ANATOMY AND PLASTICITY IN LARGE-SCALE BRAIN MODELS

EDITED BY : Markus Butz, Wolfram Schenck and Arjen van Ooyen

frontiers Research Topics

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**
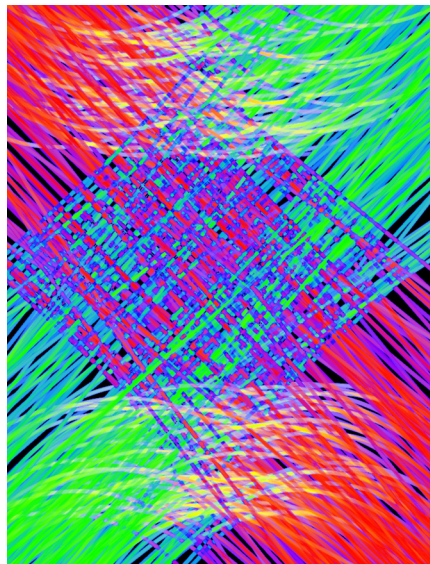
ANATOMY AND PLASTICITY IN LARGE-SCALE BRAIN MODELS

Topic Editors:
Markus Butz, Jülich Research Center, Germany
Wolfram Schenck, Bielefeld University of Applied Sciences, Germany
Arjen van Ooyen, VU University Amsterdam, Netherlands

3D-PLI based fiber orientation map of a section through simulated fibers modeling the optic chiasm of a hooded seal (Axer et al., this E-book)

Supercomputing facilities are becoming increasingly available for simulating activity dynamics in large-scale neuronal networks. On today's most advanced supercomputers, networks with up to a billion of neurons can be readily simulated. However, building biologically realistic, full-scale brain models requires more than just a huge number of neurons. In addition to network size, the detailed local and global anatomy of neuronal connections is of crucial importance. Moreover, anatomical connectivity is not fixed, but can rewire throughout life (structural plasticity)—an aspect that is missing in most current network models, in which plasticity is confined to changes in synaptic strength (synaptic plasticity).

The papers in this Ebook, which may broadly be divided into three themes, aim to bring together high-performance computing with recent experimental and computational research in neuroanatomy. In the first theme (fiber connectivity), new methods are described for measuring and data-basing microscopic and macroscopic connectivity. In the second theme (structural plasticity), novel models are introduced that incorporate morphological plasticity and rewiring of anatomical connections. In the third theme (large-scale simulations), simulations of large-scale neuronal networks are presented with an emphasis on anatomical detail and plasticity mechanisms. Together, the articles in this Ebook make the reader aware of the methods and models by which large-scale brain networks running on supercomputers can be extended to include anatomical detail and plasticity.

Citation: Butz, M., Schenck, W., van Ooyen, A., eds. (2016). Anatomy and Plasticity in Large-Scale Brain Models. Lausanne: Frontiers Media. doi: 10.3389/978-2-88945-065-7

Table of Contents

frontiers
in Neuroanatomy

Editorial: Anatomy and Plasticity in Large-Scale Brain Models

Markus Butz[1], Wolfram Schenck[2] and Arjen van Ooyen[3]*

[1] Simulation Laboratory Neuroscience, Bernstein Facility for Simulation and Database Technology, Institute for Advanced Simulation, Jülich Aachen Research Alliance, Jülich Research Center, Jülich, Germany, [2] Faculty of Engineering and Mathematics, Bielefeld University of Applied Sciences, Bielefeld, Germany, [3] Department of Integrative Neurophysiology, VU University Amsterdam, Amsterdam, Netherlands

The Editorial on the Research Topic

Anatomy and Plasticity in Large-Scale Brain Models

INTRODUCTION

OPEN ACCESS

Edited and reviewed by:
Javier DeFelipe,
Cajal Institute, Spain

*Correspondence:
Arjen van Ooyen
arjen.van.ooyen@gmail.com

Supercomputing facilities are becoming increasingly available for simulating electrical activity in large-scale neuronal networks. On today's most advanced supercomputers, networks with up to a billion of neurons can be readily simulated. However, building biologically realistic, full-scale brain models requires more than just a huge number of neurons. In addition to network size, the detailed local and global anatomy of neuronal connections is of crucial importance. Moreover, anatomical connectivity is not fixed, but can rewire throughout life (structural plasticity; Butz et al., 2009)—an aspect that is missing in most current network models, in which plasticity is confined to changes in synaptic strength (synaptic plasticity).

The papers in this research topic, which may broadly be divided into three themes, aim to bring together high-performance computing with recent experimental and computational research in neuroanatomy. In the first theme (fiber connectivity), new methods are described for measuring and data-basing microscopic and macroscopic connectivity. In the second theme (structural plasticity), novel models are introduced that incorporate morphological plasticity and rewiring of anatomical connections. In the third theme (large-scale simulations), simulations of large-scale neuronal networks are presented with an emphasis on anatomical detail and plasticity mechanisms. Together, the papers in this research topic contribute to extending high-performance computing in neuroscience to encompass anatomical detail and plasticity.

FIBER CONNECTIVITY

Investigating the brain's connectivity requires multiscale approaches and hence strategies for integrating data across different spatial scales. Axer et al. demonstrate how to bridge microscopic visualizations of fibers obtained by 3D-PLI (polarized light imaging; Axer et al., 2011) to meso- or macro-scopic fiber orientations based on dMRI (diffusion magnetic resonance imaging). A relatively new technique, 3D-PLI is applicable to microtome sections of postmortem brains and uses birefringence of brain tissue, induced by optical anisotropy of the myelin sheaths around axons, to derive a 3D description of the underlying fiber architecture. To be able to link 3D-PLI to dMRI measurements, the authors introduce fiber orientation distribution functions (ODFs) extracted from 3D-PLI. They demonstrate the validity of their approach with simulated 3D-PLI data as well as real 3D-PLI data from the human brain and the brain of a hooded seal.

Capturing different aspects of brain organization, such as connectivity and molecular composition, necessitates the use of different neuroimaging techniques. To subsequently integrate the multiscale and multimodal data into a complete 3D brain model requires an accurate definition of the spatial positions of structural entities. Defined by MRI, the Waxholm Space (WHS) (http://software.incf.org/software/waxholm-space) provides such a reference space for rodent brain data. The aim of the study by Schubert et al. was to extend the WHS rat brain atlas with information about cytoarchitecture, receptor expression and spatial orientation of fiber tracts, derived from autoradiography and PLI images. To incorporate these distinct classes of information into the WHS, the authors improved currently available registration algorithms to align sections and to correct for deformations. The extended WHS rat brain atlas now enables combined studies on receptor and cell distributions as well as fiber densities in the same anatomical structures at microscopic scales. Furthermore, the methods developed facilitate future integration of data of other modalities.

STRUCTURAL PLASTICITY

According to the long-standing connectionists' dogma, information is stored in the connection weights of neural networks. However, in biological neural networks, a synaptic connection is much more than just a weight factor, and a wealth of biological mechanisms can bring about changes in functional and structural connectivity. The following papers deal with models that go beyond the traditional modeling concept of plasticity as merely up- or down-regulating synaptic connection strength (synaptic plasticity).

Fauth and Tetzlaff review experimental and modeling studies that address the formation and deletion of synapses (structural plasticity) and how this is regulated by electrical activity. Whereas adapting existing synapses in a Hebbian manner predominantly serves memory consolidation, the *de novo* formation of synapses and deletion of existing synapses can, together with other synaptic mechanisms, grant stability to a network facing continuously changing inputs. The authors point out that these different plasticity mechanisms may be involved in different neural functions and, remarkably, may respond very differently to changes in neuronal activity.

Changes in the anatomical layout of connections (structural plasticity) are crucially dependent on morphological changes in individual neurons. A lack of understanding of how structural changes affect network function is partly due to the absence of tools for studying the impact of neuronal morphology on neuronal function. Therefore, Bezelos et al. introduce a simulation tool for systematically varying the morphology of any type of neuron, which will help investigate the role of neuronal morphology and morphological changes in large-scale neuronal networks.

Diaz-Pier et al. go beyond the level of single neurons and describe an approach based on a recent model of homeostatic structural plasticity (Butz and van Ooyen, 2013), which for the first time enables growing from scratch the neuronal network connectivity of a cortical column *in silico*. The resulting connectivity shows remarkable similarities with a real cortical column.

In biologically realistic, sparsely connected cortical networks, Knoblauch and Sommer study what the computational contribution of structural plasticity is to memory formation based on synaptic plasticity. They show that neuronal networks with structural plasticity can more efficiently adapt their connectivity to the computational problem at hand. As a consequence, structural plasticity may significantly increase the number of stably maintained memory items and may even account for psychological phenomena such as the spacing effect in rehearsal learning (Greene, 1989).

LARGE-SCALE SIMULATIONS

The third theme concerns large-scale simulations, with a special focus on anatomic detail and/or plasticity mechanisms. Because of the high computational demands of such simulations, it is crucial to have the right hardware and software tools available. Breit et al., Gosui and Yamazaki, and Knight et al. report on recent developments in this area. Furthermore, Zhou et al. show how to successfully employ established neuro-simulators for this purpose, and Kozloski how to move up to full brain-scale models.

Breit et al. developed the novel simulation framework NeuroBox and applied it to the investigation of the interplay between synapse loss and signaling synchrony. NeuroBox provides a link between the description of the morphology of neurons and networks on the one hand and the simulation of such networks in the simulation framework UG4 (Vogel et al., 2013) on the other hand. This combination enables the simulation of anatomically detailed models of large networks on supercomputers with good scaling behavior.

Gosui and Yamazaki tackle the problem of how to carry out long-term simulations within a manageable time frame. They studied the long-term gain adaptation of optokinetic response eye movements over a period of 5 days of simulated time in a large-scale cerebellar model. By implementing the simulation software on highly parallel processors (graphics processing units; GPUs), the simulation could be carried out in real-world time. This approach opens up new opportunities for studying long-term plasticity in neural networks, e.g., in the context of memory consolidation.

Knight et al. demonstrate that large-scale models that incorporate complex and demanding plasticity mechanisms (in this case, based on the Bayesian Confidence Propagation Neural Network paradigm; Lansner and Holst, 1996) can be simulated in a highly power-efficient way on the SpiNNaker neuromorphic architecture. Furthermore, in this study the largest plastic neural network ever was simulated on neuromorphic hardware. This shows that (some) neuromorphic hardware is already up to the task of simulating detailed anatomy and plasticity in large-scale brain models.

Zhou et al. rely in their work on the established neuro-simulator NEURON (Hines and Carnevale, 2001). They implemented a large-scale model of the olfactory bulb that

includes plasticity mechanisms. They investigated how prior odor experience influences the representation of new odor inputs. The results show that prior experience changes the pattern in which sparse responses occur in different sub-networks within the olfactory bulb. From a methodological point of view, this study demonstrates how to set up a large-scale simulation of a plastic neural network with rather detailed anatomy.

Kozloski positions his modeling and simulation work at brain-scale, comprising the neocortex, basal ganglia and thalamus. The proposed model is based on the principle of information-based exchange and does not cover anatomical details. Nevertheless, it is grounded in neuroanatomical observations and accounts for various forms of plasticity. Modeling at this rather abstract level enables long-term simulations at brain-scale and extensive parameter variations. In this specific study, variations in dynamic set points and modulations were investigated, leading to a theory about the emergence of neurodegenerative diseases.

CONCLUSION

Biologically realistic large-scale brain models require, besides a huge number of neurons, the right layout of local and global connections. Moreover, they need to capture the plasticity mechanisms operative in the adult brain. These mechanisms include, in addition to changes in the strength of synapses, structural modifications in anatomical connections. With the articles in this research topic, we hope that the reader will become aware of the methods and models by which large-scale brain networks running on supercomputing facilities may be extended to include anatomical detail and the full plastic potential of the brain.

AUTHOR CONTRIBUTIONS

MB, WS, and AvO wrote the Editorial.

REFERENCES

Axer, M., Amunts, K., Grässel, D., Palm, C., Dammers, J., Axer, H., et al. (2011). A novel approach to the human connectome: ultra-high resolution mapping of fiber tracts in the brain. *Neuroimage* 54, 1091–1101. doi: 10.1016/j.neuroimage.2010.08.075

Butz, M., and van Ooyen, A. (2013). A simple rule for dendritic spine and axonal bouton formation can account for cortical reorganization after focal retinal lesions. *PLoS Comput. Biol.* 9:e1003259. doi: 10.1371/journal.pcbi.1003259

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

Greene, R. (1989). Spacing effects in memory: evidence for a two-process account. *J. Exp. Psychol. Learn. Mem. Cogn.* 15, 371–377.

Hines, M. L., and Carnevale, N. T. (2001). NEURON: a tool for neuroscientists. *Neuroscientist* 7, 123–135. doi: 10.1177/107385840100700207

Lansner, A., and Holst, A. (1996). A higher order Bayesian neural network with spiking units. *Int. J. Neural Syst.* 7, 115–128. doi: 10.1142/s0129065796000816

Vogel, A., Reiter, S., Rupp, M., Nägel, A., and Wittum, G. (2013). UG 4: a novel flexible software system for simulating PDE based models on high performance computers. *Comput. Vis. Sci.* 16, 165–179. doi: 10.1007/s00791-014-0232-9

Estimating Fiber Orientation Distribution Functions in 3D-Polarized Light Imaging

Markus Axer[1]*[†], Sven Strohmer[2,3][†], David Gräßel[1], Oliver Bücker[2], Melanie Dohmen[1], Julia Reckfort[1], Karl Zilles[1,4,5] and Katrin Amunts[1,6]

[1] *Research Centre Jülich, Institute of Neuroscience and Medicine, Jülich, Germany, [2] Jülich Supercomputing Centre, Institute for Advanced Simulation, Research Centre Jülich, Jülich, Germany, [3] Research Centre Jülich, Simulation Lab Neuroscience, Bernstein Facility for Simulation and Database Technology, Institute for Advanced Simulation, Jülich, Germany, [4] Department of Psychiatry, Psychotherapy, and Psychosomatics, RWTH Aachen University, Aachen, Germany, [5] JARA Jülich-Aachen Research Alliance, Translational Brain Medicine, Aachen, Germany, [6] C. and O. Vogt Institute for Brain Research, Heinrich-Heine-University Düsseldorf, Düsseldorf, Germany*

Research of the human brain connectome requires multiscale approaches derived from independent imaging methods ideally applied to the same object. Hence, comprehensible strategies for data integration across modalities and across scales are essential. We have successfully established a concept to bridge the spatial scales from microscopic fiber orientation measurements based on 3D-Polarized Light Imaging (3D-PLI) to meso- or macroscopic dimensions. By creating orientation distribution functions (pliODFs) from high-resolution vector data via series expansion with spherical harmonics utilizing high performance computing and supercomputing technologies, data fusion with Diffusion Magnetic Resonance Imaging has become feasible, even for a large-scale dataset such as the human brain. Validation of our approach was done effectively by means of two types of datasets that were transferred from fiber orientation maps into pliODFs: simulated 3D-PLI data showing artificial, but clearly defined fiber patterns and real 3D-PLI data derived from sections through the human brain and the brain of a hooded seal.

Keywords: connectome, fiber architecture, human brain, polarized light imaging, 3D-PLI, ODF

INTRODUCTION

The repertoire of neuroimaging tools that are able to target neuronal connectivity in both the living and the post mortem brain, is continuously growing. Technological developments in particular in the field of microscopy (Osten and Margrie, 2013), new preparation and labeling methods (Chung et al., 2013; Costantini et al., 2015) and a better understanding of how to process the collected data (Amunts et al., 2013; Silvestri et al., 2015), facilitates this advancement (for recent review Amunts and Zilles, 2015). Several of these techniques address either cellular or even molecular dimensions, e.g., light sheet microscopy, or they provide data at meso- to macroscopic scales, e.g., Diffusion Magnetic Resonance Imaging (dMRI). Thus, the data output generated by different technical approaches and imaging techniques results in different data types, formats, and sizes, and is obtained on different spatial scales. As a consequence, comprehensible comparison across modalities and across scales evolves into a basic necessity for the neuroscience community.

In the present study, we demonstrate how to bridge the spatial scales from microscopic post mortem fiber visualization and orientation measurements based on 3D-Polarized Light Imaging (3D-PLI; Axer et al., 2011a,b) to meso- or macroscopic dimensions as targeted by dMRI. Our approach enables the propagation of the entire information on the microscopic fiber architecture within individual voxels by means of sophisticated data fusion. Based on the commonly used strategy to integrate vector data into a comprehensive description by employing Orientation Distribution Functions (ODFs, Bunge, 1982), we introduce the *pliODF* derived from 3D-PLI. pliODFs benefit from the unique property of 3D-PLI to extract high-resolution 3D vector fields indicating the spatial orientation of single fibers and fiber tracts in unstained brain sections. By transferring the entire vector data within a certain compartment, the so-called *super-voxel*, into a 3D statistical description (often visualized in form of a glyph), an efficient downscaling of high-resolution vector-like data becomes feasible. Considering the pure size of a 3D-PLI dataset that covers a whole human brain (i.e., 2500 sections scanned at 1.3 microns pixel size sum up to at least 500 TByte), the development of a method that reliably resamples large-scale microscopic data is of particular importance.

The choice of using ODF-like statistical descriptions of multiple-fiber compartments is based on the fact, that recent dMRI methods such as Diffusion Spectrum Imaging, Q-ball Imaging or Spherical Deconvolution (Tuch et al., 2002; Alexander, 2005; Wedeen et al., 2005; Dell'Acqua et al., 2013) also approximate the distribution of fiber orientations within an MRI-voxel by means of fiber Orientation Distribution Functions (fODF; Alexander et al., 2002; Hess et al., 2006; Rathi et al., 2009; Assemlal et al., 2011). Using a similar mathematical description is clearly beneficial for multimodal comparisons. Approaches have been reported aiming at the extraction of textural 2D-information, i.e., local fiber orientations, from images of histological brain sections stained for myelin in order to create 2D structural ODFs (Leergaard et al., 2010; Budde and Frank, 2012). In this context, small regions of interest were successfully compared to fODFs obtained from dMRI measurements. These studies represent important steps toward a region-based comparison of 2D fiber architecture obtained from different modalities. However, the 3D fiber architecture across large datasets has not been addressed yet.

Here, we benefit from the three-dimensional nature of the polarized light imaging approach, which provides measurements of the 3D fiber structures at the level of individual brain sections. Consequently, a super-voxel—and also a pliODF—can arbitrarily be composed of compartments within a section, but also across aligned neighboring sections without the requirement to change the software tools. To demonstrate the validity of our approach, two types of 3D-PLI datasets were transformed into pliODFs: (i) simulated 3D-PLI data showing synthetic, but clearly defined fiber-like patterns and (ii) real 3D-PLI data derived from sections through the human brain and the brain of a hooded seal. The human brain data were selected to highlight the gain of high-resolution imaging of brain regions with challenging fiber compositions such as the complex fiber crossings in the corona radiata or low fiber density regions in the cortex. The chiasm of the hooded seal with its nearly perpendicularly decussating fiber tracts (cf. Dohmen et al., 2015) appeared to be well suited to show the transition from simulation-based ODF generation to the most simple crossing fiber tract constellation observable in real brain tissue.

MATERIALS AND METHODS

3D-PLI (Axer et al., 2011a,b; Zeineh et al., 2016) has demonstrated its unique capabilities (i) to reveal fiber structures at multiple scales, such as long-range connections and even single fibers and crossings within unstained histological brain sections, and (ii) to determine spatial fiber orientations (i.e., 3D unit vectors down to the scales of fiber diameters (0.4– 15 μm). 3D-PLI is applicable to unstained microtome sections of post mortem brains and utilizes the optical birefringence of nerve fibers, which basically arises from the highly ordered arrangement of lipid molecules in the myelin sheath surrounding most of the axons in the brain. Polarimetric setups (e.g., a polarizing microscope) are employed to carry out birefringence measurements and to give contrast to individual nerve fibers and their tracts. Supported by fundamental principles of optics and dedicated simulation approaches (Dohmen et al., 2015; Menzel et al., 2015), the measured signals are additionally interpreted in terms of spatial fiber orientations by means of unit orientation vector descriptions (**Figure 1**). The algorithms used for the fiber orientation interpretation have been implemented as an automated 3D-PLI analysis workflow suitable for distributed supercomputing, as described by Axer et al. (2011b); Amunts et al. (2014).

The Fiber Orientation Map

Application of the 3D-PLI analysis workflow results in a unit vector-based description of the fiber orientation determined for each tissue voxel, referred to as a native voxel. The native voxel dimensions are defined by the image pixel size provided by the optical setup and the thickness of the studied histological brain section (70 μm in the present study). The generation of pliODFs was exclusively performed on 3D-PLI datasets with native voxel dimensions of $1.3 \times 1.3 \times 70$ μm. Each orientation vector reflects the net effect of all fibers comprised within a voxel. The assembly of all unit vectors represents the fiber orientation map (*FOM*).

The orientation unit vector $\vec{u}_i$ at voxel location i can be parameterized by two angles, i.e., by spherical coordinates: the direction angle φ_i, which represents the projection of the principal fiber axis within the sectioning plane, and the inclination angle α_i, which is the angle between the principal fiber axis and the sectioning plane (**Figure 1** and Equation 1).

$$\vec{u}_i = \begin{pmatrix} cos\alpha_i \cdot cos\varphi_i \\ cos\alpha_i \cdot sin\varphi_i \\ sin\alpha_i \end{pmatrix} \qquad (1)$$
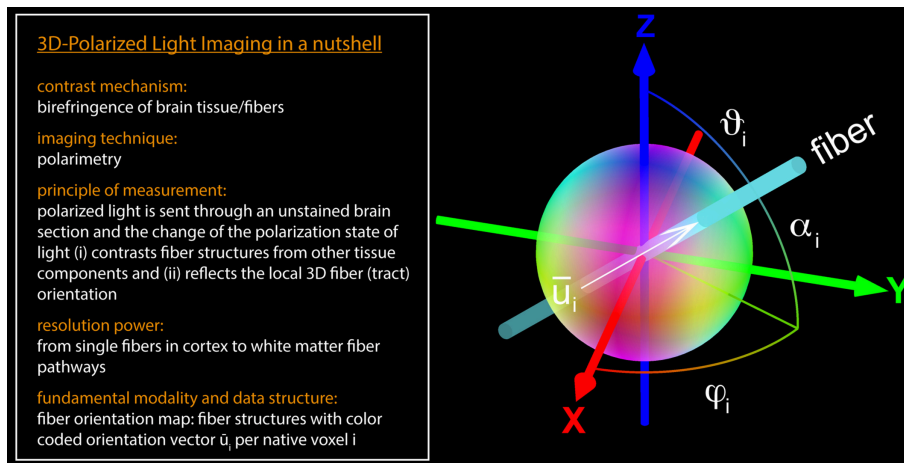
FIGURE 1 | 3D-PLI in a nutshell. The fiber orientation vector $\vec{u}_i$ located at the native voxel location i is defined by the direction angle φ_i and the inclination angle α_i. The brain section is aligned to the x-y-plane (at $z = 0$). In general, a unit vector can be described as a distinct point on a unit sphere. For 3D-PLI, the determined unit vector represents an axis in space (in the mathematical and the anatomical sense), since the initial and the terminal points cannot be assigned uniquely. Therefore, the point reflection across the center of the sphere has to be assigned to the same vector. For visualization purposes, the vectors were encoded in RGB color space: the x-component of $\vec{u}_i$ is represented by the red color channel (R), the y-component by the green color channel (G), and the z-component by the blue color channel (B).

Image Generation and Data Acquisition
Simulated 3D-PLI Data

The simulation software tool SimPLI (Dohmen et al., 2015; Menzel et al., 2015) was used to generate two synthetic 3D-PLI datasets with known configurations of fiber-like structures. In SimPLI, three main steps of simulation are implemented: (i) the generation of an arbitrary spatial arrangement of synthetic fibers and the discretization into a three-dimensional fiber orientation vector field with a certain resolution (e.g., 70 μm isotropic), (ii) the calculation of the transmitted light intensity based on the Jones matrix calculus (Jones, 1941) yielding a synthetic 3D-PLI image series, and (iii) the simulation of environmental clutter arising from the camera and the tissue in the optical path by adding blurring and noise effects.

In order to validate the different methodological steps employed to transfer a FOM into a set of orientation distribution functions, a well-defined template providing unambiguous structural macroscopic and microscopic features in terms of left/right, top/down and in-plane/out-of-plane orientations, was required. This dataset generated by means of SimPLI is shown in **Figures 2A–D**. It is composed of a stack of 18 images and comprises birefringent "fibers" forming human readable structures ("fiber bundles"), such as the capital letter "R" and the "±" sign. The line thickness of the letters (or the thickness of the "fiber bundles") was chosen to be 20 pixels on average. The fiber inclination angles in "R" were all set to $\alpha = 0°$, while the inclination angles were set to $\alpha = +45°$ and $\alpha = -45°$ for the "+" and "−" sign, respectively. The direction angles φ were aligned with the local structures using a right-handed coordinate system, i.e., the horizontal components (e.g., the "−" sign) are identified by $\varphi = 0°$ while the vertical components are represented by $\varphi = 90°$. The diagonal element of the "R" has a direction of $\varphi = 135°$. The background is composed of 90° inclined fibers corresponding to light intensity variations equal to

zero. This dataset was subjected to the 3D-PLI analysis workflow to extract the corresponding FOM (**Figures 2E, 3A**).

A second, more realistic dataset was generated with SimPLI and subjected to 3D-PLI analysis (cf. **Figure 5F**), and has already been described in detail in Dohmen et al. (2015): the model of a chiasm of the hooded seal. This data bridges to the successive section on real 3D-PLI data sets.

3D-PLI Data Obtained from Histological Sections of the Hooded Seal and the Human Brain

Two human brains and the optic chiasm of a hooded seal were immersed in 4% paraformaldehyde. After cryoprotection (with a 20% glycerin solution), the brain tissue was deep frozen at $-70°$C and stored at the same temperature till further processing. A whole human brain and the occipital lobe of a human brain were coronally sectioned and optic chiasm was axially sectioned using a large-scale cryostat microtome (Polycut CM 3500, Leica, Germany), and eventually coverslipped with a 20% glycerin solution. The chosen section thickness was 70 μm. During sectioning, each blockface was imaged using a CCD camera mounted above the brain in order to obtain an undistorted reference for each section (cf. **Figures 5A, 6A, 7A**). No staining was applied. In each case, this procedure resulted in a complete series of sections through large tissue samples, which enables a 3D reconstruction. The brains were acquired in accordance with local legal and ethical requirements.

The brain sections were measured in two custom-made polarimetric setups, the large-area polarimeter and the polarizing microscope, providing images with pixel dimensions of 64 × 64 μm and 1.3 × 1.3 μm, respectively (for technical details refer to Axer et al., 2011b). The datasets were passed to the 3D-PLI analysis workflow to extract the corresponding FOMs. For pliODF generation only the high-resolution data were used, while the 64 μm-sized FOMs were used for plausibility checks and reference measures.
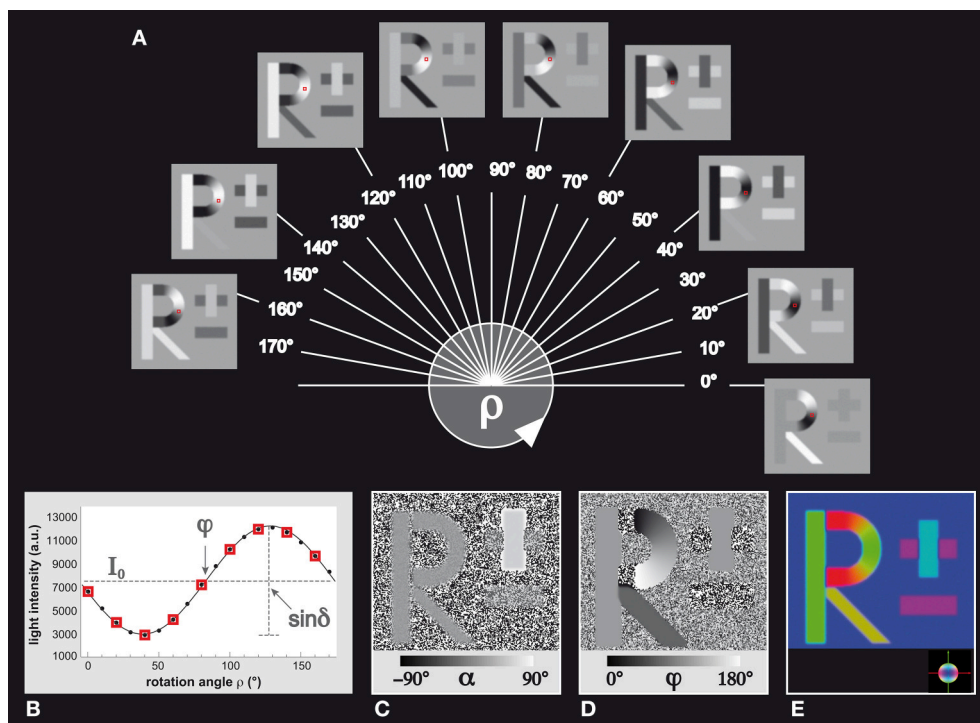
FIGURE 2 | Simulated 3D-PLI reference dataset. (A) The standard 3D-PLI measurement yields 18 images corresponding to equidistant rotation angles ρ between 0° and 170°. Here, a selection of nine generated images of simulated birefringent structures (the letter R and the ± sign) is shown. Each image has a size of 200 × 200 pixels. The varying pixel intensities are comparable to observed signals in measurements of brain sections. The red squares indicate a native voxel of interest that is displayed (B) in terms of the observed light intensity as a function of the rotation angle. The physical model that underlies 3D-PLI provides a sinusoidal description of the simulation (continuous black line), and relates (C) its amplitude to the inclination angle α via the retardation value $sin\delta$ and (D) its phase to the direction angle φ. The introduced effects of blurring and noise are evident: the minus sign in the direction map (D), for example, shows direction angles that are spread around the initial direction φ = 0° by±2.5°. In a π-periodic system, this is equivalent to an angle range between 177.5° and 2.5°. (E) Visualization of the FOM with the determined vectors $\vec{u}$ encoded in RGB color space (see color sphere for the relation between orientation and color-coding).

Estimation of Fiber Orientation Distribution Functions (pliODFs)

The fundamental aim was to generate a mathematical description -the orientation distribution function (pliODF)– that quantifies the spatial distribution of fiber orientations determined by 3D-PLI within a rectangular compartment. The compartment can be defined in a single FOM or in a series of FOMs and is referred to as super-voxel. A super-voxel is composed of ($r \times c \times s = rows \times columns \times sections$) native voxels. In order to calculate pliODFs, fundamental approaches from material science (texture analysis in crystallography; Bunge, 1982) and directional statistics (Mardia and Jupp, 2000) were applied and adopted to the needs of 3D-PLI.

The implemented procedure was based on three steps (cf. **Figure 3**):

(A) the definition of super-voxels (rectangular compartments) in high-resolution FOMs to resample the dataset,

(B) the collection of all fiber orientation vectors comprised in a super-voxel to derive the discretized distribution of orientations by creation of a normalized directional histogram on a unit sphere, and

(C) the approximation of the orientation probability distribution density by fitting the directional histogram with a series expansion using spherical harmonics.

In the following, steps (A) to (C) will be explained in more detail. As mentioned above (**Figure 1**), a 3D-PLI orientation vector can be expressed in spherical coordinates (the polar angle $\vartheta = \frac{\pi}{2} - \alpha$ and the azimuth angle φ) and defines two points on the unit sphere S^2. This feature was used to construct a directional histogram aiming at a statistical description of the fiber orientations contained in a super-voxel. To discretize the distribution of fiber orientations, the surface of a unit sphere was subdivided into planar bins (bin centers characterized by latitude and longitude) with a total bin number n calculated by

$$ n = \left(\# \ of \ latitudes \right) + \left(\# \ of \ longitudes \right) + 2 \ polar \ caps. \quad (2) $$

n was adopted to the specific requirements of the used datasets. In a final step, the number of vectors falling into each bin was determined (**Figure 3B**).

Normalization of the directional histogram enabled to mathematically describe the empirical orientation probability distribution density $p(\omega)$ by
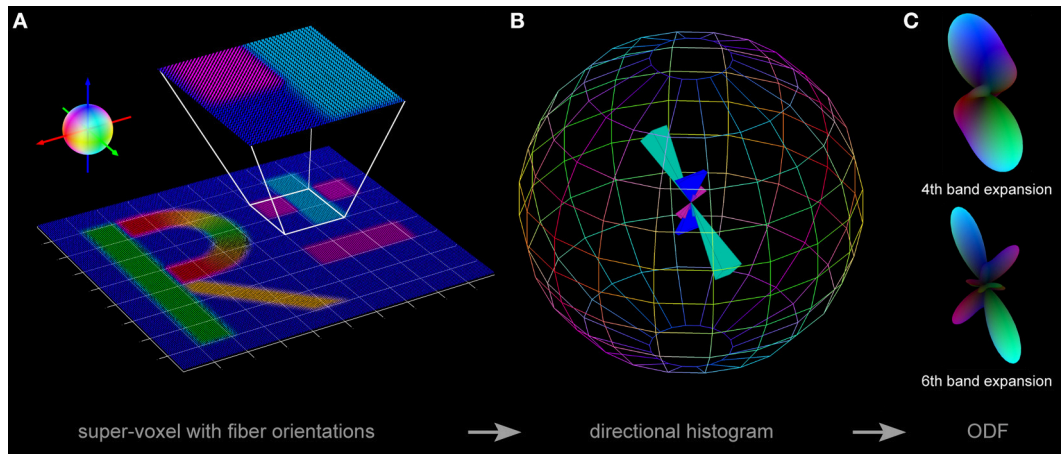
FIGURE 3 | Three steps toward pliODF generation. (A) First, a FOM is divided into regular domains or super-voxels. The exemplarily enlarged super-voxel contains 40 × 40 × 1 native voxels representing three predominant fiber orientations, which show a relative frequency of occurrence of ∼¼(blue color), ∼¼(magenta color) and ∼½(cyan color), respectively. The color sphere defines the relation between orientation and color-coding. **(B)** Second, a normalized directional histogram with a discretized binning on a unit sphere is created for each super-voxel. The relative fraction of fiber orientations assigned to a particular bin is reflected by the length of the colored solid angle originating from the middle of the sphere. The symmetry of the histogram with respect to point reflection across the center of the sphere is evident. Here, the total number of bins distributed over the sphere was set to 164 and the three predominant input fiber orientations are still preserved. **(C)** Third, a spherical harmonics expansion is used to approximate each directional histogram. Depending on the selected depth of expansion (e.g., to the 4th or the 6th band), orientation distribution features might become occluded by interpolation.

$$\frac{dN}{Z} = p(\omega)\, d\Omega, \tag{3}$$

with Z being a normalization factor ($Int\left(\frac{dN}{Z}\right) = 1$), d$\Omega$ the dihedral angle differential, and ω the spatial direction. $p(\omega)$ was expanded into a series of generalized spherical harmonics:

$$p(\omega) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{n=-l}^{l} C_l^{mn} T_l^{mn}(\omega)$$

$$= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{n=-l}^{l} C_l^{mn} e^{I\varphi_2} P_l^{mn}(\Phi)\, e^{I\varphi_1}. \tag{4}$$

The P_l^m represent the associated Legendre polynomials. Due to the existing rotational symmetry (i.e., the determined fiber orientations are invariant with respect to rotation around the axis defined by the unit vector), the general description of the series simplified to an expansion in terms of spherical harmonics $Y_l^m(\vartheta, \varphi)$:

$$p(\vartheta, \varphi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} C_l^m Y_l^m(\vartheta, \varphi), \tag{5}$$

with l and m denoting the band index and sub-band index, respectively. A series expansion to the 6th band, for example, means $l = 0, \ldots, 6$ and $-l \le m \le l$ (with $1 + 3 + 5 + 7 + 9 + 11 + 13 = 49$ members).

The expansion of the empirical orientation probability distribution $p(\vartheta, \varphi)$ into real valued symmetric spherical

harmonics series reads as:

$$p(\vartheta, \varphi) = \sum_{l=0}^{\infty} \sum_{m=-2l}^{2l} C_{2l}^m y_{2l}^m(\vartheta, \varphi) \approx \sum_{l=0}^{\hat{L}} \sum_{m=-2l}^{2l} C_{2l}^m y_{2l}^m(\vartheta, \varphi), \tag{6}$$

with

$$\frac{1}{2}(L+1)(L+2) = \left(2\hat{L}+1\right)\left(\hat{L}+1\right) \tag{7}$$

coefficients (with $2\hat{L} = L$).

To generate a pliODF, the coefficients C_{2l}^m of the expansion up to the l^{th} band had to be determined. This was realized using a least square fit algorithm, because of its numerical stability and the efficiency of the numerical implementation. Due to discretization and truncated series expansion, the pliODF only approximates the empirical orientation probability distribution p.

Computing

Both, the size of the processed data and the computationally intensive algorithms to determine the expansion coefficients required a supercomputing environment. For this reason, we used the Juelich Dedicated GPU Environment (JuDGE), hosted by the Jülich Supercomputing Center (JSC), Germany. It was equipped with 206 compute nodes, where each node consisted of two Intel Xeon Westmere 6-core processors operating at 2.66 GHz. Each compute node contained 96 GB of main memory. Per node there were either two NVIDIA Tesla M2050 (Fermi) GPUs or two NVIDIA Tesla M2070 (Fermi) GPUs integrated.
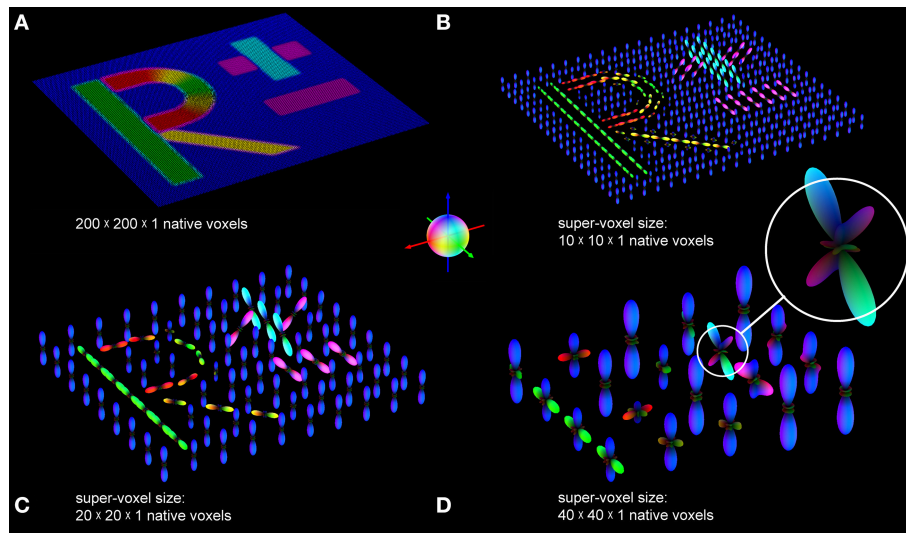
FIGURE 4 | pliODFs based on different super-voxel sizes. (A) The FOM of the simulated dataset (cf. **Figure 2**) was divided into super-voxels composed of **(B)** 10 × 10 × 1, **(C)** 20 × 20 × 1, and **(D)** 40 × 40 × 1 native voxels. pliODFs were generated by means of series expansions to the 6th band for the different super-voxels. The color sphere defines the relation between orientation and color-coding.

The inter- and intra-node communication was realized by the message passing interface (MPI) protocol.

Runtime measurements were performed for a region of interest in a FOM gained from the human brain tissue [i.e., ROI (1) as depicted in **Figure 6B**], which was composed of 3712 × 4576 orientation vectors. The number of compute cores was set to 72 in order to keep the compute time in an acceptable range, but to highlight the differences in run-time properly at the same time. The number of bands and the size of super-voxels were taken as parameters.

Visualization

ODFs are typically visualized either (i) by means of a textured sphere, where the color of a point on the surface represents the probability of its corresponding orientation, or (ii) by simply scaling the surface with the probability p. To visualize the generated pliODFs, the second option was applied in combination with the color-coding scheme also used to visualize different fiber orientations in a FOM and the surface intensity increasing with the probability (**Figure 3C**). The scheme is based on the RGB color space with the red channel representing the x-component, the green channel representing the y-component, and the blue channel representing the z-component of the fiber orientation in the reference coordinate frame. The peaks of the pliODF shape and the color-coding reflect the most common fiber directions such that both the pliODFs and the FOMs are visually comparable.

Results
Simulated 3D-PLI Data

The FOM obtained from the simulated dataset was divided into super-voxels of 10 × 10 × 1, 20 × 20 × 1, and 40 × 40 × 1

native voxels, respectively, to create pliODFs. These clusters were chosen (i) to demonstrate the accurate interpretation of the input vectors including the conservation of the coordinate system and (ii) to investigate how different super-voxel dimensions affect the shape of the resulting pliODF representations. Therefore, super-voxel dimensions close to the (fiber) structure dimensions (here: about 20 pixels, equivalent to 36 microns; cf. **Figure 3A**) were of specific interest. The binning of the directional histogram was set to (9 latitudes × 18 longitudes + 2 polar caps) = 164 bins.

As demonstrated in **Figures 3A–C** on the basis of an exemplary 40 × 40 × 1 super-voxel, the implemented resampling procedure of high-resolution FOMs provided a comprehensible description of the distribution of fiber orientations, in form of a directional histogram or a pliODF. The comparison of the pliODFs based on different levels of expansion with the corresponding directional histogram suggested that the approximation of the probability density function with spherical harmonics in the studied cases with the selected bin sizes should be confined (at least) to the 6th band, in order to reliably resolve orthogonal contributions (e.g., cyan and magenta).

The resampling results for different super-voxel dimensions are shown in **Figures 4B–D**. Compared to the original unit vector description of the fiber orientations, the peaks of the pliODFs obtained from the small super-voxels (**Figures 4B,C**) reflect the main underlying (microscopic) fiber orientations corroborated by the matching colors. In addition, the general (macroscopic) orientations of the letters agree with the orientations of the input structures. As expected, the complexity of the pliODF shapes increases in larger samples (**Figure 4D**), maintaining the major portions of fiber orientations.
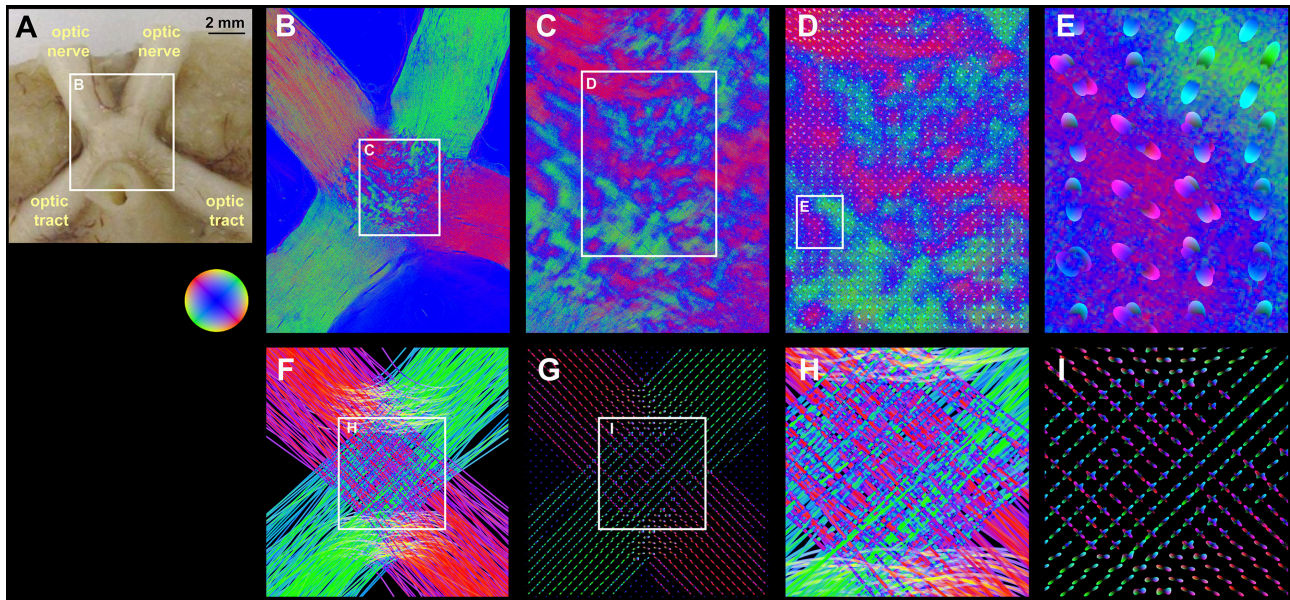
FIGURE 5 | Real and simulated brain section from the hooded seal. (A) Blockface image of the optic chiasm of the hooded seal before sectioning. **(B)** Fiber orientation map of a medial section through the optic chiasm. Optic nerves and optic tracts appear as massive and rather homogeneous fiber bundles. Most fiber tracts from the optic nerves decussate to the contralateral optic tract. **(C)** The decussation zone in the center (i.e., the chiasm) is characterized by a patch pattern produced by small fiber tracts (red and green color; exemplary orientations are indicated by black lines) and fiber crossings characterized by signal attenuation (blue color; exemplary highlighted by white arrow). Based on this FOM, pliODFs were created for super-voxel dimensions of 40 × 40 × 1 native voxels. **(D,E)** demonstrate different enlargements of the field of pliODFs overlaid with the input FOM. **(F)** FOM of a simulated section through the optic chiasm and **(G)** corresponding pliODFs for super-voxel dimensions of 40 × 40 × 1 native voxels. **(H)** Zoom into the FOM of the fiber decussation zone and **(I)** corresponding pliODFs. The effects of crossing and bending fibers on the ODF shapes are obvious.

3D-PLI Data Obtained from a Hooded Seal and a Human Brain

Brain Sections of a Hooded Seal

FOMs taken from hooded seal brain tissue show the optic nerve traversing through the optic chiasm into the optic tract (**Figures 5A,B**). The center of the optic chiasm reveals decussate fiber populations alternating with blue dots caused by signal attenuation (white arrow, **Figure 5C**). A pattern of crossing fiber tracts was used to prove the functionality of our implementation in a realistic setting. The FOMs were sampled using a histogram binning of (50 latitudes × 100 longitudes + 2 polar caps) = 5002 bins with a super-voxel size of $52 \times 52 \times 70\,\mu m^3$, which is equivalent to 40 × 40 × 1 native voxels. The series expansion of the pliODFs was confined to the 6^{th} band. The fused images of the high-resolution FOMs with the corresponding pliODFs demonstrated a sound resampling (**Figures 5D,E**). The simulated chiasm of the hooded seal was analyzed accordingly and showed concordant results (**Figures 5H,I**).

Human Brain Sections

Three high-resolution FOMs of selected regions of interest from a coronal section through the human occipital lobe (**Figure 6A**) were resampled at different super-voxel dimensions (**Figure 6B**), but with fixed histogram binning (50 latitudes × 100 longitudes + 2 polar caps = 5002 bins). The targeted super-voxel sizes of $26 \times 26 \times 70\,\mu m^3$, $52 \times 52 \times 70\,\mu m^3$, and $260 \times 260 \times 70\,\mu m^3$

correspond to $20 \times 20 \times 1$, $40 \times 40 \times 1$, and $200 \times 200 \times 1$ native voxels, respectively. The series expansion was confined to the 6th band. The pliODFs were compared both with the underlying high-resolution FOMs acquired with the polarizing microscope (**Figure 6C**) and FOMs obtained with the large-area polarimeter (**Figure 6D**).

In addition, two high-resolution FOMs of selected regions of interest from a coronal section through the human brain (**Figure 7**) at the level of the central region were resampled at different super-voxel dimensions (**Figures 7C,D**), but with fixed histogram binning (50 latitudes × 100 longitudes + 2 polar caps = 5002 bins). The targeted super-voxel sizes of $65 \times 65 \times 70\,\mu m^3$, and $260 \times 260 \times 70\,\mu m^3$ correspond to $50 \times 50 \times 1$, and $200 \times 200 \times 1$ native voxels, respectively. The series expansion was confined to the 6th band.

The following observations were made:

(i) The overall fiber architecture in all regions of interest were preserved (**Figures 6B–D**, **7C,D**) by pliODFs. This holds true for deep white matter tracts (e.g., in the stratum sagittale, **Figure 6D**, or the corpus callosum, **Figure 7D**). In cortical regions, the integrity of the fiber architecture is maintained by the 6^{th} band expansion even at increasing super-voxel sizes beyond $250\,\mu m$ (**Figure 6B**).

(ii) The peaks of the pliODFs, i.e., the prevailing directions, agree with the FOMs gained from the large-area polarimeter
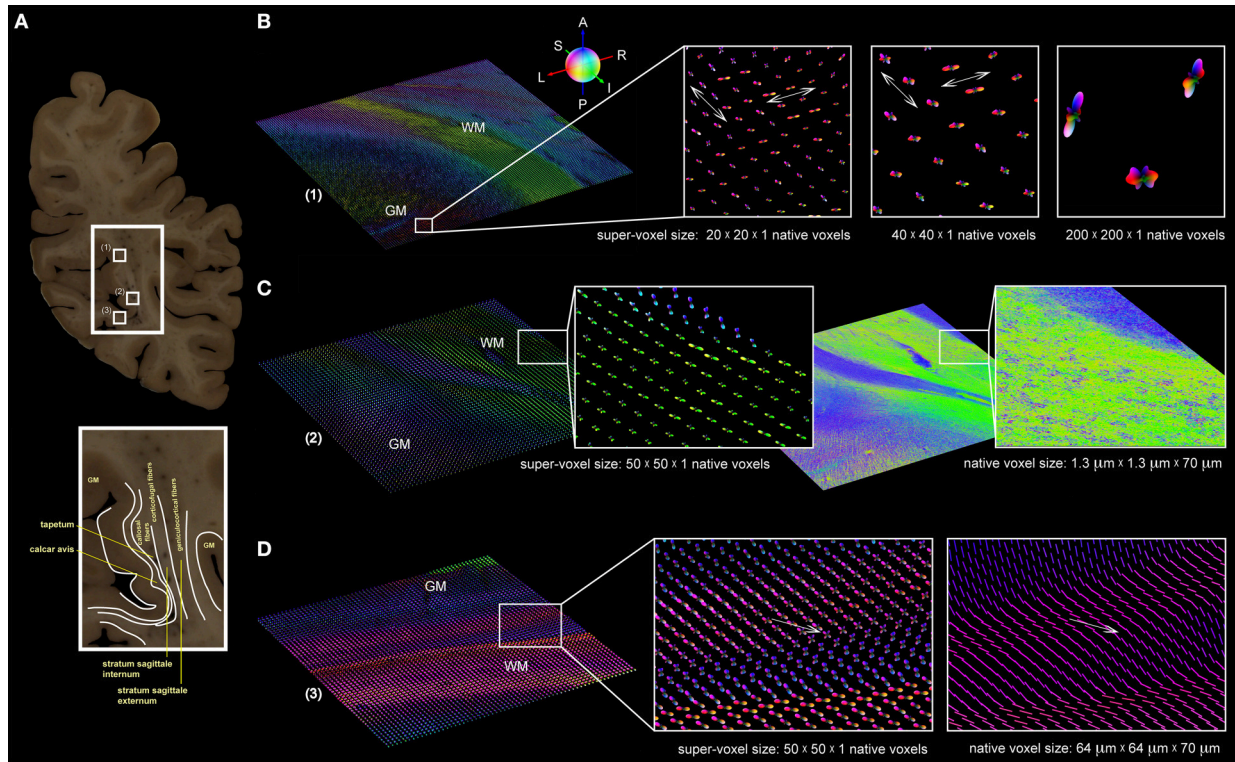
FIGURE 6 | Brain section from the human occipital lobe. (A) Segmented blockface image acquired from the surface of the frozen human occipital lobe during the sectioning process. The small white rectangles (1) to (3) indicate the selected regions of interest for which pliODFs were determined (cf. **B–D**). The enlarged extract shows the delineation of anatomical structures, such as the tapetum, the calcar avis, and the stratum sagittale. **(B)** pliODF representations in region (1) with super-voxel dimensions of $20 \times 20 \times 1$, $40 \times 40 \times 1$ and $200 \times 200 \times 1$ native voxels; the magnified images show the same cortical region, which is characterized by crossing fibers (indicated by the white arrows). The largest super-voxel size is equivalent to $260 \times 260 \times 70\,\mu m^3$ and corresponds approximately to the level of high-resolution post mortem dMRI measurements. **(C)** Region (2) demonstrates for a super-voxel dimension of $50 \times 50 \times 1$ native voxels the preservation of the overall fiber structure in comparison with the original high-resolution FOM obtained with the polarizing microscope. Zooming into the data reveals pliODFs with multiple fiber orientations in inhomogeneous white matter regions. **(D)** For region (3), pliODFs (super-voxel dimension of $50 \times 50 \times 1$ native voxels) are opposed to the vector-based representation of the FOM of the same brain region measured with the large-area polarimeter at $64 \times 64 \times 70\,\mu m^3$ voxel size. The white arrows indicate a crossing zone of fibers.

in tissue regions basically composed of parallel fibers (**Figure 6D**).

(iii) Super-voxels localized in transition areas of adjacent or crossing fiber tracts preserved the corresponding fiber orientations in the pliODF with high fidelity (**Figures 6D, 7C**), while the same regions scanned with the large-area polarimeter showed attenuated signals or averaged orientations (e.g., **Figure 6D**).

Runtime Behavior

Runtime measurements (for 3712×4576 orientation vectors) performed on the JuDGE supercomputer showed that (i) with increasing super-voxel size the overall runtime decreased and (ii) with increasing depth of expansion the overall runtime increased (cf. **Figure 8**).

DISCUSSION

General Concept

3D-Polarized Light Imaging has been demonstrated in previous neuroanatomical studies (Axer et al., 2011a,b;

Caspers et al., 2015; Dohmen et al., 2015; Zilles et al., 2015; Zeineh et al., 2016) to provide unique high-resolution data on the brains' fiber architecture of various species, such as mouse, rat, seal, vervet monkey, and human. The color-coded fiber orientation map (FOM) was the fundamental image modality for each of these studies, since it comprised both the highlighted fiber structures and their local 3D orientations. Hence, each FOM enabled a comprehensive delineation and identification of neuroanatomical fiber structures across different scales covering the micrometer (i.e., cortical and sub-cortical fibers) to the centimeter range (i.e., long-range fiber tracts across entire brain sections).

In the present study, we have successfully implemented a dedicated methodology of integrating high-resolution vector data obtained from FOMs into a comprehensive statistical description, the orientation distribution functions or pliODFs, respectively. By this means, efficient downscaling of high-resolution FOMs was achieved and cross-scale, cross-modality comparisons were enabled. The general concept of pliODF generation can be summarized in three steps:
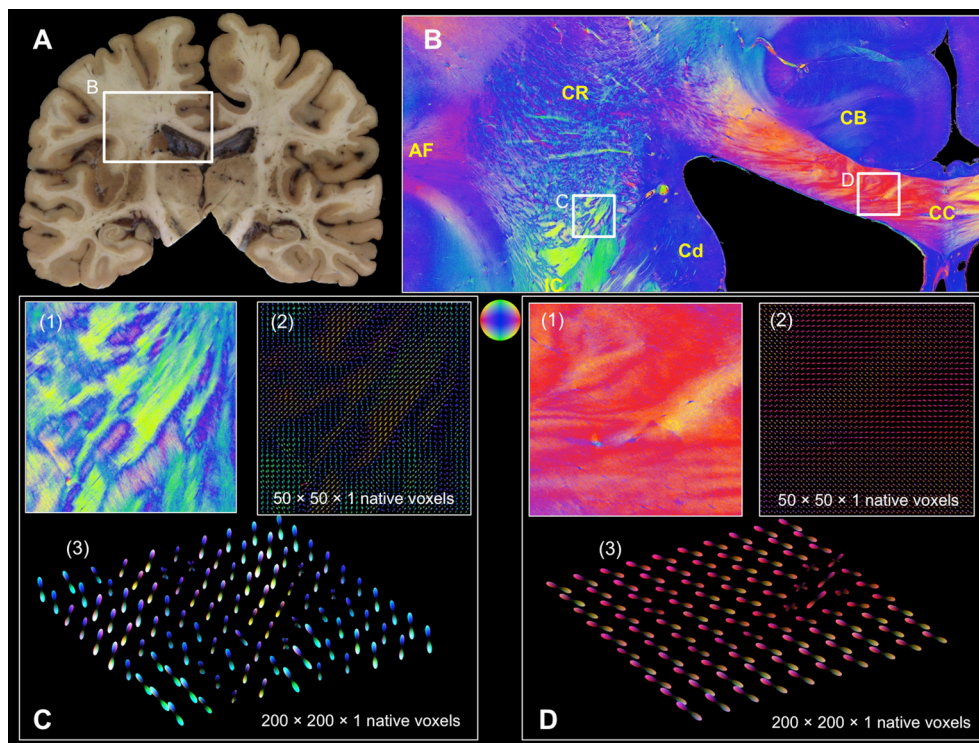
FIGURE 7 | Coronal section from the central human brain region. (A) Segmented blockface image acquired from the surface of the frozen human brain. **(B)** Fiber orientation map comprising arcuate fascicle (AF), cingulum bundle (CB), corpus callosum (CC), caudate nucleus (Cd), corona radiata (CR), internal capsule (IC). Fiber orientations are RGB color-coded (see sphere). The white rectangles indicate FOMs (C(1) and D(1)) that were transferred into pliODFs. **(C)** pliODF representations in the region of CR/IC with super-voxel dimensions of (2) 50 × 50 × 1 and (3) 200 × 200 × 1 native voxels (from different views). Patches of crossing fiber bundles are clearly visible in the FOM and the pliODF maps. **(D)** pliODF representations in the region of CC with super-voxel dimensions of (2) 50 × 50 × 1 and (3) 200 × 200 × 1 native voxels (from different views). Although a predominant fiber direction is observable, small wriggling fiber bundles cause local inhomogeneities along the CC.

(A) the definition of super-voxels (rectangular compartments) in high-resolution FOMs to resample the dataset,

(B) the collection of all fiber orientation vectors comprised in a super-voxel to derive the discretized distribution of orientations by creation of a normalized directional histogram on a unit sphere, and

(C) the approximation of the orientation probability distribution density by fitting the directional histogram with a series expansion using spherical harmonics. This yields a 3D glyph per super-voxel.

Validation Strategy

To prove the concept, we applied our implementation to simulated and real 3D-PLI datasets that reflected different characteristics of fiber compositions: well-defined parallel fibers (letter simulation, **Figure 4**), perpendicular fiber tract crossings (measurement and simulation the hooded seal chiasm, **Figure 5**), and complex fiber architectures (measurements in human brain sections, **Figures 6**, **7**). That way, the feasibility of our approach became traceable from the most simple to the most challenging cases.

The reliability of the steps required to derive a pliODF in form of a glyph was successfully demonstrated on basis of a simulated 3D-PLI dataset. The simulation tool SimPLI (Dohmen et al., 2015; Menzel et al., 2015) has specifically been developed to model the effects of birefringence in brain tissue as well as the results of polarimetric measurements. However, one of the generated synthetic datasets (cf. **Figure 2**) was not destined for realistic simulation of fiber architecture, but to enable systematic testing of the conservation of the involved coordinate systems, and fiber-like/bundle-like orientations, respectively. In combination with the same color-coding scheme applied to FOMs and pliODFs, this turned out to be an excellent approach for demonstration and validation purposes. This becomes evident in the study of directional histograms (i.e., the probability distributions of orientation vectors on a sphere) and their fitting with a series expansion with spherical harmonics (**Figure 3**).

The directional histogram is a proximate reflection of the main fiber orientations in a super-voxel as provided by the FOMs. However, the discrete arbitrary binning of the histogram affects the orientation information and, consequently, the final approximation of the density function. The different bin widths for simulation and experimental datasets, which were used for this study resulted from trials. For future studies, an automated data driven binning in S^2, depending on input and target data, is essential. The creation of directional histograms from high-statistics vector data is not computationally intensive and

FIGURE 8 | Runtime measurements. Runtime as a function of the super-voxel dimension for the series expansion up to the 6th and the 8th band. Note, that the left time axis is scaled logarithmically, while the right time axis is scaled linearly.
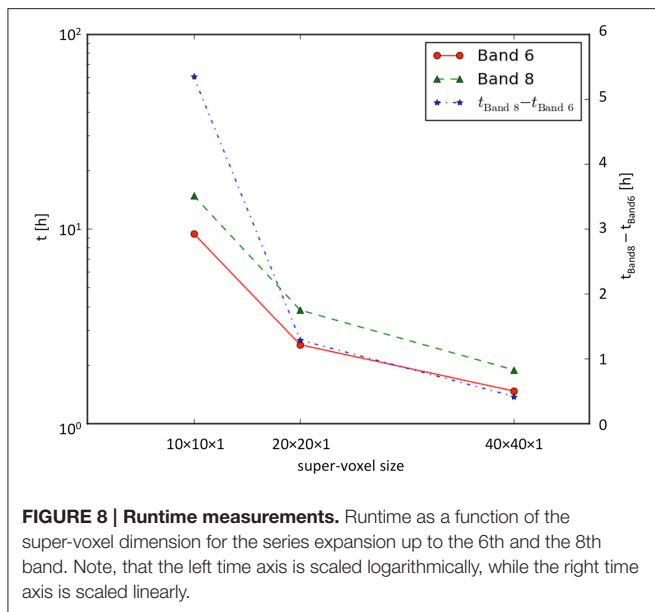
represents a very fast and robust method to derive the prevalent fiber directions in super-voxels. The precision of the fiber directions obtained from the directional histogram (defined by the bin centers) is basically limited to the size of the dihedral angle spanning a bin.

This limitation can be overcome by fitting the directional histogram with a series expansion using spherical harmonics to generate a pliODF, but at the expense of computing time that is increasing drastically. As derived from runtime measurements, the resampling of a FOM composed of 3712×4576 orientation vectors, the computationally most intensive step was the determination of the expansion coefficients. According to Equation 7, the number of coefficients to be determined increases non-linearly with increasing expansion depth. Hence, the series expansion has to be truncated. For the experiments done in this study, the expansion to the 6th band appeared to be sufficient to assess the functionality of our implementations, but, this level of resolution has to be reviewed in case of addressing different questions as to data precision, data size or computing time. By increasing the super-voxel size, the computation time is decreased in a non-linear way, due to the lower number of pliODFs to be computed.

What does this mean for whole human brain imaging and comparison? A human brain volume of $1200\,cm^3$ translates into 10^{10} native voxels provided by high-resolution 3D-PLI measurements. A voxel size of 2 mm isotropic, such as provided by standard clinical *in vivo* dMRI technologies (e.g., Bastiani and Roebroeck, 2015), leads to 150,000 dMRI voxels for a human brain. Resampling of 3D-PLI to this dMRI scale means integrating 65,000 orientation vectors into a single super-voxel or pliODF, respectively. For the runtime measurements performed on the 3712×4576 pixel-sized FOM, about 170,000 pliODFs with a super-voxel dimension of $10 \times 10 \times 1$ native voxels had to be computed. The computation time (expansion to the 6th band) was about 9.5 h utilizing only 72 compute cores. Conclusively,

computation of pliODFs for an entire human brain at 2 mm resolution is feasible. Post mortem dMRI has recently progressed to study the structural organization of the entire human brain at a voxel size of 0.7 mm isotropic (Miller et al., 2011), which results in 3.5 million target (super-)voxels, which is a factor of 20 more than provided by standard *in vivo* dMRI measurements. However, adjusting pliODFs to submillimeter dMRI data for a whole human brain is still in the realm of the feasible, taking into account that supercomputing facilities provide thousands of computing cores.

Scope of Application

We demonstrated that the pliODF generation workflow enables the study of scaling effects and related partial volume effects efficiently, already at the level of single brain sections for both real and simulated data (cf. **Figures 5–7**). 3D-PLI measurements of the same tissue at two distinct resolutions (native voxel sizes of $64 \times 64 \times 70\,\mu m$ and $1.3 \times 1.3 \times 70\,\mu m$) were beneficial in this context. The low-resolution FOMs were used as independent references for prevailing fiber orientations derived from pliODFs computed from super-voxel dimensions that matched the native voxel dimensions of the low-resolution measurements (e.g., **Figure 6D**). While the principal orientations of distinct fiber tracts from both types of fiber orientation descriptions agreed, the benefit of pliODFs became evident at transition zones between differently oriented fiber tracts. In the latter case, pliODFs preserved details about the complex fiber population, which was not observed for the low-resolution measurements. This is due to the fact, that in a 3D-PLI measurement a native voxel collects birefringence effects from multiple fibers resulting in a measurement of superimposed sinusoidal signals (Reckfort et al., 2015), each with fiber orientation specific amplitude and phase. As demonstrated by Dohmen et al. (2015), the derived fiber orientation vector for a native voxel significantly depends on the complexity of the underlying fiber population. For the low-resolution measurements this means that about 50–100 myelinated fiber contribute to the measured signal, if fiber diameters between 0.4 and $15\,\mu m$ (Aboitiz et al., 1992) are assumed. Future studies will further elaborate on scaling effects by combining sophisticated simulation approaches (Dohmen et al., 2015; Menzel et al., 2016) with measurements across scales (Reckfort et al., 2015). This will open up new avenues to derive observer independent quality measures for 3D-PLI measurements, beyond neuroanatomical inspection.

Even though the concept of the pliODF generation has been demonstrated for FOMs of individual brain sections, its application is not limited to section-like data, but it can also be extended to a volume of FOMs. In the latter case, the brain section images have to be re-aligned into a coherent 3D brain volume, which requires application of complex non-linear image registration techniques (Palm et al., 2010; Amunts et al., 2013). Assuming a whole human brain reconstruction from 2500 serial coronal sections scanned at $1.3\,\mu m$ pixel size with a single section image size of $70,000 \times 100,000$ pixels on average, it becomes evident, that the utilization of distributed high performance computing on a supercomputing environment is essential. The process of registration aims at correcting for

up to cm-sized tissue deformation introduced during brain sectioning and tissue mounting. This poses a major challenge when addressing 3D brain reconstruction at the μm-scale, as required for long-range pixel-wise tracing of fiber tracts across hundreds to thousands of brain sections. By integrating many orientation vectors into pliODFs, local inaccuracies in section alignment are likely to be polished and, therefore, tractography of long distance fiber pathways become feasible for large-scale 3D-PLI datasets, but at the expense of resolution. As a benefit, advanced tractography algorithms exploiting multiple direction information of individual dMRI voxels (Sotiropoulos et al., 2010; Reisert et al., 2011) become applicable to 3D-PLI data and, vice versa, pliODFs (or, alternatively, directional histograms) are suitable to validate various methods of tractography.

As pointed out by Hubbard and Parker (2009), "it is important to … test not only the ability of the tractography algorithms to track fibers from voxel to voxel, but to observe the details of the voxel-scale information and independently quantify the ability of dMRI to assess fiber orientation. The integrity of this orientation information is paramount to the validity of the reconstructed tract." 3D-PLI with its pliODFs particularly opens up the avenue to align with dMRI measurements by crossing the scales using common data formats, and to provide sub-voxel information on the underlying fiber architecture based on an independent technology complementary to the dMRI approach. Based on the pliODF generation, comparisons of 3D-PLI and dMRI can now be conducted at the level of individual voxels of the same size. Describing the local distribution of fiber orientations by means of spherical harmonics opens up the possibility to utilize methods being developed in the scope of computer vision to assess the similarity of datasets via shape descriptors and shape metrics, for example. This appears to be a promising approach to derive observer independent quality measures for 3D-PLI measurements, which will be evaluated in future studies.

CONCLUSIONS

The future of research about the brain connectome will depend on multiscale approaches validated by independent imaging methods applied to the same object simultaneously. We have successfully established a concept to bridge the spatial scales from microscopic fiber orientation measurements based on 3D-PLI to macroscopic dimensions by means of creating orientation distribution functions (pliODFs) from high-resolution vector data. With pliODFs the fusion and comparison with dMRI data becomes feasible even for whole human brains. The key-technology of supercomputing, that is inevitable for addressing real big data challenges as posed by the human brain, has

been applied to achieve this goal. Though, our implementation does not only limit the software's application to supercomputing environments, but it was also demonstrated to run successfully on local linux cluster systems. The established software package to generate pliODFs from 3D-PLI datasets described in this work will be made available through an ICT portal currently being developed by the human brain project consortium.

AUTHOR CONTRIBUTIONS

MA coordinated and substantially contributed to the conception and design of the study as well as to the analysis and interpretation of the 3D-PLI data. He drafted the manuscript. SS participated in the design of the study and was accountable for the mathematical description of the pliODFs and the software implementation. He helped to draft the manuscript. DG contributed to the interpretation of the 3D-PLI data and was accountable for the visualization of the directional histograms and the pliODFs. He helped to draft the manuscript. OB implemented the HPC-based analysis workflow of 3D-PLI data and generated the pliODFs for the whole human brain section. MD created and analyzed the simulated data set. JR conducted the 3D-PLI measurements and helped to compare the high- and low-resolution data sets. KZ contributed to the anatomical content of the study and helped with the interpretation of the pliODFs. KA contributed to the anatomical content and substantially assisted to the conception of the study. All authors read and revised the final manuscript and gave approval for publication.

REFERENCES

Aboitiz, F., Scheibel, A. B., Fisher, R. S., and Zaidel, E. (1992). Fiber composition of the human corpus callosum. *Brain Res.* 598, 143–153. doi: 10.1016/0006-8993(92)90178-C

Alexander, D. C. (2005). Multiple-fiber reconstruction algorithms for diffusion MRI. *Ann. N. Y. Acad. Sci.* 1064, 113–133. doi: 10.1196/annals.1340.018

Alexander, D. C., Barker, G. J., and Arridge, S. R. (2002). Detection and modeling of non-Gaussian apparent diffusion coefficient profiles in human brain data. *Magn. Res. Med.* 48, 331–340. doi: 10.1002/mrm.10209

Amunts, K., Bücker, O., and Axer, M. (2014)."Towards a multiscale, high-resolution model of the human brain," in *Brain-Inspired Computing*, vol. 8603, eds L. Grandinetti, T. Lippert, and N. Petkov (Cetraro: Springer International Publishing), 3–14. doi: 10.1007/978-3-319-12084-3_1

Amunts, K., Lepage, C., Borgeat, L., Mohlberg, H., Dickscheid, T., Rousseau, M. É., et al. (2013). BigBrain: an ultrahigh-resolution 3D human brain model. *Science* 340, 1472–1475. doi: 10.1126/science.1235381

Amunts, K., and Zilles, K. (2015). Architectonic mapping of the human brain beyond Brodmann. *Neuron* 88, 1086–1107. doi: 10.1016/j.neuron.2015.12.001

Assemlal, H.-E., Tschumperlé, D., Brun, L., and Siddiqi, K. (2011). Recent advances in diffusion MRI modeling: angular and radial reconstruction. *Med. Img. Anal.* 15, 369–396. doi: 10.1016/j.media.2011.02.002

Axer, M., Amunts, K., Gräßel, D., Palm, C., Dammers, J., Axer, H., et al. (2011a). A novel approach to the human connectome: ultra-high resolution mapping of fiber tracts in the brain. *Neuroimage* 54, 1091–1101. doi: 10.1016/j.neuroimage.2010.08.075

Axer, M., Gräßel, D., Kleiner, M., Dammers, J., Dickscheid, T., Reckfort, J., et al. (2011b). High-resolution fiber tract reconstruction in the human brain by means of three-dimensional polarized light imaging. *Front. Neuroinform.* 5:34. doi: 10.3389/fninf.2011.00034

Bastiani, M., and Roebroeck, A. (2015). Unraveling the multiscale structural organization and connectivity of the human brain: the role of diffusion MRI. *Front. Neuroanat.* 9:77. doi: 10.3389/fnana.2015.00077

Budde, M. D., and Frank, J. A. (2012). Examining brain microstructure using structure tensor analysis of histological sections. *Neuroimage* 63, 1–10. doi: 10.1016/j.neuroimage.2012.06.042

Bunge, H. J. (1982). *Texture Analysis in Materials Science: Mathematical Methods.* London, Boston, Sydney, Wellington, Durban, Toronto: Butterworths.

Caspers, S., Axer, M., Caspers, J., Jockwitz, C., Jütten, K., Reckfort, J., et al. (2015). Target sites for transcallosal fibers in human visual cortex - a combined diffusion and polarized light imaging study. *Cortex* 72, 40–53. doi: 10.1016/j.cortex.2015.01.009

Chung, K., Wallace, J., Kim, S. Y., Kalyanasundaram, S., Andalman, A. S., Davidson, T. J., et al. (2013). Structural and molecular interrogation of intact biological systems. *Nature* 497, 332–339. doi: 10.1038/nature12107

Costantini, I., Ghobril, J. P., Di Giovanna, A. P., Mascaro, A. L. A., Silvestri, L., Müllenbroich, M. C., et al. (2015). A versatile clearing agent for multi-modal brain imaging. *Sci. Rep.* 5:9808. doi: 10.1038/srep09808

Dell'Acqua, F., Simmons, A., Williams, S. C., and Catani, M. (2013). Can spherical deconvolution provide more information than fiber orientations? Hindrance modulated orientational anisotropy, a true-tract specific index to characterize white matter diffusion. *Hum. Brain Mapp.* 34, 2464–2483. doi: 10.1002/hbm.22080

Dohmen, M., Menzel, M., Wiese, H., Reckfort, J., Hanke, F., Pietrzyk, U., et al. (2015). Understanding fiber mixture by simulation in 3D polarized light imaging. *Neuroimage* 111, 464–475. doi: 10.1016/j.neuroimage.2015.02.020

Hess, C. P., Mukherjee, P., Han, E. T., Xu, D., and Vigneron, D. B. (2006). Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. *Magn. Res. Med.* 56, 104–117. doi: 10.1002/mrm.20931

Hubbard, P. L., and Parker, G. J. M. (2009). "Validation of tractography", in *Diffusion MRI*, eds K. Johansen-Berg and T. Behrens (Oxford: Elsevier), 353–375.

Jones, R. C. (1941). A new calculus for the treatment of optical systems. *J. Opt. Soc. Am.* 31, 488–493. doi: 10.1364/JOSA.31.000488

Leergaard, T. B., White, N. S., de Crespigny, A., Bolstad, I., D'Arceuil, H., Bjaalie, J. G., et al. (2010). Quantitative histological validation of diffusion MRI fiber orientation distributions in the rat brain. *PLoS ONE* 5:e8595. doi: 10.1371/journal.pone.0008595

Mardia, K. V., and Jupp, P. E. (2000). *Directional Statistics.* Chichetser, New York, NY: Wiley.

Menzel, M., Axer, M., De Raedt, H., and Michielsen, K. (2016). Finite-difference time-domain simulation for three-dimensional polarized light imaging. arXiv:1601.06290 [physics.med-ph].

Menzel, M., Michieslen, K., De Raedt, H., Reckfort, J., Amunts, K., and Axer, M. (2015). A Jones matrix formalism for simulating three-dimensional polarized light imaging of brain tissue. *J. R. Soc. Interface* 12, 1–14. doi: 10.1098/rsif.2015.0734

Miller, K. L., Stagg, C. J., Douaud, G., Jbabdi, S., Smith, S. M., and Behrens, T. E. (2011). Diffusion imaging of whole, post-mortem human brains on a clinical MRI scanner. *Neuroimage* 57, 167–181. doi: 10.1016/j.neuroimage.2011.03.070

Osten, P., and Margrie, T. W. (2013). Mapping brain circuitry with a light microscope. *Nat. Methods* 10, 515–523. doi: 10.1038/nmeth.2477

Palm, C., Axer, M., Gräßelel, D., Dammers, J., Lindemeyer, J., Zilles, K., et al. (2010). Towards ultra-high resolution fibre tract mapping of the human brain - registration of polarised light images and reorientation of fibre vectors. *Front. Hum. Neurosci.* 4:9. doi: 10.3389/neuro.09.009.2010

Rathi, Y., Michailovich, O., Shenton, M. E., and Bouix, S. (2009). Directional functions for orientation distribution estimation. *Med. Image Anal.* 13, 432–444. doi: 10.1016/j.media.2009.01.004

Reckfort, J., Wiese, H., Pietrzyk, U., Zilles, K., Amunts, K., and Axer, M. (2015). A multiscale approach for the reconstruction of the fiber architecture of the human brain based on 3D-PLI. *Front. Neuroanat.* 9:118. doi: 10.3389/fnana.2015.00118

Reisert, M., Mader, I., Anastasopoulos, C., Weigel, M., Schnell, S., and Kiselev, V. (2011). Global fiber reconstruction becomes practical. *Neuroimage* 54, 955–962. doi: 10.1016/j.neuroimage.2010.09.016

Silvestri, L., Paciscopi, M., Soda, P., Biamonte, F., Iannello, G., Frasconi, P., et al. (2015). Quantitative neuroanatomy of all Purkinje cells with light sheet microscopy and high-throughput image analysis. *Front. Neuroanat.* 9:68. doi: 10.3389/fnana.2015.00068

Sotiropoulos, S. N., Bai, L., Morgan, P. S., Constantinescu, C. S., and Tench, C. R. (2010). Brain tractography using Q-ball imaging and graph theory: improved connectivities through fiber crossings via a model-based approach. *Neuroimage* 49, 2444–2456. doi: 10.1016/j.neuroimage.2009.10.001

Tuch, D. S., Reese, T. G., Wiegell, M. R., Makris, N., Belliveau, J. W., and Wedeen, V. J. (2002). High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magn. Reson. Med.* 48, 577–582. doi: 10.1002/mrm.10268

Wedeen, V. J., Hagmann, P., Tseng, W. Y., Reese, T. G., and Weisskoff, R. M. (2005). Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magn. Reson. Med.* 54, 1377–1386. doi: 10.1002/mrm.20642

Zeineh, M. M., Palomero-Gallagher, N., Axer, M., Gräßel, D., Goubran, M., Wree, A., et al. (2016). Direct visualization and mapping of the spatial course of fiber tracts at microscopic resolution in the human hippocampus. *Cereb. Cortex.* bhw010. doi: 10.1093/cercor/bhw010. [Epub ahead of print].

Zilles, K., Palomero-Gallagher, N., Gräßel, D., Schlömer, P., Cremer, M., Woods, R., et al. (2015). "High-resolution fiber and fiber tract imaging using polarized light microscopy in the human, monkey, rat and mouse brain," in *Axons and Brain Architecture* ed K. S. Rockland (London: Elsevier Academic Press), 369–389.

3D Reconstructed Cyto-, Muscarinic M$_2$ Receptor, and Fiber Architecture of the Rat Brain Registered to the Waxholm Space Atlas

Nicole Schubert[1]*, Markus Axer[1]*, Martin Schober[1], Anh-Minh Huynh[1], Marcel Huysegoms[1], Nicola Palomero-Gallagher[1], Jan G. Bjaalie[2], Trygve B. Leergaard[2], Mehmet E. Kirlangic[1], Katrin Amunts[1,3] and Karl Zilles[1,4,5]

[1] Institute of Neuroscience and Medicine 1, Research Centre Jülich, Jülich, Germany, [2] Institute of Basic Medical Sciences, University of Oslo, Oslo, Norway, [3] C. and O. Vogt Institute for Brain Research, Heinrich-Heine University Düsseldorf, Düsseldorf, Germany, [4] Translational Brain Medicine, Jülich-Aachen Research Alliance, Aachen, Germany, [5] Department of Psychiatry, Psychotherapy, and Psychosomatics, RWTH Aachen University, Aachen, Germany

High-resolution multiscale and multimodal 3D models of the brain are essential tools to understand its complex structural and functional organization. Neuroimaging techniques addressing different aspects of brain organization should be integrated in a reference space to enable topographically correct alignment and subsequent analysis of the various datasets and their modalities. The Waxholm Space (http://software.incf.org/software/waxholm-space) is a publicly available 3D coordinate-based standard reference space for the mapping and registration of neuroanatomical data in rodent brains. This paper provides a newly developed pipeline combining imaging and reconstruction steps with a novel registration strategy to integrate new neuroimaging modalities into the Waxholm Space atlas. As a proof of principle, we incorporated large scale high-resolution cyto-, muscarinic M$_2$ receptor, and fiber architectonic images of rat brains into the 3D digital MRI based atlas of the Sprague Dawley rat in Waxholm Space. We describe the whole workflow, from image acquisition to reconstruction and registration of these three modalities into the Waxholm Space rat atlas. The registration of the brain sections into the atlas is performed by using both linear and non-linear transformations. The validity of the procedure is qualitatively demonstrated by visual inspection, and a quantitative evaluation is performed by measurement of the concordance between representative atlas-delineated regions and the same regions based on receptor or fiber architectonic data. This novel approach enables for the first time the generation of 3D reconstructed volumes of nerve fibers and fiber tracts, or of muscarinic M$_2$ receptor density distributions, in an entire rat brain. Additionally, our pipeline facilitates the inclusion of further neuroimaging datasets, e.g., 3D reconstructed volumes of histochemical stainings or of the regional distributions of multiple other receptor types, into the Waxholm Space. Thereby, a multiscale and multimodal rat brain model was created in the Waxholm Space atlas of the rat brain. Since the registration of these multimodal high-resolution datasets into the same coordinate system is an indispensable requisite for

multi-parameter analyses, this approach enables combined studies on receptor and cell distributions as well as fiber densities in the same anatomical structures at microscopic scales for the first time.

1. INTRODUCTION

Virtual high-resolution multiscale and multimodal 3D models of the brain are essential tools to visualize and understand the complex structural and functional organization of the brain. To capture different aspects of brain organization, such as the long-range fiber tracts connecting different brain regions, intracortical connectivity, and differences in molecular compositions, complementary neuroimaging techniques should be utilized. In order to interpret and compare measurements derived from different experimental techniques, all brain data sets should be integrated into a standard reference space.

This integrative approach leading to a multimodal and multiscale brain model is a major challenge because of the enormous structural complexity of the brain. The different brain regions differ not only by their cytoarchitecture, i.e., the varying densities of cells between the different layers within and between brain areas, but also by the expression of neurotransmitter receptors and gene expression. This microstructural diversity leads to a segregation of the cerebral cortex and the subcortical regions into hundreds of well definable entities with complex spatial arrangements (Toga et al., 2006; Zilles and Amunts, 2010; Amunts et al., 2013; Amunts and Zilles, 2015). Moreover, the different entities are connected by long range and short range fiber tracts, which also show an enormous spatial complexity (Zilles and Amunts, 2012). Therefore, an accurate definition of the spatial positions of structural entities is an indispensable requirement, particularly for multimodal and multiscale data sets. This is far from trivial (Bjaalie, 2002), because it often requires the registration of data collected from different brains, with different spatial resolutions, different dimensions of methodically introduced structural deformations and artifacts, and structures of considerable intersubject variability. Therefore, a multiscale and multimodal analysis must be based on an integration of the various data in a common stereotaxic brain atlas framework.

The International Neuroinformatics Coordinating Facility (INCF) Digital Atlasing Project created such a standardized framework, i.e., Waxholm Space, that operates as a connection point between miscellaneous rodent brain data. The Waxholm Space (WHS) is a common open access (http://software.incf.org/software/waxholm-space) 3D reference space based on high resolution magnetic resonance imaging (MRI) data anchored in a standardized spatial coordinate system. It also supports infrastructure for data exchange. The WHS of the mouse brain (Hawrylycz et al., 2011) was extended amongst others with neuroanatomic atlases (Goldowitz, 2010), gene expression databases and MRI and diffusion tensor imaging (DTI) (Johnson et al., 2010). Papp et al. (2014)

introduced and implemented the WHS atlas of the Sprague Dawley rat brain. The WHS rat brain atlas currently only contains high resolution MRI and DTI images, which served as basis for the delineation of the 79 major anatomical structures it depicts (Papp et al., 2014, 2015; Kjonigsen et al., 2015).

Aim of the present study was to complement the WHS rat brain atlas with information of cytoarchitecture, receptor expression and spatial orientation of fiber tracts. Thus, we processed entire postmortem brains of the Wistar rat for three different neuroimaging techniques: microscopic analysis of histological cell body stained serial sections for cytoarchitectonic analysis, *in vitro* receptor autoradiography to demonstrate muscarinic M_2 receptor density distributions (Zilles et al., 2002; Zilles and Amunts, 2010), and 3D Polarized Light Imaging (PLI) for high resolution visualization of fiber tracts (Axer et al., 2011a,b). Imaging of the cell body stained sections enables a precise microscopical identification of cytoarchitectonically definable areas and the visualization of the spatial distribution of neurons. Quantitative *in vitro* receptor autoradiography is a well established technique to visualize the topographically heterogeneous distribution of neurotransmitter receptors, the key molecules of signal transmission. 3D PLI has been introduced recently and has opened up new avenues to analyze the complex architecture of nerve fibers and fiber tracts in postmortem brains at a microscopic resolution.

All of the above mentioned techniques require brain sectioning and mounting on glass slides, and this approach results in a loss of spatial alignment between neighboring sections. To obtain 3D brain models, the sections have to be aligned, the artificial deformations must be corrected, and 3D reconstructions must be performed. Therefore, it was necessary to improve currently available registration algorithms and adapt them to the specific requirements inherent to images obtained from receptor autoradiography and PLI. Finally, the nature of the data acquired in the present study also required the development of a novel registration strategy which enables integration of large scale high-resolution images of into the 3D MRI volume of the WHS rat atlas.

The here provided data complementing the Waxholm Space rat brain atlas will provide a multiscale and multimodal rat brain model enabling for the first time combined studies on receptor and cell distributions as well as fiber densities in the same anatomical structures at microscopic scales. Furthermore it will be publicly accessible through the Human Brain Project (HBP) portal, intended for multi-parameter analyses, refinement of the atlas labels, or further expansion via the proposed registration strategies.

2. MATERIALS AND METHODS

2.1. Tissue Processing and Image Acquisition

2.1.1. Tissue Sectioning and Blockface Imaging

All animal procedures were approved by the institutional animal welfare committee at the Research Centre Jülich, and were in accordance with European Union (National Institutes of Health) guidelines for the use and care of laboratory animals. The brain of one adult male Wistar rat was used for the visualization of cell bodies and of muscarinic M_2 receptors. It is referred to as the receptor brain. The brain of a second adult male Wistar rat, which we refer to as the PLI brain, was processed for the visualization of nerve fibers and fiber tracts. The receptor brain was immediately deep frozen in isopentane at $-50°C$ and serially sectioned in the coronal plane at $20 \mu m$ thickness using a cryostat microtome (Leica Microsystems, Germany). The ensuing 1362 sections were thaw-mounted onto glass slides and organized in series of adjoining triplets of which one section was used for visualization of cell bodies, and the other two sections were processed for quantitative in vitro receptor autoradiography. The PLI brain was immersion fixed in 4% buffered formaldehyde. After two cryoprotection steps (10% glycerin for 3 days, followed by 20% glycerin for 14 days at $+4°C$), the brain was deep frozen in isopentane at $-50°C$ and serially sectioned in the coronal plane at $60 \mu m$ thickness using the same cryostat microtome (Leica Microsystems, Germany). The 446 ensuing sections were placed on glass slides and stored at $-80°C$ in airtight plastic bags until further processing. They were thaw mounted and coverslipped with 20% glycerin the day before image acquisition took place. During sectioning of both brains, blockface images of every section were taken with a CCD camera (AVT Oscar F-810 C, 3272 × 2469 pixels, $15 \mu m \times 15 \mu m$, RGB) which was installed vertically above the cryostat, in order to obtain undistorted reference images. Spatial resolution in the z-direction was $20 \mu m$ for images obtained from the receptor brain, and $60 \mu m$ for images obtained from the PLI brain. A total of 1361 blockface images were taken for the receptor brain, and 446 for the PLI brain.

2.1.2. Receptor Brain

A total of 452 sections from the receptor brain were stained with a silver staining technique after (Merker, 1983). It results in a staining of all cell bodies, which is different from the widely used cresyl-violet stain of Nissl substance by its higher contrast and more intense visualization of cytoarchitecture. The sections processed for quantitative in vitro receptor autoradiography were used to demonstrate the densities (in fmol/mg protein) of two different receptor binding sites of the cholinergic muscarinic M_2 receptor, i.e., the agonist and the antagonistic binding site, according to previously published protocols (Zilles et al., 2002; Palomero-Gallagher et al., 2013). Sections were incubated with 1,7 nM 3H-oxotremorine-M (PerkinElmer, USA) or with 5 nM 3H-AF-DX 384 (PerkinElmer, USA) to visualize the agonistic and antagonistic binding sites of the M_2 receptor, respectively. Binding assays were preceded by a preincubation in the respective buffer to eliminate the endogenous transmitters

and finalized by a washing step. The labeled sections were exposed together with plastic scales of increasing and known radioactivity concentrations against beta-radiation (tritium) - sensitive films, which were developed after 15 weeks.

The ensuing 430 autoradiographs of the agonistic binding site of the M_2 receptor as well as the 452 cell body stained histological sections were digitized using a high resolution camera (Zeiss) with an in-plane resolution of $5 \mu m \times 5 \mu m$ (4164 × 3120 pixels, RGB). Since each of these sections was $20 \mu m$ thick and sections had been organized into triplets, the resulting spatial resolution in the z-direction was $60 \mu m$ for both the digitized autoradiographs and the digitized histological sections. For further details of quantification of receptor density in fmol/mg protein and color coding see Zilles et al. (2002).

2.1.3. PLI Brain

The 446 sections from the brain cut at $60 \mu m$ were used to acquire 3D-PLI data reflecting the fiber architecture in gray and white matter regions (cf. Axer et al., 2011b; Dohmen et al., 2015; Menzel et al., 2015; Reckfort et al., 2015 for technical details). Briefly, 3D-PLI utilizes the optical birefringence of brain tissue, which is basically induced by the optical anisotropy of myelin sheaths wrapped around axons. By passing linearly polarized light through brain sections and by detecting the local changes in the polarization state of light, a 3D description of the underlying fiber architecture is derived. The imaging system used is a polarimeter. The sections were successively scanned with a large-area polarimeter (LAP), and subjected to an analysis workflow, which comprises calibration, independent component analysis, polarization analysis and calculation of fiber orientation maps (FOMs). FOMs are the fundamental data structure provided by 3D-PLI and have an in-plane resolution of $64 \mu m \times 64 \mu m$, and, since each section was $60 \mu m$ thick, a spatial resolution in the z-direction of $60 \mu m$. They contain a single 3D fiber orientation vector per voxel, that is interpreted as the spatial orientation of the fibers in this voxel.

2.2. 3D Reconstruction

2.2.1. 3D Reconstruction of Blockface Images

Non-linear deformations introduced by brain sectioning, mounting and staining were corrected using blockface images as undistorted references for the spatial alignment of histological, autoradiographic and PLI images. Hence, in a first step the blockface images had to be 3D reconstructed. In short, the here applied robust and efficient reconstruction method consisted of a two-phase registration: a marker-based alignment of the blockface images and a refinement of the pre-reconstructed volume using 3D information. First, the coordinates of markers (circles) labeled on the microtome chuck were extracted. The centers of the circles of neighboring images were aligned to each other by means of a translation transformation. Processing all images leads to an almost smoothly reconstructed 3D stack of blockface images of the brain. However, this approach causes perspective errors due to the different heights of the sectioning plane and microtome chuck with the markers, and thus their different distances to the camera lens. Therefore, in the second part of the method the median along the z-direction of the

marker-based reconstructed blockface volume was calculated to eliminate the outliers caused by perspective errors. The number of the images used by the voxelwise computation of this median volume can be specified by the radius of the median. In a next step, the marker-based reconstructed volume was aligned slice-by-slice onto the median volume using a translation transform estimated by an intensity based image registration algorithm with sum of squared differences as metric. By using this technique we took advantage of 3D information in an actually 2D slice-by-slice registration method. This led to an accurately aligned volume of blockface images that was an important reference to recover the spatial coherence of the non-linearly deformed sections corresponding to the blockface images. The procedure of 3D reconstruction of blockface images was introduced by Schober et al. (2015) with modified markers for the reconstruction. Finally, the reconstructed blockface volumes of the receptor and PLI brains were separated from the surrounding by means of a 3D watershed algorithm. The 3D reconstruction was carried out separately for images obtained from the receptor brain and for those of the PLI brain. Thus, we obtained two distinct blockface coordinate spaces (BCS): the BCS of the receptor brain (BCS^R), and the BCS of the PLI brain (BCS^{PLI}).

2.2.2. 3D Reconstruction of Cyto- and Receptor Architecture Images

After reconstruction of the blockface volume of the receptor brain, each histological and autoradiographic image was aligned to its corresponding blockface image. Due to the highly different information each modality comprises (i.e., cell body distribution patterns vs. M_2 receptor densities), it was necessary to establish different registration strategies. Each histological section was rigidly aligned with its corresponding blockface image. The centers of gravity of the brain tissue in blockface and histological image were calculated and superimposed. To determine the center of gravity, the separation of the brain tissue from the background is required. This was done by means of thresholding, extracting the largest connected component and morphological operations. After alignment of the centers of gravity a brute force optimizer tested all rotation angles with the sum of squared differences as metric. Details are described in Schubert et al. (2015). Reconstruction of the autoradiographic images is considerably more challenging due to the fact that receptors of a given type, in our case the muscarinic M_2 receptor, are not necessarily expressed in all brain regions. Furthermore, when present, they can occur at very different concentrations throughout the brain. Thus, the range of gray values present in an autoradiographic image is much larger than that of a histological section. Therefore, the registration has to compensate for "empty regions," i.e., regions without information in the images because that part of the brain does not express the receptor in question. First of all, an intra-stack registration matched consecutive autoradiographic images by means of a scale-invariant feature transform algorithm that detected characteristic points in the images based on their gradient information. Afterwards, these points were rigidly aligned by minimizing the Euclidean distance between them. With this pre-registered autoradiographic volume we were able

to use a landmark based method to align the autoradiographic images to their corresponding blockface images. In every 30th image anatomical landmarks were manually set and the rigid transformation between these landmarks was calculated. Between these 30 images the transformations were interpolated and applied to the autoradiographic images (Huynh et al., 2015). Data acquisition and 3D reconstruction of the cyto- and receptor architecture are illustrated in **Figure 1**. At the end of this procedure, the 3D reconstructed histological volume and the 3D reconstructed M_2 receptor volume were each in the BCS^R.

2.2.3. 3D Reconstruction of Fiber Architecture Images

The 3D reconstruction of the PLI data consists of two steps: a rigid slice-by-slice registration of the PLI images to the corresponding blockface images and a non-rigid refinement method. The first step is based on estimating a transformation of the PLI images to the corresponding image of the reconstructed blockface volume by image registration. To align the PLI images to the blockface images, the masks of the brain tissue of both data sets are required. For that, the reconstructed blockface volume was segmented by means of a 3D watershed algorithm and the PLI images were manually segmented. Using the segmented images the centers of gravity of the corresponding brain masks were calculated and aligned. Based on this initial transformation, an intensity based rigid registration was performed using mutual information as metric. The second step, the refinement, was done by means of a slice-by-slice B-Spline registration with sum of squared differences as metric and a grid size of 5×6. At the end of this procedure, the 3D reconstructed fiber volume was in the BCS^{PLI}.

2.3. Data Intergration into a Reference Space

2.3.1. Waxholm Space Atlas

The WHS atlas of the Sprague Dawley rat brain is an open access atlas based on a high resolution MRI and DTI template in which both WHS and stereotaxic coordinates are defined. The T2*-weighted anatomical MRI ($512 \times 1024 \times 512$ pixels) with an isotropic spatial resolution of $39 \, \mu m$ was acquired ex vivo by means of a 7T small animal MRI system. Anatomical delineations in the atlas are based on image contrast observed in T2*-weighted images and diffusion tensor images. Technical details are described in Papp et al. (2014). The latest version of the atlas contains 79 structures with new and updated delineations of the hippocampal formation and parahippocampal region, as described in Kjonigsen et al. (2015). The atlas is available from the INCF Software Center (http://software.incf.org/software/waxholm-space-atlas-of-the-sprague-dawley-rat-brain).

2.3.2. Data Integration into the Waxholm Space Atlas

In order to achieve an accurate analysis of the multimodal data sets, we aligned the atlas data to the coordinate space of each reconstructed data sets, i.e., to the BCS^R and the BCS^{PLI} by means of an advance image registration. In the literature several methods were proposed to integrate brain sections into 3D data volumes. Strategies which rely on successively increasing the degrees of freedom of the transformation demonstrated
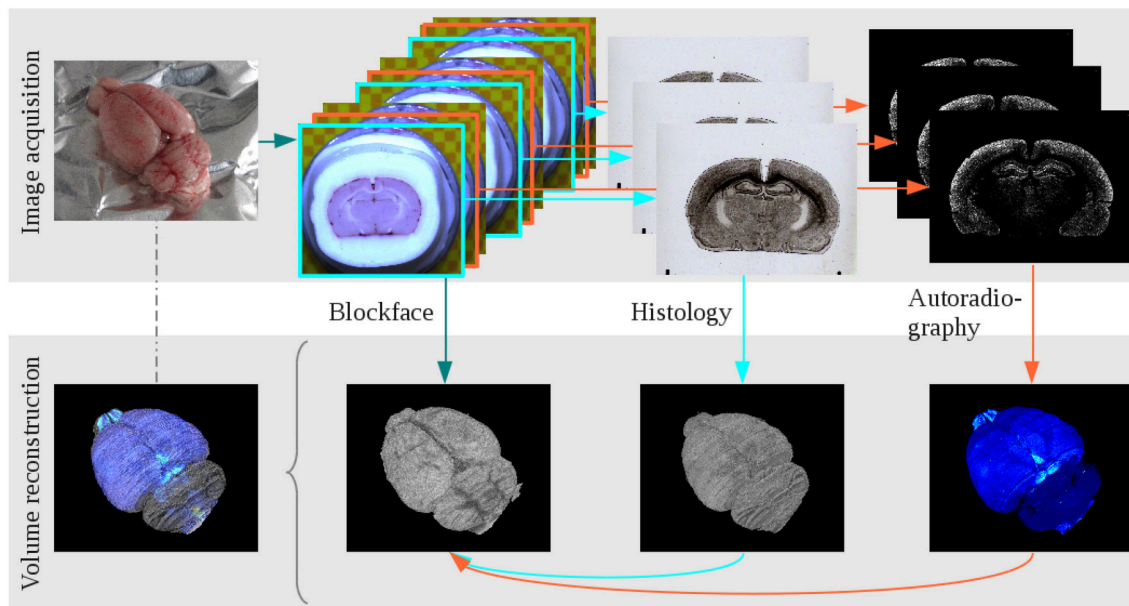
FIGURE 1 | The main steps of data acquisition and 3D reconstruction at the example of the receptor brain. Image acquisition: The brain was removed from the head, blockface photographs were taken prior to sectioning, alternating sections were either used for cell body staining or receptor autoradiography. Volume reconstruction: The blockface sections were reconstructed, the 3D reconstructed blockface volume served as undistorted reference for the reconstruction of histological and autoradiographical sections. All modalities combined build a multimodal 3D model.

the best results. For instance, Dauguet et al. (2007) and Li et al. (2009) suggested a two step procedure consisting of a rigid transformation followed by non-linear transformations. However, this approach has several drawbacks, since a rigid transformation only aligns brain orientation, so that the brains after this transformation still differ in size and shape. Non-linear registrations work locally and, therefore, need well aligned volumes as starting point, which a rigid transformation cannot garantee. Lebenberg et al. (2010) added an affine transformation between the rigid and non-linear transformations. This step is also important for our data, due to the fact that the affine transformation is able to align the size and shape of the brains by scaling and shearing. However, since Lebenberg et al. (2010) not only registered a single mouse hemisphere, but also excluded the olfactory bulb and cerebellum, they were not confronted with the challenges posed by trying to register the whole brains. Therefore, modifications of their strategy are essential to enable accurate registration of structures such as the olfactory bulbus, or of even of the hemispheres, since they are independent of each other at levels rostral of (or caudal to) the corpus callosum. We used for all three transformations (rigid, affine and non-linear) a pyramidal method, i.e., a coarse to fine approach, to align initially large structures followed by aligning small and fine structures, whereas Lebenberg et al. (2010) only applied this multi-resolution approach to the affine transformation. Furthermore, we computed the similarity of the brains by means of Mutual Information, which is the best metric for multimodal registration tasks (Rueckert et al., 1999), in all three transformation steps, while Lebenberg et al. (2010) used Correlation Coefficient in the affine registration step. Finally, we

studied the influence of the grid spacing used for the non-linear transformation to achieve best possible results for the whole brains, which was not tested by Lebenberg et al. (2010).

The T2*-weighted atlas MRI was aligned to the respective reconstructed blockface volume. The estimated transformation was then applied to the digital atlas delineations. To compensate the variability between the atlas data and the rat brain data sets (i.e., the 3D reconstructed histological, M_2 receptor and fiber volumes), the registration strategy consists of the successive steps explained below (**Figure 2**). Initially, image parts belonging to the rat brain of the MRI data set were separated from the rest of the head using the atlas template. All voxels in the MRI with their corresponding atlas label unequal to 0 were marked as "brain," the other voxels were marked as background. The resulting masked MRI containing only the brain volume was used for the registration. Note that it was necessary to perform two separate registrations, since we have reconstructed blockface images in two separate spaces, namely BCS^R and BCS^{PLI}. The procedure is the same in both cases and encompasses the following registration steps: The T2*-weighted MRI was manually aligned to a few selected images from the reconstructed blockface volume using the anchor a custom tool for affine registration of histological images to brain atlas space (Moene et al., 2011; Papp etal., 2016) in Navigator3. Then, the transformations were automatically propagated to the remaining images. These steps can be iterated with different parameters to prioritize specific boundaries or structures. Afterwards, the manually aligned MRI was automatically re-oriented to match the spatial orientation of the blockface volume. A global 3D affine registration initialized with the previously computed parameters
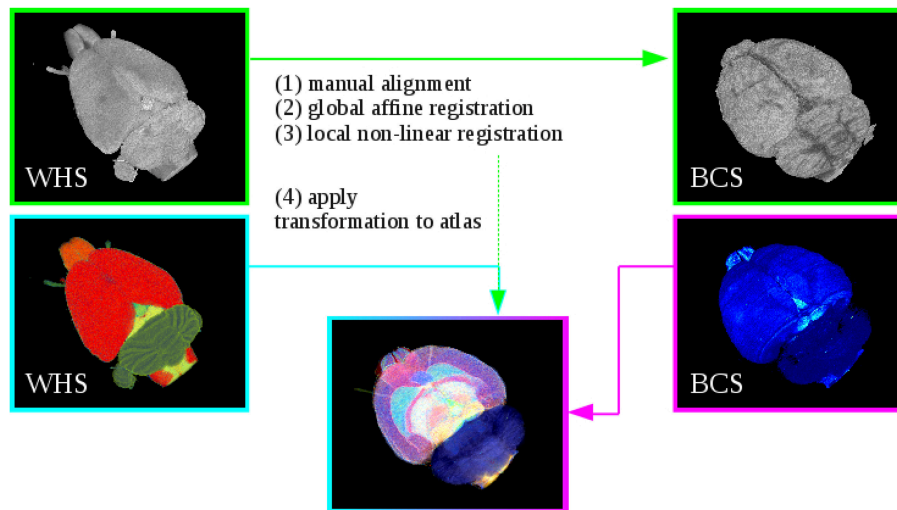
FIGURE 2 | Registration of Waxholm Space atlas with rat brain data in blockface coordinate space (BCS) at the example of the receptor brain. Steps (1)–(3) performs the alignment of the WHS MRI to the blockface volume. Step (4) applies the estimated transformation from Steps (1)–(3) to the atlas. This yields the 3D atlas model of multimodal rat brain data.

was then optimized with normalized mutual information as similarity measure. To speed up the registration and prevent local minima a coarse to fine multiresolution approach was used which consisted of six levels of Gaussian smoothing pyramids. Finally, a 3D non-linear registration based on cubic B-Splines was used to refine the former parameters and local discrepancies. Again, normalized mutual information was chosen as similarity measure, and a pyramidal approach was used. The manual anchoring as well as the automatic 3D affine and non-linear transformations were directly applied to the atlas template with one exception. Instead of cubic B-Spline interpolation we used nearest neighbor interpolation to preserve accurate label boundaries and avoid gaps. We employed elastix (Klein et al., 2010) for the 3D registration. After registration to the respective blockface volumes, the resulting volume dimensions of the MRI and atlas template equal the dimensions of the blockface volume coordinate space: $996 \times 1356 \times 1361$ pixels with a spatial resolution of $15\,\mu m \times 15\,\mu m \times 20\,\mu m$ for the receptor brain (BCS^R), and $588 \times 723 \times 446$ pixels with a spatial resolution of $22\,\mu m \times 22\,\mu m \times 60\,\mu m$ for the PLI brain BCS^{PLI}.

2.3.3. Evaluation of the Registration Results

Qualitative evaluation of the registration results in terms of anatomical accuracy was done by superimposing the aligned MRI and the blockface volume to compare the external borders and internal structures, as well as superimposing the atlas contours with the cyto-, muscarinic M_2 receptor and fiber architecture.

Quantitative evaluation was performed by computing the quality of the alignment of atlas based segmented and receptor or PLI based segmented structures. Two structures which are defined in the WHS atlas were used: the pial brain surface, and the hippocampal formation. In our datasets different strategies were used to generate the pial surface contour and that of the hippocampal formation: The surface of the entire brains was

segmented with a 3D watershed algorithm, and the hippocampal formation was manually delineated on the original images by two of the coauthors (NP-G and KZ). A comprehensive evaluation suggests the use of three uncorrelated measures that cover different aspects of the segmentation: an overlap measure (e.g., Dice coefficient), the Hausdorff distance and the average surface distance (Handels, 2000; Heimann et al., 2004). The Dice coefficient (DC) (Dice, 1945) assesses the overall overlap of the segments, it is sensitive to misplacement of the segments, but gives less weight to outliers. The average surface distance (ASD) and the Hausdorff distance (HD) determine the discrepancy of the surface of the segments. ASD is defined as the average error of all distances. A small ASD indicates a small error and variance between the segments. The HD returns the maximum distance between the segments, and therefore the maximum error. It is sensitive to outliers. The measurements were determined between an atlas based segmented structure A and the corresponding receptor or PLI based segmented structure B. The DC calculates the spatial overlap accuracy of two segmented structures A and B, whereby 0 is the result of disjunct segments and 1 is the result of a perfect agreement of the segments. With $|\cdot|$ denotes the number of voxels in the respective segmented structure, the DC is:

$$DC(A, B) = \frac{2|A \cap B|}{|A| + |B|} \tag{1}$$

The ASD determines the minimal distance in mm of one segment to the other and vice versa. This value is 0 for a perfect registration.

$$ASD(A, B) = \frac{\sum_{a \in A} \min_{b \in B} d(a, b) + \sum_{b \in B} \min_{a \in A} d(b, a)}{|A| + |B|} \tag{2}$$

The HD is defined as the maximum distance in mm of a segment to the nearest point in another segment and vice versa. A low HD indicates a good match.

$$HD(A, B) = \max(h(A, B), h(B, A)) \text{ with } h(A, B)$$
$$= \max_{a \in A} \min_{b \in B} d(a, b) \qquad (3)$$

with the Euclidean distance d between point a and b.

Note, that the aim of this procedure was not to prove that our definition of pial surface or hippocampal formation is better than that of the WHS atlas. We only wanted to determine whether the overlap could be improved (i.e., differences could be reduced) with different algorithms.

2.4. Hard- and Software

The processing was partially done using high-performance computing tools and supercomputing facilities of the Jülich Supercomputing Centre, Germany [Juelich Dedicated GPU Environment (JuDGE)], as well as the in-house Solaris computer cluster. Custom C++ software programs using ITK, `elastix`, OpenCV, MPI, OpenMP, QT and OpenGL performed the 3D reconstruction of the postmortem rat brains, the data integration into the WHS atlas, and furthermore the evaluation and the visualization of the results.

3. RESULTS

The registration of the atlas MRI volume to the respective blockface volume was done in three subsequent steps: rigid, affine and non-linear B-Spline based registration. All three steps used the Adaptive Stochastic Gradient Descent approach for optimization. A multi-resolution registration with six levels was used to overcome local minima problems. The brain volumes were downsampled by a factor of 2 compared to the next resolution level. The similarity of the intensity values of blockface and MRI data was determined with the Mutual Information metric, which was specifically developed for multimodal data sets (Viola and Wells, 1997). As expected, the matching of brain structures from the different modalities improved considerably with increasing degrees of freedom of the transformations. Considerable differences were found after the rigid registration. The affine registration improved the matching, but still differences existed. The application of the non-linear registration led to a high matching quality. It took 39 min for the PLI brain (rigid 5 min 16 s, affine 5 min 16 s, B-Spline 28 min 30 s) and 66 min for the receptor brain (rigid 11 min 33 s, affine 12 min 8 s, B-Spline 42 min 18 s).

3.1. Qualitative Evaluation

The results after each step are illustrated as checkerboard images of blockface and MRI volumes in three orthogonal views (coronal, horizontal and sagittal, c.f. **Figure 3**). Regarding the differences in size and shape of the brains, it was recognizable that these differences nearly disappeared from rigid, affine to non-linear registration transformation. Depending on the actual registration method, considerable to minor differences could be easily detected at three sites: the outer surface of the entire brain,

the olfactory bulb and the cerebellum. The blockface volume of the entire brain was wider, that of the olfactory bulb was deflected and laterally displaced, and the cerebellum of the blockface volume was more flattened compared to the MRI volume. The difference between the outer surfaces (**Figure 3**, green circles) decreased after affine registration and disappeared after the non-linear registration. The position of cerebella (**Figure 3**, blue circles) and olfactory bulbs (**Figure 3**, yellow circles) considerably differed between blockface and MRI volumes. This could not be compensated by linear and global transformations. Only non-linear registration fitted these structures. In conclusion, rigid transformation sufficiently centered the brains, affine registration compensated differences in size and shape of the brains, and finally, non-linear registration aligned both small local mismatches and also large differences (c.f. olfactory bulb and cerebellum).

The cytoarchitectonic, M_2 receptor distribution, and fiber orientation volumes were superimposed with the atlas to demonstrate the quality of the match between the reconstructed and the atlas volumes (**Figures 4, 5**).

In both brains, the overall matching quality was high, particularly at the anterior commissure (**Figure 4**, magenta arrow) and the corpus callosum (**Figure 5**, green arrow). In the receptor brain, the artificial gap between the hemispheres was caused by the removal of the unfixed brain from the skull, which resulted in an anti-clockwise rotation and lateral displacement of the hemispheres and thereby local mismatches at the mesial cortical surface, the border between the retrosplenial cortex and the underlying white matter, and the medial protrusion of the neocortex in direction to the hippocampus (**Figure 4**, yellow arrows). The differences between the position and shape of the cerebella in the atlas and the reconstructions led to further mismatches (**Figure 4**, red arrows), which were not compensated by the registration. In the PLI brain, minor mismatches were found at the outer surface of the brain (**Figure 5**, red arrows) and the cerebellum (**Figure 5**, yellow arrow). The better match of the PLI brain is understandable, because the receptor brain was not fixed and therefore, more prone to distortions whereas the PLI brain was fixed before deep freezing.

3.2. Quantitative Evaluation

The quality of alignment between the atlas and the reconstructed volumes was estimated by comparing the topography of the surface of the entire brain and of the hippocampal complex using three different measures introduced in Section 2.3.3. The surface of the entire brains was segmented with a 3D watershed algorithm. The atlas labels were used to extract the surface of the entire brain and hippocampal complex. The outer contour of the hippocampal complex was manually traced in the M_2 and the FOM sections by experienced neuroanatomists. The hippocampal complex comprises the Cornu Ammonis regions 1, 2, and 3, the dentate gyrus and the subicular complex with the subiculum, presubiculum and parasubiculum. The hippocampal complex spans a wide portion of the brain in both the rostrocaudal and dorso-ventral directions, and can be used to demonstrate the registration quality for inner anatomical structures. Since the
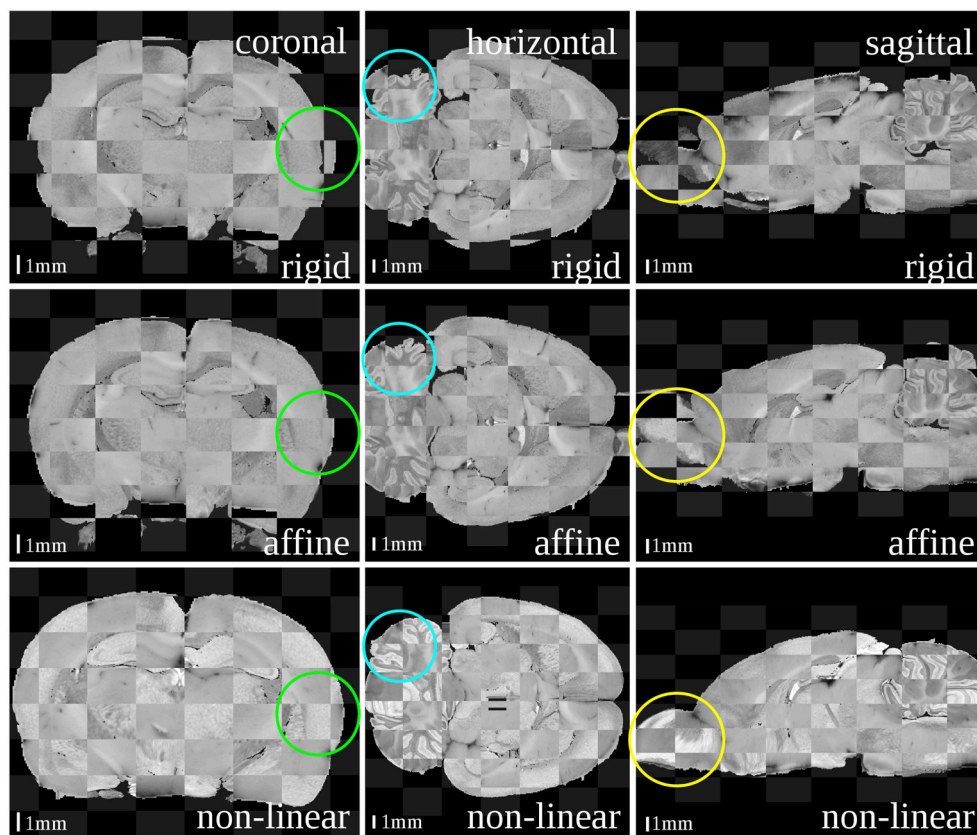
FIGURE 3 | Checkerboard images of the blockface volume of the receptor brain and the T2*-weighted atlas MRI template after each registration step: rigid (first row), affine (second row) and non-linear (third row). Considering size and shape of the brain and especially its outer surface (green circles), cerebella (blue circles) and olfactory bulbs (yellow circles) the matching quality increased with increasing degree of freedoms of the transformations.
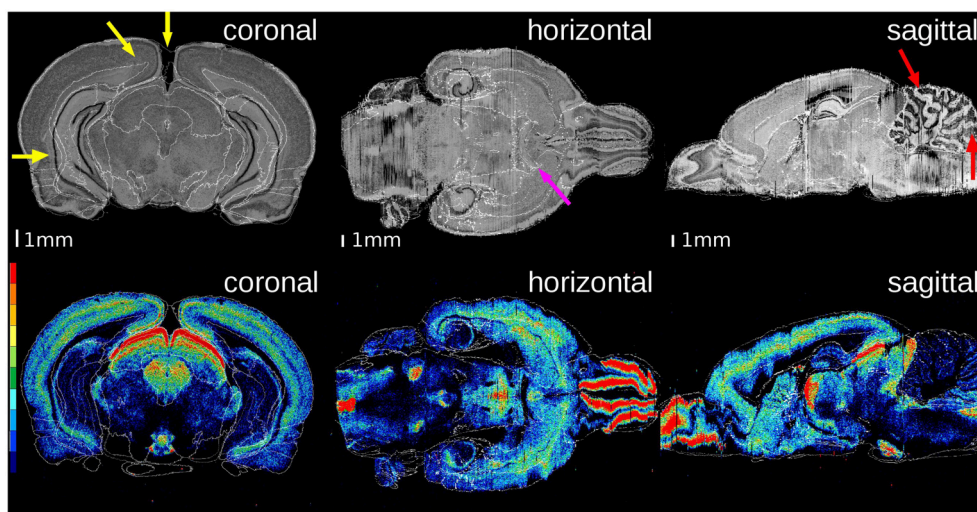
FIGURE 4 | Superimposition of the atlas structures (white contours) on the histological volume of the receptor brain (upper row) and on the M_2 receptor density distribution volume (lower row). The color legend of the lower row denotes the distribution of M_2 receptor densities (red: high, black: low). The overall matching quality is high, especially at the anterior commissure (magenta arrow). The gap between the hemispheres results in small mismatched boundaries (yellow arrows). The high differences of the cerebella in location and shape yields some small discrepancies (red arrows).
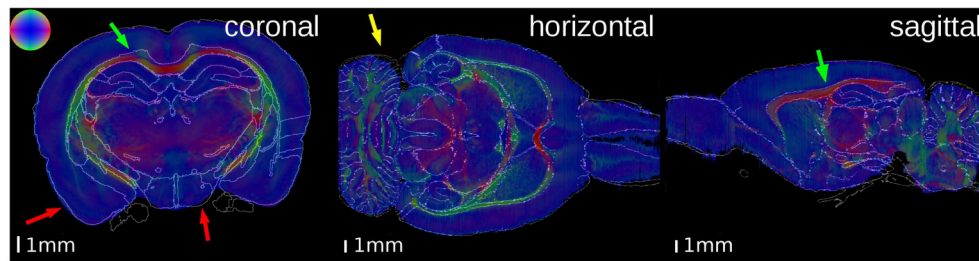
FIGURE 5 | Superimposition of the contours of the atlas delineations on the Fiber Orientation Maps (FOMs) of the PLI brain. The color sphere indicates the direction of the fiber orientation. A high matching quality is observable, especially at the corpus callosum (green arrows). Some small discrepancies are visible at the outer surface of the brain (red arrows) and the cerebellum (yellow arrow).

spatial location of the olfactory bulbs differed between MRI and blockface volumes, and caused far-reaching effects on the registration quality of many other brain structures, particularly at rostral levels, the quantitative evaluation was carried out with registrations with or without the olfactory bulbs. An important registration parameter was the spacing of the control points in the B-Spline grid, which indicates the flexibility of the transformation. A low spacing guaranties a high flexibility of the transformation.

The results of the comparisons after each registration step using different spacing of the control points are shown in **Figure 6**. Using the Dice coefficient, the affine registration was sufficient to reach a high matching of the entire brain and the hippocampal complex well above 0.7, which is commonly accepted as a limit for a good match (Zijdenbos et al., 1994). However, the measures improved significantly after the non-linear B-Spline registration. The Dice coefficient and the average surface distance reached their optimum at middle flexibility of the transformation grid for the entire brain comparison. The Hausdorff distance reached an optimum at high flexibility. This is caused by the fact that the artificial displacements of the olfactory bulbs were compensated. However, this led to undesired transformations of brain structures at rostral levels. That was also reflected in the quantitative evaluation of the hippocampal complex. Here, the best results were achieved by the relatively less flexible affine registration (receptor brain), or after a non-linear registration at low flexibility levels (PLI brain). Comparing the results after registration without inclusion of the olfactory bulbs, the measures improved significantly for the hippocampal complex in the receptor brain. In contrast to the receptor brain the results between the registration with or without the olfactory bulbs did not considerably differ. In **Table 1** the results are summarized.

4. DISCUSSION

The study aimed at integrating multimodal (i.e., cytoarchitectonic, muscarinic M_2 receptor distribution and fiber orientation data) and multiscale (i.e., mesoscopic resolution of blockface images and MR data of the 3D digital WHS Atlas, and microscopic resolution of sections) data in a common

stereotaxic reference space. This was achieved by linear and non-linear registration. The qualitative and quantitative evaluations demonstrated a good matching of all data sets. We selected the whole brain, and additionally the hippocampal complex, that spans a wide distance within the brain in both the rostro-caudal and dorso-ventral directions, as examples to prove the quality of the methods of registration.

4.1. Methodic Challenges
4.1.1. 3D Reconstruction
The 3D reconstruction of rodent brains is often carried out by means of rigid or affine registration transformations guided by blockface images (Ourselin et al., 2000; Lebenberg et al., 2010) or by volumes obtained from MRI (Li et al., 2009; Yang et al., 2012). We used blockface images for the reconstruction, due to the fact that they provide largely undistorted reference images of the brain sections. Furthermore, a reconstructed blockface volume is an excellent reference template, particularly if a real 3D volume (e.g., MRI volume) is missing or the resolution of the respective MRI volume is not appropriate (Schober et al., 2015). A particular challenge of the 3D reconstruction was the differential deformation of the brains inevitably caused by different tissue processing techniques. While the PLI brain was fixed and deep frozen before sectioning, the receptor brain was just deep frozen to maintain the receptor architecture. The fixation and deep freezing of the PLI brain introduces less deformations compared to the native brain size and shape than the deep freezing of the unfixed receptor brain (e.g., location of olfactory bulb). The sectioning procedure and mounting of sections from fixed brains also results in less deformations than that of sections from unfixed brains. To compensate these deformations, we performed, in addition to the common linear reconstruction strategies, non-linear transformations, which results in largely good reconstruction results. However, the anti-clockwise rotation of the hemispheres during the mounting procedure together with the more fragile nature of these unfixed cryostat microtome sections could not be completely eliminated at some sites (**Figure 4**, yellow arrows). A further challenge was the different information provided by the different modalities, e.g., "empty regions" (c.f. Section 2.2.2) in sections of the receptor brain. Therefore, particular reconstruction strategies were necessary for each modality,

FIGURE 6 | Diagrams of the quantitative analyses of the Dice coefficient (DC), Hausdorff distance (HD) and Average Surface distance (ASD) from rigid, affine to non-linear transformations, whereby numbers 20–100 describes the flexibility of the grid (20 -high flexibility, 100-low flexibility). The upper row illustrates the registration results of the receptor brain and the lower row demonstrates the registration results of the PLI brain. Diagrams in the first column indicate the results of the analyses of the whole brains (wb), the diagrams in the second column indicate the results of the analyses of the hippocampal complex (hc). The last two columns show diagrams of the analyses of the wb (3rd column) and the hc (4rd column) after a registration excluding olfactory bulb (ob). The best results of the quantitative and the qualitative analyses are marked with yellow rectangles.

TABLE 1 | The quantitative evaluation is measured by Dice coefficients (DC), Hausdorff distances (HD) and average surface distances (ASD) between receptor and PLI based and atlas based segmented structures, considering the whole brain and an internal structure, the hippocampal complex.

			Cyto- and receptor architecture			Fiber architecture		
			Rigid	**Affine**	**Non-linear**	**Rigid**	**Affine**	**Non-linear**
Entire brain	Whole brain	DC	0.852	0.887	**0.981**	0.898	0.919	**0.976**
		HD	3.112	3.331	**1.000**	2.291	2.138	**1.139**
		ASD	1.085	0.952	**0.592**	0.810	0.767	**0.588**
	Hippocampal complex	DC	0.642	**0.776**	0.75	0.757	0.783	**0.84**
		HD	1.158	**0.92**	1.553	1.279	**0.877**	0.982
		ASD	0.332	**0.226**	0.238	0.235	0.186	**0.137**
Brain without olfactory bulb	Whole brain	DC	0.848	0.907	**0.978**	0.912	0.932	**0.971**
		HD	2.853	2.572	**1.124**	2.348	2.141	**0.946**
		ASD	1.208	0.929	**0.645**	0.808	0.788	**0.661**
	Hippocampal complex	DC	0.655	0.766	**0.824**	0.768	0.795	**0.841**
		HD	1.177	1.010	**0.961**	1.205	0.961	**0.956**
		ASD	0.323	0.235	**0.163**	0.224	0.174	**0.135**

The Dice coefficient ranges between 0 and 1, 1 indicates full overlap. Lowest Hausdorff distance and average surface distance values indicate best alignments. The best results are labeled bold.

e.g., landmarks were interactively introduced in the receptor images, which was not required in the cytoarchitectonic and fiber tract images, because the latter images do not contain empty regions. In particular the strategy developed to solve problem of empty regions in receptor autoradiograhps represents a crucial step forwards in the reconstruction of future datasets coding for the regional and laminar distribution patterns of receptors, this is a recurrent problem, but the brain structures that do not express a certain type of receptors, or do so only

at extremely low densities, vary considerably between receptor types.

4.1.2. Data Integration
Although many studies described the registration of 3D data, e.g., MRI volumes, to 3D atlases of rodent (Sergejeva et al., 2015) or human brains (Collins and Evans, 1997), only a few studies registered postmortem rodent brain sections to an MRI volume based atlas (Lebenberg et al., 2010, 2011; Abdelmoula et al.,

2014; Sergejeva et al., 2015). Lebenberg et al. (2010) published the alignment of autoradiographic and histological data of one hemisphere of the mouse brain into a 3D digital MRI based atlas by means of a three step strategy containing rigid, affine and non-linear (elastic) transformations using the reconstructed blockface volume as intermediate modality between atlas and postmortem data. Abdelmoula et al. (2014) used a similar method to transfer mass spectrometry data into the Allen Mouse Brain atlas (Goldowitz, 2010) via affine and non-rigid B-Spline based registration. Sergejeva et al. (2015) identified anatomical landmarks in MRI, blockface or histological images for a landmark based affine registration of these data to the WHS rodent atlases (Johnson et al., 2010).

Since the 3D reconstructed receptor and PLI brains as well as the WHS Atlas brain slightly differed in size and shape of the entire brain and its inner structures, and the use of a non-linear registration is indispensable, we improved the three step registration strategy published in Lebenberg et al. (2010) by the constant use of a multi-resolution registration, i.e., application of a pyramidal method to all three transformation steps and of a similarity criterion (Mutual Information) as an optimal metric for multimodal registration tasks, and testing the influence of the grid spacing used by the non-linear transformation on the registration results. The combination of linear and non-linear transformations of the brains, and the use of the blockface volume as intermediate modality between fiber, cyto- and receptor architectonic and MRI data provided a maximal concordance of the brains. Rigid and affine transformations optimized the matching of the position of the different brains and compensated global shearing and scaling misalignments. Local structural adaptations were done with the non-linear B-Spline based registration. A crucial parameter was the flexibility of the B-Spline grid. With higher flexibility the algorithm generally works more accurately, but unrealistic deformations can be induced. This is illustrated by increasing the grid flexibility, which led to best overlap and distance results with the atlas brain (**Figure 6**). The overlap of the receptor or PLI brains and the MRI brain of the atlas was nearly perfect (Dice coefficients of 0.98 for M2 receptor brain; 0.97 for the PLI brain), but structures within the forebrain, the olfactory bulb and the cerebellum were unrealistically deformed. To overcome this problem, a lower grid flexibility was chosen, although this led to a lower overlap of the entire brains, with special focus in the region of the olfactory bulb. Since the position of the bulb in the receptor and the PLI brain does not reflect its natural position, but is extremely deformed by the necessary preparation steps for receptor autoradiography and PLI measurement, the sections through the bulb were excluded from the registration. This led to a much better overlap and improved distance measurements of the hippocampal complex, particularly in the receptor brain, and still to a good matching of the entire brains well above (Dice coefficient of 0.84 for both brains). Although the quantitative evaluation was based on automatically extracted contour in the case of the whole brain and on manually defined contours in the case of the hippocampus, a comparison of the results obtained for both structures reveals a high consistency. Likewise, although the quantitative evaluation of the matching of the hippocampal

complex was based on independent delineations in the receptor, PLI and MRI brains by different experienced neuroanatomists, the results demonstrated high congruence of the different delineations of the hippocampal complex. This registration strategy was very effective in most brain regions. However, the remaining anti-clockwise rotation of the hemispheres in the fragile sections of the unfixed receptor brain and the mismatch between the neocortex and hippocampus in the center of the section (**Figure 4**) would only be compensated with an extremely high flexibility of the B-Spline grid. This would introduce large undesired artificial deformations in adjoining brain regions, which are biologically unrealistic.

4.2. Limitations and Applications

We are aware that there are a series of putative limitations in the present study. One of them is that different sectioning thickness had to be used for processing of the receptor and PLI brains due to technical constraints. The PLI method requires fixation of brain tissue as well as a minimal section thickness in order to enable extraction of information concerning the direction of the fibers, and previous studies from our group have shown $60\,\mu$m to be an optimal thickness (Axer et al., 2011a,b). Quantitative in vitro receptor autoradiography requires usage of unfixed deep-frozen brains, since the method is based on the fact that the receptors to be visualized must maintain their ability to bind the radioactively labeled ligand present in the incubation buffer (Zilles et al., 2002). Unfortunately, it is technically not possible to obtain 60μm thick sections tissue preprocessed in this manner, and, therefore, we used $20\,\mu$m thick sections. However, the different section thicknesses were accounted for during 3D reconstruction, so we do not think this poses a problem for the registration of our different image modalities to the WHS atlas. Quite the contrary, the methods developed here to overcome these differences in section thickness will facilitate future inclusion of multiscale data into the WHS rat atlas or the WHS mouse atlas.

Our cyto-, M_2 receptor, and fiber architectonic datasets were obtained from adult Wistar rat brains, whereas the WHS atlas is based on an MRI scan of an adult Sprague Dawley rat (Papp et al., 2014, 2015; Kjonigsen et al., 2015), and the fact that brains from different rat strains have been used may also be viewed as a putative problem. However, this issue has been addressed in the past and is not thought to constitute a problem, since comparison of the cyto- and chemoarchitecture of the hippocampal formation in different rat strains has shown it to be a highly conserved brain structure (Kjonigsen et al., 2015).

A large variety of 3D digital atlases based on MRI or reconstructed histological or histochemical sections are available for rodent brains (Goldowitz, 2010; Dorr et al., 2008; Li et al., 2009, 2010; Johnson et al., 2010), nonhuman primate brains (Paxinos et al., 2000; Calabrese et al., 2015) and human brains (Hawrylycz et al., 2012; Shen et al., 2012; Amunts et al., 2013; Amunts and Zilles, 2015). Compared to the current available atlases the WHS (Hawrylycz et al., 2011; Papp et al., 2014) is an unique framework operating as a hub of an infrastructure connecting rodent brain data and reference
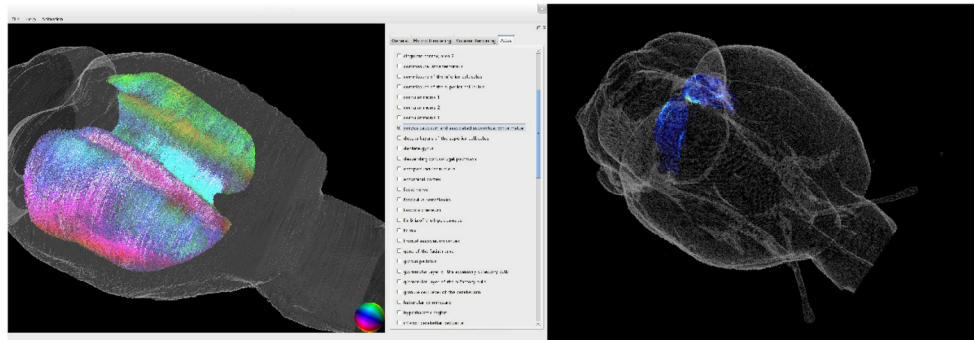
FIGURE 7 | 3D visualization of the fiber architecture in the corpus callosum (left) and of the receptor architecture of the superior colliculus (right). Both structures were extracted using the delineations of the Waxholm Space atlas.

spaces. To enrich this framework with cyto-, M_2 receptor and fiber architecture provides a valuable extension to master analyses of the enormous structural complexity of the brain data.

4.3. Conclusion

We developed a tool to register multiscale and multimodal rat brain data to the WHS atlas brain. It enables retrieval of detailed information of volume densities of cell bodies, of neurotransmitter receptor densities, and of fiber tract architecture and orientation in microscopically identified brain regions (c.f. **Figure 7**).

Therefore, our results considerably expand the data base of the WHS. Furthermore, the methods developed in the present study enable future integration of data of other modalities, which can further enhance the neuroscientific impact of the atlas. The 3D reconstructions of the cyto-, receptor and fiber architectonic images registered to WHS will be publicly accessible through the Human Brain Project (HBP) portal.

REFERENCES

Abdelmoula, W., Carreira, R., Shyti, R., Balluff, B., van Zeijl, R., Tolner, E., et al. (2014). Automatic registration of mass spectrometry imaging data sets to the allen brain atlas. *Anal. Chem.* 86, 3947–3954. doi: 10.1021/ac500148a

Amunts, K., Lepage, C., Borgeat, L., Mohlberg, H., Dickscheid, T., Rousseau, M., et al. (2013). BigBrain: an ultrahigh-resolution 3D human brain model. *Science* 340, 1472–1475. doi: 10.1126/science.1235381

Amunts, K., and Zilles, K. (2015). Architectonic mapping of the human brain beyond Brodmann. *Neuron.* 88, 1086–1107. doi: 10.1016/j.neuron.2015.12.001

Axer, M., Amunts, K., Gräßel, D., Palm, C., Dammers, J., Axer, H., et al. (2011a). A novel approach to the human connectome: ultra-high resolution mapping of fiber tracts in the brain. *NeuroImage* 54, 1091–1101. doi: 10.1016/j.neuroimage.2010.08.075

Axer, M., Gräßel, D., Kleiner, M., Dammers, J., Dickscheid, T., Reckfort, J., et al. (2011b). High-resolution fiber tract reconstruction in the human brain by means of polarized light imaging (3D-PLI). *Front. Neuroinform.* 5:34. doi: 10.3389/fninf.2011.00034

Bjaalie, J. (2002). Localization in the brain: new solutions emerging. *Nat. Rev. Neurosci.* 3, 322–325. doi: 10.1038/nrn790

Calabrese, E., Badea, A., Coe, C., Lubach, G., Shi, Y., Styner, M., et al. (2015). A diffusion tensor MRI atlas of the postmortem rhesus

macaque brain. *NeuroImage* 117, 408–416. doi: 10.1016/j.neuroimage.2015.05.072

Collins, D. L., and Evans, A. C. (1997). Animal: validation and applications of nonlinear registration-based segmentation. *Int. J. Patt. Recogn. Artif. Intell.* 11, 1271–1294. doi: 10.1142/S0218001497000597

Dauguet, J., Delzescaux, T., Condé, F., Mangin, J.-F., Ayache, N., Hantraye, P., et al. (2007). Three-dimensional reconstruction of stained histological slices and 3D non-linear registration with *in-vivo* MRI for whole baboon brain. *J. Neurosci. Methods* 164, 191–204. doi: 10.1016/j.jneumeth.2007.04.017

Dice, L. R. (1945). Measures of the amount of ecologic association between species. *Ecology* 26, 297–302. doi: 10.2307/1932409

Dohmen, M., Menzel, M., Wiese, H., Reckfort, J., Hanke, F., Pietrzyk, U., et al. (2015). Understanding fiber mixture by simulation in 3D polarized light imaging. *NeuroImage* 111, 464–475. doi: 10.1016/j.neuroimage.2015.02.020

Dorr, A., Lerch, J., Spring, S., Kabani, N., and Henkelman, R. (2008). High resolution three-dimensional brain Atlas using an average magnetic resonance image of 40 adult C57Bl/6J mice. *NeuroImage* 42, 60–69. doi: 10.1016/j.neuroimage.2008.03.037

Goldowitz, D. (2010). "Allen Reference Atlas. A Digital Color Brain Atlas of the C57BL/6J Male Mouse," in *Genes, Brain and Behavior*, Vol 9, ed H. W. Dong (John Wiley & Sons Ltd.), 128. doi: 10.1111/j.1601-183X.2009.00552.x

Handels, H. (2000). *Medizinische Bildverarbeitung*. Stuttgart: Vieweg+Teubner Verlag.

Hawrylycz, M., Baldock, R. A., Burger, A., Hashikawa, T., Johnson, G. A., Martone, M., et al. (2011). Digital atlasing and standardization in the mouse brain. *PLoS Comput. Biol.* 7:e1001065. doi: 10.1371/journal.pcbi.1001065

Hawrylycz, M., Lein, E., Guillozet-Bongaarts, A., Shen, E., Ng, L., Miller, J., et al. (2012). An anatomically comprehensive atlas of the adult human brain transcriptome. *Nature* 489, 391–399. doi: 10.1038/nature11405

Heimann, T., Thorn, M., Kunert, T., and Meinzer, H. (2004). "Empirische Vergleichsmaße für die Evaluation von Segmentierungsergebnissen," in *Proceedings of Bildverarbeitung für die Medizin 2004* (Berlin: Springer Berlin Heidelberg), 165–170.

Huynh, A.-M., Kirlangic, M. E., Schubert, N., Schober, M., Amunts, K., Zilles, K., et al. (2015). "Reconstructing a series of auto-radiographic images in rat brains," in *Proceedings of Bildverarbeitung für die Medizin*, (Lübeck), 167–172.

Johnson, G. A., Badea, A., Brandeburg, J., Cofer, G., Fubara, B., Liu, S., et al. (2010). Waxholm Space: an image-based reference for coordinating mouse brain research. *NeuroImage* 53, 365–372. doi: 10.1016/j.neuroimage.2010.06.067

Kjonigsen, L., Lillehaug, S., Bjaalie, J. G., Witter, M., and Leergaard, T. B. (2015). Waxholm Space atlas of the rat brain hippocampal region: three-dimensional delineations based on magnetic resonance and diffusion tensor imaging. *NeuroImage* 108, 441–449. doi: 10.1016/j.neuroimage.2014.12.080

Klein, S., Staring, M., Murphy, K., Viergever, M., and Pluim, J. (2010). elastix: a toolbox for intensity based medical image registration. *IEEE Trans. Med. Imaging* 29, 196–205. doi: 10.1109/TMI.2009.2035616

Lebenberg, J., Hérard, A.-S., Dubois, A., Dauguet, J., Frouin, V., Dhenain, M., et al. (2010). Validation of MRI-based 3D digital atlas registration with histological and autoradiographic volumes: an anatomofunctional transgenic mouse brain imaging study. *NeuroImage* 51, 1037–1046. doi: 10.1016/j.neuroimage.2010.03.014

Lebenberg, J., Hérard, A.-S., Dubois, A., Dhenain, M., Hantraye, P., and Delzescaux, T. (2011). A combination of atlas-based and voxel-wise approaches to analyze metabolic changes in autoradiographic data from Alzheimer's mice. *NeuroImage* 57, 1447–1457. doi: 10.1016/j.neuroimage.2011.04.059

Li, A., Gong, H., Zhang, B., Wang, Q., Yan, C., Wu, J., et al. (2010). Micro-optical sectioning tomography to obtain a high-resolution atlas of the mouse brain. *Science* 330, 1404–1408. doi: 10.1126/science.1191776

Li, X., Yankeelov, T. E., Rosen, G. D., Goreb, J. C., and Dawan, B. M. (2009). Enhancement of histological volumes through averaging and their use for the analysis of magnetic resonance images. *Magn. Reson. Imaging* 27, 401–416. doi: 10.1016/j.mri.2008.07.016

Menzel, M., Michielsen, K., DeRaedt, H., Reckfort, J., Amunts, K., and Axer, M. (2015). A Jones matrix formalism for simulating three-dimensional polarized light imaging of brain tissue. *J. R. Soc. Interface* 12:20150734. doi: 10.1098/rsif.2015.0734

Merker, B. (1983). Silver staining of cell bodies by means of physical development. *J. Neurosci. Methods* 9, 235–241. doi: 10.1016/0165-0270(83)90086-9

Moene, I., Darine, D., Ramachandran, M., Leergaard, T. B., and Bjaalie, J. G. (2011). Rodent brain navigator: database and atlas system for microscopy and imaging data. *Front. Neuroinform. Conference Abstract: 4th INCF Congress of Neuroinformatics (Boston, MA: USA)*. doi: 10.3389/conf.fninf.2011.08.00159

Ourselin, S., Roche, A., Subsol, G., Pennec, X., and Ayache, N. (2000). Reconstructing a 3D structure from serial histological sections. *Image Vision Comput.* 19, 25–31. doi: 10.1016/S0262-8856(00)00052-4

Palomero-Gallagher, N., Zilles, K., Schleicher, A., and Vogt, B. (2013). Cyto- and receptor architecture of area 32 in human and macaque brain. *J. Comp. Neurol.* 512, 3272–3286. doi: 10.1002/cne.23346

Papp, E., Leergaard, T. B., Calabrese, E., Johnson, G. A., and Bjaalie, J. G. (2015). Addendum to "Waxholm space atlas of the Sprague Dawley rat brain" [NeuroImage 97 (2014) 374-386]. *NeuroImage* 105, 561–562. doi: 10.1016/j.neuroimage.2014.10.017

Papp, E. A., Leergaard, T. B., Calabrese, E., Johnson, G. A., and Bjaalie, J. G. (2014). Waxholm space atlas of the Spraque Dawley rat brain. *NeuroImage* 97, 374–386. doi: 10.1016/j.neuroimage.2014.04.001

Papp, E. A., Leergaard, T. B., Csucs, G., and Bjaalie J. G. (2016). Brain-wide mapping of axonal connections: workflow for automated detection and spatial analysis of labeling in microscopic sections. *Front. Neuroinform.* 10. doi: 10.3389/fninf.2016.00011

Paxinos, G., Huang, X.-F., and Toga, A. W. (2000). *The Rhesus Monkey Brain in Stereotaxic Coordinates*. San Diego, CA: Academic Press.

Reckfort, J., Wiese, H., Pietrzyk, U., Zilles, K., Amunts, K., and Axer, M. (2015). A multiscale approach for the reconstruction of the fiber architecture of the human brain based on 3D-PLI. *Front. Neuroanat.* 9:118. doi: 10.3389/fnana.2015.00118

Rueckert, D., Sonoda, L., Hayes, C., Hill, D., Leach, M., and Hawkes, D. (1999). Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Trans. Med. Imaging* 18, 712–721. doi: 10.1109/42.796284

Schober, M., Schlömer, P., Cremer, M., Mohlberg, H., Huynh, A.-M., Schubert, N., et al. (2015). "Reference volume generation for subsequent 3D reconstruction of histological sections," in *Proceedings of Bildverarbeitung für die Medizin*, (Lübeck), 143–148.

Schubert, N., Kirlangic, M. E., Schober, M., Huynh, A.-M., Amunts, K., Zilles, K.,et al. (2015). "3D reconstruction of histological rat brain images," in *Proceedings of Bildverarbeitung für die Medizin*, (Lübeck), 149–154. doi: 10.1007/978-3-662-46224-9_27

Sergejeva, M., Papp, E. A., Bakker, R., Gaudnek, M. A., Okamura-Oho, Y., Boline, J., et al. (2015). Anatomical landmarks for registration of experimental image data to volumetric rodent brain atlasing templates. *J. Neurosci. Methods* 240, 161–169. doi: 10.1016/j.jneumeth.2014.11.005

Shen, E., Overly, C., and Jones, A. (2012). The allen human brain atlas: comprehensive gene expression mapping of the human brain. *Trends Neurosci.* 35, 711–714. doi: 10.1016/j.tins.2012.09.005

Toga, A., Thompson, P., Mori, S., Amunts, K., and Zilles, K. (2006). Towards multimodal atlases of the human brain. *Nat. Rev. Neurosci.* 7, 952–966. doi: 10.1038/nrn2012

Viola, P., and Wells, W. (1997). Alignment by maximization of mutual information. *Int. J. Comput. Vis.* 24, 137–154. doi: 10.1023/A:1007958904918

Yang, Z., Richards, K., Kurniawan, N. D., Petrou, S., and Reutens, D. C. (2012). MRI-guided volume reconstruction of mouse brain from histological sections. *J. Neurosci. Methods* 211, 210–217. doi: 10.1016/j.jneumeth.2012.08.021

Zijdenbos, A., Dawant, B., Margolin, R., and Palmer, A. (1994). Morphometric analysis of white matter lesions in MR images: method and validation. *IEEE Trans. Med. Imaging* 13, 716–724. doi: 10.1109/42.363096

Zilles, K., and Amunts, K. (2010). Centenary of Brodmanns map conception and fate. *Nat. Rev. Neurosci.* 11, 139–145. doi: 10.1038/nrn2776

Zilles, K., and Amunts, K. (2012). Segregation and wiring in the brain. *Science* 335, 1582–1584. doi: 10.1126/science.1221366

Zilles, K., Schleicher, A., Palomero-Gallagher, N., and Amunts, K. (2002). "Quantitative analysis of cyto- and receptor architecture of the human brain," in *Brain Mapping: The Methods*, eds J.C. Mazziotta and A. W. Toga (San Diego, CA: Elsevier), 573–602.

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.

Opposing Effects of Neuronal Activity on Structural Plasticity

Michael Fauth[1,2]* and Christian Tetzlaff[2,3]

[1] Department of Computational Neuroscience, Third Institute of Physics - Biophysics, Georg-August University, Göttingen, Germany, [2] Bernstein Center for Computational Neuroscience, Göttingen, Germany, [3] Max Planck Institute for Dynamics and Self-Organization, Göttingen, Germany

The connectivity of the brain is continuously adjusted to new environmental influences by several activity-dependent adaptive processes. The most investigated adaptive mechanism is activity-dependent functional or synaptic plasticity regulating the transmission efficacy of existing synapses. Another important but less prominently discussed adaptive process is structural plasticity, which changes the connectivity by the formation and deletion of synapses. In this review, we show, based on experimental evidence, that structural plasticity can be classified similar to synaptic plasticity into two categories: (i) Hebbian structural plasticity, which leads to an increase (decrease) of the number of synapses during phases of high (low) neuronal activity and (ii) homeostatic structural plasticity, which balances these changes by removing and adding synapses. Furthermore, based on experimental and theoretical insights, we argue that each type of structural plasticity fulfills a different function. While Hebbian structural changes enhance memory lifetime, storage capacity, and memory robustness, homeostatic structural plasticity self-organizes the connectivity of the neural network to assure stability. However, the link between functional synaptic and structural plasticity as well as the detailed interactions between Hebbian and homeostatic structural plasticity are more complex. This implies even richer dynamics requiring further experimental and theoretical investigations.

Keywords: structural plasticity, architectural plasticity, timescales, synaptic plasticity, network topology

INTRODUCTION

Information from the environment leads to the activation of neural subnetworks in the brain. The connectivity of these neural subnetworks, i.e., the existence and strength of synapses between neurons, influences the neuronal activation and, thereby, determines the way environmental information is processed. Accordingly, the long-term storage of information is related to activity-dependent (long-lasting) changes in connectivity (Hebb, 1949; Morris et al., 1986; Rioult-Pedotti et al., 1998; Leuner et al., 2003; Pastalkova et al., 2006; Whitlock et al., 2006; reviewed, e.g., in Martin et al., 2000; Chklovskii et al., 2004; Dudai, 2004; Hübener and Bonhoeffer, 2010). Basically two types of activity-dependent mechanisms yield such changes: synaptic or functional plasticity and structural plasticity. Structural or architectural plasticity determines the formation and removal of synapses. On the other hand, synaptic or functional plasticity changes the electrochemical transmission efficacy of synapses by altering, for instance, the receptor configuration of the postsynaptic site. Note, as we will show, this functional synaptic plasticity is associated with structural changes at existing synapses (size, postsynaptic density, etc.) and these changes are sometimes summarized as structural plasticity (Lamprecht and LeDoux, 2004). However, here we

restrict structural plasticity to changes of the number of synapses (and of axonal/dendritic trees) and refer the long-term functional changes at existing synapses as synaptic plasticity.

The alterations of the transmission efficacy by synaptic plasticity depend on the level of neuronal activation. However, the mapping between activity level and triggered synaptic changes is not unique. In general, they are categorized into two classes: Hebbian and homeostatic synaptic plasticity. Hebbian synaptic plasticity yields an increase in synaptic efficacy given high neuronal activities (long-term potentiation; LTP; Bliss and Lomo, 1973; Lynch et al., 1983; Bliss and Collingridge, 1993; see Feldman, 2009 for a review), while low levels of activity induce a decrease (long-term depression; LTD; Lynch et al., 1977; Dudek and Bear, 1992; Mulkey and Malenka, 1992; see Collingridge et al., 2010 for a review). Thus, Hebbian synaptic plasticity basically maps the neuronal activation onto the synaptic efficacies or rather connectivity (high activity → stronger connections; low activity → weaker connections; Hebb, 1949; Bliss and Lomo, 1973; Dudek and Bear, 1992; Kirkwood et al., 1996). These changes in the connectivity, in turn, influence the neuronal activities. Along these lines, theoretical studies show (Rochester et al., 1956; Riedel and Schild, 1992; Gerstner and Kistler, 2002; Kolodziejski et al., 2010) that Hebbian synaptic plasticity alone induces a positive feedback loop leading to unrestricted synaptic (and thus neuronal) dynamics. On the other hand, homeostatic synaptic plasticity, as synaptic scaling (Turrigiano et al., 1998), act conversely to Hebbian synaptic plasticity. If neuronal activities are high, synaptic efficacies are decreased, while, if activities are low, efficacies are increased (high activity → weaker connections; low activity → stronger connections; Turrigiano et al., 1998; Hou et al., 2008, 2011; Ibata et al., 2008). Thereby, homeostatic synaptic plasticity alone induces a negative feedback loop and, thus, stabilizes the dynamics. As several theoretical results indicate (Tetzlaff et al., 2011; Zenke et al., 2013; Toyoizumi et al., 2014), the combination of both plasticity processes lead to desired, stable dynamics.

We will argue in this review that, analogous to functional synaptic plasticity, structural plasticity can also be categorized into two different classes of activity-dependency: (i) One class of structural changes maps features of the neuronal activity onto the connectivity, such that the connectivity is strengthened with high activity levels and vice versa. These changes will be referred to as Hebbian structural plasticity (Hebb, 1949; Helias et al., 2008). (ii) The other class of structural changes weakens (strengthens) the connectivity given high (low) neuronal activities and, thus, stabilizes the dynamics. This class is named homeostatic structural plasticity (Butz et al., 2009).

Note, this classification is phenomenological. Changes in connectivity (synaptic as well as structural) are not directly linked to neuronal activity. Neuronal activity initiates such changes by triggering secondary processes as molecular signaling cascades, which lead to the corresponding changes. For the here discussed plasticity processes, these underlying signaling cascades can have different degrees of similarity, which we will not consider in detail. The focus of this review is to systematize the qualitative links between the neuronal activity level and resulting connectivity changes.

Moreover, we focus on morphological changes of connections between excitatory neurons only. The dynamics of inhibitory synapses has been reviewed, for instance, by Vogels et al. (2013) for inhibitory synaptic plasticity and by Flores and Méndez (2014) for inhibitory structural plasticity. Further non-synaptic homeostatic mechanisms stabilizing neural network dynamics have been reviewed in Turrigiano and Nelson (2004), Marder and Goaillard (2006), or Yin and Yuan (2014).

In the following, as structural and synaptic plasticity are linked to each other, we first briefly outline the main findings for synaptic plasticity. Then, we review the morphological changes of synapses induced by synaptic plasticity and relate these changes to the dynamics of synapses and, thus, to structural plasticity. Following this, we summarize the experimental evidence of activity-dependent structural changes and categorize these, similar to synaptic plasticity, into the two classes of Hebbian and homeostatic structural plasticity. We also briefly review indications of Hebbian and homeostatic processes occurring during development. Finally, we sort theoretical investigations studying the dynamics of structural plasticity by this categorization and, based on their results, arrive at conclusions about the different functional roles of Hebbian and homeostatic structural plasticity.

ACTIVITY-DEPENDENT SYNAPTIC PLASTICITY

The most investigated long-term plasticity in neuronal systems is synaptic plasticity. This mechanism adapts synaptic efficacies (by, e.g., altering the number of AMPA receptors at the postsynaptic site) between neurons dependent on the neuronal activation. One distinguishes between two different forms of synaptic plasticity: (i) Hebbian synaptic plasticity and (ii) homeostatic synaptic plasticity.

(i) Hebbian synaptic plasticity adapts the synaptic efficacies seconds or minutes after onset of a stimulus-induced neuronal activation. In general, neuronal activity induces a calcium influx into the postsynaptic site inducing a complex molecular cascade which changes, amongst others, the number of AMPA receptors determining the synaptic efficacy (Kauer et al., 1988; Muller and Lynch, 1988; Shi et al., 1999; reviewed, e.g., in Malenka and Bear, 2004). Many experiments show that a low calcium level (thus a low neuronal activity level) leads to a decrease of the number of AMPA receptors (long-term depression: LTD; Lynch et al., 1977; Dudek and Bear, 1992; Mulkey and Malenka, 1992; Beattie et al., 2000; see Collingridge et al., 2010 for a review) while a high calcium level yields an insertion of new ones resulting in a stronger synaptic efficacy (long-term potentiation: LTP; Bliss and Lomo, 1973; Lynch et al., 1983; Malenka et al., 1992; Bliss and Collingridge, 1993; see Feldman, 2009 for a review). Thus, after several minutes, Hebbian synaptic plasticity maps the strength of the stimulus onto the strength of the synaptic transmission. Note, synapses from several input sources connecting to the same postsynaptic neuron can interact with each other

yielding cooperative and competitive dynamics (Miller, 1996). Moreover, the change of the synaptic efficacy can also depend on the relative timing of pre- and postsynaptic action potentials (spike-timing-dependent plasticity: STDP; Levy and Steward, 1983; Gerstner et al., 1996; Markram et al., 1997b; Bi and Poo, 1998; see Markram et al., 2011 for a review), such that also temporal correlations might be mapped onto the synaptic efficacies. However, as several theoretical studies indicate, Hebbian synaptic plasticity alone induces a positive feedback loop leading to unrestricted growth of the synaptic efficacy (Rochester et al., 1956; Riedel and Schild, 1992; Gerstner and Kistler, 2002; Kolodziejski et al., 2010). In other words, if a stimulus drives the firing of the postsynaptic neuron, LTP potentiates the corresponding synaptic efficacy and, by this, induces a stronger input drive which, in turn, generates more potentiation and so forth. Thus, Hebbian synaptic plasticity alone would yield unstable, divergent dynamics of the synaptic efficacies.

(ii) Another process adapting the transmission strength of a synapse is homeostatic synaptic plasticity. Several different homeostatic processes dampen the dynamics of neuronal systems on various levels (Zhang and Linden, 2003; Turrigiano and Nelson, 2004; Marder and Goaillard, 2006; Turrigiano, 2011; Yin and Yuan, 2014). Thus, it is reasonable that homeostatic processes, like synaptic scaling, also adapt synaptic efficacies (Turrigiano et al., 1998; Hengen et al., 2013; Keck et al., 2013). Amongst others, this mechanism depends mainly on the average postsynaptic activity (Ibata et al., 2008). Here, in contrast to Hebbian synaptic plasticity, if the neuronal activity is high, the synaptic efficacies are decreased and, if the activities are low, the efficacies are increased (Turrigiano et al., 1998; Burrone et al., 2002; Kim et al., 2012). Hereby, synaptic scaling is unspecific, i.e., it scales all synapses onto a postsynaptic neuron preserving relative differences between synaptic efficacies induced by Hebbian plasticity (Turrigiano, 2008). However, several experiments indicate (e.g., Turrigiano et al., 1998; Hengen et al., 2013; Keck et al., 2013, but see also Ibata et al., 2008) that, compared to Hebbian synaptic plasticity, this process is much slower (hours to days) which complicates the analysis of both processes within the same experimental setup (Vitureira and Goda, 2013). Nevertheless, theoretical investigations show that synaptic scaling is one way to solve the problem of unrestricted growth discussed above (Tetzlaff et al., 2011; Zenke et al., 2013; Toyoizumi et al., 2014). Please note that there are also other solutions proposed to solve this problem (von der Malsburg, 1973; Sejnowski, 1977a,b; Bienenstock et al., 1982; Oja, 1982).

In summary, investigations in the field of synaptic plasticity show that, at least, two classes of processes adapt synaptic efficacies: Hebbian synaptic plasticity and homeostatic synaptic plasticity. Hereby, Hebbian synaptic plasticity maps neuronal activities onto the synaptic efficacies (high act. → stronger connect.; low act. → weaker connect.), which are, in turn, stabilized by homeostatic processes (high act. → weaker connect.; low act. → stronger connect.).

ACTIVITY-DEPENDENT STRUCTURAL PLASTICITY

Activity-dependent structural plasticity basically influences two different physical substrates: On the one hand, *neurites* (i.e., dendrites and axons) grow and retract dependent on the level of neuronal activation (Cohan and Kater, 1986; van Huizen and Romijn, 1987). These growth processes determine the basic shape of a neuron and its regions of afferent and efferent connections. On the other hand, *synapses* (i.e., dendritic spines and axonal boutons) are continuously formed and deleted. Although an axon and a dendrite lie close together and the gap could be bridged by a synapse, the existence of a synapse is not guaranteed (Kalisman et al., 2005). In fact, the formation and deletion of a synapse also depend on the neuronal activation of both neurons (see e.g., Annis et al., 1994; Nägerl et al., 2007; Kwon and Sabatini, 2011; Hill and Zito, 2013).

As the majority of cortical synapses resides on dendritic spines (Yuste, 2010), many studies applied time-lapse imaging of the dynamics of dendritic spines for analyzing the structural dynamics or structural plasticity of single synapses. This implies the problem that the existence of a dendritic spine does not guarantee the existence of a functional synapse. However, several experiments provide evidence that, at least after a few hours after spine formation, new born spines are structurally and functionally equivalent to mature spines hosting a synapse (Trachtenberg et al., 2002; Knott et al., 2006; Nägerl et al., 2007; Zito et al., 2009). Similarly, also the emergence and stabilization of axonal terminals or boutons seems to involve synapse formation and maturation (Friedman et al., 2000; Ruthazer et al., 2006). Thus, the existence of a spine or bouton is a good indicator for the existence of a functional synapse.

Link between Structural and Synaptic Plasticity

The dynamics of synapses is determined by the dynamics of dendritic spines. Accordingly, structural plasticity depends on the morphology of spines as their sizes and shapes (Nägerl et al., 2008; Tønnesen et al., 2011, 2014). Experiments indicate that the volume of a dendritic spine correlates with the synaptic efficacy of the corresponding synapse (Matsuzaki et al., 2001; Knott et al., 2006; Zito et al., 2009) which, in turn, is influenced by synaptic plasticity. Accordingly, stimuli causing long-term potentiation (LTP) also cause spine enlargements (Fifková and Van Harreveld, 1977; Okamoto et al., 2004; Yang et al., 2008, for a review see Yuste and Bonhoeffer, 2001) while stimuli causing long-term depression (LTD) induce spine shrinkage (Okamoto et al., 2004; Zhou et al., 2004; Oh et al., 2013). Hereby, synaptic and structural changes rely on distinct signaling cascades, which are triggered by the same signals (Matsuzaki et al., 2004; Zhou et al., 2004). Thus, blocking synaptic plasticity, for instance, by blocking NMDA-receptors also prevents changes in the spine volume. Several experiments indicate that the spine head volume is correlated to the lifetime or stability of the spine (Grutzendler et al., 2002; Majewska et al., 2006; Yasumatsu et al., 2008; Loewenstein et al., 2015). Thus, the spine stability or removal of a synapse is, in turn,

correlated to the synaptic efficacy of the corresponding synapse, which also has been directly observed in several experiments (Holtmaat et al., 2005; Le Bé and Markram, 2006, reviewed, e.g., in Kasai et al., 2003). In combination with STDP, this relation between synaptic weight, spine volume and spine stability could give rise to a spike-timing-dependent structural plasticity (Helias et al., 2008; Deger et al., 2012), which still has to be experimentally verified. Interestingly, the stability of a synapse is also influenced by the reliability of signal transmission of the synapse (Wiegert and Oertner, 2013) which is also altered by synaptic plasticity (Stevens and Wang, 1994). Thus, for Hebbian-like changes, structural and synaptic plasticity are linked with each other by the morphology of spines or properties of the synapse (Segal, 2005).

Some evidence indicates a similar link for homeostatic changes: *in vitro* (Murthy et al., 2001) and *in vivo* (Keck et al., 2013) studies show that changes of the spine volume also go along with the activity-dependent homeostatic scaling of synaptic efficacies. Given the aforementioned correlation between spine volume and spine stability, we expect that structural plasticity is also linked to homeostatic synaptic plasticity.

In the following, we will summarize experimental results indicating the different aspects of activity-dependent structural plasticity in more detail. We will classify these aspects according to Hebbian (high act. → stronger connect.; low act. → weaker connect.) and homeostatic (high act. → weaker connect.; low act. → stronger connect.) structural plasticity in adult networks. In addition, we will show that many of these experiments support the here discussed link between synaptic and structural plasticity. We will also provide a brief survey of structural dynamics during development. Finally, we will discuss experimental evidence of the interaction of Hebbian and homeostatic structural plasticity in the same neural system.

Evidence for Hebbian Structural Plasticity
LTP-Stimuli
The induction of LTP by a strong neuronal activation is mainly associated with the increase of the synaptic efficacy (e.g., number of AMPA receptors) of existing synapses (Malenka and Bear, 2004; Feldman, 2009). However, already in the 1980s first studies (Lee et al., 1980; Chang and Greenough, 1984) indicate that 15–20 min after applying the strong stimulus the number of synapses is enhanced, too. In addition, one also observes an increase in the number of filopodia (Lee et al., 1980; Chang and Greenough, 1984), which seem to be the precursors of dendritic spines (Ziv and Smith, 1996). Accordingly, about 30 min after stimulation, an increased number of dendritic spines can be observed (Moser et al., 1994; Trommald et al., 1996; Engert and Bonhoeffer, 1999, but see also Desmond and Levy, 1990). The strength of the effect and the detailed timescale, however, depend strongly on the used tissue and preparation method (Sorra and Harris, 1998; Dunaevsky et al., 1999; Kirov et al., 1999; Bourne et al., 2007; Bourne and Harris, 2011), but most studies report timescales between 5 and 30 min.

This increase in the number of spines after an LTP-stimulus provides further support for an interaction between Hebbian

synaptic plasticity and Hebbian structural plasticity: A strong neuronal activation will induce an increase in synaptic efficacies or rather in the spine volumes implying the stabilization of these enlarged dendritic spines. Given a continuous formation of new spines, this also implies that the new and small spines, which would be pruned without stimulation, will be enlarged and stabilized by synaptic plasticity. Together with the already existing (and further stabilized) spines, the stabilization of new spines by the strong stimulus would lead to an increase of the the number of spines as observed experimentally. For this, the rate of forming new spines could be independent of the neuronal activity and stay constant. This potential explanation of the increase in spine number is supported by a recent study demonstrating that LTP stabilizes nascent spines (Hill and Zito, 2013). Accordingly, blocking the signals inducing LTP (by blocking the NMDA-channels) also prevents the increase in the number of dendritic spines and also of axonal boutons (Engert and Bonhoeffer, 1999; Maletic-Savatic et al., 1999; Toni et al., 1999; Nikonenko et al., 2003). Thus, the dynamics of dendritic spines can be explained by the link between Hebbian structural plasticity and synaptic plasticity.

Note, although an increase of the number of spines could be explained by assuming a constant rate of forming new spines, the LTP-dependent appearance of more filopodia (Lee et al., 1980; Chang and Greenough, 1984; Maletic-Savatic et al., 1999) suggests that the formation rate changes, too. Thus, further experiments are required to clarify whether the formation rate of dendritic spines (and also of axonal boutons) stays constant or whether it is adapted by the level of neuronal activity.

Also at the presynaptic neuron an LTP-inducing stimulus triggers a structural remodeling: the number of axonal boutons increases. This effect arises already 15 min after the stimulation (Nikonenko et al., 2003; Ninan et al., 2006). The fact, that both the numbers of dendritic spines and of axonal boutons are enhanced, suggests that new synapses are formed by these new elements. In addition, recent findings indicate also an LTP-dependent increase in the probability that a bouton hosts one or more functional synapses (Medvedev et al., 2014). Thus, newly formed spines have a very high chance of connecting to a new or old bouton and, hence, forming a new synapse.

LTD-Stimuli
A link between the dynamics of Hebbian structural plasticity and LTD-inducing stimuli has been established, too. Several experimental studies show that the induction of an LTD-stimulus (low frequency) yields a separation of pre- and postsynaptic terminals (Bastrikova et al., 2008) and a loss of dendritic spines (Nägerl et al., 2004; Wiegert and Oertner, 2013). Thus, similar to the dynamics triggered by an LTP-stimulus, due to the induction of a low frequency stimulation, the synaptic efficacy is decreased, spines shrink and decrease their stability, and the removal rates of dendritic spines are increased (Segal, 2005). This is supported by experiments showing that the prevention of LTD by blocking NMDA-channels impedes the structural effects (Nägerl et al., 2004; Yu et al., 2013). Thus, also these results indicate

that structural plasticity is linked to synaptic plasticity which influences the stability of the corresponding dendritic spines. The timescale of spine shrinkage and removal seems to depend on the experimental conditions: some experiments report spine shrinkage after about 20 min of LTD-induction (Oh et al., 2013) while other studies report no significant changes in spine volume or stability up to 30 min after the induction of LTD (Wiegert and Oertner, 2013). Like for LTP-induced dynamics, also the presynaptic site is also influenced by a low level of activation as it increases the turnover of axonal boutons (Paola et al., 2006; Stettler et al., 2006) resulting in a loss of synapses (Becker et al., 2008).

In summary, stimulation protocols inducing Hebbian synaptic plasticity change the stability and number of synapses. A strong activation induces the formation of more synapses while a low activation induces a loss of synapses. These variations in the number of synapses seem to depend on changes in the stability of the corresponding dendritic spines and axonal boutons correlated to the actual synaptic efficacy adapted by Hebbian synaptic plasticity. However, it is still not clear whether the rate of newly formed spines and boutons is changed, too. Furthermore, the data about the dynamics induced by LTD-stimuli are less comprehensive than the data for LTP-stimuli.

Most of the above discussed structural dynamics happens on a timescale of the order of several minutes to one hour. On this timescale the dendritic trees and axons hosting spines and boutons remain quite stable (Ziv and Smith, 1996; Grutzendler et al., 2002; Trachtenberg et al., 2002; Lee et al., 2006; Paola et al., 2006; Stettler et al., 2006). Hence, fast Hebbian changes of the network structure must be mainly implemented by the growth or removal of dendritic spines. On slower timescales, also changes of the dendrites and axons take place. However, as we will discuss in the following, such changes are mainly triggered by homeostatic processes.

Evidence for Homeostatic Plasticity

As already mentioned above, the connectivity of neural networks is not only adapted by Hebbian-like changes. Similar to synaptic plasticity, also structural changes show homeostatic dynamics, i.e., a decrease of connectivity with high neuronal activities and an increase with low activities. Typically, these homeostatic dynamics are observed under chronically altered conditions of neuronal activity and, thus, also at slower timescales. In general, the resulting structural changes seem to counterbalance the altered conditions and, thereby, regulate the activity back to an intermediate level (for a complete review of homeostatic structural processes see Butz et al., 2009). Like Hebbian structural plasticity, homeostatic structural changes are determined by the dynamics of dendritic spines and axonal boutons. However, under extreme conditions, as in epilepsy or after lesions, also changes of the dendritic and axonal trees are observed.

Already in the year 1978, Wolff et al. (1978) observed *in vivo* the growth of protrusions and thickenings on the dendritic tree after decreasing neuronal activity. For this, they applied the inhibitory neurotransmitter GABA for 3–7 days. Further studies verified that chronic blockage of neuronal activity can yield an increase in the number of spines after

approximately 8 h (Dalva et al., 1994; Rocha and Sur, 1995; McAllister et al., 1996; Kirov and Harris, 1999) indicating a slower timescale for homeostatic structural changes as compared to Hebbian ones. Hereby, already the blockage of NMDA channels leads to an increase in spine number (Yu et al., 2013; Chen et al., 2015) or prevents spine elimination (Bock and Braun, 1999). Note that during development blocking activity or NMDA receptors can show the opposite effect (Annis et al., 1994; Collin et al., 1997). However, the newly formed spines often host silent synapses needing synaptic plasticity to be converted to functional synapses (Nakayama et al., 2005). On the other hand, persistent depolarization of neurons leads to a loss of dendritic spines (Müller et al., 1993; Drakew et al., 1996). Already the application of high levels of NMDA induces a spine loss by the destabilization of the spine actin scaffold (Halpain et al., 1998). Thus, the number of spines is adapted in an activity-dependent homeostatic manner.

Furthermore, the changes in the number of spines also depend on the calcium level (Kirov and Harris, 1999; Kirov et al., 2004; Tian et al., 2010). Accordingly, it has been proposed that dendritic spines follow a calcium-dependent homeostasis (Segal et al., 2000). As the postsynaptic calcium level is largely influenced by neuronal activity (Spruston et al., 1995; Helmchen et al., 1996; reviewed, e.g., in Higley and Sabatini, 2008), the calcium-dependent homeostasis could, in turn, imply an activity-dependent homeostasis as described above. However, the detailed relation between calcium, activity, and spine dynamics is more complex, as the calcium level is also regulated by other signals as neurotrophins (Stoop and Poo, 1996) or cell adhesion molecules (Bixby et al., 1994). Furthermore, in contrast to the postsynaptic activity, calcium is a local signal allowing different dynamics at different branches of the dendritic tree. Accordingly, by comparing different branches of the same dendrite, where each branch receives stimuli from other brain regions, such different spine dynamics are observed (Mattson, 1988; Bravin et al., 1999; Lohmann et al., 2005; Deller et al., 2006; Vuksic et al., 2011, see also Yu and Goda, 2009; Vlachos et al., 2012a, 2013 for evidence on local homeostasis of synaptic efficacies). However, in summary, these experiments indicate that the number of spines or synapses is adapted by activity-dependent homeostatic structural plasticity.

Evidence from Networks in Extreme Situations

Further evidence for homeostatic dynamics are obtained in more complex settings which we summarize in the following. Note that under these conditions dynamics of dendritic and axonal trees are observed, too.

For example, homeostatic regulation of connectivity is found in animal models of epilepsy. Epileptic seizures are network states of high and synchronous activity. Given a homeostatic dynamic, this would lead to a decrease in the number of spines which, indeed, was found in animal models of epilepsy (Scheibel et al., 1974; Paul and Scheibel, 1986; Geinisman et al., 1990; Isokawa and Levesque, 1991; Isokawa, 1998). These changes are likely signs of structural plasticity rather than mere damages by the epileptic seizures, as the number of spines recovers after several

days without seizures (Müller et al., 1993; Isokawa, 1998). The spine loss is only visible at least 5 h after the seizure (Mizrahi et al., 2004), which implies that, also under these conditions, the timescale of homeostatic structural plasticity is typically slower than for Hebbian structural plasticity described above. Interestingly, after several days with reoccurring seizures also changes of the neuronal morphology, like retraction of dendritic branches, are measured (Colling et al., 1996; Jiang et al., 1998).

In contrast to the elevated activities during epilepsy, phenomena like strokes, lesions, or deprivations typically lead to lowered activity levels in a group of neurons. For instance, for deprived neurons a homeostatic dynamic would increase the number of spines. Indeed, experiments show that, after 4 days of monocular deprivation, the number of newly formed spines in the binocular cortex of adult mice doubles compared to control conditions (Hofer et al., 2009). Interestingly, a second phase of monocular deprivation at the same eye does not lead to an increased formation of new spines. Now, the synaptic efficacies of spines formed during the first phase are strengthened by (presumably homeostatic) synaptic plasticity counterbalancing the lost input (Hofer et al., 2009). These findings support the link between (homeostatic) structural and synaptic plasticity.

However, as shown by Keck et al. (2008), also smaller interferences, like small lesions of the retina, lead to more new spines (in the lesion projection zone in the visual cortex). In this experiment, although more new spines are formed, the spine density is comparable to control conditions after 3 days. Another experiment shows that trimming the whiskers of rats leads to an increased number of spines and an outgrowth of dendritic trees in the input-receiving layer in the barrel cortex (Vees et al., 1998; other layers might be affected differently, see Chen et al., 2015). Along this line, one observes massive reorientation of the dendritic trees of adult rats after whisker removal, while the system regains the pre-removal dendritic lengths and spine densities (Tailby et al., 2005). Note, however, already the retraction of dendrites from denervated areas can increase the exitability of neurons, such that activity-homeostasis can be reached without regaining the pre-lesion dendritic length (Platschek et al., 2016).

Interestingly, not only the dendrites of the neurons with lesioned afferents, but also axons of neighboring neurons contribute to regain homeostasis. Although these neighboring neurons are not directly affected by the lesions, they also can be expected to experience altered activity levels. This triggers, after a few days, the growth of axons from the neighboring neurons toward the deprived region (Darian-Smith and Gilbert, 1994; Yamahachi et al., 2009; Marik et al., 2010). Furthermore, damaged axons can grow out and form new synapses, similar to growth dynamics during development (Canty et al., 2013). In summary, we find that lesions trigger the formation of new spines and the outgrowth of dendrites, which, together with new innervation from neighboring neurons, presumably form new synapses and restore the activity level.

Thus, very high or low activity levels occurring in extreme situations like epilepsy, lesions, or stroke are counterbalanced by structural changes on the timescale of several hours to days, thereby, contributing to activity-dependent homeostasis.

Activity-Dependent Structural Plasticity during Development

As already mentioned above, apart from networks in extreme situations, many experiments in adult networks observe very small or no changes of the axonal or dendritic arborization. This is different during the development of neural networks, when these dendritic and axonal trees are formed. Interestingly, also during this phase activity-dependent structural processes contribute to the network dynamics. In the following we will briefly discuss these experiments.

Homeostatic Structural Changes during Development

Single, isolated neurons in culture typically start growing axons and dendrites. This initial process could already be a homeostatic mechanism, as such neurons typically exhibit only weak activities (Kater et al., 1989). The further outgrowth of neurites also seems to be homeostatically regulated: On the one hand, the application of the inhibitory neurotransmitter GABA, which normally decreases activity, triggers an increased outgrowth (Mattson and Kater, 1989). On the other hand, excitatory neurotransmitters as glutamate (but not NMDA, see Mattson et al., 1988), which normally yield an increased activity, induce the degeneration of the dendritic structures (Haydon et al., 1984, 1987; Mattson, 1988; Mattson et al., 1988; Mattson and Kater, 1989). The strength of this effect is dose-dependent (Mattson et al., 1988). Note, during early developmental phases, GABA is an excitatory neurotransmitter (Barker et al., 1998; Ben-Ari, 2002). Still, in the above studies GABA shows the inverse effect of the excitatory neurotransmitters. Furthermore, in these experiments, changes in the axonal dynamics are initiated only at very high doses and lead to a retraction of the axon.

Further experiments targeted downstream signals of these neurotransmitters. Also here, several indications show that especially the postsynaptic calcium level seems to trigger dendritic changes: on a slower timescale, an increased level of calcium induces a retraction of dendrites while a decrease of calcium leads to an outgrowth of dendrites (Mattson and Kater, 1987, for a similar effect for CaMKII see Wu and Cline, 1998). These dynamics are summarized in the calcium-dependent homeostasis hypothesis for dendrites (Kater et al., 1989; Lipton and Kater, 1989). Furthermore, recent experiment suggest that also the dynamics of filopodia are regulated dependent on local calcium currents (Lohmann et al., 2005). As discussed above, the calcium level is mainly influenced by the neuronal activity. Therefore, we suppose that the calcium-dependent homeostasis hypothesis implies an activity-dependent homeostasis, i.e., neurons grow and retract their dendrites to find an optimal level of input which, in turn, assures a medium activity level. This hypothesis is supported by experiments showing that increased activity, due to electrical stimulation, prevents dendritic outgrowth (Cohan and Kater, 1986; Fields et al., 1990 but see Garyantes and Regehr, 1992), whereas blocking activity yields enhanced growth of dendrites (van Huizen and Romijn, 1987; Fields et al., 1990). Note, these experiments demonstrate a relation between activity and the dendritic outgrowth describing only the potential connectivity between neurons and not the realized connectivity between neurons.

Whether this also yields the formation of more functional synapses remains unclear.

Further evidence for homeostatic structural changes during development are coming from experiments analyzing the time course of the developmental process of neuronal networks. During development, neural networks evolve from an initial unconnected state to a connected matured state. Initially, neurons have very low activities (Ramakers et al., 1990; Chiappalone et al., 2006; Wagenaar et al., 2006), which could trigger the outgrowth of neurites and formation of synapses in a homeostatic manner (van Ooyen, 2011). Before reaching the matured state, neural networks typically pass through a phase of extreme build up of synapses followed by a phase of synapse pruning (so-called overshoot)—dependent on the level of neuronal activity (Feldman and Dowd, 1975; Huttenlocher et al., 1982; Huttenlocher, 1984; van Huizen et al., 1985, 1987; van Huizen and Romijn, 1987; van Pelt et al., 1996; Bock and Braun, 1999; Hua and Smith, 2004; Zuo et al., 2005a,b). Such an overshoot in synapse number is typical for neural networks with homeostatically regulated connectivity (van Ooyen, 2003, 2011).

Hebbian Structural Changes during Development

During development some structural changes of axonal and dendritic trees also show Hebbian-like dynamics as described in the following: Neurites grow by constantly adding and removing branches (Wu and Cline, 1998; Sin et al., 2002; Wong and Ghosh, 2002; Portera-Cailliau et al., 2003). Hereby, only a few branches become stable and form the axonal or dendritic tree, while others are removed on the timescale of minutes to hours (Wu and Cline, 1998). Thereby, the activation of receptors and local calcium transients are necessary to stabilize and maintain dendritic branches (Lohmann et al., 2002; Vaillant et al., 2002; Hutchins and Kalil, 2008). Accordingly, in animals experiencing four hours of increased neuronal activity due to visual stimulation, one observes significantly more stabilized dendritic branches as compared to animals left in the dark (Sin et al., 2002). Similarly, the blockage of neuronal activities yields much less complex dendritic trees (Groc et al., 2002).

Interestingly, the stabilization of dendritic and axonal branches also depends on the connectivity, more precisely, on the existence and maturation of synaptic contacts on the branch (Haas et al., 2006; Ruthazer et al., 2006). As shown above, in adult networks, the activity-stability relationship of synapses implements Hebbian changes in connectivity. Thus, if the dynamics underlying the stabilization of synaptic contacts are similar during development and in adult networks, the activity-dependent stabilization of spines and, therefore, of branches would indicate a Hebbian-component of the growth of dendritic trees.

EVIDENCE FOR THE INTERACTION OF HEBBIAN AND HOMEOSTATIC STRUCTURAL PLASTICITY

As we discussed above, for adult networks, the alteration in neuronal activity causes two different directions of structural changes (see **Figure 1**). On a fast timescale (minutes to hours)

the number of dendritic spines goes along with the change in neuronal activity in a Hebbian manner. On a typically slower timescale (hours to days), the dynamics of dendrites and dendritic spines homeostatically counterbalance the change in activity and regulate it back to an intermediate target regime. Obviously, in experiments, chronic changes in neuronal activity should trigger both processes which, then, interact with each other. With these two mechanisms and their typically different timescales at hand, in the following, we will discuss direct conclusions about the dynamics of structural changes during a period of altered activity.

For example, when neurons start to receive reduced or LTD-inducing inputs, the corresponding synapses will be depressed and, therefore, more likely to be removed due to Hebbian structural plasticity—the spine density is reduced (**Figure 1**, bottom center). Later on, due to the reduced activity of the neuron, homeostatic structural plasticity yields the formation of new synapses—the spine density will increase (**Figure 1**, bottom right). Note, as the homeostatic changes are unspecific, very likely these new synapses connect to other, more active inputs. Thus, when the neural network has again reached its homeostatic level and assuming that the synaptic efficacies are, on the long run, similar to those before the activity alteration, the spine density is probably at the same level as before receiving the LTD-inducing inputs. Thus, as a direct consequence from the interaction of Hebbian and homeostatic structural plasticity in the same neural network, we expect in general a transient decrease in the spine density.

Such transient changes have been observed already in the 1970s (Parnavelas et al., 1974; Goldowitz et al., 1979). In these studies, the transsection of afferent hippocampal axons yields a strong decrease in spine density around 4 days after deafferentiation and a restoration of the initial spine density after 10–50 days (Parnavelas et al., 1974; Goldowitz et al., 1979; Vuksic et al., 2011). Strikingly, this transient change in spine density does not result from changes in the spine formation rate, but rather from changing the elimination rate or the stability of the spines (Vlachos et al., 2012b). Similarly, one observes changes in the spine elimination rate in barrel cortex after whisker trimming also leading to a transient decrease of the spine density (Zuo et al., 2005a; Miquelajauregui et al., 2015).

These results are consistent with the correlation between spine stability, spine volume, and synaptic efficacy governing the interaction of synaptic and structural plasticity: First, Hebbian synaptic plasticity would decrease the efficacies and the stability of spines, such that their density decreases. Later, synaptic scaling would scale up the synaptic efficacies of both old and new synapses and, thereby, stabilizes them and increases spine density. Interestingly, at the same time, Hebbian synaptic plasticity can induce competitive effects between newly formed and up-scaled preexisting spines, which destabilizes the newly formed synapses and, thereby, protracts the recovery of the system (Vlachos et al., 2013).

On the other side, paradigms which supposedly trigger higher neural activities, such as motor learning or an enriched environment, have been demonstrated to elicit a transient increase in the number of spines after 2–3 days of stimulation. After 7 days the number of spines reaches control level again

FIGURE 1 | Summary of experimental findings. On the one hand, neural systems, which experience high activities, quickly form new spines and strengthen old ones (Hebbian changes). If high activities persist for longer time, spines are removed and dendrites start to retract. This reduces the input and drives the postsynaptic activities back to lower levels (homeostatic changes). On the other hand, low activities lead to spine removal and shrinkage. On the long run, however, new spines are created and dendrites start growing out such that the neurons acquire more inputs and increase their activity levels. Thus, structural plasticity shows Hebbian-type changes and homeostatic changes.

(Xu et al., 2009; Yang et al., 2009, see **Figure 1**, upper row). Also during these experiments, the number of new filopodia remains constant, which suggest a constant formation rate of new spines. Interestingly, the repeated training selectively stabilizes mainly the newly formed spines, while the stability of preexisting spines drops (Xu et al., 2009). Also for these types of experiments the interaction of structural and synaptic plasticity provides a potential explanation for the observed dynamics. We expect that Hebbian synaptic plasticity leads to a selective potentiation and, thus, a stabilization of the synapses which are important for learning (especially the ones hosted by newly formed spines, which are important for the task performance, see Xu et al., 2009; Yang et al., 2009). This, in turn, leads to an increased spine number and higher neuronal activities. In the long run, this increased activity triggers unspecific homeostatic synaptic plasticity decreasing the stability of synapses and inducing their pruning. Remarkably, in experiments, when training is stopped earlier, the newly formed spines are less stable than the preexisting ones (Xu et al., 2009). Following our reasoning, this could imply that learning was not long enough to trigger sufficient potentiation to stabilize the newly formed synapses.

The interaction of Hebbian and homeostatic mechanisms could also be used to explain a detailed EM-study conducted by Bourne and Harris (2011). This study shows that, 5–30 min after

a typically LTP-inducing tetanic burst stimulation, a transient increase in the number of stubby spines, shaft synapses, and nonsynaptic protrusions can be observed. However, already after 2 h these structures are not present anymore. In addition, the number of small spines is decreased compared to pre-stimulation, whereas the postsynaptic densities of all remaining spines have been enlarged such that the PSD (postsynaptic density) area per micrometer dendrite is the same as for controls (Bourne and Harris, 2011). This suggests a strong and, possibly, fast homeostatic mechanism (the authors argue for a resource homeostasis of the polyribosomes which are used for spine creation and enlargement). Thus, probably a group of synapses is selectively stabilized by Hebbian synaptic plasticity increasing neuronal activity. At the same time, homeostatic synaptic and structural plasticity counterbalance these changes and decrease the stability of all synapses leading to the removal of small, unpotentiated synapses. These dynamics are similar to the dynamics during motor learning described above.

These examples demonstrate that Hebbian and homeostatic as well as synaptic and structural plasticity are strongly interweaved and jointly adapt the connectivity of the neural network according to alterations in neuronal activity. To understand these complex interactions in more detail, further experiments are needed. However, to assess also the general principles, theoretical

network models are required. In the following, we will discuss the state of the art of theoretical models of structural plasticity.

THEORETICAL MODELS OF STRUCTURAL PLASTICITY

In this section, we will summarize theoretical and computational studies analyzing the dynamics and functional consequences of structural plasticity. As models of structural plasticity basically adapt the connectivity, they enable predictions about properties of the connectivity in neural networks. These properties range from statistical (e.g., the statistics of subnetwork structures (motifs) or the probability distribution of the number of synapses between two neurons) to graph theoretical features (e.g., small-worldness or shortest path lengths) which can be compared to biological data. Many studies also investigate functional consequences of structural plasticity as, for instance, the influence on the storage capacity or the ability to classify different inputs. The majority of studies focuses on either Hebbian or homeostatic structural plasticity, however, at the end of this section, we will provide an overview of the few studies combining both processes of structural plasticity.

Hebbian Structural Plasticity

As discussed above, Hebbian structural plasticity is mainly realized by the dynamics of dendritic spines. Thus, models of Hebbian structural plasticity typically describe the dynamics of dendritic spines. Synapses in these models appear and disappear at predefined potential synaptic locations with certain probabilities influenced by neuronal activities, synaptic efficacies and/or other hidden variables. As activities and efficacies depend on synaptic plasticity, Hebbian structural plasticity and Hebbian synaptic plasticity are strongly interconnected and the majority of models of Hebbian structural plasticity also incorporate the dynamics of Hebbian synaptic plasticity and some even homeostatic synaptic plasticity.

The simplest neural network to study the influence of Hebbian structural plasticity on the network's dynamics and connectivity is a postsynaptic neuron receiving input from one presynaptic neuron. Several experiments show that the connectivity between such pairs of neurons (the probability distribution of the number of synapses) is non-trivial (Markram et al., 1997a; Feldmeyer et al., 1999, 2002, 2006; Hardingham et al., 2010): these neurons are either unconnected (no synapse) or connected by *multiple* synapses (four to five synapses). This finding does not depend on the detailed anatomy of neurons, as the number of potential synapse location is much higher than the number of realized synapses (Fares and Stepanyants, 2009). However, as theoretical models show (Deger et al., 2012; Fauth et al., 2015b), Hebbian structural plasticity yields the formation of such multi-synaptic connections in a broad range of activity levels. By changing the activity level, the number of synapses between the neurons can be adjusted providing a way to change connectivity and, thus, store information in an activity-dependent manner (Fauth et al., 2015a,b). Furthermore, although the storage capacity per synapse is decreased, information, stored in such structures, can persist

for timescales much longer than the lifetime of a single synapse, as the storage is collectively implemented by all synapses and does not rely on the existence of single ones (Fauth et al., 2015a).

Instead of considering a system consisting of one postsynaptic neuron, which receives inputs from one presynaptic neuron by multiple synapses, other studies considered a slightly more complex system: a postsynaptic neuron receiving inputs from several presynaptic neurons (note that in this system each presynaptic neuron is considered to be connected by only one synapse to the postsynaptic neuron). Here, the stability of synapses depends on the activity-dependent calcium influx; a high calcium influx causes stabilization of synapses and a low influx implies destabilization of synapses. Similar, as for the multi-synaptic connections, high neuronal activities lead to a stabilization of all synapses (Helias et al., 2008). However, for intermediate activity levels only correlated inputs are stabilized. Thus, the information stored in the connectivity could also be the information about the correlations between different inputs (Helias et al., 2008). In addition, synapses from uncorrelated inputs are pruned or deleted and lose their (noisy) influence on the postsynaptic neuron (Helias et al., 2008). Thus, Hebbian structural plasticity might help to prune synapses which are unimportant for the dynamics of the neural network.

Accordingly, also in more complex and biologically more reasonable systems, as large recurrent networks (Bourjaily and Miller, 2011; Zheng et al., 2013; Miner and Triesch, 2016), synaptic pruning preferentially removes synapses which only weakly contribute to synaptic transmission. These models use synaptic plasticity rules which typically yield a bimodal distribution of the electrical transmission efficacies with many efficacies close to zero. In combination with synaptic pruning, however, synapses with small efficacies are removed leading to the emergence of a unimodal distribution as observed in the cortex (Song et al., 2005). Accordingly, the continuous pruning and creation of synapses can also be interpreted as a process of stochastic inference, in which the network continuously tests and evaluates the "usefulness" of synapses to process or represent external stimuli (Kappel et al., 2015). Thus, synaptic pruning might minimize the resources for synaptic maintenance while preserving important dynamics.

Further advantages of pruning or deletion of uncorrelated or unimportant synapses have been revealed for simpler feedforward neuronal networks, which are typically used to study associative memory: the storage capacity of these networks is increased (Knoblauch et al., 2009, 2014). Considering a Willshaw or Hopfield network (Willshaw et al., 1969; Hopfield, 1982), the deletion of the weak or unimportant synapses increases the storage capacity per synapse without perturbing the stored patterns (Knoblauch et al., 2009). Furthermore, pruning prevents the occurrence of catastrophic forgetting and could explain phenomena as retrograde amnesia or the difference between spaced- and block-learning (Knoblauch, 2009; Knoblauch et al., 2014). Intuitively, the increase in storage capacity per synapse contradicts the finding of multi-synaptic connections described above (Deger et al., 2012; Fauth et al., 2015b). However, the influence of multi-synaptic connections on memory has to be

further investigated as models at the network level are so far missing.

So far, theoretical studies of structural plasticity in recurrent networks mostly investigated storage capacity and compared the properties of the resulting connectivity with the properties of biological measured connectivities as, for instance, the statistics of the so-called motifs (Milo et al., 2002), i.e., configurations of the connectivity in small subnetworks. In cortical networks, groups of strongly connected neurons show an increased appearance (compared to random networks; Markram et al., 1997a; Feldmeyer et al., 1999; Song et al., 2005; Perin et al., 2011). As groups of strongly connected neurons typically show strongly correlated activities, which, in turn lead to stabilization of the corresponding connections, this increased appearance is naturally reproduced by Hebbian structural plasticity interacting with synaptic plasticity (Bourjaily and Miller, 2011; Miner and Triesch, 2016, but see also Zheng and Triesch, 2014). Remarkably, with the formation of more strongly connected subgroups of neurons the network's performance in discriminating different inputs increases (Bourjaily and Miller, 2011).

In summary, these results show that Hebbian structural plasticity improves several properties of neural networks compared to networks adapted only by synaptic plasticity. Especially, the storage of memories is improved in storage lifetime, capacity, and noise robustness. Furthermore, perhaps related to these improvements in memory storage, also the ability to discriminate inputs is enhanced. However, further investigations are needed to understand the influence of Hebbian structural plasticity on the dynamics of neural networks.

Homeostatic Structural Plasticity

As already described above, homeostatic structural plasticity adapts dendrites and axons dependent on the neuronal activity to reach and sustain an intermediate activity regime (Butz et al., 2009). The slow timescale of homeostatic structural plasticity implies that its influences are basically observed after long durations, as during development, or in networks under extreme activity conditions as after lesions. Thus, also theoretical models investigating the dynamics of homeostatic structural plasticity concentrate mainly on these two paradigms.

During the development of a neural network from a naive initial state to a matured network, it passes through an overshoot phase of building up many synapses followed by a pruning phase until the network settles in the ground state (van Huizen et al., 1985, 1987; van Huizen and Romijn, 1987; van Ooyen, 2003, 2011). Such dynamics are already seen in a pure excitatory network model governed by homeostatic structural plasticity without the differentiation between axons and dendrites (van Ooyen and van Pelt, 1994). Introducing also inhibition further pronounces this overshoot effect and can lead to oscillatory and bursting neuronal activities (van Ooyen et al., 1995; van Ooyen and van Pelt, 1996). Assuming different homeostatic dynamics for axons and dendrites results in even more complex activity dynamics matching cell culture data (Tetzlaff et al., 2010). The resulting network state is the so-called critical state which is predestined for maintaining stability (Bak et al., 1987; Bak,

1996). Thus, the complex interactions between all these different homeostatic processes are important to bring the whole system into a stable state showing dynamics matching experimental data.

All of these developmental models consider the dynamics of axons and dendrites. However, as described above, also the dynamics of dendritic spines and axonal boutons are determined by homeostatic structural plasticity. Theoretical network models from the 1980s (Dammasch et al., 1986, 1988; Cromme and Dammasch, 1989) already showed that also such detailed models of homeostatic structural plasticity self-organize to reach a desired activity regime. Again, the resulting system is quite stable such that even the insertion of new neurons (by, for instance, neurogenesis in the hippocampus) does not perturb the global network state (Butz et al., 2008). Furthermore, by introducing a distance-dependency for forming new synapses, the network develops into a small-world network (Butz et al., 2014b).

The dynamics of these models can also be compared to *in vivo* measurements after input lesions or stroke-induced lesions (Butz et al., 2009; Butz and van Ooyen, 2013; Butz et al., 2014a). Interestingly, this comparison between *in vivo* and model dynamics enables conclusions on the activity-dependency of the different homeostatic processes. For instance, after a retinal lesion, neurons in the lesion projection zone (which have lost their external input) start to connect with active neurons at the border of the zone (Keck et al., 2008). In network models, this dynamics can only be seen if for small neuronal activations basically new dendritic spines are formed and axonal boutons are pruned (Butz and van Ooyen, 2013; Butz et al., 2014a). In contrast, if for small activities boutons are formed and spines are deleted, the system still reaches homeostasis, but the neurons in the lesion projection zone predominantly connect with each other and, thus, the whole zone decouples from the rest of the network (Butz and van Ooyen, 2013; Butz et al., 2014a).

Further predictions from these models are, for instance, that similar effects arise in networks with lesions after stroke (Butz et al., 2009, 2014a). Neurons affected by the deafferentiation (nearby the lesion zone) have problems in regaining their activity-homeostasis when the rest of the network is in homeostasis. This problem can be solved if, after stroke, the neurons, which are still in homeostasis, receive an external stimulation to trigger homeostatic structural plasticity and, thus, encourage rewiring. After this stimulation the whole network has reached homeostasis (Butz et al., 2009). Thus, studying the effects of structural plasticity also helps to gain insights in new potential medical treatments.

Models of the Interaction of Hebbian and Homeostatic Structural Plasticity

So far only a few models investigated the interaction between Hebbian and homeostatic structural plasticity (Levy and Desmond, 1985; Adelsberger-Mangan and Levy, 1993, 1994; Levy, 2004; Thomas et al., 2015). Basically, these models include weight-dependent synapse removal (Hebbian structural plasticity) with an activity-dependent synapse formation (homeostatic structural plasticity). The combination of these processes in a feed-forward network optimizes the information

transfer from input to output layer and also supports the separation of information in the output layer while keeping the homeostasis (Adelsberger-Mangan and Levy, 1993, 1994). As already seen in developmental models, also in the combined models the assumption of different dynamics for axons and dendrites in homeostatic structural plasticity increases the overall performance of the network (Adelsberger-Mangan and Levy, 1994). In other words, this combination of Hebbian and homeostatic structural plasticity provides an unsupervised way to transfer, compress, and store information (Hebbian structural plasticity) in an energy efficient representation, i.e., with a low number of needed neurons and low firing rates (homeostatic structural plasticity). Along this line, each post-synaptic neuron becomes selective or tuned to a specific input pattern. The number of neurons tuned to one pattern grows with the occurrence of this pattern (Thomas et al., 2015). This could, in principle, be a solution to the problem of memory allocation or rather allocation of inputs to specific groups of neurons in the brain (Rogerson et al., 2014; Tetzlaff et al., 2015). These results provide first insights into the complex dynamics

resulting from the interaction between Hebbian and homeostatic structural plasticity. However, further theoretical investigations are needed.

CONCLUSIONS AND OPEN QUESTIONS

In this review, we showed that structural plasticity can be classified into two categories (for a schematic summary, see **Figure 2**; italic and bold fonts indicate key references for experimental and theoretical studies respectively): (i) Hebbian structural plasticity leads to an increase (decrease) of the number or density of dendritic spines and contacts with axonal boutons during phases of high (low) activity (**Figure 2**, first column, orange). (ii) When these alterations in activity persist, homeostatic structural plasticity balances these changes by removing (adding new) synapses (**Figure 2**, second column, orange) and, after days, even by retracting (growing out) the dendrites themselves (**Figure 2**, second column, green).

In addition, we showed that there is a strong interaction between structural plasticity and synaptic plasticity. Both

FIGURE 2 | Schematic overview of the literature. *Rectangular boxes:* The different effects of activity-dependent plasticity for high activity (upper row) and low activity (lower row) in a Hebbian (first column) and a homeostatic manner (second column). *Colors* indicate, whether synaptic efficacies (blue), dendritic spines (orange), or axons and dendrites (green) are affected. *Names and years* outside the circles indicate key experimental (italic) and theoretical (bold) studies for the respective effect. Studies which target both activity regimes and/or plasticity types are placed in-between them.

have been demonstrated to depend on (local) calcium levels. Even more strikingly, the synaptic transmission efficacies are related to the stability of the synapses. Thus, structural plasticity and its influences on the dynamics of neural networks can only be understood in conjunction with synaptic plasticity.

These complex interactions and their functional implications are best understood by theoretical models. For instance, Hebbian structural plasticity seems to remove and create synapses selectively. This selectivity leads to experimentally measured local connectivities and, furthermore, enhances memory lifetime, storage capacity, and robustness. For this, especially the pruning of non-needed synapses plays an important role. However, as for synaptic plasticity, these Hebbian dynamics lead to a positive feedback between connectivity and activity and, thus, to increasing neuronal activities. Thus, also structural plasticity requires a homeostatic process regulating activities back to an intermediate level. Accordingly, theoretical studies show that homeostatic structural plasticity organizes the connectivity of the network to maintain network stability. The combination of Hebbian and homeostatic structural plasticity preserves and improves network functions, as memory storage and input discrimination, and, in parallel, stabilizes the global dynamics in a resource efficient manner.

Still there are many open questions summarized in the following. On the experimental side, it is, for example, still unclear whether the increased number of spines after LTP-inducing stimuli results from increased stabilization or from an increased spine formation. Also, the relation between structural and synaptic plasticity is still not completely understood. Along these lines, especially the relation between homeostatic structural plasticity and synaptic scaling has not been completely unraveled yet.

Furthermore, we argued that the interaction of Hebbian and homeostatic structural plasticity will lead to a transient increase (decrease) of, e.g., the number of spines in a system which undergoes prolonged phases of enhanced (decreased) activity. Such transient increases are observed in experiments. However, experimental investigations, whether the dynamics occur due to this interaction, are still missing. In general, the interaction of Hebbian and homeostatic processes in the same system is difficult to tackle and has been addressed only by a few studies.

Accordingly, also most of the theoretical studies have focused on either Hebbian or homeostatic structural plasticity. The interaction of both mechanisms, especially in recurrent networks, is widely unknown. Moreover, theoretical studies are often restricted to either reproducing biological data or investigating functional consequences of structural plasticity. Therefore, more studies are needed to link experimentally obtained connectivity features to functional predictions.

The here reviewed theoretical models mostly considered point-neurons. However, the actual position or location of a synapse on the dendritic tree strongly influences the details of synaptic plasticity (Sjöström and Häusser, 2006). In addition, also neighborly relations between synapses influence via, for instance, calcium currents synaptic plasticity (Oh et al., 2015). Obviously, due to the interactions between synaptic and structural plasticity, these local influences on synaptic plasticity also affect structural plasticity. On the other side, structural plasticity might select synapses with certain synaptic plasticity rules and remove others. Thereby, structural plasticity could act as a metaplasticity-like process (Abraham, 2008) which adds another level of complexity to the interaction of the different plasticities.

Taken together, we already have a decent understanding of the basic mechanisms governing Hebbian and homeostatic structural changes. Yet, their interaction with each other and with synaptic plasticity, as well as their functional relevance, still leave many open questions.

AUTHOR CONTRIBUTIONS

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

FUNDING

ACKNOWLEDGMENTS

REFERENCES

Abraham, W. C. (2008). Metaplasticity: tuning synapses and networks for plasticity. *Nat. Rev. Neurosci.* 9, 387–399. doi: 10.1038/nrn2356

Adelsberger-Mangan, D. M., and Levy, W. B. (1993). Adaptive synaptogenesis constructs networks that maintain information and reduce statistical dependence. *Biol. Cybern.* 70, 81–87. doi: 10.1007/BF00202569

Adelsberger-Mangan, D. M., and Levy, W. B. (1994). The influence of limited presynaptic growth and synapse removal on adaptive synaptogenesis. *Biol. Cybern.* 71, 461–468. doi: 10.1007/BF00198922

Annis, C. M., O'Dowd, D. K., and Robertson, R. T. (1994). Activity-dependent regulation of dendritic spine density on cortical pyramidal neurons in organotypic slice cultures. *J. Neurobiol.* 25, 1483–1493.

Bak, P. (1996). *How Nature Works: The Science of Self-Organized Criticality, 1st Edn.* New York, NY: Springer.

Bak, P., Tang, C., and Wiesenfeld, K. (1987). Self-organized criticality: an explanation of 1/f noise. *Phys. Rev. Lett.* 59, 381–384. doi: 10.1103/PhysRevLett.59.381

Barker, J. L., Behar, T., Li, Y. X., Liu, Q. Y., Ma, W., Maric, D., et al. (1998). GABAergic cells and signals in CNS development. *Perspect. Dev. Neurobiol.* 5, 305–322.

Bastrikova, N., Gardner, G. A., Reece, J. M., Jeromin, A., and Dudek, S. M. (2008). Synapse elimination accompanies functional plasticity in hippocampal neurons. *Proc. Natl. Acad. Sci. U.S.A.* 105, 3123–3127. doi: 10.1073/pnas.0800027105

Beattie, E. C., Carroll, R. C., Yu, X., Morishita, W., Yasuda, H., von Zastrow, M., et al. (2000). Regulation of AMPA receptor endocytosis by a signaling mechanism shared with LTD. *Nat. Neurosci.* 3, 1291–1300. doi: 10.1038/81823

Becker, N., Wierenga, C. J., Fonseca, R., Bonhoeffer, T., and Nägerl, U. V. (2008). LTD induction causes morphological changes of presynaptic boutons and reduces their contacts with spines. *Neuron* 60, 590–597. doi: 10.1016/j.neuron.2008.09.018

Ben-Ari, Y. (2002). Excitatory actions of GABA during development: the nature of the nurture. *Nat. Rev. Neurosci.* 3, 728–739. doi: 10.1038/nrn920

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.

Bixby, J. L., Grunwald, G. B., and Bookman, R. J. (1994). Ca2+ influx and neurite growth in response to purified N-cadherin and laminin. *J. Cell Biol.* 127, 1461–1475. doi: 10.1083/jcb.127.5.1461

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

Bliss, T. V. P., and Collingridge, G. L. (1993). A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361, 31–39. doi: 10.1038/361031a0

Bock, J., and Braun, K. (1999). Blockade of N-methyl-D-aspartate receptor activation suppresses learning-induced synaptic elimination. *Proc. Natl. Acad. Sci. U.S.A.* 96, 2485–2490. doi: 10.1073/pnas.96.5.2485

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

Bourne, J. N., Kirov, S. A., Sorra, K. E., and Harris, K. M. (2007). Warmer preparation of hippocampal slices prevents synapse proliferation that might obscure LTP-related structural plasticity. *Neuropharmacology* 52, 55–59. doi: 10.1016/j.neuropharm.2006.06.020

Bravin, M., Morando, L., Vercelli, A., Rossi, F., and Strata, P. (1999). Control of spine formation by electrical activity in the adult rat cerebellum. *Proc. Natl. Acad. Sci. U.S.A.* 96, 1704–1709. doi: 10.1073/pnas.96.4.1704

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

Butz, M., Steenbuck, I. D., and van Ooyen, A. (2014a). Homeostatic structural plasticity can account for topology changes following deafferentation and focal stroke. *Front. Neuroanat.* 8:115. doi: 10.3389/fnana.2014.00115

Butz, M., Steenbuck, I. D., and van Ooyen, A. (2014b). Homeostatic structural plasticity increases the efficiency of small-world networks. *Front. Synaptic Neurosci.* 6:7. doi: 10.3389/fnsyn.2014.00007

Butz, M., Teuchert-Noodt, G., Grafen, K., and van Ooyen, A. (2008). Inverse relationship between adult hippocampal cell proliferation and synaptic rewiring in the dentate gyrus. *Hippocampus* 18, 879–898. doi: 10.1002/hipo.20445

Butz, M., and van Ooyen, A. (2013). A simple rule for dendritic spine and axonal bouton formation can account for cortical reorganization after focal retinal lesions. *PLoS Comput. Biol.* 9:e1003259. doi: 10.1371/journal.pcbi.1003259

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

Canty, A. J., Huang, L., Jackson, J. S., Little, G. E., Knott, G., Maco, B., et al. (2013). *In-vivo* single neuron axotomy triggers axon regeneration to restore synaptic density in specific cortical circuits. *Nat. Commun.* 4:2038. doi: 10.1038/ncomms3038

Chang, F. L., and Greenough, W. T. (1984). Transient and enduring morphological correlates of synaptic activity and efficacy change in the rat hippocampal slice. *Brain Res.* 309, 35–46. doi: 10.1016/0006-8993(84)91008-4

Chen, C.-C., Bajnath, A., and Brumberg, J. C. (2015). The impact of development and sensory deprivation on dendritic protrusions in the mouse barrel cortex. *Cereb. Cortex* 25, 1638–1653. doi: 10.1093/cercor/bht415

Chiappalone, M., Bove, M., Vato, A., Tedesco, M., and Martinoia, S. (2006). Dissociated cortical networks show spontaneously correlated activity patterns during *in vitro* development. *Brain Res.* 1093, 41–53. doi: 10.1016/j.brainres.2006.03.049

Chklovskii, D. B., Mel, B. W., and Svoboda, K. (2004). Cortical rewiring and information storage. *Nature* 431, 782–788. doi: 10.1038/nature03012

Cohan, C. S., and Kater, S. B. (1986). Suppression of neurite elongation and growth cone motility by electrical activity. *Science* 232, 1638–1640. doi: 10.1126/science.3715470

Collin, C., Miyaguchi, K., and Segal, M. (1997). Dendritic spine density and LTP induction in cultured hippocampal slices. *J. Neurophysiol.* 77, 1614–1623.

Colling, S. B., Man, W. D., Draguhn, A., and Jefferys, J. G. (1996). Dendritic shrinkage and dye-coupling between rat hippocampal CA1 pyramidal cells in the tetanus toxin model of epilepsy. *Brain Res.* 741, 38–43. doi: 10.1016/S0006-8993(96)00884-0

Collingridge, G. L., Peineau, S., Howland, J. G., and Wang, Y. T. (2010). Long-term depression in the CNS. *Nat. Rev. Neurosci.* 11, 459–473. doi: 10.1038/nrn2867

Cromme, L. J., and Dammasch, I. E. (1989). Compensation type algorithms for neural nets: stability and convergence. *J. Math. Biol.* 27, 327–340. doi: 10.1007/BF00275816

Dalva, M. B., Ghosh, A., and Shatz, C. J. (1994). Independent control of dendritic and axonal form in the developing lateral geniculate nucleus. *J. Neurosci.* 14, 3588–3602.

Dammasch, I. E., Wagner, G. P., and Wolff, J. R. (1986). Self-stabilization of neuronal networks. I. The compensation algorithm for synaptogenesis. *Biol. Cybern.* 54, 211–222. doi: 10.1007/BF00318417

Dammasch, I. E., Wagner, G. P., and Wolff, J. R. (1988). Self-stabilization of neuronal networks. II. Stability conditions for synaptogenesis. *Biol. Cybern.* 58, 149–158. doi: 10.1007/BF00364134

Darian-Smith, C., and Gilbert, C. D. (1994). Axonal sprouting accompanies functional reorganization in adult cat striate cortex. *Nature* 368, 737–740. doi: 10.1038/368737a0

Deger, M., Helias, M., Rotter, S., and Diesmann, M. (2012). Spike-timing dependence of structural plasticity explains cooperative synapse formation in the neocortex. *PLoS Comput. Biol.* 8:e1002689. doi: 10.1371/journal.pcbi.1002689

Deller, T., Bas Orth, C., Vlachos, A., Merten, T., Del Turco, D., Dehn, D., et al. (2006). Plasticity of synaptopodin and the spine apparatus organelle in the rat fascia dentata following entorhinal cortex lesion. *J. Comp. Neurol.* 499, 471–484. doi: 10.1002/cne.21103

Desmond, N. L., and Levy, W. B. (1990). Morphological correlates of long-term potentiation imply the modification of existing synapses, not synaptogenesis, in the hippocampal dentate gyrus. *Synapse* 5, 139–143. doi: 10.1002/syn.890050208

Drakew, A., Müller, M., Gähwiler, B. H., Thompson, S. M., and Frotscher, M. (1996). Spine loss in experimental epilepsy: quantitative light and electron microscopic analysis of intracellularly stained CA3 pyramidal cells in hippocampal slice cultures. *Neuroscience* 70, 31–45. doi: 10.1016/0306-4522(95)00379-W

Dudai, Y. (2004). The neurobiology of consolidations, or, how stable is the engram? *Annu. Rev. Psychol.* 55, 51–86. doi: 10.1146/annurev.psych.55.090902.142050

Dudek, S. M., and Bear, M. F. (1992). Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade. *Proc. Natl. Acad. Sci. U.S.A.* 89, 4363–4367. doi: 10.1073/pnas.89.10.4363

Dunaevsky, A., Tashiro, A., Majewska, A., Mason, C., and Yuste, R. (1999). Developmental regulation of spine motility in the mammalian central nervous system. *Proc. Natl. Acad. Sci. U.S.A.* 96, 13438–13443. doi: 10.1073/pnas.96.23.13438

Engert, F., and Bonhoeffer, T. (1999). Dendritic spine changes associated with hippocampal long-term synaptic plasticity. *Nature* 399, 66–70. doi: 10.1038/19978

Fares, T., and Stepanyants, A. (2009). Cooperative synapse formation in the neocortex. *Proc. Natl. Acad. Sci. U.S.A.* 106, 16463–16468. doi: 10.1073/pnas.0813265106

Fauth, M., Wörgötter, F., and Tetzlaff, C. (2015a). Formation and maintenance of robust long-term information storage in the presence of synaptic turnover. *PLoS Comput. Biol.* 11:e1004684. doi: 10.1371/journal.pcbi.1004684

Fauth, M., Wörgötter, F., and Tetzlaff, C. (2015b). The formation of multi-synaptic connections by the interaction of synaptic and structural plasticity and their functional consequences. *PLoS Comput. Biol.* 11:e1004031. doi: 10.1371/journal.pcbi.1004031

Feldman, D. E. (2009). Synaptic mechanisms for plasticity in neocortex. *Annu. Rev. Neurosci.* 32, 33–55. doi: 10.1146/annurev.neuro.051508.135516

Feldman, M. L., and Dowd, C. (1975). Loss of dendritic spines in aging cerebral cortex. *Anat. Embryol. (Berl.)* 148, 279–301. doi: 10.1007/BF00319848

Feldmeyer, D., Egger, V., Lübke, J., and Sakmann, B. (1999). Reliable synaptic connections between pairs of excitatory layer 4 neurones within a single 'barrel' of developing rat somatosensory cortex. *J. Physiol.* 521(Pt 1):169–190. doi: 10.1111/j.1469-7793.1999.00169.x

Feldmeyer, D., Lübke, J., and Sakmann, B. (2006). Efficacy and connectivity of intracolumnar pairs of layer 2/3 pyramidal cells in the barrel cortex of juvenile rats. *J. Physiol.* 575(Pt 2):583–602. doi: 10.1113/jphysiol.2006.105106

Feldmeyer, D., Lübke, 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.* 538(Pt 3):803–822. doi: 10.1113/jphysiol.2001.012959

Fields, R. D., Neale, E. A., and Nelson, P. G. (1990). Effects of patterned electrical activity on neurite outgrowth from mouse sensory neurons. *J. Neurosci.* 10, 2950–2964.

Fifková, E., and Van Harreveld, A. (1977). Long-lasting morphological changes in dendritic spines of dentate granular cells following stimulation of the entorhinal area. *J. Neurocytol.* 6, 211–230. doi: 10.1007/BF01261506

Flores, C. E., and Méndez, P. (2014). Shaping inhibition: activity dependent structural plasticity of GABAergic synapses. *Front. Cell. Neurosci.* 8:327. doi: 10.3389/fncel.2014.00327

Friedman, H. V., Bresler, T., Garner, C. C., and Ziv, N. E. (2000). Assembly of new individual excitatory synapses: time course and temporal order of synaptic molecule recruitment. *Neuron* 27, 57–69. doi: 10.1016/S0896-6273(00)00009-X

Garyantes, T. K., and Regehr, W. G. (1992). Electrical activity increases growth cone calcium but fails to inhibit neurite outgrowth from rat sympathetic neurons. *J. Neurosci.* 12, 96–103.

Geinisman, Y., Morrell, F., and deToledo Morrell, L. (1990). Increase in the relative proportion of perforated axospinous synapses following hippocampal kindling is specific for the synaptic field of stimulated axons. *Brain Res.* 507, 325–331. doi: 10.1016/0006-8993(90)90291-I

Gerstner, W., Kempter, R., van Hemmen, J. L., and Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature* 383, 76–78. doi: 10.1038/383076a0

Gerstner, W., and Kistler, W. M. (2002). Mathematical formulations of Hebbian learning. *Biol. Cybern.* 87, 404–415. doi: 10.1007/s00422-002-0353-y

Goldowitz, D., Scheff, S. W., and Cotman, C. W. (1979). The specificity of reactive synaptogenesis: a comparative study in the adult rat hippocampal formation. *Brain Res.* 170, 427–441. doi: 10.1016/0006-8993(79)90962-4

Groc, L., Petanjek, Z., Gustafsson, B., Ben-Ari, Y., Hanse, E., and Khazipov, R. (2002). *In vivo* blockade of neural activity alters dendritic development of neonatal CA1 pyramidal cells. *Eur. J. Neurosci.* 16, 1931–1938. doi: 10.1046/j.1460-9568.2002.02264.x

Grutzendler, J., Kasthuri, N., and Gan, W.-B. (2002). Long-term dendritic spine stability in the adult cortex. *Nature* 420, 812–816. doi: 10.1038/nature01276

Haas, K., Li, J., and Cline, H. T. (2006). AMPA receptors regulate experience-dependent dendritic arbor growth *in vivo*. *Proc. Natl. Acad. Sci. U.S.A.* 103, 12127–12131. doi: 10.1073/pnas.0602670103

Halpain, S., Hipolito, A., and Saffer, L. (1998). Regulation of F-actin stability in dendritic spines by glutamate receptors and calcineurin. *J. Neurosci.* 18, 9835–9844.

Hardingham, N. R., Read, J. C. A., Trevelyan, A. J., Nelson, J. C., Jack, J. J. B., and Bannister, N. J. (2010). Quantal analysis reveals a functional correlation between presynaptic and postsynaptic efficacy in excitatory connections from rat neocortex. *J. Neurosci.* 30, 1441–1451. doi: 10.1523/JNEUROSCI.3244-09.2010

Haydon, P. G., McCobb, D. P., and Kater, S. B. (1984). Serotonin selectively inhibits growth cone motility and synaptogenesis of specific identified neurons. *Science* 226, 561–564. doi: 10.1126/science.6093252

Haydon, P. G., McCobb, D. P., and Kater, S. B. (1987). The regulation of neurite outgrowth, growth cone motility, and electrical synaptogenesis by serotonin. *J. Neurobiol.* 18, 197–215. doi: 10.1002/neu.480180206

Hebb, D. (1949). *The Organization of Behavior: A Neuropsychological Theory. Wiley Book in Clinical Psychology.* New York, NY: Wiley.

Helias, M., Rotter, S., Gewaltig, M.-O., and Diesmann, M. (2008). Structural plasticity controlled by calcium based correlation detection. *Front. Comput. Neurosci.* 2:7. doi: 10.3389/neuro.10.007.2008

Helmchen, F., Imoto, K., and Sakmann, B. (1996). Ca2+ buffering and action potential-evoked Ca2+ signaling in dendrites of pyramidal neurons. *Biophys. J.* 70, 1069–1081. doi: 10.1016/S0006-3495(96)79653-4

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

Higley, M. J., and Sabatini, B. L. (2008). Calcium signaling in dendrites and spines: practical and functional considerations. *Neuron* 59, 902–913. doi: 10.1016/j.neuron.2008.08.020

Hill, T. C., and Zito, K. (2013). LTP-induced long-term stabilization of individual nascent dendritic spines. *J. Neurosci.* 33, 678–686. doi: 10.1523/JNEUROSCI.1404-12.2013

Hofer, S. B., Mrsic-Flogel, T. D., Bonhoeffer, T., and Hübener, M. (2009). Experience leaves a lasting structural trace in cortical circuits. *Nature* 457, 313–317. doi: 10.1038/nature07487

Holtmaat, A. J. G. D., Trachtenberg, J. T., Wilbrecht, L., Shepherd, G. M., Zhang, X., Knott, G. W., et al. (2005). Transient and persistent dendritic spines in the neocortex *in vivo*. *Neuron* 45, 279–291. doi: 10.1016/j.neuron.2005.01.003

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.* 79, 2554–2558. doi: 10.1073/pnas.79.8.2554

Hou, Q., Gilbert, J., and Man, H. Y. (2011). Homeostatic regulation of AMPA receptor trafficking and degradation by light-controlled single-synaptic activation. *Neuron* 72, 806–818. doi: 10.1016/j.neuron.2011.10.011

Hou, Q., Zhang, D., Jarzylo, L., Huganir, 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

Hua, J. Y., and Smith, S. J. (2004). Neural activity and the dynamics of central nervous system development. *Nat. Neurosci.* 7, 327–332. doi: 10.1038/nn1218

Hübener, M., and Bonhoeffer, T. (2010). Searching for engrams. *Neuron* 67, 363–371. doi: 10.1016/j.neuron.2010.06.033

Hutchins, B. I., and Kalil, K. (2008). Differential outgrowth of axons and their branches is regulated by localized calcium transients. *J. Neurosci.* 28, 143–153. doi: 10.1523/JNEUROSCI.4548-07.2008

Huttenlocher, P. R. (1984). Synapse elimination and plasticity in developing human cerebral cortex. *Am. J. Ment. Defic.* 88, 488–496.

Huttenlocher, P. R., de Courten, C., Garey, L. J., and Van der Loos, H. (1982). Synaptogenesis in human visual cortex–evidence for synapse elimination during normal development. *Neurosci. Lett.* 33, 247–252. doi: 10.1016/0304-3940(82)90379-2

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

Isokawa, M. (1998). Remodeling dendritic spines in the rat pilocarpine model of temporal lobe epilepsy. *Neurosci. Lett.* 258, 73–76. doi: 10.1016/S0304-3940(98)00848-9

Isokawa, M., and Levesque, M. F. (1991). Increased NMDA responses and dendritic degeneration in human epileptic hippocampal neurons in slices. *Neurosci. Lett.* 132, 212–216. doi: 10.1016/0304-3940(91)90304-C

Jiang, M., Lee, C. L., Smith, K. L., and Swann, J. W. (1998). Spine loss and other persistent alterations of hippocampal pyramidal cell dendrites in a model of early-onset epilepsy. *J. Neurosci.* 18, 8356–8368.

Kalisman, N., Silberberg, G., and Markram, H. (2005). The neocortical microcircuit as a tabula rasa. *Proc. Natl. Acad. Sci. U.S.A.* 102, 880–885. doi: 10.1073/pnas.0407088102

Kappel, D., Habenschuss, S., Legenstein, R., and Maass, W. (2015). Network plasticity as Bayesian inference. *PLoS Comput. Biol.* 11:e1004485. doi: 10.1371/journal.pcbi.1004485

Kasai, H., Matsuzaki, M., Noguchi, J., Yasumatsu, N., and Nakahara, H. (2003). Structure-stability-function relationships of dendritic spines. *Trends Neurosci.* 26, 360–368. doi: 10.1016/S0166-2236(03)00162-0

Kater, S. B., Mattson, M. P., and Guthrie, P. B. (1989). Calcium-induced neuronal degeneration: a normal growth cone regulating signal gone awry (?). *Ann. N. Y. Acad. Sci.* 568, 252–261. doi: 10.1111/j.1749-6632.1989.tb12514.x

Kauer, J. A., Malenka, R. C., and Nicoll, R. A. (1988). A persistent postsynaptic modification mediates long-term potentiation in the hippocampus. *Neuron* 1, 911–917. doi: 10.1016/0896-6273(88)90148-1

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

Keck, T., Mrsic-Flogel, T. D., Vaz Alfonso, M., Eysel, U. T., Bonhoeffer, T., and Hübener, M. (2008). Massive restructuring of neuronal circuits during functional reorganization of adult visual cortex. *Nat. Neurosci.* 11, 1162–1167. doi: 10.1038/nn.2181

Kim, J., Tsien, R. W., and Alger, B. E. (2012). An improved test for detecting multiplicative homeostatic synaptic scaling. *PLoS ONE* 7:e37364. doi: 10.1371/journal.pone.0037364

Kirkwood, A., Rioult, M. G., and Bear, M. F. (1996). Experience-dependent modification of synaptic plasticity in visual cortex. *Nature* 381, 526–528. doi: 10.1038/381526a0

Kirov, S. A., Goddard, C. A., and Harris, K. M. (2004). Age-dependence in the homeostatic upregulation of hippocampal dendritic spine number during blocked synaptic transmission. *Neuropharmacology* 47, 640–648. doi: 10.1016/j.neuropharm.2004.07.039

Kirov, S. A., and Harris, K. M. (1999). Dendrites are more spiny on mature hippocampal neurons when synapses are inactivated. *Nat. Neurosci.* 2, 878–883. doi: 10.1038/13178

Kirov, S. A., Sorra, K. E., and Harris, K. M. (1999). Slices have more synapses than perfusion-fixed hippocampus from both young and mature rats. *J. Neurosci.* 19, 2876–2886.

Knoblauch, A. (2009). "The role of structural plasticity and synaptic consolidation for memory and amnesia in a model of cortico-hippocampal interplay," in *Connectionist Models of Behavior and Cognition II: Proceedings of the 11th Neural Computation and Psychology Workshop* (Singapore: Wold Scientific Publishing), 79–90. doi: 10.1142/9789812834232_0007

Knoblauch, A., Körner, E., Körner, U., and Sommer, F. T. (2014). Structural synaptic plasticity has high memory capacity and can explain graded amnesia, catastrophic forgetting, and the spacing effect. *PLoS ONE* 9:e96485. doi: 10.1371/journal.pone.0096485

Knoblauch, A., Palm, G., and Sommer, F. T. (2009). Memory capacities for synaptic and structural plasticity. *Neural Comput.* 22, 289–341. doi: 10.1162/neco.2009.08-07-588

Knott, G. W., Holtmaat, A., Wilbrecht, L., Welker, E., and Svoboda, K. (2006). Spine growth precedes synapse formation in the adult neocortex *in vivo*. *Nat. Neurosci.* 9, 1117–1124. doi: 10.1038/nn1747

Kolodziejski, C., Tetzlaff, C., and Wörgötter, F. (2010). Closed-form treatment of the interactions between neuronal activity and timing-dependent plasticity in networks of linear neurons. *Front. Comput. Neurosci.* 4:134. doi: 10.3389/fncom.2010.00134

Kwon, H.-B., and Sabatini, B. L. (2011). Glutamate induces de novo growth of functional spines in developing cortex. *Nature* 474, 100–104. doi: 10.1038/nature09986

Lamprecht, R., and LeDoux, J. (2004). Structural plasticity and memory. *Nat. Rev. Neurosci.* 5, 45–54. doi: 10.1038/nrn1301

Le Bé, J.-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

Lee, K. S., Schottler, F., Oliver, M., and Lynch, G. (1980). Brief bursts of high-frequency stimulation produce two types of structural change in rat hippocampus. *J. Neurophysiol.* 44, 247–258.

Lee, W.-C. A., Huang, H., Feng, G., Sanes, J. R., Brown, E. N., So, P. T., et al. (2006). Dynamic remodeling of dendritic arbors in GABAergic interneurons of adult visual cortex. *PLoS Biol.* 4:e29. doi: 10.1371/journal.pbio.0040029

Leuner, B., Falduto, J., and Shors, T. J. (2003). Associative memory formation increases the observation of dendritic spines in the hippocampus. *J. Neurosci.* 23, 659–665. Available online at: http://www.jneurosci.org/content/23/2/659.long

Levy, W. B. (2004). Contrasting rules for synaptogenesis, modification of existing synapses, and synaptic removal as a function of neuronal computation. *Neurocomputing* 58, 343–350. doi: 10.1016/j.neucom.2004.01.065

Levy, W. B., and Desmond, N. L. (1985). "The rules of elemental synaptic plasticity," in *Synaptic Modification, Neuron Selectivity and Nervous System Organisation*, eds W. B. Levy, J. A. Anderson, and S. Lehmkuhle (Hillsdale, NJ: Lawrence Erlbaum Associates), 105–121.

Levy, W. B., and Steward, O. (1983). Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience* 8, 791–797. doi: 10.1016/0306-4522(83)90010-6

Lipton, S. A., and Kater, S. B. (1989). Neurotransmitter regulation of neuronal outgrowth, plasticity and survival. *Trends Neurosci.* 12, 265–270. doi: 10.1016/0166-2236(89)90026-X

Loewenstein, Y., Yanover, U., and Rumpel, S. (2015). Predicting the dynamics of network connectivity in the neocortex. *J. Neurosci.* 35, 12535–12544. doi: 10.1523/JNEUROSCI.2917-14.2015

Lohmann, C., Finski, A., and Bonhoeffer, T. (2005). Local calcium transients regulate the spontaneous motility of dendritic filopodia. *Nat. Neurosci.* 8, 305–312. doi: 10.1038/nn1406

Lohmann, C., Myhr, K. L., and Wong, R. O. L. (2002). Transmitter-evoked local calcium release stabilizes developing dendrites. *Nature* 418, 177–181. doi: 10.1038/nature00850

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

Lynch, G. S., Larson, J., Kelso, S., Barrionuevo, G., and Schottler, F. (1983). Intracellular injections of EGTA block induction of hippocampal long-term potentiation. *Nature* 305, 719–721. doi: 10.1038/305719a0

Majewska, A. K., Newton, J. R., and Sur, M. (2006). Remodeling of synaptic structure in sensory cortical areas *in vivo*. *J. Neurosci.* 26, 3021–3029. doi: 10.1523/JNEUROSCI.4454-05.2006

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

Malenka, R. C., Lancaster, B., and Zucker, R. S. (1992). Temporal limits on the rise in postsynaptic calcium required for the induction of long-term potentiation. *Neuron* 9, 121–128. doi: 10.1016/0896-6273(92)90227-5

Maletic-Savatic, M., Malinow, R., and Svoboda, K. (1999). Rapid dendritic morphogenesis in CA1 hippocampal dendrites induced by synaptic activity. *Science* 283, 1923–1927. doi: 10.1126/science.283.5409.1923

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

Marik, S. A., Yamahachi, H., McManus, J. N. J., Szabo, G., and Gilbert, C. D. (2010). Axonal dynamics of excitatory and inhibitory neurons in somatosensory cortex. *PLoS Biol.* 8:e1000395. doi: 10.1371/journal.pbio.1000395

Markram, H., Lübke, J., Frotscher, M., Roth, A., and Sakmann, B. (1997a). Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex. *J. Physiol.* 500(Pt 2), 409–440. doi: 10.1113/jphysiol.1997.sp022031

Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997b). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–215. doi: 10.1126/science.275.5297.213

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

Martin, S. J., Grimwood, P. D., and Morris, R. G. (2000). Synaptic plasticity and memory: an evaluation of the hypothesis. *Annu. Rev. Neurosci.* 23, 649–711. doi: 10.1146/annurev.neuro.23.1.649

Matsuzaki, M., Ellis-Davies, G. C., Nemoto, T., Miyashita, Y., Iino, M., and Kasai, H. (2001). Dendritic spine geometry is critical for AMPA receptor expression in hippocampal CA1 pyramidal neurons. *Nat. Neurosci.* 4:1086–1092. doi: 10.1038/nn736

Matsuzaki, M., Honkura, N., Ellis-Davies, G. C. R., and Kasai, H. (2004). Structural basis of long-term potentiation in single dendritic spines. *Nature* 429, 761–766. doi: 10.1038/nature02617

Mattson, M. P. (1988). Neurotransmitters in the regulation of neuronal cytoarchitecture. *Brain Res.* 472, 179–212. doi: 10.1016/0165-0173(88)90020-3

Mattson, M. P., Dou, P., and Kater, S. B. (1988). Outgrowth-regulating actions of glutamate in isolated hippocampal pyramidal neurons. *J. Neurosci.* 8, 2087–2100.

Mattson, M. P., and Kater, S. B. (1987). Calcium regulation of neurite elongation and growth cone motility. *J. Neurosci.* 7, 4034–4043.

Mattson, M. P., and Kater, S. B. (1989). Excitatory and inhibitory neurotransmitters in the generation and degeneration of hippocampal neuroarchitecture. *Brain Res.* 478, 337–348. doi: 10.1016/0006-8993(89)91514-X

McAllister, A. K., Katz, L. C., and Lo, D. C. (1996). Neurotrophin regulation of cortical dendritic growth requires activity. *Neuron* 17, 1057–1064. doi: 10.1016/S0896-6273(00)80239-1

Medvedev, N. I., Dallérac, G., Popov, V. I., Rodriguez Arellano, J. J., Davies, H. A., Kraev, I. V., et al. (2014). Multiple spine boutons are formed after long-lasting LTP in the awake rat. *Brain Struct. Funct.* 219, 407–414. doi: 10.1007/s00429-012-0488-0

Miller, K. D. (1996). Synaptic economics: competition and cooperation in synaptic plasticity. *Neuron* 17, 371–374. doi: 10.1016/S0896-6273(00)80169-5

Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D., and Alon, U. (2002). Network motifs: simple building blocks of complex networks. *Science* 298, 824–827. doi: 10.1126/science.298.5594.824

Miner, D., and Triesch, J. (2016). Plasticity-driven self-organization under topological constraints aaccount for nonrandom features of cortical synaptic wiring. *PLoS Comput. Biol.* 12:e1004759. doi: 10.1371/journal.pcbi.1004759

Miquelajauregui, A., Kribakaran, S., Mostany, R., Badaloni, A., Consalez, G. G., and Portera-Cailliau, C. (2015). Layer 4 pyramidal neurons exhibit robust dendritic spine plasticity *in vivo* after input deprivation. *J. Neurosci.* 35, 7287–7294. doi: 10.1523/JNEUROSCI.5215-14.2015

Mizrahi, A., Crowley, J. C., Shtoyerman, E., and Katz, L. C. (2004). High-resolution *in vivo* imaging of hippocampal dendrites and spines. *J. Neurosci.* 24, 3147–3151. doi: 10.1523/JNEUROSCI.5218-03.2004

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

Moser, M. B., Trommald, M., and Andersen, P. (1994). An increase in dendritic spine density on hippocampal CA1 pyramidal cells following spatial learning in adult rats suggests the formation of new synapses. *Proc. Natl. Acad. Sci. U.S.A.* 91, 12673–12675. doi: 10.1073/pnas.91.26.12673

Mulkey, R. M., and Malenka, R. C. (1992). Mechanisms underlying induction of homosynaptic long-term depression in area CA1 of the hippocampus. *Neuron* 9, 967–975. doi: 10.1016/0896-6273(92)90248-C

Muller, D., and Lynch, G. (1988). Long-term potentiation differentially affects two components of synaptic responses in hippocampus. *Proc. Natl. Acad. Sci. U.S.A.* 85, 9346–9350. doi: 10.1073/pnas.85.23.9346

Müller, M., Gähwiler, B. H., Rietschin, L., and Thompson, S. M. (1993). Reversible loss of dendritic spines and altered excitability after chronic epilepsy in hippocampal slice cultures. *Proc. Natl. Acad. Sci. U.S.A.* 90, 257–261. doi: 10.1073/pnas.90.1.257

Murthy, V. N., Schikorski, T., Stevens, C. F., and Zhu, Y. (2001). Inactivity produces increases in neurotransmitter release and synapse size. *Neuron* 32, 673–682. doi: 10.1016/S0896-6273(01)00500-1

Nägerl, U. V., Eberhorn, N., Cambridge, S. B., and Bonhoeffer, T. (2004). Bidirectional activity-dependent morphological plasticity in hippocampal neurons. *Neuron* 44, 759–767. doi: 10.1016/j.neuron.2004.11.016

Nägerl, U. V., Köstinger, G., Anderson, J. C., Martin, K. A. C., and Bonhoeffer, T. (2007). Protracted synaptogenesis after activity-dependent spinogenesis in hippocampal neurons. *J. Neurosci.* 27, 8149–8156. doi: 10.1523/JNEUROSCI.0511-07.2007

Nägerl, U. V., Willig, K. I., Hein, B., Hell, S. W., and Bonhoeffer, T. (2008). Live-cell imaging of dendritic spines by STED microscopy. *Proc. Natl. Acad. Sci. U.S.A.* 105, 18982–18987. doi: 10.1073/pnas.0810028105

Nakayama, K., Kiyosue, K., and Taguchi, T. (2005). Diminished neuronal activity increases neuron-neuron connectivity underlying silent synapse formation and the rapid conversion of silent to functional synapses. *J. Neurosci.* 25, 4040–4051. doi: 10.1523/JNEUROSCI.4115-04.2005

Nikonenko, I., Jourdain, P., and Muller, D. (2003). Presynaptic remodeling contributes to activity-dependent synaptogenesis. *J. Neurosci.* 23, 8498–8505. Available online at: http://www.jneurosci.org/content/23/24/8498.long

Ninan, I., Liu, S., Rabinowitz, D., and Arancio, O. (2006). Early presynaptic changes during plasticity in cultured hippocampal neurons. *EMBO J.* 25, 4361–4371. doi: 10.1038/sj.emboj.7601318

Oh, W. C., Hill, T. C., and Zito, K. (2013). Synapse-specific and size-dependent mechanisms of spine structural plasticity accompanying synaptic weakening. *Proc. Natl. Acad. Sci. U.S.A.* 110, E305–E312. doi: 10.1073/pnas.1214705110

Oh, W. C., Parajuli, L. K., and Zito, K. (2015). Heterosynaptic structural plasticity on local dendritic segments of hippocampal CA1 neurons. *Cell Rep.* 10, 162–169. doi: 10.1016/j.celrep.2014.12.016

Oja, E. (1982). A simplified neuron model as a principal component analyzer. *J. Math. Biol.* 15, 267–273. doi: 10.1007/BF00275687

Okamoto, K.-I., Nagai, T., Miyawaki, A., and Hayashi, Y. (2004). Rapid and persistent modulation of actin dynamics regulates postsynaptic reorganization underlying bidirectional plasticity. *Nat. Neurosci.* 7, 1104–1112. doi: 10.1038/nn1311

Paola, V. D., Holtmaat, A., Knott, G., Song, S., Wilbrecht, L., Caroni, P., et al. (2006). Cell type-specific structural plasticity of axonal branches and boutons in the adult neocortex. *Neuron* 49, 861–875. doi: 10.1016/j.neuron.2006.02.017

Parnavelas, J. G., Lynch, G., Brecha, N., Cotman, C. W., and Globus, A. (1974). Spine loss and regrowth in hippocampus following deafferentation. *Nature* 248, 71–73. doi: 10.1038/248071a0

Pastalkova, E., Serrano, P., Pinkhasova, D., Wallace, E., Fenton, A. A., and Sacktor, T. C. (2006). Storage of spatial information by the maintaince mechanism of LTP. *Science* 313, 1141–1144. doi: 10.1126/science.1128657

Paul, L. A., and Scheibel, A. B. (1986). Structural substrates of epilepsy. *Adv. Neurol.* 44, 775–786.

Perin, R., Berger, T. K., 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

Platschek, S., Cuntz, H., Vuksic, M., Deller, T., and Jedlicka, P. (2016). A general homeostatic principle following lesion induced dendritic remodeling. *Acta Neuropathol. Commun.* 4:19. doi: 10.1186/s40478-016-0285-8

Portera-Cailliau, C., Pan, D. T., and Yuste, R. (2003). Activity-regulated dynamic behavior of early dendritic protrusions: evidence for different types of dendritic filopodia. *J. Neurosci.* 23, 7129–7142. Available online at: http://www.jneurosci.org/content/23/18/7129.long

Ramakers, G. J., Corner, M. A., and Habets, A. M. (1990). Development in the absence of spontaneous bioelectric activity results in increased stereotyped burst firing in cultures of dissociated cerebral cortex. *Exp. Brain Res.* 79, 157–166. doi: 10.1007/BF00228885

Riedel, H., and Schild, D. (1992). The dynamics of Hebbian synapses can be stabilized by a nonlinear decay term. *Neural Netw.* 5, 459–463. doi: 10.1016/0893-6080(92)90007-6

Rioult-Pedotti, M.-S., Friedman, D., Hess, G., and Donoghue, J. P. (1998). Strengthening of horizontal cortical connections following skill learning. *Nat. Neurosci.* 1, 230–234. doi: 10.1038/678

Rocha, M., and Sur, M. (1995). Rapid acquisition of dendritic spines by visual thalamic neurons after blockade of N-methyl-D-aspartate receptors. *Proc. Natl. Acad. Sci. U.S.A.* 92, 8026–8030. doi: 10.1073/pnas.92.17.8026

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

Rogerson, T., Cai, D. J., Frank, A., Sano, Y., Shobe, J., Lopez-Aranda, M. F., et al. (2014). Synaptic tagging during memory allocation. *Nat. Rev. Neurosci.* 15, 157–169. doi: 10.1038/nrn3667

Ruthazer, E. S., Li, J., and Cline, H. T. (2006). Stabilization of axon branch dynamics by synaptic maturation. *J. Neurosci.* 26, 3594–3603. doi: 10.1523/J NEUROSCI.0069-06.2006

Scheibel, M. E., Crandall, P. H., and Scheibel, A. B. (1974). The hippocampal-dentate complex in temporal lobe epilepsy. A golgi study. *Epilepsia* 15, 55–80. doi: 10.1111/j.1528-1157.1974.tb03997.x

Segal, I., Korkotian, I., and Murphy, D. D. (2000). Dendritic spine formation and pruning: common cellular mechanisms? *Trends Neurosci.* 23, 53–57. doi: 10.1016/S0166-2236(99)01499-X

Segal, M. (2005). Dendritic spines and long-term plasticity. *Nat. Rev. Neurosci.* 6, 277–284. doi: 10.1038/nrn1649

Sejnowski, T. J. (1977a). Statistical constraints on synaptic plasticity. *J. Theor. Biol.* 69, 385–389. doi: 10.1016/0022-5193(77)90146-1

Sejnowski, T. J. (1977b). Storing covariance with nonlinearly interacting neurons. *J. Math. Biol.* 4, 303–321. doi: 10.1007/BF00275079

Shi, S.-H., Hayashi, Y., Petralia, R. S., Zaman, S. H., Wenthold, R. J., Svoboda, K., et al. (1999). Rapid spine delivery and redistribution of AMPA receptors after synaptic NMDA receptor activation. *Science* 284, 1811–1816. doi: 10.1126/science.284.5421.1811

Sin, W. C., Haas, K., Ruthazer, E. S., and Cline, H. T. (2002). Dendrite growth increased by visual activity requires NMDA receptor and Rho GTPases. *Nature* 419, 475–480. doi: 10.1038/nature00987

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

Song, S., Sjöström, P. J., Reigl, M., Nelson, S., and Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol.* 3:e68. doi: 10.1371/journal.pbio.0030068

Sorra, K. E., and Harris, K. M. (1998). Stability in synapse number and size at 2 hr after long-term potentiation in hippocampal area CA1. *J. Neurosci.* 18, 658–671.

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

Stettler, D. D., Yamahachi, H., Li, W., Denk, W., and Gilbert, C. D. (2006). Axons and synaptic boutons are highly dynamic in adult visual cortex. *Neuron* 49, 877–887. doi: 10.1016/j.neuron.2006.02.018

Stevens, C. F., and Wang, Y. (1994). Changes in reliability of synaptic function as a mechanism for plasticity. *Nature* 371, 704–707. doi: 10.1038/371704a0

Stoop, R., and Poo, M. M. (1996). Synaptic modulation by neurotrophic factors: differential and synergistic effects of brain-derived neurotrophic factor and ciliary neurotrophic factor. *J. Neurosci.* 16, 3256–3264.

Tailby, C., Wright, L. L., Metha, A. B., and Calford, M. B. (2005). Activity-dependent maintenance and growth of dendrites in adult cortex. *Proc. Natl. Acad. Sci. U.S.A.* 102, 4631–4636. doi: 10.1073/pnas.0402747102

Tetzlaff, C., Dasgupta, S., Kulvicius, T., and Wörgötter, F. (2015). The use of Hebbian cell assemblies for nonlinear computation. *Sci. Rep.* 5:12866. doi: 10.1038/srep12866

Tetzlaff, C., Kolodziejski, 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

Tetzlaff, C., Okujeni, S., Egert, U., Wörgötter, F., and Butz, M. (2010). Self-organized criticality in developing neuronal networks. *PLoS Comput. Biol.* 6:e1001013. doi: 10.1371/journal.pcbi.1001013

Thomas, B. T., Blalock, D. W., and Levy, W. B. (2015). Adaptive synaptogenesis constructs neural codes that benefit discrimination. *PLoS Comput. Biol.* 11:e1004299. doi: 10.1371/journal.pcbi.1004299

Tian, X., Kai, L., Hockberger, P. E., Wokosin, D. L., and Surmeier, D. J. (2010). MEF-2 regulates activity-dependent spine loss in striatopallidal medium spiny neurons. *Mol. Cell. Neurosci.* 44, 94–108. doi: 10.1016/j.mcn.2010.01.012

Toni, N., Buchs, P. A., Nikonenko, I., Bron, C. R., and Muller, D. (1999). LTP promotes formation of multiple spine synapses between a single axon terminal and a dendrite. *Nature* 402, 421–425. doi: 10.1038/46574

Tønnesen, J., Katona, G., Rózsa, B., and Nägerl, U. V. (2014). Spine neck plasticity regulates compartmentalization of synapses. *Nat. Neurosci.* 17, 678–685. doi: 10.1038/nn.3682

Tønnesen, J., Nadrigny, F., Willig, K. I., Wedlich-Söldner, R., and Nägerl, U. V. (2011). Two-color STED microscopy of living synapses using a single laser-beam pair. *Biophys. J.* 101, 2545–2552. doi: 10.1016/j.bpj.2011.10.011

Toyoizumi, 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

Trachtenberg, J. T., Chen, B. E., Knott, G. W., Feng, G., Sanes, J. R., Welker, E., et al. (2002). Long-term *in vivo* imaging of experience-dependent synaptic plasticity in adult cortex. *Nature* 420, 788–794. doi: 10.1038/nature01273

Trommald, M., Hulleberg, G., and Andersen, P. (1996). Long-term potentiation is associated with new excitatory spine synapses on rat dentate granule cells. *Learn. Mem.* 3, 218–228. doi: 10.1101/lm.3.2-3.218

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. (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., 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

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

Vaillant, A. R., Zanassi, P., Walsh, G. S., Aumont, A., Alonso, A., and Miller, F. D. (2002). Signaling mechanisms underlying reversible, activity-dependent dendrite formation. *Neuron* 34, 985–998. doi: 10.1016/S0896-6273(02)00717-1

van Huizen, F., and Romijn, H. J. (1987). Tetrodotoxin enhances initial neurite outgrowth from fetal rat cerebral cortex cells *in vitro*. *Brain Res.* 408, 271–274. doi: 10.1016/0006-8993(87)90386-6

van Huizen, F., Romijn, H. J., and Habets, A. M. (1985). Synaptogenesis in rat cerebral cortex cultures is affected during chronic blockade of spontaneous bioelectric activity by tetrodotoxin. *Brain Res.* 351, 67–80. doi: 10.1016/0165-3806(85)90232-9

van Huizen, F., Romijn, H. J., Habets, A. M., and van den Hooff, P. (1987). Accelerated neural network formation in rat cerebral cortex cultures chronically disinhibited with picrotoxin. *Exp. Neurol.* 97, 280–288. doi: 10.1016/0014-4886(87)90089-6

van Ooyen, A. (ed.). (2003). *Modeling Neural Development*. Cambridge: MIT Press.

van Ooyen, A. (2011). Using theoretical models to analyse neural development. *Nat. Rev. Neurosci.* 12, 311–326. doi: 10.1038/nrn3031

van Ooyen, A., and van Pelt, J. (1994). Activity-dependent outgrowth of neurons and overshoot phenomena in developing neural networks. *J. Theor. Biol.* 167, 27–43. doi: 10.1006/jtbi.1994.1047

van Ooyen, A., and van Pelt, J. (1996). Complex periodic behaviour in a neural network model with activity-dependent neurite outgrowth. *J. Theor. Biol.* 179, 229–242. doi: 10.1006/jtbi.1996.0063

van Ooyen, A., van Pelt, J., and Corner, M. A. (1995). Implications of activity dependent neurite outgrowth for neuronal morphology and network development. *J. Theor. Biol.* 172, 63–82. doi: 10.1006/jtbi.1995.0005

van Pelt, J., van Ooyen, A., and Corner, M. A. (1996). Growth cone dynamics and activity-dependent processes in neuronal network development. *Prog. Brain Res.* 108, 333–346. doi: 10.1016/S0079-6123(08)62550-9

Vees, A. M., Micheva, K. D., Beaulieu, C., and Descarries, L. (1998). Increased number and size of dendritic spines in ipsilateral barrel field cortex following unilateral whisker trimming in postnatal rat. *J. Comp. Neurol.* 400, 110–124.

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

Vlachos, A., Becker, D., Jedlicka, P., Winkels, R., Roeper, J., and Deller, T. (2012a). Entorhinal denervation induces homeostatic synaptic scaling of excitatory postsynapses of dentate granule cells in mouse organotypic slice cultures. *PLoS ONE* 7:e32883. doi: 10.1371/journal.pone.0032883

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

Vlachos, A., Orth, C. B., Schneider, G., and Deller, T. (2012b). Time-lapse imaging of granule cells in mouse entorhino-hippocampal slice cultures reveals changes

in spine stability after entorhinal denervation. *J. Comp. Neurol.* 520, 1891–1902. doi: 10.1002/cne.23017

Vogels, T. P., Frömke, 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

von der Malsburg, C. (1973). Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik* 14, 85–100. doi: 10.1007/BF00288907

Vuksic, M., Turco, D. D., Vlachos, A., Schuldt, G., Müller, C. M., Schneider, G., et al. (2011). Unilateral entorhinal denervation leads to long-lasting dendritic alterations of mouse hippocampal granule cells. *Exp. Neurol.* 230, 176–185. doi: 10.1016/j.expneurol.2011.04.011

Wagenaar, D. A., Pine, J., and Potter, S. M. (2006). An extremely rich repertoire of bursting patterns during the development of cortical cultures. *BMC Neurosci.* 7:11. doi: 10.1186/1471-2202-7-11

Whitlock, J. R., Heynen, A. J., Shuler, M. G., and Bear, M. F. (2006). Learning induces long-term potentiation in the hippocampus. *Science* 313, 1093–1097. doi: 10.1126/science.1128134

Wiegert, J. S., and Oertner, T. G. (2013). Long-term depression triggers the selective elimination of weakly integrated synapses. *Proc. Natl. Acad. Sci. U.S.A.* 110, E4510–E4519. doi: 10.1073/pnas.1315926110

Willshaw, D. J., Buneman, O. P., and Longuet-Higgins, H. C. (1969). Non-holographic associative memory. *Nature* 222, 960–962. doi: 10.1038/222960a0

Wolff, J. R., Joó, F., and Dames, W. (1978). Plasticity in dendrites shown by continuous GABA administration in superior cervical ganglion of adult rat. *Nature* 274, 72–74. doi: 10.1038/274072a0

Wong, R. O. L., and Ghosh, A. (2002). Activity-dependent regulation of dendritic growth and patterning. *Nat. Rev. Neurosci.* 3, 803–812. doi: 10.1038/nrn941

Wu, G. Y., and Cline, H. T. (1998). Stabilization of dendritic arbor structure *in vivo* by CaMKII. *Science* 279, 222–226. doi: 10.1126/science.279.5348.222

Xu, T., Yu, X., Perlik, A. J., Tobin, W. F., Zweig, J. A., Tennant, K., et al. (2009). Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature* 462, 915–919. doi: 10.1038/nature08389

Yamahachi, H., Marik, S. A., McManus, J. N. J., Denk, W., and Gilbert, C. D. (2009). Rapid axonal sprouting and pruning accompany functional reorganization in primary visual cortex. *Neuron* 64, 719–729. doi: 10.1016/j.neuron.2009.11.026

Yang, G., Pan, F., and Gan, W.-B. (2009). Stably maintained dendritic spines are associated with lifelong memories. *Nature* 462, 920–924. doi: 10.1038/nature08577

Yang, Y., Wang, X. B., Frerking, M., and Zhou, Q. (2008). Spine expansion and stabilization associated with long-term potentiation. *J. Neurosci.* 28, 5740–5751. doi: 10.1523/JNEUROSCI.3998-07.2008

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

Yin, J., and Yuan, Q. (2014). Structural homeostasis in the nervous system: a balancing act for wiring plasticity and stability. *Front. Cell. Neurosci.* 8:439. doi: 10.3389/fncel.2014.00439

Yu, L. M. Y., and Goda, Y. (2009). Dendritic signalling and homeostatic adaptation. *Curr. Opin. Neurobiol.* 19, 327–335. doi: 10.1016/j.conb.2009.07.002

Yu, X., Wang, G., Gilmore, A., Yee, A. X., Li, X., Xu, T., et al. (2013). Accelerated experience-dependent pruning of cortical synapses in ephrin-A2 knockout mice. *Neuron* 80, 64–71. doi: 10.1016/j.neuron.2013.07.014

Yuste, R. (2010). *Dendritic Spines*. Camebridge, MA: The MIT Press.

Yuste, R., and Bonhoeffer, T. (2001). Morphological changes in dendritic spines associated with long-term synaptic plasticity. *Annu. Rev. Neurosci.* 24, 1071–1089. doi: 10.1146/annurev.neuro.24.1.1071

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, 886–900. 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

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

Zhou, Q., Homma, K. J., and Poo, M. M. (2004). Shrinkage of dendritic spines associated with long-term depression of hippocampal synapses. *Neuron* 44, 749–757. doi: 10.1016/j.neuron.2004.11.011

Zito, K., Scheuss, V., Knott, G., Hill, T., and Svoboda, K. (2009). Rapid functional maturation of nascent dendritic spines. *Neuron* 61, 247–258. doi: 10.1016/j.neuron.2008.10.054

Ziv, N. E., and Smith, S. J. (1996). Evidence for a role of dendritic filopodia in synaptogenesis and spine formation. *Neuron* 17, 91–102. doi: 10.1016/S0896-6273(00)80283-4

Zuo, Y., Lin, A., Chang, P., and Gan, W.-B. (2005a). Development of long-term dendritic spine stability in diverse regions of cerebral cortex. *Neuron* 46, 181–189. doi: 10.1016/j.neuron.2005.04.001

Zuo, Y., Yang, G., Kwon, E., and Gan, W.-B. (2005b). Long-term sensory deprivation prevents dendritic spine loss in primary somatosensory cortex. *Nature* 436, 261–265. doi: 10.1038/nature03715

REMOD: A Tool for Analyzing and Remodeling the Dendritic Architecture of Neural Cells

Panagiotis Bozelos[1,2], Stefanos S. Stefanou[1,3†], Georgios Bouloukakis[1,4†], Constantinos Melachrinos[1] and Panayiota Poirazi[1]*

[1] Institute of Molecular Biology and Biotechnology (IMBB), Foundation for Research and Technology-Hellas (FORTH), Crete, Greece, [2] Department of Molecular Biology and Genetics, Democritus University of Thrace, Crete, Greece, [3] Department of Biology, University of Crete, Crete, Greece, [4] Computer Science Department, University of Crete, Crete, Greece

Dendritic morphology is a key determinant of how individual neurons acquire a unique signal processing profile. The highly branched dendritic structure that originates from the cell body, explores the surrounding 3D space in a fractal-like manner, until it reaches a certain amount of complexity. Its shape undergoes significant alterations under various physiological or neuropathological conditions. Yet, despite the profound effect that these alterations can have on neuronal function, the causal relationship between the two remains largely elusive. The lack of a systematic approach for remodeling neural cells and their dendritic trees is a key limitation that contributes to this problem. Such causal relationships can be inferred via the use of large-scale neuronal models whereby the anatomical plasticity of neurons is accounted for, in order to enhance their biological relevance and hence their predictive performance. To facilitate this effort, we developed a computational tool named REMOD that allows the structural remodeling of any type of virtual neuron. REMOD is written in Python and can be accessed through a dedicated web interface that guides the user through various options to manipulate selected neuronal morphologies. REMOD can also be used to extract meaningful morphology statistics for one or multiple reconstructions, including features such as sholl analysis, total dendritic length and area, path length to the soma, centrifugal branch order, diameter tapering and more. As such, the tool can be used both for the analysis and/or the remodeling of neuronal morphologies of any type.

Keywords: neuron, dendrite, dendritic remodeling, computational tool, statistical analysis

INTRODUCTION

The morphological complexity of dendrites has been documented since the times of Ramón y Cajal (1911) and is generally considered as an important factor for the proper functioning of the nervous system. Dendritic morphology also demonstrates a great deal of variation across different neuronal cell types (Ascoli, 2006), thus imposing an extra layer of complexity onto the deeply tangled relationship between structure and function (Mainen and Sejnowski, 1996; Krichmar et al., 2002; Schaefer et al., 2003). This structural diversity has been suggested to play a key role in shaping the mode of connectivity between neurons (Sholl, 1956; Kalisman et al., 2003; Chklovskii, 2004) as well as the information processing and signal integration capabilities of neural cells (reviewed in Spruston, 2008). Hence, the morphology of the dendritic tree directly influences

the functionality of neural tissue, both at the single-cell and network levels, in complex ways.

To deliver this structural intrication neurons appear to follow an innate growth program (McAllister, 2000; Libersat and Duch, 2004) while responding adaptively to extrinsic guidance cues supplied by the extracellular matrix (for detailed reviews, see Wong and Wong, 2000; Procko and Shaham, 2010). Moreover, calcium signaling events induced by the electrophysiological activity of neurons are regarded as indispensable part of the morphogenesis process (Wu and Cline, 1998; Wong and Ghosh, 2002; Lohmann et al., 2005). Importantly, dendrites appear to be dynamic structures, even after the post-developmental period (Magariños et al., 1996; Stranahan et al., 2007) and these morphological changes ought to be related to neuronal function.

A considerable amount of research has accumulated to validate this notion, by quantifying the dendritic alterations that happen in response to either physiological or neuropathological stimuli. For instance, in Alzheimer's disease, besides the well-known histopathological hallmarks of extracellular amyloid plaques and intracellular neurofibrillary tangles (NFTs), dendritic atrophy is consistently found among hippocampal and cortical pyramidal neurons (Yamada et al., 1988; Moolman et al., 2004). In another study, when rats were reared under chronic stress conditions, abnormal morphological changes to dendritic trees were seen across at least three different brain regions: the hippocampus, amygdala and prefrontal cortex (Vyas et al., 2002; Shansky and Morrison, 2009). Interestingly, neurons in these areas seemed to respond completely different to the same type of stress. In the CA3 subregion of the hippocampus, as well as in the medial prefrontal cortex, the dendritic arbor retracted to itself. On the contrary, dendritic trees in the basolateral amygdala (BLA) exhibited excessive growth as evident by their increased total length and number of branches.

Dendritic remodeling is also happening under physiological conditions. For instance, when there is increased need for receiving, integrating and encoding complex spatiotemporal patterns about newly encountered environments, the dendritic arbors of rat hippocampal neurons grow (Tronel et al., 2010); perhaps to accommodate new synapses and/or counterbalance the impending overexcitability, that otherwise would be toxic to the cell. Despite the profound effect that such dendritic alterations can have on neural processing, a causal relationship between structural changes and neuronal output has not been established yet.

A great deal of neuroanatomical research effort has been devoted to providing a direct link between structure and function, at both the single-cell and network levels. During the past few decades, the advent of intracellular labeling techniques and the application of various visualization methods have led to a dramatic increase in high resolution dendritic morphology data. A tremendous progress in imaging methods and automation is also expected to pave the way for an exponential growth of the data acquisition in the forthcoming years (Peng et al., 2015). The acquired morphological reconstructions are stored in dedicated databases such as the NeuroMorpho (Ascoli et al., 2007), the Fly Circuit (Chiang et al., 2011) and the Cell-Centered

Database (Martone et al., 2003), and enable the quantitative analysis of neuronal shapes by the use of parameters relevant to the metrical and topological properties of the cell. Investigators interested in deciphering the role of neuroanatomy to the proper functioning of the nervous system have developed and provided the community with commercial and freeware tools such as Neurolucida Explorer (Glaser and Glaser, 1990), L-Measure (Scorcioni et al., 2008) BTMORPH (Torben-Nielsen, 2014), and NLMorphologyViewer[1] to trace, analyze and visualize the shape of the reconstructed neurons. However, most of these tools need to be locally installed and often require at least some basic programming skills, making them difficult to use by non-experts (for a detailed review, see Parekh and Ascoli, 2013).

More advanced approaches to the remodeling of dendritic trees incorporate the de-novo generation of neuronal morphologies, based either on intrinsic correlations between morphometric parameters (e.g., the NETMORPH tool by Koene et al., 2009), on principles of neuronal material conservation vs. conduction times (e.g., the TREES toolbox by Cuntz et al., 2011), or the dynamic and competitive interplay of retraction and outgrowth processes (e.g., the CX3D tool by Zubler et al., 2013; Hjorth et al., 2014). Others implement a hybrid of local and global generative algorithms to efficiently reproduce the anatomical variability of various neuronal classes (e.g., the L-Neuron and ArborVitae tools by Ascoli et al., 2001). Recently, the NeuroMac tool introduced an interesting context-aware approach in which developing neurons interact with the surrounding brain substrate (Torben-Nielsen and De Schutter, 2014). On the core of all the above algorithms lays the use of a limited set of statistical descriptors and assumptions on growth rules to generate stochastic neuronal structures that are statistically indistinguishable from the real neurons of the same morphological class. There is little evidence as to whether these assumptions and rules would hold under different physiological and/or pathological conditions. While the importance of these tools is widely recognized by the neuroanatomical community, to our knowledge, there is currently no tool available to implement targeted, assumption-free alterations on specific dendrites or branches of already-grown morphologies. Given that dendritic remodeling is happening quite frequently in nature, as part of the ever-adapting neuronal function, the necessity of a tool to remodel the dendritic morphology of digital reconstructions is long overdue. Moreover, the recent use of detailed biophysical and morphologically realistic large-scale brain models to investigate microcircuit structure and function (Markram et al., 2015) highlights the need to understand how different anatomical features and their plasticity shape brain function and dysfunction. Towards this goal, a systematic methodology that allows efficient remodeling of any type of neuronal morphology is a prerequisite.

In this work, we introduce REMOD, a computational tool that allows the structural remodeling of any type of digitally reconstructed neuron. The algorithm focuses on the simulation

---

[1]http://neuronland.org/NLMorphologyViewer/NLMorphologyViewer.html

of the end-result of a dendritic remodeling process, without explicitly implementing any growth and/or retraction rules. It should be noted that such an end result can be achieved by several different sets of manipulations, some of which may not be biologically relevant. The important advantage of REMOD is that it provides the flexibility to choose which manipulations are considered relevant for each experimenter, condition and cell type, thus allowing a full exploration of the parameter space. This flexibility is particularly important when investigating the effects of specific morphological manipulations on a given response pattern. Such manipulations can be used to tease out the contribution of distinct morphological factors from other processes such as biophysical mechanisms.

The tool is written in Python and can be accessed through a dedicated web interface that guides the user through various options to manipulate selected neuronal morphologies. More explicitly, it provides the ability to upload one or more morphology files (in SWC, the most widely used format) and choose specific dendrites or dendritic regions on which to operate one of the following actions: shrink, remove, extend, branch or scale. The user retains complete control over the extent of each alteration and if a chosen action is not possible due to pre-existing structural constraints, appropriate warnings are generated. It is worth mentioning that REMOD can also be used to extract a plethora of descriptive statistics for one or multiple neurons, such as sholl analysis, total dendritic length and area, path length to the soma, centrifugal branch order, diameter tapering and more. As such, the tool can be used both for the analysis and/or the remodeling of neuronal morphologies of any type.

METHODS

Recent advances in high-throughput single-neuron imaging techniques are expected to stimulate a morphological data "explosion" that will revolutionize the computational neuroanatomy field. Importantly, the aforementioned data "explosion" needs to be accompanied by the development of software tools, designed to meet the future requirements of parsing and manipulating tons of 3D neuron reconstruction data in a transparent, reliable and highly-consistent manner (Peng et al., 2015). In order to achieve this longer-term goal, an appropriate approach needs to be adopted by the community to properly address this "big data" challenge. Therefore, it is imperative for the new generation of software packages to be optimized for this data volume, to handle it in a user friendly way and with the added benefit to be open-source and publicly available, allowing ease of use, sharing and active development by avid researchers in the field.

The proposed tool is developed according to this philosophy. The software implementation is in Python 2.7 and the tool is open source and available for use as a web-application at www.remod.gr/. The source code and documentation of the tool is available on the Github repository hosting service at github.com/bozelosp/remod, where it can be downloaded as a standalone application that retains almost all of its online capabilities and/or further developed.

The software package contains flexible and extendable modular blocks of code that can be generally classified in three categories: (i) parsing the morphology reconstruction data encoded in standard SWC files; (ii) performing morphometric and statistical analysis of the provided dendritic trees; and (iii) executing specific remodeling actions and exporting the new morphologies in the SWC format. The following paragraphs explain the user-action workflow between the abovementioned categories.

First, the user uploads one or more morphology reconstructions, properly formatted as SWC files. REMOD parses the geometric information of the specific 3D structure and topology and provides a rotating 3D visualization of the corresponding tree, as well as many depictions of the extracted morphometric statistics in the form of tables and charts. Next, the interface guides the user through a range of remodeling options that can be implemented on the selected morphologies as shown in **Figure 1**.

For example, the researcher can manually select specific dendrites to impose a morphological alteration or more conveniently select a whole region of the neuron (i.e., the basal or apical tree), or even a randomly selected portion of it. Remodeling actions can be implemented on any and/or multiple selected dendrites, irrespectively of whether they are terminal branches or not.

Next, the researcher decides the remodeling action to apply. If the plan is to further grow the dendritic tree, then the appropriate options to choose might be additive extension or branching, i.e., adding two new dendrites stemming from an existing parental one. The researcher can specify the extent of growth in terms of a percentage of the previous dendritic length or by defining the desirable length in micrometers. In this case, the algorithm will generate a series of somewhat randomly directed cylinders that radiate away from the parental dendritic segment and the soma as shown in **Figure 2**. The added segments have random lengths sampled from a realistic distribution and mimic the way dendrites explore the available 3D space in a quasi-random, quasi-directed way.

Similarly, to simulate a dendritic retraction, the user decides between two options: shrinking or complete removal of the selected dendrites. In both actions, cylinders are removed from the SWC file to match the desirable reduction in length, without altering any other parts of the morphology file. An example of each case is shown in **Figure 3**.

The scaling action can be used to enlarge or reduce the dimensions of either the entire dendritic tree, or the basal/apical regions of the tree, independently. Remodeled morphologies are exported in the SWC standard format and can be downloaded and/or sent to the user via email as shown in **Figure 4**.

REMOD is also able to extract useful statistics for the morphological and topological features of the neuronal reconstructions. Metrics such as total dendritic length, path length, surface and volume are available for both the basal and the apical regions of the tree. Sholl analysis is implemented either for the number of branch points, the number of intersections or the sum of dendritic lengths included in defined radius steps from the soma. The tool also supports
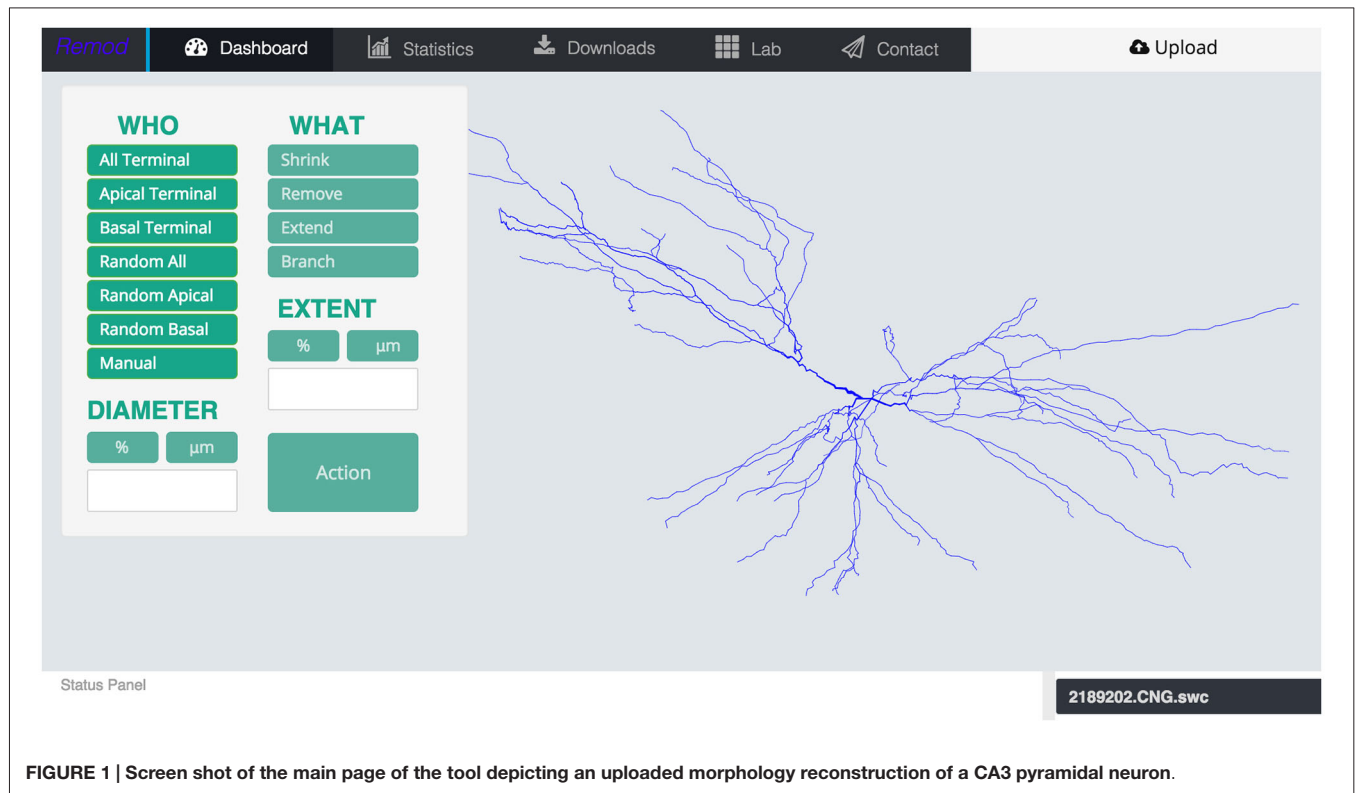
FIGURE 1 | Screen shot of the main page of the tool depicting an uploaded morphology reconstruction of a CA3 pyramidal neuron.

FIGURE 2 | Examples of the extending and branching REMOD actions on a PFC pyramidal neuron. The induced morphological changes (added segments) are shown in green. **(A)** All apical terminal dendrites are extended by 50% of their initial length. **(B)** New branches are added to the tip of the stem of all basal terminal dendrites. The produced length of each one of the branches is 80% of the parent length.

the comparison between two groups of morphologically different neurons using the abovementioned metrics, as shown in **Figure 5**. This utility can be especially useful for both experimentalists and modelers wishing to identify anatomical differences between two particular groups of neurons.

RESULTS

To demonstrate the capabilities of REMOD we implemented morphological changes that led to dendritic atrophy and enhanced dendritic arborization, in hippocampal and amygdaloid pyramidal neurons, respectively. Such changes

FIGURE 3 | Examples of the shrinking and removing REMOD actions on a PFC pyramidal neuron. The induced morphological changes (removed segments) are shown in red. **(A)** All apical terminal dendrites are shrunk by 60% of their initial dendritic length. **(B)** All basal terminal dendrites are completely removed.
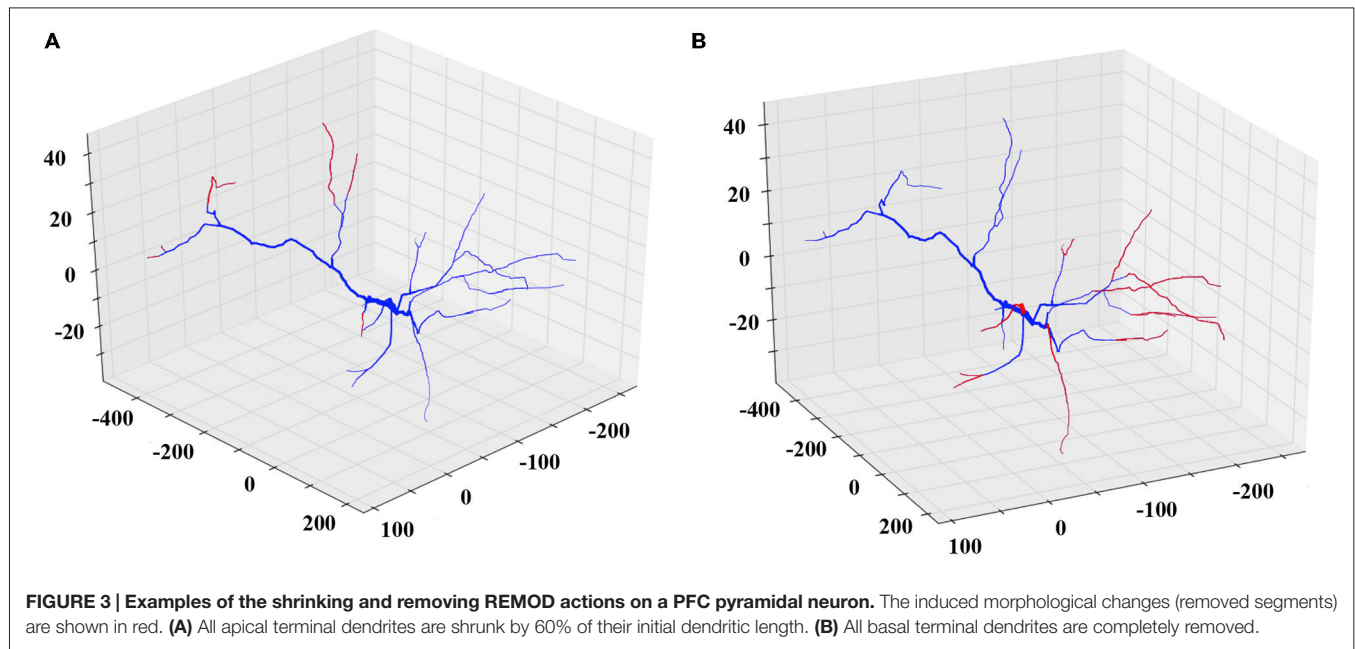
were reported to occur as a result of stress by Vyas et al. (2002). The modifications performed aimed at reproducing the percentage of change in various dendritic features as reported experimentally and were applied to a large number of morphologies downloaded from the NeuroMorpho database.

According to Vyas et al. (2002) chronic immobilization stress (CIS) led to a significant decrease of the total dendritic length and the number of branch points in hippocampal CA3 pyramidal neurons as compared with neurons of the same class in control animals. The marked shrinkage was evident in both apical as well as basal dendrites of the examined pyramidal cells. Specifically, the total dendritic length of apical dendrites was decreased by 29%, and the number of apical branch points was reduced by 31%. The same analysis applied to basal dendrites found their total dendritic length shrunk by 16% and their number of branch points reduced by 16%, respectively. Based on the observation that the number of branch points is affected, it's easy to deduce that a simple scaling of the entire dendritic tree length would not reproduce the observed phenotype. With the goal of simulating the overall shrinkage effect and reproducing the reported percentages, we took the following approach:

- Twenty five reconstructed neuronal morphologies of the hippocampal CA3 pyramidal class were obtained from the NeuroMorpho repository.
- Subsequently, the parameters that were used to remodel the morphologies were selected in such a way as to simulate the experimentally observed results. The total dendritic length of the apical tree was reduced by randomly selecting 17% of the terminal dendrites for complete removal and thereafter randomly choosing an additional 10% of the remaining terminal dendrites to shrink by 18%.

- Accordingly, the basal tree was pruned by randomly selecting 15% of its terminal dendrites for complete removal and shrinking another 5% of them by 5% of their initial length. The results of this processing are shown in **Tables 1, 2** and **Figure 6**.

It should be noted that other combinations of dendrite removal/pruning may also reproduce the reported averaged changes in dendritic length. Due to the lack of information as to which specific morphological features were altered in the experiments, we used a randomized approach.

Interestingly, the same CIS paradigm that caused dendritic atrophy in CA3 pyramidal neurons induced the opposite effect in BLA pyramidal neurons. The total dendritic length of BLA pyramidal neurons under conditions of CIS increased by 25% and the number of branch points by 15% compared to the control cells. With the goal of reproducing the experimentally obtained results of dendritic extension in the amygdala we implemented the following series of actions in REMOD:

- Twenty six reconstructed neuronal morphologies of the BLA pyramidal class of neurons were obtained from the NeuroMorpho repository.
- Subsequently, 10 percent of the terminal dendrites were selected for branching at 80 percent of their length and thereafter a 20% of all the available terminal dendrites were selected for extension by 70% of their initial length. The results are shown in **Tables 3, 4**.
- To further our analysis, we performed segmental processing to track the changes in dendritic length and number of branch points as functions of radial distance from the cell soma. The results are shown in **Figure 7**.

FIGURE 4 | Screen shot illustrating a list of downloadable files that contain various statistics calculated by REMOD for the uploaded CA3 pyramidal neuron depicted in Figure 1.

Overall, despite the fact that the reconstructed morphologies analyzed by Vyas et al. (2002) were significantly smaller than the neurons of the same class downloaded from the NeuroMorpho database, we reliably reproduced the percentage morphometric changes as well as the overall distributions of changes between the control and CIS-treated cells in both neuronal classes.

DISCUSSION

Over the past decade, computational simulations of the electrophysiological behavior of neurons have become fairly common as neuroscientists strive to develop a comprehensive understanding of the nervous system (Bower, 2013; Dudai and Kathinka, 2014). At the same time, the rapid
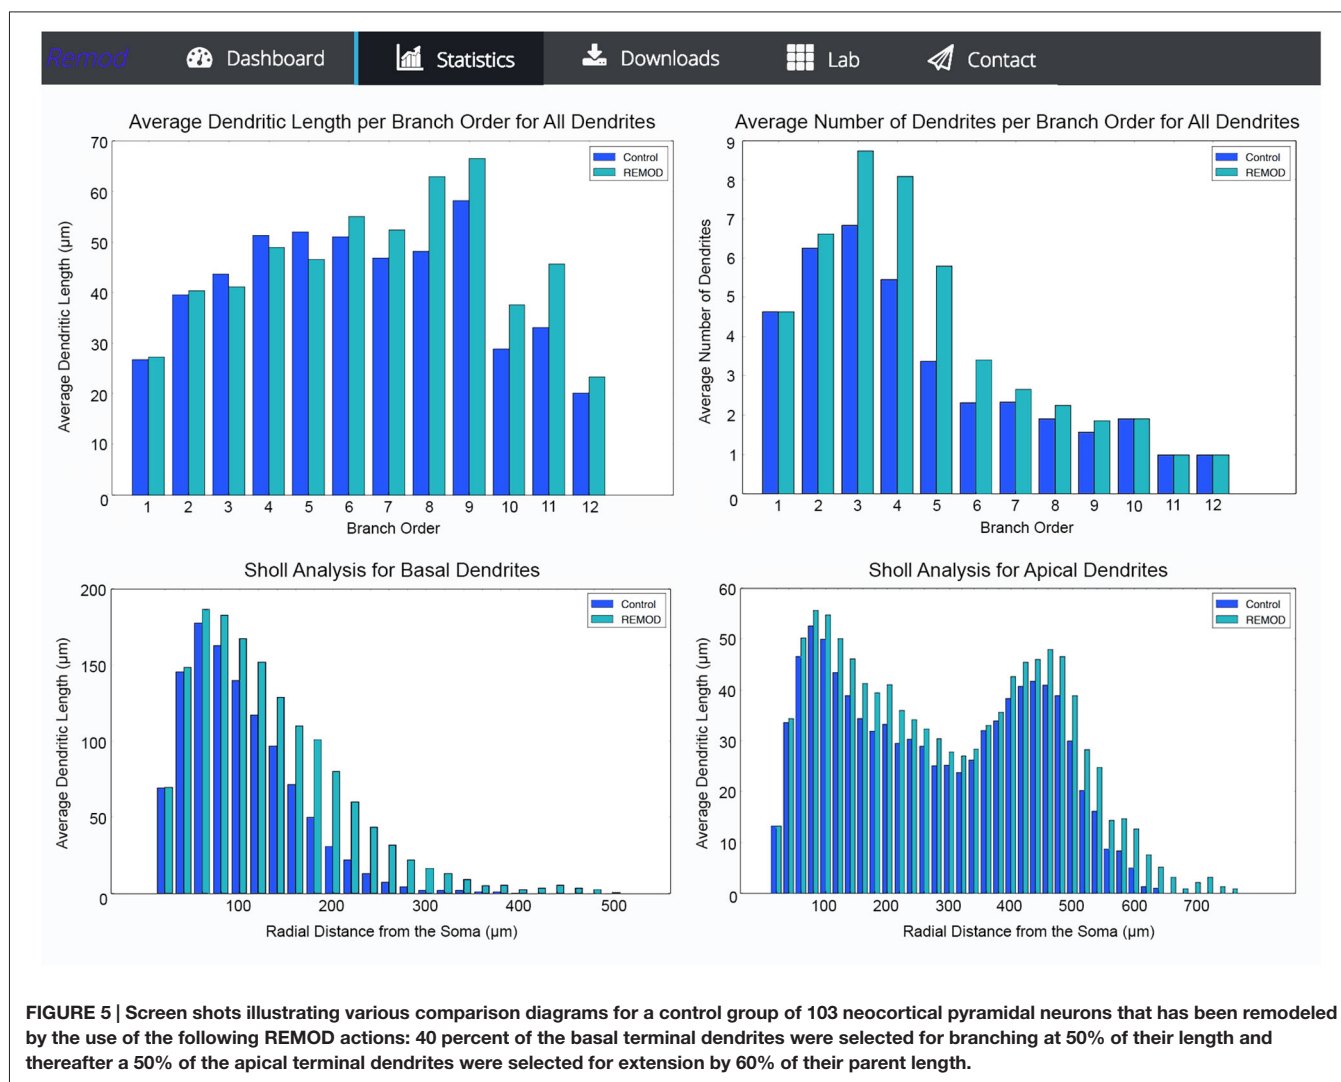
FIGURE 5 | Screen shots illustrating various comparison diagrams for a control group of 103 neocortical pyramidal neurons that has been remodeled by the use of the following REMOD actions: 40 percent of the basal terminal dendrites were selected for branching at 50% of their length and thereafter a 50% of the apical terminal dendrites were selected for extension by 60% of their parent length.

TABLE 1 | Difference in total dendritic length of CA3 pyramidal neurons as reported experimentally after CIS and via remodeling with REMOD.

Total dendritic length	Control (experiment)	CIS	% Change	Control (NeuroMorpho)	Remodeled	% Change
Apical	2113	1498	−29	5471	3922	−28.31
Basal	1827	1527	−16	4302	3623	−15.78

TABLE 2 | Difference in the number of branch points of CA3 pyramidal neurons as reported experimentally after CIS and via remodeling with REMOD.

Number of branch points	Control (experiment)	CIS	% Change	Control (NeuroMorpho)	Remodeled	% Change
Apical	14.7	10.2	−31	31.14	21.18	−31.98
Basal	16.2	13.6	−16	29.22	24.62	−15.74

advancement of imaging techniques has greatly enhanced the acquisition of high-resolution 3D neuronal reconstructions, resulting in anatomically precise and biophysically detailed compartmental models that render realistic simulation of neuronal behavior possible under virtual experimental conditions. Thus, it is expected that multi-compartmental cable models with ever-increasing accuracy in detail will become critically important to better understand the

rich repertoire of computational operations performed by dendrites.

In addition, substantial progress has been achieved by experimental neuroscientists into revealing how neural plasticity might relate to development, learning and disease and how its influence could be projected far beyond the dynamics of synapses (Butz et al., 2009; Yau et al., 2011). Along with motile spines and axonal branches, the dendritic architecture of neural cells is also subject to significant morphological readjustments (for a detailed review, see Emoto, 2011). Thus, many challenging questions arise about the structural plasticity

of the nervous system, including the rules that govern the dendritic morphological configuration, the adaptability of the neural connectivity and the preservation of functional continuity. Hence, a software tool that is able to induce structural alterations upon virtual neurons and statistically quantify the changes between acquired and remodeled morphologies could prove an invaluable resource for modelers and neuroanatomists alike, complementary to the existing approaches.

In this paper, we have presented REMOD, a novel open-access software tool for the analysis and remodeling of dendritic

FIGURE 6 | Comparison of the experimentally induced (CIS) shrinkage in dendritic structure of the CA3 cells (control cells, *n* = 50; CIS cells, *n* = 45) with the same remodeling effect induced by REMOD on 25 CA3 neuronal reconstructions downloaded from NeuroMorpho database. (A) Side by side comparison of the diagram showing the dendritic length as a function of radial distance from the soma as taken from Vyas et al. (2002, left), with the corresponding diagram of the remodeling example using REMOD (right). Changes in apical and basal dendrites are shown separately. **(B)** Comparison of the diagram showing the number of branch points as a function of radial distance from the soma as taken from the paper published by Vyas et al. (2002, left), with the corresponding diagram using REMOD.

TABLE 3 | Difference in the total dendritic length of BLA pyramidal neurons as reported experimentally after CIS and via remodeling with REMOD.

Total dendritic length	Control (experiment)	CIS	% Change	Control (NeuroMorpho)	Remodeled	% Change
All	1330	1666	25	4040	5067	25.42

TABLE 4 | Difference in the number of branch points of BLA pyramidal neurons as reported experimentally after CIS and via remodeling with REMOD.

Number of branch points	Control (experiment)	CIS	% Change	Control (NeuroMorpho)	Remodeled	% Change
All	13.5	15.0	11	25.98	28.825	10.95

morphologies. The tool is written in the non-commercial Python programming language and is publicly available on the GitHub platform, under the open-source MIT license to encourage active development by the neuroanatomical community. The tool is accompanied by an easily accessible front-end interface that anyone can use via a web browser (http://www.remod.gr), thus eliminating the need to download and configure additional packages and/or restricted libraries locally on the desktop.

REMOD is designed to fulfill the emerging need among computational neuroscientists to simulate the dynamic nature of the dendritic structure. Other toolboxes have been developed for the special purposes of de novo generating synthetic neurons or neuronal networks, or even configuring the connectivity between them (Kalisman et al., 2003; Koene et al., 2009). These methods are exploiting experimentally-substantiated principles, such as the conservation of neuronal material vs. conduction time constraints, the dynamic negotiation between retraction and outgrowth processes and the context-awareness of neuronal arborization, in order to enforce local or global generative algorithms that efficiently replicate the

extent of morphological variability manifested by various neuronal types (Ascoli et al., 2001; Cuntz et al., 2011; Zubler et al., 2013; Hjorth et al., 2014; Torben-Nielsen and De Schutter, 2014). Still, the modeling of structural plasticity in the brain is arguably hindered by a multitude of challenges. These challenges predominantly include the complex interaction between numerous interdependent factors, such as genetic, metabolic and behavioral influences, that cannot be easily dissected for integration into the current predictive models (Bestman Da Silva and Cline, 2008; Polleux and Ghosh, 2008).

To overcome this limitation, the researcher can approach morphology from a different methodological perspective, that of remodeling an already-existing neuronal structure ad arbitrium. Principally, this approach provides a wider flexibility to choose the structural manipulations that are considered more relevant for each experimenter, physiological condition and cell type, albeit allowing for a greater exploration of the available parameter space. This type of freedom could prove particularly important when investigating the effects
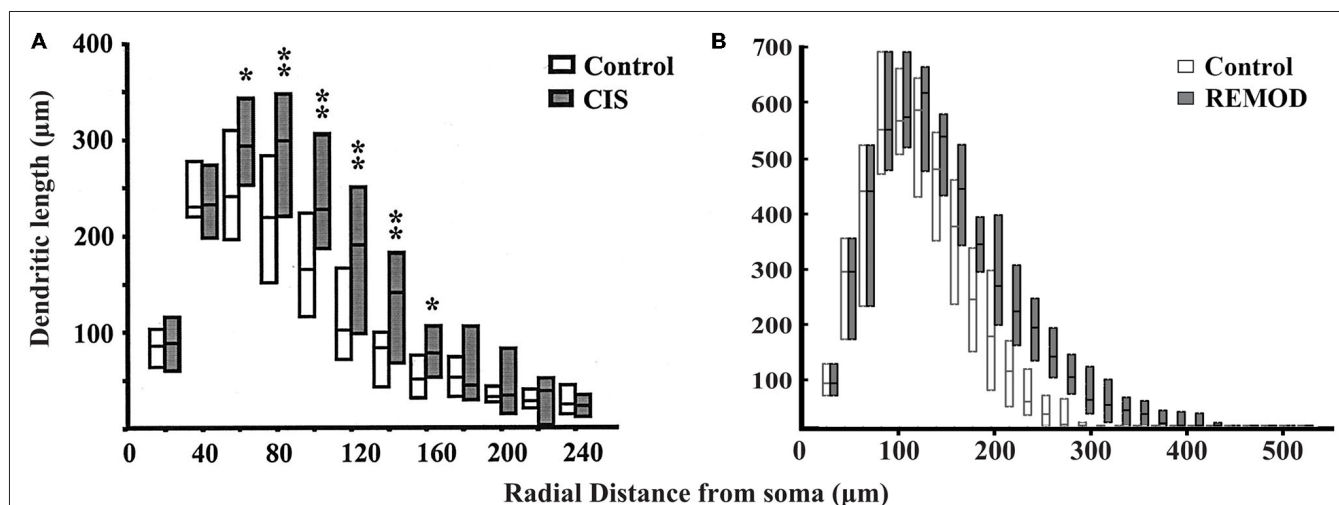
FIGURE 7 | Comparison of the experimentally induced (CIS) increase in dendritic arborization of the basolateral amygdala (BLA) cells (control cells, *n* = 18; CIS cells, *n* = 22) with the same remodeling effect induced by REMOD on 26 BLA neuronal reconstructions downloaded from NeuroMorpho database. Side by side comparison of the diagram showing the dendritic length as a function of radial distance from the soma as taken from the paper published by Vyas et al. (2002) **(A)** with the corresponding diagram of the remodeling example using REMOD **(B)**.

of certain morphological alterations on a given neuronal response pattern, even if they are unlikely to occur. Such manipulations can for instance be used to assess and tease out the individual contribution of distinct morphological factors from other processes, like the electrophysiological properties of ion channels. Furthermore, REMOD offers the ability to perform statistical morphometric analyses of the user-provided morphologies via uploading to our server, an additional utility that may also be of great service to experimentalists.

In this work, we sought to demonstrate the capabilities of our toolbox by reproducing dendritic remodeling in the hippocampus and the amygdala brain regions of the rat under stress conditions, as reported by Vyas et al. (2002). Using REMOD, we faithfully reproduced the percentage of morphological changes reported in this paper, by imposing dendritic remodeling to 3D reconstructions of the same neuronal types downloaded from the NeuroMorpho repository. Unfortunately we did not have access to the original morphologies in order to perform a direct comparison.

The reproduction of the experimental results (with respect to the percentage of changes observed), delivered using REMOD, reflects the ability of the tool to simulate complicated neuronal phenomena that may occur under physiological or pathological conditions. We believe that the functions provided within the first release of the tool are flexible and efficient enough

to simulate any type of dendritic remodeling, capturing the variant expression of neuronal adaptability that has been extensively documented, but still insufficiently explained in current theoretical models.

To the best of our knowledge, REMOD is the first software package that allows the remodeling of existing, already-grown, detailed neuronal morphologies, in parallel with the effortless extraction of morphological descriptive statistics. Thus, REMOD allows the implementation of a systematic approach for altering virtual neuronal morphologies, which is likely to promote further research into understanding the hidden associations between critical neuroanatomical characteristics and the distinct electrophysiological patterns that individual neurons, as well as neural networks, exhibit. Exploiting the benefits and the capabilities of dendritic remodeling will aid the transition from investigating a "rigid" neuronal function to a refined exploration of the intricate effect of morphology to dendritic function and neuronal processing.

ACKNOWLEDGMENTS

REFERENCES

Ascoli, G. A., Krichmar, J. L., Scorcioni, R., Nasuto, S. J., and Senft, S. L. (2001). Computer generation and quantitative morphometric analysis of virtual neurons. *Anat. Embryol.* 204, 283–301. doi: 10.1007/s004290100201

Ascoli, G. A. (2006). Mobilizing the base of neuroscience data: the case of neuronal morphologies. *Nat. Rev. Neurosci.* 7, 318–324. doi: 10.1038/nrn1885

Ascoli, G. A., Donohue, D. E., and Halavi, M. (2007). NeuroMorpho.Org: a central resource for neuronal morphologies. *J. Neurosci.* 27, 9247–9251. doi: 10.1523/jneurosci.2055-07.2007

Bestman Da Silva, J. S. and Cline, H. T (2008). "Dendrite development," in *Dendrites*, eds G. Stuart, N. Spruston, and M. Hausser (New York, NY: Oxford University Press), 69–94.

Bower, J. M. (2013) *20 Years Of Computational Neuroscience.* New York, NY: Springer.

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

Chiang, A. S., Lin, C. Y., Chuang, C. C., Chang, H. M., Hsieh, C. H., Yeh, C. W., et al. (2011). Three-dimensional reconstruction of brain-wide wiring networks in drosophila at single-cell resolution. *Curr. Biol.* 21, 1–11. doi: 10.1016/j.cub.2010.11.056

Chklovskii, D. B. (2004). Synaptic connectivity and neuronal morphology: two sides of the same coin. *Neuron* 43, 609–617. doi: 10.1016/s0896-6273(04)00498-2

Cuntz, H., Forstner, F., Borst, A., and Häusser, M. (2011). The TREES toolbox–probing the basis of axonal and dendritic branching. *Neuroinformatics* 9, 91–96. doi: 10.1007/s12021-010-9093-7

Dudai, Y., and Kathinka, E. (2014). To simulate or not to simulate: what are the questions? *Neuron* 84, 254–261. doi: 10.1016/j.neuron.2014.09.031

Emoto, K. (2011). Dendrite remodeling in development and disease. *Dev. Growth Differ.* 53, 277–286. doi: 10.1111/j.1440-169x.2010.01242.x

Glaser, J. R., and Glaser, E. M. (1990). Neuron imaging with Neurolucida-a PC-based system for image combining microscopy. *Comput. Med. Imaging Graph.* 14, 307–317. doi: 10.1016/0895-6111(90)90105-k

Hjorth, J. J. J., van Pelt, J., Mansvelder, H. D., and van Ooyen, A. (2014). Competitive Dynamics during Resource-Driven Neurite Outgrowth. *PLoS One* 9:e86741. doi: 10.1371/journal.pone.0086741

Kalisman, N., Silberberg, G., and Markram, H. (2003). Deriving physical connectivity from neuronal morphology. *Biol. Cybern.* 88, 210–218. doi: 10.1007/s00422-002-0377-3

Koene, R. A., Tijms, B., van Hees, P., Postma, F., de Ridder, A., and Ramakers, G. J. A.et al. (2009). NETMORPH: a framework for the stochastic generation of large scale neuronal networks with realistic neuron morphologies. *Neuroinformatics* 7, 195–210. doi: 10.1007/s12021-009-9052-3

Krichmar, J. L., Nasuto, S. J., Scorcioni, R., Washington, S. D., and Ascoli, G. A. (2002). Effects of dendritic morphology on CA3 pyramidal cell electrophysiology: a simulation study. *Brain Res.* 941, 11–28. doi: 10.1016/s0006-8993(02)02488-5

Libersat, F., and Duch, C. (2004). Mechanisms of dendritic maturation. *Mol. Neurobiol.* 29, 303–320. doi: 10.1385/mn:29:3:303

Lohmann, C., Finski, A., and Bonhoeffer, T. (2005). Local calcium transients regulate the spontaneous motility of dendritic filopodia. *Nat. Neurosci.* 8, 305–312. doi: 10.1038/nn1406

Magariños, M., McEwen, B. S., Flügge, G., and Fuchs, E. (1996). Chronic psychosocial stress causes apical dendritic atrophy of hippocampal CA3 pyramidal neurons in subordinate tree shrews. *J. Neurosci.* 16, 3534–3540.

Mainen, Z. F., and Sejnowski, T. J. (1996). Influence of dendritic structure on firing pattern in model neocortical neurons. *Nature* 382, 363–366. doi: 10.1038/382363a0

Markram, H., Muller, E., Ramaswamy, S., Reimann, M. W., Abdellah, M., Sanchez, C. A., et al. (2015). Reconstruction and Simulation of Neocortical Microcircuitry. *Cell* 163, 456–492. doi: 10.1016/j.cell.2015.09.029

Martone, M. E., Zhang, S., Gupta, A., Qian, X., He, H., Price, D. L., et al. (2003). The cell-centered database: a database for multiscale structural and protein localization data from light and electron microscopy. *Neuroinformatics* 1, 379–395. doi: 10.1385/ni:1:4:379

McAllister, A. K. (2000). Cellular and molecular mechanisms of dendrite growth. *Cereb. Cortex* 10, 963–973. doi: 10.1093/cercor/10.10.963

Moolman, D. L., Vitolo, O. V., Vonsattel, J. P. G., and Shelanski, M. L. (2004). Dendrite and dendritic spine alterations in Alzheimer models. *J. Neurocytol.* 33, 377–387. doi: 10.1023/b:neur.0000044197.83514.64

Parekh, R., and Ascoli, G. (2013). Neuronal morphology goes digital: a research hub for cellular and system neuroscience. *Neuron* 77, 1017–1038. doi: 10.1016/j.neuron.2013.03.008

Peng, H., Hawrylycz, M., Roskams, J., Hill, S., Spruston, N., Meijering, E., et al. (2015). BigNeuron: Large-Scale 3D Neuron Reconstruction from Optical Microscopy Images. *Neuron* 87, 252–256. doi: 10.1016/j.neuron.2015.06.036

Polleux, F., and Ghosh, A. (2008). "Molecular determinants of dendrite and spine development," in *Dendrites*, eds G. Stuart, N. Spruston, and M. Hausser (New York, NY: Oxford University Press), 95–115.

Procko, C., and Shaham, S. (2010). Assisted morphogenesis: glial control of dendrite shapes. *Curr. Opin. Cell Biol.* 22,560–565. doi: 10.1016/j.ceb.2010.07.005

Ramón y Cajal, S. (1911). Histologie du Système Nerveux de l'Homme et des Vertébrés. Paris: Maloine.

Schaefer, A. T., Larkum, M. E., Sakmann, B., and Roth, A. (2003). Coincidence detection in pyramidal neurons is tuned by their dendritic branching pattern. *J. Neurophysiol.* 89, 3143–3154. doi: 10.1152/jn.00046.2003

Scorcioni, R., Polavaram, S., and Ascoli, G. A. (2008). L-Measure: a web- accessible tool for the analysis, comparison and search of digital reconstructions of neuronal morphologies. *Nat. Protoc.* 3, 866–876. doi: 10.1038/nprot.2008.51

Shansky, R. M., and Morrison, J. H. (2009). Stress-induced dendritic remodeling in the medial prefrontal cortex: effects of circuit, hormones and rest. *Brain Res.* 1293, 108–113. doi: 10.1016/j.brainres.2009.03.062

Sholl, D. A. (1956). The measurable parameters of the cerebral cortex and their significance in its organization. *Prog. Neurobiol.* 2, 324–333.

Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. *Nat. Rev. Neurosci.* 4, 206–221. doi: 10.1038/nrn2286

Stranahan, A. M., Khalil, D., and Gould, E. (2007). Running induces widespread structural alterations in the hippocampus and entorhinal cortex. *Hippocampus* 17, 1017–1022. doi: 10.1002/hipo.20348

Torben-Nielsen, B. (2014). An efficient and extendable python library to analyze neuronal morphologies. *Neuroinformatics* 12, 619–622. doi: 10.1007/s12021-014-9232-7

Torben-Nielsen, B., and De Schutter, E. (2014). Context-aware modeling of neuronal morphologies. *Front. Neuroanat.* 8:92. doi: 10.3389/fnana.2014.00092

Tronel, S., Fabre, A., Charrier, V., Oliet, S. H. R., Gage, F. H., and Abrous, D. N. (2010). Spatial learning sculpts the dendritic arbor of adult-born hippocampal neurons. *Proc. Natl. Acad. Sci. U S A* 107, 7963–7968. doi: 10.1073/pnas.0914613107

Vyas, A., Mitra, R., Rao, B. S. S., and Chattarji, S. (2002). Chronic stress induces contrasting patterns of dendritic remodeling in hippocampal and amygdaloid. *J. Neurosci.* 22, 6810–6818.

Wong, R. O. L., and Ghosh, A. (2002). Activity-dependent regulation of dendritic growth and patterning. *Nat. Rev. Neurosci.* 3, 803–812. doi: 10.1038/nrn941

Wong, W. T., and Wong, R. O. (2000). Rapid dendritic movements during synapse formation and rearrangement. *Curr. Opin. Neurobiol.* 10, 118–124. doi: 10.1016/s0959-4388(99)00059-8

Wu, G. Y., and Cline, H. T. (1998). Stabilization of dendritic arbor structure *in vivo* by CaMKII. *Science* 279, 222–226. doi: 10.1126/science.279.5348.222

Yamada, M., Wada, Y., Tsukagoshi, H., Otomo, E., and Hayakawa, M. (1988). A quantitative Golgi study of basal dendrites of hippocampal CA1 pyramidal cells in senile dementia of Alzheimer type. *J. Neurol. Neurosurg. Psychiatr.* 51, 1088–1090. doi: 10.1136/jnnp.51.8.1088

Yau, S.-Y., Lau, B. W. M., Tong, J. B., Wong, R., Ching, Y.-P., Qiu, G., et al. (2011). Hippocampal Neurogenesis and Dendritic Plasticity Support Running-Improved Spatial Learning and Depression-Like Behaviour in Stressed Rats. *PLoS One* 6:e24263. doi: 10.1371/journal.pone.0024263

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

frontiers
in Neuroanatomy

Automatic Generation of Connectivity for Large-Scale Neuronal Network Models through Structural Plasticity

Sandra Diaz-Pier[1]*, Mikaël Naveau[1,2], Markus Butz-Ostendorf[1] and Abigail Morrison[1,3,4]

[1] Simulation Laboratory Neuroscience – Bernstein Facility for Simulation and Database Technology, Institute for Advanced Simulation, Jülich Aachen Research Alliance, Jülich Research Center, Jülich, Germany, [2] Serine Proteases and Pathophysiology of the Neurovascular Unit, Institut National de la Santé et de la Recherche Médicale UMR-S U919, Caen Normandy University, Groupement d'Intérêt Public (GIP) CYCERON, Caen, France, [3] Institute of Neuroscience and Medicine (INM-6), Computational and Systems Neuroscience, Jülich Research Centre, Jülich, Germany, [4] Faculty of Psychology, Institute of Cognitive Neuroscience, Ruhr-University Bochum, Bochum, Germany

With the emergence of new high performance computation technology in the last decade, the simulation of large scale neural networks which are able to reproduce the behavior and structure of the brain has finally become an achievable target of neuroscience. Due to the number of synaptic connections between neurons and the complexity of biological networks, most contemporary models have manually defined or static connectivity. However, it is expected that modeling the dynamic generation and deletion of the links among neurons, locally and between different regions of the brain, is crucial to unravel important mechanisms associated with learning, memory and healing. Moreover, for many neural circuits that could potentially be modeled, activity data is more readily and reliably available than connectivity data. Thus, a framework that enables networks to wire themselves on the basis of specified activity targets can be of great value in specifying network models where connectivity data is incomplete or has large error margins. To address these issues, in the present work we present an implementation of a model of structural plasticity in the neural network simulator NEST. In this model, synapses consist of two parts, a pre- and a post-synaptic element. Synapses are created and deleted during the execution of the simulation following local homeostatic rules until a mean level of electrical activity is reached in the network. We assess the scalability of the implementation in order to evaluate its potential usage in the self generation of connectivity of large scale networks. We show and discuss the results of simulations on simple two population networks and more complex models of the cortical microcircuit involving 8 populations and 4 layers using the new framework.

Keywords: structural plasticity, large scale neural networks, high performance computing, homeostatic growth, self-organizing network

1. INTRODUCTION

Models of large scale neural networks are an important tool for understanding the mechanics of the brain (De Garis et al., 2010; Helias et al., 2012; Eliasmith and Trujillo, 2014). Such models are created based on experimental information that has been collected for years by neuroscientists and combine mathematical methods with algorithms to reproduce observed behavior. It is known that

the connectivity of the network plays an essential role in defining the way function is achieved at higher levels of activity. Nevertheless, obtaining accurate measurements of connectivity is complex, even with the most advanced experimental techniques, due to the resolution of sensors and difficult access to the target areas. The dynamics of the connectivity are also not yet well understood, although it has been shown that synaptic plasticity is fundamental for understanding how learning and memory work. Non invasive techniques such as DTI imaging and fMRI scans can provide a glimpse to the real complexity of the problem in structure and function. Higher resolution techniques like electron microscopy (Gray, 1959), photostimulation (Dantzker and Callaway, 2000) and electrophysiological recordings (Thomson et al., 2002) provide more detailed connectivity information of specific regions. Regardless, creating an exact connectivity map of even a small region of the brain is extremely challenging (Deco et al., 2008; Essen et al., 2012; Van Essen and Ugurbil, 2012; Reckfort et al., 2013). This poses a significant problem for the modeling approach, as connectivity must be specified. For small networks, parameter scans can be carried out with respect to the unknown or imprecisely known connection probabilities between populations. For larger networks, which are more costly to simulate and also potentially have many more unknown connectivity parameters, this approach is hardly feasible.

One way to address the issue of modeling connectivity within a neural network is to allow a network model to determine its own suitable connectivity to achieve target activity patterns, e.g., experimental measurements of the spiking frequency, which is easier to measure accurately than connectivity. In addition to addressing the problem of network model specification, a framework that accounts for the appearance and disappearance of synapses on the basis of network activity can provide insight into how connectivity is generated during development and learning or even on how healing after lesions takes place (De Paola et al., 2006). It can also help understand how certain structures arise as a result of exposition to adequate external stimuli during critical periods in the development of the brain (Hensch, 2005) and the mechanisms underlying experience dependent structural synaptic plasticity (Holtmaat and Svoboda, 2009).

An appropriate model of structural plasticity that incorporates the dynamic generation, deletion and rewiring of synapses within a network was presented by Butz and van Ooyen (2013). In this model, synapses are represented as connections between pre and a post synaptic elements. The growth or diminishment of these synaptic elements is an independent process for each neuron. The model is based on the idea that plasticity in cortical networks is mainly driven by the need of individual neurons to homeostatically maintain their average electrical activity. As a consequence, if activity is lower than a desired set-point, neurons will form synaptic elements, and remove them when activity becomes too high. Additionally, a minimum level of activity is needed to form synaptic elements at all. If activity falls below this level the neuron will remove synaptic

elements, too. Results show that small networks of hundreds or thousands of neurons robustly grow toward a stable homeostatic equilibrium of activity and connectivity. An important advance on earlier work is that all cell types had different desired average firing rates (achieved by different homeostatic set-points) and developed connectivity accordingly. It was shown that these local rules for structural plasticity can account for network rewiring after a partial loss of external input (deafferentation) and shows remarkable similarities with biological data from network rewiring in the primary visual cortex after focal retinal lesions (Keck et al., 2008; Yamahachi et al., 2009). Further analysis by Butz et al. (2014) of changes in network topology revealed that betweenness centrality could be used as an indicator of successful brain repair, in the sense that it is related to the ability of the neurons to restore their electrical activity by network rewiring. It was concluded by the authors that structural plasticity may account for network reorganization on different spatial scales.

In this work, we provide a complete description of how the structural plasticity model proposed by Butz and van Ooyen (2013) could be implemented in the neuronal network simulator NEST (Gewaltig and Diesmann, 2007) in order to create self-organizing large scale neural networks. We evaluate the scalability of the implementation and assess the performance of the model on two use cases. We demonstrate that our implementation is capable of self-organizing the connectivity within a cortical microcircuit model consisting of $100,000$ neurons in total, starting with a fully disconnected setup. We also show the scenario where partial information of the connectivity is given as initial condition and an stable connectivity pattern is obtained in the end.

The structural plasticity extension to NEST is included in release 2.10.0 (Bos et al., 2015) and creates a novel possibility for setting up large-scale neuronal networks. While supercomputers are required for very large-scale simulation, we show that smaller networks can also be run on a personal workstation or laptop according to the NEST development philosophy. This is a fundamental advantage of this implementation of structural plasticity in terms of capacity to test different configurations, as it provides high flexibility and portability for the neuroscientist.

The corresponding extension of the Python interface of NEST (PyNEST) allows the user to set up their own structural plasticity experiments for large scale networks.

The rest of this work is divided into three major parts. The first describes the major elements of the structural plasticity algorithm and the set of tests that were designed in order to measure the performance of an implementation of this algorithm. We also present some use cases for the structural plasticity framework. In the second part we provide the results of the technical implementation in NEST and describe how the design matches the memory and speed requirements for large scale simulations. We also present results for the use cases described in the previous section. In the third part, we discuss the results of the implementation and performance tests.

Some of this material has previously been presented in abstract form (Naveau and Butz, 2014).

2. MATERIALS AND METHODS

2.1. The Algorithm of Structural Plasticity

The original formulation of the structural plasticity algorithm defined in Butz and van Ooyen (2013) consists of three repeating steps which are described as follows:

(1) Update in electrical activity and intracellular calcium concentration. The electrical activity is calculated for each neuron on a millisecond timescale. The time-averaged level of the neuron's electrical activity drives changes in neuronal morphology. Intracellular calcium concentration is updated according to the electrical activity as follows:

$$\frac{d\text{Ca}}{dt} = \begin{cases} -\frac{Ca(t)}{\tau} + \beta & \text{if the neuron fires} \\ -\frac{Ca(t)}{\tau} & \text{otherwise} \end{cases} \quad (1)$$

where τ is the calcium decay constant and β is the calcium intake constant which indicates how much calcium is accumulated each time the neuron fires. Calcium concentration is linearly proportional to average firing rate and thus is the measure that is used to guide the growth dynamics of the synaptic elements.

(2) Update in synaptic elements. The detailed morphology of the synaptic elements is abstracted, and is represented in this formulation only by the number of possible synaptic contacts on axons (axonal elements representing axonal boutons: senders of synaptic activity) and on dendrites (dendritic elements representing dendritic spines: receivers of synaptic activity) collectively called synaptic elements. Synaptic elements are created or deleted according to a homeostatic rule. In general, the homeostatic rule will create synaptic elements when the activity is lower than the desired setpoint and delete them when the activity is higher until the desired activity level is achieved. This homeostasis is represented by a curve which defines how quickly new elements are created or deleted according to the current level of electrical activity. The original work considers two types of growth curve, linear and Gaussian:

Linear:

$$\frac{dz}{dt} = \nu\left(1 - \frac{1}{\epsilon}Ca(t)\right)$$

where ν is the growth rate and ϵ is the target level of calcium concentration that the neuron should achieve.

Gaussian:

$$\frac{dz}{dt} = \nu\left(2\exp\left(-\frac{Ca(t) - \xi}{\zeta}\right) - 1\right)$$

where $\xi = (\eta + \epsilon)/2$, $\zeta = (\epsilon - \eta)/2\sqrt{\ln 2}$ and ν is the growth rate as before. In this Gaussian growth curve, η represents the minimum amount of calcium concentration that the neuron must have in order to start creating new synaptic elements. Same as in the linear growth, ϵ represents the target level of calcium concentration that the neuron should achieve.

A synaptic element is formed (or deleted) when the rounded down z value increase (or decrease) by one. Newly-formed synaptic elements are initially vacant and available for synapse formation.

(3) Update in connectivity. In every connectivity update, available synaptic elements allow the formation of new synapses and deleted synaptic elements dictate synapse breaking. Every available synaptic element has the same probability to be randomly chosen for a new connection. Synaptic elements to be deleted are also chosen in a uniformly random manner out from the pool of already connected elements. It is important to notice that in this algorithm when a synapse breaks due to the deletion of one synaptic element, the counterpart remains and becomes vacant again. This remaining counterpart can form a new synapse at the next update in connectivity. This effect models network rewiring by re-routing of axons or dendrites.

An important characteristic of this algorithm, is that it relies on global communication to update the connectivity in the network, as available compatible synaptic elements must be matched during the simulation to create new connections. This must be taken into consideration for the design of any implementation of this model.

2.2. Scalability

To assess the scalability of the framework, we designed strong and weak scaling tests of the structural plasticity implementation. For all tests, networks with 80% excitatory and 20% inhibitory neurons were created. The growth rate for synaptic elements in the simulation was set to 4.0×10^{-4} elements/ms for the excitatory elements of the inhibitory population and 1.0×10^{-4} elements/ms for all the other elements. The set point for desired calcium concentration in the excitatory population was defined as 0.05 Ca^{2+}, while in the inhibitory population it was set to 0.2 Ca^{2+}. The calcium concentration intake constant was set to $\beta = 0.001$ and the calcium concentration decay constant to $\tau = 10000.0$ for all neurons. The post synaptic amplitude of individual synapses was set to 1.0 mV. External input was provided using a Poisson generator with a frequency of 10^4 Hz. The post synaptic amplitude of individual synaptic input was set to 0.01 mV. The simulation was run for 100 s, with a step size for the numerical integration of 0.1 ms. The updates in the network connectivity were performed every 10 ms. These values were chosen as they proved to be one parameter combination that allowed for stable self-organizing growth of the network toward the homeostatic equilibrium (See Section 3.3.1 for additional comments on the selection of this parameter set).

Weak scaling tests were performed for networks with 5000 neurons per node and settings of 1, 2, 4, 8, and 16 nodes, each node using 28 cores. Strong scaling tests were performed with a network of $100,000$ on the same hardware configurations as the weak scaling tests. Only physical cores were used, no simultaneous multithreading was enabled. A hybrid optimization approach was chosen, in which MPI is used for communication between nodes and OpenMP for intra node communication. All measurements were performed on the JUROPATEST cluster,

which provides up to 70 nodes (T-Platforms V210s Blades), each with $2 \times$ Intel(R) Xeon(R) CPU E5 $-$ 2695 v3 (Haswell) with 14-core processors (2.30 GHz) and 128 GB DDR memory, running with Scientific Linux release 6.5 (Carbon).

2.3. Use Cases for the Structural Plasticity Framework

The main objective of the structural plasticity framework is to provide the user with a tool to model the dynamic creation and deletion of synapses between neurons of a neural network in a scalable manner. There are several applications in which structural plasticity can be used. In this section we detail two use cases as examples. The first use case shows the basic functionality of the framework and how it can be used to study the relationship between connectivity and activity. We also show how this simple set-up can model critical development periods in the network connectivity. The second example is a more complicated case with several populations, where the objective is to show how connectivity can be self-generated in a network by using the synaptic element growth curves as connectivity fitness rules. All simulations were carried out with NEST version 2.8.0 extended by our structural plasticity implementation.

2.3.1. A Simple Two Population Network

In this initial use case, we generate a network with a total of 1000 leaky integrate and fire neurons, 80% excitatory and 20%

inhibitory. For the excitatory neurons, $\eta = 0.0$, $\epsilon = 0.05$ and $\nu = 1.0 \times 10^{-4}$ elements/ms. For the inhibitory neurons, $\eta = 0.0$, $\epsilon = 0.2$ and $\nu = 1.0 \times 10^{-4}$ elements/ms, except for the excitatory elements which had $\nu = 4.0 \times 10^{-4}$ elements/ms. The connectivity in the system was allowed to evolve using a Gaussian growth curve for 3000 s, with an integration step of 0.1 ms and a delay of the connectivity update equal to 100 integration steps. The simulations were performed on a workstation with 8 Intel core i7 $-$ 4770@3.4 GHz CPUs running openSUSE 13.1.

2.3.2. The Cortical Microcircuit Network

In this second use case, we create a four layer network based on the model of the cortical microcircuit proposed by Potjans and Diesmann (2014). Each layer contains one inhibitory and one excitatory population of leaky integrate and fire neurons. In the simulations presented here, the network starts with the same number of neurons in each population as in the previous study, but without any synaptic connections. For each population, we define a level of desired mean electrical activity based on experimental literature and a growth curve which defines the dynamics of the variation in the number of pre- and post-synaptic elements. These are Gaussian shaped curves with two intersections with the x-axis that determine the minimum amount of electrical activity required to form any synapse (η), and the target mean calcium concentration for the neuron (ϵ). The curves are illustrated in **Figure 1.**
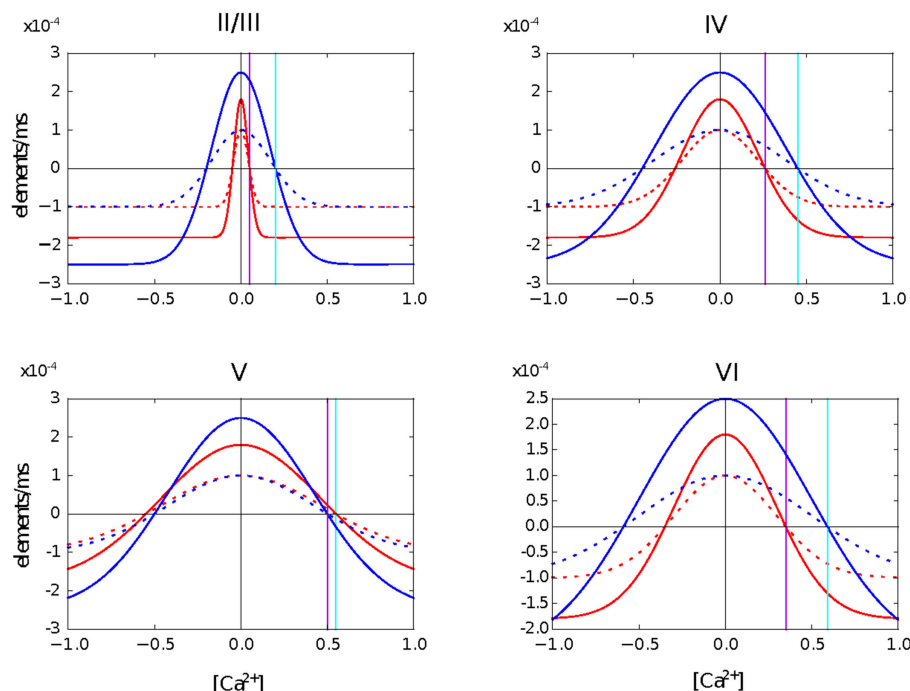
FIGURE 1 | Growth curves for each synaptic element in each layer of the cortical microcircuit model. The growth curves define the rate at which synaptic elements are created depending on the amount of calcium concentration in the cell at the moment. Red curves are for neurons in the excitatory population. Blue curves are for neurons in the inhibitory population. Solid lines are for the excitatory synaptic elements and dotted lines represent inhibitory synaptic elements. The vertical purple line defines the target level of calcium concentration for excitatory neurons and the vertical cyan line represents the target level of calcium concentration for inhibitory neurons. It is important to highlight that all synaptic elements of the same neuron must have a growth curve with the same target level of calcium concentration, otherwise equilibrium will never be reached.

In the first example, we tune the growth rate to achieve an stable growth regime for the network connectivity. This means that the structural plasticity algorithm will stop creating and deleting synaptic connections when the desired mean activity is reached, and that this mean activity is actually reached on average in each population. In a second example, the growth rate provided leads to an unstable connectivity pattern, where the target mean electrical activities are never reached by all populations. A table containing the parameters for both cases can be seen in Appendix in Supplementary Material.

A third example was run to illustrate a more common situation where there are some assumption about the connectivity in a network and where the structural plasticity framework can help to find a suitable balance between excitation and inhibition in the network. Here we used the original model of Potjans and Diesmann (2014) and enable the structural plasticity after an initial stabilization period of 30 s.

Simulations were performed on JUROPATEST (70 nodes with 2×14-core processors Intel(R) Xeon(R) CPU E5 − 2695 v3 (Haswell) at 2.30 GHz and 128 GB DDR memory, running with Scientific Linux release 6.5) and JURECA (with 260 compute nodes with Intel Xeon E5 − 2680 v3 Haswell CPUs with 2×12 cores per CPU, 128 GB of RAM per node and running on CentOS 7 Linux distribution).

3. RESULTS

3.1. Implementation of the Structural Plasticity Model into NEST

The implementation of the structural plasticity algorithm described in this work is based on the version 2.8 of the NEST software Eppler et al. (2015). In accordance with the original formalization described in Section 2.1, the algorithm consists of three repeating parts which can be visualized in a general form in **Figure 2** and described as follows:

FIGURE 2 | Diagram of the implementation of the structural plasticity model in NEST. In (1) the number of synaptic elements is calculated depe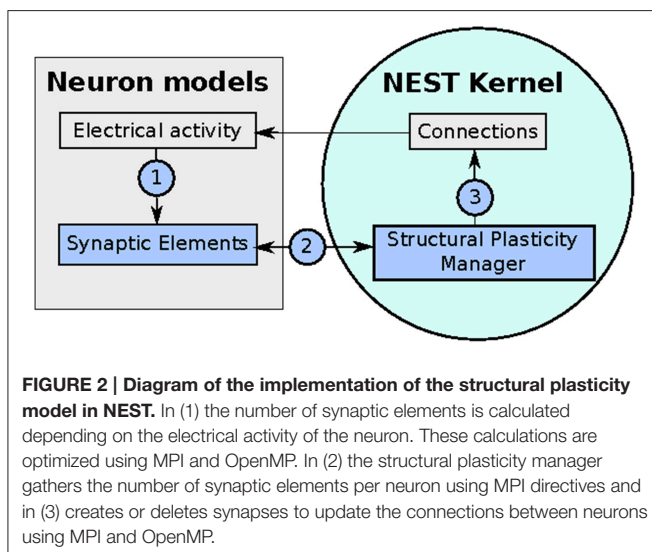nding on the electrical activity of the neuron. These calculations are optimized using MPI and OpenMP. In (2) the structural plasticity manager gathers the number of synaptic elements per neuron using MPI directives and in (3) creates or deletes synapses to update the connections between neurons using MPI and OpenMP.

1. Update in electrical activity and intracellular calcium concentration. The *Archiving_Node* class, which is the general interface for all neurons, was modified and new variables to store the of the intracellular calcium concentration, the calcium decay and calcium intake constants were added. The *Archiving_Node::set_spiketime* method was also modified to update the calcium concentration according to the first case defined in (1), which is performed at every time the neuron spikes. The Node class was modified by adding the method *Archiving_Node::update_synaptic_element*. This method updates the calcium concentration according to the second case defined in (1). This method is called by the *Scheduler* class when every synaptic update interval is reached.

2. Update in synaptic elements. The first step taken in order to design and develop a framework for synaptic elements (e.g., axonal boutons and dendritic spines) was to redefine synapses in such a way that they can now be described using connection elements. This description can be applied to every available neuron model in NEST for generating electrical activity. The design also considers that the users can define their own synaptic elements and their corresponding growth dynamics. The class *SynapticElement* was created in order to represent the connection points for the neurons. The class *GrowthCurve* was also created in order to define the homeostatic rules which guide the creation and deletion of synaptic elements. Currently, the available growth dynamics are based on either a linear or a Gaussian growth curve. The linear growth curve uses an exact integration method to update the number of synaptic elements, while the Gaussian growth curve uses a forward Euler integration method. The framework can be further extended by the user to incorporate more complex element growth dynamics models. An example of such curves is shown in **Figure 1**, where independent dynamics for each type of element in a network of 8 populations (see use case on cortical microcircuit) have been defined. Synaptic elements are used as a discrete value, the actual number of available synaptic elements is an integer truncated from the float variable used to represent them. The *Archiving_Node* class now incorporates a map data structure to store the synaptic elements. The method *Archiving_Node::update_synaptic_element* takes care of updating the number of each of the synaptic elements in the map using the value of the calcium concentration at the time of the call and the corresponding growth curve. The method *Archiving_Node::decay_synaptic_element_vacant* takes care of deleting a percentage of the unused or vacant synaptic elements on every call. Both methods are called by the *Scheduler* at the end of every synapse update interval.

3. Update in connectivity. To coordinate the changes in the structure of the network, a new *StructuralPlasticityManager* class was implemented. At the end of every synapse update interval, the *Scheduler* calls the the newly implemented structural plasticity connectivity manager via the *Network* class. The *StructuralPlasticityManager* determines, for each neuron, how many vacant synaptic elements are available for new synapse formation and how many deleted synaptic elements caused synapse breaking. Then it makes use of the *ConnBuilder* in order to create or delete connections. For this,

the *ConnBuilder* was extended to include the new methods *ConnBuilder::sp_disconnect_* and *ConnBuilder::sp_connect_*. Once new synapses are formed, synaptic elements are tagged from "vacant" to "connected." It is important to notice that when a synapse breaks due to the deletion of one synaptic element, the counterpart remains and becomes vacant again. This remaining counterpart can form a new synapse at the next update in connectivity. This effect preserves the network rewiring capabilities of the original formulation. A detailed diagram of how the new calls are integrated into the normal simulation flow of NEST can be seen in **Figure 3**.

An important feature that we implemented to simulate structural plasticity in NEST is the ability to create and delete synapses during the simulation time. Our new implementation of the connection management overcomes the limitation of the NEST simulator that currently models networks with a fixed connectivity. We have implemented the dynamic creation and deletion of synapses using the new connection framework released in version 2.6.0. The new connection framework improves memory usage to store connection data and reduces the computation time needed to create a connection.

The main limitation of the structural plasticity algorithm described by Butz and van Ooyen (2013) is that it requires global knowledge of the synaptic elements of the entire network. Fortunately, the MPI global communications, also used by the NEST kernel to communicate the electrical activity between the neurons during the simulation, do not pose a substantial bottleneck since changes in connectivity are assumed to take place on average around a factor of 100 times slower than changes in electrical activity. Therefore selecting a biologically realistic growth rate of around 10^{-4} elements/ms will result in an exchange of data that is sufficiently low rate so as not to impact the scalability of the simulator as a whole. At the end of each connectivity update step, the number of created/deleted synaptic elements per neuron are communicated to all MPI processes and a global shuffle subsequently assigns the new pairs of neurons that should be connected, and likewise chooses existing connections for deletion. In the current implementation, no topological constraints are taken into account while deciding which neurons will be connected. The probability of two neurons connecting to each other depends solely on the number of available compatible synaptic elements between them. The actual creation and deletion of the synapses is finally done in parallel using the NEST connection framework. As stated before, a single update in connectivity should not produce a major modification of the network. That means that only a small part of the neurons should create or delete a synaptic element between two updates in connectivity.

It is important to highlight that the usage of global communication is a characteristic of the technical implementation of the algorithm and is not related to the functionality of the model. If topology was to be taken into account, the ability of a neuron to connect to any other would be limited by the constraints imposed by its relative position to others. Global communication would still be used by the implementation, but only relevant information would be taken into account to define the connectivity. The local homeostatic rules only define the creation or deletion of synaptic elements
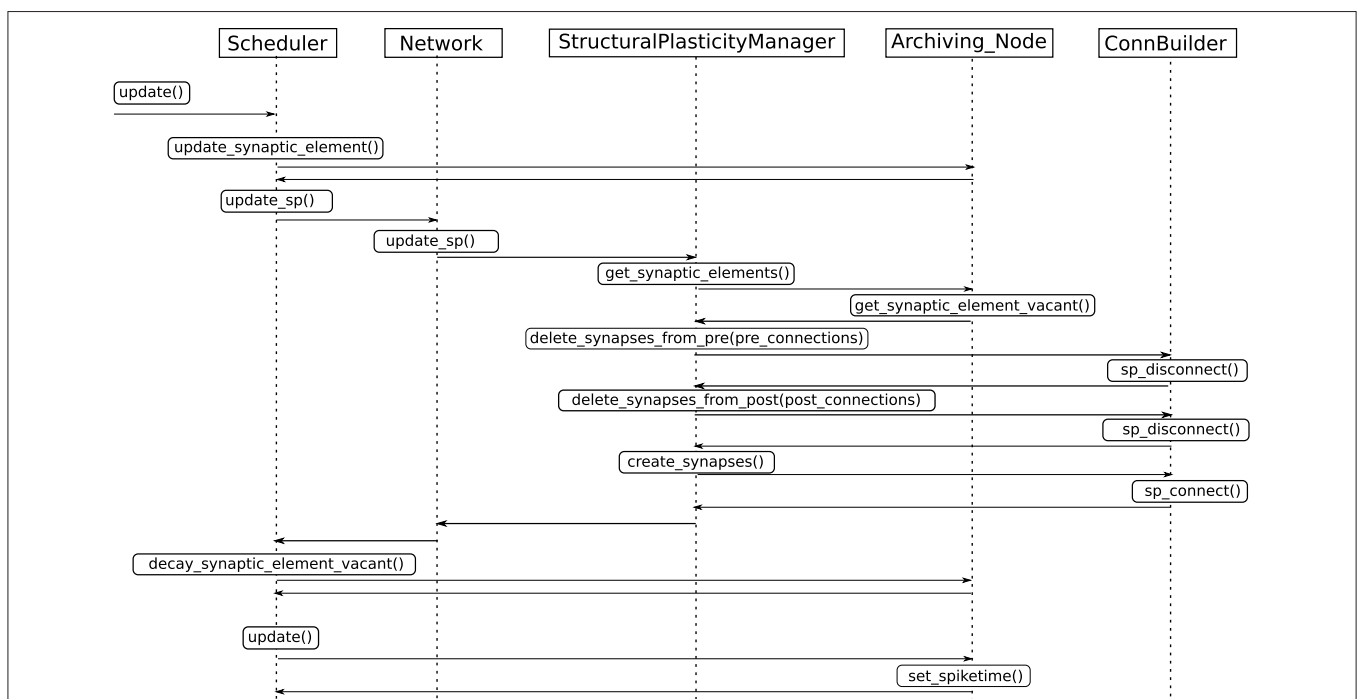
FIGURE 3 | Integration of the new structural plasticity calls into the normal simulation flow of NEST.

per neuron. The number of available synaptic elements is transmitted globaly and the synaptic plasticity manager takes care of forming new synapses or deleting existing ones based on this information.

The update of electrical activity and of the number of synaptic elements is performed by every individual neuron and therefore benefits from the parallel framework already implemented in NEST. Indeed, the NEST software has already demonstrated its high scaling properties on supercomputer, including the JUQUEEN system (Helias et al., 2012; Kunkel et al., 2014).

Finally, the Python interface of NEST (PyNEST) was extended to allow users to easily set up the structural plasticity parameters. It is important to highlight that the user can enable structural plasticity inside the simulation and then disable it when the network has achieved a desired connectivity pattern or activity level. The user can now also delete synapses even without enabling structural plasticity, in a similar way as the connect functions work in NEST.

3.1.1. Setting up a Network in NEST with Structural Plasticity

In this section we will introduce the high level functions that are introduced into NEST with the structural plasticity framework using PyNEST.

In order to set up the network using structural plasticity, one first needs to define the time at which updates in the structure of the network should take place as follows:

```
nest.SetStructuralPlasticityStatus({

    'structural_plasticity_update_
        interval':

        update_interval,

})
```

The next step is to define the synapses which can be dynamically modified by the structural plasticity manager during the simulation. This is achieved by:

```
nest.SetStructuralPlasticityStatus({

    'Structural_plasticity_synapses': {

        'structural_plasticity_synapse_
            ex': {

            'model': 'structural_
                plasticity_synapse_ex',
            'post_synaptic_element'
                : 'Den_ex',
            'pre_synaptic_element'
                : 'Axon_ex',
        },
        'structural_plasticity_synapse_
            in': {

            'model': 'structural_
                plasticity_synapse_in',
            'post_synaptic_element'
                : 'Den_in',
            'pre_synaptic_element'
                : 'Axon_in',
        },
    }

})
```

Here, two types of synapses are being defined, one for the excitatory synapses and another one for the inhibitory synapses. It is important to notice that in this definition, a name

FIGURE 4 | Results of the scalability tests performed with structural plasticity. (A) Efficiency as a function of the number of nodes for 5000 neurons in the weak scaling test. The network was allowed to grow synapses following the structural plasticity rules during a simulation of 100 s of biological time. **(B)** Simulation time (red curve) as a function of the number of nodes for a network of 100, 000 neurons and in the strong scaling test. The blue curve indicates ideal linear scaling. **(C)** Efficiency as a function of the number of nodes for a network of 100, 000 neurons in the strong scaling test. The peak scaling efficiency is marked with a star.

for the post and pre synaptic elements is also specified. This allows the structural plasticity manager to create new synapses of the type specified in model when synaptic elements related to this label become available. This way of setting up the dynamic synapses also allows the user to define static connectivity constraints in the network. This can be achieved by using one synapse model which is not registered for structural plasticity to define this fixed connectivity. For the moment, no other constraints in connectivity like indegree or outdegree ranges can be specified. Nevertheless, thanks to its flexible design, the model can be extended to add new constraints.

Next step involves defining the growth curves for the synaptic elements defined above. This is done as follows:

```
growth_curve_e_e = {

    'growth_curve': "gaussian",
    'growth_rate': 0.0001,
    'continuous': False,
    'eta': 0.0,
    'eps': 0.05,

}
```

This is an example of a Gaussian growth curve where the minimum level of calcium concentration required to start generating synaptic elements is $\eta = 0.0$ Ca^{2+}, and the desired calcium concentration is set to $\epsilon = 0.5$ Ca^{2+}. Finally, the rate at which the synaptic elements will grow is $\nu = 1 \times$
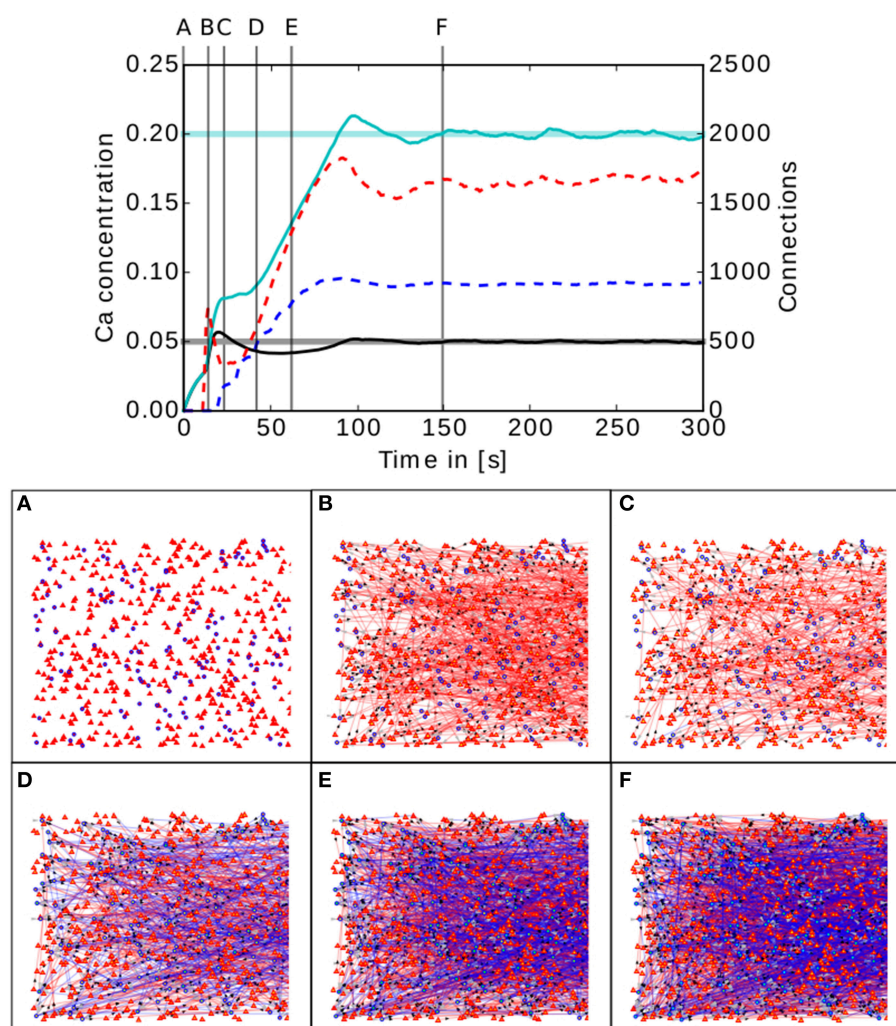
FIGURE 5 | Upper panel: Calcium concentration and numbers of connections as functions of time in a simple two population network. The cyan and black curves show the calcium concentration measured in the inhibitory and excitatory populations, respectively. The paler horizontal lines indicate the corresponding target levels ϵ. The blue and red dashed curves indicate the total number of connections in the inhibitory and excitatory populations, respectively. Vertical gray lines indicate the times of the snapshots displayed in the lower panel. Lower panel (A–F): Evolution of the connectivity in the two population network visualized using MSPViz . Images show half of the total amount of neurons in the network, where triangles represent excitatory neurons and circles inhibitory neurons. Red lines indicate excitatory connections while blue lines indicate inhibitory connections.

10^{-4} elements/ms. Independent growth curves can be created for each synaptic element.

Now that we have defined the growth curve, we can assign this growth curve to the synaptic elements that each neuron will be able to grow. After that, we create the neurons and let NEST know that these synaptic elements are linked to the neurons:

```
synaptic_elements = {

    'Den_ex': growth_curve_e_e,
    'Den_in': growth_curve_e_i,
    'Axon_ex': growth_curve_e_e,

}
nodes = nest.Create('iaf_neuron',
    number_excitatory_neurons)
nest.SetStatus(nodes, 'synaptic_
    elements', synaptic_elements)
```

In this case we are creating the neurons pertaining to the excitatory population. Each neuron has three types of synaptic elements, one dendritic excitatory, one dendritic inhibitory and one axonal excitatory.

The final step is to enable structural plasticity and simulate:

```
nest.EnableStructuralPlasticity()
nest.Simulate(t_sim)
```

A complete PyNEST example which describes how to create a network with two populations, enable structural plasticity and simulate the network is available as Supplementary Material for this paper.

3.2. Scalability

While the update in electrical activity has been proven to scale up to 10^9 neurons, it is important to verify that updating the number of elements and the deletion and formation of synapses does not restrict the expected scaling, at least in the desired regime of up to 10^6 neurons. Updates in synaptic elements and connectivity make use of MPI's "AllGather" communication scheme to communicate the data. This collective communication is also used by the NEST kernel to communicate the spiking

activity between the neurons during the simulation. Although AllGather implements communication between all processes, it is very unlikely that a huge amount of data has to be communicated when a reasonable growth rate of around 10^{-4} elements/ms because updating the number of synaptic elements and the connectivity are very slow processes compared to the update in electrical activity.

3.2.1. Weak Scaling

Figure 4A shows the efficiency, defined as the speed-up divided by the number of nodes, of the implementation as measured by a weak scaling test with 28 OMP threads running on each node. It is visible that, as the number of neurons increases, so does the total number of synapses. The presence of new synapses leads to an increase of communication between neurons, which leads to a decrease in the efficiency of the simulation.

3.2.2. Strong Scaling for a Network of 100,000 Neurons

Figure 4B shows the computation times of the strong scaling tests for a network of $100,000$ neurons, and **Figure 4C** shows the efficiency, defined as speed-up divided by the number of nodes, of the strong scaling test. The peak efficiency is achieved with 4 nodes and 112 cores. These results show supra-linear scaling for this network. In Morrison et al. (2005) and Plesser et al. (2007), supra linear scaling for biological neural networks on NEST was demonstrated due to increasingly efficient caching.

These results show that the introduction of the new structural plasticity framework into NEST has no impact in the scalability of the simulation up to a network size close to that of a cortical column if a suitably low growth rate is selected.

3.3. Performance on the Use Cases
3.3.1. A Simple Two Population Network

The upper panel of **Figure 5** shows the evolution of the calcium concentration and total number of connections for the two population model described in 2.3.1.

The lower panel of **Figure 5** shows a graphical representation of the evolution of the connectivity in the network. During the first 30 s of the simulation, mostly excitatory connections
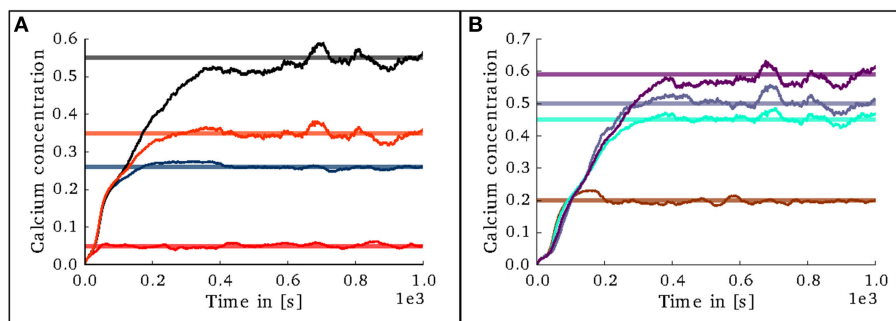
FIGURE 6 | Evolution of calcium concentration in each layer of the cortical microcircuit model. Pale horizontal lines indicate the target concentration of the corresponding population. **(A)** Excitatory populations in layers II/III (red), IV (blue), V (black), and VI (orange). **(B)** Inhibitory populations in layers II/III (brown), IV (cyan), V (gray,) and VI (purple).

are created (**Figure 5A**). This allows the calcium concentration to increase in both populations. When the target mean electrical activity is reached and overshoots in the excitatory population (**Figure 5B**), the number of excitatory connections starts to decrease (**Figure 5C**) until the desired level of calcium concentration is achieved and stabilized in the excitatory population. However, both pre- and post-synaptic elements in the inhibitory population are still being created because it has not yet reached its target mean electrical activity . It is important to remember that neurons have no information regarding the

global status of the network and the evolution of their synaptic elements depends solely on the predefined homeostatic local rules. At around 40 s (**Figure 5D**), an increment in excitatory connections is triggered by the enhanced levels of inhibition. This leads to a complete rewiring of the network (**Figure 5E**). The trend is preserved until the mean electrical activity in the inhibitory population is also reached (**Figure 5F**).

In this network setting, the inhibitory population has a higher level of activity than the excitatory population. It is important to remember that the probability of two neurons connecting
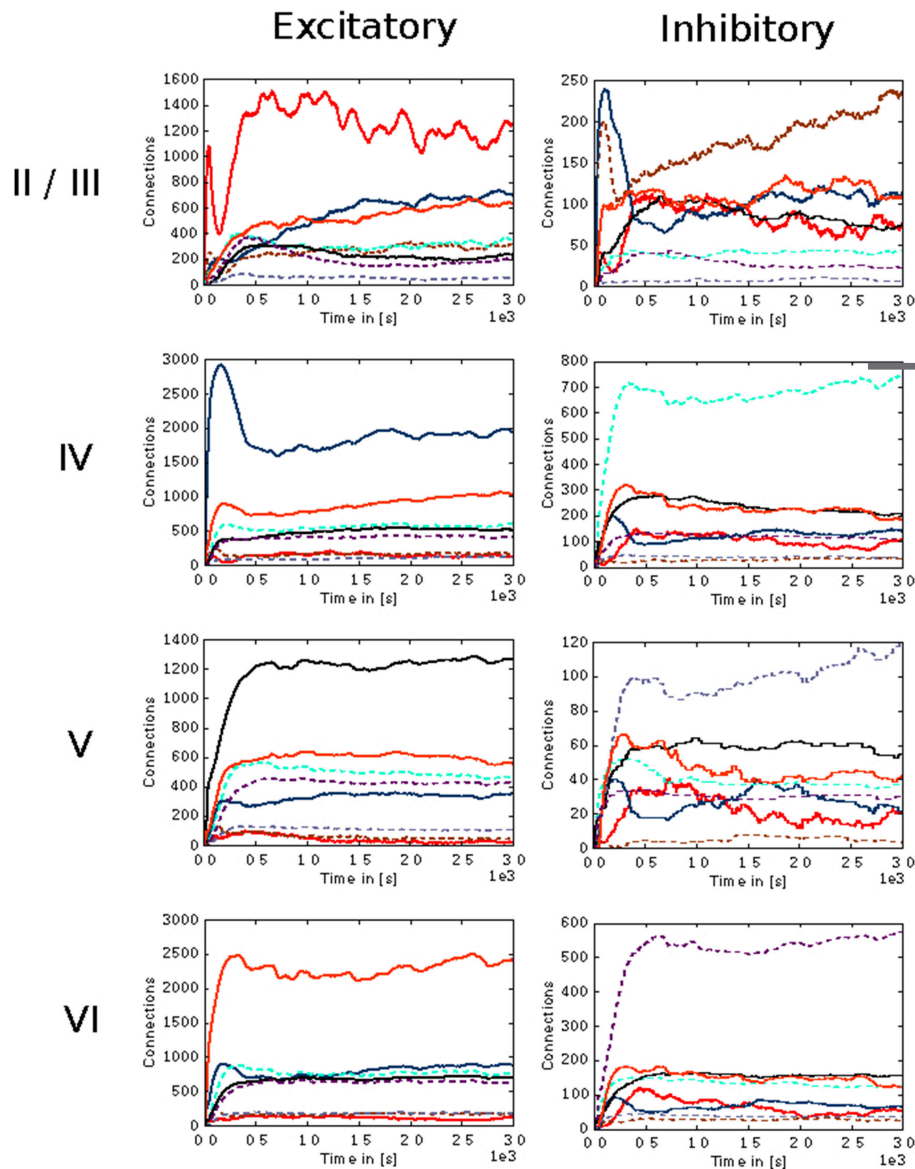
FIGURE 7 | Evolution of connectivity in the microcircuit model resolved by source and target population. Panels on the left and right illustrate efferent connections from excitatory and inhibitory populations, respectively, while the vertical arrangement indicates the layer of the source neurons. In each panel, the numbers of connections to each of the eight population in the model are shown as a function of time. The color of the curves indicates the target population, as in **Figure 6**: connections to excitatory populations in layers II/III (red), IV (blue), V (black), and VI (orange) are shown as solid curves, and connections to inhibitory populations in layers II/III (brown), IV (cyan), V (gray), and VI (purple) are shown as dashed curves.

depends only on the number of available compatible synaptic elements between them. At the start of the simulation, the inhibitory population must offer more post-synaptic elements for excitatory synapses than the excitatory population, otherwise the excitatory population would reach equilibrium first and cease to create excitatory pre-synaptic elements. As a result, not enough excitatory synapses would be created to the inhibitory population and it would never reach the desired level of activity. It is important to remember that the structural plasticity parameter space is broad and a certain amount of exploration is required to discover combinations for each synaptic element which take the network to equilibrium. However, there is in general no unique combination of parameters leading to equilibrium, and different equilibrium combinations will typically produce different connectivity patterns. At this point, biological constraints must be applied to choose between them.

3.3.2. The Cortical Microcircuit Network

In the case of the cortical microcircuit model described in 2.3.2, **Figure 6** shows the changes in calcium concentration, while **Figure 7** shows the evolution of connectivity among layers as the simulation runs. In this case, parameters which lead to stable network connectivity were chosen. Reaching stable connectivity in the networks takes around 700 biological seconds of simulation, which takes 24 h using 25 nodes and 24 cores per node in the JURECA cluster to simulate. It is visible that during the first 20–30 s of simulation, connectivity highly increases on every layer. After the initial overshoot, a smoother approximation toward the desired activity levels is achieved. As seen only from the calcium concentration diagram, the evolution of the network appears to be quite stable. Regardless, the connectivity plots show a continuous dynamical reorganization. While neurons on some layers might start deleting connections due to excess of activity, the post-synaptic neurons must then create new connections in order to compensate for missing activity in case they have not reached their setpoint yet. This leads to a continuous search for compensating excitation and inhibition which must satisfy the requirements of all 8 populations. From **Figure 7** it can be seen that outgoing connections from excitatory populations on layers IV, V, and VI are quite stable. On the other hand, layer II/III exhibits the highest amount of reorganization, both from the excitatory and inhibitory populations. This might be due to the fact that their reduced target levels of activity might be easily influenced by variations in all other layers. Inhibitory populations on all layers in general exhibit a higher degree of reorganization during the whole simulation.

The search space of connectivity parameters for this model of the cortical microcircuit is large as each setup requires 64 values to be defined. If a brute force exploration would be performed on these parameters by simulating each combination for 1 biological second, only $1 - 2$ values per parameter could be considered before more biological seconds would be simulated than using the structural plasticity approach. When adequate synaptic element growth curves are defined, the structural plasticity framework allows a progressive exploration of the space in which the dynamics of the the 8 populations are balanced

at every step, thus providing an efficient way to find stable connectivity combinations.

Figure 8 presents a comparison between the proportional values of connectivity among layers between the results obtained from the simulation using structural plasticity and the original values reported by Potjans and Diesman. The average error in percentual connectivity is of 1.058 ± 1.175.

A second case was also explored, in which parameters lead to unstable network activity are chosen. **Figure 9** shows the evolution of connectivity among layers and **Figure 10** shows the changes in calcium concentration in each layer for this scenario. Overshoots in the connectivity, are originated by a choice of higher rate in the creation of synaptic elements. The system behaves as a feedback control system, with a delay which is defined by the time between updates in connectivity and the synaptic element creation rate. The synaptic element growth rate determines the steepness of the growth curve, and influences the speed at which control changes are made. The instability in the connectivity is reflected in the calcium concentration, never reaching the desired levels. A stable setting involves finding a suitable balance between the speed in the creation of excitatory and inhibitory connections related to the desired level of activity for each layer.

To study the behavior of the structural plasticity algorithm on partially pre-connected networks, another simulation was set in which the initial conditions in connectivity for the structural plasticity algorithm were those specified in the original model of Potjans and Diesman. The network was simulated without plasticity for an initial period of 30 s in order to allow the calcium concentration reach an initial stable value. The evolution of the calcium concentration in all layers after the structural plasticity algorithm was enabled can be seen in **Figures 11A,B**. The stability point is reached a lot faster than in the scenario with no initial connections, at around 400 s. A final simulation was set
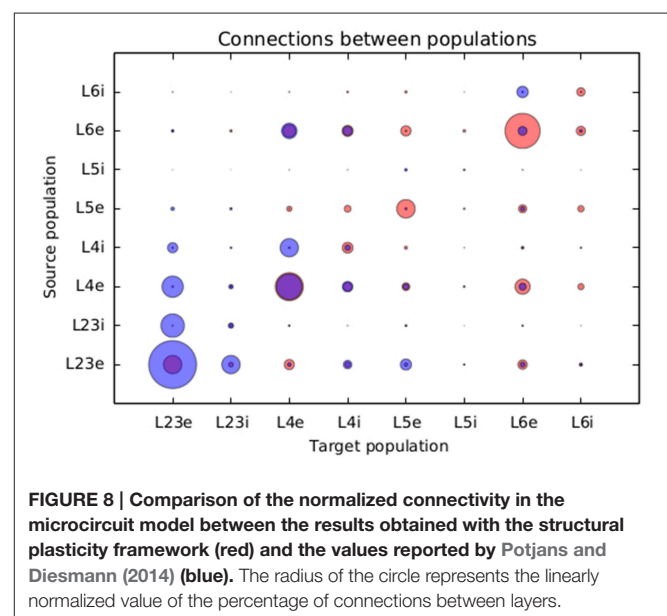
FIGURE 8 | Comparison of the normalized connectivity in the microcircuit model between the results obtained with the structural plasticity framework (red) and the values reported by Potjans and Diesmann (2014) **(blue).** The radius of the circle represents the linearly normalized value of the percentage of connections between layers.
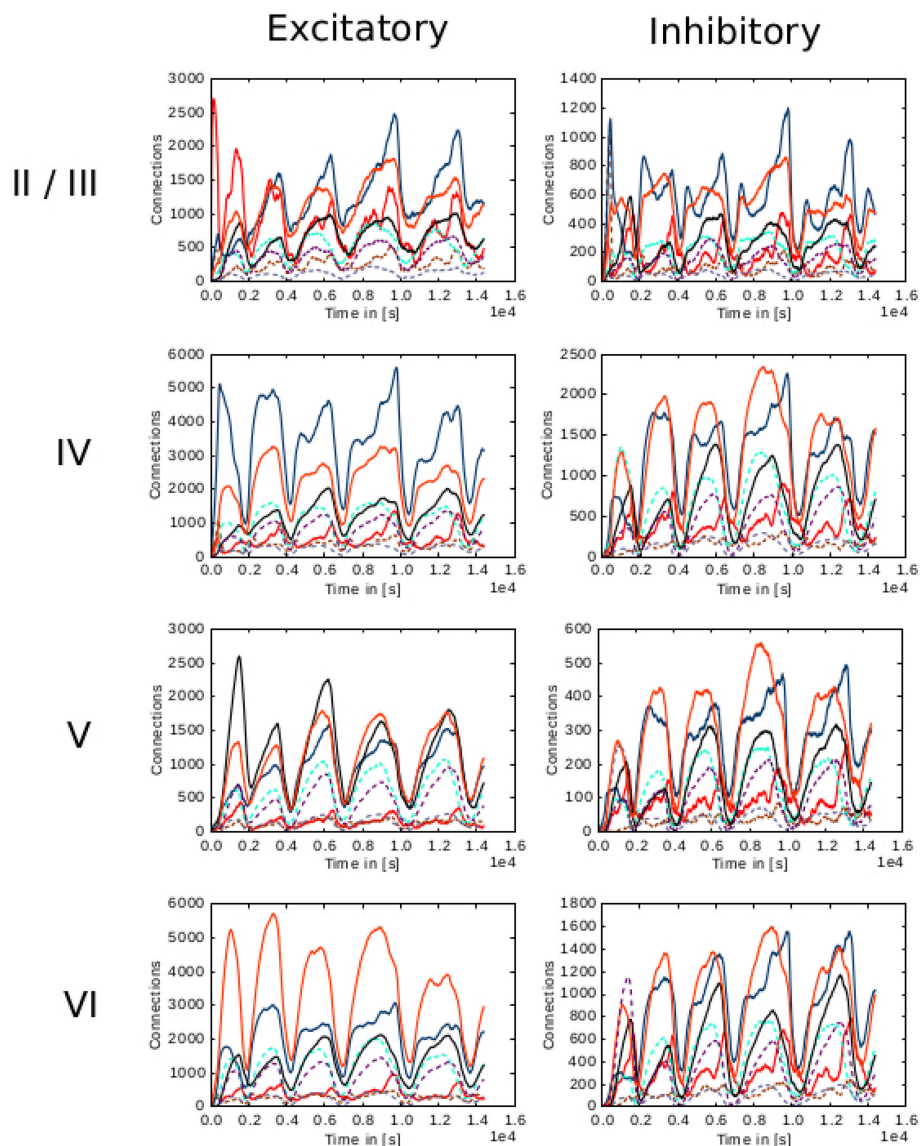
FIGURE 9 | Evolution of connectivity through time for each layer in the cortical microcircuit model with an unstable set of parameters. Panels on the left illustrate connections incoming from excitatory populations. On the right, connections incoming from inhibitory populations. Rows show connections incoming from layers II/III, IV, V, and VI from top to bottom, respectively. On every panel, connections to excitatory populations in layers II/III (red), IV (blue), V (black), and VI (orange) are shown as solid curves, and connections to inhibitory populations in layers II/III (brown), IV (cyan), V (gray), and VI (purple) are shown as dashed curves.

in which the connectivity was specified with a 10% error margin from the original setup reported by Potjans and Diesman. The evolution of the calcium concentration in all layers after plasticity was enabled can be seen in **Figures 11C,D**. The structural plasticity algorithm is able to find a suitable balance between excitation and inhibition. The initial overshoot in electrical activity is a reflection of the initial stronger reconfigurations of the network connectivity. It is important to highlight that a suitable growth scheme is required for the algorithm to reach this stability. Not all setups will become stable or find a solution, this depends on the initial conditions, the desired set points, the shape of the growth curve and the growth rate.

4. DISCUSSION

In this paper we have described the implementation of a framework of structural plasticity for the neural network simulator NEST. We show that the framework is scalable and can be used to model the dynamical creation and deletion of synapses inside a large scale network guided by simple homeostatic rules.

This work also presents some use cases for the framework and some of its potential applications. Researchers can now use structural plasticity in NEST to generate the connectivity of a network from scratch, defining homeostatic rules, in form of synaptic element growth curves, which may vary according to
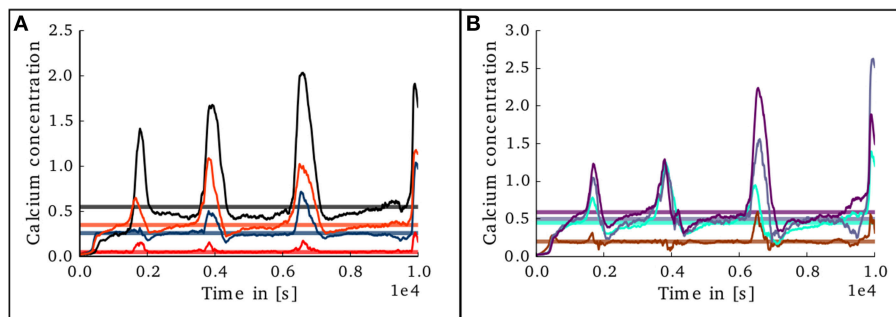
FIGURE 10 | Evolution of calcium concentration in each layer of the cortical microcircuit model with an unstable set of parameters. (A) Shows the calcium concentration in the excitatory populations in layers II/III (red), IV (blue), V (black), and VI (orange). **(B)** Shows the calcium concentration in the inhibitory populations in layers II/III (brown), IV (cyan), V (gray), and VI (purple).
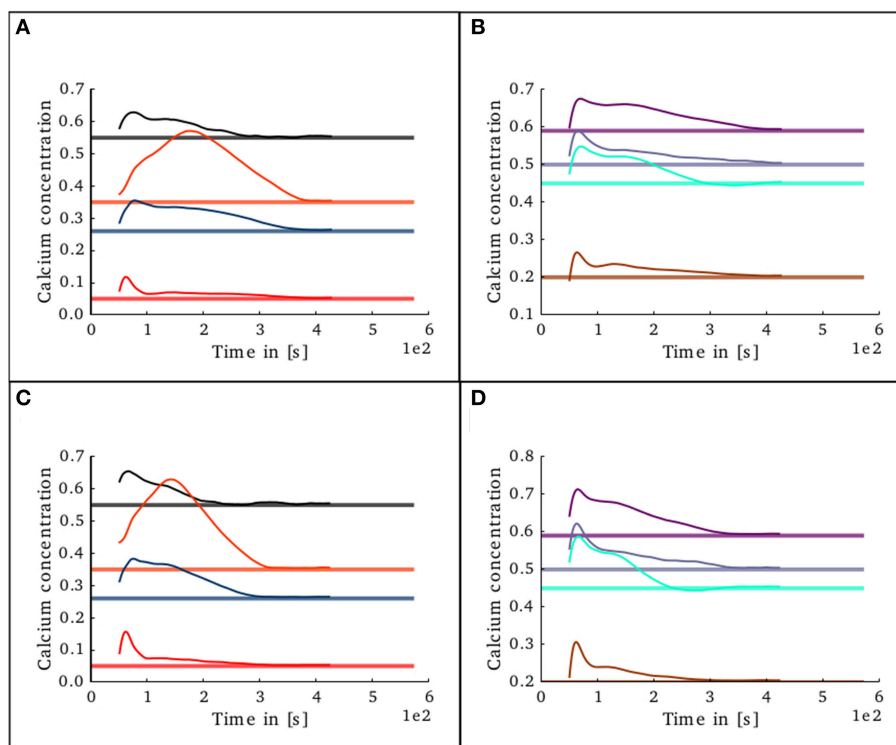
FIGURE 11 | Evolution of calcium concentration in each layer of the cortical microcircuit model with partially pre-connected initial conditions. Pale horizontal lines indicate the target concentration of the corresponding population. Left pannels show excitatory populations in layers II/III (red), IV (blue), V (black), and VI (orange). Right pannels show inhibitory populations in layers II/III (brown), IV (cyan), V (gray), and VI (purple). **(A,B)** Show the scenario where the network was started with the connectivity as specified in the original work by Potjans and Diesman. **(C,D)** Show the same scenario but with a 10% error in the initial connectivity setup.

their needs. The shape of the growth curve defines the speed with which new synaptic elements are created, and as a result, defines the acceleration at which calcium is stored inside the neuron. The relationship between the growth speed at certain level of calcium concentration of excitatory and inhibitory elements is fundamental to achieve stable setups under the model of structural plasticity. As is has been shown, some parameter combinations lead to unstable activity in the network. There are cases where the desired average electrical activity will never be reached by the system. In other cases the average electrical activity will oscillate continuously or suddenly go out of bounds. This relationship depends also on the size of the network and the neuron model used. As a consequence, some care is required in navigating the parameter space in order to achieve desired results.

The example of the two population network illustrates how this framework can be used to understand the interaction

between activity and the creation of synapses. The behavior observed in the simulation can be used to model how inhibition triggers critical periods of connectivity during development of neural networks (Hensch, 2005). During this window, external stimuli can also be used to shape the formation of the new connections. Together with the performance measurements, these results show that our implementation of structural plasticity is suitable to study the development of connectivity patterns inside a neural network in an efficient and scalable manner.

In the specific case of the cortical microcircuit presented in this work, we are able to see some similarities and differences between the results obtained by simulating with the structural plasticity framework and the data reported in Potjans and Diesmann (2014). One of the most visible differences is the smaller amount of recurrent connections generated in the simulation for layer 2/3. This layer has a very low target electrical activity, which is initially almost reached by external input. This means that very few synapses are required to reach this target. This fact limits the creation of synapses for this layer. Note that the results shown in this paper were obtained only by defining target activity levels; no other connectivity constraints were specified. A more elaborate simulation could incorporate tailored growth curves for each layer, and implement additional connectivity restrictions which promote recurrent connections and other connectivity patterns that do not emerge naturally from the current approach.

Another visible difference is that the excitatory population of layers 5 and 6 show a higher number of connections than the ones shown in the original work. On the other hand, connections from and to the inhibitory population of layer 5 and layer 2/3 are well fit. Except from connections between the inhibitory and excitatory populations of layer 4, connections from and to layer 4 are also well predicted.

In this paper we describe a framework which can be used to study structural network dynamics. The focus of this paper is on the technical implementation. It is not the scope of the present work to perform a deep analysis of the biological results that can be obtained using this framework. However, we provide some examples of how the framework can be used, its capacities and limitations. This implementation gives researchers flexibility to explore complex connectivity dynamics by extending the synaptic elements growth rules. As our implementation is integrated into NEST, simulations using structural plasticity can also be combined with other features available in the simulator. For example, the user may take into account dynamic synaptic

weights by mixing this framework with synaptic plasticity. The framework can also be further extended using the current topology framework in NEST in order to constrain connectivity by relative position.

We also show that the structural plasticity algorithm is able to solve the complex balance of interaction between layers with different levels of electrical activity when partial information of the connectivity is available. This result is very promising, as it shows that given the right growth rules, it would now be possible to reconstruct connectivity inside a network without having exact anatomical information. We therefore conclude that our approach represents a novel and useful technique to close the current gaps in information about the connectivity in certain regions of the brain.

AUTHOR CONTRIBUTIONS

SD worked on the implementation of the structural plasticity framework in NEST, simulated the use cases and performed the scalability tests. MN did most of the implementation of the structural plasticity framework in NEST and assessed the simulations. MB gave the theoretical guidance regarding the structural plasticity algorithm and the use cases. AM gave scientific and theoretical guidance on the simulation and implementation of the framework.

FUNDING

ACKNOWLEDGMENTS

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: http://journal.frontiersin.org/article/10.3389/fnana. 2016.00057

REFERENCES

Bos, H., Morrison, A., Peyser, A., Hahne, J., Helias, M., Kunkel, S., et al. (2015). *NEST* 2.10.0. doi: 10.5281/zenodo.44222

Butz, M., Steenbuck, I. D., and van Ooyen, A. (2014). Homeostatic structural plasticity increases the efficiency of small-world networks. *Front. Synaptic Neurosci.* 6:7. doi: 10.3389/fnsyn.2014.00007

Butz, M., and van Ooyen, A. (2013). A simple rule for dendritic spine and axonal bouton formation can account for cortical reorganization after focal retinal lesions. *PLoS Comput. Biol.* 9:e1003259. doi: 10.1371/journal.pcbi.1003259

Dantzker, J. L., and Callaway, E. M. (2000). Laminar sources of synaptic input to cortical inhibitory interneurons and pyramidal neurons. *Nat. Neurosci.* 3, 701–707. doi: 10.1038/76656

De Garis, H., Shuo, C., Goertzel, B., and Ruiting, L. (2010). A world survey of artificial brain projects, part i: Large-scale brain simulations. *Neurocomputing* 74, 3–29. doi: 10.1016/j.neucom.2010.08.004

De Paola, V., Holtmaat, A., Knott, G., Song, S., Wilbrecht, L., Caroni, P., et al. (2006). Cell type-specific structural plasticity of axonal branches and boutons in the adult neocortex. *Neuron* 49, 861–875. doi: 10.1016/j.neuron.2006. 02.017

Deco, G., Jirsa, V. K., Robinson, P. A., Breakspear, M., and Friston, K. (2008). The dynamic brain: from spiking neurons to neural masses and cortical fields. *PLoS Comput. Biol.* 4:e1000092. doi: 10.1371/journal.pcbi.1000092

Eliasmith, C., and Trujillo, O. (2014). The use and abuse of large-scale brain models. *Curr. Opin. Neurobiol.* 25, 1–6. doi: 10.1016/j.conb.2013.09.009

Eppler, J. M., Pauli, R., Peyser, A., Ippen, T., Morrison, A., Senk, J., et al. (2015). *NEST 2.8.0*. doi: 10.5281/zenodo.32969

Essen, D. V., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T., Bucholz, R., et al. (2012). The human connectome project: a data acquisition perspective. *Neuroimage* 62, 2222–2231. doi: 10.1016/j.neuroimage.2012.02.018

Gewaltig, M.-O., and Diesmann, M. (2007). NEST (NEural Simulation Tool). *Scholarpedia* 2:1430. doi: 10.4249/scholarpedia.1430

Gray, E. G. (1959). Axo-somatic and axo-dendritic synapses of the cerebral cortex: an electron microscope study. *J. Anat.* 93(Pt 4), 420.

Helias, M., Kunkel, S., Masumoto, G., Igarashi, J., Eppler, J. M., Ishii, S., et al. (2012). Supercomputers ready for use as discovery machines for neuroscience. *Front. Neuroinform.* 6:26. doi: 10.3389/fninf.2012.00026

Hensch, T. K. (2005). Critical period plasticity in local cortical circuits. *Nat. Rev. Neurosci.* 6, 877–888. doi: 10.1038/nrn1787

Holtmaat, A., and Svoboda, K. (2009). Experience-dependent structural synaptic plasticity in the mammalian brain. *Nat. Rev. Neurosci.* 10, 647–658. doi: 10.1038/nrn2699

Keck, T., Mrsic-Flogel, T. D., Vaz Afonso, M., Eysel, U. T., Bonhoeffer, T., and Hubener, M. (2008). Massive restructuring of neuronal circuits during functional reorganization of adult visual cortex. *Nat. Neurosci.* 11, 1162–1167. doi: 10.1038/nn.2181

Kunkel, S., Schmidt, M., Eppler, J. M., Plesser, H. E., Masumoto, G., Igarashi, J., et al. (2014). Spiking network simulation code for petascale computers. *Front. Neuroinform.* 8:78. doi: 10.3389/fninf.2014.00078

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

Naveau, M., and Butz-Ostendorf, M. (2014). Simulating structural plasticity of large scale networks in NEST. *BMC Neurosci.* 15(Suppl. 1):P194. doi: 10.1186/1471-2202-15-S1-P194

Plesser, H. E., Eppler, J. M., Morrison, A., Diesmann, M., and Gewaltig, M.-O. (2007). "Efficient parallel simulation of large-scale neuronal networks on clusters of multiprocessor computers," in *Euro-Par 2007 Parallel Processing* (Berlin; Heidelberg: Springer), 672–681.

Potjans, T. C., and Diesmann, M. (2014). The cell-type specific cortical microcircuit: relating structure and activity in a full-scale spiking network model. *Cereb. Cortex* 24, 785–806. doi: 10.1093/cercor/bhs358

Reckfort, J., Wiese, H., Dohmen, M., Grässel, D., Pietrzyk, U., Zilles, K., et al. (2013). "Extracting the inclination angle of nerve fibers within the human brain with 3D-PLI independent of system properties," in *Proceeding SPIE, Vol. 8873, Polarization Science and Remote Sensing VI, 88730F* (San Diego, CA). doi: 10.1117/12.2023198

Thomson, A. M., West, D. C., Wang, Y., and Bannister, A. P. (2002). Synaptic connections and small circuits involving excitatory and inhibitory neurons in layers 2–5 of adult rat and cat neocortex: triple intracellular recordings and biocytin labelling *in vitro*. *Cereb. Cortex* 12, 936–953. doi: 10.1093/cercor/12.9.936

Van Essen D. C., and Ugurbil, K. (2012). The future of the human connectome. *Neuroimage* 62, 1299–1310. doi: 10.1016/j.neuroimage.2012.01.032

Yamahachi, H., Marik, S. A., McManus, J. N., Denk, W., and Gilbert, C. D. (2009). Rapid axonal sprouting and pruning accompany functional reorganization in primary visual cortex. *Neuron* 64, 719–729. doi: 10.1016/j.neuron.2009.11.026

Structural Plasticity, Effectual Connectivity, and Memory in Cortex

Andreas Knoblauch[1]* and Friedrich T. Sommer[2]

[1] *Informatics Faculty, Albstadt-Sigmaringen University, Albstadt, Germany,* [2] *Redwood Center for Theoretical Neuroscience, University of California at Berkeley, Berkeley, CA, USA*

Learning and memory is commonly attributed to the modification of synaptic strengths in neuronal networks. More recent experiments have also revealed a major role of structural plasticity including elimination and regeneration of synapses, growth and retraction of dendritic spines, and remodeling of axons and dendrites. Here we work out the idea that one likely function of structural plasticity is to increase "effectual connectivity" in order to improve the capacity of sparsely connected networks to store Hebbian cell assemblies that are supposed to represent memories. For this we define effectual connectivity as the fraction of synaptically linked neuron pairs within a cell assembly representing a memory. We show by theory and numerical simulation the close links between effectual connectivity and both information storage capacity of neural networks and effective connectivity as commonly employed in functional brain imaging and connectome analysis. Then, by applying our model to a recently proposed memory model, we can give improved estimates on the number of cell assemblies that can be stored in a cortical macrocolumn assuming realistic connectivity. Finally, we derive a simplified model of structural plasticity to enable large scale simulation of memory phenomena, and apply our model to link ongoing adult structural plasticity to recent behavioral data on the spacing effect of learning.

Keywords: synaptic plasticity, effective connectivity, transfer entropy, learning, potential synapse, memory consolidation, storage capacity, spacing effect

1. INTRODUCTION

Traditional theories attribute adult learning and memory to Hebbian modification of synaptic weights (Hebb, 1949; Bliss and Collingridge, 1993; Paulsen and Sejnowski, 2000; Song et al., 2000), whereas recent evidence suggests also a role for network rewiring by structural plasticity including generation of synapses, growth and retraction of spines, and remodeling of dendritic and axonal branches, both during development and adulthood (Raisman, 1969; Witte et al., 1996; Engert and Bonhoeffer, 1999; Chklovskii et al., 2004; Butz et al., 2009; Holtmaat and Svoboda, 2009; Xu et al., 2009; Yang et al., 2009; Fu and Zuo, 2011; Yu and Zuo, 2011). One possible function of structural plasticity is effective information storage, both in terms of space and energy requirements (Poirazi and Mel, 2001; Chklovskii et al., 2004; Knoblauch et al., 2010). Indeed, due to space and energy limitations, neural networks in the brain are only sparsely connected, even on a local scale (Abeles, 1991; Braitenberg and Schüz, 1991; Hellwig, 2000). Moreover, it is believed that the energy consumption of the brain is dominated by the number of postsynaptic potentials or, equivalently, the number of functional non-silent synapses (Attwell and Laughlin, 2001; Laughlin and Sejnowski, 2003; Lennie, 2003). Together this implies a pressure to minimize the number

and density of functional (non-silent) synapses. It has therefore been suggested that the function of structural plasticity "moves" the rare expensive synapses to the most useful locations, while keeping the mean number of synapses on a constant low level (Knoblauch et al., 2014). By this, sparsely connected networks can have computational abilities that are equivalent to densely connected networks. For example, it is known that memory storage capacity of neural associative networks scales with the synaptic density, such that networks with a high connectivity can store many more memories than networks with a low connectivity (Buckingham and Willshaw, 1993; Bosch and Kurfess, 1998; Knoblauch, 2011). For modeling structural plasticity it is therefore necessary to define different types of "connectivity," for example, to be able to distinguish between the actual number of anatomical synapses per neuron and the "potential" or "effectual" synapse number in an equivalent network with a fixed structure (Stepanyants et al., 2002; Knoblauch et al., 2014).

In this work we develop substantial new analytical results and insights focusing on the relation between network connectivity, structural plasticity, and memory. First, we work out the relation between "effectual connectivity" in structurally plastic networks and functional measures of brain connectivity such as "effective connectivity" and "transfer entropy." Assuming a simple model of activity propagation between two cortical columns or areas, we argue that effectual connectivity is basically equivalent to the functional measures, while maintaining a precise anatomical interpretation. Second, we give improved estimates on the information storage capacity of a cortical macrocolumn as a function of effectual connectivity (cf., Stepanyants et al., 2002; Knoblauch et al., 2010, 2014). For this we develop exact methods (Knoblauch, 2008) to analyze associative memory in sparsely connected cortical networks storing random activity patterns by structural plasticity. Moreover, we generalize our analyses that are reasonable only for very sparse neural activity, to a recently proposed model of associative memory with structural plasticity (Knoblauch, 2009b, 2016) that is much more appropriate for moderately sparse activity deemed necessary to stabilize cell assemblies or synfire chains in networks with sparse connectivity (Latham and Nirenberg, 2004; Aviel et al., 2005). Third, we point out in more detail how effectual connectivity may relate to cognitive phenomena such as the spacing effect that learning improves if rehearsal is distributed to multiple sessions (Ebbinghaus, 1885; Crowder, 1976; Greene, 1989). For this, we analyze the temporal evolution of effectual connectivity and optimize the time gap between learning sessions to compare the results to recent behavioral data on the spacing effect (Cepeda et al., 2008).

2. MODELING

2.1. Memory, Cell Assemblies and Synapse Ensembles

Memories are commonly identified with patterns of neural activity that can be revisited, evoked and/or stabilized by appropriately modified synaptic connections (Hebb, 1949; Bliss and Collingridge, 1993; Martin et al., 2000; Paulsen and Sejnowski, 2000; for alternative views see Arshavsky, 2006). In the simplest case such a memory corresponds to a group of neurons that fire at the same time and, according to the Hebbian hypothesis that "what fires together wires together" (Hebb, 1949) develop strong mutual synaptic connections (Caporale and Dan, 2008; Clopath et al., 2010; Knoblauch et al., 2012). Such groups of strongly connected neurons are called *cell assemblies* (Hebb, 1949; Palm et al., 2014) and have a number of properties that suggest a function for associative memory (Willshaw et al., 1969; Marr, 1971; Palm, 1980; Hopfield, 1982; Knoblauch, 2011): For example, if a stimulus activates a subset of the cells, the mutual synaptic connections will quickly activate the whole cell assembly which is thought to correspond to the retrieval or completion of a memory. In a similar way, a cell assembly in one brain area u can activate an associated cell assembly in another brain area v. We call the set of synapses that supports retrieval of a given set of memories their *synapse ensemble S*. Memory consolidation is then the process of consolidating the synapses S.

Formally, networks of cell assemblies can be modeled as associative networks, that is, single layer neural networks employing Hebbian-type learning. **Figure 1** illustrates a simple associative network with clipped Hebbian learning (Willshaw et al., 1969; Palm, 1980; Knoblauch et al., 2010; Knoblauch, 2016) that associates binary activity patterns $u^1, u^2, \ldots$ and $v^1, v^2, \ldots$ within neuron populations u and v having size $m = 7$ and $n = 8$, respectively: Here synapses are binary, where a weight W_{ij} may increase from 0 to 1 if both presynaptic neuron u_i and postsynaptic neuron v_j have been synchronously activated for at least θ_{ij} times,

$$W_{ij} = \begin{cases} 1, & \omega_{ij} := \sum_{\mu=1}^{M} R(u_i^{\mu}, v_j^{\mu}) \geq \theta_{ij} \\ 0, & \text{otherwise} \end{cases}. \quad (1)$$

where M is the number of stored memories, ω_{ij} is called the synaptic potential, R defines a local learning rule, and θ_{ij} is the threshold of the synapse. In the following we will consider the special case of Equation (1) with Hebbian learning, $R(u_i^{\mu}, v_j^{\mu}) = u_i^{\mu} \cdot v_j^{\mu}$, and minimal synaptic thresholds $\theta_{ij} = 1$, which corresponds to the well-known Steinbuch or Willshaw model (**Figure 1**; cf., Steinbuch, 1961; Willshaw et al., 1969). Further, we will also investigate the recently proposed general "zip net" model, where both the learning rule R and synaptic thresholds θ_{ij} may be optimized for memory performance (Knoblauch, 2016): For R we assume the optimal homosynaptic or covariance rules, whereas synaptic thresholds θ_{ij} are chosen large enough such that the chance $p_1 := \text{pr}[W_{ij} = 1]$ of potentiating a given synapse is 0.5 to maximize entropy of synaptic weights (see Appendix A.3 for further details). In general, we can identify the synapse ensemble S that supports storage of a memory set $\mathfrak{M}$ by those neuron pairs ij with a sufficiently large synaptic potential $\omega_{ij} \geq \theta_{ij}$ where θ_{ij} may depend on $\mathfrak{M}$. For convenience we may represent S as a binary matrix (with $S_{ij} = 1$ if $ij \in S$ and $S_{ij} = 0$ if $ij \notin S$) similar as the weight matrix W_{ij}.

After learning a memory association $u^{\mu} \rightarrow v^{\mu}$, a noisy input $\tilde{u}$ can retrieve an associated memory content $\hat{v}$ in a single
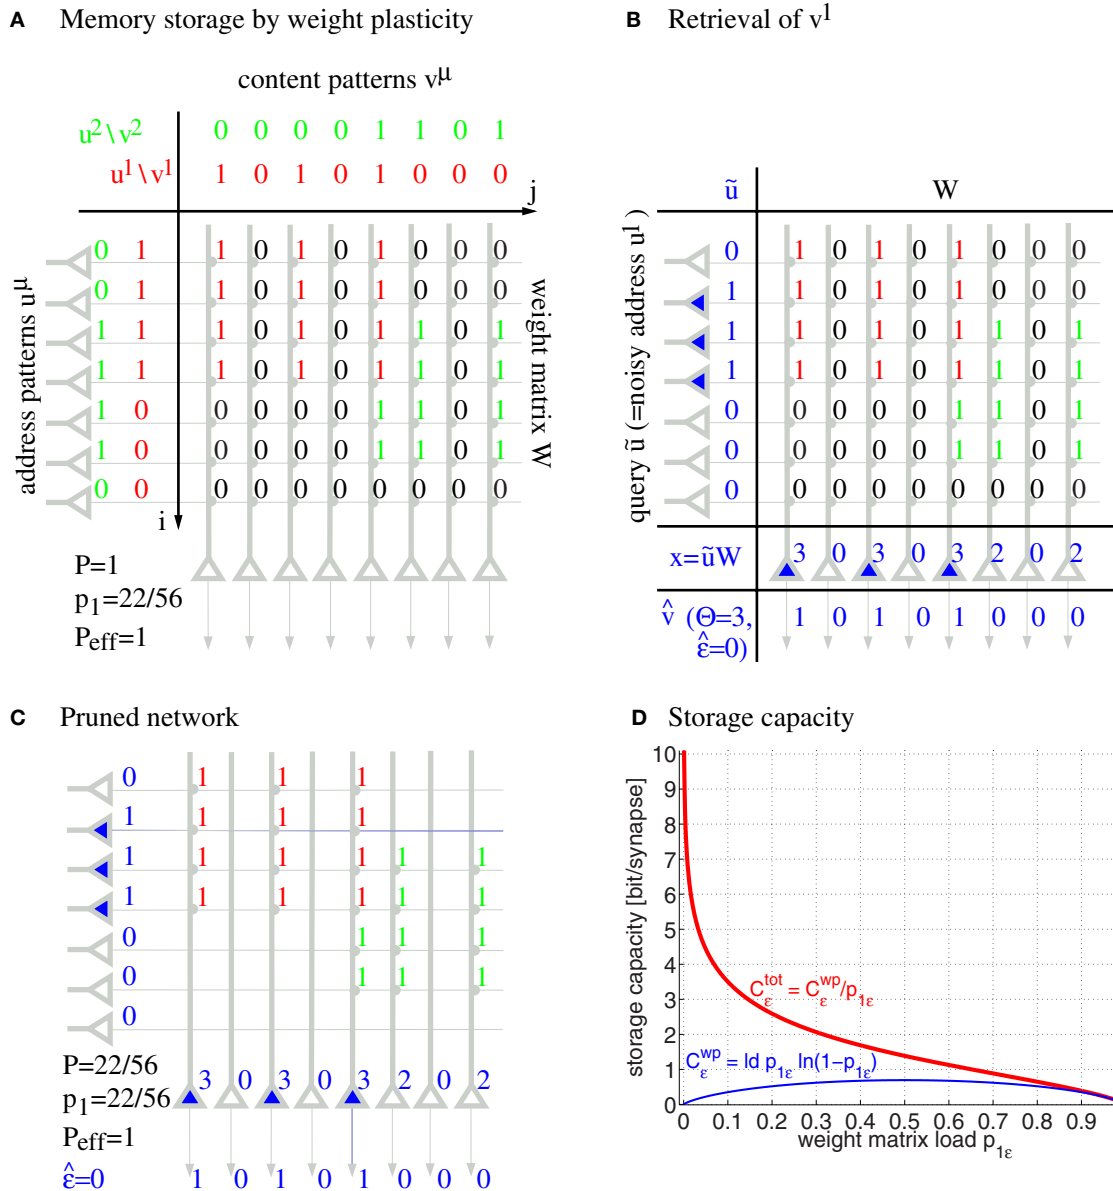
FIGURE 1 | Willshaw model for associative memory. Panels show learning of two associations between activity patterns u^μ and v^μ **(A)**, retrieval of the first association **(B)**, pruning of irrelevant silent synapses **(C)**, and the asymptotic storage capacity in bit/synapse as a function of the fraction p_1 of potentiated synapses **(D)** for networks with and without structural plasticity (C^{tot} vs. C^{wp}; computed from Equations (49, 50, 47) for $P_{eff} = 1$; subscripts ϵ refer to maximized values at output noise level ϵ). Note that networks with structural plasticity can have a much higher storage capacity in sparsely potentiated networks with small fractions $p_1 \ll 1$ of potentiated synapses.

processing step by

$$\hat{v}_j = \begin{cases} 1, & x_j = \left(\sum_{i=1}^{m} \tilde{u}_i W_{ij} + \mathcal{N}_j\right) \geq \Theta_j \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

for appropriately chosen neural firing thresholds Θ_j. The model may include random variables $\mathcal{N}_j$ to account for additional synaptic inputs and further noise sources, but for most analyses and simulations (except Section 3.1) we assume $\mathcal{N}_j = 0$ such that

retrieval depends deterministically on the input $\tilde{u}$. In **Figure 1B**, stimulating with a noisy input pattern $\tilde{u} \approx u^1$ perfectly retrieves the corresponding output pattern $\hat{v} = v^1$ for thresholds $\Theta_j = 2$. In the literature, input and output patterns are also called address and content patterns, and the (noisy) input pattern used for retrieval is called query pattern. In the illustrated completely connected network, the thresholds can simply be chosen according to the number of active units in the query pattern, whereas in biologically more realistic models, firing

thresholds are thought to be controlled by recurrent inhibition, for example, regulating the number of active units to a desired level l being the mean activity of a content pattern (Knoblauch and Palm, 2001). Thus, a common threshold strategy in the more abstract models is to simply select the l most activated "winner" neurons having the largest dendritic potentials x_j. In general, the retrieval outputs may have errors and the retrieval quality can then be judged by the output noise

$$\hat{\epsilon} = \frac{\sum_{j=1}^{n} |\hat{v}_j - v_j^\mu|}{l} \tag{3}$$

defined as the Hamming distance between $\hat{v}$ and v^μ normalized to the mean number l of active units in an output pattern. Similarly, we can define input noise $\tilde{\epsilon}$ as the Hamming distance between $\tilde{u}$ and u^μ normalized to the mean number k of active units in an input pattern.

In the illustrated network u and v are different neuron populations corresponding to hetero-association. However, all arguments will also apply to auto-association when u and v are identical (with $m = n$, $k = l$), and cell assemblies correspond to cliques of interconnected neurons. In that case output activity can be fed back to the input layer iteratively to improve retrieval results (Schwenker et al., 1996). Stable activation of a cell assembly can then expected if output noise $\hat{\epsilon}$ after the first retrieval step is lower than input noise $\tilde{\epsilon}$.

Capacity analyses show that each synapse can store a large amount of information. For example, even without any structural plasticity, the Willshaw model can store $C^{wp} = 0.69$ bit per synapse by weight plasticity (wp) corresponding to a large number of about $n^2/\log^2 n$ small cell assemblies, quite close to the theoretical maximum of binary synapses (Willshaw et al., 1969; Palm, 1980). However, unlike in the illustration, real networks will not be fully connected, but, on a local scale of macrocolumns, the chance that two neurons are connected is only about 10% (Braitenberg and Schüz, 1991; Hellwig, 2000). In this case it is still possible to store a considerable number of memories, although maximal M scales with the number of synapses per neuron, and cell assemblies need to be relatively large in this case (Buckingham and Willshaw, 1993; Bosch and Kurfess, 1998; Knoblauch, 2011).

By including structural plasticity, for example, through pruning the unused silent synapses after learning in a network with high connectivity (**Figure 1C**), the total synaptic capacity of the Willshaw model can even increase to $C^{tot} \sim \log n \gg 1$ bit per (non-silent) synapse, depending on the fraction p_1 of potentiated synapses (**Figure 1D**; see Knoblauch et al., 2010). Moreover, the same high capacity can be achieved for networks that are sparsely connected at any time, if the model includes ongoing structural plasticity and repeated memory rehearsal or additional consolidation mechanisms involving memory replay (Knoblauch et al., 2014).

In Section 3.2 we precisely compute the maximal number of cell assemblies that can be stored in a Willshaw-type cortical macrocolumn. As the Willshaw model is optimal only for extremely small cell assemblies with $k \sim \log n$ (Knoblauch, 2011), we will extend these results also for the general "zip model" of

Equation (1) that performs close to optimal Bayesian learning even for much larger cell assemblies (Knoblauch, 2016).

2.2. Anatomical, Potential, and Effectual Connectivity

As argued in the introduction, connectivity is an important parameter to judge performance. However, network models with structural plasticity need to consider different types of connectivity, in particular, anatomical connectivity P, potential connectivity P_{pot}, effectual connectivity P_{eff}, and target connectivity as measured by consolidation load P_{1S} (see **Figure 2**; cf., Krone et al., 1986; Braitenberg and Schüz, 1991; Hellwig, 2000; Stepanyants et al., 2002; Knoblauch et al., 2014),

$$P := \frac{\#\text{actual synaptic connections}}{mn}, \tag{4}$$

$$P_{pot} := \frac{\#\text{potential synaptic connections}}{mn}, \tag{5}$$

$$P_{eff} := \frac{\sum_{i=1}^{m} \sum_{j=1}^{n} H(W_{ij} S_{ij})}{\sum_{i=1}^{m} \sum_{j=1}^{n} H(S_{ij}^2)}, \tag{6}$$

$$P_{1S} := \frac{\sum_{i=1}^{m} \sum_{j=1}^{n} H(S_{ij}^2)}{mn}, \tag{7}$$

where H is the Heaviside function (with $H(x) = 1$ if $x > 0$ and 0 otherwise) to include the general case of non-binary weights and synapse ensembles ($W_{ij}, S_{ij} \in \mathbb{R}$).

First, *anatomical connectivity P* is defined as the chance that there is an actual synaptic connection between two randomly chosen neurons (**Figure 2A**)[1]. However, for example in the pruned network of **Figure 1C**, the anatomical connectivity P equals the fraction p_1 of potentiated synapses (before pruning) and, thus, conveys only little information about the true (full) connectivity within a cell assembly. Instead, it is more adequate to consider potential and effectual connectivity (**Figures 2B,C**).

Second, *potential connectivity P_{pot}* is defined as the chance that there is a potential synapse between two randomly chosen neurons, where a potential synapse is defined as a cortical location ij where pre- and postsynaptic fibers are close enough such that a synapse could potentially be generated or has already been generated (Stepanyants et al., 2002).

Third, *effectual connectivity P_{eff}* defined as the fraction of "required synapses" that have already been realized is most interesting to judge the functional state of memories or cell assemblies during ongoing learning or consolidation with structural plasticity. Here we call the synapse ensemble S_{ij} required for stable storage of a given memory set also the *consolidation signal*. If ij corresponds to an actual synapse, we

---

[1]More precisely, this means the presence of *at least one synapse* connecting the first to the second neuron. This definition is motivated by simplifications employed by many theories for judging how many memories can be stored. These simplifications include, in particular, (1) point neurons neglecting dendritic compartments and non-linearities, and (2) ideal weight plasticity such that any desired synaptic strength can be realized. Then having two synapses with strength 1 would be equivalent to a single synapse with strength 2. The definition is further justified by experimental findings that the number of actual synapses per connection is narrowly distributed around small positive values (Fares and Stepanyants, 2009; Deger et al., 2012).

may identify the case $S_{ij} > 0$ with tagging synapse ij for consolidation (Frey and Morris, 1997). In case of simple binary network models such as the Willshaw or zip net models, the S_{ij} simply equal the optimal synaptic weights in a fully connected network after storing the whole memory set (Equation 1). Intuitively, if a set of cell assemblies or memories has a certain effectual connectivity P_{eff}, then retrieval performance will be as if these memories would have been stored in a structurally static network with anatomical connectivity P_{eff}, whereas true P in the structurally plastic network may be much lower than P_{eff}.

Last, *target connectivity* or *consolidation load* P_{1S} is the fraction of neuron pairs ij that require a consolidated synapse as specified by S_{ij}. This means that P_{1S} is a measure of the learning load of a consolidation task.

Note that our definitions of P_{eff} and P_{1S} apply as well to network models with gradual synapses ($W_{ij}, S_{ij} \in \mathbb{R}$). More generally, by means of the consolidation signal S_{ij}, we can abstract from any particular network model or application domain. Our theory is therefore not restricted to models of associative memory, but may be applied as well to other connectionist domains, given that the "required" synapse ensembles $\{ij|S_{ij} \neq 0\}$ and their weights can be defined properly by S_{ij}. The following provides a minimal model to simulate the dynamics of effectual connectivity during consolidation.

2.3. Modeling and Efficient Simulation of Structural Plasticity

Figure 3A illustrates a minimal model of a "potential" synapse that can be used to simulate the dynamics of ongoing structural plasticity (Knoblauch, 2009a; Deger et al., 2012; Knoblauch et al., 2014). Here a potential synapse ij^ν is the possible location of a real synapse connecting neuron i to neuron j, for example, a cortical location where axonal and dendritic branches of neurons i and j are close enough to allow the formation of a novel connection by

spine growth and synaptogenesis (Krone et al., 1986; Stepanyants et al., 2002). Note that there may be multiple potential synapses per neuron pair, $\nu = 1, 2, \ldots$. The model assumes that a synapse can be either potential but not yet realized (state π), realized but still silent (state and weight 0), or realized and consolidated (state and weight 1).

For real synapses, state transitions are modulated by the consolidation signal S_{ij} specifying synapses to be potentiated and consolidated Then *structural plasticity* means the transition processes between states π and 0 described by transition probabilities $p_g := \text{pr}[\text{state}(t + 1) = 0|\text{state}(t) = \pi]$ and $p_{e|s} := \text{pr}[\text{state}(t + 1) = \pi|\text{state}(t) = 0, S_{ij} = s]$. Similarly, *weight plasticity* means the transitions between states 0 and 1 described by probabilities $p_{c|s} := \text{pr}[\text{state}(t + 1) = 1|\text{state}(t) = $

FIGURE 3 | Two simple models (A,B) of a potential synapses that can be used for simulating ongoing structural plasticity. State π corresponds to potential but not yet realized synapses. State 0 corresponds to unstable silent synapses not yet potentiated or consolidated. State 1 corresponds to potentiated and consolidated synapses. Transition probabilities of actual synapses (state 0 or 1) depend on a consolidation signal $s = S_{ij}$ that may be identified with the synaptic tags (Frey and Morris, 1997) marking synapses required to be consolidated for long-term memory storage. Thus, typically $p_{c|1} > p_{c|0}$ for synaptic consolidation $0 \rightarrow 1$ and $p_{e|1} < p_{e|0}, p_{d|1} < p_{d|0}$ for synaptic elimination $0 \rightarrow \pi$ and deconsolidation $1 \rightarrow 0$. All simulations assume synaptogenesis $\pi \rightarrow 1$ (by p_g) in homeostatic balance with synaptic elimination such that network connectivity P is constant over time.

● = actual synapse　　　**□ = potential synapse**　　　**● = requested synapse of synapse ensemble S (supportive for storing a memory set)**

FIGURE 2 | Illustration of different types of "connectivity" corresponding to actual (A), potential (B), and requested synapses (C). The requested synapses in **(C)** correspond to the synapse ensemble S required to store the memory patterns in **Figure 1**.

0, $S_{ij} = s$] and $p_{d|s} := \text{pr}[\text{state}(t + 1) = 0 | \text{state}(t) = 1, S_{ij} = s]$. For simplicity, we do not distinguish between long-term potentiation (LTP) and synaptic consolidation (or L-LTP), both corresponding to the transition from state 0 to 1. In accordance with the state diagram of **Figure 3A**, the evolution of synaptic states can then be described by probabilities $p_{\text{state}}^{(s)}(t)$ that a given potential synapse receiving $S_{ij} = s$ is in a certain state $\in \{\pi, 0, 1\}$ at time step $t = 0, 1, 2, \ldots$,

$$p_1^{(s)}(t) = (1 - p_{d|s(t)})p_1^{(s)}(t - 1) + p_{c|s(t)}p_0^{(s)}(t - 1) \quad (8)$$

$$p_0^{(s)}(t) = (1 - p_{c|s(t)} - p_{e|s(t)})p_0^{(s)}(t - 1) + p_{d|s(t)}p_1^{(s)}(t - 1) + p_g p_\pi^{(s)}(t - 1) \quad (9)$$

$$p_\pi^{(s)}(t) = (1 - p_g)p_\pi^{(s)}(t - 1) + p_{e|s(t)}p_0^{(s)}(t - 1)$$
$$= 1 - p_1^{(s)}(t) - p_0^{(s)}(t) , \quad (10)$$

where the consolidation signal $s(t) = S_{ij}(t)$ may depend on time.

The second model variant (**Figure 3B**) can be described in a similar way except that $p_{d|s}$ describes the transition from state 1 to state π. Model B is more convenient to analyze the spacing effect. We will see that, in relevant parameter ranges, both model variants behave qualitatively and quantitatively very similar. However, in most simulations we have used model A.

Note that a binary synapse in the original Willshaw model (Equation 1, **Figures 1A,B**) is a special case of the described potential synapse ($p_g = p_{e|s} = p_{d|s} = 0, p_{c|s} = s \in \{0, 1\}, S_{ij} = W_{ij}$ as in Equation 1). Then pruning following a (developmental) learning phase (**Figure 1C**) can be modeled by the same parameters except increasing $p_{e|s} > 0$ to positive values. Finally, adult learning with ongoing structural plasticity can be modeled by introducing a homeostatic constraint to keep P constant (cf., Equation 69 in Appendix B.1; cf., Knoblauch et al., 2014), such that in each step the number of generated and eliminated synapses are about the same. **Figure 4** illustrates such a simulation for $p_{e|s} = 1 - s$ and a fixed consolidation signal S_{ij} corresponding to the same memories as in **Figure 1**. Here the instable silent (state 0) synapses take part in synaptic turnover until they grow at a tagged location ij with $S_{ij} = 1$ where they get consolidated (state 1) and escape further turnover. This process of increasing effectual connectivity (see Equation 70 in Appendix B.2) continues until all potential synapses with $S_{ij} = 1$ have been realized and consolidated (**Figure 4**, $t = 4$) or synaptic turnover comes to an end if all silent synapses have been depleted.

Microscopic simulation of large networks of potential synapses can be expensive. We have therefore developed a method for efficient simulation of structural plasticity on a macroscopic level: Instead of the lower case probabilities (Equations 8–10) we consider additional memory-specific upper-case connectivity variables $P_{\text{state}}^{(s)}$ defined as the fractions of neuron pairs ij that receive a certain consolidation signal $s(t) = S_{ij}(t)$ and are in a certain state $\in \{\emptyset, \pi, 0, 1\}$ (where $\emptyset$ denotes neuron pairs without any potential synapses). In general it is

$$P_1^{(s)}(t) = P_{\text{pot}}^{(s)} \sum_{\mathfrak{n}=1}^{\infty} \mathfrak{p}(\mathfrak{n}) \left(1 - (1 - p_1^{(s)}(t))^{\mathfrak{n}}\right) \quad (11)$$

$$P_\pi^{(s)}(t) = P_{\text{pot}}^{(s)} \sum_{\mathfrak{n}=1}^{\infty} \mathfrak{p}(\mathfrak{n}) \left(p_\pi^{(s)}(t)\right)^{\mathfrak{n}} \quad (12)$$

$$P_0^{(s)}(t) = P_{\text{pot}}^{(s)} - P_1^{(s)}(t) - P_\pi^{(s)}(t) \quad (13)$$

where $p_1^{(s)}$ and $p_\pi^{(s)}$ are as in Equations (8, 10); $P_{\text{pot}}^{(s)}$ is the fraction of neuron pairs receiving s that have at least one potential synapse; and $\mathfrak{p}(\mathfrak{n})$ is the conditional distribution of potential synapse number $\mathfrak{n}$ per neuron pair having at least one potential synapse. Thus, we define a pre-/postsynaptic neuron pair ij to be in state 1 iff it has at least one state-1 synapse; in state 0 iff it does not have a state-1 synapse but at least one state-0 synapse; and in state π if it is neither in state 1 nor state 0 but has at least one potential synapse. See Fares and Stepanyants (2009) for neuroanatomical estimates of $\mathfrak{p}(\mathfrak{n})$ in various cortical areas.

Summing over s we obtain further connectivity variables P_1, P_0, P_π from which we can finally determine the familiar network connectivities defined in the previous section,

$$P_{\text{state}}(t) = \sum_s P_{\text{state}}^{(s)}(t) \quad \text{for state} \in \{\emptyset, \pi, 0, 1\} \quad (14)$$

$$P(t) = P_0(t) + P_1(t) \quad (15)$$

$$P_{\text{pot}}(t) = P_\pi(t) + P_0(t) + P_1(t) \quad (16)$$

$$P_{1S} = \sum_{s \neq 0} \sum_{\text{state} \in \{\emptyset, \pi, 0, 1\}} P_{\text{state}}^{(s)}(t) \quad (17)$$

$$P_{\text{eff}}(t) = \frac{\sum_{s \neq 0} P_1^{(s)}(t)}{P_{1S}} . \quad (18)$$

In general, the consolidation signal $s = s(t) = S_{ij}(t)$ will not be constant but may be a time-varying signal (e.g., if different memory sets are consolidated at different times). To efficiently simulate a large network of many potential synapses, we can partition the set of potential synapses in groups that receive the same signal $s(t)$. For each group we can calculate the temporal evolution of state probabilities $p_\pi^{(s)}(t), p_0^{(s)}(t), p_1^{(s)}(t)$ of individual synapses from Equations (8–10). From this we can then compute from Equations (11–13) the group-specific macroscopic connectivity variables $P_\pi^{(s)}(t), P_0^{(s)}(t), P_1^{(s)}(t)$, and finally from Equations (14–18) the temporal evolution of the various network connectivities $P_\pi(t), P_0(t), P_1(t), P(t)$ as well as effectual connectivity $P_{\text{eff}}(t)$ for certain memory sets. For such an approach the computational cost of simulating structural plasticity scales only with the number of different groups corresponding to different consolidation signals $s(t)$ (instead of the number of potential synapses as for the microscopic simulations).

Moreover, this approach is the basis for further simplifications and the analysis of cognitive phenomena like the spacing effect described in Appendix B. For example, for simplicity, the following simulations and analyses assume that each neuron pair ij can have at most a single potential synapse [i.e., $\mathfrak{p}(1) = 1$]. In previous works we have simulated also a model variant
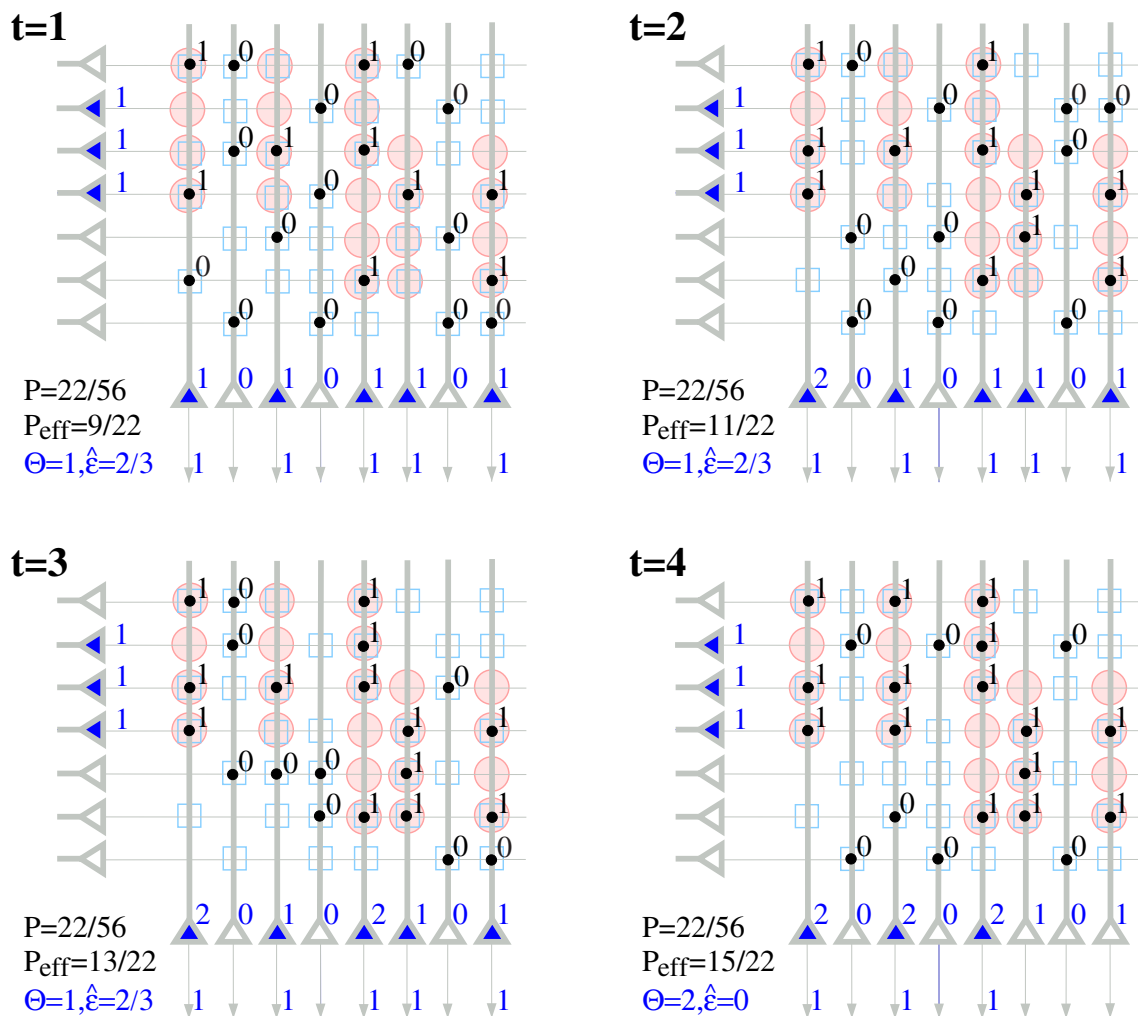
FIGURE 4 | Ongoing structural plasticity maintaining a constant anatomical connectivity *P* = 22/56 for the memory patterns of Figure 1 with actual, potential and requested synapses as in Figure 2 and assuming only single potential synapses per neuron pair ($p(1) = 1$, $p_{e|s} = 1 - s$, $p_{c|s} = s$, $p_{d|s} = 0$). Note that P_{eff} increases with time from the anatomical level $P_{eff} = 9/22 \approx P$ at *t* = 1 toward the level of potential connectivity with $P_{eff} = 15/22 \approx P_{pot}$ at *t* = 4. Correspondingly, output noise $\hat{\epsilon}$ decreases with increasing P_{eff}. At each time firing threshold Θ is chosen maximally to activate at least *l* = 3 neurons corresponding to the mean cell assembly size in the output population.

allowing multiple synapses per neuron pair, where we observed very similar results as for single synapses (Knoblauch et al., 2014). As synapse number per connected neuron pair has sometimes been reported to be narrowly distributed around a small number (e.g., $\mathfrak{n} = 4$; cf., Fares and Stepanyants, 2009), one may also identify each single synapse in our model with a group of about 4 real cortical synapses (see Section 4).

This assumption is actually justified by evidence that $\mathfrak{n}$ is narrowly distributed around a small number, e.g., $\mathfrak{n} = 4$ (Fares and Stepanyants, 2009). This means that two neurons are either unconnected or connected by a group of about four synapses (which is actually a very surprising finding as it is unclear how the neurons can regulate $\mathfrak{n}$; cf., Deger et al., 2012). This situation is well consistent with our modeling assumption $\mathfrak{p}(1) = 1$ if we

identify each model synapse with such a group of about 4 real synapses.

3. RESULTS

3.1. Information Storage Capacity, Effectual Connectivity and its Relation to Functional Measures of Brain Connectivity

For an information-theoretic evaluation, associative memories are typically viewed as memory channels that transmit the original content patterns ν^μ and retrieve corresponding retrieval output pattern $\hat{\nu}^\mu$ (see **Figure 5A**). Thus, the absolute amount of transmitted or stored information C_{abs} of all *M* memories

FIGURE 5 | Relation between effectual connectivity P_{eff}, information storage capacity C, and output noise $\hat{\epsilon}$. (A) Processing model for computing storage capacity $C := C_{abs}/Pmn$ for M given memory associations between input patterns u^μ and output patterns v^μ stored in the synaptic weights (Equation 1; $p := pr[u^\mu = 1]$, $q := pr[v^\mu = 1]$; k and l are mean cell assembly sizes in neuron populations u and v). During retrieval noisy address inputs $\tilde{u}^\mu$ with component errors $p_{ab} := pr[\tilde{u}_i^\mu = b | u_i^\mu = a]$ and input noise $\tilde{\epsilon} := p_{10} + (1/q - 1)p_{01}$ are propagated through the network (Equation 2) yielding output patterns $\hat{v}^\mu$ with component errors $q_{ab} := pr[\hat{v}_j^\mu = b | v_j^\mu = a]$ and output noise $\hat{\epsilon} = q_{10} + (1/q - 1)q_{01}$. The retrieved information is then the transinformation between v^μ and $\hat{v}^\mu$. To simplify analysis, we assume independent transmission of individual (i.i.d.) memory bits v_j^μ over a binary channel with transmission errors q_{01}, q_{10}. **(B)** Information storage capacity $C(P_{eff})$ (blue curve), and output noise $\hat{\epsilon}(P_{eff})$ (red curve) as functions of effectual connectivity P_{eff} for a structurally plastic Willshaw network (similar to **Figure 4**) of $m = n = 100,000$ neurons storing $M = 10^6$ cell assemblies of sizes $k = l = 50$ and having anatomical connectivity $P = 0.1$ assuming zero input noise ($\tilde{\epsilon} = 0$). Data have been computed similar to Equation (37) using Equations (44–46) for $0 \leq P_{eff} \leq P/p_1$.

equals the transinformation or mutual information (Shannon and Weaver, 1949; Cover and Thomas, 1991)

$$C_{abs} := T(\hat{V}; V) := \sum p(\hat{V}, V) \log_2 \frac{p(\hat{V}, V)}{p(\hat{V}) \cdot p(V)} \quad (19)$$

where $V := (v^1, v^2, \ldots v^M)$ and $\hat{V} := (\hat{v}^1, \hat{v}^2, \ldots, \hat{v}^M)$ correspond to the sets of original and retrieved content patterns, and $p(.)$ to their probability distributions. If all M memories and n neurons have independent and identically distributed (i.i.d) activities (e.g., same fraction q of active units per pattern and component transmission error probabilities q_{01}, q_{10}), we can approximate this memory channel by a simple binary channel transmitting $M \cdot n$ memory bits $v_j^\mu \mapsto \hat{v}_j^\mu$ as assumed in Appendix A. Then

$$C_{abs} \approx M \cdot T(\hat{v}^\mu; v^\mu) \approx M \cdot n \cdot T(q, q_{01}, q_{10}) \quad (20)$$

where $T(\hat{v}^\mu; v^\mu)$ is the transinformation for single memory patterns and $T(q, q_{01}, q_{10})$ is the transinformation of a single bit (see Equation 38). From this we obtain the normalized information storage capacity C per synapse after dividing C_{abs} by the number of synapses Pmn (similar to Equation 37).

In our first experiment we have investigated the relation between information storage capacity and effectual connectivity P_{eff} during ongoing structural plasticity. For this we have assumed a larger network of size $m = n = 100000$ with anatomical connectivity $P = 0.1$ and larger cell assemblies with sizes $k = l = 50$, but otherwise a similar setting as for the toy example illustrated by **Figure 4**. **Figure 5B** shows output noise $\hat{\epsilon}$ and normalized capacity C as functions of effectual connectivity P_{eff} for a given number of $M = 10^6$ random

memories. Interestingly, both $\hat{\epsilon}$ and C turn out to be monotonic functions of P_{eff} because output errors decrease with increasing P_{eff} (see Equations 45, 46). Therefore, also output noise $\hat{\epsilon}(P_{eff})$ decreases with increasing P_{eff} whereas, correspondingly, stored information per synapse $C(P_{eff})$ increases with P_{eff}. Because monotonic functions are invertible, we can thus conclude that effectual connectivity P_{eff} is an equivalent measure of information storage capacity or the transinformation (=mutual information) between the activity patterns of two neuron populations u and v. As can be seen from our data, $C(P_{eff})$ tends to be even linear over a large range, $C \sim P_{eff}$, until saturation occurs if $\hat{\epsilon} \rightarrow 0$ approaches zero corresponding to high-fidelity retrieval outputs.

Next, based on the this equivalence between P_{eff} and C, we work out the close relationship between P_{eff} and commonly used functional measures of brain connectivity. Recall that we have introduced "effectual connectivity" as a measure of memory related synaptic connectivity (**Figure 2C**) that shares with other definitions of connectivity (such as anatomical and potential connectivity) the idea that any "connectivity" measure should correspond to the chance of finding a connection element (such as an actual or potential synapse) between two cells. By contrast, in brain imaging and connectome analysis (Friston, 1994; Sporns, 2007) the term "connectivity" has a more heterogeneous meaning ranging from patterns of synaptic connections (anatomical connectivity) and correlations between neural activity (functional connectivity) to causal interactions between brain areas. The latter is also referred to as "effective connectivity" although usually measured in information theoretic terms (bits) such as delayed mutual information or transfer entropy (Schreiber, 2000). For example, in the simplest case the transfer entropy between activities $u(t)$ and $v(t)$ measured in two brain areas u

and v is defined as

$$T_{u \to v} := \sum p(v(t+1), u(t), v(t)) \log_2 \frac{p(v(t+1)|u(t), v(t))}{p(v(t+1)|v(t))} \tag{21}$$

where $p(.)$ denotes the distribution of activity patterns (see Equation 4 in Schreiber, 2000)[2] . Such ideas of effective connectivity come from the desire to extract directions of information flow between two brain areas from measured neural activity, contrasting with (symmetric) correlation measures that can neither detect processing directions nor distinguish between causal interactions and correlated activity due to a common cause.

To see the relation between these functional measures of "effective connectivity" and P_{eff}, first, note that transfer entropy equals the well-known conditional transinformation or conditional mutual information between $v(t+1)$ and $u(t)$ given $v(t)$ (Dobrushin, 1959; Wyner, 1978),

$$T(v(t+1); u(t)|v(t)) :=$$

$$\sum p(v(t+1), u(t), v(t)) \log_2 \frac{p(v(t+1), u(t)|v(t))}{p(v(t+1)|v(t)) \cdot p(u(t)|v(t))} \tag{22}$$

$$= \sum p(v(t+1), u(t), v(t)) \log_2$$

$$\frac{p(v(t+1)|u(t), v(t)) \cdot p(u(t)|v(t))}{p(v(t+1)|v(t)) \cdot p(u(t)|v(t))} = T_{u \to v} . \tag{23}$$

Second, we may apply this to one-step retrieval in an associative memory (Equation 2). Then $u(t) = \tilde{u}^\mu$ is a noisy input, and the update $v(t+1) = F(u(t)) = \hat{v}^\mu$ produces the corresponding output pattern, where the mapping F corresponds to activity propagation through the associative network. As here the update does not depend on the old state $v(t)$, we may approximate transfer entropy by the regular transinformation or mutual information

$$T_{u \to v} = T(v(t+1); u(t)|v(t)) \approx T(F(u(t)); u(t)) \tag{24}$$

$$= I(u(t)) - I(u(t)|F(u(t))) \tag{25}$$

$$= I(F(u(t))) - I(F(u(t))|u(t)) \tag{26}$$

where $I(X) := -\sum_x p(x) \log p(x)$ is the Shannon information of a random variable X, and $I(X|Y) := -\sum_{x,y} p(x, y) \log p(x|y)$ the conditional information of X given Y (Shannon and Weaver, 1949; Cover and Thomas, 1991). Thus, up to normalization, transfer entropy $T_{u \to v} \approx T(F(u(t)); u(t)) = T(\hat{v}^\mu; \tilde{u}^\mu)$ has a very similar form as storage capacity C_{abs} in Equation (20). If $F(u)$ is deterministic, the second term in Equation (26) vanishes and transfer entropy equals the output information $I(F(u(t))) \leq I(u(t))$. If $F(u)$ is also invertible, the second term in Equation (25) would vanish and $T_{u \to v} = I(u(t)) = I(F(u(t))) = C_{abs}/M$. However, in the associative memory application many

input patterns are (ideally) mapped to one memory and $F(u)$ is noninvertible and thus $T_{u \to v} = I(F(u(t))) < I(u(t))$. Moreover, in more realistic cortex models F is also nondeterministic as $v(t+1)$ will depend not only on activity $u(t)$ from a single input area, but also on inputs from further cortical and subcortical areas as well as on numerous additional noise sources. Thus, in fact it will be $T_{u \to v} < I(F(u(t)))$.

Third, we can compare $T_{u \to v}$ to information storage capacity (Equation 20) by normalizing to single memory patterns,

$$\frac{C_{abs}}{M} = \frac{CPmn}{M} = T(\hat{v}^\mu; v^\mu) = T(F(u(t); v^{\mu(u(t))}) \tag{27}$$

$$= I(F(u(t))) - I(F(u(t))|v^{\mu(u(t))}) \tag{28}$$

where $\mu(u(t))$ is a function determining the memory index of the input pattern $u^{\mu(u(t))}$ best matching the current input $\tilde{u} = u(t)$. Thus, comparing Equation (26) to Equation (28) yields generally

$$T_{u \to v} - \frac{C_{abs}}{M} = I(F(u(t))|v^{\mu(u(t))}) - I(F(u(t))|u(t)) \geq 0 . \tag{29}$$

where the bound is true as $v^{\mu(u(t))}$ is a deterministic function of $u(t)$. In particular, for deterministic F, transfer entropy $T_{u \to v} = \frac{C_{abs}}{M} + I(F(u(t))|v^{\mu(u(t))})$ typically exceeds normalized capacity $\frac{C_{abs}}{M}$, whereas equality follows for $I(F(u(t))|v^{\mu(u(t))}) = I(F(u(t))|u(t))$, for example, error-free retrieval with $F(u(t)) = v^{\mu(u(t))}$. Appendix A.4 shows that equality holds generally as well for nondeterministic propagation of activity (e.g., Equation 2 with $\mathcal{N}_j \neq 0$) if we assume that component retrieval errors occur independently with probabilities $q_{01} := pr[\hat{v}_j^\mu = 1|v_j^\mu = 0] \approx pr[\hat{v}_j^\mu = 1|v_j^\mu = 0, \tilde{u}] = pr[\hat{v}_j^\mu = 1|\tilde{u}]$ and $q_{10} := pr[\hat{v}_j^\mu = 0|v_j^\mu = 1] \approx pr[\hat{v}_j^\mu = 0|v_j^\mu = 1, \tilde{u}] = pr[\hat{v}_j^\mu = 0|\tilde{u}]$ corresponding to the same (nondeterministic, i.i.d.) processing model as we have presumed in our capacity analysis (**Figure 5A**; see also Appendix A, Equations 42–43 or Equations 45–46 for Willshaw networks). Then normalizing transfer entropy TE and information capacity CN per output unit yields (see Equations 53, 38)

$$TE := \frac{T_{u \to v}}{n} \overset{>}{\approx} T(q, q_{01}, q_{10}) \approx CN := \frac{C_{abs}}{Mn} = \frac{CPm}{M} . \tag{30}$$

Thus, "effective connectivity" as measured by transfer entropy becomes (up to normalization) equivalent to the information storage capacity C of associative networks (see Equation 37 with Equation 38).

Figure 6 shows upper bounds $TE \leq OE := I(v_j^\mu)$ and lower bounds TE≥CN of transfer entropy as functions of output noise level $\hat{\epsilon} = qq_{10} + (1-q)q_{01}$ for different activities q of output patterns (cf., Equations 26, 29, 30). For low output noise ($\hat{\epsilon} \to 0$) both $T_{u \to v}$ and C approach the full information content of the stored memory set. In general both TE and CN are monotonic functions of $\hat{\epsilon}$ for relevant (sufficiently low) noise levels $\hat{\epsilon}$. While TE increases with $\hat{\epsilon}$ for deterministic retrieval ($\mathcal{N}_j = 0$; cf. Equation 2), TE becomes a decreasing function of $\hat{\epsilon}$ already for low levels of intrinsic noise ($\mathcal{N}_j$ on the order of single synaptic inputs; see panel D). Similar decreases are obtained even without
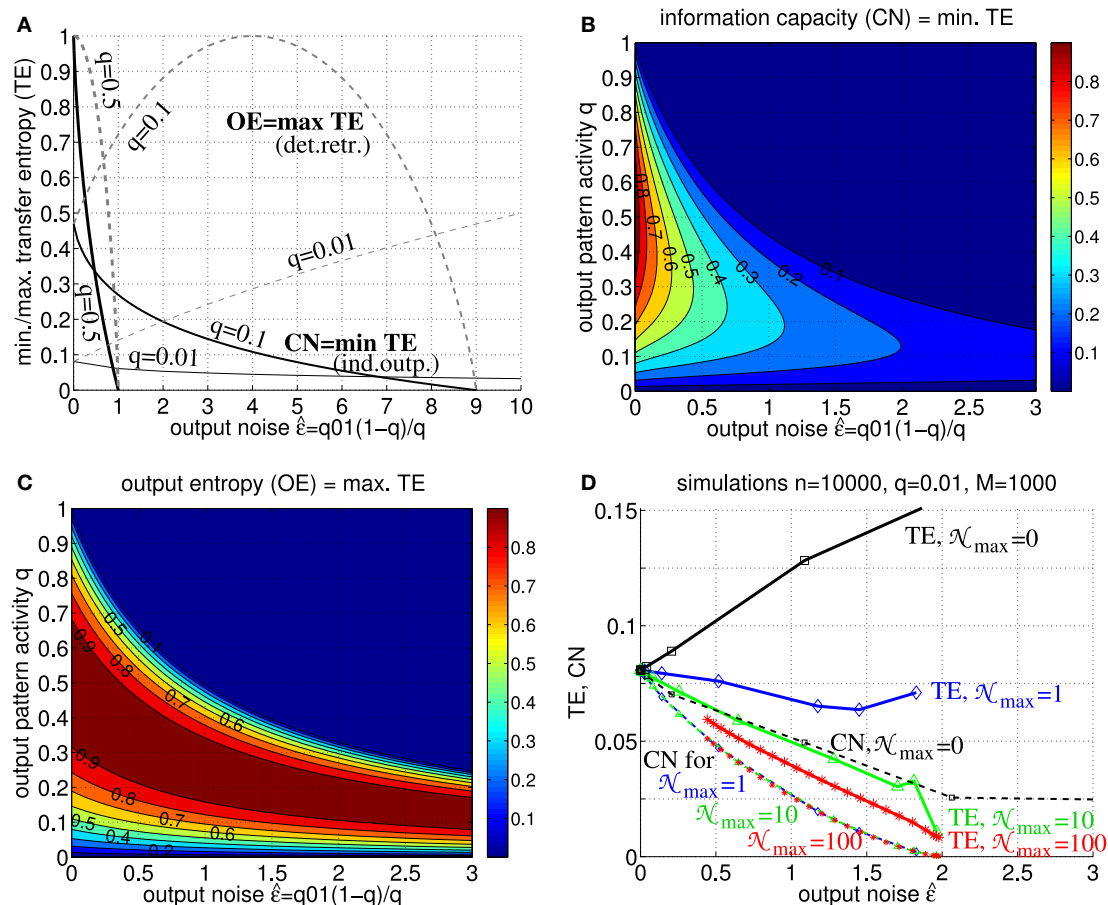
FIGURE 6 | Transfer entropy, output entropy and information capacity. (A) Normalized transfer entropy (TE $:= T_{u \to v}/n$) is bounded by normalized information storage capacity (solid; CN $:= CPm/M \leq$ TE; see Equation 30 with Equation 38) and output entropy (dashed; OE $:= I(\hat{v}_j^{\mu}) \geq$ TE), where TE = OE for deterministic retrieval and TE = CN for non-deterministic retrieval with independent output noise (see text for details). The curves show TE,CN,OE as functions of output noise $\hat{\epsilon} = (1 - q)q_{01}$ assuming only add noise $q_{01} = \text{pr}[\hat{v}_j = 1|v_j = 0]$ but no miss noise $q_{10} = \text{pr}[\hat{v}_j = 0|v_j = 1] = 0$ (e.g., as it is the case for optimal "pattern part" retrieval; see Equation 46 in Appendix A.2). Different curves correspond to different fractions q of active units in a memory pattern (thick, medium, and thin lines correspond to $q = 0.5$, $q = 0.1$, and $q = 0.01$, respectively). **(B)** Contour plot of CN = min TE as function of output noise $\hat{\epsilon}$ and activity parameter q for $q_{10} = 0$. **(C)** Contour plot of OE=max TE as function of output noise $\hat{\epsilon}$ and activity parameter q for $q_{10} = 0$. **(D)** TE (thick solid) and CN (thin dashed) as functions of $\hat{\epsilon}$ for simulated retrieval (zero input noise $\tilde{\epsilon} = 0$) in Willshaw networks of size $n = 10,000$ storing $M = 1000$ cell assemblies of size $k = 100$ ($q = 0.01$) and increasing P_{eff} from 0 to 1 (markers correspond to $P_{eff} = 0.001, 0.01, 0.1, 0.15, 0.2, \ldots, 0.95, 1$). Each data point corresponds to averaging over 10 networks each performing 10,000 retrievals of 100 memories (see Equations 51, 52). Different curves correspond to different levels of intrinsic noise $\mathcal{N}_j$ in output neurons v_j (see Equation 2; $\mathcal{N}_j$ uniformly distributed in $[0; \mathcal{N}_{max}]$ for $\mathcal{N}_{max} = 0, 1, 10, 100$ as indicated by black, blue, green, red lines). Note that, already for low noise levels, retrieval is non-deterministic such that TE becomes monotonic decreasing in $\hat{\epsilon}$ and, thus, similar or even equivalent to CN (and effectual connectivity P_{eff}; see **Figure 5B** and Equation 49; cf. **Figures 7, 8**).

intrinsic noise, $\mathcal{N}_j = 0$, if the target assembly v^{μ} receives (noisy) synaptic inputs from multiple cortical populations (data not shown; cf. Braitenberg and Schüz, 1991).

Our results thus show that, at least for realistic intrinsic noise and/or inter-columnar synaptic connectivity, transfer entropy $T_{u \to v}$ becomes equivalent to information capacity C. Because of the monotonic (or even linear) dependence of C on P_{eff} (see **Figure 5B** and Equation 49; cf. **Figures 7, 8**), transfer entropy is equivalent also to effectual connectivity P_{eff}. Thus, we may interpret effectual connectivity P_{eff} as an essentially equivalent measure of "effective connectivity" as previously defined for functional brain imaging. Still, due to its anatomical definition, P_{eff} can only measure a *potential* causal interaction. For example, if both the synaptic connections from brain area u to v and

the reverse connections from v to u have high P_{eff}, we will not be able to infer the direction of information flow in a certain memory task unless we measure the actual neural activity.

3.2. Storage Capacity of a Macrocolumnar Cortical Network

A typical cortical macrocolumn comprises on the order of $n = 10^5$ neurons below about 1 mm^2 cortex surface, where the anatomical connectivity is about $P = 0.1$ and the potential connectivity about $P_{pot} = 0.5$ corresponding to a filling fraction $f := P/P_{pot} = 0.2$ (Braitenberg and Schüz, 1991; Hellwig, 2000; Stepanyants et al., 2002). Sizes of cell assemblies have been estimated to be somewhere between 50 and 500 in entorhinal

Willshaw model
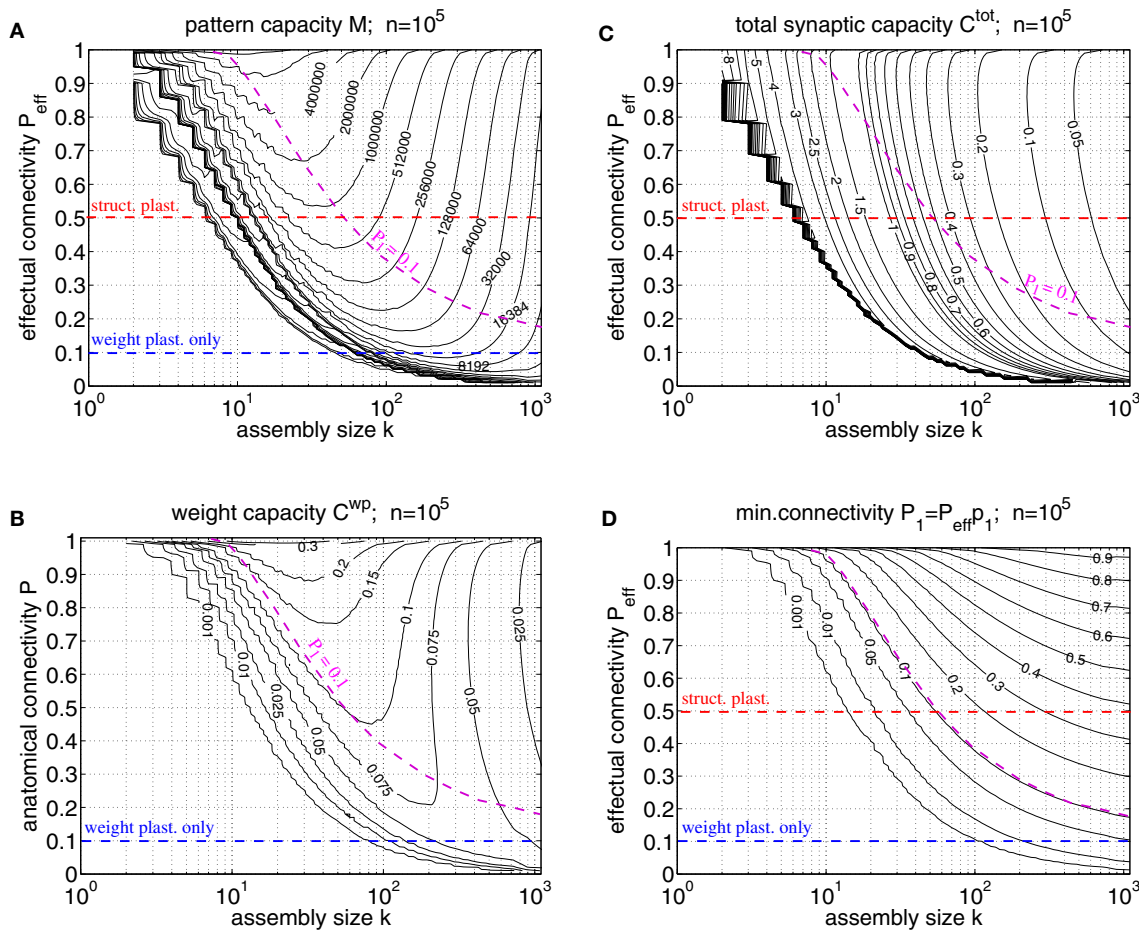
FIGURE 7 | Exact storage capacities for a finite Willshaw network having the size of a cortical macrocolumn ($n = 10^5$). **(A)** Contour plot of pattern capacity M_ϵ (maximal number of stored memories or cell assemblies) as a function of assembly size k (number of active units in a memory pattern) and effectual network connectivity P_{eff} assuming output noise level $\epsilon = 0.01$ and no input noise ($\tilde{u} = u^\mu$). **(B)** Weight capacity C_ϵ^{wp} (in bit/synapse) corresponding to maximal M_ϵ in **(A)** for networks without structural plasticity. **(C)** Total storage capacity C_ϵ^{tot} (in bit/non-silent synapse) corresponding to maximal M_ϵ in **(A)** for networks with structural plasticity. Note that C^{tot} may increase even further if less than the maximum M_ϵ memories are stored (see text for details). **(D)** Minimal anatomical connectivity $P_1 = p_1 P_{\text{eff}} \leq P$ required to achieve the data in **(A-C)**. Data computed as described in Appendix A.1. Red and blue dashed lines correspond to plausible values of P_{eff} for networks with and without structural plasticity (assuming $P = 0.1$, $P_{\text{pot}} = 0.5$). Note that only the area below the magenta dashed line ($P_1 = 0.1$) is consistent with $P = 0.1$. Our exact data is in good agreement with earlier approximative data (Knoblauch et al., 2014, **Figure 5**) unless k is very small (e.g., $k < 50$).

cortex (Waydo et al., 2006). Given these data we can try to estimate the number M of local cell assemblies or memories that can be stored in a macrocolumn (Sommer, 2000). In a previous work (Knoblauch et al., 2014, **Figure 5**) we have estimated the storage capacity for the Willshaw model (**Figures 1**, **4**) by approximating dendritic potential distributions by Gaussians. However, this approximation can be off as, in particular, for sparse activity dendritic potentials can strongly deviate from Gaussians. We have therefore developed a method to compute the exact storage capacity for the Willshaw model storing random memories (see Appendix A). **Figure 7** shows corresponding contour plots of pattern capacity M_ϵ, weight capacity C_ϵ^{wp}, total synaptic capacity C_ϵ^{tot}, and the required minimal anatomical connectivity P_1 (assuming that all silent synapses have been

pruned in the end). We can make several observations: First, the exact results can significantly deviate from the approximations (cf., Knoblauch et al., 2014, **Figure 5**). In particular, for extremely sparse activity ($k < 10$) the Gaussian assumption seems violated and the true capacities are significantly lower than estimated previously. Still, for larger more realistic $50 < k < 500$ the new data is in good agreement with the previous Gaussian estimates, and for even larger $k > 500$ the true capacities even slightly exceed the previous estimates. Second, the previous conclusions, therefore, largely hold: Without structural plasticity ($P_{\text{eff}} = P = 0.1$) the storage capacity would be generally very low and only a small number of memories could be stored. For very sparse $k \approx 50$ not even a single memory could be stored and thus, the cell assembly hypothesis would be inconsistent with experimental
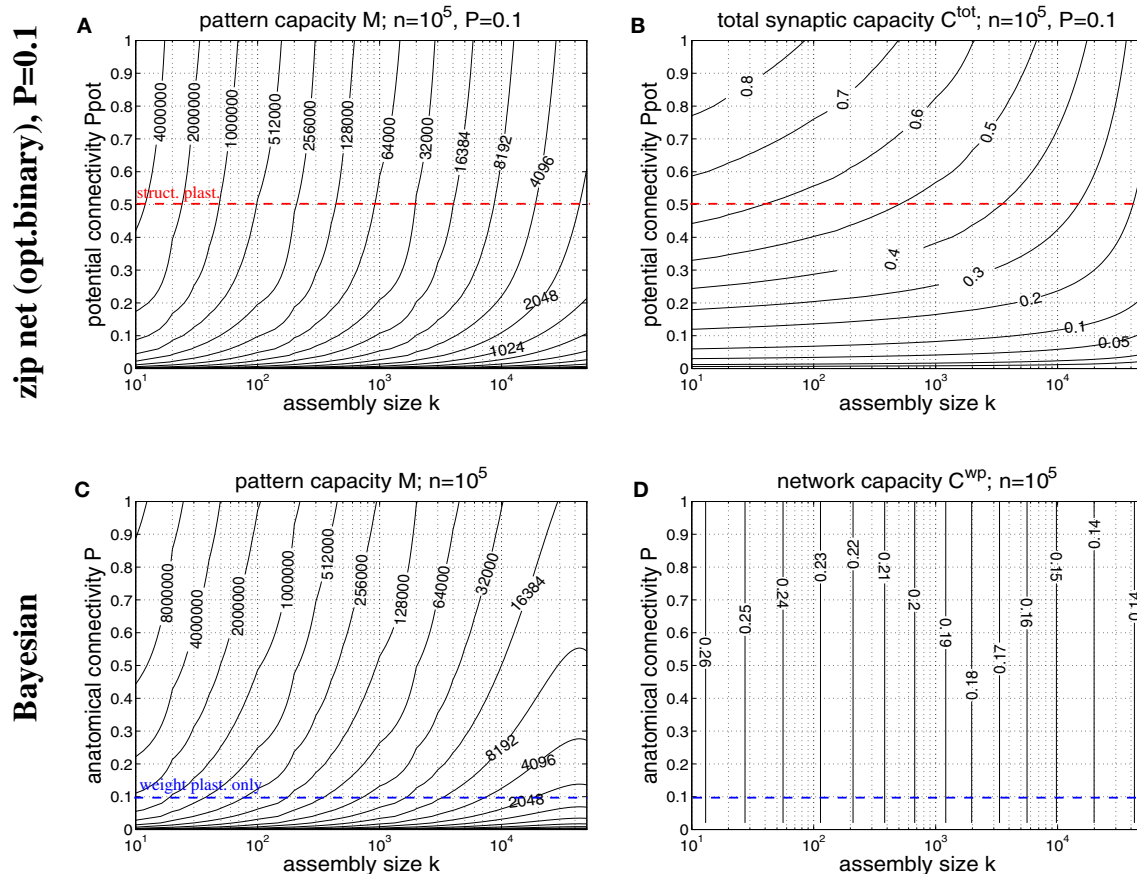
FIGURE 8 | Storage capacities for binary zip nets (A,B) and Bayesian neural networks (C,D) having the size of a cortical macrocolumn ($n = 10^5$). (A) Contour plot of the pattern capacity M_ϵ of an optimal binary zip net (employing the optimal covariance or homosynaptic learning rule; see Knoblauch, 2016) with $P = P_1 = 0.1$ as a function of cell assembly size k and potential network connectivity P_{pot} (which is here an upper bound on the achievable effectual connectivity P_{eff}). **(B)** Total storage capacity C_ϵ^{tot} for zip nets including structural plasticity for the setting of **(A)**. **(C)** Contour plot of the pattern capacity M_ϵ of an optimal Bayesian associative network (Knoblauch, 2011) without structural plasticity as a function of cell assembly size k and anatomical network connectivity P. **(D)** Weight capacity C_ϵ^{wp} for the Bayesian net for the setting of **(C)**. Other parameters are as assumed for **Figure 7** ($\epsilon = 0.01$, $\tilde{u} = u^\mu$). Data computed as described in Appendix A.3. Red and blue dashed lines correspond to plausible values for P_{pot} and P, respectively.

estimates of k. Third, by contrast, networks including structural plasticity increasing P_{eff} from $P = 0.1$ to $P_{pot} \approx 0.5$ can store many more memories: For example, for $k = 50$, the pattern capacity increases from $M \approx 0$ to about $M \approx 800,000$. For $k = 500$, there is still an increase from $M \approx 13,000$ to $M \approx 45,000$. Fourth, correspondingly, networks without structural plasticity would have only a very small weight capacity C^{wp}: For example, at $P_{eff} = P = 0.1$ it is $C^{wp} \approx 0$bps for $k \leq 50$ and still $C^{wp} < 0.07$bps for $k = 500$. Fifth, by contrast, networks with structural plasticity have a much higher total synaptic capacity C^{tot}, i.e., they can store much more information per actual synapse and are therefore also much more energy-efficient, in particular for sparse activity: Although the very high values $C^{tot} \to \log n$ are approached only for unrealistically low k and high P_{eff}, they can still store $C^{tot} \approx 0.5$bps for realistic $P_{eff} = 0.5$ and $k = 50$. This high value appears to decrease, however, to only $C^{tot} \approx 0.06$bps for $k = 500$ which would suggest that, for relatively large cell assemblies with $k = 500$, a

network without structural plasticity (at $P = 0.1$) would be more efficient than a network with structural plasticity (at $P_{eff} = 0.5$). However, as the Willshaw model is known to be sub-optimal for relatively large $k \gg \log n$, we will re-discuss this issue below for a more general network model. Sixth, another weakness of the Willshaw model is that the fraction $p_1 := 1 - (1 - \frac{k^2}{n^2})^M$ of 1-synapses is coupled both to cell assembly size k and number of stored memories M (due to the fixed synaptic threshold $\theta = 1$, cf., Equation 1). Therefore, the residual (minimal) anatomical connectivity of a pruned network $P_1 = p_1 P_{eff}$ depends also on k, M, and we can obtain $P_1 \approx P = 0.1$ consistent with physiology only in a limited range of the k-P_{eff}-planes of **Figure 7**. At least, physiological $k \approx 50$ and $P_{eff} \approx 0.5$ match physiological $P_1 = 0.1$, whereas larger $k \gg 50$ would require P_1 being larger than the anatomical connectivity $P = 0.1$. As many cortical areas comprise significant fractions $P_0 > 0$ of silent synapses we may as well allow for smaller $P_1 < P = 0.1$ satisfying $P_0 + P_1 = P$ (where C^{tot} would become a measure only of energy efficiency,

but no longer of space efficiency), but the very high values of $C^{\text{tot}} \gg 1$ can generally be reached only for tiny fractions of 1-synapses.

In order to overcome some weaknesses of the Willshaw model we have recently proposed a novel network model (so called binary "zip nets") where the fraction p_1 of potentiated 1-synapses is no longer coupled to cell assembly size k and number M (Knoblauch, 2009b, 2010b, 2016). Instead, the model assumes that synaptic thresholds θ_{ij} (see Equation 1) are under homeostatic control to maintain a constant fractions p_1 (or P_1) of potentiated 1-synapses. We have shown for the limit $Mpq \to \infty$ that this model can reach for $p_1 = 0.5$ up to a "zip" factor $\zeta \approx 0.64$ almost the same high storage capacities M_ϵ and C_ϵ^{wp} as the optimal Bayesian neural network (Kononenko, 1989; Lansner and Ekeberg, 1989; Knoblauch, 2011), although requiring only binary synapses. Moreover, if compressed by structural plasticity, zip nets can also reach $C_\epsilon^{\text{tot}} \to \log n$ for $p_1 \to 0$, similar to the Willshaw model. As the Willshaw model is optimal only for extremely sparse activity ($k \le \log n$) it is thus interesting to evaluate the performance gain of structural plasticity for physiological k using the zip net instead of the Willshaw model. **Figure 8** shows data from evaluating storage capacity of a cortical macrocolumn of size $n = 10^5$ both for the zip net model (upper panels) and the Bayesian model (lower panels), the latter being a benchmark for the optimal network without structural plasticity (Knoblauch, 2011). In order to compute the capacity of the zip net we have assumed physiological anatomical connectivity $P = P_1 = 0.1$ where structural plasticity "moves" the $P_1 n^2$ relevant 1-synapses to the most useful locations within the limits given by potential connectivity P_{pot} (as P_1 is fixed, unlike to the Willshaw model, final P_{eff} after learning may be lower than P_{pot} in zip nets; see Appendix A.3 for methodological details). We can make the following observations: First, as expected, for high connectivity and very sparse activity (e.g., $k \ll 100$) the zip nets may perform worse than the Willshaw model (because the Willshaw model then performs close to the optimal Bayesian net). Second, for more physiological parameters $P_{\text{pot}} \le 0.5$, $k \ge 50$ the zip net can store significantly more memories than the Willshaw model, for example, for $P_{\text{pot}} = 0.5$ the zip net reaches $M \approx 1000000$ for $k = 50$ and still $M \approx 120,000$ for $k = 500$. Third, also the total synaptic capacity C^{tot} is higher than for the Willshaw network, for example for $P_{\text{pot}} = 0.5$, it is $C^{\text{tot}} \approx 0.6$ for $k = 50$ and still $C^{\text{tot}} \approx 0.5$ for $k = 500$ (remember that the corresponding value for the Willshaw model required unphysiological $P_1 > 0.1$). Fourth, although the Bayesian network can store significantly more memories M it has only a moderate storage capacity below $C^{\text{wp}} = 0.25$. In fact, for plausible cell assembly sizes, the binary synapses of the zip net with structural plasticity at $P = 0.1$ and $P_{\text{pot}} = 0.5$ achieve more than double the capacity of the optimal (but biologically implausible) Bayesian network with real-valued synapses at $P = 0.1$.

In summary, the new data confirms our previous conclusion that structural plasticity strongly increases space and energy efficiency of associative memory storage in neural networks under physiological conditions (Knoblauch et al., 2014).

## 3.3. Structural Plasticity and the Spacing Effect

In previous works we have linked structural plasticity and cognitive effects like retrograde amnesia, absence of catastrophic forgetting, and the spacing effect (Knoblauch, 2009a; Knoblauch et al., 2014). Here we focus on a more detailed analysis of the spacing effect that learning is most efficient if learning is distributed in time (Ebbinghaus, 1885; Crowder, 1976; Greene, 1989). For example, learning a list of vocabularies in two sessions each lasting 10 min is more efficient than learning in a single session of 20 min. We have explained this effect by slow ongoing structural plasticity and fast synaptic weight plasticity: Thus, spaced learning is useful because during the (long) time gaps between two (or more) learning sessions structural plasticity can grow many novel synapses that are potentially useful for storing new memories and that can quickly be potentiated and consolidated by synaptic weight plasticity during the (brief) learning sessions (Knoblauch et al., 2014, Section 7.3).

Appendix B.2 develops a simplified theory of the spacing effect that is based on model variant B of a potential synapse (which can more easily be analyzed than model A; see **Figure 3**) and the concept and methods proposed in Section 2.3. In particular, with (Equations 73–75) we can easily compute the temporal evolution of effectual connectivity $P_{\text{eff}}(t)$ for arbitrary rehearsal sequences of a novel set of memories to be learned. As output noise $\hat{\epsilon}$ is a decreasing function of P_{eff} (see **Figure 5B**), we can use P_{eff} as a measure of retrieval performance.

To illustrate the effect of spaced vs. non-spaced rehearsal (or consolidation) on P_{eff}, and to verify the theory in Appendix B.2, **Figure 9** shows the temporal evolution of $P_{\text{eff}}(t)$ for different models and synapse parameters. It can be seen that for high potential connectivity $P_{\text{pot}} \approx 1$ and low deconsolidation probability $p_{d|s} \approx 0$ the spacing effect is most pronounced and the network easily realizes high-performance long-term memory (with high P_{eff}; see panel A). Larger $p_{d|0} > 0$ is plausible to model short-term memory, whereas realizing long-term memory would then require repeated consolidation steps (panels B–D). Significant spacing effects are visible for any parameter set. Comparing the microscopic simulations of both synapse models from **Figure 3** to the macroscopic simulations using the methods of Section 2.3 and Appendix B.2, it can be seen that all model and simulation variants behave qualitatively and quantitatively very similar. This justifies to use the theory of Appendix B.2 in the following analysis of recent psychological experiments exploring the spacing effect.

For example, Cepeda et al. (2008) describe an internet-based learning experiment investigating the spacing effect over longer time intervals of more than a year (up to 455 days). The structure of the experiment followed **Figure 10**. The subjects had to learn a set of facts in an initial study session. After a gap interval (0–105 days) without any learning the subjects restudied the same material. After a retention interval (RI; 7–350 days) there was the final test.

These experiments showed that the final recall performance depends both on the gap and the RI showing the following characteristics: First, for any gap duration, recall performance
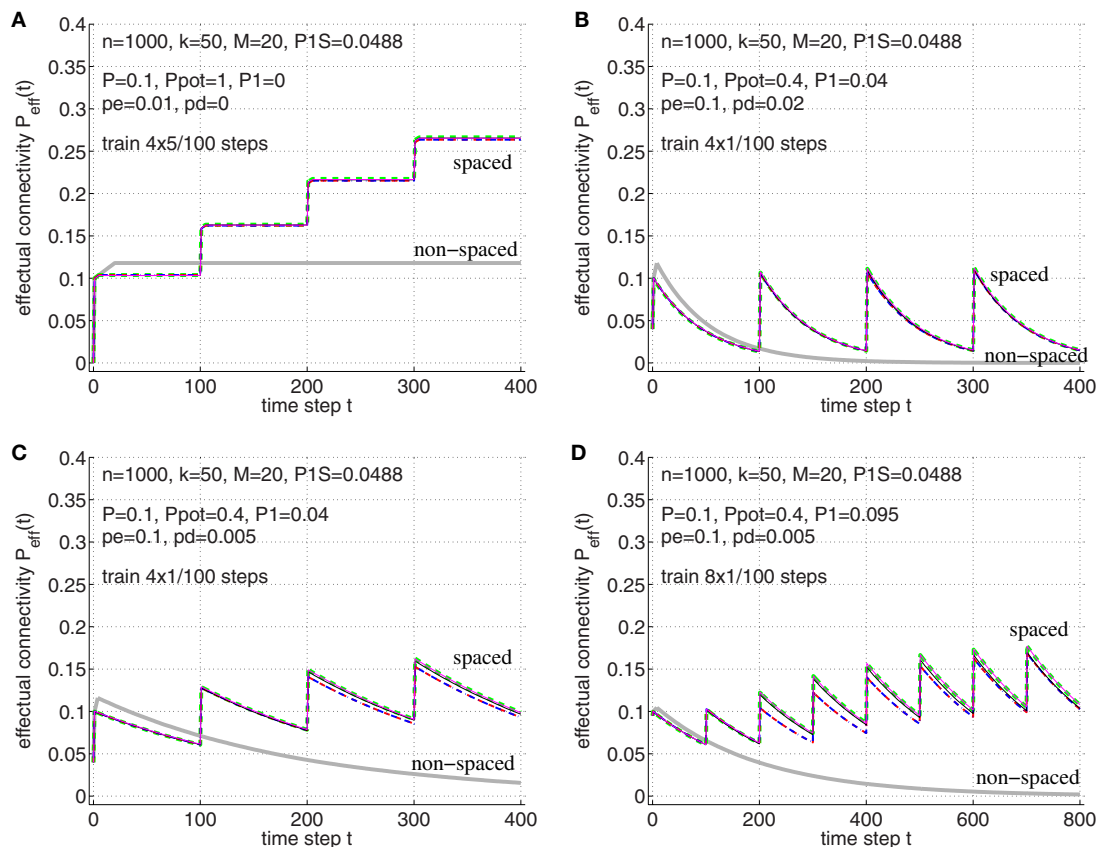
FIGURE 9 | Verification of the theoretical analyses of the spacing effect in Section B.2 in Appendix. Each curve shows effectual connectivity P_{eff} over time for different network and learning parameters. Thin solid lines correspond to simulation experiments of synapse model A (magenta; see **Figure 3A**) and synapse model B (black; see **Figure 3B**), where both variants assume that at most one synapse can connect a neuron pair ($\mathbf{p}(1) = 1$). Green dashed lines correspond to the theory of synapse model A in Appendix B.1 (see Equations 54–56). Blue dash-dotted lines correspond to the theory of synapse model B in Appendix B.2 (see Equations 71–72) and, virtually identical, red-dashed lines correspond to the final theory of model B (see Equations 73–75). For comparison, thick light-gray lines correspond to non-spaced rehearsal of the same total duration as the spaced rehearsal sessions (using model A). **(A)** Spaced rehearsal of a set of $M = 20$ memories at times $t = 0 - 4$, $100 - 104$, $200 - 204$, and $300 - 304$. Each memory had $k = l = 50$ active units out of $m = n = 1000$ neurons corresponding to a consolidation load $P_{1S} \approx 0.0488$. Further we used anatomical connectivity $P = 0.1$, potential connectivity $P_{pot} = 1$, initial fraction of consolidated synapses of $P_1 = 0$ and $p_{e|1} = p_{d|1} = 0$, $p_{c|s=s}$. In each simulation step a fraction $p_e := p_{e|0} = 0.01$ of untagged silent synapses was replaced by new synapses at other locations, but there was no deconsolidation $p_d := p_{d|0} = 0$. **(B)** Similar parameters as before, but $P_{pot} = 0.4$, $P_1 = 0.04$, $p_e = 0.1$, and $p_d = 0.02$. Memories were rehearsed for a single time step $t = 0$, $t = 100$, $t = 200$, and $t = 300$. **(C)** Similar parameters as for panel B, but smaller $p_d = 0.05$. **(D)** Similar parameters as for panel C, but larger $P_1 = 0.095$, i.e., 95 percent of real synapses are initially consolidated. Rehearsal times were $t = 0, 100, 200, \ldots, 700$. Note that the theoretical curves for model A closely match the experimental curves (magenta vs. green). The theory for model B is still reasonably good (black vs. blue/red), although panel D shows some deviations to the simulation experiments. Such deviations may be due to the small number of unstable silent synapses (P_1 near P). In any case, synapse models A and B behave very similar.
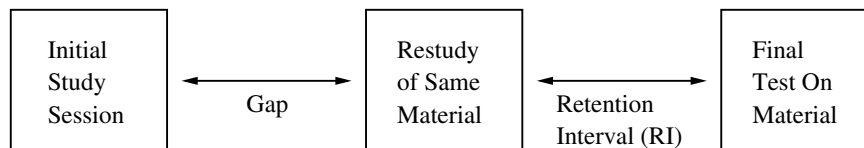
FIGURE 10 | Structure of a typical study of spacing effects on learning. Study episodes are separated by a varying gap, and the final study episode and test are separated by a fixed retention interval. Figure modified from Cepeda et al. (2008).

decline as a function of RI in a negatively accelerated fashion, which corresponds to the familiar "forgetting curve." Second, for any RI greater than zero, an increase in study gap causes

recall to first increase and then decrease. Third, as RI increases, the optimal gap increases, whereas that ratio of optimal gap to RI declines. The following shows that our simple associative

memory model based on structural plasticity can explain most of these characteristics.

It is straight-forward to model the experiments of Cepeda et al. (2008) by applying our model of structural plasticity and synaptic consolidation. **Figure 11** illustrates $P_{\text{eff}}(t)$ for a learning protocol as employed in the experiments: In an initial study session facts are learned until time $t^{(1)}$ when some desired performance level $P_{\text{eff}}^{(1)}$ is reached. After a gap the facts are rehearsed briefly at time $t^{(2)}$ reaching a performance equivalent to $P_{\text{eff}}^{(2)}$. After the retention interval at time $t^{(3)}$ performance still corresponds to an effectual connectivity $P_{\text{eff}}^{(3)}$.

Similar to Cepeda et al. (2008), we want to optimize the gap duration in order to maximize $P_{\text{eff}}^{(3)}$ for a given retention interval RI. After the second rehearsal at time $t^{(2)}$, P_{eff} decays exponentially by a fixed factor $1 - p_{d|0}$ per time step (Equation 74). Therefore, $P_{\text{eff}}^{(3)} = P_{\text{eff}}^{(2)}(1-p_{d|0})^{t^{(3)}-t^{(2)}}$ is a function of $P_{\text{eff}}^{(2)}$ that decreases with the retention interval length $t^{(3)} - t^{(2)}$. We can therefore equivalently maximize $P_{\text{eff}}^{(2)}$ with respect to the gap length $\Delta t := t^{(2)} - t^{(1)}$. For $p_{c|s} = s$, $p_{e|1} = p_{d|1} = 0$, a good approximation of $P_{\text{eff}}^{(2)}$ follows from Equation (73),

$$
P_{\text{eff}}^{(2)} \approx \frac{\begin{array}{c} PP_{\text{pot}} + [(P_{\text{pot}} - P)P_{\text{eff}}^{(1)} - P_{\text{pot}}P_1^{(t1)}](1 - p_{d|0})^{\Delta t} \\ -P_{\text{pot}}(P - P_1^{(t1)})(1 - p_{e|0})^{\Delta t} \end{array}}{P_{\text{pot}} - P_1^{(t1)}(1 - p_{d|0})^{\Delta t} - [P - P_1^{(t1)}](1 - p_{e|0})^{\Delta t}} , \quad (31)
$$

where $P_1^{(t1)} := P_1^{(t0)}(1 - P_{1S})(1 - p_{d|0})^{t^{(1)}} + P_{1S}P_{\text{eff}}^{(1)}$ with $P_1^{(t0)}$ denoting the initial fraction of consolidated synapses at time 0.[3] Since $P_{\text{eff}}^{(2)}$ does not depend on the RI we can already see that the optimal gap interval Δt depends on the RI neither (which contrasts with the experiments reporting that optimal Δt increases with RI). Optimizing Δt yields the optimality criterion (see Appendix B.3)

$$
\frac{P_{\text{eff}}^{(1)} - P_1^{(t1)}}{P - P_1^{(t1)}} + (\alpha - 1)\frac{P_{\text{eff}}^{(1)}}{P_{\text{pot}}}x^{\alpha} - \alpha x^{\alpha-1} = 0 . \quad (32)
$$

with

$$
x := (1 - p_{d|0})^{\Delta t} = e^{\Delta t \ln(1 - p_{d|0})} \Leftrightarrow \quad \Delta t = \frac{\ln x}{\ln(1 - p_{d|0})} \quad (33)
$$

$$
\alpha := \frac{\ln(1 - p_{e|0})}{\ln(1 - p_{d|0})} , \quad (34)
$$

which can easily be evaluated using standard Newton-type numerical methods. Note that Equation (32) can be used to link neuroanatomical and neurophysiological to psychological data.

---

[3] Note that a constant (instead of decaying) "background" consolidation P_1 can be modeled, for example, by using $P_1^{(t0)} = 0$ and then excluding the initially consolidated synapses from further simulation. This means to simulate a network with anatomical connectivity $P' = P - P_1$, potential connectivity $P'_{\text{pot}} = P_{\text{pot}} - P_1$, no initial consolidation with $P'_1 = 0$, and otherwise same parameters as the original network. Then the effectual connectivity can be computed from $P_1^{(1)} = P_1^{(1)'} + P_{1S}P_1$ using Equation (18) where $P_1^{(1)'} mn$ is obtained from the simulation.
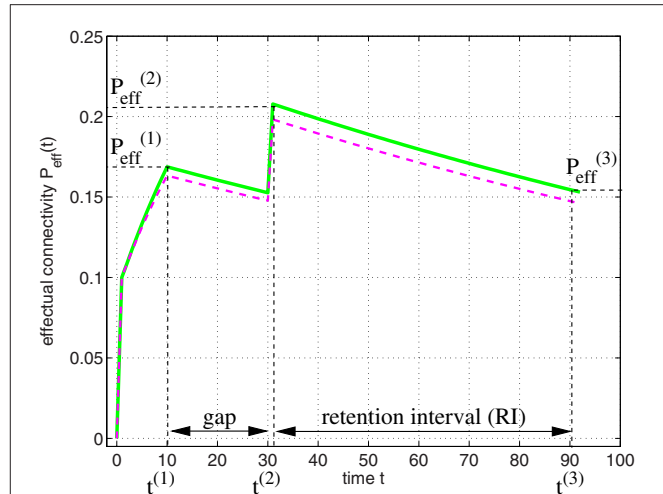
FIGURE 11 | Modeling the spacing effect experiment of Cepeda et al. (2008) **as illustrated by Figure 10.** Curves show effectual connectivity P_{eff} as function of time t according to the theory of synapse model A (green solid; **Figure 3A**; see Appendix B.1) and synapse model B (magenta dashed; **Figure 3B**; see Equations 73–75). In an initial study session, facts are learned until some desired performance level $P_{\text{eff}}^{(1)}$ is reached at time $t^{(1)} = 10$. After a gap the facts are rehearsed briefly at time $t^{(2)} = 30$ reaching a performance equivalent to $P_{\text{eff}}^{(2)}$. After the retention interval at time $t^{(3)} = 90$ performance has decreased corresponding to an effectual connectivity $P_{\text{eff}}^{(3)}$. Parameters were $P = 0.1$, $P_{\text{pot}} = 0.4$, $P_1 = 0$, $P_{1S} = 0.1$, $p_{c|s} = s$, $p_{e|0} = 0.1$, $p_{d|0} = 0.005$, and $p_{e|1} = p_{d|1} = 0$.

For example, given the optimal gap Δt_{opt} from psychological experiments, Equation (32) gives a constraint on the remaining network and learning parameters. Alternatively, we can solve Equation (32) to determine the optimal gap Δt_{opt} given the remaining parameters.

We have verified Equation (32) by simulations illustrated in **Figure 12** (compare simulation data to Cepeda et al., 2008, **Figure 3**). For these simulations we chose physiologically plausible model parameters: Similarly as before we used $P_{\text{pot}} = 0.4$ (Stepanyants et al., 2002; DePaola et al., 2006), $P = 0.1$ (Braitenberg and Schüz, 1991; Hellwig, 2000). Further, we used $P_1^{(t0)} = 0.02$ as neurophysiological experiments investigating two-state properties of synapses suggest that about 20% of synapses are in the "up" state (Petersen et al., 1998; O'Connor et al., 2005)[4] . Then we chose a small consolidation load $P_{1S} = 0.001$ assuming that the small set of novel facts is negligible compared to the presumably large set of older memories. As before, we assumed p_g in homeostatic balance to maintain a constant anatomical connectivity $P(t)$ (Equation 69) and binary consolidation signals $s = S_{ij} \in \{0, 1\}$ with $p_{c|s} = s$ and $p_{d|1} = p_{e|1} = 0$ for any synapse ij. For the remaining learning

---

[4] It may be more realistic that the total number of "up"-synapses is kept constant by homeostatic processes (i.e., $P_1/P = 0.2$). However, here we were more interested in verifying our theory which assumes exponential decay of "up"-synapses. To account for homeostasis with constant P_1 one may proceed as described in footnote 3. Nevertheless, the qualitative behavior of the model does not strongly depend on P_1 or $P_1^{(t0)}$ unless their values being close to P which would strongly impair learning.
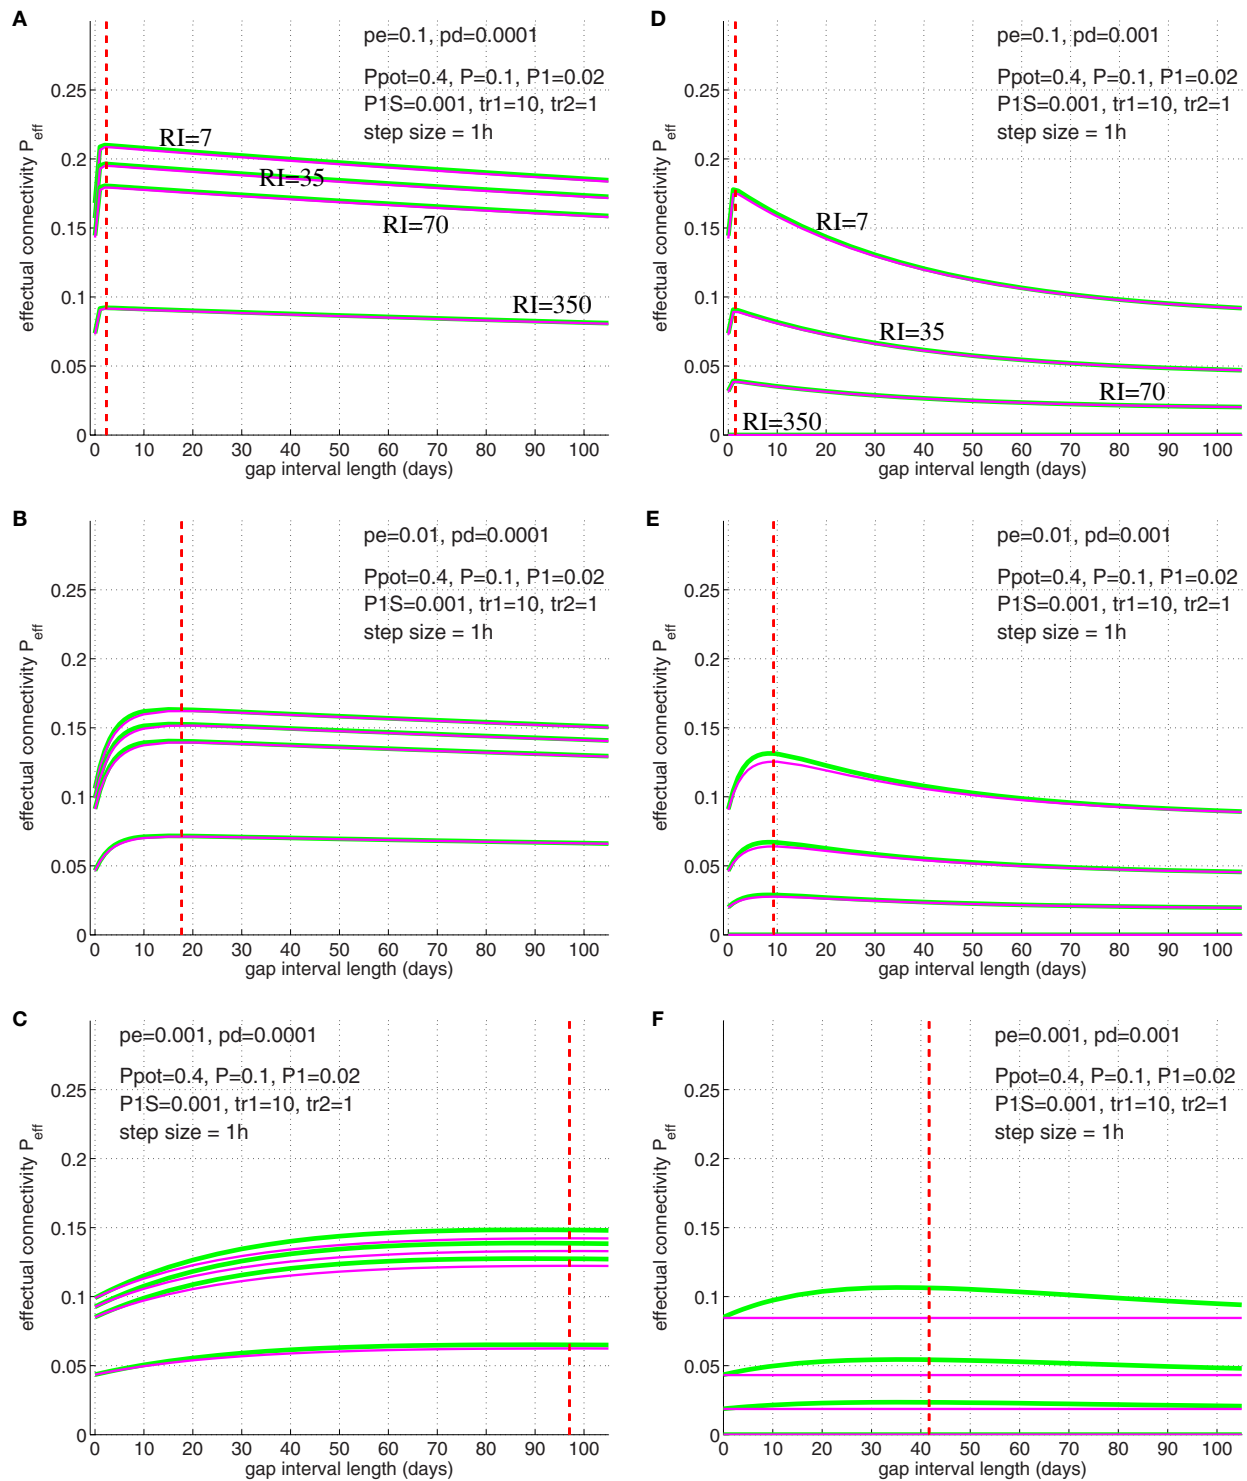
FIGURE 12 | Simulation of the spacing effect described by Cepeda et al. (2008, Figure 3) using synapse model variant A (green lines) and B (magenta lines; see Figure 3). Each curve shows final effectual connectivity $P_{eff} = P_{eff}^{(3)}$ as a function of rehearsal gap Δt for different retention intervals (RI = 7, 35, 70, 350 days) assuming an experimental setting as in illustrated in **Figures 10, 11**. Initially, memory facts were rehearsed for tr1=10 time steps (1 time step = 1 h). After the gap, memory facts were rehearsed again for a single time step (tr2=1). Finally, after RI steps the resulting effectual connectivity was tested. Red dashed lines indicate optimal gap interval length for synapse model **B** as computed from solving Equation (32). Different panels correspond to different synapse parameters $p_{e|0}$ and $p_{d|0}$: Elimination probabilities are $p_{e|0} = 0.1$ (top panels **A,D**), $p_{e|0} = 0.01$ (middle panels **B,E**), and $p_{e|0} = 0.001$ (bottom panels **C,F**). Deconsolidation probabilities are $p_{d|0} = 0.0001$ (left panels **A–C**) and $p_{d|0} = 0.001$ (right panels **D–F**). Remaining model parameters are described in the main text.

parameters $p_{e|0}$ and $p_{d|0}$ we have chosen several combinations to test their relevance for fitting the model to the observed data.

The simulation results of **Figure 12** imply the following conclusions: First, the simulations show that the optimal gap determined by Equation (32) closely matches the simulation results, for both synapse models (**Figure 3**). Second, for fixed deconsolidation $p_{d|0}$, larger $p_{e|0}$ implies smaller optimal gaps Δt_{opt}. Thus, faster synaptic turnover implies smaller optimal gaps. Third, for fixed turnover $p_{e|0}$, larger $p_{d|0}$ implies smaller Δt_{opt}. Thus, faster deconsolidation implies also smaller optimal gaps. Fourth, together this means that faster (weight and structural) plasticity implies smaller optimal gaps. Fifth, although model variants A and B (**Figure 3**) behave very similar for most parameters settings, they can differ significantly for some parameter combinations. For example, for $p_{e|0} = p_{d|0} = 0.001$ (panel F) the peak in P_{eff} of model A is more than a third larger than the peak of model B. In fact, there the curve of model B is almost flat. Still, even here, the optimal gap interval length is very similar for the two models. An obvious reason why model A sometimes performs better than model B is that deconsolidation of a synapse in model A does not necessarily imply elimination as in model B (see **Figure 3**). Sixth, our simple model already satisfies two of the three characteristics of the spacing effect mentioned above: Both the forgetting effect and the existence of an optimal time gap can be observed in a wide parameter range. Best fits to the experimental data occurred for $p_{e|0} = 0.01$ and $p_{d|0} = 0.0002$ (between parameters of panels B,C; data not shown). Last, however, our simple model cannot reproduce the third characteristic: As argued above, the optimal gap interval length Δt_{opt} does not depend on the retention interval RI. This is in contrast to the experiments of Cepeda et al. (2008) reporting that Δt_{opt} increases with RI.

Nevertheless, we have shown in some preliminary simulations that a slight extension of the model can easily resolve the latter discrepancy (Knoblauch, 2010a): By mixing two populations of synapses having different plasticity parameters corresponding to a small and large optimal gap (or fast and slow plasticity), respectively, it is possible to obtain a dependence of optimal spacing as in the experiments.

4. DISCUSSION

In this theoretical work we have identified roles of structural plasticity and effectual connectivity P_{eff} for network performance, measuring brain connectivity, and optimizing learning protocols. Analyzing how many cell assemblies or memories can be stored in a cortical macrocolumn (of size $1\,\text{mm}^3$), we find a strong dependence of storage capacity on P_{eff} and cell assembly size k (see **Figures 7**, **8**). We find that, without structural plasticity, when cell assemblies would have a connectivity close to the low anatomical connectivity $P \approx 0.1$, only a small number of relatively large cell assemblies could be stably stored (Latham and Nirenberg, 2004; Aviel et al., 2005) and, correspondingly, retrieval would not be energy efficient (Attwell and Laughlin, 2001; Laughlin and Sejnowski, 2003; Lennie, 2003; Knoblauch et al., 2010; Knoblauch, 2016). It thus appears that storing and

efficiently retrieving a large number of small cell assemblies as observed in some areas of the medial temporal lobe (Waydo et al., 2006) would require structural plasticity increasing P_{eff} from the low anatomical level toward the much larger level of potential connectivity $P_{\text{pot}} \approx 0.5$ (Stepanyants et al., 2002). Similarly, our model predicts ongoing structural plasticity for any cortical area that exhibits sparse neural activity and high capacity.

Moreover, we have shown a close relation between our definition of effectual connectivity P_{eff} and previous measures of functional brain connectivity. While the latter, for example transfer entropy, are solely based on correlations between neural activity in cortical areas (Schreiber, 2000), our definition of P_{eff} as the fraction of realized required synapses has also a clear anatomical basis (**Figure 2**). Via the link of memory channel capacity $C(P_{\text{eff}})$ used to measure storage capacity of a neural network, we have shown that P_{eff} is basically an equivalent measure of functional connectivity as transfer entropy. By this, it may become possible to establish an anatomically grounded link between structural plasticity and functional connectivity. For example, this could enable predictions on which cortical areas exhibit strong ongoing structural plasticity during certain cognitive tasks.

Further, as one example linking cognitive phenomena to its potential anatomical basis, we have more closely investigated the spacing effect that learning becomes more efficient if rehearsal is distributed to multiple sessions (Crowder, 1976; Greene, 1989; Cepeda et al., 2008). In previous works we have already shown that the spacing effect can easily be explained by structural plasticity and that, therefore, structural plasticity may be the common physiological basis of various forms of the spacing effect (Knoblauch, 2009a; Knoblauch et al., 2014). Here we have extended these results to explain some recent long-term memory experiments investigating the optimal time gap between two learning sessions (Cepeda et al., 2008). For a given retention interval, our model, if fitted to neuroanatomical data, can easily explain the profile of the psychological data, in particular, the existence of an optimal gap that maximizes memory retention. It is even possible to analyze this profile, linking the optimal gap to parameters of the synapse model, in particular, the rate of deconsolidation $p_{d|0}$ and elimination $p_{e|0}$. Our results show that small optimal gaps correspond to fast structural and weight plasticity with a high synaptic turnover rate $p_{e|0}$ and relative large $p_{d|0}$ with a high forgetting rate, whereas large gaps correspond to slow plasticity processes. This result has two implications: First, it may be used to explain the remaining discrepancy that in the psychological data the time gap depends on the retention interval, whereas in our model it does not: As preliminary simulations indicate, the experimental data could be reproduced by mixing (at least) two synapse populations with different sets of parameters, where they could be both within the same cortical area (stable vs. unstable synapses; cf., Holtmaat and Svoboda, 2009) or distributed to different areas (e.g., fast plasticity in the medial temporal lobe, and slower plasticity in neocortical areas). Moreover, as the temporal profile of optimal learning depends on parameters of structural plasticity, it may become possible in future experiments to link behavioral data

on memory performance to physiological data on structural plasticity in cortical areas where these memories are finally stored.

Although we have concentrated on analyzing one-step retrieval in feed-forward networks, our results apply as well to recurrent networks and iterative retrieval (Hopfield, 1982; Schwenker et al., 1996; Sommer and Palm, 1999): Obviously, all results on the temporal evolution of P_{eff} (including the results on the spacing effect) depend only on synapses having proper access to consolidation signals S_{ij} by either repeated rehearsal or memory replay, and therefore hold independently of network and retrieval type. However, linking P_{eff} to output noise (Equation 3) needs to assume a particular retrieval procedure. At least one-step retrieval is known to be almost equivalent for both feedforward and recurrent networks yielding almost identical output noise and pattern capacity M_ϵ (Knoblauch, 2008). Estimating retrieved information for pattern completion in auto-associative recurrent networks, however, requires to subtract the information already provided by the input patterns $\tilde{u}^\mu$. Here information storage capacity C is maximal if $\tilde{u}^\mu$ contains half of the one-entries (or information) of the original pattern u^μ, which leads to factor $1/2$ and $1/4$ decreases of M and C compared to hetero-association (cf., Equations 48, 49 for $\lambda = 1/2$; Palm and Sommer, 1992). Nevertheless, up to such scaling, our results demonstrating C increasing with P_{eff} are still valid. Similarly, our capacity analyses of M_ϵ and C_ϵ can also be applied to iterative retrieval by requiring that the one-step output noise level ϵ is smaller than the initial input noise $\tilde{\epsilon}$. As typically output noise $\hat{\epsilon}$ steeply decreases with input noise $\tilde{\epsilon}$ (cf. Equation 45), additional retrieval steps will drive $\hat{\epsilon}$ toward zero, with activity quickly converging to the memory attractor.

Our theory depends on the assumption that potential connectivity P_{pot} is significantly larger than anatomical connectivity P. This assumption may be challenged by experimental findings suggesting that cortical neuron pairs are either unconnected or have multiple (e.g., 4 or 5) instead of single synapses (Fares and Stepanyants, 2009) and the corresponding theoretical works to explain these findings (Deger et al., 2012; Fauth et al., 2015b). For example, Fauth et al. (2015a) predict that narrow distributions of synapse numbers around 4 or 5 follow from a regulatory interaction between synaptic

and structural plasticity, where connections having a smaller synapse number cannot stably exist. If true this would mean that most potential synapses could never become stable actual synapses because the majority of potentially connected neuron pairs have less than 4 potential synapses (e.g., see Fares and Stepanyants, 2009, **Figure 1**). As a consequence, actual P_{pot} would be significantly lower than assumed in our work, perhaps only slightly larger than P, strongly limiting a possible increase of effectual connectivity P_{eff} by structural plasticity. On the other hand, the data of Fares and Stepanyants (2009) are based only on neuron pairs having very low distances ($< 50\,\mu m$), whereas our model rather applies to cortical macrocolumns where most neuron pairs have much larger distances. Thus, unlike Fauth et al. (2015a), our theory of structural plasticity increasing effectual connectivity and synaptic storage efficiency predicts that neuron pairs within a macrocolumn should typically be connected by a much smaller synapse number (e.g., 1 or perhaps 2).

AUTHOR CONTRIBUTIONS

Conceived, designed, and performed experiments: AK. Analyzed the data: AK, FS. Contributed simulation/analysis tools: AK. Wrote the paper: AK, FS.

FUNDING

ACKNOWLEDGMENTS

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: http://journal.frontiersin.org/article/10.3389/fnana.2016.00063

REFERENCES

Abeles, M. (1991). *Corticonics: Neural Circuits of the Cerebral Cortex.* Cambridge: Cambridge University Press.

Arshavsky, Y. (2006). "The seven sins" of the hebbian synapse: can the hypothesis of synaptic plasticity explain long-term memory consolidation? *Progress Neurobiol.* 80, 99–113. doi: 10.1016/j.pneurobio.2006.09.004

Attwell, D., and Laughlin, S. (2001). An energy budget for signaling in the grey matter of the brain. *J. Cereb. Blood Flow Metabol.* 21, 1133–1145. doi: 10.1097/00004647-200110000-00001

Aviel, Y., Horn, D., and Abeles, M. (2005). Memory capacity of balanced networks. *Neural Comput.* 17, 691–713. doi: 10.1162/0899766053019962

Bliss, T., and Collingridge, G. (1993). A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361, 31–39.

Bosch, H., and Kurfess, F. (1998). Information storage capacity of incompletely connected associative memories. *Neural Netw.* 11, 869–876.

Braitenberg, V., and Schüz, A. (1991). *Anatomy of the Cortex. Statistics and Geometry.* Berlin: Springer-Verlag.

Buckingham, J., and Willshaw, D. (1992). Performance characteristics of the associative net. *Network* 3, 407–414.

Buckingham, J., and Willshaw, D. (1993). On setting unit thresholds in an incompletely connected associative net. *Network* 4, 441–459.

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

Caporale, N., and Dan, Y. (2008). Spike timing-dependent plasticity: a hebbian learning rule. *Ann. Rev. Neurosci.* 31, 25–46. doi: 10.1146/annurev.neuro.31.060407.125639

Cepeda, N., Vul, E., Rohrer, D., Wixted, J., and Pashler, H. (2008). Spacing effects in learning: a temporal ridgeline of optimal retention. *Psychol. Sci.* 19, 1095–1102. doi: 10.1111/j.1467-9280.2008.02209.x

Chklovskii, D., Mel, B., and Svoboda, K. (2004). Cortical rewiring and information storage. *Nature* 431, 782–788. doi: 10.1038/nature03012

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

Cover, T., and Thomas, J. (1991). *Elements of Information Theory.* New York, NY: Wiley.

Crowder, R. (1976). *Principles of Learning and Memory.* Oxford: Lawrence Erlbaum.

Deger, M., Helias, M., Rotter, S., and Diesmann, M. (2012). Spike-timing dependence of structural plasticity explains cooperative synapse formation in the neocortex. *PLoS Comput. Biol.* 8:e1002689. doi: 10.1371/journal.pcbi.1002689

DePaola, V., Holtmaat, A., Knott, G., Song, S., Wilbrecht, L., Caroni, P., and Svoboda, K. (2006). Cell type-specific structural plasticity of axonal branches and boutons in the adult neocortex. *Neuron* 49, 861–875. doi: 10.1016/j.neuron.2006.02.017

Dobrushin, R. (1959). General formulation of shannon's main theorem in information theory. *Ushepi Mat. Nauk.* 14, 3–104.

Ebbinghaus, H. (1885). *Über das Gedächtnis: Untersuchungen zur Experimentellen Psychologie.* Leipzig: Duncker & Humblot.

Engert, F., and Bonhoeffer, T. (1999). Dendritic spine changes associated with hippocampal long-term synaptic plasticity. *Nature* 399, 66–70.

Fares, T., and Stepanyants, A. (2009). Cooperative synapse formation in the neocortex. *Proc. Natl. Acad. Sci. U.S.A.* 106, 16463–16468. doi: 10.1073/pnas.0813265106

Fauth, M., Wörgötter, F., and Tetzlaff, C. (2015a). Formation and maintenance of robust long-term information storage in the presence of synaptic turnover. *PLoS Comput. Biol.* 11:e1004684. doi: 10.1371/journal.pcbi.1004684

Fauth, M., Wörgötter, F., and Tetzlaff, C. (2015b). The formation of multi-synaptic connections by the interaction of synaptic and structural plasticity and their functional consequences. *PLoS Comput. Biol.* 11:e1004031. doi: 10.1371/journal.pcbi.1004031

Frey, U., and Morris, R. (1997). Synaptic tagging and long-term potentiation. *Nature* 385, 533–536.

Friston, K. (1994). Functional and effective connectivity. *Hum. Brain Mapp.* 2, 56–78.

Fu, M., and Zuo, Y. (2011). Experience-dependent structural plasticity in the cortex. *Trends Neurosci.* 34, 177–187. doi: 10.1016/j.tins.2011.02.001

Greene, R. (1989). Spacing effects in memory: evidence for a two-process account. *J. Exp. Psychol.* 15, 371–377.

Hebb, D. (1949). *The Organization of Behavior. A Neuropsychological Theory.* New York, NY: Wiley.

Hellwig, B. (2000). A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biol. Cybernet.* 82, 111–121. doi: 10.1007/PL00007964

Holtmaat, A., and Svoboda, K. (2009). Experience-dependent structural synaptic plasticity in the mammalian brain. *Nat. Rev. Neurosci.* 10, 647–658. doi: 10.1038/nrn2699

Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.* 79, 2554–2558.

Knoblauch, A. (2008). Neural associative memory and the Willshaw-Palm probability distribution. *SIAM J. Appl. Mathemat.* 69, 169–196. doi: 10.1137/070700012

Knoblauch, A. (2009a). "The role of structural plasticity and synaptic consolidation for memory and amnesia in a model of cortico-hippocampal interplay," in *Connectionist Models of Behavior and Cognition II: Proceedings of the 11th Neural Computation and Psychology Workshop*, eds J. Mayor, N. Ruh and K. Plunkett (Singapore: World Scientific Publishing), 79–90.

Knoblauch, A. (2009b). "Zip nets: Neural associative networks with non-linear learning," in *HRI-EU Report 09-03, Honda Research Institute Europe GmbH, D-63073* (Offenbach/Main).

Knoblauch, A. (2010a). Bimodal structural plasticity can explain the spacing effect in long-term memory tasks. *Front. Neurosci. Conference*

Abstract: Computational and Systems Neuroscience. 2010, doi: 10.3389/conf.fnins.2010.03.00227

Knoblauch, A. (2010b). "Efficient associative computation with binary or low precision synapses and structural plasticity," in *Proceedings of the 14th International Conference on Cognitive and Neural Systems (ICCNS)* (Boston, MA: Center of Excellence for Learning in Education, Science, and Technology (CELEST)), 66.

Knoblauch, A. (2010c). "Structural plasticity and the spacing effect in willshaw-type neural associative networks," in *HRI-EU Report 10-10, Honda Research Institute Europe GmbH, D-63073* (Offenbach/Main).

Knoblauch, A. (2011). Neural associative memory with optimal bayesian learning. *Neural Comput.* 23, 1393–1451. doi: 10.1162/NECO_a_00127

Knoblauch, A. (2016). Efficient associative computation with discrete synapses. *Neural Comput.* 28, 118–186. doi: 10.1162/NECO_a_00795

Knoblauch, A., Hauser, F., Gewaltig, M.-O., Körner, E., and Palm, G. (2012). Does spike-timing-dependent synaptic plasticity couple or decouple neurons firing in synchrony? *Front. Comput. Neurosci.* 6:55. doi: 10.3389/fncom.2012.00055

Knoblauch, A., Körner, E., Körner, U., and Sommer, F. (2014). Structural plasticity has high memory capacity and can explain graded amnesia, catastrophic forgetting, and the spacing effect. *PLoS ONE* 9:e96485. doi: 10.1371/journal.pone.0096485

Knoblauch, A., and Palm, G. (2001). Pattern separation and synchronization in spiking associative memories and visual areas. *Neural Netw.* 14, 763–780. doi: 10.1016/S0893-6080(01)00084-3

Knoblauch, A., Palm, G., and Sommer, F. (2010). Memory capacities for synaptic and structural plasticity. *Neural Comput.* 22, 289–341. doi: 10.1162/neco.2009.08-07-588

Kononenko, I. (1989). Bayesian neural networks. *Biol. Cybernet.* 61, 361–370.

Krone, G., Mallot, H., Palm, G., and Schüz, A. (1986). Spatiotemporal receptive fields: a dynamical model derived from cortical architectonics. *Proc. R. Soc. Lond. B* 226, 421–444.

Lansner, A., and Ekeberg, O. (1989). A one-layer feedback artificial neural network with a Bayesian learning rule. *Intern. J. Neural Syst.* 1, 77–87.

Latham, P., and Nirenberg, S. (2004). Computing and stability in cortical networks. *Neural Comput.* 16, 1385–1412. doi: 10.1162/089976604323057434

Laughlin, S., and Sejnowski, T. (2003). Communication in neuronal networks. *Science* 301, 1870–1874. doi: 10.1126/science.1089662

Lennie, P. (2003). The cost of cortical computation. *Curr. Biol.* 13, 493–497. doi: 10.1016/S0960-9822(03)00135-0

Marr, D. (1971). Simple memory: a theory for archicortex. *Philos. Trans. R. Soc. Lond. Ser. B* 262, 24–81.

Martin, S., Grimwood, P., and Morris, R. (2000). Synaptic plasticity and memory: an evaluation of the hypothesis. *Ann. Rev. Neurosci.* 23, 649–711. doi: 10.1146/annurev.neuro.23.1.649

O'Connor, D., Wittenberg, G., and Wang, S.-H. (2005). Graded bidirectional synaptic plasticity is composed of switch-like unitary events. *Proc. Natl. Acad. Sci. U.S.A.* 102, 9679–9684. doi: 10.1073/pnas.0502332102

Palm, G. (1980). On associative memories. *Biol. Cybernet.* 36, 19–31.

Palm, G., Knoblauch, A., Hauser, F., and Schüz, A. (2014). Cell assemblies in the cerebral cortex. *Biol. Cybernet.* 108, 559–572. doi: 10.1007/s00422-014-0596-4

Palm, G., and Sommer, F. (1992). Information capacity in recurrent McCulloch-Pitts networks with sparsely coded memory states. *Network* 3, 177–186.

Paulsen, O., and Sejnowski, T. (2000). Natural patterns of activity and long-term synaptic plasticity. *Curr. Opin. Neurobiol.* 10, 172–179. doi: 10.1016/S0959-4388(00)00076-3

Petersen, C., Malenka, R., Nicoll, R., and Hopfield, J. (1998). All-or-none potentiation at CA3-CA1 synapses. *Proc. Natl. Acad. Sci. U.S.A.* 95, 4732–4737.

Poirazi, P., and Mel, B. (2001). Impact of active dendrites and structural plasticity on the memory capacity of neural tissue. *Neuron* 29, 779–796. doi: 10.1016/S0896-6273(01)00252-5

Raisman, G. (1969). Neuronal plasticity in the septal nuclei of the adult rat. *Brain Res.* 14, 25–48.

Schreiber, T. (2000). Measuring information transfer. *Phys. Rev. Lett.* 85, 461–464. doi: 10.1103/PhysRevLett.85.461

Schwenker, F., Sommer, F., and Palm, G. (1996). Iterative retrieval of sparsely coded associative memory patterns. *Neural Netw.* 9, 445–455.

Shannon, C., and Weaver, W. (1949). *The Mathematical Theory of Communication.* Urbana/Chicago: University of Illinois Press.

Sommer, F. (2000). On cell assemblies in a cortical column. *Neurocomputing* 32–33, 517–522. doi: 10.1016/S0925-2312(00)00207-1

Sommer, F., and Palm, G. (1999). Improved bidirectional retrieval of sparse patterns stored by Hebbian learning. *Neural Netw.* 12, 281–297.

Song, S., Miller, K., and Abbott, L. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neurosci.* 3, 919–926. doi: 10.1038/78829

Sporns, O. (2007). Brain connectivity. *Scholarpedia* 2:4695. doi: 10.4249/scholarpedia.4695

Steinbuch, K. (1961). Die Lernmatrix. *Kybernetik* 1, 36–45.

Stepanyants, A., Hof, P., and Chklovskii, D. (2002). Geometry and structural plasticity of synaptic connectivity. *Neuron* 34, 275–288. doi: 10.1016/S0896-6273(02)00652-9

Waydo, S., Kraskov, A., Quiroga, R., Fried, I., and Koch, C. (2006). Sparse representation in the human medial temporal lobe. *J. Neurosci.* 26, 10232–10234. doi: 10.1523/JNEUROSCI.2101-06.2006

Willshaw, D., Buneman, O., and Longuet-Higgins, H. (1969). Non-holographic associative memory. *Nature* 222, 960–962.

Witte, S., Stier, H., and Cline, H. (1996). *In vivo* observations of timecourse and distribution of morphological dynamics in Xenopus retinotectal axon arbors. *J. Neurobiol.* 31, 219–234.

Wyner, A. (1978). A definition of conditional mutual information for arbitrary ensembles. *Inform. Control* 38, 51–59.

Xu, T., Yu, X., Perlik, A., Tobin, W., Zweig, J., Tennant, K., Jones, T., and Zuo, Y. (2009). Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature* 462, 915–919. doi: 10.1038/nature08389

Yang, G., Pan, F., and Gan, W.-B. (2009). Stably maintained dendritic spines are associated with lifelong memories. *Nature* 462, 920–924. doi: 10.1038/nature08577

Yu, X., and Zuo, Y. (2011). Spine plasticity in the motor cortex. *Curr. Opin. Neurobiol.* 21, 169–174. doi: 10.1016/j.conb.2010.07.010

Anatomically Detailed and Large-Scale Simulations Studying Synapse Loss and Synchrony Using NeuroBox

Markus Breit[1], Martin Stepniewski[1], Stephan Grein[1,2], Pascal Gottmann[1], Lukas Reinhardt[1] and Gillian Queisser[1,2]*

[1] Computational Neuroscience, Department for Computer Science and Mathematics, Goethe Center for Scientific Computing, Goethe University, Frankfurt am Main, Germany, [2] Department of Mathematics, Temple University, Philadelphia, PA, USA

The morphology of neurons and networks plays an important role in processing electrical and biochemical signals. Based on neuronal reconstructions, which are becoming abundantly available through databases such as *NeuroMorpho.org*, numerical simulations of Hodgkin-Huxley-type equations, coupled to biochemical models, can be performed in order to systematically investigate the influence of cellular morphology and the connectivity pattern in networks on the underlying function. Development in the area of synthetic neural network generation and morphology reconstruction from microscopy data has brought forth the software tool NeuGen. Coupling this morphology data (either from databases, synthetic, or reconstruction) to the simulation platform UG 4 (which harbors a neuroscientific portfolio) and VRL-Studio, has brought forth the extendible toolbox NeuroBox. NeuroBox allows users to perform numerical simulations on hybrid-dimensional morphology representations. The code basis is designed in a modular way, such that e.g., new channel or synapse types can be added to the library. Workflows can be specified through scripts or through the VRL-Studio graphical workflow representation. Third-party tools, such as ImageJ, can be added to NeuroBox workflows. In this paper, NeuroBox is used to study the electrical and biochemical effects of synapse loss vs. synchrony in neurons, to investigate large morphology data sets within detailed biophysical simulations, and used to demonstrate the capability of utilizing high-performance computing infrastructure for large scale network simulations. Using new synapse distribution methods and Finite Volume based numerical solvers for compartment-type models, our results demonstrate how an increase in synaptic synchronization can compensate synapse loss at the electrical and calcium level, and how detailed neuronal morphology can be integrated in large-scale network simulations.

Keywords: HPC, large-scale neuronal networks, synaptic plasticity, electrical scale, anatomy, reconstruction, simulation, cable equation

1. INTRODUCTION

The structure of neurons and networks in the brain is known to change continuously over time. Cellular growth, synapse formation or synapse loss, reorganization of intracellular architecture constantly make changes to the overall cellular and network anatomy (Hughes, 1958; Abbott and Nelson, 2000; Sheng and Hoogenraad, 2007; Shepherd and Huganir, 2007; Tai et al., 2008; Colon-Ramos, 2009; Branco et al., 2010; Zeltser et al., 2012; Tyagarajan and Fritschy, 2014). These changes in geometric layout can be interpreted as a strong indicator that the anatomy of the (sub)cellular and network level is deeply involved on various functional levels. Neuroscientific research has always been devoted to the interplay between morphology and function on various functional levels. Experimental research draws from microscopy techniques that can make morphology and spatio-temporal signals visible (Spacek and Harris, 1997; Arellano et al., 2007; Chen et al., 2008), theoretical work in Computational Neuroscience has brought forth an abundant spread of cellular and network models, many of them rely on a spatial representation of neurons and networks (Bower and Beeman, 1997; Hines and Carnevale, 1997; Balls et al., 2004; Gewaltig and Diesmann, 2007; Andrews et al., 2010). General purpose simulators such as NEURON or Genesis couple electrical and biochemical models to graph-representations of neurons and synaptically connected networks. The importance of neuronal morphology used in such simulations can be seen in reconstruction projects, such as the database project *NeuroMorpho* (cf. Ascoli, 2006). Currently more than 30,000 cell reconstructions are freely available on this platform.

Reconstructing morphology from microscopy data is a further example of how deeply structure is integrated in the brain. Semi-manual or fully automated reconstruction methods are being developed in research groups around the world (e.g., Jungblut et al., 2011; Popov et al., 2011; Burette et al., 2012), trying to unravel the filigreed multi-level organization of the brain. This dedication has advanced the field significantly, still many of the anatomical questions are currently unresolved. To leverage the power of large-scale network simulations, synthetic neuron morphology tools have been developed (Wolf et al., 2013). These algorithms are capable of generating synthetic networks with realistic morphology statistics which can be used within detailed functional simulations. In order to use these large data sets in detailed and large network simulations high performance computing platforms become an inevitable component of the process. While most of the available network simulators were originally conceived to run serially, there has been effort to parallelize and optimize the code for ever growing computing power.

In this paper, we present an approach focusing on the topic of cellular and network anatomy within a large-scale computing context. Building on scalable numerical methods in a flexible and parallelized discretization and solver framework for general ordinary and partial differential equation systems, this unified approach does not make use of the NEURON simulation environment (Hines and Carnevale, 1997) used in similar projects (Markram, 2015; Ramaswamy, 2015; Reimann, 2015). We introduce some of the authors' contributions in morphology reconstruction as well as artificial construction, hybrid-dimensional modeling and simulation of coupled biochemical and electrical signals, and link these to newly developed algorithms for massively parallel simulation of cable equation models and synapse distribution on cells. The latter can be used to simulate healthy and disease state neurons with different synapse numbers and distributions.

The Materials and Methods section of this paper discusses the tool NeuGen (Eberhard et al., 2006; Wolf et al., 2013) and how it ties into a generalized simulation framework. Our model for simulating electrical signals builds upon the known cable theory and is briefly summarized. We introduce our methods for handling synapse types and synapse distributions and introduce a new way of numerically discretizing the resulting model equations and computational domains, ultimately resulting in a system that can be solved on massively parallel computing architectures. These methods are compiled in the toolbox NeuroBox which is developed on top of the numerics engine UG 4 (cf. Vogel et al., 2013) that has been used in several detailed studies of structure-function interplay (Xylouris et al., 2007; Hansen et al., 2008; Nägel et al., 2008, 2009; Wittmann et al., 2009; Grillo et al., 2010; Muha et al., 2011).

To study this anatomy-high-performance framework we present a study of synapse loss vs. signal synchronicity and the influence on somatic calcium signals as well as simulations of large and detailed network simulations (10,000 neurons, each neuron containing 574–586 degree of freedom) of a neocortical column synthetically generated with NeuGen. In these studies we show that synapse loss, which is a major factor in neurodegenerative diseases, can be partially compensated by an increase in synaptic synchronicity, while somatic calcium signals rely strongly on the activation and frequency of action potentials. We further show that wave activation in neocortical networks is clearly driven by synapse density and that our employed simulation framework scales well on JUQUEEN, one of the high-performance computers at the German Jülich Supercomputing Center. This in turn demonstrates that large-scale network simulations do not necessarily have to come at the cost of anatomy anymore.

2. MATERIALS AND METHODS

In this section we will introduce the tools and methods used for the simulations performed in Section 3. A combination of neuron and network generating tools (Section 2.1), synapse distribution algorithms, a new approach for numerical discretization of the network topology and a parallel computing framework (Section 2.2) forms the basis of our detailed anatomical and large-scale network simulations and is integrated in a new and extendible simulation toolbox, NeuroBox (Section 2.3).

2.1. Generating Large and Anatomically Detailed Networks

The generation of large neural networks (containing more than 10,000 neurons) is accomplished with the neural network

generator NeuGen (Eberhard et al., 2006). NeuGen uses anatomical fingerprints, i.e., experimental morphology data and standard deviations to generate anatomically consistent neurons that fit experimental mean and standard deviation. NeuGen thus generates non-identical neurons of various types—e.g., pyramidal cells and spiny stellate cells of the neocortex and hippocampus—and synaptically connects these to form neural networks. The topology of the network is described in terms of graph theory as an undirected, connected graph containing edges and vertices in three-dimensional coordinate space. NeuGen algorithms sample parameter values from experimental data distributions and incorporates two categories of synapses: *Primary* synapses representing external stimulation of the network; as well as *interconnecting* synapses which represent chemical synapses between neurons present in the network, typically formed by a presynaptic axon and a postsynaptic dendrite. The anatomy of the network can be exported to a 3D graphics format for visualization and various discrete morphology file formats that can be used in simulators such as NEURON (Hines and Carnevale, 1997) or UG 4 (Vogel et al., 2013). NeuGen is intended to provide anatomically accurate large network topologies for general purpose neuron network simulators.

The algorithm, which is not a growth-based algorithm, is summarized by the following steps (cf. **Figure 1**):

– Generate sections for each neuron based on anatomical fingerprints

– Interconnect sections of individual neurons
– Generate synapses based on a distance criterion and attach functional parameters

It is worth highlighting two parameters when discussing anatomical detail. To regulate the number of vertices for each neuron (which represents the level of detail at which neuron morphology is represented), one may adjust a parameter termed `section_length`, the average compartment length in μm. In cases where memory consumption is a constraint, choosing an increased section length permits the creation of and simulation on larger networks (with less anatomical detail) using the same amount of memory. Secondly, the number of synapses inserted into the network may be adjusted by a global threshold parameter termed `dist_synapse`. If and only if the euclidean distance between two sections falls below the threshold specified by this parameter, these sections will be marked as potential synaptic contact points. Whether or not a synapse will be placed in the network depends on the type of pre- and postsynaptic neurons. A connectivity matrix specifies which classes of neurons are interconnected by synapses (Wolf et al., 2013).

Subsequent simulations need to refer to the compartments contained in the grid for simulation control setup. Therefore, an alphanumerical identifier is stored within the grid too. The identifier is a string and composed out of the cell type (e.g., pyramidal or stellate) and the compartment type (e.g., axon or dendrite) and groups all edges and vertices belonging to a given cell and compartment type (cf. **Figure 2**). If desired one

FIGURE 1 | Flowchart sketching the streamlined pipeline for the generation and subsequent transformation of neuronal network morphologies to grids suitable for large-scale network simulations. In case of three-dimensional simulations, where one-dimensional point/line reconstructions are used to generate three-dimensional representations (see Grein et al., 2014) quality assessment of the generated grid can be performed in a semi-automatic way to allow for the best possible preparation for the subsequent numerical simulations, for instance, we check for intersecting dendrites introduced during neuron tracing.

FIGURE 2 | An exemplary neocortical network consisting of a total of ten individual cells synthetically generated with NeuGen. The cellular composition of the displayed network is identical to the networks described in **Table 1** above. Visualization comprises the different cell types distinguishable by color: axons of L4 spiny stellate (lavender), axons of L2/3 pyramidal (red), axons of L5A pyramidal (bright green), axons of L5B pyramidal (blue), dendrites of L4 spiny stellate (yellow), dendrites of L2/3 pyramidal (cyan), dendrites of L5A pyramidal (rose), dendrites of L5B pyramidal (light green), somata of L4 spiny stellate (ocher), somata of L2/3 pyramidal (orange), somata of L5A pyramidal (light blue), and L5B pyramidal (brown).

TABLE 1 | Network creation statistics sorted by size, i.e., by number of contained cells within the network, in ascending order.

Vertices	Sections	Cells	Elapsed time [s]	Grid size [mb]
1403	418	12	0.01	0.14
15535	4382	120	0.02	1.60
156596	43892	1200	0.50	16.9
1644260	465262	12000	44.7	65.6
2840213	1212108	120000	300	221.1

The networks are composed of L5A and L5B pyramidal cells ($\approx$ 16% each), of L4 spiny stellate cells ($\approx$ 42%) as well as L2/3 pyramidal neurons ($\approx$ 26%). The smallest network contains 12 and the largest network 120,000 cells in total. To create even larger networks with the same memory resources, one can decrease the number of compartments using the `section_length` *parameter.*

can request the identifier to group edges and vertices based also on the section number of the compartment resulting in a fine-grained access of the network (not shown).

The network can be exported to a variety of formats including a format suitable for large neural network grid generation, e.g., a custom sparse data format based on a file format derived from TXT (plain text or compressed plain text) or a more convenient XML-based file format.

To use NeuGen in conjunction with the simulation framework UG 4 (Vogel et al., 2013, cf. Section 3), the exported morphology is exported to the UG 4 geometry format UGX (an xml-based file format). To that end, topology information of the exported network, consisting of the raw nodes and vertices, is enriched by grid attachments such as diameter information and synapses, together with their parametrizations. This procedure is implemented as a plugin for UG 4 and produces large neural networks ($\geq$ 10, 000 neurons) in the matter of seconds (cf. **Table 1**).

In addition to directly writing UGX-files from NeuGen, it is possible to convert the following formats to UGX: SWC (commonly used in the *NeuroMorpho.org* database, Ascoli et al., 2007), HOC (widespread format utilized by NEURON, Hines and Carnevale, 1997), TXT and NeuroML. The last three file formats can be exported directly by NeuGen. NeuGen and the corresponding UG 4-plugins thus form an efficient pipeline for integrating large and anatomically realistic neural networks and publicly accessible anatomical neuron reconstructions into neuron and network simulation frameworks.

2.2. Simulating Electrical and Biochemical Signals

Having established methods for generating network topologies in the previous section, we now focus on the steps from modeling

electrical signals, handling membrane transport mechanisms, including synapses to discretizing the model equations by means of a new approach via finite volumes. Lastly we summarize parallel methods for efficiently solving large-scale networks.

2.2.1. Model Equations for Membrane Potential and Ion Species

We follow the well established cable theory (cf. Thompson, 1854; Scott, 1975) to model electrical signals on spatially resolved neuron morphologies. A neuron's morphology is given as a graph consisting of vertices in a three-dimensional space and edges connecting them. Common file formats for neuronal morphologies (such as SWC or HOC) contain radius or diameter values assigned to each vertex. We make use of this diameter in the most simplistic way, i.e., by supposing the morphology to be piecewise tubular, each piece being located around a vertex and with the radius associated to this vertex. With only very few modifications, we also implemented compartments shaped like truncated cones resulting in a continuous radius along the neurites, however, we restrict ourselves to the case of tubular compartments in the following description for the sake of simplicity. In each of the compartments, we impose the following equation expressing the membrane's role as an ideal capacitor:

$$C_m \frac{\partial V}{\partial t} = I_{ax} + I_m, \tag{1}$$

where V is the membrane potential, C_m is the capacitance of the compartment and I_{ax}, I_m are the axial and transmembrane (inward) electric currents, respectively. The compartment's capacitance C_m depends on its shape and can be expressed in terms of a membrane-specific constant c_m,

$$C_m = c_m \cdot 2\pi\, al,$$

where a and l are the radius and the axial length of the compartment, respectively.

Axial currents need to be calculated at both ends of a compartment, at the interface with the neighboring compartments. They are assumed to be purely ohmic in

nature and are expressed in terms of voltage between the two vertices associated with the neighboring compartments:

$$I_{ax}^{x_2 \to x_1} = \frac{V(x_2) - V(x_1)}{\frac{r_c}{\pi} \int_{x_1}^{x_2} (a(x))^{-2}\, dx} = \frac{V(x_2) - V(x_1)}{\frac{r_c}{2\pi} \left(a_1^{-2} + a_2^{-2}\right) |x_2 - x_1|},$$

where r_c is a material constant, the specific resistance of the cytosol, x is the axial coordinate, and x_1, x_2 as well as a_1, a_2 are depicted in **Figure 3**. Note that the former equation implicitly assumes that the extracellular potential is constant in space.

Finally, the transmembrane current I_m into the compartment is expressed in terms of electrical flux density i_m as

$$I_m = i_m \cdot 2\pi a l \qquad (2)$$

and depends on transport mechanisms (e.g., Hodgkin-Huxley-type channels, Na/K pumps, leakage), synapses and electrodes definable on the membrane.

In order to track individual ion species, concentrations for K^+, Na^+, and Ca^{2+} or any other ion type can be added to the model. Each of the species satisfies a diffusion-convection equation in axial direction and is coupled to transport mechanisms in the plasma membrane.

Note that as these ions are charged, they are affected by potential gradients in reality—and conversely, for the same reason, their concentrations directly affect the potential. A physically more accurate model of ionic movement in neurons incorporating both electric and diffusive properties of individual ion species is electro-diffusion. It has ben demonstrated that the modeling error introduced by using the cable equation can be prominent in thin compartments (Qian and Sejnowski, 1989) or where three-dimensional structural detail is concerned (Lopreore et al., 2008).

2.2.2. Membrane Transport Mechanisms

What is truly at the heart of most neuronal simulations is transport across membranes. We have defined an interface allowing the addition of arbitrary transport mechanisms to the electrical model in the transmembrane current density term i_m of Equation (2). These transport mechanisms are granted access to the underlying grid as well as to the unknowns of the voltage and ion species equations. Thus, they are able to declare and calculate

their own sets of states, which may depend on given ones and vary in space and in time—like the gating parameters m, n and h in classical Hodgkin-Huxley-type channels governed by ordinary differential equations in time which depend on the membrane potential (Hodgkin and Huxley, 1952). As the dependence of inner states of membrane transport systems on the potential and on ion concentrations is typically strongly non-linear, we have decided (in the interest of fast computation) to include transmembrane currents only by an explicit scheme, i.e., inner states are updated before any time step of the solution process using only the solution from the previous time step.

The concept is not unlike the NMODL model description language for NEURON by Hines and Carnevale (1997, 2000). In fact, we have developed an automated translation unit that can convert existing NMODL files to C++ source code compilable in our framework.

2.2.3. Synapses

Glutamate being the primary excitatory neurotransmitter in most synapses of the central nervous system, we define excitatory synaptic input localized at dendrites as the postsynaptic response of AMPA or NMDA receptors to presynaptic glutamate signals. AMPA and NMDA receptors, cation channels that become permeable in glutamate-bound state and thereby exhibit a conductance change in direct response to incoming presynaptic spikes, induce transmembrane flux of sodium, potassium and calcium ions causing a local excitatory depolarization of the membrane potential.

In our simulations we distinguish two general categories of synapses: *Primary* synapses connected to dendrites as the postsynaptic side,—they are used to initialize activity in single cells as well as networks and represent connections to other neurons not included in the simulation. The second category are synapses connecting dendrites and axons both present within a network morphology. We call these *interconnecting* synapses.

As there is no information on the presynaptic side of primary synapses, the common and simple approach of *alpha functions* provides a reasonable approximation to model postsynaptic conductance profiles (Roth and van Rossum, 2009):

$$g(t) = g_{max} \frac{t - t_{onset}}{\tau} \exp\left(-\frac{t - t_{onset} - \tau}{\tau}\right), \qquad (3)$$
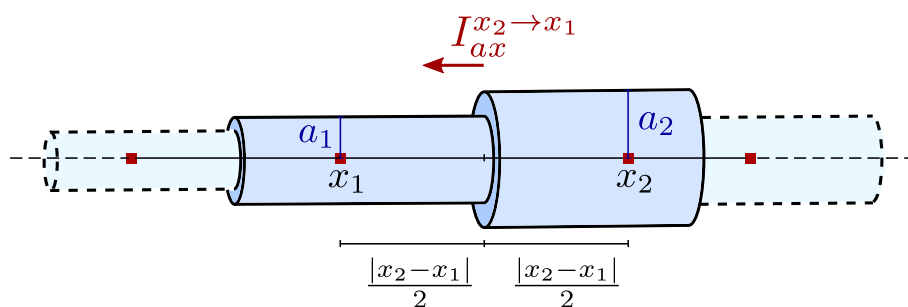
FIGURE 3 | Illustration of the piece-wise tubular compartments of the Finite Volume cable equation model and the definition of axial ohmic flux.

where g_{max} denotes the maximal conductance, τ the rise/decay and t_{onset} the arrival time of a single presynaptic spike. Note that g_{max} occurs at $t = t_{onset} + \tau$. The synaptic current $I_{ps}(t)$ is then defined by

$$I_{ps}(t) = g(t)(V(t) - E_{rev}), \tag{4}$$

with $g(t)$ given by (3) for $t_{onset} \leq t \leq t_{onset} + 6\tau$ and $g(t) = 0$ otherwise. $V(t)$ denotes the current postsynaptic membrane potential and E_{rev} a reversal potential. For glutamatergic synapses, we use $E_{rev} \approx 0\,mV$ (Purves et al., 2001).

Interconnecting synapses are activated upon rise of the presynaptic membrane potential above a threshold V_{th} and the following current $I_{is}(t)$ to the postsynaptic end is modeled according to a bi-exponential activity function:

$$t_{max} = \frac{\tau_1 \tau_2}{\tau_2 - \tau_1} \log\left(\frac{\tau_2}{\tau_1}\right), \tag{5}$$

$$n = \left(\exp\left(-\frac{t_{max}}{\tau_2}\right) - \exp\left(-\frac{t_{max}}{\tau_1}\right)\right)^{-1}, \tag{6}$$

$$I_{is}(t) = g_{max}(V - E_{rev})\, n\left(\exp\left(-\frac{t}{\tau_2}\right) - \exp\left(-\frac{t}{\tau_1}\right)\right), \tag{7}$$

where g_{max} is the maximal conductance; E_{rev} is a reversal potential; τ_1 and τ_2 are constants regulating rise and decay time of the conductance; t_{max} designates the point in time (after initial activation) at which the conductance is maximal, and the factor n normalizes the conductance such that its value is g_{max} at t_{max}.

Synaptic currents—like all other trans-membrane currents—are evaluated using the solution for the potential of the previous time step only. This has significant benefits in parallel computation, as there is no direct coupling of solutions for the next time step between cells connected to one another by synapses.

2.2.4. Activation Patterns of Primary Synapses

Our implementation provides a method to set generic activation patterns for a given set of input synapses in the computational domain. To achieve that, we introduce the continuous random variables X_{onset} and X_τ for the timing parameters t_{onset} and τ [cf. Equation (3)], respectively. Both of which we assume to be normally distributed, i.e., $X_{onset} \sim \mathcal{N}(\mu_{onset}, \sigma_{onset}^2)$ and $X_\tau \sim \mathcal{N}(\mu_\tau, \sigma_\tau^2)$ with probability density functions given by:

$$f_{\mathcal{N}}(x_\xi, \mu_\xi, \sigma_\xi^2) = \frac{1}{\sigma_\xi \sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x_\xi - \mu_\xi}{\sigma_\xi}\right)^2}, \quad \xi \in \{onset, \tau\} \tag{8}$$

After specification of a peak conductance g_{max}, a mean onset time μ_{onset} and duration μ_τ of synaptic activity as well as corresponding standard deviation values σ_{onset} and σ_τ, the parameters t_{onset} and τ are set to random values drawn from the above normal distributions $\mathcal{N}(\mu_{onset}, \sigma_{onset}^2)$ and $\mathcal{N}(\mu_\tau, \sigma_\tau^2)$, respectively.

2.2.5. Spatial Distribution of Primary Synapses

Given neuron morphologies (defined as graphs in three-dimensional coordinate space), we attach all information parameterizing synapses to the dendritic edges they are associated to. The distribution is managed by the C++ class `SynapseDistributor`. It provides methods to create new or delete existing ones to user-specified statistical distributions.

In our studies, we assume a uniform distribution of $n_{syn} \in \mathbb{N}$ synapses on the edge sample space

$$S_{edge} := \{e_i \mid i = 1, ..., n_{edge}\} \tag{9}$$

of the basal and apical dendrites. For this purpose, we consider a discrete random variable X_{syn}^i, $i \in \{1, ..., n_{syn}\}$, for the i-th synapse to be distributed. With every draw X_{syn}^i can thereby take one of the edge indices $j \in \{1, ..., n_{edge}\}$ as value, i.e., $X_{syn}^i = x_j^i := j$. To account for the heterogenous edge lengths every edge index is assigned an associated probability given by the following *probability mass function*:

$$P(X_{syn}^i = x_j^i) = p_j^i := \frac{\|e_j\|_2}{\sum_{k=1}^{n_{edge}} \|e_k\|_2} \tag{10}$$

The exact location of the i-th synapse x_j^i on the j-th edge is then drawn from a continuous uniform distribution in the range $(0, 1)$.

2.2.6. Discretization and Solution

We use a first-order (vertex-centered) Finite Volume (FV) scheme. This type of discretization method is well-suited for any type of problem resulting from a conservation law. In a FV scheme, one typically has a conservation formulation like the following:

$$\frac{\partial \rho}{\partial t} = -\text{div}\,\vec{j} \qquad \text{on the domain } \Omega, \tag{11}$$

where ρ is the density of a conserved quantity, $\vec{j}$ is a flux density of the same quantity. In our case, ρ represents the charge density for the voltage equation and the ionic concentration for the species equations; the flux densities are given by the electric current density and the ionic flux density, respectively. The conservation equation is then transformed into a system of ordinary—i.e., non-differential—equations by partitioning the domain on which the equation holds into so-called control volumes (in our case, those are exactly the compartments as defined above),

$$\Omega = \bigcup_i \Omega_i,$$

then by integrating Equation (11) on each control volume (thus ensuring local conservation)

$$\int_{\Omega_i} \frac{\partial \rho}{\partial t}\, dx = \int_{\Omega_i} -\text{div}\,\vec{j}\, dx \left(= -\int_{\partial \Omega_i} \vec{j} \cdot \vec{n}_i\, dS\right) \quad \forall i,$$

and finally by assuming the unknown function to be part of some finite-dimensional space (in our case: piecewise linear) in order to be able to represent it by a finite number of unknowns which can be used to express the integrals explicitly in a system of ordinary equations,

$$\rho = \sum_k \lambda_k \rho_k,$$

where $\{\rho_k\}$ are a known basis of the finite-dimensional function space; while $\{\lambda_k\}$ are the coefficients in the corresponding representation of ρ and, at the same time, the unknowns of the resulting system of equations.

Time discretization is achieved by an Euler scheme, backwards with respect to axial fluxes and forward with respect to radial fluxes. The latter treatment results in a step size requirement for the time integration, the numerically well-known Courant-Friedrichs-Lewy (CFL) condition (Courant et al., 1928). In this particular case this condition states: The more trans-membrane flux there is, the smaller the time step has to be chosen. If the requirement is not met (i.e., if the time step size is chosen too big) the solution will "explode," meaning that it will tend to infinity very rapidly. In order to prevent such instability, we calculate and use an estimate for the allowed step size. Thus, our time step is neither too big ("explosion") nor too small (inefficiency).

Discretization is performed using the numerical framework UG 4 (Vogel et al., 2013). It is written in C++ and simulations can be set up and run using the widespread scripting language Lua, which makes this framework easy to use without learning a highly specialized language of its own.

Solution of the symmetric system of linear equations emerging from the discretization is also done within the UG 4 framework. The tree structure of neurons allows for an efficient usage (i.e., with linear runtime complexity in terms of the number of unknowns) of a direct solver if the unknowns are numbered in such a way that, in each line of the matrix, there is at most one non-zero entry to the right of the diagonal. We use a Cuthill-McKee (Cuthill and McKee, 1969) ordering to guarantee this. We solve by calculating the LU decomposition in a sparse matrix format.

2.2.7. Parallelization

As UG 4 comes with full MPI support for parallel calculations, the inevitable usage of large-scale computer facilities for the simulation of large networks is straight forward. Partitioning of the domain can be performed using METIS 5.0 (Karypis and Kumar, 1998) and can be achieved on two levels:

In large networks, whole neurons can be assigned to the processors (as described for NEURON in Migliore et al., 2006), resulting in an "embarrassing" parallelism, since there is no direct coupling between the neurons if synaptic events triggered on the presynaptic side in one time step are taken into account only in the next time step on the postsynaptic side. If whole neurons can be distributed in such a way that the processors' workloads are well balanced, this will be the preferred way of parallelizing, as the solution of the problem works exactly like in the serial case and communication is only needed at active synapses.

On a second parallelization level, it is also possible to cut neurons and assign their parts to different processors. The process of solving the system of equations is a little bit more involved then. Assuming the system to be solved on a processor is

$$Ax = b,$$

then the iterative solving process on each processor is defined by the following pseudo-code:

$x_0 =$ solution from the precedent time step
$d = d_0 = b - Ax_0$ ("defect" vector)
while $|d| < |d_0| \cdot$ reductionFactor on any processor **do**
    $c = A^{-1}d$ (calculate correction)
    Sum up (over all processors) the corrections in all cutting points and store back in c.
    $x = x + c$ (update solution)
    $d = b - Ax$ (update defect)
end while

In order for this to work, the process-wise matrices A need to be stored "additively," i.e., the entries of the global system matrix must be equal to the sum of the corresponding entries in the process-wise matrices (where existent).

It usually takes about five to fifteen iterations until convergence is achieved, depending on how many neurons are cut and at which locations. Of course, in the case where no neuron is cut by the distribution of the network, the iteration will converge in one step. The gain in computation time from parallelizing on this level is not as big as from distributing whole neurons, obviously—however, it can still provide some speedup as it is not *solving* the system which takes the most time, but *setting it up* in the first place.

2.3. Simulation Workflow

The efficiency of simulating large and complex systems in neuroscience strongly depends on the scaling properties of code on high performance computers (Section 2.2.7). Additional aspects when looking at efficiency are the time invested for setting up a model, the computational tools, compiling and visualizing data and finally accessibility to an extendible code basis.

The simulation toolbox NeuroBox focusses on these aspects by allowing users to compile visual or script-based workflows. Workflows can define models, numerical tools and include third-party tools, such as ImageJ (Schneider et al., 2012). The multi-level design, founded on the multi-physics engine UG 4 (Vogel et al., 2013) and the Visual Reflection Library (VRL, Hoffer et al., 2013), allows non-experts intuitive access to advanced numerical methods for solving anatomically detailed biophysical models. NeuroBox is an open-source project hosted on github and thus is conceived as a modular and extendible C++ framework, where new biological components such as ion channels, receptors, synapse types etc. can be added manually or through an NMODL importer. This section briefly introduces script-based and visual workflow design and examples of the extendibility of NeuroBox, which as a platform is capable of hosting large multi-domain workflows.

2.3.1. Using Lua Scripts

The complete process of setting up and solving a problem in parallel is handled internally by UG 4. In order to use its functionality, we developed our code as a UG 4 plugin and compile against the UG 4 libraries. We register our classes and functions at the UG 4 registry (this is done in the C++ code) in order to make them available at the Lua script level, where a simulation can then be formulated using the registered functionality (in addition to any valid Lua command; see **Figure 4**). A schematic representation of what a typical simulation workflow looks like is shown in **Figure 5**, an example script with extensive comments is provided in Listing 1 in Supplementary Material.

FIGURE 4 | Sketch of the NeuroBox framework, using an ion channel model as an example. Functionality of the channel is implemented in a C++ class deriving from a pre-defined interface. The implementation can be automatically generated by conversion from NMODL model file. Registering the class at the UG 4 registry and compiling makes the channel available for usage in simulations defined either by a Lua script or by a graphical workflow representation using VRL-Studio.

FIGURE 5 | Illustration of the simulation workflow. After creation of a neuronal (network) morphology, the system of linear equations emerging from the cable equation is assembled by the central class `CableEquation`; synapse handling (i.e., activation, calculation of fluxes, parallel coordination) is taken care of by the class `SynapseHandler`, while all trans-membrane fluxes are handled by individual classes which all derive from a common interface known to the `CableEquation` class. The system is solved using UG4 solvers and parallelism.

2.3.2. Using Graphical Workflows

We take advantage of the open source software VRL-Studio (Hoffer et al., 2013) to represent simulation workflows graphically. Each class and function registered in the UG 4 registry can be represented in VRL-Studio. This allows any user to put together a simulation by dragging and dropping the graphical representations of involved objects (like instances of the cable equation discretization, the channel and pump mechanisms or the synapse handler) and adding application of their methods with only a few clicks. Scripts are not necessary but possible. The important aspect is that VRL-Studio can combine textual and visual programming in a single interactive development environment. For some aspects, script-based development has many advantages. Therefore, VRL-Studio provides access to the UG 4 APIs. Lua-scripts can be integrated into the visual workflow. A Lua editor with advanced autocompletion support allows for intuitive Lua-based development. Even more important is the fact that VRL-Studio workflows can integrate any Java library, such as ImageJ and JFreeChart. Automatic GUI generation works for these external libraries as well. Users can easily extend existing workflows with custom Groovy scripts, e.g., for pre- and post-processing. Custom scripts are also available as graphical components. Using external libraries in custom scripts is a powerful tool for adding domain-specific knowledge to the NeuroBox platform.

Typically the following steps can be followed to set up a new NeuroBox workflow (a screenshot of a simple graphical simulation workflow created in this way is depicted in **Figure 6**):

1. The first step is the definition of the computational domain (the neuronal morphology) and the unknown functions (membrane potential, ion concentrations) to be computed. This is done by adding an instance of `DomainAndFunctionDefinition` to the canvas and selecting the grid file as well as names for the unknown functions and subsets of the domain they are supposed to be defined on (subsets defined in the geometry file can be chosen from a list).
2. The following step in the workflow is the definition of all membrane transport mechanisms (channels and pumps) as well as a synapse handler (if any synapses are present in the domain). All of them may be drag-and-dropped onto the canvas and then parameterized as needed.
3. An instance of the central `CableEquation` class is added to the canvas. All defined membrane transport mechanisms as well as the synapse handler (if applicable) need to be connected by (gray) data connections. Initial conditions need to be supplied.
4. A solver is added to the canvas. The time stepping parameters need to be set as required. Output options can be specified.
5. Workflow connections (yellow) are drawn establishing the unique order in which the objects are created and their methods called. Any object receiving data input from another object needs to be behind that object in the workflow, i.e., the objects need to be in the order indicated by the enumeration here.
6. The simulation may be started, output can be observed in the log window, recorded and visualized.

2.3.3. Adding Functionality

As there is an abundance of membrane transport mechanisms and even more models trying to describe them, it is hardly
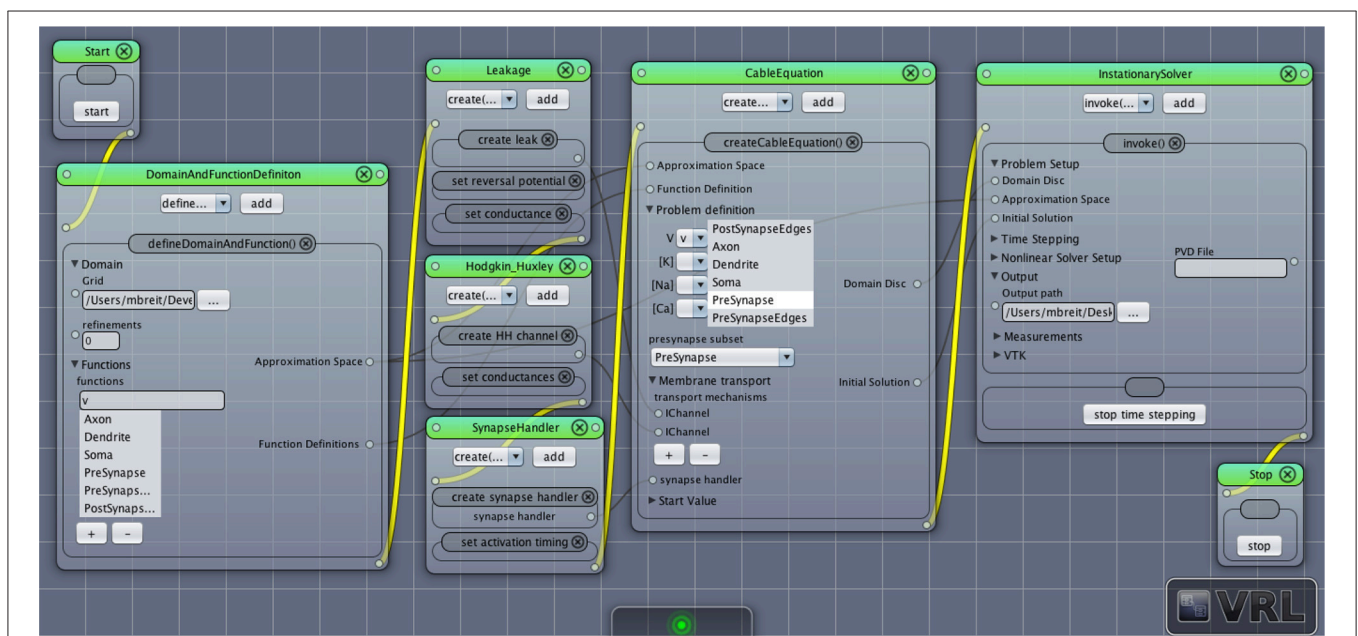
FIGURE 6 | Screenshot of a simple network simulation workflow assembled in VRL-Studio. Each window represents an object, the named panels in the window represent a method call with the contents of the panel as parameters. The control flow is defined by the yellow connections, data transfer between objects is marked by gray connections.

possible to implement all of them in advance. In order to support a large pool of available models, we wrote a file converter that will produce C++ code suitable to be compiled with our UG 4 implementation from any model file conforming to Neuron's NMODL description language (Hines and Carnevale, 2000). Of course, membrane transport models can also be implemented directly on the C++ level, implementing the required methods of a pre-defined interface class. This requires writing code for the initialization and updating (typically: evolving some kind of gating variables, expressed in terms of ordinary differential equations) of a model as well as code for the computation of the ion and charge flux through the membrane effectuated by the model. After registration of a new model at the UG 4 registry and compilation of the corresponding code, the model can be used on the Lua script level or on the graphical workflow level in VRL-Studio. The whole process is depicted schematically in **Figure 4**.

2.4. Setups for Our Simulations
2.4.1. Synapse Loss Simulations
We conducted *in silico* experiments investigating the impact of synapse loss in various activation patterns, particularly focussing on the effects it has on the formation of action potentials and the somatic calcium signal. For the simulations we chose a layer 3 pyramidal cell from the rat neocortex reconstructed by Radman et al. (2009), which was well suited to serve as reference cell for further studies as its reconstruction comprised the complete description of soma, dendrites and axon. The corresponding neuronal morphology is publicly available in the SWC file format as part of the *NeuroMorpho.org* database (Ascoli et al., 2007) under the name 13-L3pyr-77. It was converted to the UGX file format to meet UG 4 format specifications.

Subject to the discrete probability distribution specified in section 2.2.5, $N = 100$ distributions of $n_{syn} = 1000$ synapses each were drawn from the sample space S_{edge} defined in (9). We simulated synapse loss by successively removing portions of the previously created synapses uniformly from the neuron.

Regarding synapse activity we used a maximal conductance of $g_{max} = 1.2$ nS and a constant rise/decay time of $\tau = 0.4$ ms representing a fast AMPA receptor channel parameterization (Gabbiani et al., 1994) throughout the simulations. We compared three levels of input pattern synchrony, namely: complete synchrony ($\sigma_{onset} = 0$), moderate asynchrony ($\sigma_{onset} = 5$ ms), and high asynchrony ($\sigma_{onset} = 10$ ms).

A fraction of 0.2–4% of the current through AMPA receptor channels is carried by calcium ions, depending mainly on the exact AMPA subtype (Burnashev et al., 1995; Garaschuk et al., 1996). As we did not consider calcium buffer (calmodulin, calbindin) reactions in our simulations, we reduced this amount to 0.1% in order to (roughly) represent fast binding of free calcium to these buffers. Calcium dynamics were also regulated by N-type voltage-dependent calcium channels modeled according to Borg-Graham (1999) and NCX and PMCA pump mechanisms (first-order, second-order Hill-type model, resp.). A leakage term was added to ensure zero-flux for the equilibrium state.

2.4.2. Network Simulations
For the simulations in Section 3.2.2, we used NeuGen to create five neocortex geometries composed of 3500 L2/3 pyramidal; 3500 L4 spiny stellate; 1500 L5A and L5B pyramidal cells each whose somata were contained in a box with extensions of about 0.5 mm × 0.5 mm × 1 mm (length × width × depth), resulting in a cell density which is of the same order of magnitude as reported by Rockel et al. (1980). In each of the five geometries, NeuGen distributed an average of 30 primary synapses per L4 spiny stellate cell and an average of 25 per L5B pyramidal (cf. Constantinople and Bruno, 2013) for thalamic input. Interconnecting synapses were created wherever axon and dendrite from compatible neuron types came close enough, with the critical distance dist_synapse (cf. Section 2.1) being 1 μm for the first network, 2 μm for the second and so on. The numbers of synapses thus created show a cubical dependence on the critical creation distance (cf. **Table 2**), which is to be expected, as the sphere around any dendritic point within which axonal points eligible for connection through a synapse are located grows cubically in volume with increasing radius.

Axonal, dendritic and somatic membranes contained classical Hodgkin-Huxley-type sodium and potassium channels. Their flux density is described by

$$i_{hh} = c(T)\left(g_K n^4(V - E_K) + g_{Na} m^3 h\left(V - E_{Na}\right)\right), \quad (12)$$

$$\frac{\partial n}{\partial t} = c(T)\left(\alpha_n(V)(1-n) - \beta_n(V)n\right), \quad (13)$$

$$\frac{\partial m}{\partial t} = c(T)\left(\alpha_m(V)(1-m) - \beta_m(V)m\right), \quad (14)$$

$$\frac{\partial h}{\partial t} = c(T)\left(\alpha_h(V)(1-h) - \beta_h(V)h\right), \quad (15)$$

where $c(T)$ is a temperature-dependent constant with a value of about 3.2 at 37 °C (roughly taken from Collins and Rojas, 1982; Tiwari and Sikdar, 1999); g_K, g_{Na} are (location-specific) conductance constants; E_K and E_{Na} Nernst potentials; and the rate functions α and β are taken from the original Hodgkin and Huxley publication (Hodgkin and Huxley, 1952).

We used a leakage flux density to achieve zero net flux at resting potential:

$$i_l = c(T) g_l (V - E_l), \quad (16)$$

where g_l is the leakage flux conductance and E_l an (artificial) reversal potential calibrated to ensure zero membrane net flux at resting potential.

TABLE 2 | Number of synapses connecting specific cell types for various critical creation distances dist_synapse (averaged w.r.t. postsynaptic cell type) in the 10,000 cell networks created by NeuGen.

Synapse creation distance [μm]	1	2	3	4	5
L4 → L2/3	7.5	59	200	474	926
L4 → L5A	2.7	22	74	173	338
L2/3 → L5A	1.4	11	39	92	179
L2/3 → L5B	1.2	9.5	33	77	150
L2/3 → L2/3	0.48	3.8	13	30	58

TABLE 3 | Base synaptic conductance values for connections between different types in units of nS.

Synaptic type: pre- \post-	L2/3 pyr.	L4 stell.	L5A pyr.	L5B pyr.
L2/3 pyr.	1.0	–	0.8	0.3
L4 stell.	0.7	1.6	0.6	–
L5A pyr.	0.5	–	2.0	–
L5B pyr.	–	–	–	1.3

Connections that are not created by NeuGen are marked by a dash.

For initialization, the membrane potential was set to the resting potential of -0.065 V globally, voltage-dependent potassium and sodium channels were also set to their resting states. At the beginning of the simulation, thalamic input synapses were activated using an alpha function [cf. Equation (3)] with t_{onset}, τ drawn from normal distributions with $(\mu_{onset}, \sigma_{onset}) = (5\,\text{ms}, 2.5\,\text{ms})$, $(\mu_\tau, \sigma_\tau) = (2.5\,\text{ms}, 0.1\,\text{ms})$, respectively, and $g_{max} = 1.2$ nS.

Synapses between cells of the network were exclusively excitatory glutamatergic in nature and modeled as described in Section 2.2.3 using a parameterization which represents a fast AMPA receptor channel (Gabbiani et al., 1994). The maximal conductance parameter of synapse S with a presynaptic neuron of type T_1 and a postsynaptic neuron of type T_2 is calculated by NeuGen according to the formula

$$g_{max}(S) = \left(1 + 0.001 \cdot \text{dsd}(S)\right) \cdot g_s(T_1, T_2), \quad (17)$$

with dsd(S) being the post-synapse's distance to the soma in μm; and a type-specific base conductance the values of which are summed up in **Table 3**. All other synaptic parameters were the same for each synapse. No delay through neuro-transmitter release and diffusion was considered.

All parameter values for the network simulations are summed up in **Table 4**.

Simulations were performed on 160 processors for a simulated time period of 20 ms and took about two hours. Parallel scaling results for this type of problem are presented in Section 3.2.1.

3. RESULTS

3.1. Influence of Synapse Loss on Formation of Action Potentials and Somatic Calcium Signal

The human brain is one of the most complex structures known in the universe. It consists of nearly 100 billion nerve cells, each of which is entangled in a dense and constantly adapting network of massive information exchange. On average, a single neuron is linked with 10,000 to 100,000 other neuronal or non-neuronal cells via synapses (Cragg, 1975). Brain function relies essentially on those highly dynamic synaptic connections.

In this part of our study, we investigate the three-dimensional spatial distribution and activity pattern in time of glutamatergic

TABLE 4 | Parameters for the large-scale network simulation.

Parameter	Meaning	Value	Unit
r_c	Specific resistance of the cytosol	1.5	Ω m
c_m	Specific capacitance of the membrane	1×10^{-2}	F m^{-2}
g_K^a	Specific potassium channel conductance of the axonal membrane	4×10^2	S m^{-2}
g_K^s	Specific potassium channel conductance of the somatic membrane	2×10^2	S m^{-2}
g_K^d	Specific potassium channel conductance of the dendritic membrane	3×10^1	S m^{-2}
g_{Na}^a	Specific sodium channel conductance of the axonal membrane	3×10^4	S m^{-2}
g_{Na}^s	Specific sodium channel conductance of the somatic membrane	1.5×10^3	S m^{-2}
g_{Na}^d	Specific sodium channel conductance of the dendritic membrane	4×10^1	S m^{-2}
g_l^a	Specific leak conductance of the axonal membrane	2×10^2	S m^{-2}
g_l^s	Specific leak conductance of the somatic membrane	1	S m^{-2}
g_l^d	Specific leak conductance of the dendritic membrane	1	S m^{-2}
E_K	Potassium Nernst potential	-0.09	V
E_{Na}	Sodium Nernst potential	0.06	V
E_l^a	Leak reversal potential of the axonal membrane	-0.066148458	V
E_l^s	Leak reversal potential of the somatic membrane	-0.030654022	V
E_l^d	Leak reversal potential of the dendritic membrane	-0.057803624	V
V_r	Resting potential (global)	-0.065	V
$c(T)$	Temperature factor for HH channel activity	3.21	
τ_1	Time constant in interconnecting synapses	2×10^{-4}	s
τ_2	Time constant in interconnecting synapses	1.7×10^{-3}	s
E_{rev}	Reversal potential for interconnecting synapse influx	0	V
V_{th}	Threshold potential for interconnecting synapse activation	-0.01	V

synapses in neurons of the cerebral cortex. Both are key factors to the integrative properties of the cell. For this purpose, we have developed a tool for automatic placement of synaptic functionality onto neuron morphologies. We apply this tool to systematically assess the impact of activation patterns on the signal processing in single neurons. In particular, we perform *in silico* experiments where we successively knock out synapses at dendritic locations. We thus investigate situations where synapse loss contributes to pathological states e.g., Alzheimer's disease (Scheff et al., 2006). At the same time, we address the question under which circumstances the neuron will sustain its integrative capability. More precisely, how does impulse conductance and especially the initiation of action potentials at the axon hillock depend on the number of input synapses and

their signal synchrony? Does a higher input signal synchrony sustain action potential initiation during increasing synapse loss? The degree of synchrony is defined by the size of the standard deviation from a given mean value. In our experiments, we vary the standard deviation of the start time σ_{onset} of synaptic excitations.

In the following sections, we present the results of a series of *in silico* experiments on a layer 3 pyramidal cell from the rat neocortex (cf. Section 2.4.1), in which we compare three levels of input pattern synchrony, namely: complete synchrony ($\sigma_{onset} = 0$), moderate asynchrony ($\sigma_{onset} = 5$ ms), and high asynchrony ($\sigma_{onset} = 10$ ms). We randomly distributed 1000 excitatory synapses on the geometry in 100 sample configurations. In each of these 100 configurations, we gradually increased synapse loss and analyzed the neuron's capability of creating action potentials, and at the same time, recorded corresponding calcium levels within the soma.

3.1.1. Generation of Action Potentials
Both moderate ($\mu_{onset} = 15$ ms, $\sigma_{onset} = 5$ ms) and high ($\mu_{onset} = 30$ ms, $\sigma_{onset} = 10$ ms) asynchrony cases show a strong action potential spike train response to the initial synapse distribution. The number of spikes ranges from two to three in the moderate asynchrony case and from one to three in the high asynchrony case (**Figure 7**). The synchronous setup, however, produced exactly one action potential for the initial distribution of 1000 synapses in all samples. Only the cation influx at new synapses perpetually being active in the asynchronous cases can induce the repetitive spiking, while cation influx through all synapses is completely compensated by potassium efflux during hyper-polarization in the synchronous case. The number of action potentials decreased with increasing synapse loss in both asynchronous cases until complete signal breakdown (in at least 90 % of the sample patterns) at 75% synapse loss in the moderately and at 60% in the highly asynchronous case. In contrast, synchronous activation patterns sustained generation of

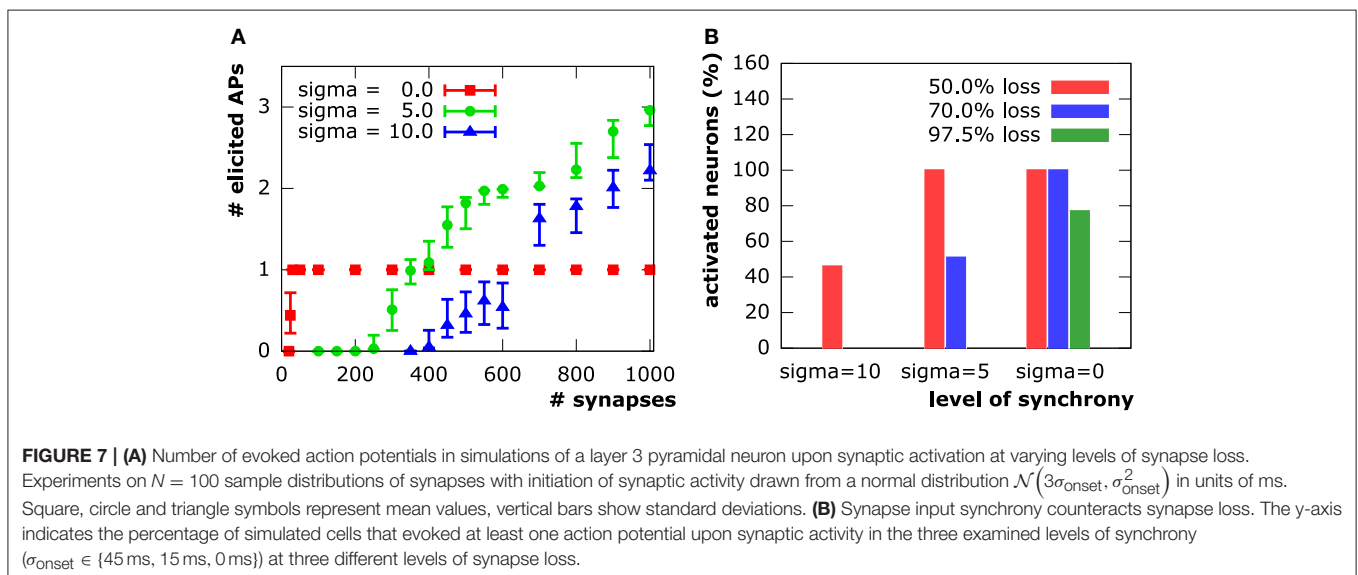an action potential up to a loss of about 97.7% (corresponding to 23 synapses).

3.1.2. Calcium Signaling
In all setups, synchronous as well as asynchronous, calcium levels at the soma exclusively depend on whether or not an action potential is elicited. We see step increases in the calcium concentration with every action potential. Calcium diffusion, however, is only able to transport calcium within a very local vicinity of its original point of entry at active synapses. After termination of electrical signaling, calcium levels exponentially decay to equilibrium levels due to the activity of NCX and PMCA pumps. This shows a direct correspondence between synapse loss and somatic calcium levels through the number of action potentials elicited in a neuron. Sample evolution of membrane potential and calcium concentration at the soma (from the moderately asynchronous setting) for various levels of excitatory synapse loss are depicted in **Figure 8**.

3.2. Large-Scale Network Simulations with Detailed Anatomy
3.2.1. Parallel Scaling
In order to test the parallel scaling properties of our network simulation implementation, we created six neocortical geometries containing 320, 640, 1280, 2560, 5120, and 10,240 neurons, respectively. The average number of compartments per neuron in the six geometries ranged from 574 to 586. We defined a random thalamic activation pattern, where synapse activation times and durations for the thalamic input synapses created by NeuGen were drawn from the same normal distribution for all geometries. We then performed one thousand time steps using 32, 64, 128, 256, 512, and 1024 processors of the Jülich supercomputer JUQUEEN on the geometries, respectively—thus in each simulation, a processor would be assigned approximately the same amount of work ("weak scaling"). We profiled the

FIGURE 7 | (A) Number of evoked action potentials in simulations of a layer 3 pyramidal neuron upon synaptic activation at varying levels of synapse loss. Experiments on $N = 100$ sample distributions of synapses with initiation of synaptic activity drawn from a normal distribution $\mathcal{N}\left(3\sigma_{onset}, \sigma_{onset}^2\right)$ in units of ms. Square, circle and triangle symbols represent mean values, vertical bars show standard deviations. **(B)** Synapse input synchrony counteracts synapse loss. The y-axis indicates the percentage of simulated cells that evoked at least one action potential upon synaptic activity in the three examined levels of synchrony ($\sigma_{onset} \in \{45$ ms, 15 ms, 0 ms$\}$) at three different levels of synapse loss.
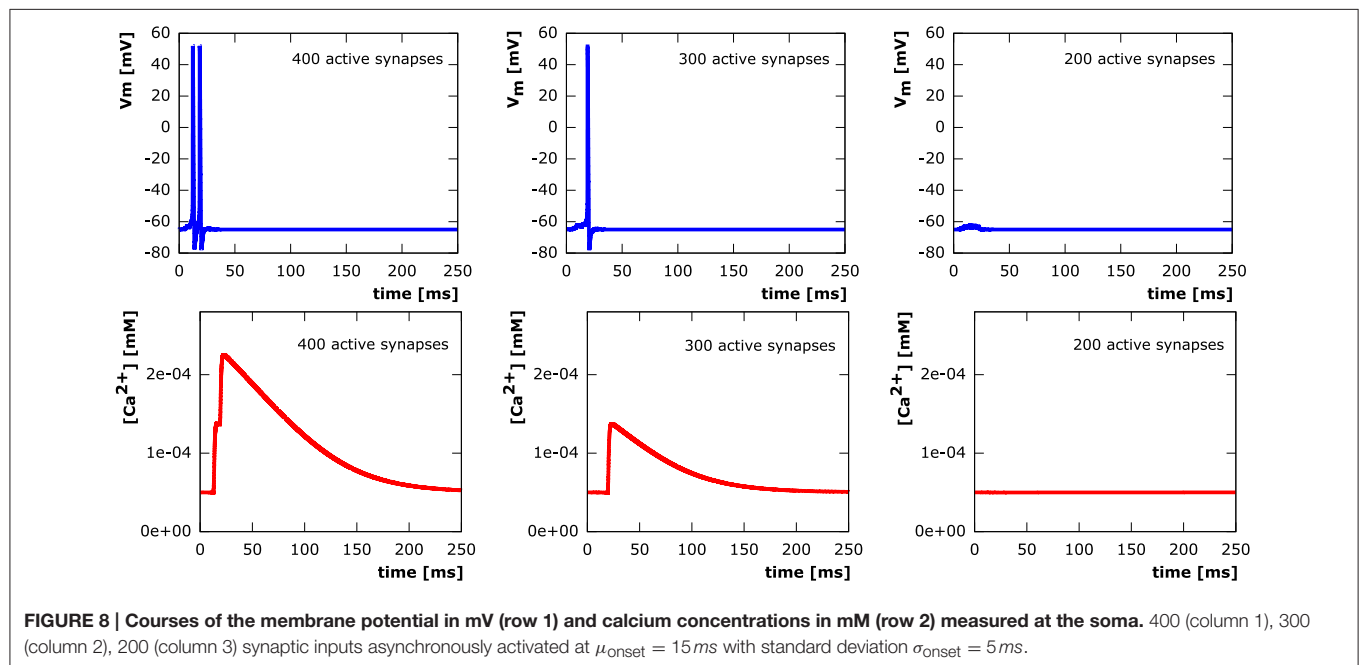
execution of the program to obtain the amount of time spent in the main components of the simulation. **Table 5** shows the results. Leaving out the loading of the geometry into memory and its distribution to the involved processors (both are inherently serial), we achieve good scaling. The times spent for preparing the channel mechanisms and synapses before a time step, for assembling, for factorizing the matrix and applying the inverse remain approximately constant. As a typical network simulation will have more than 1000 time steps, the loading and distribution of the domain (which is only performed once, i.e., at the start of the parallel simulation) will have much less of an impact on scaling behavior than in this particular study. We thus

demonstrated that our code is suitable to be used efficiently for simulations of large-scale networks of neurons.

3.2.2. Network Connectivity Affects Network Activity

When a neuronal network is created by NeuGen, synapses connect presynaptic axons to postsynaptic dendrites (if the involved neuron types allow this) where axon and dendrite are sufficiently close to each other (cf. Wolf et al., 2013). The maximal distance `dist_synapse` for which synapses are placed can be chosen by the user. This criterion, albeit not representative of an actual model of synaptogenesis (NeuGen does not reproduce neuronal growth, but only a fully grown state), might be

FIGURE 8 | Courses of the membrane potential in mV (row 1) and calcium concentrations in mM (row 2) measured at the soma. 400 (column 1), 300 (column 2), 200 (column 3) synaptic inputs asynchronously activated at $\mu_{onset} = 15\,ms$ with standard deviation $\sigma_{onset} = 5\,ms$.

TABLE 5 | Weak parallel scaling results obtained by code profiling.

| Problem size (#neurons) | 320 | | 640 | | 1280 | | 2560 | | 5120 | | 10240 |
#Processors	32	×	64	×	128	×	256	×	512	×	1024
Loading domain	12.8	2.08	26.6	2.03	54.1	2.09	113	2.17	246	2.29	562
Domain distribution	4.89	2.12	10.4	2.06	21.4	2.11	45.3	2.10	95.0	2.07	197
Determining step size	42.0	1.03	43.5	1.02	44.2	0.98	43.4	1.00	43.4	1.02	44.1
Preparing time step	52.9	1.01	53.5	0.99	53.2	1.02	54.2	1.02	55.5	1.13	62.5
Assembling system	325	1.02	330	1.03	338	0.97	329	0.99	327	1.02	333
Applying solver	24.0	1.03	24.8	1.01	25.0	0.94	23.4	1.04	24.3	0.92	22.3
Rest	6.41	1.90	12.2	0.33	4.00	2.18	8.7	2.05	17.8	1.13	20.1
Total time	468	1.07	501	1.08	540	1.14	617	1.31	809	1.54	1241
Total w/o load & distribute	450	1.03	464	1.00	465	0.99	458	1.02	468	1.03	483

Timings show almost ideal scaling for setting up the system of equations as well as solving it. Loading and distributing the domain to the involved processors is inherently serial as it is done by a single processor in every simulation—however, the time percentage of the two tasks is much lower in typical simulations as they run much longer and loading and distributing is only performed once at the beginning. The remaining variance of the simulation runs is mainly due to differences in the quality of domain distribution. All time values in units of seconds.

considered as a parameterization of the agility of filopodia and growth cones during synaptogenesis (Munno and Syed, 2003). In any case, it has a direct effect on the connectivity properties of the network.

We conducted simulations on five neocortical networks, each composed of the same 10,000 neurons, but with the connection distance ranging from 1 μm to 5 μm in steps of 1 μm. This resulted in networks with increasing numbers of synapses and connected neurons (**Table 3**). As previously described in Wanner (2007), we initialized network activity by depolarizing L4 spiny stellate cells via primary thalamic input synapses, activity then spread out through the cortical layers due to interconnecting synapses. Analysis of the time courses of the membrane potential at the somata in conjunction with activity data from the interconnecting synpases (**Figure 9**) reveals significant impact of the connectedness on the overall qualitative (and quantitative) behavior following the same thalamic input pattern in all five simulations.

In the least connected network, the number of synapses connecting thalamically activated L4 spiny stellate cells to L2/3

pyramidal cells (only 7.5 per L2/3 pyramidal cell on average) does not suffice to lead to the depolarization of a single L2/3 cell in the network. Obviously, this means there can be no active synapses connecting L2/3 to L5A and although there are also synapses connecting L4 to L5A directly, there is no activity in L5A, either. While in the network next in synapse number, considerable depolarization of layer 2/3 pyramidal neuron somata manifests itself due to 7.5-fold increase in average number of active synapses from L4 to L2/3, there is still practically no signal in L5A. Only in the networks created with synapse creation distance parameters ≥ 3 μm are action potentials elicited at the somata of L5A. The same networks exhibit the formation of a second action potential in some of the initially activated L2/3 somata, the two most connected networks also show the occurrence of a second action potential in some of the L5A cells. These second action potentials are the combined effect of (i) charge from previous synaptic inputs that has not yet been cleared and (ii) additional influx at the re-activated synapses. It is noteworthy that somatic activity in both L2/3 and L5A pyramidal cells peaks higher (L2/3: 0.0, 0.56, 0.84, 0.87, 0.87; L5A: 0.0, 0.0, 0.39, 0.78, 0.87) and earlier
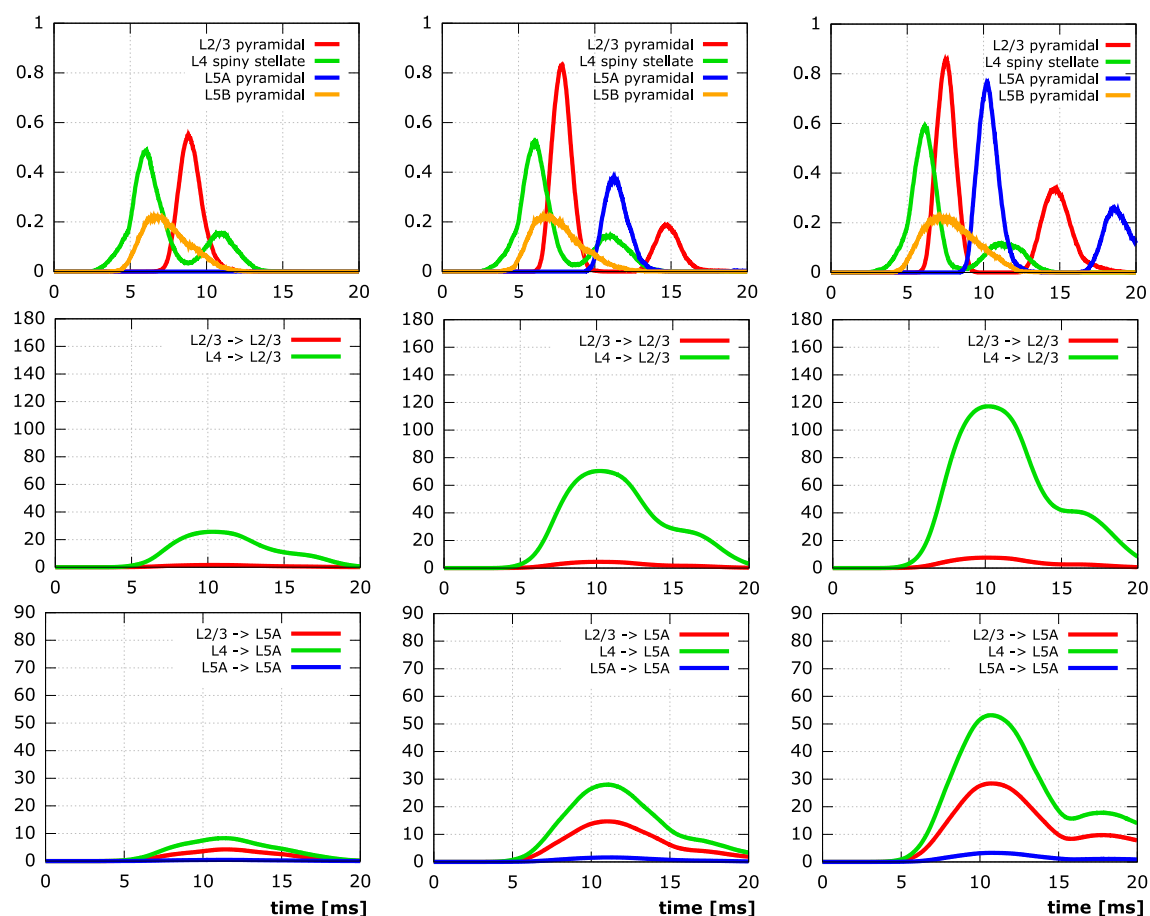
FIGURE 9 | Simulation results for networks of 10,000 neurons. The columns contain time plots for networks created with a synapse creation distance of 2 μm, 3 μm, 4 μm (FLTR), respectively. Row 1: Relative number of active somata, i.e., $V \geq -45$ mV, in different levels (L2/3 red, L4 green, L5A blue, L5B orange). Row 2: Number of active synapses at L2/3 pyramidal neurons originating from different levels (L2/3 red, L4 green). Row 3: Number of active synapses at L5A pyramidal neurons originating from different levels (L2/3 red, L4 green, L5A blue). Initial activation of L4 spiny stellate and L5B pyramidal cells by the same thalamic input pattern in all simulations.
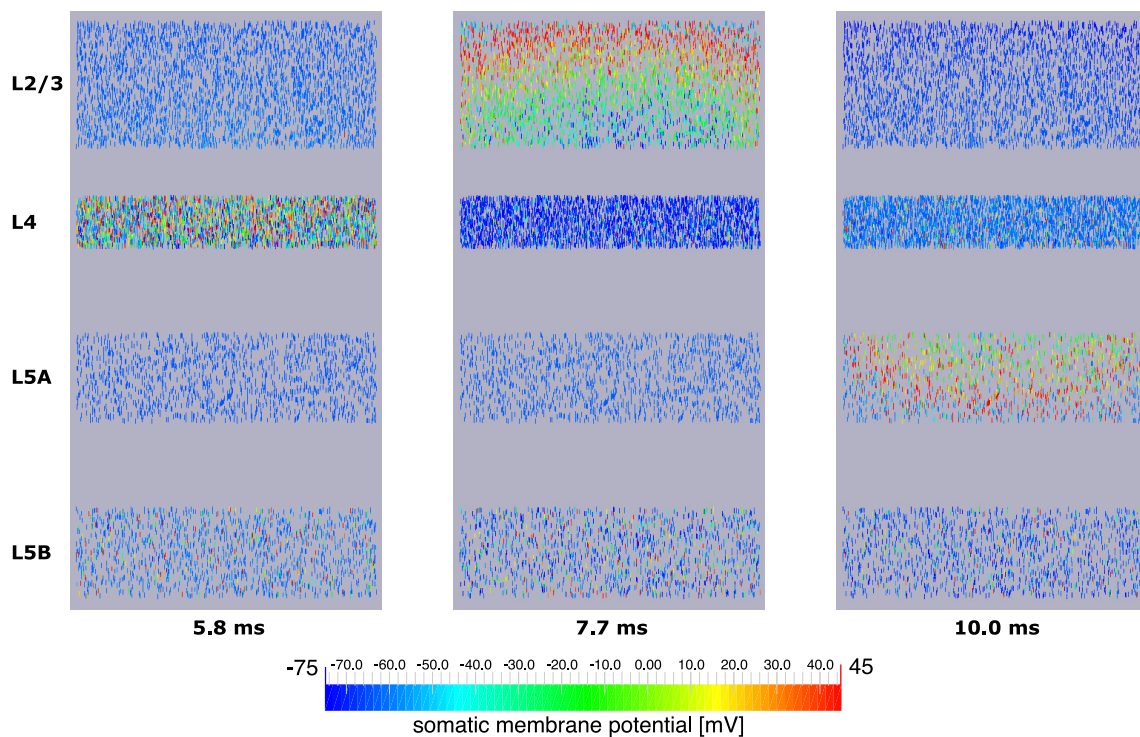
FIGURE 10 | Signal propagation through a neocortical slice of 10,000 neurons. Only somata are visualized. Primary excitation of L4 spiny stellate cells (left) is followed by activation of somata of L2/3 pyramidal cells (middle), before the signal propagates to pyramidal neurons of layer 5A (right). Note how the depolarization of somata expands through layers 2/3 and 5A like a wavefront indicating increased signal run time to neurons more remote with respect to the signal origin.

(L2/3: –, 8.8 ms, 7.9 ms, 7.6 ms, 7.6 ms; L5A: –, –, 11.4 ms, 10.2 ms, 9.7 ms) the more synaptic connections there are in the respective cortical layers.

Explicit influence of spatial extensions of the neural network can be identified in **Figure 10**: Somatic depolarization and hyperpolarization expands through the layers of L2/3 and L5A pyramidal cells like a wave, activating the neurons in the order imposed by the distance to the respective origin of that activation.

4. DISCUSSION

In this paper we presented studies of electrical and biochemical signals in single cells and networks to investigate the interplay between synapse loss and signaling synchrony. A major focus was the anatomically realistic representation of cells and networks, for which a novel simulation toolbox NeuroBox was developed.

The synapse distribution studies on the layer 3 pyramidal cell from the rat neocortex show a significant impact of the activation pattern (in space and time) on the signal conductance capabilities of the cell. Two effects are apparent: (1) The more asynchronous the input signals are, the more spikes can be generated by this input—up to a point where the asynchrony begins to affect the likelihood of generating a single spike. (2) The more synchronous input signals are, the higher the cell's resilience is to synapse loss with regard to its capability of generating action potentials in response to synaptic input.

In the context of the study of synaptic input patterns, we also conducted simulations of calcium dynamics, including Ca^{2+} influx through synaptic AMPA-R channels as well as voltage-dependent calcium channels, NCX and PMCA pump mechanisms distributed throughout the membranes of dendrites and soma. Results showed that the somatic calcium concentration, key factor in the control of gene expression (Hardingham et al., 1997) and thus development and survival of cells, is directly coupled to the number of action potentials initiated in the cell, each action potential leading to a step increase in calcium levels. However, we neglected effects of internal calcium stores and also correct consideration of calcium buffers here. Especially the large amounts of calcium releasable through ryanodine and IP$_3$ receptor channels in the membrane of the endoplasmic reticulum need to be taken into account in a detailed three-dimensional simulation in order to achieve a more accurate description of calcium signaling, possibly including calcium waves (cf. Berridge, 1998, among others). A method of coupling the one-dimensional simulation of the membrane potential to a detailed three-dimensional simulation of calcium signals has previously been developed by the authors (Grein et al., 2014) and may be applied here.

Using NeuGen for the generation of a neocortical column, we have shown that our implementation of a compartment model for the cable equation and trans-membrane current mechanisms in neural networks is adequate for large-scale applications and scales well with the number of neurons involved. It is reasonable

to assume that simulations on even larger neural networks can successfully and efficiently be conducted on high-performance computers with the help of our implementation.

The neural network simulations we performed were very basic in nature. We only considered four of the diverse neuron types present in the neocortex. Unlike (e.g., Anderson et al., 2007; Vierling-Claassen et al., 2010; Neymotin et al., 2011), we did not take into account inhibitory synapses and their role in regulating cortical signal processing. Unlike the three aforementioned contributions, however, we created neural networks whose spatial resolution—about 500 compartments per neuron in our simulations as compared to 3, 16 and 1 in theirs—allowed for a realistic spatial positioning of synapses. We utilized the simple (yet not unreasonable) distance rule of NeuGen to create synapses instead of putting experimental projection data (as extensively reviewed for excitatory neurons by Feldmeyer, 2012) to good use. Incorporation of experimental findings into the existing framework, however, is not difficult. The addition of inhibitory synapses, for instance, is merely a question of re-parameterization in a preprocessing step. All that considered, our network simulations make it possible to examine the impact of intra- and trans-laminar synaptic connections on each level and can therefore serve as a valuable tool to decipher the functional role of detailed anatomy in cortical information processing.

With a focus on accessible workflow control that includes high-performance numerical methods, a modular neuroscientific repository and the option of including third-party tools, we developed the toolbox NeuroBox and used it to perform all simulations in this paper. NeuroBox is an open source project hosted on github with the intent to offer its full functional scope to a broad community. Visual workflow design and control through VRL-Studio makes NeuroBox projects easy to use and share with experts and non-experts alike. This feature is highly beneficial for rapid prototyping and offers an efficient pathway from *in silico* experiment design to full implementation thereof.

The possibility to integrate third-party tools, such as ImageJ, anatomical reconstructions (e.g., neuromorpho.org) and the automated import of NMODL models, integrates NeuroBox ideally into ongoing endeavors in the computational neuroscience field. Due to the modular design, this toolbox is easily extendible through various pathways discussed in Section 2 and thus can grow with continued research. As problem sizes typically increase alongside growing high-performance computing power, NeuroBox was built with links to UG 4, a general purpose package for solving partial differential equations. Advanced numerical methods with time and space adaptivity, error estimation and parallel communication layer advance the possibilities for solving anatomically realistic large-scale network problems.

AUTHOR CONTRIBUTIONS

Cable equation modeled and implemented by MB and PG. Synapse handling designed and implemented by MB, SG, LR, MS. Neuron and network morphology preparation by SG. Network simulations carried out by MB and PG, synapse loss and calcium simulations by MB, PG, LR, and MS. GQ co-designed all methods and experiments and analyzed data. MB, SG, GQ, and MS wrote the manuscript.

FUNDING

ACKNOWLEDGMENTS

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: http://journal.frontiersin.org/article/10.3389/fnana.2016.00008

REFERENCES

Abbott, L. F., and Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nat. Neurosci.* 3, 1178–1183. doi: 10.1038/81453

Ahrens, J., Geveci, B., and Law, C. (2005). "ParaView: an end-user tool for large data visualization," in *Visualization Handbook*, eds C. D. Hansen and C. R. Johnson (Burlington: Butterworth-Heinemann), 717–731. doi: 10.1016/B978-012387582-2/50038-1

Anderson, W. S., Kudela, P., Cho, J., Bergey, G. K., and Franaszczuk, P. J. (2007). Studies of stimulus parameters for seizure disruption using neural network simulations. *Biol. Cybern.* 97, 173–194. doi: 10.1007/s00422-007-0166-0

Andrews, S. S., Addy, N. J., Brent, R., and Arkin, A. P. (2010). Detailed simulations of cell biology with Smoldyn 2.1. *PLoS. Comput. Biol.* 6:e1000705. doi: 10.1371/journal.pcbi.1000705

Arellano, J. I., Benavides-Piccione, R., DeFelipe, J., and Yuste, R. (2007). Ultrastructure of dendritic spines: correlation between synaptic and spine morphologies. *Front. Neurosci.* 1, 131–143. doi: 10.3389/neuro.01.1.1.010.2007

Ascoli, G. A. (2006). Mobilizing the base of neuroscience data: the case of neuronal morphologies. *Nat. Rev. Neurosci.* 7, 318–324. doi: 10.1038/nrn1885

Ascoli, G. A., Donohue, D. E., and Halavi, M. (2007). NeuroMorpho.org: a central resource for neuronal morphologies. *J. Neurosci.* 27, 9247–9251. doi: 10.1523/JNEUROSCI.2055-07.2007

Balls, G. T., Baden, S. B., Kispersky, T., Bartol, T. M., and Sejnowski, T. J. (2004). "A large scale Monte Carlo simulator for cellular microphysiology," in *Parallel and Distributed Processing Symposium, 2004. Proceedings of the 18th International*. doi: 10.1109/IPDPS.2004.1302957

Berridge, M. J. (1998). Neuronal calcium signaling. *Neuron* 21, 13–26. doi: 10.1016/S0896-6273(00)80510-3

Borg-Graham, L. J. (1999). "Interpretations of data and mechanisms for hippocampal pyramidal cell models" in *Models of Cortical Circuits*, eds P. S. Ulinski, E. G. Jones, and A. Peters (Boston, MA: Springer US), 19–138. doi: 10.1007/978-1-4615-4903-1_2

Bower, J., and Beeman, D. (1997). *The book of GENESIS: Exploring realistic neural models with the GEneral NEural SImulation System*. New York, NY: Springer.

Branco, T., Clark, B. A., and Häusser, M. (2010). Dendritic discrimination of temporal input sequences in cortical neurons. *Science* 329, 1671–1675. doi: 10.1126/science.1189664

Burette, A. C., Lesperance, T., Crum, J., Martone, M., Volkmann, N., Ellisman, M. H., et al. (2012). Electron tomographic analysis of synaptic ultrastructure. *J. Comp. Neurol.* 520, 2697–2711. doi: 10.1002/cne.23067

Burnashev, N., Zhou, Z., Neher, E., and Sakmann, B. (1995). Fractional calcium currents through recombinant GluR channels of the NMDA, AMPA and kainate receptor subtypes. *J. Physiol.* 485, 403–418. doi: 10.1113/jphysiol.1995.sp020738

Chen, X., Winters, C., Azzam, R., Li, X., Galbraith, J. A., Leapman, R. D., et al. (2008). Organization of the core structure of the postsynaptic density. *Proc. Natl. Acad. Sci. U.S.A.* 105, 4453–4458. doi: 10.1073/pnas.0800897105

Collins, C. A., and Rojas, E. (1982). Temperature dependence of the sodium channel gating kinetics in the node of Ranvier. *Q. J. Exp. Physiol.* 67, 41–55. doi: 10.1113/expphysiol.1982.sp002623

Colón-Ramos, D. A. (2009). Synapse formation in developing neural circuits. *Curr. Top. Dev. Biol.* 87, 53–79. doi: 10.1016/S0070-2153(09)01202-2

Constantinople, C. M., and Bruno, R. M. (2013). Deep cortical layers are activated directly by thalamus. *Science* 340, 1591–1594. doi: 10.1126/science.1236425

Courant, R., Friedrichs, K., and Lewy, H. (1928). Über die partiellen Differenzengleichungen der mathematischen Physik. *Math. Ann.* 100, 32–74. doi: 10.1007/BF01448839

Cragg, B. G. (1975). The density of synapses and neurons in normal, mentally defective and ageing human brains. *Brain* 98, 81–90. doi: 10.1093/brain/98.1.81

Cuthill, E., and McKee, J. (1969). "Reducing the bandwidth of sparse symmetric matrices," in *Proceedings of the 1969 24th National Conference, ACM'69* (New York, NY: ACM), 157–172. doi: 10.1145/800195.805928

Eberhard, J., Wanner, A., and Wittum, G. (2006). NeuGen: a tool for the generation of realistic morphology of cortical neurons and neural networks in 3D. *Neurocomputing* 70, 327–342. doi: 10.1016/j.neucom.2006.01.028

Feldmeyer, D. (2012). Excitatory neuronal connectivity in the barrel cortex. *Front. Neuroanat.* 6:24. doi: 10.3389/fnana.2012.00024

Gabbiani, F., Midtgaard, J., and Knöpfel, T. (1994). Synaptic integration in a model of cerebral granule cells. *J. Neurophysiol.* 72, 999–1009. Available online at: http://jn.physiology.org/content/72/2/999

Garaschuk, O., Schneggenburger, R., Schirra, C., Tempia, F., and Konnerth, A. (1996). Fractional Ca2+ currents through somatic and dendritic glutamate receptor channels of rat hippocampal CA1 pyramidal neurones. *J. Physiol.* 491, 757–772. doi: 10.1113/jphysiol.1996.sp021255

Gewaltig, M. O., and Diesmann, M. (2007). NEST (NEural Simulation Tool). *Scholarpedia* 2, 1430. doi: 10.4249/scholarpedia.1430

Grein, S., Stepniewski, M., Reiter, S., Knodel, M. M., and Queisser, G. (2014). 1D-3D hybrid modelling – From multi-compartment models to full resolution models in space and time. *Front. Neuroinform.* 8:68. doi: 10.3389/fninf.2014.00068

Grillo, A., Logashenko, D., Stichel, S., and Wittum, G. (2010). Simulation of density-driven flow in fractured porous media. *Adv. Water Resour.* 33, 1494–1507. doi: 10.1016/j.advwatres.2010.08.004

Hansen, S., Henning, A., Nägel, A., Heisig, M., Wittum, G., Neumann, D., et al. (2008). *In-silico* model of skin penetration based on experimentally determined input parameters. Part I: experimental determination of partition and diffusion coefficients. *Eur. J. Pharm. Biopharm.* 68, 352–367. doi: 10.1016/j.ejpb.2007.05.012

Hardingham, G. E., Chawla, S., Johnson, C. M., and Bading, H. (1997). Distinct functions of nuclear and cytoplasmic calcium in the control of gene expression. *Nature* 385, 260–265. doi: 10.1038/385260a0

Hines, M. L., and Carnevale, N. T. (1997). The NEURON simulation environment. *Neural Comput.* 9, 1179–1209. doi: 10.1162/neco.1997.9.6.1179

Hines, M. L., and Carnevale, N. T. (2000). Expanding NEURON's repertoire of mechanisms with NMODL. *Neural Comput.* 12, 995–1007. doi: 10.1162/089976600300015475

Hodgkin, A. L., and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* 117, 500–544. doi: 10.1113/jphysiol.1952.sp004764

Hoffer, M., Poliwoda, C., and Wittum, G. (2013). Visual reflection library: a framework for declarative GUI programming on the Java platform. *Comput. Vis. Sci.* 16, 181–192. doi: 10.1007/s00791-014-0230-y

Hughes, J. R. (1958). Post-tetanic potentiation. *Physiol. Rev.* 38, 91–113.

Jungblut, D., Queisser, G., and Wittum, G. (2011). Inertia based filtering of high resolution images using a GPU cluster. *Comput. Vis. Sci.* 14, 181–186. doi: 10.1007/s00791-012-0171-2

Karypis, G., and Kumar, V. (1998). A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comput.* 20, 359–392. doi: 10.1137/S1064827595287997

Thompson, W. (1854). On the theory of the electric telegraph. *Proc. R. Soc.* 7, 382–399.

Lopreore, C. L., Bartol, T. M., Coggan, J. S., Keller, D. X., Sosinsky, G. E., Ellisman, M. H., et al. (2008). Computational modeling of three-dimensional electrodiffusion in biological systems: application to the node of Ranvier. *Biophys. J.* 95, 2624–2635. doi: 10.1529/biophysj.108.132167

Markram, H. (2015). Reconstruction and simulation of neocortical microcircuitry. *Cell* 163, 456–492. doi: 10.1016/j.cell.2015.09.029

Migliore, M., Cannia, C., Lytton, W., Markram, H., and Hines, M. L. (2006). Parallel network simulations with NEURON. *J. Comput. Neurosci.* 21, 119–129. doi: 10.1007/s10827-006-7949-5

Muha, I., Nägel, A., Stichel, S., Grillo, A., Heisig, M., and Wittum, G. (2011). Effective diffusivity in membranes with tetrakaidekahedral cells and implications for the permeability of human stratum corneum. *J. Memb. Sci.* 368, 18–25. doi: 10.1016/j.memsci.2010.10.020

Munno, D. W., and Syed, N. I. (2003). Synaptogenesis in the CNS: an odyssey from wiring together to firing together. *J. Physiol.* 552, 1–11. doi: 10.1113/jphysiol.2003.045062

Nägel, A., Hansen, S., Neumann, D., Lehr, C. M., Schaefer, U. F., Wittum, G., et al. (2008). *In-silico* model of skin penetration based on experimentally determined input parameters. Part II: mathematical modelling of *in-vitro* diffusion experiments. identification of critical input parameters. *Eur. J. Pharm. Biopharm.* 68, 368–379. doi: 10.1016/j.ejpb.2007.05.018

Nägel, A., Heisig, M., and Wittum, G. (2009). A comparison of two- and three-dimensional models for the simulation of the permeability of human stratum corneum. *Eur. J. Pharm. Biopharm.* 72, 332–338. doi: 10.1016/j.ejpb.2008.11.009

Neymotin, S. A., Jacobs, K. M., Fenton, A. A., and Lytton, W. W. (2011). Synaptic information transfer in computer models of neocortical columns. *J. Comp. Neurosci.* 30, 69–84. doi: 10.1007/s10827-010-0253-4

Popov, V. I., Kleschevnikov, A. M., Klimenko, O. A., Stewart, M. G., and Belichenko, P. V. (2011). Three-dimensional synaptic ultrastructure in the dentate gyrus and hippocampal area CA3 in the Ts65Dn mouse model of down syndrome. *J. Comp. Neurol.* 519, 1338–1354. doi: 10.1002/cne.22573

Purves, D., Augustine, G. J., and Fitzpatrick, D. (2001). *Neuroscience*. Sunderland, MA: Sinauer Associates.

Qian, N., and Sejnowski, T. (1989). An electro-diffusion model for computing membrane potentials and ionic concentrations in branching dendrites, spines and axons. *Biol. Cybern.* 62, 1–15. doi: 10.1007/BF00217656

Radman, T., Ramos, R., Brumberg, J., and Bikson, M. (2009). Role of cortical cell type and morphology in sub- and suprathreshold uniform electric field stimulation. *Brain Stimul.* 2, 215–228. doi: 10.1016/j.brs.2009.03.007

Ramaswamy, S. (2015). The neocortical microcircuit collaboration portal: a resource for rat somatosensory cortex. *Front. Neural Circuits* 9:44. doi: 10.3389/fncir.2015.00044

Reimann, M. (2015). An algorithm to predict the connectome of neural microcircuits. *Front. Comput. Neurosci.* 9:120. doi: 10.3389/fncom.2015.00120

Reiter, S. (2012). ProMesh – Meshing of unstructured grids in 1, 2, and 3 dimensions. Available online at: http://promesh3d.com

Reiter, S. (2014). *Effiziente Algorithmen und Datenstrukturen für die Realisierung von Adaptiven, Hierarchischen Gittern auf Massiv Parallelen Systemen*. Dissertation. Universität Frankfurt, Germany.

Rockel, A. J., Hiorns, R. W., and Powell, T. P. S. (1980). The basic uniformity in structure of the neocortex. *Brain* 103, 221–244. doi: 10.1093/brain/103.2.221

Roth, A., and van Rossum, M. C. W. (2009). "Modeling synapses," in *Computational Modeling Methods for Neuroscientists*, ed E. De Schutter (Cambridge, MA: MIT Press), 139–159. doi: 10.7551/mitpress/9780262013277.003.0007

Scheff, S. W., Price, D. A., Schmitt, F. A., and Mufson, E. J. (2006). Hippocampal synaptic loss in early Alzheimer's disease and mild cognitive impairment. *Neurobiol. Aging* 27, 1372–1384. doi: 10.1016/j.neurobiolaging.2005.09.012

Schneider, C. A., Rasband, W. S., and Eliceiri, K. W. (2012). NIH Image to ImageJ: 25 years of image analysis. *Nat. Methods* 9, 671–675. doi: 10.1038/nmeth.2089

Scott, A. C. (1975). The electrophysics of a nerve fiber. *Rev. Mod. Phys.* 47, 487–533. doi: 10.1103/revmodphys.47.487

Sheng, M., and Hoogenraad, C. C. (2007). The postsynaptic architecture of excitatory synapses: a more quantitative view. *Ann. Rev. Biochem.* 76, 823–847. doi: 10.1146/annurev.biochem.76.060805.160029

Shepherd, J. D., and Huganir, R. L. (2007). The cell biology of synaptic plasticity: AMPA receptor trafficking. *Ann. Rev. Cell Dev. Biol.* 23, 613–643. doi: 10.1146/annurev.cellbio.23.090506.123516

Spacek, J., and Harris, K. M. (1997). Three-dimensional organization of smooth endoplasmic reticulum in hippocampal CA1 dendrites and dendritic spines of the immature and mature rat. *J. Neurosci.* 17, 190–203.

Tai, C. Y., Kim, S. A., and Schuman, E. M. (2008). Cadherins and synaptic plasticity. *Curr. Opin. Cell Biol.* 20, 567–576. doi: 10.1016/j.ceb.2008.06.003

Tiwari, J. K., and Sikdar, S. K. (1999). Temperature-dependent conformational changes in a voltage-gated potassium channel. *Eur. Biophys. J.* 28, 338–345. doi: 10.1007/s002490050216

Tyagarajan, S. K., and Fritschy, J.-M. (2014). Gephyrin: a master regulator of neuronal function. *Nat. Rev. Neurosci.* 15, 141–156. doi: 10.1038/nrn3670

Vierling-Claassen, D., Cardin, J. A., Moore, C. I., and Jones, S. R. (2010). Computational modeling of distinct neocortical oscillations driven by cell-type selective optogenetic drive: separable resonant circuits controlled by low-threshold spiking and fast-spiking interneurons. *Front. Hum. Neurosci.* 4:198. doi: 10.3389/fnhum.2010.00198

Vogel, A., Reiter, S., Rupp, M., Nägel, A., and Wittum, G. (2013). UG 4: a novel flexible software system for simulating PDE based models on high performance computers. *Comput. Vis. Sci.* 16, 165–179. doi: 10.1007/s00791-014-0232-9

Wanner, A. (2007). *Ein effizientes Verfahren zur Berechnung der Potentiale in kortikalen neuronalen Kolumnen*. diploma thesis, Fakultät für Mathematik und Informatik, Universität Heidelberg.

Williams, T., Kelley, C., and many others (2010). *Gnuplot 4.4: An Interactive Plotting Program*. Available online at: http://gnuplot.sourceforge.net/

Wittmann, M., Queisser, G., Eder, A., Wiegert, J. S., Bengtson, C., Hellwig, A., et al. (2009). Synaptic activity induces dramatic changes in the geometry of the cell nucleus: interplay between nuclear structure, histone H3 phosphorylation, and nuclear calcium signaling. *J. Neurosci.* 29, 14687–14700. doi: 10.1523/JNEUROSCI.1160-09.2009

Wolf, S., Grein, S., and Queisser, G. (2013). Employing NeuGen 2.0 to automatically generate realistic morphologies of hippocampal neurons and neural networks in 3D. *Neuroinformatics* 11, 137–148. doi: 10.1007/s12021-012-9170-1

Xylouris, K., Queisser, G., and Wittum, G. (2007). A three-dimensional mathematical model of active signal processing in axons. *Comput. Vis. Sci.* 13, 409–418. doi: 10.1007/s00791-011-0155-7

Zeltser, L. M., Seeley, R. J., and Tschop, M. H. (2012). Synaptic plasticity in neuronal circuits regulating energy balance. *Nat. Neurosci.* 15, 1337–1342. doi: 10.1038/nn.3219

Real-World-Time Simulation of Memory Consolidation in a Large-Scale Cerebellar Model

Masato Gosui[1] and Tadashi Yamazaki[1, 2, 3]*

[1] Department of Communication Engineering and Informatics, Graduate School of Informatics and Engineering, The University of Electro-Communications, Tokyo, Japan, [2] Neuroinformatics Japan Center, RIKEN Brain Science Institute, Saitama, Japan, [3] Artificial Intelligence Research Center, National Institute of Advanced Industrial Science and Technology, Ibaraki, Japan

We report development of a large-scale spiking network model of the cerebellum composed of more than 1 million neurons. The model is implemented on graphics processing units (GPUs), which are dedicated hardware for parallel computing. Using 4 GPUs simultaneously, we achieve realtime simulation, in which computer simulation of cerebellar activity for 1 s completes within 1 s in the real-world time, with temporal resolution of 1 ms. This allows us to carry out a very long-term computer simulation of cerebellar activity in a practical time with millisecond temporal resolution. Using the model, we carry out computer simulation of long-term gain adaptation of optokinetic response (OKR) eye movements for 5 days aimed to study the neural mechanisms of posttraining memory consolidation. The simulation results are consistent with animal experiments and our theory of posttraining memory consolidation. These results suggest that realtime computing provides a useful means to study a very slow neural process such as memory consolidation in the brain.

Keywords: cerebellum, model, memory consolidation, optokinetic response, realtime simulation, graphics processing unit

1. INTRODUCTION

Memory formation has two stages: memory acquisition and memory consolidation (Dudai, 2004). A single session of training forms a type of memory which is fragile and persists only a short period up to minutes to hours. This phase is called memory acquisition. After the training, the learned memory, a short-term memory, decays spontaneously and quickly within a day. Meanwhile, repeated training with a sufficient rest between training sessions gradually form another type of memory, a long-term memory, which is robust and persists for days and weeks. This phase is called memory consolidation. Memory consolidation occurs after training but not during training. That is, when we take a rest after training, the brain still continues working to consolidate the learned memory. This posttraining memory consolidation is thought to be the basis of spacing effect (Ebbinghaus, 1885), in which a massed training is inferior to repeated training to form a robust long-term memory, even if the total training time is equal. Therefore, it is important to study how the brain works after training as well as during training to elucidate the memory mechanisms and our behaviors.

In cerebellar motor learning, both memory formation and consolidation occur within the cerebellum. In gain adaptation of vestibulo-ocular reflex (VOR) and optokinetic response (OKR), parallel fiber-Purkinje cell (PF-PC) synapses in the cerebellar cortex store short-term memory,

whereas mossy fiber-vestibular nuclear cell (MF-VN) synapses in the brain stem store long-term memory (Kassardjian et al., 2005; Shutoh et al., 2006). OKR is an oculomotor reflex in which the eye moves to the same direction of the visual world's movement to reduce the slip of the retinal image. In OKR adaptation, the amplitude of eye movement, called gain, changes by training. By a single 1-h training, the gain increases quickly, which corresponds to memory acquisition. After the training, the gain decreases naturally to the original level within a day. By repeating the 1-h training everyday, the gain increases gradually throughout 1 week (Shutoh et al., 2006), which represents memory consolidation. Moreover, injection of muscimol, a γ-Aminobutyric acid (GABA) receptor agonist, to the cerebellar cortex immediately after the training disrupts memory consolidation (Okamoto et al., 2011), indicating that training alone is not sufficient for memory consolidation. Accumulating evidence suggests that posttraining memory consolidation of OKR gain takes the following steps (Shutoh et al., 2006). By a single 1-h training, PF-PC synapses undergo LTD induced by conjunctive activation of PFs and the CF innervated to the same PCs (Ito, 1989), and thereby the OKR gain increases. After the training, PFs gradually recover from the LTD, which erase the memory of learned OKR gain in the cortex. On the other hand, because inhibition exerted by PCs to VN is weakened due to the LTD, the VN is deporalized tonically. This deporalization, combined with presynaptic MF activation, induces LTP at MF-VN synapses (McElvain et al., 2010; Person and Raman, 2010), and thereby forming the memory of OKR gain in the nucleus. In this way, while the cortical memory is erased gradually after the training, the nuclear memory forms simultaneously as a long-term memory, as if the learned cortical memory is transferred to the nucleus and consolidated there.

We have proposed a theory of the cerebellar posttraining memory consolidation in OKR adaptation (Yamazaki et al., 2015). The theory captures an essence of the macroscopic dynamics of synaptic mechanisms underlying the posttraining memory consolidation. On the other hand, the theory does not provide insights on mesoscopic cellular/synaptic dynamics on the posttraining memory consolidation. For example, the theory does not tell us about spatiotemporal spike patterns of individual neurons. To study the detailed cellular/synaptic dynamics, an elaborated, realistic cerebellar model is necessary. A problem of such elaborated models, however, is that they would spend too much computational time. Typically, computer simulation of large-scale spiking network models is 10–100 times slower than the real-world time (Nageswaran et al., 2009). This means, if we wanted to carry out a computer simulation of memory consolidation for 1 week, and the computer simulation was 100 times slower than real time, the simulation would spend about 2 years in total to complete. This is practically impossible.

In this study, we adopted high-performance computing (HPC) technology to solve these problems. We used graphics processing units (GPUs) to calculate equations of neurons in parallel, which could speed up the numerical simulation drastically. Specifically, we built a very large-scale spiking network model of the cerebellum composed of 1 million neurons, which is a model of 1 mm^3 of cats' cerebellum. Moreover,

owing to the parallel computing on GPUs, we were able to conduct the computer simulation fast enough to complete a very long computer simulation in a practical time, Eventually, we achieved realtime simulation, which means that computer simulation of the cerebellar activity for 1 s completes within 1 s in the real-world time (Igarashi et al., 2011; Yamazaki and Igarashi, 2013). This is essential for computer simulation of the cerebellar posttraining memory consolidation, because the memory consolidation takes days or even weeks. Using the present cerebellar model, we performed computer simulation of long-term OKR adaptation of training for 5 days, and obtained qualitatively the same results with experiments (Shutoh et al., 2006) and our previous theoretical model (Yamazaki et al., 2015). We also examined the detailed spike patterns of neurons, which was abstracted and therefore ignored in our theory.

2. MATERIALS AND METHODS

2.1. Model

Our cerebellar model is built based on a 1 mm^3 of the cerebellar corticonuclear microcomplex (**Figure 1**) of cats, which is thought to be a functional module of the cerebellum (Ito, 1984, 2012). The original model had 100,000 granule cells, which is 10 times smaller than cats' cerebellum (Ito, 1984), and was already reported elsewhere (Yamazaki and Tanaka, 2007; Yamazaki and Nagao, 2012; Yamazaki and Igarashi, 2013). In this study, we extended the previous model as follows. First, the present model includes 1 million granule cells, thereby the model includes the same number of neurons with 1mm^3 of the cats' cerebellum. Second, the present model has synaptic plasticity at mossy fiber-vestibular nuclear cell (MF-VN) synapses, as well as parallel fiber-Purkinje cell (PF-PC) synapses. Except the number of granule cells and MF-VN synaptic plasticity, the previous and present models are the same. Therefore, we summarize the model specification only briefly below. The details are found in our previous papers (Yamazaki and Tanaka, 2007; Yamazaki and Nagao, 2012; Yamazaki and Igarashi, 2013).

The present model is composed of 1,048,576 (= 1024 × 1024) granule cells, 1024 Golgi cells, 16 PCs, 16 basket cells, 1 inferior olivary cell, and 1 VN, connected according to cats anatomical data (Yamazaki and Tanaka, 2007). Neurons are modeled as conductance-based integrate-and-fire units (Gerstner et al., 2014):

$$C\frac{du}{dt} = -g_{leak}(u(t) - E_{leak}) - g_{ex:AMPA}(t)(u(t) - E_{ex})$$
$$- g_{ex:NMDA}(t)(u(t) - E_{ex}) - g_{inh:GABA}(t)(u(t) - E_{inh})$$
$$- g_{ahp}(t)(u(t) - E_{ahp}) + I_{ext}(t), \quad (1)$$

where $u(t)$ is the membrane potential at time t, C is the capacitance, g_{leak} and E_{leak} are the conductance and reversal potential of the leak current, respectively, $g_{ex:AMPA}(t)$, $g_{ex:NMDA}(t)$, $g_{inh:GABA}(t)$ are synaptic conductances of excitatory α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid (AMPA) and N-methyl-D-aspartic acid (NMDA), and inhibitory GABA synapses, E_{ex} and E_{inh} are reversal potentials, $g_{ahp}(t)$ and E_{ahp} are the conductance and reversal potential of after-hyper
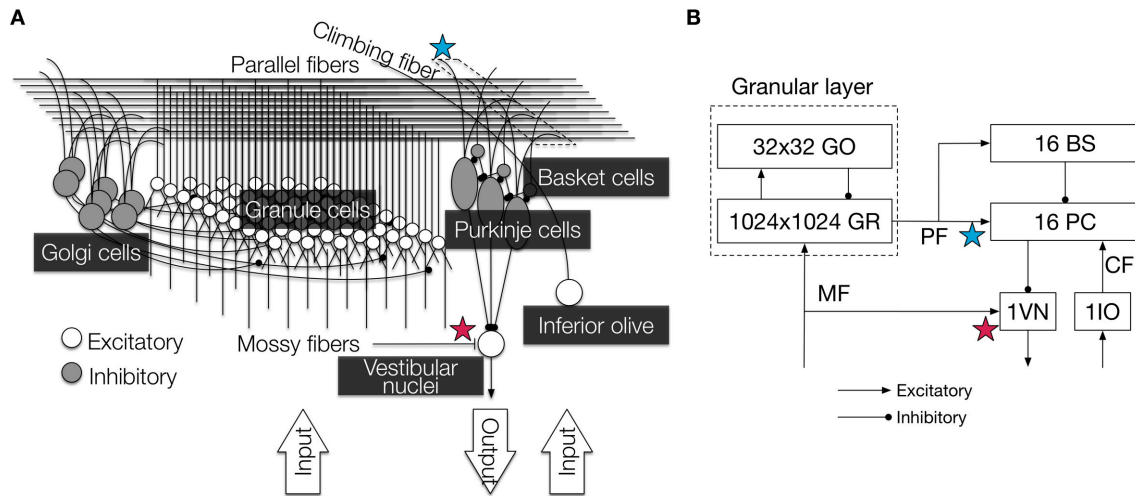
FIGURE 1 | Illustration of the cerebellar circuit implemented in this study. (A) Detailed 3D diagram with locations of synaptic plasticity. We implemented six major types of neurons: granule cells, Golgi cells, PCs, basket cells, inferior olivary cell, and VN. They were connected according to anatomical data, and cell parameters were taken from electrophysiological data. Contextual information and error information were conveyed by MFs and a CF, respectively, and the VN provided the final output of the circuit. PF-PC synapses (blue star) and MF-VN synapses (red star) underwent plastic change. The figure was reproduced from Yamazaki and Igarashi (2013). **(B)** 2D diagram of the connectivity and the number of neurons. PF-PC synapses (blue star) and MF-VN synapses (red star) underwent plastic change as in **(A)**. Arrowheads represent types of synaptic connections (triangle, excitatory; circle, inhibitory). GR, granule cell; GO, Golgi cell; PC, Purkinje cell; BS, basket cell; VN, vestibular nucleus; IO, inferior olive; MF, mossy fiber; CF, climbing fiber; PF, parallel fiber.

polarization, respectively, and $I_{ext}(t)$ is an external current. When $u(t)$ exceeds a threshold θ_{spike} at time t, the neuron elicits a spike at time t. Cell parameters are taken from turtles and rodents electrophysiological data (Yamazaki and Tanaka, 2007). The values used in this study are summarized in **Table 1**. Synaptic conductance $g_x(t)$ for type x is calculated as a convolution of presynaptic spike events with an exponential kernel as

$$g_x(t) = \sum_j w_j \cdot \bar{g}_x \sum_{f \in S_j} \exp_x\left(-\left(t - t^{(f)}\right)\right) \Theta\left(t - t^{(f)}\right), \quad (2)$$

where $\bar{g}_x$ is the peak conductance, w_j is the synaptic weight which is constant, S_j is the set of spikes elicited by presynaptic cell j, $t^{(f)}$ is the spike time for the fth spike, $\exp_x(t)$ is the exponential kernel, and $\Theta(t)$ is the Heaviside step function. The exponential kernels used in the present study are summarized in **Table 2**, whereas the synaptic weights are shown in **Table 3**.

The model has two distinct synaptic plasticity sites. One is PF-PC synapses, which undergo long-term depression (LTD) by conjunctive activation of granule-cell axons called parallel fibers (PFs) and a climbing fiber (CF) innervating to the same PC (Ito, 1989), and long-term potentiation (LTP) as well by sole activation of PFs (Lev-Ram et al., 2003). We modeled these bidirectional plasticity as follows:

$$\tau_w \frac{dw_{ij}}{dt} = -w_{ij}(t) + x_{ij}(t)$$

$$\tau_x \frac{dx_{ij}}{dt} = -x_{ij}(t) - c_{LTD} \sum_{s=0}^{50ms} PF_{ij}(t-s)CF(t) + c_{LTP}PF_{ij}(t),$$

$$(3)$$

TABLE 1 | Summary of cell parameters.

Parameter	Neuron type					
	GR	GO	PC	BS	VN	IO
θ_{spike} (mV)	−35.0	−52.0	−55.0	−55.0	−38.8	−50.0
C (pF)	3.1	28.0	107.0	107.0	122.3	10.0
g_{leak} (nS)	0.43	2.3	2.32	2.32	1.63	0.67
E_{leak} (mV)	−58.0	−55.0	−68.0	−68.0	−56.0	−60.0
$\bar{g}_{ex:AMPA}$ (nS)	0.18	45.5	0.7	0.7	50.0	1.0
$\bar{g}_{ex:NMDA}$ (nS)	0.025	30.0	-	-	25.8	-
E_{ex} (mV)	0	0	0	0	0	0
$\bar{g}_{inh}$ (nS)	0.028	-	1.0	-	30.0	0.18
E_{inh} (mV)	−82.0	-	−75.0	-	−88.0	−75.0
$\bar{g}_{ahp}$ (nS)	1.0	20.0	0.1	0.1	50.0	1.0
E_{ahp} (mV)	−82.0	−72.7	−70.0	−70.0	−70.0	−75.0
τ_{ahp} (ms)	5.0	5.0	5.0	2.5	2.5	10.0
I_{ext} (nA)	-	-	0.25	0.1	0.8	−

GR, granule cell; GO, Golgi cell; PC, Purkinje cell; BS, basket cell; VN, vestibular nuclear neuron; IO, inferior olivary neuron; -, nonexistent.

where $w_{ij}(t)$ is the synaptic weight between PC i and PF j, τ_w and τ_x are time constants where $\tau_w \ll \tau_x$, $x_{ij}(t)$ is an internal variable, c_{LTD} and c_{LTP} are constants, $PF_{ij}(t)$ is 1 if PF j on PC i elicits a spike at time t and 0 otherwise, and $CF(t)$ is 1 if the climbing fiber elicits a spike at time t and 0 otherwise. The term $\sum_{s=0}^{50ms} PF_{ij}(t-s)CF(t)$ means that PFs that elicit spikes 0–50 ms earlier than the time when the climbing fiber elicits a spike undergo LTD. If n spikes travel along a PF during 50 ms, the weight change becomes n times c_{LTD}. Transmission delay

TABLE 2 | Summary of synaptic functions.

Neuron type	Equation
GR	$\exp_{\text{ex:AMPA}}(t) = e^{-t/1.2}$
	$\exp_{\text{ex:NMDA}}(t) = e^{-t/52.0}$
	$\exp_{\text{inh}}(t) = 0.43 \times e^{-t/7.0} + 0.57 \times e^{-t/59.0}$
GO	$\exp_{\text{ex:AMPA}}(t) = e^{-t/1.5}$
	$\exp_{\text{ex:NMDA}}(t) = 0.33 \times e^{-t/31.0} + 0.67 \times e^{-t/170.0}$
PC	$\exp_{\text{ex:AMPA}}(t) = e^{-t/8.3}$
	$\exp_{\text{inh}}(t) = e^{-t/10.0}$
BS	$\exp_{\text{ex:AMPA}}(t) = e^{-t/8.3}$
VN	$\exp_{\text{ex:AMPA}}(t) = e^{-t/9.9}$
	$\exp_{\text{ex:NMDA}}(t) = e^{-t/30.6}$
	$\exp_{\text{inh}}(t) = e^{-t/42.3}$
IO	$\exp_{\text{ex:AMPA}}(t) = e^{-t/10.0}$

Abbreviations as in **Table 1**.

TABLE 3 | Summary of synaptic weights.

Presynaptic neuron	Postsynaptic neuron					
	MF	GR	GO	PC	BS	VN
MF	-	4.0	-	-	-	0.2
GR	-	-	0.00004	0.00075	0.00015	-
GO	-	10.0	-	-	-	-
PC	-	-	-	-	-	0.05
BS	-	-	-	5.3	-	-
IO	-	-	-	1.0	-	-

Abbreviations as in **Table 1**.

of PF spikes might be essential for plasticity (Knoblauch et al., 2014). The conduction velocity of PFs has been experimentally estimated as 0.24 m/s (Vranesic et al., 1994). This results in the transmission delay of 1 mm PF is maximally about 4.2 ms, which could be negligible as long as we assume 50 ms time window for LTD. Therefore, we did not model transmission delays of PF spikes. On the other hand, we do not exactly describe the biological counterpart of x_{ij}. A potential interpretation of x_{ij} would be intracellular concentration of some kinases involving PKC-MAPK positive feedback loop, which plays an essential role in maintenance of induced LTD (Kuroda et al., 2001). The initial values of w and x were set at 1.0 and 0.0, respectively.

The other plasticity is MF-VN synapses, which undergo bidirectional plasticity by a modified Hebbian mechanism. The original equation was proposed by our previous theoretical model (Yamazaki et al., 2015) based on Zhang and Linden (2006); Person and Raman (2010); McElvain et al. (2010) as follows:

$$\tau_v \frac{dv}{dt} = -v(t)\langle \text{MF}(t) \rangle + \langle \text{MF}(t)\left(\text{VN}(t) - \theta(t)\right)\rangle, \quad (4)$$

where τ_v is time constant, $v(t)$ is the synaptic weight at MF-VN synapses at time t, $\text{MF}(t)$ is the activity of MFs, $\text{VN}(t)$ is the activity of VN, $\langle \cdot \rangle$ is the temporal average over a certain time window (we assumed 6 s), and $\theta(t)$ is a running average

TABLE 4 | Summary of learning parameters.

Parameter	Value
τ_w (min)	20.0
τ_x (min)	240.0
c_{LTD}	0.005
c_{LTP}	0.1

of $\text{VN}(t)$, namely $\theta(t) = \langle \text{VN}(t) \rangle$. The left-hand side represents the temporal increment of $v(t)$. The 1st term in the right-hand side represents LTD by sole activation of MFs, and 2nd term represents the Hebbian mechanism, where the weight change correlates with the correlated activity of pre- and postsynaptic neurons. Here, the term $\theta(t)$ acts as a threshold; only when the postsynaptic neuron is activated strongly to exceed $\theta(t)$, the synapses undergo LTP, otherwise LTD or no change. In this way, $\theta(t)$ determines the direction of synaptic change. Moreover, the value of $\theta(t)$ itself changes temporally depending on the temporal history of $\text{VN}(t)$. Higher $\theta(t)$ value makes the synapse harder to undergo LTP. The initial value of v was set at 1.0. The parameters for w and v are summarized in **Table 4**.

As far as we have tested, the general network dynamics does not change so largely over a wide range of parameter settings. We have found three points that are necessary to achieve robust learning. First, granule-Golgi cell recurrent network should be tuned so as to generate the population code of granule cells robustly. Second, basket cell→ PC synaptic connections should not be so strong; otherwise, PCs would be silent completely. Third, PC→VN synaptic connections should not be so strong; otherwise, VN would be silent completely. If we satisfy these three points, the network, as far as we have tested, works robustly.

2.2. Simulation Paradigm

We conducted computer simulation of long-term OKR gain adaptation as in Shutoh et al. (2006). Specifically, we repeated a 1-h simulated OKR training followed by 23-h rest 5 times corresponding to 5-days training. In each OKR training, simulated optokinetic stimulus is fed to MFs, and retinal slip is fed to a CF. Both optokinetic stimulus and retinal slip are modeled as Poisson spikes with the following firing rate:

$$f_{\text{MF}_{\text{train}}}(t) = \overline{\text{MF}_{\text{train}}}\left(1 + \sin\frac{2\pi t}{T}\right) \quad \text{(for MFs)}$$

$$f_{\text{CF}_{\text{train}}}(t) = \overline{\text{CF}_{\text{train}}}\left(1 + \sin\frac{2\pi t}{T}\right) \quad \text{(for a CF),}$$

$$(5)$$

where $f_{\text{MF}_{\text{train}}}(t)$ and $f_{\text{CF}_{\text{train}}}(t)$ are the firing rate of MFs and a CF, respectively, $\overline{\text{MF}_{\text{train}}}$ and $\overline{\text{CF}_{\text{train}}}$ are the mean activity of MFs and a CF, which are set at 15 spikes/s and 1.5 spikes/s, respectively. T is a period of a cycle of optokinetic stimulus, which is assumed to be rotated sinusoidally in front of animal subjects. We set $T = 6$ s consistently with the experiments (Shutoh et al., 2006). Because one cycle is 6 s, daily 1-h training consists of 600 cycles of simulated optokinetic stimulus. On the other hand, after training,

we assumed that both MFs and a CF elicited spikes spontaneously with the following firing rate:

$$f_{\mathrm{MF_{rest}}}(t) = \overline{\mathrm{MF_{rest}}} \quad \text{(for MFs)}$$
$$f_{\mathrm{CF_{rest}}}(t) = \overline{\mathrm{CF_{rest}}} \quad \text{(for a CF)}, \tag{6}$$

where $\overline{\mathrm{MF_{rest}}}$ and $\overline{\mathrm{CF_{rest}}}$ are set at 5 spikes/s and 1 spikes/s, respectively.

Once we define the firing rate of MF and CF as above, and assume that the activity of a simulated neuron (e.g., firing rate) reflects the strength of input signals to the neuron almost linearly as in the case of integrate-and-fire neurons used in this study (Gerstner et al., 2014), we could estimate the activity of VN as a linear sum of excitatory MF activity and inhibitory PC activity. The PC activity could be estimated as a linear sum of PF activity and basket cell activity, and further by solely MF activity. By substituting the MF and VN activities for MF(t) and VN(t) in Equation (4), we could obtain the following simplified equation for v. The detailed derivation is found in our previous paper (Yamazaki et al., 2015).

$$\tau_v \frac{dv}{dt} = -w(t) + w_c, \tag{7}$$

where $w(t)$ is the average synaptic weight of all PF-PC synapses, and w_c is a constant that defines the initial weight of PF-PC synapses, namely, 1.0. We used Equation (7) rather than Equation (4) for simplicity to update $v(t)$.

2.3. Data Analysis

We conducted computer simulation of the 5-days OKR training, and obtained spike data of all individual neurons and synaptic weight data of PF-PC synapses and MF-VN synapses. The total simulation time was $5 \times 24 \times 60 \times 60 \times 1000 = 4.32 \times 10^8$ ms, with temporal resolution of 1 ms.

We analyzed how the OKR gain changed before and after training for each day. To do so, before training for each day, we fed 10 cycles of simulated optokinetic stimulus to the network and obtained the spike data of VN. We made a spike histogram with bin size of 100 ms, fitted the data with a cosine function with the period of 6 s, and calculated the modulation amplitude. We defined the modulation amplitude as the OKR gain before training. We made the same procedure to obtain the OKR gain after training for each day.

We also examined how the granule cells transmit mossy fiber signals robustly against noise in Poisson spike trains. Granule cells must produce almost identical spike pattern in response to the same optokinetic stimulus with different noise across cycles; otherwise, learning at Purkinje cells would fail. To quantify the reproducibility of the granule cell spike pattern in response to the same simulated optokinetic stimulus, we calculated the reproducibility index at time t defined as the normalized cross correlation as follows:

$$R(t) = \frac{\sum_j z_j^{(i)}(t) z_j^{(i+1)}(t)}{\sqrt{\sum_j z_j^{(i)}(t)} \sqrt{\sum_j z_j^{(i+1)}(t)}}, \tag{8}$$

where $z_j^{(i)}(t)$ is the activity of granule cell j at cycle i of simulated optokinetic stimulus at time t, which was calculated by convolution of the spikes with a causal exponential:

$$z_j^{(i)}(t) = \sum_{f \in S_j^{(i)}} \exp\left(-\frac{t - t^{(f)}}{\tau}\right) \Theta\left(t - t^{(f)}\right), \tag{9}$$

where $S_j^{(i)}$ is the set of spikes elicited by granule cell j at cycle i, $t^{(f)}$ is the spike time for the fth spike, τ is the time constant of 8.3 ms, and $\Theta(t)$ is the Heaviside step function. Intuitively, $z_j^{(i)}(t)$ is a temporal trace of EPSPs of PF j on a PC at cycle i, and $\tau = 8.3$ ms is the time constant of AMPA receptor-mediated PF-EPSPs at a PC (Llano et al., 1991). We calculated the average and standard deviation of the reproducibility index among 10 pairs of two successive cycles.

2.4. Numerical Method

All equations that govern the network dynamics are solved numerically. Specifically, differential equations describing membrane potentials are solved by 2nd-order Runge-Kutta method with temporal resolution (Δt) of 1 ms. The simulation program is written in C with CUDA (Common Unified Device Architecture) (NVIDIA, 2015) and most of the calculation is made on GPUs.

In our previous study (Yamazaki and Igarashi, 2013), we used only 1 GPU (NVIDIA GeForce GTX580) to simulate 100,000 granule cells in realtime. On the other hand, the present model has 10 times more granule cells, which makes computer simulation far slower than realtime. The most time-consuming part is to calculate synaptic conductances of Golgi cells, basket cells and PCs, where these cells receive excitatory inputs from granule cells via PFs. Due to the large number of granule cells, the calculation spends too much time. To address this issue, we decomposed the granular layer network composed of granule cells and Golgi cells into 4 identical subnetworks and calculated the dynamics in parallel on 4 GPUs (2 boards of NVIDIA GeForce GTX TITAN Z, each contains 2 GPUs). In the following, we explain how to decompose the network and calculate the conductance on 4 GPUs.

Figure 2A illustrates a part of the granular layer of our model. The granular layer is composed of 1024×1024 granule cells and 32×32 Golgi cells arranged regularly on a two-dimensional grid. Granule cells are further divided as 32×32 clusters, where each cluster consists of 32×32 granule cells. Due to short dendrites of granule cells, we assumed that the granule cells in the same cluster shared inhibitory inputs from the same Golgi cells. On the other hand, granule cells receive 4 excitatory MF inputs. We assumed that granule cells receive 4 MF inputs independently of the other granule cells. This structure allows us to decompose the granular layer network into 4 identical subnetworks composed of 512×512 granule cells and 32×32 Golgi cells, where granule cells are further divided into 32×32 clusters in which each cluster contained $16 \times 16 = 256$ granule cells as shown in **Figure 2B**. We conducted simulation of each subnetwork on a GPU, thereby we employed 4 GPUs for
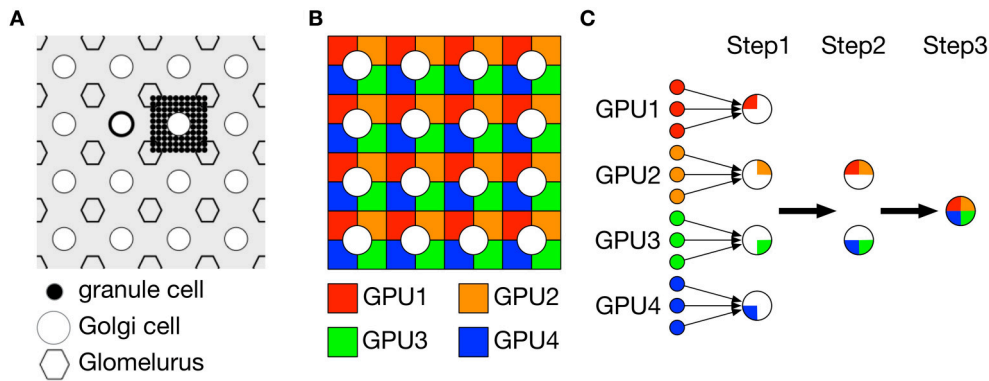
FIGURE 2 | Decomposition of granule cell population into four subpopulations for parallel simulation on 4 GPUs. (A) Schematic of the model granular layer. Granule cells (black dot), Golgi cells (circle) and fictious glomeruli (hexagon) were arranged regularly on a two-dimensional grid. Each Golgi cell was surrounded by 4 glomeluri, and within the rectangle, $32 \times 32 = 1024$ granule cells were located and constituted a granule-cell cluster. These granule cells were assumed to share the same inhibitory inputs from Golgi cells and receive excitatory inputs for 4 independent mossy fibers, so the granule cells were functionally identical. **(B)** Decomposition of granule-cell clusters. We decomposed each granule-cell cluster composed of $32 \times 32 = 1024$ granule cells into 4 subclusters composed of $16 \times 16 = 256$ granule cells shown by 4 colors. **(C)** Calculation of synaptic conductances for Golgi cells, PCs and basket cells from granule cells (left small dots). Each postsynaptic neuron must sum up the postsynaptic potentials of all granule cells with certain synaptic weights. At step 1, for each postsynaptic neuron, a quarter of the synaptic conductance was calculated from granule cells on each GPU, which was illustrated by a pie-shape color where a large circle represents a postsynaptic neuron. At step 2, the calculated partial conductances were reduced between 2 GPUs in parallel to obtain a half of the conductance. At step 3, the partial conductances were further reduced and the full conductance was calculated. **(A)** was reproduced from Yamazaki and Tanaka (2007).

simulation of 4 subnetworks. In each subnetwork, we calculated quarter of synaptic conductance for each Golgi cell from granule cells in the same subnetwork. We then exchanged the partial conductances across subnetworks over GPUs and obtained the full conductance by summing up the partial values (**Figure 2C**). This is made by direct memory exchange between 2 GPUs called peer access, which is much faster than conventional memory exchange via CPUs. Because calculation of synaptic conductance is linear, our split-reduction method over 4 GPUs provides the same result with the conventional method. The same method is used to calculate synaptic conductances of basket cells and PCs as well.

3. RESULTS

3.1. Simulation Time

First, we measured how the simulation time was accelerated by using multi GPUs. Using only 1 GPU, we found that computer simulation of the cerebellar activity for 6 s, corresponding to 1 cycle of simulated optokinetic stimulus, spends 17.7 s. Using 2 GPUs, 9.10 s are spent. Finally, using 4 GPUs, we achieved 5.33 s for 6 s simulation, indicating realtime simulation. Therefore, we used 4 GPUs for further simulation.

3.2. Long-term OKR Gain Change

We conducted computer simulation of long-term OKR adaptation for 5 days. For each day, we performed a simulated 1-h OKR training. During the training, MFs convey simulated optokinetic stimuli, whereas a CF conveys simulated retinal slip error signals. After the training, both MFs and the CF elicit Poisson spikes spontaneously with a constant firing rate, respectively.

Figure 3A plots the OKR gain obtained in our long-term OKR training simulation for 5 days. By daily 1-h training, the OKR gain increases during training, and after the training, the learned OKR gain almost disappears. This indicates memory acquisition of OKR gain. On the other hand, throughout the 5 days, OKR gain gradually increases, indicating memory consolidation. The present numerical result is qualitatively consistent with previous experimental and theoretical results (Shutoh et al., 2006; Yamazaki et al., 2015).

Figure 3B plots the daily increment of learned OKR gain by 1-h training. The increment becomes larger day by day, suggesting that repeated daily training accelerates the memory acquisition. This result is consistent with previous experiments (Shutoh et al., 2006).

3.3. Change of Synaptic Weights

Figure 4 plots the change of weights at PF-PC synapses (w) and MF-VN synapses (v) throughout the 5 days training. For w, we calculated the average of all PF-PC synaptic weights with respect to PFs and PCs. Similarly for v, we calculated the average of all MF-VN synaptic weights with respect to MFs. PF-PC synapses undergo LTD during training, and slowly return to the original weight value after training spontaneously. PF-PC synapses repeat the same temporal change 5 times for 5 days, suggesting that PF-PC synapses store only short-term memory of OKR gain for hours. On the other hand, MF-VN synapses change little during training, and slowly increase after training. The synaptic weight accumulates every day after training, suggesting that MF-VN synapses store long-term memory of OKR gain.

The overall dynamics is as follows. First, memory of OKR gain is formed in the cerebellar cortex by PF-PC LTD during training. Second, after training, learned cortical memory is
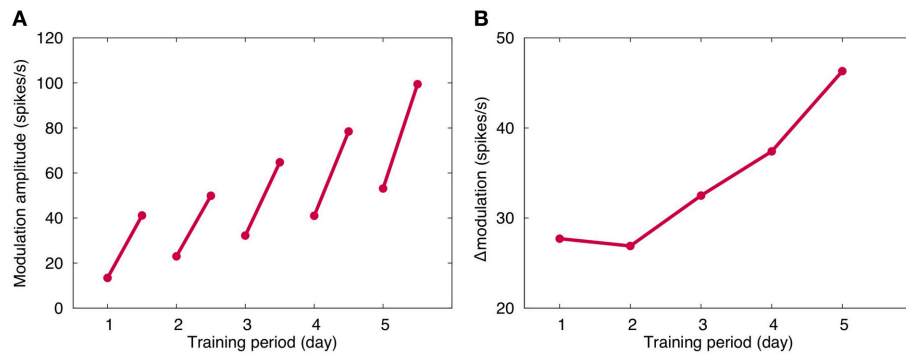
FIGURE 3 | Simulated OKR gain. (A) Modulation amplitude of VN, which corresponded to OKR gain. Horizontal axis represents time (day) and vertical axis the firing rate (spikes/s). For each day, the modulation amplitude increased by 1-h training, and decreased after training until the next training in the next day. Throughout the 5 days, the amplitude gradually increased, suggesting consolidation of memory of learned OKR gain. (B) Increase of OKR gain before and after daily 1-h training. Conventions as in (A). Although the same 1-h training was performed, the increase became larger day by day.
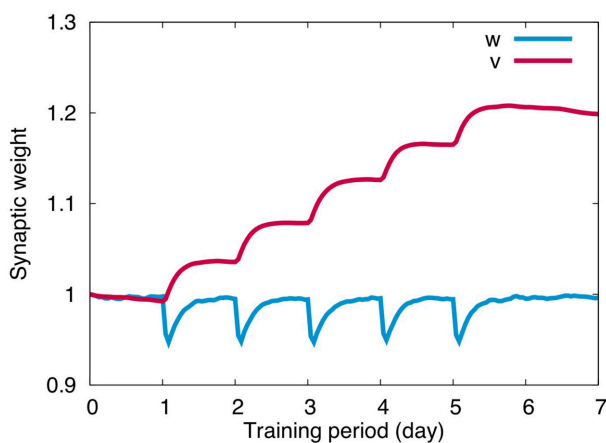
FIGURE 4 | Synaptic weight change. PF-PC synaptic weights (blue) repeated quick decrease during training and slow recovery after training for each day, whereas MF-VN synaptic weight started to increase mainly after training and was accumulated throughout 5 days. Horizontal axis represents training period (day) and the vertical axis represents the weight value. For PF-PC synaptic weights, the average value on all presynaptic granule cells and postsynaptic PCs was plotted.

decayed slowly and disappears completely by the next day, and finally, during the slow decay of the cortical memory, memory is formed in the vestibular nucleus by MF-VN LTP, as if the cortical memory is transferred to the nucleus and consolidated. The present numerical result is consistent with the previous theoretical results (Yamazaki et al., 2015).

3.4. Change of Eye Movement Trajectory

So far, both the current numerical and previous theoretical studies show qualitatively the same results. A benefit of our numerical study is that we could obtain detailed data of individual neurons such as membrane potential and spike trains with a fine temporal resolution of 1 ms, which were abstracted in our theoretical model (Yamazaki et al., 2015).

Figure 5 plots the firing rate of VN in response to simulated sinusoidal optokinetic stimulus before and after training at the 1st day (A) and the 5th day (B). The firing rate modulates sinusoidally as the input signals. The modulation amplitude increases by daily 1-h training, and the amplitude also increases gradually throughout 5 days. On the other hand, the baseline firing rate does not change largely from 30–50 spikes/s. Here, the modulation amplitude of VN represents the OKR gain (Shutoh et al., 2006), suggesting that the OKR gain becomes larger by repeated daily training. These results also suggest that realtime simulation allows us to study both macroscopic behaviors of a neural network such as OKR gain, and mesoscopic dynamics of individual neurons in the neural network such as a membrane potential and spike trains.

3.5. Robust Signal Transmission by the Enormous Number of Granule Cells

Granule cells must transmit information conveyed by mossy fibers to Purkinje cells and interneurons faithfully against input noise, otherwise, learning at Purkinje cells would fail. In OKR, mossy fibers convey information on visual world movement, and granule cells produce a spatiotemporal spike pattern that represents the stimulus reliably. For this purpose, the almost identical spike pattern of granule cells must be produced across cycles of the optokinetic stimulus.

Here, we examined how the enormous number of granule cells help them to transmit mossy fiber information faithfully and robustly against input noise. Specifically, we calculated the reproducibility index (Equation 8) that quantifies the reproducibility of the spike pattern of granule cells across cycles of the simulated optokinetic stimulus on different cycles, while changing the number of granule cells in the network.

Figure 6A plots an example of the spike pattern of granule cells, whereas Figure 6B plots the reproducibility. As can be seen, the reproducibility is better when 1 million granule cells were employed than 0.1 million granule cells. This result suggests that a functional role of the enormous number of granule cells is
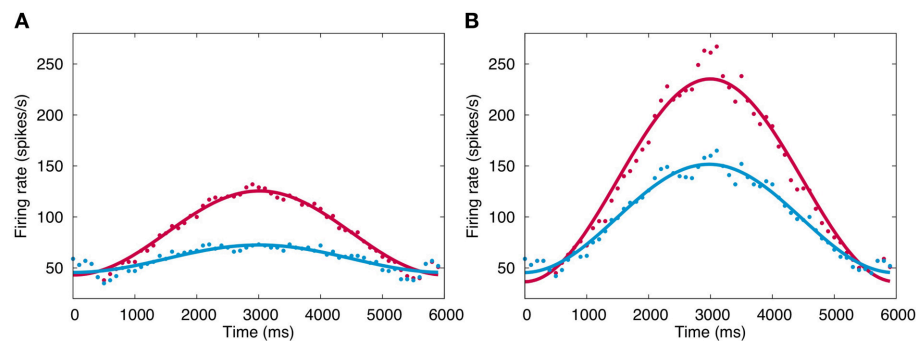
FIGURE 5 | Firing rate change of VN. (A) Firing rate in response to simulated optokinetic stimulus modulating sinusoidally in time with the period of 6 s before and after training (blue and red, respectively) in the 1st day. Horizontal axis represents time in ms, and vertical axis the firing rate (spikes/s). The data points are fitted with a cosine function and the fitted curves are also plotted to show the modulation clearly. **(B)** The same firing rate in the 5th day. Conventions as in **(A)**.
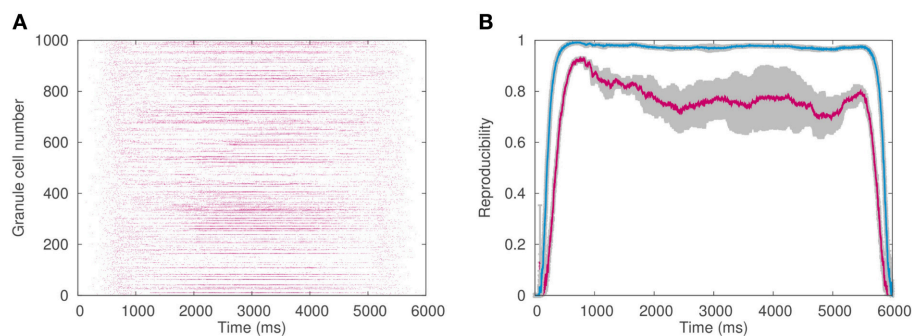
FIGURE 6 | Granule cell activity. (A) Spike pattern of 1024 out of 1,048,576 granule cells chosen randomly in response to simulated optokinetic stimulus, which modulated sinusoidally with the period of 6 s. Horizontal axis represents time (ms), whereas the vertical axis the neuron number. **(B)** Reproducibility of the granule-cell spike pattern across different cycles of the simulated stimulus. The reproducibility with 1 million granule cells (blue) was higher than that with 0.1 million cells (red). Horizontal axis represents time (ms), whereas the vertical axis the reproducibility index. The average values on 10 pairs of cycles were plotted in color, while gray regions represented the standard deviation.

robust transmission of mossy fiber signals to PCs against input noise.

4. DISCUSSION

4.1. Understanding Memory Consolidation Mechanisms

Memory consolidation is a slow process that takes days and weeks. To study the neural mechanisms of memory consolidation, two ways are possible: either conducting experiments or making theoretical models. A theoretical model is a mathematical description of a specific phenomenon. To make such model, we ignore most of experimental details and capture the essence of the phenomenon. For example, in our theoretical model of posttraining memory consolidation in the cerebellum (Yamazaki et al., 2015), we abstracted all detailed physiology of individual neurons, detailed anatomical structure, and detailed input stimuli. This provides a clear view of how the memory consolidates after training, but we still do not know the detailed neuronal process during the memory consolidation. Large-scale, realistic spiking network models are appropriate for

this purpose, but the computational time would be problematic instead.

HPC technology solves this problem. The advantage is two-folds. First, the technology allows us to build a larger-scale model composed of more neurons and synapses with more detailed morphology and biophysical properties than conventional models. Very large-scale functional brain models have been built (Izhikevich and Edelman, 2008; Eliasmith et al., 2012). Notably, The Blue Brain Project and Human Brain Project attempt to build a realistic whole brain model, and they recently published a very detailed cortical microcolumn model (Markram et al., 2015). Second, the technology allows us to carry out computer simulations much faster than that on a single-threaded CPU. The latter makes the above-mentioned long-term computer simulation possible in a reasonable time. For instance, if the computer simulation runs in real time, a simulation of memory consolidation for 1 week completes in 1 week. In our study, we adopted GPUs. Using our large-scale, detailed spiking network model of the cerebellum implemented on multi GPUs, we were able to simulate the detailed temporal dynamics of individual neurons, while observing the slow memory consolidation process

simultaneously. The present study is, as far as the author knows, a first demonstration of a very long time computer simulation of an elaborated spiking network model for days. We were able to replicate our previous theoretical results (Yamazaki et al., 2015), and further examined detailed neuronal and synaptic dynamics during memory consolidation. In cerebellar motor learning, location of motor memory and the role of LTD at PF-PC synapses have been a matter of debate for more than 30 years (Mellvill-Jones, 2000). The present study could provide an answer from the modeling view point.

4.2. Realtime Simulation and the Programming

The present cerebellar model consists of more than 1 million spiking neurons. In general, computer simulation of such large-scale model takes very long time. The simulation could be 10–100 times slower than the real-world time (Nageswaran et al., 2009). However, owing to HPC technology, we were able to conduct computer simulation in realtime, where simulation of cerebellar activity for 1 s completes within 1 s in the real-world time. This allowed us to conduct a complete computer simulation of long-term OKR adaptation training for 5 days in a practical time.

We used 4 GPUs simultaneously to perform realtime simulation of 1 million neurons. To do so, we had to write the simulation program in C with CUDA, a platform for GPU computing, and employed some parallel algorithms to use GPUs efficiently. Specifically, we used some algorithms to compute synaptic conductances of Golgi cells, basket cells and PCs that receive excitatory inputs from many granule cells. This is quite technical and difficult, and so there should be a more simple way to adopt the power of parallel computing in neuroscience. One potential way would be to develop a neural simulator primarily designed for GPUs and some other accelerators. Naveros et al. (2014) has reported development of such a spiking neuron network simulator on a GPU. Some groups have used the software for realtime robot control (Garrido et al., 2013; Casellato et al., 2015).

Realtime simulation is only a milestone, and we expect even faster computer simulation. Scalability, however, would be a problem. Generally speaking, using more GPUs would employ more latency for communications and overhead of communication operations, which could easily be a bottle neck.

4.3. Advantages of Large-Scale Models over Theoretical Models

Although the present study reproduced qualitatively the same results with our previous theoretical model (Yamazaki et al., 2015), some results are slightly different. First, the MF-VN synaptic weight in the present model tends to decay spontaneously, whereas that in the theoretical model did not. This is because in the theoretical model, the decay term was canceled out and removed by a mathematical treatment. In experiments (Shutoh et al., 2006), the learned long-term OKR gain almost vanishes after 2 weeks from the last training, suggesting that it is natural for the synaptic weight to decay spontaneously. Second, the increase of modulation amplitude

of VN before and after the 1-h training gradually becomes larger throughout 5 days in the current study (**Figure 2B**), whereas the change is constant in the theoretical model. The same experiments demonstrate that the increase becomes larger gradually day by day. This result suggests that the present large-scale model captures the detailed dynamics of long-term OKR gain adaptation better than the theoretical model.

Moreover, the present model allows us to study the detailed temporal dynamics of individual neurons with a fine temporal resolution. We were able to obtain detailed spike data of PCs and VN, and analyzed the firing patterns as in **Figure 6**. This is an advantage of an elaborated spiking network model over theoretical models, which abstract detailed temporal dynamics of individual neurons. We will be able to go into the details of molecular mechanisms of memory acquisition and consolidation (Abel and Lattal, 2001; Ito, 2002), if the HPC technology advances further.

We were also able to examine how the number of neurons could affect the stability of the network dynamics. In the present model, we incorporated more than 1 million granule cells, because the cats' cerebellum has 1 million granule cells per 1 mm^3 (Ito, 1984). The cerebellar granule cells constitute the largest population in the whole brain (Azevedo et al., 2009). A question arises: why does the cerebellum have such an enormous number of granule cells? A theoretical study has demonstrated that incorporating more granule cells makes the network more reliable for controlling hardware robots (Pinzon-Morales and Hirata, 2015). In the present study, we demonstrated that the enormous number of granule cells makes signal transmission from MFs to PFs more robust as in **Figure 6**.

In summary, combination of large-scale, detailed spiking network models with HPC technology for realtime simulation will provide a strong means to study mesoscopic, detailed neural mechanisms for macroscopic behavioral phenomenon that could take very long time for days and weeks such as memory formation.

4.4. Data Sharing

We will release the source code of the model used in this study under an opensource license upon publication, to facilitate open collaboration and ensure scientific reproducibility, on Cerebellar Platform (https://cerebellum.neuroinf.jp/).

REFERENCES

Abel, T., and Lattal, K. (2001). Molecular mechanisms of memory acquisition, consolidation and retrieval. *Curr. Opin. Neurobiol.* 11, 180–187. doi: 10.1016/S0959-4388(00)00194-X

Azevedo, F. A., Carvalho, L. R., Grinberg, L. T., Farfel, J. M., Ferretti, R. E., Leite, R. E., et al. (2009). Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *J. Comp. Neurol.* 513, 532–541. doi: 10.1002/cne.21974

Casellato, C., Antonietti, A., Garrido, J. A., Ferrigno, G., D'Angelo, E., and Pedrocchi, A. (2015). Distributed cerebellar plasticity implements generalized multiple-scale memory components in real-robot sensorimotor tasks. *Front. Comput. Neurosci.* 9:24. doi: 10.3389/fncom.2015.00024

Dudai, Y. (2004). The neurobiology of consolidations, or, how stable is the engram? *Ann. Rev. Psychol.* 55, 51–86. doi: 10.1146/annurev.psych.55.090902.142050

Ebbinghaus, H. (1885). *Memory: A Contribution to Experimental Psychology*. New York, NY: Dover.

Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., et al. (2012). A large-sale model of the functioning brain. *Science* 338, 1202–1205. doi: 10.1126/science.1225266

Garrido, J. A., Luque, N. R., D'Angelo, E., and Ros, E. (2013). Distributed cerebellar plasticity implements adaptable gain control in a manipulation task: a closed-loop robotic simulation. *Front. Neural Circuits* 7:159. doi: 10.3389/fncir.2013.00159

Gerstner, W., Kistler, W. M., Naud, R., and Paninski, L. (2014). *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge: Cambridge University Press. doi: 10.1017/cbo9781107447615

Human Brain Project (2013). Available online at: https://www.humanbrain project.eu/

Igarashi, J., Shouno, O., Fukai, T., and Tsujino, H. (2011). Real-time simulation of a spiking neural network model of the basal ganglia circuitry using general purpose computing on graphics processing units. *Neural Netw.* 24, 950–960. doi: 10.1016/j.neunet.2011.06.008

Ito, M. (1984). *The Cerebellum and Neural Control*. New York, NY: Raven Press.

Ito, M. (1989). Long-term depression. *Ann. Rev. Neurosci.* 12, 85–102. doi: 10.1146/annurev.ne.12.030189.000505

Ito, M. (2002). The molecular organization of cerebellar long-term depression. *Nat. Rev. Neurosci.* 3, 896–902. doi: 10.1038/nrn962

Ito, M. (2012). *The Cerebellum: Brain for the Implicit Self*. Upper Saddle River, NJ: FT Press.

Izhikevich, E. M., and Edelman, G. M. (2008). Large-scale model of mammalian thalamocortical systems. *Proc. Natl. Acad. Sci. U.S.A.* 105, 3593–3598. doi: 10.1073/pnas.0712231105

Kassardjian, C. D., Tan, Y. F., Chung, J. Y., Heskin, R., Peterson, M. J., and Broussard, D. M. (2005). The site of a motor memory shifts with consolidation. *J. Neurosci.* 25, 7979–7985. doi: 10.1523/JNEUROSCI.2215-05.2005

Knoblauch, A., Hauser, F., Gewaltig, M.-O., Körner, E., and Palm, G. (2014). Does spike-timing-dependent synaptic plasticity couple or decouple neurons firing in synchrony? *Front. Comput. Neurosci.* 6:55. doi: 10.3389/fncom.2012.00055

Kuroda, S., Schweighofer, N., and Kawato, M. (2001). Exploration of signal transduction pathways in cerebellar long-term depression by kinetic stimulation. *J. Neurosci.* 21, 5693–5702.

Lev-Ram, V., Mehta, S. B., Kleinfeld, D., and Tsien, R. Y. (2003). Reversing cerebellar long-term depression. *Proc. Natl. Acad. Sci. U.S.A.* 100, 15989–15993. doi: 10.1073/pnas.2636935100

Llano, I., Marty, A., Armstrong, C. M., and Konnerth, A. (1991). Synaptic- and agonist-induced excitatory currents of purkinje cells in rat cerebellar slices. *J. Physiol. (Lond.)* 434, 183–213. doi: 10.1113/jphysiol.1991.sp018465

Markram, H., Muller, E., Ramaswamy, S., Reimann, M. W., Abdellah, M., Sanches, C. A., et al. (2015). Reconstruction and simulation of neocortical microcircuitry. *Cell* 163, 456–492. doi: 10.1016/j.cell.2015.09.029

McElvain, L. E., Bagnall, M. W., Sakatos, A., and du Lac, S. (2010). Bidirectional plasticity gated by hyperpolarization controls the gain of postsynaptic firing

responses at central vestibular nerve synapses. *Neuron* 68, 763–775. doi: 10.1016/j.neuron.2010.09.025

Mellvill-Jones, G. (2000). "Posture," in *Principles of Neural Science, 4th Edn.*, eds E. Kandel, J. Schwartz, and T. Jessell (New York, NY: McGraw-Hill), 816–831.

Nageswaran, J. M., Dutt, N., Krichmar, J. L., Nicolau, A., and Veidenbaum, A. (2009). A configurable simulation environment for the efficient simulation of large-scale spiking neural networks on graphics processors. *Neural Netw.* 22, 791–800. doi: 10.1016/j.neunet.2009.06.028

Naveros, F., Luque, N. R., Garrido, J. A., Carrillo, R. R., Anguita, M., and Ros, E. (2014). A spiking neural simulator integrating event-driven and time-driven computation schemes using parallel cpu-gpu co-processing: a case study. *IEEE Trans. Neural Netw. Learn. Syst.* 26, 1567–1574. doi: 10.1109/TNNLS.2014.2345844

NVIDIA (2015). *CUDA* Programming Guide 7.0.

Okamoto, T., Shirao, T., Shutoh, F., Suzuki, T., and Nagao, S. (2011). Post-training cerebellar cortical activity plays an important role for consolidation of memory of cerebellum-dependent motor learning. *Neurosci. Lett.* 504, 53–56. doi: 10.1016/j.neulet.2011.08.056

Person, A. L., and Raman, I. M. (2010). Deactivation of L-type Ca current by inhibition controls LTP at excitatory synapses in the cerebellar nuclei. *Neuron* 66, 550–559. doi: 10.1016/j.neuron.2010.04.024

Pinzon-Morales, R.-D., and Hirata, Y. (2015). A realistic bi-hemispheric model of the cerebellum uncovers the purpose of the abundant granule cells during motor control. *Front. Neural Circuits* 9:18. doi: 10.3389/fncir.2015.00018

Shutoh, F., Ohki, M., Kitazawa, H., Itohara, S., and Nagao, S. (2006). Memory trace of motor learning shifts transsynaptically from cerebellar cortex to nuclei for consolidation. *Neuroscience* 139, 767–777. doi: 10.1016/j.neuroscience.2005.12.035

The Blue Brain Project (2005). Available online at: http://bluebrain.epfl.ch/

Vranesic, I., Iijima, T., Ichikawa, M., Matsumoto, G., and Knöpfel, T. (1994). Signal transmission in the parallel fiberpurkinje cell system visualized by high-resolution imaging. *Proc. Natl. Acad. Sci. U.S.A.* 91, 13014–13017. doi: 10.1073/pnas.91.26.13014

Yamazaki, T., and Igarashi, J. (2013). Realtime cerebellum: a large-scale spiking network model of the cerebellum that runs in realtime using a graphics processing unit. *Neural Netw.* 47, 103–111. doi: 10.1016/j.neunet.2013.01.019

Yamazaki, T., and Nagao, S. (2012). A computational mechanism for unified gain and timing control in the cerebellum. *PLoS ONE* 7:e33319. doi: 10.1371/journal.pone.0033319

Yamazaki, T., Nagao, S., Lennon, W., and Tanaka, S. (2015). Modeling memory consolidation during posttraining periods in cerebellovestibular learning. *Proc. Natl. Acad. Sci. U.S.A.* 112, 3541–3546. doi: 10.1073/pnas.1413798112

Yamazaki, T., and Tanaka, S. (2007). A spiking network model for passage-of-time representation in the cerebellum. *Eur. J. Neurosci.* 26, 2279–2292. doi: 10.1111/j.1460-9568.2007.05837.x

Zhang, W., and Linden, D. (2006). Long-term depression at the mossy fiber-deep cerebellar nucleus synapse. *J. Neurosci.* 26, 6935–6944. doi: 10.1523/JNEUROSCI.0784-06.2006

Large-Scale Simulations of Plastic Neural Networks on Neuromorphic Hardware

James C. Knight[1]*, Philip J. Tully[2,3,4], Bernhard A. Kaplan[5], Anders Lansner[2,3,6] and Steve B. Furber[1]

[1] Advanced Processor Technologies Group, School of Computer Science, University of Manchester, Manchester, UK,
[2] Department of Computational Biology, Royal Institute of Technology, Stockholm, Sweden, [3] Stockholm Brain Institute, Karolinska Institute, Stockholm, Sweden, [4] Institute for Adaptive and Neural Computation, School of Informatics, University of Edinburgh, Edinburgh, UK, [5] Department of Visualization and Data Analysis, Zuse Institute Berlin, Berlin, Germany, [6] Department of Numerical analysis and Computer Science, Stockholm University, Stockholm, Sweden

SpiNNaker is a digital, neuromorphic architecture designed for simulating large-scale spiking neural networks at speeds close to biological real-time. Rather than using bespoke analog or digital hardware, the basic computational unit of a SpiNNaker system is a general-purpose ARM processor, allowing it to be programmed to simulate a wide variety of neuron and synapse models. This flexibility is particularly valuable in the study of biological plasticity phenomena. A recently proposed learning rule based on the Bayesian Confidence Propagation Neural Network (BCPNN) paradigm offers a generic framework for modeling the interaction of different plasticity mechanisms using spiking neurons. However, it can be computationally expensive to simulate large networks with BCPNN learning since it requires multiple state variables for each synapse, each of which needs to be updated every simulation time-step. We discuss the trade-offs in efficiency and accuracy involved in developing an event-based BCPNN implementation for SpiNNaker based on an analytical solution to the BCPNN equations, and detail the steps taken to fit this within the limited computational and memory resources of the SpiNNaker architecture. We demonstrate this learning rule by learning temporal sequences of neural activity within a recurrent attractor network which we simulate at scales of up to 2.0×10^4 neurons and 5.1×10^7 plastic synapses: the largest plastic neural network ever to be simulated on neuromorphic hardware. We also run a comparable simulation on a Cray XC-30 supercomputer system and find that, if it is to match the run-time of our SpiNNaker simulation, the super computer system uses approximately $45\times$ more power. This suggests that cheaper, more power efficient neuromorphic systems are becoming useful discovery tools in the study of plasticity in large-scale brain models.

Keywords: SpiNNaker, learning, plasticity, digital neuromorphic hardware, Bayesian confidence propagation neural network (BCPNN), event-driven simulation, fixed-point accuracy

1. INTRODUCTION

Motor, sensory and memory tasks are composed of sequential elements and are therefore thought to rely upon the generation of temporal sequences of neural activity (Abeles et al., 1995; Seidemann et al., 1996; Jones et al., 2007). However it remains a major challenge to learn such functionally meaningful dynamics within large-scale models using biologically plausible synaptic and neural plasticity mechanisms. Using SpiNNaker, a neuromorphic hardware platform for simulating large-scale spiking neural networks, and BCPNN, a plasticity model based on Bayesian inference, we demonstrate how temporal sequence learning could be achieved through modification of recurrent cortical connectivity and intrinsic excitability in an attractor memory network.

Spike-Timing-Dependent Plasticity (Bi and Poo, 1998) (STDP) inherently reinforces temporal causality which has made it a popular choice for modeling temporal sequence learning (Dan and Poo, 2004; Caporale and Dan, 2008; Markram et al., 2011). However, to date, all large-scale neural simulations using STDP (Morrison et al., 2007; Kunkel et al., 2011) have been run on large cluster machines or supercomputers, both of which consume many orders of magnitude more power than the few watts required by the human brain. Mead (1990) suggested that the solution to this huge gap in power efficiency was to develop an entirely new breed of "neuromorphic" computer architectures inspired by the brain. Over the proceeding years, a number of these neuromorphic architectures have been built with the aim of reducing the power consumption and execution time of large neural simulations.

Large-scale neuromorphic systems have been constructed using a number of approaches: NeuroGrid (Benjamin et al., 2014) and BrainScaleS (Schemmel et al., 2010) are built using custom analog hardware; True North (Merolla et al., 2014) is built using custom digital hardware and SpiNNaker (Furber et al., 2014) is built from software programmable ARM processors.

Neuromorphic architectures based around custom hardware, especially the type of sub-threshold analog systems which Mead (1990) proposed, have huge potential to enable truly low-power neural simulation, but inevitably the act of casting algorithms into hardware requires some restrictions to be accepted in terms of connectivity, learning rules, and control over parameter values. As an example of these restrictions, of the large-scale systems previously mentioned, only BrainScaleS supports synaptic plasticity in any form implementing both short-term plasticity and pair-based STDP using a dedicated mixed-mode circuit.

As a software programmable system, SpiNNaker will require more power than a custom hardware based system to simulate a model of a given size (Stromatias et al., 2013). However this software programmability gives SpiNNaker considerable flexibility in terms of the connectivity, learning rules, and ranges of parameter values that it can support. The neurons and synapses which make up a model can be freely distributed between the cores of a SpiNNaker system until they fit within memory; and the CPU and communication overheads taken in advancing the simulation can be handled within a single simulation time step.

This flexibility has allowed the SpiNNaker system to be used for the simulation of large-scale cortical models with up to 5.0×10^4 neurons and 5.0×10^7 synapses (Sharp et al., 2012, 2014); and various forms of synaptic plasticity (Jin et al., 2010; Diehl and Cook, 2014; Galluppi et al., 2015; Lagorce et al., 2015). In the most recent of these papers, Galluppi et al. (2015) and Lagorce et al. (2015) demonstrated that Sheik et al.'s (2012) model of the learning of temporal sequences from audio data can be implemented on SpiNNaker using a voltage-gated STDP rule. However, this model only uses a small number of neurons and Kunkel et al.'s (2011) analysis suggests that STDP alone cannot maintain the multiple, interconnected stable attractors that would allow spatio-temporal sequences to be learnt within more realistic, larger networks. This conclusion adds to growing criticism of simple STDP rules regarding their failure to generalize over experimental observations (see e.g., Lisman and Spruston, 2005, 2010; Feldman, 2012 for reviews).

We address some of these issues by implementing spike-based BCPNN (Tully et al., 2014)—an alternative to phenomenological plasticity rules which exhibits a diverse range of mechanisms including Hebbian, neuromodulated, and intrinsic plasticity—all of which emerge from a network-level model of probabilistic inference (Lansner and Ekeberg, 1989; Lansner and Holst, 1996). BCPNN can translate correlations at different timescales into connectivity patterns through the use of locally stored synaptic traces, enabling a functionally powerful framework to study the relationship between structure and function within cortical circuits. In Sections 2.1–2.3, we describe how this learning rule can be combined with a simple point neuron model as the basis of a simplified version of Lundqvist et al.'s (2006) cortical attractor memory model. In Sections 2.4, 2.5, we then describe how this model can be simulated efficiently on SpiNNaker using an approach based on a recently proposed event-driven implementation of BCPNN (Vogginger et al., 2015). We then compare the accuracy of our new BCPNN implementation with previous non-spiking implementations (Sandberg et al., 2002) and demonstrate how the attractor memory network can be used to learn and replay spatio-temporal sequences (Abbott and Blum, 1996). Finally, in Section 3.3, we show how an anticipatory response to this replay behavior can be decoded from the neurons' sub-threshold behavior which can in turn be used to infer network connectivity.

2. MATERIALS AND METHODS

2.1. Simplified Cortical Microcircuit Architecture

We constructed a network using connectivity based on a previously proposed cortical microcircuit model (Lundqvist et al., 2006) and inspired by the columnar structure of neocortex (Mountcastle, 1997). The network consists of N_{HC} hypercolumns arranged in a grid where each hypercolumn consists of 250 inhibitory basket cells and 1000 excitatory pyramidal cells evenly divided into 10 minicolumns. Within each hypercolumn, the pyramidal cells send AMPA-mediated

connections to the basket cells with a connection probability of 10 % and a weight of 0.4 nA (defined as a postsynaptic current (PSC) amplitude). The basket cells then send GABAergic connections back to the pyramidal cells with a connection probability of 10 % and a weight of 2 nA. The basket cells are also recurrently connected through GABAergic connections, again with a connection probability of 10 % and a connection weight of 2 nA. The functional outcome of this local connectivity (excitatory to inhibitory and vice versa) is to enable winner-take-all (WTA) dynamics within each hypercolumn. While the strength of the local synapses remains fixed, all pyramidal cells in the network are also recurrently connected to each other through global AMPA and NMDA connections using plastic BCPNN synapses (see Section 2.2): also with a connection probability of 10 %. All connections in the network have distance-dependent synaptic delays such that, between two cells located in hypercolumns H_{xy}^{pre} and H_{xy}^{post}, the delay is calculated based on the Euclidean distance between the grid coordinates of the hypercolumns (meaning that all local connections have delays of 1 ms):

$$t_d^{H_{xy}^{pre} H_{xy}^{post}} = \frac{d_{norm}\sqrt{\left(H_x^{post} - H_x^{pre}\right)^2 + \left(H_y^{post} - H_y^{pre}\right)^2}}{V} + 1 \quad (1)$$

Where conduction velocity $V = 0.2 \, \text{mm ms}^{-1}$ and $d_{norm} = 0.75 \, \text{mm}$.

2.2. Synaptic and Intrinsic Plasticity Model

The spike-based BCPNN learning rule is used to learn the strengths of all global synaptic connections and the intrinsic excitabilities of all pyramidal cells in the network described in Section 2.1. The goal of the learning process is to estimate the probabilities of pre- and postsynaptic neurons firing (P_i and P_j respectively), along with the probability of them firing together (P_{ij}). Then, as Lansner and Holst (1996) describe, these probabilities can be used to calculate the synaptic strengths and intrinsic excitabilities of the network allowing it to perform Bayesian inference. Tully et al. (2014) developed an approach for estimating these probabilities based on pre- and postsynaptic spike trains (S_i and S_j respectively), defined as summed Dirac delta functions $\delta(\cdot)$ where $t_{i,j}^f$ represent the times of spikes:

$$S_i(t) = \sum_{t_i^f} \delta(t - t_i^f) \qquad S_j(t) = \sum_{t_j^f} \delta(t - t_j^f) \quad (2)$$

These spike trains are then smoothed using exponentially weighted moving averages to calculate the Z traces:

$$\tau_{z_i} \frac{dZ_i}{dt} = \frac{S_i}{f_{max}\Delta t} - Z_i \qquad \tau_{z_j} \frac{dZ_j}{dt} = \frac{S_j}{f_{max}\Delta t} - Z_j \quad (3)$$

Here, the maximum allowed firing rate f_{max} and spike duration $\Delta t = 1$ ms combine with the lowest attainable probability estimate $\epsilon = \frac{1000}{f_{max}\tau_p}$ introduced in Equation (5) to maintain a

linear mapping from neuronal spike rates to probabilities. For more details on the Bayesian transformation entailed by these equations, see Tully et al. (2014). The Z trace time constants τ_{z_i} and τ_{z_j} determine the time scale over which correlations can be detected and are inspired by fast biological processes such as Ca^{2+} influx via NMDA receptors or voltage-gated Ca^{2+} channels. The Z traces are then fed into the P traces, where a coactivity term is introduced:

$$\tau_p \frac{dP_i}{dt} = Z_i - P_i \quad \tau_p \frac{dP_{ij}}{dt} = Z_i Z_j - P_{ij} \quad \tau_p \frac{dP_j}{dt} = Z_j - P_j \quad (4)$$

The P trace time constant τ_p models long-term memory storage events such as gene expression or protein synthesis. It can be set higher to more realistically match these slow processes, but since simulation time increases with higher τ_p values, in this work we keep them just long enough to preserve the relevant dynamics. Estimated levels of activity in the P traces are then combined to compute a postsynaptic bias membrane current I_{β_j} and synaptic weight between pre- and postsynaptic neurons w_{ij}:

$$I_{\beta_j} = \beta_{gain} \log(P_j + \epsilon) \quad w_{ij} = w_{gain}^{syn} \log \frac{P_{ij} + \epsilon^2}{(P_i + \epsilon)(P_j + \epsilon)} \quad (5)$$

Here, β_{gain} is used to scale the BCPNN bias into an intrinsic input current to the neuron which is used to model an A-type K+ channel (Jung and Hoffman, 2009) or other channel capable of modifying the intrinsic excitability of a neuron (Daoudal and Debanne, 2003). Similarly, w_{gain}^{syn} is used to scale the BCPNN weight into a current-based synaptic strength.

2.3. Neuronal Model

We model excitatory and inhibitory cells as IAF neurons with exponentially decaying PSCs (Liu and Wang, 2001; Rauch et al., 2003). The sub-threshold membrane voltage V_m of these neurons evolves according to:

$$\tau_m \frac{dV_m}{dt} = -V_m + R_m \left(I_s + I_a + I_{\beta_j}\right) \quad (6)$$

The membrane time constant τ_m and capacitance C_m determine the input resistance $R_m = \frac{\tau_m}{C_m}$ through which input currents from the afferent synapses (I_s), spike-frequency adaption mechanism (I_a) and the intrinsic input current from the BCPNN learning rule (I_{β_j}) – described in Section 2.2—are applied. When V_m reaches the threshold V_t a spike is emitted, V_m is reset to V_r and α is added to the adaption current I_a. We use Liu and Wang's (2001) model of spike-frequency adaption with the adaption time constant τ_a:

$$\tau_a \frac{dI_a}{dt} = -I_a \quad (7)$$

The synaptic input current to postsynaptic neuron j (I_{s_j}) is modeled as a sum of exponentially shaped PSCs from other presynaptic neurons in the network:

$$\tau_{syn} \frac{dI_{s_j}}{dt} = -I_{s_j} + \sum_{syn} \sum_{i=0}^{n} w_{ij}^{syn} \sum_{t_i^f} \delta(t - t_i^f) \quad (8)$$

w_{ij}^{syn} indicates the weight of the connection between neurons i and j [where $syn \in$ (AMPA, GABA, NMDA) denotes the synapse type], t_i^f represents the arrival time of spikes from presynaptic neuron i (where there are n neurons in the network), and τ_{syn} is the synaptic time constant.

2.4. Simulating Spiking Neural Networks on SpiNNaker

SpiNNaker is a digital neuromorphic architecture designed for the simulation of spiking neural networks. Although systems built using this architecture are available in sizes ranging from single boards to room-size machines, they all share the same basic building blocks—the SpiNNaker chip (Furber et al., 2014). Each of these chips is connected to its six immediate neighbors using a chip-level interconnection network with a hexagonal mesh topology. Each SpiNNaker chip contains 18 ARM cores connected to each other through a network-on-chip, and connected to an external network through a multicast router. Each core has two small tightly-coupled memories: 32 KiB for instructions and 64 KiB for data; and shares 128 MiB of off-chip SDRAM with the other cores on the same chip. Although this memory hierarchy is somewhat unusual, the lack of global shared memory means that many of the problems of simulating large spiking neural networks on a SpiNNaker system are shared with more typical distributed computer systems. Morrison et al. (2005) and Kunkel et al. (2012) developed a collection of approaches for mapping such networks onto large distributed systems in a memory-efficient manner while still obtaining supra-linear speed-up as the number of processors increases. The SpiNNaker neural simulation kernel employs a very similar approach where, as shown in **Figure 1**, each processing core is responsible for simulating between 100 and 1000 neurons and their afferent synapses. The neurons are simulated using a time-driven approach with their state held in the tightly-coupled data memory. Each neuron is assigned a 32 bit ID and, when a simulation step results in a spike, it sends a packet containing this ID to the SpiNNaker router. These "spike" packets are then routed across the network fabric to the cores that are responsible for simulating these synapses. Biological neurons have in the order of $10^3 – 10^4$ afferent synapses, so updating all of these every time step would be extremely computationally intensive. Instead, as individual synapses only receive spikes at relatively low rates, they can be updated only when they transfer a spike as long as their new state can be calculated from:

1. The synapse's previous state.
2. The time since the last spike was transferred.
3. Information available from the time-driven simulation of the postsynaptic neuron.

Using this event-driven approach on SpiNNaker is also advantageous as, due to their sheer number, synapses need to be stored in the off-chip SDRAM which has insufficient bandwidth for every synapse's parameters to be retrieved every simulation time step (Painkras and Plana, 2013). Instead, on receipt of a "spike" packet, a core retrieves the row of the connectivity matrix associated with the firing neuron from SDRAM. Each
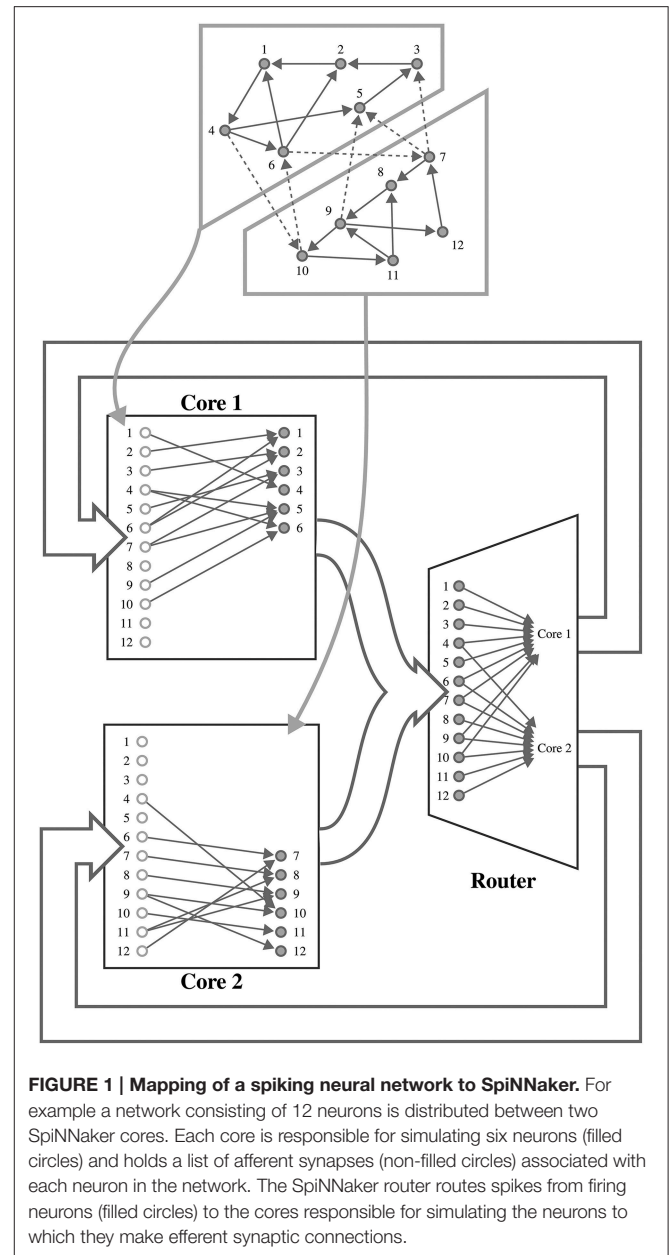
FIGURE 1 | Mapping of a spiking neural network to SpiNNaker. For example a network consisting of 12 neurons is distributed between two SpiNNaker cores. Each core is responsible for simulating six neurons (filled circles) and holds a list of afferent synapses (non-filled circles) associated with each neuron in the network. The SpiNNaker router routes spikes from firing neurons (filled circles) to the cores responsible for simulating the neurons to which they make efferent synaptic connections.

of these rows describes the parameters associated with the synapses connecting the firing neuron to those simulated on the core. Once a row is retrieved the weights are inserted into an input ring-buffer, where they remain until any synaptic delay has elapsed and they are applied to the neuronal input current.

In addition to enabling large-scale simulations with static synapses, this event-driven approach can, in theory, be extended to handle any type of plastic synapse that meets the 3 criteria outlined above. However, simulating plastic synapses has additional overheads in terms of memory and CPU load—both of which are very limited resources on SpiNNaker. Several different approaches have been previously developed that aim to minimize

Algorithm 1 Algorithmic Implementation of STDP

function processRow(t)
  for each j **in** postSynapticNeurons **do**
    history $\leftarrow$ getHistoryEntries(j, t_{old}, t)

    for each (t_j, s_j) **in** history **do**
      $w_{ij} \leftarrow$ applyPostSpike($w_{ij}, t_j, t_{old}, s_i$)

    $(t_j, s_j) \leftarrow$ getLastHistoryEntry(t)
    $w_{ij} \leftarrow$ applyPreSpike(w_{ij}, t, t_j, s_j)
    addWeightToRingBuffer(w_{ij}, j)

    $s_i \leftarrow$ addPreSpike(s_i, t, t_{old})
    $t_{old} \leftarrow t$

memory usage (Jin et al., 2010), reduce CPU load (Diehl and Cook, 2014) or offload the processing of plastic synapses to dedicated cores (Galluppi et al., 2015). Morrison et al. (2007) also extended their work on distributed spiking neural network simulation to include synaptic plasticity, developing an algorithm for simulating plastic synapses in an event-driven manner, using a simplified model of synaptic delay to reduce CPU and memory usage. In this work, we combine elements of Diehl and Cook's (2014) and Morrison et al.'s (2007) approaches, resulting in **Algorithm 1** which is called whenever the connectivity matrix row associated with an incoming "spike" packet is retrieved from the SDRAM. As well as the weights of the synapses connecting the presynaptic neuron to the postsynaptic neurons simulated on the local core (w_{ij}), the row also has a header containing the time at which the presynaptic neuron last spiked (t_{old}) and its state at that time (s_i). The exact contents of the state depends on the plasticity rule being employed, but as only the times of presynaptic spikes are available at the synapse, the state often consists of one or more low-pass filtered versions of this spike train.

The algorithm begins by looping through each postsynaptic neuron (j) in the row and retrieving a list of the times (t_j) at which that neuron spiked between t_{old} and t and its state at that time (s_j). In the SpiNNaker implementation, these times and states are stored in a fixed-length circular queue located in the tightly-coupled data memory to which a new entry gets added whenever a local neuron fires. Next, the effect of the interaction between these postsynaptic spikes and the presynaptic spike that occurred at t_{old} is applied to the synapse using the applyPostSpike function. The synaptic update is then completed by applying the effect of the interaction between the presynaptic spike that instigated the whole process and the most recent postsynaptic spike to the synapse using the applyPreSpike function before adding this weight to the input ring buffer. Finally, the header of the row is updated by calling the addPreSpike function to update s_i and setting t_{old} to the current time.

2.5. An Event-based, SpiNNaker Implementation of Bayesian Learning

Equations (3)–(5) cannot be directly evaluated within the event-driven synaptic processing scheme outlined in Section 2.4, but

as they are simple first-order linear ODEs, they can be solved to obtain closed-form solutions for $Z(t)$ and $P(t)$. These equations then need only be evaluated when spikes occur. Vogginger et al. (2015) converted this resultant system of equations into a spike-response model (Gerstner and Kistler, 2002) which, as it only consists of linear combinations of $e^{\frac{-t}{\tau_z}}$ and $e^{\frac{-t}{\tau_p}}$, can be re-framed into a new set of new state variables $Z_i^*, Z_j^*, P_i^*, P_j^*,$ and P_{ij}^*. These, like the state variables used in many STDP models are simply low-pass filtered versions of the spike-trains and can be evaluated when a spike occurs at time t:

$$Z_i^*(t) = Z_i^*(t^{last})e^{-\frac{\Delta t}{\tau_{z_i}}} + S_i(t) \quad P_i^*(t) = P_i^*(t^{last})e^{-\frac{\Delta t}{\tau_p}} + S_i(t)$$
(9)

Z^* and P^* can now be stored in the pre and postsynaptic state (s_i and s_j) and updated in the addPreSpike function called from algorithm 1; and when postsynaptic neurons fire. The correlation trace, P_{ij} can similarly be re-framed in terms of a new state variable:

$$P_{ij}^*(t) = P_{ij}^*(t^{last})e^{-\frac{\Delta t}{\tau_p}} + S_i(t)Z_j^*(t)$$
(10)

P_{ij}^* can now be stored alongside the synaptic weight w_{ij} in each synapse and evaluated in the applyPreSpike and applyPostSpike functions called from algorithm 1. The final stage of the event-based implementation is to obtain the P_i, P_j and P_{ij} values required to evaluate (Equation 5) from the new state variables and thus obtain w_{ij} and β_j.

$$P_i(t) = a_i \left(Z_i^*(t) - P_i^*(t) \right)$$
(11)

$$P_{ij}(t) = a_{ij} \left(Z_i^*(t)Z_j^*(t) - P_{ij}^*(t) \right)$$
(12)

With the following coefficients used for brevity:

$$\tau_{z_{ij}} = \left(\frac{1}{\tau_{z_i}} + \frac{1}{\tau_{z_j}} \right)^{-1} \quad a_i = \frac{1}{f_{max}\left(\tau_{z_i} - \tau_p\right)}$$
$$a_{ij} = \frac{1}{f_{max}^2 \left(\tau_{z_j} + \tau_{z_i}\right)\left(\tau_{z_{ij}} - \tau_p\right)}$$
(13)

This approach makes implementing spike-based BCPNN on SpiNNaker feasible from an algorithmic point of view, but limitations of the SpiNNaker architecture further complicate the problem. The most fundamental of these limitations is that, as Moise (2012, p. 20) explains, for reasons of silicon area and energy efficiency, SpiNNaker has no hardware floating point unit. While floating point operations can be emulated in software, this comes at a significant performance cost meaning that performance-critical SpiNNaker software needs to instead use fixed-point arithmetic. Hopkins and Furber (2015) discussed the challenges of using fixed-point arithmetic for neural simulation on the SpiNNaker platform in detail but, in the context of this work, there are two main issues of particular importance. Firstly the range of fixed-point numeric representations is static so, to attain maximal accuracy, the optimal representation for storing each state variable must be chosen ahead of time. Vogginger et al.

(2015) investigated the use of fixed-point types for BCPNN as a means of saving memory and calculated that, in order to match the accuracy of a time-driven floating point implementation, a fixed-point format with 10 integer and 12 fractional bits would be required. However, not only is the model described in Section 2.2 somewhat different from the reduced modular model considered by Vogginger et al. (2015), but the ARM architecture only allows 8, 16, or 32 bit types to be natively addressed. Therefore, we re-evaluated these calculations for the SpiNNaker implementation and chose to use 16 bit types for two reasons:

1. In order to implement the *getLastHistoryEntry* and *getHistoryEntries* functions used in algorithm 1, each neuron needs to store a history of Z_j^* and P_j^* values in the tightly-coupled data memory, therefore minimizing the size of these variables is important.
2. The SpiNNaker CPU cores can perform multiplication operations on signed 16 bit values faster than it can on 32 bit values, allowing more spikes to be transferred each time-step.

Based on a total of 16 bit, the number of bits used for the integer and fractional parts of the fixed-point representation needs to be determined based on the range of the state variables. As all of the Z^* and P^* state variables are linear sums of exponential spike responses and P^* has the largest time constant, it decays slowest meaning that it will reach the highest value. Therefore we can calculate the maximum value which our fixed-point format must be able to represent in order to handle a maximum spike frequency of f_{max} as follows:

$$P_{max}^* = \frac{1}{1 - e^{-\frac{1}{f_{max} \times \tau_p}}} \quad (14)$$

In order to match the firing rates of pyramidal cells commonly observed in cortex, low values of the maximum firing rate (f_{max}, e.g., 20 or 50 Hz) are often used with the BCPNN model described in Section 2.2. On this basis, by using a signed fixed-point format with 6 integer and 9 fractional bits, if $f_{max} = 20$ Hz, traces with $\tau_p < 3.17$ s can be represented and, if $f_{max} = 50$ Hz, traces with $\tau_p < 1.27$ s can be represented.

The second problem caused by the lack of floating point hardware is that there is no standard means of calculating transcendental functions for fixed-point arithmetic. This means that the exponential and logarithm functions required to implement BCPNN must be implemented by other means. While it is possible to implement approximations of these functions using, for instance a Taylor series, the resultant functions are likely to take in the order of 100 CPU cycles to evaluate (Moise, 2012), making them too slow for use in the context of BCPNN where around ten of these operations will be performed every time a spike is transferred. Another approach is to use pre-calculated lookup tables (LUTs). These are particularly well suited to implementing functions such as $e^{\frac{-t}{\tau}}$ where t is discretized to simulation time steps and, for small values of τ, the function decays to 0 after only a small number of table entries. While the $\log(x)$ function has neither of these ideal properties, x can be normalized into the form $x = y \times 2^n : n \in \mathbb{Z}, y \in [1, 2)$ so a LUT is only required to cover the interval $[1, 2)$ within which $\log(x)$ is relatively linear.

3. RESULTS

3.1. Validating BCPNN Learning on SpiNNaker with Previous Implementations

In this section we demonstrate that the implementation of BCPNN we describe in Section 2.5 produces connection weights and intrinsic excitabilities comparable to those learned by previous models. To do this we used the procedure developed by Tully et al. (2014) and the network described in **Table 1** to compare two neurons, connected with a BCPNN synapse, modeled using both our spiking BCPNN implementation and as abstract units with simple, exponentially smoothed binary activation patterns (Sandberg et al., 2002). We performed this comparison by presenting the neurons with five patterns of differing relative activations, each repeated for ten consecutive 200 ms trials. Correlated patterns meant both neurons were firing at f_{max} Hz or ϵ Hz each trial; independent patterns meant uniform sampling of f_{max} Hz and ϵ Hz patterns for both neurons in each trial; anti-correlated patterns meant one neuron fired at f_{max} Hz and the other at ϵ Hz or vice-versa in each trial; both muted meant both neurons fired at ϵ Hz in all trials; and post muted meant uniform sampling of presynaptic neuron activity while the postsynaptic neuron fired at ϵ Hz in all trials.

As **Figure 2** shows, during the presentation of patterns in which both units are firing, the responses from the abstract model fall well within the standard deviation of the SpiNNaker model's responses, but as units are muted, the two models begin to diverge. Further investigation into the behavior of the individual state variables shows that this is due to the P^* term of Equation (11) coming close to underflowing the 16 bit fixed-point format when a long time has passed since the last spike. This inaccuracy in the P^* term is then further amplified when the weights and intrinsic excitabilities are calculated using (Equation 5) as for small values of x, $\log(x)$ approaches its vertical asymptote. The standard deviations visible in **Figure 2** reflect the fact that for the spiking learning rule, the realization of Poisson noise that determined firing rates was different for each trial, but with a rate modulation that was repeated across trials.

3.2. Learning Sequential Attractors using Spike-Based BCPNN

In this section we consider a functional use case of the the modular attractor network described in Section 2.1 involving learning temporal sequences of attractors. With asymmetrical BCPNN time constants, it was previously proposed that this network could self-organize spontaneously active sequential attractor trajectories (Tully et al., 2014). We built a suitable network using the neuron and plasticity models described in Sections 2.2, 2.3; and the parameters listed in **Table 2**. Using this network we employed a training regime—a subset of which is shown in **Figure 3A**—in which we repeatedly stimulated all cells in a mutually exclusive sequence of minicolumns for 50 training epochs. Each minicolumn was stimulated for 100 ms, such that the neurons within it fired at an average rate of f_{max} Hz. During training we disabled the term in Equation (8) that incorporates input from the plastic AMPA and NMDA synapses meaning that, while the weights were learned online, the dynamics of the network did not disturb the training regime. A recall phase

TABLE 1 | Model description of the BCPNN validation network.

(A) Model summary

Populations	Presynaptic, postsynaptic, presynaptic input, postsynaptic input
Connectivity	One-to-one
Neuron model	Leaky integrate-and-fire with exponential-shaped synaptic current inputs and spike-frequency adaption (Liu and Wang, 2001)
Synapse model	Current-based with exponential-shaped PSCs
Plasticity	BCPNN AMPA synapses
Input	Externally generated Poisson spike trains
Measurements	Intrinsic bias current and synaptic weights

(B) Populations

Name	Elements	Size
Presynaptic	Leaky IAF	1
Postsynaptic	Leaky IAF	1
Presynaptic input	External spike source	1
Postsynaptic input	External spike source	1

(C) Connectivity

Source	Target	Pattern	Weight
Presynaptic input	Presynaptic	One-to-one	2 nA
Postsynaptic input	Postsynaptic	One-to-one	2 nA
Presynaptic	Postsynaptic	One-to-one	Plastic

(D) Neuron and synapse model

|	
Type	Leaky integrate-and-fire with exponential-shaped synaptic current inputs and spike-frequency adaption (Liu and Wang, 2001) as described in Section 2.3
Parameters	$\tau_m = 10$ ms membrane time constant
	$C_m = 250$ pF membrane capacitance
	$V_t = -55.4$ mV threshold voltage
	$V_r = -70$ mV reset voltage
	$\alpha = 0.0$ nA adaption current (disabled)
	$\tau_{AMPA} = 2.5$ ms AMPA synapse time constant

(E) Plasticity

Type	BCPNN AMPA synapses as described in Section 2.2
Parameters	$f_{max} = 50$ Hz maximum spiking frequency
	$\tau_{z_i} = 10$ ms presynaptic primary trace time constant
	$\tau_{z_j} = 10$ ms postsynaptic primary trace time constant
	$\tau_p = 1000$ ms probability trace time constant
	$w_{gain}^{syn} = 1$ nA weight gain
	$\beta_{gain} = 1$ nA intrinsic bias gain

(F) Input

Type	Description
Externally generated Poisson spike trains	As described in Section 3.1

After Nordlie et al. (2009).
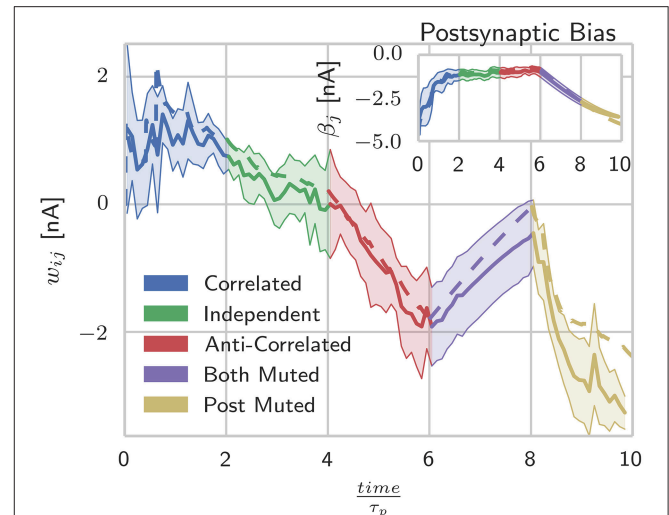
FIGURE 2 | Spike-based BCPNN estimates abstract BCPNN for different input patterns. Comparing weight and bias (inset) development under different protocols when using abstract (dotted) and SpiNNaker (solid) versions of the learning rule. SpiNNaker simulations were repeated 10 times and averaged, with standard deviations illustrated by the shaded regions.

followed this learning phase in which a 50 ms stimulus of f_{max} Hz was applied to all cells in the first minicolumn of the learned sequence. During both the training and recall phases we provided background input to each cell in the network from an independent 65 Hz Poisson spike source. These Poisson spike sources are simulated on additional SpiNNaker cores to those used for the neural simulation algorithm described in Section 2.4.

We found that the training regime was able to produce the recurrent connectivity required to perform temporal sequence recall in the same serial order that patterns were presented during training as shown in **Figure 3B**. Each sequence element replayed as a learned attractor state that temporarily stifled the activity of all other cells in the network due to WTA and asymmetrically projected NMDA toward neurons of the subsequent sequence element, allowing a stable trajectory to form. Activity within attractor states was sharpened and stabilized by learned auto-associative AMPA connectivity; and sequential transitions were jointly enabled by neural adaptation and inter-pattern heteroassociation via NMDA synapses.

Because of the modular structure of the network described in Section 2.1, this temporal sequence learning can be performed using networks of varying scales by instantiating different number of hypercolumns and linearly scaling the w_{gain}^{syn} parameter of the connections between them. By doing this, we investigated how the time taken to simulate the network on SpiNNaker scales with network size. **Figure 4** shows how these times are split between the training and testing phases; and how long is spent generating data on the host computer, transferring it to and from SpiNNaker and actually running the simulation. As the SpiNNaker simulation always runs at a fixed fraction of real-time (for this simulation 0.5×), the simulation time remains constant

TABLE 2 | Parameters for the modular attractor network.

(A) Model summary

Populations and connectivity	Modular structure described in Section 2.1
Neuron model	Leaky integrate-and-fire with exponential-shaped synaptic current inputs and spike-frequency adaption (Liu and Wang, 2001)
Synapse model	Current-based with exponential-shaped PSCs
Plasticity	BCPNN AMPA and NMDA synapses
Input	Externally generated Poisson spike trains and independent fixed-rate Poisson spike trains
Measurements	Spiking activity, membrane voltages, intrinsic bias current and synaptic weights

(B) Neuron and synapse model

Type	Leaky integrate-and-fire with exponential-shaped synaptic current inputs and spike-frequency adaption (Liu and Wang, 2001) as described in Section 2.3
Parameters	$\tau_m = 20$ ms membrane time constant
	$C_m = 250$ pF membrane capacitance
	$V_t = -50$ mV threshold voltage
	$V_r = -70$ mV reset voltage
	$\alpha = 0.15$ nA adaption current
	$\tau_a = 300$ ms adaption time constant
	$\tau_{AMPA} = 5$ ms AMPA synapse time constant
	$\tau_{GABA} = 5$ ms GABA synapse time constant
	$\tau_{NMDA} = 150$ ms NMDA synapse time constant

(C) Plasticity

Type	BCPNN AMPA synapses as described in Section 2.2
Parameters	$f_{max} = 20$ Hz maximum spiking frequency
	$\tau_{z_i} = 5$ ms presynaptic primary trace time constant
	$\tau_{z_j} = 5$ ms postsynaptic primary trace time constant
	$\tau_p = 2000$ ms probability trace time constant
	$w_{gain}^{syn} = \frac{0.546}{N_{HC}}$ nA weight gain
Type	BCPNN NMDA synapses as described in Section 2.2
Parameters	$f_{max} = 20$ Hz maximum spiking frequency
	$\tau_{z_i} = 5$ ms presynaptic primary trace time constant
	$\tau_{z_j} = 150$ ms postsynaptic primary trace time constant
	$\tau_p = 2000$ ms probability trace time constant
	$w_{gain}^{syn} = \frac{0.114}{N_{HC}}$ nA weight gain
	$\beta_{gain} = 0.05$ nA intrinsic bias gain

(D) Input

Type	Description
Externally generated Poisson spike trains	As described in Section 3.2
Independent fixed-rate Poisson spike trains	As described in Section 2.1

After Nordlie et al. (2009).

as the network grows, but the times required to generate the data and to transfer it grow significantly, meaning that when $N_{HC} = 16$ (2.0×10^4 neurons and 5.1×10^7 plastic synapses), the total simulation time is 146 min. However, the amount of time spent in several phases of the simulation is increased by limitations of the current SpiNNaker toolchain. 84 min is spent downloading the learned weight matrices and re-uploading them for the testing: A process that is only required because the changing of parameters (in this case, whether learning is enabled or not) mid-simulation is not currently supported. Additionally, the current implementation of the algorithm outlined in Section 2.4 only allows neurons simulated on one core to have afferent synapses with a single learning rule configuration. This means that we have to run the training regime twice with the same input spike trains, once for the AMPA synapses and once for the NMDA synapses: Doubling the time taken to simulate the training network.

Previous supercomputer simulations of modular attractor memory networks have often used more complex neuron models and connectivity (Lundqvist et al., 2010), making simulation times difficult to compare with our SpiNNaker simulation due to the simplifications we outlined in Section 2.1. In order to present a better comparison, we built a network model with the same connectivity as our SpiNNaker model and simulated it on a Cray XC-30 supercomputer system using NEST version 2.2 (Gewaltig and Diesmann, 2007) with the spike-based BCPNN implementation developed by Tully et al. (2014). NEST does not include the adaptive neuron model we described in Section 2.3 so we used the adaptive exponential model (Brette and Gerstner, 2005): a simple point neuron model with spike-frequency adaption.

As previously discussed SpiNNaker runs at a fixed-fraction of real-time so we distribute our NEST simulations across increasing numbers of Cray XC-30 compute nodes (each consisting of two 2.5 GHz Intel Ivy Bridge Xeon processors) until the simulation completed in the same time as those shown in **Figure 4** for our SpiNNaker simulations. **Table 3** shows the result of both these supercomputer simulations and a second set with the time taken for the mid-simulation downloading and re-uploading of weights—currently required by the SpiNNaker software—removed. Due to this redundant step and because NEST parallelizes the generation of simulation data across the compute nodes, at all three scales, our modular attractor network can be simulated using 2 compute nodes. However, if we remove the time spent downloading and re-uploading the weights, 9 compute nodes are required to match the run-time of the SpiNNaker simulation when $N_{HC} = 16$.

While a more in-depth measurement of power usage is out of the scope of this work, we can also derive approximate figures for the power usage of our simulations running on both systems based on the 1 W peak power usage of the SpiNNaker chip and the 30 kW power usage of a Cray XC-30 compute rack (Cray, 2013). While these figures ignore the power consumed by the host computer connected to the SpiNNaker system; the power consumed by the "blower" and storage cabinets connected to the Cray XC-30; and assume that all CPUs are running at peak power usage, they show that even in the worst case, SpiNNaker uses 45× less power than the Cray XC-30 and, if the limitations of the current SpiNNaker software are addressed, this can be improved to 200×.

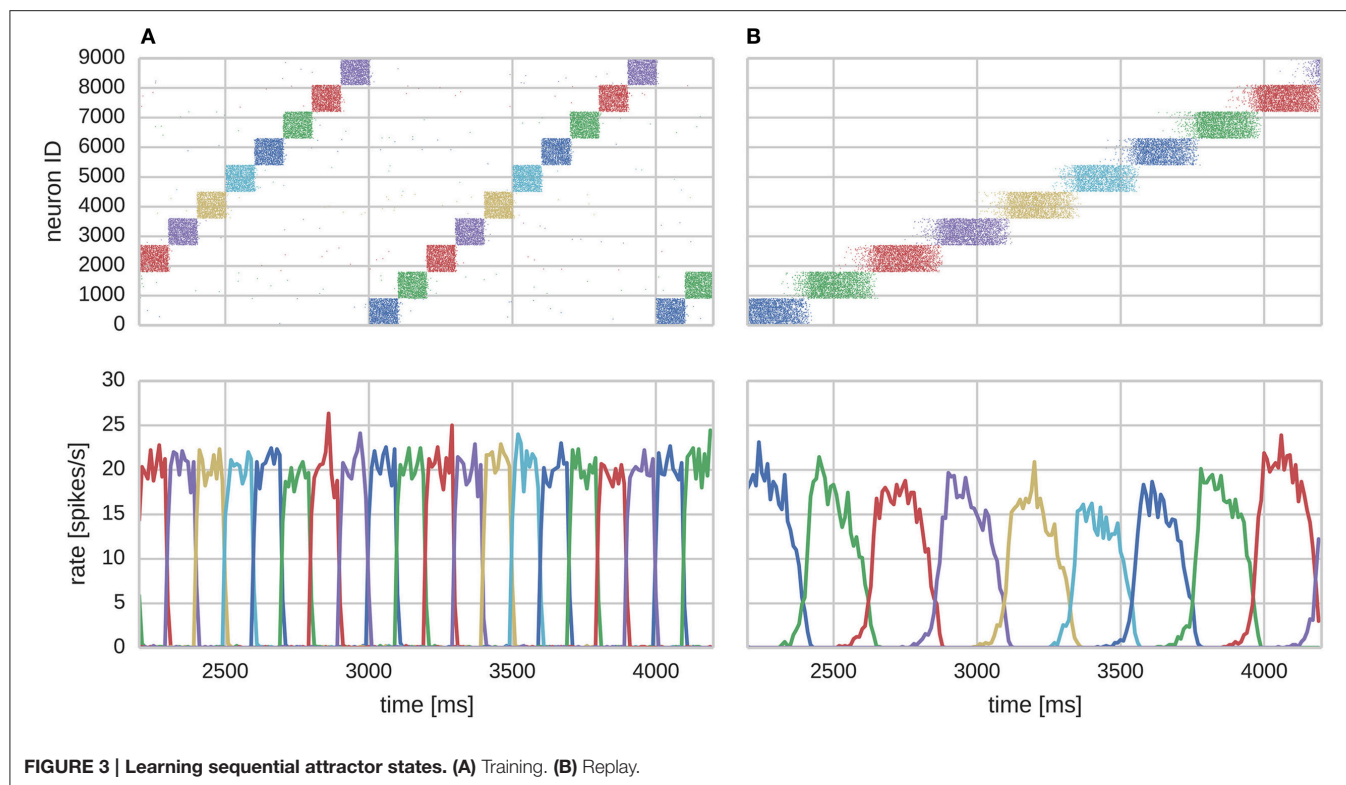**FIGURE 3 | Learning sequential attractor states. (A)** Training. **(B)** Replay.
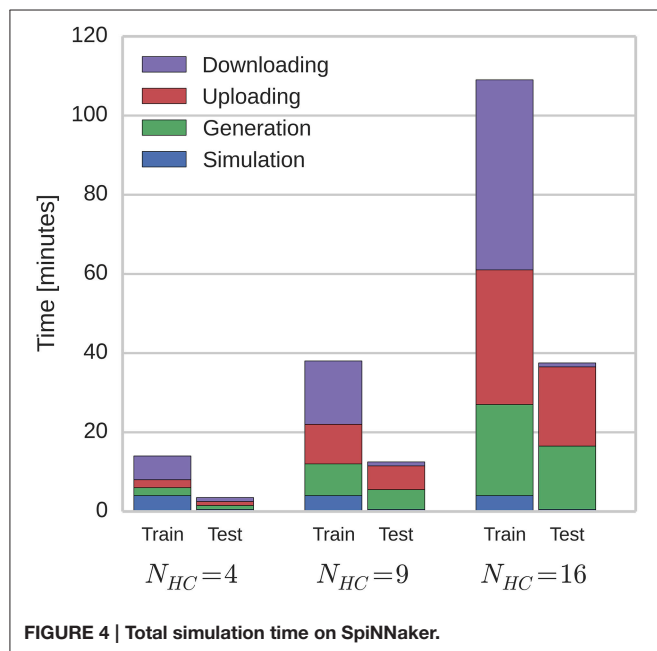

FIGURE 4 | Total simulation time on SpiNNaker.

3.3. Connectivity Patterns Show Different Signatures in Membrane Potentials

The purpose of this section is to study how learning parameters influence the resulting connectivity patterns and the effect of learned connectivity on membrane dynamics during sequence

replay. For this purpose we vary two parameters of the learning rule that control the time window within which correlations are detected $-$ τ_{z_i} on the pre- and τ_{z_j} on the postsynaptic side. The network is trained using the same regime described in Section 3.2 and two different configurations, one with $\tau_{z_i} = \tau_{z_j}$ on NMDA synapses, and one with $\tau_{z_i} \neq \tau_{z_j}$. If τ_{z_i} and τ_{z_j} are equal, the Z_i and Z_j traces evolve in the same manner, meaning that, as their dynamics propagate through the P traces to the synaptic weights, the forward and reciprocal connections between minicolumns develop symmetrically as shown in the top row of **Figure 5**. However, when $\tau_{z_i} \neq \tau_{z_j}$, the Z_i and Z_j traces evolve differently and, as the bottom row of **Figure 5** shows, asymmetrical connections develop between minicolumns. It is important to note that the spiking activity during the training regime is the same in both configurations and the shape of the resulting connectivity results only from the learning time-constants τ_{z_i} and τ_{z_j}.

In order to analyze the effect of the different learned connectivity patterns shown in **Figure 5**, we studied the impact of the two connectivity kernels on the subthreshold dynamics of neurons during sequence replay. As described in Section 3.2, after training, the trained sequence can be replayed by applying a 50 ms stimulus of f_{max} Hz into all cells of the first minicolumn in the learned sequence. Later, in the sequence replay when the stimulus has been removed, we recorded the membrane potential of all of the neurons in the network and stored the point in time when the firing rate of the respective minicolumn was maximal. We then align the membrane potential traces to this point in time and average them over all cells in a minicolumn. Interestingly,

TABLE 3 | Comparison of power usage of modular attractor network simulations running on SpiNNaker with simulations distributed across enough compute nodes of a Cray XC-30 system to match SpiNNaker simulation time.

Simulation		SpiNNaker		Cray XC-30	
N_{HC}	time [min]	# chips	Peak CPU power usage [W]	# compute nodes	Peak CPU power usage [W]
4	17	6	6	2	938
9	50	12	12	2	938
16	146	21	21	2	938
4	9	6	6	4	1875
9	23	12	12	14[a]	6563
16	62	21	21	9	4219

Cray XC-30 power usage is based on the 30 kW power usage of an entire Cray XC-30 compute rack (Cray, 2013). SpiNNaker power usage is based on the 1 W peak power usage of the SpiNNaker chip (Furber et al., 2014).

Top: SpiNNaker simulation times include downloading of learned weights and re-uploading required by current software.

Bottom: Time taken to download learned weights, re-generate and re-upload model to SpiNNaker have been removed.

[a]We are unsure why more supercomputer compute nodes are required to match the SpiNNaker simulation times when $N_{HC} = 9$ than when $N_{HC} = 16$. We assume this is an artifact of the different scaling properties of the two simulators, but further investigation is outside of the scope of this work.
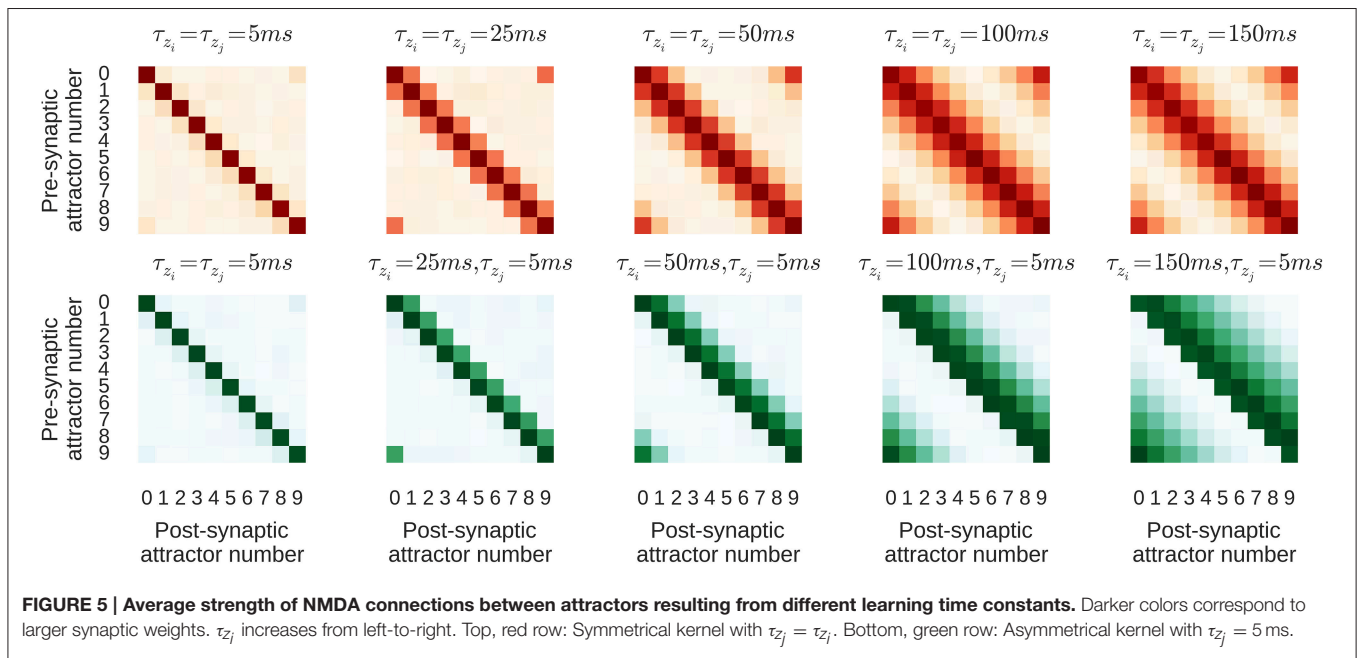
FIGURE 5 | Average strength of NMDA connections between attractors resulting from different learning time constants. Darker colors correspond to larger synaptic weights. τ_{z_i} increases from left-to-right. Top, red row: Symmetrical kernel with $\tau_{z_j} = \tau_{z_i}$. Bottom, green row: Asymmetrical kernel with $\tau_{z_j} = 5$ ms.

as **Figure 6** illustrates, these averaged and aligned membrane responses show different characteristics for the network models built on symmetric and asymmetric connectivity. Both network types show similar membrane characteristics before the sequence arrives at the minicolumn, but, the network with symmetric connectivity shows a significantly slower decrease in membrane potential after the sequence has passed. In contrast, the network with asymmetric connectivity shows a strong after-stimulus hyperpolarization due to the increased inhibitory input originating from minicolumns later in the sequence which get subsequently activated. The slower decrease in the mean membrane potential in the symmetric network can be explained by the excitatory projections in both directions of the sequence providing excitatory current flow to previously activated neurons. The implications of this experiment and interpretations of these different characteristics is discussed in Section 4.2.

4. DISCUSSION

The contribution of this study is threefold. Firstly, we have shown that BCPNN can be efficiently implemented within the constraints of the SpiNNaker neuromorphic architecture. Secondly, we have shown how BCPNN can be used in a functionally meaningful context to perform near real-time learning of temporal sequences within a large-scale modular attractor network—the largest plastic neural network ever to
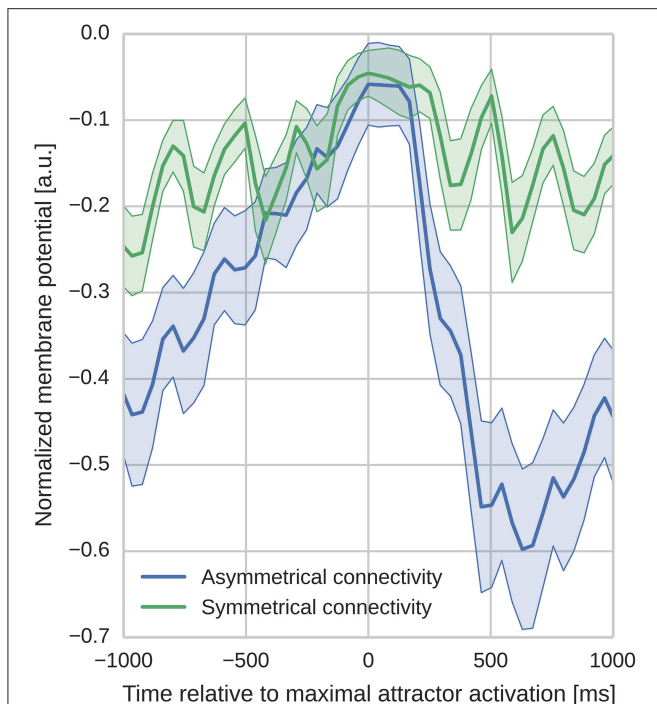
FIGURE 6 | Aligned average membrane potentials during sequence replay for two different connectivities. The membrane potentials have been recorded from all neurons in the trained network during sequence replay. These membrane voltages have then been averaged and aligned to the time of peak activity in the temporal domain (bold lines represent the mean, shaded areas represent the standard deviation). The y-axis has been normalized to improve visibility (0 corresponds to V_t and -1 corresponds to the minimal membrane voltage in the sample). In the network with asymmetric connectivity the mean membrane response shows a pronounced drop after the peak response, whereas the network with symmetric connectivity does not. Oscillatory behavior originates from switches between discrete attractor states alternated by phases of inhibitory feedback.

be simulated on neuromorphic hardware. Finally, we have demonstrated the value of SpiNNaker as a tool for investigating plasticity within large-scale brain models by exploring how, by changing a single parameter in the BCPNN learning rule, both symmetric and asymmetric connectivity can be learned, which in turn influence underlying membrane potential characteristics.

4.1. Learning Temporal Sequences in Cortical Microcircuits

The total duration of temporal sequences is longer than the time courses of individual cellular or synaptic processes and therefore, such sequences are thought to be driven by circuit-level phenomena although the intricacies of this relationship have yet to be fully explored. The massively recurrent and long-range nature of cortical connectivity, taken together with the emergence of temporal sequences at fine scales and distributed over spatial areas, suggests the presence of generic cortical microcircuit mechanisms. The model presented here is modularly organized into hypercolumns, each implementing WTA dynamics (Douglas and Martin, 2004). This modular structure also allowed us to vary the number of hypercolumns the

network contained without effecting its functionality (Djurfeldt et al., 2008). Such distributed systems generally exhibit an improved signal-to-noise ratio, error resilience, generalizability and a structure suitable for Bayesian calculations (McClelland et al., 1986; Barlow, 2001). Like their uniformly interconnected counterparts, they can also exhibit high variability in their spike train statistics (Litwin-Kumar and Doiron, 2012; Lundqvist et al., 2012). Moreover, due to their capacity to exhibit a rich repertoire of behaviorally relevant activity states, these topologies are also well suited for information processing and stimulus sensitivity (Lundqvist et al., 2010; Wang et al., 2011).

Previous investigations have shown that the attractors which emerge within such modular networks reproduce features of local UP states (Lundqvist et al., 2006). This observation remains consistent with the extension considered here since, *in vivo*, UP state onsets are accompanied by the sequential activation of cortical neurons (Luczak et al., 2007). This redundant neural coding scheme should not necessarily be viewed in terms of anatomical columns, but rather functional columns consisting of subgroups of neurons with similar receptive fields that are highly connected (Yoshimura et al., 2005) and co-active (Cossart et al., 2003). Similar stereotypical architectures have been used as building blocks for other unified mathematical frameworks of the neocortex (Johansson and Lansner, 2007; George and Hawkins, 2009; Bastos et al., 2012).

The dynamics of the model consists of attractors, whose activations produce self-sustaining spiking among groups of neurons spread across different hypercolumns. Activity within attractors is sharpened by the fast dynamics of the AMPA receptor, until the network transitions to a subsequent attractor due to neural adaptation and asymmetrical NMDA connectivity, both of which have longer time constants of activation. In this work we have shown how these dynamics could be learned using BCPNN, a learning rule which offers an alternative to phenomenological STDP rules that often require complementary mechanisms due to their prevailing instability (Kempter et al. 2001; Babadi and Abbott 2010; but see Gütig et al. 2003).

4.2. Sequence Anticipation and Asymmetric Connectivity as Observed in the Membrane Potential Dynamics

In both the symmetric and asymmetric networks, the stimulus-aligned mean membrane potential traces show a similar rise prior to sequence arrival which can be interpreted as a form of anticipation of the impending activity peak. By anticipation we mean the premature build-up of a neuronal response which is becoming increasingly similar to the response when the actual stimulus is present and represents the neural signature of expectation or prediction of future input. Anticipation is an important function of neural circuits and is observed not only in early sensory structures such as the retina (Berry et al., 1999; Hosoya et al., 2005; Vaney et al., 2012), but also in downstream structures and cortical areas (Rao and Ballard, 1999; Enns and Lleras, 2008) which are involved in more abstract

cognitive tasks (Riegler, 2001; Butz et al., 2003). Anticipation can also be regarded as a form of prediction of future events: something which Bar (2007) and Bubic et al. (2010) argue is a fundamental function of the brain. This predictive capability can improve sensory perception (Yoshida and Katz, 2011; Rohenkohl et al., 2012) and is important for other modalities such as motor control and learning (Shadmehr et al., 2010; Schlerf et al., 2012). However, the connectivity which implements this predictive or anticipatory function, and the mechanisms which give rise to it, are not well understood. We believe that BCPNN learning helps fill this gap—as we discussed in Section 3.2, it can learn functional connectivity at a network scale and, as previously argued in this section, it exhibits anticipatory behavior.

We studied the network response by looking at the membrane potential dynamics prior to and after a stimulus and compared the response of two network connectivities trained with different learning parameters. As membrane potential dynamics are the result of a multitude of parameters, we constructed identical experimental settings in terms of input, to make sure that the differences in the membrane potential dynamics can be linked as closely as possible to the differences in the underlying connectivity. That is, the only major difference between the two settings is the characteristic shape of the connectivity (either being symmetric or asymmetric, see **Figure 5**) resulting from different learning parameters. Since the two models implement a different flow of recurrent excitation, the gain parameters in both networks have been adjusted so that both operate in a similar activity regime in order to enable a meaningful comparison of the temporal characteristics introduced by the connectivity shape. The voltage traces arising from the different network connectivities shown in **Figure 6** exhibit different post-stimulus characteristics during sequence replay, with a faster hyperpolarization happening in networks with asymmetric connectivity. Hence we propose that by aligning the average membrane voltage of a population of neurons—in a perceptual context, to its preferred stimulus and, in a task-related context, to its peak activity—and then analyzing the post-stimulus characteristics of this average voltage, the population's afferent connectivity can be inferred.

4.3. Asymmetric Connectivity Supports Motion Preference

In the context of visual perception of motion, asymmetric connectivity has been found to play an important role in direction-sensitive ganglion cells in the retina (Kim et al., 2008; Vaney et al., 2012). A previous study by Kaplan et al. (2013) proposed asymmetric connectivity as a means of extrapolating the trajectory of a moving stimulus in the absence of a stimulus: similar, in many respects, to the experiment presented here. Recently, this hypothesized tuning property-based connectivity has been confirmed by the observation of neuronal modules in mouse V1 that exhibit similar motion direction preference (Wertz et al., 2015). The model we present here uses a Hebbian-Bayesian rule to explain how such feature-selective connectivity between neurons tuned to similar stimulus features

could arise. It could therefore serve not only as a framework for modeling observed connectivity patterns and helping to understand their functional implications, but also as a means of linking experimentally observed connectivity with earlier modeling studies (Kaplan et al., 2013) by explaining how asymmetric connectivity can arise through learning.

The question of how a preferred sequence direction could be learned and replayed is not only relevant for sensory systems, but also other systems where sequence learning, generation and replay are important (Luczak et al., 2015). We addressed this question by training networks with both symmetric and asymmetric connectivity using a single sequence direction. We then triggered sequence replay in both networks in a similar way to experiments by Gavornik and Bear (2014) and Xu et al. (2012) which studied sequence learning in early visual cortices. Models with symmetric connectivity can show sequence replay in both directions, not only in the trained one. The intuition being that if one were to employ the same training protocol described in Section 3.2, one could replay the sequence forwards or backwards by presenting a cue to the first or last attractor. Instead of being directed by asymmetrical connectivity, the preferred sequence trajectory would evolve according to adaptation. Hence, the direction of the sequence during training alone is not sufficient to create a preferred replay direction as observed in experiments (Xu et al., 2012). Instead, we argue that the asymmetric connectivity caused by a difference in the learning parameters, i.e., an unequal temporal correlation time window, is necessary to replay sequences in only the trained, and therefore preferred, direction.

4.4. Results in Context of Anatomical Data

The presented model addresses the question of how connectivity emerges at a cellular and network level in response to temporally varying stimuli. Through usage of different learning time constants, connectivity kernels of varying widths develop as shown in **Figure 5**. There exists a large body of anatomical evidence reporting regional variations in cortical circuitry in terms of structural features such as dendritic morphology and the density of dendritic spines (see e.g., Jacobs and Scheibel, 2002; Elston, 2003 for reviews). In the visual system the hierarchical organization of areas (Riesenhuber and Poggio, 1999) is reflected in their varying dendritic complexity. When compared to areas such as V1, V2, and V4 which respond to simpler visual features, areas associated with more complex functionality also exhibit more complex dendritic morphologies and have a higher number of connections (Jacobs and Scheibel, 2002; Elston and Fujita, 2014).

It stands to reason that the structural and electrophysiological differences observed in both pyramidal cells and interneurons influences activity on a cellular level (Spruston, 2008), shaping the way in which information is integrated and therefore the functional roles of both the individual cells and the entire circuit (Elston, 2003). These regional variations appear to be consistent across species and to change during development (see Elston and Fujita, 2014 for a recent review). Pyramidal cells in V1 reduce their dendritic complexity and those in inferotemporal and prefrontal areas grow larger dendritic structures over

the first months and years of development. In the light of the presented model, these observations could lead to the interpretation that reducing the dendritic extent of pyramidal cells mirrors an improved perceptual precision in V1 as finer temporal correlations are detected (represented by short learning time constants $\tau_{z_{i,j}}$ and smaller dendritic extent as shown in the panels on the left side of **Figure 5**). In contrast, as more abstract associations are learned, pyramidal cells in higher areas grow more spines over larger dendritic territories. This allows these cells to integrate information from more diverse sources, requiring integration and association to occur over longer time scales (larger $\tau_{z_{i,j}}$). In this context, it is important to note that the learning time constants may not necessarily equal the synaptic time constants (which are determined by receptor and channel kinetics), but could vary depending on the area and with it the function or task to be learned.

Despite the fact that our model uses point neurons and thus does not directly represent the dendritic field, we argue that the learning time-constants determine a neuron's capability to integrate information over time which – given a topographic stimulus representation such as that seen in V1—could be linked to the size of the dendritic field of a neuron. Hence, the presented learning framework offers the possibility to study these arguments in more quantitative detail.

4.5. Scaling the Modular Attractor Model

In Section 3.2 we presented simulations of the modular attractor network model described in Section 2.1 with up to 16 hypercolumns, connected using sparse, random 10 % global connectivity. At this scale each pyramidal cell in the network receives 4.0×10^3 afferent excitatory synapses but—if the model were scaled up to, for example, the scale of the mouse neocortex with approximately 1.6×10^7 neurons (Braitenberg and Schüz, 2013)—each pyramidal cell would receive 1.3×10^6 afferent synapses. As we discussed in Section 4.4, pyramidal cell connectivity varies widely across the layers and areas of the neocortex. However, in this section we base our discussion of the scaling properties of our model on the assumption that each pyramidal cell receives 8.0×10^3 afferent synapses. This number is consistent with averages calculated across cortical layers and areas in mice (Braitenberg and Schüz, 2013), cats (Beaulieu and Colonnier, 1989) and humans (Pakkenberg et al., 2003). The reason this number is significantly lower than the one obtained by naïvely scaling our current model is because of the "patchy" nature of long-range cortical connectivity (Goldman and Nauta, 1977; DeFelipe et al., 1986; Gilbert and Wiesel, 1989; Bosking et al., 1997). Specifically, each pyramidal cell only connects to around 10, approximately hypercolumn-sized, clusters of neurons located within a radius of a few millimeters. Additionally, while each hypercolumn in our model contains 10 minicolumns, biological hypercolumns typically have closer to 100 (Mountcastle, 1997; Buxhoeveden and Casanova, 2002). This means that, because of the winner-take-all dynamics within each hypercolumn, while 10 % of neurons in our model are active at any given time, only 1 % would be active in a more realistic model.

As Sharp and Furber (2013) discuss, when simulating spiking neural networks on SpiNNaker, the majority of CPU time is spent within the event-driven synaptic processing stage, making the CPU load highly dependent on the rate of incoming *synaptic events* (a single spike innervating a single synapse). The combined effect of the more realistic global connectivity and sparser activity discussed in the previous paragraph would be to reduce the rate of incoming synaptic events by a factor of 5 when compared to our current model. This means that a model with more realistic connectivity could actually run faster than the current model on SpiNNaker - Potentially in biological real-time rather than the $0.5\times$ real-time we use in this work.

However, as we discussed in Section 3.2, the time spent actually running simulations on SpiNNaker is often dwarfed by the time spent generating simulation data on the host computer and transferring it to and from the SpiNNaker system. One way of reducing the time taken to generate the simulation data and upload it to SpiNNaker would be to perform some of the data generation on SpiNNaker itself. The most obvious target for this approach would be the generation of the connectivity matrices as, not only do these represent the bulk of the uploaded data, but they are typically defined probabilistically meaning that they could be generated based on a very small uploaded definition. While this approach would undoubtedly reduce the time taken to generate and upload the simulation data, even the 1 min currently taken to download the results at the end of the simulation would grow to several hours if the network was scaled up to the size of even a mouse's neocortex. These slow upload and download times are due to current simulations all having been run on single board SpiNNaker systems, connected to the host computer through a single ethernet link. While the theoretical bandwidth of this link is $100\,\mathrm{Mbit\,s^{-1}}$, inefficiencies in the current SpiNNaker system software reduce the effective bandwidth to only a few $\mathrm{MiB\,s^{-1}}$.

Not only is work currently underway to improve the bandwidth of the ethernet links, but in the case of large-scale network simulations running across multiple SpiNNaker boards, if the host computer is powerful enough and connected to the SpiNNaker system through a sufficiently fast network, data can be transferred to multiple SpiNNaker boards in parallel. Furthermore, if still more bandwidth is required, each SpiNNaker board also has several high-speed serial connectors which could be used for transferring data to and from SpiNNaker at the full $1\,\mathrm{Gbit\,s^{-1}}$ bandwidth of the chip-level interconnect network. Together, the improvements to the scalability of the model discussed in this section would also act to further increase the power efficiency of SpiNNaker when compared to traditional super computer systems that we briefly discuss in Section 3.2.

4.6. Extensions of BCPNN on SpiNNaker and other Future Considerations

Since we have shown that BCPNN learning is possible on SpiNNaker, the implementation we describe in Section 2.5 could be extended to support spike-based reinforcement learning (Izhikevich, 2007) by adding an extra level of E (i.e., "eligibility") traces with time constants between those of the Z and P traces (Tully et al., 2014). Representing downstream

cellular processes that interact with increased intracellular Ca^{2+} concentrations (Yagishita et al., 2014), E traces propagate into the P traces at a rate previously described as κ (Tully et al., 2014). The κ parameter models the delivery of delayed reward signals in the form of interactions with global neuromodulatory systems, which have been linked to the emergence of sequential activity (Gavornik and Bear, 2014; Ikegaya, 2004). Using this extended BCPNN model, the modular attractor memory model we describe in Section 2.1 could be extended to include basal ganglia input (Berthet et al., 2012), allowing it to switch between behavioral sequences when this might be a beneficial strategy for successful task completion (Ponzi and Wickens, 2010).

Similarly, Vogginger et al.'s (2015) original event-driven BCPNN model includes a set of E^* state variables which are used to represent the components of the spike-response model arising from E trace dynamics. Though omitted here, the SpiNNaker BCPNN implementation could be extended to include these traces at the cost of some extra computational cost, and the memory required to store an additional 16 bit trace with each synapse and with each entry in the postsynaptic history structure. In Section 3.1 we showed that by using a 16 bit fixed-point representation for the Z^* and P^* state variables, we can produce results comparable to previous floating-point implementations when both τ_p and f_{max} are relatively small. However, this approach doesn't scale to the type of model described by Fiebig and Lansner (2014) where learning time constants span many orders of magnitude. In these situations, it may be necessary to use a 32 bit fixed-point representation for the P^* traces, further increasing the memory and computational cost of the learning rule.

As spikes from neuromodulator-releasing populations can arrive at the synapse at any time, integrating spike-based reinforcement learning into an event-driven, distributed simulation requires incorporating the times of modulatory as well as postsynaptic spikes into algorithm 1. Because entire populations of neuromodulator-releasing neurons can affect the modulatory input received by a single synapse, the per-neuron history structure discussed in Section 2.4 is not a viable means of storing them. Potjans et al. (2010) extend Morrison et al.'s (2007) STDP algorithm to support neuromodulated learning by introducing "volume transmitter" populations which handle all the incoming modulatory input to a virtual "volume." These populations maintain a spike-history of all incoming modulatory spikes which they deliver to the synapses of neuronal populations within this volume, both at presynaptic spike times and after a fixed period so as to 'flush out' the spike-history data structure and allow it to be kept relatively small. This approach has the potential to map well to the SpiNNaker architecture and could be used as the basis of a future SpiNNaker implementation of spike-based reinforcement learning using BCPNN.

A benefit of the model proposed here is its robustness and flexibility. Non-sequential attractor networks without learning have previously been emulated on a neuromorphic microchip (Pfeil et al., 2013) and on a simulated version of the BrainScaleS system (Petrovici et al., 2014). Though not shown here, the connectivity required by these types of randomly hopping attractor networks can also be learned. Variations of this network run on supercomputers have been shown to account for disparate cognitive phenomena including perceptual rivalry and completion (Kaplan and Lansner, 2013); attentional blink (Lundqvist et al., 2006; Silverstein and Lansner, 2011); and diverse oscillatory regimes (Lundqvist et al., 2010). But our model was a reduced version of previous detailed ones insofar that we did not utilize Hodgkin-Huxley neurons with calcium-dependent potassium channels or regular spiking non-pyramidal cells; nor did we explicitly model connections among basket cells, saturating synapses, a Vm-dependent Mg^{2+} blockade or short-term depression.

A problem not stressed by the aforementioned models is how the connectivity required for stable activity propagation might be learned (Wörgötter and Porr, 2005; Kunkel et al., 2011), despite the biochemical (Peters et al., 2014) and metabolic (Picard et al., 2013) changes accompanying learned sequential behaviors. Several promising approaches have been developed (Sussillo and Abbott, 2009; Laje and Buonomano, 2013; Hennequin et al., 2014), albeit with biological motivations driven more from the perspective of algorithmic optimization, rather than from bottom-up neural processing. Here, we have shown that activity could propagate through recurrent cortical microcircuits as a result of a probabilistic learning rule based on neurobiologically plausible time courses and dynamics. The model predicts that the interaction between several learning and dynamical processes constitute a compound mnemonic engram that can flexibly generate step-wise sequential increases of activity within pools of excitatory neurons. We have shown that this large-scale learning model can be efficiently simulated at scale using neuromorphic hardware and our simulations suggest that flexible systems such as SpiNNaker offer a promising tool for the study of collective dynamical phenomena emerging from the complex interactions occurring between individual neurons and synapses whose properties change over time.

AUTHOR CONTRIBUTIONS

ACKNOWLEDGMENTS

REFERENCES

Abbott, L. F., and Blum, K. I. (1996). Functional significance of long-term potentiation for sequence learning and prediction. *Cereb. Cortex* 6, 406–416. doi: 10.1093/cercor/6.3.406

Abeles, M., Bergman, H., Gat, I., Meilijson, I., Seidemann, E., Tishby, N., et al. (1995). Cortical activity flips among quasi-stationary states. *Proc. Natl. Acad. Sci. U.S.A.* 92, 8616–8620. doi: 10.1073/pnas.92.19.8616

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

Bar, M. (2007). The proactive brain: using analogies and associations to generate predictions. *Trends Cogn. Sci.* 11, 280–289. doi: 10.1016/j.tics.2007.05.005

Barlow, H. (2001). Redundancy reduction revisited. *Network* 12, 241–253. doi: 10.1080/net.12.3.241.253

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., and Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron* 76, 695–711. doi: 10.1016/j.neuron.2012.10.038

Beaulieu, C., and Colonnier, M. (1989). Number and size of neurons and synapses in the motor cortex of cats raised in different environmental complexities. *J. Comp. Neurol.* 289, 178–187. doi: 10.1002/cne.902890115

Benjamin, B. V., Gao, P., McQuinn, E., Choudhary, S., Chandrasekaran, A. R., Bussat, J. M., et al. (2014). Neurogrid: a mixed-analog-digital multichip system for large-scale neural simulations. *Proc. IEEE* 102, 699–716. doi: 10.1109/JPROC.2014.2313565

Berry, M. J. II, Brivanlou, I. H., Jordan, T. A., and Meister, M. (1999). Anticipation of moving stimuli by the retina. *Nature* 398, 334–338. doi: 10.1038/18678

Berthet, P., Hellgren-Kotaleski, J., and Lansner, A. (2012). Action selection performance of a reconfigurable basal ganglia inspired model with hebbian–bayesian go-nogo connectivity. *Front. Behav. Neurosci.* 6:65. doi: 10.3389/fnbeh.2012.00065

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.

Bosking, W. H., Zhang, Y., Schofield, B., and Fitzpatrick, D. (1997). Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *J. Neurosci.* 17, 2112–2127.

Braitenberg, V., and Schüz, A. (2013). *Cortex: Statistics and Geometry of Neuronal Connectivity*. Berlin; Heidelberg; New York, NY: Springer Science & Business Media.

Brette, R., and Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J. Neurophysiol.* 94, 3637–3642. doi: 10.1152/jn.00686.2005

Bubic, A., von Cramon, D. Y., and Schubotz, R. I. (2010). Prediction, cognition and the brain. *Front. Hum. Neurosci.* 4:25. doi: 10.3389/fnhum.2010.00025

Butz, M. V., Sigaud, O., and Gérard, P. (2003). "Internal models and anticipations in adaptive learning systems," in *Anticipatory Behavior in Adaptive Learning Systems*, eds M. V. Butz, O. Sigaud, and P. Gérard (Heidelberg; Berlin: Springer), 86–109. doi: 10.1007/978-3-540-45002-3_6

Buxhoeveden, D. P., and Casanova, M. F. (2002). The minicolumn hypothesis in neuroscience. *Brain* 125, 935–951. doi: 10.1093/brain/awf110

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

Cossart, R., Aronov, D., and Yuste, R. (2003). Attractor dynamics of network up states in the neocortex. *Nature* 423, 283–288. doi: 10.1038/nature01614

Cray (2013). Cray XC30-ACø Supercomputer. Technical Report, Cray Inc.

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

Daoudal, G., and Debanne, D. (2003). Long-term plasticity of intrinsic excitability: learning rules and mechanisms. *Learn. Mem.* 10, 456–465. doi: 10.1101/lm.64103

DeFelipe, J., Conley, M., and Jones, E. (1986). Long-range focal collateralization of axons arising from corticocortical cells in monkey sensory-motor cortex. *J. Neurosci.* 6, 3749–3766.

Diehl, P. U., and Cook, M. (2014). "Efficient Implementation of STDP Rules on SpiNNaker Neuromorphic Hardware," in *Neural Networks (IJCNN), The 2014 International Joint Conference on* (Beijing), 4288–4295. doi: 10.1109/IJCNN.2014.6889876

Djurfeldt, M., Lundqvist, M., Johansson, C., Rehn, M., Ekeberg, O., and Lansner, A. (2008). Brain-scale simulation of the neocortex on the IBM Blue Gene/L supercomputer. *IBM J. Res. Dev.* 52, 31–41. doi: 10.1147/rd.521.0031

Douglas, R. J., and Martin, K. A. (2004). Neuronal circuits of the neocortex. *Annu. Rev. Neurosci.* 27, 419–451. doi: 10.1146/annurev.neuro.27.070203.144152

Elston, G. N. (2003). Cortex, cognition and the cell: new insights into the pyramidal neuron and prefrontal function. *Cereb. Cortex* 13, 1124–1138. doi: 10.1093/cercor/bhg093

Elston, G. N., and Fujita, I. (2014). Pyramidal cell development: postnatal spinogenesis, dendritic growth, axon growth, and electrophysiology. *Front. Neuroanat.* 8:78. doi: 10.3389/fnana.2014.00078

Enns, J. T., and Lleras, A. (2008). What's next? new evidence for prediction in human vision. *Trends Cogn. Sci.* 12, 327–333. doi: 10.1016/j.tics.2008.06.001

Feldman, D. E. (2012). The spike-timing dependence of plasticity. *Neuron* 75, 556–571. doi: 10.1016/j.neuron.2012.08.001

Fiebig, F., and Lansner, A. (2014). Memory consolidation from seconds to weeks : a three-stage neural memory consolidation from seconds to weeks : a three-stage neural network model with autonomous reinstatement dynamics. *Front. Comput. Neurosci.* 8:64. doi: 10.3389/fncom.2014.00064

Furber, S. B., Galluppi, F., Temple, S., and Plana, L. A. (2014). The SpiNNaker project. *Proc. IEEE* 102, 652–665. doi: 10.1109/JPROC.2014.2304638

Galluppi, F., Lagorce, X., Stromatias, E., Pfeiffer, M., Plana, L. A., Furber, S. B., et al. (2015). A framework for plasticity implementation on the SpiNNaker neural architecture. *Front. Neurosci.* 8:429. doi: 10.3389/fnins.2014.00429

Gavornik, J. P., and Bear, M. F. (2014). Learned spatiotemporal sequence recognition and prediction in primary visual cortex. *Nat. Neurosci.* 17, 732–737. doi: 10.1038/nn.3683

George, D., and Hawkins, J. (2009). Towards a mathematical theory of cortical micro-circuits. *PLoS Comput. Biol.* 5:e1000532. doi: 10.1371/journal.pcbi.1000532

Gerstner, W., and Kistler, W. M. (2002). *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge: Cambridge University Press. doi: 10.1017/cbo9780511815706

Gewaltig, M.-O., and Diesmann, M. (2007). NEST (NEural simulation tool). *Scholarpedia* 2:1430. doi: 10.4249/scholarpedia.1430

Gilbert, C. D., and Wiesel, T. N. (1989). Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *J. Neurosci.* 9, 2432–2442.

Goldman, P. S., and Nauta, W. J. (1977). Columnar distribution of cortico-cortical fibers in the frontal association, limbic, and motor cortex of the developing rhesus monkey. *Brain Res.* 122, 393–413. doi: 10.1016/0006-8993(77)90453-x

Gütig, R., Aharonov, R., Rotter, S., and Sompolinsky, H. (2003). Learning input correlations through nonlinear temporally asymmetric hebbian plasticity. *J. Neurosci.* 23, 3697–3714.

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

Hopkins, M., and Furber, S. (2015). Accuracy and efficiency in fixed-point neural ODE solvers. *Neural Comput.* 27, 2148–2182. doi: 10.1162/NECO_a_00772

Hosoya, T., Baccus, S. A., and Meister, M. (2005). Dynamic predictive coding by the retina. *Nature* 436, 71–77. doi: 10.1038/nature03689

Ikegaya, Y. (2004). Synfire chains and cortical songs: temporal modules of cortical activity. *Science* 304, 559–564. doi: 10.1126/science.1093173

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

Jacobs, B., and Scheibel, A. (2002). "Regional dendritic variation in primate cortical pyramidal cells," in *Cortical Areas: Unity Diversity*, eds A. Schuez and R. Miller (London; New York, NY: Tailor and Francis), 111–131. doi: 10.1201/9780203299296.pt2

Jin, X., Rast, A., Galluppi, F., Davies, S., and Furber, S. (2010). "Implementing spike-timing-dependent plasticity on SpiNNaker neuromorphic hardware," in *The 2010 International Joint Conference on Neural Networks (IJCNN)*, Number 1 (Barcelona: IEEE), 1–8. doi: 10.1109/ijcnn.2010.5596372

Johansson, C., and Lansner, A. (2007). Towards cortex sized artificial neural systems. *Neural Netw.* 20, 48–61. doi: 10.1016/j.neunet.2006.05.029

Jones, L. M., Fontanini, A., Sadacca, B. F., Miller, P., and Katz, D. B. (2007). Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc. Natl. Acad. Sci. U.S.A.* 104, 18772–18777. doi: 10.1073/pnas.0705546104

Jung, S.-C., and Hoffman, D. A. (2009). Biphasic somatic a-type k+ channel downregulation mediates intrinsic plasticity in hippocampal ca1 pyramidal neurons. *PLoS ONE* 4:e6549. doi: 10.1371/journal.pone.0006549

Kaplan, B. A., and Lansner, A. (2013). A spiking neural network model of self-organized pattern recognition in the early mammalian olfactory system. *Front. Neural Circuits* 8:5. doi: 10.3389/fncir.2014.00005

Kaplan, B. A., Lansner, A., Masson, G. S., and Perrinet, L. U. (2013). Anisotropic connectivity implements motion-based prediction in a spiking neural network. *Front. Comput. Neurosci.* 7:112. doi: 10.3389/fncom.2013.00112

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

Kim, I.-J., Zhang, Y., Yamagata, M., Meister, M., and Sanes, J. R. (2008). Molecular identification of a retinal cell type that responds to upward motion. *Nature* 452, 478–482. doi: 10.1038/nature06739

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

Kunkel, S., Potjans, T. C., Eppler, J. M., Plesser, H. E., Morrison, A., and Diesmann, M. (2012). Meeting the memory challenges of brain-scale network simulation. *Front. Neuroinform.* 5:35. doi: 10.3389/fninf.2011.00035

Lagorce, X., Stromatias, E., Galluppi, F., Plana, L. A., Liu, S.-C., Furber, S. B., et al. (2015). Breaking the millisecond barrier on SpiNNaker: implementing asynchronous event-based plastic models with microsecond resolution. *Front. Neurosci.* 9:206. doi: 10.3389/fnins.2015.00206

Laje, R., and Buonomano, D. V. (2013). Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci.* 16, 925–933. doi: 10.1038/nn.3405

Lansner, A., and Ekeberg, Ö. (1989). A one-layer feedback artificial neural network with a bayesian learning rule. *Int. J. Neural Syst.* 1, 77–87. doi: 10.1142/S0129065789000499

Lansner, A., and Holst, A. (1996). A higher order Bayesian neural network with spiking units. *Int. J. Neural Syst.* 7, 115–128. doi: 10.1142/s0129065796000816

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. Synaptic Neurosci.* 2:140. doi: 10.3389/fnsyn.2010.00140

Litwin-Kumar, A., and Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat. Neurosci.* 15, 1498–1505. doi: 10.1038/nn.3220

Liu, Y. H., and Wang, X. J. (2001). Spike-frequency adaptation of a generalized leaky integrate-and-fire model neuron. *J. Comput. Neurosci.* 10, 25–45. doi: 10.1023/A:1008916026143

Luczak, A., Barthó, P., Marguet, S. L., Buzsáki, 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

Luczak, A., McNaughton, B. L., and Harris, K. D. (2015). Packet-based communication in the cortex. *Nat. Rev. Neurosci.* 16, 745–755. doi: 10.1038/nrn4026

Lundqvist, M., Compte, A., and Lansner, A. (2010). Bistable, irregular firing and population oscillations in a modular attractor memory network. *PLoS Comput. Biol.* 6:e1000803. doi: 10.1371/journal.pcbi.1000803

Lundqvist, M., Herman, P., and Lansner, A. (2012). Variability of spike firing during theta-coupled replay of memories in a simulated attractor network. *Brain Res.* 1434, 152–161. doi: 10.1016/j.brainres.2011.07.055

Lundqvist, M., Rehn, M., Djurfeldt, M., and Lansner, A. (2006). Attractor dynamics in a modular network model of neocortex. *Network* 17, 253–276. doi: 10.1080/09548980600774619

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

McClelland, J. L., Rumelhart, D. E., and Hinton, G. E. (1986). *The Appeal of Parallel Distributed Processing*. Cambridge, MA: MIT Press.

Mead, C. (1990). Neuromorphic electronic systems. *Proc. IEEE* 78, 1629–1636. doi: 10.1109/5.58356

Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., et al. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 668–673. doi: 10.1126/science.1254642

Moise, M. (2012). *A Fixed Point Arithmetic Library for SpiNNaker*. PhD thesis, The University of Manchester, Manchester.

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

Morrison, A., Mehring, C., Geisel, T., Aertsen, A. D., and Diesmann, M. (2005). Advancing the boundaries of high-connectivity network simulation with distributed computing. *Neural Comput.* 17, 1776–1801. doi: 10.1162/0899766054026648

Mountcastle, V. B. (1997). The columnar organization of the neocortex. *Brain* 120, 701–722. doi: 10.1093/brain/120.4.701

Nordlie, E., Gewaltig, M. O., and Plesser, H. E. (2009). Towards reproducible descriptions of neuronal network models. *PLoS Comput. Biol.* 5:e1000456. doi: 10.1371/journal.pcbi.1000456

Painkras, E., and Plana, L. (2013). SpiNNaker: a 1-W 18-core system-on-chip for massively-parallel neural network simulation. *IEEE J. Solid-State Circ.* 48, 1943–1953. doi: 10.1109/JSSC.2013.2259038

Pakkenberg, B., Pelvig, D., Marner, L., Bundgaard, M. J., Gundersen, H. J. G., Nyengaard, J. R., et al. (2003). Aging and the human neocortex. *Exp. Gerontol.* 38, 95–99. doi: 10.1016/S0531-5565(02)00151-1

Peters, A. J., Chen, S. X., and Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. *Nature* 510, 263–267. doi: 10.1038/nature13235

Petrovici, M. A., Vogginger, B., Müller, P., Breitwieser, O., Lundqvist M., Muller, L., et al. (2014). Characterization and compensation of network-level anomalies in mixed-signal neuromorphic modeling platforms. *PLoS ONE* 9:e108590. doi: 10.1371/journal.pone.0108590. Available online at: http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0108590

Pfeil, T., Grübl, A., Jeltsch, S., Müller, E., Müller, P., Petrovici, M. A., et al. (2013). Six networks on a universal neuromorphic computing substrate. *Front. Neurosci.* 7:11. doi: 10.3389/fnins.2013.00011

Picard, N., Matsuzaka, Y., and Strick, P. L. (2013). Extended practice of a motor skill is associated with reduced metabolic activity in m1. *Nat. Neurosci.* 16, 1340–1347. doi: 10.1038/nn.3477

Ponzi, A., and Wickens, J. (2010). Sequentially switching cell assemblies in random inhibitory networks of spiking neurons in the striatum. *J. Neurosci.* 30, 5894–5911. doi: 10.1523/JNEUROSCI.5540-09.2010

Potjans, W., Morrison, A., and Diesmann, M. (2010). Enabling functional neural circuit simulations with distributed computing of neuromodulated plasticity. *Front. Comput. Neurosci.* 4:141. doi: 10.3389/fncom.2010.00141

Rao, R. P., and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87. doi: 10.1038/4580

Rauch, A., La Camera, G., Lüscher, H.-R., Senn, W., and Fusi, S. (2003). Neocortical pyramidal cells respond as integrate-and-fire neurons to *in vivo*–like input currents. *J. Neurophysiol.* 90, 1598–1612. doi: 10.1152/jn.00293.2003

Riegler, A. (2001). "The role of anticipation in cognition," in *AIP Conference Proceedings* (Liege: IOP Institute of Physics Publishing Ltd), 534–544. doi: 10.1063/1.1388719

Riesenhuber, M., and Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nat. Neurosci.* 2, 1019–1025. doi: 10.1038/14819

Rohenkohl, G., Cravo, A. M., Wyart, V., and Nobre, A. C. (2012). Temporal expectation improves the quality of sensory information. *J. Neurosci.* 32, 8424–8428. doi: 10.1523/JNEUROSCI.0804-12.2012

Sandberg, A., Lansner, A., Petersson, K. M., and Ekeberg, O. (2002). A Bayesian attractor network with incremental learning. *Network* 13, 179–194. doi: 10.1080/net.13.2.179.194

Schemmel, J., Bruderle, D., Grubl, A., Hock, M., Meier, K., and Millner, S. (2010). "A wafer-scale neuromorphic hardware system for large-scale neural modeling," in *Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on* (Paris: IEEE), 1947–1950. doi: 10.1109/ISCAS.2010.5536970

Schlerf, J., Ivry, R. B., and Diedrichsen, J. (2012). Encoding of sensory prediction errors in the human cerebellum. *J. Neurosci.* 32, 4913–4922. doi: 10.1523/JNEUROSCI.4504-11.2012

Seidemann, E., Meilijson, I., Abeles, M., Bergman, H., and Vaadia, E. (1996). Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J. Neurosci.* 16, 752–768.

Shadmehr, R., Smith, M. A., and Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Ann. Rev. Neurosci.* 33, 89–108. doi: 10.1146/annurev-neuro-060909-153135

Sharp, T., and Furber, S. (2013). "Correctness and performance of the SpiNNaker architecture," in *Neural Networks (IJCNN), The 2013 International Joint Conference on* (Dallas, TX). doi: 10.1109/ijcnn.2013.6706988

Sharp, T., Galluppi, F., Rast, A., and Furber, S. (2012). Power-efficient simulation of detailed cortical microcircuits on SpiNNaker. *J. Neurosci. Methods* 210, 110–118. doi: 10.1016/j.jneumeth.2012.03.001

Sharp, T., Petersen, R., and Furber, S. (2014). Real-time million-synapse simulation of rat barrel cortex. *Front. Neurosci.* 8:131. doi: 10.3389/fnins.2014.00131

Sheik, S., Coath, M., Indiveri, G., Denham, S. L., Wennekers, T., and Chicca, E. (2012). Emergent auditory feature tuning in a real-time neuromorphic VLSI system. *Front. Neurosci.* 6:17. doi: 10.3389/fnins.2012.00017

Silverstein, D. N., and Lansner, A. (2011). Is attentional blink a byproduct of neocortical attractors? *Front. Comput. Neurosci.* 5:13. doi: 10.3389/fncom.2011.00013

Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. *Nat. Rev. Neurosci.* 9, 206–221. doi: 10.1038/nrn2286

Stromatias, E., Galluppi, F., Patterson, C., and Furber, S. (2013). "Power analysis of large-scale, real-time neural networks on SpiNNaker," in *The 2013 International Joint Conference on Neural Networks (IJCNN)* (Dallas, TX), 1–8. doi: 10.1109/IJCNN.2013.6706927

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

Tully, P. J., Hennig, M. H., and Lansner, A. (2014). Synaptic and nonsynaptic plasticity approximating probabilistic inference. *Front. Synaptic Neurosci.* 6:8. doi: 10.3389/fnsyn.2014.00008

Vaney, D. I., Sivyer, B., and Taylor, W. R. (2012). Direction selectivity in the retina: symmetry and asymmetry in structure and function. *Nat. Rev. Neurosci.* 13, 194–208. doi: 10.1038/nrn3165

Vogginger, B., Schüffny, R., Lansner, A., Cederström, L., Partzsch, J., and Höppner, S. (2015). Reducing the computational footprint for real-time BCPNN learning. *Front. Neurosci.* 9:2. doi: 10.3389/fnins.2015.00002

Wang, S.-J., Hilgetag, C. C., and Zhou, C. (2011). Sustained activity in hierarchical modular neural networks: self-organized criticality and oscillations. *Front. Comput. Neurosci.* 5:30. doi: 10.3389/fncom.2011.00030

Wertz, A., Trenholm, S., Yonehara, K., Hillier, D., Raics, Z., Leinweber, M., et al. (2015). Single-cell–initiated monosynaptic tracing reveals layer-specific cortical network modules. *Science* 349, 70–74. doi: 10.1126/science.aab1687

Wörgötter, F., and Porr, B. (2005). Temporal sequence learning, prediction, and control: a review of different models and their relation to biological mechanisms. *Neural Comput.* 17, 245–319. doi: 10.1162/0899766053011555

Xu, S., Jiang, W., Poo, M.-M., and Dan, Y. (2012). Activity recall in a visual cortical ensemble. *Nat. Neurosci.* 15, 449–455. doi: 10.1038/nn.3036

Yagishita, S., Hayashi-Takagi, A., Ellis-Davies, G. C., Urakubo, H., Ishii, S., and Kasai, H. (2014). A critical time window for dopamine actions on the structural plasticity of dendritic spines. *Science* 345, 1616–1620. doi: 10.1126/science.1255514

Yoshida, T., and Katz, D. B. (2011). Control of prestimulus activity related to improved sensory coding within a discrimination task. *J. Neurosci.* 31, 4101–4112. doi: 10.1523/JNEUROSCI.4380-10.2011

Yoshimura, Y., Dantzker, J. L., and Callaway, E. M. (2005). Excitatory cortical neurons form fine-scale functional networks. *Nature* 433, 868–873. doi: 10.1038/nature03252

Odor Experience Facilitates Sparse Representations of New Odors in a Large-Scale Olfactory Bulb Model

Shanglin Zhou[1], Michele Migliore[2,3] and Yuguo Yu[1]*

[1] School of Life Science and The Collaborative Innovation Center for Brain Science, The Center for Computational Systems Biology, Fudan University, Shanghai, China, [2] Division of Palermo, Institute of Biophysics, National Research Council, Palermo, Italy, [3] Department of Neurobiology, Yale University School of Medicine, New Haven, CT, USA

Prior odor experience has a profound effect on the coding of new odor inputs by animals. The olfactory bulb, the first relay of the olfactory pathway, can substantially shape the representations of odor inputs. How prior odor experience affects the representation of new odor inputs in olfactory bulb and its underlying network mechanism are still unclear. Here we carried out a series of simulations based on a large-scale realistic mitral-granule network model and found that prior odor experience not only accelerated formation of the network, but it also significantly strengthened sparse responses in the mitral cell network while decreasing sparse responses in the granule cell network. This modulation of sparse representations may be due to the increase of inhibitory synaptic weights. Correlations among mitral cells within the network and correlations between mitral network responses to different odors decreased gradually when the number of prior training odors was increased, resulting in a greater decorrelation of the bulb representations of input odors. Based on these findings, we conclude that the degree of prior odor experience facilitates degrees of sparse representations of new odors by the mitral cell network through experience-enhanced inhibition mechanism.

Keywords: odor representation, prior experience, sparse representation, olfactory bulb, large scale network

INTRODUCTION

Prior sensory experience is very important for animals in learning and processing novel incoming signals. In olfaction, prior odor experience can significantly improve the ability of the animal to discriminate new odor inputs (Mandairon et al., 2006a,b,c; Mandairon and Linster, 2009; Sinding et al., 2011). The olfactory bulb is the first relay of the olfactory pathway, and encodes odor inputs as the network responses of mitral cells (Kay and Sherman, 2007; Mandairon and Linster, 2009). The olfactory bulb has been observed to encode signals in a spatiotemporally sparse and decorrelated manner (Khan et al., 2010; Yu et al., 2014). Moreover, it has been observed that mitral cells become less responsive after prior odor exposure (Buonviso et al., 1998; Buonviso and Chaput, 2000; Fletcher and Wilson, 2003; Mandairon and Linster, 2009; Kato et al., 2012). On the other hand, it has been shown that interneurons may become more (Mandairon et al., 2008) or less (Kato et al., 2012) responsive with new odors.

In previous experimental and computational studies, the numbers of prior odor experiences and new incoming odors are limited. How an animal's prior experience with odorants affects the representation by the olfactory bulb (i.e., the firing patterns of mitral and granule cells) in response

to new odors is an open question. Considering the limitations of current experimental techniques, it is nearly impossible to access the synaptic dynamics or neuronal response to odor inputs in the olfactory bulb network at a large scale. However, large-scale supercomputer simulation of realistic olfactory bulb models has been employed to carry out a series of simulations examining these issues (Yu et al., 2013, 2014; Migliore et al., 2014, 2015). Our previous reports have shown that a sparse spatial spiking representation of specific odor signals can emerge naturally from mitral-granule interactions and can be realistically implemented by our model via balanced excitatory-inhibitory synapses (Yu et al., 2013, 2014). Here, we examine how and to what extent prior odor experience modulates the excitatory and inhibitory interactions and how they shape odor representations.

To address these issues, we performed a series of simulations based on a previously established large-scale olfactory bulb model (Yu et al., 2013, 2014). The simulation results show that prior odor experience can accelerate the formation of sparseness in the mitral cell network in response to new odors. Furthermore, the sparseness of the mitral cell network is increased but the sparseness of the granule cell network is decreased with an increasing number of prior training odors. Further analysis demonstrated that this phenomenon is accompanied by a nonlinear change in the excitatory and inhibitory synaptic weighting of the network. Mitral cell network responses demonstrated a gradual increase in their intrinsic decorrelation property, suggesting an increased odor discrimination ability.

MATERIAL AND METHODS

Computational Simulations

All simulations were carried out with the NEURON simulation program v7.3 (Hines and Carnevale, 1997, 2001) on a Cray XC30 system (INCF, Sweden). All the present work was based on a previously verified scaled-up olfactory bulb model (Yu et al., 2013, 2014). Briefly, The network was composed of multi-compartment canonical models of 500 mitral and 10,000 granule cells, implemented as described in our previous studies (Migliore and Shepherd, 2008; Migliore et al., 2010).The model uses a reduced number of MCs and granule cells (glom: MC: GC = 1: 5: 100). As we have already explained in detail in a previous paper (Yu et al., 2013), the reason for this choice is that our main aim with this model is to understand the basic processes underlying lateral and feedback inhibition in a network. To this purpose the full number of cells is not needed, especially in the presence of experimental data limited to a very small subset of glomeruli; however, the relative ratio between mitral and granule cells is consistent with experimental estimations, validated against a number of experimental findings (Willhite et al., 2006; Shusterman et al., 2011). The canonical model for mitral cells was implemented with 312 compartments representing an axon, the soma, the apical dendrite, and two lateral dendrites each 1.5 mm in length, in the range indicated by anatomical measurements (Mori et al., 1981). Real mitral cells have a number of lateral dendrites that cover a relatively large, bidimensional surrounding area. From this point of view,

our simplifying choice of using only two lateral dendrites per mitral cell has the obvious limitation that, since many glomeruli are at variable distances from the single projection tract, the interactions between mitral cells belonging to specific neighboring glomeruli are not precisely represented. However, our additional choice to project the glomeruli into a single tract, results in the interactions of a given mitral cell with many nearby mitral cells still holding in a generic sense, so that the model gives a relatively accurate reflection of these population interactions within the mitral-granule network. In this way we were able to maintain the requirements for computational resources within a reasonable limit. An indirect proof of the overall quality of this model is its qualitative agreement with a number of experimental findings (Yu et al., 2013). Uniform passive properties were used, with $R_a = 150\ \Omega\cdot cm$, $\tau_m = 20\ ms$, and R_m and C_m adjusted to obtain an input resistance of about 100 MΩ. Resting potential was set at -65 mV and temperature at 35°C. Cells were modeled as regular firing cells (Migliore et al., 2005), with Na, KA, and KDR conductances uniformly distributed over the entire dendritic tree (Bischofberger and Jonas, 1997). Kinetics for the Na conductance were from hippocampal pyramidal neurons (Migliore et al., 1999), whereas those for KA and KDR were from mitral cell data (Wang et al., 1996). Granule cells were modeled with a soma and a 20 segment radial dendrite (250 μm of total length) representing the dendritic tree. Na$^+$ and KA channels were distributed throughout (Schoppa and Westbrook, 1999; Pinato and Midtgaard, 2005; Zelles et al., 2006) whereas KDR was present only in the soma (Schoppa and Westbrook, 1999).

Effective dendrodendritic coupling between granule cell synapses and mitral cell secondary dendrites was implemented by connecting a GC synapse, containing the same proportion of AMPA and NMDA channels, with the appropriate compartment of mitral cell GABA channel-containing secondary dendrites. The details of the synaptic mechanisms have been described in our previous work (Yu et al., 2013, 2014). It should be noted that we applied a generic use-dependent plasticity rule to the dendrodendritic connection. Briefly, all synaptic weights started at zero and, in response to an odor input, the components (inhibitory or excitatory) of each dendrodendritic synapse were independently modified according to local spiking activity in the lateral dendrite of the mitral cell or the granule cell synapse. After each spike, the peak conductance (w) and the state (p) of any given synapse were updated from their current value $w_{\{exc,inh\},p} = g_{max,\{exc,inh\}}\cdot S(p)$ to a new value. The new values were calculated according to the instantaneous presynaptic interspike interval (ISI) (see Migliore et al., 2007) as $w_{\{exc,inh\},p+\Delta} = g_{max,\{exc,inh\}}\cdot S(p+\Delta)$. The value of p was limited to the range 0–50, and is subjected to the classical scheme $\Delta = \{0, +1, -1\}$ (Stanton, 1996) in which $\Delta = 0$ for an ISI ≥ 250 ms (i.e., no changes for spike rates ≤ 4 Hz), $\Delta = -1$ for $33 < ISI < 250$ ms (LTD in the range of 4–30 Hz), and $\Delta = 1$ for ISI ≤ 33 ms (LTP for a spike rate ≥ 30 Hz). The sigmoidal activation function $S(p)$ was defined as S $(p) = 1/\{1 + exp[(25 - p)/3]\}$ (Haykin, 1994). In this way, the weight (i.e., the peak synaptic conductance) of any given synapse could transition from a fully depressed ($w \approx 0$, for $p = 0$) to a fully potentiated state ($w \approx g_{max}$, for $p = 50$), or vice-versa, over a span of 50 consecutive spikes of the appropriate

frequency. At the beginning of a simulation $p = 0$, the spikes resulting in values of $p < 0$ or > 50 were ignored.

It should be stressed that synaptic plasticity is fundamental to any dynamic network. Although in the mitral-granule circuit it has not been observed directly, we consider this lack of information as a shortcoming of the experimental techniques rather than a demonstration that there is no plasticity in the olfactory bulb. Indeed, recent studies have shown more or less direct evidence for long term plasticity of olfactory input in mitral cells (Ennis et al., 1998; Ma et al., 2012), and in granule cells (Patneau and Stripling, 1992; Gao and Strowbridge, 2009; Arenkiel et al., 2011). Also note that the plasticity rule used in this model has already been shown (Yu et al., 2013) to generate synaptic clusters and firing patterns in qualitative agreement with experimental findings. As discussed in detail elsewhere (Xiong and Chen, 2002; Migliore and Shepherd, 2008), the formation of synaptic clusters consistent with those observed experimentally is an extremely robust process that can be understood by considering the follow dynamics: (a) a strong odor input causes mitral cells to fire at high-frequency; (b) somatic APs backpropagate along the lateral dendrites and potentiate excitatory mitral–granule synapses along their way, activating granule cells; (c) granule cells begin to fire at high-frequency, potentiating inhibitory synapses on the lateral dendrites of mitral cells, (d) inhibition from granule cells hinders AP back-propagation as it travels far from the soma, thus reducing, locally, the firing frequency of mitral and granule cells, and (e) this finally results in the selective depression of synapses far from the soma of the active mitral cell. Therefore, as long as: (1) action potentials backpropagate along the mitral cell lateral dendrites, (2) granule cells form dendrodendritic connections, and (3) LTD and LTP are induced by different levels of synaptic activity, a column will form independently from the specific learning rule. This mechanism is robust and independent of the plasticity rule used to update the synaptic weights during a simulation (Migliore et al., 2007, 2010); we have tested it with hebbian, non-hebbian, and spike-time-dependent plasticity, obtaining in all cases the same qualitative result (i.e., the formation of a column).

It should be noted that in this paper we were interested in the results obtained for a relatively high odor concentration, which is needed to form glomerular units as observed in the experiments. The overall amount of LTP or LTD obtained in a real system, and its overall effect on the I/O properties, will of course depend from the actual plasticity rules in effect for mitral and granule cells. There are no sufficient experimental indications on these processes. However, we stress that the plasticity rule used in this model has already been shown (Yu et al., 2013) to generate synaptic clusters and firing patterns in qualitative agreement with experimental findings.
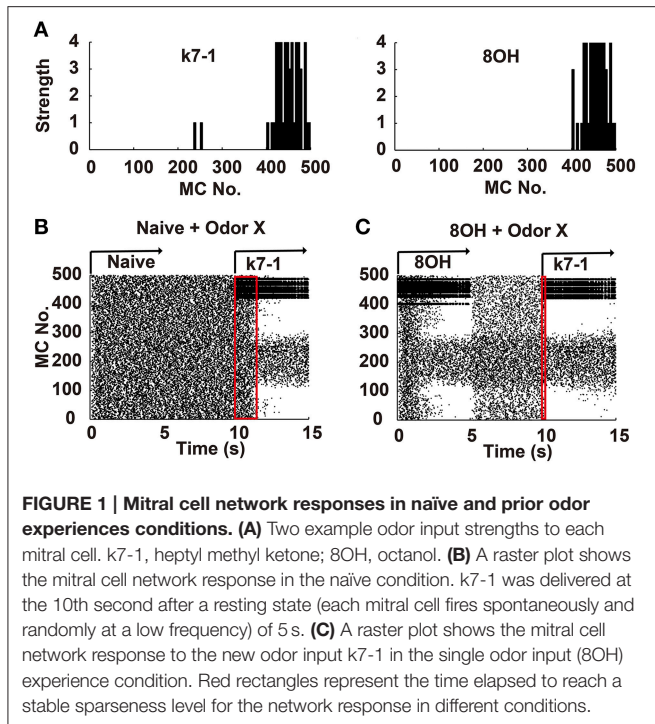
Other details of the model were identical to those described previously (Yu et al., 2013, 2014). The simulation codes used to run the simulations described in the present work are available in the ModelDB database (http://senselab.med.yale.edu/modeldb, accession number 144570), with the exception of run control files. Kinetic equations and implementation details for all ionic currents are described in these model files.

Odor Input Paradigm

In our model, the network contains 100 glomeruli, 500 mitral cells, and 10,000 granule cells. The 100 glomeruli spatially distributed within which 74 glomeruli are chosen to have active responses to represent the spatial responses to 72 different odor stimuli. Each glomerulus makes synaptic connections with five mitral cells. For those 74 glomeruli, there are 370 spatially distributed mitral cells connected to them. The other 130 mitral cells are connected to other 26 glomeruli (which could be stimulated by new odors, other than the present 72 odors). We distributed them in such a way to have a roughly uniform overall spatial distribution of glomeruli. Note that although there is no odor input feeding to those 130 mitral cells, their firing is modulated by the random background activity and by the lateral inhibition received from granule cells that are connected with odor-activated mitral cells. As described in our previous work, 72 odor inputs were used for simulations (Yu et al., 2013, 2014). The basic activation strength (0–4) for each glomerulus and each odor is taken directly from the experimental values kindly provided by Mori et al. (2006). To simulate an odor presentation, these values are multiplied by a coefficient representing the odor concentration, and that resulted in an aggregate synaptic input up to 10 nS, as explained in details in the Methods section of Yu et al. (2013).

A new model of the olfactory bulb, representing the actual 3D layout of the mitral-granule cell network, has been recently developed (e.g., Migliore et al., 2015). This model represents in a very realistic way the possible interactions between glomeruli located within the dendritic field of mitral cells, and it would be especially useful to study natural odors, which exhibit a rather broad and dense input. However, it requires large computational resources. With the particular set of inputs we are considering in this paper, i.e., single monomolecular odors with rather sparse and segregated inputs, such a model would not give results qualitatively different from those obtained with the 1D model.

To represent the range of intensities with adequate sensitivity (i.e., including the weakest concentration without saturating the network at the highest concentration), we set the peak conductance sensitivity to give suprathreshold responses to levels 3 and 4. Then, we defined strong inputs as strengths of 3 or 4 and weak inputs as strengths of 0, 1, or 2 (**Figure 1A**). All odor inputs were presented over 4–10 Hz. To address how prior odor experience interferes with the subsequent sparse representation of new odors, a series of odor inputs was used to train the network in sequence. In one odor experience condition, the first odors were presented within the first 5 s, and after a 5 s rest, another odor input was presented for the next 5 s (**Figure 1C**). In other experience conditions (for instance, five odors experience), more odor series were presented similarly to the single odor experience condition: each odor input was presented for 5 s, with a 5 s resting state between each presentation. The last odor input was denoted as the new odor input, and all prior odor inputs were defined as experienced odors, implying that in the five odors experience condition, a total of six odor inputs were used. For the control condition, only the new odor inputs were given at the time when the new odors were given in the experience

FIGURE 1 | Mitral cell network responses in naïve and prior odor experiences conditions. (A) Two example odor input strengths to each mitral cell. k7-1, heptyl methyl ketone; 8OH, octanal. **(B)** A raster plot shows the mitral cell network response in the naïve condition. k7-1 was delivered at the 10th second after a resting state (each mitral cell fires spontaneously and randomly at a low frequency) of 5 s. **(C)** A raster plot shows the mitral cell network response to the new odor input k7-1 in the single odor input (8OH) experience condition. Red rectangles represent the time elapsed to reach a stable sparseness level for the network response in different conditions.

conditions (**Figure 1B**). Unless otherwise noted, all experienced odors during training were presented in order from low level to high level of input strength.

Sparseness Calculation
The method for the sparseness calculation of network response was identical to our previous work (Yu et al., 2014). Briefly, based on previous work (Vinje and Gallant, 2000; Franco et al., 2007), the sparseness of response to a given stimulus can be calculated as follows:

$$S = \left\{ 1 - \frac{\left[\sum_{i=1}^{N} \left(\frac{r_i}{N} \right) \right]^2}{\sum_{i=1}^{N} \frac{r_i^2}{N}} \right\} \Big/ \left(1 - \frac{1}{N} \right),$$

where S is the sparseness of the network in one period of odor input (from the beginning of one input to the beginning of the next odor input); r_i is the mean firing rate of mitral cell i in that period; N is the total number of mitral cell (500). A high sparseness value in our present work indicates only a few neurons with high firing rates.

Correlation between Mitral Cell Firing in a Network
To calculate the correlation between mitral cell firing in a network, we used a coherence measure based on the normalized cross-correlation of neuronal pairs in the network. The coherence between two mitral cell i and j was measured by their cross-correlation of spike trains at zero time lag within a time bin of τ. Precisely, supposing that a long time interval T (one period of odor input) was divided into small bins of τ, and that two spike

trains (value of 0 or 1) were given by $X(l)$, $Y(l)$, with $l = 1, 2, \ldots K$ (here $T/K = \tau$), respectively, a coherence for the pair (K_{ij}) was calculated as follows (Wang and Buzsaki, 1996; Yu et al., 2014):

$$K_{ij}(\tau) = \frac{\sum_{l=1}^{K} X(l) Y(l)}{\sqrt{\sum_{l=1}^{K} X(l)} \sqrt{\sum_{l=1}^{K} Y(l)}}.$$

And then, the correlation between mitral cells across the whole network K was defined by the average of $K_{i,j}(\tau)$ over all pairs of mitral cells in the network. That is

$$K = \frac{1}{N(N-1)} \sum_{i=1}^{N} \sum_{j=1, j \neq i}^{N} K_{ij}(\tau),$$

where N is the total number of the mitral cells in the network. And in our present work, τ was taken as 20 ms through the whole analysis.

Correlation between Mitral Cell Network Responses
To compare the similarity between mitral cell network response to odor inputs x and y during an odor input period, we defined and calculated it as the correlation coefficient (C_{xy}) as follows:

$$C_{xy} = \frac{1}{N} \sum_{i=1}^{N} Corrcoef \left\{ MC_i \left[x(t) \right], MC_i \left[y(t) \right] \right\},$$

where MC_i is the i'th mitral cell; $MC_i[x(t)]$ and $MC_i[y(t)]$ are the mitral cell network response in an odor input period to odor input $x(t)$ and $y(t)$ respectively; $Corrcoef$ is to calculate the classic correlation coefficient. To investigate how prior odor experience affects the network response to the news odor inputs, we calculated the average of C_{xy} between one new odor input in the experience conditions and the other tested new odor inputs in the naïve condition.

1/2 Time of Sparseness
To test the dynamic evolution of the sparseness of the mitral cell network response, sparseness values were calculated at series time points when the odor inputs were presented. This sparseness time series could be fitted by the classic logarithmic function as follows:

$$S = A_2 + (A_1 - A_2) \Big/ \left(1 + \left(\frac{x}{x_0} \right)^{S_{1/2}} \right),$$

where S is the sparseness, and $S_{1/2}$ is the time at which S reaches the half of the maximum S ($A1$).

Correlation of Input Strengths between Different Odors
In some simulations, we quantified the similarity of two odor inputs by calculating the Pearson correlation coefficient based on their strength values for 500 mitral cells (i.e., 500 values for

each odor). A higher correlation coefficient indicates that a pair of odors is more similar.

Data were presented as mean ± SEM. Statistical significance was assessed by paired Student's *t*-test or ANOVA analysis with Tukey's multiple comparison test, and $p < 0.05$ was considered significant. Data analyses were performed using Graphpad Prism software *Version 6.0* (San Diego, USA).

RESULTS

To systematically address how prior odor experience affects the representation of new odor inputs by the olfactory bulb network, we used a previously verified olfactory bulb network model that includes 500 mitral cells and 10,000 granule cells connected through dendrodendritic synapses (Yu et al., 2013, 2014). In this model, we simulated different odor inputs to mitral cells with varied strength intensities ranging from 0 to 4 based on previous experimental results (Mori et al., 2006; **Figure 1A**). As shown in **Figure 1B**, in the naïve condition (i.e., no prior odor input experience, without odor inputs during the first 10 s), a sparse spatial spiking representation of specific odor input (k7-1 in this example) emerged naturally within several seconds of the training period from the mitral-granule cell interactions, as verified by our previous work (Yu et al., 2013, 2014). In one training paradigm, after delivery of a prior odor input (8OH) for 5 s and a 5 s resting state (no odor input), a new odor input (K7-1or other, see below) induced a different response of mitral cell network compared with that observed in the naïve condition (**Figure 1C**, compare with the mitral cell network response during the period of 10th–15th second in **Figure 1B**). From the raster plot, we observed that the response of the mitral cell network reached a stable sparseness state much faster than the naïve condition (**Figures 1B,C**, note that the red rectangle denotes the course to reach stable sparseness in **Figure 1C** that is much narrower than in **Figure 1B**). Since the sparseness of the mitral cell network reaches steady state after about 2 s of odor stimulus, we trained the network with specific odor input for 5 s in the following results. We also extended the simulation time to 10 s, and no significantly different results were found (Supplementary Figure 1). We will now present additional details describing our results.

Prior Odor Input Experience Facilitates the Evolution of the Sparseness of the Mitral Cell Network Response

Experimental and computational studies have shown that the response of the mitral cell network to odor inputs tends to be heterogeneous and spatiotemporally sparse (Yu et al., 2013, 2014). Our previous reports have shown that a sparse spatial spiking representation of specific odor signals can emerge naturally within several seconds of the training period from mitral-granule cell interactions and that the network response reaches a stable level of sparseness (Yu et al., 2013, 2014). To address how prior odor experience affects the evolution of sparseness in the mitral cell network and the final sparseness level in response to new odor inputs, we fixed the prior odor inputs to

8OH or o-Eph and then varied the new odor inputs or trained the network only with the new odor inputs (**Figures 1B,C**). Same as in our previous reports, the sparseness of the mitral cell network response gradually evolved from a relatively low sparseness level to a high sparseness level (**Figures 1**, **2A**). We found that the sparseness of the mitral cell network response to new odor inputs in the single odor experience (8OH or o-Eph) condition was larger than that in the naïve condition at all sniff points the input were given (**Figure 2A**, $n = 14$ for the number of second odors). **Figure 2B** shows that the stable sparseness levels of the mitral cell network (represented by the last sniff point of 14.8 s) in both 8OH and o-Eph experience conditions are statistically larger than those in the naïve condition (one way ANOVA analysis, $p < 0.01$, **Figure 2B**). To demonstrate this phenomenon in a more systematic way, we trained the network with additional prior odor series in a manner similar to the single odor experience condition. As shown in **Figure 2C**, the stable sparseness level (represented by the last sniff point) of the mitral cell network increases with the number of prior odors experienced (one way ANOVA, $p < 0.01$). This scenario was more significant, as shown by the sparseness at the first sniff (**Figure 2C**). As shown in Supplementary Figure 2, the prior odors were delivered from low input strength level to high strength level in the 72 odor experience condition [72 Odors (lh)]. We also reversed the training sequence [i.e., prior odors were delivered from high input strength level to low strength level, 72 Odors (hl)]. Interestingly, the final sparseness of the mitral cell network is significantly lower in the 72 odor (hl) condition than that in the 72 odor (lh) condition (paired *t*-test, $p < 0.001$, Supplementary Figure 2). We plan to address this phenomenon extensively in future work, but the present work will mainly focus on the former training sequence (i.e., all prior odors were delivered from low input strength level to high strength level). Moreover, we fixed the new odor to 8OH and varied the experienced odor inputs. We found that the final sparseness level of the mitral cell network to 8OH was negatively correlated to the correlation coefficients of input strength of experienced odors and 8OH (**Figure 2D**). Similar results were found in the cases of k7-1 and o-Eph as the second odors (Supplementary Figure 3A).

It is worthwhile to note that, the new odors we used in our model were different from the experienced odors. We also tested the case that the second odor was the same as the first odor (experienced odor) in the single experienced odor condition and found no significant difference (Supplementary Figure 4).

A previous experimental study showed that prior odor experience could increase the tuning specificity of mitral cell to a variety of odors (Fletcher and Wilson, 2003). In our model, prior odor training could decrease the response of mitral cell to weak odor input leading to a slight increase of the tuning specificity of the mitral cell (Supplementary Figure 5A). To test whether the network sparseness change observed above was due to the increase of tuning specificity of mitral cells, we arbitrarily set the responses of mitral cell receiving no input from a given new odor the same as that in the naïve condition and left the rest responses of mitral cells (receiving at least one intensity from new odor) unchanged as in **Figure 2C**; we found that the stable sparseness levels hardly changed under different conditions (one
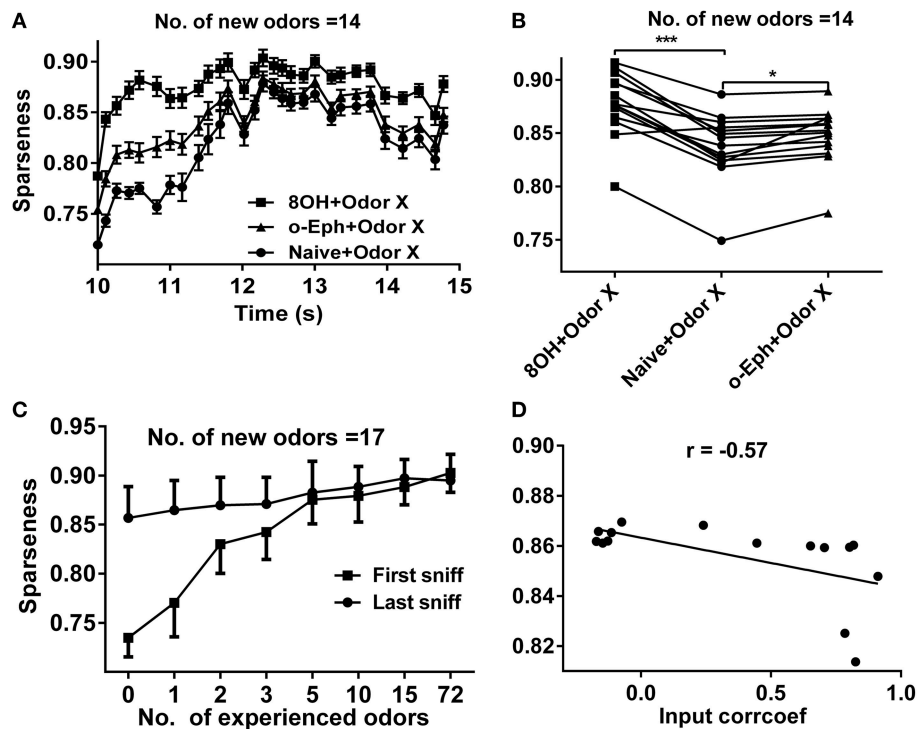
FIGURE 2 | Sparseness of the mitral cell network in different odor input conditions. (A) A time course plot shows the sparseness of the mitral cell network in response to 14 new odor inputs in the single odor input (8OH or o-Eph) experience or naive condition. **(B)** A dot aligned plot shows the sparseness of the mitral cell network response at the last sniff shown in (A). ***$p < 0.001$, *$p < 0.05$; one-way ANOVA with Tukey's *post-hoc* comparison test. **(C)** Sparseness measured in the first and last sniff period of the mitral cell network response to 17 new odor inputs in additional odor inputs experience conditions. The experienced odor inputs were delivered in sequence in ascending order of input strength. **(D)** The correlation between sparseness to 8OH under 14 different single odor input experience conditions vs. the correlation coefficients of input strength of these 14 experienced odor inputs and 8OH. The solid line represents the linear fitting curve.

way ANOVA, Supplementary Figure 5B). Such results suggest that the observed sparseness change of mitral cell network under prior odor experience condition is mainly due to the increase of the sparseness of mitral cells with no input from the new odors. And the experimentally observed tuning specificity of MCs after the odor exposure (e.g., Fletcher and Wilson, 2003) may have additional mechanisms that are beyond the present model simulation study.

To quantify how prior odor experience affects the evolution of the sparseness of the mitral cell network response to new odors, we fitted the time course of sparseness using a classical logarithmic function (**Figure 3A**). Then, based on the fitting curve, we determined the time at which the sparseness reaches half of the maximum value (denoted $S_{1/2}$). We found that $S_{1/2}$ of the network response to new odor inputs in both the 8OH and o-Eph experience conditions is less than in the naïve condition (one way ANOVA, $p < 0.001$, **Figure 3B**), implying that prior odor experience could accelerate the formation of sparse state sin the mitral cell network in response to new odor inputs. Moreover, to determine how the correlation of experience and new odor input strength affects the evolution of sparseness induced by new odor inputs, we fixed the new odor as 8OH and varied the experienced odor inputs, then measured the $S_{1/2}$ of the network response to 8OH. Interestingly, we found that $S_{1/2}$ was negatively

proportional to the correlation coefficient of the input strength of the experienced odors and 8OH ($r = 0.89$, **Figure 3C**). Similar results were found in the cases of k7-1 and o-Eph as the second odors (Supplementary Figure 3B). This may imply that the network response to new odor input requires less time to evolve to a stable sparseness state following experienced odor input more similar to the new odor.

In summary, prior odor experience could accelerate the evolution of sparseness in the mitral cell network response to new odor inputs and increases the sparseness level.

Prior Odor Input Experience Increases the Response of the Granule Cell Network

Previous experiments have shown that prior odor experience has a profound effect on the activity of the granule cell network in response to new odor inputs (Mandairon et al., 2008; Kato et al., 2012). We next tested the activity of the granule cell network in our model system. We also applied the sparseness measurement for mitral cells to quantify the activity in the granule cell network. Contrary to the sparseness in mitral cell network, the sparseness of the granule cell network decreased with the number of prior odors, implying that more prior odors leads to a larger increase of the response of the granule cell network to new odor inputs (**Figure 4**, one way ANOVA, $p < 0.01$).
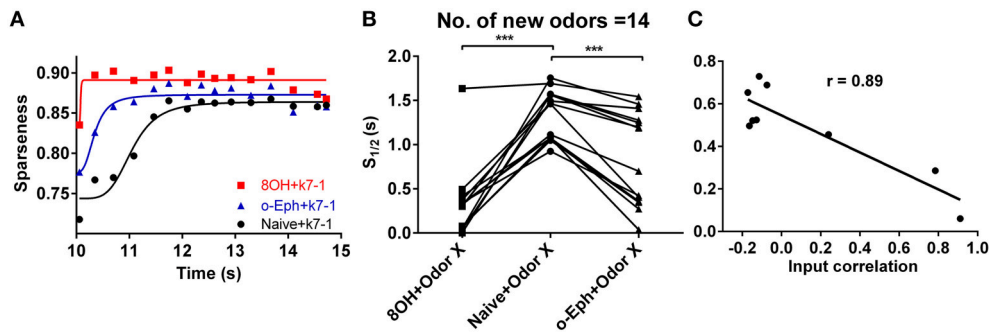
FIGURE 3 | Evolution of sparseness of mitral cell network responses in naïve and experienced odor input conditions. (A) A time course plot shows the sparseness of the mitral cell network response to odor input K7-1 in the single odor input (8OH or o-Eph) experience or naive condition. The solid line represents the fitting curve described by the classical logarithmic function (see Materials and Methods). **(B)** 1/2 time of sparseness ($S_{1/2}$) of the mitral cell network response to 17 new odor inputs in the single odor input (8OH or o-Eph) experience or naive conditions. $S_{1/2}$ is the time elapsed from the first presentation of new odor inputs to the time when the sparseness reaches half of the maximum value. ***$p < 0.001$, one-way ANOVA with Tukey's *post-hoc* comparison test. **(C)** $S_{1/2}$ of mitral cell network responses to 8OH in nine different single odor input experience conditions vs. the correlation coefficients of input strength of these nine experienced odor inputs and 8OH. The solid line represents the linear fitting curve.
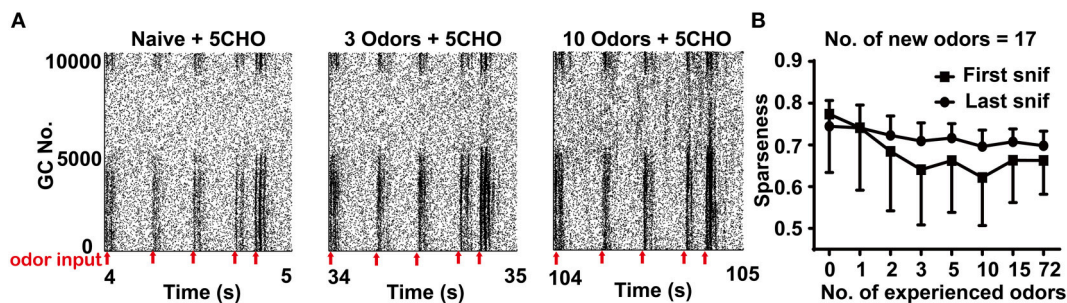
FIGURE 4 | Sparseness of the granule cell network in different odor input conditions. (A) A raster plot shows the firing of the granule cell network in response to 5CHO in the 3 (middle) or 10 (right) odors experience or naïve conditions (left) in the stable state. 5CHO: pentanal. Red arrows show the sniff points. **(B)** The sparseness measured after the first and last sniff of the granule cell network response to 17 new odor inputs in a series of prior odor inputs conditions. The series of experienced odor inputs were delivered in sequence according to ascending order of input strength.

Effects of Prior Odor Experience on Synaptic Weight in the Mitral Cell Network

Because synaptic plasticity exists in our model, the different stable sparseness of mitral or granule cell networks under different conditions may be due to the final synaptic weights in the bulb network. We tested the excitatory and inhibitory synaptic weight under different conditions in response to specific new odors. We divided the input strength into a strong group with strengths of 3 or 4 and a weak group with strengths of 0, 1, or 2. As shown in **Figure 5A**, the average excitatory synaptic weights for a mitral cell receiving weak or all inputs significantly increased with prior odor number, but decreased for mitral cells receiving strong inputs (two way ANOVA, $p < 0.01$). The same scenario applied to the average inhibitory synaptic weight (**Figure 5B**, two way ANOVA, $p < 0.01$).

Our previous studies reported that the change of sparseness of the mitral cell network resulted from the evolved dynamic changes in the synaptic weight of both excitatory and inhibitory dendrodendritic synapses (Yu et al., 2013, 2014). Such a developed dynamic change of synaptic weights has been

suggested to affect the changes in the time course of mitral cell network sparseness (Yu et al., 2013, 2014). Now we would like to examine how the synaptic weights will further evolve with the continuous training of prior odor experience, and then examine how this prior experience could modulate the response sparseness of the mitral network to new odor inputs. We analyzed the time course of average excitatory weight (G_{ex}) and inhibitory weight (G_{in}) for mitral cells receiving strong inputs and weak inputs. Similar to our previous results (Yu et al., 2013, 2014), the time courses of response sparseness are tightly correlated with the changes of synaptic weight (especially excitatory synaptic weight) during the response to 8OH, both for naïve (Supplementary Figure 6) and single odor input experience (k7-1, Supplementary Figure 6). In the naïve condition, strong excitatory synaptic inputs gradually increased from 0.06 nS to a steady state of ~0.47 nS after 2 s of 8OH input (Supplementary Figures 6A,B). However, in the k7-1 odor experience condition, the sparseness of the mitral cell network reached a maximum level immediately in response to new 8OH input, and the strong excitatory and inhibitory synaptic inputs also reached a maximum value at
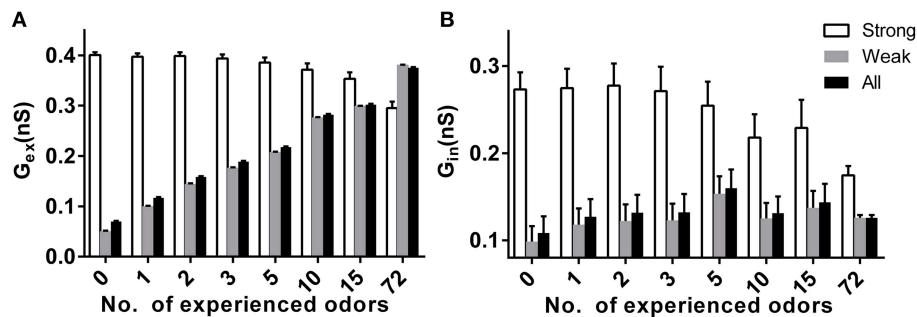
FIGURE 5 | Effects of prior odor experience on synapticweightin the mitral cell network. (A) The average excitatory synaptic weight for mitral cells receiving strong (strength 3 or 4), weak (strength 0, 1, or 3), or all inputs in a variety of prior odor experience or naïve conditions. **(B)** Same as in **(A)**, but for the inhibitory synaptic weight.

the beginning of the 8OH input period (Supplementary Figures 6A,B). Therefore, these results imply that prior odor experience accelerates the evolution of synaptic weight in the mitral cell network to the steady state, which in turn accelerates the evolution of sparseness in the mitral cell network in response to new odor inputs.

Prior Odor Input Experience Decreased the Correlation of the Mitral Cell Firing Pattern

Sparse coding is an efficient scheme by which an individual neuron independently encodes different properties of the input (Olshausen and Field, 1996; Vinje and Gallant, 2000). This naturally leads us to predict that when the mitral cell network reaches a high sparseness level, the correlation level among responses of mitral cells in the network should reach a low level. In fact, we have verified this prediction in our previous work (Yu et al., 2014). We therefore tried to determine whether prior odor experience would also affect the evolution of the decorrelated state among mitral cells in the network to new coming odor inputs. We quantified this correlation by averaging the correlation coefficients of all possible pairs of 500 mitral cell responses at each sniff point during new odor delivery. Similar to the evolution of sparseness in the mitral cell network, the correlation among mitral cells gradually evolved from a relatively high level to a low level (**Figure 6A**). We found that the correlation of mitral cell firing in the network in response to new odor inputs in one odor experience (8OH or o-Eph) condition was lower than that in the naïve condition at all the sniff points tested (**Figure 6A**). **Figure 6B** shows the stable correlation level of mitral cell firing in the network (represented by the last sniff point) for both 8OH and o-Eph experience conditions were statistically lower than that in the naïve condition (paired t-test, $p < 0.05$, **Figure 6B**). We then tested the correlation in conditions with more prior odor inputs. As shown in **Figure 6C**, the correlations of mitral cell firing in the network in response to new odor inputs at first and last sniff following prior odor experience both decreased as the number of prior odors increased (one way ANOVA, $p < 0.01$). Moreover, we fixed the new odor to 8OH and varied the experienced odor inputs; we found that the correlation of mitral cell firing in network in response to the

new odor 8OH at last sniff was weakly linearly correlated to the correlation coefficients of the input strength of experienced odors and 8OH (**Figure 6D**). Similar results were found in the cases of k7-1 and o-Eph as the second odors (Supplementary Figure 3C). This result implies that the mitral cell firing response tends to be more decorrelated if the input strength of new coming odor differs more greatly from that of the prior experienced odor.

Prior Odor Input Experience Decrease Correlation of Mitral Cell Network Response

We already tested the effects of prior odor experience on mitral cell firing pattern in response to the corresponding new odor inputs. A more direct way to measure the coding efficiency of the mitral cell network in response to different odor inputs is to calculate the similarity of the network response to different odor inputs, especially to similar odor inputs (Yu et al., 2014). We therefore investigated how prior odor experience affects the mitral cell network response to different new odor inputs. Odors 7OH and 6OH are two very similar odor inputs (**Figure 7A**). The response of the mitral cell network to 6OH in one odor input (8OH) experience condition was more different than that in the naïve condition to the network response to 7OH in naïve condition after the training process (**Figures 7B,C**). For instance, the firing rates of mitral cells 1–100 to 6OH in the 8OH experience condition (**Figure 7C**, right) were more different than that in naïve condition (**Figure 6B**, right) from that to 7OH in naïve condition (**Figure 7B**, left). A similar result was also found when we compared the response of the mitral cell network to 7OH in the 8OH experience condition or the naïve condition with that to 6OH in the naïve condition (**Figures 7B,C**).

To quantify the similarity of the mitral cell network response to different odor inputs in different conditions, we measured the network response correlation as described in the Materials and Methods section. The correlations of mitral cell network responses evolved gradually from relatively high to low level (**Figure 8A**). Furthermore, the correlations between mitral cell network responses to new coming odor inputs in the one odor experience (8OH or o-Eph) condition were lower than in the naïve condition at all sniff points tested (**Figure 8A**).
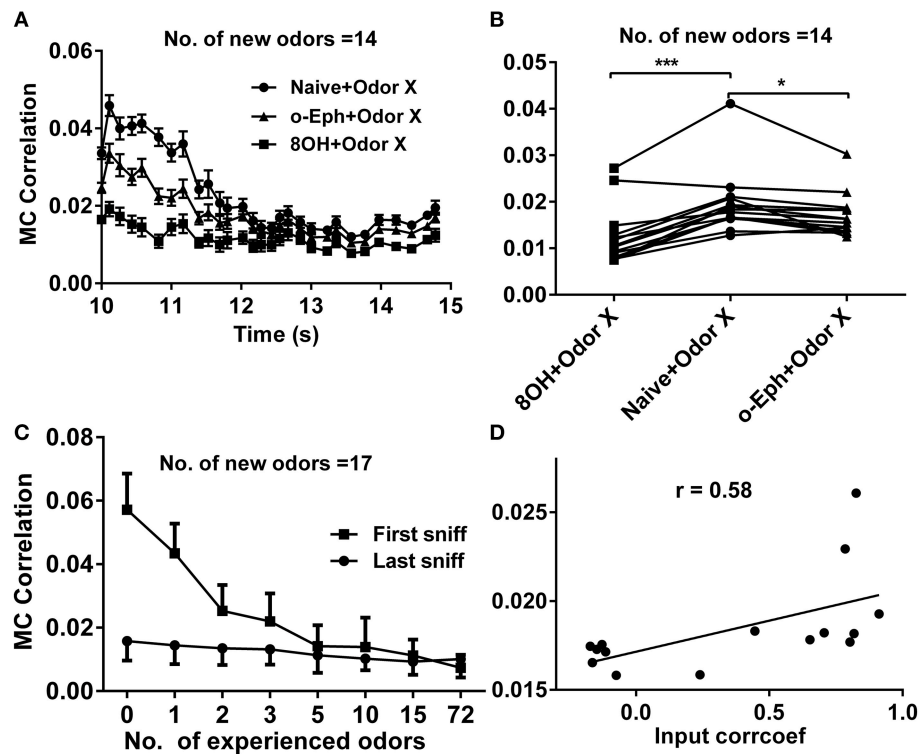
FIGURE 6 | Correlation between mitral cell firing patterns. (A) A time course of the correlations between mitral cell firing patterns in response to 14 new odor inputs in the single odor input (8OH or o-Eph) experience or naïve condition. **(B)** A dot aligned plot shows the correlation between mitral cell firing pattern in response to 14 new odor inputs at the last sniff shown in **(A)**.***$p < 0.001$, *$p < 0.05$; one-way ANOVA with Tukey's *post-hoc* comparison test. **(C)** The correlation between mitral cell firing patterns in response to 17 new odor inputs in networks measured after the first and last sniff with increased number of prior experienced odors. The odor inputs were delivered sequentially in order of ascending input strength. **(D)** The correlation between mitral cell firing patterns in response to 8OH under 14 different single odor input experience conditions vs. the correlation coefficients of input strength of these 14 experienced odor inputs and 8OH. The solid line represents the linear fitting curve.

Figure 8B shows the stable correlation level of the mitral cell network response (represented by the last sniff point) in both 8OH and o-Eph experience conditions were statistically lower than in the naïve condition (one way ANOVA, $p < 0.05$, **Figure 8B**). We then tested the correlation in conditions with more series of prior odor inputs. As shown in **Figure 8C**, the correlations of mitral cell network response to new coming odor inputs in prior odor experience at the first and last sniff point decreased while the number of prior odors increased (one way ANOVA, $p < 0.01$). Similar to **Figure 6D**, we fixed the new odor to 8OH, and varied the experienced odor inputs and we found that the correlation of mitral cell network response to new coming odor 8OH was weakly linearly correlated to correlation coefficients of the input strength of experienced odors and 8OH (**Figure 8D**). Similar results were also found in the cases of k7-1 and o-Eph as the second odors (Supplementary Figure 3D).

DISCUSSION

Using a scaled up mitral-granule cell network model and a set of odors with relatively strong input strength, we systematically investigated how prior odor input affects the coding paradigm of mitral cells to new incoming odor inputs. The following findings were observed: (1) when increasing the number of prior odors, the activity of the mitral cell network decreased and the granule cell network increased, gradually reaching an equilibrium level. (2) prior odor experience accelerated the formation of a stable sparseness level of the mitral cell network to new odors; (3) increasing prior odor experience also facilitated the mitral cell network to evolve to a more decorrelated state; (4) prior odor experience decreased the correlation of the mitral cell network response to new odors and this effect is more obvious after training the network with a larger number of prior odors. Note that all the changes gradually reach an equilibrium level that does not change with additional odor experience. All changes may be attributed to two key factors: (1) the continuous LTP effect for those mitral cells receiving a sustained strong input; (2) more and more mitral cells are activated when more odors are presented to the network, inducing more dynamic changes in the excitatory and inhibitory synaptic weights of the dendrodendritic synapses. An equilibrium is reached when most of the relevant mitral cells have been activated during the past odor experiences.
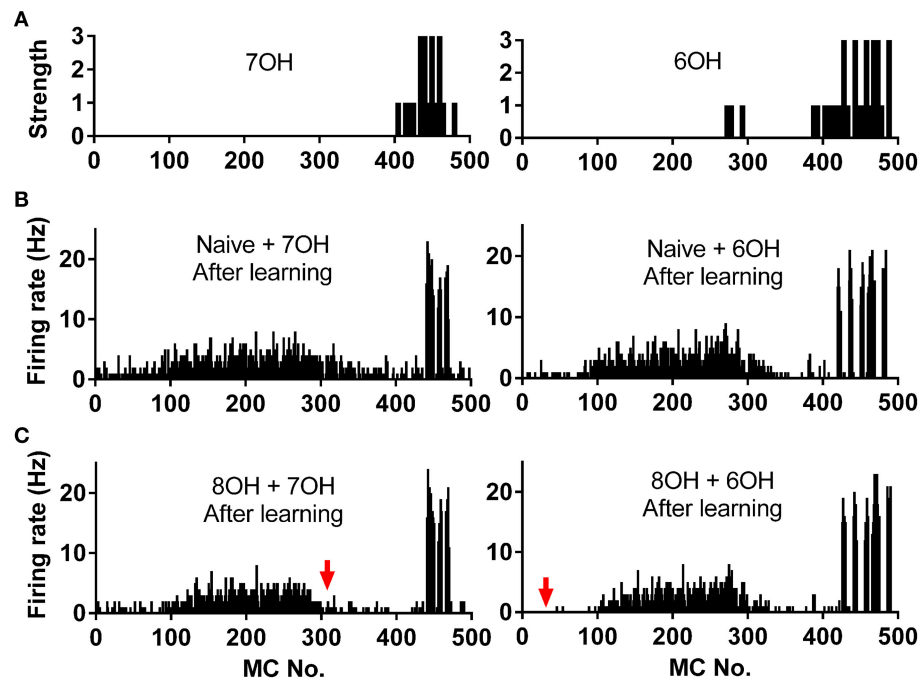
FIGURE 7 | Spatial firing rate patterns of 500 mitral cells. (A) The strength of odor inputs 7OH (left) and 6OH (right) to each mitral cell. **(B)** The spatial firing pattern of 500 mitral cells during the 4th second after training in response to 7OH (left) or 6OH (right) in the naïve condition. **(C)** The spatial firing pattern of 500 mitral cells during the 4th second after training in response to 7OH (left) or 6OH (right) in the 8OH experience condition. The red arrow (right) indicates that the firing pattern of the network in response to 7OH in the 8OH experience condition **(C**, left) differs more than that of the naïve condition **(B**, left) from the firing pattern of the network to 6OH in the naïve condition **(B**, right). Similarly, the red arrow (left) indicates that the firing pattern of the network in response to 6OH in the 8OH experience condition **(C**, right) differs more than that of the naïve condition **(B**, right) from the firing pattern of the network to 7OH in the naïve conditions **(B**, left).
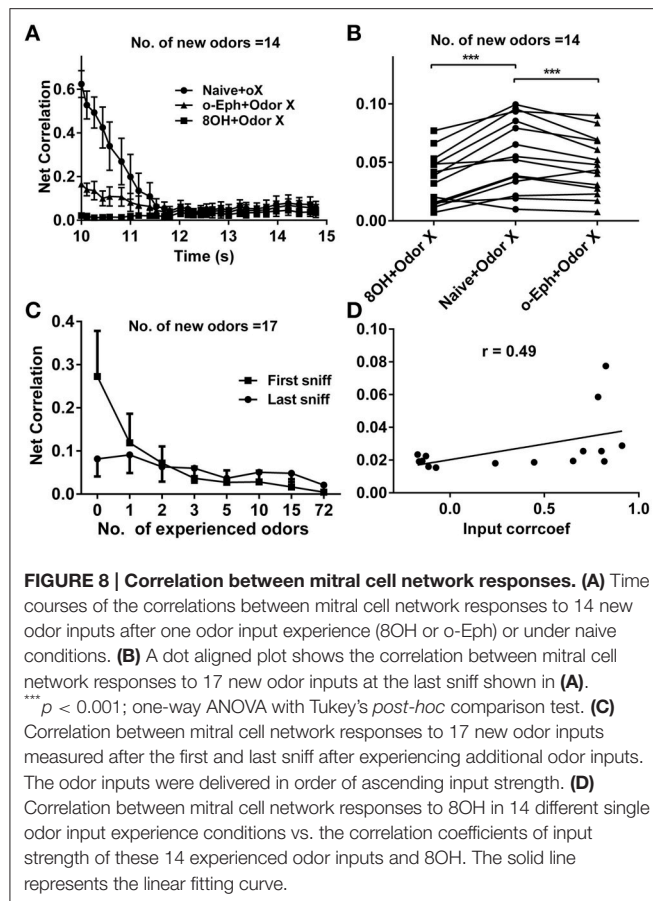
Sparseness of Mitral and Granule Cell in Prior Odor Experience Conditions

Sparse coding has been suggested as an efficient way to code the sensory inputs (Olshausen and Field, 1996; Rinberg et al., 2006; Davison and Katz, 2007; Koulakov and Rinberg, 2011). Previous studies have found that the response of mitral cell network to new coming odor inputs decreases after prior odor experience (Buonviso et al., 1998; Buonviso and Chaput, 2000; Fletcher and Wilson, 2003; Kato et al., 2012). We systematically tested several prior odors inputs in our large scale olfactory bulb model, and we found the sparseness of the mitral cell network to new coming odors increases with the number of prior odors (**Figure 2C**). Moreover, we found the sparseness of the granule cell network to new coming odors decrease with the number of prior odors (**Figure 4B**). In fact, previous experimental results also found that prior odor experience could increase the activity of interneuron (Mandairon et al., 2008). We further found the average of excitatory and inhibitory synaptic weight both increase along with the number of prior odor inputs (**Figure 5**). As the mitral cell is the main target of granule cell and granule cell is also the main target of mitral cell, we can infer that the increase of excitatory synaptic weight leads to the increase of granule cell activity, and the increase of inhibitory synaptic weight combined with the increase of granule cell activity leads to the decrease of the mitral cell activity.

Previous experimental study also showed that prior odor experience could decrease the granule cell activity (Kato et al., 2012). In our simulations, prior odor experience decreases the response of mitral cell network, which tends to decrease the activity of granule cells. By contrast, prior odor experience increases the average excitatory synaptic weight to granule cell, which tends to increase the activity of granule cells. These two contrary effects of prior odor experience on granule cell activity might be the cause for the varied results of granule cell activity change to prior odor experience observed in experiments.

It should be noted that an experimental study shows that mitral cell responses decreases more after the same prior odor exposure than with different odor experience (Kato et al., 2012). In our work, the response of mitral cells hardly changed when it was in the same prior odor conditions. This suggests the phenomenon observed by Kato et al. (2012) might be attributed to a different mechanism that is not involved in the current model network.

Previous experimental results have also shown that the first sniff after odor input is very important for odor discrimination behavior (Uchida and Mainen, 2003; Cury and Uchida, 2010). In addition to analyzing olfactory bulb responses at equilibrium, we also analyzed the response properties right after the first sniff in all conditions. We found the sparseness change (or correlation of mitral cell network response change) induced by

FIGURE 8 | Correlation between mitral cell network responses. (A) Time courses of the correlations between mitral cell network responses to 14 new odor inputs after one odor input experience (8OH or o-Eph) or under naive conditions. (B) A dot aligned plot shows the correlation between mitral cell network responses to 17 new odor inputs at the last sniff shown in (A). *** $p < 0.001$; one-way ANOVA with Tukey's *post-hoc* comparison test. (C) Correlation between mitral cell network responses to 17 new odor inputs measured after the first and last sniff after experiencing additional odor inputs. The odor inputs were delivered in order of ascending input strength. (D) Correlation between mitral cell network responses to 8OH in 14 different single odor input experience conditions vs. the correlation coefficients of input strength of these 14 experienced odor inputs and 8OH. The solid line represents the linear fitting curve.

prior odor experience is more significant for first sniff cycle than that for the last sniff cycle. The more experienced odors result in less response difference in the sparseness level (or response correlation) between the first and last sniff cycles. This may suggest that the response in the first sniff may contain important information for odor discrimination that can be enhanced by the experienced odors.

It is worthwhile to note that the increasing rate of sparseness of mitral cell network tend to decrease when the number of prior odors increases (**Figure 2C**). For instance, the increase of sparseness between three and five odors conditions (increase of only two odors) is larger than that between 15 and 72 odors conditions (increase of 57 odors). We may infer that the sparseness level of the mitral network would saturate after a certain number of prior odor experience. Another interesting phenomenon needed to further test is that the different training sequence of series of prior odors would have significantly different effect on the response of mitral cell network to new coming odors.

Sparseness Evolution in Mitral Cell Network

Previous studies have shown a sparse spiking representation of specific odor can emerge naturally after several seconds of a learning period (with certain odor input frequency)

from the mitral-granule cell synaptic connections (Yu et al., 2013, 2014). And this phenomenon may be corresponding to the learning process of animal to a new odor inputs. In the one odor experience condition, we found the prior odor experience accelerates the process to reach the stable sparse state (**Figure 3**). Furthermore, we found that such acceleration of the sparseness to reach stable sparse state was well correlated to the acceleration of the excitatory and inhibitory synaptic weight to reach the maximum value (Supplementary Figure 6). And further experiments are needed to confirm such findings.

We also found that the more similar the new odor was to the experienced odor, the faster the mitral cell network reached a stable sparseness level, which may be due to less time needed to train more overlapping synaptic interactions to reach steady state. On the other hand, more disparate prior experienced odors lead to a higher stable sparseness level of mitral cell to the new odor, which may be due to the increase of the overall inhibitory synaptic weight resulting from the activation of more mitral and granule cells by more experienced odors. We may infer that the rate of formation of stable sparseness and the sparseness level itself are two different aspects of the odor representation of the mitral cell network.

Decorrelation of the Mitral Cell Network Response

Our previous study has shown that the response of the mitral cell network tends to be decorrelated and accompanied by sparseness (Yu et al., 2014). Our current work shows correlations within the mitral cell network to new odors decrease with the number of prior odors (**Figure 6C**). A similar correlation exists for mitral cell network responses to different odor inputs—a direct way to quantify coding efficiency under different conditions (**Figure 8C**). This may partially give an explanation to why enrichment could increase the ability of an animal to discriminate different odors (Mandairon et al., 2006a,b,c; Sinding et al., 2011).

Previous experimental and computational results have extensively shown the importance of granule cell activity and inhibitory synaptic weight for representation of odor inputs in olfactory bulb network (Mandairon et al., 2006a, 2008; Koulakov and Rinberg, 2011; Kato et al., 2012). We infer that such a decorrelated state between mitral cell firing in a specific network and the network response to different odor inputs is due to an increase in granule cell activity and inhibitory synaptic weight after odor experience.

In summary, using a scaled up olfactory bulb model, we systematically investigated how prior odor experience affects the sparse representation of new odor inputs by the olfactory bulb network. In conclusion, the gradual increased inhibitory weight of granule cells together with the slightly increased firing rates of gradual cell populations promote the response sparseness and decorrelated state of mitral populations to new odor inputs. These results may help to better explain how prior sensory experience affects the behavior of animals in response to new odor inputs.

AUTHOR CONTRIBUTIONS

ACKNOWLEDGMENTS

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: http://journal.frontiersin.org/article/10.3389/fnana.2016.00010

REFERENCES

Arenkiel, B. R., Hasegawa, H., Yi, J. J., Larsen, R. S., Wallace, M. L., Philpot, B. D., et al. (2011). Activity-induced remodeling of olfactory bulb microcircuits revealed by monosynaptic tracing. *PLoS ONE* 6:e29423. doi: 10.1371/journal.pone.0029423

Bischofberger, J., and Jonas, P. (1997). Action potential propagation into the presynaptic dendrites of rat mitral cells. *J. Physiol.* 504, 359–365. doi: 10.1111/j.1469-7793.1997.359be.x

Buonviso, N., and Chaput, M. (2000). Olfactory experience decreases responsiveness of the olfactory bulb in the adult rat. *Neuroscience* 95, 325–332. doi: 10.1016/S0306-4522(99)00450-9

Buonviso, N., Gervais, R., Chalansonnet, M., and Chaput, M. (1998). Short-lasting exposure to one odour decreases general reactivity in the olfactory bulb of adult rats. *Eur. J. Neurosci.* 10, 2472–2475. doi: 10.1046/j.1460-9568.1998.00266.x

Cury, K. M., and Uchida, N. (2010). Robust odor coding via inhalation-coupled transient activity in the mammalian olfactory bulb. *Neuron* 68, 570–585. doi: 10.1016/j.neuron.2010.09.040

Davison, I. G., and Katz, L. C. (2007). Sparse and selective odor coding by mitral/tufted neurons in the main olfactory bulb. *J. Neurosci.* 27, 2091–2101. doi: 10.1523/JNEUROSCI.3779-06.2007

Ennis, M., Linster, C., Aroniadou-Anderjaska, V., Ciombor, K., and Shipley, M. T. (1998). Glutamate and synaptic plasticity at mammalian primary olfactory synapses. *Ann. N.Y. Acad. Sci.* 855, 457–466.

Fletcher, M. L., and Wilson, D. A. (2003). Olfactory bulb mitral-tufted cell plasticity: odorant-specific tuning reflects previous odorant exposure. *J. Neurosci.* 23, 6946–6955.

Franco, L., Rolls, E. T., Aggelopoulos, N. C., and Jerez, J. M. (2007). Neuronal selectivity, population sparseness, and ergodicity in the inferior temporal visual cortex. *Biol. Cybern.* 96, 547–560. doi: 10.1007/s00422-007-0149-1

Gao, Y., and Strowbridge, B. W. (2009). Long-term plasticity of excitatory inputs to granule cells in the rat olfactory bulb. *Nat. Neurosci.* 12, 731–733. doi: 10.1038/nn.2319

Haykin, S. (1994). *Neural Networks: A Comprehensive Foundation.* New York, NY: Macmillan Publishing.

Hines, M. L., and Carnevale, N. T. (1997). The NEURON simulation environment. *Neural Comput.* 9, 1179–1209. doi: 10.1162/neco.1997.9.6.1179

Hines, M. L., and Carnevale, N. T. (2001). NEURON: a tool for neuroscientists. *Neuroscientist* 7, 123–135. doi: 10.1177/107385840100700207

Kato, H. K., Chu, M. W., Isaacson, J. S., and Komiyama, T. (2012). Dynamic sensory representations in the olfactory bulb: modulation by wakefulness and experience. *Neuron* 76, 962–975. doi: 10.1016/j.neuron.2012.09.037

Kay, L. M., and Sherman, S. M. (2007). An argument for an olfactory thalamus. *Trends Neurosci.* 30, 47–53. doi: 10.1016/j.tins.2006.11.007

Khan, A. G., Parthasarathy, K., and Bhalla, U. S. (2010). Odor representations in the mammalian olfactory bulb. *Wiley Interdiscip. Rev. Syst. Biol. Med.* 2, 603–611. doi: 10.1002/wsbm.85

Koulakov, A. A., and Rinberg, D. (2011). Sparse incomplete representations: a potential role of olfactory granule cells. *Neuron* 72, 124–136. doi: 10.1016/j.neuron.2011.07.031

Ma, T. F., Zhao, X. L., Cai, L., Zhang, N., Ren, S. Q., Ji, F., et al. (2012). Regulation of spike timing-dependent plasticity of olfactory inputs in mitral cells in the rat olfactory bulb. *PLoS ONE* 7:e35001. doi: 10.1371/journal.pone.0035001

Mandairon, N., Didier, A., and Linster, C. (2008). Odor enrichment increases interneurons responsiveness in spatially defined regions of the olfactory bulb correlated with perception. *Neurobiol. Learn. Mem.* 90, 178–184. doi: 10.1016/j.nlm.2008.02.008

Mandairon, N., and Linster, C. (2009). Odor perception and olfactory bulb plasticity in adult mammals. *J. Neurophysiol.* 101, 2204–2209. doi: 10.1152/jn.00076.2009

Mandairon, N., Stack, C., Kiselycznyk, C., and Linster, C. (2006a). Broad activation of the olfactory bulb produces long-lasting changes in odor perception. *Proc. Natl. Acad. Sci. U.S.A.* 103, 13543–13548. doi: 10.1073/pnas.0602750103

Mandairon, N., Stack, C., Kiselycznyk, C., and Linster, C. (2006b). Enrichment to odors improves olfactory discrimination in adult rats. *Behav. Neurosci.* 120, 173–179. doi: 10.1037/0735-7044.120.1.173

Mandairon, N., Stack, C., and Linster, C. (2006c). Olfactory enrichment improves the recognition of individual components in mixtures. *Physiol. Behav.* 89, 379–384. doi: 10.1016/j.physbeh.2006.07.013

Migliore, M., Cavarretta, F., Hines, M. L., and Shepherd, G. M. (2014). Distributed organization of a brain microcircuit analyzed by three-dimensional modeling: the olfactory bulb. *Front. Comput. Neurosci.* 8:50. doi: 10.3389/fncom.2014.00050

Migliore, M., Cavarretta, F., Marasco, A., Tulumello, E., Hines, M. L., and Shepherd, G. M. (2015). Synaptic clusters function as odor operators in the olfactory bulb. *Proc. Natl. Acad. Sci. U.S.A.* 112, 8499–8504. doi: 10.1073/pnas.1502513112

Migliore, M., Hines, M. L., Mctavish, T. S., and Shepherd, G. M. (2010). Functional roles of distributed synaptic clusters in the mitral-granule cell network of the olfactory bulb. *Front. Integr. Neurosci.* 4:122. doi: 10.3389/fnint.2010.00122

Migliore, M., Hines, M. L., and Shepherd, G. M. (2005). The role of distal dendritic gap junctions in synchronization of mitral cell axonal output. *J. Comput. Neurosci.* 18, 151–161. doi: 10.1007/s10827-005-6556-1

Migliore, M., Hoffman, D. A., Magee, J. C., and Johnston, D. (1999). Role of an A-type K+ conductance in the back-propagation of action potentials in the dendrites of hippocampal pyramidal neurons. *J. Comput. Neurosci.* 7, 5–15. doi: 10.1023/A:1008906225285

Migliore, M., Inzirillo, C., and Shepherd, G. M. (2007). Learning mechanism for column formation in the olfactory bulb. *Front. Integr. Neurosci.* 1:12. doi: 10.3389/neuro.07.012.2007

Migliore, M., and Shepherd, G. M. (2008). Dendritic action potentials connect distributed dendrodendritic microcircuits. *J. Comput. Neurosci.* 24, 207–221. doi: 10.1007/s10827-007-0051-9

Mori, K., Nowycky, M. C., and Shepherd, G. M. (1981). Analysis of a long-duration inhibitory potential in mitral cells in the isolated turtle olfactory bulb. *J. Physiol.* 314, 311–320. doi: 10.1113/jphysiol.1981.sp013709

Mori, K., Takahashi, Y. K., Igarashi, K. M., and Yamaguchi, M. (2006). Maps of odorant molecular features in the Mammalian olfactory bulb. *Physiol. Rev.* 86, 409–433. doi: 10.1152/physrev.00021.2005

Olshausen, B. A., and Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381, 607–609. doi: 10.1038/381607a0

Patneau, D. K., and Stripling, J. S. (1992). Functional correlates of selective long-term potentiation in the olfactory cortex and olfactory-bulb. *Brain Res.* 585, 219–228. doi: 10.1016/0006-8993(92)91210-6

Pinato, G., and Midtgaard, J. (2005). Dendritic sodium spikelets and low-threshold calcium spikes in turtle olfactory bulb granule cells. *J. Neurophysiol.* 93, 1285–1294. doi: 10.1152/jn.00807.2004

Rinberg, D., Koulakov, A., and Gelperin, A. (2006). Sparse odor coding in awake behaving mice. *J. Neurosci.* 26, 8857–8865. doi: 10.1523/JNEUROSCI.0884-06.2006

Schoppa, N. E., and Westbrook, G. L. (1999). Regulation of synaptic timing in the olfactory bulb by an A-type potassium current. *Nat. Neurosci.* 2, 1106–1113. doi: 10.1038/16033

Shusterman, R., Smear, M. C., Koulakov, A. A., and Rinberg, D. (2011). Precise olfactory responses tile the sniff cycle. *Nat. Neurosci.* 14, 1039–1044. doi: 10.1038/nn.2877

Sinding, C., Thomas-Danguin, T., Crepeaux, G., Schaal, B., and Coureaud, G. (2011). Experience influences elemental and configural perception of certain binary odour mixtures in newborn rabbits. *J. Exp. Biol.* 214, 4171–4178. doi: 10.1242/jeb.063610

Stanton, P. K. (1996). LTD, LTP, and the sliding threshold for long-term synaptic plasticity. *Hippocampus* 6, 35–42.

Uchida, N., and Mainen, Z. F. (2003). Speed and accuracy of olfactory discrimination in the rat. *Nat. Neurosci.* 6, 1224–1229. doi: 10.1038/nn1142

Vinje, W. E., and Gallant, J. L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287, 1273–1276. doi: 10.1126/science.287.5456.1273

Wang, X. J., and Buzsaki, G. (1996). Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *J. Neurosci.* 16, 6402–6413.

Wang, X. Y., Mckenzie, J. S., and Kemm, R. E. (1996). Whole-cell K+ currents in identified olfactory bulb output neurones of rats. *J. Physiol.* 490(Pt 1), 63–77. doi: 10.1113/jphysiol.1996.sp021127

Willhite, D. C., Nguyen, K. T., Masurkar, A. V., Greer, C. A., Shepherd, G. M., and Chen, W. R. (2006). Viral tracing identifies distributed columnar organization in the olfactory bulb. *Proc. Natl. Acad. Sci. U.S.A.* 103, 12592–12597. doi: 10.1073/pnas.0602032103

Xiong, W. H., and Chen, W. R. (2002). Dynamic gating of spike propagation in the mitral cell lateral dendrites. *Neuron* 34, 115–126. doi: 10.1016/S0896-6273(02)00628-1

Yu, Y. G., McTavish, T. S., Hines, M. L., Shepherd, G. M., Valenti, C., and Migliore, M. (2013). Sparse distributed representation of odors in a large-scale olfactory bulb circuit. *PLoS Comput. Biol.* 9:e1003014. doi: 10.1371/journal.pcbi.1003014

Yu, Y. G., Migliore, M., Hines, M. L., and Shepherd, G. M. (2014). Sparse coding and lateral inhibition arising from balanced and unbalanced dendrodendritic excitation and inhibition. *J. Neurosci.* 34, 13701–13713. doi: 10.1523/JNEUROSCI.1834-14.2014

Zelles, T., Boyd, J. D., Hardy, A. B., and Delaney, K. R. (2006). Branch-specific Ca2+ influx from Na+-dependent dendritic spikes in olfactory granule cells. *J. Neurosci.* 26, 30–40. doi: 10.1523/JNEUROSCI.1419-05.2006

frontiers
in Neuroanatomy

Closed-Loop Brain Model of Neocortical Information-Based Exchange

James Kozloski *

IBM Research Division, Computational Biology Center, IBM T.J. Watson Research Center, Yorktown Heights, NY, USA

Here we describe an "information-based exchange" model of brain function that ascribes to neocortex, basal ganglia, and thalamus distinct network functions. The model allows us to analyze whole brain system set point measures, such as the rate and heterogeneity of transitions in striatum and neocortex, in the context of neuromodulation and other perturbations. Our closed-loop model is grounded in neuroanatomical observations, proposing a novel "Grand Loop" through neocortex, and invokes different forms of plasticity at specific tissue interfaces and their principle cell synapses to achieve these transitions. By implementing a system for maximum information-based exchange of action potentials between modeled neocortical areas, we observe changes to these measures in simulation. We hypothesize that similar dynamic set points and modulations exist in the brain's resting state activity, and that different modifications to information-based exchange may shift the risk profile of different component tissues, resulting in different neurodegenerative diseases. This model is targeted for further development using IBM's Neural Tissue Simulator, which allows scalable elaboration of networks, tissues, and their neural and synaptic components toward ever greater complexity and biological realism.

Keywords: neocortex, thalamus, basal ganglia, information-based exchange, brain model

1. INTRODUCTION

Synaptic plasticity regulates neuronal responses to patterns of inputs impinging on dendritic arbors from multiple presynaptic sources. Resulting input selectivity at the single neuron level is often associated with learning and memory in models of cognition. At the circuit level, synaptic plasticity can serve more complex functions over arbitrary inputs, from selecting fixed points in recurrent networks (Hopfield, 1982), to implementing optimizations such as information maximization in artificial neural networks (Linsker, 1997), to dynamically encoding inputs in winnerless networks (Rabinovich et al., 2001). A challenge to analyzing the role of any neuron or circuit that implements these functions for cognition is that of modeling appropriate, naturalistic neuronal and circuit inputs, which in real brains derive from tens of thousands to millions of other neurons.

Here we present a closed-loop brain model, including component models of several neural tissues that we hypothesize implement some of these functions. Synapses and plasticity connecting components at principle cell interfaces together create a set of closed neuroanatomical loops. Without extrinsic inputs or stochastic intrinsic drivers, our model avoids the challenges and

assumptions of modeling naturalistic inputs separately, and instead derives them exclusively from the dynamics of upstream neurons and tissues. The challenge then is model validation, which we won't address in this report. Instead the aim here is to delineate hypotheses and a theory of brain resting state function using the model and its simulations. We propose that models implemented similarly constitute a class of "brain models," and are distinct from component "neural tissue models," which instead assume an arbitrary set of inputs or stochastic processes to drive intrinsic tissue dynamics. By avoiding these assumptions, a coarse but consistent model of global brain function may be useful for better constraining the most detailed neural tissue simulations.

We introduce the term "traversal" to refer to a "synfire chain" as defined by Abeles (1991), but with additional neuroanatomical constraints defining a minimum set of neocortical regions traversed by the event. The cortico-cortical feedback loop in our model acts as a substrate for combined traversals of sensory, limbic, and motor areas, which we propose together drive behavior in the organism. The cortico-thalamo-cortical feed forward loop acts to maximize the entropy of these global traversals and to maximize information about the environment relayed as inputs to the loop. Lastly, the striato-nigro-striatal loop provides a means to select subsequent configurations by monitoring changes in ongoing traversals and signaling them with dopamine to alter routing within the feed forward entropy maximizing network. We propose this function as the substrate for reward learning in the organism. Each of these loops therefore has both a closed-loop function (global traversal, traversal entropy maximization, and traversal change monitoring and rerouting) and an organismal input-output function (behavior generation, sensory processing, and behavior selection based on reward learning).

The objective of this report is to describe the closed-loop model and simulations of it. Perturbations to the closed loop that alter dynamic set points will also be described. We hypothesize about dynamic disease mechanisms and progression based on the model, and describe methods that use neural tissue simulation (Kozloski and Wagner, 2011) for modeling treatments that alter brain disease risks. To summarize our overall approach and long-term research goals, the driving hypotheses relating our closed-loop model to brain disorder and disease states are: (1) The primary disease and disorder risk is a disturbance in plasticity that critically maintains brain system dynamic set points; (2) Compensatory circuit dynamics achieves near-normal set points despite genetic or environmental perturbations, but with increased secondary risk of neuronal dysfunction, damage, or loss; (3) Secondary risk correlates with feed forward destruction or dysfunction of neural tissues because with each neuron function lost, maintaining system set points requires even greater secondary risks; and (4) Slowing progression may therefore lie in mitigating the primary risk's effect on system set points or in limiting secondary risks incurred by inherent compensatory dynamics.

## 2. INFORMATION-BASED EXCHANGE BRAIN MODEL

2.1. Cytoarchitectonics of Bidirectional Neocortical Projections: The "Grand Loop"

We propose a model that emphasizes a specific cortico-cortical connectivity across the major sensory, limbic, and motor categories of Brodmann areas. This emphasis derives from several observations. First, we note the importance of signals traversing all three categories of cortical representations in order to produce a stable basis for perception and behavior by integrating information about the environment, internal needs, and behavioral opportunities of the organism. While many loops have been discovered in studies of the neocortical connectome, none provide the directed graph (feed forward vs. feedback) needed to identify a system to support such traversals. Instead, we note that the cytoarchitectonic granularity of neocortical areas provides one means to interpret feed forward (i.e., more granular to less granular) and feedback (i.e., less granular to more granular) connections between cortical areas (Rempel-Clower and Barbas, 2000) and therefore a means to identify a backbone for global brain traversals (**Figure 1A**).

Granularity refers to the density of punctate Nissl bodies in stained layer 4 of neocortex. The granularity across all of neocortex was studied and mapped extensively by von Economo (1929), and we reproduce his illustrations and some key findings in **Figures 1A,B**. Note that granular cortices typically have smaller diffuse Nissl bodies in layer 5, and agranular cortices have very large diffuse layer 5 Nissl bodies. Tiling in von Economo's map shows that regions of cortex with similar granularity are adjacent, with key exceptions at the boundaries between primary motor (M1) and primary somatosensory (S1), hippocampus (HC) and retrosplenial granular areas (RGA), and subgenual anterior cingulate (ACC) and prefrontal (PFC) cortices. Each of these three pairs of Brodmann areas are interconnected, and in our model represent key boundaries in the backbone for traversing the sensory-limbic (HC-RGA), limbic-motor (ACC-PFC), and motor-sensory (M1-S1) cortices (**Figure 1B**, black arrows). To complete a "Grand Loop" backbone, we join each pair of areas by an area in their adjacent dysgranular neocortical regions: the secondary somatosensory (S2), posterior cingulate (PCC), and supplemental motor (SMA) areas (**Figure 1C**). While others have noted that organizing principles for intrinsic microcircuits may be derived from combining von Economo's observations with those regarding granularity and the direction of cortico-cortical projections (Beul and Hilgetag, 2015), none to our knowledge have proposed a Grand Loop that traverses all of neocortex according to these principles.

2.2. Cortico-Cortical and Cortico-Thalamo-Cortical Functional Pathways

Having defined the feed forward neocortical Grand Loop, we'll now embellish this structural model with additional components based on observations regarding feed forward projections and
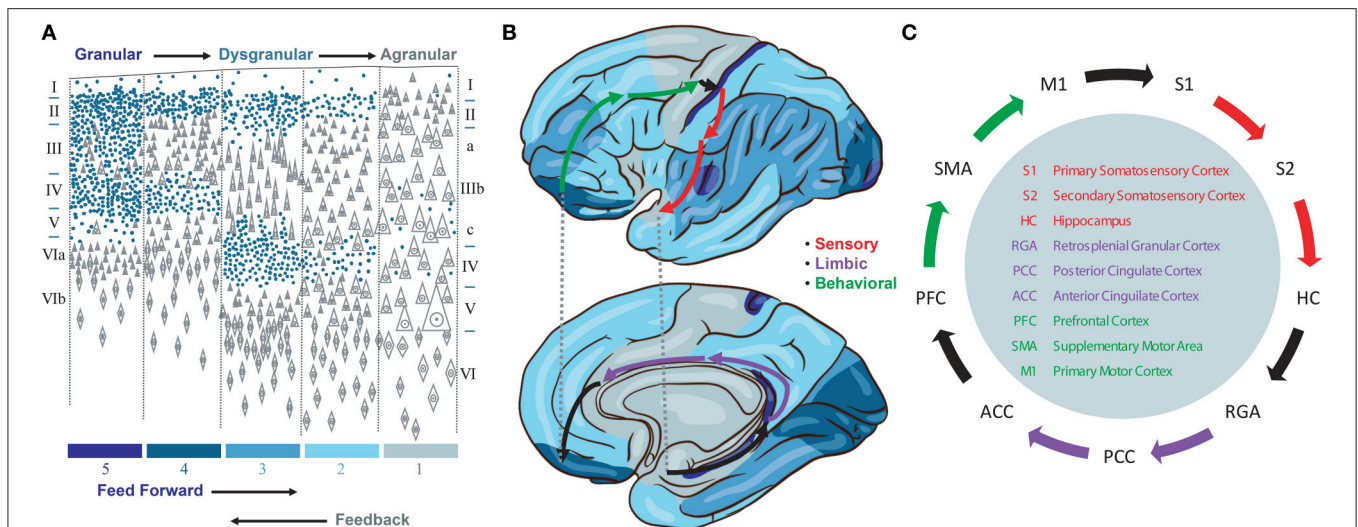
FIGURE 1 | (A) Granularity of different neocortical areas, adapted from von Economo (1929). Colors at bottom correspond to the map in **(B)**. **(B)** von Economo's neocortical tiling based on the granularity of large regions of neocortex spanning multiple Brodmann areas. The location of three Brodmann areas per stage are waypoints along a feed forward Grand Loop (arrows). **(C)** These Brodmann areas are connected based on projection data. Evidence that feed forward connections progress from granular to agranular areas provides directionality. The reciprocal feedback loop is not shown.
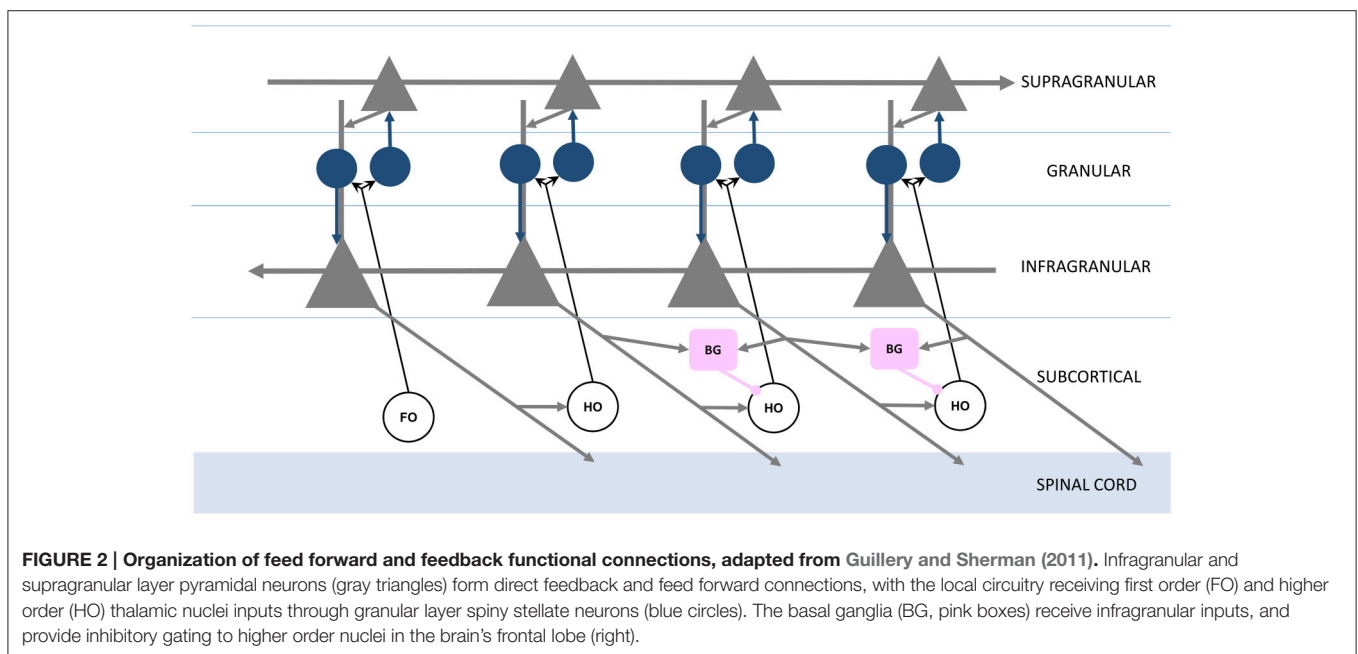
FIGURE 2 | Organization of feed forward and feedback functional connections, adapted from Guillery and Sherman (2011). Infragranular and supragranular layer pyramidal neurons (gray triangles) form direct feedback and feed forward connections, with the local circuitry receiving first order (FO) and higher order (HO) thalamic nuclei inputs through granular layer spiny stellate neurons (blue circles). The basal ganglia (BG, pink boxes) receive infragranular inputs, and provide inhibitory gating to higher order nuclei in the brain's frontal lobe (right).

signaling between neocortical areas. Sherman and Guillery emphasized different roles for direct cortico-cortical feed forward projections, which join one cortical area to another through their supragranular layers, and indirect cortico-thalamo-cortical projections, which join infragranular layers of the same original area to the granular layer of the same target area (**Figure 2**; Guillery and Sherman, 2011). In Sherman and Guillery's model, direct cortico-cortical projections are "modulatory," providing restricted activation to the target area, and indirect cortico-thalamo-cortical projections are "driving," providing activation

across all layers of the target area. **Figure 2** represents Sherman and Guillery's model (based on a simplification of their schematic). We will now describe how this model may be integrated into the Grand Loop.

Recall that each station of the loop in **Figure 1** is coupled in the feed forward direction. These connections, largely through the supragranular layers, are mirrored in the feedback direction by connections through infragranular layers (**Figure 2**; Rempel-Clower and Barbas, 2000). Thus, the Grand Loop represents two reciprocal loops, one in the feed forward direction and one in

the feedback direction. Furthermore, according to Sherman and Guillery, higher order thalamic nuclei provide at every stage a redundant relay for driving inputs over each feed forward connection. The local cortical circuit then receives signals from these thalamic nuclei and mixes the otherwise independent direct feed forward modulation loop and feedback traversal loop, primarily at layer 4's synaptic connections onto supragranular layers, and at supranular layers' onto the infragranular layers' apical dendrites.

We proposed previously that layers 2/3 of neocortex implement a network for maximizing mutual information between thalamic inputs and cortical responses (Kozloski et al., 2007). Entropy maximization in these layers (equivalent to information maximization when noise in the inputs is assumed to be negligible) would require a dense lateral network (Linsker, 1997), which fits well with the high proportion (∼22%) of total cortical synapses dedicated to intralaminar 2/3 connections (Binzegger et al., 2004). Given this role for the supragranular layers, we now propose that the role of cortico-thalamo-cortical driving inputs in Sherman and Guillery's model is to provide inputs both from first order thalamic nuclei about the environment and from feedback traversals through higher order thalamic nuclei about the behavioral state of the organism. A global supragranular network then extracts maximally informative features from combinations of these inputs. In addition, we propose that these features become conditional modulators on feedback traversals by boosting or reducing the gain on proximal inputs to layer 5 neurons by means of synaptic inputs onto their apical dendrites from layer 2/3 neurons.

2.3. Basal Ganglia Gating of Feed Forward Functional Pathways

In Sherman and Guillery's model, thalamic relay neurons in both first order and second order nuclei are subject to modulation. Modulation may derive from direct cortico-thalamic feedback from layer 6, inhibition from the thalamic reticular nucleus, or from neuromodulatory inputs such as norepinepherine from the locus coeruleus. Sherman and Guillery's model derives largely from their studies of sensory cortices and feed forward pathways through them, projecting from more granular to less granular regions. Here we extend the discussion of thalamic relay neuron modulation to include a role for inhibitory inputs from the basal ganglia to thalamic nuclei that act as relays in the frontal lobe between more granular limbic and motor areas to less granular areas in these regions.

The basal ganglia (including ventral limbic and dorsal motor) are in a privileged position to influence traversals by means of their inhibitory inputs onto thalamic relay neurons within the Grand Loop. These inputs derive from nucleus inominata in the ventral limbic subpallium and from the globus pallidus in the dorsal motor subpallium. Neurons in the ventral pallidus (nucleus inominata) receive inhibition from medium spiny neurons (MSNs) in the nucleus accumbens (ventral striatum) and those in the dorsal globus pallidus from those in the dorsal striatum. These neurons then either directly disinhibit thalamic relay neurons or indirectly inhibit thalamic relay neurons

through an additional stage of inhibitory neurons (in globus pallidus, this is organized as direct and indirect projections through the external and internal segments). Spiking models of inhibitory pallidothalamic gating have focused on the bird song system (Goldberg et al., 2012), where gating inputs to thalamic relay neurons serve the role of transitioning syllables of the organism's vocalizations. Here we propose a more generic role for this gating in selecting and deselecting different pathways for internal traversals of the pallium.

Inputs to these direct and indirect pathways through basal ganglia derive from neocortical layer 5 neurons' projections onto MSNs, and their corticostriatal synapses undergo spike-timing dependent plasticity (STDP) which is modulated differentially by dopamine depending on the selective expression of either D1 dopamine receptors in the direct or D2 dopamine receptors in the indirect pathways (Pawlak and Kerr, 2008; **Figure 3**). Each layer 5 neuron's collaterals then include a branch descending to the brainstem or spinal cord, a branch descending to thalamus (Guillery and Sherman, 2011), and additionally a branch descending to striatum (Lévesque et al., 1996). A recent review of additional types of layer 5 projection neurons and the role of corticostriatal connectivity in disease provides a thorough examination and schematic of these pathways (Shepherd, 2013), and our model of thalamic gating, for now and for simplicity, includes only the "Pyramidal Tract" layer 5 neurons and their projections to basal ganglia and thalamus for the function of thalamic gating.
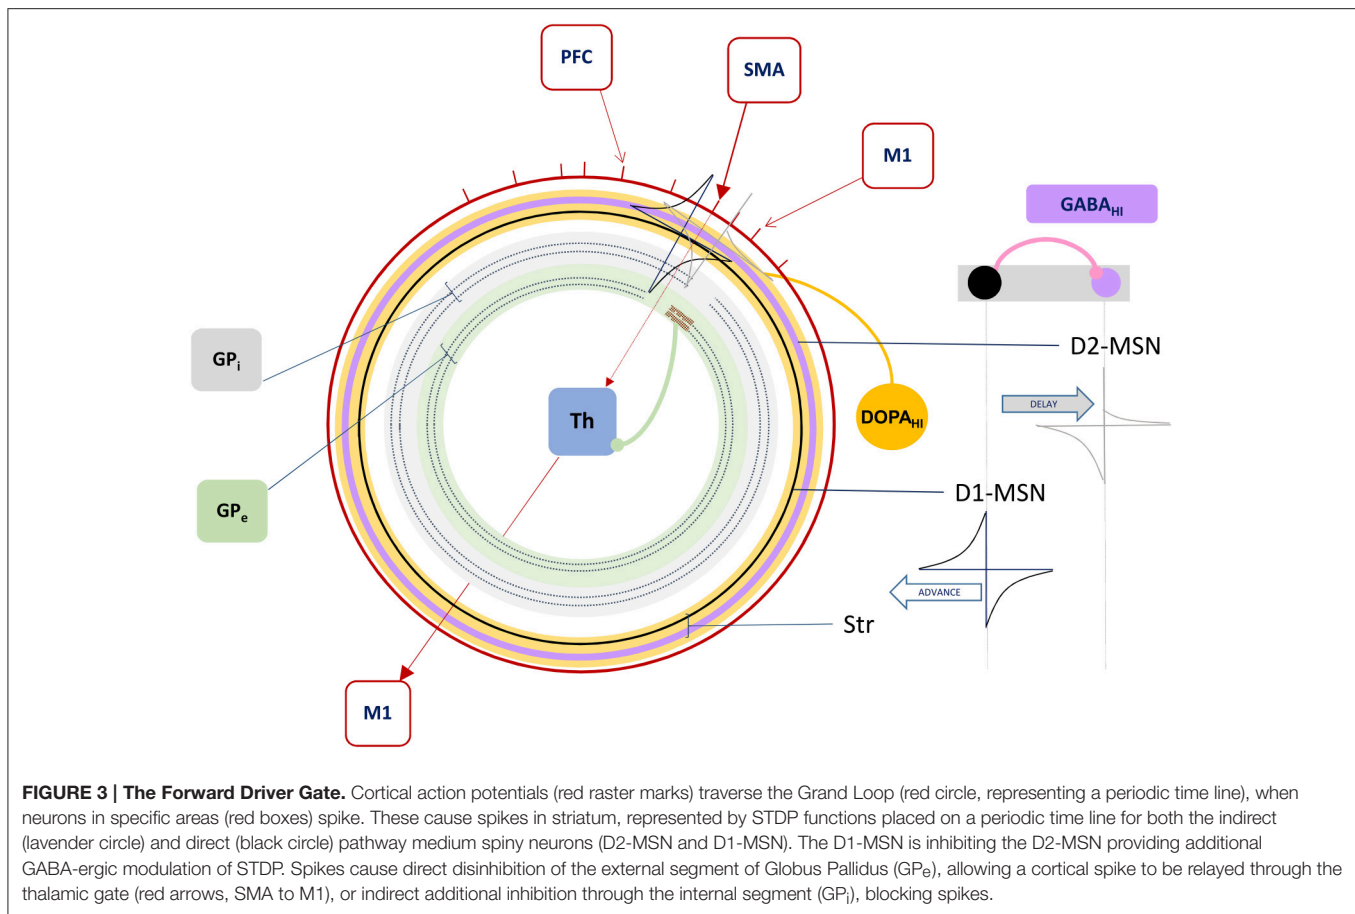
In summary, our model extends Sherman and Guillery's model of cortico-thalamo-cortical gating of driving, feed forward inputs to include modulation from striatal and pallidal neurons in both the direct and indirect pathways (**Figure 3**). MSNs in our model receive convergent layer 5 collaterals from all layer 5 neurons that send convergent collaterals onto a specific thalamic relay neuron. This relay neuron is then gated by the same MSNs, indirectly through globus pallidus (the gate opens for the direct pathway, and closes for the indirect pathway). Such a scheme does not preclude so called "closed loops" that originate and terminate in the same cortical area (Kelly and Strick, 2004), but downplays their significance as only partial regulators of feed forward thalamic gating (**Figure 4A**).

The basal ganglia in our model is then a "forward driver gate" for all feed forward driving signals relayed through the frontal lobe's cortico-thalamo-cortical functional pathways. Because these pathways relay layer 5 traversals through thalamus to the granular and supragranular layers of cortex, they can indirectly control the routing of feedback signals and the selection of certain traversals over others through the Grand Loop, as we describe in the next section. Additional area to area cortico-thalamo-cortical pathways not on the main loop backbone (such as the visual system) are then available for additional modulation and traversals of the global layer 5 behavioral network, possibly including loops requiring reafference from the environment.

2.4. Information Based Gating of Feedback Traversals

Our model provides two distinct functional signaling pathways through the Grand Loop: feed forward for driving the

FIGURE 3 | The Forward Driver Gate. Cortical action potentials (red raster marks) traverse the Grand Loop (red circle, representing a periodic time line), when neurons in specific areas (red boxes) spike. These cause spikes in striatum, represented by STDP functions placed on a periodic time line for both the indirect (lavender circle) and direct (black circle) pathway medium spiny neurons (D2-MSN and D1-MSN). The D1-MSN is inhibiting the D2-MSN providing additional GABA-ergic modulation of STDP. Spikes cause direct disinhibition of the external segment of Globus Pallidus (GP$_e$), allowing a cortical spike to be relayed through the thalamic gate (red arrows, SMA to M1), or indirect additional inhibition through the internal segment (GP$_i$), blocking spikes.
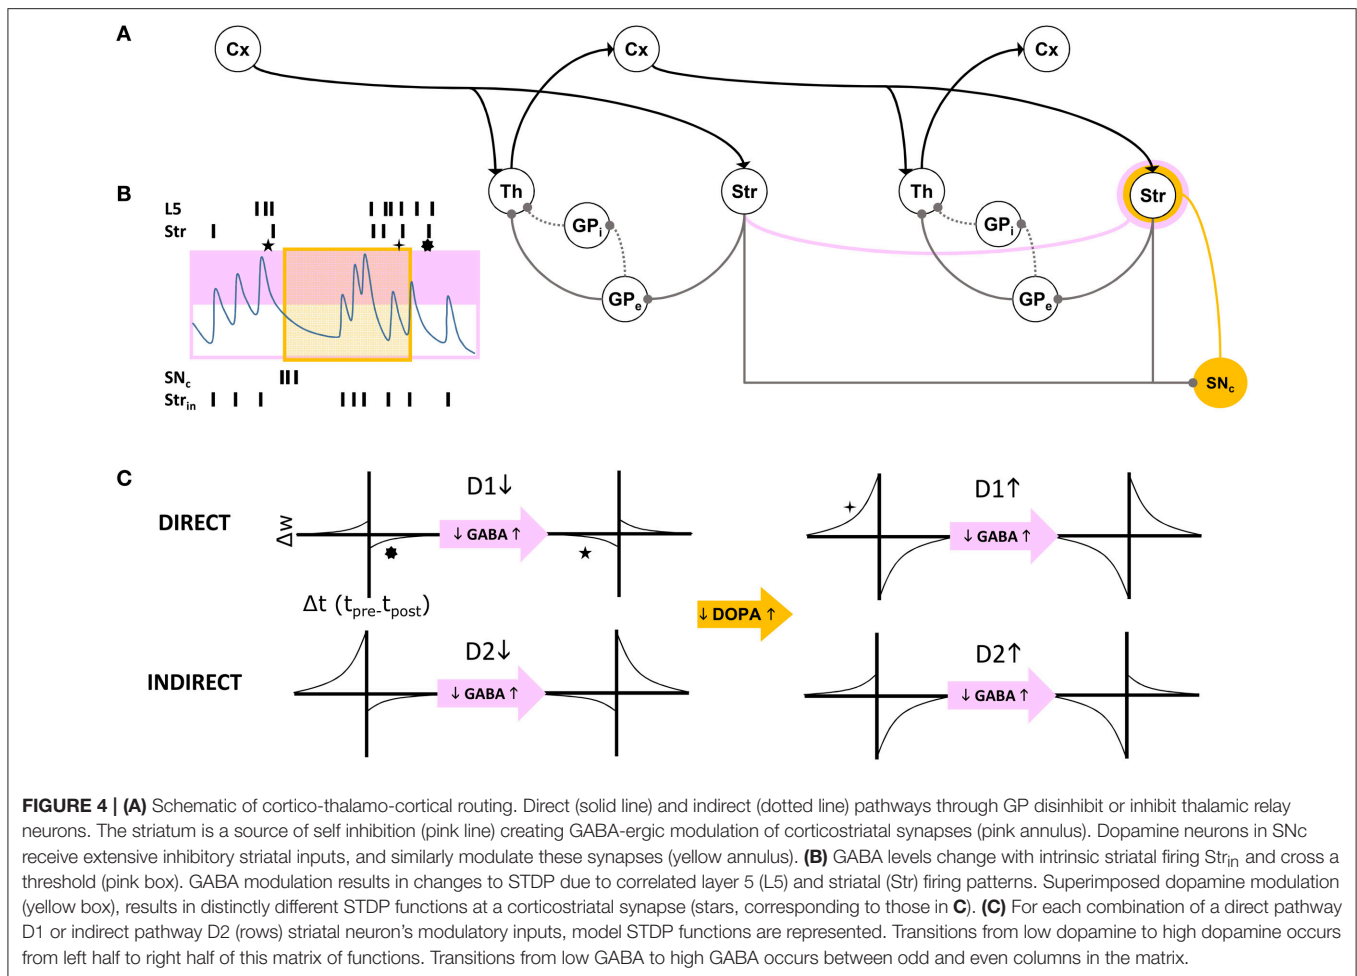
supragranular entropy maximizing network, and feedback for traversal of the infragranular behavior generation network. The latter, in our implementation of the model, propagates synfire events through a loop, as described by Zheng and Triesch in their model of "synfire ring" formation and propagation (Zheng and Triesch, 2014). Restricting synfire activity to the feedback direction is a key aspect of our model. Unlike other models of feedback, which ascribe to it solely a sensory processing "top down" function, we model the propagation of feedback activity as potentially independent of feed forward activity (for example when a coupling parameter between these two networks is zero). Specifically, the emergence of activations in the supragranular layers are rate coded, while activations in the infragranular layers are spike timing based in order to support synfire events. (We won't speculate here on how these distinct coding schemes are implemented and maintained by the neocortical microcicuit, but it would seem there are ample mechanisms available.)

Conditional coupling between features, extracted by information maximization in the supragranular layers, and spike propagation in the infragranular layers, is then under the control of a parameter that models cholinergic modulation in neocortex. Acetylcholine enhances the influence of sensory inputs on pyramidal cell firing relative to their processing of intrinsic signals within neocortical circuits (Hasselmo and

Giocomo, 2006). We model this modulatory parameter as changing the slope and dynamic range of a gain function. The function sets the gain on feedback integration within the synfire ring based on the level of activity in the corresponding functional units (e.g., orientation columns) in layer 2/3. Thus, feature encoding acts as a gate for synfire propagation, and we hypothesize this gain function may be implemented by layer 2/3 inputs to layer 5 neurons' apical dendrites. Varying cholinergic modulation of these inputs in cortex then controls the slope and dynamic range of the mapping from layer 2/3 activity to layer 5 feedback integration gain. The result is that propagation of synfire activity through a column of cortex is *informed* by the categorization of thalamic inputs to that same area. Information maximization among responses in the supragranular areas over environmental inputs becomes entropy maximization of synfire propagation pathways through the infragranular layers, provided that coupling between these is strong (i.e., cholinergic modulation is high). It is because of this coupling that we have named our model an *information-based exchange network*.

2.5. The Forward Driver Gate: Bursting, Modulation, and Plasticity

Having proposed a central cortico-thalamo-cortical routing function for striatal MSNs by means of their directly disinhibiting

FIGURE 4 | (A) Schematic of cortico-thalamo-cortical routing. Direct (solid line) and indirect (dotted line) pathways through GP disinhibit or inhibit thalamic relay neurons. The striatum is a source of self inhibition (pink line) creating GABA-ergic modulation of corticostriatal synapses (pink annulus). Dopamine neurons in SNc receive extensive inhibitory striatal inputs, and similarly modulate these synapses (yellow annulus). **(B)** GABA levels change with intrinsic striatal firing Str$_{in}$ and cross a threshold (pink box). GABA modulation results in changes to STDP due to correlated layer 5 (L5) and striatal (Str) firing patterns. Superimposed dopamine modulation (yellow box), results in distinctly different STDP functions at a corticostriatal synapse (stars, corresponding to those in **C**). **(C)** For each combination of a direct pathway D1 or indirect pathway D2 (rows) striatal neuron's modulatory inputs, model STDP functions are represented. Transitions from low dopamine to high dopamine occurs from left half to right half of this matrix of functions. Transitions from low GABA to high GABA occurs between odd and even columns in the matrix.

or indirectly inhibiting thalamic relay neurons, we will now propose on what basis a striatal MSN adapts to perform this function in the context of system set points. We call this the forward driver gate's "routing function." Our model of MSN firing includes constraints from a weak, assymetric lateral inhibitory network giving rise to "winnerless competition" (Rabinovich et al., 2001), and closely matching the periods of striatal bursting lasting hundreds of milliseconds observed *in vivo* (Miller et al., 2008). Ponzi and Wickens have similarly used this network to model spiking properties of striatum (Ponzi and Wickens, 2010), and have shown that at transition points in the lateral network configuration (from low, ∼10%, to high, ∼20%, rates of connectivity), an optimal balance is achieved that facilitates winnerless encoding of variations in driving inputs from neocortex (Ponzi and Wickens, 2013). To achieve this balance, our model instead varies the strength of cortical inputs dynamically by a dual source of modulation of STDP at the corticostriatal synapse.

The first dynamic modulator of STDP at the corticostriatal synapse in our model is GABA inhibition from the lateral network, itself responsible for "turn-taking" among MSNs and their bursts, characteristic of the winnerless network. We assume that both direct and indirect pathways show STPD reversal under

GABA inhibition (Fino et al., 2010; Paille et al., 2013), and we model winnerless competition between striatal neurons as the source for this inhibition (**Figure 4B**).

The second dynamic modulator of STDP at the corticostriatal synapse is dopamine. Given the routing function's potential as a critical determiner of the emergence of behavior, affect, and cognition in the organism via its direct control over traversals of the layer 5 network, reward-based learning of this function is ultimately required. For now, we simulate our brain model of information-based exchange with dopamine-based learning serving only a closed-loop function, separate from the environment and therefore independent of reward encoding. This closed-loop function is sensitive to system set points and monitors traversals. It is equivalent to so-called "tonic firing" in dopamine neurons, which can also include bursts. We propose that the intrinsic dynamics of dopamine neuron membrane currents implements this closed-loop function by measuring time and the abruptness of changes to system states, with bursts generated under specific conditions summarized below. Dopamine provides a potent modulation of STDP at the corticostriatal synapse (Pawlak and Kerr, 2008), and further modulates it differentially at the inputs to D1-MSNs and D2-MSNs. In our model this differential modulation, combined

with GABA modulation, produces the complex routing function summarized in **Figure 4C**.

Dopamine neurons have been shown to fire bursts of action potentials to signal basal reward inputs to the organism encoded as strong excitatory inputs to medial tegmentum and substantia nigra pars compacta (SNc) neurons. Because of these responses, the dopamine system has been extensively modeled as recapitulating reinforcement learning and operant conditioning in the organism. We propose here for the first time an additional closed-loop role for dopamine neurons in learning routing functions and selecting traversals. Specifically, we propose that dopamine neurons signal changes to traversals, and thereby influence the subsequent emergence of new traversals. The basis for this proposal derives from recent connectomics studies, which demonstrate that 70% of dopamine neuron inputs are inhibitory, and that most of this inhibition arises from striatum (Watabe-Uchida et al., 2012). Our closed-loop role for dopamine modulation depends on this inhibition, and this proportion and source suggests that closed-loop responses to inhibitory inputs, not open loop responses to basal excitation and reward, may be the predominant operating mode of the dopaminergic system.

Dopamine neurons exhibit heterogeneous combinations of intrinsic I_h and I_A currents (Amendola et al., 2012), as well as T-type calcium currents, which together generate post inhibitory rebound bursting in slice preparations. These currents' role *in vivo* has not yet been demonstrated, but our model assumes that the dynamical criteria for dopamine neuron bursting (and subsequent learning of routing at corticostriatal synapses) are implemented at least in part by rebound bursting. Other models have explored rebound bursting in dopamine neurons (Lobb et al., 2011), but not in the context of a closed-loop regulatory function. In our model, if the duration and abruptness of removal of striatal inhibition to dopamine neurons is appropriate, a rebound burst occurs. This aspect of the model indirectly imposes the additional criterion that inputs from layer 5 to striatum that transition the MSN winnerless network should be similarly matched to the duration and abruptness of change required for rebound spiking in dopamine neurons. In this way the striato-nigro-striatal loop monitors changes in traversals and alters routing within the feed forward entropy maximizing network by modulating corticostriatal STDP. We therefore propose that MSNs learn this routing based in part on their ability to recognize patterns of spiking in layer 5 that remain stable for a minimum duration of time then fade rapidly, a property expected during traversals of the Grand Loop.

3. SIMULATION METHODS

We simulated the model to explore its dynamics, characterize preliminary set points for measurement and analysis, and study traversal behavior and its regulation under different modulatory conditions. The five major components of the model to be simulated included cortical layers 2/3, 5, thalamus, striatum, and dopamine neurons. Meeting this challenge at the detailed level of neural tissue simulation is beyond the scope of this report, and without a good understanding of target model set points, likely impossible. We therefore aimed to draw upon four simplified abstractions of the key behaviors we ascribe to principle cells in these structures (Linsker, 1997; Rabinovich et al., 2001; Mihalas and Niebur, 2009; Zheng and Triesch, 2014). With four base component models replicated from other studies, we then coupled them across novel interfaces, realizing the closed, functioning Grand Loop, complete with its subcortical regulators.

3.1. Component Models

Four component models from the literature were targeted here to capture the functions of cortical layers 2/3 and 5, striatum, and dopamine neurons in the brain model. These four met sufficient requirements to implement information-based exchange (**Algorithm 1**), with very few changes to published parameters. We list the models below and describe the requirements they satisfy. Parameters defined in the original references for each component model are found in **Table 1**. Because thalamic relay neurons were implemented as a simple set of sums over inputs, they are described as an interface between component models in the subsequent section.

- **Neocortex, layer 2/3**: The model applies the "Infomax" algorithm of Bell and Sejnowski (1995) to thalamic relay neuron inputs. A neural network implementation of the same optimization (Linsker, 1997), based entirely on a local learning rule, establishes the biological plausibility of this function for this tissue (Kozloski et al., 2007). The three stage network modifies a weight matrix C that couples thalamic relay neurons to cortical units (e.g., at layer 4) based on microcircuit feedback from layer 2/3. Stage one receives the rate vector $\widehat{x}_I$ from an ensemble of time averaged thalamic spike trains and computes the zero mean input vector $x_I = \widehat{x}_I - x_0$, where $x_0 \approx \langle \widehat{x}_I \rangle$ is learned at the learning rate β_{x_0}. Stage two computes the sum of weighted inputs to stage three, $u \equiv C x_I$. In addition, each stage two unit computes an element of the output vector y, $y_i = \sigma(u_i)$, where $\sigma(\cdot)$ denotes a nonlinear squashing function, here the logistic transfer function, $y = 1/1 + e^{-(u+w_0)}$, where w_0 is an adaptive output bias vector learned by $\Delta w_0 = \beta_{w_0}[1 - 2y]$. The output vector y maximizes the mutual information over the input ensemble and provides regulatory microcircuit feedback to the model of layer 5 described below.

  Stage three then computes an entropy maximizing learning vector, which is fed back to stage two and applied by Hebbian learning to modify C at the learning rate β_{C_0}. Derived by Linsker (1997), this learning vector when applied in this way precisely yields the Infomax anti-redundancy term of Bell and Sejnowski (1995), $(C')^{-1}$ (ie., the inverse of the transpose of the input weight matrix) which for simplicity may also be computed directly. In this model of layer 2/3, the entropy maximizing learning vector emerges from a fully connected lateral network, whose weight matrix $\widehat{Q}$ undergoes Hebbian learning according to $\Delta \widehat{Q} = \beta_Q[uu' - \widehat{Q}]$, such that $\widehat{Q} \approx Q = \langle uu' \rangle$. For a given input presentation, these lateral connections evolve an auxiliary vector v according to $v_t = v_{t-1} + u - \alpha \widehat{Q} v_{t-1}$. Regardless of initial v, and assuming the scalar α is chosen so that v converges, the Infomax anti-redundancy term can be approximated by iterating the lateral network

TABLE 1 | Component models and parameters.

| Component	References	Parameter name	Parameter value
Neocortex, layers 4 and 2/3	Bell and Sejnowski, 1995;	β_{x_0}	0.00002
	Linsker, 1997;	β_{w_0}	0.0007
	Kozloski et al., 2007	β_C	0.0007
		β_Q	0.0007
Neocortex, layer 5	Zheng and Triesch, 2014	η_{IP}	0.01
		T^E_{max}	1.0
		T^I_{max}	0.5
		μ_{IP}	0.1
		σ_{HIP}	0
		η_{inhib}	0.001
		σ^2_ξ	0.01
		η_{STDP}	0.004
		η_{iSTDP}	11.0
Striatum	Rabinovich et al., 2001	g_{max}	0.25
		g_{min}	0
		a	0.7
		b	0.8
		τ_1	0.08
		τ_2	4.1
		ν	−1.5
		x_0	−1.2
		y_0	−1.62
		z_0	0
Dopamine neurons	Mihalas and Niebur, 2009	b	1.0
		G/C	50
		k_1	200
		k_2	20
		Θ_{inf}	−0.05
		R_1	0
		R_2	1.0
		E_L	−0.07
		V_R	−0.07
		Θ_r	−0.06
		a	1.0
		A_1	5.0
		A_2	−0.3

and applying its output by Hebbian learning, since $(C')^{-1} = Q^{-1}C\langle x_I x_I'\rangle$, and by substitution, $(C')^{-1} = \alpha\langle v_\infty x_I'\rangle$ (Linsker, 1997).

- **Neocortex, layer 5**: The model evolves from a self-organizing recurrent network (SORN) of binary spiking units through application of homeostatic plasticity, weight normalization, and STDP learning rules, together with synaptic pruning and synaptogenesis (Zheng and Triesch, 2014). This biologically consistent set of synaptic modifications creates distributions

of synaptic densities and weights that evolve over time to closely match data from developing neocortex. The weight matrix W also develops robust feed forward motifs and synfire activity similar to the model of Kozloski and Cecchi (2010), but with the remarkable topological feature of a closed, global loop of distinct propagation layers (**Figure 5**), which together engender "synfire rings." We evolved this network for 200, 000 time steps ($\Delta t = 1$ ms) to create four areas of cortex, which were then embedded into the larger model as two frontal lobe (M_1, M_{sup}) and two sensory lobe (S_1, S_{sec}) areas. Weights close to zero were held at zero for the remainder of all simulations. Propagating activity is maintained in the excitatory network, satisfying the requirement for layer 5 traversals. An inhibitory network that undergoes biologically plausible inhibitory STDP at its synapses onto excitatory neurons, together with homeostatic plasticity in the excitatory network, maintains spiking activity, $s(t)$, in the synfire ring at a nominal firing rate of 100 spikes/s. The inhibitory network imposes global, persistent competition across the network of excitatory layers. We propose this inhibition as an approximate functional model of inhibition from the thalamic reticular nucleus, which also integrates activity from across the thalamocortical system.

- **Striatum**: The model creates activation paths within the state space of a weakly connected, asymmetric inhibitory network to give rise to "winnerless competition" (Rabinovich et al., 2001), and alternating bouts of activity (i.e., "turn-taking") among the different neurons in the network (**Figure 6**). These bouts have been used by others to model the intrinsic dynamics of the striatum (Ponzi and Wickens, 2010), and together represent global attractor states that encode the modulatory and reorganizing influence of excitatory inputs to the network from cortical layer 5. Using a FitzHugh-Nagumo model, MSNs are represented by three dynamic variables. First, $x_f(t)$ in the model represents the "burst potential" of the neuron, with a positive transient in this potential representing a ∼350 ms burst. Computed using the same time step as the binary spiking layer 5 model, this coarse resolution model of the neuron's membrane potential is appropriate given the dominant bursting mode of firing in MSNs, and the observation that activity is often observed as alternating series of bursts of bursts (Miller et al., 2008). The remaining variables $y_f(t)$ represent a recovery from inhibition and $z_f(t)$ the inhibitory synaptic current received by the neuron, summed over the inhibitory inputs from other neurons through a Heaviside step function and the inhibitory weights W_{Str}.

- **Dopamine Neurons**: The model is that of a leaky integrate and fire (IAF) neuron. Four state variables are computed: a membrane potential $V(t)$, a variable threshold $\Theta(t)$, and two intrinsic currents $I_1(t)$ and $I_2(t)$, each integrated over the same time step as the previous two models. Because spikes in this model are represented by instantaneous resets of each variable at $V(t) > \Theta(t)$, the time step ($\Delta t = 1$ ms) is sufficient to integrate the neuron's spiking dynamics. Based on the published model, we derived an instance of a "rebound burst" model, and satisfied the requirements of dopamine neurons in the closed striato-nigro-striatal loop. Specifically,
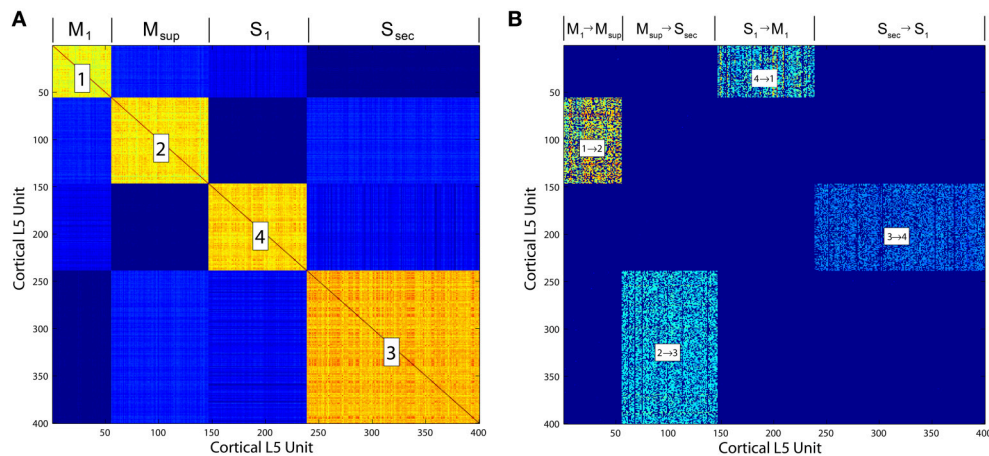
FIGURE 5 | (A) Correlation matrix computed over the final 10, 000 iterations of a simulation of the Layer 5 model based on Zheng and Triesch (2014). The four self-organized layers of this cortico-cortical topology are correlated in firing. **(B)** The cortico-cortical feedback weight matrix, showing clear dominance of the feed forward area to area connections over all others. This self-organized topology supports synfire ring activity, which in the current model is referred to as traversal activity.

the voltage-dependence of $\Theta(t)$ permits the model to generate rebound spiking under conditions when the neuron has been hyperpolarized deeply, or for a prolonged period (**Figure 7**). Due to the independent spike-induced current R_2, each rebound event generates a burst of four action potentials. This simplification's phenomenology also approximates that generated by other more complex models of rebound firing in dopamine neurons (Lobb et al., 2011).

3.2. Component Model Interfaces

The interfaces between component models that implement the integrated brain model of information-based exchange summarized in **Algorithm 1**, are now listed and described.

- **Feedback Cortico-cortical**: Layer 5 feedback inputs to a cortical area layer 5 are modeled as in the self-organizing recurrent network of Zheng and Triesch (2014) and implement the traversal network. Inputs are also categorized by the layer 2/3 model as an input vector of sums of binary spike trains over a time window τ_X. This vector $\widehat{x}_{FB}(t)$ comprises the elements $\widehat{x}_{FB_i}(t) \leftarrow \sum_{t-\tau_X}^{t} s_{FB_i}(T)$, where $s_{FB_i}(T) \in \{0, 1\}$ is the spike train from one unit in the upstream layer 5 area. Because of the full rank requirement of the information maximizing algorithm, the model of layer 2/3 includes a fixed first stage random mixing matrix M_{FB}, drawn from a lognormal distribution with unit mean and unit standard deviation, and linearly combining the elements of $\widehat{x}_{FB}$ to create the feedback input vector $x_{FB} \leftarrow M_{FB} \cdot \widehat{x}_{FB}$.

- **Feed Forward Cortico-thalamo-cortical**: Inputs to a thalamic relay neuron j projecting to a cortical area are modeled as a vector of sums over a time window τ_X of binary spike trains from layer 5 units in the cortical area projecting in the feed forward direction to the same area. This vector $\widehat{x}_{FF}(t)$ then comprises elements $\widehat{x}_{FF_i}(t) \leftarrow \sum_{t-\tau_X}^{t} s_{FF_i}(T)$, and is similarly transformed by a mixing matrix such that the thalamic relay neuron's activity $\theta_{FF} \leftarrow M_{FF} \cdot \widehat{x}_{FF}$. Each element θ_{FF_j} is then subjected to the forward driver gating function $G(\cdot)$ over the

gating vector $\widehat{G}$, such that elements of the feed forward input vector are $x_{FF_j} \leftarrow \widehat{G}_j \cdot \theta_{FF_j}$.

- **Layer 4 Thalamic and Feedback Inputs** : Feed forward thalamic inputs to layer 4 are combined with feedback inputs, such that the input vector to information maximization in layer 2/3, $\widehat{x}_I \leftarrow x_{FB} + x_{FF}$. It is at this stage also that sensory inputs from a simulated environment may be added to the model.

- **Layer 2/3 to Layer 5**: The Layer 2/3 output vector y provides an input to a gain function for layer 5's integration of binary spikes from feedback traversals of the Grand Loop. This gain function is a model of the layer 5 neuron's apical dendrite, and is parameterized by the term $Ach \in [0, 1]$, a proxy for the level of cholinergic modulation in neocortex. The gain on inputs to layer 5 unit j is then $U_j = [1 - Ach(1 - y_j)]/[1 - Ach/2]$, which at $Ach = 0$, preserves unitary gain regardless of y, and at $Ach = 1$ provides a gain $U \in (0, 2)$ for $y \in (0, 1)$. In this way, assuming information maximization divides the population into different halves of active and inactive units, the total synaptic input to the network will remain constant, since the U_j will always have a mean of 1, and is applied multiplicatively to the excitatory synaptic integration function of each layer 5 neuron as in Zheng and Triesch (2014).

- **Globus Pallidus to Thalamus**: The forward driver gating function G applied to thalamic relay neurons in the feed forward cortico-thalamo-cortical pathway models the final output of basal ganglia, a transient increase (via the indirect pathway) or decrease (via the direct pathway) in inhibition. $\widehat{G}$ is computed using a modified pallido-thalamic adjacency matrix D, comprising 1 for all direct pathway pallido-thalamic inputs, -1 for all indirect pathway pallido-thalamic inputs, and 0 for all unconnected pallidal to thalamic relay neurons. The bursting outputs of MSNs are represented by the half wave rectification function P, of the burst potential variable $x_f(t)$, and the gating vector is then $\widehat{G} = \mathcal{H}[D \cdot P(x_f(t))]$, where $\mathcal{H}$ is the Heaviside function.
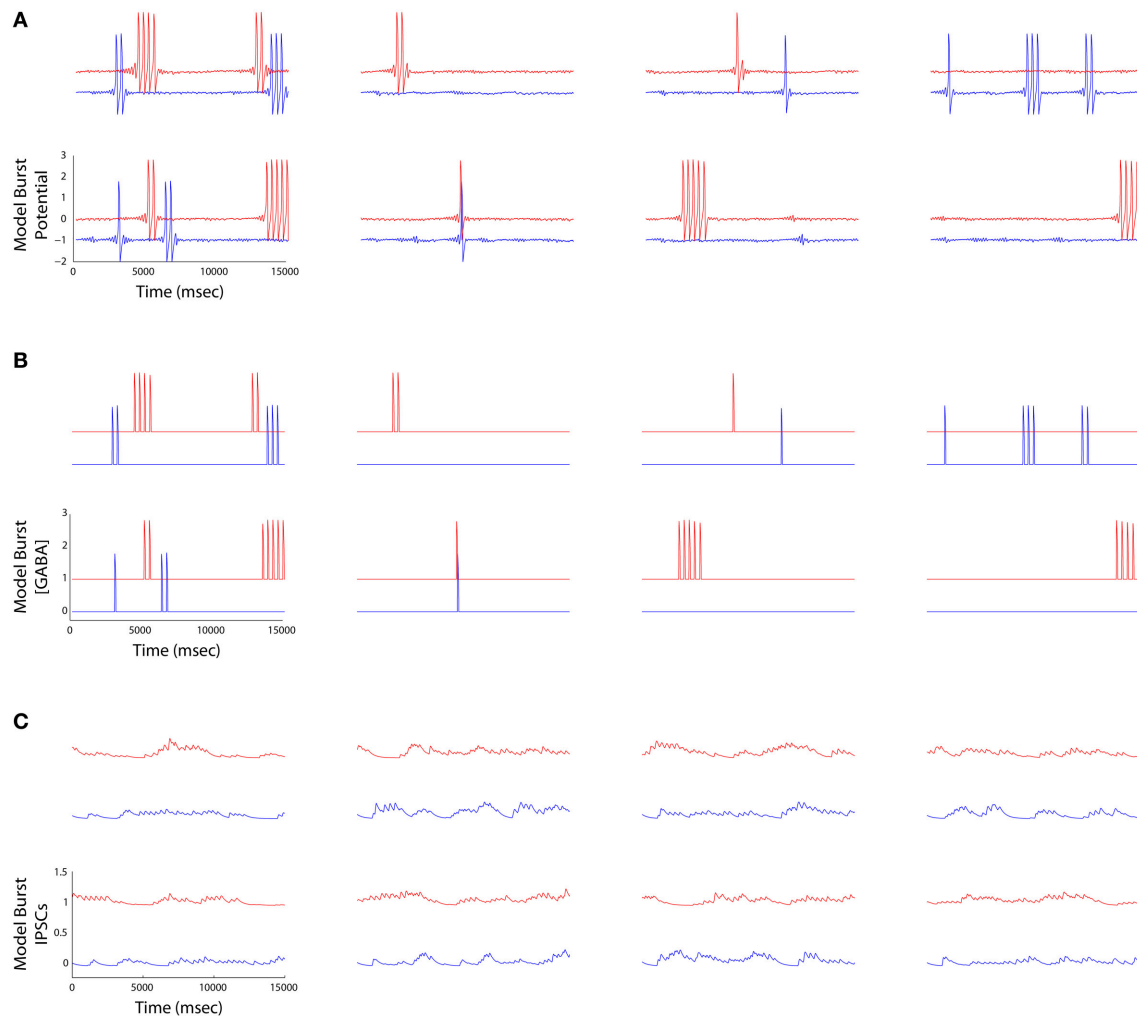
FIGURE 6 | (A) Striatal burst potential time series derived from FitzHugh-Nagumo models of four MSNs (columns) plotted at the beginning (blue) and the end (red) of a simulation. Each burst potential represents a series of MSN spikes fired in a burst. **(B)** The time aligned burst output traces represent the half wave rectified version of the potentials in **(A)**. **(C)** IPSCs received by each MSN (outward currents plotted as a positive deflection). Currents are maximal when the neurons bursting ceases, and are low when it is bursting, reflecting the operation of winnerless competition in the lateral inhibitory network.

- **Layer 5 to Striatum**: The inputs from the Layer 5 model to an MSN in the Striatum model are drawn from all layer 5 neurons in the cortical area for which the MSN gates inputs at the thalamus, and from those in the areas connected to it in either the feed forward or feedback directions. These Layer 5 inputs may also be directed to motor outputs of the model to a simulated environment (as in the Pyramidal tract). Corticostriatal synapses are subjected to STDP that differentially adjusts weights, W_{CStr}, based on correlation between cortical spiking and the derivative of the burst outputs of MSNs, $P'(x_f(t))$. Pre-post pairing is defined as when a cortical spike occurs and this derivative is positive, and post-pre pairing when a cortical spike occurs and it is negative. Each kind of pairing is computed separately and subjected to the modulatory conditions at the synapse, as illustrated in **Figure 4**. Briefly, depending on 1. the identity of the MSN (D1-

or D2-type), 2. whether dopamine is or is not present at the synapse, and 3. whether the inhibitory synaptic current $z_f(t)$ at the MSN exceeds a threshold ($z_f(t) > 0.00707$), each pairing value may be either 1 or 0, and the adjustment to the weight a multiple of this value and a learning rate of 0.002. As in Zheng and Triesch (2014), weights are normalized such that the sum of all inputs to an MSN cannot exceed 0.1. When weights reach zero they are pruned, and a single new connection may then be formed during a time step with probability 0.2.

- **Striatum to Dopamine Neurons**: The input to each dopamine neuron in the model, I_e (Mihalas and Niebur, 2009), is computed by summing all burst outputs from those MSNs projecting to the dopamine neuron, multiplied by a constant weight of -2.25.

- **Dopamine Release to Corticostriatal Synapses**: Unlike all other projections in the model's interfaces, the dopamine
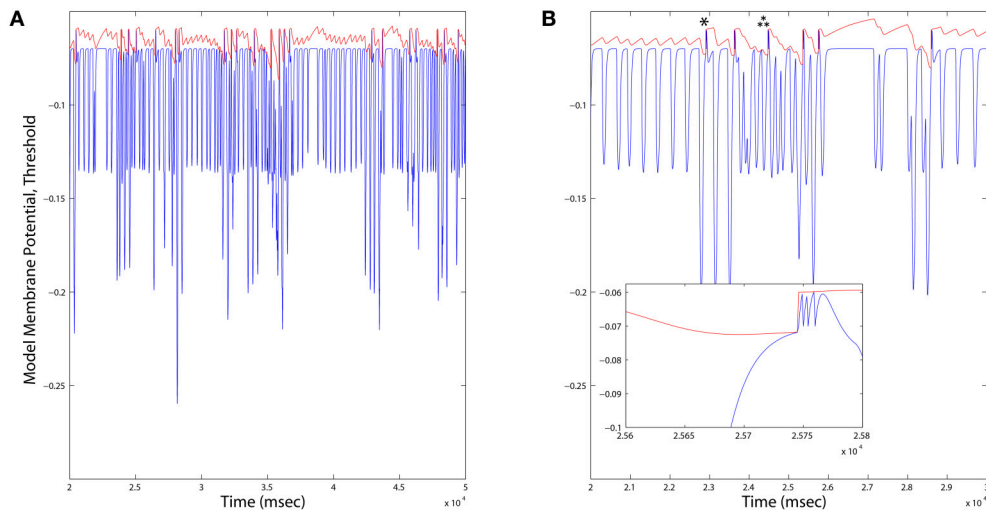
FIGURE 7 | (A) Long duration time series plots of Dopamine Neuron model variables. When the membrane potential (blue) reaches the variable threshold (red), a spike reset occurs. Firing rates of Dopamine neurons across all simulations were consistently on average ~1.6 Hz, and varied locally depending up the ongoing integration of dynamic inputs. **(B)** An expanded time scale reveals bursts (inset) occurring in response to deep hyperpolarization (single star), or prolonged weaker hyperpolarization (triple star) events.

neuron projection is to a synapse, not a neuron. Specifically, dopamine spiking results in a persistent dopamine modulation of STDP at a specific set of corticostriatal synapses. Dopamine neurons are assigned randomly without replacement to corticostriatal synapses onto each MSN. The duration of dopamine modulation following a Dopamine Neuron model spike persists at the synapse for a time τ_{DA}.

3.3. Simulation Materials and Experiments

We simulated the model to explore the rate and heterogeneity of transitions in traversals and in subcortical modulators of these traversals. The configuration (**Table 2**) allowed for a rapid prototyping in Matlab because of the simulation's small size. Following initialization of the cortico-cortical Grand Loop network of four areas, we simulated the full model for an additional 500, 000 iterations using an Intel Xeon E5-2640 v3 Processor (20 MB Cache, 2.60 GHz), requiring 2.5 hours of compute time. The first 50, 000 iterations were used to adjust the biases of the layer 2/3 model, during which time *Ach* modulation of layer 5 was drawn from the positive half of a zero mean normal distribution with standard deviation of 1. All plots, except where noted, show the final iterations of the 500, 000 total. Reported are experiments wherein the parameter τ_{DA} was set at either 25 or 100 ms, and *Ach* at 0.25, 0.5, or 0.75. All plots except **Figure 12** show results for $\tau_{\text{DA}} = 100$ ms.

4. SIMULATION RESULTS

4.1. Coordinated Behavior Among Component Models

Behavior of the model may be analyzed first based on inspection of various raster plots from different components of the model.

Algorithm 1: Information-Based Spike Exchange

1 **while** *Simulation Running* **do**
2   **for** $(M_1, M_{\text{sup}}, S_{\text{sec}}, S_1)$ **do**
3     $x_{\text{FB}}(t) \leftarrow M_{\text{FB}} \cdot \sum_{t-\tau_X}^{t} s_{\text{FB}}(T)$;
4     $\theta_{\text{FF}}(t) \leftarrow M_{\text{FF}} \cdot \sum_{t-\tau_X}^{t} s_{\text{FF}}(T)$;
5   **if** (M_1, M_{sup}) **then**
6       $x_{\text{FF}}(t) \leftarrow G(\theta_{\text{FF}}(t))$;
7   **else**
8       $x_{\text{FF}}(t) \leftarrow \theta_{\text{FF}}(t)$;
9   **end**
10     $\widehat{x}_I \leftarrow x_{\text{FB}} + x_{\text{FF}}$;
11     $y \leftarrow \sigma(Cx_I)$;
12     $\Delta C \leftarrow \texttt{Infomax}(x_I, C, y, \ldots)$;
13     $U \leftarrow [1 - Ach(1 - y_j)]/[1 - Ach/2]$;
14     $s(t) := \texttt{SORN}(s_{\text{FB}}(t - 1), W, U, \ldots)$;
15     $x_f(t) := \texttt{Winnerless}(s(t), W_{\text{CStr}}, P(x_f(t - 1)), W_{\text{Str}}, \ldots)$;
16     $V(t) := \texttt{IAF}(P(x_f(t)), \Theta(t - 1), I_1(t - 1), I_2(t - 1), \ldots)$;
17     $\Delta W_{\text{CStr}} \leftarrow \texttt{STDP}(s(t), P(x_f(t)), P(x_f(t - 1)), W_{\text{Str}}, \sum_{t-\tau_{\text{DA}}}^{t} V(T), \ldots)$;
18   **if** (M_1, M_{sup}) **then**
19       $\widehat{G} \leftarrow \mathcal{H}[D \cdot P(x_f(t))]$
20   **end**
21   **end**
22 **end**

In this way coordination between the different components is apparent. We first observed that cortico-cortical traversals through the feedback layer 5 network occur without subcortical

TABLE 2 | Configuration parameters specific to model simulation.

Parameter name	Description	Parameter value
N_A	Number of cortical areas	4
N_E	Number of thalamocortical units (layer 5, 2/3 pyramidal, thalamic relay neurons)	400
τ_X	Layer 2/3 integration window for spiking to rate code transformation (ms)	100
N_I	Number of thalamic reticular inhibitory neurons	80
$N_{Frontal}$	Number of frontal cortical areas under striatal gating	2
N_{Str}	Number of striatal MSNs	100
$N_{Cx,Str}$	Number of cortical neurons projecting to a striatal neuron	20
$W_{0Cx,Str}$	Initial corticostriatal weight	0.005, (0.1/$N_{Cx,Str}$)
$N_{Str,Th}$	Number of striatal neurons projecting to a thalamic relay neuron	11
N_{DA}	Number of dopamine neurons	20
$N_{Str,DA}$	Number of striatal neurons projecting to a dopamine neuron	20
τ_{DA}	Dopamine modulation window (ms)	25, 100

regulation, and were similar to the synfire events reported by Zheng and Triesch (2014). There are two main regulators of these traversals in our model: (1) an information based gain on layer 5 feedback inputs provided from layer 2/3, and (2) basal ganglia gating of cortico-thalamo-cortical feed forward inputs to layer 2/3 information maximization by the forward driver gate.

Upon introducing these regulators, we noted that traversals became structured into long bouts of smoothly alternating and repeating patterns of activity across the different cortical layers' raster plots. Each pattern persisted for ∼400 ms (**Figure 8A**), and sequences of patterns, while similar over each cycle, were not identical. The *Ach* parameter provides a means to adjust the influence of categories learned by layer 2/3 on traversals. For this initial experiment, *Ach* = 0.25 provided a gain $U \in (0.86, 1.14)$ for $y \in (0, 1)$.

Information maximization creates maximal entropy in the ensemble of output vectors over an input ensemble, and because of the logistic function, activity in each layer 2/3 neuron was typically close to zero or one. We interpret these values as cortical up and down states, which have both an extrinsic and intrinsic origin in the local cortical microcircuit.

Maximizing entropy of the ensemble of gain functions, applied to layer 5 inputs in the feedback traversal network, had the interesting effect of creating more irregularity in the patterns of activity across all of cortex as *Ach* increased. At *Ach* = 0.5, $U \in (0.67, 1.33)$, (**Figure 8B**), pattern combinations became varied, even though average global firing rates imposed by homeostatic plasticity in the network were consistently maintained (100 spiked/s). Finally, at *Ach* = 0.75, $U \in (0.4, 1.6)$, traversal transition rates increase significantly, and patterns were highly varied (**Figure 8C**).

Inspecting the information-bearing up and down states in layer 2/3 directly in state rasters from all four cortical areas

also reveals coordination between areas and with transitions in traversals. In **Figure 9A**, under *Ach* = 0.25, the rate of state changes among layer 2/3 units appeared coordinated, especially in the secondary sensory area. This coordination is less regularly transitioned than in the traversals, and occurs at a higher rate. At higher *Ach* = 0.5 (**Figure 9B**) up and down state coordination with traversals increases, while coordination across layer 2/3 is weakened. At *Ach* = 0.75 (**Figure 9C**) states become synchronized in the secondary sensory area and more coordinated with traversals overall, even though traversals themselves become more heterogeneous. Note that the heterogeneity in traversals due to increased control by the information maximizing network is not due to a lack of convergence in the weights of the networks. Weights among both the layer 2/3 Infomax input weights *C* and layer 5 feedback weights *W* converged during these simulations.

MSN bursts generated by the model were ongoing, as in the winnerless network and the model of Ponzi and Wickens (2010). These bursts appeared in fast sequences, which were of longer duration in D1-type MSNs than D2-type (**Figure 10**). Variability in burst rate between MSNs was also observed, with some not firing at all, likely because of inhibition from the active network. Increasing *Ach* had only a small effect on the raster appearance, and so we began our quantitative analysis by examining coordination between the Striatum model and the Layer 5 model.

4.2. Measurements of Information-Based Exchange

To quantify coordination between striatum bursting and cortical layer 5 spiking, we computed pairwise linear correlation coefficients between each cortical spike train and striatal burst train. We plotted each using a color scale (red, more correlated; blue, less correlated) in a matrix showing how different areas of cortex fired in relation to D1- and D2-type MSN bursts (**Figure 11**, left column). Only significant correlations were plotted, and all others were represented by zero. We also show that the mean of each distribution of correlation values (**Figure 11**, right column) for both D1- (blue) and D2-type (red) MSNs differed. Most coefficient distributions of D1 vs. D2 burst correlation with cortical spiking were significantly different($p < 0.05$), based on pairwise student *t*-tests. More striking is the difference in sign for each mean coefficient of correlation to each cortical area as *Ach* increases. Positive correlation coefficients dominated at low *Ach* and negative at high. At the intermediate level, M1 in particular showed a divergence in sign between mean correlation coefficients for D1-type (positive) and D2-type (negative) MSNs.

Finally, to quantify information-based exchange directly, we measured the entropy of cortical spiking and dopamine neuron spiking, and the mutual information between cortical and dopamine neuron spiking (Strong et al., 1998). Instead of measuring entropy and information among spike trains of individual neurons however, which quantifies the distribution of patterns of spikes over time, we measured entropy and information in population spiking, which
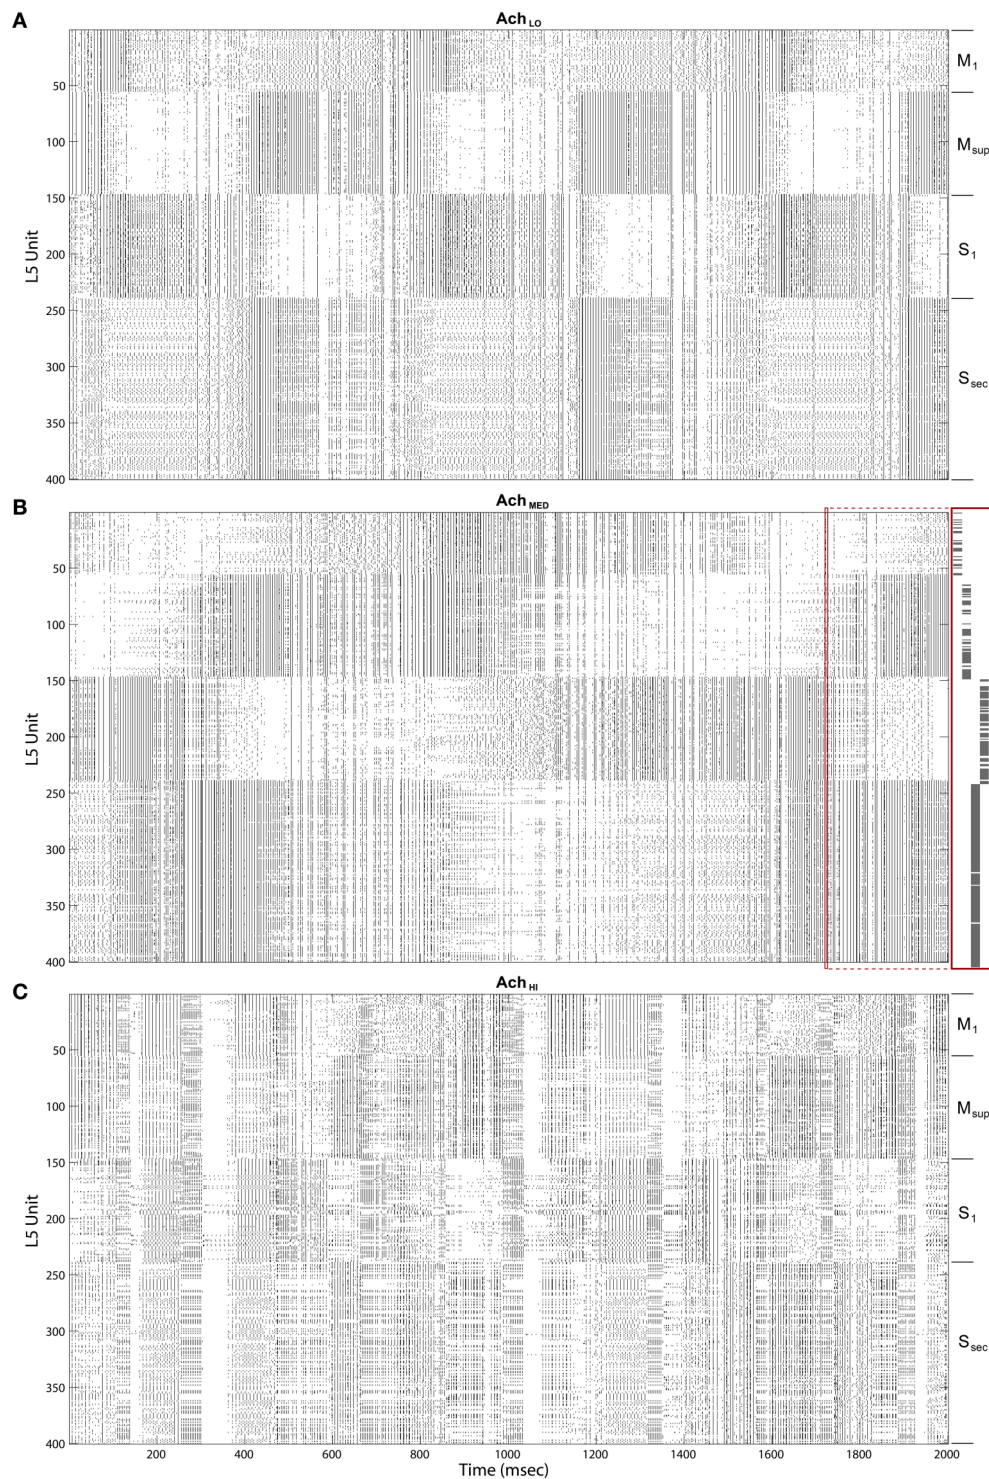
FIGURE 8 | Raster plots of the cortical Layer 5 model outputs. 400 spike trains from the final 2 s of simulated time. Cortical areas noted on right. **(A)** Under low *Ach* (0.25) patterns of traversals are long lasting (~400 ms) and smooth. **(B)** Under moderate *Ach* (0.5) patterns of traversals become briefer and choppy, with example of a single traversal of the Grand Loop expanded on right (red boxes). **(C)** Under high *Ach* (0.75) patterns of traversals are brief (100–200 ms) and heterogeneous.

quantifies the distribution of patterns of spikes over the population for single time steps. The method was aimed at asking if traversals themselves show entropy maximization

based on increased modulation from layer 2/3. Synfire events are encoded by the sets of units that participate at every stage of the chain or ring propagation. Therefore,
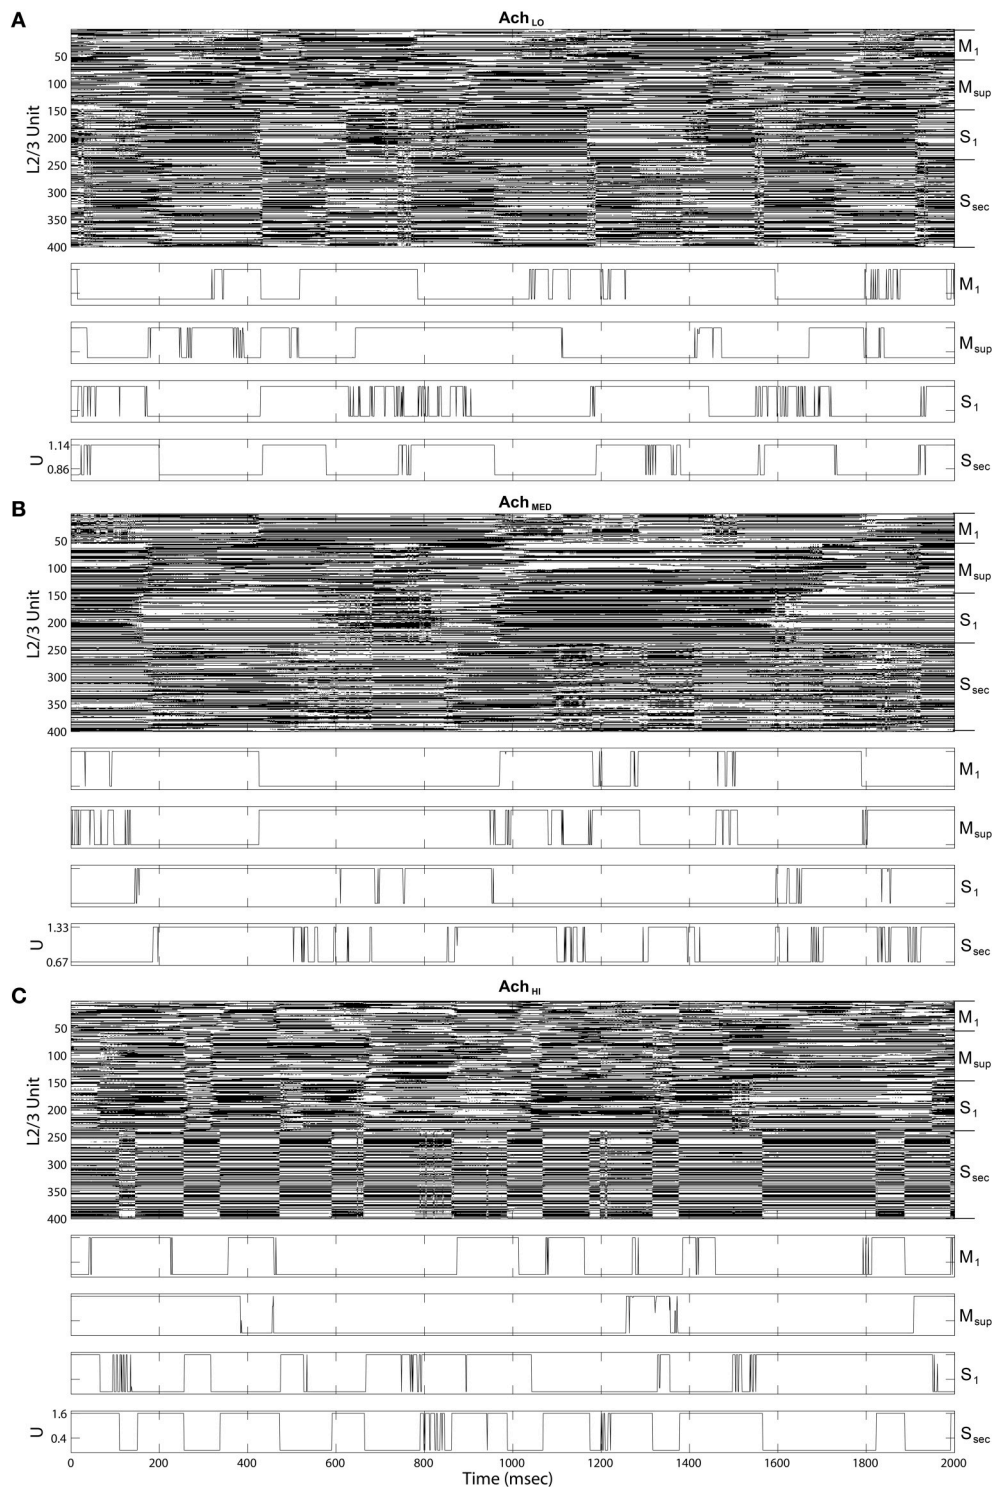
FIGURE 9 | Up state (black) raster plots of Layer 2/3 model outputs (upper panels) and example time series of gain *U* on traversal inputs to each area (lower panels). (A) Under low *Ach* (0.25), states transition more quickly than traversals from **Figure 8**. **(B)** Under moderate *Ach* (0.5), states transition more slowly in sensory areas. **(C)** Under high *Ach* (0.75), up and down states become synchronized in the secondary sensory area.

if the entropy of synfire population spiking increases, it can be concluded that the synfire chain entropy itself has increased.

We found that entropy in cortical layer 5 population spiking increased as *Ach* increased (**Figure 12A**). We also show that as the window of dopamine integration τ_{DA} increased, the
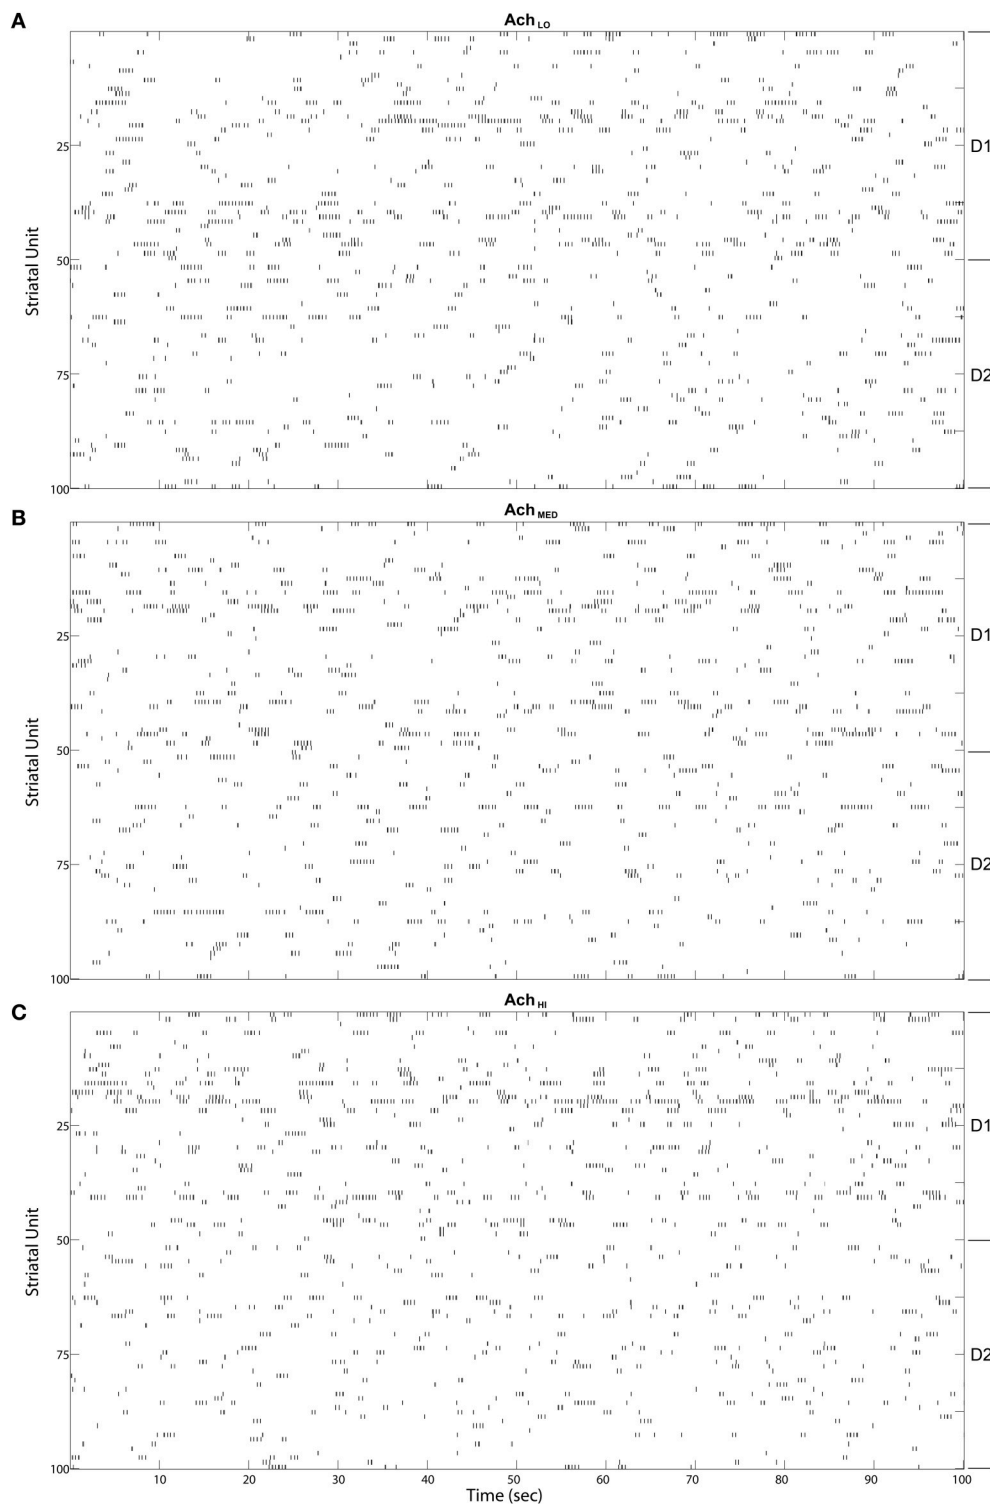
FIGURE 10 | Striatal model MSN burst raster plots over the final 100 s of simulated time. Bursts come in alternating bursts of bursts across the population due to winnerless competition. D1-type and D2-type of MSN noted at right. Bursts of bursts are longer in duration among D1-type MSNs under **(A)** Low (0.25), **(B)** Moderate (0.5), and **(C)** High (0.75) *Ach*.
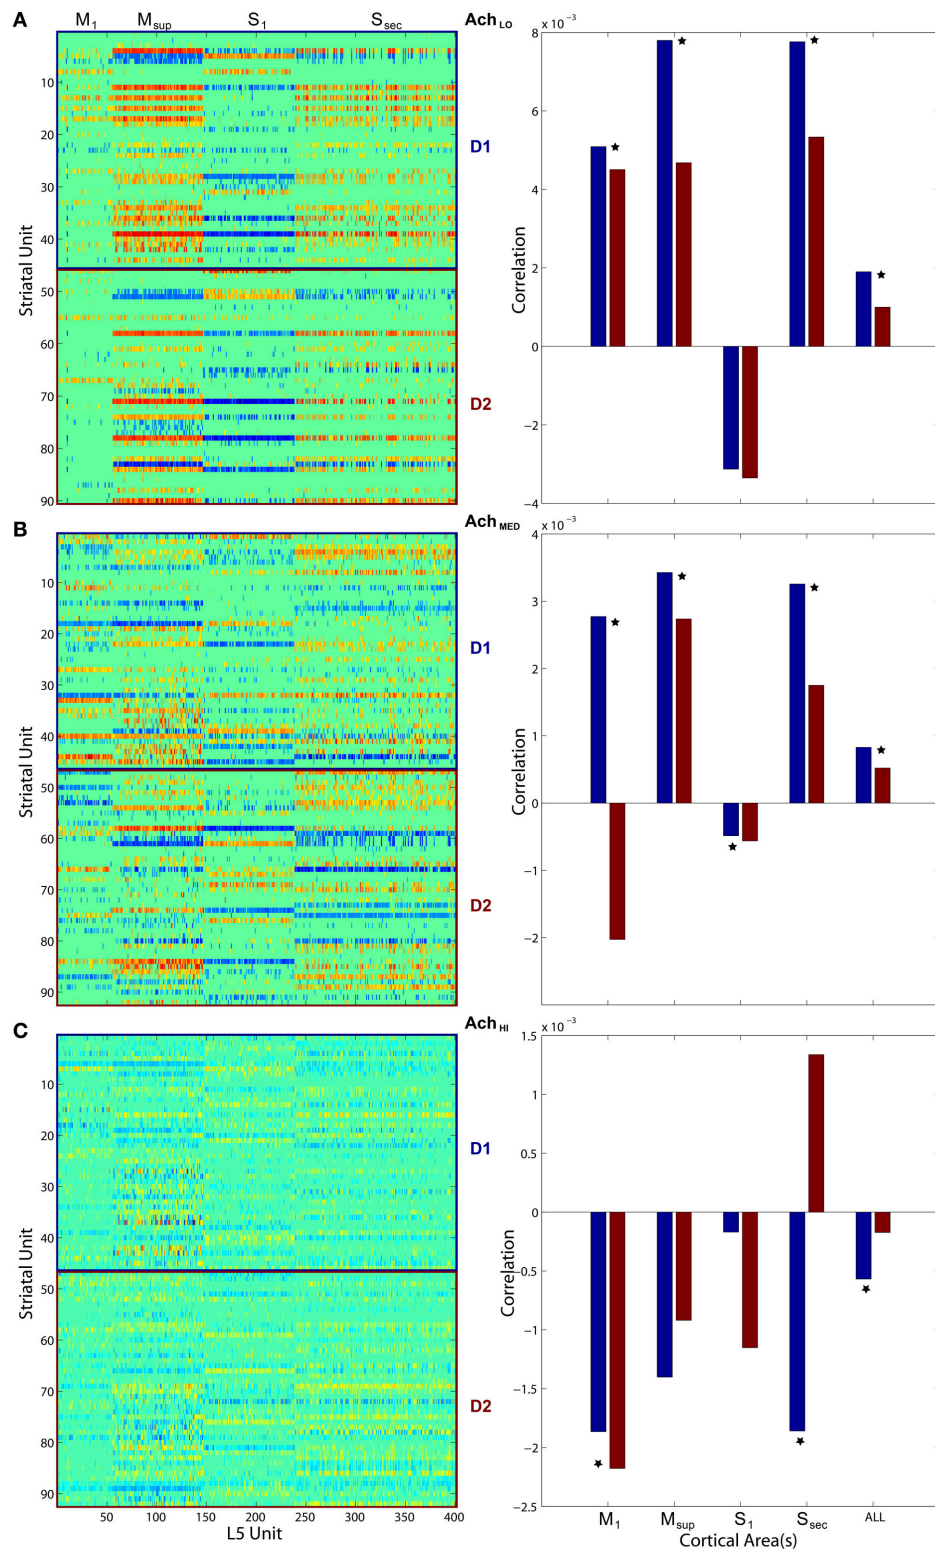
FIGURE 11 | Pairwise linear correlations coefficients between cortical spike trains and striatal bursts from Figures 8, 10. Only significant correlations are shown in the matrices plotted on left (D1- and D2-types noted along right) with others shown as zero. Mean correlation values are plotted on right, for D1- (blue) and D2-type MSNs (red). A star indicates significant difference ($p < 0.05$) in the distributions of coefficients (pairwise student t-test). **(A)** Under low *Ach* (0.25), correlations are positive across motor areas and all areas combined. **(B)** Under moderate *Ach* (0.5), correlations become positive for primary motor D1 and negative for D2. **(C)** Under high *Ach* (0.75), correlations are negative across motor areas and all areas combined.
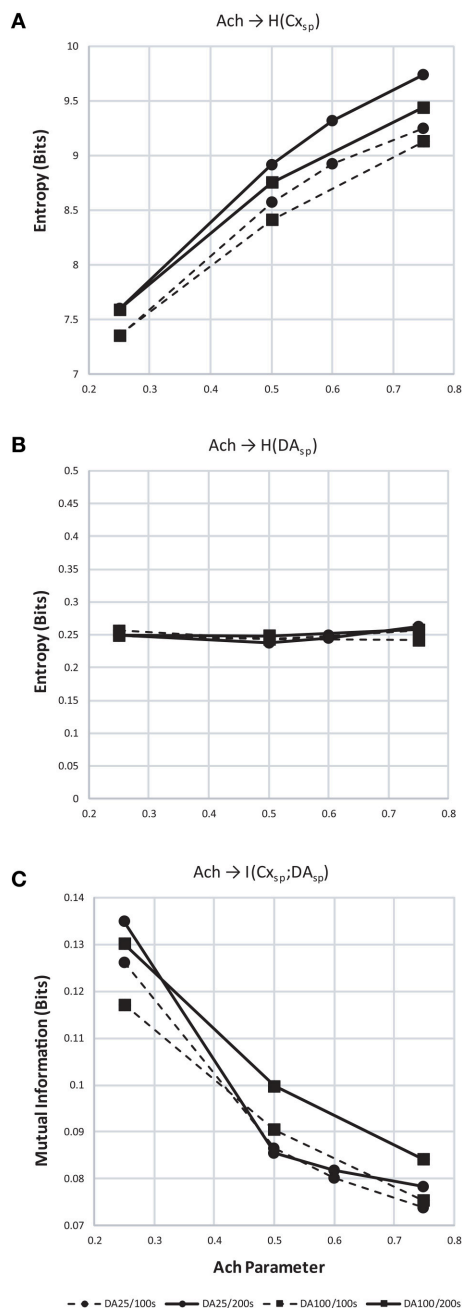
FIGURE 12 | Entropy and mutual information computed for cortical and dopamine neuron population spiking. (A) Entropy (*H*) of cortical population spiking for final 100 (dashed) and 200 (solid line) seconds of simulated time under increasing *Ach*. Circles plot simulations with τ_{DA} equal to 25 ms and squares 100 ms. In all conditions *H* increases with *Ach* and τ_{DA}. **(B)** Same as **(A)** but showing entropy of dopamine neuron population spiking. **(C)** With increasing cortical population spiking entropy **(A)**, the mutual information between cortical and dopamine neuron population spiking decreased.

entropy depends on dopamine population spiking, we measured the mutual information between these two populations, and found it to decrease as *Ach* increased (**Figure 12C**).

5. DISCUSSION

We discuss the brain model of information-based exchange in three contexts: brain evolution and development, brain resting state networks, and new approaches to the study of brain disorders such as neurodegenerative diseases.
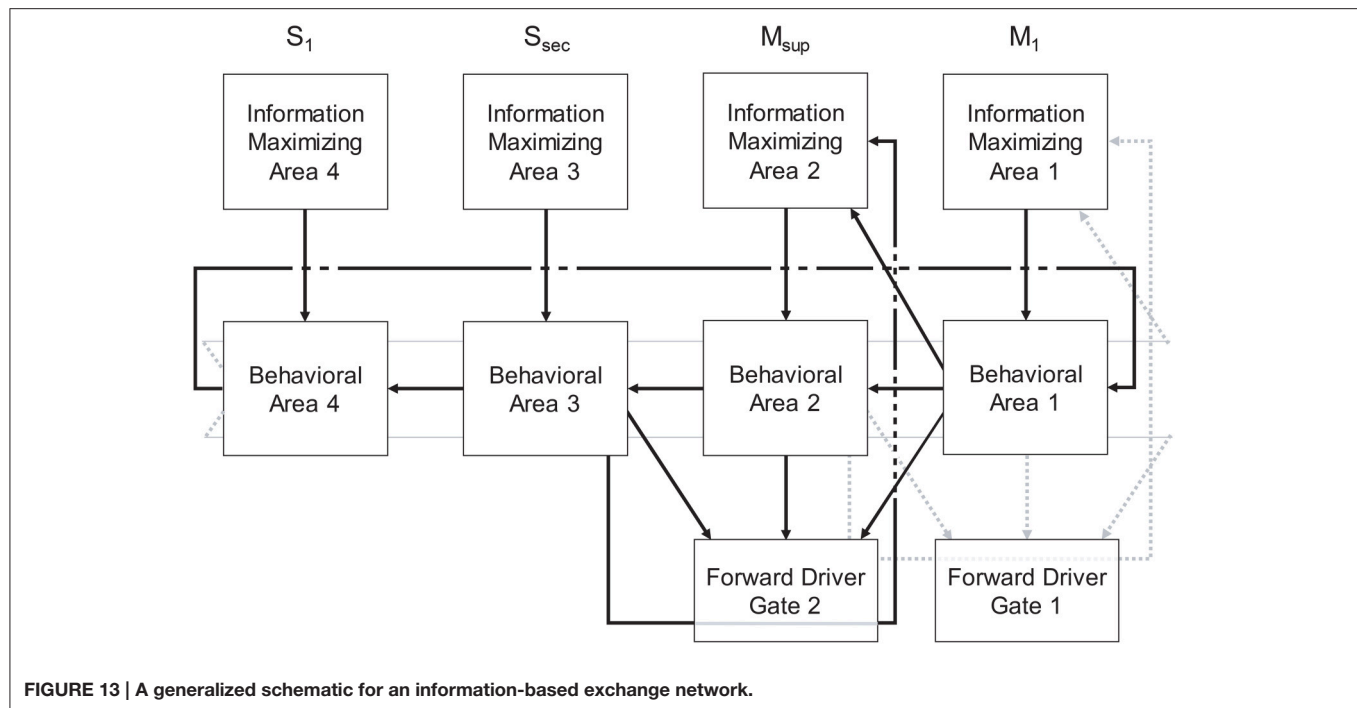
5.1. Brain Evolution and Development

We propose that the Grand Loop, spanning sensory, limbic, and motor cortices, and specifically traversing in our model *somatosensory* cortices, is prototypical and embryonic in origin, since other modalities develop fully only after birth and do not share a granular-agranular tiling boundary in von Economo's map. The topological relationship between other modalities and this backbone may then provide alternative pathways for completing a full traversal and rapidly binding percepts, needs, and behaviors. Finally, the tight coupling between somatosensory inputs and limbic states (i.e., tissue damage, pain) and motor states (i.e., sensorimotor feedback, proprioception) argues that this loop is likely preeminent in both brain evolution, organization, and development.

This model additionally provides insights into those organisms lacking cortices, wherein the stages of the proposed traversals may not be segregated anatomically (e.g., into Brodmann areas), but instead may be nucleated (e.g., in the birdsong system), or even superimposed within the same pallial regions (e.g., in fish and amphibians). Synfire ring development is robust given the synaptic modifications proposed by Zheng and Triesch (2014). It furthermore does not require anatomical segregation between layers to emerge or for synfire activity to propagate (e.g., for **Figure 5**, we sorted each matrix after areas developed in order to illustrate them clearly and connect subcortical structures to each).

Synfire rings may represent a prototypical substrate for behavior generation (**Figure 13**), and through subpallial regulatory inputs from thalamus and basal ganglia as described herein, for behavior selection. In such a scenerio, the evolution of a multilaminate neocortex to support such rings may have solved the problem of entropy maximization over the ensemble of synfire events in very large networks. Since the neural network implementation of Infomax requires a dense lateral network, to optimize each stage of a synfire ring and traversals in general would necessitate both the segregation of stages and a registered information maximizing network (**Figure 13**). This solution to the problem would then support rapid expansion of the synfire ring substrate by evolution, given that redundancy in large networks could suddenly be managed and eliminated by information maximization.

5.2. Resting State Networks

The challenge of modeling resting state activity in the brain has presented itself based on observations that distinct networks spanning multiple cortical areas appear in imaging studies to

entropy of layer 5 population spiking increased slightly as well. Surprisingly, the entropy of dopamine neuron population spiking (**Figure 12B**) remained constant while both parameters in the model were altered. Finally, to measure how increasing traversal

FIGURE 13 | A generalized schematic for an information-based exchange network.

serve either active or inactive states of the organism (Fox et al., 2005). Inactivity correlated networks appear even under anesthesia (Vincent et al., 2007), and these areas have very high metabolic rates, tipping the brain's energy budget toward a large investment in the organism's doing nothing.

What this costly outlay accomplishes may be explained by our model's use of closed-loop activity in the information-based exchange network to increase entropy over the ensemble of traversals. In an evolutionary context, this activity may be viewed as preadapting the brain to selecting novel behaviors in novel contexts by maximizing such a quantity first, before engaging with the environment, then using the preadapted diverse traversals to explore it and seek reward.

While others have noted that resting state dynamics may represent a "constant state of inner exploration" (Deco et al., 2011), our model is the first to assign a quantitative measure to the fruits of this brain activity, providing a new way to reason about the trade-off between evolutionary pressures toward latent adaptive behaviors and the large metabolic cost of resting state network activity.

5.3. Dynamic Disease Risk

We hypothesize that basic controls are required to establish "cognitive homeostasis," i.e., a process by which variables that change brain dynamics are carefully regulated so that properties of brain state transitions (and thus brain information processing and behavioral dynamics) remain relatively stable under constant neuromodulatory conditions. We refer to these stable properties as "set points," i.e., targeted norms for critical system variables supporting normal behavior, percepts, affect, and cognition. In our model, these controls are based on a consistent set of parameters that yield consistent spiking and bursting patterns,

even when the network undergoes reorganization (e.g., when *Ach* was modified, the system adjusted and produced stable traversals). Stable ranges of firing among burst rates and traversals, coefficients of correlated firing and bursts, and entropy and mutual information among population spiking and bursting have been our initial targets for describing these system set points using the brain model.

In real brains, given evolutionary pressures for robust self-regulation and behavior, the system is certainly replete with controls aimed at maintaining these set points. The challenge of studying brain disorders such as neurodegenerative disease is sorting primary and secondary risks from the multitude of compensatory mechanisms, each of which manifests itself as a deviation from normal brain and neuronal function given some primary genetic or injury risk. Researchers have shown, for example, that mutant Huntingtin protein disturbs NMDA receptor localization, densities, and currents at the corticostriatal synapse in mouse models of the disease (Cepeda et al., 2001). Knowing how this change arises and perturbs circuit dynamics, plasticity, and system set points may provide a better understanding of why certain neurons succumb and others don't when subjected to the same mutant protein.

We propose that perturbations in our model may result in stable dynamics, but with measurable risks related to stressors on normal neuronal function. If these deviations are extreme in our model, and therefore difficult to compensate for in biological tissue, a cascade of neuronal dysfunction may result. Neurodegenerative diseases such as Huntington's, Parkinson's, and Alzheimer's, may then be understood as cascading failures given initial stressors derived from plasticity abnormalities at the corticostriatal synapse, within the striato-nigro-striatal loop, and over the process of entropy maximization in layer 2/3,

respectively. For example, subtle changes to STDP or homeostatic plasticity may result in increased synaptic competition or cycling in the space of possible weights, which is then difficult to compensate for locally, given that traversals entail global brain states. If these risks increase when stressed neurons are removed from a simulation, the model may then be used to predict disease progression.

Implementation of the current brain model of information-based exchange forms a framework for the analysis of cognitive homeostasis and disease using IBM's scalable approach to structural and neurophysiological modeling of neocortex and brain nuclei (Kozloski and Wagner, 2011). Here we extend this approach and that of many brain modeling projects, which seem focused on validating complex local circuit and tissue models at the expense of validating tissue inputs. Minimal complexity brain models, in our case an information-based exchange network, may be necessary to capture brain dynamics and provide validatable inputs to complex tissue models. The current model has now been reimplemented in the same model graph simulation infrastructure in which IBM's Neural Tissue Simulator was implemented, and thus will allow direct coupling between these in a single scalable, extensible program.

With this new approach, inputs and models of the various components may be simulated and compared to *in vivo* experimental observations. Furthermore, simulations over very long time scales can be used to stress the model and its set points in physiologically and clinically realistic ways. Additional perturbations to the model may include physiological stimulation, such as simulated deep brain stimulation (DBS) in simulated neural tissue, simulated drugs with known targets in the detailed model, and different simulated disease states with hypothesized mechanisms at the level of gene, protein, regulatory network, etc. Stimulation, drug effects, and disease mechanisms can then be targeted to test certain hypotheses about

modifications to dynamic disease risk, and to study the wider system's behavior. Increasing complexity of perturbation sets (targets and combinations) may be designed to validate the model under different therapeutic conditions, and to test for phenotypic outcomes (e.g., symptomatology). Therefore, elaboration of these simulations within each modeled neural tissue might allow for *in silico* study of therapeutic interventions in living brain tissue.

In the above discussion, a model of several brain circuit components and their global set points has been proposed as a means to test disease mechanisms and therapeutic inputs such as DBS and drugs. The implicit assumption of these tests is that risks can be inferred from outlier variables that maintain system set points, and that these outliers may then be implicated as causes of phenotypic symptoms such as abnormal behavior at the organismal, circuit, neuronal, or synapse level. Targeting these variables in real world systems is one approach we propose for novel therapeutic design and discovery using brain modeling combined with neural tissue simulation.

REFERENCES

Abeles, M. (1991). *Corticonics.* Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511574566

Amendola, J., Woodhouse, A., Martin-Eauclaire, M., and Goaillard, J. (2012). Ca^{2+}/cAMP-sensitive covariation of I(A) and I(H) voltage dependences tunes rebound firing in dopaminergic neurons. *J. Neurosci.* 32, 2166–2181. doi: 10.1523/JNEUROSCI.5297-11.2012

Bell, A. J., and Sejnowski, T. J. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural Comput.* 7, 1129–1159. doi: 10.1162/neco.1995.7.6.1129

Beul, S., and Hilgetag, C. (2015). Towards a "canonical" agranular cortical microcircuit. *Front. Neuroanat.* 8:165. doi: 10.3389/fnana.2014.00165

Binzegger, T., Douglas, R., 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

Cepeda, C., Ariano, M., Calvert, C., Flores-Hernndez, J., Chandler, S., Leavitt, B., et al. (2001). NMDA receptor function in mouse models of Huntington's disease. *J. Neurosci.* 66, 525–539. doi: 10.1002/jnr.1244

Deco, G., Jirsa, V., and McIntosh, A. (2011). Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nat. Rev. Neurosci.* 12, 43–56. doi: 10.1038/nrn2961

Fino, E., Paille, V., Cui, Y., Morera-Herreras, T., Deniau, J.-M., and Venance, L. (2010). Distinct coincidence detectors govern the corticostriatal

spike timing-dependent plasticity. *J. Physiol.* 588, 3045–3062. doi: 10.1113/jphysiol.2010.188466

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., and Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc. Natl. Acad. Sci. U.S.A.* 102, 9673–9678. doi: 10.1073/pnas.0504136102

Goldberg, J. H., Farries, M. A., and Fee, M. S. (2012). Integration of cortical and pallidal inputs in the basal ganglia-recipient thalamus of singing birds. *J. Neurophysiol.* 108, 1403–1429. doi: 10.1152/jn.00056.2012

Guillery, R. W., and Sherman, S. M. (2011). Branched thalamic afferents: what are the messages that they relay to the cortex. *Brain Res. Rev.* 66, 205–219. doi: 10.1016/j.brainresrev.2010.08.001

Hasselmo, M. E., and Giocomo, L. M. (2006). Cholinergic modulation of cortical function. *J. Mol. Neurosci.* 30, 1559–1166. doi: 10.1385/JMN:30:1:133

Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.* 79, 2554–2558. doi: 10.1073/pnas.79.8.2554

Kelly, R., and Strick, P. (2004). Macro-architecture of basal ganglia loops with the cerebral cortex: use of rabies virus to reveal multisynaptic circuits. *Prog. Brain Res.* 143, 449–459. doi: 10.1016/s0079-6123(03)43042-2

Kozloski, J., and Cecchi, G. (2010). A theory of loop formation and elimination by spike timing-dependent plasticity. *Front. Neural Circuits* 4:7. doi: 10.3389/fncir.2010.00007

Kozloski, J., Cecchi, G., Peck, C., and Rao, A. (2007). Topographic infomax in a neural multigrid. *Lect. Notes Comput. Sci.* 4492, 500–509. doi: 10.1007/978-3-540-72393-6-60

Kozloski, J., and Wagner, J. (2011). An ultrascalable solution to large-scale neural tissue simulation. *Front. Neuroinform.* 5:15. doi: 10.3389/fninf.2011.00015

Lévesque, M., Gagnon, S., Parent, A., and Deschênes, M. (1996). Axonal arborizations of corticostriatal and corticothalamic fibers arising from the second somatosensory area in the rat. *Cereb. Cortex* 6, 759–770. doi: 10.1093/cercor/6.6.759

Linsker, R. (1997). A local learning rule that enables information maximization for arbitrary input distributions. *Neural Comput.* 9, 1661–1665. doi: 10.1162/neco.1997.9.8.1661

Lobb, C., Troyer, T., Wilson, C., and Paladini, C. (2011). Disinhibition bursting of dopaminergic neurons. *Front. Syst. Neurosci.* 5:25. doi: 10.3389/fnsys.2011.00025

Mihalas, S., and Niebur, E. (2009). A generalized linear integrate-and-fire neural model produces diverse spiking behaviors. *Neural Comput.* 21, 704–718. doi: 10.1162/neco.2008.12-07-680

Miller, B. R., Walker, A. G., Shah, A. S., Barton, S. J., and Rebec, G. V. (2008). Dysregulated information processing by medium spiny neurons in striatum of freely behaving mouse models of Huntington's disease. *J. Neurophysiol.* 100, 2205–2216. doi: 10.1152/jn.90606.2008

Paille, V., Fino, E., Du, K., Morera-Herreras, T., Perez, S., Kotaleski, J. H., et al. (2013). Gabaergic circuits control spike-timing-dependent plasticity. *J. Neurosci.* 33, 9353–9363. doi: 10.1523/JNEUROSCI.5796-12.2013

Pawlak, V., and Kerr, J. (2008). Dopamine receptor activation is required for corticostriatal spike-timing-dependent plasticity. *J. Neurosci.* 28, 2435–2446. doi: 10.1523/JNEUROSCI.4402-07.2008

Ponzi, A., and Wickens, J. (2010). Sequentially switching cell assemblies in random inhibitory networks of spiking neurons in the striatum. *J. Neurosci.* 30, 5894–5911. doi: 10.1523/jneurosci.5540-09.2010

Ponzi, A., and Wickens, J. (2013). Optimal balance of the striatal medium spiny neuron network. *PLoS Comput. Biol.* 9:e1002954. doi: 10.1371/journal.pcbi.1002954

Rabinovich, M., Volkovskii, A., Lecanda, P., Huerta, R., Abarbanel, H. D., and Laurent, G. (2001). Dynamical encoding by networks of competing neuron groups: winnerless competition. *Phys. Rev. Lett.* 87:068102. doi: 10.1103/PhysRevLett.87.068102

Rempel-Clower, N., and Barbas, H. (2000). The laminar pattern of connections between prefrontal and anterior temporal cortices in the rhesus monkey is related to cortical structure and function. *Cereb. Cortex* 10, 851–865. doi: 10.1093/cercor/10.9.851

Shepherd, G. (2013). Corticostriatal connectivity and its role in disease. *Nat. Rev. Neurosci.* 14, 278–291. doi: 10.1038/nrn3469

Strong, S., Korberle, R., de Ruyter van Steveninck, R., and Bialek, W. (1998). Entropy and information in neural spike trains. *Phys. Rev. Lett.* 80, 197. doi: 10.1103/PhysRevLett.80.197

Vincent, J. L., Patel, G. H., Fox, M. D., Snyder, A. Z., Baker, J. T., Van Essen, D. C., et al. (2007). Intrinsic functional architecture in the anaesthetized monkey brain. *Nature* 447, 83–86. doi: 10.1038/nature05758

von Economo, C. (1929). *Die Encephalitis Lethargica, ihre Nachkrankheiten und Ihre behandlung.* Berlin: Urban und Schwarzenberg.

Watabe-Uchida, M., Zhu, L., Ogawa, S., Vamanrao, A., and Uchida, N. (2012). Whole-brain mapping of direct inputs to midbrain dopamine neurons. *Neuron* 74, 858–873. doi: 10.1016/j.neuron.2012.03.017

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

Advantages of publishing in Frontiers

OPEN ACCESS

Articles are free to read, for greatest visibility

COLLABORATIVE PEER-REVIEW

Designed to be rigorous – yet also collaborative, fair and constructive

FAST PUBLICATION

Average 85 days from submission to publication (across all journals)

COPYRIGHT TO AUTHORS

No limit to article distribution and re-use

TRANSPARENT

Editors and reviewers acknowledged by name on published articles

SUPPORT

By our Swiss-based editorial team

IMPACT METRICS

Advanced metrics track your article's impact

GLOBAL SPREAD

5'100'000+ monthly article views and downloads

LOOP RESEARCH NETWORK

Our network increases readership for your article

Find us on