Plasticity-Driven Self-Organization under Topological Constraints Accounts for Non-Random Features of Cortical Synaptic Wiring.

Daniel Miner[1,◕], Jochen Triesch[1,‡],
1 Frankfurt Institute for Advanced Studies, Frankfurt am Main, Germany
◕This author performed the programming, analysis, and writing, and contributed to the conceptualization.
‡This author provided significant expertise and consultation, provided initial editing, and contributed to the conceptualization.
*** E-mail: miner@fias.uni-frankfurt.de**

Abstract

Understanding the structure and dynamics of cortical connectivity is vital to understanding cortical function. Experimental data strongly suggest that local recurrent connectivity in the cortex is significantly non-random, exhibiting, among other properties, above-chance bidirectionality, an overrepresentation of certain triangular motifs, and a heavy-tailed distribution of synaptic efficacies. Additional evidence suggests significant distance dependency to connectivity over a scale of several hundred micrometers, and particular patterns of synaptic turnover dynamics. It is currently not understood what processes give rise to this combination of features of cortical wiring. We present a spiking network model of a cortical slice culture which, via the interactions of a few simple biologically motivated intrinsic, synaptic, and structural plasticity mechanisms, qualitatively reproduces these non-random effects when combined with simple topological constraints. Our model suggests that mechanisms of self-organization arising from a small number of plasticity rules provide a parsimonious explanation for numerous experimentally observed non-random features of recurrent cortical wiring. Interestingly, similar mechanisms have been shown to endow recurrent networks with powerful learning abilities, suggesting that these mechanism are central to understanding both structure and function of cortical synaptic wiring.

Author Summary

The problem of how the brain wires itself up has important implications for the understanding of both brain development and cognition. We present a self-organizing neural network model incorporating just a handful of plausible mechanisms of change and development alongside topological constraints. The model gives rise to numerous non-random features which have been observed in experiments, but never before simultaneously produced by a single model. The results imply that only a few simple mechanisms and constraints are required to produce, at least to the first approximation, various characteristic features of a typical fragment of brain microcircuitry. In the absence of any of these mechanisms, simultaneous production of all desired features fails, suggesting a minimal set of necessary mechanisms for their production.

Introduction

The patterns of synaptic connectivity in our brains are thought to be the neurophysiological substrate of our memories, and framework upon which our cognitive functions are computed. It is believed that a small population of strong synapses forms a relatively stable backbone in recurrent cortical networks – perhaps the basis of long-term memories – while a larger population of weaker connections forms a more dynamic pool with a high rate of turnover [1–3]. It has been shown that much of the lateral recurrent connectivity of the layers of the cortex is significantly non-random [4–6], with a focus on layer 5 (L5), as this is more conventionally

examined via slice studies. It is an open question which non-random features are developed as a result of direct genetic programming, neural plasticity under structured input, and spontaneous self-organization. We examine here several noted non-random features of recurrent cortical wiring that we believe can be explained as the result of spontaneous self-organization — specifically, self-organization driven by the interaction of multiple neural plasticity mechanisms under physiological constraints. The features we will examine are the heavy-tailed, log-normal-like distribution of synaptic efficacies or dendritic spine sizes [6–10] and their associated synaptic dynamics, and the overrepresentation of bidirectional connectivity and certain triangular graph motifs [6]. Understanding the development of micro-structure in the cortex has significant implications for the understand of both developmental and cognitive / computational processes as well. Such insight would be invaluable in understanding the root causes of cognitive and developmental impairments, as well as understanding better the nature of the computational mechanisms employed by the brain.

The interaction of multiple plasticity mechanisms, such as synaptic scaling and Hebbian plasticity has been studied before [11, 12], with results suggesting that the interactions for such mechanisms are useful for the formation and stability of patterns of representation. However, we desire a more detailed look at how such self-organization might take place in the cortex. The predecessor to the model we use to address these issues is the Self-Organizing Neural Network, or SORN [13]. The SORN is a recurrent network model of excitatory and inhibitory binary neurons which incorporates both Hebbian and homeostatic plasticity mechanisms. Specifically, it incorporates binarized spike timing dependent plasticity (STDP), synaptic normalization (SN), and intrinsic homeostatic plasticity (IP). It had been demonstrated to be computationally powerful and flexible for unsupervised sequence and pattern learning, presenting apparent approximate Bayesian inference and sampling-like behavior [14]. Additionally, it has been used to reproduce synaptic weight distributions and growth dynamics in the cortex [15].

In this paper, we formally introduce the LIF-SORN, a leaky integrate-and-fire based SORN-inspired network model that incorporates a spatial topology with a distance-dependent connection probability, in addition to more biologically plausible variants of and extensions to the plasticity mechanisms of the SORN. The LIF-SORN models a recurrently connected network of excitatory and inhibitory neurons in L5 of the neocortex, or a slice thereof. This new model is the first to reproduce numerous elements of the synaptic phenomena examined in [10], [16], and [15] in combination with the sort of non-random graph connectivity phenomena observed in [6]. The simultaneous reproduction of all these elements with a minimal set of plasticity mechanisms and constraints, and furthermore using only point neurons, represents an unprecedented success in explaining noted features of the cortical micro-connectome in terms of self-organization.

Materials and Methods

Simulation Methods.

We randomly populate a 1000 × 1000 μm grid with 400 LIF neurons with intrinsic Ornstein-Uhlenbeck membrane noise as the excitatory pool, and a similar (though faster refracting) population of 80 noisy LIF neurons as the inhibitory pool. All synapses are populated with a gaussian distance-dependent connection probability profile with a half-width of 200 μm. This particular profile is chosen as a middle ground between the results of [6], which finds no distance dependence up to a scale of 80 - 100 μm, and the results of [5], which finds an exponential distance dependence at a scale of 200 - 300 μm. Recurrent excitatory synapses are not populated, as they will be grown "naturally" with the structural plasticity. Excitatory to inhibitory and inhibitory to excitatory synapses are populated to a connection fraction of 0.1 and inhibitory recurrent synapses to a connection fraction of 0.5, in approximate accordance with L5 experimental data [17]. Excitatory to inhibitory, inhibitory to excitatory, and inhibitory to inhibitory connections are given fixed efficacies and connectivities. Recurrent excitatory connectivity is begun empty and is to be grown in the course of the simulation. The relevant parameters are summarized in Tables 1 and 2.

We use the Brian spiking neural network simulator [18]. The neuron model is a leaky integrate-and-fire (LIF) neuron, the behavior of which is defined by

$$\frac{dV}{dt} = -\frac{V - E_l}{\tau} + \frac{\sigma\xi}{\sqrt{\tau}} \quad , \tag{1}$$

where V is the membrane potential, E_l is the resting membrane potential, τ is the membrane time constant,

Table 1. Basic network parameters.

parameter	value
N_{exc}	400
N_{inh}	80
sheet size	1000×1000 μm
connection probability profile	200 μm half-width Gaussian

Table 2. Basic connectivity parameters. * indicates growth via plasticity

parameter	EE	EI	IE	II
connection fraction	target 0.1*	0.1	0.1	0.5
connection strength	insertion at 0.0001 mV*	1.5 mV	-1.5 mV	-1.5 mV
conduction delay	1.5 ms	0.5 ms	1.0 ms	1.0 ms

σ is the standard deviation of the intrinsic membrane noise, and ξ is the Ornstein-Uhlenbeck process which drives the noise. When V reaches a threshold V_T, the neuron spikes, and the membrane potential V is returned to V_{reset} (which may be lower than E_l in order to provide effective refractoriness). The parameters used are given in Table 3.

Table 3. LIF neuron parameters.

parameter	value
E_l	-60 mV
τ	20 ms
$V_{\text{reset}}^{\text{exc}}$	-70 mV
$V_{\text{reset}}^{\text{inh}}$	-60 mV
σ	$\sqrt{5}$ mV
V_T	variable via IP

All parameters are shared between excitatory and inhibitory units unless otherwise denoted by superscripts "exc" and "inh."

A simple transmitting synapse model is used, connecting neuron i to neuron j. When neuron i spikes, the synaptic weight W_{ij} is added to the membrane potential V of neuron j following the conduction delay for the type of connection (as in Table 2).

As in the original binary SORN, we include multiple plasticity mechanisms. The first is exponential spike timing dependent plasticity (STDP), which is executed at a biologically realistic timescale [19–24]. This defines the weight change to a synapse caused by a pair of pre- and post-synaptic spikes as in Equations 2, 3, and 4:

$$\Delta w_j = \sum_{f=1}^{N_f} \sum_{n=1}^{N_n} W\left(t_i^n - t_j^f\right) \tag{2}$$

$$W(x) = A_+ \exp\left(-x/\tau_+\right), \qquad x > 0 \tag{3}$$

$$W(x) = A_- \exp\left(x/\tau_-\right), \qquad x < 0. \tag{4}$$

Here, j indexes the synapse, f indexes presynaptic spikes, and n indexes postsynaptic spikes. A_+ and A_- are the maximal amplitudes of the weight changes, and τ_+ and τ_- are the time constants of the decay windows. Values are set to approximate experimental data; in particular, round numbers were chosen that roughly approximate data in [19] and [20], with $\tau_+ = 15$ ms, $A_+ = 15$ mV, $\tau_- = 30$ ms, and $A_- = 7.5$ mV. We use the "nearest neighbor" convention in order to efficiently implement this online, in which only the closest pairs of pre- and post-synaptic spikes are used.

Synaptic normalization [25–28] serves to ensure that a neuron receives a roughly constant input in the face of changing relative input strengths by homeostatic regulation of incoming connections to a neuron so that their sum approaches a constant value over a long timescale. Anatomically this is manifested in the behavior of synaptic areas on dendrites. [26] demonstrated, for example, that the summed synaptic area per μm of dendrite stays roughly constant over long-term potentiation, but the area per synapse increases while the number of synapses per μm of dendrite decreases. [28] suggests a multiplicative scaling dynamic manifests on a slow timescale. Most biological evidence suggest that synaptic normalization is not an instantaneous process, and may even involve multiple timescales or mechanisms working in concert. We use a simple multiplicative model with a fast (but variable) timescale to keep things computationally simple:

$$W_i \rightarrow \; W_i \left(1 + \eta_{\mathrm{SN}} \left(\frac{W_{\mathrm{total}}}{\sum_j^N W_{ij}} - 1\right)\right). \tag{5}$$

Here, W_i is the vector of incoming weights for any neuron i, W_{ij} are the weights of the individual synapses, W_{total} is the target total input for each neuron, and η_{SN} is a rate variable which, together with the size of the timestep, determines the timescale of the normalization. W_{total} is calculated before the simulation run for each of the four types of synapse (E to E, E to I, I to E, and I to I) by multiplying the connection fraction for that type of connection by the mean synapse strength by the size of the incoming neuron population (or, in the case of the E to E connections, the target values thereof). The timescale we use is on the order of seconds and therefore accelerated from biology; corresponding to an application of the process once per second and $\eta_{\mathrm{SN}} = 1.0$. The accelerated timescale is sufficiently separated from that of the STDP, which operates on the order to tens of milliseconds, to avoid unwanted interactions while decreasing the necessary simulation time.

Neuronal excitability is regulated by various mechanisms and over different time scales. On a very fast milliseconds time scale, a neuron's refractory mechanism prevents it from exhibiting excessive activity in response to very strong inputs [29]. This is inherently included in the neuron model we use. At a somewhat slower time scale, spike rate adaptation reduces a neuron's response to continuous drive [30]. Given that our model lacks any strong external drive, we neglect this. At very slow time scales of hours to days, intrinsic plasticity mechanisms change a neuron's excitability through the modification of voltage gated ion channels that can modify its firing threshold and the slope of its frequency-current curve in a homeostatic fashion [31,32]. Here we use a simple form of homeostatic firing threshold regulation implemented at discrete time steps in the following way:

$$V_T \rightarrow \; V_T + \eta_{\mathrm{IP}} \left(N_{\mathrm{spikes}} - h_{\mathrm{IP}}\right) \tag{6}$$

$$N_{\mathrm{spikes}} \rightarrow \; 0. \tag{7}$$

Here, V_T is the threshold for an individual neuron, η_{IP} is a learning rate, h_{IP} is the target number of spikes per update interval, and N_{spikes} is the number of times a neuron has spiked since the last time a homeostatic plasticity step was executed. This operation is performed at a biologically accelerated timescale. The desired target rate is chosen to be 3.0 Hz, so $h_{\mathrm{IP}} = 3.0\ \mathrm{Hz} \times 0.1\ \mathrm{ms} = 0.0003$ and η_{IP} is set to 0.1 mV. The operation is performed at timesteps of 0.1 ms, which is the basic timestep of the simulation. Like the SN process, the accelerated (relative to biology) timescale is sufficiently separated from the timescale of the STDP to avoid unwanted interactions while decreasing the necessary simulation time.

We implement structural plasticity via simultaneous synaptic pruning and synaptic growth. Synaptic pruning is implemented in a direct fashion in which synapses whose strength has been driven too close to zero (the threshold is 0.0001 mV) by the other plasticity mechanisms are eliminated. At the same time, new synapses are stochastically added, according to the distance-dependent per-pair connection probabilities, at a regular rate. This is done at an accelerated timescale by adding a random number of synapses (drawn from an appropriately scaled normal distribution) once a second. A mean growth rate is hand-tuned to lead to the desired excitatory-excitatory connection fraction. In this case, the growth rate is 800 synapses per second and the target connection fraction is 0.1. The synapses are added according to pre-calculated connection probabilities determined by the gaussian connectivity profile described in the first paragraph of this section. Like the previous two plasticity mechanisms, the acceleration of the timescale from biology is justified by the principle of separation of timescales.

Analysis.

As the network runs, new recurrent excitatory synapses are allowed to grow and, if their strengths are driven to zero, be pruned. The network first enters a growth phase, which lasts 100-200 seconds of simulation time, and then a stable phase in which the growth rate balances the pruning rate, as in Figure 1. The growth rate has been hand-selected so that the connection fraction in the stable phase corresponds with experimental observations for small slices (0.1, [6,17]). The network is allowed to run for 500 seconds and the state of the excitatory connectivity and the dynamics of the connection changes during the final epoch are then examined.

We first examine, alongside the growth of the network, the prevalence of bidirectional connections as compared to chance, a phenomenon noted to be significantly above-chance in [4] and [6], as shown in Figure 1. We observe for the total connection fraction a reliable value of 0.1, as selected. We observe a stable phase value of 0.021 for the bidirectional connection fraction, translating to a factor of 2.05 above chance. Our control for chance is the expected number of bidirectional connections for an Erdős-Rényi graph containing the same number of nodes and edges as the simulated network. For comparison purposes, a value of approximately 4 times chance is observed in [6]. We note that an otherwise equivalent non-topological network, in which the probability of connection between neurons is uniform rather than distance-dependent, produces a slightly below chance overrepresentation of bidirectional connections, reinforcing the well-known expectation that classical STPD, in the absence of other factors, favors unidirectional connectivity.

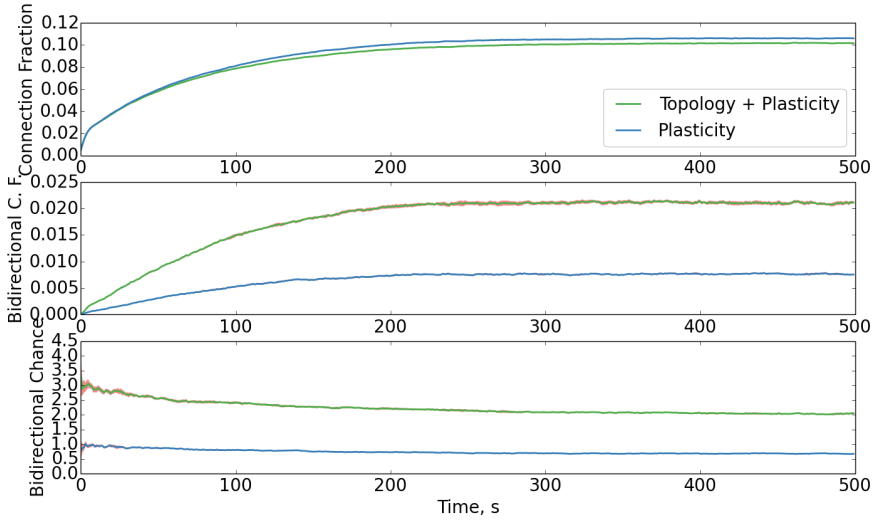

Figure 1. Evolution of total and bidirectional connection fraction with simulation time. Connection fraction evolution for plastic networks with and without topology. (a) Growth and subsequent stabilization of the connection fraction of the network with simulation time. (b) Growth of the bidirectional connection fraction. (c) Evolution of the bidirectional connection fraction with time as it relates to chance level. Data averaged over ten trials; standard deviation in red.

Regarding the growth of the network and the stabilization of its activity, we note two additional things. Firstly, in Figure 2, we note that the distribution of interspike intervals (ISIs) and their coefficients of variation (CVs) follow the properties of an approximately Poisson-like spiking with an effective refractory period, as is observed in cortical circuits. That is to say, the distribution of ISIs follows an exponential decay with a distortion, induced by the refractory period, at the low end, and that the CVs of the ISIs tend to be close to one. Secondly, we note the reproduction of some of the results of [16], specifically that during network growth there is a tendency for larger synaptic weights to be more likely to shrink than smaller synaptic weights, as seen in Figure 3.

We next observe the distribution of synaptic weights via histogramming, as previously stated, in Figure 4. This is in qualitative agreement with the heavy-tailed, log-normal-like shape typically observed in experimental data [6–10]. We believe that this is primarily due multiplicative-like behavior under the interaction between the (additive) spike timing dependent plasticity and the synaptic normalization while network activity is

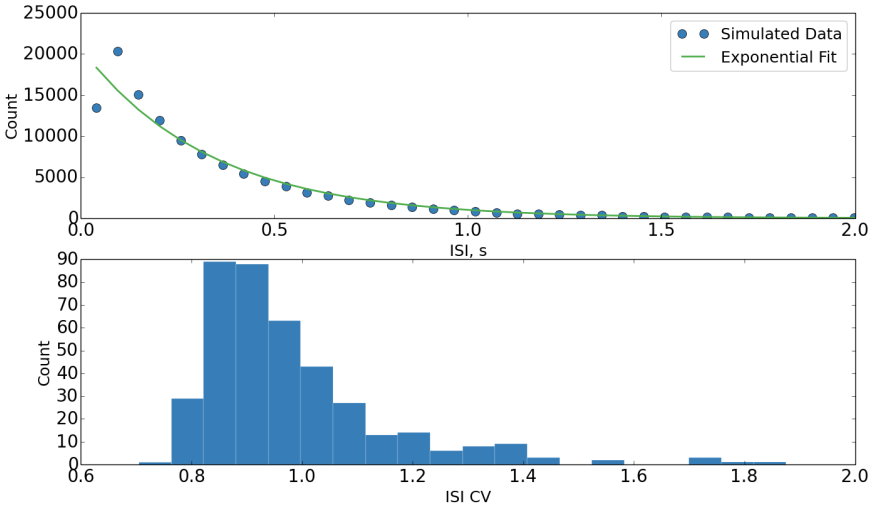

Figure 2. Distributions of ISIs and CVs thereof during stabilized network activity. (a) Pooled distribution of ISIs with exponential fit, suggesting Poisson-like behavior with a refractory period. Individual neuron distributions have been tested to be similar. (b) Distribution of CVs of ISIs, suggesting Poisson-like behavior. Single trial data.

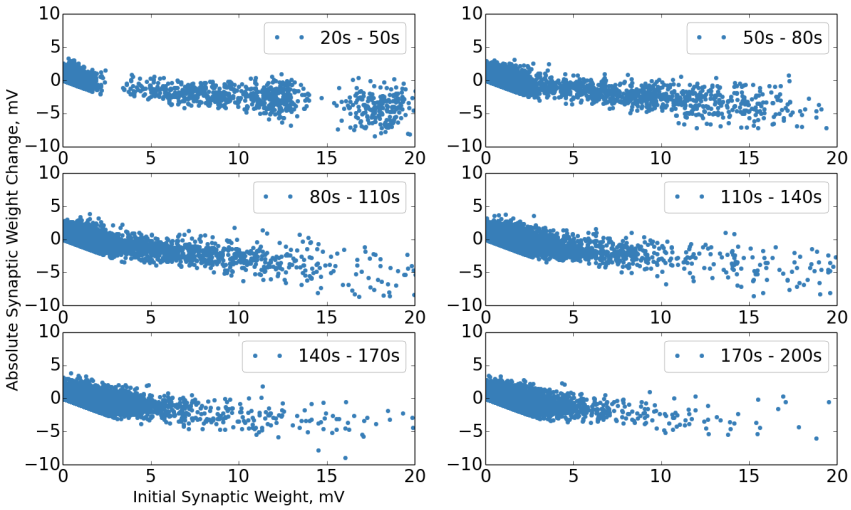

Figure 3. Synaptic change dynamics during network growth. Synaptic change dynamics during network growth epochs, before stabilization. "Bunching" in earliest epoch is an artifact of normalization under a small number of synapses. Single trial data.

stabilized by the intrinsic plasticity. We note that the topology of the network seems to have a minimal effect on this result, as would be expected from the results of [15].

We observe next the synaptic change dynamics in the stable phase of the network. We follow the format used in [10], comparing initial synaptic weight during a test epoch to both absolute and relative changes in synaptic weight, and demonstrate in Figure 5 that strong synaptic weights are less likely to fluctuate over time, as experimentally observed [10]. Additionally, this serves to reinforce the earlier success of [15] in modeling such synaptic dynamics as the result of self-organization, and demonstrates that such results carry over into a biologically more realistic model.

We examine, as well, the distribution of synaptic lifetimes. It has been predicted that the lifetimes of

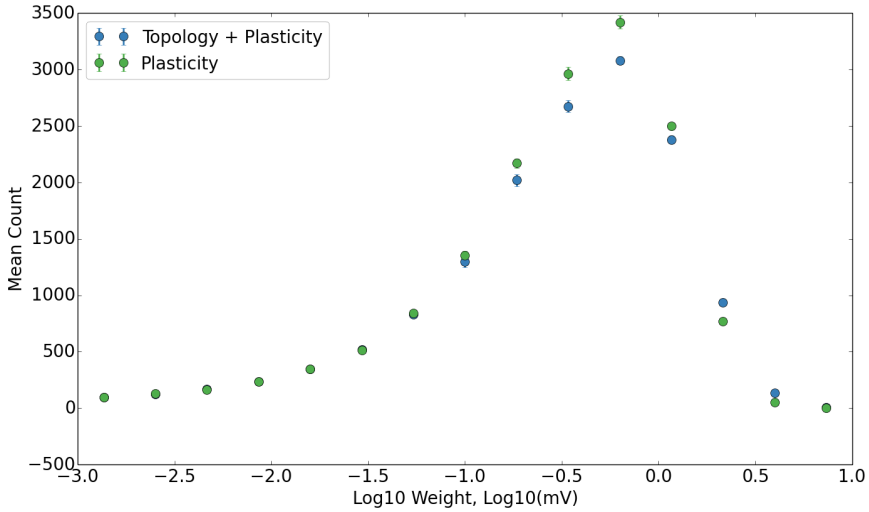

Figure 4. Log distribution of synaptic weights. The distribution of the base ten logarithm of synaptic weights for plastic networks with and without topology. Data averaged over ten trials; error bars are standard deviation.

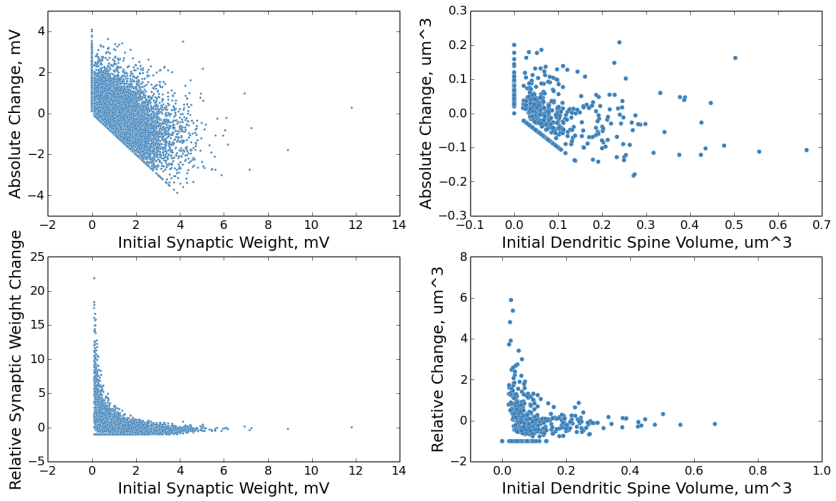

Figure 5. Change in synaptic weight as a function of initial synaptic weight. The above plots show the distributions of change in synaptic weight as a function of initial synaptic weight. The plots on the left are from the simulated network and are in electrophysiological units. The plots on the right are from experiment [10] and are in units of volume as determined by fluorescence data. The plots on the top show the absolute change in synaptic weight. The plots on the bottom show the relative change in synaptic weight. Single trial data.

fluctuating synapses may follow a power law distribution [15]; our model makes this prediction as well. However, we expand upon previous predictions with two highly notable observations. In its current form, our model produces a slope of approximately 5/3 in the stable phase. This decreases slightly in the growth phase. Secondly, we have observed as well that the slope can be modified by adjusting the balance of potentiation and depression in the STDP rule, varying between values between 1 and greater than 2, depending on the chosen parameters. For example, doubling the amplitude of the depression term in the STDP leads to a slope of approximately 5/2, while halving it leads to a slope of approximately 5/4. This is, in retrospect,

an intuitive phenomenon. A preponderance of potentiation will lead to synapses being depressed to a value below the pruning threshold less frequently, thereby lengthening the typical lifetime and decreasing the slope of the power law. Similarly, in a depression-dominated scenario, synapses will be driven below the pruning threshold more frequently, and, as a result, have shorter typical lifetimes, leading to a higher power law slope. Returning to the slight decrease in slope during the growth phase, this makes sense, as a reduction in the effective pruning rate is necessary for the network to continue to grow. We believe that with a more extensive investigation of the effects of other model parameters on the power law, the slope of this distribution could be used as a meaningful measure of the potentiation-depression balance in a recurrent cortical network.

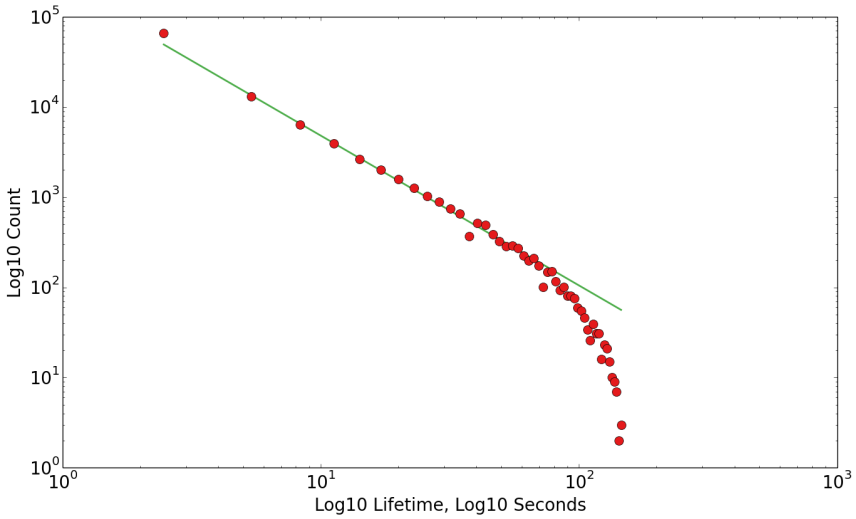

Figure 6. Distributions of synaptic lifetimes. The above plots show the distributions of of synaptic lifetimes during the stable (right) phases. Slope is approximately 5/3. The equivalent slope in the growth phase is slightly less. Here, we define entries in the growth phase as having synaptic end times of less than 150 seconds, and entries in the stable phase as having synaptic start times of greater than 350 seconds. Slopes are approximated via linear regression to the data points before the drop-off. Single trial data.

We subsequently examine the prevalence of triadic motifs in the graph of the simulated network. An overrepresentation of certain motifs was noted in [6]. We used a script written for the NetworkX Python module [33, 34] to acquire a motif count for the graph of the simulated network. As the overrepresentation of bidirectional connections will trivially lead to an overrepresentation of graph motifs containing bidirectional edges, the control for chance is, in this case, a modified Erdős-Rényi graph with the same number of nodes, same number of unidirectional edges, and same number of bidirectional edges as the graph of the simulated network, with the unidirectional and bidirectional edges being independently populated. A similar control is used in [6]. We observe a similar pattern of "closed loop" triadic motifs being overrepresented in Figure 7, as as experimentally observed in [6]. We note that the results for a non-topological plastic network with classical STDP, in the absence of additional factors, does not, relatively speaking, strongly select for any particular family of motifs. We similarly note that while distance-dependent topology does select for the observed family of motifs, it does not do so at the experimentally observed level. It is only the combination of topology and plasticity that strongly selected for the desired family of motifs.

Discussion

The problem of how the non-random micro-connectivity of the cortex arises is a nontrivial one with significant implications for the understanding of both cognition and development. We attempt, in this paper, to provide insight into this problem by presenting a plausible model by which such non-random connectivity arises as the self-organized result of the interaction of multiple plasticity mechanisms under physiological constraints. Some models attempt to describe elements of the graph structure of the micro-connectome in purely physiological

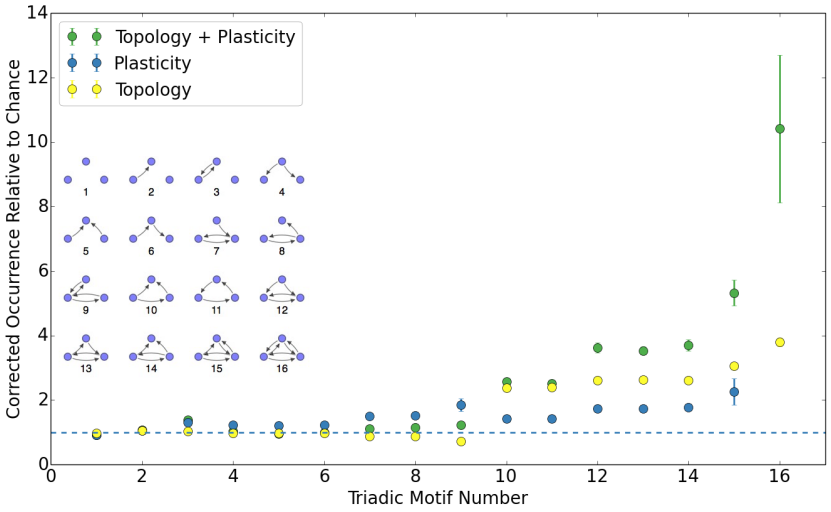

Figure 7. Triadic motif counts as a multiple of chance, corrected for bidirectional overrepresentation. Triadic motif counts (in the same order as [6]) for a simulated network as a multiple of chance value. The counts have been corrected for the observed overrepresentation of bidirectional connections. Results are shown for a complete network, a purely topological construction, and an equivalent network with no topology. For the latter, the count of motif 16 is out of range due to the extremely low expected count after bidirectionality corrections. Data averaged over ten trials; error bars are standard deviation.

and topological terms [35]. However, such models necessarily lack an active network, and are thus unable to simultaneously account for synaptic dynamics, as our model does. Our model is, of course, a simple model, but the degree to which it accounts for observed non-random features of the typical cortical microcircuit without detailed structural features, metabolic factors, or structured input to drive the plasticity in a particular fashion is highly suggestive in terms of what is necessary at a bare minimum to drive the development and maintenance of the complex microstructure of the brain. Additionally, as mentioned in the introduction, it is hypothesized that a small backbone of strong synapses may form the stable backbone of long-term memory. The fact that in our model, strong weights remain stable in the presence of ongoing plasticity and despite significant fluctuations of smaller weights (which has been modeled as a stochastic Kesten process [36]), and the naturalness with which such a dynamic arises out of the interactions of known plasticity mechanisms, is both suggestive and supportive of this theory.

An additional noted non-random feature of cortical recordings that has been passed over in this model is the observed log-normal distribution of cortical firing rates. Our intrinsic plasticity mechanism necessarily negates this feature, which may be self-organized via mechanisms not included in our model, such as diffusive homeostasis [37]. In order to maximize simplicity, a single target firing rate is chosen for all neurons. This also permits pooling of the ISIs for analysis. Additional testing in which the target firing rate is drawn from a log-normal distribution produces minimal qualitative effects on the observed features (except, trivially, the ISI distribution, and a slight increase to the peak overrepresentation, though with no change to the overall pattern (S2)). Another issue is that as things stand, the exact statistics of the micro-connectome are difficult to discern – though strong inferences can be made in the right direction – due to inherent sampling biases in paired patch-clamp reconstructions of limited size [38]. It is our hope and belief that advances in fluorescence imaging, automated electron microscopy reconstruction [39, 40], and massive multi-unit array recordings will help to alleviate these biases. One might imagine that additional biases may be caused by the relatively small model size of 400 excitatory neurons, when realistic cortical densities would result in thousands of neurons in such an equivalent volume. We have tested the network at much larger sizes and found no significant qualitative change to our observed results (the higher density does appear to lead to a higher peak motif overrepresentation, without changing the overall pattern of overrepresentation (Figure S1), so we maintained

a relatively small network size to increase computational ease.

Some might view the fact that, in this model, the primary driver behind the overrepresentation of bidirectional connections is topology, as a shortcoming. The authors do not view this as such; after all, topology exists in the cortex and the rest of the study's results suggest it is an important factor in the self-organization of cortical circuits. There are mechanisms utilizing non-classical STDP, such as the so-called triplet and voltage rules [41, 42], which, in the presence of high-frequency activity, are capable of producing and maintaining bidirectional connections. Introducing such mechanisms into a similar model would be a welcome and interesting future study, and could potentially lead to an even stronger and more precise motif selectivity.

Often, models of cortical microcircuits are described as random graphs, such as the classical random balanced network [43]. However, experiments have demonstrated that the structure of the cortical microcircuit is significantly non-random [5, 6], suggesting that random networks are not sufficient platforms for modeling cortical development and activity. Lacking in structural plasticity of topology, such random graph based balanced networks are incapable of producing the sort of results we have observed. Having provided a mechanism with which one may generate a cortex-like non-random structure, it would be enlightening to determine if said structure provides any significant computational or metabolic advantage as compared to a random graph. Similarly, limitations in online plasticity capabilities significantly hinder the use of such random networks and their relatives in reservoir computing [44] for unsupervised learning and inference tasks, while earlier studies with the original SORN model [13, 14] suggest that the particular combination of plasticity mechanisms in our model can endow networks with impressive learning and inference capabilities. It is additionally our desire to further study learning and inference in such networks with more biologically complete models. It is the belief of the authors that the future of neural network-based computation and modeling of biological processes lies in the incorporation of multiple plasticity and homeostatic mechanisms under simple sets of constraints.

Acknowledgments

We would like to thank the authors of [10] for sharing their data. We would like as well to thank Christoph Hartmann and Pengsheng Zheng for their valuable consultation.

Supporting Information

Supporting Figures

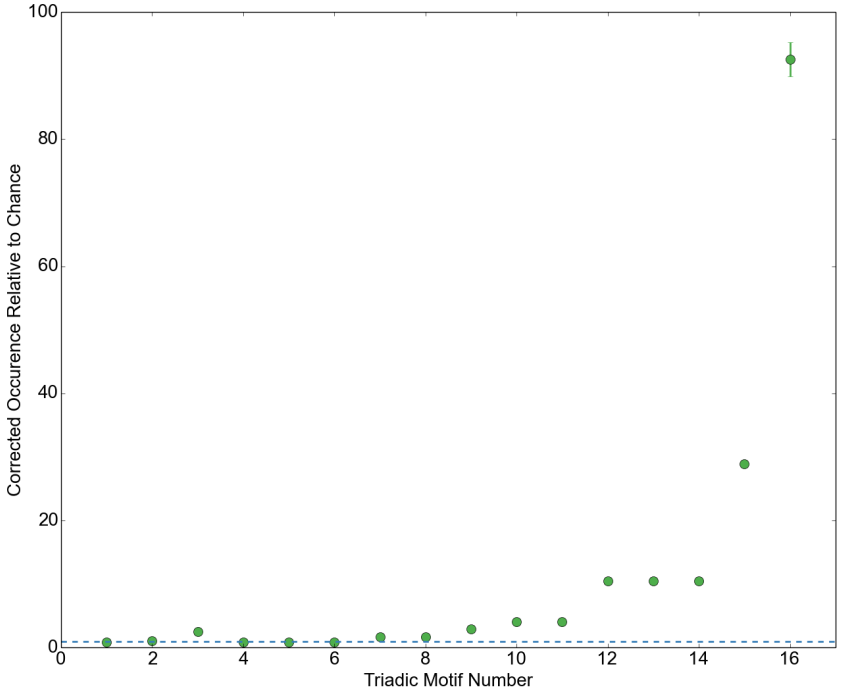

Figure S1. Triadic motif counts as a multiple of chance, corrected for bidirectional overrepresentation. Triadic motif counts (in the same order as [6]) for a simulated network as a multiple of chance value. The counts have been corrected for the observed overrepresentation of bidirectional connections. Results are shown for a complete network of 2000 neurons. Other parameters remain the same, aside from scaling of IP target to maintain mean per-neuron input spike rate and growth rate to obtain stable phase connection fraction of 0.1. Error bars are standard deviation.

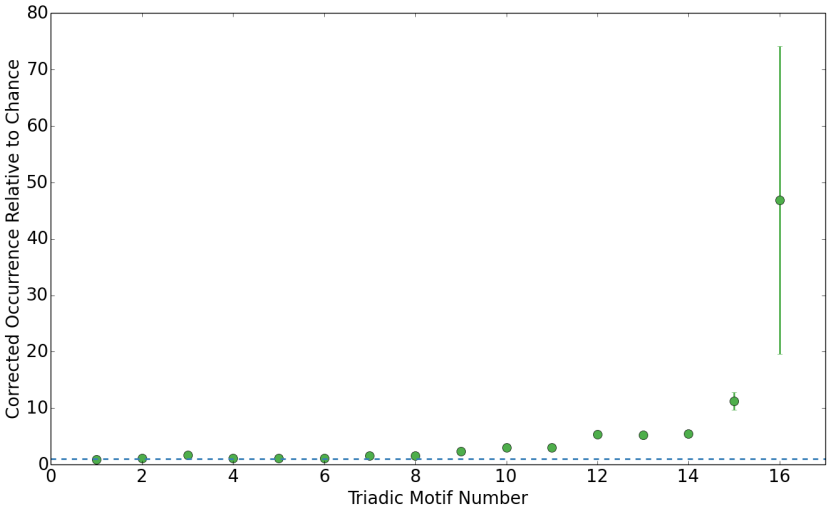

Figure S2. Triadic motif counts as a multiple of chance, corrected for bidirectional overrepresentation. Triadic motif counts (in the same order as [6]) for a simulated network as a multiple of chance value. The counts have been corrected for the observed overrepresentation of bidirectional connections. Results are shown for a complete network with IP target rates drawn from a log-normal distribution instead of a single value. Other parameters remain the same, aside from scaling of growth rate to obtain stable phase connection fraction of 0.1. Error bars are standard deviation.

References

1. Gilson M, Fukai T (2011) Stability versus neuronal specialization for STDP: long-tail weight distributions solve the dilemma. PloS one 6: e25339.
2. Grutzendler J, Kasthuri N, Gan W (2002) Long-term dendritic spine stability in the adult cortex. Nature 420.
3. Trachtenberg JT, Chen BE, Knott GW, Feng G, Sanes JR, et al. (2002) Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. Nature 420: 788–94.
4. Markram H (1997) A network of tufted layer 5 pyramidal neurons. Cerebral cortex (New York, NY : 1991) 7: 523–33.
5. Perin R, Berger TK, Markram H (2011) A synaptic organizing principle for cortical neuronal groups. Proceedings of the National Academy of Sciences of the United States of America 108: 5419–24.
6. Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB (2005) Highly nonrandom features of synaptic connectivity in local cortical circuits. PLoS biology 3: e68.
7. Harris KM, Stevens JK (1989) Dendritic spines of CA 1 pyramidal cells in the rat hippocampus: serial electron microscopy with reference to their biophysical characteristics. The Journal of neuroscience : the official journal of the Society for Neuroscience 9: 2982–97.
8. Lisman JE, Harris KM (1993) Quantal analysis and synaptic anatomy–integrating two views of hippocampal plasticity. Trends in neurosciences 16: 141–7.
9. Thomson AM, Deuchars J, West DC (1993) Large, deep layer pyramid-pyramid single axon EPSPs in slices of rat motor cortex display paired pulse and frequency-dependent depression, mediated presynaptically and self-facilitation, mediated postsynaptically. Journal of neurophysiology 70: 2354–69.
10. Yasumatsu N, Matsuzaki M, Miyazaki T, Noguchi J, Kasai H (2008) Principles of long-term dynamics of dendritic spines. The Journal of neuroscience : the official journal of the Society for Neuroscience 28: 13592–608.
11. Tetzlaff C, Kolodziejski C, Timme M, Wörgötter F (2011) Synaptic scaling in combination with many generic plasticity mechanisms stabilizes circuit connectivity. Frontiers in Computational Neuroscience 5: 1–15.
12. Tetzlaff C, Kolodziejski C, Timme M, Wörgötter F (2012) Analysis of Synaptic Scaling in Combination with Hebbian Plasticity in Several Simple Networks. Frontiers in Computational Neuroscience 6: 1–17.
13. Lazar A, Pipa G, Triesch J (2009) SORN: a self-organizing recurrent neural network. Frontiers in computational neuroscience 3: 23.
14. Lazar A, Pipa G, Triesch J (2011) Emerging Bayesian priors in a self-organizing recurrent network. Neural Networks and Machine Learning : 1–8.
15. Zheng P, Dimitrakakis C, Triesch J (2013) Network self-organization explains the statistics and dynamics of synaptic connection strengths in cortex. PLoS computational biology 9: e1002848.
16. Minerbi A, Kahana R, Goldfeld L, Kaufman M, Marom S, et al. (2009) Long-term relationships between synaptic tenacity, synaptic remodeling, and network activity. PLoS biology 7: e1000136.
17. Thomson AM, West DC, Wang Y, Bannister aP (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. Cerebral cortex (New York, NY : 1991) 12: 936–53.
18. Goodman D, Brette R (2008) Brian: a simulator for spiking neural networks in python. Frontiers in neuroinformatics 2: 5.

19. Bi GQ, Poo MM (1998) Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. The Journal of neuroscience : the official journal of the Society for Neuroscience 18: 10464–72.

20. Froemke R, Poo Mm, Dan Y (2005) Spike-timing-dependent synaptic plasticity depends on dendritic location. Nature 2033: 2032–2033.

21. Gerstner, W, Kempter, R, van Hemmen, J L, & Wagner H (1996) A neuronal learning rule for sub-millisecond temporal coding. Nature 383: 76–78.

22. Kempter R, Gerstner W, van Hemmen J (1999) Hebbian learning and spiking neurons. Physical Review E 59: 4498–4514.

23. Song S, Miller KD, Abbott LF (2000) Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. Nature neuroscience 3: 919–26.

24. Zhang LI, Tao HW, Holt CE, Harris WA, Poo Mm (1998) A critical window for cooperation and competition among developing retinotectal synapses. Nature 395: 37–44.

25. Abbott LF, Nelson SB (2000) Synaptic plasticity: taming the beast. Nature neuroscience 3 Suppl: 1178–83.

26. Bourne J, Harris K (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: 354–73.

27. Ibata K, Sun Q, Turrigiano GG (2008) Rapid synaptic scaling induced by changes in postsynaptic firing. Neuron 57: 819–26.

28. Turrigiano GG, Leslie KR, Desai NS, Rutherford LC, Nelson SB (1998) Activity-dependent scaling of quantal amplitude in neocortical neurons. Nature 391: 892–6.

29. Hill A (1936) Excitation and accommodation in nerve. Proceedings of the Royal Society of London Series B - Biological Sciences 119: 305–355.

30. Benda J, Herz AVM (2003) A universal model for spike-frequency adaptation. Neural computation 15: 2523–64.

31. Desai NS, Rutherford LC, Turrigiano GG (1999) Plasticity in the intrinsic excitability of cortical pyramidal neurons. Nature neuroscience 2: 515–20.

32. Zhang W, Linden DJ (2003) The other side of the engram: experience-driven changes in neuronal intrinsic excitability. Nature reviews Neuroscience 4: 885–900.

33. Hagberg AA, Schult DA, Swart PJ (2008) Exploring network structure, dynamics, and function using NetworkX. In: Proceedings of the 7th Python in Science Conference (SciPy2008). Pasadena, CA USA, pp. 11–15.

34. Levenson A, van Liere D (2011). triadic census. `https://networkx.lanl.gov/trac/ticket/190`.

35. Aćimović J, Mäki-Marttunen T, Linne ML (2015) The effects of neuron morphology on graph theoretic measures of network connectivity: the analysis of a two-level statistical model. Frontiers in Neuroanatomy 9.

36. Statman A, Kaufman M, Minerbi A, Ziv NE, Brenner N (2014) Synaptic Size Dynamics as an Effectively Stochastic Process. PLoS computational biology 10: e1003846.

37. Sweeney Y, Hellgren Kotaleski J, Hennig MH (2015) A Diffusive Homeostatic Signal Maintains Neural Heterogeneity and Responsiveness in Cortical Networks. PLOS Computational Biology 11: e1004389.

38. Miner DC, Triesch J (2014) Slicing, sampling, and distance-dependent effects affect network measures in simulated cortical circuit structures. Frontiers in Neuroanatomy 8: 1–9.

39. Chklovskii DB, Vitaladevuni S, Scheffer LK (2010) Semi-automated reconstruction of neural circuits using electron microscopy. Current Opinion in Neurobiology 20: 667–675.

40. Plaza SM, Scheffer LK, Chklovskii DB (2014) Toward large-scale connectome reconstructions. Current Opinion in Neurobiology 25: 201–210.

41. Pfister JP, Gerstner W (2006) Triplets of spikes in a model of spike timing-dependent plasticity. The Journal of neuroscience : the official journal of the Society for Neuroscience 26: 9673–82.

42. Clopath C, Büsing L, Vasilaki E, Gerstner W (2010) Connectivity reflects coding: a model of voltage-based STDP with homeostasis. Nature neuroscience 13: 344–52.

43. van Vreeswijk C, Sompolinsky H (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science (New York, NY) 274: 1724–6.

44. Lukoševičius M, Jaeger H (2009) Reservoir computing approaches to recurrent neural network training. Computer Science Review 3: 127–149.