

Feedback Inhibition Shapes Emergent Computational Properties of Cortical Microcircuit Motifs

Zeno Jonke,*  Robert Legenstein,* Stefan Habenschuss, and  Wolfgang Maass

Institute for Theoretical Computer Science, Graz University of Technology, Inffeldgasse 16b/I, 8010 Graz, Austria

Cortical microcircuits are very complex networks, but they are composed of a relatively small number of stereotypical motifs. Hence, one strategy for throwing light on the computational function of cortical microcircuits is to analyze emergent computational properties of these stereotypical microcircuit motifs. We are addressing here the question how spike timing-dependent plasticity shapes the computational properties of one motif that has frequently been studied experimentally: interconnected populations of pyramidal cells and parvalbumin-positive inhibitory cells in layer 2/3. Experimental studies suggest that these inhibitory neurons exert some form of divisive inhibition on the pyramidal cells. We show that this data-based form of feedback inhibition, which is softer than that of winner-take-all models that are commonly considered in theoretical analyses, contributes to the emergence of an important computational function through spike timing-dependent plasticity: The capability to disentangle superimposed firing patterns in upstream networks, and to represent their information content through a sparse assembly code.

Key words: computational function; cortical microcircuits; divisive inhibition; feedback inhibition; synaptic plasticity; winner-take-all

Significance Statement

We analyze emergent computational properties of a ubiquitous cortical microcircuit motif: populations of pyramidal cells that are densely interconnected with inhibitory neurons. Simulations of this model predict that sparse assembly codes emerge in this microcircuit motif under spike timing-dependent plasticity. Furthermore, we show that different assemblies will represent different hidden sources of upstream firing activity. Hence, we propose that spike timing-dependent plasticity enables this microcircuit motif to perform a fundamental computational operation on neural activity patterns.

Introduction

A promising strategy for understanding the computational function of a cortical column was proposed by Douglas et al. (1989), Shepherd (2004), Grillner and Graybiel (2006), and others. To probe computational properties of prominent network motifs of a cortical column, commonly referred to as microcircuit motifs. We are addressing computational properties of one of the most prominent microcircuit motifs: densely interconnected populations of excitatory and inhibitory neurons. We focus on motifs in layer 2/3, where parvalbumin-positive (PV⁺) inhibitory neurons (often characterized as fast-spiking interneurons, in particular basket cells) are interconnected with nearby

pyramidal cells with very high connection probability in both directions (see, e.g., Packer and Yuste, 2011; Avermann et al., 2012; Fino et al., 2013). One usually refers to this type of inhibition as lateral or feedback inhibition. The dynamics of this microcircuit motif has frequently been examined *in vivo* (Wilson et al., 2012; Petersen and Crochet, 2013; Pala and Petersen, 2015) and modeled in (Avermann et al., 2012). We examine computational properties that emerge in model M for this microcircuit motif under spike timing-dependent plasticity (STDP).

One cannot model this microcircuit motif by the frequently considered winner-take-all (WTA) model because this model would require that the firing of a single pyramidal cell (the “winner”) can suppress firing of other pyramidal cells in the motif. But experimental data show that several pyramidal cells need to fire to engage feedback inhibition through PV⁺ cells (Isaacson and Scanziani, 2011; Avermann et al., 2012). Divisive inhibition has been proposed as a more realistic mathematical model for this softer type of inhibition (Carandini and Heeger, 2011; Wilson et al., 2012). Our goal is to understand the impact of this softer type of inhibition on neural codes and computational properties that emerge under STDP. There exists a large number of preceding studies of emergent computational properties of WTA-like mi-

Received June 27, 2016; revised July 18, 2017; accepted July 23, 2017.

Author contributions: Z.J., R.L., S.H., and W.M. designed research; Z.J., R.L., and W.M. performed research; Z.J., R.L., and S.H. analyzed data; Z.J., R.L., and W.M. wrote the paper.

This work was supported in part by Human Brain Project of the European Union #604102 and #720270, and the Austrian Science Fund I3251-N33. We thank Carl Petersen for helpful discussions.

The authors declare no competing financial interests.

*Z.J. and R.L. contributed equally to the work.

Correspondence should be addressed to Dr. Robert Legenstein, Institute for Theoretical Computer Science, Graz University of Technology, Inffeldgasse 16b/I, 8010 Graz, Austria. E-mail: robert.legenstein@igi.tugraz.at.

DOI:10.1523/JNEUROSCI.2078-16.2017

Copyright © 2017 the authors 0270-6474/17/378511-13\$15.00/0

microcircuit motifs, from Rumelhart and Zipser (1985) to Nessler et al. (2013). But they were based on the assumption of strong WTA-like lateral inhibition.

The functional role of inhibition in this microcircuit motif can be better approximated by a variation of the k -WTA model (Maass, 2000), where several (k) winners can emerge simultaneously from a competition of pyramidal cells for firing. We show that this softer competition leads to the emergence of shared feature selectivity of pyramidal cells, as in the experimental data of Lee et al. (2012), where small subsets of pyramidal cells (assemblies), instead of single neurons, respond to specific input features (see Fig. 2).

We also show that an important computational operation, blind source separation (Földiák, 1990), also referred to as independent component analysis (ICA) (Hyvärinen et al., 2004), emerges in this microcircuit motif through STDP. This operation enables a network to disentangle and sparsely represent superimposed spike inputs that may result from separate sources in the environments or upstream neural networks. This modular coding scheme avoids a combinatorial explosion of the number of neurons that are needed to encode superimposed sources because they become encoded by superpositions of neural codes (assemblies) for each of the sources, rather than by a separate neural code for every superposition that occurs. An example is given in Figures 4 and 5 for the case of arbitrarily superimposed vertical and horizontal bars, a well-known benchmark task for blind source separation (Földiák, 1990). This distributed coding scheme also supports intracortical communication and computation based on spike patterns or spike packets as proposed by Luczak et al. (2015) (see Figure 3).

Materials and Methods

Definition of a data-based microcircuit motif model M

We consider in this article a model for interacting populations of pyramidal cells with PV⁺ inhibitory neurons on layer 2/3 that is based on data from the Petersen laboratory (Avermann et al., 2012) and refer to this specific model as the microcircuit motif model M .

The microcircuit motif model M consists of two reciprocally connected pools of neurons: an excitatory pool and an inhibitory pool. Inhibitory network neurons are recurrently connected. Excitatory network neurons receive additional excitatory synaptic input from a pool of N input neurons. Figure 1A summarizes the connectivity structure of the model together with connection probabilities. Connection probabilities have been chosen according to the experimental data described by Avermann et al. (2012) and listed in Table 1 together with connection-type specific synaptic parameters. For a connection probability p between two pools, each individual pair of neurons from these two pools is randomly chosen to be connected by a synapse with probability p .

Input neurons emit Poisson spike trains with time-varying rates. We tested several temporal profiles of these rates in different simulations as described below in the corresponding sections. Let $t_i^{(1)}, t_i^{(2)}, \dots$ denote the spike times of input neuron i . The output trace $\tilde{y}_i(t)$ of input neuron i is given by the temporal sum of unweighted postsynaptic potentials (PSPs) arising from input neuron i as follows:

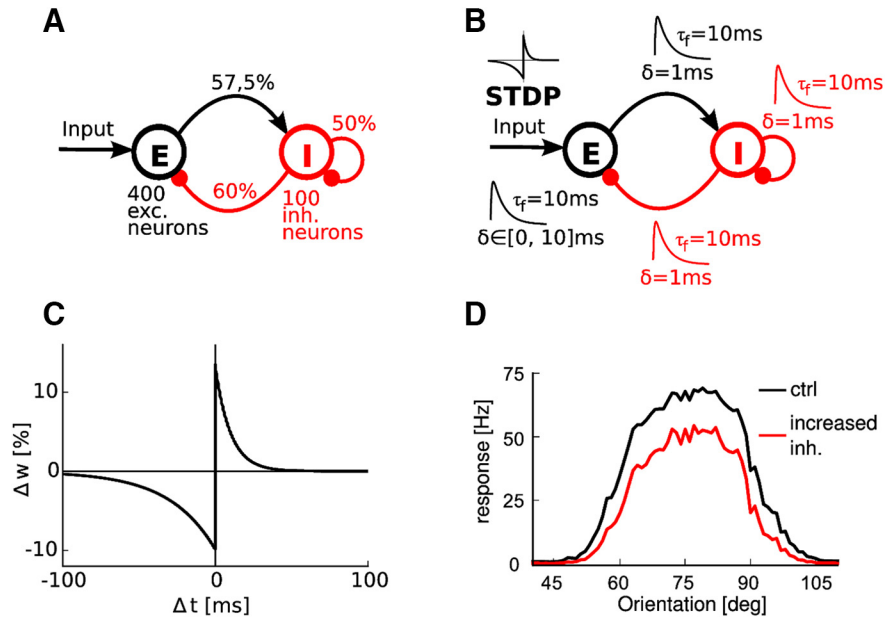


Figure 1. A data-based microcircuit motif model M . **A**, Network anatomy. Circles represent excitatory (black) and inhibitory (red) pools of neurons. Black arrows indicate excitatory connections. Red lines with dots indicate inhibitory connections. Numbers above connections indicate corresponding connection probabilities. **B**, Network physiology. Same as in **A**, but connection delays δ and PSP shapes with decay time constant τ_r are indicated for synaptic connections. Input synapses are subject to STDP. **C**, Standard STDP curve that is used in M . Shown is the change of the synaptic efficacy in our model for 10 pre-post pairings in dependence on the time difference $\Delta t = t_{\text{post}} - t_{\text{pre}}$ between a postsynaptic spike at time t_{post} and a presynaptic spike at time t_{pre} . **D**, Divisive normalization in model M . The response of an excitatory neuron in the circuit to a visual bar-stimulus at various orientations (for details, see Materials and Methods) in control condition (black) and for simulated increased firing of inhibitory neurons (red). Note the divisive nature of inhibition (stronger responses are more strongly depressed in absolute terms). Compare with Wilson et al. (2012, their Figs. 2e,3f).

Table 1. Neuron-type specific synaptic connection parameters in M^a

Connection	Connection probability (%)	Symbol	Synaptic weight (a.u.)	Synaptic delay (ms)
Input \rightarrow E	100	w	0.01, 1	0, 10
E \rightarrow I	57.5	w^{EI}	13.57	1
I \rightarrow E	60	w^{IE}	1.86	1
I \rightarrow I	55	w^{II}	13.57	1

^aConnection probability, synaptic weight, and synaptic delay. Weights from inputs to excitatory network neurons are plastic and bounded to the given range. The corresponding delays are uniformly distributed in the given range. Synaptic efficacies can be motivated from a theoretical perspective (see Discussion).

$$\tilde{y}_i(t) = \sum_j \epsilon(t - t_i^{(j)}), \quad (1)$$

where ϵ is the synaptic response kernel (i.e., the shape of the PSP). It is given by a double-exponential function as follows:

$$\epsilon(s) = \begin{cases} c_e(e^{-s/\tau_r} - e^{-s/T_e}), & \text{if } 0 \leq s \leq T_e \\ 0, & \text{otherwise} \end{cases}, \quad (2)$$

with the rise time constant $\tau_r = 1$ ms, a fall time constant $\tau_f = 10$ ms, and a cutoff after $T_e = 50$ ms (see also Figure 1B). The constant $c_e = 1.435$ was chosen to assure a peak value of 1. All synapses in the network have the same response kernel ϵ . For given spike times, output traces of excitatory network neurons and inhibitory network neurons are defined analogously and denoted by $\tilde{z}_m(t)$ and $I_j(t)$, respectively.

The network consists of 400 excitatory neurons, modeled as stochastic spike response model neurons (Jolivet et al., 2006) that we define in the following. The stochasticity of the model stems from its stochastic spike generation, where spikes are generated according to a Poisson process with a time-varying rate (the instantaneous firing rate of the neuron). The instantaneous firing rate ρ_m of a neuron m depends exponentially on its current membrane potential u_m as follows:

$$\rho_m(t) = \frac{1}{\tau} \exp(\gamma \cdot u_m(t)), \quad (3)$$

where $\tau = 10$ ms and $\gamma = 2$ are scaling parameters that control the shape of the response function. After emitting a spike, the neuron enters an absolute refractory period for 10 ms during which the neuron cannot spike again. These excitatory neurons project to and receive inputs from a pool of inhibitory neurons. Thus, the membrane potential of excitatory neuron m is given by the sum of external inputs, inhibition from inhibitory neurons, and its excitability α as follows:

$$u_m(t) = \sum_i w_{im} \bar{y}_i(t) - \sum_{j \in \mathcal{J}_m} w^{IE} I_j(t) + \alpha, \quad (4)$$

where \mathcal{J}_m denotes the set of indices of inhibitory neurons that project to neuron m , and w^{IE} denotes the weight of these inhibitory synapses. $I_j(t)$ and $\bar{y}_i(t)$ denote synaptic input (output traces) from inhibitory neurons and input neurons, respectively (see above). We used $\alpha = -5.57$. These parameter values can be motivated from a theoretical perspective (Legenstein et al., 2017).

Apart from excitatory neurons there are 100 inhibitory neurons in the network. While Jolivet et al. (2006) provide a stochastic model for pyramidal cells, no such model is available for PV⁺ inhibitory neurons. Experimental data indicate that, in these neurons, the relationship between the synaptic drive and the firing rate, that is, the frequency-current (f-I) curve, is rather linear over a large range of input strengths (Ho et al., 2012; Ferguson et al., 2013). We therefore modeled inhibitory neurons as stochastic spike response neurons with an instantaneous firing rate given by the following:

$$\rho_m^{\text{inh}}(t) = \sigma_{\text{rect}}(u_m^{\text{inh}}(t)), \quad (5)$$

where σ_{rect} denotes the linear rectifying function $\sigma_{\text{rect}}(u) = u$ for $u \geq 0$ and 0 otherwise. The absolute refractory period of inhibitory neurons in the model is 3 ms. Inhibitory neurons receive excitatory inputs from excitatory network neurons as well as connections from other inhibitory neurons. The membrane potentials of inhibitory neurons are thus given by the following:

$$u_m^{\text{inh}}(t) = \sum_{i \in \mathcal{E}_m} w^{\text{EI}} \bar{z}_i(t) - \sum_{j \in \mathcal{J}_m} w^{\text{II}} I_j(t) + u_{\text{opt}}, \quad (6)$$

where $\bar{z}_i(t)$ denotes synaptic input (output trace) from excitatory network neuron i , \mathcal{E}_m (\mathcal{J}_m) denotes the set of indices of excitatory (inhibitory) neurons that project to inhibitory neuron m , w^{EI} (w^{II}) denotes the excitatory (inhibitory) weight to inhibitory neurons, and u_{opt} denotes an external optogenetic activation of inhibitory neurons. u_{opt} was set to 0 in all simulations except for the simulation shown in Figure 1D, where optogenetic activation was modeled by setting $u_{\text{opt}} = 50$ (arbitrary units). The synaptic weights from excitatory network neurons to inhibitory neurons imply that a single spike in the excitatory pool induces a spike in a given postsynaptic inhibitory neuron with a probability of 0.17, consistent with experimental findings that several excitatory neurons have to be active to induce robust spiking in PV⁺ interneurons (Avermann et al., 2012).

Synaptic connections from input neurons to excitatory network neurons are subject to STDP. A standard version of STDP is used with an exponential weight dependency for potentiation (Habenschuss et al., 2013b) (Fig. 1C). All input weights w_{ij} are updated as follows. For each postsynaptic spike at time t_{post} , all presynaptic spikes in the preceding 100 ms are considered. For each such pre-before-after spike pair with time difference $t_{\text{post}} - t_{\text{pre}}$, the weight is increased by the following:

$$\Delta w_{ij}(t_{\text{post}} - t_{\text{pre}}) = \eta e^{-w_{ij}+1} e^{-\frac{t_{\text{post}} - t_{\text{pre}}}{\tau_+}}, \quad (7)$$

with $\tau_+ = 10$ ms. The learning rate η is 0.01, except for Figure 4 where $\eta = 0.02$ to speed up learning. For each presynaptic spike at time t , all postsynaptic spikes in the preceding 100 ms are considered. For each such postspike-before-prespike pair with time difference $t_{\text{pre}} - t_{\text{post}}$, the weight change is given by the following:

$$\Delta w_{ij}(t_{\text{pre}} - t_{\text{post}}) = -\eta e^{-\frac{t_{\text{pre}} - t_{\text{post}}}{\tau_-}}, \quad (8)$$

with $\tau_- = 25$ ms. Synaptic weights are clipped to $w_{\text{min}} = 0.01$ and $w_{\text{max}} = 1$. In all simulations, initial input weights were drawn from a uniform distribution in the interval $(w_{\text{min}}, w_{\text{max}})$. This concludes the definition of the microcircuit motif model M.

Details to computer simulations of model M

Here, we provide details to the computer simulations reported in Results. It is recommended that the reader skip this section at first reading. References to the individual subsections are given at the appropriate places in Results.

All simulations were performed in PCSIM, a spiking neural network simulator written in C++ that provides a Python interface, which was extended to support simulation of the model. All simulations were performed with a discretization time step Δt of 1 ms. Simulation code is available at <https://github.com/zjonke/EImotif>.

Details to simulations for Figure 1

For Figure 1D, the input to M was given by simulated visual bars stimuli at various orientations (see Details to simulations for Fig. 2). Orientation-tuned neurons emerged in a learning phase that lasted 400 s of simulated biological time. The tuning curve of one excitatory neuron was evaluated in the original circuit. Then, optogenetic stimulation of inhibitory neurons was mimicked by setting the external activation u_{opt} in Equation 6 to $u_{\text{opt}} = 50$ (arbitrary units) in all inhibitory neurons, with the effect of increasing the total rate of inhibition. The tuning curve of the same excitatory neuron was then evaluated in this modified network. All procedures in the learning and evaluation phase were the same as described below in Details to simulations for Figure 2.

Details to simulations for Figure 2

Here, we tested the behavior of M on an input distribution that mimics visual bar patterns of various orientations. In this simulation, network inputs were generated from 180 2D binary pixel arrays of size 20×20 . A prototypical horizontal bar of width 2 pixels centered on the array was rotated in steps of 1 degree to obtain 180 pixel arrays that span the space of possible bar orientations. These pixel arrays were then transformed into 400-dimensional rate vectors where each entry had a rate of 75 Hz if the corresponding pixel was on (i.e., the bar covered that pixel) and 1 Hz otherwise. During a simulation, a rate vector was chosen randomly (uniformly out of the 180 vectors). The i^{th} component of this rate vector then defined the firing rate of input neuron i to the network. One rate vector was presented to the network for 50 ms. During this time, input neurons produced Poisson spike trains with the rate as defined in the corresponding entry of the chosen rate vector. Between the presentation of two consecutive bar patterns, all input neurons spiked with a rate of 2 Hz for a duration drawn from a geometric distribution with a mean of 50 simulation time steps Δt , corresponding to 50 ms simulated biological time.

During the learning phase, the network was presented with such patterns for 400 s. In a testing phase, STDP in the network was disabled and input patterns were presented to the network in the same manner as in the training phase for 100 h of simulated time. Average firing rates of excitatory and inhibitory neurons were computed conditioned on specific bar orientations for Figure 2E–G.

In the simulation for H, the same network input was presented to a WTA network model proposed by Nessler et al. (2013). This model was termed spike-based expectation maximization (SEM) network. We used a model consisting of 400 neurons that competed in a WTA-like manner (for details on the model, see Nessler et al., 2013). The SEM network was simulated with a time step of 1 ms with rectangular PSPs of length 10 ms, a total output rate of 100 Hz, initial weights chosen from a uniform distribution in $[-0.5, 0.5]$, nonadaptive biases of 0, and a learning rate of $\eta = 0.02$ (for details on the simulated SEM model, see Habenschuss et al., 2013b). The learning phase and the testing phase were performed in the same manner as for model M.

Details to simulations for Figure 3

Here we tested our microcircuit motif model M on input that was created by the nonlinear superposition of 150-ms-long spatiotemporal patterns.

Creation of basic rate patterns. Input spike trains to the circuit were created by the superposition of two basic rate patterns. We first describe the creation of basic patterns; the superposition of these patterns will be discussed below.

Let \mathbf{R}^i denote the i^{th} basic rate pattern. Formally, a rate pattern \mathbf{R}^i is a matrix $\mathbf{R}^i = [r_{n,s}^i]_{n=1, \dots, N; s=1, \dots, S}$ with $r_{n,s}^i$ denoting the firing rate of the pattern in channel n at frame s and S is the number of frames of the pattern. Each frame defines the firing rates of channels (corresponding to the rates of input neurons) for a discrete time bin of length $\Delta t = 1$ ms.

For Figure 3, we defined a set of two basic rate patterns $\mathbf{R}^1, \mathbf{R}^2$, each consisting of $N = 200$ channels with $S = 150$ frames (i.e., the length of basic patterns was 150 ms). The firing rate for each channel was obtained by an Ornstein–Uhlenbeck (OU) process drawn independently for each pattern and each channel. More precisely, it was calculated as $r_{j,s}^i = 1.5 \exp(x_{j,s}^i \Delta t)$, where $x_{j,t}^i$ is given by a maximum-bounded OU process. The maximum-bounded OU process for a variable x_t is defined as $dx_t = \Theta_{OU}(\mu_{OU} - x_t)dt + \sigma_{OU}dW_t$ if $x_t < \log(50)$ and $dx_t = 0$ otherwise. Here, t denotes continuous time, $\Theta_{OU} > 0$ is the changing rate (speed), $\mu_{OU} > 0$ is the mean, $\sigma_{OU} > 0$ is the noise variance, and W_t is the standard Wiener process. The parameters were $\mu_{OU} = 0$, $\Theta_{OU} = 5$, and $\sigma_{OU} = 0.5$. The initial values for x_t in the OU process were drawn from a normal distribution with zero mean and unit variance. The first 50 ms of the OU process were discarded.

Superposition of basic rate patterns. Input spike trains were created by the superposition of a number of patterns, or more precisely their rates, from the set of basic patterns $P = \{\mathbf{R}^1, \mathbf{R}^2, \dots\}$. Because the procedure will below also be used for the superposition of bar patterns, we describe it here for an arbitrary number of basic patterns.

We first describe the procedure that determines which basic patterns to be superimposed at which times (i.e., the timing of bars in Fig. 3B, top). Given is a set of basic patterns $P = \{\mathbf{R}^1, \mathbf{R}^2, \dots\}$, each pattern of length S time steps. Let n_{\max} denote the maximum number of basic patterns that can be superimposed at any time t . We define n_{\max} registers $v_1, \dots, v_{n_{\max}}$. Each register holds at any time step t either no pattern (empty register) or one basic pattern, with the constraint that the registers hold different patterns at any given time step t . The following procedure ensures that at any time, the probability that a given register holds some pattern is p_{loaded} . At time step t , each empty register v_i is loaded with some pattern independently from other time steps and other registers

with probability $\frac{1}{1 + S(1 - p_{\text{loaded}})/p_{\text{loaded}}}$. If a register is loaded at time step t , the basic pattern to be loaded to this register is chosen uniformly from the set of basic patterns that are currently not held by any register. This basic pattern is then kept in the register for the length of its duration S (afterward the register is empty, but can be loaded again right away). Whether a basic pattern is in register v_i or v_j at some time t is irrelevant with respect to the produced superimposed patterns.

This defines for any time step t , which basic patterns are to be superimposed and also the frame at which each of these patterns is at that time. Superposition of basic patterns is then accomplished as described above to obtain the rate for each input neuron. Poisson spike trains are drawn from the resulting rates.

Patterns were first superimposed linearly; then a nonlinearity was applied. Consider a time t when a set of patterns should be superimposed (these patterns overlap at this time point). For the linear superposition, the rate of a particular channel in the superposition at time t is given by the sum of the rates of this channel in all patterns that overlap at time t . More formally, let $S(t)$ denote the set of indices of patterns that overlap at time t and let $s_i(t)$ denote the frame at which pattern i is at time t (if the pattern presentation started at time t' , the pattern is in frame $s_i(t) = \frac{t - t'}{\Delta t} + 1$ at time t , with Δt being the discretization time step). Then the linearly superimposed rate $r_j^{\text{linear}}(t)$ for channel j is given by the following:

$$r_j^{\text{linear}}(t) = \sum_{i \in S(t)} r_{j,s_i(t)}^i \quad (9)$$

In the final nonlinear step, the firing rate $r_j(t)$ of input neuron j at time t is squashed by a sigmoidal nonlinearity as follows:

$$r_j(t) = \frac{f_H}{1 + \exp\left(-\frac{2\kappa}{f_H}(r_j^{\text{linear}}(t) - 0.5f_H)\right)}, \quad (10)$$

where $f_H = 75$ Hz is the maximum attainable rate and $\kappa = 5$ sets the width of the sigmoidal function. To avoid completely silent periods in the input between pattern presentations, the rate of each input neuron is set to 2 Hz at times t when no patterns are superimposed (i.e., $S(t) = \{\}$).

For Figure 3, we used 2 basic patterns with a maximum number of superimposed basic patterns of $n_{\max} = 2$ and each register had a load probability of $p_{\text{loaded}} = 0.5$ (i.e., each register was loaded with some basic pattern half of the time).

Pattern selectivity and activity plots in Figure 3B, C. Network activity was analyzed after a learning period of 400 s of simulated biological time. Neurons were classified as preferring pattern 1 (green pattern), as preferring pattern 2 (blue pattern), or as nonselective based on a procedure similar to the one used by Harvey et al. (2012). First, an activity trace for each neuron was obtained by convolving its spike response with a double exponential kernel Equation 2 with $\tau_r = 1$ ms, $\tau_f = 20$ ms, and cutoff time $T_\epsilon = 200$ ms. A neuron was considered to be active if it had at least 2 spikes during the simulation time. We classified a neuron as pattern modulated if it was an active neuron and if it had a twice as high average activity trace during presentations of patterns than during the times without patterns. From the pattern modulated neurons, a neuron was classified as pattern selective if it had significantly different activity traces during presentation of blue and green patterns. This was determined by a two-tailed t test with significance value set at $p < 0.05$. If a neuron was pattern selective, its pattern preference was decided based on the average activity trace during blue and green pattern presentations: the preferred pattern was defined as the pattern for which the mean of the activity trace is higher. Finally, we call a neuron that is not pattern selective a nonselective neuron.

For average activity plots in Figure 3C, the green and blue patterns were presented to the network in isolation, 200 presentations per pattern. Activity traces of all pattern selective neurons were averaged over all presentations of the pattern and subsequently normalized to their peak average activity. Neurons were then sorted by the time of their peak average activity at presentation of their preferred pattern, and average activity was plotted in the sorted order for both patterns. In Figure 3B, spike trains were also plotted separately for green pattern preferring neurons, blue pattern preferring neurons, and nonselective neurons. The sorting of the former two groups was the same as in Figure 3C.

Details to simulations for Figure 4

For this simulation, basic rate patterns were superimposed as described above for Figure 3. The experiment, however, differed in the number and choice of basic patterns. Network responses, precision measures, and synaptic weight vectors were evaluated and plotted after a learning period of 400 s simulated biological time.

Creation of basic rate patterns. Basic patterns consisted of 64 channels that were representing horizontal and vertical bars in a 2D pixel array of size 8×8 pixels. The pattern length was 50 ms (50 frames); and in contrast to the basic patterns for Figure 3, the rate in each individual channel was constant over the period of the pattern (i.e., $r_{n,s}^i = r_n^i$ for $s = 1, \dots, S$). Each of the 64 channels, r_n^i in a basic pattern \mathbf{R}^i corresponded to one pixel in an 8×8 pixel array. We defined 16 basic patterns in total, corresponding to all possible horizontal and vertical bars of width 1 in this pixel array. For a horizontal (vertical) bar, all pixels of a row (column) in the array attained the value 75, whereas all other pixels were set to 0. The channel rates r_n^i were then defined by the values of the corresponding pixels in the array.

Superposition of basic rate patterns. Basic rate patterns were superimposed as described above for Figure 3. A maximum of $n_{\max} = 3$ basic patterns were allowed to be superimposed at any time with a load probability of $p_{\text{loaded}} = 0.9$ (for a definition, see Details to simulations for Fig. 3). In addition to the rates defined by this superposition, a noise rate of $r_{\text{noise}}(t) = 3(3 - n_{\text{pat}}(t))$ Hz was added to each channel, where $n_{\text{pat}}(t)$ denotes the number of patterns that are superimposed at time t .

Precision measure. Our aim was to quantify how well neurons prefer

particular basic patterns (i.e., are tuned to particular basic patterns). To measure tuning properties, we computed the precision measure (van Rijsbergen, 1974) Precision_{ij} for each pair of excitatory neuron i and basic pattern (bar) j . To this end, we say that a neuron i indicates the presence of a pattern whenever the neuron spikes. The precision Precision_{ij} is then the fraction between the number of times the presence of pattern j is correctly indicated by neuron i divided by the number of times that neuron i indicates that pattern. Hence, the precision measures how well one can predict the presence of a pattern j , given that neuron i spiked. The precision measures whether the pattern is present whenever there is a spike, and not whether the neuron spikes whenever the pattern is present. Because many neurons are jointly representing a pattern, the latter question does not make sense on the individual neuron level (it will be quantified later in Fig. 5C on the population level).

Figure 4C shows for each pair of excitatory neuron i and basic pattern (bar) j the precision measure Precision_{ij} . Formally, the precision measure is defined according to van Rijsbergen (1974) as follows:

$$\text{Precision}_{ij} = \frac{\text{TP}_{ij}}{\text{TP}_{ij} + \text{FP}_{ij}}, \quad (11)$$

where TP_{ij} denotes the true-positive count and FP_{ij} denotes the false-positive count for that pair. The true-positive count TP_{ij} is given by the number of times that neuron i spikes while basic pattern j is present in the input. A pattern that starts at time t is defined to be present in the interval $[t, t + S\Delta t + \tau]$. Here, $S\Delta t$ is the length of the pattern and $\tau = 10$ ms corrects for PSPs that increase the firing rates of excitatory neurons, even after the pattern disappeared. The false-positive count FP_{ij} denotes the number of times that neuron i spikes when pattern j is not present.

We say that a neuron i prefers basic pattern j if the neuron has maximum precision for pattern j and this precision is larger or equal to 0.8, and if the second largest precision that neuron i has for any other pattern is <0.7 . A neuron is said to be pattern-selective if it prefers some pattern and nonselective otherwise. In Figure 4B, C, pattern-selective neurons are shown and sorted according to their preferred basic pattern.

Details to simulations for Figure 5

Figure 5 shows the behavior of M over the course of learning in the overlapping bars task (same setup as Fig. 4). Ten independent simulation runs were performed, each for 1000 s of simulated biological time.

In Figure 5A, a neuron is considered to be recruited if it is pattern selective. In Figure 5B, a pattern is considered to be represented if at least one neuron prefers that pattern. Because several excitatory neurons in the circuit can specialize on a given basic pattern, the network performance shown in Figure 5C was evaluated over ensembles of neurons, where ensemble E_i is given by the set of neurons that prefer basic pattern i . To quantify how well basic pattern i is represented by ensemble E_i , we computed the F1 measure (Van Rijsbergen, 2004). The F1 measure is at its maximum value of 1 if the following holds true: a neuron in the ensemble E_i is active if and only if basic pattern i is present in the input. False-positives (i.e., some neuron in the ensemble is active in the absence of the basic pattern) and false-negatives (i.e., the basic pattern is present in the input, but no neuron of the ensemble is active) reduce the measure, and the minimum possible value of the measure is 0. Hence, the F1 measure for basic pattern i measures how well this basic pattern is represented by the ensemble E_i . Formally, we computed the F1 measure (Van Rijsbergen, 2004) for ensemble E_i defined as follows:

$$\text{F1}_i = \frac{2\text{TP}_i}{2\text{TP}_i + \text{FN}_i + \text{FP}_i}, \quad (12)$$

where FN_i denotes the false-negative count. The true-positive count TP_i is given by the number of times that basic pattern i is present and detected by ensemble E_i , where the pattern active during $[t, t + S\Delta t]$ is detected if some neuron of the ensemble fires at least one spike within $[t, t + S\Delta t + \tau]$. The false-negative count FN_i denotes the number of times when the pattern i is active, but there is not a single spike from ensemble E_i . To calculate the false-positive count FP_i , we split the time between two presentations of the pattern (time without pattern i) into periods of $[t, t + S\Delta t + \tau]$ (where the last period can be shorter). Then the false-positive

count FP_i denotes the number of such periods during which there is at least one spike from ensemble E_i . In Figure 5C, the mean F1 measure over all 16 basic patterns is plotted, thus indicating how well all the patterns are represented by the network. For comparison, a SEM network as used for Figure 2H was trained on the same input for 2000 s simulated time.

Details to simulations for Figure 6

Figure 6 shows analysis regarding the temporal relation between excitation and inhibition in M in the experiment of Emergent computation on spike patterns (Fig. 3) after learning. Figure 6A depicts the mean firing rate of excitatory neurons in the network (blue represents mean taken over all excitatory neurons) and the mean firing rate of inhibitory neurons (red represents smoothed through a 100 ms boxcar filter) as well as the scaled mean firing rate of inhibitory neurons (dashed green). The mean firing rate of inhibitory neurons is scaled down by the ratio of the average excitatory and inhibitory firing rates (average taken over the whole simulation time) to facilitate comparison.

The lag between excitation and inhibition was quantified in a similar manner as by Okun and Lampl (2008). The cross-correlation function between excitatory and scaled inhibitory firing rate for a duration of 10 s was computed (plotted in Fig. 6B). The lag was then given by the offset of the peak in the cross-correlation function (plotted in Fig. 6C) from 0. To evaluate the influence of connections between inhibitory neurons, we performed the same simulations without I-to-I connections and quantified the lag in the same manner (Fig. 6C). To facilitate a fair comparison, in addition to removing I-to-I connections we also scaled down synaptic weights of I-to-E connections by factor of 0.1155 to obtain the same excitatory firing rate as in the case with I-to-I connections.

Results

A data-based model for a microcircuit motif consisting of excitatory and inhibitory neurons

We analyze computational properties of densely interconnected populations of excitatory and inhibitory neurons. In particular, we analyze a model for interacting populations of pyramidal cells with PV^+ inhibitory neurons on layer 2/3 that is based on data from the Petersen laboratory (Avermann et al., 2012) (Fig. 1A, B). We refer to this specific model as the microcircuit motif model M .

The excitatory pool in M consists of M stochastic spiking neurons, for which we use a stochastic version of the spike response model that has been fitted to experimental data in (Jolivet et al., 2006). In this model, the instantaneous firing rate $\rho_m(t)$ of neuron m is approximated by the exponential function applied to the current membrane potential (see Eq. 4). These excitatory neurons project to and receive inputs from a pool of inhibitory neurons, which are also interconnected among themselves, with connections probabilities taken from Avermann et al. (2012). Each excitatory neuron m in the network also receives excitatory synaptic inputs $\tilde{y}_1(t), \dots, \tilde{y}_N(t)$ from external input neurons, whose contribution to its membrane potential at time t depends on the synaptic efficiency w_{im} between the input neuron i and neuron m . We assume that these afferent connections are subject to a standard form of STDP (Fig. 1C, definition of a data-based microcircuit motif model M ; see Materials and Methods).

Negative (inhibitory) contributions $\sum_{j \in \mathcal{I}_m} w^{IE} I_j(t)$ to the membrane potential of pyramidal cell m have according to the neuron model a divisive effect on its firing activity because its instantaneous firing rate ρ_m can be written (by substituting Eq. 4 in Eq. 3) as follows:

$$\rho_m(t) = \frac{1}{\tau} \frac{\exp(\gamma \sum_i w_{im} \tilde{y}_i(t) + \gamma \alpha)}{\exp(\gamma \sum_{j \in \mathcal{I}_m} w^{IE} I_j(t))}. \quad (13)$$

Here, the numerator includes all excitatory contributions to the firing rate $\rho_m(t)$, that is, the synaptic inputs (unweighted sum of

EPSPs) $\tilde{y}_i(t)$ from input neurons weighted by the corresponding synaptic weights w_{im} . α denotes the neuronal excitability, and τ and γ are scaling parameters that control the shape of the response function of the neuron. The denominator in this equation for the firing rate describes inhibitory contributions, thereby reflecting divisive inhibition (Carandini and Heeger, 2011). Here, $I_j(t)$ denotes synaptic input from inhibitory neuron j weighted by some common weight w^{IE} (\mathcal{J}_m denotes the set of all inhibitory neurons that connect to neuron m).

Divisive inhibition has been shown to be characteristic for the interaction of pyramidal cells with PV⁺ inhibitory neurons (Wilson et al., 2012). To test also on a functional level the divisive character of inhibition in the model, we artificially increased the firing rate of inhibitory neurons in the circuit by a constant, corresponding to the *in vivo* experiment described by Wilson et al. (2012), where activity of the PV⁺ neurons was increased through optogenetic stimulation. The response of pyramidal neurons to this increased inhibition in M resembles the experimental data (Fig. 1D).

Emergent neural codes: from WTA to k -WTA

In our first test of emergent computational properties of this microcircuit motif model M, we examined the emergence of orientation selectivity. We provided as external spike inputs pixel-wise representations of bars in numerous random orientations with superimposed noise (Fig. 2A). Bars were transformed into high-dimensional spike inputs by representing each black pixel of an oriented bar for 50 ms through a Poisson input neuron with a Poisson rate of 75 Hz, whereas all other input neurons had a Poisson rate of 1 Hz (for a typical resulting spike input pattern, see Fig. 2B). The initial network response is shown in Figure 2C, and the network response after applying STDP for 400 s to all synapses from input neurons to excitatory neurons in Figure 2D. One clearly sees in Figure 2D the emergence of assembly codes for oriented bars. A closer look at the resulting tuning curves of excitatory neurons in Figure 2E, F shows a dense covering of orientations by Gaussian-like tuning curves similar to experimental data from orientation pinwheels (Ohki et al., 2006, their Fig. 2d,e). In contrast, inhibitory neurons did not become orientation-selective (Fig. 2G) in accordance with experimental data (Kerlin et al., 2010; Isaacson and Scanziani, 2011).

The tuning curves of excitatory neurons in Figure 2E, F demonstrate a clear difference between the impact of divisive inhibition in this data-based model M and previously considered idealized strong inhibition in WTA circuits (Nessler et al., 2013) (for emergent computational properties, see Fig. 2H). In the data-based model M, several (on average $k = 17$) neurons respond to each orientation with an increased firing rate. This suggests that the emergent computational operation of the layer 2/3 microcircuit motif with divisive inhibition is better described as a k -WTA computation, where k winners may emerge simultaneously from the competition. In contrast, for the WTA model with idealized strong inhibition (Nessler et al., 2013), at most, a single neuron could fire at any moment of time and, as a result, at most two neurons responded after a corresponding learning protocol with an increased firing rate to a given orientation (Fig. 2H) (Nessler et al., 2013, their Fig. 5).

From the perspective of computational complexity theory, the k -WTA computation is known to be for $k > 1$ more powerful than the simple WTA computation (Maass, 2000). However, the number k of winners in this microcircuit motif is not fixed. It depends on synaptic weights and the external input.

Hence, one can best describe its computation as an adaptive k -WTA operation.

Emergent computation on spike patterns

Simultaneous recordings from large numbers of neurons demonstrate the prominence of large-scale activity patterns in networks of neurons (Luczak et al., 2015). They are commonly referred to as assemblies, assembly sequences, or assembly phase sequences. Since Hebb (1949), they have been proposed to reflect tokens of brain computations that connect the fast time-scale of spikes (milliseconds) to the slower time-scale of cognition and behavior (100s of milliseconds). But their precise role in neural coding and computation has remained unknown. Luczak et al. (2015) proposed that they serve as basic information components in global cortical communication, where each of these activity patterns is initiated by a particular cortical region and broadcast to all areas to which it projects. We show here that our microcircuit motif model M is able to perform a computational operation on large-scale activity patterns that is fundamental for such a global communication scheme. It can demix superimposed spike patterns that impinge on a generic cortical area and represent the presence of each pattern in their input stream through the firing of separate populations of neurons. This suggests that the layer 2/3 microcircuit motif has an inherent capability to solve the well-known cocktail party problem (blind source separation) (Cherry, 1953) on the level of larger activity patterns. This capability emerges automatically through STDP, as demonstrated in Figure 3 for our data-based model M.

The input to the microcircuit motif model M is generated in Figure 3 by 200 spiking neurons. Two repeating activity patterns (green and blue patterns) are superimposed for the generation of Poisson spike trains (shown for every second neuron in the top row of Fig. 3B). These two large-scale activity patterns consist of two time-varying rate patterns for the 200 input neurons (Fig. 3A, middle) that are nonlinearly superimposed with random offsets in the continuous spike input to our model. Despite these random offsets and the large trial-to-trial variability of spike times in each of the two patterns (Fig. 3A, right panels), STDP in the synaptic connections from inputs to the excitatory neurons in the model was produced after 400 s two assemblies (Fig. 3B, middle row, green and blue spikes). Each responded to just one of the two input patterns and represented its temporal progress through a stereotypical sequential firing pattern (Fig. 3C). This effect occurs even if none of the two input patterns is ever presented in isolation during learning, as shown for illustration purposes for test inputs after learning on the right side of Figure 3B. Such emergent demixing of superimposed spike patterns in the layer 2/3 microcircuit motif could enable downstream neurons to selectively respond to just one of the patterns. Furthermore, the sequential activation of the two assemblies can also inform downstream networks through the firing of specific neurons about the current phase of each of the two input patterns.

Emergent modular sparse coding

Földiák (1990) suggested that complex objects or scenes are encoded in the brain through a sparse modular code, where each neuron signals through its firing the presence of a particular feature in the network input. In this way, a combinatorial explosion of the number of neurons is avoided, which would be required if each complex external object or scene is encoded as a whole by separate neurons. Földiák (1990) proposed to use superpositions of bars (lines), as in Figure 4A (top), as benchmark inputs to test sparse modular coding capabilities of neural network models. A

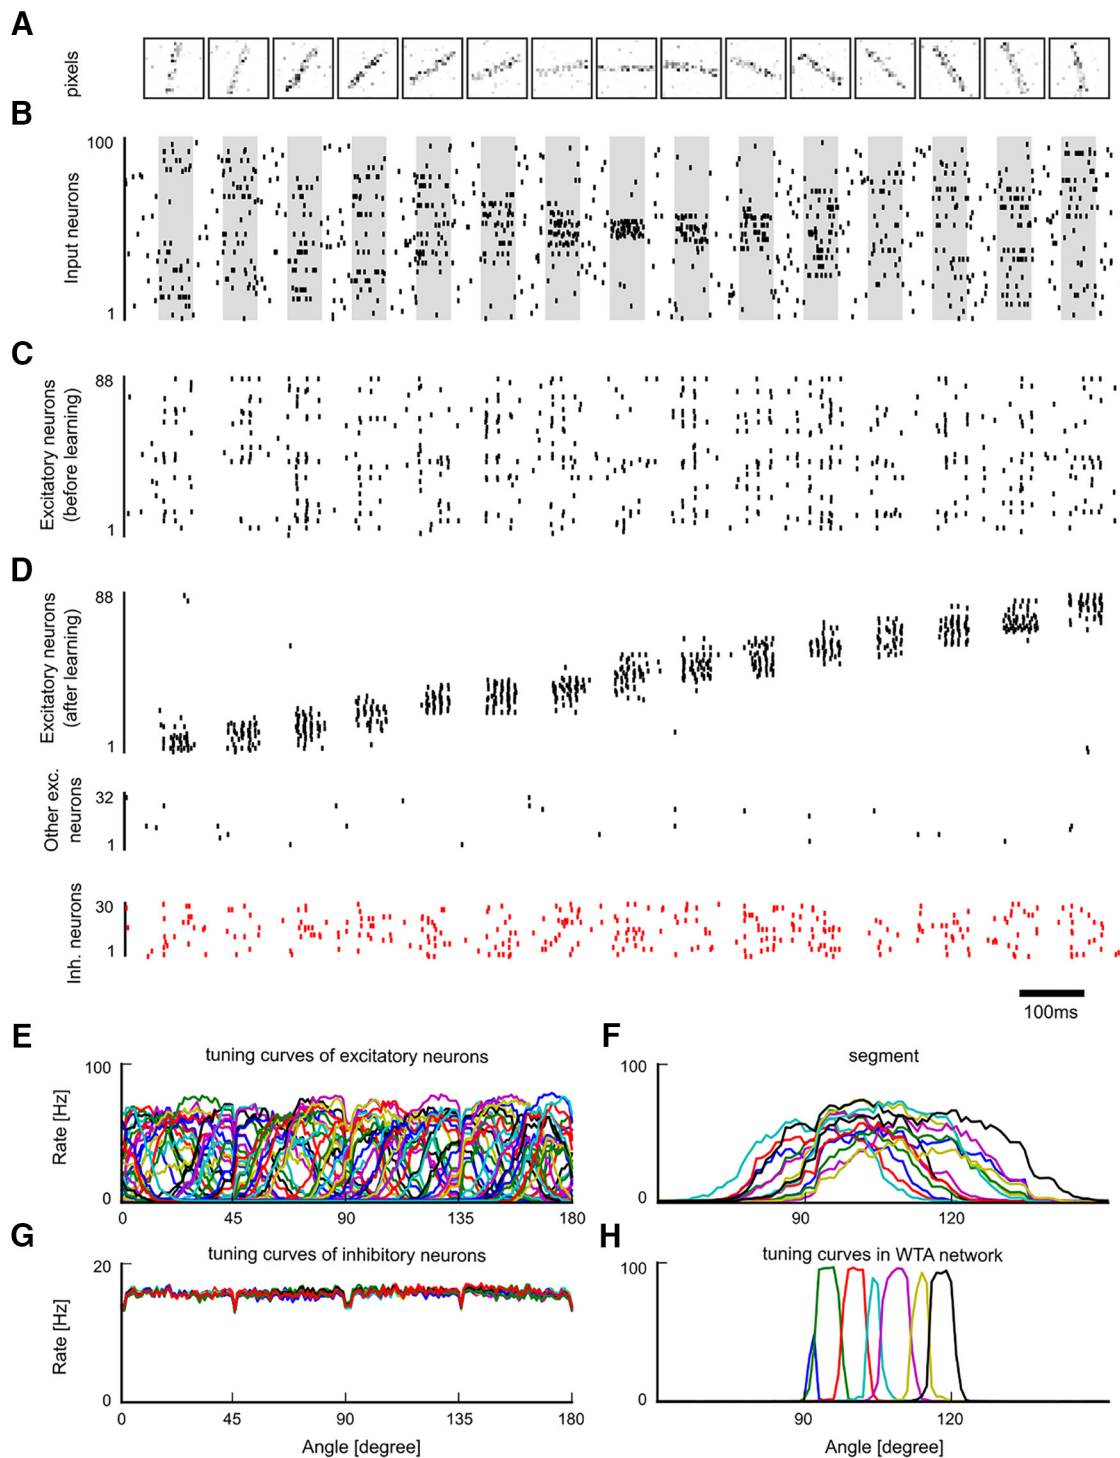


Figure 2. Emergent neural codes in the microcircuit motif model. **A**, Bars at various orientations serve as network inputs. Shown are network inputs arranged in 2D for clarity. Gray-level of each pixel indicates the resulting effective network input $\tilde{y}_i(t)$ (see Eq. 1) at some time point t . **B**, Resulting spike pattern of input neurons (every fourth neuron shown) for different bar orientations. Gray shading represents the presence of a bar in the input with orientation indicated in **A**. **C**, Example spike pattern of a subset of excitatory neurons in the circuit to this input before learning. **D**, Spiking activity of the same neurons for the same input after applying STDP to all synapses from input neurons to excitatory neurons for 400 s. Only responses of orientation-selective neurons are shown, sorted by preferred orientation. Spiking activity of a random subset of non-orientation-selective neurons and inhibitory neurons to the same input is shown below. **E**, Emergent tuning curves of orientation-selective excitatory neurons. **F**, The same as in **E**, but zoomed in on orientations between 90 and 120 degrees. **G**, Inhibitory neurons are not orientation-selective. **H**, Emergent tuning curves of neurons in a previously considered WTA model (Habenschuss et al., 2013b; Nessler et al., 2013).

neural network is able to avoid the combinatorial explosion of the number of neurons that are needed to encode such complex inputs if it learns to represent them in a modular fashion, where each neuron encodes the presence of one of the bars (in a particular location) in the composed input. A key question is how such

codes can emerge in a network autonomously if only composite images (consisting of several superimposed bars) are presented as network inputs. A WTA circuit is not able to develop a good modular code because it does not allow inputs to be represented through the firing of more than one neuron. Hence, a natural

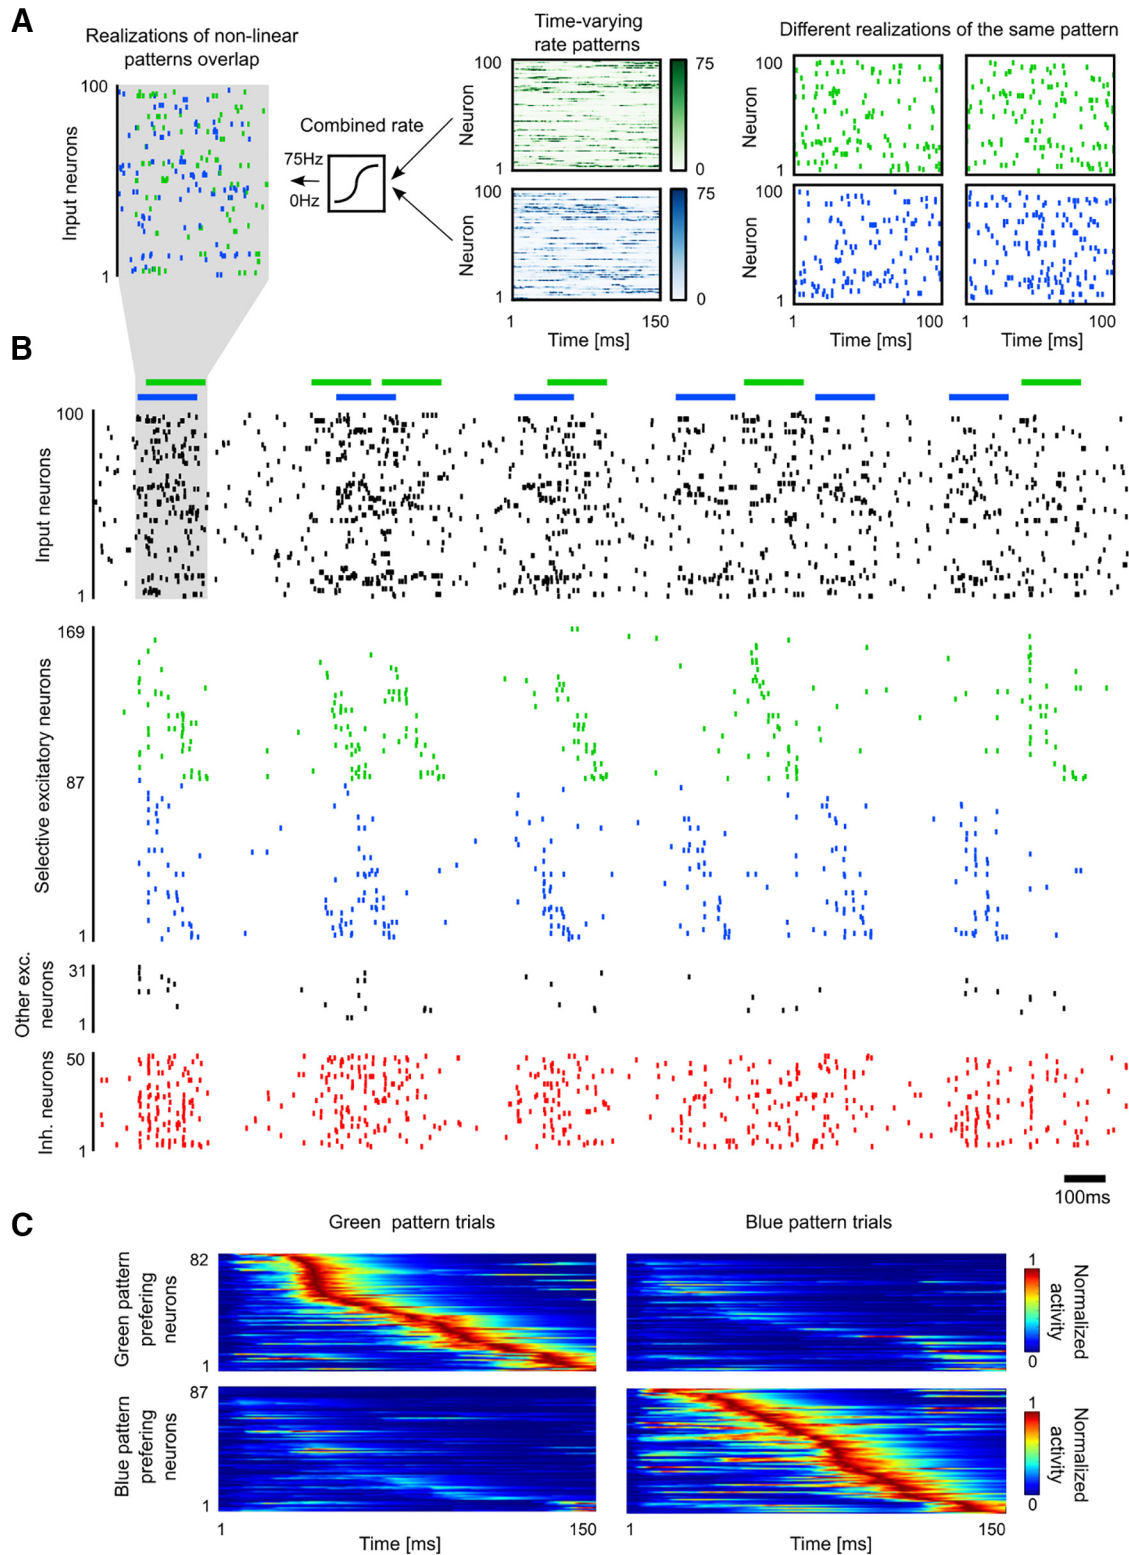


Figure 3. Emergent computation on large-scale spike patterns. **A**, Two spatiotemporal patterns. Each pattern consists of 200 time-varying firing rates over 150 ms generated by an OU process (middle, only every second channel shown for clarity). These rate patterns give rise to highly variable spike patterns, as shown on the right. Basic rate patterns are superimposed nonlinearly with arbitrary relative timing. Left, One realization of superimposed patterns for a time segment in **B**. Spikes are colored according to the basic pattern that most probably caused the spike (i.e., the one with the higher rate at that time). **B**, Firing response of the neurons in our model *M* for a test input stream after letting STDP be active for synaptic connections from input neurons to excitatory neurons in *M*. Two subpopulations emerged (green and blue), where each neuron specialized on a specific pattern and on a particular time segment within this pattern. Spiking activity of a subset of nonselective neurons (black) and inhibitory neurons (red) are shown below. **C**, Average firing rate of neurons preferring the green (top) and blue (bottom) input pattern when the green (left) and blue pattern (right) is shown in isolation. Neurons are ordered according to their peak firing rate for the preferred pattern as in **B**. Resulting selective firing responses are qualitatively similar to data from sensory cortices (Luczak et al., 2015) and higher cortical areas (Harvey et al., 2012).

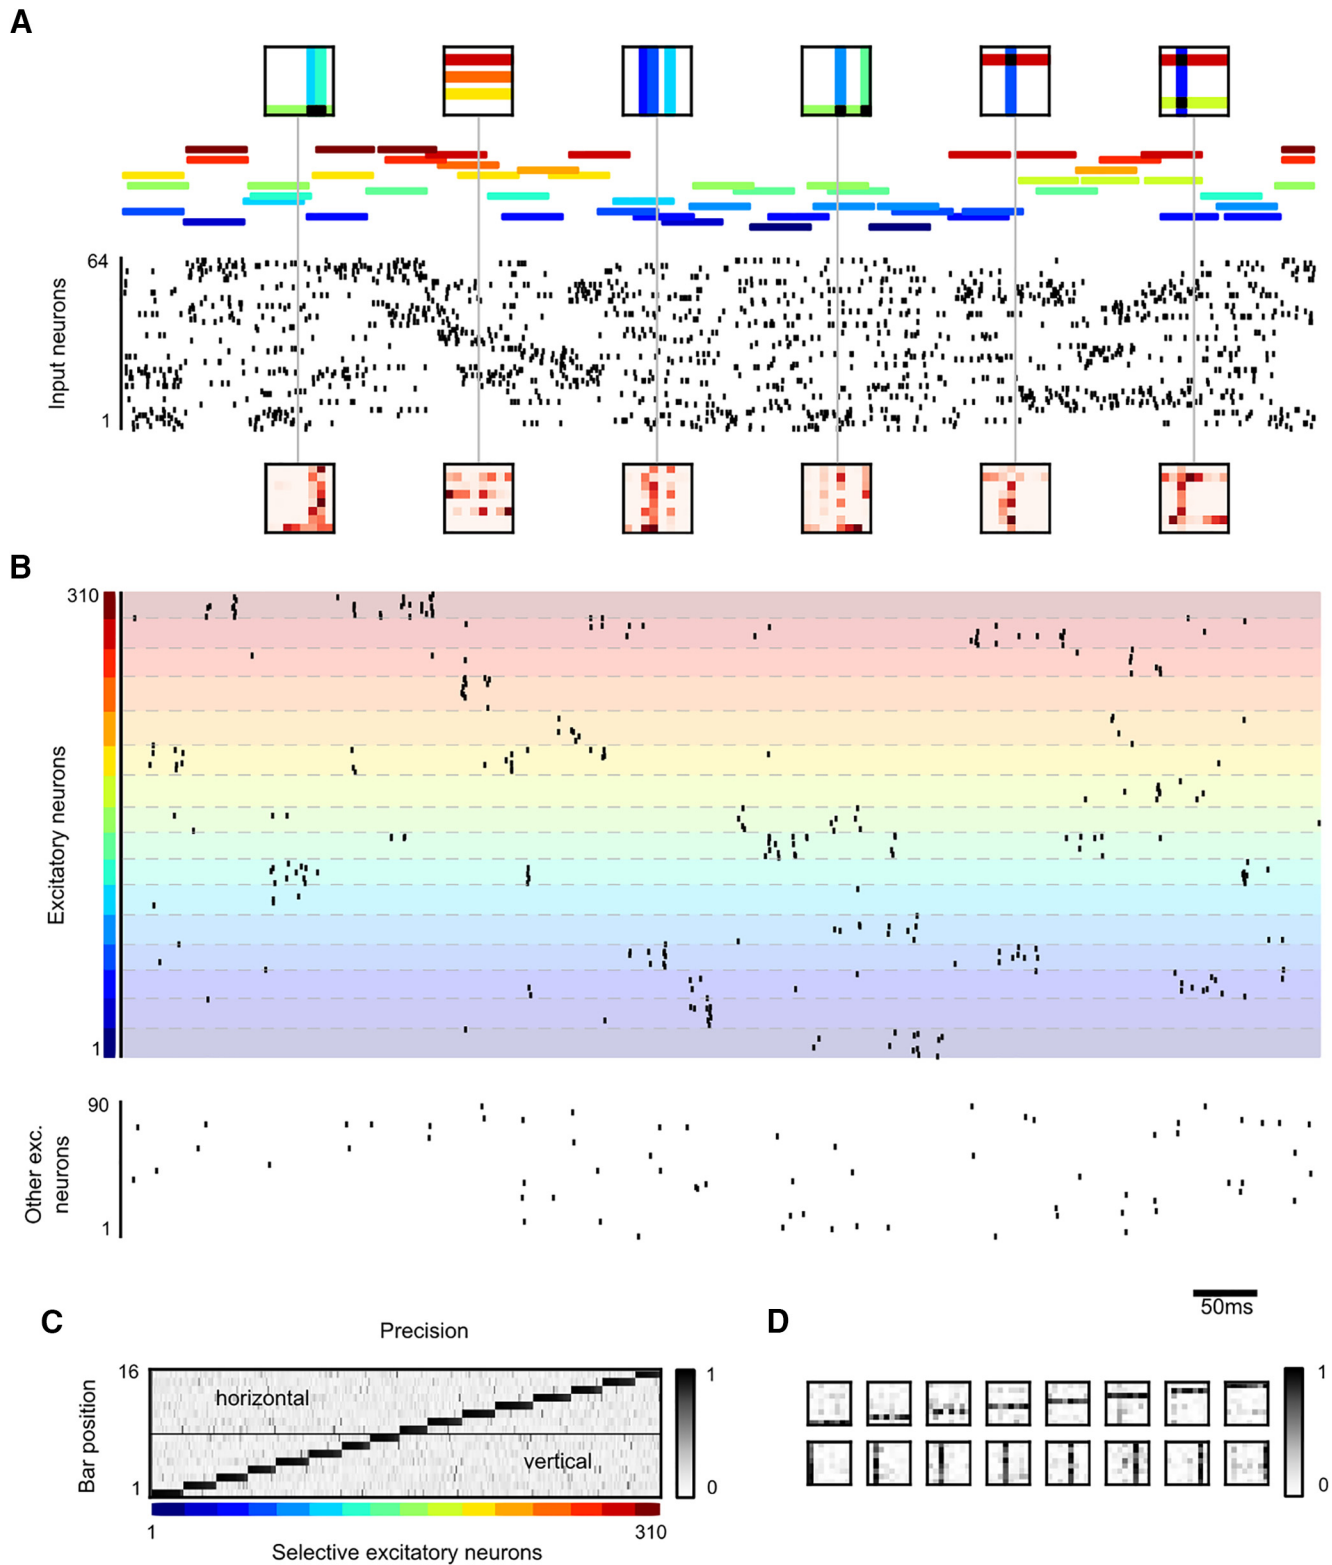


Figure 4. Test of the emergence of modular sparse codes for a common benchmark test. **A**, A difficult version of superposition-of-bars problem with asynchronously varying numbers of up to 3 superimposed bars (Földiák, 1990). Each of the 16 bar positions is indicated by a separate color, and its presence in the resulting spike input stream (see middle row) is indicated by a horizontal colored line above the spike raster. Six of 696 possible composed input patterns are shown for arbitrarily chosen time points indicated by gray vertical lines. The pattern at the bottom of each line indicates the effective spike input that the network receives at that moment in time (see Eq. 1). **B**, Emergent modular assembly codes in model M after 400 s. A small assembly of neurons emerges for each of the 16 bar positions (color code at left axis and background shading). Activity of nonselective excitatory neurons is shown at the bottom. **C**, Quantitative analysis of the precision of the emergent assembly codes measured according to van Rijsbergen (1974) for each of the 310 neurons from **B** (top) on the x-axis. Dark shading represents high precision for encoding the bar position plotted on the y-axis. **D**, Typical weight vectors of neurons from the 16 assemblies that had emerged.




Figure 5. Quantitative analysis of emergent coding properties in the benchmark task of Figure 4. **A**, Evolution of the number of bar selective (blue) and nonselective (red) neurons during learning. **B**, The number of bar positions represented by the network rises rapidly during learning. A bar position is considered to be represented if at least one excitatory neuron is selective for it. **C**, Average F1 measure of pattern-selective neural ensembles during learning. High F1 measure (maximum is 1) indicates emergence of highly selective assemblies of neurons for all bar positions (see Materials and Methods). In all plots, saturated colors represent mean, and light colored shading represents SD over 10 runs.

question is whether biologically more realistic softer lateral inhibition, as implemented in our model M, supports the emergence of sparse modular codes through STDP. Figure 3 demonstrated already some weak form of modular coding for superpositions of two spatiotemporal patterns in the input.

Emergent neural codes for the superposition-of-bars problem (Földiák, 1990) are examined in Figures 4 and Figure 5 for our model M with 400 excitatory and 100 inhibitory neurons as before. Superpositions of up to 3 bars were presented through 64 spiking input neurons in a pixel-wise encoding. Each Poisson input neuron signaled for 50 ms through an increased firing rate if the corresponding pixel was covered by a bar (each bar covered 8 horizontal or 8 vertical pixels in an 8×8 pixel array). Each of the 16 possible bar positions is indicated in Figure 4A through a different color. Composed network inputs were created by randomly drawing superposition of bars from the pool of 696 combinations of up to 3 bars. Obviously, our model M would not be able to represent each of these input patterns by a separate neuron. Nevertheless, a complete and noise robust modular code emerged in M.

A typical spike input stream from the 64 input neurons is shown in Figure 4A (middle row). The 6 squares at the bottom of Figure 4A show for 6 representative time points (indicated by gray vertical lines) the resulting pixel-wise code that represents the network input (darkness of red color indicates the output trace of that input neuron at that time, i.e., its output spike train convolved with the synaptic response kernel). After representing such a continuously varying input stream for 400 s to the network, subpopulations of a few neurons (19.4 on average) emerged that each indicated through their firing the presence of a bar at a particular position in a noise robust manner through the firing of several neurons (bar position indicated in Fig. 4B, left, and through a corresponding shading in the background of the spike raster). In this way, each composite input image is represented through an emergent sparse modular neural code. This holds despite the fact that the image presentations were not synchronized (i.e., individual bars appeared and disappeared at random time points) and the number of simultaneously present bars varied.

We quantified the learning performance of our model M in extended simulations where the network was exposed to this input for 1000 s of simulated biological time. The evaluation based on 10 runs with independently drawn initial synaptic weight settings and input patterns is shown in Figure 5 (for details, see Materials and Methods). Figure 5A shows the number of neurons recruited for modular neural coding during learning. Figure 5B

shows that the network rapidly and robustly learns to represent all 16 bar positions. In Figure 5C, network coding performance is plotted against learning time in terms of the F1 measure (Van Rijsbergen, 2004). This measure is suitable for analyzing the reliability of assembly codes, where several neurons in an assembly can become selective for the same feature (here: bar position) in the network input. The F1 measure was separately computed for each bar position. An F1 measure of 1 for a bar position indicates that the bar is correctly reported by those neurons that are selective for a bar at this position (i.e., at least one of the neurons in the corresponding emergent assembly is active if this bar is present and all are inactive otherwise). Hence, a high F1 measure indicates a robust encoding of bar positions by excitatory neurons in M. In Figure 5C, the average F1 measure over all bar positions is plotted. After, 1000 s of learning, an average F1 measure of 0.87 was attained. The network already represents the input very well after about 200 s of learning (Fig. 5C), although only ~ 200 neurons have become pattern selective at this point (Fig. 5A). Subsequently, the ensembles that represent basic patterns become larger, but this has only a small impact on network performance. For comparison, a WTA network with idealized strong inhibition (as used for Fig. 2H) was trained on the same input. In the WTA circuit, neurons did not develop a modular code but specialized on combinations of bars.

Comparing the resulting temporal dynamics of inhibition with experimental data

We analyzed the resulting temporal dynamics of inhibition in model M after the learning experiment of Emergent computation on spike patterns (Fig. 3). Figure 6A shows the time courses of the average firing rates of excitatory neurons (blue) and inhibitory neurons (red represents scaled inhibitory rate in green for better comparison) during an example time interval of 500 ms. Consistent with experimental findings (Okun and Lampl, 2008), inhibition tracks excitation quite precisely, with a small time lag.

We quantified the lag between excitation and inhibition as in Okun and Lampl (2008) as the temporal offset of the peak of the cross-correlation function between the excitatory and scaled inhibitory firing rates (Fig. 6B, plotted as black line). The resulting lag of 3 ms is comparable with the measured mean lag of 3.5 ms *in vivo* (Okun and Lampl, 2008).

In accordance with the data in Avermann et al. (2012), we included inhibitory connections within the pool of inhibitory neurons (I-I connections) in the microcircuit motif model M. We found that these connections play an important role because they

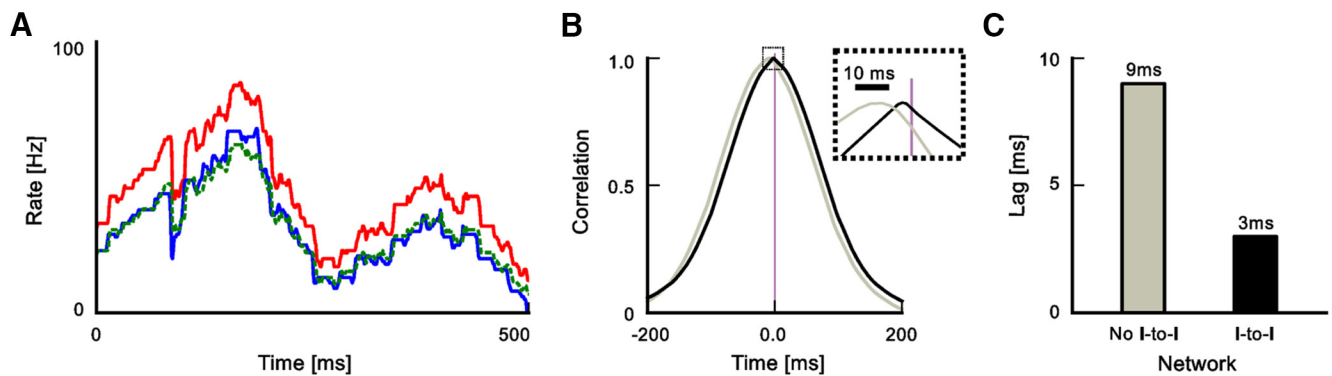


Figure 6. Time course of excitation and inhibition in model M. **A**, Time course of the average firing rate of excitatory (blue) and inhibitory neurons (red) during 500 ms after learning in the experiment shown in Figure 3. Dashed green line indicates a scaled version of the average inhibitory rate for better comparison. **B**, Cross-correlation function between the excitatory and inhibitory rate reveals a small lag of ~ 3 ms between excitation and inhibition, comparable with *in vivo* data. Shown is the cross-correlation with intact connections among inhibitory neurons (black) and without these connections (gray). Inset, Zoom into the dotted rectangle. **C**, Quantification of the lag between excitation and inhibition from **B**. Intact inhibition among inhibitory neurons significantly reduces the lag between excitation and inhibition.

decrease the lag between excitation and inhibition. This is quantified in Figure 6B, where the black line indicates the resulting correlation between excitation and inhibition in the model, and the gray line for a variation of the model where all I-I connections have been deleted. The average lag between excitation and inhibition increased through this deletion from 3 ms (Fig. 3C, black bar) to 9 ms (Fig. 3C, gray bar). With intact I-I connections, inhibition is sharpened because early inhibitory responses to excitation reduce subsequent inhibitory spikes with a larger lag. For further details on these experiments, see Details to simulations for Figure 6.

Discussion

We have investigated the computational properties of interconnected populations of pyramidal cells and PV⁺ interneurons in layer 2/3 (Fig. 1), one of the most prominent motifs in cortical neural networks. Our analysis was based on data from the Petersen laboratory for layer 2/3 of mouse barrel cortex as summarized by Avermann et al. (2012). We have shown that the dynamics of inhibition in a simple model M for this microcircuit motif is consistent with additional experimental data. Figure 1D shows that the resulting feedback inhibition is consistent with data from Wilson et al. (2012). Furthermore, inhibition follows excitation in our model with a lag of ~ 3 ms (Fig. 6A), a value that is close to the experimentally measured mean lag of 3.5 ms (Okun and Lampl, 2008). Model M has produced, in addition in Figure 6B, C, a hypothesis for the functional role of synaptic interconnections among PV⁺ cells in this context. It suggests that these connections contribute to the small value of this lag.

We found that the role of inhibition in this microcircuit motif cannot be captured adequately by a WTA model. We are proposing to consider instead a variation of the *k*-WTA model, where the *k* most excited neurons are allowed to fire. The *k*-WTA model is well known in computational complexity theory and tends to produce more computational power than the simple WTA model (Maass, 2000). A closer look shows that the dynamics of the microcircuit motif can be captured even better by an adaptive *k*-WTA model. In this model, the actual number of neurons that fire in response to a network input may vary.

We have investigated the computational properties that emerge in model M under STDP for spike input streams that contain superimposed firing patterns. We found the emergent capability to disentangle these patterns and represent the occur-

rence of each pattern by a separate sparse assembly of neurons (Figs. 2–4). Hence, we propose that the ubiquitous microcircuit motif of densely interconnected populations of excitatory and inhibitory neurons provides an important atomic computational operation to large-scale distributed brain computations. Through this operation, each network module may produce one of a small repertoire of stereotypical firing patterns, commonly referred to as assemblies, assembly sequences, or packets of information (Luczak et al., 2015). If these assembly activations are fundamental tokens of global cortical computation and communication, as proposed by Luczak et al. (2015), then cortical columns have to solve a particular instance of the well-known cocktail party problem (Cherry, 1953). They have to recognize and separately represent spike inputs from different assemblies that are superimposed in their network input stream.

The existence of blind source separation mechanisms of this type had already been postulated by Földiák (1990) as a prerequisite for avoiding a combinatorial explosion in the number of neurons that are needed to represent the information contained in complex spike input streams. We have shown in Figures 3–5 that blind source separation for spike patterns emerges automatically in the microcircuit motif model M through STDP. This holds even for a more demanding version of the benchmark task that (Földiák, 1990) had proposed. Disentangling and representing superpositions of bars not only for a fixed number, but also for varying numbers of superimposed bars.

Relation to theoretical models for cortical microcircuit motifs

It is natural to ask whether a theoretical analysis can be performed to better understand the emergence of this fundamental computational capability. Unfortunately, the analysis from Nessler et al. (2013) and Habenschuss et al. (2013b) in terms of mixture distributions is only applicable to WTA circuits. A novel theoretical analysis in Legenstein et al. (2017) shows that one can relate some parameters of model M, such as the neural excitability α and various synaptic efficacies in the network, to parameters of a probabilistic model for softer divisive inhibition. Synaptic efficacy parameters and neural excitabilities used in the simulations for this article can be related to this probabilistic model.

Related work

Learning in networks of excitatory and inhibitory neurons was also studied by Litwin-Kumar and Doiron (2014). However, they did not

study plasticity of synaptic connections from inputs to the network. Consequently, their model could not learn to perform any feature extraction from input patterns, which is the primary emergent computational property of model M. Rather, self-organization led in the model of Litwin-Kumar and Doiron (2014) to an associative memory-like network behavior. An interesting feature of their model was the use of a fast Hebbian STDP rule for synaptic connections from inhibitory to excitatory neurons (iSTDP), which was in their model essential for maintaining a balance of excitation and inhibition. We did not find a need for such fast inhibitory plasticity. Instead, we set the strengths of inhibitory connections to fixed values. However, it would be interesting to study which types of iSTDP would lead to a self-organization of inhibitory dynamics that also supports blind source separation.

A soft WTA model for cortical circuits with lateral inhibition was previously studied by de Almeida et al. (2009). Consistent with our model, the authors arrived at the conclusion that lateral inhibition in cortical circuits gives rise to an adaptive *k*-WTA mechanism, rather than a strict *k*-WTA computation. However, because it was essential for their study that the circuit operates in the limit of no noise, their model is hard to compare with the stochastic model that we have examined. Further, the authors did not incorporate synaptic plasticity into their model, which is the focus of this paper.

Model M is also somewhat similar to other models (Habenschuss et al., 2013a; Nessler et al., 2013; Kappel et al., 2014). However, these studies did not model inhibition through feedback from inhibitory neurons. Instead, inhibition was provided in a symbolic manner as a normalization of network activity, leading to strict WTA behavior.

The emergent computational operation in our model, the extraction of superimposed components of input patterns, is closely related to ICA (Hyvärinen et al., 2004). Previous work in this direction includes the classical work by Földiák (1990) and implementations of ICA in artificial neural networks (Hyvärinen, 1999). It was shown by Bell and Sejnowski (1997) that ICA predicts features of neural tuning in primary visual cortex. A more recent model for a similar computational goal was proposed by Lücke and Eggert (2010). This model is more abstract and only loosely connected to cortical microcircuit motifs. ICA with spiking neurons was previously considered by Savin et al. (2010). The authors derived theoretical rules for intrinsic plasticity (i.e., rules for homeostasis of neurons) which, when combined with input normalization, weight scaling, and STDP, enable each neuron to extract one of a set of independent components of inputs. Although closely related in terms of the computational function, the data-based form of inhibition in our model M has quite different features. In Savin et al. (2010), the main purpose of inhibition is to decorrelate neuronal activity so that different neurons extract different features. Sparse activity is enforced there by intrinsic plasticity. Intrinsic plasticity in their model is thus required to work on a fast time-scale (the time-scale of input presentations). In contrast, sparse network activity in our data-based model M is enforced by inhibition. It is known that feedback inhibition is very fast and precise (Okun and Lampl, 2008), although it is unclear whether this is also true for intrinsic plasticity (Turrigiano and Nelson, 2004).

Experimentally testable predictions of our model

A main prediction of our model (Fig. 3) is the emergence of blind source separation of superimposed spike patterns. In addition, our model predicts that each of the identified basic patterns of the spike inputs becomes represented through some separate assem-

bly of pyramidal cells. Our model predicts that this effect takes place for any type of network input (e.g., also for artificially generated stimuli that the organism is never exposed to in a natural environment). This hypothesis can be tested experimentally (e.g., through optogenetic control).

In addition, our model predicts a specific role of synaptic connections among PV⁺ inhibitory cells (Fig. 6). They contribute to the experimentally found small time lag of just a few milliseconds by which inhibition trails excitation. This prediction can be tested experimentally by silencing synaptic connections among PV⁺ cells and measuring the impact on the lag between excitation and inhibition.

References

- Avermann M, Tomm C, Mateo C, Gerstner W, Petersen CC (2012) Microcircuits of excitatory and inhibitory neurons in layer 2/3 of mouse barrel cortex. *J Neurophysiol* 107:3116–3134. [CrossRef Medline](#)
- Bell AJ, Sejnowski TJ (1997) The ‘independent components’ of natural scenes are edge filters. *Vision Res* 37:3327–3338. [CrossRef Medline](#)
- Carandini M, Heeger DJ (2011) Normalization as a canonical neural computation. *Nat Rev Neurosci* 13:51–62. [CrossRef Medline](#)
- Cherry EC (1953) Some experiments on the recognition of speech, with one and with two ears. *J Acoust Soc Am* 25:975–979. [CrossRef](#)
- de Almeida L, Idiart M, Lisman JE (2009) A second function of gamma frequency oscillations: a E%-max winner-take-all mechanism selects which cells fire. *J Neurosci* 29:7497–7503. [CrossRef Medline](#)
- Douglas RJ, Martin KA, Whitteridge D (1989) A canonical microcircuit for neocortex. *Neural Comput* 1:480–488. [CrossRef](#)
- Ferguson KA, Huh CY, Amilhon B, Williams S, Skinner FK (2013) Experimentally constrained cal fast-firing parvalbumin-positive interneuron network models exhibit sharp transitions into coherent high frequency rhythms. *Front Comput Neurosci* 7:144. [CrossRef Medline](#)
- Fino E, Packer AM, Yuste R (2012) The logic of inhibitory connectivity in the neocortex. *Neuroscientist* 19:228–237. [CrossRef Medline](#)
- Földiák P (1990) Forming sparse representations by local anti-Hebbian learning. *Biol Cybern* 64:165–170. [CrossRef Medline](#)
- Grillner S, Graybiel A (2006) Microcircuits: the interface between neurons and global brain function, Vol 93. Cambridge, MA: Massachusetts Institute of Technology.
- Habenschuss S, Jonke Z, Maass W (2013a) Stochastic computations in cortical microcircuit models. *PLoS Comput Biol* 9:e1003311. [CrossRef Medline](#)
- Habenschuss S, Puh H, Maass W (2013b) Emergence of optimal decoding of population codes through STDP. *Neural Comput* 25:1371–1407. [CrossRef Medline](#)
- Harvey CD, Coen P, Tank DW (2012) Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484:62–68. [CrossRef Medline](#)
- Hebb DO (1949) *The organization of behavior*. New York: Wiley.
- Ho EC, Strüber M, Bartos M, Zhang L, Skinner FK (2012) Inhibitory networks of fast-spiking interneurons generate slow population activities due to excitatory fluctuations and network multistability. *J Neurosci* 32:9931–9946. [CrossRef Medline](#)
- Hyvärinen A (1999) Fast and robust fixed-point algorithms for independent component analysis. *IEEE Trans Neural Netw* 10:626–634. [CrossRef Medline](#)
- Hyvärinen A, Karhunen J, Oja E (2004) *Independent component analysis*. New York: Wiley.
- Isaacson JS, Scanziani M (2011) How inhibition shapes cortical activity. *Neuron* 72:231–243. [CrossRef Medline](#)
- Jolivet R, Rauch A, Lüscher HR, Gerstner W (2006) Predicting spike timing of neocortical pyramidal neurons by simple threshold models. *J Comput Neurosci* 21:35–49. [CrossRef Medline](#)
- Kappel D, Nessler B, Maass W (2014) Stdp installs in winner-take-all circuits an online approximation to hidden Markov model learning. *PLoS Comput Biol* 10:e1003511. [CrossRef Medline](#)
- Kerlin AM, Andermann ML, Berezovskii VK, Reid RC (2010) Broadly tuned response properties of diverse inhibitory neuron subtypes in mouse visual cortex. *Neuron* 67:858–871. [CrossRef Medline](#)
- Lee SH, Kwan AC, Zhang S, Phoumthippavong V, Flannery JG, Masmanidis SC, Taniguchi H, Huang ZJ, Zhang F, Boyden ES, Deisseroth K, Dan Y

- (2012) Activation of specific interneurons improves V1 feature selectivity and visual perception. *Nature* 488:379–383. [CrossRef Medline](#)
- Legenstein R, Jonke Z, Habenschuss S, Maass W (2017) A probabilistic model for learning in cortical microcircuit motifs with data-based divisive inhibition. [arXiv 1707.05182](#).
- Litwin-Kumar A, Doiron B (2014) Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nat Commun* 5:5319. [CrossRef Medline](#)
- Lücke J, Eggert J (2010) Expectation truncation and the benefits of preselection in training generative models. *J Machine Learn Res* 9999:2855–2900.
- Luczak A, McNaughton BL, Harris KD (2015) Packet-based communication in the cortex. *Nat Rev Neurosci* 16:745–755. [CrossRef Medline](#)
- Maass W (2000) On the computational power of winner-take-all. *Neural Comput* 12:2519–2535. [CrossRef Medline](#)
- 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. [CrossRef Medline](#)
- Ohki K, Chung S, Kara P, Hübener M, Bonhoeffer T, Reid RC (2006) Highly ordered arrangement of single neurons in orientation pinwheels. *Nature* 442:925–928. [CrossRef Medline](#)
- Okun M, Lampl I (2008) Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat Neurosci* 11:535–537. [CrossRef Medline](#)
- Packer AM, Yuste R (2011) Dense, unspecific connectivity of neocortical parvalbumin-positive interneurons: a canonical microcircuit for inhibition? *J Neurosci* 31:13260–13271. [CrossRef Medline](#)
- Pala A, Petersen CC (2015) In vivo measurement of cell-type-specific synaptic connectivity and synaptic transmission in layer 2/3 mouse barrel cortex. *Neuron* 85:68–75. [CrossRef Medline](#)
- Petersen CC, Crochet S (2013) Synaptic computation and sensory processing in neocortical layer 2/3. *Neuron* 78:28–48. [CrossRef Medline](#)
- Rumelhart DE, Zipser D (1985) Feature discovery by competitive learning. *Cogn Sci* 9:75–112. [CrossRef](#)
- Savin C, Joshi P, Triesch J (2010) Independent component analysis in spiking neurons. *PLoS Comput Biol* 6:e1000757. [CrossRef Medline](#)
- Shepherd GM (2004) *The synaptic organization of the brain*. New York: Oxford UP.
- Turrigiano GG, Nelson SB (2004) Homeostatic plasticity in the developing nervous system. *Nat Rev Neurosci* 5:97–107. [CrossRef Medline](#)
- Van Rijsbergen CJ (2004) *The geometry of information retrieval*, Vol 157. Cambridge, United Kingdom: Cambridge UP.
- van Rijsbergen CJ (1974) Foundation of evaluation. *J Documentation* 30:365–373. [CrossRef](#)
- Wilson NR, Runyan CA, Wang FL, Sur M (2012) Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature* 488:343–348. [CrossRef Medline](#)