

Developmental Self-Construction and -Configuration of Functional Neocortical Neuronal Networks

Roman Bauer^{1,2*}, Frédéric Zubler^{1,3}, Sabina Pfister¹, Andreas Hauri¹, Michael Pfeiffer¹, Dylan R. Muir^{1,4}, Rodney J. Douglas¹

1 Institute of Neuroinformatics, University/ETH Zürich, Zürich, Switzerland, **2** School of Computing Science, Newcastle University, Newcastle upon Tyne, United Kingdom, **3** Department of Neurology, Inselspital Bern, Bern University Hospital, University of Bern, Bern, Switzerland, **4** Biozentrum, University of Basel, Basel, Switzerland



Abstract

The prenatal development of neural circuits must provide sufficient configuration to support at least a set of core postnatal behaviors. Although knowledge of various genetic and cellular aspects of development is accumulating rapidly, there is less systematic understanding of how these various processes play together in order to construct such functional networks. Here we make some steps toward such understanding by demonstrating through detailed simulations how a competitive co-operative ('winner-take-all', WTA) network architecture can arise by development from a single precursor cell. This precursor is granted a simplified gene regulatory network that directs cell mitosis, differentiation, migration, neurite outgrowth and synaptogenesis. Once initial axonal connection patterns are established, their synaptic weights undergo homeostatic unsupervised learning that is shaped by wave-like input patterns. We demonstrate how this autonomous genetically directed developmental sequence can give rise to self-calibrated WTA networks, and compare our simulation results with biological data.

Citation: Bauer R, Zubler F, Pfister S, Hauri A, Pfeiffer M, et al. (2014) Developmental Self-Construction and -Configuration of Functional Neocortical Neuronal Networks. *PLoS Comput Biol* 10(12): e1003994. doi:10.1371/journal.pcbi.1003994

Editor: Olaf Sporns, Indiana University, United States of America

Received: July 24, 2014; **Accepted:** October 9, 2014; **Published:** December 4, 2014

Copyright: © 2014 Bauer et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability: The authors confirm that all data underlying the findings are fully available without restriction. All relevant data are within the paper and its Supporting Information files.

Funding: This work was supported by the EU grant 216593 SECO (<http://www.seco-project.eu>). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

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

* Email: roman.bauer@newcastle.ac.uk

Introduction

In this paper we address the question of how progenitor cells of the neocortical subplate can give rise to large functional neuronal sub-networks in the developed cortex. We choose winner-take-all (WTA) [1,2] connectivity as the target of this self-construction and -configuration process because these sub-networks are consistent with the observed physiology [3,4] and connectivity [5,6] of neurons in the superficial layers of neocortex, and because they are powerful elements of computation [7,8]. WTA networks actively select the strongest of multiple input signals, while suppressing the weaker ones. This fundamental characteristic is applicable in various contexts, and so many studies modeling cortical function are based on WTA modules [8–15].

The idealized WTA network architecture is shown in Fig. 1A. Excitatory neurons are recurrently connected to each other and also with one or more inhibitory neurons, which project back to the excitatory neurons. This architecture does not in itself guarantee WTA functionality. The degree of recurrent excitation, excitation of inhibitory neurons, and inhibition of excitatory neurons need all to lie within preferred ranges [8] in order for the network to exhibit effective WTA behavior. The appropriate neural architecture must be grown, and then the weights of the many synapses must be tuned to fall within the necessary ranges. Such neuronal growth and synapse formation are subject to variability (1B,C), for which the homeostatic learning mechanisms must compensate.

The behavior of a WTA network depends on the ratios of the effects of its various excitatory and inhibitory connection paths. In its high excitatory gain regime a WTA network will report only the strongest of its feed-forward inputs, and suppress the remainder of the excitatory neurons, which are weakly activated. In a more relaxed regime (soft-WTA, sWTA) the network will return a pattern of winners that best conforms to its input. In this sense the sWTA performs a pattern based signal restoration, which is a crucial mechanism for resisting degradation of processing in neural systems across their many computational steps. In this paper we choose to have the developmental process grow and tune these sWTA networks.

Our goal is to demonstrate how plausible genetic developmental mechanisms can combine with homeostatic synaptic tuning to bring networks of neurons into sWTA functionality (Fig. 1). Our demonstration is based on simulations of the development and growth of neural tissue in 3D physical space using *Cx3D* [16]. The simulation begins with a single precursor cell. This cell encodes gene-like instructions that are sequentially and conditionally expressed through a gene regulatory network (GRN). By controlling the expression of different genes, this GRN gives rise to pools of differentiated excitatory and inhibitory neurons. These neurons, which are placed randomly in 3D space, extend axons and dendrites and make synapses according to a proximity rule. This process results in a synaptically connected network that matches well experimentally obtained connectivity statistics.

Author Summary

Models of learning in artificial neural networks generally assume that the neurons and approximate network are given, and then learning tunes the synaptic weights. By contrast, we address the question of how an entire functional neuronal network containing many differentiated neurons and connections can develop from only a single progenitor cell. We chose a winner-take-all network as the developmental target, because it is a computationally powerful circuit, and a candidate motif of neocortical networks. The key aspect of this challenge is that the developmental mechanisms must be locally autonomous as in Biology: They cannot depend on global knowledge or supervision. We have explored this developmental process by simulating in physical detail the fundamental biological behaviors, such as cell proliferation, neurite growth and synapse formation that give rise to the structural connectivity observed in the superficial layers of the neocortex. These differentiated, approximately connected neurons then adapt their synaptic weights homeostatically to obtain a uniform electrical signaling activity before going on to organize themselves according to the fundamental correlations embedded in a noisy wave-like input signal. In this way the precursor expands itself through development and unsupervised learning into winner-take-all functionality and orientation selectivity in a biologically plausible manner.

During this neurite outgrowth, the synaptic weights calibrate themselves homeostatically using experimentally established synaptic scaling [17] and BCM learning rules [18]. This synaptic learning is conditioned by coarsely patterned neuronal activity similar to that of retinal waves or cortico-thalamic loops [19–23].

We compare these grown networks with biological data, and demonstrate WTA functionality. This comparison is done also in the context of cortical functionality, such as orientation selectivity. Importantly, the overall behavior stems solely from local processes, which are instantiated from internally encoded and developmental primitives [24]. Hence, we provide a model that explains the developmental self-construction and -configuration of a neocortical WTA network in a biologically plausible way.

Results

Development of Differentiated Neurons Based on a Gene-Regulatory-Network

Cell proliferation and differentiation into different cell types is specified implicitly in the genetic code of a single precursor cell. This code determines how a given number of excitatory and inhibitory neurons is produced. During the unfolding process of this code, each cell contains the same genetic code, but because of its local environment can follow different developmental trajectories.

We model the molecular mechanisms that regulate cell differentiation by a dynamical gene regulatory network (GRN). This GRN is defined by a set of 5 variables (G_0 , G_1 , G_2 , GE , and GI) that represent substance concentrations, where each substance is the expression of a gene. Importantly, all cells have their own instantiations of these variables. The secretion, interaction, and decay of substances, is regulated by the laws of kinetics. The differential equations specifying these dynamics are shown in Methods.

During the evolution of the substance concentrations, also cell growth and division is simulated. The cell cycle time and model

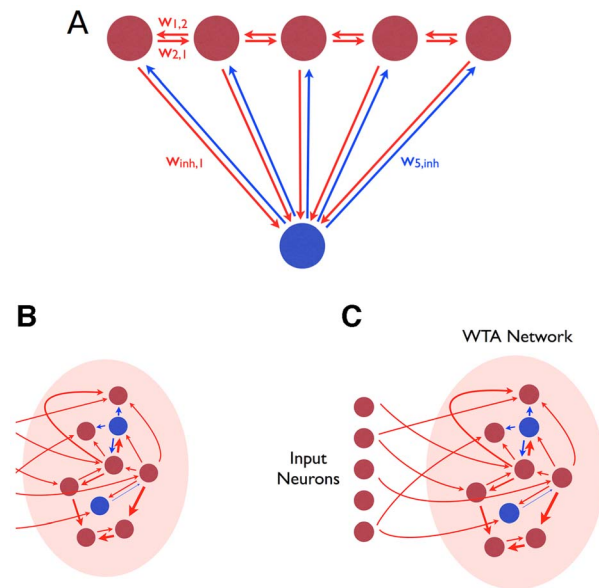


Figure 1. Winner-take-all architecture. (A) Architecture of an idealised winner-take-all-network. Several excitatory neurons (red) excite a single shared inhibitory neuron, or a shared population of inhibitory neurons (blue). Each excitatory neuron receives inhibitory feedback in proportion to the average activity of the excitatory population. (B) The WTA architecture is embedded in the field of recurrent connections between a population of excitatory and inhibitory neurons. (C) Once the WTA architecture has formed, coarsely structured synaptic input drives synaptic refinement of the recurrent connections within the network. doi:10.1371/journal.pcbi.1003994.g001

parameters of the differential equations are fixed and independent of the substance dynamics.

Initially, all concentrations are set to zero. At this stage, only the “starter” substance G_0 is produced, which reaches high concentration levels in the first time step, and triggers the production of a second gene G_1 . G_1 is produced according to a prespecified intrinsic production constant α . This value determines how many cell divisions will occur until the concentration of G_1 reaches a value of 0.99. When this threshold is reached, a probabilistic decision is induced: GE or GI , responsible for activating the excitatory and inhibitory cell phenotypes, are triggered with probability $p_E = 0.8$ or $p_I = 0.2$, respectively.

Such a GRN network configuration would enable us to generate 2^n cells, where n is the number of symmetric divisions. However, the target number of cells might not be an exponential of 2. Therefore, we have introduced a second gene G_2 that is (probabilistically) activated by high concentrations of G_1 , and that leads to a second round of symmetric division. As for G_1 , G_2 activates GE or GI in a probabilistic manner. The probability to enter into this secondary cell cycle is given by p_2 , which is computed based on the target number of cells. The evolution of the GRN across cell types is depicted in Fig. 2.

By setting the production rate constant α of gene G_1 and the probabilistic activation of G_2 , we can control the final number of cells produced. The equations for computing the probabilities for either differentiating into neurons by G_1 induction (p_1) or by G_2 induction (p_2), depending on the target number of cells, are shown in Methods.

Overall, the GRN is designed so that a desired total number of cells is reached, and that the distribution of excitatory *vs.* inhibitory cells follows the approximate 4:1 ratio observed in




Figure 2. Gene Regulatory Network. (A) Schematic representation of the GRN, composed of five interacting genes that give rise to excitatory and inhibitory neurons. The identity of a neuron is determined by the genes GE and GI for excitatory or inhibitory neurons, respectively. Arrows indicate a positive effect on gene expression. (B) Lineage tree. Nodes indicate cells; boxes indicate gene expression patterns. G0 triggers the expression of G1, which characterizes the undifferentiated state of progenitor cells. After a series of symmetric divisions, G1 reaches a concentration threshold. According to fixed probabilities, G1 can then activate the differentiation toward excitatory (red) or inhibitory (blue) neurons. Alternatively, a small proportion of cells probabilistically undergoes a second round of cell division and activates gene G2, which again promotes the differentiation toward excitatory or inhibitory neurons by the expression of GE or GI. The probabilistic activation of inhibitory or excitatory genes is a simplification, but guarantees the production of a homogeneously mixed population of neurons.

doi:10.1371/journal.pcbi.1003994.g002

cortex [25–27] (S1 Figure). Fig. 3(A–D) shows the evolution of an initial cell giving rise to a number of cells which eventually grow out neurites based solely on their genetic encoding.

Neurite Growth and Synapse Formation

Neurite growth and arborization is caused by growth cone traction and bifurcation. The growth cone is able to sense the presence and gradient of morphogens and other signal molecules, and also able to actively explore the local extracellular space. Importantly, neurite growth is steered via a growth cone model instantiated at the tip of the axon or dendrite, and so is a local process.

Diffusible signal molecules are secreted by the cell somata. In these simulations excitatory and inhibitory neurons secrete two characteristic signals, that enable excitatory and inhibitory axons to find inhibitory and excitatory neurons, respectively. The axonal growth cones initially grow out of the somata in random directions. However, they retract whenever the concentration they sense falls below a threshold. The retraction stops and growth recommences when a second higher threshold is exceeded. In this way the axons remain close to substance secreting sources. Retraction is an efficient strategy for establishing connections because axons grow only into regions containing a potential target, and is commonly observed in developing neurons [28–31]. A video of a developing neural network with axonal retraction (simulated in Cx3D) is included in the Supporting Information (S1 Video) and on Youtube (<http://www.youtube.com/watch?v=il2uc-ZUZQ4>).

Axons deploy boutons. Whenever these boutons are sufficiently close to a potential post-synaptic site on a dendrite a synapse is created between them. Consequently, the final synaptic network connectivity depends on the nearly stochastic arrangement of regions of spatial proximity of the outgrowing axons and dendrites.

We adapted the parameters of the neurite outgrowth (see Table 1) so that the connectivity of the simulated neuronal growth matched our experimental observations in layers II/III of cat visual cortex [5,32] (see Fig. 4A). Overall, we found that connectivity was robust to reasonable variation of the growth

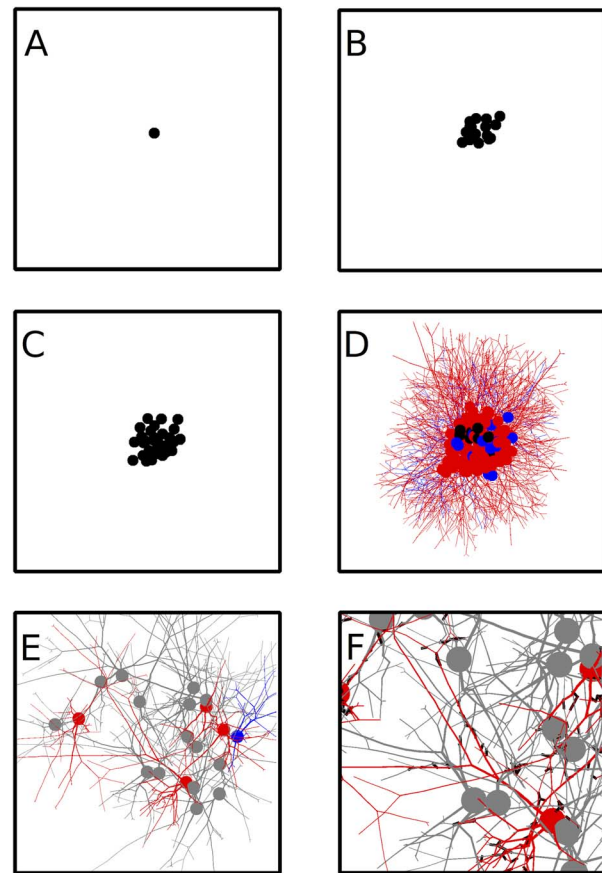


Figure 3. Developmental process for building a competitive network. A single precursor cell (A) contains the genetic code specifying the entire developmental process. (B) The precursor cell first undergoes repeated division to increase the pool of neuronal precursors (black). (C) Precursor neurons then differentiate into excitatory and inhibitory cell classes. (D) Neurite outgrowth begins to provide a scaffold for synaptic connections. (E) A network of differentiated neurons (grey) after neurite outgrowth has finished. For better visualization, examples of excitatory and inhibitory neurons are colored in red and blue, respectively. (F) Synapses (black rectangles) can form at appositions between axons and dendrites.

doi:10.1371/journal.pcbi.1003994.g003

parameters and the random location of somata. The absolute numbers of synapses simulated here are smaller than observed in biology, due to constraints on computational resources. However, there is no inherent restriction on scalability using our methods, and so we expect that realistic numbers of cells and synapses could if necessary be simulated using supercomputers.

Fig. 4B shows the distribution of the percentage of excitatory input synapses to the neurons, across the whole population. The average percentage of excitatory inputs to a neuron in this network is 84%, which is in good agreement with the experimental data. This result is consistent with observations across species and cortical areas that some 15% of all the synapses are GABAergic [5,33–35], irrespective of neuronal densities. Importantly, this good agreement arises naturally out of the growth model, and did not require extensive tuning of the model parameters.

Electrophysiology

The self-configuration of electrophysiological processing depends on the tuning of network synaptic weights and neuronal activity. In order to simulate this aspect of the developing

Table 1. Parameters for simulating axonal and dendritic growth.

Parameter	Value (Ex. Axon/Inh. Axon/Ex. Dendrite/Inh. Dendrite)
d_{minimal} (minimal diameter)	0.2/0.2/0.3/0.3
r_{move} (diameter reduction when moving)	0.004/0.012/0.02/0.042
r_{fork} (diameter reduction when bifurcating)	0.12/0.105/0.14/0.12
p_{base} (baseline probability for bifurcation)	0.05/0.08/0.04/0.05
$p_{\text{substance}}$ (substance dependent probability for bifurcation)	0.005/0.05/0.0/0.0
speed of growth	100/100/100/100
speed of retraction	5/5/(no retraction)/(no retraction)
θ_1 (concentration threshold triggering retraction)	1e-8/1e-8/(no retraction)/(no retraction)
θ_2 (concentration threshold stopping retraction)	0.036/0.036/(no retraction)/(no retraction)
w_{prev} (weight of previous growth direction)	0.75/0.75/0.75/0.75
w_{noise} (weight of random direction)	0.25/0.25/0.25/0.25
s_{max} (neurite discretization size)	7/7/7/7

Growth parameters are dependent on the type of the neurite, as well as the neuron type. In order to qualitatively match biological observations, we modeled axons to be longer than dendrites and inhibitory (basket) cells to have smaller spatial extent than excitatory neurons [56,132–134]. In our model, axons direct their growth based on extracellular substance concentrations.

doi:10.1371/journal.pcbi.1003994.t001

networks, we must model also the electrical activity of neurons. However, the time scales of morphological growth and electrophysiological dynamics are many orders of magnitude different, and this difference makes for substantial technical problems in simulation.

For simplicity, and for minimizing computational demands we have used a rate-based approach to modeling neuronal activity. We approximate the neuronal activation by a linear-threshold function [36] that describes the output action potential discharge rate of the neuron as a function of its input. This type of neuronal activation function is a good approximation to experimental observations of the adapted current discharge relation of neurons [37,38] and has been used in a wide range of modeling works [8,39–41].

The linear-threshold activation function is:

$$\tau \frac{dx_i}{dt} = -x_i + \max(s_i + I_i + \sum_j w_{ij}x_j - T, 0) \quad (1)$$

where x_i denotes the firing rate of a neuron with index i , τ is the neuronal time constant, s_i is the spontaneous activity, I_i is the feed-forward input to neuron i , w_{ij} is the weight of the connection from neuron j to neuron i (can be positive or negative, depending on the presynaptic neuron's type), and T is the neuron's threshold. For simplicity, τ and T are set to 1 and 0. Exploratory simulations where $\tau_{\text{inh}} \neq \tau_{\text{ex}}$ yielded very similar results.

For computational efficiency, the electrophysiology simulator is implemented as a global process that acts on the total weight matrix of the neuronal network, rather than performing these frequent computations locally. We chose this global methodology because it leads to a significant speed-up compared with a local version that had been used initially. The total weight matrix is obtained by summation of the weights of all synapses in the Cx3D simulation. Using these connection weights, neuronal activity is computed as described in Eq. 1. Connection weight changes resulting from the learning and adaptation (explained below) are computed based on this summed weight matrix and the activities of the two respective connected neurons, which are saved at each electrophysiology time step.

The same connection weights (and neuronal activities) would be computed if only local processes at the synapses were simulated, because the synaptic learning and adaptation dynamics (Eq. 2 and 3) are dependent on the (locally available) neuronal activities, and linearly dependent on the synaptic weight. Hence, the dynamics of the summed synaptic weights match the sum of the individual synapse weight changes.

For reasons of biological plausibility, the electrophysiology simulator incorporates a maximum connection weight. This maximum weight for the functional connection strength between two neurons is determined by counting the number of synapses involved. This number, multiplied by the maximal weight of a single synapse, is defined as the maximum of the total connection weight. Hence, neurons that are connected by few synapses can not establish a strong functional link.

In our model, self-configuration of the weights towards sWTA functionality occurs during sequential developmental phases. Sequential phases of electrical adaptation and learning during development have been observed experimentally [42,43], and have also been applied in previous models [44,45].

During the first, homeostatic phase neurons adapt the synaptic weights of their own input in order to maintain a target output activity. The effect of this phase is to bring the neuronal firing rates into a balanced regime, and so allow for a reliable synaptic learning without interference by unresponsive neurons or runaway excitation. During the second, specification phase the neurons structure their individual responses by correlation-based learning on their inputs.

Homeostatic phase. During this first phase of activity-dependent adaptation, neurite outgrowth, synapse formation and homeostatic adaptation of neuronal activity occur simultaneously. Neurons implement the synaptic scaling rule [17,46,47], whereby they scale their synaptic input weights to achieve a preferred average output firing rate. Thus, when their average output activity exceeds a given target, neurons scale their excitatory and inhibitory inputs down and up respectively. The opposite effect occurs when the average activity has fallen below the target. Since there is no correlation-based learning during this phase the population of neurons can converge towards stable average levels of activity, but there is no input learning.

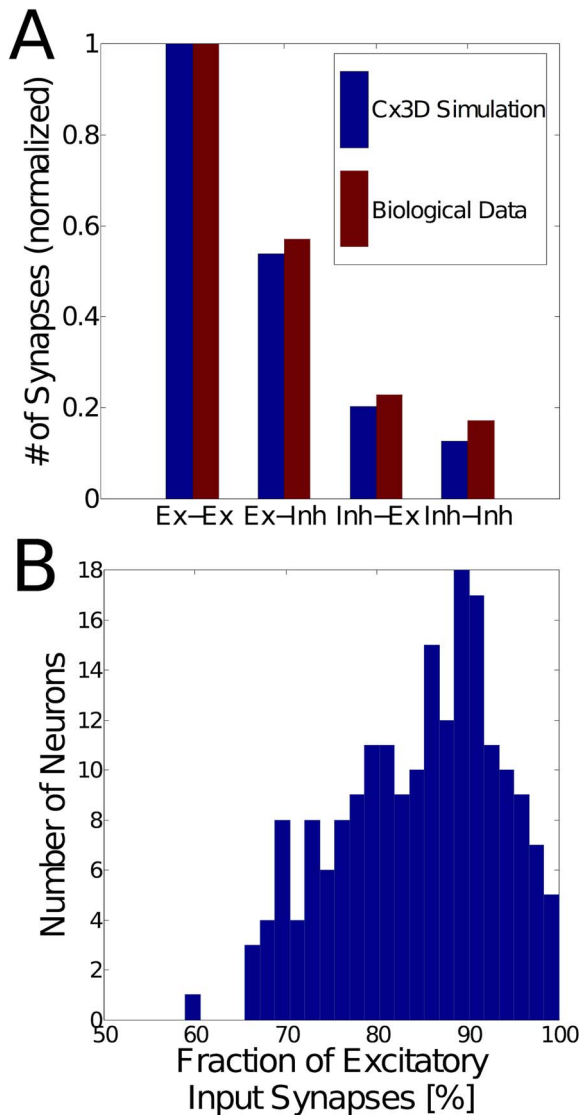


Figure 4. Connectivity after simulated neurite outgrowth. (A) Comparison of connectivity statistics from Cx3D simulations (blue) with experimental data (red) from [5]. Indicated on the vertical axis are the numbers (normalized with respect to the first bar) of synapses onto a single neuron. The individual bars show the values for the different pre- and postsynaptic neuron pairs (excitatory or inhibitory synapses onto an excitatory or inhibitory postsynaptic neuron). The numbers match in proportion, while the absolute quantities are higher in the biological data (approximately 155 vs. 3500 excitatory synapses onto a single excitatory neuron in the simulated and biological connectivity, respectively). This particular simulation consists of 250 neurons (200 excitatory and 50 inhibitory), which are randomly arranged in 3D space. (B) Histogram of the percentage of excitatory input synapses across the simulated network from (A). Each bar indicates the number of neurons that have a particular percentage of excitatory input synapses (after neurite growth and synapse formation have ended). The final distribution has a mean of 84%, which is in line with experimental assessments [5,33–35]. doi:10.1371/journal.pcbi.1003994.g004

The equations for synaptic scaling are given by [48,49]:

$$\tau_{SS} \frac{dw_{ij}}{dt} = w_{ij}(A_i - \langle x_i \rangle) \quad (2)$$

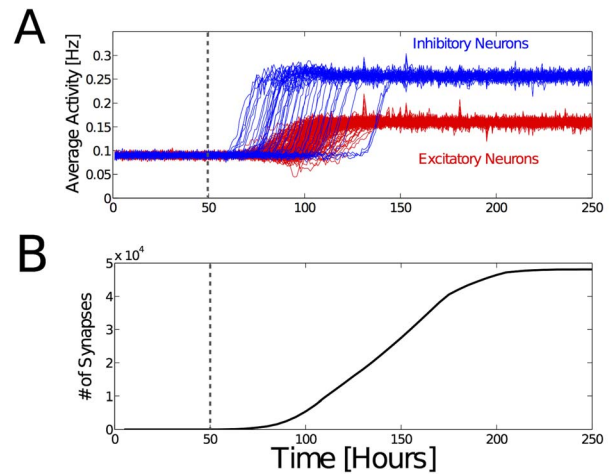


Figure 5. Homeostatic adaptation of neuronal firing rates during establishment of synaptic connectivity. (A) Synaptic scaling during neurite outgrowth leads to robust average activities of both excitatory (red) and inhibitory (blue) neurons. The network consists of 250 neurons that are randomly arranged in 3D space. The horizontal axis indicates the estimated real-time when taking into account that the time constant of synaptic scaling is in the order of several hours [17]. At $t = 50$ (dashed line), the neurite outgrowth begins. Average firing rates of layer II/III pyramidal neurons have been shown to be smaller than 1 Hz *in-vivo* [128,129]. Experimental data indicates that inhibitory neurons have higher activities A_i (Eq. 2) than excitatory neurons [68,100,130,131]. In this simulation there are not yet any input projections, so the activity originates solely from internally generated and random activity. (B) Total (excitatory and inhibitory) number of synapses in the network during development. New synapses are formed also after the neurons reach the target average activities, without disrupting the homeostatic adaptation process or bringing the network out of balance. These simulation results demonstrate the robustness of the synaptic scaling process during network growth. doi:10.1371/journal.pcbi.1003994.g005

where w_{ij} is the connection strength from neuron j to neuron i , τ_{SS} is the time constant of the learning rule (usually hours or days), A_i is desired average activity of postsynaptic neuron i , and $\langle x_i \rangle$ is the actual average activity of neuron i . Fig. 5 shows that this synaptic scaling permits the simulated network to reach a stable state with robust excitatory and inhibitory firing rates.

Post-synaptic scaling is not the only mechanism that can be used for neuronal activity homeostasis. For example, [49] has described an extended version of synaptic scaling: The presynaptic-dependent synaptic scaling (PSD) rule. We also implemented that PSD rule, but obtained results which differed only slightly from traditional synaptic scaling.

In the later stages of this first phase, input neurons (that are not part of the growing network) are added to the model (see Fig. 1C). These input neurons could correspond, for example, to thalamic or cortical layer IV neurons. They are initially fully connected to neurons of the grown network, and their projection efficacies are randomly drawn from a uniform distribution. Importantly, there is a neighborhood relationship amongst the input neurons: Input populations can be topologically close to, or distant from one another. The input neurons provide coarsely patterned input activity to the grown network. We chose hill-shaped patterns of activity centered on a given population, and decaying with topological distance from its center. The centers of these patterns move periodically in a noisy wave-like fashion (see Methods). This patterning of the electrical activity in the input layer can be interpreted as, for example, the retinal waves in early development

[19–23], that can induce correlations within the activities of downstream neural subpopulation.

By the end of this homeostatic phase, neurons and synapses have reached their final structural configuration. Overall, this phase prepares the network for the next phase of correlation-based learning of input stimuli.

Specification phase. In this phase synapses onto excitatory postsynaptic neurons obey the Bienenstock-Cooper-Munro (BCM) learning rule [18,50,51], rather than synaptic scaling. The BCM rule is composed of a Hebbian term, and a homeostatic term which determines whether the Hebbian synapse grows stronger or weaker.

The BCM learning rule is:

$$\tau_{BCM} \frac{dw_{ij}}{dt} = x_i \cdot x_j \cdot (x_i - \theta_i) \quad (3)$$

where x_i, x_j denote the discharge rates of post- and pre-synaptic neurons i, j ; θ_i is the averaged square of neuron i 's firing rate, multiplied by a constant ($\theta_i = c_i x_i^2$). The constant c_i determines the average firing rate that the neuron converges towards in the stationary state; the condition $\frac{dw_{ij}}{dt} = 0$ is met in the non-trivial case where $x_i = \theta_i$. Let A_{ex} and A_{inh} denote the target average firing rates of excitatory and inhibitory neurons, respectively. Then in order for the neurons to converge to these firing rates, c_i is set to $\frac{1}{A_{ex}}$ if neuron i is excitatory, or $\frac{1}{A_{inh}}$ if it is inhibitory.

All synapses (excitatory and inhibitory) made onto excitatory neurons follow the BCM learning rule, while those onto inhibitory neurons follow the synaptic scaling (Eq. 2) rule. While learning is commonly attributed to excitatory synapses, inhibitory synapses can also undergo long-term potentiation (LTP) as well as long-term depression (LTD) [52–55].

The lack of a correlative term for synapses onto inhibitory postsynaptic neurons is, as shown below, necessary to match experimental data on orientation selectivity of excitatory and inhibitory neurons in mouse visual cortex. We therefore hypothesize that basket cells in the superficial layers of cortex homeostatically adapt their input synapses, in contrast to pyramidal neurons, which also use correlational information.

We have also explored the case in which the same learning rule is used by all synapses. This case also yields WTA functionality (see below). Given that there are many different classes of inhibitory neurons [56], which differ also in their developmental characteristics [57], it is possible that different interneuron types follow different learning rules.

Functional Properties

Self-organization of WTA functionality. As a consequence of the synaptic learning in the second developmental phase, the network learns the topology of its inputs. Those neurons which are excited by a common input, become more strongly connected with one another. Because of the competition that is inherent to the BCM learning rule, excitatory neurons become progressively more connected to only particular input neurons (those which evoke their strongest response), while decreasing their affinity to the others. Fig. 6A shows that the final functional connectivity of excitatory neurons indeed exhibits a strong neighborhood relationship: The connection weights are stronger around the diagonal, so that the neurons are close to or distant from one another in weight space. This connection topology reflects the (1-dimensional) topology of the input patterns.

The inhibitory neurons do not integrate into this topology because the synapses onto inhibitory neurons follow the non-Hebbian synaptic scaling rule, and so their input correlations can not be learned. Fig. 6B,D show examples of the final soft-WTA functionality, after the network has learned the input topology. The excitatory neurons receiving the largest input are predominantly enhanced due to recurrent excitation with one another. The inhibitory neurons reflect the overall activity, and reduce the losing neurons' activity, more than they are enhanced by excitatory inputs. From a functional point of view, this active selection of the winning population improves the signal to noise ratio, and confirms their sWTA properties.

Unsupervised clustering. WTA networks are able to perform pattern recognition and classification, i.e. that neurons cluster functionally and respond to patterns in a discriminative and classifying manner. We explored whether this property can arise in a biological setting, as captured by our developmental model. To do this, the processes of connectivity establishment and synaptic homeostasis were simulated as described before. However, during the learning phase input patterns consisting of discrete bars of different position and orientation (Fig. 7A) were presented to the network. In this input regime there are no continuous orderings between individual patterns (which is the case for the retinal-wave like activation patterns).

Learning the discrete input stimuli causes the population to partition into sub-populations, or assemblies, as shown in Fig. 7C. We demonstrated the generality of this learning by simulating the clustering in response to presentation of only 4 input stimuli, using the same network and simulation parameters as in the case with a full range of stimuli (S2A Figure). We also examined the scenario in which all synapses (including those onto inhibitory neurons) follow the same BCM learning rule. As anticipated, this case yielded networks in which inhibitory neurons cluster along with the excitatory populations (compare S2B Figure with Fig. 7C).

The clustered functional connectivity allows the network to decorrelate its inputs, so that even noisy signals can be reliably differentiated. We quantified this ability by testing the network response to a particular pattern U , by comparison with pattern V . This comparison was measured using the scalar product between the activities in the network after presentation of the different input patterns. Let a_U and a_V be the n -dimensional vectors of the neuronal activities in a network of n neurons, in response to the stimuli U and V , respectively. The scalar product $s = a_U \cdot a_V$ then quantifies whether the responses to the two stimuli U and V are very different (s close to 0) or correlated (s close to 1). To demonstrate that the results are valid under more biologically plausible conditions, noisy stimuli were used. A noisy input stimulus is defined as:

$$I_{k,U} = \begin{cases} M + \epsilon & \text{input population } k \text{ is active in pattern } U, \\ \epsilon & \text{otherwise.} \end{cases} \quad (4)$$

where k is the index of the input population and U is the stimulus identifier. M is the amplitude of the active populations in the input (which we set to 10 Hz in this case), and ϵ is uniformly distributed noise in the range $[0, 0.3 \cdot M]$. Fig. 6D, E shows the correlations of the network's responses for 8 different input stimuli before (D) and after (E) learning under noisy *vs.* not noisy ($\epsilon = 0$) conditions. The off-diagonal elements in the correlation matrix are much lower after learning than before. These results demonstrate the decorrelation of the network's activity, and the robustness to input noise.




Figure 6. Winner-take-all functionality. (A) Weight matrix of 117 excitatory neurons in a WTA network. After learning the network exhibits a 1-dimensional neighborhood topology, as shown by the predominantly strong weights around the diagonal. This topology mirrors the neighborhood relationship of the input stimuli, which are continuously and periodically moving hills of activity. Only the excitatory connections are shown here, because the inhibitory neurons do not integrate into the neighborhood topology (see text). (B) Demonstration of WTA functionality on the network connectivity shown in (A). Neurons are ordered here such that adjacent neurons connect most strongly. The input to the network (x_{in} ; top row) has a hill shape, with added noise. The network response (x_{out} ; middle row) is a de-noised version of the input with the bump in the same location. The neuronal gain ($\frac{x_{out}}{x_{in}}$; bottom row) is high for neurons receiving the strongest input, and low (or zero) for neurons distant from the main input to the network. The dashed horizontal line indicates a gain of 1. (C) Activity of a winning neuron (blue, solid), during presentation of its feedforward input (blue, dashed) in the same simulation as shown in (B). Recurrent connectivity amplifies the response of the neuron for the duration of the stimulus ($t \in [0; 0.5s]$). In contrast, a losing neuron (green, solid) receives non-zero feedforward input (green, dashed), but is suppressed due to the WTA functionality of the network. (D) Response of the same network to a different feedforward input. The recovery of a bump shaped activity can occur anywhere in the network topology.
doi:10.1371/journal.pcbi.1003994.g006

Competition between states. In addition to the decorrelation of responses, clustering provides competition between inputs. This property is computationally interesting because it forces the network to make a decision based on its input (Fig. 8A). We demonstrated this competition by presenting simultaneously two competing patterns (after the network had learned 4 different patterns). The relative proportion of these patterns in the input was gradually varied between the first and second pattern.

The results show that the stronger stimulus non-linearly dominates responses in the WTA network, so that the masked stimulus evokes an activity pattern that resembles that evoked by the strong stimulus alone. These results are in accord with experimental studies in visual cortex [58–60] and auditory cortex [61].

The nature of the competition between the states is dependent on the functional connectivity. Strong recurrent excitation (i.e. a high gain) yields strong inhibition, which results in a marked switching behavior between the different populations. This is because the competition is strong, and so the switch from one state to the other is more evident. A high slope of the transition reflects a functionality similar to a bistable switch. More specifically, the

slope of the transition (middle part of the interpolation in Fig. 8A) increases with the gain of the WTA network. This gain can be adjusted via the homeostatic average activities: Higher target activities lead to more recurrent excitation, which increases the gain. Such differently graded competition is seen in Fig. 8B. Bistability is also interesting from a computational point of view, because discrete states can be represented reliably. This kind of reliability is useful for performing computation with states based on analog elements [8,62]. Competition also develops when synapses onto excitatory as well as inhibitory neurons follow the BCM learning rule, as shown in S3 Figure.

Correspondence with Orientation Selectivity of Excitatory and Inhibitory Neurons

We investigated whether our developmental model can account for experimental findings on orientation selectivity in visual cortex; for example, differences in tuning between excitatory and inhibitory neurons. In order to address this question, we assumed that the hills of activity in the input layer correspond to oriented stimuli (e.g. bars), which are smoothly and periodically rotating between 0 and 180 degrees. As anticipated from the previous




Figure 7. Clustering and decorrelation of representations. (A–C) Discrete input patterns give rise to clusters in the functional connectivity of the WTA network. (A) Input stimuli used in the learning process. Filled and empty spheres indicate strongly and weakly active populations, respectively. (B,C) Visualization of the network structure before and after learning. Strongly-coupled neurons are drawn close together; excitatory synaptic connections are indicated by grey links. Excitatory neurons are coloured according to their preferred input pattern (colours in A); inhibitory neurons (square) are drawn in yellow. (B) Before learning, no clustering of synaptic connections is present. (C) After learning, neurons with the same preferred stimulus are strongly interconnected. See S2 Video. (D) Before learning, the response of the network is similar across all stimuli. Shown is the scalar product between the vectors of neuronal responses to pairs of stimuli $\langle x_i, x_j + \text{noise} \rangle$. The noise was added in order to assess the sensitivity of the network's activity to a perturbation of the input signal (see text). The high values and uniformity of scalar products in (D) indicates that network responses poorly distinguish between stimuli. (E) After learning, responses to noisy stimulus presentations are highly similar (high values of scalar product; black diagonal), whereas responses to different stimuli are decorrelated (low values of scalar product; light shading). doi:10.1371/journal.pcbi.1003994.g007

results, excitatory neurons become highly orientation selective (Fig. 9), in contrast to inhibitory neurons. These results are in line with biological data. For example, [63] have analyzed orientation selectivity of excitatory and inhibitory neurons in mouse visual cortex. They report inhibitory neurons to be more broadly tuned and hence less selective than excitatory, pyramidal neurons. Similar findings were reported by [64–68].

We also quantified the orientation tuning based on the orientation selectivity index (OSI), which specifies the degree to which a neuron is selective for orientation. The value of this index lies between 0 (non-selective) and 1 (selective to a single, specific orientation). Fig. 9B shows the distribution of the OSI for excitatory and inhibitory neurons in a WTA network, demonstrating the discrepancy of orientation selectivity also on a population level. We conducted additional simulations, which demonstrated that when inhibitory neurons follow the same

learning rule as excitatory neurons, they exhibit more narrowly tuned orientation selectivity (Fig. 9C). Hence, experimental findings of orientation selective inhibitory neurons in cat visual cortex [69–72] can also be accounted for by our model.

Inhibition of Excitatory Neurons

We have analyzed the consequences of our model on the nature of the inhibition of excitatory neurons. As mentioned above, inhibitory synapses onto excitatory neurons are subject to the BCM learning rule (Eq. 3).

The competition between excitatory neurons depends on the common input that they all receive from inhibitory neurons. This common input must reflect the overall activity of the network, so that the competition is suitably normalized. However, the inhibition of the excitatory neurons stems from multiple inhibitory neurons, which should partition their common inhibitory task

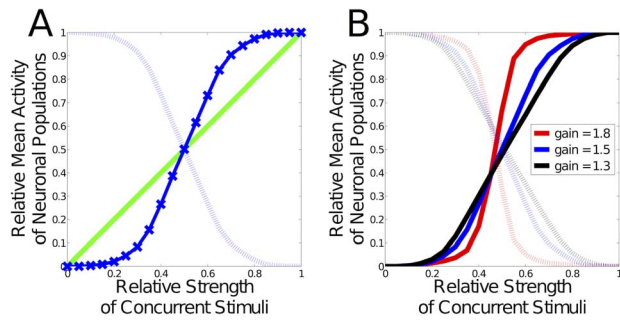


Figure 8. Stimuli are represented by competing subpopulations. (A) Competition for representation of a mixture of 2 concurrent stimuli. Shown is the normalized average activity of two subpopulations, in response to mixtures of the preferred stimuli of the two populations. For mixtures containing predominately one stimulus (mixture proportions close to 0 and 1), the populations are strongly in competition, and the network represents exclusively the stronger of the two stimuli (responses near 0 and 1). For intermediate mixture proportions, competition causes a rapid shift between representations of the two stimuli (deviation from diagonal reference line). (B) Increasing the gain of the network $\frac{\bar{x}_{out}}{\bar{x}_{in}}$ (black line: 1.3, blue: 1.5, red: 1.8) increases the stability of representations, and increases the rate of switching between representations due to stronger competition. doi:10.1371/journal.pcbi.1003994.g008

amongst each other in a self-organizing way. We investigated this partitioning, and how an excitatory neuron is inhibited during stimulation.

In order to quantify the impact of a neuron j on another neuron i for a given stimulus, we calculate a value that we will call the recursively effective exertion (REE). It is obtained by multiplying the activity of neuron j (under a given stimulus U) with the total connection weight w_{ij} from neuron j to i :

$$REE_{ij}(U) = x_{j,U} \cdot w_{ij} \quad (5)$$

The REE value is therefore stimulus-dependent, and dependent on the recurrent network connectivity. Fig. 10 shows that inhibition is distributed non-uniformly: A few inhibitory neurons dominate the suppression of an excitatory neuron. This dominance is due to the BCM learning by inhibitory synapses: Strongly and weakly correlated inhibitory connections to excitatory neurons are strengthened or weakened, respectively. These inhibitory connection strengths converge because of the homeostatic activity regulation, which is part of the BCM learning rule.

The nature of inhibition of excitatory neurons is interesting in the context of the anatomy of inhibitory basket cells. These neurons predominantly target locations close to the soma or the proximal dendrites, where they can strongly influence the excitatory neuron [73]. Therefore, it is plausible that the recruitment of a small number of inhibitory neurons is sufficient to inhibit an excitatory neuron. Electrophysiological experiments could in principle validate this hypothesis by showing that only a small proportion of the inhibitory neurons projecting to a pyramidal neuron are predominantly responsible for its suppression.

Discussion

In this paper we have demonstrated by simulation of physical development in a 3D space, how an autonomous gene regulatory

network can orchestrate the self-construction and -calibration of a field of soft-WTA neural networks, able to perform pattern restoration and classification on their input signals. The importance of this result is that it demonstrates in a systematic and principled way how genetic information contained in a single precursor cell can unfold into a functional network of neurons with highly organized connections and synaptic weights.

The principles of morphological and functional development captured in our model are necessarily simplified with respect to the boundless detail of biology. Nevertheless, these principles are both strongly supported by experimental data, and sufficiently rich in their collective expression to explain coherently the complex process of expansion of a genotype to a functional phenotypic neuronal circuit. In this way our work offers a significant advance over previous biological and modeling studies which have focused either on elements of neuronal development, or on learning in networks whose initial connectivity is given. Therefore we expect that methods and results of the kind reported here will be of interest both to developmental biologists seeking a modeling approach to exploring system level processes, as well as to neuronal learning theorists who usually neglect the genetic-developmental and homeostatic aspects of detailed learning in favor of an initial network that serves as a basic scaffold for subsequent learning [74–76].

It is relatively easy to express a well-characterized biological process through an explicit simulation. That is, one in which the simulation simply recapitulates the process by expanding some data through a simple model, without regard for physical and mechanistic constraints. By contrast, the simulation methods [16] that we have used here are strictly committed to physical realities such as 3D space, forces, diffusion, gene-expression networks, cellular growth mechanisms, etc. Our methods are also committed to local agency: All active processes are localized to cells, can only have local actions, and have access to only local signals. There is no global controller with global knowledge, able simply to paint the developmental picture into a 3D space. Instead, the ability of a precursor cell to expand to a functional network is the result of collective interaction between localized cellular processes. And overall, the developmental process is the expression of an organization that is encoded only implicitly, rather than explicitly, in the GRN of the precursor cell. Thus, our GRN encodes constraints and methods rather than explicit behaviors.

In previous work [77,78] we have shown how this approach can be used to explain the development of neocortical lamination and connectivity. In that case we did not consider also the electrophysiological signaling between cells and so the self-configuration of their computational roles, as we have done here. However, the incorporation of electrophysiological signaling into the growth model brings substantial technical difficulties, such as those arising out of the large differences in spatio-temporal scales between cellular developmental and electrophysiological signaling processes, as well as the supply and management of sufficient computational resources. Therefore we have chosen to keep these problems tractable in this first functional study, by restricting our question to a sub-domain of cortical development: How could neuronal precursors expand into functional circuits, at all. Even then, we must be satisfied for the moment with a rate based model of neuronal activity, rather than a fully spiking one.

The emphasis of this paper is on the process whereby a precursor expands to some useful network function. The particular function is less relevant, and in any case the functional/computational details of cortical circuits are as yet not fully understood. We have chosen to induce WTA-like function because our previous work has been focused on the likely similarity




Figure 9. Excitatory neurons are strongly tuned; inhibitory neurons are poorly tuned. Tuning properties of excitatory and inhibitory neurons. (A) Representative tuning curves for 3 excitatory (red, 1-3) and 3 inhibitory (blue, 4-6) neurons in a WTA network after the learning process. Excitatory neurons exhibit strong and narrowly tuned preference for certain inputs, in contrast to inhibitory neurons. (B) Distribution of the orientation selectivity index (OSI) across all excitatory and inhibitory neurons in a WTA network, demonstrating the discrepancy of tuning on a population level. (C) Simulation of the same learning rule for synapses onto excitatory as well as inhibitory neurons yields orientation-tuned neurons in both populations.

doi:10.1371/journal.pcbi.1003994.g009

between the WTA motif and the neuronal types and their interconnectivity in the superficial layers of cortex [6]. Moreover these WTA networks are intriguing from both the biological, and

computational perspective [3,6–15,41]. The strong recurrent excitation available in the superficial layers of cortex, and their critical dependence on feedback inhibition has been clearly




Figure 10. Inhibition of excitatory neurons. Excitatory neurons are predominantly inhibited by subsets of the inhibitory neurons that project to them. **(A)** Representative examples of the inhibition to excitatory neurons in a learned WTA network, during presentation of a stimulus. The vertical axis indicates the percentage of the total inhibitory REE (see definition of REE in text) that an individual inhibitory neuron delivers to this particular excitatory neuron. Few (usually 2 or 3) inhibitory neurons provide the major part of the inhibition. **(B)** Histogram of all the REE contributions (in %) from inhibitory neurons, across all excitatory neurons in the WTA network. The distribution shows that few inhibitory neurons provide the major part of the inhibitory REE on an excitatory neuron. This specialization is a result of the BCM learning rule, which is followed also by inhibitory synapses onto excitatory neurons.

doi:10.1371/journal.pcbi.1003994.g010

demonstrated by intracellular recordings in the presence of ionophoretic manipulation of GABA agonists and antagonists [4]. These relationships are crucial for WTA-like processing, because they offer the network induced gain that is crucial for providing the signal restoration, signal selection, and process control that support systematic computation. Recent optogenetic studies appear to confirm the presence of circuit induced gain, in the input layers of mouse visual cortex [79,80]. Taken together these experimental and theoretical results support the hypothesis

that at least some fundamental WTA functionality is embedded in the processing architecture of superficial neuronal circuits, and so makes the WTA motif a worth target of the developmental process that we have described here.

Our model predicts that neurons form specific subgroups, or cell assemblies [81,82]. There is indeed strong evidence from biological data for this clustered connectivity [83–85], which (as in our simulations), appears to be grounded in the similarity of functional selectivity [86].

We did not allow dynamic rearrangement of synapses in these first simulations. However, it is plausible that weak synapses are pruned away, freeing synaptic resource to explore for more correlated partners.

Peters' rule [87–89] proposes that connectivity can be estimated by the product of the random overlap of pre- and postsynaptic sites. This rule may be true for average connectivity, but specific functionality obviously calls for more specific low level connectivity within the average. One opinion is that such specificity is explicitly genetic, and so accounts for example for the diversity of cortical interneurons [90,91]. Instead, our result speaks for an implicit rather than explicit genetic specificity. That is, the apparently specific wiring of the WTA network arises by neurons collectively satisfying genetically expressed constraints. This concept is in stark contrast to the view that network functionality emerges from individual processes that do not coordinate with potential interaction partners. In our simulations, a neuron's morphology and the functional strengths of its synapses depend on the collective behavior of the other neurons. Hence, the structure and function of a neuron grown in isolation is different from a neuron with the same genetic code, but that interacts and coordinates with its environment during development.

Our learning rule requires that input projections are ordered in such a way that their collective input patterns provide at least a coarsely structured signal against which the presumptive WTA layer of neurons can successfully deploy a BCM-like learning mechanism. This ordering is not a stringent requirement. For example, provided that there is some degree of coherent axonal mapping of axons from input neurons of the subplate or thalamus into the target WTA layer, then even metabolically induced travelling waves of activity across the developing input population could provide a sufficiently structured signal for learning.

Traditionally, many modeling studies have been based on the assumption that the limited lateral extent of the neuronal axonal and dendritic tree naturally leads to a properly configured 2D neighborhood topology [92–94]. However, it is unclear how more realistic anatomical properties (anisotropy, variation of neurite extent, irregular locations of somata etc.) affect these topologies. Our work addresses this problem by demonstrating how neurons can self-calibrate in a stimulus-induced way, within a non-uniform and irregular neuronal setting. Hence, our work provides a better understanding of how developmental mechanisms can generate a neighborhood topology, and so is complementary to the classical approach.

As development of input neurons proceeds, the degree of structuring is likely to improve also, so that input neurons projecting to the same targets share similar features (for example, their ON- and OFF-subfields). This is in line with studies on thalamo-cortical projections [95], as well as cortico-cortical projections from layer IV to II/III [96]. However, it should be noted that this input specificity does not play onto inhibitory targets, which is in accordance with our work. Since the input to the neurons shapes the functional connectivity in the network, it follows from our model that neurons which receive common input are more likely to connect with each other (assuming that

structural connectivity is adjusting to functional connectivity). The studies of [96] and [97] provide evidence for this input-dependent intra-network specificity.

Our results predict that only a few inhibitory neurons provide the major part of WTA-relevant inhibition, i.e. a relatively small subset of all the inhibitory basket cells projecting to a single pyramidal cell is responsible for its WTA suppression. These results suggest that WTA inhibition might not be very redundant, so that de-activation of only a few inhibitory neurons could result in very different electrophysiological behavior of pyramidal cells.

Our networks employ Hebbian-type learning for both excitatory and inhibitory synapses onto excitatory postsynaptic neurons. It is known that inhibitory synapses can undergo long-term potentiation (LTP) as well as long-term depression (LTD) [52–54], and learning by inhibitory synapses has been used in previous modeling studies [45,98].

Non-Hebbian synaptic scaling of synapses onto inhibitory neurons results in orientation-nonselective inhibitory neurons. This distinction with respect to pyramidal neurons has been observed in mouse visual cortex, where the tuning of inhibitory neurons is broader than that of excitatory neurons [63–67,99–101]. There is evidence for broadly tuned thalamo-cortical input to inhibitory neurons [95], as well as cortico-cortical input to those of layer II/III of mouse cortex [100]. Therefore we propose that at least some types of inhibitory neurons (e.g. fast-spiking (FS), PV-expressing interneurons) do not selectively adjust their inputs, but uniformly adapt the electrophysiological properties of their inputs for homeostasis.

Orientation-selective inhibitory neurons are found in cat visual cortex [69–72]. Since we do not model orientation maps, our findings are not directly applicable to the cat. However, we argue that it is the spatial location in the orientation map that determines the tuning curve of inhibitory neurons. Most cortical interneurons have a small horizontal dendritic extent [56], and so they likely receive inputs from similarly tuned excitatory neurons within an orientation map. Inhibitory neurons located close to orientation pinwheels are expected to have relatively broad orientation tuning, as reported in the above studies. The unbiased pooling of surrounding activity by inhibitory neurons is also supported by experimental results across species and sensory modalities [63,101]. By contrast, we have shown that inhibitory neurons become orientation-selective when they follow the same (BCM) learning rule as excitatory neurons.

Our learning model provides a computational explanation for why most interneurons are smooth, i.e. have very few dendritic spines. It is believed that spines, by compartmentalizing biochemical signals, provide the molecular isolation required for independent synaptic learning [102–104]. The nonspecific and homogeneous adaptation of inhibitory neurons, which in our model are homogeneously scaling the input efficacies, is therefore well in line with this suggested function of dendritic spines. This model also provides an explanation for the finding that inhibitory, but not excitatory neurons exhibit structural remodeling of dendrites in the adult rat [105]. Changes in excitatory morphology at the level of dendritic branches (rather than spines) could have detrimental effects on already consolidated memories. Inhibitory neurons may retain their potential for dendritic restructuring, because their homeostatic adaptation does not interfere with learning of sensory experience.

We believe our findings to be robust also with respect to models incorporating spikes, because the main features of the adaptation and learning behavior have been demonstrated also on this more detailed level of electrophysiology. Along these lines, the studies of [106,107] have explored spike-based WTA network functionality. Spike-dependent plasticity (STDP) is a Hebbian learning rule [108] and can yield synaptic homeostasis [109]. In particular, the

BCM learning rule has been related to STDP mechanisms [109–112].

The robust self-organization of the WTA network is remarkable in that it arises out a single precursor cell, by simple genetically encoded rules. In future, this genetic developmental approach to functional circuit construction could be extended to larger networks composed of multiple WTA networks. For example, it has been hypothesized that by cooperation of multiple WTA circuits, the superficial layers of cortex could perform context-dependent processing [8]. Along these lines, [78] provide a model for the development of long-range projections connecting multiple columns, arranged on an hexagonal grid, as is observed in the superficial patch system [113–116]. It also remains to integrate these computational aspects into the context of a laminated cortical structure, which has already been simulated in Cx3D [24,77].

Materials and Methods

Simulation

Cx3Dp. The growth simulations were conducted with the open-source package Cx3Dp, the parallelized version of Cx3D [16] (available at <http://www.ini.uzh.ch/projects/cx3d/>). As in the non-parallel version, neurons in Cx3Dp are decomposed into discrete spherical or cylindrical *physical elements* emulating the physical properties of developing tissues, whereas the biological properties derive from *modules*, that is, smaller programs expressed within specific physical elements. This local instantiation forces simulations to be based only on local interactions, without any global control of the developmental processes.

The default parameters for physical objects in Cx3D were initially chosen so that (1) density of cells/neurites can not be infinite, (2) objects do have a minimal adhesive property ensuring tissue integrity and (3) viscosity and not mass opposes to movement (see [16]). For the present study, we did not have to modify these default parameters.

Computer specifications. We used a rack computer with two 12-core AMD Opteron 6168 processors (1.9 GHz, 64 GB of RAM), running under Ubuntu 12.04 LTS.

Gene Regulatory Network (GRN)

The GRN is defined by a set of variables x that represent genes and the corresponding substance concentrations. Changes in substance concentration are described by the rate equation:

$$\dot{x}_i = k_1 \mathcal{F}_i[\mathbf{Z}(\mathbf{x})] - k_2 x_i \quad (6)$$

where x_i is the concentration of a protein encoded by the gene i (i.e. $G0, G1, G2, GE$ or GI), and \mathbf{x} the corresponding concentration vector. The function \mathcal{F}_i expresses how the synthesis rate of the protein encoded by gene i depends on the cooperative binding of all the substances, and k_1, k_2 represent the production and degradation rates ($k_1, k_2 \geq 0$). $\mathbf{Z}(\mathbf{x}) = [Z(x_1), Z(x_2), \dots, Z(x_n)]$ is a vector of Hill functions, which compute the binding probability of a substance x_i to a regulatory region given the affinity constant θ , cooperativity m and binding bias b :

$$Z(x_i, b, \theta, m) = \frac{(x_i + b)^m}{\theta^m + (x_i + b)^m} \quad (7)$$

Gene substances can regulate gene expression by binding to specific sites in the genomic cis-regulatory regions. Substances that

regulate each others' transcription are called transcription factors. Many genes are controlled by a number of different transcription factors and different arrangements of binding sites can compute logic operations on multiple inputs. Here, the function \mathcal{F}_i takes the form of a logical combination of interacting substances and is defined by the elementary operations:

$$\text{AND}(x_1, x_2) = x_1 \cdot x_2 \quad (8)$$

$$\text{OR}(x_1, x_2) = x_1 + x_2 - \text{AND}(x_1, x_2) \quad (9)$$

$$\text{NOT}(x) = 1 - x \quad (10)$$

More information on this description of GRN dynamics can be found in [117,118]. Although abstract, this formalism can be directly translated into the corresponding mechanistic, kinetic differential equations. For our computational model based on 5 genes, we have used the following equations:

$$G0'[t] = 1 - G0[t] \quad (11)$$

$$G1'[t] = \alpha(\text{THR}[G0[t]] - G1[t]) \quad (12)$$

$$G2'[t] = \text{THR}[G1[t]] - G2[t] \quad (13)$$

$$\begin{aligned} GE'[t] = & \text{OR}[\text{AND}[\text{THR}[G1[t]], \\ & \text{NOT}[\text{THR}[G2[t]]], \text{THR}[G2[t]]] - GE[t] \end{aligned} \quad (14)$$

$$\begin{aligned} GI'[t] = & \text{OR}[\text{AND}[\text{THR}[G1[t]], \\ & \text{NOT}[\text{THR}[G2[t]]], \text{THR}[G2[t]]] - GI[t] \end{aligned} \quad (15)$$

with:

$$\text{THR}[x] = (x + 0.01)^4 / ((x + 0.01)^4 + 1) \quad (16)$$

The probabilities of either differentiating into neurons by $G1$ induction (p_1) or by $G2$ induction (p_2) are computed as follows:

$$u = \lfloor \frac{\log N_{\text{Target}}}{\log 2} \rfloor \quad (17)$$

$$v = \lfloor N_{\text{Target}} - 2^u \rfloor \quad (18)$$

$$p_1 = 1 - \frac{v}{2^u} \quad (19)$$

$$p_2 = 1 - p_1 \quad (20)$$

where u is the number of divisions in the first division cycle, v is the difference between the target number of neurons (N_{Target}) and the number of neurons resulting from the first division cycle, and $\lfloor \cdot \rfloor$ denotes the floor function for rounding to integers.

The intrinsic production constant α determines the number of cell divisions until differentiation into excitatory and inhibitory neurons can occur. The higher it is, the faster the $G1$ gene reaches the threshold of 0.99. α was adjusted manually in order for u divisions to occur in the $G1$ cycle.

Development of Neuronal Morphology

Initially, neuronal cell bodies are assigned uniformly random positions in 3D unprepared space. In Cx3D, these cell bodies are modeled as physical spheres. The neuronal cell density was in agreement with values derived from experimental data, i.e. in the range of 40'000 to 86'900 per mm^3 [119–121]. We found 250 neurons (200 excitatory and 50 inhibitory) to be appropriate for the available computer resources. For the establishment of neuronal connectivity, the somata were placed randomly in a cube with side length 160 μm . A smaller network of 150 neurons in a cube with side length 140 μm was used for simulations where the second developmental phase was simulated, in order to decrease simulation time. 3 of these 150 neurons did not get inhibitory inputs after the initial outgrowth and were not included for the simulation of learning, such that the analyzed network consisted of 117 excitatory and 30 inhibitory neurons. Standard Cx3D parameters for the physical properties of the cells (e.g. mass or adherence) were used [16]. The somatic diameters were set to 8 μm . Variation of these parameters had only minimal effects on the simulation results.

Axonal and dendritic growth were encoded with the instruction language G-code [24]. We used the following mechanisms, which are executed by such G-code “modules” located in the growth cone, for axonal and dendritic growth, as well as synapse formation:

Axonal growth. The axon is initially extended from the cell body in a random direction, and is dependent on the concentration of an extracellular substance that is secreted by the somata of the neurons participating in the WTA network. The growth cones of excitatory and inhibitory neurons sense the substance secreted by inhibitory and excitatory neurons, respectively. As long as the concentration of this substance is higher than a given threshold θ_1 , the axonal growth continues (see below). If the growth cone enters a zone where the concentration is below the threshold θ_1 , then the axon retracts until the concentration is above another threshold θ_2 (with $\theta_2 > \theta_1$), before resuming elongation. During elongation, the direction at each time step is a weighted sum of the direction from the previous time step, and of a random perturbation (i.e. $\mathbf{d} = w_{\text{prev}}\mathbf{d}_{\text{prev}} + w_{\text{noise}}\mathbf{d}_{\text{noise}}$). At each time step, the axon bifurcates with probability $p_{\text{base}} + p_{\text{substance}} \cdot c_{\text{substance}}$, where p_{base} and $p_{\text{substance}}$ are constants and $c_{\text{substance}}$ is the substance concentration at the current location of the growth cone. When the axon elongates or branches, its diameter is reduced by a factor r_{move} or r_{fork} , respectively. The outgrowth stops when the axonal diameter falls below d_{minimal} .

Dendritic growth. Each soma produces three dendrites. As for the axon, the initial sprouting direction for dendrites is random. The subsequent elongation direction is also chosen as a weighted sum of the previous direction and a random perturbation. The major difference is that dendrites are not sensitive to extracellular cues (and do not retract). Branching and stopping mechanisms are implemented as in the axonal case. As a result, the overall dendritic morphology develops isotropically.

Synapse formation. In Cx3D, synapses are modeled as connections between excrescences on neurite elements of axons (boutons) and dendrites (spines). For simplification, spines represent potential postsynaptic densities in general, such that we do not model the absence of spines in most classes of inhibitory neurons [56]. During elongation of the neurite, an excrescence is instantiated in the middle of the discretized element. Whenever the excrescence of the tip of a growing axon or dendrite is close to another excrescence, it can check whether they are of complementary types. This local process of synapse formation is done by means of a module that is instantiated in the most distal neurite element, which corresponds to the growth cone's location. If this condition of complementarity is fulfilled and the excrescences are close enough (i.e. the distance $|\mathbf{x}_{\text{bouton}} - \mathbf{x}_{\text{spine}}| \leq 2 \mu\text{m}$, where \mathbf{x} denotes the location where the excrescence is attached to the segment of the axon or dendrite), a synapse is formed. Synapse formation between neurites belonging to the same neuron (i.e. autapses) is prohibited, because their number has been reported to be relatively small [122,123], and their electrophysiological significance is unclear.

In our Cx3D implementation, synapse formation also implies the establishment of a physical bond between the two excrescences on the axon and dendrite. This bond is approximated as a spring which reacts linearly to the force to which it is subject to. Therefore, connected neurite segments are kept close to each other, except when a certain repulsive force is exceeded. Once a bond is overstretched, it is released and the synapse destroyed. This ensures that synapses do not over-restrict the neuronal morphology.

The synaptic weights are assigned at synapse formation: Excitatory and inhibitory synapses are initialized with weights 0.001 and 0.01 respectively, in qualitative accordance with the estimate of [124]. The overall behavior of the simulations was not sensitive to these initial weights, because of the homeostatic adaptation processes.

Electrophysiology

The computation of the electrical activity was implemented in Java, to allow a direct interfacing with Cx3D. All the synaptic weights in the Cx3D simulation are summed up, which yields a weight matrix. Based on this weight matrix, the input activity and the spontaneous activity, the firing rate of a neuron is computed according to Eq. 1. The numerical solution of the differential equations was computed using the explicit Euler integration method. The network's activity is computed with 3000 iterations and integration step $dt=0.01$. The maximal firing rate is set to 250 Hz.

Learning. In all the simulations of WTA learning, two different scenarios of input stimulation were conducted. In one case, the input neurons were activated in the form of a hill of activity. This hill was centered around an individual input population, and decayed with the distance from this center. The center population of the hill of activity was active with a rate of 1.4 Hz, the immediate neighboring populations at $0.5 \cdot 1.4$ Hz, the second at $0.5^2 \cdot 1.4$ Hz, and the third at $0.5^3 \cdot 1.4$ Hz. If the distance of the input population to the center population was bigger than 3, the activity of the input population was chosen to be a random number between 0 and 0.06. The hill of activity in the input layer moved in a periodic fashion (i.e. the peak restarts in the first input population of the input layer after having reached the last). This first scenario represents retinal waves, or orientation stimuli between 0 and 180 degrees. In the second case, discrete input patterns (Fig. 7A) are presented to the network. The stimulated input populations are active at a rate of 2.1 Hz, while the

unstimulated ones are (as in the first scenario) active with a rate drawn uniformly from the interval $[0, 0.06]$ Hz. In both scenarios, the target average activity for excitatory neurons (A_{ex}) was identical for all excitatory neurons, and chosen between 0.14 Hz and 0.68 Hz. The inhibitory target activity was set to $1.6 \cdot A_{ex}$. For the simulations shown in Fig. 8, target average activities of excitatory neurons were 0.4, 0.55 and 0.68 for the gains 1.3, 1.5 and 1.8, respectively. During learning, the input populations were active or not active in an alternating fashion. In the non-active case, only spontaneous and random activity generated electrical activity. This non-active input mode was introduced to demonstrate that correlated, instructive input can be intermittent in time. The neurons in the network have a random spontaneous activity that is drawn uniformly from the interval $[0.06, 0.12]$ Hz.

Usually, the simulations took around 1 day to develop networks that exhibit WTA behaviour. The main computational bottleneck is the computation of the average activity, which relies on a large number of samples of neuronal activity. During an entire simulation, electrical activity in response to the input is computed around 1'000'000 times, and in the order of 100'000 learning steps (synaptic scaling and BCM learning) are performed.

The average activities used for the BCM learning (Eq. 3) are computed as the arithmetic mean of the neuronal activities from the last N inputs. N is set to $2 \times L_{\text{inp}}$, with L_{inp} being the number of input populations. The factor 2 comes from the fact that the inputs are alternating between two different modi: active and non-active input neurons (as described before). In order for the WTA neurons to keep track of their activities, N has to be long enough to allow averaging the activities evoked by all possible inputs. Analogously, for the learning of the clustering based on discrete input patterns (second scenario), N was set to $2 \times L_{\text{stim}}$, with L_{stim} being the number of different input stimuli presented to the network. The time constant of the average activity is long compared to the time constant of the instantaneous firing rate, and assumed to be in the range of several hours to days. Based on this assumption, we assessed the real time of learning, and so some figures in this work indicate the estimated real-time of the learning process.

As a standard, we chose 20 input populations (or neurons) in the case of learning continuous input patterns, and the input strengths were initialized randomly and uniformly distributed between 0 and 0.15. In the case of discrete input stimuli, we used 3×3 input populations. The initial connection strengths (from each of these input populations to each of the neurons in the grown network) were distributed in the range of $[0, 0.25]$. We found little change of behavior when varying these numbers. During unsupervised learning, retinal-wave like activity or discrete input patterns are presented to the network in a periodic fashion. After every period, the new values of average activities are obtained and learning is simulated. The same synaptic learning time constants are used for input, excitatory and inhibitory lateral synapses. In the case of retinal waves, learning is simulated only after the hill of activity has passed every input population. In the case of learning discrete input patterns, learning is simulated each time a pattern is presented. The time duration for synaptic scaling after a retinal wave was chosen equal to the duration of learning one pass of all input patterns. Because of the 8 input patterns, τ_{SS} was set 8 times larger in the scenario of discrete patterns ($\tau_{SS}=250$ for retinal waves and $\tau_{SS}=2000$ for discrete input patterns). The same time constants for the BCM learning rule were chosen, but exploratory simulations showed that this is not a necessary condition. In order to demonstrate that the WTA learning is not sensitive to changes in these time constants, simulations of WTA networks learning 4 stimuli were conducted with the same time constants as with 8 stimuli (S2A Figure). The weight change per learning step is

constrained to maximally 3 % of the current value, in order to prevent large or destabilizing disruptions at the initial stages of development. The weight of an individual synapse between two neurons is approximated by dividing the total connection weight by the number of synapses involved in this particular connection. This absolute synaptic weight is constrained to a maximal value of 0.1 [125]. Overall, variation of the learning parameters do not have a strong effect on our results.

Network Analysis and Visualization

Analyses of the simulated networks were performed with MATLAB (Mathworks Inc.). In order to assess WTA functionality, electrical activity was computed in the same way as in the Java implementation of the Cx3Dp simulation, namely using the rate-based model (Eq. 1) and the explicit Euler method. The integration step was decreased to $dt=0.001$ for minimizing integration errors.

The ordering of neurons for visualization, such as for Fig. 6A, was done using the genetic algorithm “ga.m” from the Global Optimization Toolbox of MATLAB. The energy to be minimized was defined as the sum of weighted topological distances between neurons, i.e. $E = \sum_{i,j} d(i,j) \cdot (w_{ij} + w_{ji})$, where w_{ij} are the summed

synaptic weights from neuron j to neuron i . The topological distances $d(i,j)$ are inferred from a discrete 1-dimensional position vector of the neurons, which is initialized randomly and optimized. The ordering for the matrix visualization is then given by the locations of the neurons in this vector (i.e. neighbors in this vector are also neighbors in the matrix ordering). Note that the topological position is unrelated to the physical position of the neurons, and is only used for the optimization process. The visualization of the clustering was done with CytoScape [126], an open-source framework that is downloadable from <http://www.cytoscape.org/>. We used the “dynnetwork” plugin implemented by Sabina Pfister, which clusters weighted networks based on the Kamada-Kawai algorithm [127].

Neurite Outgrowth Parameters

The neurite outgrowth has several parameters, which depend on the neuronal type (excitatory/inhibitory) and also on the neurite type (axon or dendrite). Table 1 lists all these parameters. The 2 substances which are secreted by the cell bodies and used by the axons as guidance cues both have a diffusion coefficient of 50 and a degradation constant of 5.

Supporting Information

Figure S1 Histograms of resulting numbers of neurons after simulation of the GRN. The intrinsic instructions of the precursor cell in an unprepared environment lead to multiple neurons of two types (excitatory and inhibitory, other types like for example glia cells could facultatively be added). We conducted 100 trials of a GRN, that was set to give rise to 100 neurons, of which 80 are excitatory and 20 inhibitory. These results demonstrate that the (probabilistic) GRN produces approximately the desired number and proportion of neurons.
(TIFF)

Figure S2 Visualization of network connectivity in weight space, after learning 4 input patterns. (A) The locations of the neurons are determined using a clustering algorithm, such that strongly connected neurons are close to each other. Different colors indicate different preferred patterns of the neurons. The preferred pattern of a neuron was assessed by determining the pattern that evokes the largest electrical response.

Inhibitory neurons are colored yellow and rectangular-shaped. The same network as in Fig. 7B and 7C is simulated, but after learning 4 input stimuli (horizontal, vertical and both diagonally oriented bars) instead of 8. The 4 clusters defined by the spatially proximal assemblies of neurons are visible. Importantly, the same parameters (time constants of synaptic scaling and BCM learning) were used, demonstrating the robustness of the learning scheme. (B) Network connectivity after using the same BCM learning rule both for excitatory and inhibitory neurons. As in (A), the locations of the neurons are determined using a clustering algorithm, such that strongly connected neurons are close to each other. Different colors indicate different selectivities of the neurons. In contrast to the simulations where synapses onto inhibitory neurons were following the synaptic scaling rule, here we used exactly the same learning dynamics for both types of neurons. This influences the clustering, such that also the inhibitory neurons become selective for the learning input stimuli.
(TIFF)

Figure S3 WTA competition between two populations after correlation-based BCM learning of excitatory and inhibitory neurons. WTA populations compete for representation of a mixture of 2 concurrent stimuli. In contrast to the simulations for Fig. 8, synapses onto excitatory as well as inhibitory neurons followed the BCM rule during learning. Also in this case, the populations with different preferred stimuli compete and mutually suppress each other. The blue crosses indicate samples of the relative activity of a WTA population selective for one of the two concurrent stimuli. The continuous blue line is the interpolation of these samples. The dashed blue line indicates the relative activity of the competing population. The green line is the angle bisector given by $x = y$. The horizontal and vertical axes show the relative contribution of two concurrent stimuli (two orthogonal orientations) and the corresponding populations (see legend of Fig. 8 for a detailed description). If there was no competition between the populations, simulation samples would lie on the green line, because then the network simply mirrors its input.
(TIFF)

Video S1 Neurite outgrowth in the 2-dimensional plane, with concentration-dependent axonal retraction. Excitatory and inhibitory neurons (colored red and purple, respectively) are initially randomly positioned on a 2-dimensional unprepared environment. The somata of both types secrete different substances, which are sensed by the growth cones at the tip of the axons. Whenever the sensed concentration falls below a predefined threshold, axons retract until they reach a high enough concentration (this retraction is indicated in green). This behavior is iteratively instantiated, allowing the network to project more efficiently, because axons do not grow into regions where no potential targets are located.
(WMV)

Video S2 Clustering of functional connectivity in a WTA network. The presence of functional connections among excitatory and inhibitory neurons (red and blue respectively) are indicated with arrows. For clearer visualization, the strength is not shown. A clustering algorithm was applied to move the nodes such that strong connections are more probable to be close to each other. Therefore, the video does not show any physical movement, but only the arrangements performed by the clustering algorithm in weight space. 4 input stimuli referring to horizontal, vertical and both diagonal orientations are presented to the network. In the first part of the video (until 0:07 min), all neurons do synaptic scaling. Subsequently, synapses onto excitatory neurons become

subject to the BCM learning rule, which has impact on the clustering of the functional connectivity: 4 clusters emerge for excitatory neurons, in contrast to the inhibitory neurons. This discrepancy is because of the different learning rule simulated after the first part, which is BCM learning for synapses onto excitatory and synaptic scaling for synapses onto inhibitory postsynaptic neurons.
(MOV)

References

- Yuille A, Grzywacz N (1989) A winner-take-all mechanism based on presynaptic inhibition feedback. *Neural Comput* 1: 334–347.
- Yuille A, Geiger D (2003) Winner-take-all networks. *The handbook of brain theory and neural networks*: 1228–1231.
- Douglas R, Martin K, Whitteridge D (1989) A canonical microcircuit for neocortex. *Neural Comput* 1: 480–488.
- Douglas RJ, Martin K (1991) A functional microcircuit for cat visual cortex. *J Physiol* 440: 735–769.
- Binzegger T, Douglas R, Martin K (2004) A quantitative map of the circuit of cat primary visual cortex. *J Neurosci* 24: 8441–8453.
- Douglas R, Martin K (2004) Neuronal circuits of the neocortex. *Annu Rev Neurosci* 27: 419–451.
- Maass W (2000) On the computational power of winner-take-all. *Neural Comput* 12: 2519–2535.
- Rutishauser U, Douglas R (2009) State-dependent computation using coupled recurrent networks. *Neural Comput* 21: 478–509.
- Indiveri G (2000) Modeling selective attention using a neuromorphic analog VLSI device. *Neural Comput* 12: 2857–2880.
- Hahnloser R, Douglas R, Hepp K (2002) Attentional recruitment of inter-areal recurrent networks for selective gain control. *Neural Comput* 14: 1669–1689.
- Johansson C, Rehn M, Lansner A (2006) Attractor neural networks with patchy connectivity. *Neurocomputing* 69: 627–633.
- Johansson C, Lansner A (2007) Towards cortex sized artificial neural systems. *Neural Networks* 20: 48–61.
- Jitsev J, von der Malsburg C (2009) Experience-driven formation of parts-based representations in a model of layered visual memory. *Front Comput Neurosci* 3: 15.
- Nessler B, Pfeiffer M, Buesing L, Maass W (2013) Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity. *PLoS Comput Biol* 9: e1003037.
- Pfeiffer M, Nessler B, Douglas R, Maass W (2010) Reward-modulated hebbian learning of decision making. *Neural Comput* 22: 1399–1444.
- Zubler F, Douglas R (2009) A framework for modeling the growth and development of neurons and networks. *Front Comput Neurosci* 3: 25.
- Turrigiano G, Leslie K, Desai N, Rutherford L, Nelson S (1998) Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391: 892–896.
- Bienenstock E, Cooper L, Munro P (1982) Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci* 2: 32–48.
- Wong R (1999) Retinal waves and visual system development. *Annu Rev Neurosci* 22: 29–47.
- Weliky M, Hall M (2000) Correlated neuronal activity minireview and visual cortical development. *Neuron* 27: 427–430.
- Butts D (2002) Retinal waves: implications for synaptic learning rules during development. *Neuroscientist* 8: 243–253.
- Huberman A, Feller M, Chapman B (2008) Mechanisms underlying development of visual maps and receptive fields. *Annu Rev Neurosci* 31: 479–509.
- Kirkby LA, Sack GS, Firl A, Feller MB (2013) A role for correlated spontaneous activity in the assembly of neural circuits. *Neuron* 80: 1129–1144.
- Zubler F, Hauri A, Pfister S, Whatley A, Cook M, et al. (2011) An instruction language for self-construction in the context of neural networks. *Front Comput Neurosci* 5: 57.
- Gabbott P, Somogyi P (1986) Quantitative distribution of GABA-immunoreactive neurons in the visual cortex (area 17) of the cat. *Exp Brain Res* 61: 323–331.
- Somogyi P, Tamas G, Lujan R, Buhl E (1998) Salient features of synaptic organisation in the cerebral cortex. *Brain Res Rev* 26: 113–135.
- Li J, Schwark H (1994) Distribution and proportions of GABA-Immunoreactive neurons in cat primary somatosensory cortex. *J Comp Neurol* 343: 353–361.
- LaMantia A, Rakic P (1990) Axon overproduction and elimination in the corpus callosum of the developing rhesus monkey. *J Neurosci* 10: 2156–2175.
- Godement P, Wang L, Mason C (1994) Retinal axon divergence in the optic chiasm: dynamics of growth cone behavior at the midline. *J Neurosci* 14: 7024–7039.
- Luo L, O'Leary D (2005) Axon retraction and degeneration in development and disease. *Annu Rev Neurosci* 28: 127–156.
- Portera-Cailliau C, Weimer R, De Paola V, Caroni P, Svoboda K (2005) Diverse modes of axon elaboration in the developing neocortex. *PLoS Biol* 3: e272.
- Binzegger T, Douglas R, Martin K (2010) An axonal perspective on cortical circuits. In: Feldmeyer, D. and Lübke, JHR. *New aspects of axonal structure and function*. New York (NY): Springer, 117–139 pp.
- Beaulieu C, Kisvarday Z, Somogyi P, Cynader M, Cowey A (1992) Quantitative distribution of GABA-immunopositive and-immunonegative neurons and synapses in the monkey striate cortex (area 17). *Cereb Cortex* 2: 295–309.
- Beaulieu C, Campistrone G, Crevier C (1994) Quantitative aspects of the GABA circuitry in the primary visual cortex of the adult rat. *J Comp Neurol* 339: 559–572.
- De Felipe J, Marco P, Fairen A, Jones E (1997) Inhibitory synaptogenesis in mouse somatosensory cortex. *Cereb Cortex* 7: 619–634.
- Dayan P, Abbott L (2001) *Theoretical neuroscience*, volume 31. MIT press Cambridge, MA.
- Azouz R, Gray C, Nowak L, McCormick D (1997) Physiological properties of inhibitory interneurons in cat striate cortex. *Cereb Cortex* 7: 534–545.
- Ahmed B, Anderson J, Douglas R, Martin K, Whitteridge D (1998) Estimates of the net excitatory currents evoked by visual stimulation of identified neurons in cat visual cortex. *Cereb Cortex* 8: 462–476.
- Ben-Yishai R, Bar-Or R, Sompolinsky H (1995) Theory of orientation tuning in visual cortex. *Proc Natl Acad Sci* 92: 3844–3848.
- Carandini M, Ringach D (1997) Predictions of a recurrent model of orientation selectivity. *Vision Res* 37: 3061–3071.
- Hahnloser R, Seung H, Slotine J (2003) Permitted and forbidden sets in symmetric threshold-linear networks. *Neural Comput* 15: 621–638.
- Hensch TK (2005) Critical period plasticity in local cortical circuits. *Nature Rev Neurosci* 6: 877–888.
- Espinosa JS, Stryker MP (2012) Development and plasticity of the primary visual cortex. *Neuron* 75: 230–249.
- Bednar JA, Miikkulainen R (2004) Prenatal and postnatal development of laterally connected orientation maps. *Neurocomputing* 58: 985–992.
- Srinivasan N, Jiang Q (2013) Stable learning of functional maps in self-organizing spiking neural networks with continuous synaptic plasticity. *Front Comput Neurosci* 7: 10.
- Turrigiano G, Nelson S (2004) Homeostatic plasticity in the developing nervous system. *Nature Rev Neurosci* 5: 97–107.
- Turrigiano G (2011) Too many cooks? Intrinsic and synaptic homeostatic mechanisms in cortical circuit refinement. *Annu Rev Neurosci* 34: 89–103.
- Van Rossum MC, Bi G, Turrigiano G (2000) Stable hebbian learning from spike timing-dependent plasticity. *J Neurosci* 20: 8812–8821.
- Liu J (2011) Learning rule of homeostatic synaptic scaling: Presynaptic dependent or not. *Neural Comput* 23: 3145–3161.
- Cooper LN, Intrator N, Blais BS, Shouval HZ (2004) *Theory of cortical plasticity*. World Scientific.
- Cooper L, Bear M (2012) The BCM theory of synapse modification at 30: interaction of theory with experiment. *Nature Rev Neurosci* 13: 798–810.
- Haas J, Nowotny T, Abarbanel H (2006) Spike-timing-dependent plasticity of inhibitory synapses in the entorhinal cortex. *J Neurophysiol* 96: 3305–3313.
- Ivshitz M, Segal M (2006) Simultaneous NMDA-dependent long-term potentiation of EPSCs and long-term depression of IPSCs in cultured rat hippocampal neurons. *J Neurosci* 26: 1199–1210.
- Saar D, Reuveni I, Barkai E (2012) Mechanisms underlying rule learning-induced enhancement of excitatory and inhibitory synaptic transmission. *J Neurophysiol* 107: 1222–1229.
- Vogels TP, Froemke RC, Doyon N, Gilson M, Haas JS, et al. (2013) Inhibitory synaptic plasticity: spike timing-dependence and putative network function. *Front Neural Circuits* 7: 119.

Acknowledgments

The authors wish to thank Matthew Cook, Nuno da Costa, Arko Ghosh, Dennis Goehlsdorf, Kevan Martin, Gabriela Michel, Emre Neftci, Ueli Rutishauser, Andreas Steimer, Christoph von der Malsburg and Adrian Whatley for stimulating discussions and useful comments.

Author Contributions

Conceived and designed the experiments: RB SP RJD. Performed the experiments: RB SP. Analyzed the data: RB SP. Contributed reagents/materials/analysis tools: RB FZ SP AH. Wrote the paper: RB FZ SP MP DRM RJD.

56. Markram H, Toledo-Rodriguez M, Wang Y, Gupta A, Silberberg G, et al. (2004) Interneurons of the neocortical inhibitory system. *Nature Rev Neurosci* 5: 793–807.
57. Batista-Brito R, Fishell G (2009) The developmental integration of cortical interneurons into a functional network. *Curr Top Dev Biol* 87: 81–118.
58. Busse L, Wade A, Carandini M (2009) Representation of concurrent stimuli by population activity in visual cortex. *Neuron* 64: 931.
59. Brouwer GJ, Heeger DJ (2011) Cross-orientation suppression in human visual cortex. *J Neurophysiol* 106: 2108–2119.
60. Tsai JJ, Wade AR, Norcia AM (2012) Dynamics of normalization underlying masking in human visual cortex. *J Neurosci* 32: 2783–2789.
61. Bathellier B, Ushakova L, Rumpel S (2012) Discrete neocortical dynamics predict behavioral categorization of sounds. *Neuron* 76: 435–449.
62. Von Neumann J (1958) *The computer and the brain*. Yale Univ. Press.
63. Kerlin A, Andermann M, Berezovskii V, Reid R (2010) Broadly tuned response properties of diverse inhibitory neuron subtypes in mouse visual cortex. *Neuron* 67: 858–871.
64. Niell C, Stryker M (2008) Highly selective receptive fields in mouse visual cortex. *J Neurosci* 28: 7520–7536.
65. Hofer S, Ko H, Pichler B, Vogelstein J, Ros H, et al. (2011) Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nat Neurosci* 14: 1045–1052.
66. Zariwala H, Madisen L, Ahrens K, Bernard A, Levin E, et al. (2010) Visual tuning properties of genetically identified layer 2/3 neuronal types in the primary visual cortex of cre-transgenic mice. *Front Syst Neurosci* 4: 162.
67. Kuhlman S, Tring E, Trachtenberg J (2011) Fast-spiking interneurons have an initial orientation bias that is lost with vision. *Nat Neurosci* 14: 1121–1123.
68. Mruczek R, Sheinberg D (2012) Stimulus selectivity and response latency in putative inhibitory and excitatory neurons of the primate inferior temporal cortex. *J Neurophysiol* 108: 2725–2736.
69. Azouz R, Gray C, Nowak L, McCormick D (1997) Physiological properties of inhibitory interneurons in cat striate cortex. *Cereb Cortex* 7: 534–545.
70. Hirsch J, Martinez L, Pillai C, Alonso J, Wang Q, et al. (2003) Functionally distinct inhibitory neurons at the first stage of visual cortical processing. *Nat Neurosci* 6: 1300–1308.
71. Cardin J, Palmer L, Contreras D (2007) Stimulus feature selectivity in excitatory and inhibitory neurons in primary visual cortex. *J Neurosci* 27: 10333–10344.
72. Nowak L, Sanchez-Vives M, McCormick D (2008) Lack of orientation and direction selectivity in a subgroup of fast-spiking inhibitory interneurons: cellular and synaptic mechanisms and comparison with other electrophysiological cell types. *Cereb Cortex* 18: 1058–1078.
73. Somogyi P, Martin K (1985) Cortical circuitry underlying inhibitory processes in cat area 17. *Models of the Visual Cortex*: 514–523.
74. Grabzka-Barwińska A, von der Malsburg C (2008) Establishment of a scaffold for orientation maps in primary visual cortex of higher mammals. *J Neurosci* 28: 249–257.
75. Lazar A, Pipa G, Triesch J (2009) SORN: a self-organizing recurrent neural network. *Front Comput Neurosci* 3: 23.
76. Bednar JA (2012) Building a mechanistic model of the development and function of the primary visual cortex. *J Physiol Paris* 106: 194–211.
77. Zubler F, Hauri A, Pfister S, Bauer R, Anderson JC, 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.
78. Bauer R, Zubler F, Hauri A, Muir D, Douglas R (2014) Developmental origin of patchy axonal connectivity in the neocortex: A computational model. *Cereb Cortex* 24: 487–500.
79. Li Yt, Ibrahim LA, Liu Bh, Zhang LI, Tao HW (2013) Linear transformation of thalamocortical input by intracortical excitation. *Nat Neurosci* 16: 1324–1330.
80. Lien AD, Scanziani M (2013) Tuned thalamic excitation is amplified by visual cortical circuits. *Nat Neurosci* 16: 1315–1323.
81. Hebb D (1949) *The organization of behavior*, volume 70. Wiley, New York, 71–72 pp.
82. Harris KD (2012) Cell assemblies of the superficial cortex. *Neuron* 76: 263–265.
83. Song S, Sjöström P, Reigl M, Nelson S, Chklovskii D (2005) Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol* 3: e68.
84. Yoshimura Y, Callaway E (2005) Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity. *Nat Neurosci* 8: 1552–1559.
85. Perin R, Berger T, Markram H (2011) A synaptic organizing principle for cortical neuronal groups. *Proc Natl Acad Sci* 108: 5419–5424.
86. Ko H, Hofer S, Pichler B, Buchanan K, Sjöström P, et al. (2011) Functional specificity of local synaptic connections in neocortical networks. *Nature* 473: 87–91.
87. Peters A, Feldman M (1976) The projection of the lateral geniculate nucleus to area 17 of the rat cerebral cortex. I. general description. *J neurocytol* 5: 63–84.
88. Braitenberg V, Schüz A (1991) *Anatomy of the cortex: Statistics and geometry*. Springer-Verlag Publishing, 109–112 pp.
89. Peters A, Payne B (1993) Numerical relationships between geniculocortical afferents and pyramidal cell modules in cat primary visual cortex. *Cereb Cortex* 3: 69–78.
90. Monyer H, Markram H (2004) *Interneuron Diversity series: Molecular and genetic tools to study GABAergic interneuron diversity and function*. *Trends Neurosci* 27: 90–97.
91. Wonders CP, Anderson SA (2006) The origin and specification of cortical interneurons. *Nature Rev Neurosci* 7: 687–696.
92. Beurl RL (1956) Properties of a mass of cells capable of regenerating pulses. *Philos T Roy Soc B* 240: 55–94.
93. Griffith J (1963) A field theory of neural nets: I. Derivation of field equations. *B Math Biophys* 25: 111–120.
94. Griffith JS (1965) A field theory of neural nets: II. properties of the field equations. *B Math Biophys* 27: 187–195.
95. Alonso J, Swadlow H (2005) Thalamocortical specificity and the synthesis of sensory cortical receptive fields. *J Neurophysiol* 94: 26–32.
96. Yoshimura Y, Dantzker J, Callaway E (2005) Excitatory cortical neurons form fine-scale functional networks. *Nature* 433: 868–873.
97. Wang L, Kloc M, Gu Y, Ge S, Maffei A (2013) Layer-specific experience-dependent rewiring of thalamocortical circuits. *J Neurosci* 33: 4181–4191.
98. Vogels T, Sprekeler H, Zenke F, Clopath C, Gerstner W (2011) Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* 334: 1569–1573.
99. Sohya K, Kameyama K, Yanagawa Y, Obata K, Tsumoto T (2007) GABAergic neurons are less selective to stimulus orientation than excitatory neurons in layer II/III of visual cortex, as revealed by in vivo functional Ca2+ imaging in transgenic mice. *J Neurosci* 27: 2145–2149.
100. Liu B, Li P, Li Y, Sun Y, Yanagawa Y, et al. (2009) Visual receptive field structure of cortical inhibitory neurons revealed by two-photon imaging guided recording. *J Neurosci* 29: 10520–10532.
101. Bock D, Lee W, Kerlin A, Andermann M, Hood G, et al. (2011) Network anatomy and in vivo physiology of visual cortical neurons. *Nature* 471: 177–182.
102. Bloodgood B, Giessel A, Sabatini B (2009) Biphasic synaptic Ca influx arising from compartmentalized electrical signals in dendritic spines. *PLoS Biol* 7: e1000190.
103. Kasai H, Hayama T, Ishikawa M, Watanabe S, Yagishita S, et al. (2010) Learning rules and persistence of dendritic spines. *Eur J Neurosci* 32: 241–249.
104. Chiu CQ, Lur G, Morse TM, Carnevale NT, Ellis-Davies GC, et al. (2013) Compartmentalization of gabaergic inhibition by dendritic spines. *Science* 340: 759–762.
105. Lee WCA, Huang H, Feng G, Sanes JR, Brown EN, et al. (2005) Dynamic remodeling of dendritic arbors in GABAergic interneurons of adult visual cortex. *PLoS Biol* 4: e29.
106. Oster M, Douglas R, Liu SC (2009) Computation with spikes in a winner-take-all network. *Neural Comput* 21: 2437–2465.
107. Chen Y, McKinstry JL, Edelman GM (2013) Versatile networks of simulated spiking neurons displaying winner-take-all behavior. *Front Comput Neurosci* 7: 16.
108. Caporale N, Dan Y (2008) Spike timing-dependent plasticity: a hebbian learning rule. *Annu Rev Neurosci* 31: 25–46.
109. Clopath C, Büsing L, Vasilaki E, Gerstner W (2010) Connectivity reflects coding: a model of voltage-based stdp with homeostasis. *Nature Neurosci* 13: 344–352.
110. Pfister J, Gerstner W (2006) Triplets of spikes in a model of spike timing-dependent plasticity. *J Neurosci* 26: 9673–9682.
111. Bush D, Philippides A, Husbands P, O'Shea M (2010) Reconciling the stdp and bcm models of synaptic plasticity in a spiking recurrent neural network. *Neural Comput* 22: 2059–2085.
112. Gjorgieva J, Clopath C, Audet J, Pfister JP (2011) A triplet spike-timing-dependent plasticity model generalizes the biennstock-cooper-munro rule to higher-order spatiotemporal correlations. *Proc Natl Acad Sci* 108: 19383–19388.
113. Rockland K, Lund J (1982) Widespread periodic intrinsic connections in the tree shrew visual cortex. *Science* 215: 1532–1534.
114. Rockland K, Lund J, Humphrey A (1982) Anatomical banding of intrinsic connections in striate cortex of tree shrews (*Tupaia glis*). *J Comp Neurol* 209: 41–58.
115. Lund J, Angelucci A, Bressloff P (2003) Anatomical substrates for functional columns in macaque monkey primary visual cortex. *Cereb Cortex* 13: 15–24.
116. Muir DR, Douglas RJ (2011) From neural arbors to daisies. *Cereb Cortex* 21: 1118–1133.
117. Alon U (2006) *An introduction to systems biology: design principles of biological circuits*. CRC Press.
118. Pfister S (2013) *Inference of developmental motifs in the developing mouse cerebral cortex*. Ph.D. thesis, ETH Zuerich.
119. Warren M, Bedi K (1984) A quantitative assessment of the development of synapses and neurons in the visual cortex of control and undernourished rats. *J Comp Neurol* 227: 104–108.
120. Leuba G, Garey L (1989) Comparison of neuronal and glial numerical density in primary and secondary visual cortex of man. *Exp Brain Res* 77: 31–38.
121. Miki T, Fukui Y, Itoh M, Hisano S, Xie Q, et al. (1997) Estimation of the numerical densities of neurons and synapses in cerebral cortex. *Brain Res Prot* 2: 9–16.
122. Lübke J, Markram H, Frotscher M, Sakmann B (1996) Frequency and dendritic distribution of autapses established by layer 5 pyramidal neurons in

- the developing rat neocortex: comparison with synaptic innervation of adjacent neurons of the same class. *J Neurosci* 16: 3209–3218.
123. Hellwig B (2000) A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biol Cybern* 82: 111–121.
 124. Binzegger T, Douglas R, Martin K (2009) Topology and dynamics of the canonical circuit of cat V1. *Neural Networks* 22: 1071–1078.
 125. Trappenberg TP (2010) *Fundamentals of computational neuroscience*. Oxford University Press.
 126. Smoot M, Ono K, Ruscheinski J, Wang P, Ideker T (2011) Cytoscape 2.8: new features for data integration and network visualization. *Bioinformatics* 27: 431–432.
 127. Kamada T, Kawai S (1989) An algorithm for drawing general undirected graphs. *Inform Process Lett* 31: 7–15.
 128. Greenberg D, Houweling A, Kerr J (2008) Population imaging of ongoing neuronal activity in the visual cortex of awake rats. *Nat Neurosci* 11: 749–751.
 129. De Kock C, Sakmann B (2009) Spiking in primary somatosensory cortex during natural whisking in awake head-restrained rats is cell-type specific. *Proc Natl Acad Sci* 106: 16446–16450.
 130. Swadlow HA (1989) Efferent neurons and suspected interneurons in S-1 vibrissa cortex of the awake rabbit: receptive fields and axonal properties. *J Neurophysiol* 62: 288–308.
 131. Povysheva N, Zaitsev A, Rotaru D, Gonzalez-Burgos G, Lewis D, et al. (2008) Parvalbumin-positive basket interneurons in monkey and rat prefrontal cortex. *J Neurophysiol* 100: 2348–2360.
 132. Fiala JC, Harris KM (1999) Dendrite structure. *Dendrites*: 1–34.
 133. Binzegger T, Douglas RJ, Martin KA (2005) Axons in cat visual cortex are topologically self-similar. *Cereb Cortex* 15: 152–165.
 134. Stepanyants A, Martinez LM, Ferecskó AS, Kisvárdy ZF (2009) The fractions of short-and long-range connections in the visual cortex. *Proc Natl Acad Sci* 106: 3555–3560.