—layer 2 (*E*) to layer 3 (*E*) shown in magenta; G—all these connections shuffled simultaneously) are compared against the original network with learned connectivity (A) shown in green. The DI scores represent an average obtained after shuffling 10 times for each control experiment. The plot shows that altering the connectivity by shuffling the connections between layer 1 to *E* neurons in layer 2 affects discriminability severely while shuffling connections within layer 2 does not (see text for more details). **(B)** A control experiment was also performed to assess the importance of the specific locations of the

firing activity within the readout code. This was achieved by first randomly shuffling the components of the readout code 10 different times. The *DI* scores were then computed for each case and then averaged. The results shows that the *DI* score is much lower (red trace) compared to the original network (green trace) suggesting that the locations of the firing activity within the readout code caused due to learning is also very important. **(C)** The average firing rate of the network for the various control experiments in (A) is shown here. **(D)** The average firing rate for the control experiment in (B) is shown here. The two firing rates overlap completely (red and green overlap completely). This is because there is no change in the connections between layer 2 and layer 3 neurons after STDP potentiates all of them over time early within the first hour. Furthermore, all connections go to its max value (i.e., unimodal distribution). So, swapping the outputs does not make a difference to the firing rates.

learned receptive fields. This implies that the learned receptive fields between the source layer *E* neurons and the *E* neurons in the reservoir are very important for pattern discrimination in this network. However, interestingly, the learned connections between the neurons in the reservoir (the second through fourth cases) or the connections between *E* neurons in the reservoir and the *E* neurons in layer 3 (the fifth case) did not have a major effect in the *DI* scores.

For the second through fourth cases, the effect of shuffling the connections within the reservoir did not affect the *DI* score very much irrespective of whether the connections were from *E* or *I* neurons. The lateral connectivity between *I* → *I* or between *I* → *E* neurons only serves to regulate the balance of currents. Furthermore, all the synapses from *I* neurons are fully potentiated (i.e., unimodal distribution as shown in Figure 5C). Similarly, the most of the connections from *E* → *I* neurons are also fully potentiated (Figure 5C) while most of the *E* → *E* connections are weak (bimodal distribution as shown in Figure 5C). So swapping

connections between strong *I* connections (for the third case) or between weak *E* → *E* or between strong *E* → *I* connections (for the fourth case) does not seem to change the network performance much as well. Since the second case is a combination of the third and fourth cases, the result is similar. For the fifth case, the connections between *E* neurons in the reservoir and the *E* neurons in layer 3 become fully potentiated (Figure 5C). So, swapping connections between strong *E* connections does not affect the *DI* score. Since the sixth case includes the first case as a subset, the *DI* score is severely affected much like in the first case.

The average firing rate of the network was initially most affected for the third and sixth control experiments while the first control experiment reduced the average firing rate of the network (Figure 8C) relative to the original network. For the third case, initially swapping the connections between *I* → *I* or between *I* → *E* affects the current balance as the synapses have not had the time to fully potentiate to its peak values. This in turn can affect the firing rates by causing them to be high. However,

as learning proceeds, the balance is restored due to I-STDP causing these inhibitory synapses to become fully potentiated (Figure 5C) and the firing rates fall as expected (Figure 8C). Since the sixth case contains the third as a subset, it follows the same trend as the third. For the first case, the firing rates fall below the original network because swapping the connections between the E neurons in the source layer and the E neurons in the reservoir results in disturbing the learned receptive fields. Thus, on average, the E neurons in the reservoir do not have good matches with the input patterns because of the shuffling resulting in lower average firing rates.

In order to assess the importance of location of the active nodes in the firing rate code on the discrimination ability of the network, another control experiment was performed. The firing rate code was shuffled by shifting the location of the active nodes in the sink layer and this was repeated 10 times. Once every 10 s, the DI was computed with the shuffled firing rate code and then averaged to produce a DI that was compared against the DI generated by the original network. The shuffling of the active nodes results in much lower DI compared to the original network (Figure 8B). The firing rate of the network is unaltered by shuffling (Figure 8D) because the connectivity from E neurons in layer 2 to the E neurons in layer 3 has a unimodal distribution (Figure 5C) with all synapses being fully potentiated and thus being immune to the shuffling.

STABILITY OF LEARNING

The network was analyzed for stability of learning by studying the change in DI as more inputs were presented after 1 h of training.

To test this, the duration of presentation of the inputs was doubled from 1 to 2 h. During this time, the inputs were once again sampled at random from the training set and presented for a duration that was selected from an exponential distribution of 30 ms.

The readout codes for all patterns in the training set after 1 h and after 2 h were compared. The ternary codes for each pattern were compared and the codes do not match at all. This change is partially explained by subtle changes in the receptive fields (Figure 9A) compared to after 1 h (Figure 4A). The DI is very stable and hovers around 0.8 (Figure 9B) throughout the extra hour of training. To measure the change in synapses more precisely, the relative synaptic distance Equation (19) between layer 1 and layer 2 synapses after 1 and 2 h of training. This relative distance $\phi_{12}^{rel}(3600, t)$ was tracked once every 10 s from 1 to 2 h (Figure 9C). The plot shows that the distance slowly changes during the first hour and stabilizes to ~ 0.6 . This implies that the synaptic weight changes in a more meaningful fashion compared to changes due to pure chance. Furthermore, the rate of change of the relative synaptic distance during the hour is slow and thus implies a stable regime of synaptic adaptation.

The selectivity of the readout neurons to a subset of the reservoir neurons emerges from E-STDP based on pattern of firing in the reservoir. Once the selectivity is established for the readout neurons, it does not change very much during the second hour of training. This is evident from the observation that the synaptic distance between the two layers does not change (Figure 5C) and that all the synapses to readout neurons become fully potentiated with a unimodal distribution of synaptic strengths. The

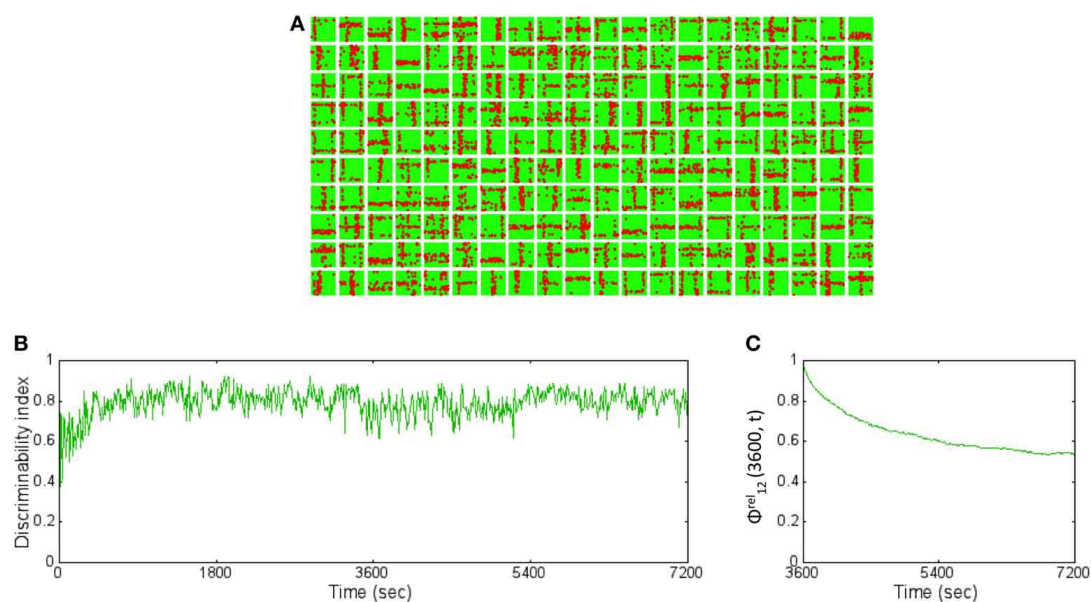


FIGURE 9 | The basis vectors and basis set learning after 2 h of training on initial training data set. (A) The synapses between the source neurons and each E neuron in the reservoir form receptive fields (similar to the ones shown in Figure 4B) but slightly modified after 2 h of training compared to after 1 h of training. **(B)** The DI is steady throughout the extra 1 h of training hovering around 0.8. This implies that the learning has stabilized causing the

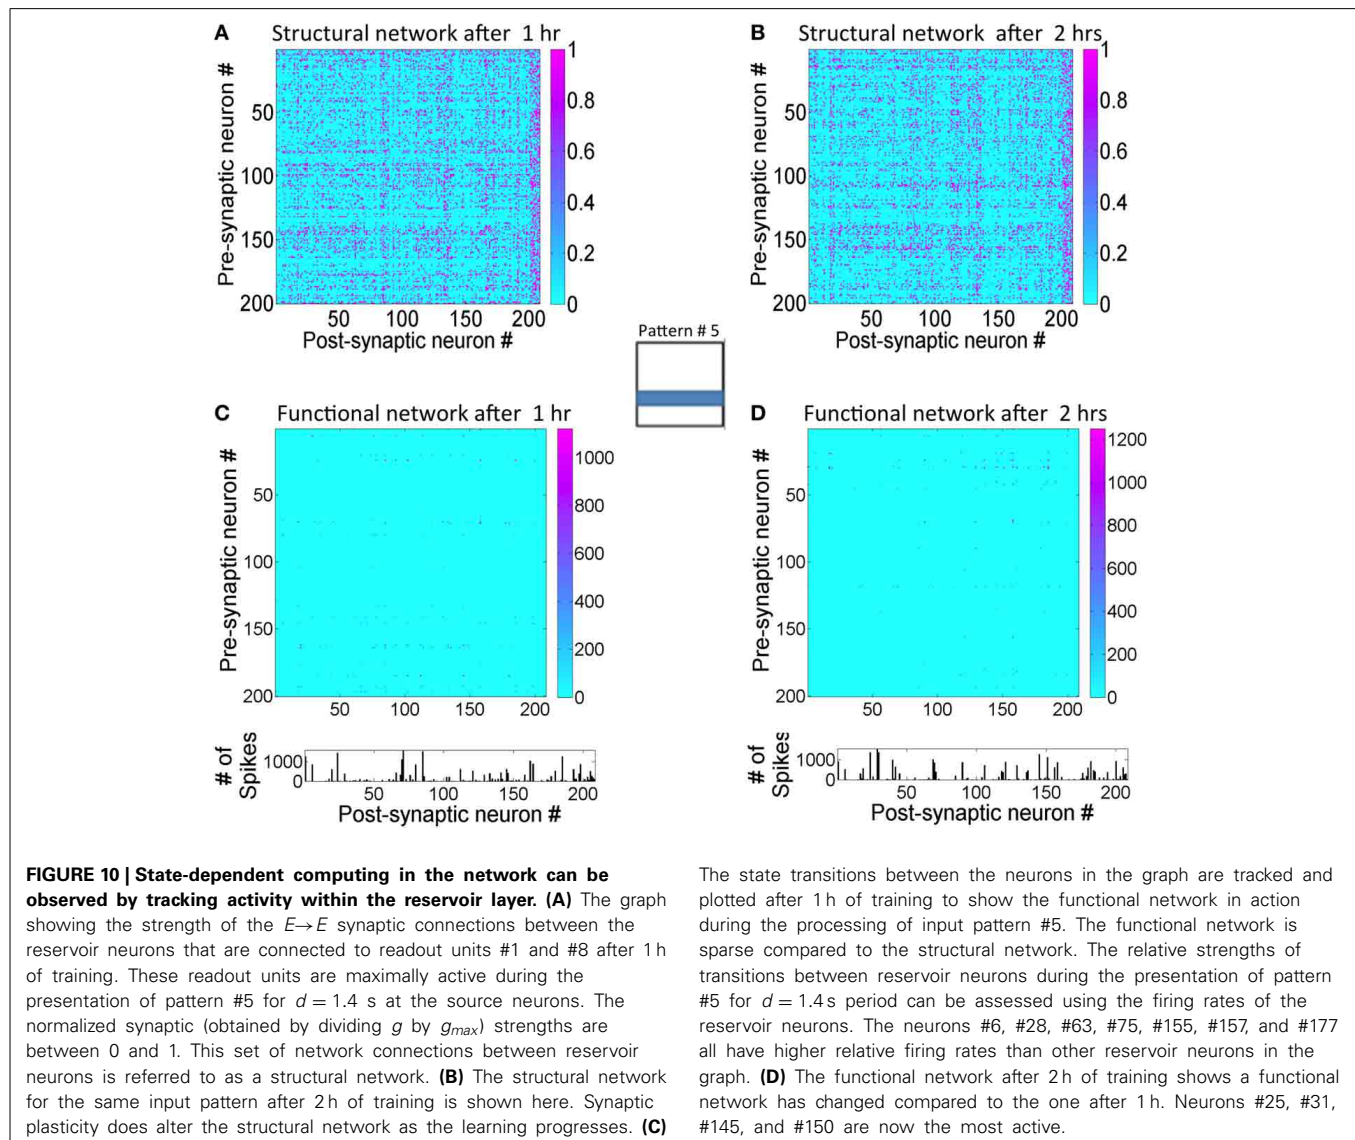
discriminability to be stable as well. **(C)** The relative synaptic distance between $E \rightarrow E$ synapses from source layer neurons to the reservoir neurons was compared to randomly shuffled synapses or synapses formed due to chance. The slope of the distance trace slowly decreases suggesting stability in learning while the final value of 0.6 suggests that the learning stabilizes the network to a state that is far from chance (see text for further details).

synaptic strengths between the reservoir neurons and the readout neurons are unimodal by the end of 1 h of training due to E-STDTP and continue to remain stable during the extra hour of training. This is reflected in $\phi_{23}^{rel}(3600, t)$ that is not defined since the $\phi_{23}(3600, t) = \phi_{23}^{shuffled}(3600, t)$ for all t after 3600 s. This is because the synapses between the E neurons in the reservoir and the E neurons in the sink layer do not change (Figure 5C) for reasons explained earlier. This means that $\phi_{23}^{rel}(3600, t)$ is not defined for all t .

The lateral connectivity in the reservoir causes a state dependent firing regime (Buonamano and Maass, 2009). To understand this better for the proposed network, used a *back-trace* approach was adopted as follows. The readout neurons that fired with a maximum firing rate for a given input pattern were first selected. The neurons back-trace to the reservoir from these readout neurons were then identified. For example, the readout neurons #1 and #8 fired with a ternary code of 2 were first selected. The synaptic connections between the reservoir neurons

are represented in the graph (Figure 10A after 1 h and Figure 10B after 2 h). It should be noted that the strong connections between the reservoir neurons only depict the anatomical or structural aspect of the network. The resulting set of reservoir neurons and their connections between each other and the two readout units is referred to as a *structural* network. The strong connection between the E reservoir neurons are not necessarily unique to the selected readout neurons since other readout neurons that are active for other inputs may also be connected to some of the same reservoir neurons found in the structural network. Similarly, the readout neurons are not unique to the input pattern since the code in the sink layer is distributed. This means that the same readout neuron could fire as part of another readout code that represents a different input pattern.

When the structural network is tracked temporally for the duration of input presentation (i.e., for $d = 1.4$ s), the network dynamics shows that a select subset of reservoir neurons fire in a complex spatiotemporal sequence. A state transition graph can



be plotted from this firing sequence (**Figure 10C**). This graph is referred to as the *functional* network that is very sparse when compared to the structural network. After further training for an additional hour, the functional network for the same pattern changes (**Figure 10D**). Here the transition frequencies between some of the neurons are somewhat reduced compared to the functional network after 1 h. This implies that the network is able to sharpen the neural trajectory further with training.

The state-transitions at the reservoir combined with stable connections between the reservoir and output neurons means that the ternary code at the readout neurons will change based on these transitions in the reservoir. These changes in the neural trajectory cause the ternary code in the readout units to be different even for identical inputs. This implies that repeatable readout neuron activity (in response to input patterns) is not achievable in this network. However, the *relative* codes between a pattern and the rest as computed by the *DI* are stable (**Figure 9B**).

PLASTICITY TO NEW INPUTS AFTER INITIAL LEARNING

In order to study the capacity of the network to learn new inputs, a second training set with new input patterns was added to the initial training data set (**Figure 11A**). The network that was trained with the initial training data for 1 h was presented with the both old and new inputs for an additional hour. The receptive fields show slow adaptation to the new features while retaining features

from the old patterns as well (**Figure 11B**). For example, the learned receptive for reservoir neuron #4 (fourth box from the left on top row in **Figure 11B**) at the end of 2 h shows a diagonal set of synaptic connections that was absent after training the network with the initial training set (**Figure 4B**). On the other hand, the neuron #18 (third box from the right on the top row in **Figure 11B**) remains very similar to the receptive field learned after 1 h (**Figure 4B**).

The network was also studied for their ability to learn new patterns when they were presented with only new patterns during the second hour of training. This test was more stringent than the first experiment above and provides a more precise picture of the network's ability to retain past information while learning new information. The changes in the receptive fields reflect rapid re-learning with adaptations to features found in the new inputs (**Figure 11C**) after 2 h of training compared to receptive fields after training for 1 h with the initial training set (**Figure 4B**). For example many of the receptive fields show diagonal and circular features. It is noteworthy that very few receptive fields now reflect the initial training set.

The *DI* was computed for the above two cases of training. When the network is presented with both old and new patterns during the second hour of training, the *DI* was very stable relative at ~ 0.8 when tested on all the 26 patterns after 2 h of training relative to the first hour of training (the green trace in **Figure 11D**). This implies that the network is able to discriminate both the old

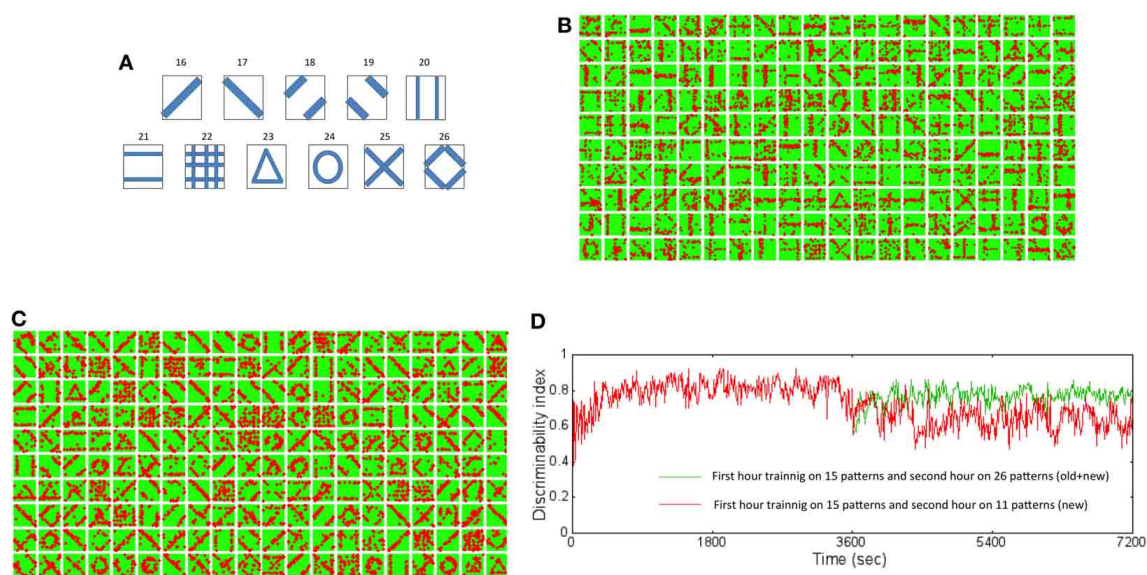


FIGURE 11 | New training set data set and resulting receptive fields formed due to learning for two different training regimes. (A) The new training data set is composed of 11 new patterns not in the original data set. These patterns were added to the original data set and then used for training the network to test the ability of the network to learn new information without forgetting old information. **(B)** The new receptive fields formed by training for an additional hour with both old and new data after initial training on the original data set for an hour is shown here. The new receptive fields show newly learned features that incorporate features such as a diagonal line (for example, neuron #4 and #18). **(C)** The new receptive fields formed by training for an additional hour with only new data after initial training on the

original data set for an hour is shown here. The new receptive fields change dramatically from the original set (see **Figure 4B**) with mostly features that reflect the new patterns and very little from the old patterns. **(D)** The *DI* was compared for the two training regimes. The *DI* was retained at a high value of 0.8 (green trace) when the network was exposed to both old and new patterns in the second hour of training. The network however was found to have a lower *DI* value of 0.62 (red trace) on average suggesting forgetting of old information. However, the interesting aspect is that *DI* decreases gradually suggesting that the network does not lose the ability to discriminate between old patterns or between old and new patterns abruptly or “catastrophically” but in a more graceful manner.

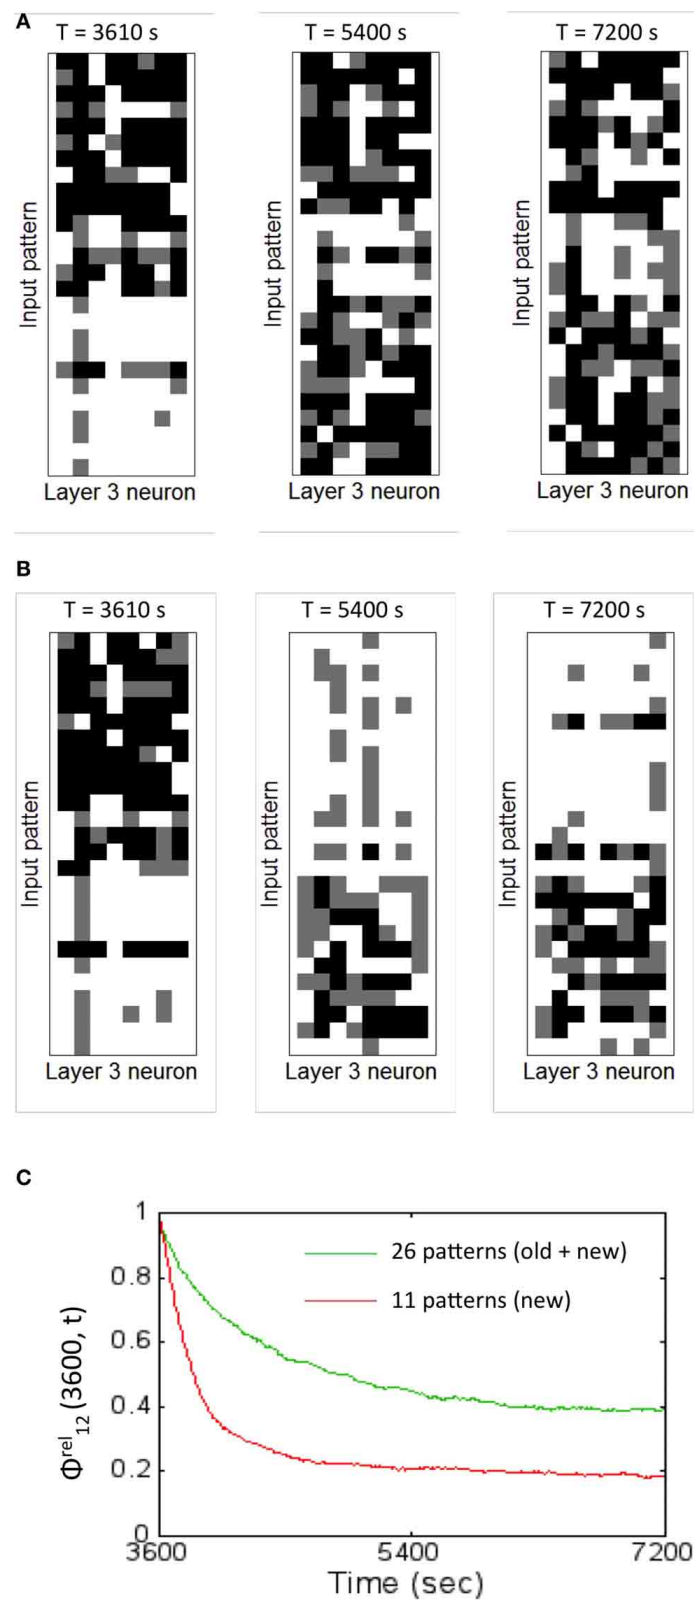


FIGURE 12 | The readout codes generated by the network for two different training regimes shows the network is plastic to new inputs while also being stable to old information. (A) The readout codes

generated for the network when trained on all the 26 patterns during the second hour of training. The testing was performed with all the 26 patterns.

(Continued)

FIGURE 12 | Continued

The readout code is washed out in the beginning for the new patterns but slowly is assimilated by the network generating a rich readout code for all the 26 patterns after 2 h. The resulting *DI* is stable (as shown in **Figure 11**) suggesting robust discrimination. **(B)** The readout codes generated for the network when trained on only on the 11 new patterns during the second hour of training. The testing was performed with all the 26 patterns. The readout code is washed out in the beginning for the new patterns but the network rapidly learns the new features suggesting the highly plastic nature of the network. The readout code starts to wash out for the 15 old patterns after 2 h. The resulting *DI* is

still strong enough (as shown in **Figure 11**) suggesting slow forgetting of old information. **(C)** The relative synaptic distance between $E \rightarrow E$ synapses from source layer neurons to the reservoir neurons was compared to randomly shuffled synapses or synapses formed due to chance for the two training regimes are shown. The slope of the relative distance trace slowly decreases (green trace) for the case where the old and new patterns are presented suggesting stability in learning. However, the slope changes more dramatically (red trace) for the case when the network is trained only on the new patterns. The final value of 0.5 (old + new) and 0.2 (for new only) suggests that the learning stabilizes the network to a state that is different from pure chance.

and new patterns after 2 h of training. In comparison, when the network was trained only with new patterns, the *DI* falls to a lower value of ~ 0.6 (the red trace in **Figure 11D**). This implies that the network is not as discriminatory as in the first case implying that the network forgets. However, the network does not exhibit catastrophic forgetting (French, 1994). Catastrophic forgetting occurs when the network abruptly (i.e., in a few time steps) and completely (i.e., with very poor discrimination) forgets previously learned patterns in exchange for learning new patterns. Since the *DI* only degrades from ~ 0.8 after 1 h to ~ 0.6 after an additional hour of training only with new patterns, the network does not abruptly forget the old patterns. This graceful and slow degradation in *DI* shows that the network gradually forgets previously learned information but not catastrophically.

The readout codes for the two cases provide some more insight into the network performance. The network begins at the same starting point (i.e., after training with 15 patterns for 1 h). The readout codes are very different for the two cases after 2 h. When the network is trained with both old and new patterns, the readout code for the 26 patterns appears with much more variations (**Figure 12A**) compared to the case when trained only with new patterns. In the latter case, the network appears to have more washed out codes for the old patterns compared to the new patterns (**Figure 12B**). This confirms that the network forgets the old patterns in the second case compared to the first case. This is to be expected to some extent because the network is plastic and is expected to learn the new inputs as it experiences that more than the old patterns. However, it is noteworthy that the network does not exhibit catastrophic forgetting as discussed above. This shows that the network is able to assimilate the old information along with the new information to create a new readout codes such that the resulting *DI* is sufficient for discrimination between all patterns (old and new) at least for some time (in this case for about an hour). Understanding how this could be extended is a subject for future study.

The relative synaptic distance $\phi_{12}^{rel}(3600, t)$ was tracked once every 10 s from 1 to 2 h between the receptive fields for the two cases of training. The distance changes slowly for the case when the network is presented with old and new patterns (the green trace in **Figure 12C**). In comparison, the weight changes are far more drastic (the red trace in **Figure 12C**). Here the rate of change is steep implying that the network undergoes sharp changes during the early learning phase in the second hour but then stabilizes to a non-zero value. This implies that the network undergoes synaptic changes due to learning driven by STDP based on new

training data as opposed to changes due to pure chance. The synapses between the *E* neurons in the reservoir and the *E* neurons in the sink layer do not change (similar to **Figure 5C**) for reasons explained earlier. This means that $\phi_{23}^{rel}(3600, t)$ is not defined for all *t*.

In order to understand how this occurs, the lateral connectivity of the graph was analyzed. A pattern from the new data set was selected for analysis (**Figure 13**) and presented to the source neurons for $d = 1.4$ s. Since the first hour of training is based on the initial training set, the network was never exposed to this new pattern before. The readout neurons (#4 and #8) that fired maximally for the new pattern was identified and its structural network was identified at the 1 h mark (**Figure 13A**).

The structural network selected by these two neurons changes between 1 and 2 h when the network is trained with both old and new patterns (**Figure 13B**). When the functional network was extracted after the first hour of training, the network dynamics shows that many reservoir neurons are accessed in a complex spatiotemporal sequence (**Figure 13C**). This is because the network attempts to process the unknown input pattern using as many receptive fields as possible. Once the network is trained for an hour more with both old and new patterns, the network is able to readout using a relatively sparse neural trajectory (compare **Figures 13C,D**). This is because the basis set adapts to incorporate new features in the new input pattern data set during the second hour of training. This modifies the functional network due to changes in firing rates of some neurons in the graph (see in **Figures 12C,D**) as well as changes in the functional network. The functional network for the case when trained only on new patterns in the second hour is similar to **Figure 12D** except that the spatiotemporal trajectories in the network are biased toward the new inputs as opposed to the network that is exposed to both the old and new patterns.

DISCUSSION

The interaction between E-STDP and I-STDP enables spiking neuronal networks to learn to discriminate patterns in a self-organized fashion. A hallmark of self-organizing systems is a composition of relatively “dumb” units connected together and constrained by “interaction dominant dynamics” (Ihlen and Vereijken, 2010). In the case of the simulated network presented here, the connection strength between neurons is always altered by synaptic plasticity, effectively changing the network topology. The structure of the network is tuned so as to enable uncorrelated neurons that are randomly connected to become correlated

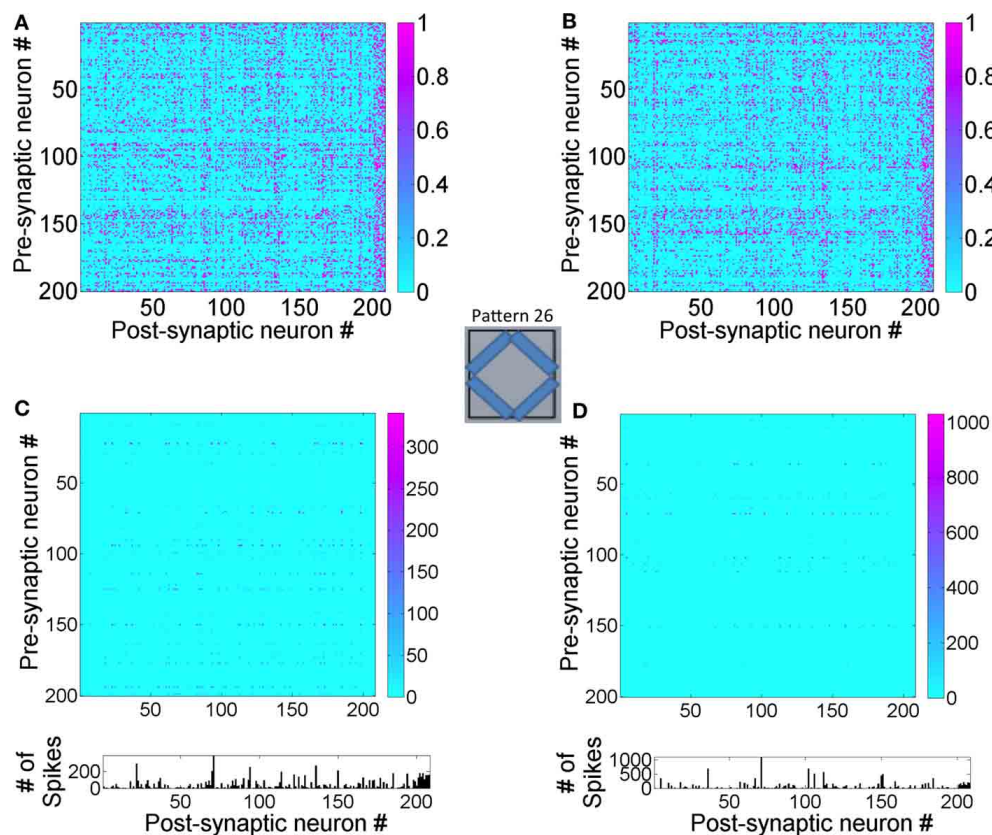


FIGURE 13 | State-dependent computing in the reservoir layer is affected by new training data. (A) The structural network connected to readout units #4 and #8 after 1 h of training. These readout units that are maximally active during the presentation of the new pattern #26 for $d = 1.4$ s. **(B)** The structural network for the combined training data of old and new input patterns after 2 h of training is shown here. Synaptic plasticity alters the structural network as expected. **(C)** The transitions

between $E \rightarrow E$ neurons in the reservoir after 1 h of training. This functional network shows firing activity at many more neurons and this is manifested by the novelty of the new pattern that the network has never processed before. **(D)** This functional network after 2 h of training is sparser compared to the one shown after 1 h since the network is able to learn the new pattern by adapting its receptive fields and generating new readout codes that are more discriminatory.

in a balanced way so as to produce meaningful network-level behavior.

Balance in the network is at the level of excitatory and inhibitory currents. These currents are observed to balance each other, leaving the resultant current near zero. In the present model, changing synaptic conductance via STDP for inhibitory and excitatory synapses helps achieve this balance. There are other models (Vogels et al., 2011; Srinivasa and Jiang, 2013) that also have explored the effects of interaction between these two types of STDP for memory formation and stability but have not explored the question of unsupervised discrimination of patterns.

In a recent model, the synaptic efficacy via short-term plasticity (Klompf and Maass, 2013) was used to achieve stable and balanced networks but that work also did not look into unsupervised pattern learning. In other biologically plausible networks, synaptic connections can be created or destroyed also known as structural plasticity (Leuner and Gould, 2010). But without that option, plasticity is left as the only possible mechanism for change within the network.

Networks without I-STDP fail to reach a balanced state for any of a large set of possible parameters. It is not only a practical

matter that inhibitory STDP is required, but there are deep connections to self-organizing systems as well. Self-organization is usually the result of two opposing effects. In the strong cases, these opposing effects are some mutually-referring function of each other (Nicolis and Prigogine, 1977; Witten and Sander, 1981). Here, excitatory and inhibitory STDP play these roles, and together produce various forms of compensatory feedback (Luz and Shamir, 2012). It should be noted that the obtained results are based on one of many possible I-STDP functions found in the brain (Vogels et al., 2013) and not all of them will necessarily result in a current balance. A recent article explored the distinct I-STDP window shapes in tuning neuronal responses (Kleberg et al., 2014). The exact role of each shape of I-STDP function on brain function remains to be explored in the future.

This memory trace or neural trajectories in the reservoir evolve both in space and time (Rabinovich et al., 2008; Buonomano and Maass, 2009; Buzsáki, 2010). It is known that discriminating between several trajectories requires complex mechanisms with many dedicated readers (Jortner et al., 2007; Masquelier et al., 2009; Buzsáki, 2010). Our approach proposes an algorithm for computing DI that can help discriminate between input

patterns but this is still not a biologically plausible mechanism. Using the *DI* metric to probe network dynamics shows that the proposed network can discriminate between patterns that form complex neural trajectories. Furthermore, this discrimination is not susceptible to catastrophic forgetting.

It should be noted the order of presentation of the input patterns to the network at the source neurons provides different contexts and changes the state-dependent firing patterns in the network. This causes the ternary code in the readout units for identical inputs to be different. However, the *relative* firing rate code as computed by the *DI* metric is invariant to the order of input pattern presentation (not shown).


The *DI* measure derived in this work is related to information theory. Information theory informs us about the amount of information a neural response (by sink layer neurons) carries about the stimulus (source layer neurons). In this theory, information is quantified using entropy measures. The *DI* measure captures the amount of information about the stimulus in the neural response and is thus closely associated to mutual information (Borst and Theunissen, 1999). In the *DI* measure, separability is closely linked to entropy measures. For example, D_{inter} is associated to noise entropy since it provides a measure of variability in neural response to the same input stimulus while D_{intra} is associated to response entropy as it measures variability in neural response to all the stimulus types presented to the network. Uniqueness on the other hand is directly linked to response entropy as it measures the variations in all possible neural responses to all possible stimulus conditions. This could be a useful future direction for further investigation.

The network design proposed in this paper for unsupervised discrimination has two key features that enable fault-tolerant properties in a manner similar to our previous work on self-supervised learning of spatiomotor transformations (Srinivasa and Cho, 2012). The first feature is the reduction of number of spiking neurons from layer 1 to layer 2 (i.e., $K > N$). This allows

the network to compress the input features into an encoding consisting of a smaller subset of neurons in layer 2. The absence of spiking activity from some input neurons can still be tolerated due to inputs from neighboring input neurons within the reservoir. The second is the recurrent STDP connections between neurons within layer 2. With this feature, the spiking activity due to neighboring neurons within layer 2 enables STDP to eventually strengthen synapses between the neurons that receive inputs from layer 1 and weaken those that do not receive any inputs from neurons in layer 1. This feature might provide robustness against the complete loss of spiking activity within neurons because recurrent connections within the reservoir might enable neurons that do not receive any feed-forward input from layer 1 to still propagate spiking activity to layer 3. Thus, the network could exhibit tolerance to complete loss of spiking activity in the input neurons.

It may be possible to extend the proposed architecture for unsupervised discrimination to learn using supervisory labels. Two possible mechanisms are briefly explored here. In the first case, each readout neuron in layer 3 can be stimulated by an externally provided spike train that corresponds to a label for the input pattern presented in layer 1 (Figure 14A). Here each readout neuron uniquely codes for one label thus requiring as many readout neurons as labels. This external input can be considered as the top-down (TD) input while the spikes from the reservoir neurons to the readout neurons can be considered as bottom-up (BU) inputs.

When an external label is available, it can cause TD stimulation of the appropriate readout neuron even if there are no BU inputs since it is assumed that these TD inputs are provided via very strong synapses (reflected by thick green arrows in Figure 14A). The TD inputs increase the inhibition to the readout neurons thus creating network dynamics in the readout akin to a winner take-all network. It should be noted that the inhibitory neurons in layer 3 is constantly stimulated by weak external spiking inputs (think



arrow to the red circle in **Figure 14A**). When the readout neurons are stimulated by both BU and TD inputs, the synapses from the reservoir neurons to the readout neurons could be strengthened due to E-STDP. At a later time, if the label is removed, the readout neurons can generate peak activity in the readout neuron corresponding to the correct label. But as observed before, the readout codes change slowly and constantly due to plasticity in the reservoir. This can result in misclassification errors. These errors can however be fixed by periodically providing the correct labels that can trigger STDP based learning to quickly correct the mistakes by retuning the synapses from the reservoir to the readout neurons.

In the second case, the readout neurons could represent a distributed code (**Figure 14B**). In this case, the joint spiking activity of the entire population of readout neurons represents a label for each input. This network can also be trained via labels at the appropriate readout neurons as described above and the learning process due to BU and TD stimulation of the readout neurons can cause the readout neurons to generate correct answers even when the teaching labels are removed. This network is however also susceptible to constant forgetting thus requiring periodic stimulation by the external supervisory sources to correct the mistakes made by the readout neurons in classifying input patterns.

The functional network response of the network to inputs via neural trajectories in the reservoir indicates the use of a distributed code with many reservoir neurons being activated during the input presentation. Normally the relative firing rates between reservoir neurons (i.e., the number of times a neurons fires relative to other neurons in the reservoir) is not high. However, in some cases a single neuron in the reservoir may exhibit a high relative firing rate. Thus, some neurons in the reservoir can encode for the entire input in some cases while at the same time require neural trajectories to encode other inputs. This flexibility is a hallmark of neural systems where single neurons (also known as *grandmother cells*) are known to encode for objects (Perrett et al., 1982; Rolls, 1984; Yamane et al., 1988; Quiroga et al., 2005) while there are other concepts that require distributed codes with complex neural trajectories (Rabinovich et al., 2008; Buzsáki, 2010).

The most interesting aspect of our network is that it is able to discriminate between old patterns already presented in an initial training session while also adapting to new patterns without losing its ability to discriminate among old patterns. The learned connectivity especially between the source layer neurons and the *E* neurons in the reservoir appears to be critical for this capability as exemplified from the control experiments. The ternary code appears to produce the highest *DI* metric compared to other codes since the inter-pattern distance is the lowest for the ternary code compared to the case with other codes. The network does not exhibit catastrophic forgetting and is more robust if exposed to old patterns occasionally during the learning of new patterns. Incorporating the means to achieve a homeostatic balance due to an interaction between inhibitory and excitatory STDP in all these networks may well enable self-organized discrimination of patterns while exhibiting the requisite dynamics to address the stability vs. plasticity dilemma.

ACKNOWLEDGMENTS

The authors acknowledge the support for this work by HRL Shared Research funding SR12201. We would like to thank the reviewers for their insightful comments that greatly helped improve the quality of the manuscript.

REFERENCES

- Bernacchia, A., Seo, H., Lee, D., and Wang, X. J. (2011). A reservoir of time constants for memory traces in cortical neurons. *Nat. Neurosci.* 14, 366–372. doi: 10.1038/nn.2752
- Bi, Q. Q., and Poo, M. M. (1998). Activity-induced synaptic modification in hippocampal culture, dependence on spike timing, synaptic strength and cell type. *J. Neurosci.* 18, 10464–10472.
- Borst, A., and Theunissen, F. E. (1999). Information theory and neural coding. *Nat. Neurosci.* 2, 947–957. doi: 10.1038/14731
- Brader, J. M., Senn, W., and Fusi, S. (2007). Learning real-world stimuli in a neural network with spike-driven synaptic dynamics. *Neural Comput.* 19, 2881–2912. doi: 10.1162/neco.2007.19.11.2881
- Buonamano, D., and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nat. Rev. Neurosci.* 10, 113–125. doi: 10.1038/nrn2558
- Buzsáki, G. (2010). Neural syntax: cell assemblies, synapse ensembles, and readers. *Neuron* 68, 362–385. doi: 10.1016/j.neuron.2010.09.023
- Chacron, M. J., Longtin, A., and Maler, L. (2001). Negative interspike interval correlations increase the neuronal capacity for encoding time-varying stimuli. *J. Neurosci.* 21, 5328–5343.
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13, 369–378. doi: 10.1038/nn.2501
- Crowe, D. A., Averbeck, B. B., and Chafee, M. Y. (2010). Rapid sequences of population activity patterns dynamically encode task-critical spatial information in parietal cortex. *J. Neurosci.* 30, 11640–11653. doi: 10.1523/JNEUROSCI.0954-10.2010
- Dan, Y., and Poo, M. M. (2006). Spike timing dependent plasticity: from synapse to perception. *Physiol. Rev.* 86, 1033–1048. doi: 10.1152/physrev.00030.2005
- Eden, U. T., and Kramer, M. A. (2010). Drawing inferences from Fano factor calculations. *J. Neurosci. Methods* 190, 149–152. doi: 10.1016/j.jneumeth.2010.04.012
- Freeman, W. J. (2001). *How Brains Make Up Their Minds*. New York, NY: Columbia University Press.
- Freeman, W. J., and Schneider, W. S. (1982). Changes in spatial patterns of rabbit olfactory EEG with conditioning to odors. *Psychophysiology* 19, 44–56. doi: 10.1111/j.1469-8986.1982.tb02598.x
- French, R. M. (1994). Catastrophic forgetting in connectionist networks: causes, consequences and solutions. *Trends Cogn. Sci.* 3, 128–135.
- Garcia-Sanchez, M., and Huerta, R. (2004). Neural networks with Hebbian learning do not outperform random ones in fan-out system. *Neurocomputing* 58, 337–342. doi: 10.1016/j.neucom.2004.01.064
- Grossberg, S. (1980). How does a brain build a cognitive code? *Psychol. Rev.* 87, 1–51. doi: 10.1037/0033-295X.87.1.1
- Grossberg, S. (2012). Adaptive resonance theory: how a brain learns to consciously attend, learn, and recognize a changing world. *Neural Netw.* 37, 1–47. doi: 10.1016/j.neunet.2012.09.017
- Hahnloser, R. H. R., Kozhevnikov, A. A., and Fee, M. S. (2002). An ultra-sparse code underlies the generation of neural sequence in a songbird. *Nature* 419, 65–70. doi: 10.1038/nature00974
- Ihlen, E. A., and Vereijken, B. (2010). Interaction-dominant dynamics in human cognition: beyond 1/f fluctuation. *J. Exp. Psychol. Gen.* 139, 436. doi: 10.1037/a0019098
- Jaeger, H., and Haas, H. (2004). Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science* 304, 78–80. doi: 10.1126/science.1091277
- Jortner, R. A., Farivar, S. S., and Laurent, G. (2007). A simple connectivity scheme for sparse coding in an olfactory system. *J. Neurosci.* 27, 1659–1669. doi: 10.1523/JNEUROSCI.4171-06.2007
- Kanerva, P. (1988). *Sparse Distributed Memory*. Cambridge, MA: MIT Press.

- Katchalsky, A., and Kedemo, O. (1962). Thermodynamics of flow processes in biological systems. *Biophys. J.* 2, 53–78. doi: 10.1016/S0006-3495(62)86948-3
- Kello, C. T. (2013). Critical branching neural networks. *Psychol. Rev.* 120, 230–254. doi: 10.1037/a0030970
- Klampafl, S., David, S. V., Yin, P., Shamma, S. A., and Maass, W. (2012). A quantitative analysis of information about past and present stimuli encoded by spikes of A1 neurons. *J. Neurophysiol.* 108, 1366–1380. doi: 10.1152/jn.00935.2011
- Klampafl, S., and Maass, W. (2013). Emergence of dynamic memory traces in cortical microcircuit models through STDP. *J. Neurosci.* 33, 11515–11529. doi: 10.1523/JNEUROSCI.5044-12.2013
- Kleberg, F. I., Fukai, T., and Gilson, M. (2014). Excitatory and inhibitory STDP jointly tune feedforward neural circuits to selectively propagate correlated spiking activity. *Front. Comput. Neurosci.* 8:53. doi: 10.3389/fncom.2014.00053
- Laje, R., and Buonomano, D. (2013). Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci.* 16, 925–936. doi: 10.1038/nn.3405
- Leuner, B., and Gould, E. (2010). Structural plasticity and hippocampal function. *Ann. Rev. Psychol.* 61, 111–140. doi: 10.1146/annurev.psych.093008.100359
- Long, M. A., Jin, D. Z., and Fee, M. S. (2010). Support for a synaptic chain model of neuronal sequence generation. *Nature* 468, 394–399. doi: 10.1038/nature09514
- Luz, Y., and Shamir, M. (2012). Balancing feed-forward excitation and inhibition via hebbian inhibitory synaptic plasticity. *PLoS Comput. Biol.* 8:e1002334. doi: 10.1371/journal.pcbi.1002334
- Maass, W. (2010). “Liquid state machines: motivation, theory and applications,” in *Computability in Context: Computation and Logic in the Real World*, eds N. Cooper and A. Sorbi (London: Imperial College Press), 275–296.
- Maass, W., Natschlager, T., and Markram, H. (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* 14, 2531–2560. doi: 10.1162/089976602760407955
- Markram, H., Gerstner, W., and Sjöström, P. J. (2011). A history of spike-timing-dependent plasticity. *Front. Synaptic Neurosci.* 3:4. doi: 10.3389/fnsyn.2011.00004
- Markram, H., Lubke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–215. doi: 10.1126/science.275.5297.213
- Masquelier, T., Guyonnet, R., and Thorpe, S. J. (2009). Competitive STDP-based spike pattern learning. *Neural Comput.* 21, 1259–1276. doi: 10.1162/neco.2008.06-08-804
- McLachlan, G. J. (2004). *Discriminant Analysis and Statistical Pattern Recognition*. Hoboken, NJ: Wiley Interscience.
- Minkovich, K., Thibault, C. M., Nogin, A., Cho, Y. K., O’Brien, M. J., and Srinivasa, N. (2014). HRLSim: a high performance spiking neural network simulator for GPGPU clusters. *IEEE Trans. Neural Netw. Learn. Syst.* 25, 316–331. doi: 10.1109/TNNLS.2013.2276056
- Nicolis, G., and Prigogine, I. (1977). *Self-Organization in Nonequilibrium Systems*. New York, NY: Wiley.
- Nikolic, D., Hausler, S., Singer, W., and Maass, W. (2009). Distributed fading memory for stimulus properties in the primary visual cortex. *PLoS Biol.* 7:e1000260. doi: 10.1371/journal.pbio.1000260
- Olyphar, A. V., Lytton, W. W., and Prinz, A. A. (2012). Input-to-output transformation in a model of the rat hippocampal CA1 network. *Front. Comput. Neurosci.* 6:57. doi: 10.3389/fncom.2012.00057
- Pascual-Leone, A., Amedi, A., Fregni, F., and Marabet, L. B. (2005). The plastic human brain cortex. *Ann. Rev. Neurosci.* 28, 377–401. doi: 10.1146/annurev.neuro.27.070203.144216
- Perrett, D. I., Rolls, E. T., and Caan, W. (1982). Visual neurons responsive to faces in the monkey temporal cortex. *Exp. Brain Res.* 47, 329–342. doi: 10.1007/BF00239352
- Quiroga, R., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature* 435, 1102–1107. doi: 10.1038/nature03687
- Rabinovich, M. I., Huerta, R., Varona, P., and Afraimovich, V. S. (2008). Transient cognitive dynamics, metastability, and decision making. *PLoS Comput. Biol.* 4:e1000072. doi: 10.1371/journal.pcbi.1000072
- Ratnam, R., and Nelson, M. E. (2000). Nonrenewal statistics of electrosensory afferent spike trains: implications for the detection of weak sensory signals. *J. Neurosci.* 20, 6672–6683.
- Rieke, F., Warland, D., van Steveninck, R. D. R., and Bialek, W. (1996). *Spikes: Exploring the Neural Code*. Cambridge MA: MIT Press.
- Rolls, E. T. (1984). Neurons in the cortex of the temporal lobe and in the amygdala of the monkey with responses selective for faces. *Hum. Neurobiol.* 3, 209–222.
- Song, S., Miller, K., and Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent plasticity. *Nat. Neurosci.* 3, 919–925. doi: 10.1038/78829
- Srinivasa, N., and Cho, Y. K. (2012). A self-organizing spiking neural model for learning fault-tolerant spatio-motor transformations. *IEEE Trans. Neural Netw. Learn. Syst.* 23, 1526–1538. doi: 10.1109/TNNLS.2012.2207738
- Srinivasa, N., and Jiang, Q. (2013). Stable learning of functional maps in self-organizing spiking neural networks with continuous synaptic plasticity. *Front. Comput. Neurosci.* 7:10. doi: 10.3389/fncom.2013.00010
- Sussillo, D., and Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron* 63, 554–557. doi: 10.1016/j.neuron.2009.07.018
- Swenson, R., and Turvey, M. (1991). Thermodynamics reasons for action-perception cycles. *Ecol. Psychol.* 3, 317–348. doi: 10.1207/s15326969eco0304_2
- Tiesinga, P., Fellous, J. M., and Sejnowski, T. J. (2008). Regulation of spike timing in visual cortical circuits. *Nat. Rev. Neurosci.* 9, 97–107. doi: 10.1038/nrn2315
- Tsodyks, M., and Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci. U.S.A.* 94, 719–723. doi: 10.1073/pnas.94.2.719
- Tsodyks, M., Pawelzik, K., and Markram, H. (1998). Neural networks with dynamic synapses. *Neural Comput.* 10, 821–835. doi: 10.1162/089976698300017502
- Van Rullen, R., Guyonnet, R., and Thorpe, S. J. (2005). Spike-times make sense. *Trends Neurosci.* 28, 1–4. doi: 10.1016/j.tins.2004.10.010
- Victor, J. D., and Purpura, K. P. (1996). Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J. Neurophysiol.* 76, 1310–1326.
- Vogels, T. P., Froemke, R. C., Doyon, N., Gilson, M., Haas, J. S., Liu, R., et al. (2013). Inhibitory synaptic plasticity: spike-timing dependence and putative network function. *Front. Neural Circuits* 7:119. doi: 10.3389/fncir.2013.00119
- Vogels, T. P., Rajan, K., and Abbott, L. F. (2005). Neural network dynamics. *Ann. Rev. Neurosci.* 28, 357–376. doi: 10.1146/annurev.neuro.28.061604.135637
- Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory Plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* 334, 1569–1573. doi: 10.1126/science.1211095
- Witten, T. Jr., and Sander, L. M. (1981). Diffusion-limited aggregation, a kinetic critical phenomenon. *Phys. Rev. Lett.* 47, 1400. doi: 10.1103/PhysRevLett.47.1400
- Woodin, M. A., Ganguly, K., and Poo, M. M. (2003). Coincident pre- and postsynaptic activity modifies GABAergic synapses by postsynaptic changes in Cl⁻ transporter activity. *Neuron* 39, 807–820. doi: 10.1016/S0896-6273(03)00507-5
- Yamane, S., Kaji, S., and Kawano, K. (1988). What facial features activate face neurons in the inferotemporal cortex of the monkey? *Exp. Brain Res.* 73, 209–214. doi: 10.1007/BF00279674

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 03 June 2014; accepted: 18 November 2014; published online: 15 December 2014.

Citation: Srinivasa N and Cho Y (2014) Unsupervised discrimination of patterns in spiking neural networks with excitatory and inhibitory synaptic plasticity. *Front. Comput. Neurosci.* 8:159. doi: 10.3389/fncom.2014.00159

This article was submitted to the journal *Frontiers in Computational Neuroscience*.

Copyright © 2014 HRL Laboratories LLC. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

APPENDIX

EXAMPLES FOR DI COMPUTATION

Let us assume that the coding level $C = 3$ for both the examples below and further let us assume that there are two readout units in the sink layer of the network. This allows us to visualize the readout codes in 2-D for simplicity. This analysis however readily extends to other coding levels.

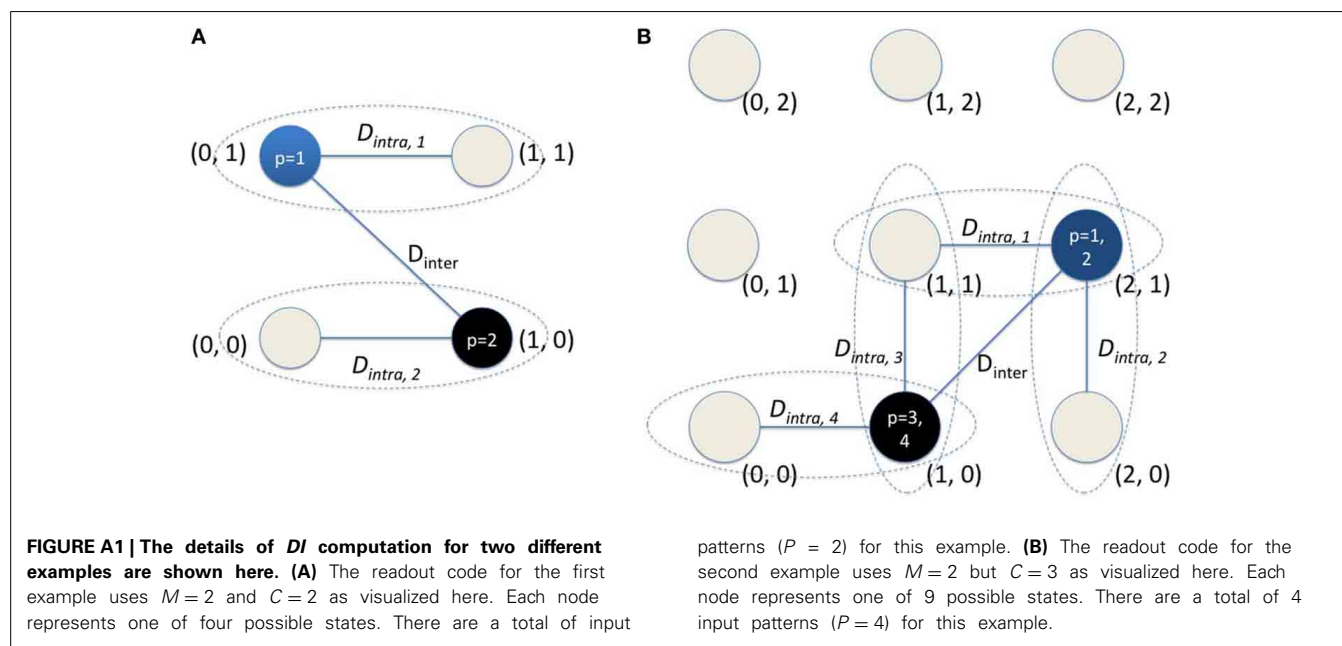
Example 1: Assume that there are two test patterns ($P = 2$) and each of them is presented 10 times resulting in 10 readout codes for each test pattern as follows: $S^1 = \{(0, 1); (0, 1); (0, 1); (0, 1); (1, 0); (1, 1); (1, 1); (0, 1); (0, 1); (0, 1); (0, 1)\}$ and $S^2 = \{(1, 0); (1, 0); (0, 0); (0, 0); (1, 0); (1, 0); (0, 0); (1, 0); (0, 0); (1, 0)\}$. Since there are two patterns and $C = 2$, there are four possible readout codes in general (Figure A1A). But based on S^1 and S^2 , the two readout codes cluster around $(0, 1)$ for $p = 1$ and $(1, 0)$ for $p = 2$.

The various values for DI can be computed as follows. Using Equation (10) and the readout code S^1 , $D_{intra,1} = 6*4/(10*9/2) = 24M/45$ and $D_{intra,2} = 8*2/(10*9/2) = 16M/45$; thus the average $D_{intra} = 20M/45$. Since the readout codes cluster around $(0, 1)$ and $(1, 0)$ for the two test patterns, $D_{inter} = 2M$. Using these D_{inter} and D_{intra} , separability can be computed as $\varepsilon = 1 - 10/45 = 35/45$. Since there are two unique codes for the two test patterns, uniqueness $\gamma = 1$. Thus, the discriminability index for this example will be $DI = 35/45 = 0.78$.

Example 2: Assume that there are four test patterns ($P = 4$) and each of them is presented 10 times resulting in 10 readout codes for each test pattern as follows: $S^1 = \{(2, 1); (2, 1); (2, 1); (2, 1); (2, 1); (1, 1); (1, 1); (2, 1); (2, 1); (2, 1)\}$, $S^2 = \{(2, 1); (2, 1); (2, 0); (2, 0); (2, 1); (2, 0); (2, 1); (2, 1); (2, 1); (2, 1)\}$, $S^3 = \{(1, 0); (1, 0); (0, 0); (0, 0); (1, 0); (1, 0); (0, 0); (1, 0); (0, 0); (1, 0)\}$, $S^4 = \{(1, 0); (1, 0); (1, 1); (1, 0); (1, 0); (1, 1); (1, 0); (1, 0); (1, 1); (1, 0)\}$. This scenario can be visualized using four possible readout codes in general (Figure A1B). Based on four readout codes, they cluster around $(2, 1)$ for $p = 1, 2$ and $(1, 0)$ for $p = 3, 4$.

The various values for DI can be computed as follows. Using Equation (10) and the readout code S^1 , $D_{intra,1} = 16M/45$, $D_{intra,2} = 21M/45$, $D_{intra,3} = 24M/45$, $D_{intra,4} = 21M/45$. Thus, the average D_{intra} can be computed as $41M/90$. Since the readout codes cluster around $(0, 1)$ and $(1, 0)$ for the four test patterns, the $D_{inter} = 2*2*M/(4*3/2) = 2M/3$. Using these D_{inter} and D_{intra} , separability can be calculated as $\varepsilon = 1 - 41/60 = 19/60$. Since there are only two unique codes for the four test patterns, uniqueness $\gamma = 1/2$. Thus, DI for this example will be $DI = (19/60)*(1/2) = 0.16$. This is lower compared to Example 1 since the four patterns are less separable compared to the two test pattern case and the readout codes are also not unique enough compared to Example 1.

These examples illustrate the basics of how the DI is computed and can be readily extended to deal with networks that have larger M and are coded with different coding levels.





Learning and stabilization of winner-take-all dynamics through interacting excitatory and inhibitory plasticity

Jonathan Binas^{1*}, Ueli Rutishauser^{2,3}, Giacomo Indiveri¹ and Michael Pfeiffer¹

¹ Institute of Neuroinformatics, University of Zurich and ETH Zurich, Zurich, Switzerland

² Department of Neurosurgery and Department of Neurology, Cedars-Sinai Medical Center, Los Angeles, CA, USA

³ Computation and Neural Systems Program, Division of Biology and Biological Engineering, California Institute of Technology, Pasadena, CA, USA

Edited by:

Cristina Savin, IST Austria, Austria

Reviewed by:

Yanqing Chen, The Neurosciences Institute, USA

Cristina Savin, IST Austria, Austria

*Correspondence:

Jonathan Binas, Institute of Neuroinformatics, University of Zurich and ETH Zurich, Winterthurerstrasse 190, Zurich 8057, Switzerland
e-mail: jbinas@ini.ethz.ch

Winner-Take-All (WTA) networks are recurrently connected populations of excitatory and inhibitory neurons that represent promising candidate microcircuits for implementing cortical computation. WTAs can perform powerful computations, ranging from signal-restoration to state-dependent processing. However, such networks require fine-tuned connectivity parameters to keep the network dynamics within stable operating regimes. In this article, we show how such stability can emerge autonomously through an interaction of biologically plausible plasticity mechanisms that operate simultaneously on all excitatory and inhibitory synapses of the network. A weight-dependent plasticity rule is derived from the triplet spike-timing dependent plasticity model, and its stabilization properties in the mean-field case are analyzed using contraction theory. Our main result provides simple constraints on the plasticity rule parameters, rather than on the weights themselves, which guarantee stable WTA behavior. The plastic network we present is able to adapt to changing input conditions, and to dynamically adjust its gain, therefore exhibiting self-stabilization mechanisms that are crucial for maintaining stable operation in large networks of interconnected subunits. We show how distributed neural assemblies can adjust their parameters for stable WTA function autonomously while respecting anatomical constraints on neural wiring.

Keywords: winner-take-all, competition, plasticity, self-organization, contraction theory, canonical microcircuits, inhibitory plasticity

1. INTRODUCTION

Competition through shared inhibition is a powerful model of neural computation (Maass, 2000; Douglas and Martin, 2007). Competitive networks are typically composed of populations of excitatory neurons driving a common set of inhibitory neurons, which in turn provide global negative feedback to the excitatory neurons (Amari and Arbib, 1977; Douglas and Martin, 1991; Hertz et al., 1991; Coultrip et al., 1992; Douglas et al., 1995; Hahnloser et al., 2000; Maass, 2000; Rabinovich et al., 2000; Yuille and Geiger, 2003; Rutishauser et al., 2011). Winner-take-all (WTA) networks are one instance of this circuit motif, which has been studied extensively. Neurophysiological and anatomical studies have shown that WTA circuits model essential features of cortical networks (Douglas et al., 1989; Mountcastle, 1997; Binzegger et al., 2004; Douglas and Martin, 2004; Carandini and Heeger, 2012). An individual WTA circuit can implement a variety of non-linear operations such as signal restoration, amplification, filtering, or max-like winner selection, e.g., for decision making (Hahnloser et al., 1999; Maass, 2000; Yuille and Geiger, 2003; Douglas and Martin, 2007). The circuit plays an essential role in both early and recent models of unsupervised learning, such as receptive field development (von der Malsburg, 1973; Fukushima, 1980; Ben-Yishai et al., 1995), or map formation (Willshaw and Von Der Malsburg, 1976; Amari, 1980; Kohonen, 1982; Song and Abbott, 2001). Multiple WTA instances can be

combined to implement more powerful computations that cannot be achieved with a single instance, such as state dependent processing (Rutishauser and Douglas, 2009; Neftci et al., 2013). This modularity has given rise to the idea of WTA circuits representing *canonical microcircuits*, which are repeated many times throughout cortex and are modified slightly and combined in different ways to implement different functions (Douglas and Martin, 1991, 2004; Rutishauser et al., 2011).

In most models of WTA circuits the network connectivity is given a priori. In turn, little is known about whether and how such connectivity could emerge without precise pre-specification. In this article we derive analytical constraints under which local synaptic plasticity on all connections of the network tunes the weights for WTA-type behavior. This is challenging as high-gain WTA operation on the one hand, and stable network dynamics on the other hand, impose diverging constraints on the connection strengths (Rutishauser et al., 2011), which should not be violated by the plasticity mechanism. Previous models like Jug et al. (2012) or Bauer (2013) have shown empirically that functional WTA-like behavior can arise from an interplay of plasticity on excitatory synapses and homeostatic mechanisms. Here, we provide a mathematical explanation for this phenomenon, using a mean-field based analysis, and derive conditions under which biologically plausible plasticity rules applied to all connections of a network of randomly connected inhibitory and excitatory units produce

a functional WTA network with structured connectivity. Due to plastic inhibitory synapses, convergence of the model does not rely on constant, pre-defined inhibitory weights or other common assumptions for WTA models. We prove that the resulting WTA circuits obey stability conditions imposed by contraction analysis (Lohmiller and Slotine, 1998; Rutishauser et al., 2011). This has important implications for the stability of larger networks composed of multiple interconnected WTA circuits, and thus sheds light onto the mechanisms responsible for the emergence of both local functional cortical microcircuits and larger distributed coupled WTA networks.

This article is structured as follows: We first define the network and plasticity models in sections 2.1 to 2.3. Our main analytical results are given in sections 2.4 and 2.5, and illustrated with simulation results in section 2.6. The results are discussed in section 3, and detailed derivations of the analytical results can be found in section 4.

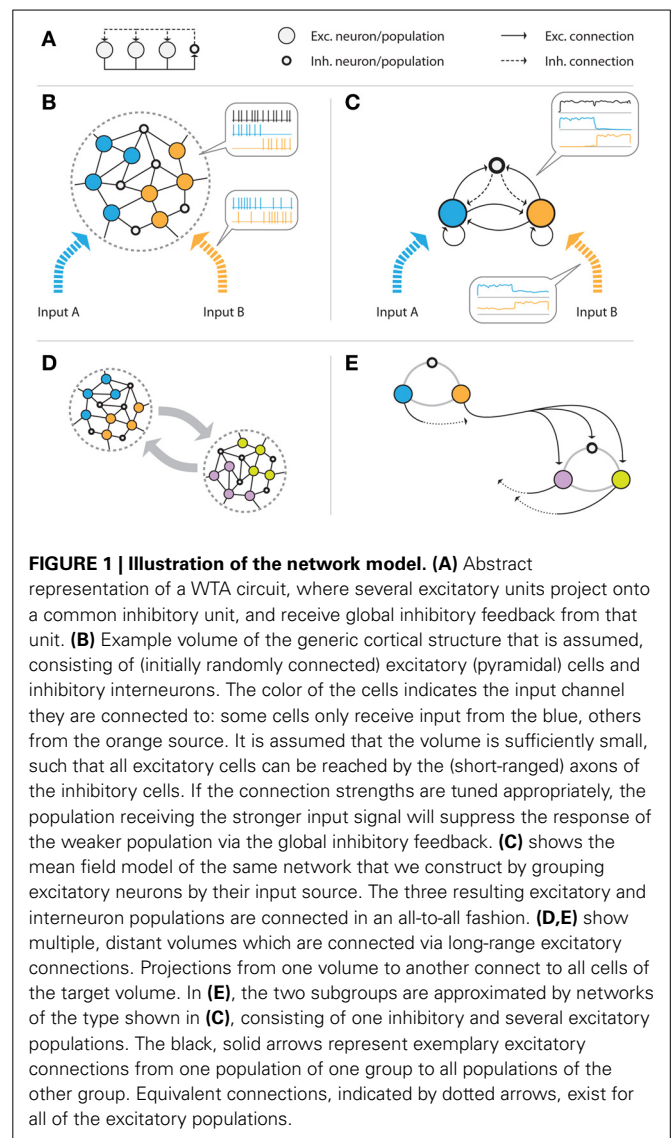
2. RESULTS

2.1. NETWORK TOPOLOGY

In its simplest abstract form, a WTA circuit (**Figure 1A**) consists of a number of excitatory units that project onto a common inhibitory unit. This unit, in turn, provides recurrent inhibitory feedback to all excitatory units. Given appropriate connection strengths, such inhibition makes the excitatory units compete for activation in the sense that the unit receiving the strongest input signal will suppress the activation of all other units through the inhibitory feedback loop, and “win” the competition.

We design a biologically plausible network by taking into account that inhibitory feedback is local, i.e., it only affects cells within a cortical volume that is small enough such that the relatively short inhibitory axonal arbors can reach their targets. We assume excitatory and inhibitory neurons in this volume to be connected randomly (see **Figure 1B**). Furthermore, we assume that there are a finite number of different input signals, each activating a subset of the excitatory cells in the volume. We construct a mean-field model by grouping the excitatory neurons for each driving input stimulus, summarizing the activity of each group of cells by their average firing rate. This results in a simplified population model of the network which—in the case of two different input signals—consists of two excitatory populations (one for each input), and one inhibitory population (see **Figure 1C**). We assume full recurrent connectivity between all populations. This scheme can easily be extended toward more input groups. In particular, if an excitatory group receives multiple inputs, it can be modeled as a new class.

Since inhibitory axons are (typically) short-range, distant populations can communicate only via excitatory projections. We combine multiple local circuits of the form shown in **Figures 1B,C** by introducing excitatory long-range connections between them, as illustrated in **Figure 1D**. Specifically, we add projections from the excitatory populations of one local group to all excitatory and inhibitory populations of the other group. A similar connectivity scheme for implementing distributed WTA networks has been proposed by Rutishauser et al. (2012). Unlike their model, our network does not require specific wiring, but rather targets any potential cell in the other volume. We will



show in section 2.4.4 that this is sufficient to achieve competition between units of spatially distributed WTA circuits.

2.2. NETWORK DYNAMICS

The activation of a neural populations x_i , which can be excitatory or inhibitory, is described by

$$\tau_i \dot{x}_i(t) = -x_i(t) + \left[\sum_j w_{ij} x_j(t) + I_{\text{ext},i}(t) - T_i \right]_+, \quad (1)$$

where τ_i is the time constant of the population, w_{ij} is the weight of the incoming connection from the j th population, $I_{\text{ext},i}(t)$ is an external input given to the population, and T_i is the activation threshold. Furthermore, $[v]_+ := \max(0, v)$ is a half-wave rectification function, preventing the firing rates from taking negative values. Assuming identical time constants for all populations, i.e., $\tau_i = \tau$ for all i , the dynamics of the full system can be written as

$$\tau \dot{\mathbf{x}}(t) = -\mathbf{x}(t) + [\mathbf{W}\mathbf{x}(t) + \mathbf{I}_{\text{ext}}(t) - \mathbf{T}]_+, \quad (2)$$

where $\mathbf{x} = (x_1, \dots, x_N)$ are the firing rates of the respective populations (excitatory and inhibitory), \mathbf{W} is the connectivity matrix (describing local excitatory, local inhibitory, and long-range excitatory connections), $\mathbf{I}_{\text{ext}}(t)$ is a vector of external inputs, and $\mathbf{T} = (T_1, \dots, T_N)$ are the activation thresholds of the populations. For the single local microcircuit shown in **Figure 1C**, for example, \mathbf{W} would be a 3-by-3 matrix with all entries w_{ij} non-zero except for the inhibitory to inhibitory coupling. For two coupled microcircuits as in **Figure 1E**, the connectivity matrix consists of 4 blocks, with the diagonal blocks describing local connectivity, and the off-diagonal blocks describing long-range projections from excitatory units to the other circuit.

2.3. PLASTICITY MECHANISMS AND WEIGHT DYNAMICS

In our model, we assume that all connections w_{ij} in Equation (2) are plastic, and are subject to the following weight update rule:

$$\dot{w} = \tau_s^2 x_{\text{pre}} x_{\text{post}} (x_{\text{post}}(w_{\text{max}} - w) - (\Theta_w + A_w x_{\text{pre}})w). \quad (3)$$

Here, x_{pre} and x_{post} are the pre- and postsynaptic firing rates, respectively, w_{max} is the maximum possible weight value, and Θ_w , A_w , and τ_s are positive constants, which we set to values that are compatible with experimental findings (see **Table 1**). The learning rate is determined by τ_s , and Θ_w and A_w determine the point at which the rule switches between depression (LTD) and potentiation (LTP). We will show that in a plastic network, global stability and circuit function are determined exclusively by those plasticity parameters. The plasticity rule is derived from the mean-field approximation of the triplet STDP rule by Pfister and Gerstner (2006), which we augment with a weight-dependent term, effectively limiting the weight values to the interval $[0, w_{\text{max}}]$. A more detailed derivation of the learning rule can be found in the Methods (section 4.1). The parameters Θ_w and A_w are set differently for excitatory and inhibitory connections, leading to two types of simultaneously active plasticity mechanisms and weight dynamics, even though the same learning equation is used. We set $\Theta_w = \Theta_{\text{exc}}$ and $A_w = A_{\text{exc}}$ for all excitatory connections, and $\Theta_w = \Theta_{\text{inh}}$ and $A_w = A_{\text{inh}}$ for all inhibitory connections. In particular, we assume A_{inh} to take very low values and set $A_{\text{inh}} = 0$ in our analysis, effectively eliminating any dependence of the fixed point of inhibitory weights on the presynaptic rate. According to fits of the parameters to experimental data (see **Table 1**), this is a plausible assumption. For the sake of simplicity, we also assume

the maximum possible weight value w_{max} to be the same for all excitatory and inhibitory connections. **Figure 2** illustrates the weight change as a function of the pre- and postsynaptic activity.

2.4. STABILITY ANALYSIS

The WTA circuit is assumed to function correctly if it converges to a stable state that represents the outcome of the computation it is supposed to perform. Conditions under which these networks converge to their (single) attractor state exponentially fast were previously derived by Rutishauser et al. (2011). Here, we extend those results to plastic networks and express stability criteria in terms of global learning rule parameters, rather than individual weight values. We first describe criteria for the stabilization of the network and learning rule dynamics, then derive from them conditions on the learning rule parameters. Our analysis leads to very simple sufficient conditions that ensure the desired stable WTA behavior.

The dynamics of the network activation and the weights are given by Equations (2) and (3), respectively. In the following, we will denote them by \mathbf{f} and \mathbf{g} , so the full dynamics can be written as a coupled dynamical system

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{w}), \quad (4)$$

$$\dot{\mathbf{w}} = \mathbf{g}(\mathbf{x}, \mathbf{w}), \quad (5)$$

where \mathbf{f} corresponds to the right hand side of Equation (2), and \mathbf{g} combines the update rules for all weights (with different sets of parameters for excitatory and inhibitory connections) in one vector-valued function. We first restrict our analysis to the simplest case of a single winning excitatory population and derive conditions under which the plastic network converges to its fixed point. Later, we extend our analysis to larger systems of multiple coupled excitatory populations.

2.4.1. Analysis of single-node system

Let us first consider a simplified system, in which only one excitatory population is active, e.g., because one population receives much more external input than all others, and the inhibitory feedback suppresses the other populations. As silent populations neither contribute to the network dynamics nor to the weight dynamics, they can be excluded from the analysis. We can therefore reduce the description of the system to a single excitatory population x_E , and an inhibitory population x_I , together with the connections $w_{E \rightarrow E}$, $w_{E \rightarrow I}$, and $w_{I \rightarrow E}$ between them.

Table 1 | Learning rule parameters A_2^\pm , and A_3^\pm from Pfister and Gerstner (2006).

Model	A_2^+	A_3^+	A_2^-	A_3^-	Θ	A	τ_s^2
All-to-all	5×10^{-10}	6.2×10^{-3}	7×10^{-3}	2.3×10^{-4}	18.19	0.06	1.3×10^{-5}
Nearest spike	8.8×10^{-11}	5.3×10^{-2}	6.6×10^{-3}	3.1×10^{-3}	6.24	2.09	3.6×10^{-5}

The values of Θ , A , and τ_s for the plasticity rule Equation (3) have been computed using Equations (14) to (16). The data corresponds to fits of the triplet Spike-Timing Dependent Plasticity (STDP) model with all-to-all spike interactions (first row) and with nearest spike interactions (second row) to recordings from plasticity experiments in rat visual cortex. Note that the time constants $\tau_{x,y}$ and τ_\pm are not reproduced here. However, they all are of the order of hundreds of milliseconds and can be found in Pfister and Gerstner (2006). In our simulations, we use parameters very similar to the “all-to-all” parameters for inhibitory connections, while for excitatory connections we use ones that are close to the “nearest spike” parameters.

For a given set of (fixed) weights \mathbf{w}_c , Rutishauser et al. (2011) have shown by means of contraction theory (Lohmiller and Slotine, 1998) that the system of network activations $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{w}_c)$ converges to its fixed point \mathbf{x}^* exponentially fast if its generalized Jacobian is negative definite. In our case, this condition reduces to

$$\text{Re}\left(w_{E \rightarrow E} - 2 + (w_{E \rightarrow E}^2 - 4 w_{I \rightarrow E} w_{E \rightarrow I})^{1/2}\right) < 0. \quad (6)$$

If condition (6) is met, the system is called contracting and is guaranteed to converge to its attractor state

$$x_E^* = \Lambda I_{\text{ext}}, \quad (7)$$

$$x_I^* = \Lambda w_{E \rightarrow I} I_{\text{ext}}, \quad (8)$$

exponentially fast for any constant input I_{ext} , where the contraction rate is given by the left hand side of (6), divided by 2τ . Here, $\Lambda = (1 - w_{E \rightarrow E} + w_{E \rightarrow I} w_{I \rightarrow E})^{-1}$ corresponds to the network gain. A more detailed derivation of the fixed point can be found in section 4.2. Note that we have set the activation threshold T equal to zero and provide external input I_{ext} to the excitatory population only. This simplifies the analysis but does not affect our results qualitatively.

2.4.2. Decoupling of network and weight dynamics

In the following, we assume that the population dynamics is contracting, i.e., that condition (6) is met, to show that the plasticity dynamics Equation (5) drives the weights \mathbf{w} to a state that is consistent with this condition. Essentially, our analysis has to be self-consistent with respect to the contraction of the activation dynamics. If we assume \mathbf{f} and \mathbf{g} to operate on very different timescales, we can decouple the two systems given by Equations (4) and (5). This is a valid assumption since neural (population) dynamics vary on timescales of tens or hundreds of milliseconds

(see Figure 5 for typical timescales of our system), while synaptic plasticity typically acts on timescales of seconds or minutes. This means that from the point of view of the weight dynamics \mathbf{g} the population activation is at its fixed point \mathbf{x}^* almost all the time, because it converges to that point exponentially fast. We can thus model the activation dynamics as a quasi-static system, and approximate the learning dynamics as a function of the fixed point of the activation instead of the instantaneous activation.

$$\mathbf{g}(\mathbf{x}, \mathbf{w}) \approx \mathbf{g}(\mathbf{x}^*, \mathbf{w}), \quad (9)$$

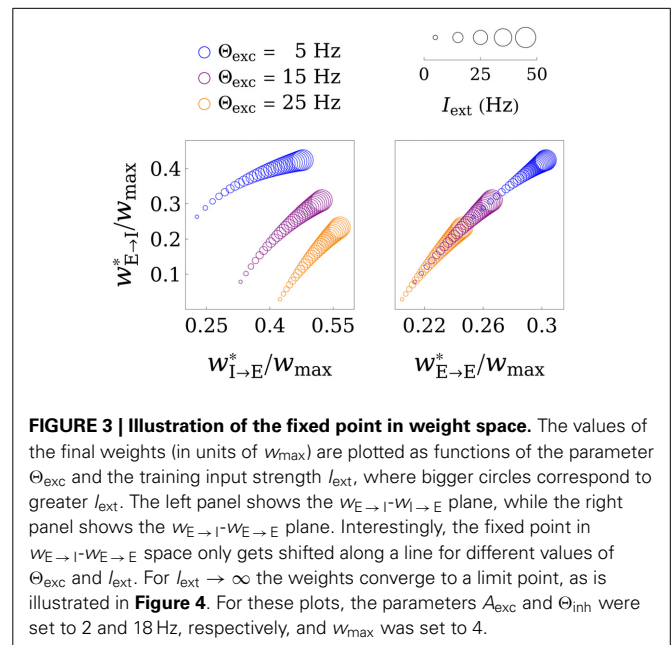

The fixed point of this simplified system is found by setting $\mathbf{g}(\mathbf{x}^*, \mathbf{w}) = \mathbf{0}$, and according to Equation (3) is given by

$$w^* = \frac{w_{\text{max}} x_{\text{post}}^*}{\Theta_w + A_w x_{\text{pre}}^* + x_{\text{post}}^*}. \quad (10)$$

Combining this expression with Equations (7) and (8) leads to a system of non-linear equations that can be solved for the fixed point weights $w_{E \rightarrow E}^*$, $w_{E \rightarrow I}^*$, $w_{I \rightarrow E}^*$, and activations x_E^* , and x_I^* . These values solely depend on the learning rule parameters Θ_w , A_w , w_{max} , and the external (training) input I_{ext} .

Figure 3 shows the fixed points of the weight dynamics as a function of Θ_{exc} , and the input strength I_{ext} . Notably, $w_{E \rightarrow E}^*$ and $w_{E \rightarrow I}^*$ lie on a fixed line in the $w_{E \rightarrow E}$ - $w_{E \rightarrow I}$ plane for all parameters Θ_w and A_w . As the weight values are bounded by 0 and w_{max} , the weights converge to a finite value for $I_{\text{ext}} \rightarrow \infty$. This is also illustrated in Figure 4, which shows the final weight values as a function of w_{max} , both for a finite training input and in the limit $I_{\text{ext}} \rightarrow \infty$.

Importantly, the function of a WTA circuit critically depends on the strength of the recurrent connection $w_{E \rightarrow E}$ (Rutishauser et al., 2011). If $w_{E \rightarrow E} > 1$, the network operates in “hard” mode, where only one unit can win at a time and the activation of all



other units is zero. On the other hand, if $w_{E \rightarrow E}$ is smaller than 1, the network implements “soft” competition, which means that multiple units can be active at the same time. From Equation (27) (Methods) it follows that $w_{E \rightarrow E} > 1$ is possible only if $w_{\max} > A + 1$. As we will show in the following section, this condition is necessarily satisfied by learning rules that lead to stable WTA circuits.

2.4.3. Parameter regimes for stable network function

We can now use the fixed points found in the previous section to express the condition for contraction given by condition (6) in terms of the learning rule parameters. In general, this new condition does not assume an analytically simple form. However, we can find simple sufficient conditions which still provide a good approximation to the actual value (see Methods section 4.2 for details). Specifically, as a key result of our analysis we derive the following sufficient condition: Convergence to a point in weight space that produces stable network dynamics is guaranteed if

$$A_{\text{exc}} + b < w_{\max} < 2(1 + A_{\text{exc}}), \quad (11)$$

where b is a parameter of the order 1, which is related to the minimum activation x_E (or the minimum non-zero input I_{ext}) during training for which this condition should hold. If the minimum input I_{\min} that the network will be trained on is known, then b can be computed from the fixed point $x_{E,\min}^* = x_E^*(I_{\text{ext}} = I_{\min})$, and set to $b = \Theta_{\text{exc}}/x_{E,\min}^*$. This will guarantee contracting dynamics for the full range of training inputs $I_{\text{ext}} \in [I_{\min}, \infty)$. In typical scenarios, b can be set to a number of the order 1. This is due to the fact that the network activation is roughly of the same order as the input strength. Setting Θ_{exc} to a value of similar order leads to $b = \Theta_{\text{exc}}/x_{E,\min}^* \approx 1$.




FIGURE 4 | Limit behavior of the fixed point of the weights for weak and strong inputs. The final weight values (in units of w_{\max}) are plotted as a function of w_{\max} , both for $I_{\text{ext}} = 15$ Hz (solid lines) and in the limit of very large inputs $I_{\text{ext}} \rightarrow \infty$ (dashed lines). In the limit case, $w_{E \rightarrow E}$ and $w_{I \rightarrow E}$ converge to expressions that are linear in w_{\max} , while $w_{E \rightarrow I}$ increases superlinearly. The learning rule parameters were set to $\Theta_{\text{exc}} = 6$ Hz, $\Theta_{\text{inh}} = 18$ Hz, and $A_{\text{exc}} = 2$.

Note that condition (11) is independent of Θ_{exc} and Θ_{inh} . This is due to a simplification that is based on the assumption $A_{\text{exc}} + b \gg 1$, which can be made without loss of generality. If b and A_{exc} are set to very low values, the full expressions given by 38 and (39) (see Methods section 4.2) apply instead. **Figure 5** shows the the region defined by (11) for different b together with the exact

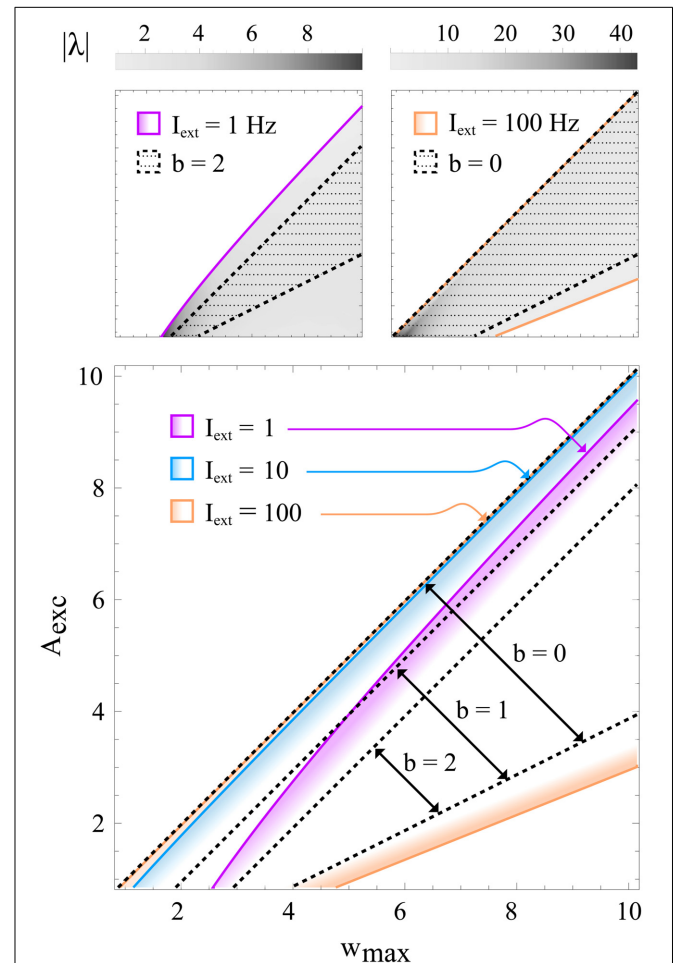


FIGURE 5 | Regions in learning rule parameter space that lead to a stable, contracting network. All panels show the regions of stability in w_{\max} - A_{exc} space for different training input strengths. Colored lines correspond to exact solutions, while black, dotted lines correspond to the sufficient condition (11) for different values of b . The top panels illustrate that relatively small values of b (e.g., 2) roughly approximate the exact solution even for very small inputs (e.g., $I_{\text{ext}} = 1$ Hz; left), whereas b can be set to lower values (e.g., $b = 0$; right) if the input is larger. The gray-scale value represents the convergence rate $|\lambda|$ (in units of s^{-1}) of the activation dynamics for $\tau = 10$ ms. The bottom panel shows in color the exact regions of contraction for inputs $I_{\text{ext}} = 1, 10, 100$ Hz and the approximation given by condition (11) for $b = 0, 1, 2$. Some of the colored regions (and dotted lines) correspond to the ones shown in the upper panels. It can be seen that for higher input strengths the upper bound on A_{exc} (or equivalently, the lower bound on w_{\max}) quickly converges to the $b = 0$ diagonal, which represents the asymptotic condition for $I_{\text{exc}} \rightarrow \infty$. For these plots, the learning rule parameters were set to $\Theta_{\text{exc}} = 6$ Hz and $\Theta_{\text{inh}} = 18$ Hz.

condition for contraction, indicating that (11) is indeed sufficient and that b can safely be set to a value around 1 in most cases.

2.4.4. Extension to multiple units

So far, we have only studied a small network that can be regarded as a single subunit of a larger, distributed WTA system. However, our results can be generalized to larger systems without much effort. In our model, as illustrated in **Figures 1D,E**, different localized WTA circuits can be coupled via excitatory projections. These projections include excitatory-to-inhibitory connections, as well as reciprocal connections between distant excitatory units. In order to demonstrate the effects of this coupling, we consider two localized subsystems, $\mathbf{x} = (x_E, x_I)$ and $\mathbf{x}' = (x'_E, x'_I)$, consisting of one excitatory and one inhibitory unit each. Furthermore, we add projections from x_E to x'_E and x'_I , as required by our model. We denote by $w_{E \rightarrow E'}$ the strength of the long-range excitatory-to-excitatory connection, while we refer to the long-range excitatory-to-inhibitory connection as $w_{E \rightarrow I'}$. Note that for the sake of clarity we only consider the unidirectional case $\mathbf{x} \rightarrow \mathbf{x}'$ here, while the symmetric case $\mathbf{x} \leftrightarrow \mathbf{x}'$ can be dealt with analogously.

We first look at the excitatory-to-inhibitory connections. If only x_E is active and x'_E is silent, then x_I and x'_I are driven by the same presynaptic population (x_E), and $w_{E \rightarrow I'}$ converges to the same value as $w_{E \rightarrow I}$. Thus, after convergence, both inhibitory units are perfectly synchronized in their activation when x_E is active, and an equal amount of inhibition can be provided to x_E and x'_E .

Besides synchronization of inhibition, proper WTA functionality also requires the recurrent excitation $w_{E \rightarrow E'}$ (between the excitatory populations of the different subunits) to converge to sufficiently low values, such that different units compete via the synchronized inhibition rather than exciting each other through the excitatory links. As pointed out by Rutishauser et al. (2012), the network is stable and functions correctly if the recurrent excitation between populations is lower than the recurrent self-excitation, i.e., $w_{E \rightarrow E'} < w_{E \rightarrow E}$.

We now consider the case where x_E and x'_E receive an external input I_{ext} . Whenever x'_E alone receives the input, there is no interaction between the two subunits, and the recurrent self-connection $w_{E' \rightarrow E'}$ converges to the value that was found for the simplified case of a single subunit (section 2.4.2). The same is true for the connection $w_{E \rightarrow E}$ if x_E alone receives the input. However, in this case x_E and x'_E might also interact via the connection $w_{E \rightarrow E'}$, which would then be subject to plasticity. As \mathbf{x} projects to \mathbf{x}' , but not vice versa, we require $x_E > x'_E$ if both x_E and x'_E receive the same input I_{ext} , because x_E should suppress x'_E via the long-range competition mechanism. In terms of connection strengths, this means that $w_{E \rightarrow I'}^* > w_{E' \rightarrow E'}^*$, i.e., the inhibitory input to x'_E that is due to x_E must be greater than the excitatory input x'_E receives from x_E . In the Methods (section 4.3), we show that a sufficient condition for this to be the case is

$$w_{\text{max}} > A + b + 1, \quad (12)$$

which alters our results from section 2.4.3 only slightly, effectively shifting the lower bound on w_{max} by an offset of 1, as can

be seen by comparing conditions (11) and (12). On the other hand, making use of the fact that $x'_E < x_E$, it can be shown that $w_{E \rightarrow E'}$ converges to a value smaller than $w_{E' \rightarrow E'}$ (see Methods section 4.3), as required by the stability condition mentioned above.

2.5. GAIN CONTROL AND NORMALIZATION

In the previous section, we showed how synaptic plasticity can be used to drive the connection strengths toward regimes which guarantee stable network dynamics. Since the actual fixed point values of the weights change with the training input, this mechanism can as well be used to tune certain functional properties of the network. Here we focus on controlling the gain of the network, i.e., the relationship between the strength of the strongest input and the activation of the winning excitatory units within the recurrent circuit, as a function of the training input.

In the case of a single active population, the gain is given by $\Lambda = x_E / I_{\text{ext}} = (1 - w_{E \rightarrow E}^* + w_{E \rightarrow I}^* w_{I \rightarrow E}^*)^{-1}$, as can be inferred from Equation (7). Depending on the gain, the network can either amplify ($\Lambda > 1$) or weaken ($\Lambda < 1$) the input signal.

Figure 6 shows how the gain varies as a function of the learning rule parameters and the training input strength I_{ext} . Low average input strengths cause the weights to converge to values that lead to an increased gain, while higher training inputs lower the gain. This can be regarded as a homeostatic mechanism, acting to keep the network output within a preferred range. This provides a mechanism for the network to adapt to a wide range of input strengths, while still allowing stable WTA competition.

2.6. SIMULATION RESULTS

As a final step, we verify the analytical results in software simulations of a distributed, plastic WTA network, as illustrated in **Figures 1D,E**. Note that here we consider the case where two subgroups are coupled bidirectionally via excitatory long-range projections, while in section 2.4.4, for the sake of clarity, we focus on the unidirectional case. The desired functionality of




FIGURE 6 | Control of the network gain. The network gain is plotted for different learning rule parameters Θ_{exc} , A_{exc} as a function of the training input strength I_{ext} . Θ_{inh} was set to 18 Hz and w_{max} was set to 4.

the resulting network is global competition between the excitatory populations, i.e., the population that receives the strongest input should suppress activation of the other populations, even if the excitatory populations are not directly competing via the same, local inhibitory population. We consider a network with two groups, each consisting of two excitatory populations and one inhibitory population (see **Figure 1E**). While the excitatory populations are connected in an all-to-all manner, inhibitory populations can only target the excitatory populations within their local groups, but do not form long-range projection. Initially, all connection weights (excitatory and inhibitory) are set to random values between 0.3 and 1.8. Note that those values could potentially violate the conditions for contraction defined in (6), but we will show empirically that the plasticity mechanism can still drive the weights toward stable regimes. As training input, we present 1000 constant patterns for 2 s each. In every step, four input values in the ranges 5 ± 2 Hz, 10 ± 2 Hz, 15 ± 2 Hz, and 20 ± 2 Hz are drawn from uniform distributions and applied to the four excitatory units. The different input signals are randomly assigned to the populations in every step, such that a randomly chosen population receives the strongest input. Thereby, each population only receives one of the four inputs.

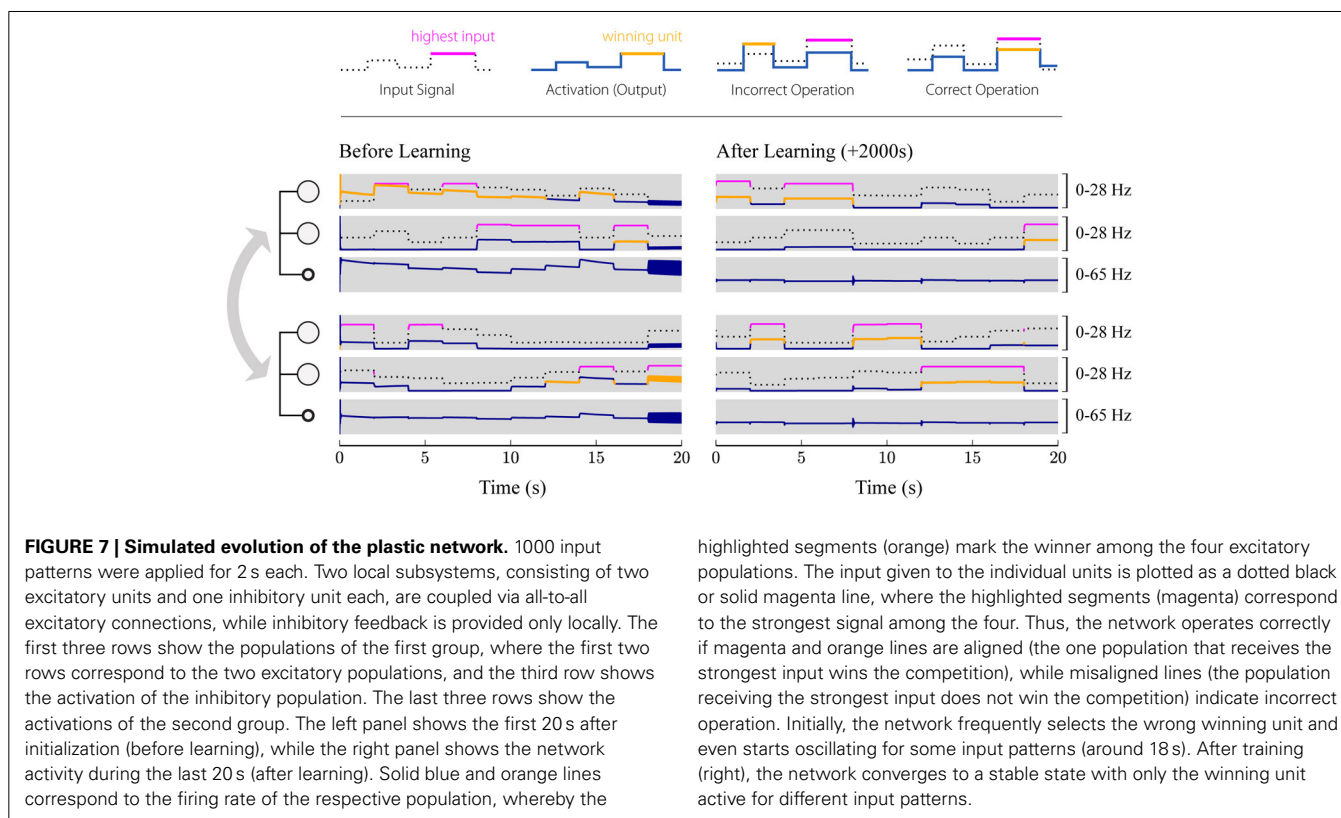
Figure 7 shows the activation of the different populations before and after learning. Before learning (left), the network does not necessarily implement stable competition between the different excitatory populations. Instead, it may end up in an oscillating state or amplify the wrong winning unit. However, after training (**Figure 7**, right), the network always converges to a stable state representing the winner of the competition. Furthermore, it can

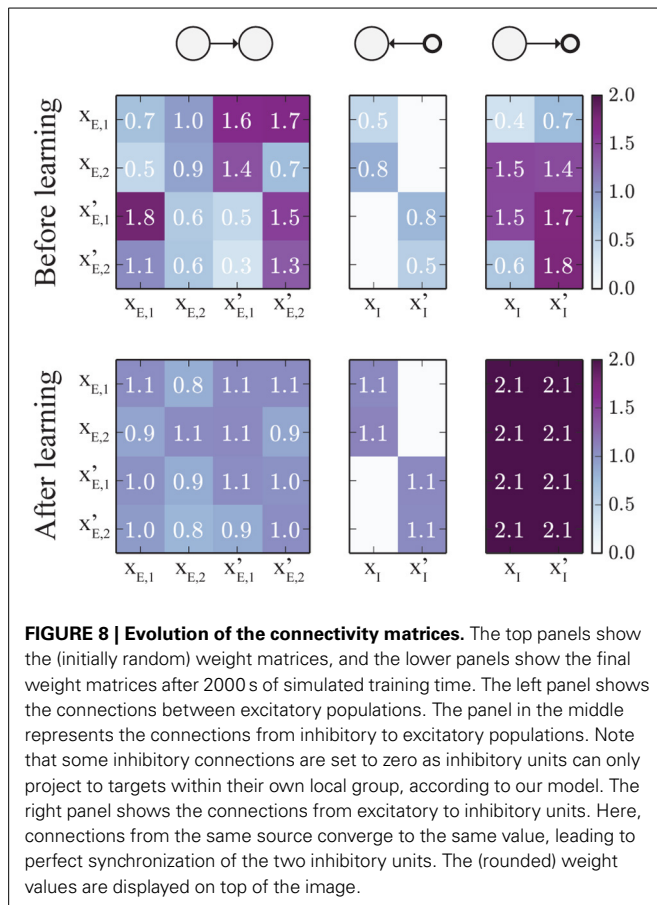
be seen that the inhibitory populations perfectly synchronize, as described in section 2.4.4.

The change of weights is illustrated in **Figure 8**: Initially (top), all weights were set to random values in the range $[0.3, 1.8]$. Since all populations receive the same average input, the weight matrices should converge to symmetric states. For the specific set of learning rule parameters we chose in this example, and the specific input rates described above, $w_{E \rightarrow E}$ converges to a value around 1, which means that the network is at the edge of the transition between hard and soft WTA behavior. The weights $w_{E \rightarrow I}$, connecting excitatory to inhibitory units, converge to values around 2. Furthermore, the weights $w_{I \rightarrow E}$, which connect inhibitory to excitatory units all converge to very similar values (around 1.1), such that inhibition is synchronized across the whole network. Note that not all connections between excitatory populations have converged to the same value. This is because as soon as the network is close to the hard WTA regime, some connections cannot change anymore as only one excitatory unit is active at a time, and the weight change is zero if either the pre- or the post-synaptic unit is inactive.

3. DISCUSSION

We have shown how neural circuits of excitatory and inhibitory neurons can self-organize to implement stable WTA competition. This is achieved through an interplay of excitatory and inhibitory plasticity mechanisms operating on all synapses of the network. As a key result, we provide analytical constraints on the learning rule parameters, which guarantee emergence of the desired network function.





Although constraints on the weights for stable competition in recurrent excitatory-inhibitory circuits have been derived before (Xie et al., 2002; Hahnloser et al., 2003; Rutishauser and Douglas, 2009; Rutishauser et al., 2012), it has remained unclear how a network can self-tune its synaptic weights to comply with these conditions. The presented model achieves this and provides important insights regarding the mechanisms responsible for this self-tuning. Our results predict a relationship between the maximum synaptic weight w_{\max} in a circuit and the learning rule parameter A_{exc} , which controls the contribution of the presynaptic rate to the shifting of the threshold between potentiation and depression. Furthermore, our model predicts a relationship between the network gain and the amount of excitatory input into the circuit during development or training (see Figure 6), indicating that high gain (amplification) should be expected for weak inputs, and low gain for strong inputs, which is in accordance with common assumptions about homeostatic mechanisms (Turrigiano, 2011).

From a developmental perspective, the self-configuration of functional WTA circuits through plasticity has the advantage of requiring a smaller number of parameters to be encoded genetically to obtain stable and functional network structures (Zubler et al., 2013). With self-tuning mechanisms like the ones suggested here, only the parameters for the two different types of plasticity in excitatory and inhibitory synapses, rather than the

strengths of all synaptic connections, need to be specified, and the network can adapt to the statistics of inputs it receives from its environment and from other brain regions.

Besides guaranteeing stability, it is also desirable to control functional properties of the circuit, such as its gain. Experimental data suggests that cortical recurrent circuits often operate in a high gain regime and with strong (larger than unity) recurrent excitatory feedback (Douglas et al., 1995). The strength of this feedback determines whether the WTA is “soft” (multiple excitatory units can be active at the same time) or “hard” (only one unit can be active at a time, i.e., the network operates in a nonlinear regime) (Rutishauser et al., 2011). Many interesting computations that can be realized with these types of networks rely on the non-linearities introduced by such strong recurrent excitation (e.g., Vapnik, 2000), therefore it is important that similar conditions can be achieved with our model. In addition, various forms of learning rely on balanced WTA competition (Masquelier et al., 2009; Habenschuss et al., 2012; Nessler et al., 2013), which requires an adaptation of the gain as the excitatory connections into the circuit undergo plasticity. In our network, the resulting network gain is a function of both the learning rule parameters and the strength of the training input signals. As a consequence, our system can switch between high and low gain, and hard or soft WTA behavior simply by receiving input stimuli of different (average) strengths. Thus, different parts of the network might develop into different functional modules, depending on the inputs they receive.

Our model does not specifically address the question of how the network structure, which leads to our results (essentially random all-to-all connectivity) might develop in the first place. For instance, if certain long-range connections between multiple subcircuits do not exist initially, they will never be established by our model, and the units of the different subcircuits can never compete. On the one hand, this might be a desired effect, e.g., to construct hierarchies or asymmetric structures for competition, in which some parts of the network are able to suppress other parts, but not vice-versa. On the other hand, structural plasticity could account for the creation of missing synaptic connections, or the removal of ineffective connections if the desired stable function cannot be achieved with the anatomical substrate. There is increasing evidence for activity dependent synapse formation and elimination in both juvenile and adult brains (Butz et al., 2009), in particular a coordinated restructuring of inhibitory and excitatory synapses for functional reorganization (Chen and Nedivi, 2013). Another approach, recently investigated in simulations by Bauer (2013), is to set up the right network topology by developmental self-construction processes in a first step, and then tune the network using synaptic plasticity in a second step.

Our model is based on a weight-dependent variation of the learning rule proposed by Pfister and Gerstner (2006), but this is by no means the only learning rule capable of the self-calibration effect we describe in this article. By changing its parametrization, the rule can subsume a wide variety of commonly used Hebbian, STDP-like, and homeostatic plasticity mechanisms. Indeed, further experiments, which are not presented in this manuscript, indicate that a whole class of learning rules with depression at

low and potentiation at high postsynaptic firing rates would lead to similar results. We chose the triplet rule to demonstrate our findings as its parameters have been mapped to experiments, and also because it can be written in an analytically tractable form. We have assumed here a specific type of inhibitory plasticity, which analytically is of the same form as the simultaneous excitatory plasticity, but uses different parameters. With the parameters we chose for the inhibitory plasticity rule, we obtain a form that is very similar to the one proposed by Vogels et al. (2011). By introducing inhibitory plasticity it is no longer necessary to make common but biologically unrealistic assumptions, like pre-specified constant and uniform inhibitory connection strengths (Oster et al., 2009), or more abstract forms of summing up the excitatory activity in the circuit (Jug et al., 2012; Nessler et al., 2013), because inhibitory weights will automatically converge toward stable regions. Inhibitory plasticity has received more attention recently with the introduction of new measurement techniques, and has revealed a great diversity of plasticity mechanisms, in line with the diversity of inhibitory cell types (Kullmann and Lamsa, 2011; Kullmann et al., 2012). Our model involves only a single inhibitory population per local sub-circuit, which interacts with all local excitatory units. Not only is this a common assumption in most previous models, and greatly simplifies the analysis, but also is in accordance with anatomical and electrophysiological results of relatively unspecific inhibitory activity in sensory cortical areas (Kerlin et al., 2010; Bock et al., 2011). However, recent studies have shown more complex interactions of different inhibitory cell types (Pfeffer et al., 2013), making models based on diverse cell types with different properties an intriguing target for future studies. The assumption of a common inhibitory pool that connects to all excitatory units is justified for local circuits, but violates anatomical constraints on the length of inhibitory axons if interacting populations are far apart (Binzegger et al., 2005). Our results easily generalize to the case of distributed inhibition, by adapting the model of Rutishauser et al. (2012) (see **Figure 1E**). Our contribution is to provide the first learning theory for these types of circuits.

Since our model is purely rate-based, a logical next step is to investigate how it translates into the spiking neural network domain. Establishing similar constraints on spike-based learning rules that enable stable WTA competition remains an open problem for future research, although Chen et al. (2013) have shown empirically that WTA behavior in a circuit with topologically ordered input is possible under certain restrictions on initial synapse strengths, and in the presence of STDP and short-term plasticity. Spiking WTA circuits can potentially utilize the richer temporal dynamics of spike trains in the sense that the order of spikes and spike-spike correlations have an effect on the connectivity.

Potential practical applications of our model, and future spiking extensions, lie in neuromorphic VLSI circuits, which have to deal with the problem of device mismatch (Indiveri et al., 2011), and can thus not be precisely configured a priori. Our model could provide a means for the circuits to self-tune and autonomously adapt to the peculiarities of the hardware.

4. MATERIALS AND METHODS

4.1. DERIVATION OF THE PLASTICITY MECHANISM

The learning rule given by Equation (3) is based on the triplet STDP rule by Pfister and Gerstner (2006). Since we are interested in the rate dynamics, we use the mean-field approximation of this rule, which is provided by the authors and leads to an expected weight change of

$$\dot{w} = x_{\text{pre}}x_{\text{post}}(A_2^+\tau_+ - A_2^-\tau_- + A_3^+\tau_+x_{\text{post}} - A_3^-\tau_-x_{\text{pre}}), \quad (13)$$

where x_{pre} , x_{post} are the pre- and postsynaptic activations and A_2^\pm , A_3^\pm , τ_\pm , $\tau_{x,y}$ are parameters that determine the amplitude of weight changes in the triplet STDP model. All of the parameters are assumed to be positive. Through a substitution of constants given by

$$\tau_s^2 := A_3^+\tau_+, \quad (14)$$

$$\Theta_w := (A_2^-\tau_- - A_2^+\tau_+)/\tau_s^2, \quad (15)$$

$$A_w := A_3^-\tau_-/\tau_s^2, \quad (16)$$

the rule in Equation (13) can be written in the simpler form

$$\dot{w} = \tau_s^2 x_{\text{pre}}x_{\text{post}}(x_{\text{post}} - (\Theta_w + A_w x_{\text{pre}})), \quad (17)$$

where Θ_w is in units of a firing rate and A_w is a unitless constant. The terms in parentheses on the right of Equation (17) can be divided into a positive (LTP) part that depends on x_{post} , and a negative (LTD) part that depends on x_{pre} . In order to constrain the range of weights, we add weight-dependent terms $m_+(w)$ and $m_-(w)$ to the two parts of the rule, which yields

$$\dot{w} = \tau_s^2 x_{\text{pre}}x_{\text{post}}(x_{\text{post}}m_+(w) - (\Theta_w + A_w x_{\text{pre}})m_-(w)). \quad (18)$$

Throughout this manuscript, we use a simple, linear weight dependence $m_+ = w_{\text{max}} - w$ and $m_- = w$, which effectively limits the possible values of weights to the interval $[0, w_{\text{max}}]$. We chose this form, which is described by a single parameter, for reasons of analytical tractability and because it is consistent with experimental findings (Gütig et al., 2003). In Pfister and Gerstner (2006), values for the parameters $\tau_{x,y}$, τ_\pm , and $A_{2,3}^\pm$ of the rule Equation (13) were determined from fits to experimental measurements in pyramidal cells in visual cortex (see **Table 1**) and hippocampal cultures (Bi and Poo, 1998, 2001; Sjöström et al., 2001; Wang et al., 2005). We used these values to calculate plausible values for Θ_w , A_w , and τ_s using Equations (14) to (16). In our simulations, we use parameters very similar to the experimentally derived values in **Table 1**. Specifically, for inhibitory connections we use parameters very similar to the ones found from fits of experimental data to the triplet STDP model with all-to-all spike interactions. On the other hand, we choose parameters for the excitatory plasticity rules which are close to fits of the triplet STDP rule with nearest-neighbor spike interactions. The parameters that were used in software simulations and to obtain most of the numeric results are listed in **Table 2**. Note that for the

Table 2 | Model parameters used in software simulation.

Parameter	Value	Description
Θ_{exc}	6 Hz	Learning rule parameter
Θ_{inh}	18 Hz	Learning rule parameter
A_{exc}	2	Learning rule parameter
w_{max}	4	Maximum weight value
$\tau_{\text{s,exc}}^2$	3.6 ms ²	Exc. connection learning rate parameter
$\tau_{\text{s,inh}}^2$	1.3 ms ²	Inh. connection learning rate parameter
τ_{exc}	5 ms	Exc. population time constant
τ_{inh}	1 ms	Inh. population time constant

weight-dependent rule in Equation (18) we have assumed that the parameter Θ_w influences only the LTD part. According to the definition in Equation (15), this is the case if $A_2^- \gg A_2^+$, or $\Theta_w \approx A_2^- \tau_- / \tau_s$, respectively. Otherwise Θ_w contains both a potentiating (A_2^+) and a depressing (A_2^-) component, and Equation (18) should be replaced with a more complex expression of the form of Equation (13).

4.2. DERIVATION OF THE STABILITY CRITERIA

In section 2.4, we outlined how the fixed points and stability criteria for the WTA system can be found. In this section, we provide the detailed derivations that led to these results.

As described in section 2.4, we first consider a simplified system of one excitatory and one inhibitory population, x_E and x_I , which yield an activation vector $\mathbf{x} = (x_E, x_I)^T$. They are coupled recurrently through a weight matrix $\mathbf{W} = \begin{bmatrix} w_{E \rightarrow E} & w_{I \rightarrow E} \\ w_{E \rightarrow I} & 0 \end{bmatrix}$, receive external inputs $I_{\text{ext}}(t)$ with weights μ_E and μ_I respectively, and have thresholds T_E , T_I . Assuming that both units are active, i.e., their total synaptic input is larger than their thresholds, their dynamics are described by

$$\tau_{\text{exc}} \dot{x}_E = -x_E + w_{E \rightarrow E} x_E - w_{I \rightarrow E} x_I + \mu_E I_{\text{ext}} - T_E, \quad (19)$$

$$\tau_{\text{inh}} \dot{x}_I = -x_I + w_{E \rightarrow I} x_E + \mu_I I_{\text{ext}} - T_I, \quad (20)$$

where τ_{exc} , τ_{inh} are the population time constants. The fixed points of the activations can be found by setting $\dot{x}_E = \dot{x}_I = 0$. If we assume, for simplicity, that $T_E = T_I = 0$ this yields the fixed points

$$x_E^* = \Lambda I_{\text{ext}} (\mu_E - w_{I \rightarrow E} \mu_I), \quad (21)$$

$$x_I^* = \Lambda I_{\text{ext}} (w_{E \rightarrow I} \mu_E - (w_{E \rightarrow E} - 1) \mu_I). \quad (22)$$

where

$$\Lambda = (1 - w_{E \rightarrow E} + w_{E \rightarrow I} w_{I \rightarrow E})^{-1} \quad (23)$$

is the network gain. Furthermore, we can make the assumption that $\mu_I = 0$ and $\mu_E = 1$, effectively disabling the external input to the inhibitory population. This reduces Equations (21) and (22) to

$$x_E^* = \Lambda I_{\text{ext}}, \quad (24)$$

$$x_I^* = \Lambda w_{E \rightarrow I} I_{\text{ext}}. \quad (25)$$

These simplifications do not change the results of our analysis qualitatively and can be made without loss of generality.

Approximating x_{pre} and x_{post} by their fixed point activities (as described in section 2.4), and setting $\dot{w} = 0$ in the learning rule Equation (18), the fixed point of the weight dynamics (with $w > 0$) takes the form

$$w^* = \frac{w_{\text{max}} x_{\text{post}}^*}{\Theta_w + A_w x_{\text{pre}}^* + x_{\text{post}}^*}. \quad (26)$$

Note that this fixed point in weight space always exists for any given x_{pre} and x_{post} , and is stable for the weight dependence $m_+(w) = w_{\text{max}} - w$; $m_-(w) = w$ that we chose in Equation (18). In fact, this is true for all choices of the weight dependence satisfying $\partial m_+ / \partial w < 0$ and $\partial m_- / \partial w > 0$, as can be shown by means of a linear stability analysis.

We now derive the fixed points for the weights $w_{E \rightarrow E}$, $w_{E \rightarrow I}$, and $w_{I \rightarrow E}$ of the simplified system. For $w_{E \rightarrow E}$, Equation (26) can be simplified by noting that $x_{\text{pre}}^* = x_{\text{post}}^* = x_E^*$, leading to an expression that depends on the activation of the excitatory population x_E^* :

$$w_{E \rightarrow E}^* = \frac{w_{\text{max}}}{\Theta_{\text{exc}} / x_E^* + A_{\text{exc}} + 1}. \quad (27)$$

Similarly, we can compute the fixed point of $w_{E \rightarrow I}$ as a function of x_E^* , noting that $x_{\text{post}}^* = x_I^* = w_{E \rightarrow I} x_E^*$ [see Equations (24) and (25)]:

$$w_{E \rightarrow I}^* = w_{\text{max}} - \Theta_{\text{exc}} / x_E^* - A_{\text{exc}}. \quad (28)$$

Finally, using the relationship $x_I^* = w_{E \rightarrow I} x_E^*$ from Equations (24) and (25), and the previously computed value of $w_{E \rightarrow I}$ from Equation (28) with the fixed point equation for $w_{I \rightarrow E}$, we obtain

$$w_{I \rightarrow E}^* = \frac{w_{\text{max}}}{\Theta_{\text{inh}} / x_E^* - A_{\text{inh}} (\Theta_{\text{exc}} / x_E^* + (A_{\text{exc}} - w_{\text{max}})) + 1}. \quad (29)$$

In the following, we set $A_{\text{inh}} = 0$, as described in section 2.3. An exact solution for the activation x_E^* at the fixed point of the system is obtained by inserting $w_{E \rightarrow E}^*$, $w_{E \rightarrow I}^*$, and $w_{I \rightarrow E}^*$ into Equation (24), and solving the resulting fixed-point problem $x_E^* = f(x_E^*)$. This corresponds to finding the roots of the third order polynomial

$$P(x) = a_0 + a_1 x + a_2 x^2 + a_3 x^3 = 0 \quad (30)$$

with coefficients

$$a_0 = \Theta_{\text{exc}} \Theta_{\text{inh}} I_{\text{ext}}, \quad (31)$$

$$a_1 = -\Theta_{\text{exc}} \Theta_{\text{inh}} + \Theta_{\text{exc}} I_{\text{ext}} + \Theta_{\text{inh}} I_{\text{ext}} + \Theta_{\text{inh}} A_{\text{exc}} I_{\text{ext}} + \Theta_{\text{exc}}^2 w_{\text{max}}, \quad (32)$$

$$a_2 = -\Theta_{\text{exc}} - \Theta_{\text{inh}} - \Theta_{\text{inh}} A_{\text{exc}} + I_{\text{ext}} + A_{\text{exc}} I_{\text{ext}} + \Theta_{\text{exc}} w_{\text{max}} + \Theta_{\text{inh}} w_{\text{max}} + 2\Theta_{\text{exc}} A_{\text{exc}} w_{\text{max}} - \Theta_{\text{exc}} w_{\text{max}}^2, \quad (33)$$

$$a_3 = -1 - A_{\text{exc}} + w_{\text{max}} + A_{\text{exc}} w_{\text{max}} + A_{\text{exc}}^2 w_{\text{max}} - w_{\text{max}}^2 - A_{\text{exc}} w_{\text{max}}^2. \quad (34)$$

The activation of the excitatory population x_E at the fixed point is then given by the positive, real root of Equation (30).

The fixed point of the activation x_E^* , and thus the fixed points of the weights, are monotonic functions of the training input strength I_{ext} (see **Figure 3**, for example). In the following, we investigate the behavior of the fixed point weight values for very large and very small external inputs during training, respectively. This helps us to find conditions on the learning rule parameters that lead to stable dynamics (of the network activation) for any training input strength. We define a positive constant $b := \Theta_{\text{exc}}/x_E^*$, and plug it into Equations (27)–(29). This yields

$$w_{E \rightarrow E}^* = \frac{w_{\text{max}}}{A_{\text{exc}} + b + 1}, \quad (35)$$

$$w_{E \rightarrow I}^* = w_{\text{max}} - A_{\text{exc}} - b, \quad (36)$$

$$w_{I \rightarrow E}^* = \frac{w_{\text{max}} \Theta_{\text{exc}}}{b \Theta_{\text{inh}} + \Theta_{\text{exc}}}. \quad (37)$$

Inserting Equations (35)–(37) into the condition for contraction of the activation dynamics given by (6), we can describe the condition in terms of the learning rule parameters, and a new constant $\tilde{\Theta} := \Theta_{\text{exc}}/(\Theta_{\text{exc}} + b \Theta_{\text{inh}})$:

$$\frac{1}{(1 + A_{\text{exc}} + b)} < \tilde{\Theta} < 1, \quad (38)$$

$$(A_{\text{exc}} + b) \left(1 + \frac{1}{\tilde{\Theta}(1 + A_{\text{exc}} + b)^2 - 1} \right) < w_{\text{max}} < 2(1 + A_{\text{exc}} + b), \quad (39)$$

Assuming $A_{\text{exc}} + b \gg 1$ (note that we can always set A_{exc} to a sufficiently large value), the conditions reduce to

$$0 < \tilde{\Theta} < 1, \quad (40)$$

$$A_{\text{exc}} + b < w_{\text{max}} < 2(1 + A_{\text{exc}} + b), \quad (41)$$

whereby the first condition can be dropped, since $\tilde{\Theta} \in [0, 1]$ always holds. The second condition still depends on b , and therefore on x_E^* . We will illustrate how to eliminate this dependence under very weak assumptions. First, in the limit of very large inputs x_E^* also takes very large values, leading to $b \rightarrow 0$ for $I_{\text{ext}} \rightarrow \infty$. In that case, condition (41) becomes independent of b and can be written as

$$A_{\text{exc}} < w_{\text{max}} < 2(1 + A_{\text{exc}}). \quad (42)$$

On the other hand, in the case of very small inputs we have to include the effects of b , as b can in principle take very large values. In typical scenarios the output of the network can be assumed to be roughly of the order of its input. If Θ_{exc} is chosen to be of the same order, then $b \approx 1$. For any finite b , we can express the stability condition that is valid for all inputs as the intersection of the conditions for large inputs, condition (42), with the one for arbitrarily small inputs, condition (41), leading to

$$A_{\text{exc}} + b < w_{\text{max}} < 2(1 + A_{\text{exc}}). \quad (43)$$

Note that this condition can be met for any finite b by choosing sufficiently large A_{exc} and w_{max} . However, as discussed above, choices of the parameter b of the order 1 should be sufficient for typical scenarios, whereas higher values would guarantee stable dynamics for very low input strengths (e.g., $I_{\text{ext}} \ll \Theta_{\text{exc}}$). This is illustrated in **Figure 5**, where the exact regions of stability as a function of w_{max} and A_{exc} are shown for different training input strengths, together with the sufficient conditions given by (43). In practice, a good starting point for picking a value b for which the stability conditions should hold is to determine the minimum non-zero input I_{min} encountered during training for which this condition should hold, and setting $b = \Theta_{\text{exc}}/x_{E,\text{min}}^*$, where $x_{E,\text{min}}^*$ is the fixed point activation for $I_{\text{ext}} = I_{\text{min}}$.

4.3. EXTENSION TO MULTIPLE UNITS

In this section, we illustrate how multiple subunits, as analyzed in the previous section, can be combined to larger WTA networks with distributed inhibition. For the sake of simplicity, we only consider the unidirectional case, where a subunit $\mathbf{x} = (x_E, x_I)$ projects onto another subunit $\mathbf{x}' = (x'_E, x'_I)$ via excitatory connections $w_{E \rightarrow E'}$ and $w_{E \rightarrow I'}$. The bidirectional case $\mathbf{x} \leftrightarrow \mathbf{x}'$ can be analyzed analogously. If x_E and x'_E receive the same input, the response of x'_E should be weaker, such that activation of x_E causes suppression of x'_E rather than excitation. This means that

$$w_{E \rightarrow E'}^* < w_{E \rightarrow I'}^* w_{I' \rightarrow E'}^* \quad (44)$$

must hold. We assume that both subsystems have been trained on inputs of the same average strength, such that their local connections have converged to the same weights, i.e., $w_{E' \rightarrow I'}^* = w_{E \rightarrow I}^*$ and $w_{I' \rightarrow E'}^* = w_{I \rightarrow E}^*$. Furthermore, we assume that condition (44) is true initially. This can be guaranteed by setting the initial value of $w_{E \rightarrow E'}$ to a sufficiently small number. Our task then is to show that condition (44) remains true for all time. The values of $w_{E' \rightarrow I'}^*$ and $w_{I' \rightarrow E'}^*$, or $w_{E \rightarrow I}^*$ and $w_{I \rightarrow E}^*$ respectively, are described by Equations (36) and (37). On the other hand, according to Equation (26), the value of $w_{E \rightarrow E'}^*$ is given by

$$w_{E \rightarrow E'}^* = \frac{w_{\text{max}} x_E'^*}{\Theta_{\text{exc}} + A_{\text{exc}} x_E^* + x_E'^*}. \quad (45)$$

Plugging all this into condition (44) and simplifying the expression, leads to the condition

$$w_{\text{max}} > A_{\text{exc}} + b + \frac{x_E'}{A_{\text{exc}} x_E + x_E' + \Theta}, \quad (46)$$

which can be replaced by the sufficient condition

$$w_{\max} > A_{\text{exc}} + b + 1, \quad (47)$$

that guarantees $x_{E'}^* < x_E^*$ if both excitatory populations receive the same input. On the other hand, this result implies $w_{E \rightarrow E'}^* < w_{E' \rightarrow E'}^*$, which is required for stable network dynamics (Rutishauser et al., 2012), and can be verified by comparing the respective fixed point equations

$$w_{E \rightarrow E'}^* = x_{E'}' / (\Theta_{\text{exc}} + A_{\text{exc}} x_E + x_E'), \quad (48)$$

$$w_{E' \rightarrow E'}^* = x_{E'}' / (\Theta_{\text{exc}} + A_{\text{exc}} x_{E'}' + x_E'). \quad (49)$$

4.4. SOFTWARE SIMULATION

Software simulations of our model were implemented using custom Python code based on the “NumPy” and “Dana” packages, and run on a Linux workstation. Numerical integration of the system dynamics was carried out using the forward Euler method with a 1 ms timestep.

AUTHOR CONTRIBUTIONS

Jonathan Binas, Ueli Rutishauser, Giacomo Indiveri, Michael Pfeiffer conceived and designed the experiments. Jonathan Binas performed the experiments and analysis. Jonathan Binas, Ueli Rutishauser, Giacomo Indiveri, Michael Pfeiffer wrote the paper.

FUNDING

The research was supported by the Swiss National Science Foundation Grant 200021_146608, and the European Union ERC Grant “neuroP” (257219).

ACKNOWLEDGMENT

We thank Rodney Douglas, Peter Diehl, Roman Bauer, and our colleagues at the Institute of Neuroinformatics for fruitful discussion.

REFERENCES

- Amari, S., and Arbib, M. (1977). “Competition and cooperation in neural nets,” in *Systems Neuroscience*, ed J. Metzler (San Diego, CA: Academic Press), 119–165.
- Amari, S.-I. (1980). Topographic organization of nerve fields. *Bull. Math. Biol.* 42, 339–364. doi: 10.1007/BF02460791
- Bauer, R. (2013). *Self-Construction and -Configuration of Functional Neuronal Networks*. PhD Thesis, ETH Zürich.
- Ben-Yishai, R., Bar-Or, R. L., and Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* 92, 3844–3848. doi: 10.1073/pnas.92.9.3844
- Bi, G., and Poo, M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Bi, G., and Poo, M. (2001). Synaptic modification by correlated activity: Hebb’s postulate revisited. *Ann. Rev. Neurosci.* 24, 139–166. doi: 10.1146/annurev.neuro.24.1.139
- Binzegger, T., Douglas, R. J., and Martin, K. (2004). A quantitative map of the circuit of cat primary visual cortex. *J. Neurosci.* 24, 8441–8453. doi: 10.1523/JNEUROSCI.1400-04.2004
- Binzegger, T., Douglas, R. J., and Martin, K. A. (2005). Axons in cat visual cortex are topologically self-similar. *Cereb. Cortex* 15, 152–165. doi: 10.1093/cercor/bhh118
- Bock, D. D., Lee, W.-C. A., Kerlin, A. M., Andermann, M. L., Hood, G., Wetzell, A. W., et al. (2011). Network anatomy and *in vivo* physiology of visual cortical neurons. *Nature* 471, 177–182. doi: 10.1038/nature09802
- Butz, M., Wörgötter, F., and van Ooyen, A. (2009). Activity-dependent structural plasticity. *Brain Res. Rev.* 60, 287–305. doi: 10.1016/j.brainresrev.2008.12.023
- Carandini, M., and Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* 13, 51–62. doi: 10.1038/nrn3136
- Chen, J. L., and Nedivi, E. (2013). Highly specific structural plasticity of inhibitory circuits in the adult neocortex. *Neuroscientist* 19, 384–393. doi: 10.1177/1073858413479824
- Chen, Y., McKinstry, J. L., and Edelman, G. M. (2013). Versatile networks of simulated spiking neurons displaying winner-take-all behavior. *Front. Comput. Neurosci.* 7:16. doi: 10.3389/fncom.2013.00016
- Coultrip, R., Granger, R., and Lynch, G. (1992). A cortical model of winner-take-all competition via lateral inhibition. *Neural Netw.* 5, 47–54. doi: 10.1016/S0893-6080(05)80006-1
- Douglas, R., Koch, C., Mahowald, M., Martin, K., and Suarez, H. (1995). Recurrent excitation in neocortical circuits. *Science* 269, 981–985. doi: 10.1126/science.7638624
- Douglas, R. J., and Martin, K. A. (1991). Opening the grey box. *Trends Neurosci.* 14, 286–293. doi: 10.1016/0166-2236(91)90139-L
- Douglas, R. J., and Martin, K. A. (2004). Neuronal circuits of the neocortex. *Ann. Rev. Neurosci.* 27, 419–451. doi: 10.1146/annurev.neuro.27.070203.144152
- Douglas, R. J., and Martin, K. A. (2007). Recurrent neuronal circuits in the neocortex. *Curr. Biol.* 17, R496–R500. doi: 10.1016/j.cub.2007.04.024
- Douglas, R. J., Martin, K. A., and Whitteridge, D. (1989). A canonical microcircuit for neocortex. *Neural Comput.* 1, 480–488. doi: 10.1162/neco.1989.1.4.480
- Fukushima, K. (1980). Neocognitron: a self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol. Cybern.* 36, 193–202. doi: 10.1007/BF00344251
- Gütig, R., Aharonov, R., Rotter, S., and Sompolinsky, H. (2003). Learning input correlations through nonlinear temporally asymmetric Hebbian plasticity. *J. Neurosci.* 23, 3697–3714.
- Habenschuss, S., Bill, J., and Nessler, B. (2012). “Homeostatic plasticity in Bayesian spiking networks as Expectation Maximization with posterior constraints,” in *Proceedings of Neural Information Processing Systems (NIPS)*, 782–790.
- Hahnloser, R., Douglas, R., Mahowald, M., and Hepp, K. (1999). Feedback interactions between neuronal pointers and maps for attentional processing. *Nat. Neurosci.* 2, 746–752. doi: 10.1038/11219
- Hahnloser, R. H., Sarpeshkar, R., Mahowald, M. A., Douglas, R. J., and Seung, H. S. (2000). Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature* 405, 947–951. doi: 10.1038/35016072
- Hahnloser, R. H., Seung, H. S., and Slotine, J.-J. (2003). Permitted and forbidden sets in symmetric threshold-linear networks. *Neural Comput.* 15, 621–638. doi: 10.1162/089976603321192103
- Hertz, J., Krogh, A., and Palmer, R. (1991). *Introduction to the Theory of Neural Computation*. Redwood City, CA: Addison-Wesley.
- Indiveri, G., Linares-Barranco, B., Hamilton, T. J., Van Schaik, A., Etienne-Cummings, R., Delbruck, T., et al. (2011). Neuromorphic silicon neuron circuits. *Front. Neurosci.* 5:73. doi: 10.3389/fnins.2011.00073
- Jug, F., Cook, M., and Steger, A. (2012). “Recurrent competitive networks can learn locally excitatory topologies,” in *International Joint Conference on Neural Networks (IJCNN)*, 1–8.
- Kerlin, A. M., Andermann, M. L., Berezovskii, V. K., and Reid, R. C. (2010). Broadly tuned response properties of diverse inhibitory neuron subtypes in mouse visual cortex. *Neuron* 67, 858–871. doi: 10.1016/j.neuron.2010.08.002
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biol. Cybern.* 43, 59–69. doi: 10.1007/BF00337288
- Kullmann, D. M., and Lamsa, K. P. (2011). LTP and LTD in cortical GABAergic interneurons: emerging rules and roles. *Neuropharmacology* 60, 712–719. doi: 10.1016/j.neuropharm.2010.12.020
- Kullmann, D. M., Moreau, A. W., Bakiri, Y., and Nicholson, E. (2012). Plasticity of inhibition. *Neuron* 75, 951–962. doi: 10.1016/j.neuron.2012.07.030
- Lohmiller, W., and Slotine, J.-J. E. (1998). On contraction analysis for non-linear systems. *Automatica* 34, 683–696. doi: 10.1016/S0005-1098(98)00019-3
- Maass, W. (2000). On the computational power of winner-take-all. *Neural Comput.* 12, 2519–2536. doi: 10.1162/089976600300014827
- Masquelier, T., Guyonneau, R., and Thorpe, S. (2009). Competitive STDP-based spike pattern learning. *Neural Comput.* 21, 1259–1276. doi: 10.1162/neco.2008.06-08-804
- Mountcastle, V. B. (1997). The columnar organization of the neocortex. *Brain* 120, 701–722. doi: 10.1093/brain/120.4.701

- Neftci, E., Binas, J., Rutishauser, U., Chicca, E., Indiveri, G., and Douglas, R. J. (2013). Synthesizing cognition in neuromorphic electronic systems. *Proc. Natl. Acad. Sci. U.S.A.* 110, E3468–E3476. doi: 10.1073/pnas.1212083110
- Nessler, B., Pfeiffer, M., Buesing, L., and Maass, W. (2013). Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity. *PLoS Comput. Biol.* 9:e1003037. doi: 10.1371/journal.pcbi.1003037
- Oster, M., Douglas, R., and Liu, S.-C. (2009). Computation with spikes in a winner-take-all network. *Neural Comput.* 21, 2437–2465. doi: 10.1162/neco.2009.07.08-829
- Pfeffer, C. K., Xue, M., He, M., Huang, Z. J., and Scanziani, M. (2013). Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nat. Neurosci.* 16, 1068–1076. doi: 10.1038/nn.3446
- Pfister, J.-P., and Gerstner, W. (2006). Triplets of spikes in a model of spike timing-dependent plasticity. *J. Neurosci.* 26, 9673–9682. doi: 10.1523/JNEUROSCI.1425-06.2006
- Rabinovich, M. I., Huerta, R., Volkovskii, A., Abarbanel, H. D. I., Stopfer, M., and Laurent, G. (2000). Dynamical coding of sensory information with competitive networks. *J. Physiol. (Paris)* 94, 465–471. doi: 10.1016/S0928-4257(00)01092-5
- Rutishauser, U., and Douglas, R. J. (2009). State-dependent computation using coupled recurrent networks. *Neural Comput.* 21, 478–509. doi: 10.1162/neco.2008.03-08-734
- Rutishauser, U., Douglas, R. J., and Slotine, J.-J. (2011). Collective stability of networks of winner-take-all circuits. *Neural Comput.* 23, 735–773. doi: 10.1162/NECO-a-00091
- Rutishauser, U., Slotine, J.-J., and Douglas, R. J. (2012). Competition through selective inhibitory synchrony. *Neural Comput.* 24, 2033–2052. doi: 10.1162/NECO-a-00304
- Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32, 1149–1164. doi: 10.1016/S0896-6273(01)00542-6
- Song, S., and Abbott, L. F. (2001). Cortical development and remapping through spike timing-dependent plasticity. *Neuron* 32, 339–350. doi: 10.1016/S0896-6273(01)00451-2
- Turrigiano, G. (2011). Too many cooks? Intrinsic and synaptic homeostatic mechanisms in cortical circuit refinement. *Ann. Rev. Neurosci.* 34, 89–103. doi: 10.1146/annurev-neuro-060909-153238
- Vapnik, V. (2000). *The Nature of Statistical Learning Theory*. New York, NY: Information Science and Statistics, Springer.
- Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* 334, 1569–1573. doi: 10.1126/science.1211095
- von der Malsburg, C. (1973). Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik* 14, 85–100. doi: 10.1007/BF00288907
- Wang, H.-X., Gerkin, R. C., Nauen, D. W., and Bi, G.-Q. (2005). Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nat. Neurosci.* 8, 187–193. doi: 10.1038/nn1387
- Willshaw, D. J., and Von Der Malsburg, C. (1976). How patterned neural connections can be set up by self-organization. *Proc. R. Soc. Lond. Ser. B Biol. Sci.* 194, 431–445. doi: 10.1098/rspb.1976.0087
- Xie, X., Hahnloser, R. H., and Seung, H. S. (2002). Selectively grouping neurons in recurrent networks of lateral inhibition. *Neural Comput.* 14, 2627–2646. doi: 10.1162/089976602760408008
- Yuille, A., and Geiger, D. (2003). “Winner-take-all networks,” in *The Handbook of Brain Theory and Neural Networks*, ed M. Arbib (Cambridge, MA: MIT Press), 1228–1231.
- Zubler, F., Hauri, A., Pfister, S., Bauer, R., Anderson, J. C., Whatley, A. M., et al. (2013). Simulating cortical development as a self constructing process: a novel multi-scale approach combining molecular and physical aspects. *PLoS Comput. Biol.* 9:e1003173. doi: 10.1371/journal.pcbi.1003173

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 15 April 2014; accepted: 16 June 2014; published online: 08 July 2014.

Citation: Binas J, Rutishauser U, Indiveri G and Pfeiffer M (2014) Learning and stabilization of winner-take-all dynamics through interacting excitatory and inhibitory plasticity. *Front. Comput. Neurosci.* 8:68. doi: 10.3389/fncom.2014.00068

This article was submitted to the journal *Frontiers in Computational Neuroscience*. Copyright © 2014 Binas, Rutishauser, Indiveri and Pfeiffer. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.



Adaptation of short-term plasticity parameters via error-driven learning may explain the correlation between activity-dependent synaptic properties, connectivity motifs and target specificity

Umberto Esposito¹, Michele Giugliano^{1,2,3} and Eleni Vasilaki^{1,2,4*}

¹ Department Computer Science, University of Sheffield, Sheffield, UK

² Theoretical Neurobiology and Neuroengineering Laboratory, Department Biomedical Sciences, University of Antwerp, Antwerp, Belgium

³ Laboratory of Neural Microcircuitry, Brain Mind Institute, Swiss Federal Institute of Technology of Lausanne, École Polytechnique Fédérale de Lausanne, Switzerland

⁴ INSIGNEO Institute for in Silico Medicine, University of Sheffield, Sheffield, UK

Edited by:

Cristina Savin, Institute of Science and Technology Austria, Austria

Reviewed by:

Rui Ponte Costa, University of Oxford, UK

Alberto Bernacchia, Jacobs University Bremen, Germany

*Correspondence:

Eleni Vasilaki, Department Computer Science, University of Sheffield, Regent Court, 211 Portobello Street, S1 4DP Sheffield, UK
e-mail: e.vasilaki@sheffield.ac.uk

The anatomical connectivity among neurons has been experimentally found to be largely non-random across brain areas. This means that certain connectivity motifs occur at a higher frequency than would be expected by chance. Of particular interest, short-term synaptic plasticity properties were found to colocalize with specific motifs: an over-expression of bidirectional motifs has been found in neuronal pairs where short-term facilitation dominates synaptic transmission among the neurons, whereas an over-expression of unidirectional motifs has been observed in neuronal pairs where short-term depression dominates. In previous work we found that, given a network with fixed short-term properties, the interaction between short- and long-term plasticity of synaptic transmission is sufficient for the emergence of specific motifs. Here, we introduce an error-driven learning mechanism for short-term plasticity that may explain how such observed correspondences develop from randomly initialized dynamic synapses. By allowing synapses to change their properties, neurons are able to adapt their own activity depending on an error signal. This results in more rich dynamics and also, provided that the learning mechanism is target-specific, leads to specialized groups of synapses projecting onto functionally different targets, qualitatively replicating the experimental results of Wang and collaborators.

Keywords: short-term plasticity, long-term plasticity, learning, rate code, motifs, target-specificity

1. INTRODUCTION

It is the current belief that experiences and memories are registered in long-term stable synaptic changes. Long-term plasticity, and in particular Hebbian learning or Spike-Timing-Dependent-Plasticity (STDP), is a form of unsupervised learning that captures correlations in the neuronal input. Hence, their involvement in, for instance, the development of receptive fields (e.g., Song et al., 2000; Clopath et al., 2010) or memory and associations is long-standing knowledge. However, the variety of different long-term plasticity rules (Markram et al., 2011), indicates that the precise synaptic prescriptions of long-term plasticity mechanisms remain unclear.

On the contrary, short-term plasticity (STP) is well-described (Varela et al., 1997; Markram et al., 1998b; Le Be' and Markram, 2006; Rinaldi et al., 2008; Testa-Silva et al., 2012; Costa et al., 2013; Romani et al., 2013) in the context of specific models (Tsodyks and Markram, 1997; Hennig, 2013; Rotman and Klyachko, 2013). Its role in neuronal computation is assumed to be related to temporal processing, see for instance (Natschläger et al., 2001) or the work by Carvalho and Buonomano (2011), where STP is demonstrated to enhance the discrimination ability of a single

neuron (i.e., a tempotron, see Gütiğ and Sompolinsky, 2006), when presented with forward and reverse patterns. Synapses with STP are also optimal estimators of presynaptic membrane potentials (Pfister et al., 2010).

The investigation of the brain wiring diagram known as *connectomics* has recently made spectacular progress and generated excitement for its perspectives (Seung, 2009). Novel discoveries in molecular biology (Wickersham et al., 2007; Zhang et al., 2007; Lichtman et al., 2008), neuroanatomical methods (Denk and Horstmann, 2004; Chklovskii et al., 2010), electrophysiology (Song et al., 2005; Hai et al., 2010; Perin et al., 2011), and imaging (Friston, 2011; Minderer et al., 2012; Wedeen et al., 2012) have pushed forward the technological limits for ultimate access to neuronal connectivity. The comprehension of this level of organization of the brain (Kandell et al., 2008) is pivotal to understanding the richness of its high-level cognitive, computational and adaptive properties, as well as its dysfunctions.

At the microcircuit level (Binzegger and Douglas, 2004; Grillner et al., 2005; Silberberg et al., 2005; Douglas and Martin, 2007a,b), the non-random features of cortical connectivity have recently raised a lot of interest (Song et al., 2005; Perin et al.,

2011). The occurrence of stereotypical connectivity motifs was experimentally demonstrated and, in some cases, accompanied by physiological information on neuronal and synaptic properties (Song et al., 2005; Wang et al., 2006; Silberberg and Markram, 2007; Perin et al., 2011), on activity-dependent short-term and long-term plasticity (Buonomano and Merzenich, 1998) and rewiring (Chklovskii et al., 2004; Le Be' and Markram, 2006). Recent experimental findings obtained in young ferret cortices (Wang et al., 2006) indicate that short-term facilitation and depression correlate to specific connectivity motifs: neurons connected by synapses exhibiting short-term facilitation form predominantly reciprocal (bidirectional) motifs; neurons connected by synapses exhibiting short-term depression form unidirectional motifs. Interestingly, the same overexpression of connectivity motifs has been observed in another brain area, i.e., the excitatory microcircuitry of the olfactory bulb (Pignatelli, 2009).

Earlier work by Vasilaki and Giugliano (2012, 2014) attempted to shed light on this correlation between STP and the observed wiring diagram configuration. They demonstrate that all facilitating or all depressing networks, upon receiving the same wave-like stimulation, give rise to the experimentally observed motifs: bidirectional for facilitating synapses and unidirectional for depressing synapses. This was explained both in the context of mean field analysis and microscopic simulations as a frequency-dependent effect. This is a simple consequence of the type of input (wave like) and the choice of the STDP triplet rule (Pfister and Gerstner, 2006). Differently from the classical pair rule, the triplet rule displays a frequency-dependent behavior, which can explain some experimental results (Sjöström et al., 2001): at low frequencies the rule reveals the classic STDP and, given a wave-like input, it results in unidirectional connectivity (Clopath et al., 2010; Vasilaki and Giugliano, 2014). At high frequencies, however, it reveals “classic Hebb” behavior: neurons that fire together, wire together. Hence, the low firing network develops unidirectional connectivity, while the high firing network develops bidirectional connectivity; for details see (Vasilaki and Giugliano, 2014). However, the observed synaptic development was not associated to any particular type of learning, but was explored as the emerging structure upon receiving a wave like input: what the network learned *per se* in that context was not clear.

With the present work we aim to complement and extend on Vasilaki and Giugliano (2012, 2014). We define a learning model for STP through which a population of neurons can modify its synapses in order to adapt its own activity and then fulfill a given time-varying task. The key idea comes from an optimization perspective: neurons that are able to modify their synapses, for instance making depressing synapses more and more depressing or even turning them into facilitating ones, would allow for much more flexibility and efficacy in signal transmission. A similar argument can be found in Markram et al. (1998a), whereas for earlier but different mechanisms of STP optimization or learning we redirect to Natschläger et al. (2001) and Carvalho and Buonomano (2011).

Then, we construct a typical inverted associative learning problem (Asaad et al., 1998; Fusi et al., 2007; Vasilaki et al., 2009b) where neurons have to learn to respond with high or low frequencies, when presented with the same wave-like input signal. We use

this paradigm to show the potential of our model. In particular, not only do we provide an explanation for the correspondence motifs-synaptic properties within the context of learning both STP and STDP (triplet rule) but we also qualitatively capture, for instance, the heterogeneity in synaptic properties observed by Wang et al. (2006).

Moreover, having defined the learning model as a target-specific mechanism, we are able to obtain variability in the short-term profile of synapses innervating functionally different targets. Finally, we show that the learning model can be reduced to a minimal model where only the time constant of recovery from depression τ_{rec} needs to be learnt in order to obtain neurons firing at high or low frequency. Comparing this finding with the results from Carvalho and Buonomano (2011), we suggest that different parameters of the model describing STP might be related to different types of coding.

2. MATERIALS AND METHODS

2.1. SINGLE NEURON MODEL

Each neuron is modeled as in Carvalho and Buonomano (2011): the sub-threshold dynamics of the electrical potential V_i of the generic neuron i are described by the equation:

$$\frac{dV_i}{dt} = -g_L V_i + \sum_{j=1, j \neq i}^N g_{ij} (E_{rev} - V_i), \quad (1)$$

where $E_{rev} = 30 \text{ mV}$ is the reversal potential and $g_L = 0.1 \mu\text{S}$ is the leak conductance - both quantities are equal and fixed for all neurons. $\{g_{ij}\}_{i,j=1,\dots,N}$ is the matrix of conductances and the generic element g_{ij} represents the conductance of the synapse going from neuron j to neuron i . Upon arrival of a presynaptic action potential elicited by neuron j , each of the conductances g_{ij} with $i = 1, \dots, N$, $i \neq j$ increases by a quantity w_{ij} , called effective synaptic efficacy, and decays exponentially back to zero with a fixed time constant $\tau_g = 10 \text{ ms}$, equal for all synapses:

$$\frac{dg_{ij}}{dt} = -\frac{g_{ij}}{\tau_g} + \sum_f w_{ij} \delta(t - t_j^f), \quad (2)$$

where t_j^f is the f -th spike emitted by neuron j . The effective synaptic efficacy depends on both presynaptic and postsynaptic factors:

$$w_{ij} = r_{ij} u_{ij} A_{ij}, \quad (3)$$

where r_{ij} and u_{ij} are the presynaptic variables representing depression and facilitation in the STP model (see Subsection 2.2) and A_{ij} is the postsynaptic variable for the maximum synaptic strength (or absolute efficacy), which represents the maximum synaptic response (see Subsection 2.7). If $V_i(t) \geq 1 \text{ mV}$ a spike is elicited by neuron i and $V_i(t + dt)$ is set to 0 for the next $t_{ref} = 10 \text{ ms}$ (refractory period).

2.2. STP MODEL

Short-term synaptic plasticity is described at each synapse through the evolution of two variables, r_{ij} and u_{ij} , representing

the degree of depression and facilitation of the synapse connecting neuron j to neuron i . The time course of r_{ij} and u_{ij} is given by the following kinetic equations (Markram et al., 1998b; Maass and Markram, 2002):

$$\frac{dr_{ij}}{dt} = \frac{1 - r_{ij}}{\tau_{recij}} - \sum_{j=1, j \neq i}^N \sum_f r_{ij} u_{ij} \delta(t - t_j^f) \quad (4)$$

$$\frac{du_{ij}}{dt} = \frac{U_{ij} - u_{ij}}{\tau_{facilij}} + \sum_{j=1, j \neq i}^N \sum_f U_{ij} (1 - u_{ij}) \delta(t - t_j^f) \quad (5)$$

U_{ij} , τ_{recij} and $\tau_{facilij}$ are the parameters of the model and they represent, respectively: fraction of resources used by the first action potential, time constant of recovery from depression and time constant of synaptic facilitation. A learning rule for STP has to allow changes to (at least one of) these parameters. At each synapse, the product of r_{ij} and u_{ij} determines the presynaptic efficacy.

2.3. STDP MODEL

We use the triplet learning rule defined by Pfister and Gerstner (2006) with hard bounds: maximum weights can only vary in the interval $[A_{min}, A_{max}]$. In this model, each neuron has two presynaptic variables m^1, m^2 and two postsynaptic variables o^1, o^2 . In the absence of any activity, these variables exponentially decay toward zero with different time constants:

$$\begin{aligned} \tau_{m_1} \frac{dm_i^1}{dt} &= -m_i^1 & \tau_{m_2} \frac{dm_i^2}{dt} &= -m_i^2 \\ \tau_{o_1} \frac{do_i^1}{dt} &= -o_i^1 & \tau_{o_2} \frac{do_i^2}{dt} &= -o_i^2 \end{aligned} \quad (6)$$

whereas when the neuron elicits a spike they increase by 1:

$$m_i^1 \rightarrow m_i^1 + 1 \quad m_i^2 \rightarrow m_i^2 + 1 \quad o_i^1 \rightarrow o_i^1 + 1 \quad o_i^2 \rightarrow o_i^2 + 1. \quad (7)$$

Then, assuming that neuron i fires a spike, the STDP implementation of the triplet rule can be written as follows:

$$\begin{cases} \Delta A_{ji}^{STDP} = -\gamma o_j^1(t) [A_2^- + A_3^- m_i^2(t - \epsilon)] \\ \Delta A_{ij}^{STDP} = +\gamma m_j^1(t) [A_2^+ + A_3^+ o_i^2(t - \epsilon)] \end{cases} \quad (8)$$

where γ is the learning rate, ϵ is an infinitesimal time constant to ensure that the values of m_i^2 and o_i^2 used are the ones right before the update due to the spike of neuron i , and A_{ij} is the maximum strength of the connection from j to i . Values of STDP amplitudes are taken from Pfister and Gerstner (2006) and are listed in Table 1.

In order to set A_{min} we note that if the maximum weights connecting the input neurons to a specific output neuron all collapse to zero in the low firing rate regime, then, in the subsequent high firing rate regime, inputs were not able to “wake up” this neuron: it remained almost silent all the time. To avoid this, we set $A_{min} = 10^{-3}$. With such a small value we can still apply the symmetry measure (Esposito et al., 2014), which assumes $A_{min} = 0$, see Subsection 2.9, to evaluate the symmetry of the network.

Table 1 | Parameters used in simulations.

Symbol	Description	Value
N	Number of total neurons	{40, 80}
N_{in}	Number of input neurons	{30, 60}
N_{out}	Number of output neurons	10
E_{rev}	Reversal potential	30 mV
g_L	Decay constant of neuron potential	$0.1 \mu S$
τ_g	Decay constant of synaptic conductances	10 ms
V_{thr}	Threshold potential for spike emission	1 mV
t_{ref}	Refractory period	10 ms
v_{in}	Input firing rate	10 Hz
v_{targ}	Output target firing rate	{2, 20} Hz
τ_{facil}^{min}	Facilitation time constant - minimum value	1 ms
τ_{facil}^{max}	Facilitation time constant - maximum value	000 ms
τ_{rec}^{min}	Depression time constant - minimum value	100 ms
τ_{rec}^{max}	Depression time constant - maximum value	900 ms
U^{min}	Synaptic utilization - minimum value	0.05
U^{max}	Synaptic utilization - maximum value	0.95
$\bar{\eta}$	Fixed learning rate for $U, \tau_{rec}, \tau_{facil}$	0.1
A_2^+	Amplitude of maximum weights change - pair term in Long-Term Potentiation	4.6×10^{-3}
A_3^+	Amplitude of maximum weights change - triplet term in Long-Term Potentiation	9.1×10^{-3}
A_2^-	Amplitude of maximum weights change - pair term in Long-Term Depression	3.0×10^{-3}
A_3^-	Amplitude of maximum weights change - triplet term in Long-Term Depression	7.5×10^{-9}
τ_{m_1}	Decay constant of presynaptic indicator m_1	16.8 ms
τ_{m_2}	Decay constant of presynaptic indicator m_2	575 ms
τ_{o_1}	Decay constant of postsynaptic indicator o_1	33.7 ms
τ_{o_2}	Decay constant of postsynaptic indicator o_2	47 ms
A_{min}	Lower bound for maximum synaptic weights	0.001
A_{max}	Higher bound for maximum synaptic weights	1
γ	Learning rate for STDP and for STP-dependent Δw	{1, 2}
dt	Discretization time step	1 ms

STDP parameters are as in the nearest-spike triplet-model, described in Pfister and Gerstner (2006).

2.4. LEARNING TASK

Neurons are divided into different populations, each of them is required to fire at one of the two target firing rates: 30 Hz (*high*) or 5 Hz (*low*). To allow the populations to reach their target rate, both short- and long-term plasticity parameters are adapted via error-driven learning (see Subsection 2.6) and, in addition, the maximum synaptic strength is shaped by the STDP triplet rule (see Subsection 2.3).

2.5. INPUT SIGNAL AND INPUT NEURONS

In all simulations, the input signal is delivered only to a subset of neurons in the network, which we call input neurons N_{in} . Each of these neurons receives a pulse-like stimulus with a fixed frequency $v_{in} = 10$ Hz, whose amplitude (2 mV) is chosen to always elicit an action potential in the corresponding input neuron. The stimulus delivery, however, is not synchronous across the input neurons, but it follows a *sequential protocol*: neurons are stimulated one

after another with a fixed time delay t_{delay} and in a fixed order. We choose $t_{delay} = (v_{in} N_{in})^{-1}$ so that neurons that belong to input cyclically receive a stimulus. To further explain this, one may imagine labeling the neurons depending on the order they receive the stimulus, and therefore on the firing order, then have the firing pattern $N_1, N_2, N_3, \dots, N_{N_{in}}, N_1, N_2, N_3, \dots, N_{N_{in}}, N_1, \dots$, with each pair of adjacent spikes being separated by a time interval of t_{delay} . We can think of the N_{in} neurons as if they are organized in a ring and the stimulus as a cyclically traveling wave across this ring. To include the effect of noise, a random Gaussian variable with zero mean and standard deviation equal to $0.1 t_{delay}$ is added to the firing times. The magnitude of the standard deviation is such that there is no inversion in the firing order. With this construction, the stimulus delivered to input neurons can be thought as generated by an external (not explicitly simulated) population of neurons where each external neuron projects only onto one corresponding input neuron.

Note that, by construction, in the absence of any other signal, the firing pattern of the input neurons reflects that of the stimulus. This means that the external signal implicitly fixes a level of minimum activation for the N_{in} neurons: their firing rate cannot be smaller than v_{in} . Due to this constraint, the input neurons, despite being free to change their parameters according to STP learning rules (see Subsection 2.6), are not totally free to regulate their firing activity, which may prevent them from effectively fulfilling the task. The rest of the neurons, instead, are totally free to adapt their activity and are called output neurons. For these reasons, we read out the interesting quantities only from output neurons (we refer to Results and to **Figures 1A, 4A** for more details on the architecture).

2.6. ERROR-DRIVEN LEARNING RULE FOR STP

The task can be formulated as an optimization problem where neurons regulate their own activity in order to minimize the objective function defined as:

$$E = \left(\frac{v_{targ} - \langle v \rangle}{v_{lim}} \right)^2, \quad (9)$$

where v_{lim} is the maximum allowed frequency due to the refractory period ($v_{lim} = 1/t_{ref}$), v_{targ} is the target firing rate and $\langle v \rangle$ is the mean firing rate across a single population. To calculate firing rates of single neurons v_i we use an exponential moving average with time constant $\tau_v = 1$ s:

$$\tau_v \frac{dv_i}{dt} = -v_i + \hat{v}_i \quad (10)$$

where \hat{v}_i is the current firing rate, which basically reflects if neuron i has fired ($\hat{v}_i = 1$ Hz) or not ($\hat{v}_i = 0$ Hz). The population mean firing rate is therefore:

$$\langle v \rangle = \frac{1}{N_{pop}} \sum_{i=1}^{N_{pop}} v_i \quad (11)$$

with N_{pop} being the size of the population.

By following a standard procedure, learning rules can be derived from Equation (9) by applying the gradient descent method (Hertz et al., 1991). Since the task is not based on single neurons but it involves an entire population, we use a mean-field approach for the derivation of the learning rules. Therefore, from now on in this section, we switch from the above single neuron notation to mean-field variables, by dropping the ij indices. It is worth noting that in our formulation the target is achieved not by directly acting on the firing rates, but by tuning the STP parameters, which in turn affects the firing itself. Therefore, $\langle v \rangle = \langle v \rangle (U, \tau_{rec}, \tau_{facil})$ and by using the chain rule we can formally write the following update rule for each parameter p :

$$\Delta p = -\eta_p \frac{\partial E}{\partial p} = -\eta_p \frac{\partial E}{\partial \langle v \rangle} \frac{\partial \langle v \rangle}{\partial p} = 2\eta_p \frac{v_{targ} - \langle v \rangle}{v_{lim}^2} \frac{\partial \langle v \rangle}{\partial p}, \quad p = U, \tau_{rec}, \tau_{facil} \quad (12)$$

where η_p is the learning rate, which in principle could be different for each parameter. The form of the function $\langle v \rangle (U, \tau_{rec}, \tau_{facil})$ can be derived with a semi-heuristic procedure, following (Vasilaki and Giugliano, 2014). Whenever possible, for the mean-field variables we use the same symbols as in Vasilaki and Giugliano (2014) for consistency. Thus, we introduce the mean-field variables u , x , U , and A , respectively describing facilitation, depression, synaptic utilization and maximum strength. We assume a threshold-linear gain function between input mean current h and output mean firing rate $\langle v \rangle = a[(h - \vartheta)]_+$, for some constants a , ϑ . We can then write the dynamic mean-field equations for a population of neurons recurrently connected by short-term synapses as follows (Chow et al., 2005):

$$\begin{cases} \tau \dot{h} = -h + Aux \langle v \rangle + I_{ext} \\ \dot{x} = \frac{1-x}{\tau_{rec}} - ux \langle v \rangle \\ \dot{u} = \frac{U-u}{\tau_{facil}} - +U(1-u) \langle v \rangle \end{cases} \quad (13)$$

where I_{ext} is the mean external current and τ is a decay-constant. By imposing equilibrium conditions, $\dot{h} = \dot{x} = \dot{u} = 0$, and combining the resulting equations, we can finally write:

$$h = F(\langle v \rangle_h) = \frac{AU(\langle v \rangle^{-1} + \tau_{facil})}{\langle v \rangle^{-2} + \langle v \rangle^{-1} U \tau_{facil} + \langle v \rangle^{-1} U \tau_{rec} + U \tau_{facil} \tau_{rec} + I_{ext}} \quad (14)$$

Now we observe that by taking the limit $h \rightarrow \infty$ in $F(\langle v \rangle_h)$ we obtain an upper bound for the maximum allowed firing rate $\langle v \rangle \leq \frac{A}{\tau_{rec}} + I_{ext}$ (for more details see Vasilaki and Giugliano, 2014). We can heuristically turn the above inequality into an equality:

$$\langle v \rangle = \frac{A}{\tau_{rec}} + I_{ext} \quad (15)$$

so as by plugging Equation (15) into Equation (12) we can finally obtain an explicit form for the learning rule. In particular, since

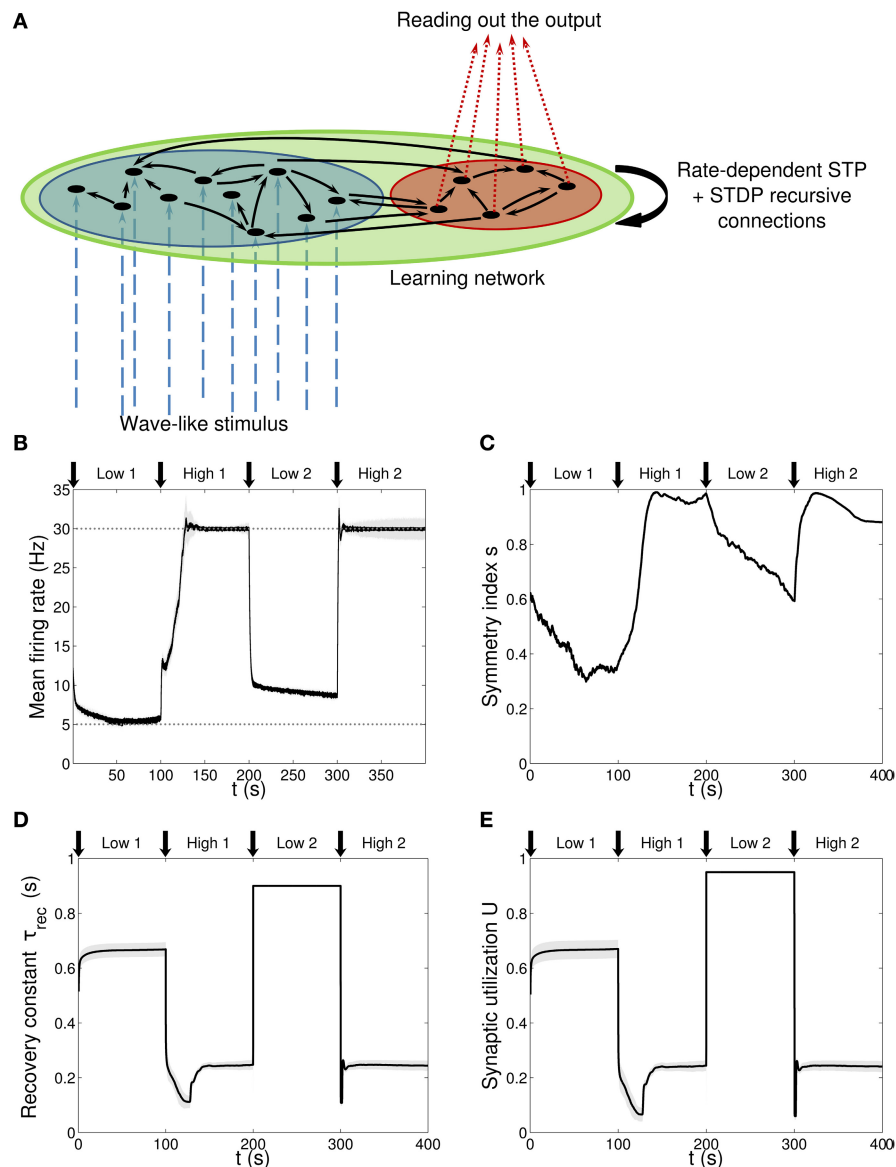


FIGURE 1 | Single population scenario: network architecture and activity, connectivity and STP parameters adaptation in the output population with (U, τ_{rec}) learning scheme (Part 1). (A) Architecture. The learning network (green) is divided into an input region (blue) and an output region (red). Connections (black arrows) are all-to-all and obey both Spike-Timing Dependent Plasticity and rate-dependent Short-Term Plasticity. Input neurons receive an external wave-like stimulus (blue dashed arrows). (B) Mean firing rate of the output population. Shaded area represents standard deviation, horizontal dotted gray lines show the two

target firing rates (high = 30 Hz, low = 5 Hz) and vertical black arrows mark the onset of the four dynamic phases alternating the target according to the sequence low-high-low-high. (C) Symmetry measure applied on the connectivity of the output population. In accordance to the target, connectivity switches between unidirectionality (low values) and bidirectionality (high values). (D,E) Mean values of the recovery time constant τ_{rec} and synaptic utilization U for the synapses projecting onto the output neurons. We observe depression (high values) at low firing rates and facilitation (low values) at high firing rates.

only one of the three parameters appears in Equation (15), we have a single rule for τ_{rec} only:

$$\Delta\tau_{rec} = -2\eta_{\tau_{rec}} (v_{target} - \langle v \rangle) \frac{A}{v_{lim}^2 \tau_{rec}^2} \quad (16)$$

Then, according to the above derivation, the only parameter that needs to be learnt is τ_{rec} . Here we adopt the view (Tsodyks and Markram, 1997; Markram et al., 1998b; Thomson, 2000; Chow et al., 2005) that facilitation/depression corresponds to small/large values of τ_{rec} and U as well. Therefore, assuming that

they apparently play a similar role, we can heuristically take a similar dependence of $\langle v \rangle$ upon U : $\langle v \rangle = \frac{A}{U} + I_{ext}$, which leads us to a similar learning rule:

$$\Delta U = -2\eta_U (v_{targ} - \langle v \rangle) \frac{A}{v_{lim}^2 U^2} \quad (17)$$

With the same heuristic argument we can also write down a relation involving τ_{facil} . Indeed, it is well-known that facilitation/depression corresponds to large/small values of τ_{facil} , so we can hypothesize a linear relation, also including the dependence on the maximum strength for similarity with the other parameters. Thus, $\langle v \rangle \propto A\tau_{facil} + I_{ext}$, which gives the following learning rule:

$$\Delta \tau_{facil} = 2\eta_{\tau_{facil}} (v_{targ} - \langle v \rangle) \frac{A}{v_{lim}^2} \quad (18)$$

Finally, based on the fact that A turns out to appear in Equation (15), and supported by experimental results showing an interaction between STP and STDP (Markram et al., 1997; Sjöström et al., 2003), we can also introduce a STP-dependent learning rule for the maximum synaptic strength:

$$\begin{aligned} \Delta A^{STP} &= -\eta_A \frac{\partial E}{\partial A} = -\eta_A \frac{\partial E}{\partial \langle v \rangle} \frac{\partial \langle v \rangle}{\partial A} \\ &= 2\eta_A (v_{targ} - \langle v \rangle) \frac{1}{v_{lim}^2 \tau_{rec}}. \end{aligned} \quad (19)$$

This synaptic modification clearly does not substitute the traditional STDP, since the two rules come from different mechanisms. Rather, we assume they both contribute to maximum weights changes (see Subsection 2.7).

2.7. SINGLE NEURON LEARNING FRAMEWORK: COMBINING STDP AND STP LEARNING MODELS

Equations (16–19) are mean field learning rules for the four parameters τ_{rec} , U , τ_{facil} , A . It is straightforward to turn them into single neuron online learning rules. From now on, we return to a single neuron notation. Similarly to STDP, we hypothesize that modifications of STP are triggered by postsynaptic events: every time neuron i elicits a spike, its current firing rate is updated as well as the mean population firing rate. Neuron i can therefore backwards regulate its incoming synapses, through the following set of equations:

$$\Delta \tau_{recij} = -2\eta_{\tau_{recij}} (v_{targ} - \langle v \rangle) \frac{A_{ij}}{v_{lim}^2 \tau_{recij}^2} \quad (20)$$

$$\Delta U_{ij} = -2\eta_{U_{ij}} (v_{targ} - \langle v \rangle) \frac{A_{ij}}{v_{lim}^2 U_{ij}^2} \quad (21)$$

$$\Delta \tau_{facilij} = 2\eta_{\tau_{facilij}} (v_{targ} - \langle v \rangle) \frac{A_{ij}}{v_{lim}^2} \quad (22)$$

$$\Delta A_{ij}^{STP} = 2\eta_{A_{ij}} (v_{targ} - \langle v \rangle) \frac{1}{v_{lim}^2 \tau_{recij}}. \quad (23)$$

The firing event of the neuron i also triggers STDP, according with Equation (8). This contribution sums up with the above STP-dependent change, so as the total modification of the maximum synaptic strength is:

$$A_{ij} \longrightarrow A_{ij} + \Delta A_{ij}^{tot}, \quad \Delta A_{ij}^{tot} = \Delta A_{ij}^{STDP} + \Delta A_{ij}^{STP}. \quad (24)$$

Note that when we converted mean field population equations into single neuron equations we kept the population mean firing rate $\langle v \rangle$, instead of turning it into the single rate v_i . This is because the task is defined at a population level. Learning rates of the three STP parameters are chosen to be equal and error-dependent:

$$\eta_{p_{ij}} = \bar{\eta} \left(1 + \frac{v_{targ} - \langle v \rangle}{v_{lim}} \right)^2, \quad p = U, \tau_{rec}, \tau_{facil}, \quad (25)$$

with $\bar{\eta} = 0.1$. The learning rate for maximum synaptic strength, instead, is fixed in time and it is the same as the one used for STDP, $\eta_{A_{ij}} \equiv \gamma$.

Now we have four single neuron rules for the STP learning model, plus an equation for STDP and an equation for combining the different rules for the maximum synaptic strength. All these six rules together, Equations (8, 20–24) form a complete learning scheme for each neuron, which is implemented in our simulations. These rules are now local, since their computation takes place separately in each neuron, but receive a global signal encoding for the task performance error.

2.8. INVESTIGATION OF DIFFERENT RULE COMBINATIONS

In the Results section we consider different learning mechanisms: in addition to STDP, that is crucial for the formation of motifs (Vasilaki and Giugliano, 2012, 2014), different combinations of the four STP rules are taken into account while the remaining parameters are kept fixed. At first we allow only two parameters to change: (i) τ_{rec} , because Equation (15) implies that for high frequencies this is the only critical parameter for adapting the firing rate of the population, and (ii) U , since it was a key parameter adopted in the work in Carvalho and Buonomano (2011). Then, we introduce the STP-dependent rule on the maximum synaptic strength, Equation (23), with the view to observe a more stable learning process. Following this, we also include τ_{facil} in the learning scheme for a full parameter adaptation (*full model*) and finally we investigate the minimal number of parameters that needs to be adapted (*minimal model*), based on Equation (15). Looking for other parameter combinations might not be meaningful, as Equation (15) indicates the key parameters that are involved in changing the mean firing of the population.

2.9. CONNECTIVITY ANALYSIS

To reveal the type of connectivity in the output population, we use a symmetry index defined as a measure of the symmetry of the connectivity matrix W (Esposito et al., 2014):

$$s = 1 - \frac{2}{N(N-1) - 2M} \sum_{i=1}^N \sum_{j=i+1}^N \frac{|A_{ij} - A_{ji}|}{A_{ij} + A_{ji}}. \quad (26)$$

Here M is the number of instances where both A_{ij} and A_{ji} are zero, i.e., there is no connection between two neurons. Since in our case connections are bounded in the interval $[10^{-3}, 1]$, $M = 0$ all the time. Equation (26) is able to capture the presence of global non-random structures in a network, returning a value included in $[0, 1]$. Values of s close to 1 reflect the presence of a global bidirectional motif, whereas when s approaches 0, a unidirectional motif is emerging. Note that, in order to apply the measure Equation (26), we assume that the lower bound for connections is 0. However, the choice of a small value such as 10^{-3} does not affect the measure.

2.10. DATA SHARING

We provide the scripts that were used to construct the main figures of the paper in the ModelDB database, accession number:169242.

3. RESULTS

3.1. SINGLE POPULATION WITH A TIME-VARYING TASK: A CONTINUUM BETWEEN FACILITATION AND DEPRESSION

First, we apply our learning model to a specific task demonstrating how synapses can change their behavior driven by an external feedback signal. The problem we study is simple: a population of neurons is presented with a stimulus and is required to produce a certain output as a response to that stimulus. Once the learning has been successful, for the same input signal the desirable output changes. In other words, neurons are trained to respond to a change in the associative paradigm (*inverted associative learning problem*), that can be due to, for instance, changes in the environmental conditions.

Let us give a concrete example of an inverted associative learning problem, taken from Asaad et al. (1998). In their work, the authors trained monkeys to associate visual stimuli (pictures) with delayed saccadic movements, left or right, with associations being reversed from time to time. Monkeys had to go beyond learning a single cue-response association: they are required to learn to associate, on alternate blocks, two cue objects with two different saccades. In other words, after having learned the relation $\{\text{object A, go right}\}$, and $\{\text{object B, go left}\}$, the associations were reversed such that now they needed to learn $\{\text{object A, go left}\}$ and $\{\text{object B, go right}\}$.

Similar to the (Asaad et al., 1998) experiment, we assume a binary problem, i.e., environmental conditions can change only between two states, and we measure the neurons' activity in terms of firing rate. This means that neurons are initially asked to fire at some rate and, after learning this task, they are asked to fire at a different rate, while keeping the same input signal all the time. Thus, the problem we defined is a simpler version of the monkey experiment, with only a single input. In order to train the neurons on the current associative paradigm, an external global signal is required, that can be considered as an error signal (see Section Methods 2.6 and 2.7).

3.1.1. Problem description and network architecture

We created a learning network of $N = 40$ conductance-based integrate-and-fire neurons (see Section Methods 2.1) all to all connected. Synaptic connections are modified by the STDP triplet rule (Pfister and Gerstner, 2006) and STP is implemented by

using the Tsodyks and Markram model (TM model) described in Markram et al. (1998b); Maass and Markram (2002).

Figure 1A shows the network architecture, composed by two non-overlapping regions: a *blue* one with $N_{in} = 30$ neurons receiving the input signal and a *red* one with $N_{out} = 10$ neurons from which we read out the quantities of interest. Note that for clarity, only a few neurons (*black circles*) and connections (*black arrows*) are drawn. The network is therefore formed by two functionally distinct populations, with the input population delivering the stimulus to the output one. Recursive connections are present within each population and across populations, and they are all plastic, in the sense of both long-term and STP. We refer to this architecture as a first or single population scenario.

The input neurons are stimulated one after the other, following a *sequential protocol*, and approximately with the same rate, $v_{in} = 10$ Hz. The amplitude of the stimulus is such that input neurons release a spike every time they receive an input (see Section Methods 2.5). This external source can be thought as an additional population of neurons, which we are not simulating here, where each "external" neuron is connected only with a corresponding neuron in the input population by means of a fixed feedforward connection (*blue dashed arrows*).

We hypothesize that the whole learning network (*green region in Figure 1A*) is presented with a sequence of two tasks while the stimulus pattern is kept fixed. The tasks are firing low (5 Hz) and firing high (30 Hz) and the sequence is *low-high-low-high*. Therefore, neurons have to repeatedly learn a new association and forget the previous one in a dynamic context divided in four phases of $t_{ph} = 100$ s. We refer to them as: low 1, high 1, low 2, high 2. As discussed at the beginning of this section, this picture is inspired by a typical inverted associative learning problem: considering the monkey experiment from Asaad et al. (1998) as a metaphor, our scenario provides a simplified version, where instead of having two different inputs, *object A* and *object B*, we have a single input. Indeed, we can think we are presenting the network with only *object A* and while doing this we switch the target association between the two states *go right* and *go left*, which correspond to our low and high firing rate targets. We call the desirable context-dependent target rate, v_{targ} . As described in Methods, the difference between v_{targ} and the current firing rate of each population $\langle v \rangle$ is the error signal that, according which our rate-dependent STP causes synapses to adapt their activity.

In all simulations, single neuron parameters $\{U, \tau_{rec}, \tau_{facil}, w\}_{ij}$ are initially drawn from uniform distributions (for $i \neq j$), respectively in $[0.05, 0.95]$, $[100, 900]$ ms, $[1, 900]$ ms, $[10^{-3}, 1]$. Synaptic variables are initialized at their equilibrium values, i.e., $r_{ij} = 1$ and $u_{ij} = U_{ij}$. All the simulations in this subsection use $\gamma = 1$ for the high rate regime and $\gamma = 2$ for the low rate regime. Values of the parameters are listed in **Table 1**.

3.1.2. Learning U and τ_{rec}

We initially studied the problem with a learning scheme involving U and τ_{rec} only, Equations (20, 21), so there is no additional change in maximum synaptic strengths due to STP. Indeed, due to Equation (15) and (Carvalho and Buonomano, 2011), we wanted to test the hypothesis that U and τ_{rec} are the only crucial

parameters that need to be learnt for adapting the firing rate of a population. The results are displayed in **Figures 1B–E**, with *vertical black arrows* marking the beginning of each of the four phases, and in **Figure 2**.

Figure 1B shows the average firing rate of the N_{out} neurons, with *shaded area* being the standard deviation. Target firing rates are shown with *gray dotted lines*. The adaptation to the new target is fast, except during the *low 2* phase, when we switch from high to low rate, where an initial fast decrease of the firing rate is followed by a much slower adaptation. Despite the fact that neurons do not reach the target rate during this phase, we observe a monotonically decreasing activity which would eventually stabilize at 5 Hz if we were allowing the simulation to run for longer. The reason for this double slope adaptation will be further discussed now.

Figure 1C shows the evolution of the symmetry index (see Section Methods 2.9). At the beginning, the value reflects the randomness in the connections (the mean value of s for a network with uniform random connections is indeed $\simeq 0.614$, see Esposito et al., 2014), whereas, as learning takes place, we observe the development of unidirectional (low values of s) and bidirectional (high values of s) motifs, depending on the set target. This can also be formalized by applying the p -value hypothesis test obtained by using mean and variance of s on a completely random network with uniform distribution of connections (Esposito et al., 2014). P -values are shown in **Table 2**. We, again, observe rather slow dynamics during the *low 2* phase that, within the fixed simulation time, prevent the system from reaching a clear connectivity configuration. However, the trend of s clearly shows that the connectivity within the output population is approaching unidirectionality.

Figures 1D,E depicts the time course of the recovery time constant τ_{rec} and synaptic utilization U averaged across the output neurons, with *shaded area* representing standard deviation. Both parameters oscillate between high values, which correspond to depressing behavior, and low values, that indicate facilitation. Note that the dynamics of τ_{rec} and U is fast in all phases, the third included. This is not surprising since STP is a fast process and leads to fast adaptation of its parameters. As a result, neurons' response to a change in the target rate takes place in a short time. However, during the *low 2* phase, synaptic parameters saturate before the neurons could fulfill the task, with STDP being the only remaining mechanism through which the output population can regulate its own activity. This results in a much slower decrease toward the target rate for two reasons: (i) STDP by itself acts on much longer time scales, (ii) switching from high to low rate is the most challenging part of the entire sequence of tasks due to the saturation of the maximum weights in the previous *high 1* phase, which slows down the process even further.

Figure 2 provides additional evidence of the alternation between the two different synaptic behaviors. Plots are organized in five rows, with each row displaying information from a different phase of the simulation. Panel A shows the initial uniform condition, panel B the end of *low 1* phase, etc. For each stage, we draw the histograms of recovery time constant (*Column 1*) and synaptic utilization (*Column 2*). According to the narrow standard deviation observed in **Figures 1D,E**, distributions peak around extreme values, reflecting two different, synaptic

behaviors. *Column 3* in **Figure 2**, displaying the single synapse traces obtained with a TM model, demonstrates the corresponding behaviors: at the end of the phases where neurons are required to fire low we observe a typical depressing response, whereas at the end of the high rate regimes synapses show a typical facilitating trace. To generate the traces, we used a 5 Hz signal to stimulate a synapse with parameters given by the mean values obtained from the corresponding histograms. Note that the synaptic trace for the initial condition, i.e., before learning shapes the parameters, already shows depression, which explains why the distributions of τ_{rec} and U at the end of the *low 1* phase are much broader than in the following phases.

Altogether, the four panels (**B–E**) in **Figure 1** and the five panels (**A–E**) in **Figure 2** show that the properties and activity of the output population oscillate between two states and that the desirable structure is formed depending on the target rate. In particular, we observe that neurons that fire at low frequency turn their synaptic properties into depressing and the connections formed are mostly unidirectional. On the other hand, when the target rate is set at a high frequency, neurons develop facilitating synapses and bidirectional connections.

3.1.3. STP-dependent modification of A enhances performance

Given the speed convergence issue in the *low 1* phase, we introduced an additional learning mechanism, i.e., the STP-dependent rule for A, Equations (23, 24). Indeed, this mechanism provides an additional way, besides the STDP, for regulating the long-term plastic synapses. In all the other aspects, the model remains as above.

Figure 3 shows simulation results, with panels A–D depicting the same quantities as panels B–E in **Figure 1** (symbols as before). A direct panel-by-panel comparison shows that the results are very similar, meaning that with this new learning configuration the output population also learns to adapt its synaptic properties in order to fulfill the current task, with subsequent motifs formation. As expected, due to the additional learning rule for A, the dynamics are faster: in particular, during the *low 1* phase, neurons reach the target rate within the simulation time, and the value of the symmetry measure is much lower than before, confirming the formation of a unidirectional motif; compare with **Figures 1C, 3B** and see **Table 2**. Note that the adaptation of the STP parameters is also faster, as they depend on the current value of the maximum synaptic strength. Thus, the STP-dependent modification of A improves the overall performance and introduces an interesting link between STP and STDP.

3.2. TWO POPULATIONS WITH A DIFFERENT TASK: SYNAPTIC DIFFERENTIATION

Now we consider a different scenario, which we refer to as the second or double population scenario. The two tasks associated with low and high targets are now simultaneously active and must be learnt by different populations, interacting via lateral connections and receiving the same stimulus source. Reasons are multiple: we want to investigate if our model allows to contemporarily encode both associative paradigms, without the need of forgetting one of the two. In addition, we want to study the possibility that target-specific STP emerges as a consequence of the target-dependent

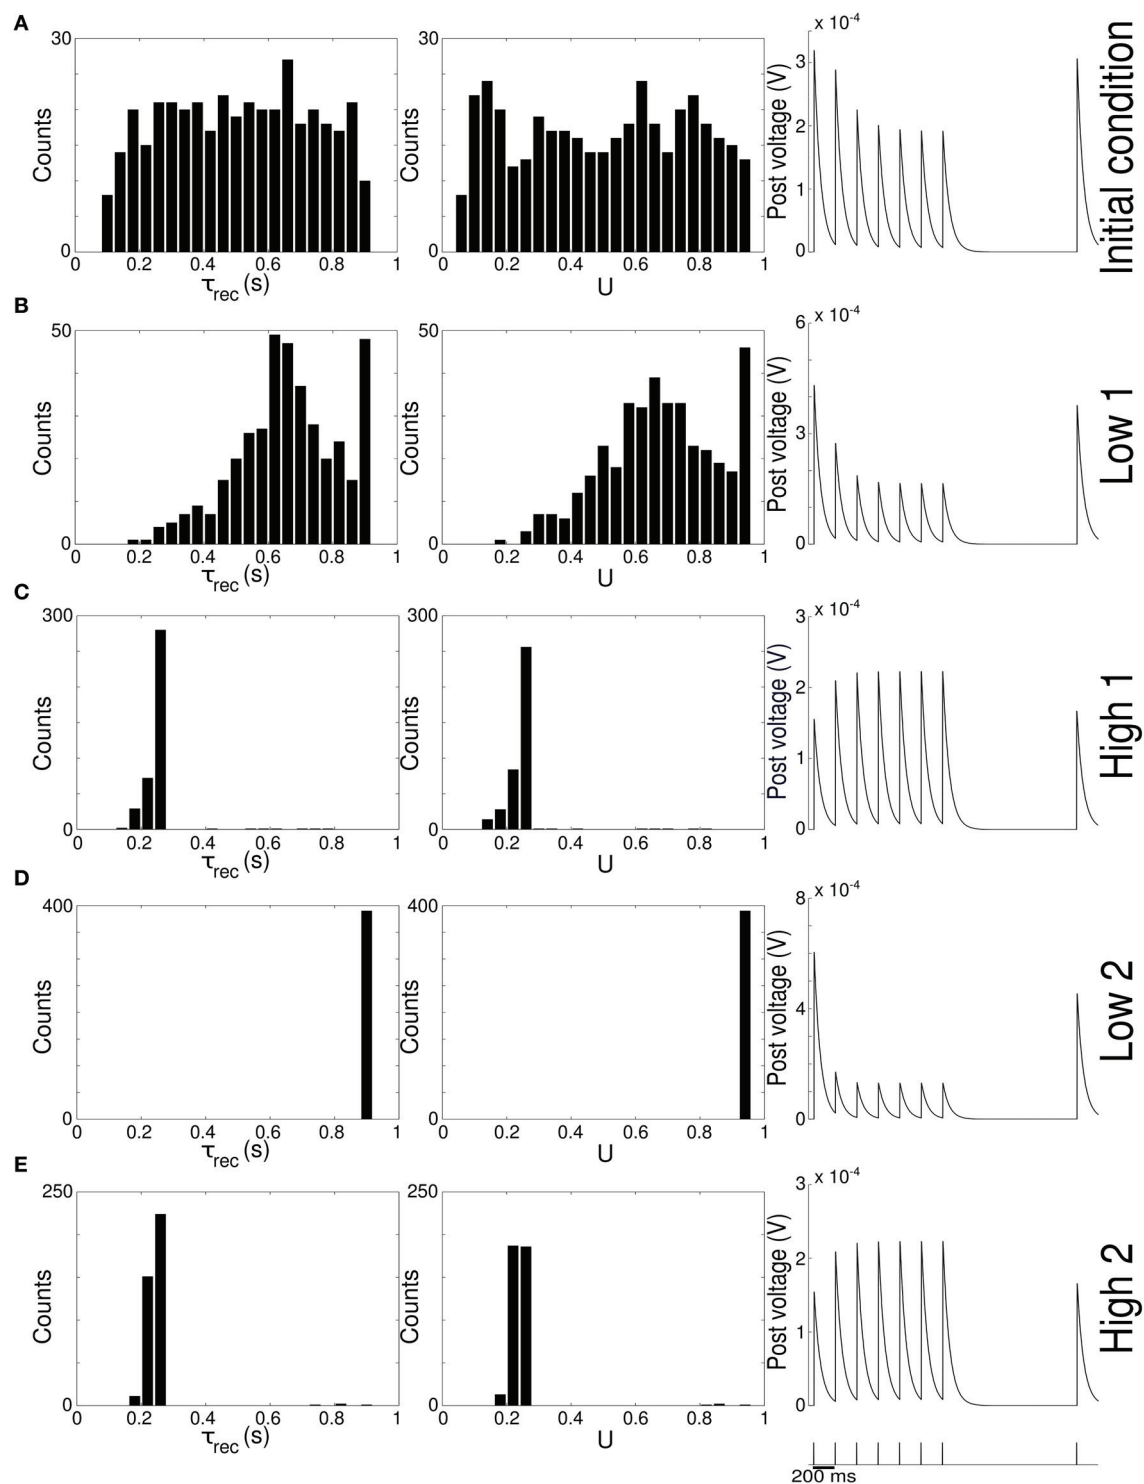


FIGURE 2 | Single population scenario: STP parameters distribution during adaptation with (U, τ_{rec}) learning scheme (Part 2). Different phases of the dynamics are represented. **(A)** Initial (uniformly random) condition. **(B)** End of *low 1* phase (target rate is 5 Hz). **(C)** End of *high 1* phase (30 Hz). **(D)** End of *low 2* phase (5 Hz). **(E)** End of *high 2* phase (30 Hz). **Columns 1, 2** Histograms of recovery time constant and synaptic

utilization of the synapses projecting onto the output neurons. Low values indicates facilitation whereas high values suggest depression. **Column 3** Single synapse traces obtained with the TM model by applying a 5 Hz stimulus. Synaptic parameters used are mean values obtained from the distributions drawn in **(A,B)**. Synapses display a clear alternation between depressing and facilitating behavior.

Table 2 | Symmetry measure and p -value for the single population scenario with $\{\tau_{rec}, U\}$ and $\{\tau_{rec}, U, A\}$ learning schemes.

Phase	(τ_{rec}, U) scheme		(τ_{rec}, U, w) scheme	
	s	p -value	s	p -value
Low 1 (0-500 s)	0.36	1.82×10^{-9}	0.28	1.13×10^{-15}
High 1 (500-1000 s)	0.98	5.47×10^{-19}	0.99	8.73×10^{-20}
Low 2 (1000-1500 s)	0.59	6.47×10^{-1}	0.41	2.19×10^{-6}
High 2 (1500-2000 s)	0.88	1.50×10^{-10}	0.82	3.63×10^{-7}

Column 1: the four phases of dynamics with the corresponding simulation time. Columns 2,3: symmetry measure and p -value for the adaptation of $\{\tau_{rec}, U\}$. Values are computed at end of each phase and by considering output neurons only. Columns 4,5 same as columns 2 and 3 except for the adaptation of $\{\tau_{rec}, U, A\}$.

learning rules we chose for our model. In particular, we want to test whether our model is able to reproduce existing experimental data, specifically that appearing in **Table 1** from the paper by Wang et al. (2006).

3.2.1. Network architecture

The new configuration is depicted in **Figure 4A** and it is obtained by mirroring the structure of the first scenario and by adding recursive connections between functionally homologous populations. This led to a network of $N = 80$ conductance-based integrate-and-fire neurons, organized in two distinct branches of 40 neurons each, with the first branch required to fire at a high rate ($\nu = 30$ Hz) and the second branch at a low rate ($\nu = 5$ Hz). Both targets remain fixed throughout the entire simulation. Each branch is a replication of the architecture we used previously, i.e., it is formed by an input and an output population recursively connected. Thus, the network is formed by four functionally different populations: ϕ_1^{in} , ϕ_2^{in} , ϕ_1^{out} , ϕ_2^{out} , with obvious meaning of symbols. Input populations in both branches receive the stimulus from the same source: a single wave-like signal is delivered to the $N_{input} = 60$ neurons with $\nu = 10$ Hz, stimulating one neuron per time (see Section Methods 2.5), first the neurons in ϕ_1^{in} and then the neurons in ϕ_2^{in} . All connections are plastic following the STDP triplet rule and TM model for STP.

Lateral connections are present between the inputs ϕ_1^{in} , ϕ_2^{in} and between the outputs ϕ_1^{out} , ϕ_2^{out} . To stress that they are functionally different, we drew their initial values from a uniform distribution in $[10^{-3}, 10^{-1}]$, but, during the evolution, synapses are allowed to grow up to $A_{max} = 1$ as any other synapse. Furthermore, cross connections between different output and input populations, i.e., between ϕ_1^{in} , ϕ_2^{out} and between ϕ_2^{in} , ϕ_1^{out} are absent. The rest of the connections - within each population and across populations belonging to the same branch - are drawn from a uniform distribution in $[10^{-3}, 1]$ and they are not allowed to exceed this interval during the simulation. STP variables are initialized as in the single population scenario and in all the simulations presented in this subsection we used $\gamma = 2$ as the learning rate.

3.2.2. Full model: adaptation of U , τ_{rec} , τ_{facil} and A

We begin by studying the behavior of the *full model*: all four parameters are modified by our rate-dependent STP, Equations (20–24). Taking into account the modifications of all three STP parameters allows us to make a direct comparison with (Wang et al., 2006). Results are displayed in **Figures 4B–C** and in **Figure 5**.

Figures 4B,C shows the time course of the mean firing rate and symmetry index in both output populations, *black lines* for ϕ_1^{out} and *light gray lines* for ϕ_2^{out} . *Shaded areas* and *dark gray dotted lines* represent standard deviation and target firing rates. Both populations ϕ_1^{out} , ϕ_2^{out} approach the target rate while developing specific connectivity: as expected, a bidirectional motif emerges in the population firing at the high rate whereas the population firing at the low rate develops mostly unidirectional connections.

Figures 5A–C shows the time evolution of the three parameters of the TM model: *black lines* and *gray lines* represent the mean value of the synapses projecting from the two output populations $\phi_1^{out} \cup \phi_2^{out}$, respectively onto ϕ_1^{out} and ϕ_2^{out} . *Shaded area* is the standard deviation. As expected from the previous simulation, we observe that the two populations develop different synaptic types: high values of τ_{facil} and low values of τ_{rec} and U , as observed in the population firing at the high rate, suggest a facilitating behavior, whereas values as the one observed in ϕ_2^{out} , characterize depressing synapses. Mean values at the end of the simulation are reported in **Table 3** rows 1,4. These results show that our model develops target-specific STP and results in good agreement with the data in Wang et al. (2006). Indeed, although single values are not identical, the qualitative synaptic behavior is represented: recalling the notation used in Wang et al. (2006), two main types of synapses are present. The group projecting from $\phi_1^{out} \cup \phi_2^{out}$ onto ϕ_1^{out} can be mapped onto the type E1 and the group projecting from $\phi_1^{out} \cup \phi_2^{out}$ onto ϕ_2^{out} that can be mapped onto the type E2.

Following Wang et al. (2006), we can also refine our classification, introducing a further distinction within each class. With this purpose, we show in **Figures 5D–F** the distributions of τ_{rec} , τ_{facil} and U at the end of the simulation within the entire output population $\phi_1^{out} \cup \phi_2^{out}$. For each histogram, data have been divided into four groups, representing the four different subtypes: ϕ_2^{out} to ϕ_2^{out} with *light gray*, ϕ_1^{out} to ϕ_2^{out} with *medium gray*, ϕ_1^{out} to ϕ_1^{out} with *dark gray*, ϕ_2^{out} to ϕ_1^{out} with *black*. While the distinction between the two synaptic types mapping onto E1 and E2 is evident, the difference between two subtypes in the same type cannot be easily seen. However, by looking at the mean values of synaptic parameters in **Table 3** rows 2, 3, 5, 6 and in particular the ratio τ_{rec}/τ_{facil} in **Table 3** column 5, the distinction into four subtypes becomes more clear. As reported in column 7 of **Table 3**, we can map the synaptic subtypes as follows: E1a corresponds to the group $\phi_1^{out} \rightarrow \phi_1^{out}$, E1b to $\phi_2^{out} \rightarrow \phi_1^{out}$, E2a to $\phi_2^{out} \rightarrow \phi_2^{out}$ and E2b to $\phi_1^{out} \rightarrow \phi_2^{out}$.

Finally, similarly to **Figure 2**, in **Figures 5G–J** we show single synapse traces for each subtype. We observe that, except for the last trace, different groups effectively show a distinctive response to the same stimulus (12 Hz) and the traces reproduce the ones of the corresponding subtypes in Wang et al. (2006).

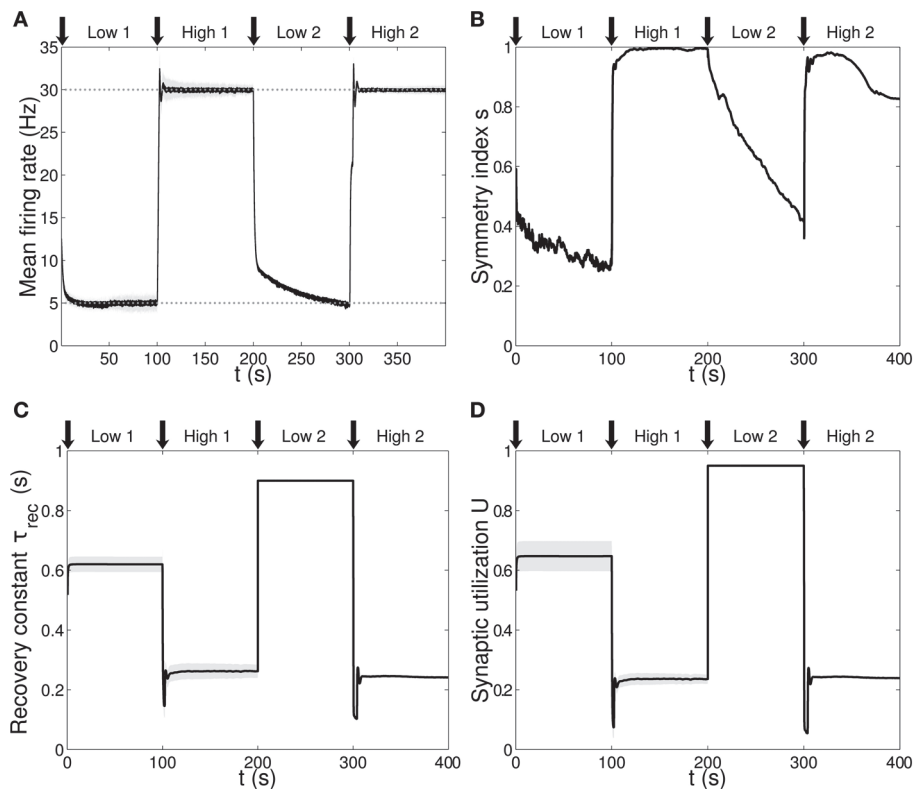


FIGURE 3 | Single population scenario: activity, connectivity and STP parameter adaptation in the output population with (U , τ_{rec} , A) learning scheme. (A) Mean firing rate of the output population. *Shaded area* represents standard deviation, *horizontal dotted gray lines* show the two target firing rates (high = 30 Hz, low = 5 Hz) and *vertical black arrows* mark the onset of the four dynamic phases alternating the targets in the sequence low-high-low-high. **(B)** Symmetry measure applied on the connectivity of the

output population. In accordance with the target rate, connectivity switches between unidirectionality (low values) and bidirectionality (high values). **(C,D)** Mean values of recovery time constant τ_{rec} and synaptic utilization U for the synapses projecting onto the output neurons. We observe depression (high values) at low firing rates and facilitation (low values) at high firing rates. Compared to **Figure 1**, we observe an improvement in the overall performance due to the inclusion of the STP-dependent modification of A .

Although in **Figures 5D–F** we present four different histograms for each parameter, we can reason on the overall distribution within the entire output population $\varphi_1^{out} \cup \varphi_2^{out}$ as the sample size is the same in all histograms. We can therefore observe that the distribution of τ_{rec} closely matches that in Wang et al. (2006), whereas the distribution of U reproduces the peak at around 0.25 but is less broad. On the other hand, the distribution of τ_{facil} is rather different, being totally shifted toward facilitating values in our case. This may be due to the fact that U is much more peaked around low values. We decided then to discard τ_{facil} from the learning scheme and run a simulation where only U , τ_{rec} and A are learnt, as we did for the single population scenario in subsection 3.1.3. We observed that the behavior of the output populations and all the results remain unchanged. We provide an explanation for this in the Discussion.

3.2.3. A minimal model for rate-dependent STP: adaptation of τ_{rec} and A

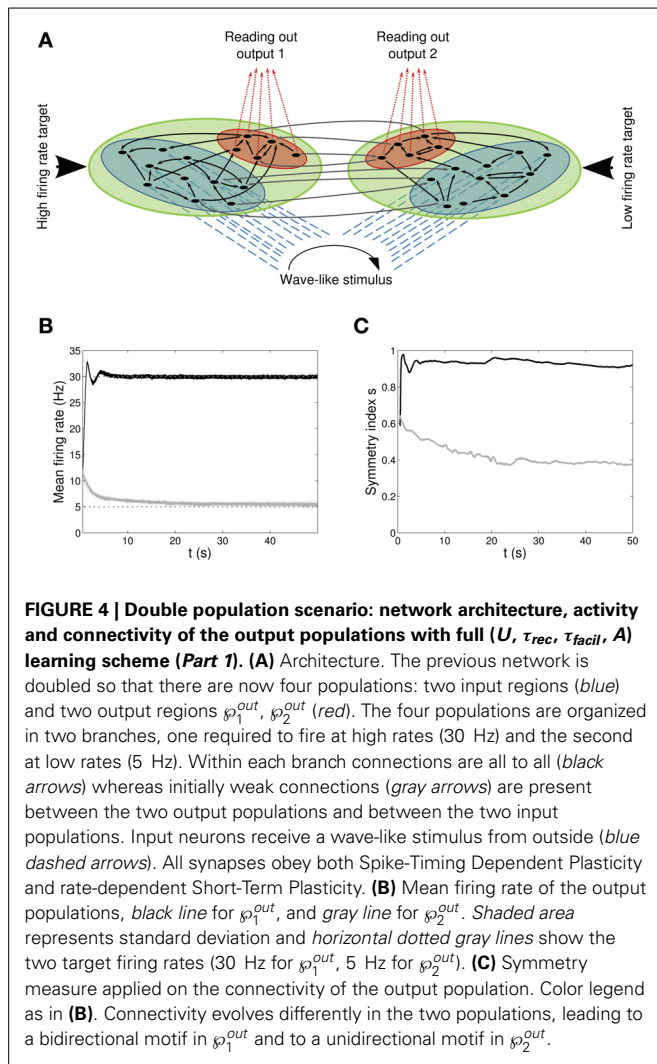
Finally, we study the *minimal model*: a model that suffices to obtain the desired behaviors by adapting as few parameters as possible. The choice of the parameters to be learnt is naturally suggested by the form of the objective function Equation

(9): τ_{rec} and A . Interestingly, this minimal model preserves two key features: (i) both a presynaptic parameter, τ_{rec} , and a postsynaptic parameter, A , participate in learning, (ii) STP and STDP are linked to each other through the STP-dependent modification of A .

In **Figure 6** we show the results of the minimal model: from A to D, respectively: mean output firing rates, symmetry index, τ_{rec} evolution and τ_{rec} distribution in the four groups of synapses. By comparing these panels with the ones from the full model simulation, we observe that output populations still efficiently fulfill the task while developing the expected connectivity motifs. Also, in **Table 4** we report the mean values of τ_{rec} for the four groups of synapses that we identified with the full model: there is still a clear distinction between them. We can therefore conclude that this minimal model is sufficient for qualitatively reproducing the main two types and also the subtypes of Wang et al. (2006).

4. DISCUSSION

It is well-known that synapses are activity-dependent connections through which neurons propagate information. STP is a mechanism that describes these phenomena in short time scales and



introduces two typical synaptic behaviors: depression and facilitation. Contrary to long-lasting modifications of maximum synaptic strengths, for example STDP, existing models of STP do not rely on any learning mechanisms, apart from very few exceptions; see for instance (Carvalho and Buonomano, 2011). Motivated by their work, it is our belief that more efficient dynamics would be possible if synapses were allowed to change their short-term behavior by tuning their own parameters, depending on one or more external controlling factors, for example, their current task. Typically, one asks which is the firing regime for which a certain type of synapse performs better (Barak and Tsodyks, 2007), whereas we are looking at the picture from a reverse perspective: we want to obtain some frequency regime, which is the most efficient way to do it from a synaptic point of view? A similar concept can be found in Natschläger et al. (2001), where the authors trained a network with a temporal structured target signal, using optimization techniques.

In our work, we developed a learning scheme for STP, and we obtained, with a semi-rigorous argument, a learning rule for only one of the three parameters of the TM model, τ_{rec} . Based on specific experimental results (Tsodyks and Markram, 1997; Markram

et al., 1998b; Thomson, 2000) and data fitting (Chow et al., 2005), we used the conjecture that STP behavior of synapses has the same functional dependence on U and τ_{rec} , which allowed us to write a similar rule for the synaptic utilization U . Interestingly, such learning rules depend on the maximum synaptic strength, and they therefore: (i) provided a natural link between STP and STDP and (ii) allowed us to derive an STP-dependent rule for the maximum synaptic strength, to be added to the STDP contribution.

The interaction between short- and long-term plasticity is largely supported by experimental evidence (Markram et al., 1997), although the exact mechanisms are still unknown. Some results (Markram and Tsodyks, 1996; Sjöström et al., 2003, 2007) suggest that synapses become more/less depressing after long-term potentiation/depression. Our rules incorporate this behavior: long-term potentiation/depression always produces larger/smaller changes in STP parameters. However, whether these modifications bring more facilitation or depression critically depends on whether the population firing rate $\langle v \rangle$ is approaching the target rate v_{targ} from above or below. Consider, for example, Equation (16): if $v_{targ} - \langle v \rangle < 0$, then long-term potentiation will produce a stronger depression, thus reproducing the experimentally observed behavior. In our simulations, this happens to the neurons that are firing at low frequencies. On the other hand, if $v_{targ} - \langle v \rangle > 0$, then an increase in A will make τ_{rec} even smaller, resulting in a less depressing synapse. In our simulations, this happens to the neurons that are firing at high frequencies. A similar argument can be formulated for the induction of long-term depression. We note that several mechanisms have been identified to compete during synaptic transmission, resulting in a more complex and less clear relationship between STP and STDP (Sjöström et al., 2007).

In Sjöström et al. (2003, 2007) the authors link the interaction between short- and long-term plasticity with the frequency of firing: at high rates, synapses tend to become stronger and more depressing, while at lower frequencies they tend to become weaker and less depressing. Our derivation, instead, suggests the opposite: if we rely on the hypothesis that large values of τ_{rec} lead to depression and small values to facilitation (Chow et al., 2005), according to Equation (15), facilitating synapses allow neurons to reach higher frequencies. These findings, together with the STDP triplet rule, from the basis of our work: they provide the theoretical basis for the experimentally observed correspondence between facilitation and bidirectionality, and between depression and unidirectionality. The behavior expressed by Equation (15) is experimentally and computationally based on previous work that relates facilitation with high frequency and rate code, and depression with low frequency and temporal code (Fuhrmann et al., 2002; Blackman et al., 2013). This is because, for example, a facilitating synapse may require several spikes to elicit an action potential, meaning that only high frequency stimulation can generate postsynaptic spikes (Matveev and Wang, 2000; Klyachko and Stevens, 2006).

We derived our rules by minimizing an error function that is equal to zero when the target and actual firing rates are equal. Alternatively, we could have defined a reward function opposite to the error function in the sense that for zero error the reward

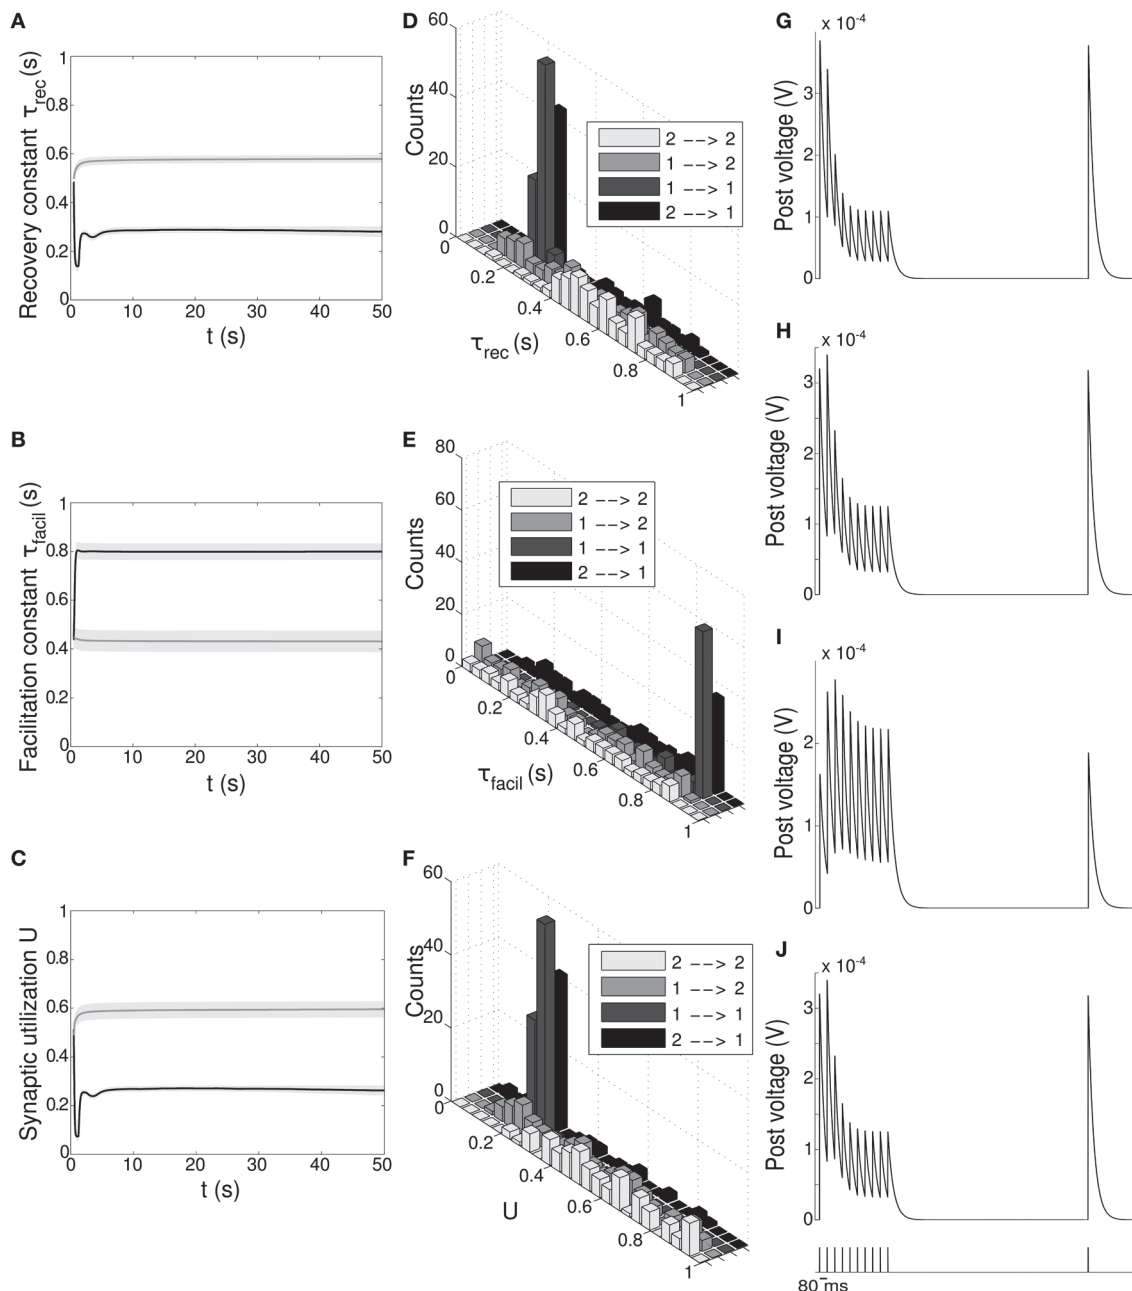


FIGURE 5 | Double population scenario: STP parameters adaptation and final distribution for the output populations with full (U , τ_{rec} , τ_{facil} , A) learning scheme (Part 1). (A–C) Mean values of recovery time constant τ_{rec} , facilitation time constant τ_{facil} and synaptic utilization U . **Black lines** represent mean values across the synapses projecting onto output population 1 from both output populations, $\varphi_1^{out} \cup \varphi_2^{out} \rightarrow \varphi_1^{out}$, whereas **gray lines** describe the synapses projecting onto output population 2 from both output populations, $\varphi_1^{out} \cup \varphi_2^{out} \rightarrow \varphi_2^{out}$. **Shaded areas** show standard deviation. We observe that the two populations develop different synaptic types, facilitating for φ_1^{out} and depressing for φ_2^{out} . **(D–F)** Corresponding histograms of the three synaptic

parameters at the end of the simulation. For each of them we show four different groups of values, mapping qualitatively to the four subtypes identified by Wang et al. (2006), see **Table 3**. **Light gray**: $\varphi_2^{out} \rightarrow \varphi_2^{out}$ (E2a). **Medium gray**: $\varphi_1^{out} \rightarrow \varphi_2^{out}$ (E2b). **Dark gray**: $\varphi_1^{out} \rightarrow \varphi_1^{out}$ (E1a). **Black**: $\varphi_2^{out} \rightarrow \varphi_1^{out}$ (E1b). **(G–J)** Single synapse traces obtained with the TM model by using a 12 Hz stimulus. Each panel represents a different subtype of synapses. **(G)** $\varphi_2^{out} \rightarrow \varphi_2^{out}$. **(H)** $\varphi_1^{out} \rightarrow \varphi_2^{out}$. **(I)** $\varphi_1^{out} \rightarrow \varphi_1^{out}$. **(J)** $\varphi_2^{out} \rightarrow \varphi_1^{out}$. Synaptic parameters used are the mean values obtained from the distributions drawn in **(D–F)**. A comparison with (Wang et al., 2006) on the basis of the traces only shows that we are able to identify three of the four subtypes.

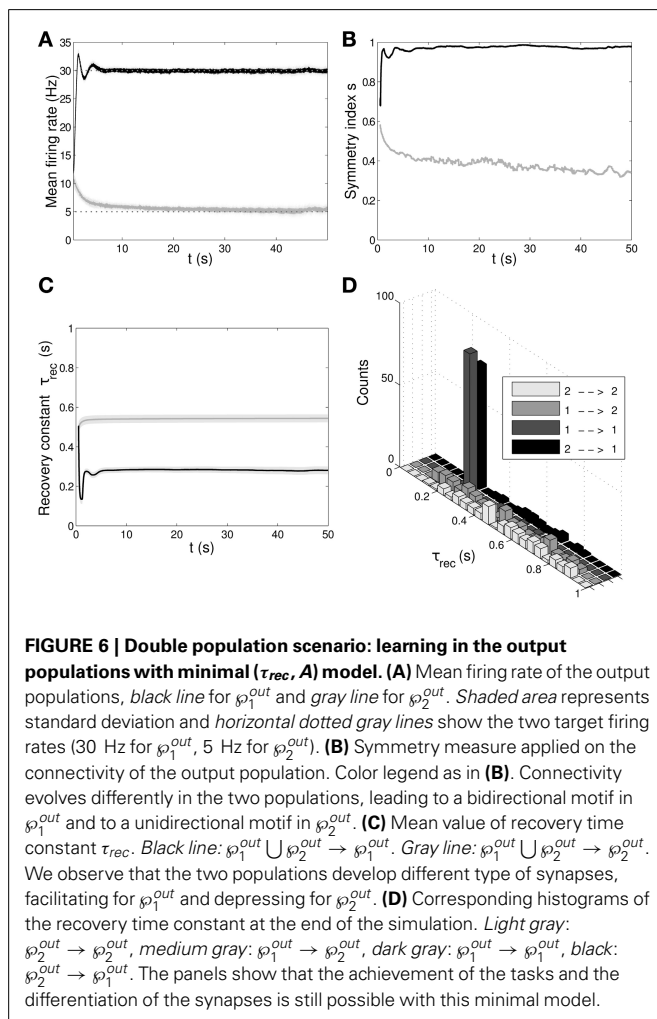
function has its maximum value, and it is equal to zero for large error. We could have then taken the gradient of the reward function instead, bringing the derived rules into the framework of policy gradient learning methods and reinterpreting the feedback

signal as a reward signal (Urbanczik and Senn, 2009; Vasilaki et al., 2009a; Richmond et al., 2011). In biological systems, dopamine is thought to act as reward signal (Schultz et al., 1997; Fiorillo et al., 2003), and its role in the context of learning

Table 3 | Types and subtypes of excitatory synapses between the two output populations in the full model (τ_{rec} , U , τ_{facil} , A).

Synaptic groups	$\tau_{rec}(ms)$	$\tau_{facil}(ms)$	U	τ_{rec}/τ_{facil}	τ_{rec}/τ_{facil} as in Wang	Wang's subtypes
$\wp_1^{out} \cup \wp_2^{out} \rightarrow \wp_1^{out}$	310 ± 11	733 ± 17	0.27 ± 0.01	0.42	0.38	E1
$\wp_1^{out} \rightarrow \wp_1^{out}$	260 ± 5	833 ± 13	0.25 ± 0.01	0.31	0.34	E1a
$\wp_2^{out} \rightarrow \wp_1^{out}$	356 ± 19	643 ± 27	0.29 ± 0.01	0.55	0.43	E1b
$\wp_2^{out} \cup \wp_1^{out} \rightarrow \wp_2^{out}$	550 ± 14	440 ± 19	0.55 ± 0.02	1.25	39.47	E2
$\wp_2^{out} \rightarrow \wp_2^{out}$	595 ± 16	436 ± 26	0.61 ± 0.02	1.36	76.88	E2a
$\wp_1^{out} \rightarrow \wp_2^{out}$	510 ± 23	443 ± 28	0.50 ± 0.03	1.15	25.55	E2b

Column 1: synaptic groups. For instance $\wp_1^{out} \cup \wp_2^{out} \rightarrow \wp_1^{out}$ includes all synapses from both output populations, \wp_1^{out} and \wp_2^{out} , to the output population firing high, \wp_1^{out} . Columns 2,3,4: mean values of STP parameters τ_{rec} , τ_{facil} , U . As in Wang et al. (2006), we provide the results in the form mean \pm s.m.e.. Column 5: ratio between the two time constants, τ_{rec}/τ_{facil} , in our simulation. Column 6: for a direct comparison, we provide the values of τ_{rec}/τ_{facil} as in Wang et al. (2006). Column 7: mapping of our subtypes onto Wang's subtypes.



associated with STDP, and more generally with Hebbian learning, has been extensively studied (Tobler et al., 2005; Izhikevich, 2007; Legenstein et al., 2008).

Each of the learning rules we proposed depends, however, on the difference between the target and the actual firing rates, computed at the population level. This implies the presence of: (i) a

Table 4 | Types and subtypes of excitatory synapses between the two output populations in the minimal model (τ_{rec} , A).

Synaptic groups	$\tau_{rec}(ms)$
$\wp_1^{out} \cup \wp_2^{out} \rightarrow \wp_1^{out}$	300 ± 9
$\wp_1^{out} \rightarrow \wp_1^{out}$	267 ± 6
$\wp_2^{out} \rightarrow \wp_1^{out}$	327 ± 15
$\wp_2^{out} \cup \wp_1^{out} \rightarrow \wp_2^{out}$	524 ± 16
$\wp_1^{out} \rightarrow \wp_2^{out}$	486 ± 23
$\wp_2^{out} \rightarrow \wp_2^{out}$	567 ± 22

Symbols are as in Table 3. Similar to Wang et al. (2006), we provide the results in the form mean \pm s.m.e.

single feedback signal encoding the population activity, which is processed outside the population and broadcasted to all neurons; (ii) an external signal bringing information about the current paradigm, i.e., the target firing rate. Similar to Urbanczik and Senn (2009), we can assume that synapses receive both signals via ambient neurotransmitter concentrations, leading to an on-line plasticity rule.

We initially tested our learning scheme by implementing the rules for τ_{rec} and U on a classical paradigm of inverting associations: keeping the stimulus fixed and varying the associations, the network had to learn to first make choice A and then unlearn it in favor of choice B. This led to a network able to periodically switch its behavior from depressing to facilitating and vice versa, closely following the change in the association paradigm. Throughout the simulation, the network formed motifs similar to those experimentally observed in Wang et al. (2006) and Pignatelli (2009), with facilitating synapses developing bidirectional motifs and depressing synapses developing unidirectional motifs. The desirable motifs were formed due to two factors: (i) the triplet rule that governed long-term potentiation and (ii) the wave-like input stimulus of the network. The form of the plasticity rule guarantees that when neurons fire at high frequency, the synaptic efficacy increases. Hence, synapses will grow up to their bounds, leading to bidirectional connections. On the contrary, when neurons fire at low frequencies, the synaptic efficacy decreases, yet the wave-like input imposes unidirectional connectivity.

We further extended this learning model by adding an STP-motivated rule for the maximum synaptic strength, and we tested it on the same invert association scenario. Results showed the same behavior as before but with faster dynamics due to the joint action of STP and STDP on the absolute efficacy.

In the second part of the paper, we extended our study. First, we considered two populations that have to fire at different frequencies (low, high). Then, we introduced a learning rule for the facilitating time constant, in order to have a full learning model involving all four parameters. The aim was twofold:

(i) *Comparison of our results with experimental data in Wang et al. (2006)*. Although the accuracy is not excellent, we were able to qualitatively reproduce the basic differentiation in the ranges of values of the STP parameters, reflecting the existence of four different synaptic subtypes. We believe that by further adapting the model, in particular learning rates and target frequencies or by considering other rule combinations, it is possible to obtain different parameter values (in principle an infinite combination of them), and thus possibly reproduce the results of Wang and collaborators even better. However, we think this may not be critical because, as a recent study (Costa et al., 2013) has pointed out, fitting techniques generally used for deriving STP parameters from experimental data may give unreliable results. Given this limitation, we think it is important that our model accounts for a large variety of parameter values in principle, and that in this specific case of Wang et al. (2006) it is able to replicate the basic distinction in the synaptic response.

(ii) *Differentiation of synaptic types innervating two functionally different populations*. The reason for this lies in the way we constructed the learning model: what triggers the synaptic modification is the spike of the postsynaptic neuron. The firing rate of the population to which this postsynaptic neuron belongs is the information used to tune the values of STP parameters. In other words, we implement a target-specific learning mechanism. This choice is based on an optimization argument: the more direct and efficient way for a neuron to influence its own activity through synaptic changes is to modify incoming synapses rather than outgoing synapses. A second scheme, a source-specific learning mechanism modifying the outgoing synapses, would have probably led to the same results within closed microcircuits, but on a much longer time scale.

Our target-specific learning mechanism is also supported by experimental evidence (see Blackman et al., 2013 for a review). Despite the fact that STP seems to be mainly a presynaptic mechanism, it has been shown that the target cells can also determine the STP dynamics. All the studies we are aware of have established such a target specificity only in the context of excitatory cells innervating other excitatory cells on one hand and inhibitory cells on the other, specially interneurons (Markram et al., 1998b; Reyes et al., 1998; Buchanan et al., 2012). It would therefore be interesting to appropriately modify the double population scenario by incorporating a population of inhibitory neurons and comparing the results with existing data. In addition, some authors (Blackman et al., 2013; Costa et al., 2013) suggested that a similar differentiation might exist within excitatory only populations. Having target-specific STP for excitatory-excitatory connections is still an open possibility that needs to be further explored. Here

we show from a theoretical point of view that such a differentiation is possible between fundamentally similar (all excitatory) but functionally different (encoding for different paradigms) targets.

The well-established existence of STP-target specificity provides us with a possible biological explanation for the learning rules we derived. Indeed, this scenario requires that the postsynaptic neuron can regulate specifically its own presynaptic compartment only, by a retrograde signal that does not affect neighboring cells. Thus, diffusive retrograde messengers, for example endocannabinoids and nitric oxide, do not appear to be the most suited agents, whereas synaptic adhesion molecules, for example cadherins (Bozdagi et al., 2004) and neuroligins (Dean and Dresbach, 2006), seem to be better candidates for playing this role. These molecules are responsible for governing the presynaptic transmitter release through many different presynaptic mechanisms (Zucker and Regehr, 2002; Blatow et al., 2003; Deng et al., 2011; Blackman et al., 2013).

We underline that the way we obtained the learning rules is based in part on heuristic evaluation. According to Equation (15), derived from a semi-rigorous argument, the key parameters seems to be τ_{rec} and A . By also including U following Carvalho and Buonomano (2011), we obtain a learning scheme involving τ_{rec} , U and A only, which we used to study the double population problem and evaluate the importance of τ_{facil} . Results remain essentially unchanged from the full model, suggesting that τ_{facil} does not play a critical role in the task we defined. This is not surprising and the reason is that our rules link facilitation with a high firing rate, and depression with a low firing rate. Indeed, even with a small facilitation time constant (small τ_{facil}), synapses are still able to fire at a high rate, as long as the stimulating frequency is high enough and recovery from depression is fast enough (low τ_{rec}). Therefore, the time constant of recovery from depression seems to be the only parameter regulating the firing frequency of the neuron for high firing rates, exactly as it comes out from the objective function (we recall that Equation 15 comes from an inequality obtained in the limit of high frequency). With our novel view of allowing synapses to modify their properties from facilitating to depressing and vice versa, we therefore suggest that τ_{rec} is the parameter that is mostly related to rate coding, whereas U to temporal coding.

This conclusion is also supported by Carvalho and Buonomano (2011). In this paper the authors described a simple problem based on temporal synchrony between two inputs that cannot be solved unless STP is learnt, together with STDP. Besides the long-lasting change in A , they introduce a temporal synaptic plasticity for U only and they showed that this indeed solves the problem. Also, they reported that changing U only was the most efficient way to solve the problem. Our work supports the hypothesis that, when dealing with rate coding tasks, the only necessary parameter that has to be learnt is τ_{rec} , whereas, based on Carvalho and Buonomano (2011), when dealing with temporal coding tasks, the only necessary parameter is U .

Another result pointing to a similar direction can be found in Natschläger et al. (2001), where the authors use optimization techniques, rather than explicit learning rules, to train a network of neurons in order to transform a time-varying input into a desired time-varying output. They show that to achieve good

performance, one needs to change at least two parameters, either A and τ_{rec} , or A and U . This confirms that learning must involve at least one presynaptic and one postsynaptic parameter, and that τ_{facil} seems not to be relevant in these types of tasks.

We finally presented results from what we call the minimal model, where only τ_{rec} and A were allowed to change, since both their corresponding update rules come directly from the gradient of the objective function we defined. Results confirmed our belief, as we were still able to learn the tasks while obtaining results similar to those from Wang and collaborators. It is in agreement with our conjecture that when we tried to apply learning on U and A only (results now shown here), the network failed to perform its task because the population that was supposed to fire high stabilized at a much lower frequency, i.e., ~ 15 Hz. Therefore, an alternative minimal model adapting U and A would be able to successfully learn only targets of a lower firing regime. We believe that specialization of parameters in the STP model depending on tasks and signal encoding may be a key ingredient toward a better understanding of synaptic and neuron functionality.

AUTHOR CONTRIBUTIONS

All authors provided materials and contributed to writing the article.

ACKNOWLEDGMENTS

This work was supported by the Royal Society (travel grant JP091330-2009/R4, <http://royalsociety.org>) and the European Commission (FP7 Marie Curie Initial Training Network “NAMASEN,” grant n. 264872, <http://cordis.europa.eu/fp7>). Michele Giugliano acknowledges additional support from the European Commission (FP7 Information and Communication Technologies, Future Emerging Technology programme, “BRAINLEAP” grant n. 306502, <http://cordis.europa.eu/fp7/ict/fet-open>) and the Flemish agency for Innovation by Science and Technology (grant n. 90455/1955, <http://www.iwt.be>); Eleni Vasilaki acknowledges additional support from the Engineering and Physical Sciences Research Council (grant n. EP/J019534/1 and e-futures award EFXD12003/EFXD12004, <http://efutures.ac.uk/xd>). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

REFERENCES

- Asaad, W. F., Rainer, G., and Miller, E. K. (1998). Neural activity in the primate prefrontal cortex during associative learning. *Neuron* 21, 1399–1407. doi: 10.1016/S0896-6273(00)80658-3
- Barak, O., and Tsodyks, M. (2007). Persistent activity in neural networks with dynamic synapses. *PLoS Comput. Biol.* 3:e35. doi: 10.1371/journal.pcbi.0030035
- Binzegger, T., and Douglas, R. J. Martin, K. (2004). A quantitative map of the circuit of cat primary visual cortex. *J. Neurosci.* 24, 8441–8453. doi: 10.1523/JNEUROSCI.1400-04.2004
- Blackman, A. V., Abrahamsson, T., Costa, R. P., Lalanne, T., and Sjöström, P. J. (2013). Target-cell-specific short-term plasticity in local circuits. *Front. Synaptic Neurosci.* 5:11. doi: 10.3389/fnsyn.2013.00011
- Blatow, M., Caputi, A., Burnashev, N., Monyer, H., and Rozov, A. (2003). Ca²⁺ buffer saturation underlies paired pulse facilitation in calbindin-d28k-containing terminals. *Neuron* 38, 79–88. doi: 10.1016/S0896-6273(03)00196-X
- Bozdagi, O., Valcin, M., Poskanzer, K., Tanaka, H., and Benson, D. L. (2004). Temporally distinct demands for classic cadherins in synapse formation and maturation. *Mol. Cell. Neurosci.* 27, 509–521. doi: 10.1016/j.mcn.2004.08.008
- Buchanan, K. A., Blackman, A. V., Moreau, A. W., Elgar, D., Costa, R. P., Lalanne, T., et al. (2012). Target-specific expression of presynaptic nmda receptors in neocortical microcircuits. *Neuron* 75, 451–466. doi: 10.1016/j.neuron.2012.06.017
- Buonomano, D. V., and Merzenich, M. (1998). Cortical plasticity: from synapses to maps. *Annu. Rev. Neurosci.* 21, 149–186. doi: 10.1146/annurev.neuro.21.1.149
- Carvalho, T., and Buonomano, D. V. (2011). A novel learning rule for long-term plasticity of short-term synaptic plasticity enhances temporal processing. *Front. Integr. Neurosci.* 5:20. doi: 10.3389/fnint.2011.00020
- Chklovskii, D., Mel, B., and Svoboda, K. (2004). Cortical rewiring and information storage. *Nature* 431, 782–788. doi: 10.1038/nature03012
- Chklovskii, D., Vitaladevuni, S., and Scheffer, L. (2010). Semi-automated reconstruction of neural circuits using electron microscopy. *Curr. Opin. Neurobiol.* 20, 667–675. doi: 10.1016/j.conb.2010.08.002
- Chow, C., Gutkin, B., Hansel, D., Meunier, C., Dalibard, J. (eds.). (2005). *Methods and Models in Neurophysics*. Chapter Synaptic Dynamics. Les Houchers: Elsevier. 245–266.
- Clopath, C., Buesing, L., Vasilaki, E., and Gerstner, W. (2010). Connectivity reflects coding: a model of voltage-based stdp with homeostasis. *Nat. Neurosci.* 13, 344–352. doi: 10.1038/nn.2479
- Costa, R. P., Sjöström, P. J., and van Rossum, M. C. W. (2013). Probabilistic inference of short-term synaptic plasticity in neocortical microcircuits. *Front. Comput. Neurosci.* 7:75. doi: 10.3389/fncom.2013.00075
- Dean, C., and Dresbach, T. (2006). Neuroligins and neuexins: linking cell adhesion, synapse formation and cognitive function. *Trends Neurosci.* 29, 21–29. doi: 10.1016/j.tins.2005.11.003
- Deng, L., Kaeser, P. S., Xu, W., and Sdhof, T. C. (2011). Rim proteins activate vesicle priming by reversing autoinhibitory homodimerization of munc13. *Neuron* 69, 317–331. doi: 10.1016/j.neuron.2011.01.005
- Denk, W., and Horstmann, H. (2004). Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *PLoS Biol.* 2:e329. doi: 10.1371/journal.pbio.0020329
- Douglas, R., and Martin, K. (2007a). Mapping the matrix: the ways of neocortex. *Neuron* 56, 226–238. doi: 10.1016/j.neuron.2007.10.017
- Douglas, R., and Martin, K. (2007b). Recurrent Neuronal circuits in the neocortex. *Curr. Biol.* 17, R496–R500. doi: 10.1016/j.cub.2007.04.024
- Esposito, U., Giugliano, M., van Rossum, M., and Vasilaki, E. (2014). Measuring symmetry, asymmetry and randomness in neural network connectivity. *PLoS ONE* 9:e100805. doi: 10.1371/journal.pone.0100805
- Fiorillo, C. D., Tobler, P. N., and Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine Neurons. *Science* 299, 1898–1902. doi: 10.1126/science.1077349
- Friston, K. (2011). Functional and effective connectivity: a review. *Brain Connect.* 1, 13–36. doi: 10.1089/brain.2011.0008
- Fuhrmann, G., Segev, I., Markram, H., and Tsodyks, M. (2002). Coding of temporal information by activity-dependent synapses. *J. Neurophysiol.* 87, 140–148. doi: 10.1152/jn.00258.2001
- Fusi, S., Asaad, W. F., Miller, E. K., and Wang, X.-J. (2007). A neural circuit model of flexible sensorimotor mapping: learning and forgetting on multiple timescales. *Neuron* 54, 319–333. doi: 10.1016/j.neuron.2007.03.017
- Gütig, R., and Sompolinsky, H. (2006). The tempotron: a neuron that learns spike timing-based decisions. *Nat. Neurosci.* 9, 420–428. doi: 10.1038/nn1643
- Grillner, S., Markram, H., De Schutter, E., Silberberg, G., and LeBeau, F. (2005). Microcircuits in action - from cpgs to neocortex. *Trends Neurosci.* 28, 525–533. doi: 10.1016/j.tins.2005.08.003
- Hai, A., Shappir, J., and Spira, M. (2010). In-cell recordings by extracellular microelectrodes. *Nat. Methods* 7, 200–202. doi: 10.1038/nmeth.1420
- Hennig, M. H. (2013). Theoretical models of synaptic short term plasticity. *Front. Comput. Neurosci.* 7:154. doi: 10.3389/fncom.2013.00154
- Hertz, J., Krough, A., and Palmer, R. (1991). *Introduction to the Theory of Neural Computation*. Redwood City, CA: Addison Wesley.
- Izhikevich, E. M. (2007). Solving the distal reward problem through linkage of stdp and dopamine signaling. *Cereb. Cortex* 17, 2443–2452. doi: 10.1093/cercor/bhl152
- Kandell, E., Schwartz, J., and Jessel, A. (2008). *Principles of Neural Science*. New York, NY: McGraw-Hill.

- Klyachko, V. A., and Stevens, C. F. (2006). Excitatory and feed-forward inhibitory hippocampal synapses work synergistically as an adaptive filter of natural spike trains. *PLoS Biol.* 4:e207. doi: 10.1371/journal.pbio.0040207
- Le Be, V., and Markram, H. (2006). Spontaneous and evoked synaptic rewiring in the neonatal neocortex. *Proc. Natl. Acad. Sci. U.S.A.* 103, 13214–13219. doi: 10.1073/pnas.0604691103
- Legenstein, R., Pecevski, D., and Maass, W. (2008). A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback. *PLoS Comput. Biol.* 4:e1000180. doi: 10.1371/journal.pcbi.1000180
- Lichtman, J., Livet, J., and Sanes, J. (2008). A technicolour approach to the connectome. *Nat. Rev. Neurosci.* 9, 417–422. doi: 10.1038/nrn2391
- Maass, W., and Markram, H. (2002). Synapses as dynamic memory buffers. *Neural Netw.* 15, 155–161. doi: 10.1016/S0893-6080(01)00144-7
- Markram, H., and Tsodyks, M. (1996). Redistribution of synaptic efficacy between neocortical pyramidal neurones. *Nature* 382, 807–809. doi: 10.1038/382807a0
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic ap and epsps. *Science* 275, 213–215. doi: 10.1126/science.275.5297.213
- Markram, H., Pikus, D., Gupta, A., and Tsodyks, M. (1998a). Potential for multiple mechanisms, phenomena and algorithms for synaptic plasticity at single synapses. *Neuropharmacology* 37:489–500. doi: 10.1016/S0028-3908(98)00049-5
- Markram, H., Wang, Y., and Tsodyks, M. (1998b). Differential signaling via the same axon of neocortical pyramidal neurons. *Proc. Nat. Acad. Sci. U.S.A.* 95, 5323–5328. doi: 10.1073/pnas.95.9.5323
- Markram, H., Gerstner, W., and Sjöström, P. J. (2011). A history of spike-timing-dependent plasticity. *Front. Synaptic Neurosci.* 3:4. doi: 10.3389/fnsyn.2011.00004
- Matveev, V., and Wang, X. J. (2000). Differential short-term synaptic plasticity and transmission of complex spike trains: to depress or to facilitate? *Cereb. Cortex* 10, 1143–1153. doi: 10.1093/cercor/10.11.1143
- Minderer, M., Liu, W., Sumanovski, L., Kügler, S., Helmchen, F., and Margolis, D. (2012). Chronic imaging of cortical sensory map dynamics using a genetically encoded calcium indicator. *J. Physiol.* 590, 99–107. doi: 10.1113/jphysiol.2011.219014
- Natschläger, T., Maass, W., and Zador, A. (2001). Efficient temporal processing with biologically realistic dynamic synapses. *Network* 12, 75–87. doi: 10.1080/net.12.1.75.87
- Perin, R., Berger, T., and Markram, H. (2011). A synaptic organizing principle for cortical neuronal groups. *Proc. Natl. Acad. Sci. U.S.A.* 108, 5419–5424. doi: 10.1073/pnas.1016051108
- Pfister, J.-P., and Gerstner, W. (2006). Triplets of spikes in a model of spike timing-dependent plasticity. *J. Neurosci.* 26, 9673–9682. doi: 10.1523/JNEUROSCI.1425-06.2006
- Pfister, J.-P., Dayan, P., and Lengyel, M. (2010). Synapses with short-term plasticity are optimal estimators of presynaptic membrane potentials. *Nat. Neurosci.* 13, 1271–1275. doi: 10.1038/nn.2640
- Pignatelli, M. (2009). *Structure and Function of the Olfactory Bulb Microcircuit*. Ph.D. thesis, École Polytechnique Fédérale de Lausanne. Available online at: <http://library.epfl.ch/en/theses/?nr=4275>.
- Reyes, A., Lujan, R., Rozov, A., Burnashev, N., Somogyi, P., and Sakmann, B. (1998). Target-cell-specific facilitation and depression in neocortical circuits. *Nat. Neurosci.* 1, 279–285. doi: 10.1038/1092
- Richmond, P., Buesing, L., Giugliano, M., and Vasilaki, E. (2011). Democratic population decisions result in robust policy-gradient learning: a parametric study with GPU simulations. *PLoS ONE* 6:e18539. doi: 10.1371/journal.pone.0018539
- Rinaldi, T., Perrodin, C., and Markram, H. (2008). Hyper-connectivity and hyper-plasticity in the medial prefrontal cortex in the valproic acid animal model of autism. *Front. Neural Circ.* 2:4. doi: 10.3389/neuro.04.004.2008
- Romani, A., Marchetti, C., Bianchi, D., Leinekugel, X., Poirazi, P., Migliore, M., et al. (2013). Computational modeling of the effects of amyloid-beta on release probability at hippocampal synapses. *Front. Comput. Neurosci.* 7:1. doi: 10.3389/fncom.2013.00001
- Rotman, Z., and Klyachko, V. A. (2013). Role of synaptic dynamics and heterogeneity in Neuronal learning of temporal code. *J. Neurophysiol.* 110, 2275–2286. doi: 10.1152/jn.00454.2013
- Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science* 275, 1593–1599. doi: 10.1126/science.275.5306.1593
- Seung, H. (2009). Reading the book of memory: sparse sampling reading the book of memory: sparse sampling versus dense mapping of connectomes. *Neuron* 62, 17–29. doi: 10.1016/j.neuron.2009.03.020
- Silberberg, G., and Markram, H. (2007). Disynaptic inhibition between neocortical pyramidal cells mediated by martinotti cells. *Neuron* 53, 735–746. doi: 10.1016/j.neuron.2007.02.012
- Silberberg, G., Grillner, S., LeBeau, F., Maex, R., and Markram, H. (2005). Synaptic pathways in neural microcircuits. *Trends Neurosci.* 28, 541–551. doi: 10.1016/j.tins.2005.08.004
- Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32, 1149–1164. doi: 10.1016/S0896-6273(01)00542-6
- Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2003). Neocortical ltd via coincident activation of presynaptic nmda and cannabinoid receptors. *Neuron* 39, 641–654. doi: 10.1016/S0896-6273(03)00476-8
- Sjöström, P. J., Turrigiano, G. G., and Nelson, S. B. (2007). Multiple forms of long-term plasticity at unitary neocortical layer 5 synapses. *Neuropharmacology* 52, 176–184. doi: 10.1016/j.neuropharm.2006.07.021
- Song, S., Miller, K., and Abbott, L. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.* 3, 919–926. doi: 10.1038/78829
- Song, S., Sjöström, P., Reigl, M., Nelson, S., and Chklovskii, D. (2005). Highly non-random features of synaptic connectivity in local cortical circuits. *PLoS Biol.* 3:e68. doi: 10.1371/journal.pbio.0030068
- Testa-Silva, G., Loebel, A., Giugliano, M., de Kock, C. P. J., Mansvelder, H. D., and Meredith, R. M. (2012). Hyperconnectivity and slow synapses during early development of medial prefrontal cortex in a mouse model for mental retardation and autism. *Cereb. Cortex* 22, 1333–1342. doi: 10.1093/cercor/bhr224
- Thomson, A. M. (2000). Facilitation, augmentation and potentiation at central synapses. *Trends Neurosci.* 23, 305–312. doi: 10.1016/S0166-2236(00)01580-0
- Tobler, P. N., Fiorillo, C. D., and Schultz, W. (2005). Adaptive coding of reward value by dopamine Neurons. *Science* 307, 1642–1645. doi: 10.1126/science.1105370
- Tsodyks, M., and Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci. U.S.A.* 94, 719–723. doi: 10.1073/pnas.94.2.719
- Urbanczik, R., and Senn, W. (2009). Reinforcement learning in populations of spiking Neurons. *Nat. Neurosci.* 12, 250–252. doi: 10.1038/nn.2264
- Varela, J., Sen, K., Gibson, J., Fost, J., Abbott, L., and Nelson, S. (1997). A quantitative description of short-term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex. *J. Neurosci.* 17, 7926–7940.
- Vasilaki, E., and Giugliano, M. (2012). “Emergence of connectivity patterns from long-term and short-term plasticities,” in *Artificial Neural Networks and Machine Learning ICANN 2012* Vol. 7552 of *Lecture Notes in Computer Science*, eds A. Villa, W. Duch, P. Rdi, F. Masulli, and G. Palm (Berlin; Heidelberg: Springer), 193–200.
- Vasilaki, E., and Giugliano, M. (2014). Emergence of connectivity motifs in networks of model neurons with short- and long-term plastic synapses. *PLoS ONE* 9:e84626. doi: 10.1371/journal.pone.0084626
- Vasilaki, E., Frémaux, N., Urbanczik, R., Senn, W., and Gerstner, W. (2009a). Spike-based reinforcement learning in continuous state and action space: when policy gradient methods fail. *PLoS Comput. Biol.* 5:e1000586. doi: 10.1371/journal.pcbi.1000586
- Vasilaki, E., Fusi, S., Wang, X.-J., and Senn, W. (2009b). Learning flexible sensorimotor mappings in a complex network. *Biol. Cybern.* 100, 147–158. doi: 10.1007/s00422-008-0288-z
- Wang, Y., Markram, H., Goodman, P., Berger, T., Ma, J., and Goldman-Rakic, P. (2006). Heterogeneity in the pyramidal network of the medial prefrontal cortex. *Nat. Neurosci.* 9, 534–542. doi: 10.1038/nn1670
- Weden, V., Rosene, D., Wang, R., Dai, G., Mortazavi, F., Hagmann, P., et al. (2012). The geometric structure of the brain fiber pathways. *Science* 335, 1628–1634. doi: 10.1126/science.1215280
- Wickersham, I., Lyon, D. C., Barnard, E., Mori, T., Finke, S., Conzelmann, K., et al. (2007). Monosynaptic restriction of transsynaptic tracing from single, genetically targeted neurons. *Neuron* 53, 639–647. doi: 10.1016/j.neuron.2007.01.033

- Zhang, F., Aravanis, A., Adamantidis, A., de Lecea, L., and Deisseroth, K. (2007). Circuit-breakers: optical technologies for probing neural signals and systems. *Nat. Rev. Neurosci.* 8, 577–581. doi: 10.1038/nrn2192
- Zucker, R., and Regehr, W. (2002). Short-term synaptic plasticity. *Annu. Rev. Physiol.* 64, 355–405. doi: 10.1146/annurev.physiol.64.092501.114547

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 15 May 2014; accepted: 31 December 2014; published online: 29 January 2015.

Citation: Esposito U, Giugliano M and Vasilaki E (2015) Adaptation of short-term plasticity parameters via error-driven learning may explain the correlation between activity-dependent synaptic properties, connectivity motifs and target specificity. Front. Comput. Neurosci. 8:175. doi: 10.3389/fncom.2014.00175

This article was submitted to the journal *Frontiers in Computational Neuroscience*. Copyright © 2015 Esposito, Giugliano and Vasilaki. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

RM-SORN: a reward-modulated self-organizing recurrent neural network

Witali Aswolinskiy^{*†} and Gordon Pipa

Institute of Cognitive Science, University of Osnabrück, Osnabrück, Germany

OPEN ACCESS

Edited by:

Cristina Savin,
Institute of Science and Technology
Austria, Austria

Reviewed by:

Jochen Triesch,
Johann Wolfgang Goethe University,
Germany
Paul Miller,
Brandeis University, USA
Florentin Wörgötter,
University Goettingen, Germany

*Correspondence:

Witali Aswolinskiy,
Institute of Cognitive Science,
University of Osnabrück,
Neuroinformatics Research Group,
Albrechtstr. 28, 49069 Osnabrück,
Germany
waswolinskiy@uos.de

†Present Address:

Witali Aswolinskiy,
CoR-Lab N-128, Bielefeld University,
Universitätsstr. 25, 33615 Bielefeld,
Germany

Received: 17 November 2014

Accepted: 04 March 2015

Published: 24 March 2015

Citation:

Aswolinskiy W and Pipa G (2015)
RM-SORN: a reward-modulated
self-organizing recurrent neural
network.
Front. Comput. Neurosci. 9:36.
doi: 10.3389/fncom.2015.00036

Neural plasticity plays an important role in learning and memory. Reward-modulation of plasticity offers an explanation for the ability of the brain to adapt its neural activity to achieve a rewarded goal. Here, we define a neural network model that learns through the interaction of Intrinsic Plasticity (IP) and reward-modulated Spike-Timing-Dependent Plasticity (STDP). IP enables the network to explore possible output sequences and STDP, modulated by reward, reinforces the creation of the rewarded output sequences. The model is tested on tasks for prediction, recall, non-linear computation, pattern recognition, and sequence generation. It achieves performance comparable to networks trained with supervised learning, while using simple, biologically motivated plasticity rules, and rewarding strategies. The results confirm the importance of investigating the interaction of several plasticity rules in the context of reward-modulated learning and whether reward-modulated self-organization can explain the amazing capabilities of the brain.

Keywords: reward-modulated STDP, intrinsic plasticity, recurrent neural networks, self-organization, plasticity, hebbian learning

Introduction

The brain is a complex, self-organizing system, where a multitude of neural plasticity mechanisms shape learning, and memory. These plasticity mechanisms are, in turn, shaped by neuro-modulators, which are often part of a reward system (Pawlak et al., 2010). *In vivo* experiments showed that rewarding behavior can change synapses and neurons selectively to achieve a rewarded goal (Fetz, 1969; Ahissar et al., 1992; Sigala and Logothetis, 2002). Several models of reward-modulated recurrent neural networks are able to partially replicate these experiments and solve simple tasks (Izhikevich, 2007; Legenstein et al., 2008; Soltoggio and Steil, 2013; Hoerzer et al., 2014). In these models, correct outputs are rewarded through the application of STDP or a hebbian learning rule, and noise is used to explore possible output sequences. Noise as a part of a model, however, makes the model non-deterministic and introduces a random, transient component that can counteract the learning of causal relations by STDP. Here, we propose an alternative to combine deterministic behavior and the ability to explore states for reward modulated learning. For this, either deterministic chaos or other complex deterministic behavior may be used. Here, we study complex behavior that is introduced by Intrinsic Plasticity (IP)—neuronal plasticity associated with homeostasis (Turrigiano et al., 1998; Desai et al., 1999). We introduce a simple binary neural network model, which learns through interaction of IP and reward-modulated STDP. Exploration of the output state space is carried out through IP and not noise.

The Reward-Modulated Self-Organizing Recurrent Network (RM-SORN), model is based on SORN: Self-Organizing Recurrent Network (Lazar et al., 2009). SORN consists of a recurrent layer with binary thresholded neurons and a readout layer. In the recurrent layer three plasticity mechanisms are applied: IP, Synaptic Normalization (SN) and Spike-Timing-Dependent Plasticity (STDP). The readout layer is trained with linear regression. The network is trained in two phases: in the first phase the recurrent layer processes the input with ongoing plasticity. In the second phase plasticity is disabled, the recurrent layer processes the input again, and the neuron activations serve to train the readout layer. Lazar showed, that all three plasticity mechanisms are necessary to create an effective representation of the input in the recurrent layer and that these representations allow SORN to outperform randomly initialized non-plastic networks. STDP forms the internal representations, IP activates silent neurons and dampens neurons with too high activity, and SN decorrelates neurons preventing seizure-like activity.

SORN was successfully applied to tasks for prediction (Lazar et al., 2009), recall and non-linear computation (Toutounji and Pipa, 2014) and artificial grammar learning (Duarte et al., 2014). The main advantage of SORN is its simplicity and the biological plausibility of the plastic recurrent layer. The biological plausibility is further underlined by the findings of Zheng et al. (2013), who added two plasticity mechanisms to the recurrent layer: structural plasticity and inhibitive STDP. The authors observed a log-normal weight distribution of the synaptic weights in the recurrent layer matching experimental findings. Additionally, the patterns of fluctuation of the weights were consistent with the dynamics of dendritic spines found in rat hippocampus.

SORN offers the possibility to study plasticity mechanisms similar to those in the brain in simple, manageable networks. The

biologically not plausible part of SORN is the linear regression readout, which is replaced here by a plastic, non-recurrent, reward-modulated neuron layer.

Materials and Methods

Both SORN and RM-SORN consist of a recurrent layer with three plasticity mechanisms and a readout or output layer. However, whereas in SORN the output is trained with linear regression, in RM-SORN, the weights from the recurrent layer to the output layer are plastic and adapted through reward-modulated STDP. The model allows, but doesn't prescribe, the application of reward-modulated STDP to the recurrent layer. In this paper, we test both versions, i.e., with and without reward-modulated STDP in the recurrent network, and explain for which conditions reward-modulated STDP applied to the weights in the recurrent network improves the computational performance.

Network Model

Figure 1 depicts the model structure. Both layers consist of binary threshold neurons. The first layer is recurrent and consists of N^E excitatory and N^I inhibitory neurons. The connectivity between the excitatory neurons is sparse (5–10%) and full between excitatory and inhibitory neurons. Self-connections are not allowed. The ratio of excitatory to inhibitory neurons is 5:1. The second layer is the output or readout layer with neurons that are not interconnected. In tasks where the network has to generate sequences, a feedback connection from the output layer to the recurrent layer is necessary. A random subset of the excitatory units in the recurrent layer receives the input: for each symbol of an input sequence, e.g., “1234,” a different subset of the units




FIGURE 1 | The Reward-Modulated Self-Organizing Recurrent Neural Network (RM-SORN). Excitatory units are depicted as red, inhibitory units as blue. The arrows symbolize directed connections between units.

Excitatory units are sparsely interconnected and excitatory and inhibitory units are fully connected between each other. Only the excitatory units project to the output layer. Input and output are sequences of predefined symbols.

receives an input of 1 and the rest 0. The units are binary with Θ being the heaviside step function, applied independently to every neuron:

$$x_i(t+1) = \Theta \left(\sum_j^{N^E} w_{ij}^{EE} x_j(t) - \sum_k^{N^I} w_{ik}^{EI} y_k(t) + u_i(t) - T_i^E(t) \right) \quad (1)$$

$$y_i(t+1) = \Theta \left(\sum_j^{N^E} w_{ij}^{IE} x_j(t) - T_i^I \right) \quad (2)$$

$$o_i(t+1) = a \left(\sum_j^{N^E} w_{ij}^{OE} x_j(t) - T_i^O(t) \right) \quad (3)$$

Neurons in the first layer are updated according to (1) and (2) while neurons in the output layer are updated according to (3). x and y represent the activity of the excitatory and inhibitory neurons in the recurrent layer and o the activity of the output neurons. a is the activation function for the output neurons. With several output neurons, the activation function is winner-takes-all (WTA), with one output neuron, the heaviside step function Θ is used. w^{AB} is the weight matrix with weights from B-units to A-units, u is the input and T^A is the threshold for the A-units. (Hence, w^{OE} are the weights from the recurrent to the output layer and T^O are the thresholds for the output neurons). Initially, the weights are drawn from a normal distribution, but change through application of three plasticity rules: IP, SN, and reward-modulated Spike-Timing-Dependent Plasticity (rm-STDP).

Plasticity Rules

IP adapts the thresholds so that on average each neuron has the firing rate μ_{IP} :

$$\Delta T_i^E(t) = \eta_{IP}(x_i(t) - \mu_{IP}) \quad (4)$$

The threshold is increased when the unit was active and decreased when the unit was inactive, leading to an asymptotic fix point of the average firing rate μ_{IP} . Thereby, IP activates neurons, which would be otherwise inactive and regulates down neurons which fire too often, enforcing the given average firing rate. During the experiments, in the recurrent layer, μ_{IP} was set to values between 0.05 and 0.25 depending on which value performed best. In the output layer, μ_{IP} was set per neuron to correspond to the expected occurrence probability of the symbol represented by the neuron. η_{IP} is the learning rate for IP.

STDP strengthens the connection from x_j to x_i when x_j was active before x_i (x_j “causes” x_i) and weakens it, when x_j was active after x_i . The main difference to SORN is the modeling of the output layer as another plastic neuron layer and the reward-modulation of STDP with the modulation m :

$$\Delta w_{ij}^{EE} = m_r * \eta_{STDP} (x_j(t-1)x_i(t) - x_j(t)x_i(t-1)) \quad (5)$$

$$\Delta w_{ij}^{OE} = m_o * \eta_{STDP} (x_j(t-1)o_i(t)) \quad (6)$$

η_{STDP} is the learning rate for STDP. m_r and m_o are the modulation factors for the recurrent and the output layer, respectively. During the simulations, m_r was either set to one (no modulation) or to the same value as m_o . m_o is determined according to a rewarding strategy, which is a function of the reward r . Both modulation and reward can be positive, negative or zero. Depending on the task, different modulation strategies can be chosen for m_o .

After application of STDP the incoming weights to a neuron are scaled to sum up to 1:

$$w_{ij}(t) = \frac{w_{ij}(t)}{\sum_j w_{ij}(t)} \rightarrow \sum_j w_{ij}(t) = 1 \quad (7)$$

The relative strength of the synapses remains the same.

Reward-Modulation Strategies

In tasks with known target values (which are all tasks except the generation task), the reward was set to 1 for correct outputs, and either 0 or -1 for wrong outputs, depending on which setting led to the highest performance. Negative reward—punishment—can lead to Anti-STDP. In the generation task, where the network has to generate a sequence of symbols without input, the network is rewarded if it produces a part of the target sequence, starting from the beginning of the sequence. The reward is proportional to the length of the correctly generated sequence part. For example, let the network be rewarded for generation of the sequence “1234.” If the network generates “1234,” it receives the full reward of unit 1 at the time when it generates the last state “4.” Analogous, it receives the reward of $\frac{3}{4}$ for the sequence “x123,” $\frac{2}{4}$ for the sequence “xx12” and only $\frac{1}{4}$ for “xxx1” (here “x” represents any symbol or state). An exception is made for “xx11”: this combination is punished to prevent the generation of the trivial sequence of repetitions of “1.” For any other sequence, no reward is given.

The reward-modulation strategy (M) determines the modulation factor m and therefore, whether STDP is applied, suppressed or inverted. The network can be modulated either directly by (8) or the modulation can be computed from the previous rewards. Particularly interesting is here the hypothesis, that dopamine neurons encode reward prediction errors. (9) defines a simple estimate of the reward prediction error.

$$M0 : m(r, t) = r(t) \quad (8)$$

$$Mk : m(r, k, t) = r(t) - \bar{r}(t, k), \quad k \in (1, 5, 10, 20) \quad (9)$$

\bar{r} is computed as the moving average of the previous k rewards. With window size $k = 1$, the modulation factor is the current reward minus the previous reward. The k -values were defined *ad-hoc* and selected for each task independently with a parameter grid search. The selected strategies for the tasks and other parameters are listed in the supplementary material.

Training and Testing

The network is trained in two phases:

1. In tasks with training data, the network is reward-modulated, while processing 20,000 steps of training data. In the generation task, the network output is fed back as input. Every 100 steps, the network is validated on 500 steps of validation data. After 20,000 steps, 200 networks are evaluated and the network with the best validation performance is chosen for the next phase.
2. Step 1 is repeated while the plasticity in the first layer is shut off, allowing to fine-tune the weights to the output layer.

After this composed training phase, the network is tested on 10,000 steps of test data or, in generation tasks, is used as a generative model to generate the desired output for 10,000 steps.

The evaluation of the network at different plasticity times is essential, since the performance of the networks fluctuates strongly during reward-modulated plasticity.

During the simulations, the network size consisted of at least 100 and at most 400 excitatory neurons. The number of inhibitory neurons was always a fifth of the excitatory neurons, as in the original SORN. From here on, the number of excitatory neurons will be referred to as the size of the network or simply N .

Model Evaluation

The performance of the model was compared to SORN, static (non-plastic) supervised trained networks and random networks. In SORN networks, the 20,000 steps of training data were processed by the plastic recurrent layer. After every 1000 steps, the weights were frozen, the network processed the 20,000 steps of training data and the readout weights were trained on the resulting network states. Thus, 20 intermediary networks were created, with the first network being subjected to plasticity for 1000 steps and the last network for 20,000 steps. The network that performed best on 500 steps of validation data was chosen for the performance evaluation on the test data. Static networks were created by taking the best SORN network, shuffling the weights 20 times and choosing the network that performed best on the validation data. Random networks were trained in the same manner as RM-SORN, but with the reward-modulated STDP-weight updates randomly distributed across all eligible weights at each step. The eligible weights in the recurrent layer are the initially non-zero weights in the sparse connectivity matrix. (Initially non-zero weights can become zero through STDP). From the recurrent layer to the output layer all weights are eligible. The performance difference between the RM-SORN and “random” networks is the difference between learning through reward-modulated STDP and lucky guessing.

Notably, during training of a network, 200 RM-SORN, but only 20 SORN and static-networks were evaluated. This discrepancy is due to the high computational cost of the linear regression in SORN and the fact, that SORN achieves in most tasks almost perfect performance - more frequent evaluation is not necessary. In the pattern recognition task, where RM-SORN was better, the same number of networks (200) was evaluated to even the odds.

All results were averaged over ten data sets and ten networks per data set—the procedure described above was applied to each of the 100 network/data set combinations. In the motion generation task, which has no input, 100 networks were evaluated.

Since the weights in RM-SORN are positive, a similar restriction was imposed onto the supervised training—instead of least squares, non-negative least squares method (Lawson and Hanson, 1995) was used.

Task Descriptions

The model was evaluated on eight different tasks, including those in Lazar et al. (2009) and Toutounji and Pipa (2014).

In the **counting** task (Lazar et al., 2009), the network receives random alternations of two words “abbb...c” and “eddd...f” with $n + 2$ letters per word and n b’s or n d’s in a word. The goal is to predict the next letter. In order to correctly predict the *last* letter in the word, the network has to “count” the b’s and the d’s. Thus, the first layer needs to learn separable representations (linearly separable when using a simple linear readout, as we use here) of the input conditions $a, b_1, b_2, \dots, b_n, e, d_1, d_2, \dots, d_n$. In the output layer, these representations must be mapped to the next letter in the word: $a \rightarrow b, b_1 \rightarrow b, b_2 \rightarrow b \dots b_{n-1} \rightarrow b, b_n \rightarrow c$ and similar for e’s and d’s. Given the random alternation of words, the first letter of a word is unpredictable and therefore excluded from the performance measure. We use two kinds of performance measures. Firstly, the overall performance that measures the match of all letters of the entire sequence, with the exception of the excluded first letter of each word. Secondly, we measure the counting performance that is the accuracy of predicting the last letter in a word. This performance reflects the capability of the network to retain and use information from previous inputs.

As a second task, we use **motion prediction** (motivated by Lazar et al., 2009). In the motion prediction task, the network receives random renewal sequences of the two words “123...n” and “n....321.” These sequences can be interpreted as movement of an object in one dimension from left to right and back, that is sensed by a line array of sensors. Therefore, the task was initially set up to mimic the learning of motion-specific visual receptive fields (Lazar et al., 2009). Comparing the counting task and the motion prediction task highlights their difference in respect to subsequence learning of the individual words. For the motion prediction task, all subsequences (e.g., 12, 23, 34, ...) of the word “123...n” can be learned independently. This is not the case in the counting task, where, for example, the input condition b_3 (“abbb”) cannot be learned before b_2 (“abb”) is learned.

As a third task, we used the **occluder** task (Lazar et al., 2009), that is a combination of the counting and motion task. With $n = 8$, the input consists of random alternations of four words: “12345678,” “87654321,” “19999998,” and “89999991.” As in the motion prediction task, they can be interpreted as the movement of an object sensed by a line of sensors. To model the occlusion of part of the sensors, two additional words “19999998” and “89999991” are used. Here, the symbol “9” represents the lacking information about the object position when the object

is occluded. Note that because the first letter occurs more than once, the second letter in a word is unpredictable, and therefore excluded from the performance measure.

As a fourth task, we measure the **memory capacity** (corresponds to the RAND task in Toutounji and Pipa, 2014). Here, the network receives a random sequence of symbols and has to reproduce the symbol from n steps back. The number of symbols used here is 6.

As a fifth task, we use the **Markov-85** task (Toutounji and Pipa, 2014). For this task, we generate an input sequence that consists of symbols generated by a first order Markov chain. The chain has six states: 1, 2, 3, 4, 5, 6. The transition probabilities for $1 \rightarrow 2$, $2 \rightarrow 3$, $3 \rightarrow 4$, $4 \rightarrow 5$, $5 \rightarrow 6$, and $6 \rightarrow 1$ are $p = 0.85$. All other transition probabilities are $p = 0.03$. The goal of the task is either to recall the inputs from n steps back or to predict the state n steps in the future.

As a sixth task, we use the **parity** task (Toutounji and Pipa, 2014). Here, the network receives a series of binary values and has to compute the parity for the current input and $n - 1$ previous inputs. This task tests the capability of the network for non-linear computation.

As a seventh task, we designed a sequence or **motion generation** task, where the network has to generate either “123... n ” or “ n ...321.” This is the only task, where the network receives no input. The number of symbols, and thus input dimensions, is n , and the longer the sequence, the harder the task. Similar to the motion prediction task, this task can be interpreted as the movement of an object in one dimension. More generally, success in this task shows that the network can generate an arbitrary symbol sequence with the same symbol distribution as in the motion words. The reinforcement of two words instead of one is more difficult, because the same output symbol can be rewarded for two different reasons.

The last task is a **pattern recognition** task, designed to highlight the effect of reward-modulated changes of synaptic weights in the recurrent network. Here, the network receives random alternations of the four words “1234,” “4321,” “4213,” “2431” and has to recognize the word “1234”: the output for every letter in this word has to be 1, and for all others, 0. In this task only one output neuron was used.

Results

The performance measures for the counting, occluder, motion prediction and motion generation tasks are shown in **Figure 2**. The results for the Markov-85, memory capacity, parity, and pattern recognition tasks are shown in **Figure 3**. In most tasks, RM-SORN achieves high accuracy, and is only slightly worse than SORN. This is remarkable, considering that RM-SORN learns in a self-organized manner through interaction of plasticity mechanisms, while SORN learns through a supervised, mathematically derived algorithm.

Reward-modulation of the recurrent layer improved performance only in the pattern recognition task: it allowed RM-SORN to outperform SORN for small network sizes. In the other

tasks, reward-modulation of the recurrent layer didn't improve performance and was not applied during the experiments.

Prediction, Recall and Non-Linear Computation

In the counting task, RM-SORN achieves a high overall performance (**Figure 2A**) in the range 95–100% and has a counting performance (**Figure 2B**) in-between static and SORN networks. Higher n increases the difficulty for the prediction of the last letter, as the network needs to remember more of the past inputs, but reduces the overall difficulty of predicting the other letters, which are either b's or d's. A system that just produces “b” when it sees “a” or “b” and “d,” when it sees “e” or “d” can achieve a high accuracy without being able to count—this happens in the random networks, which achieve, for example, 85% accuracy for $n = 20$. The difference between static and SORN networks is due to improvement of the representational capability of the recurrent layer through unmodulated plasticity.

The performance of RM-SORN in the occluder task (**Figures 2C,D**) is worse, due to the higher number of input conditions, but the task can still be solved with high accuracy for $n \leq 8$ and good overall accuracy for $n > 8$.

In the motion prediction task (**Figure 2E**), high accuracy can be achieved until very high n —the network can learn many different input condition mappings. Random networks achieve a performance slightly above the chance level of $1/n$ for higher n .

Remarkable is the high accuracy of RM-SORN in the motion generation task (**Figure 2F**), where the network never receives any input. These results will be discussed in more detail in the next section.

The performance in the memory capacity task (**Figure 3A**) is similar for static, SORN and RM-SORN networks: high for low n and low for higher n , hitting the chance level of $1/6$ in the end. Since the input is random, learning of effective representations through STDP in the recurrent layer is not possible, and static and SORN networks have similarly low network memory. RM-SORN stays slightly behind static and SORN networks for higher n , which was also observed in the other tasks.

The Markov-85 task (**Figure 3B**) offers a different picture—static, SORN and RM-SORN networks feature high performance, also for high negative n (recall—corresponds to positive n in the memory capacity task). The structure in the input allows an efficient representation in the recurrent layer and also a more efficient mapping of the representations in the output layer. The performance for prediction (positive n) is lower, since the maximal achievable performance declines exponentially with each step, because of the stochastic nature of the input sequences.

The parity task (**Figure 3C**) offers a picture similar to the memory capacity task, since its input is also unstructured. Nevertheless, the high performance of RM-SORN for small n demonstrates the capability of the model for non-linear computation.

In conclusion, RM-SORN achieves high performance in all tasks and is in most tasks in-between SORN and static networks. For complex tasks (high n), the performance deteriorates more than in SORN. A more detailed analysis of the motion generation and the pattern recognition task (**Figure 3D**) follows in the next sections.

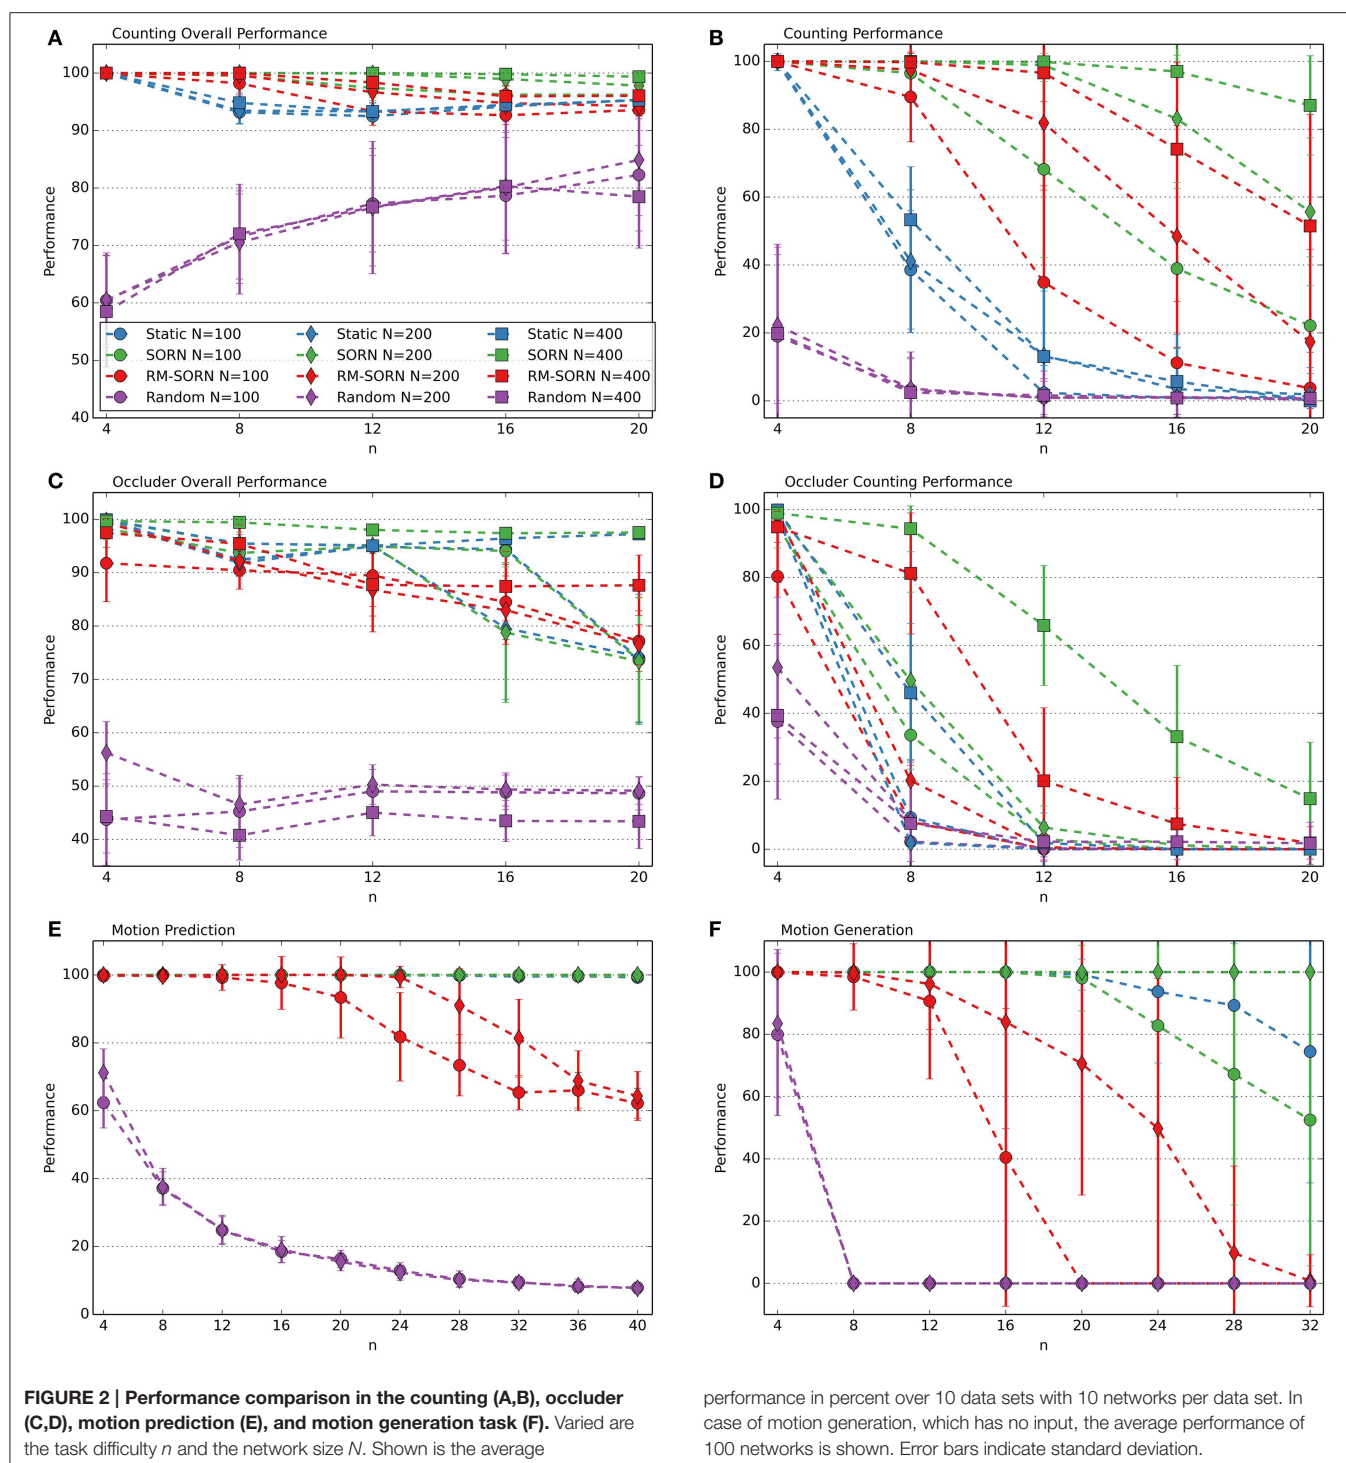


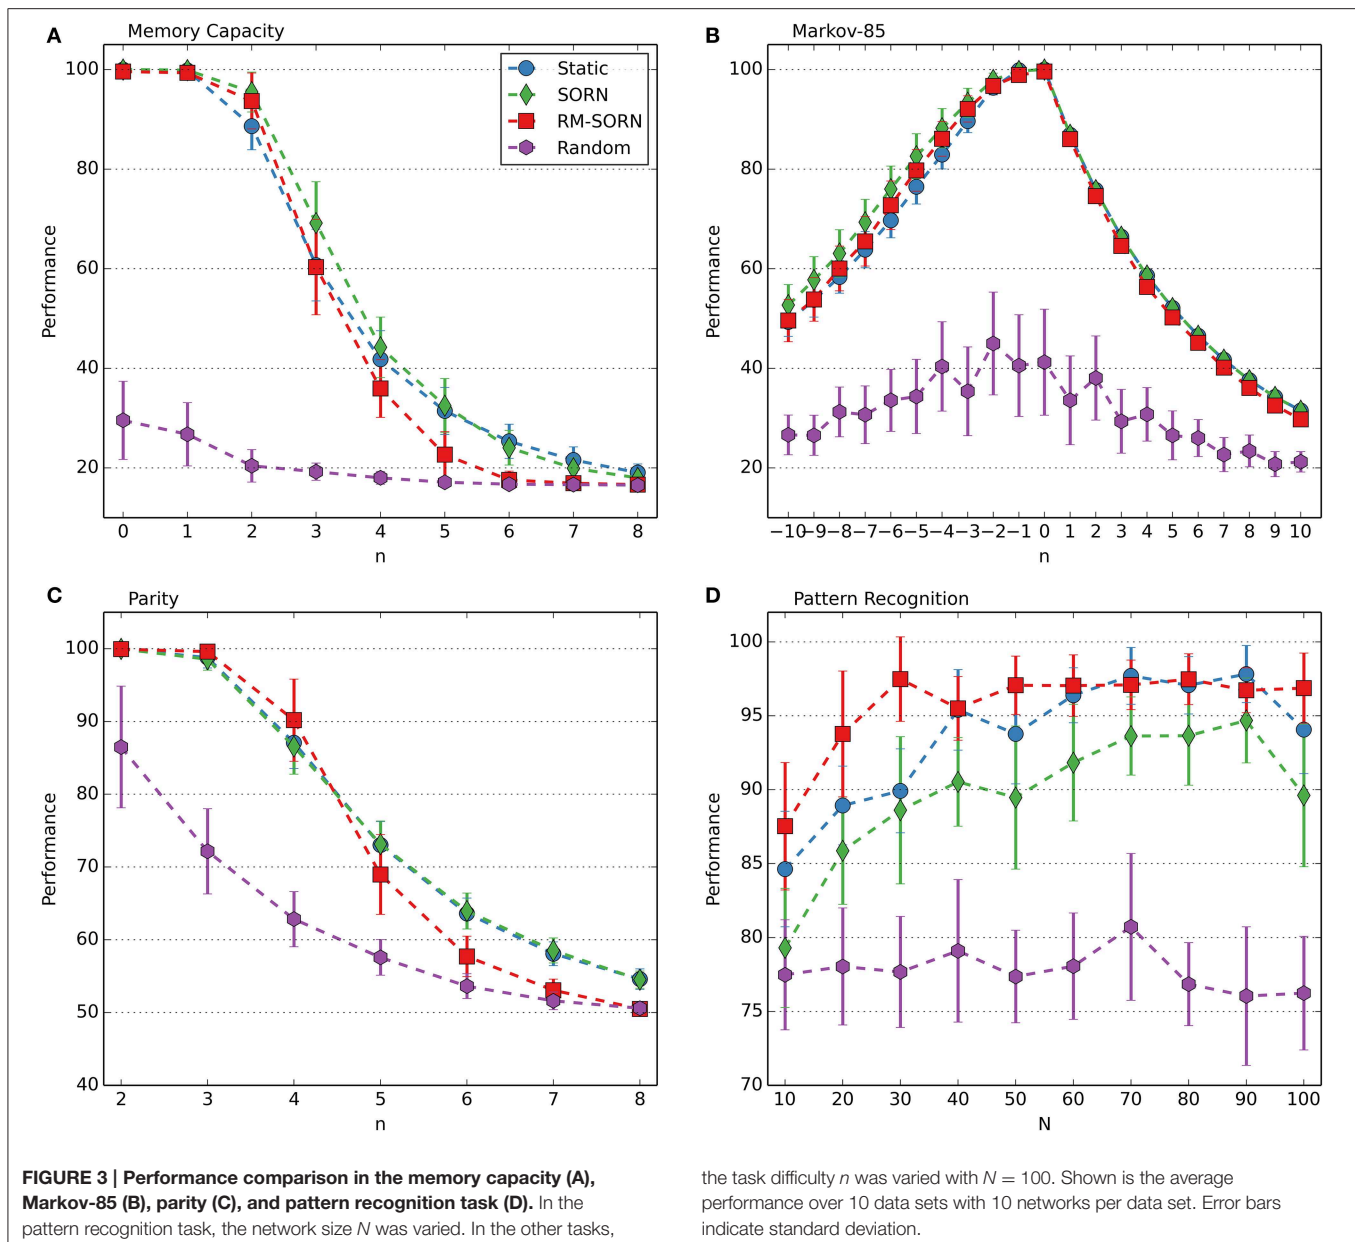
FIGURE 2 | Performance comparison in the counting (A,B), occluder (C,D), motion prediction (E), and motion generation task (F). Varied are the task difficulty n and the network size N . Shown is the average

performance in percent over 10 data sets with 10 networks per data set. In case of motion generation, which has no input, the average performance of 100 networks is shown. Error bars indicate standard deviation.

Pattern Recognition

In this task, RM-SORN achieved a higher performance than SORN, especially for small network sizes. The results are shown in **Figure 3D**. The best average performance of 97.48% was achieved with a network with only 30 excitatory neurons. Static and SORN networks of this size stay below the 90%-mark.

The reason for the better performance of RM-SORN is the reward-modulation of the recurrent layer. In order to recognize the target word, only the representations of parts of the target word are necessary, and all other symbol combinations can be ignored. SORN tries to create representations of all occurring symbol combinations and has therefore less memory for each individual combination. This leads



to a poor performance with small networks. RM-SORN, on the other hand, is only rewarded if it recognizes the target word, and its recurrent layer is plastic only during this time—it learns to represent parts of the target word exclusively.

Figure 4 visualizes the selectivity of neurons in the recurrent layer for static, SORN and RM-SORN networks with 30 neurons. For all two-symbol input sequences that occurred during testing, the probability of a neuron to spike was estimated by counting the occurred spikes. In static and SORN networks, neurons have no preferred input stimuli. The selectivity for partial sequences “43” and “21” in SORN is slightly higher than in static networks, because both occur in two patterns, while the other combinations occur only in one. In RM-SORN, neurons mostly encode

the parts of the target word (“12,” “23,” “34”), which allows for a simple and effective mapping to the output neuron and a high performance.

The pattern recognition task is also the only task, where static networks perform better than SORN networks. This is due to the training procedure, as explained in section Model Evaluation: in the pattern recognition task, static networks were created by shuffling the SORN weights and taking the best out of 200 networks (in other tasks, out of 20 because of the computational load). While all intermediary evaluated SORN networks try to map all input conditions equally, some static networks, by chance, better represent the target word parts. Thus, the best static network can achieve a better performance than the best SORN network.

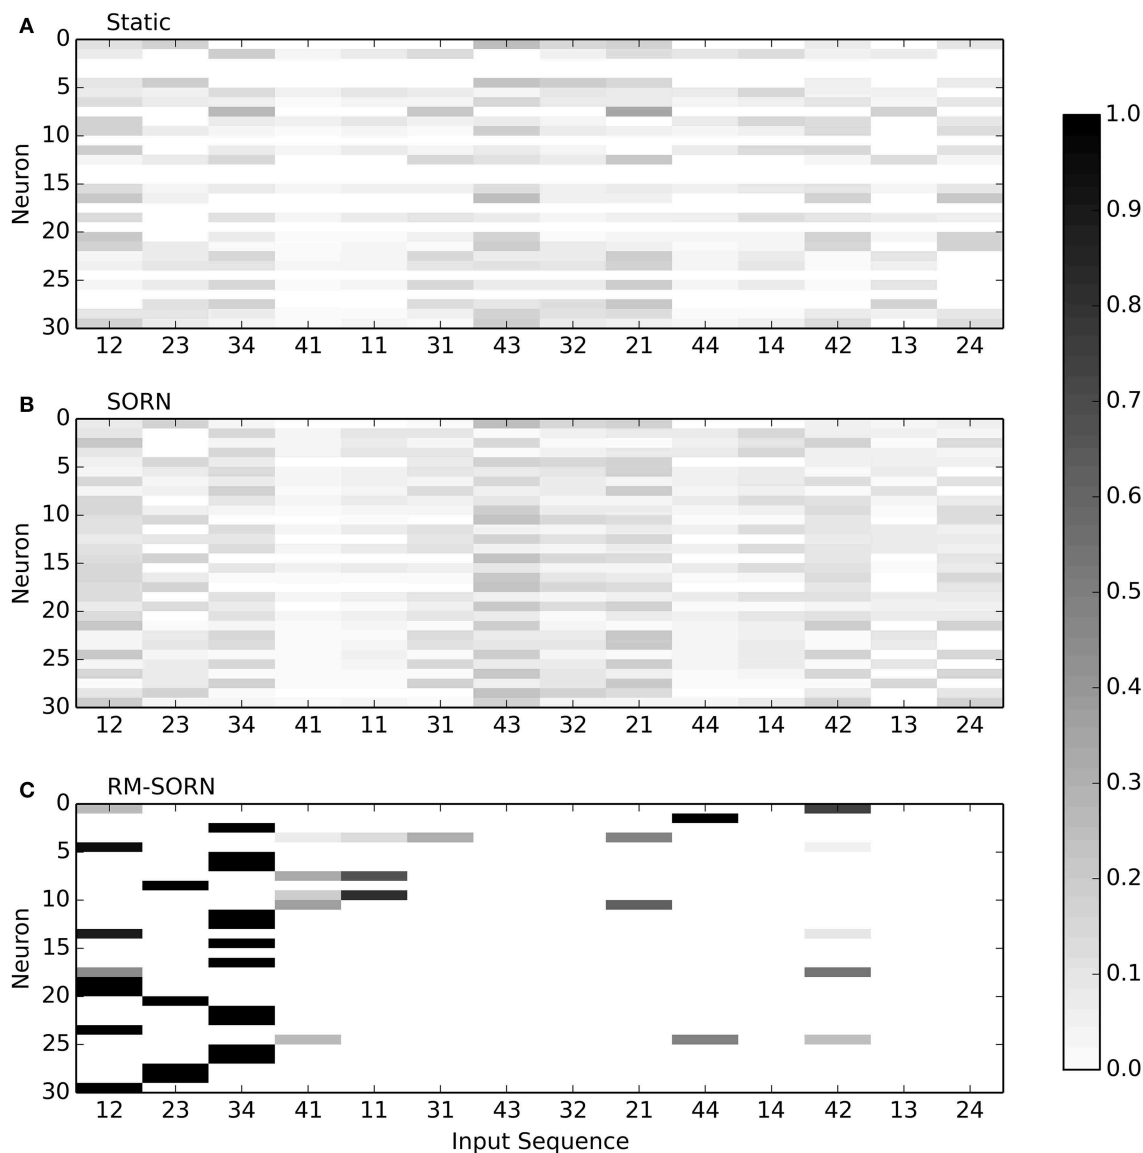


FIGURE 4 | Neuron selectivity in the recurrent layer for the pattern recognition task with 30 neurons. Shown is the probability of a neuron to spike for each possible two-symbol input sequence. For example, the 30th neuron in RM-SORN (C) is activated only by the sequence “12.” Static (A) and SORN (B) are less selective.

Motion Generation

The performance is computed as the percentage of symbols that belong to a target word. Despite rewarding both words, for $n > 4$ the network learns to generate only one word. The performance results are shown in **Figure 2F**. The performance of RM-SORN is impressive, since, in contrast to SORN, which learns with teacher-forcing (Jaeger, 2001), RM-SORN does not receive any external teaching signal, except the reward, and is still capable of generating the desired behavior.

Figure 5 visualizes the activity in the recurrent and output layer during 100 steps of reward-modulated plasticity at different time points during training. In the beginning, the recurrent and output activity is almost constant and has no resemblance to the target words “12345678” and “87654321.”

Then, reward-modulated STDP adapts the output weights, and through feedback, changes the dynamics in the recurrent layer. As can be seen from the output activity in **Figure 4B**, the network generates alternately “123” and “876”—the beginnings of the target words. After an additional 5000 steps of reward-modulation, the network settles on the generation of “87654321.”

Exploration during Motion Generation

In the motion generation task, exploration in the output layer corresponds to the production of different output sequences. **Figure 6** visualizes the extent of exploration during 20,000 steps of training. Shown is the number of unique output sequences of different lengths in the

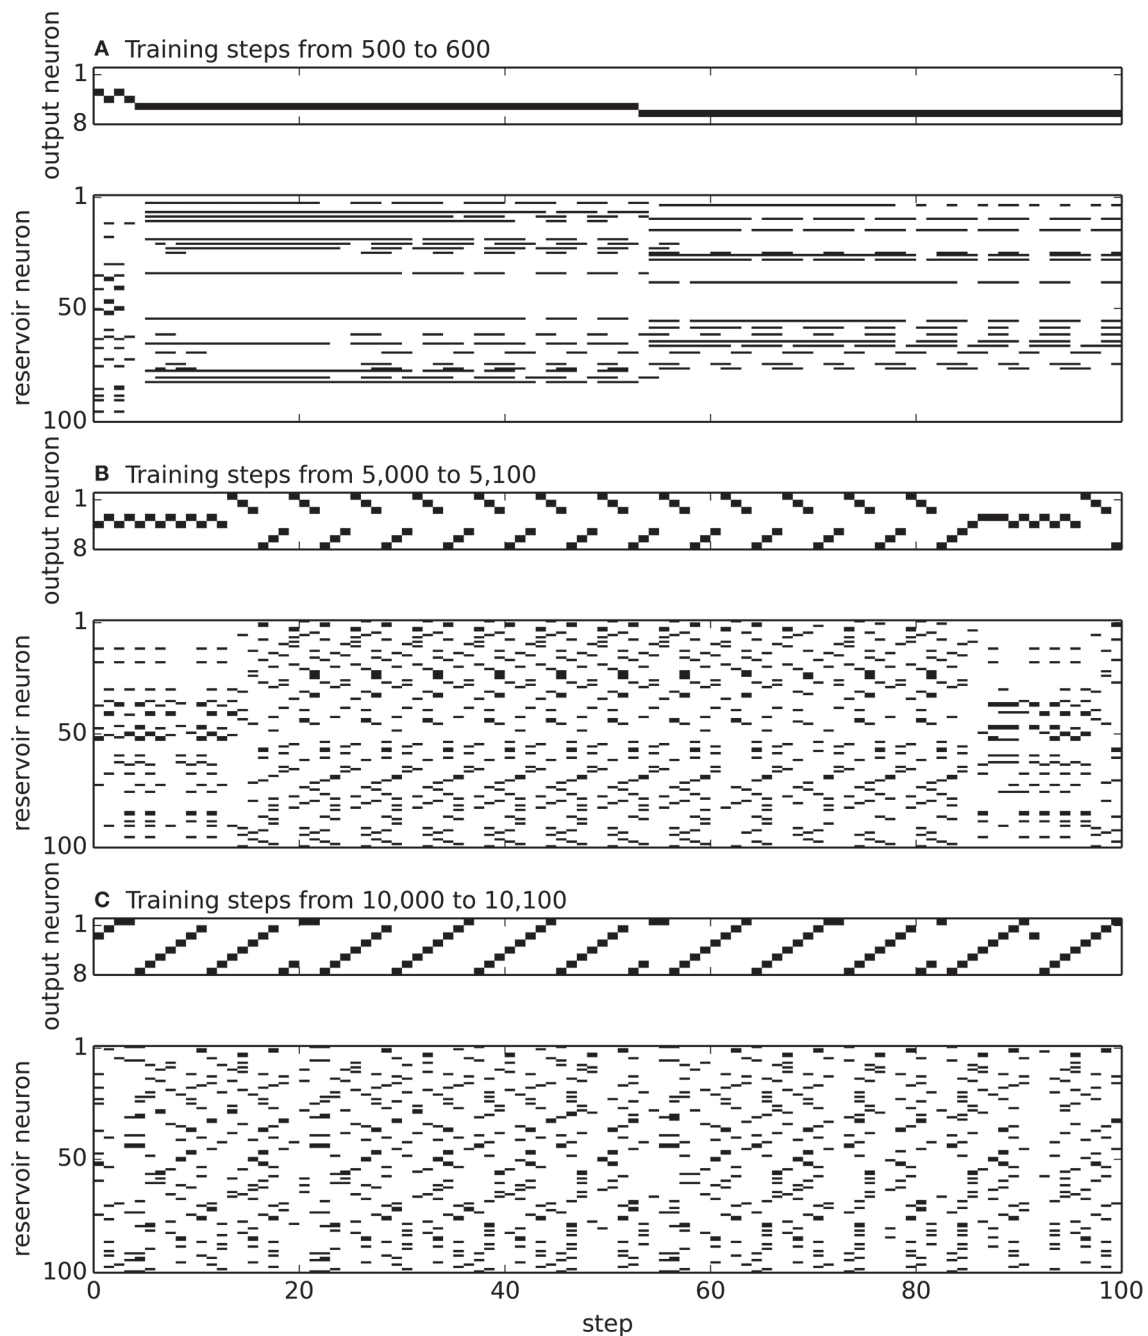


FIGURE 5 | Neural activity during reward-modulated learning in the motion generation task with $n = 8$ and $N = 100$. The narrow tile visualizes the spiking activity of the eight output neurons and the wide tile the activity of the 100 neurons in the recurrent layer. Neuron i allocates the i -th

row of the tile. In the output layer the i -th neuron represents the i -th output symbol. A diagonal in the output tile corresponds therefore either to “1234568” or “87654321.” **(A)** shows the activity from steps 500 to 600, **(B)** from 5000 to 5100 and **(C)** from 10,000 to 10,100.

previous 100 steps, which don't contain any parts of the target words. Exploration is highest during the first 10,000 steps. With time, the output sequences resemble more and more the target words and only few, short, original output sequences are produced. Notably, the exploration doesn't stop completely.

Reward-Modulation in the Recurrent Layer

The pattern recognition task was the only task in which the reward-modulation of the recurrent layer improved the performance. Punishment in the recurrent layer prevented learning of unnecessary input conditions and allowed neurons selective for the relevant input conditions to emerge.




FIGURE 6 | Exploration during training in the Motion Generation Task. Shown is the number of unique output sequences produced during the previous 100 steps, which are not part of the target words. The stronger the color, the longer the unique sequence.

However, the pattern recognition task is the only task, where only a part of the input sequences is relevant. In the other tasks, all input conditions matter. Reward-modulation of the recurrent layer (either suppression or inversion of STDP for wrong outputs) worsened the performance. Enhancement of the exploration of the recurrent layer state space through reward-modulation of the recurrent layer was not observed. The effect of reward-modulation in the recurrent layer on the performance is visualized in the supplementary material.

Effect of Synaptic Normalization of the Weights to the Output Layer

In the recurrent layer, SN is applied to the weights between the neurons, decorrelating them and preventing seizure-like activity (Lazar et al., 2009). In the output layer, SN is applied to the weights from the recurrent to the output neurons. Since in the output layer the neurons are not interconnected and only one neuron is activated at a time (WTA), correlated activity is not possible. The performance comparison of RM-SORN with and without SN, as shown in **Figure 7**, suggests another effect. In the figure, only the performance results for the counting, occluder, motion prediction and motion generation task are shown—in the other tasks the performance differences are negligible.

Depression of the weights from the recurrent to the output layer happens either through punishment—Anti-STDP or SN. The challenge in the counting task is to map the representations of long, similar sequences to the corresponding output neurons. When, through chance, such a representation is mapped correctly, STDP reinforces the weights from the active neurons in the recurrent layer to the activated output neuron. Then, SN scales the weights hereby reducing the weights from the inactive recurrent neurons to the active output neuron. Thus, SN introduces synaptic competition, that

leads to a stronger mapping of the representations of the long sequences.

In the motion tasks the focus is not on long sequences but on a high number of symbols. With increasing n , the number of symbols increases and the number of neurons representing an input decreases. Thus, less neurons represent an input condition and more neurons not related to the input condition need to be ignored and their outgoing weights depressed in order to map an input condition correctly.

The occluder task is a combination of the counting and motion task. Notably, without SN, a higher number of neurons worsens the occluder overall and the motion task performance. When there are more neurons in the recurrent layer, the output neurons have more incoming weights and can be activated more easily by the wrong input representations. In the other tasks the length of sequences, the number of mappings and the number of neurons is moderate—synaptic competition through SN does not lead to an advantage.

Intrinsic Plasticity vs. Noise for Exploration

Exploration of possible output mappings or output sequences is an essential part of reward-modulated learning. Most reward-modulated models use noise for exploration. Noise, however, is per definition transient and random—the rewarded behavior is not guaranteed to appear again, even with the same input and the same neuronal state. A correct model guess induced by noise might even be derogatory. For example, in a prediction task, if the correct target neuron is activated purely by chance, while the input representation in the recurrent layer is “bad” (non-distinctive for the input condition), through STDP, connections with the “wrong” neurons (neurons which don’t represent input conditions or represent other input conditions) are reinforced. IP, on the other hand, is deterministic—the rewarded behavior is reproducible, and when the output is correct, it is always due to the neuronal structure and not to chance. It is therefore not

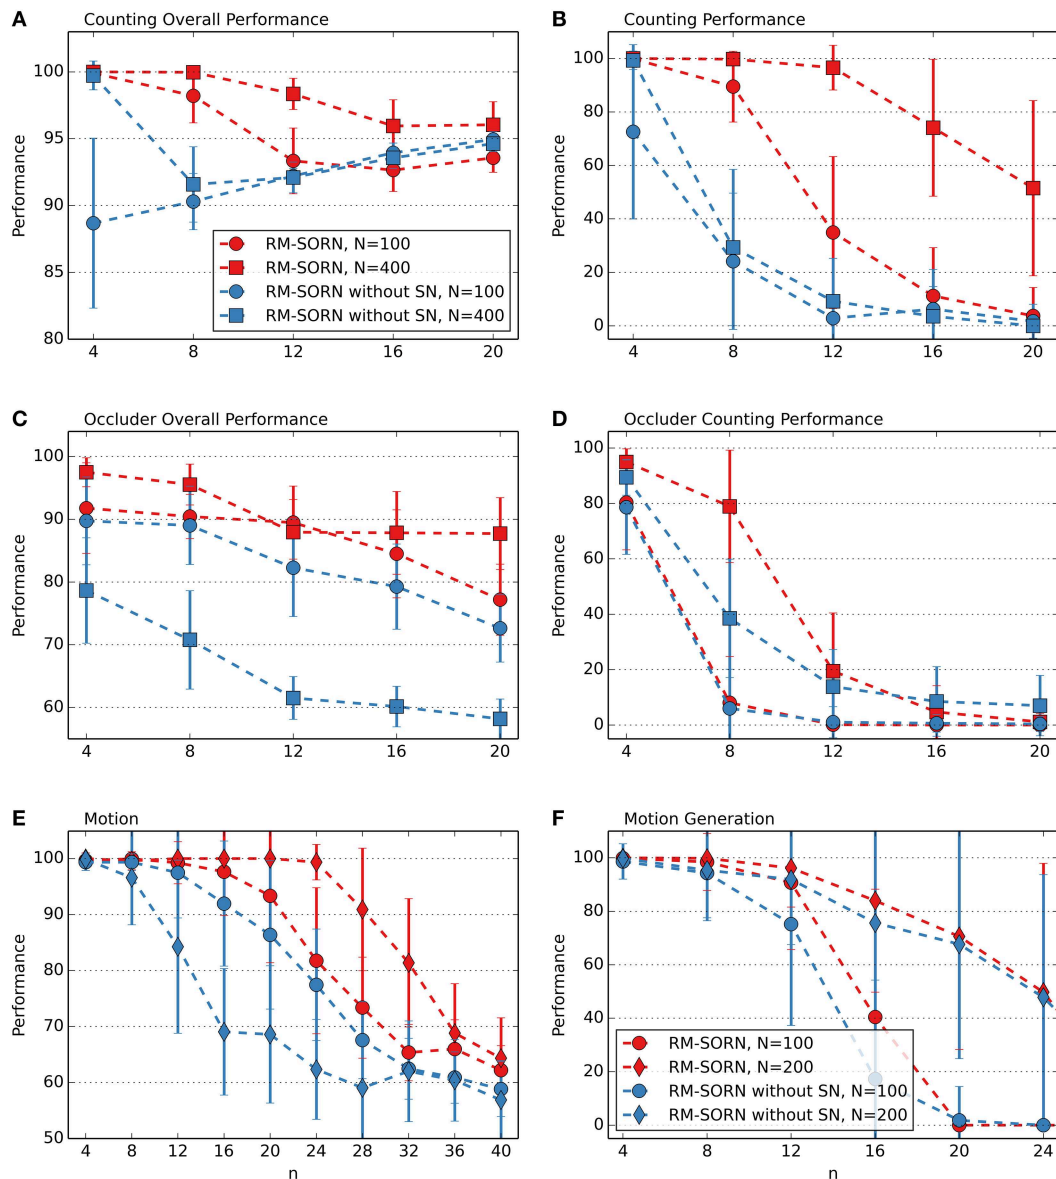


FIGURE 7 | Comparison of RM-SORN with and without SN in the counting (A,B), occluder (C,D), motion (E), and motion generation task (F). Shown is the average performance for different task complexity

values n . In the motion generation task, the results are averaged over 100 networks, in the other tasks the results are averaged over 10 data sets with 10 networks per data set. Error bars indicate standard deviation.

surprising that the performance results confirm the superiority of IP.

The performance comparison with noise was made by disabling IP in the output layer and introducing bit-flip-noise instead: at each step, with a given probability, the active output neuron was set to zero and another randomly chosen output neuron set to one. In order to compare IP and noise in their roles as exploration drives, they have to be aligned, regarding the target average firing rate. Ensuring the target firing rate with noise is not possible, but the thresholds of the output neurons can be selected to match on average the thresholds found through IP. Therefore, before the actual task, for each network, a preliminary

run with 20,000 steps with unmodulated plasticity was made and the output thresholds averaged over the values of the last 1000 steps. Then, the network was reset to its initial state, but with the threshold averages as the new threshold values.

During simulations without IP at different noise levels, the highest performance results were obtained with noise probabilities of 5, 15, and 100%. **Figure 8** compares the performance of RM-SORN with IP and RM-SORN without IP but with noise at these levels. Overall, networks with IP achieve a higher performance than networks without IP, but with noise. Particularly motion generation would not be possible with noise as exploration drive. Noise achieves slightly

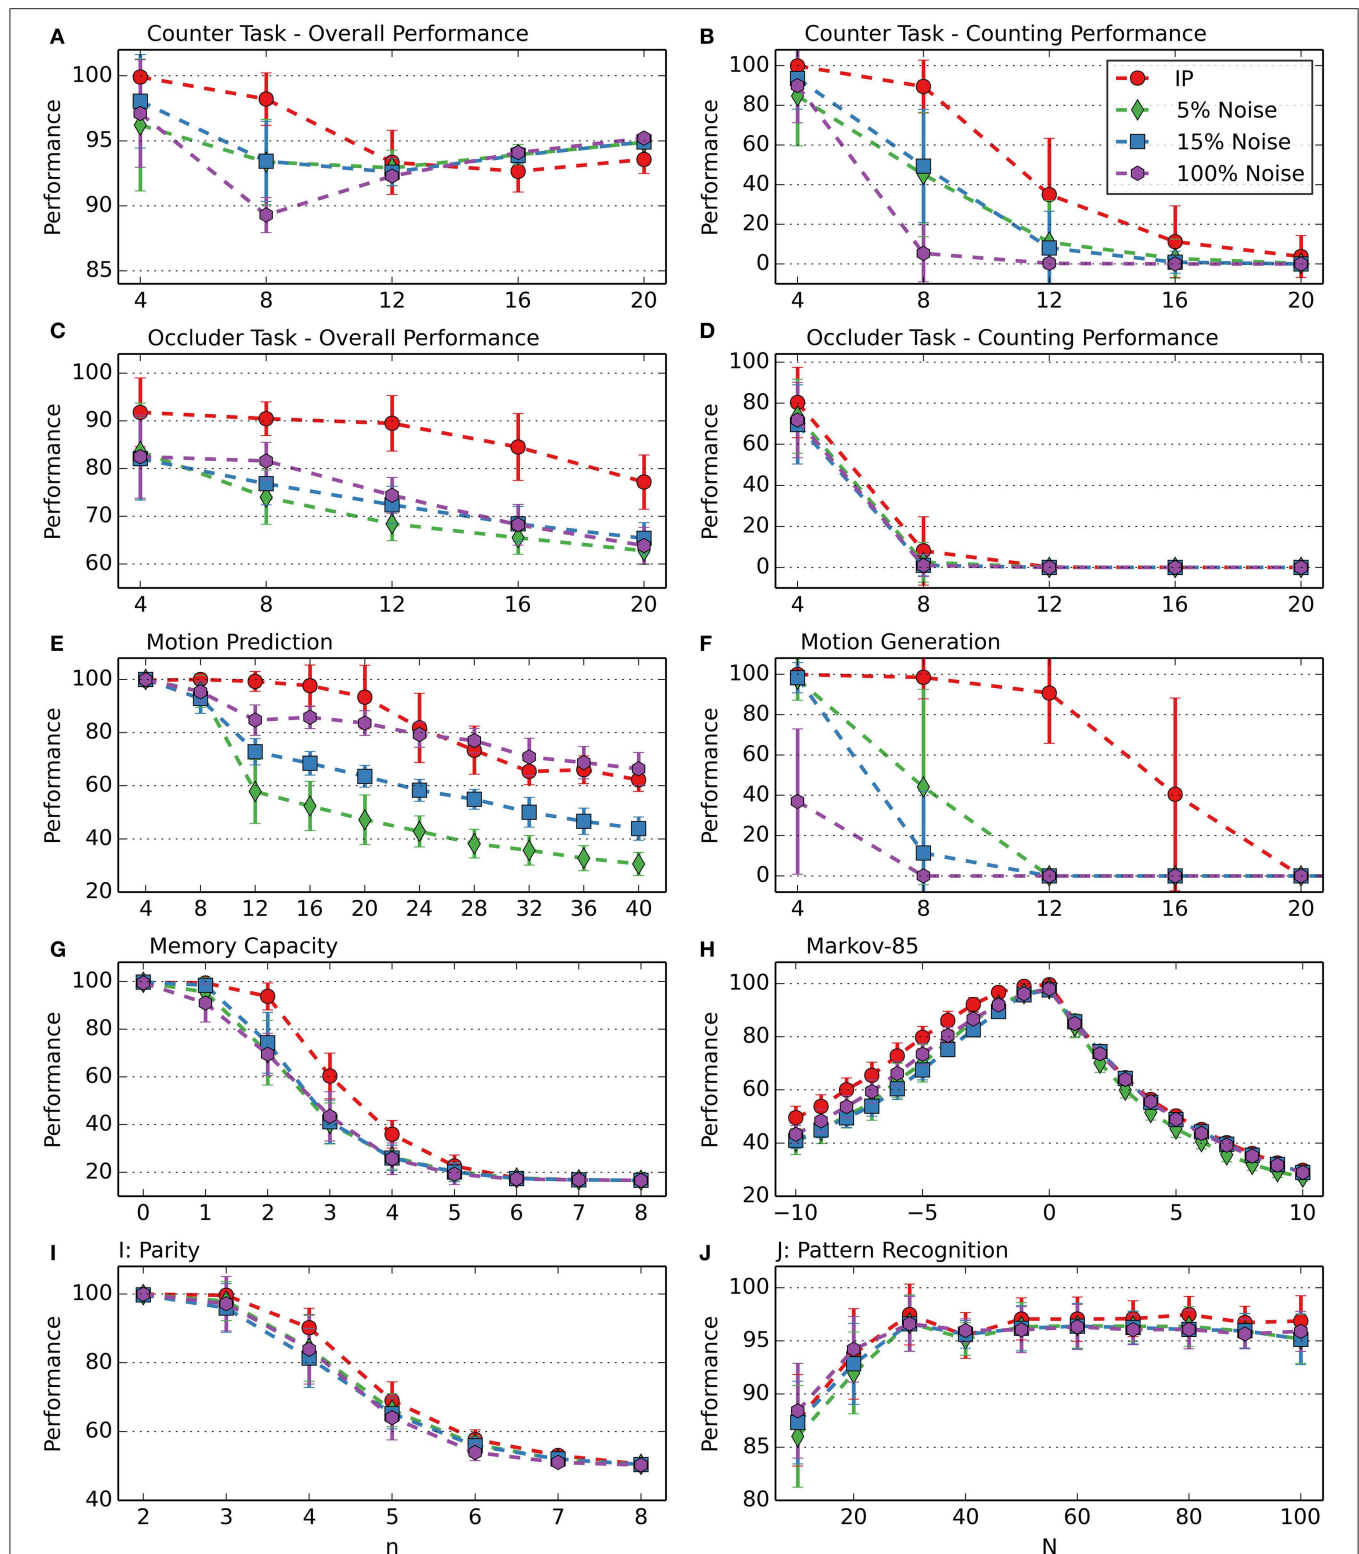


FIGURE 8 | Comparison of RM-SORN with IP and without IP but with noise at different noise levels in the counting (A,B), occluder (C,D), motion prediction (E), motion generation (F), memory capacity (G), Markov-85 (H), parity (I), and pattern recognition (J) task. The noise probability is the probability of a randomly chosen

output neuron to be activated at each training step. In the pattern recognition task, the network size N was varied. In the other tasks, the task difficulty n was varied with $N = 100$. Shown are the averages over 10 data sets with 10 networks per data set. Error bars indicate standard deviation.

higher performance in the motion task for $n > 24$, and roughly similar performance in the parity and pattern recognition task, and also in the prediction part of the Markov-85 task.

These results demonstrate the power of IP as exploration drive. Noise is a comparatively weak alternative.

Discussion

In this article, we introduced the RM-SORN, in which reward-modulated STDP replaced supervised learning in the readout of SORN and additionally, when applied to the recurrent layer in the pattern recognition task, fitted the representation of the inputs to the task goals. RM-SORN achieved high performance comparable to that of supervised trained networks in all tasks. For complex tasks (high n), the performance deteriorated more strongly than in SORN, from which we can conclude that in RM-SORN, similar representations in the recurrent layer cannot be differentiated as well as in SORN. This is not surprising, considering that supervised methods have an exact error signal, while reward-modulated learning has only a general goodness signal and works on a trial and error basis.

Reinforcement Learning and Reward-Modulation

RM-SORN and similar models are a form of reinforcement learning, where the state is defined by the network activation and the action by the readout. The weights are adapted to maximize the reward. In a recent review of reinforcement learning in cortical networks, Senn and Pfister (2014) generalize the weight update rule to follow (10), with R being the reward and PI a plasticity induction based on pre- and postsynaptic activity. The hypothesis, that synaptic plasticity is driven by the covariance between reward and neural activity was initially introduced by Loewenstein and Seung (2006).

$$\Delta w = R * PI \quad (10)$$

Senn and Pfister differentiate between policy gradient methods, where the average policy induction $\langle PI \rangle = 0$ and Temporal Difference (TD) methods, where the average reward $\langle R \rangle = 0$. The postulated purpose of these restrictions is to prevent systematic weight drift. A simple alternative to $\langle R \rangle = 0$ is to subtract the average reward from the modulating factor as is done in RM-SORN with the rewarding strategy Mk. A similar method was used by Hoerzer et al. (2014).

TD learning with SORN was implemented by Franz (2010): the readout was replaced with action neurons and the weights modulated via the TD error. The network was able to learn symbolical sequences. A more complex actor-critic network was implemented by Frémaux et al. (2013) based on simplified spike response model neurons and used to solve a version of the cartpole task. Most recently, Dasgupta et al. (2014) developed a model, consisting of a recurrent neural network critic model representing the basal ganglia and a feed-forward correlation-based learning model representing the cerebellum. This combinatorial model was validated by letting it control a robot to forage in an enclosed environment.

These increasingly complex, biologically motivated models of reward-based reinforcement learning in neural networks are able to solve complex tasks, but neglect other forms of synaptic plasticity.

Reward-Modulated STDP Models and Exploration

The core of RM-SORN is the interaction of IP and STDP: IP explores possible output mappings or output sequences, and STDP reinforces the rewarded ones. In contrast to most previous models, noise for exploration is neither necessary nor desirable, as IP is considerably superior to noise in most tasks and on the same level in the rest.

Previous reward-modulated models, with one exception, use only STDP or a hebbian rule, and noise for exploration. Legenstein et al. (2008) performed an extensive analytical and simulational analysis of reward-modulated STDP. Their networks made from noisy, leaky integrate-and-fire neurons solved several tasks: increasing the firing rate of a single neuron, learning of spike times, spike pattern discrimination and isolated digit recognition. One of their findings was that spontaneous activity is essential for reward-modulated learning in order to explore which firing patterns are rewarded.

Shortly before Legenstein, in 2007, Izhikevich created a model with spiking neurons, where the distal reward problem was solved through eligibility traces and reward-modulated STDP (Izhikevich, 2007). He validated his model on three simulations: reinforcement of a synapse between two excitatory neurons, classical (Pavlovian) conditioning, and stimulus-response-instrumental conditioning. During learning, spontaneous activity was achieved through random input, mimicking noisy miniature PSPs. Izhikevich concluded that STDP is insensitive to random firings during the waiting time for the reward, and is only triggered by precise firing patterns in the millisecond range, which are rare. He also argued that the precise timing of spikes is essential for reinforcement with STDP, and that this effect could not be reproduced with firing rate models. This statement was disproved by Soltoggio and Steil (2013).

Soltoggio and Steil reproduced most of the experiments of Izhikevich in a rate-based model, and they showed that classical and instrumental conditioning with delayed rewards can be learned without precise spike timing. Their Rare Correlation Model features a rate-based hebbian rule with a threshold that allows only the upper 1% of all correlations to be applied. Noise is added to the firing rate after the tanh-activation, to generate spontaneous activity. Beside the tasks from Izhikevich, the model was also successfully applied in robotics for classical and operant conditioning of the humanoid robot iCub (Soltoggio et al., 2013).

Another rate-based model with reward-modulated hebbian learning was created by Hoerzer et al. (2014). It consists of a recurrent layer and a linear readout with a feedback connection to the recurrent layer. In Hörzer's model, the recurrent layer is chaotic with tanh-neurons, following Sussillo and Abbott (2009), but instead of supervised learning, a reward-modulated hebbian rule is used to train the readout. The

model was successfully applied to several tasks, including periodic pattern generation and non-linear analog computations on complex input signals. During learning, noise was applied to the firing rates of neurons in the readout. The authors point out that this exploration noise is the driving force for learning, and that without it, no learning would take place.

Other notable reward-modulated spiking neuron models include Soltani and Wang (2010), where a connection between reward-modulated plasticity and probabilistic inference was observed and Bourjaily and Miller (2011), where reward-modulated STDP combined with multiplicative synaptic scaling was used to learn a XOR task.

The most recent model combines reward-modulated STDP with eligibility traces, IP and synaptic scaling in a 2-layer network with binary thresholded units (Savin and Triesch, 2014), similar to the model in this article. The model was applied to two working memory tasks: delayed response and delayed categorization. During learning, task-dependent representations beneficial for task performance emerged in the recurrent layer. The neurons were noisy, but the noise did not play any special role in learning.

The focus in these publications is on the hebbian or STDP learning rule and tasks with delayed reward. The interaction of several plasticity mechanisms and the role, that homeostatic plasticity plays in reward-modulated learning was not investigated. This article demonstrates that reward-modulated learning can achieve performance comparable to that of supervised learning methods in tasks of different nature and complexity, and that IP can serve as the exploratory drive during learning. This is at first glance surprising, since IP is a homeostatic mechanism. However, IP as implemented here ensures an average firing rate by lowering or raising the thresholds continuously and these threshold-changes alter the neuronal activity and drive the exploration.

References

- Ahissar, E., Vaadia, E., Ahissar, M., Bergman, H., Arieli, A., and Abeles, M. (1992). Dependence of cortical plasticity on correlated activity of single neurons and on behavioral context. *Science* 257, 1412–1415. doi: 10.1126/science.1529342
- Bourjaily, M. A., and Miller, P. (2011). Synaptic plasticity and connectivity requirements to produce stimulus-pair specific responses in recurrent networks of spiking neurons. *PLoS Comput. Biol.* 7:e1001091. doi: 10.1371/journal.pcbi.1001091
- Dasgupta, S., Wörgötter, F., and Manoonpong, P. (2014). Neuromodulatory adaptive combination of correlation-based learning in cerebellum and reward-based learning in basal ganglia for goal-directed behavior control. *Front. Neural Circuits* 8:126. doi: 10.3389/fncir.2014.00126
- Desai, N. S., Rutherford, L. C., and Turrigiano, G. G. (1999). Plasticity in the intrinsic excitability of cortical pyramidal neurons. *Nat. Neurosci.* 2, 515–520. doi: 10.1038/9165
- Duarte, R., Seriès, P., and Morrison, A. (2014). “Self-organized artificial grammar learning in spiking neural networks,” in *36th Annual Conference of the Cognitive Science Society*. (Freiburg).
- Fetz, E. E. (1969). Operant conditioning of cortical unit activity. *Science* 163, 955–958. doi: 10.1126/science.163.3870.955

Outlook

Supervised learning requires a precise error signal, which is probably not present in the brain. For reward-modulated learning, only a general goodness signal is necessary. Additionally, it can be applied to any network structure. In this article, only two-layer networks were investigated, which in itself is biologically not plausible. More complex network architectures, which present a challenge for supervised training methods, may, in contrast, unfold the capabilities of reward-modulated learning.

Another possible line of investigation is the effect and nature of modulation via the reward-prediction error. In all tasks, except the counting- and pattern prediction-task, modulation via the reward-prediction error was better than modulation via the reward directly. The difference in the task performances offers a starting point for a more detailed investigation.

An interesting question is also, how exploration and exploitation can be balanced. During reward-modulated plasticity the exploration diminishes, but never stops completely, which made it necessary to evaluate intermediate networks in order to get the best performance (as described in sections Model Evaluation and Exploration during Motion Generation). From a functional point of view, a mechanism that stops exploration, when a sufficient performance level is achieved, is desirable. The rewarding strategy seems to be a good place for such a mechanism.

Supplementary Material

The Supplementary Material for this article can be found online at: <http://www.frontiersin.org/journal/10.3389/fncom.2015.00036/abstract>

Effect of reward-modulation of the recurrent layer on performance, 603 parameters for the tasks, detailed performance results (means and standard deviation).

- Franz, A. (2010). *Neural Network Models of Cognitive Development in Infancy*. Frankfurt: Frankfurt Institute for Advanced Studies (FIAS).
- Frémaux, N., Sprekeler, H., and Gerstner, W. (2013). Reinforcement learning using a continuous time actor-critic framework with spiking neurons. *PLoS Comput. Biol.* 9:e1003024. doi: 10.1371/journal.pcbi.1003024
- Hoerzer, G. M., Legenstein, R., and Maass, W. (2014). Emergence of complex computational structures from chaotic neural networks through reward-modulated Hebbian learning. *Cereb. Cortex* 24, 677–690. doi: 10.1093/cercor/bhs348
- Izhikevich, E. M. (2007). Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cereb. Cortex* 17, 2443–2452. doi: 10.1093/cercor/bhl152
- Jaeger, H. (2001). The “echo state” approach to analysing and training recurrent neural networks.
- Lawson, C. L., and Hanson, R. J. (1995). *Solving Least Squares Problems*. Englewood Cliffs, NJ: SIAM. doi: 10.1137/1.9781611971217
- Lazar, A., Pipa, G., and Triesch, J. (2009). SORN: a self-organizing recurrent neural network. *Front. Comput. Neurosci.* 3:23. doi: 10.3389/fncom.2009.0023.2009
- Legenstein, R., Pecevski, D., and Maass, W. (2008). A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback. *PLoS Comput. Biol.* 4:e1000180. doi: 10.1371/journal.pcbi.1000180

- Loewenstein, Y., and Seung, H. S. (2006). Operant matching is a generic outcome of synaptic plasticity based on the covariance between reward and neural activity. *Proc. Natl. Acad. Sci. U.S.A.* 103, 15224–15229. doi: 10.1073/pnas.0505220103
- Pawlak, V., Wickens, J. R., Kirkwood, A., and Kerr, J. N. D. (2010). Timing is not Everything: Neuromodulation Opens the STDP Gate. *Front. Synaptic Neurosci.* 2:146. doi: 10.3389/fnsyn.2010.00146
- Savin, C., and Triesch, J. (2014). Emergence of task-dependent representations in working memory circuits. *Front. Comput. Neurosci.* 8:57. doi: 10.3389/fncom.2014.00057
- Senn, W., and Pfister, J.-P. (2014). “Reinforcement learning in cortical networks,” in *Encyclopedia of Computational Neuroscience*, eds D. Jaeger and R. Jung (New York, NY: Springer), 1–6.
- Sigala, N., and Logothetis, N. K. (2002). Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature* 415, 318–320. doi: 10.1038/415318a
- Soltani, A., and Wang, X.-J. (2010). Synaptic computation underlying probabilistic inference. *Nat. Neurosci.* 13, 112–119. doi: 10.1038/nn.2450
- Soltoggio, A., Lemme, A., Reinhart, F., and Steil, J. J. (2013). Rare neural correlations implement robotic conditioning with delayed rewards and disturbances. *Front. Neurobot.* 7:6. doi: 10.3389/fnbot.2013.00006
- Soltoggio, A., and Steil, J. J. (2013). Solving the distal reward problem with rare correlations. *Neural Comput.* 25, 940–978. doi: 10.1162/NECO_a_00419
- Sussillo, D., and Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron* 63, 544–557. doi: 10.1016/j.neuron.2009.07.018
- Toutounji, H., and Pipa, G. (2014). Spatiotemporal computations of an excitable and plastic brain: neuronal plasticity leads to noise-robust and noise-constructive computations. *PLoS Comput. Biol.* 10:e1003512. doi: 10.1371/journal.pcbi.1003512
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., and Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391, 892–896. doi: 10.1038/36103
- Zheng, P., Dimitrakakis, C., and Triesch, J. (2013). Network self-organization explains the statistics and dynamics of synaptic connection strengths in cortex. *PLoS Comput. Biol.* 9:e1002848. doi: 10.1371/journal.pcbi.1002848

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2015 Aswolinskiy and Pipa. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.



Emergence of task-dependent representations in working memory circuits

Cristina Savin^{1*†} and Jochen Triesch^{1,2}

¹ Frankfurt Institute for Advanced Studies, Frankfurt am Main, Germany

² Physics Department, Goethe University, Frankfurt am Main, Germany

Edited by:

Friedemann Zenke, École
Polytechnique Fédérale de
Lausanne, Switzerland

Reviewed by:

Sukbin Lim, University of Chicago,
USA

Alberto Bernacchia, Jacobs
University Bremen, Germany

*Correspondence:

Cristina Savin, Institute of Science
and Technology, Am Campus 1,
Klosterneuburg 3400, Austria
e-mail: csavin@ist.ac.at

† Present address:

Cristina Savin, Institute of Science
and Technology, Klosterneuburg,
Austria

A wealth of experimental evidence suggests that working memory circuits preferentially represent information that is behaviorally relevant. Still, we are missing a mechanistic account of how these representations come about. Here we provide a simple explanation for a range of experimental findings, in light of prefrontal circuits adapting to task constraints by reward-dependent learning. In particular, we model a neural network shaped by reward-modulated spike-timing dependent plasticity (r-STDP) and homeostatic plasticity (intrinsic excitability and synaptic scaling). We show that the experimentally-observed neural representations naturally emerge in an initially unstructured circuit as it learns to solve several working memory tasks. These results point to a critical, and previously unappreciated, role for reward-dependent learning in shaping prefrontal cortex activity.

Keywords: working memory, reward-dependent learning, STDP, intrinsic plasticity, synaptic scaling, prefrontal cortex, delayed categorization

1. INTRODUCTION

Working memory is defined as the temporary storage of stimulus-specific information during a delay period. This function has been traditionally associated with circuits in prefrontal cortex (PFC). Classic work in monkeys revealed that single neurons in this region exhibit selective persistent activity during the delay period (Miyashita, 1988; Goldman-Rakic, 1990) and its disruption (by electrical stimulation, or due to distracters) leads to a decay in performance (Funahashi et al., 1989). These early observations have been interpreted as the circuit exhibiting attractor dynamics, which enable a subset of the neurons to maintain high firing throughout the delay after a brief stimulus presentation (Amit and Brunel, 1997; Brunel and Wang, 2001). This view has been revised in recent years, as it was shown that most neurons in PFC change their firing rates during the delay (Miller et al., 1996; Chafee and Goldman-Rakic, 1998; Pesaran et al., 2002; Rainer and Miller, 2002; Barak et al., 2010), suggesting that working memory circuits rely on feedforward rather than attractor dynamics (Goldman, 2009). Still, while experiments generally agree on how information is represented in working memory circuits, i.e., using spatio-temporal patterns of neural activity, exactly what information gets encoded is less clear.

An accumulation of data across different working memory experiments paints an increasingly complex picture of the features encoded in PFC. We find neurons may represent the previous stimulus, the forthcoming action, or a more complex function of the two (Durstewitz et al., 2000). When the task requires a generalization across stimuli, neurons develop category selectivity (Freedman et al., 2001). Moreover, there is a gradual shift in these representations as the number of examples per class increases, with animals switching from a stimulus-response association strategy to representing categorical distinctions directly (Antzoulatos and Miller, 2011).

Things get even more complicated when animals need to alternate between different tasks. While PFC neurons generally represent the task to be performed (Asaad et al., 1998; Cromer et al., 2010; Roy et al., 2010; Warden and Miller, 2010; Meyer et al., 2011), they can differ significantly with respect to how the information is distributed across the population in different tasks. For instance, neurons can show task-specific changes in overall firing rate, in time-dependent response profiles and in stimulus and response selectivity (Asaad et al., 2000). In some situations, the same neurons seem to participate in encoding features related to different tasks (e.g., making different category distinctions, Cromer et al., 2010), effectively multiplexing information across contexts. In other situations, however, information is encoded in different neurons for different contexts (Roy et al., 2010), and—worse still—it is unclear when one or the other coding strategy may be employed. Generally, we are missing a unifying account for PFC representations during the delay period.

Here we hypothesize that reward-dependent learning underlies the variety in PFC representations in different working memory tasks. The data itself suggest that this may be the case: the most striking feature of the above experiments is not the diversity of neural responses, but the sheer number of neurons displaying an effect. Regardless of the actual task the monkey has been trained to carry out, a significant subset of the recorded neurons are found to exhibit selectivity to the specifics of that particular task. This is a strong indication that PFC neurons adapt their responses to reflect current cognitive demands. Indeed, PFC representations change significantly over the course of training (Rainer et al., 1998b; Rainer and Miller, 2000; Baeg et al., 2003; Kennerley and Wallis, 2009). Neural responses become increasingly sparse, the tuning of the neurons narrows, and the representation becomes more robust to input noise (Rainer and Miller, 2000). These changes in neural representation parallel behavioral learning, and

allow for a better decoding of stimuli and actions (Baeg et al., 2003). Moreover, since the training-induced changes in neural responses include changes in functional connectivity (Baeg et al., 2007), it seems likely that associative learning within the circuit is responsible—at least in part—for the refinement of neural representations with learning. The specific mechanisms involved remain unclear, however.

We assume that learning in PFC is reward-dependent. This hypothesis is consistent with the observation that dopamine, a neuromodulator associated with reward prediction error (Schultz, 1998), modulates synaptic plasticity in this circuit (Otani et al., 2003). It also explains the dependence of neural representations on the magnitude of the expected reward (Kennerley and Wallis, 2009). However, the primary reason for our assumption is computational. Working memory circuits are known to operate under strict capacity constraints (Cowan, 2001), and a circuit with limited resources cannot simply encode everything. To perform well, it needs to represent the specific aspects of the stimulus that matter for the task at hand (Duncan, 2001). Hence, reward should modulate learning so as to shift representations toward task-relevant features. This points to a critical and previously unrecognized role for reward-dependent plasticity in shaping prefrontal representations.

Can reward-dependent learning alone explain the wide variety of experimental observations on PFC encoding? To address this question, we studied the effects of reward-dependent learning on the encoding properties of neurons in a working memory circuit. More specifically, we trained a generic recurrent neural network to solve tasks similar to those employed in working memory experiments. We then investigated the neural representations emerging in the circuit and compared them to neural data. We chose a simple abstract model for the network dynamics in which the output of a neuron depends only on its instantaneous inputs. Learning was implemented by reward-dependent spike timing dependent plasticity (rSTDP) (Izhikevich, 2007), supplemented by homeostatic mechanisms that stabilized the network dynamics during learning (Lazar et al., 2009). Importantly, as individual neurons have no memory themselves, the storage of information in this circuit relies exclusively on the recurrent connectivity. While this simple model cannot capture the full complexity of the temporal dynamics in PFC, it allows us to focus specifically on the reward-dependent reorganization of recurrent connections and its effects on circuit function.

We found that our model is able to capture key aspects of neuronal dynamics during working memory tasks. Neurons in the model develop specificity in space and time and, depending on the task, they preferentially encode individual stimuli, actions, or context information. In a simple delayed-response task, neurons encode stimulus identity (Miller et al., 1996; Constantinidis and Franowicz, 2001). In a delayed-categorization task, neurons learn to preferentially encode category boundaries (Freedman et al., 2001). Lastly, when learning several tasks at the same time, the degree of neural specialization depends on the specifics of the task, mirroring experimental data. When the task involves several independent category schemes, neurons act as “multiplexers,” coding for different things in different contexts (Cromer et al., 2010); when the same stimuli need to be categorized differently

depending on behavioral context, the neurons segregate into distinct task-specific subpopulations (Roy et al., 2010). Furthermore, reward-dependent learning is critical for these results. A similar circuit trained by unsupervised learning shows a significant loss in working memory performance, paired with poorer neural representations. Taken together, our findings show reward-dependent learning could be a central force in the organization of working memory circuits.

2. MATERIALS AND METHODS

2.1. THE GENERAL TASK

The working memory tasks we investigated share a simple general structure (**Figure 1A**): at the beginning of a trial one stimulus (out of K) is briefly presented to the network. After a delay period (either fixed for a block of trials or selected at random from a given distribution) a “Go” cue is presented, after which the reward is given according to the action selected by the model (one out of M)—either +1 for a correct answer or −1, otherwise. Different tasks correspond to different mappings between stimuli and actions and each are described in detail in the corresponding Results section. To speed up learning, we adopt the same strategy employed in training animals for experiments, i.e., we start with the minimum delay version of the task and progressively increase the duration of the delay period during learning (Klingberg, 2010).

2.2. NETWORK MODEL

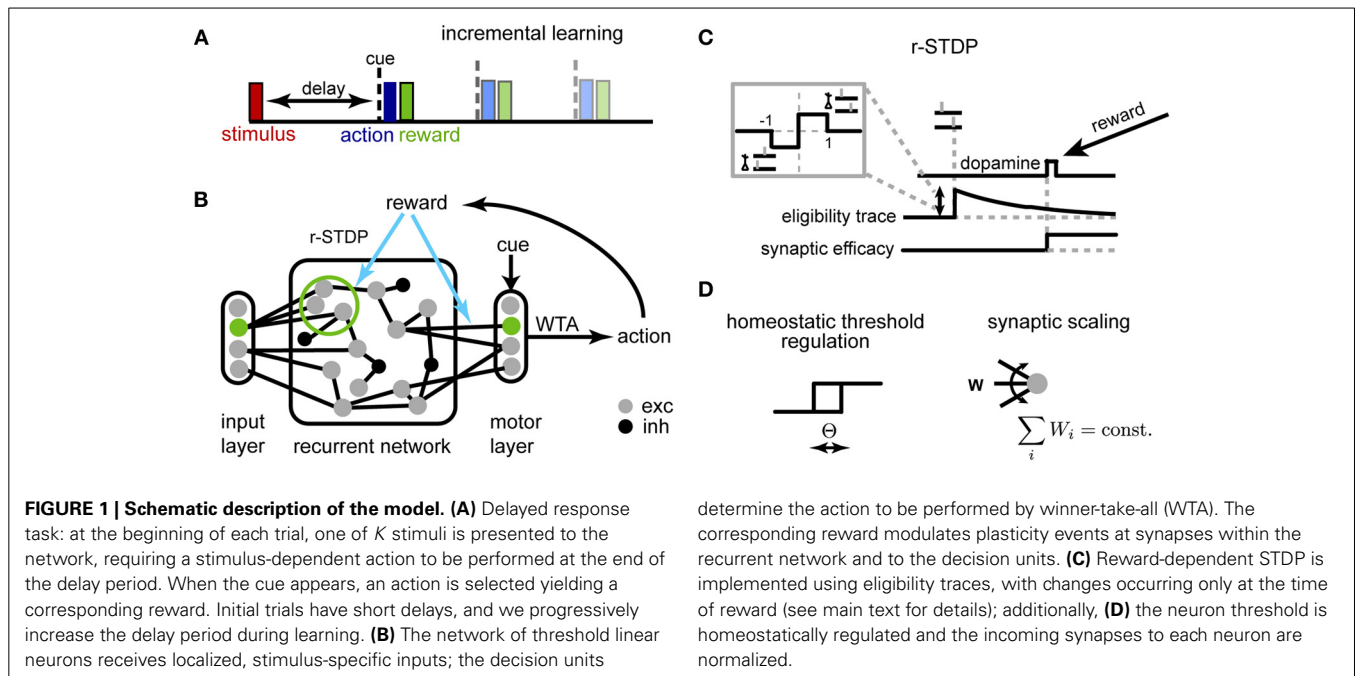
An overview of the network is shown in **Figure 1B**. The recurrent network consists of N units (unless otherwise specified, $N = 250$), 80% excitatory and 20% inhibitory, with sparse random connectivity. Input units encoding different stimuli (and possibly the context cue) activate small, non-overlapping subsets within the recurrent layer, each consisting of N_{in} excitatory neurons $N_{\text{in}} = 5$; the activation of the input unit provides a suprathreshold current which forces the corresponding subpopulation to be active for one time step. The output layer receives inputs from all excitatory units within the network and generates a decision response through a winner-take-all (WTA) mechanism. This decision outcome determines the received reward, which in turn modulates synaptic changes through r-STDP. Reward-dependent learning affects both excitatory synapses within the recurrent network and those connecting to the decision layer.

We chose a abstract model for the neural dynamics, whose simplicity allows us to focus on the essential mechanisms required for explaining the data. More specifically, we used linear threshold units to model neurons within the network, i.e., each unit has a binary output:

$$x_i(t) = I_i(t) \geq \Theta_i, \quad (1)$$

with activation depending on the total current to the neuron $I_i(t)$ and the neuron's spike threshold Θ_i (this threshold also changes over a slower time scale because of homeostatic mechanisms, see below). The activity proceeds in discrete time steps, with synchronous updates for all neurons. The input to a neuron is given by:

$$I_i = \mathbf{w}_i^T \cdot \mathbf{x} + \epsilon, \quad (2)$$



where column vectors w and x describe the synaptic weights and the activity of all presynaptic neurons, respectively. The stochastic term ϵ corresponds to an unspecific background input to each unit, modeled as independent uniform random noise, $\epsilon \in [0, 0.1]$. Importantly, since the model neuron has no memory itself, working memory can develop in the model only through the network dynamics. Hence, we can use the model to study specifically reward-dependent plasticity and its effects on information storage.

The connectivity matrix was initialized randomly at the beginning of each experiment, with weights drawn from the uniform distribution $w_{ij} \in [0, 1]$, followed by a sum-to-one weight normalization of incoming synapses. The connection probabilities were $p_{ee} = 0.1$, $p_{ei} = 0.25$, $p_{ie} = 0.4$, $p_{ii} = 0$, with indices “e” and “i” marking the excitatory and inhibitory populations, respectively.

For the decision layer, the current to each neuron is computed as before, with the WTA mechanism selecting the neuron with the strongest input as the only active unit: $I_m = w_i^T \cdot x + \epsilon$, $x_m = 1$ if $m = \arg\max_j I_j$, and $x_m = 0$, otherwise. Decision neurons were allowed to fire during the delay period without any effect on reward.

2.3. PLASTICITY MECHANISMS

2.3.1. Reward-dependent learning

We adapted a model for r-STDP from Izhikevich (2007) (Figure 1C). As in the original, each synapse has an associated eligibility trace e_{ij} :

$$\frac{de_{ij}}{dt} = -\frac{e_{ij}}{\tau_e} + x_i(t) \cdot x_j(t-1) - f \cdot x_i(t-1) \cdot x_j(t) \quad (3)$$

where x_i and x_j are the output of the pre- and post-synaptic neuron, respectively, and f is a model parameter ($f = 1$ for synapses

in the recurrent layer, and $f = 0.01$ for synapses in the motor layer).

The eligibility trace stores a history of potential weight changes at the synapse, with an exponential decay, specified by the time constant τ_e ($\tau_e = 2.5$). The individual synaptic plasticity events follow a simplified STDP window: potentiation occurs when presynaptic activity is followed by a postsynaptic spike, while the reverse pattern causes depression, with a width of 1 time step (since that is the timescale of causal interactions in our network). Additionally, weights are rectified such that $w_{ij} \geq 0$ in order to respect Dale's law.

At the time of the reward synaptic weights change proportionally to the eligibility trace e_{ij} and the reward signal r :

$$w_{ij}(t+1) = w_{ij}(t) + \eta \cdot r(t) \cdot e_{ij}(t), \quad (4)$$

with learning rate η .

For simplicity, we used the absolute reward as the signal modulating synaptic modifications instead of the reward prediction error (Schultz, 1998), as done in previous models (Izhikevich, 2007). Additionally, we assumed the reward to be either positive or negative, as biological evidence from cortico-striatal synapses suggests that dopamine can induce both potentiation or depression in response to tetanic stimulation, depending on its concentration relative to baseline (Reynolds and Wickens, 2002). Specifically, at the time of the reward delivery $r(t) = 1$, if the motor output was correct and $r(t) = -1$, otherwise; $r(t) = 0$ at all other times.

To ensure that the system is given time to exploit the emerging neural representation, we assumed that changes at synapses to the decision layer occur faster than those in the recurrent network ($\eta = 10^{-5}$ for synapses in the recurrent layer and $\eta = 10^{-4}$ for those connecting to decision neurons). These changes in learning

rate were paralleled for intrinsic plasticity to ensure that the dynamics remain stable during learning (see below).

2.3.2. Homeostatic plasticity

A critical problem when optimizing recurrent networks is how to stabilize the dynamics during learning (Turrigiano and Nelson, 2004; Lazar et al., 2009). Traditionally, working memory models with attractor dynamics circumvent this problem by keeping weights fixed and fine-tuning a limit set of gain parameters by hand (Brunel and Wang, 2001). Here, we use two distinct homeostatic mechanisms to ensure stability (Figure 1D): synaptic scaling (Turrigiano et al., 1998) and homeostatic threshold regulation (Zhang and Linden, 2003).

First, as synaptic scaling constrains the total drive received by neurons by rescaling all weights in a multiplicative fashion, we implemented this mechanism by an explicit weight normalization, $\sum_j w_{ij} = 1$. We chose this for simplicity, although a similar outcome could in principle be achieved through a local weight-dependent rule (Gerstner and Kistler, 2002). Second, intrinsic plasticity was implemented by assuming that the threshold of excitatory neurons adapts to maintain a certain mean average firing rate, $x_0 \in (0, 1)$:

$$\Delta\Theta = \lambda_{\text{exc}} (x(t) - x_0), \quad (5)$$

where λ_{exc} is the time constant for the threshold adaptation ($x_0 = 0.03$ within the recurrent network and $x_0 = 0.25$ in the decision layer). As mentioned above, the timescale of plasticity for the decision units is 10 times faster to match the more rapid synaptic plasticity ($\lambda_{\text{exc}} = 10^{-4}$ within the recurrent layer, and $\lambda_{\text{exc}} = 10^{-3}$ for the decision units).

We assumed a similar threshold regulation for controlling the excitability of the inhibitory neurons. The specific form was suggested by experimental evidence showing that the excitability of inhibitory neurons is determined by the overall activity of neighboring excitatory neurons, estimated via the release of diffusible messengers, such as BDNF (Rutherford et al., 1998; Turrigiano and Nelson, 2000). Specifically, we assume that the threshold of inhibitory neurons changes as:

$$\Delta\Theta_{\text{inh}} = -\lambda_{\text{inh}} (\langle x_{\text{exc}}(t) \rangle - x_0), \quad (6)$$

with $\langle x_{\text{exc}}(t) \rangle$ denoting the population average of the activation of all excitatory neurons at time t . This is a simplification of a more realistic input-specific regulation of excitability chosen for convenience, consistent with inhibitory neurons pooling activity across a large part of the circuit. As before, x_0 is the desired average firing rate of the excitatory neurons, and λ_{inh} is the learning rate ($\lambda_{\text{inh}} = 10^{-5}$). Although this mechanism is not strictly necessary for network stability, we find it improves memory performance and ensures a fairer distribution of neuronal resources across stimuli.

2.3.3. Other simulation parameters

All trials are assumed to have fixed duration $T_{\text{trial}} = 10$ time steps, with $2 \cdot 10^4$ trials per block. We repeat each experiment five times to quantify effects different sources of variability, such as the network initialization, internal noise, etc.

3. RESULTS

3.1. A DELAYED RESPONSE TASK

The most common experimental paradigm for exploring the circuits involved in working memory is the delayed response task, where a simple stimulus-specific response needs to be delivered after a delay (Rainer et al., 1998a; Durstewitz et al., 2000). Computational models of this function assume a circuit with distinct submodules for storing the initial stimulus (the working memory component), comparing it to the sample and deciding on the action (Engel and Wang, 2011). Here we focus on the first component, and thus assume a one-to-one mapping between stimuli and actions ($M = K$). Although we neglect the intermediate step, i.e., making same-or-different judgements, nonetheless the model preserves the nature of the underlying computation. Hence this simplification should not affect our results concerning the representation within the working memory circuit.

We used two variants for the basic setup: a fixed-and-a variable-delay version (Figure 2, right). As its name suggests, the first uses a fixed delay for all trials in a block. This version is useful for estimating the memory capacity of the network, defined as the longest delay for which performance is better than chance. However, it could potentially lead to unrealistic delay-specific representations. In the second setup, the delay for each trial is selected uniformly at random between one and a maximum delay T_{max} time steps. This version seems closer to the true constraints of the biological system, where information needs to be accessible on demand whenever the environmental conditions call for it. Hence, we used the second version of the task to investigate the emerging neural representations.

We found that the network performance is influenced by task difficulty (Figure 2). As expected, it decreases with increasing delay, due to the accumulation of noise. For intermediate delays, the fixed-delay task yields slightly better results compared to the variable-delay task, consistent with it being computationally simpler. At longer delays however, the network exhibits a sharp performance decay, which signals the network reaching its memory capacity. In the variable delay task, performance degrades more gracefully, as shorter memory spans are still rewarded. In both cases, we found that recall performance increases with network size N , and decays with the number of distinct stimuli K and that the incremental learning paradigm dramatically improves network performance (not shown).

Importantly, performance is remarkably stable within a block of trials despite the constant changes induced by the different plasticity mechanisms, with the network reaching the final performance after a small number of trials (on the order of 100 trials). The critical condition to achieve such good and stable performance is a sparse representation within the recurrent layer (enforced by intrinsic plasticity), combined with balanced rSTDP. While not strictly necessary, synaptic scaling and inhibitory plasticity improve performance; additionally we found it was beneficial to reduce the LTD component for learning in the motor layer (presumably because it limits the interference due to motor activity in the delay period). Overall, the interaction between different plasticity mechanisms is needed for the circuit to maintain stable function despite variable underlying neural “hardware.”

To examine the representation that emerges after learning, we measured both the spatial and the temporal selectivity of neural




FIGURE 2 | Circuit performance for a simple delayed response task. Recall performance as a function of the number of trials for $K = 4$ stimuli. Vertical lines mark the time when the maximum delay is incremented. Within a block the maximum delay is kept fixed, with two

variations: in the fixed delay task (blue) all trials have the same delay, while in the variable delay task (red) the delay of each trial is drawn independently from the uniform distribution $[1, \text{max}]$. Performance estimated across 100 trial blocks.

responses. For the spatial component, we computed the average neural activation during the delay period for each stimulus (Figure 3A, top). This simple measure reveals that most of the neurons respond to one of the stimuli, while remaining relatively silent for the others, as demonstrated in classic working memory experiments (Miyashita, 1988). To better quantify the effect, we used a measure called the depth of selectivity (Rainer et al., 1998b), defined as:

$$S = \frac{N_{\text{cond}} - \frac{\sum R_i}{R_{\text{max}}}}{N_{\text{cond}} - 1}, \quad (7)$$

where R_i is the firing rate corresponding to stimulus i , $R_{\text{max}} = \max\{R_i\}$ and N_{cond} is the number of different behavioral states considered, here the K stimuli. This measure takes the value zero when the neural response is identical for all objects and can reach the maximum of one when the neuron responds exclusively to one of the stimuli. Note that we will use this measure more generally in the following sections, to also measure the specificity to distinct actions or contexts. The depth of selectivity confirmed that most neurons exhibit stimulus-dependent activation (Figure 3A, bottom).

Neural responses are structured also in the temporal domain, reproducing at least qualitatively the temporal specificity in experiments (Meyers et al., 2008). A post-stimulus time histogram (PSTH) of the network responses for a given stimulus reveals that, although before learning the response is highly variable (Figure 3B, top and Figure 3C, left), after learning neuronal responses become highly reproducible (Figure 3B; note that neuron indices were reordered as a result of sorting the neuron by the time of the peak response). Moreover, neurons respond at specific times relative to stimulus onset, pointing to a synfire chain-like representation (Aertsen et al., 1996; Prut et al., 1998). Such temporal dynamics allow neurons to remain stimulus specific, while maintaining a sparse activation enforced thorough the homeostatic regulation of neural excitability. Additionally, the network dynamics reflect the details of the task (Figure 3C): the delay itself is encoded much better during the fixed-delay version of the task. A low-dimensional projection of the neural activity by principal component analysis (PCA) reveals distinct vs. overlapping stimulus-specific clusters in the fixed- and the variable delay task, respectively. This

reflects the intuition that the time since the stimulus presentation is important for the fixed delay task, whereas in the variable delay version the motor layer just needs to linearly separate the activity corresponding to different stimuli, irrespective of the delay. The time-dependent encoding is also reflected in the connectivity matrix, which becomes sparse and more feedforward (Figure 3D). More generally, learning organizes the network in largely non-overlapping feedforward chains, each starting from one of the input sub-populations and with a total size determined by the number of inputs, the size of the network, and the sparseness enforced through the homeostatic mechanisms (not shown). In summary, in a simple delayed-response task, the network uses distributed representations for encoding information about the stimuli across time and space, in a way that makes it easily accessible for decision circuits and is consistent with experiments.

3.2. A DELAYED CATEGORIZATION TASK

Neurons in PFC can encode either the initial stimulus, or the action to be taken in response to it (Brody et al., 2003). For the simple delayed-response task above there is no difference between the two, as actions simply signal stimulus identity. To investigate under which conditions the circuit learns to represent preferentially stimuli or actions, we used a delayed categorization task, inspired by experiments in monkeys, in which arbitrary categories are defined using morphed images (generated from e.g., cat and dog prototypes), see Freedman et al. (2001).

To mimic this paradigm, we constructed an arbitrary map between $K = 8$ stimuli and $M = 2$ decision outputs signaling stimulus class. Here, category boundaries are defined exclusively by the reward function (Freedman et al., 2001; Antzoulatos and Miller, 2011), unlike some experiments in which category specificity may be—to some extent—stimulus driven (Meyers et al., 2008). For illustration purposes, we define the mapping by stimulus color (Figure 4A, right), though in the model the random initialization of the connectivity makes any subdivision of non-overlapping stimuli to be equivalent.

Our network is able to successfully learn the task (75% correct for a delay of five time steps). The neural representations for this task show some novel characteristics compared to the simple delay task, which reflect the experimental data (Freedman et al., 2001). While some of the neurons still respond selectively

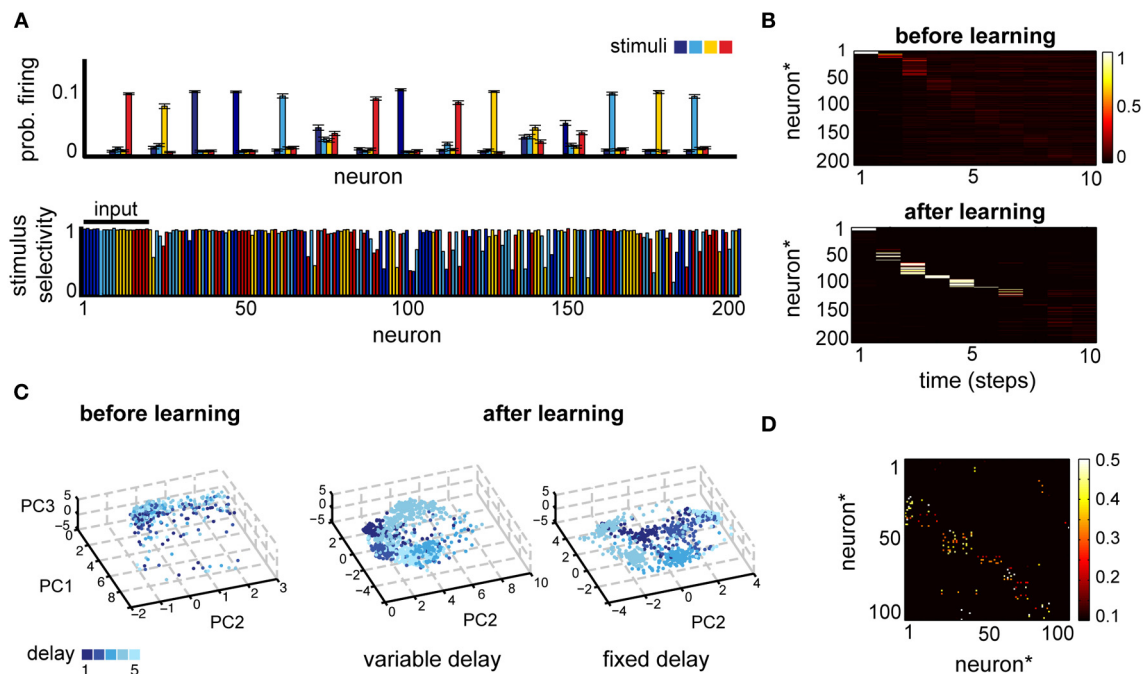


FIGURE 3 | (A) Neural selectivity in a simple delayed response task ($T_{\max} = 5$, variable delay). Top: neural responses averaged across trials where one of four stimuli (different colors) was presented, for a subset of 15 randomly selected neurons. Bottom: selectivity of neural responses across the population for one example experiment; estimated using activity in 1000 trials at the end of learning. Note that the first 20 neurons receive direct inputs from the input layer. **(B)** Comparison of the post-stimulus time histogram of neural responses before and after learning for one example stimulus. Neuron indices

have been reordered based on the time of maximum firing relative to stimulus onset. **(C)** A low-dimensional view of the population dynamics in response to the same stimulus before learning (left) or after training using either the fixed (right) or the variable delay (middle) paradigm. Individual points correspond to the state of the network projected along the first three principal components; color intensity marks the time since the stimulus presentation and different points of the same color correspond to different trials. **(D)** The corresponding weight matrix at the end of learning.




FIGURE 4 | Neural selectivity in a delayed categorization task. **(A)** Average responses to each stimulus in 10 randomly selected example neurons. **(B)** Specificity of neural responses to one of the two categories (red or blue); colors show preferred category for each neuron. Neural selectivity was estimated using activity in 1000 trials at the end of learning. $T_{\max} = 5$, variable delay.

to individual stimuli, a significant subpopulation responds now to several stimuli, and often to those belonging to the same category (Figure 4A). Using the depth of selectivity (with categories rather than stimuli as behaviorally relevant variable) enables us to quantify the category selectivity of neurons across the population (Figure 4B). Using this metric, we found that a significant fraction of the neurons (32% of excitatory neurons have

$S \geq 0.75$) exhibit category selectivity, close to the 33% reported in monkeys (Freedman et al., 2001). As in the previous experiment, their representations are time-varying; at any time, only a small fraction of neurons encode category information, with information being passed between different small subsets of neurons over the course of the trial, as shown in experiments (Meyers et al., 2008). Overall, these results confirm our hypothesis that the differences in neural selectivity in category- vs. stimulus-specific delayed response tasks could emerge due to the task-dependent reorganization of the circuit by reward-dependent learning.

3.3. MULTIPLE CATEGORY BOUNDARIES

Up to now, we have looked at representations in a circuit that specializes on one specific memory task. While this scenario is useful for describing a typical behavioral experiment in monkeys, in real-life conditions the PFC needs to flexibly (and quickly) switch across a variety of different tasks.

How exactly are multiple tasks represented in PFC circuits? The answer should not come as a surprise: “it depends on the tasks.” For tasks involving non-overlapping stimuli, in particular, two independent categorization tasks (cats vs. dogs and sedans vs. sports cars), the activity of many neurons reflects both category distinctions. Thus, the neurons multitask different types of information depending on the context (Cromer et al., 2010). In

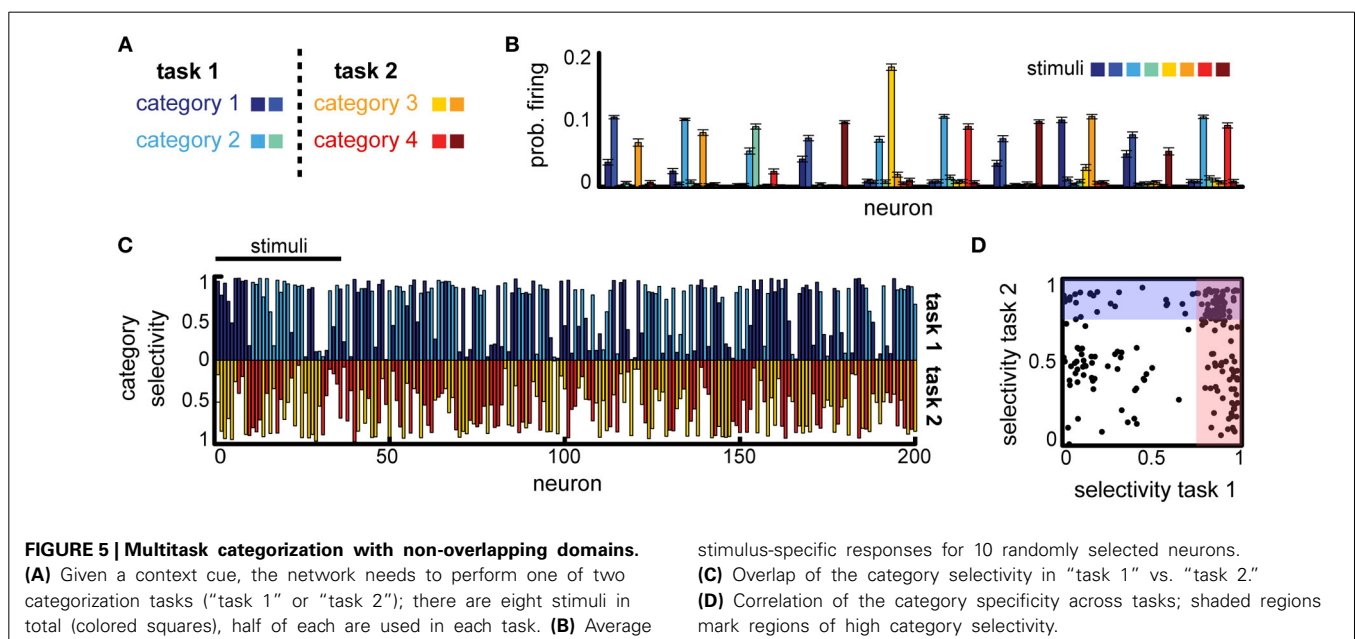
contrast, when the same stimuli need to be categorized differently depending on behavioral context, the two category boundaries are represented by largely non-overlapping neuronal populations (Roy et al., 2010).

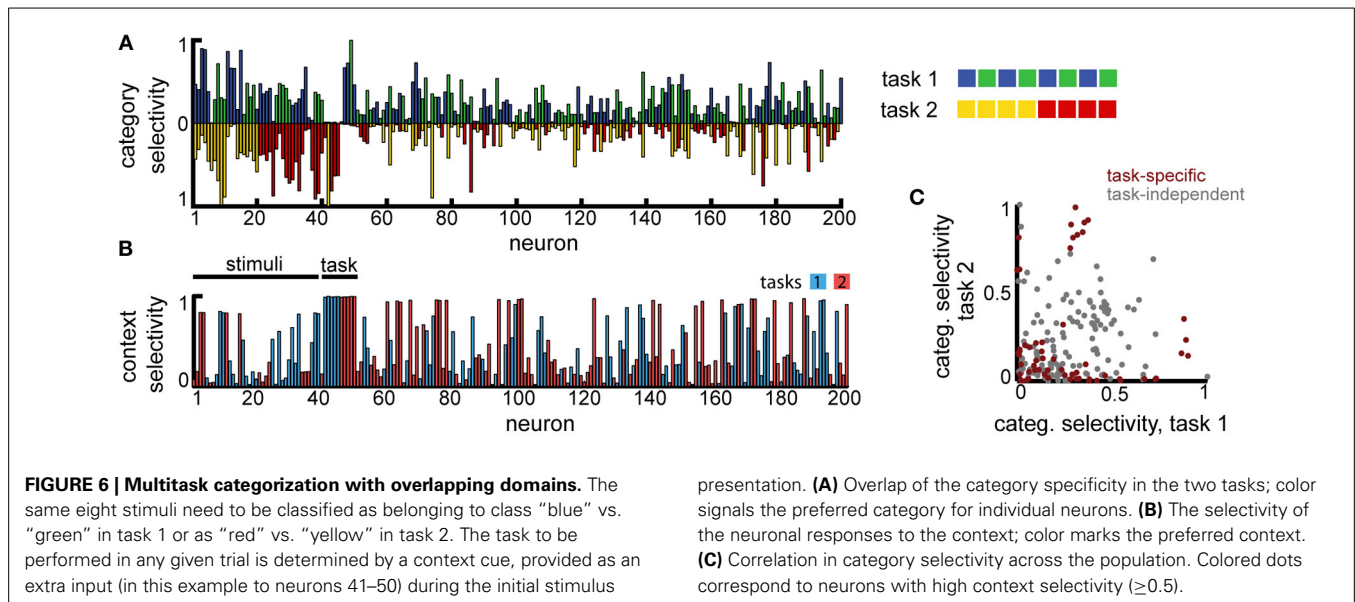
Can a difference in task constraints explain these conflicting results? To answer this question, we constructed two versions of the multi-class delayed categorization task, similar to those used experimentally. First, to implement the multiple independent categories task we used $K = 8$ stimuli and defined two non-overlapping subsets, representing the animals and cars in the original experiment. These subsets were each split in two categories, corresponding to, e.g., the cats vs. dogs distinction, see **Figure 5A**, with $M = 4$ actions, corresponding to the different category distinctions. As in the basic task, the cue signal (now two inputs) was provided directly the decision layer; the cue was active for one time step at the end of the delay period. We found that the network was able to learn this task (average performance 85% for a variable delay task, with maximum delay $T_{\max} = 5$). To assess the emerging neural representations learned for this task, we measured the average firing rate of the neurons in response to different stimuli. We found that many of the neurons responded strongly to several stimuli (**Figure 5B**). These stimuli often belonged to the same class (**Figure 5B**, e.g., for neurons 1, 3, 4, etc.), reproducing the category selectivity we have seen previously, but often neural responses are strong also for stimuli corresponding to different tasks (**Figure 5B**, e.g., the first neuron responds to category 1 and 3). Measuring the category specificity of neurons for each of the two contexts revealed that most neurons are strongly category selective (**Figure 5C**). Moreover, 33.5% of the neurons were sensitive to both category distinctions (selectivity threshold 0.75, see **Figure 5D**). This suggests that, indeed, when the tasks do not interfere with one another the circuit should multiplex information across tasks for good performance.

Second, to model the scenario involving overlapping category boundaries, we assumed $K = 8$ input stimuli that are classified, depending on the context, using two orthogonal category boundaries (**Figure 6A**, right). In this case, the context needs to be provided at the beginning of the trial, together with the stimulus (the context, i.e., which task needs to be performed in the current trial, is encoded as two non-overlapping sub-populations of the same size N_{in} , just as the stimuli). The decision layer consisted, as before, of $M = 4$ neurons, one for each category, and trials from both tasks were interleaved at random.

This version of the multiple categories experiment is significantly harder, as it requires storing information about both stimuli and the current context (because of the two extra inputs, we assumed the recurrent layer has a slightly increased firing rate $x_0 = 0.05$). Still, the network is able to perform significantly above chance (approximatively 60%, for a variable delay with $T_{\max} = 3$). In contrast to the task before, however, fewer neurons develop category specificity (19.5% as opposed to 74.5%), most represent single stimuli and several neurons encode the context itself (**Figure 6B**), suggesting that the network converges to a largely input-driven solution, in which information about stimuli and task is stored separately and combined only at the level of the decision layer. Among the neurons that exhibit category specificity, almost all are selective to only one of the category boundaries (points with high selectivity cluster close to the two axes, and more so if the neural responses are context modulated, see **Figure 6C**), unlike the previous scenario. This observation is reiterated when restricting the analysis to neurons that show task specific encoding (**Figure 6C**, dark red). Thus, the network organizes into separate task-specific subpopulations, as seen experimentally (Roy et al., 2010).

Overall, we found that reward-dependent learning can account for the differences in category representation across experiments. Moreover, the emerging representations showed a strong





task-dependent component, consistent with experimental observations (Asaad et al., 2000; Roy et al., 2010; Warden and Miller, 2010). The general match between our model and experiments is also quantitative, as shown in **Figure 7** (note that the asymmetry between tasks is a consequence of a small number of experiments). This is particularly remarkable given that these results were obtained without any tuning of the model parameters, beyond that required to obtain a good performance in the simple delayed-response task. Taken together, our results suggest task demands dramatically shape neuronal representations in working-memory circuits via reward-dependent learning.

3.4. THE IMPORTANCE OF REWARD-DEPENDENT LEARNING

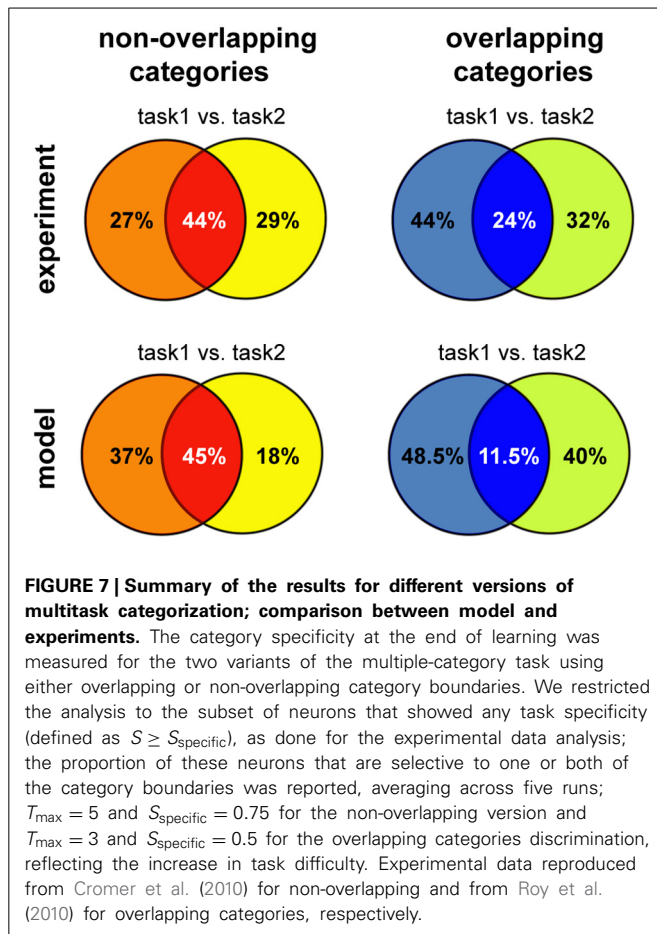
To tease apart the contribution of different plasticity mechanisms to the observed effects, we compared our model to a similarly constructed network, in which weights within the recurrent layer remain fixed, or alternatively are modified by STDP independently of the obtained reward. In both cases, the readout to the decision layer was learned by r-STDP, with all homeostatic mechanisms in place.

We found that learning within the recurrent layer is critical for good memory performance, and in particular that networks with r-STDP are consistently better than those in which recurrent connectivity is fixed (**Figure 8A**). For a simple delayed-response task ($K = 4$ stimuli), reward modulation is not strictly necessary for good performance and unsupervised learning alone can improve neural representations (the performance in unsupervised learning is indistinguishable from that using reward-dependent learning; not shown), as reported elsewhere (Lazar et al., 2009). This result is expected, since when each stimulus defines an action, it is best to represent each input as distinctly as possible, something which can be done by unsupervised learning. Indeed, the emerging representations are similar for the different learning scenarios (stimulus-specific synfire chains; not shown), such that they can be exploited for reward-dependent learning at the decision units.

Importantly, we found that simple unsupervised learning by STDP is no longer sufficient once the task difficulty is increased, by introducing more stimuli and more complex decision boundaries. Indeed, a very different picture emerges when comparing reward dependent vs. unsupervised learning in a categorization task ($K = 8$ stimuli randomly mapped into $M = 2$ categories). In this case, we find that the performance of the two differs significantly (**Figure 8B**). A possible reason for this difference is that attempting to represent each different stimulus separately, via unsupervised learning, exceeds the capacity of this particular network. Because if this, unsupervised learning results in poorer performance in this task. Furthermore we found that the outcome of unsupervised learning is less robust than that of reward-dependent learning: error levels depend on the particular instantiation of the network, leading to increased across trial variability (**Figure 8B**, shaded region in red vs. blue). This dissociation is also apparent at the level of the neural representations. After reward-independent learning, the percentage of category-specific neurons is significantly lower to both our model and the experimental data (20.5% instead of 32% for reward-dependent learning and 33% in the data; see **Figure 8C**). Furthermore, the network responses appear more noisy, suggesting that the number of stimuli exceed the capacity of the network and the reward-independent learning cannot learn a robust representations for all stimuli. All in all, this suggests that for complex tasks, when the pool of available resources is indeed a limiting factor, neuronal representations need to shift toward task-relevant features for good memory performance.

4. DISCUSSION

Prefrontal circuits are shaped by a variety of task-related variables. These representations are likely to form during extensive training prior to experimental recordings, but the mechanisms underlying this development are poorly understood. Here, we have shown that representations similar to those reported experimentally naturally emerge in an initially unstructured circuit



through reward-dependent learning. Moreover, we found that a few generic mechanisms (rSTDP and homeostasis) are sufficient to explain a range of puzzling (and seemingly complex) experimental observations. Neurons in our model developed stimulus and action specificity, both across neurons and in time, as seen experimentally (Miller et al., 1996; Chafee and Goldman-Rakic, 1998; Rainer and Miller, 2002). The same model (with no further parameter tuning) could also account for neural representations during context-dependent tasks. For tasks involving multiple independent category sets, individual neurons multiplexed information across different contexts, matching experimental observations (Cromer et al., 2010); when the same stimuli mapped into different actions depending on the context, neurons specialized to represent single category distinctions, as in Roy et al. (2010). To the best of our knowledge, our model is the first to provide an unified account of these observations.

When comparing our model to a network using reward-independent learning we found reward-dependent plasticity to be critical for solving hard tasks, such as the categorization of many stimuli. This finding is consistent with the notion that reward-dependent learning should be particularly important when resources are limited, either in terms of the amount of information that can be stored (unsupervised learning can be used to store four stimuli for the required time, but not eight), or in terms of the computations allowed for retrieving it (the readout

is linear). In such scenarios, separately representing each stimulus and then mapping the neural activity into the correct output becomes unfeasible (because the resources may not suffice for representing all stimuli individually or because reading out the answer becomes too complicated). Instead, the circuit needs to compute some parts of the map between stimuli and actions during the delay, by clustering together stimuli which should yield the same behavioral response. Given generally recognized resource limitations in working memory circuits (Cowan, 2001), this finding suggests that PFC needs to be malleable, with experience shaping the sensitivity of neurons to reflect current behavior.

Here we chose a very simple model for the network dynamics, known to have small memory capacity (Büsing et al., 2010), because we wanted to focus on the recurrent circuitry and its changes during learning. It should in principle be possible to extend the memory capacity of the network closer to the biologically-relevant range (order of seconds) by using larger networks, a more realistic model of the neural dynamics and including slow time-constants, e.g., NMDA receptors (Durstewitz et al., 2000; Brunel and Wang, 2001) or short-term facilitation (Mongillo et al., 2008). Nonetheless, as the restrictions enforced by resource limitations are likely general, we expect the main features of the representations emerging in the model to be preserved, at least qualitatively, in a detailed circuit. Thus, we predict reward-dependent learning should play a general role in the formation and task-specific tuning of working memory circuits.

From a developmental perspective, it is tempting to hypothesize that reward-dependent learning may play a role in the age-dependent improvement of working memory (estimated to be approximately four-fold between the ages of 4 and 14) (Luciana and Nelson, 1998), complementing other known factors such as the maturation of the underlying cortical architecture, a better representation of the inputs, the development of attention, or the usage of memorization strategies such as rehearsal and chunking (Gathercole, 1999). This suggestion is consistent with the known dependence of PFC function on dopamine in early life (Diamond and Baddeley, 1996). Furthermore, the same mechanisms may account for training-induced improvements in working memory in adults (Klingberg, 2010).

From a broader computational perspective, our work is also relevant in the context of reservoir computing (Lukoševičius and Jaeger, 2009). While this framework traditionally assumes fixed recurrent connectivity, recent work has increasingly argued for the importance of learning in shaping reservoir properties (Schmidhuber et al., 2007; Haeusler et al., 2009; Lazar et al., 2009). Previous work used general-purpose optimization through unsupervised learning. Here, however, the network is shaped directly by the task, which improves performance significantly compared to static networks or networks shaped by reward-independent learning. Thus, our model provides a stepping stone toward general task-specific optimization of recurrent networks.

Time-dependent representations are preferred to traditional attractor-based solutions (Amit and Brunel, 1997; Brunel and Wang, 2001; Mongillo et al., 2008) in our model, consistent with recent experimental observations (Miller et al., 1996; Chafee and Goldman-Rakic, 1998; Pesaran et al., 2002; Rainer and Miller, 2002; Barak et al., 2010) and previous theoretical predictions

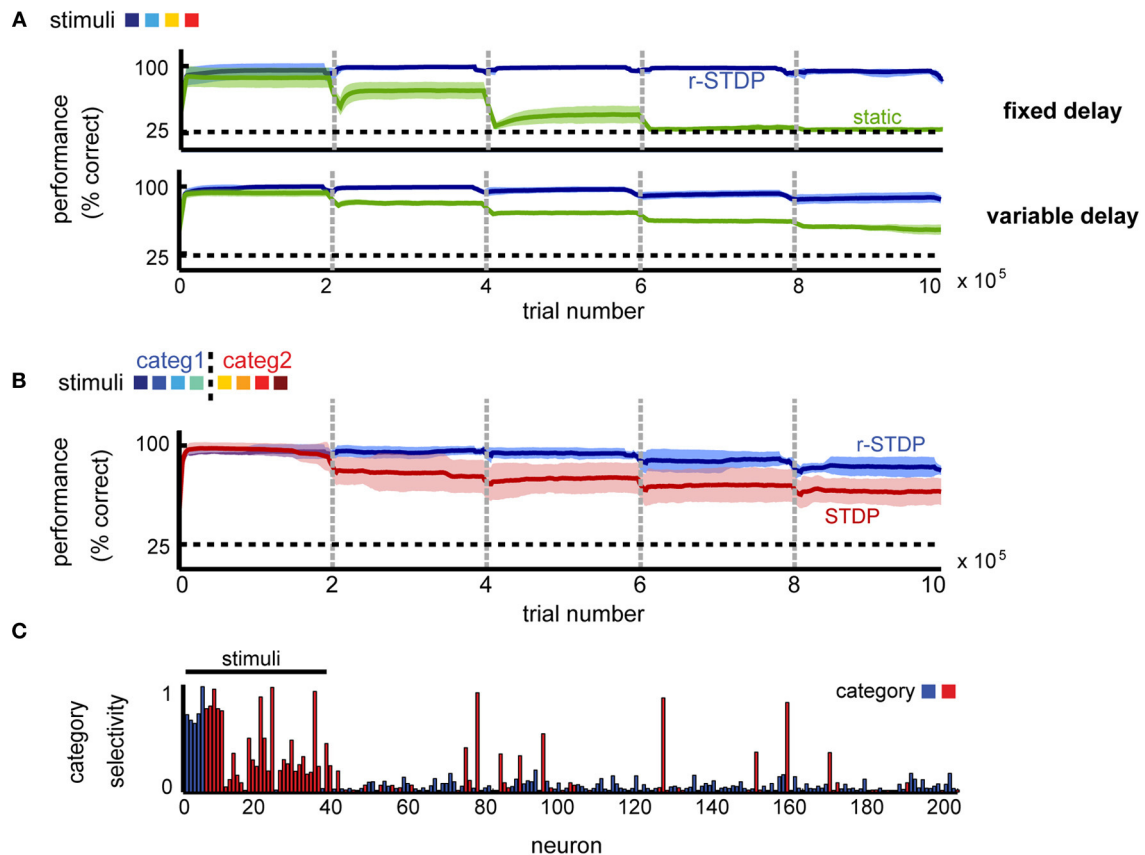


FIGURE 8 | The importance of reward-dependent learning for shaping circuit dynamics. (A) Performance comparison of networks shaped by r-STDP (blue) vs. circuits where the recurrent connectivity is static (green) in a simple delayed response task. **(B)** Performance comparison of networks shaped by r-STDP (blue) vs. circuits where the recurrent connectivity is shaped by reward-independent STDP (red) in a 2-class

categorization task; $K = 8$, fixed delay. For all conditions, the learning of the decision output was reward-modulated. Dark colors mark averages across five repetitions; light colors show the standard deviation around this mean. **(C)** Neuronal selectivity to stimulus category when learning in the recurrent circuit is reward independent; $K = 8$, fixed delay (compare to Figure 3D).

(Goldman, 2009). This effect is a consequence of intrinsic plasticity, which discourages neurons from remaining active for a long time (Horn and Usher, 1989). Given that homeostasis plays a critical role in stabilizing the circuit dynamics during learning (Turrigiano and Nelson, 2004), the fact that the emerging representation is time-varying is not really surprising. While our model emphasizes the temporal component of this representation, it is likely that the patterns of activity seen experimentally emerge through the interaction between feedforward and feedback dynamics, which would require a more detailed model of the neural dynamics. Although the homeostatic mechanisms acting in PFC circuits have yet to be characterized experimentally, it is tempting to assume that the sparsification of activity and increase in robustness observed experimentally after training (Rainer and Miller, 2000) may be signatures of the interaction between hebbian and homeostatic plasticity as shown in our model. More generally, similar mechanisms could play a role in developing feedforward dynamics in other recurrent circuits (see also Levy et al., 2001; Buonomano, 2005; Gilson et al., 2009; Fiete et al., 2010), for instance in other areas known to exhibit delay period responses, such as the perirhinal cortex, inferotemporal cortex,

or the hippocampus (Miller et al., 1993; Quintana and Fuster, 1999).

Our model combines both hebbian (r-STDP) and homeostatic (intrinsic plasticity, synaptic scaling) forms of plasticity, lending further support to the notion that the interaction between different forms of plasticity is critical for circuit computation (Triesch, 2007; Lazar et al., 2009; Savin et al., 2010). In particular, our results confirm the computational importance of intrinsic plasticity and synaptic scaling in excitatory neurons (Savin et al., 2010; Keck et al., 2012). To this, we add the role of inhibitory plasticity, which we found improved both neural representations and memory performance.

We view this model as a starting point for investigating reward-dependent learning in working memory circuits, to which many additions can be made. While the abstract network model used here allowed us to focus on the essential mechanisms underlying PFC coding, it would be important to investigate reward-dependent learning in more realistic spiking neural networks. Furthermore, the model for different plasticity mechanisms operating in the network could be refined as well. First, reward-dependent learning could be improved by using recent

extensions of r-STDP to spiking neuron populations (Urbanczik and Senn, 2009). Second, the simplistic regulation of inhibition should be replaced by realistic inhibitory plasticity (Castillo et al., 2011), which is expected to also aid network selectivity (Vogels et al., 2011). Third, activity-dependent structural plasticity could optimize the cortical connectivity to best encode the task-specific information (Savin and Triesch, 2010; Bourjaily and Miller, 2011), consistent with experimental observations that working memory training alters circuit connectivity (Takeuchi et al., 2010). Lastly, preliminary work, supported by recent observations about the effects of neuromodulation on inhibitory and homeostatic plasticity (Seamans et al., 2001; Di Pietro and Seamans, 2011), suggests that the homeostatic plasticity mechanisms themselves may be reward-dependent.

AUTHOR CONTRIBUTIONS

Designed the experiments: Cristina Savin and Jochen Triesch; Implemented the model and analyzed the data: Cristina Savin; Wrote the paper: Cristina Savin and Jochen Triesch.

ACKNOWLEDGMENTS

Supported in part by EC MEXT project PLICON and the LOEWE-Program “Neuronal Coordination Research Focus Frankfurt” (NeFF). Jochen Triesch was supported by the Quandt foundation.

REFERENCES

- Aertsen, A., Diesmann, M., and Gewaltig, M. O. (1996). Propagation of synchronous spiking activity in feedforward neural networks. *J. Physiol. Paris* 90, 243–247. doi: 10.1016/S0928-4257(97)81432-5
- Amit, D., and Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex* 7, 237–252. doi: 10.1093/cercor/7.3.237
- Antzoulatos, E. G., and Miller, E. K. (2011). Differences between neural activity in prefrontal cortex and striatum during learning of novel abstract categories. *Neuron* 71, 243–249. doi: 10.1016/j.neuron.2011.05.040
- Asaad, W. F., Rainer, G., and Miller, E. K. (1998). Neural activity in the primate prefrontal cortex during associative learning. *Neuron* 21, 1399–1407. doi: 10.1016/S0896-6273(00)80658-3
- Asaad, W. F., Rainer, G., and Miller, E. K. (2000). Task-specific neural activity in the primate prefrontal cortex. *J. Neurophysiol.* 84, 451–459.
- Baeg, E., Kim, Y., Huh, K., Mook-Jung, I., Kim, H., and Jung, M. (2003). Dynamics of population code for working memory in the prefrontal cortex. *Neuron* 40, 177–188. doi: 10.1016/S0896-6273(03)00597-X
- Baeg, E., Kim, Y., Kim, J., Ghim, J., Kim, J., and Jung, M. (2007). Learning-induced enduring changes in functional connectivity among prefrontal cortical neurons. *J. Neurosci.* 27, 909–918. doi: 10.1523/JNEUROSCI.4759-06.2007
- Barak, O., Tsodyks, M., and Romo, R. (2010). Neuronal population coding of parametric working memory. *J. Neurosci.* 30, 9424–9430. doi: 10.1523/JNEUROSCI.1875-10.2010
- Bourjaily, M. A., and Miller, P. (2011). Excitatory, inhibitory, and structural plasticity produce correlated connectivity in random networks trained to solve paired-stimulus tasks. *Front. Comput. Neurosci.* 5:37. doi: 10.3389/fncom.2011.00037
- Brody, C., Romo, R., and Kepecs, K. (2003). Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. *Curr. Opin. Neurobiol.* 13, 204–211. doi: 10.1016/S0959-4388(03)00050-3
- Brunel, N., and Wang, X. (2001). Effects of neuromodulation in a cortical model of object working memory dominated by recurrent inhibition. *J. Comput. Neurosci.* 11, 63–85. doi: 10.1023/A:1011204814320
- Buonomano, D. V. (2005). A learning rule for the emergence of stable dynamics and timing in recurrent networks. *J. Neurophysiol.* 94, 2275–2283. doi: 10.1152/jn.01250.2004
- Büsing, L., Schrauwen, B., and Legenstein, R. (2010). Connectivity, dynamics, and memory in reservoir computing with binary and analog neurons. *Neural Comput.* 22, 1272–1311. doi: 10.1162/neco.2009.01-09-947
- Castillo, P. E., Chiu, C. Q., and Carroll, R. C. (2011). Long-term plasticity at inhibitory synapses. *Curr. Opin. Neurobiol.* 21, 328–338. doi: 10.1016/j.conb.2011.01.006
- Chafee, M., and Goldman-Rakic, P. (1998). Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *J. Neurophysiol.* 79, 2919–2940.
- Constantinidis, C., and Franowicz, M. (2001). The sensory nature of mnemonic representation in the primate prefrontal cortex. *Nature* 4, 311–316. doi: 10.1038/85179
- Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav. Brain Sci.* 24, 87–114. doi: 10.1017/S0140525X01003922
- Cromer, J. A., Roy, J. E., and Miller, E. K. (2010). Representation of multiple, independent categories in the primate prefrontal cortex. *Neuron* 66, 796–807. doi: 10.1016/j.neuron.2010.05.005
- Di Pietro, N. C., and Seamans, J. K. (2011). Dopamine and serotonin interactively modulate prefrontal cortex neurons *in vitro*. *Biol. Psychiatry* 69, 1204–1211. doi: 10.1016/j.biopsych.2010.08.007
- Diamond, A., and Baddeley, A. (1996). Evidence for the importance of dopamine for prefrontal cortex functions early in life. *Philos. Trans. R. Soc. B Biol. Sci.* 351, 1483–1494. doi: 10.1098/rstb.1996.0134
- Duncan, J. (2001). An adaptive coding model of neural function in prefrontal cortex. *Nat. Rev. Neurosci.* 2, 820–829. doi: 10.1038/35097575
- Durstewitz, D., Seamans, J., and Sejnowski, T. (2000). Neurocomputational models of working memory. *Nat. Neurosci.* 3, 1184–1191. doi: 10.1038/81460
- Engel, T. A., and Wang, X.-J. (2011). Same or different? A neural circuit mechanism of similarity-based pattern match decision making. *J. Neurosci.* 31, 6982–6996. doi: 10.1523/JNEUROSCI.6150-10.2011
- Fiete, I. R., Senn, W., Wang, C. Z. H., and Hahnloser, R. H. R. (2010). Spike-Time-Dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron* 65, 563–576. doi: 10.1016/j.neuron.2010.02.003
- Freedman, D. J., Riesenhuber, M., Poggio, T., and Miller, E. K. (2001). Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 291, 312–316. doi: 10.1126/science.291.5502.312
- Funahashi, S., Bruce, C. J., and Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.* 61, 331–349.
- Gathercole, S. (1999). Cognitive approaches to the development of short-term memory. *Trends Cogn. Sci.* 3, 410–419. doi: 10.1016/S1364-6613(99)01388-1
- Gerstner, W., and Kistler, W. (2002). Mathematical formulations of Hebbian learning. *Biol. Cybern.* 87, 404–415. doi: 10.1007/s00422-002-0353-y
- Gilson, M., Burkitt, A. N., Grayden, D. B., Thomas, D. A., and Hemmen, J. L. (2009). Emergence of network structure due to spike-timing-dependent plasticity in recurrent neuronal networks IV. *Biol. Cybern.* 101, 427–444. doi: 10.1007/s00422-009-0343-4
- Goldman, M. S. (2009). Memory without feedback in a neural network. *Neuron* 61, 621–634. doi: 10.1016/j.neuron.2008.12.012
- Goldman-Rakic, P. S. (1990). Cellular and circuit basis of working memory in prefrontal cortex of nonhuman primates. *Prog. Brain Res.* 85, 325–335. doi: 10.1016/S0079-6123(08)62688-6
- Haeusler, S., Schuch, K., and Maass, W. (2009). Motif distribution and computational performance of two data-based cortical microcircuit templates. *J. Physiol. Paris* 103, 73–87. doi: 10.1016/j.jphysparis.2009.05.006
- Horn, D., and Usher, M. (1989). Neural networks with dynamical thresholds. *Phys. Rev. A* 40, 1036–1044. doi: 10.1103/PhysRevA.40.1036
- Izhikevich, E. (2007). Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cereb. Cortex* 17, 2443–2452. doi: 10.1093/cercor/bhl152
- Keck, C., Savin, C., and Lücke, J. (2012). Feedforward inhibition and synaptic scaling – two sides of the same coin? *PLoS Comput. Biol.* 8:e1002432. doi: 10.1371/journal.pcbi.1002432
- Kennerley, S., and Wallis, J. (2009). Reward-dependent modulation of working memory in lateral prefrontal cortex. *J. Neurosci.* 29, 3259–3270. doi: 10.1523/JNEUROSCI.5353-08.2009

- Klingberg, T. (2010). Training and plasticity of working memory. *Trends Cogn. Sci.* 14, 317–324. doi: 10.1016/j.tics.2010.05.002
- Lazar, A., Pipa, G., and Triesch, J. (2009). Sorn: a self-organizing recurrent neural network. *Front. Comput. Neurosci.* 3:23. doi: 10.3389/fncom.2009.0023
- Levy, N., Horn, D., Meilijson, I., and Rupp, E. (2001). Distributed synchrony in a cell assembly of spiking neurons. *Neural Netw.* 14, 815–824. doi: 10.1016/S0893-6080(01)00044-2
- Luciana, M., and Nelson, C. (1998). The functional emergence of prefrontally-guided working memory systems in four- to eight-year-old children. *Neuropsychologia* 36, 273–293. doi: 10.1016/S0028-3932(97)00109-7
- Lukoševičius, M., and Jaeger, H. (2009). Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.* 3, 127–149. doi: 10.1016/j.cosrev.2009.03.005
- Meyer, T., Qi, X.-L., Stanford, T. R., and Constantinidis, C. (2011). Stimulus selectivity in dorsal and ventral prefrontal cortex after training in working memory tasks. *J. Neurosci.* 31, 6266–6276. doi: 10.1523/JNEUROSCI.6798-10.2011
- Meyers, E. M., Freedman, D. J., Kreiman, G., Miller, E. K., and Poggio, T. (2008). Dynamic population coding of category information in inferior temporal and prefrontal cortex. *J. Neurophysiol.* 100, 1407–1419. doi: 10.1152/jn.90248.2008
- Miller, E., Erickson, C., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J. Neurosci.* 16, 5154–5167.
- Miller, E., Li, L., and Desimone, R. (1993). Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J. Neurosci.* 13, 1460–1478.
- Miyashita, Y. (1988). Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature* 335, 817–820. doi: 10.1038/335817a0
- Mongillo, G., Barak, O., and Tsodyks, M. (2008). Synaptic theory of working memory. *Science* 319, 1543–1546. doi: 10.1126/science.1150769
- Otani, S., Daniel, H., Roisin, M.-P., and Crepel, F. (2003). Dopaminergic modulation of long-term synaptic plasticity in rat prefrontal neurons. *Cereb. Cortex* 13, 1251–1256. doi: 10.1093/cercor/bhg092
- Pesaran, B., Pezaris, J., Sahani, M., Mitra, P., and Andersen, R. (2002). Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat. Neurosci.* 5, 805–811. doi: 10.1038/nn890
- Prut, Y., Vaadia, E., Bergman, H., Haalman, I., Slovlin, H., and Abeles, M. (1998). Spatiotemporal structure of cortical activity: properties and behavioral relevance. *J. Neurophysiol.* 79, 2857–2874.
- Quintana, J., and Fuster, J. (1999). From perception to action: temporal integrative functions of prefrontal and parietal neurons. *Cereb. Cortex* 9, 213–221. doi: 10.1093/cercor/9.3.213
- Rainer, G., Asaad, W. F., and Miller, E. K. (1998a). Memory fields of neurons in the primate prefrontal cortex. *Proc. Natl. Acad. Sci. U.S.A.* 95, 15008–15013. doi: 10.1073/pnas.95.25.15008
- Rainer, G., Asaad, W. F., and Miller, E. K. (1998b). Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature* 393, 577–579. doi: 10.1038/31235
- Rainer, G., and Miller, E. (2000). Effects of visual experience on the representation of objects in the prefrontal cortex. *Neuron* 27, 179–189. doi: 10.1016/S0896-6273(00)00019-2
- Rainer, G., and Miller, E. (2002). Timecourse of object-related neural activity in the primate prefrontal cortex during a short-term memory task. *Eur. J. Neurosci.* 15, 1244–1254. doi: 10.1046/j.1460-9568.2002.01958.x
- Reynolds, J., and Wickens, J. (2002). Dopamine-dependent plasticity of corticostriatal synapses. *Neural Netw.* 15, 507–521. doi: 10.1016/S0893-6080(02)00045-X
- Roy, J. E., Riesenhuber, M., Poggio, T., and Miller, E. K. (2010). Prefrontal cortex activity during flexible categorization. *J. Neurosci.* 30, 8519–8528. doi: 10.1523/JNEUROSCI.4837-09.2010
- Rutherford, L. C., Nelson, S. B., and Turrigiano, G. G. (1998). BDNF has opposite effects on the quantal amplitude of pyramidal neuron and interneuron excitatory synapses. *Neuron* 21, 521–530. doi: 10.1016/S0896-6273(00)80563-2
- Savin, C., Joshi, P., and Triesch, J. (2010). Independent component analysis in spiking neurons. *PLoS Comput. Biol.* 6:e1000757. doi: 10.1371/journal.pcbi.1000757
- Savin, C., and Triesch, J. (2010). “Structural plasticity improves stimulus encoding in a working memory model,” in *Frontiers in Computational Neuroscience. Conference Abstract: Computational and systems neuroscience* (Salt Lake City, UT).
- Schmidhuber, J., Wierstra, D., Gagliolo, M., and Gomez, F. (2007). Training recurrent networks by Evolino. *Neural Comput.* 19, 757–779. doi: 10.1162/neco.2007.19.3.757
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *J. Neurophysiol.* 80, 1–27.
- Seamans, J., Gorelova, N., Durstewitz, D., and Yang, C. (2001). Bidirectional dopamine modulation of GABAergic inhibition in prefrontal cortical pyramidal neurons. *J. Neurosci.* 21, 3628–3638.
- Takeuchi, H., Sekiguchi, A., and Taki, Y. (2010). Training of working memory impacts structural connectivity. *J. Neurosci.* 30, 3297–3303. doi: 10.1523/JNEUROSCI.4611-09.2010
- Triesch, J. (2007). Synergies between intrinsic and synaptic plasticity mechanisms. *Neural Comput.* 19, 885–909. doi: 10.1162/neco.2007.19.4.885
- Turrigiano, G., Leslie, K., Desai, N., Rutherford, L., and Nelson, S. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391, 892–896. doi: 10.1038/36103
- Turrigiano, G., and Nelson, S. (2000). Hebb and homeostasis in neuronal plasticity. *Curr. Opin. Neurobiol.* 10, 358–364. doi: 10.1016/S0959-4388(00)00091-X
- Turrigiano, G., and Nelson, S. (2004). Homeostatic plasticity in the developing nervous system. *Nat. Rev. Neurosci.* 5, 97–107. doi: 10.1038/nrn1327
- Urbanczik, R., and Senn, W. (2009). Reinforcement learning in populations of spiking neurons. *Nat. Neurosci.* 12, 250–252. doi: 10.1038/nn.2264
- Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., and Gerstner, W. (2011). Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* 334, 1569–1573. doi: 10.1126/science.1211095
- Warden, M. R., and Miller, E. K. (2010). Task-dependent changes in short-term memory in the prefrontal cortex. *J. Neurosci.* 30, 15801–15810. doi: 10.1523/JNEUROSCI.1569-10.2010
- Zhang, W., and Linden, D. (2003). The other side of the engram: experience-dependent changes in neuronal intrinsic excitability. *Nat. Rev. Neurosci.* 4, 885–900. doi: 10.1038/nrn1248

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 26 March 2014; accepted: 10 May 2014; published online: 28 May 2014.

Citation: Savin C and Triesch J (2014) Emergence of task-dependent representations in working memory circuits. *Front. Comput. Neurosci.* 8:57. doi: 10.3389/fncom.2014.00057

This article was submitted to the journal *Frontiers in Computational Neuroscience*.

Copyright © 2014 Savin and Triesch. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

ADVANTAGES OF PUBLISHING IN FRONTIERS



FAST PUBLICATION

Average 90 days
from submission
to publication



COLLABORATIVE PEER-REVIEW

Designed to be rigorous –
yet also collaborative, fair and
constructive



RESEARCH NETWORK

Our network
increases readership
for your article



OPEN ACCESS

Articles are free to read,
for greatest visibility



TRANSPARENT

Editors and reviewers
acknowledged by name
on published articles



GLOBAL SPREAD

Six million monthly
page views worldwide



COPYRIGHT TO AUTHORS

No limit to
article distribution
and re-use



IMPACT METRICS

Advanced metrics
track your
article's impact



SUPPORT

By our Swiss-based
editorial team