

RESEARCH ARTICLE

Distribution of Orientation Selectivity in Recurrent Networks of Spiking Neurons with Different Random Topologies

Sadra Sadeh, Stefan Rotter*

Bernstein Center Freiburg & Faculty of Biology, University of Freiburg, Freiburg, Germany

*stefan.rotter@biologie.uni-freiburg.de



CrossMark
click for updates

 OPEN ACCESS

Citation: Sadeh S, Rotter S (2014) Distribution of Orientation Selectivity in Recurrent Networks of Spiking Neurons with Different Random Topologies. PLoS ONE 9(12): e114237. doi:10.1371/journal.pone.0114237

Editor: Thomas Wennekers, Plymouth University, United Kingdom

Received: June 13, 2014

Accepted: November 4, 2014

Published: December 3, 2014

Copyright: © 2014 Sadeh, Rotter. This is an open-access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), 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.

Funding: Funding by the German Ministry of Education and Research (BMBF; BFNT-F*T, grant 01GQ0830) and the German Research Foundation (DFG; grant EXC 1086) is gratefully acknowledged. The article processing charge was covered by the open access publication fund of the University of Freiburg. 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.

Abstract

Neurons in the primary visual cortex are more or less selective for the orientation of a light bar used for stimulation. A broad distribution of individual grades of orientation selectivity has in fact been reported in all species. A possible reason for emergence of broad distributions is the recurrent network within which the stimulus is being processed. Here we compute the distribution of orientation selectivity in randomly connected model networks that are equipped with different spatial patterns of connectivity. We show that, for a wide variety of connectivity patterns, a linear theory based on firing rates accurately approximates the outcome of direct numerical simulations of networks of spiking neurons. Distance dependent connectivity in networks with a more biologically realistic structure does not compromise our linear analysis, as long as the linearized dynamics, and hence the uniform asynchronous irregular activity state, remain stable. We conclude that linear mechanisms of stimulus processing are indeed responsible for the emergence of orientation selectivity and its distribution in recurrent networks with functionally heterogeneous synaptic connectivity.

Introduction

When arriving at the cortex from the sensory periphery, sensory signals are further processed by local recurrent networks. Indeed, the vast majority of all the connections a cortical neuron receives are from the cortical networks within which it is embedded and only a small fraction of connections are from feedforward afferents: The fraction of recurrent connections has been estimated to be as large as 80% [1]. What is the precise role of this recurrent network in sensory processing is not yet fully clear.

In the primary visual cortex of mammals like carnivores and primates, for instance, it has been proposed that the recurrent network might be mainly responsible for the amplification of orientation selectivity [2, 3]. Only a small bias provided by the feedforward afferents would be enough, and selectivity is then amplified by a non-linear mechanism implemented by the recurrent network. This mechanism is a result of the feature-specific connectivity assumed in the model, where neurons with similar input selectivities are connected to each other with a higher probability. This, in turn, could follow from the arrangement of neurons in orientation maps [4–6], which implies that nearby neurons have similar preferred orientations. As nearby neurons are also connected with a higher likelihood than distant neurons, feature-specific connectivity is a straight-forward result in this scenario.

Feature-specific connectivity is not evident in all species, however. In rodent visual cortex, for instance, a salt-and-pepper organization of orientation selectivity has been reported, with no apparent spatial clustering of neurons according to their preferred orientations [6]. As a result, each neuron receives a heterogeneous input from pre-synaptic sources with different preferred orientations [7].

Although an over-representation of connections between neurons of similar preferred orientations has been reported in rodents [8–12], presumably as a result of a Hebbian growth process during a later stage of development [13], such feature-specific connectivity is not yet statistically significant immediately after eye opening [10]. A comparable level of orientation selectivity, however, has indeed been reported already at this stage [10]. If cortical recurrent networks make a contribution to sensory processing at this stage, random recurrent networks should be chosen as a model [14–16]. Activity-dependent reorganization of the network, however, may still refine the connectivity and improve the performance of the processing later during development.

Here we study the distribution of orientation selectivity in random recurrent networks with heterogeneous synaptic projections, i.e. networks where the recurrent connectivity does not depend on the preferred feature of the input to the neurons. We show that in structurally homogeneous networks, the heterogeneity in functional connectivity, i.e. the heterogeneity in preferred orientations of recurrently connected neurons, is indeed responsible for a broad distribution of selectivities. A linear analysis of the network operation can account quite precisely for this distribution, for a wide range of network topologies including Erdős-Rényi random networks and networks with distance-dependent connectivity.

Methods

Network Model

In this study, we consider networks of leaky integrate-and-fire (LIF) neurons. For this spiking neuron model, the sub-threshold dynamics of the membrane potential $V_i(t)$ of neuron i is described by the leaky-integrator equation

$$\tau \dot{V}_i(t) + V_i(t) = RI_i(t). \tag{1}$$

The current $I_i(t)$ represents the total input to the neuron, the integration of which is governed by the leak resistance R , and the membrane time constant $\tau = 20$ ms. When the voltage reaches the threshold at $V_{th} = 20$ mV, a spike is generated and transmitted to all postsynaptic neurons, and the membrane potential is reset to the resting potential at $V_0 = 0$ mV. It remains at this level for short absolute refractory period, $t_{ref} = 2$ ms, during which all synaptic currents are shunted.

The response statistics of a LIF neuron, which is driven by randomly arriving input spikes, can be analytically solved in the stationary case. Assuming a fixed voltage threshold, V_{th} , the solution of the first-passage time problem in response to randomly and rapidly fluctuating input yields explicit expressions for the moments of the inter-spike interval distribution [17, 18]. In particular, the mean response rate of the neuron, r , in terms of the mean, μ , and variance, σ^2 , of the fluctuating input is obtained

$$r = F(\mu, \sigma) = \left[t_{ref} + \tau \sqrt{\pi} \int_{\tilde{V}_0}^{\tilde{V}_{th}} \underbrace{e^{u^2} (1 + \operatorname{erf}(u))}_{h(u)} du \right]^{-1} \tag{2}$$

with $\tilde{V}_{th} = (V_{th} - \mu)/\sigma$ and $\tilde{V}_0 = (V_0 - \mu)/\sigma$.

Employing a mean field ansatz, the above theory can be applied to networks of identical pulse-coupled LIF neurons, randomly connected with homogeneous in-degrees, and driven by external excitatory input of the same strength. Under these conditions, all neurons exhibit the same mean firing rate, which can be determined by a straight-forward self-consistency argument [19, 20]: The firing rate r is a function of the first two cumulants of the input fluctuations, μ and σ^2 , which are, in turn, functions of the input. If s is the input (stimulus) firing rate, and r is the mean response rate of all neurons in the network, respectively, we have the relation

$$\begin{aligned} \mu(s, r) &= \tau [J_s s + J_r r N \epsilon (f - g(1 - f))], \\ \sigma^2(s, r) &= \tau [J_s^2 s + J_r^2 r N \epsilon (f + g^2(1 - f))]. \end{aligned} \tag{3}$$

Here J_s denotes the amplitude of an excitatory post-synaptic potential (EPSP) of external inputs, and J_r denotes the amplitude of recurrent EPSPs. The factor g is the inhibition-excitation ratio, which fixes the strength of inhibitory post-synaptic potentials (IPSPs) to $-gJ_r$. Synapses are modeled as δ -functions, where the pre-synaptic current is delivered to the post-synaptic neuron instantaneously, after a fixed transmission delay of $d = 1.5$ ms.

The remaining structural parameters are the total number of neurons in the network, N , the connection probability, ϵ , and the fraction f of neurons in the

network that are excitatory ($N_{exc} = fN$), implying that a fraction $1 - f$ is inhibitory ($N_{inh} = (1 - f)N$). For all networks considered here we have used $f = 0.8$ and $g = 8$. J_s was always fixed at 0.1 mV. For all network connectivities, we fix the in-degree, separately for the excitatory and the inhibitory population, respectively. That is, each neuron, be it excitatory or inhibitory, receives exactly $\epsilon_{exc}N_{exc}$ connections randomly sampled from the excitatory population and $\epsilon_{inh}N_{inh}$ connections randomly sampled from the inhibitory population. Multiple synaptic contacts and self-contacts are excluded.

In our simulations, inputs are stationary and independent Poisson processes, denoted by a vector \vec{s} of average firing rates. Its i -th entry, s_i , corresponding to the average firing rate of the input to the i -th neuron, depends on the stimulus orientation θ and the input preferred orientation (PO) of the neuron θ_i^* according to

$$s_i(\theta) = s_b[1 + m \cos(2(\theta - \theta_i^*))]. \tag{4}$$

The baseline s_b is the level of input common to all orientations, and the peak input is $(1 + m)s_b$. The input PO is assigned randomly and independently to each neuron in the population. To measure the output tuning curves in numerical simulations, we stimulated the networks for 8 different stimulus orientations, covering the full range between 0° and 180° in steps of 22.5° . The stimulation at each orientation was run for 15 s, using a simulation time step of 0.1 ms. Onset transients (the first 150 ms) were discarded.

Linearized Rate Equations

To quantify the response of a network to tuned input, we first compute its baseline (untuned) output firing rate, r_b . This procedure is described elsewhere in detail [16], and we only recapitulate the main steps and equations here. If the attenuation of the baseline and amplification of the modulation is performed by two essentially independent processing channels in the network [16], the baseline firing rate can be computed from the fixed point equation

$$r_b = F(\mu(s_b, r_b), \sigma(s_b, r_b)), \tag{5}$$

the root of which can be found numerically [16, 20].

Now we linearize the network dynamics about an operating point defined by the baseline. First, we write the full nonlinear rate equation of the network as $\vec{r} = F(\vec{\mu}, \vec{\sigma})$. Here, the mean and the variance of the input are expressed, in matrix-vector notation, as

$$\begin{aligned} \vec{\mu}(\vec{s}, \vec{r}) &= \tau[J_s \vec{s} + W \vec{r}], \\ \vec{\sigma}^2(\vec{s}, \vec{r}) &= \tau[J_s^2 \vec{s} + V \vec{r}], \end{aligned} \tag{6}$$

where \vec{s} and \vec{r} are N -dimensional column vectors of input and output firing rates, respectively, and W is the weight matrix of the network. Its entry W_{ij} , the weight

of a synaptic connection from neuron j to neuron i , is either 0 if there is no synapse, J_r if there is an excitatory synapse, or $-gJ_r$ if there is an inhibitory synapse. Matrix V is the element-wise square of W , that is $V_{ij} = W_{ij}^2$.

The extra firing rate of all neurons, $\delta\vec{r} = \vec{r}_m$ (output modulation), in response to a small perturbation of their inputs, $\delta\vec{s} = \vec{s}_m$ (input modulation), is obtained by linearizing the dynamics about the baseline, i.e. about μ_b and σ_b (obtained from Eq. (3) evaluated at r_b and s_b)

$$\delta\vec{r} = \frac{\partial F(\mu, \sigma)}{\partial \mu} \Big|_{\mu_b, \sigma_b} \delta\vec{\mu} + \frac{\partial F(\mu, \sigma)}{\partial \sigma} \Big|_{\mu_b, \sigma_b} \delta\vec{\sigma}. \tag{7}$$

The partial derivatives of $F(\mu, \sigma)$ at this operating point can be computed from Eq. (2) as

$$\begin{aligned} \alpha = \frac{\partial F(\mu, \sigma)}{\partial \mu} \Big|_{\mu_b, \sigma_b} &= -F^2(\mu_b, \sigma_b) \tau \sqrt{\pi} \frac{\partial}{\partial \mu} \left[\int_{\tilde{V}_0}^{\tilde{V}_{th}} h(u) du \right] \Big|_{\mu_b} \\ &= -\tau \sqrt{\pi} F^2(\mu_b, \sigma_b) \left[h(\tilde{V}_{th}^b) \frac{\partial \tilde{V}_{th}}{\partial \mu} \Big|_{\mu_b} - h(\tilde{V}_0^b) \frac{\partial \tilde{V}_0}{\partial \mu} \Big|_{\mu_b} \right] \\ &= \frac{\tau \sqrt{\pi}}{\sigma_b} r_b^2 [h(\tilde{V}_{th}^b) - h(\tilde{V}_0^b)] \end{aligned} \tag{8}$$

and, in a similar fashion,

$$\begin{aligned} \beta = \frac{\partial F(\mu, \sigma)}{\partial \sigma} \Big|_{\mu_b, \sigma_b} &= -\tau \sqrt{\pi} F^2(\mu_b, \sigma_b) \left[h(\tilde{V}_{th}^b) \frac{\partial \tilde{V}_{th}}{\partial \sigma} \Big|_{\sigma_b} - h(\tilde{V}_0^b) \frac{\partial \tilde{V}_0}{\partial \sigma} \Big|_{\sigma_b} \right] \\ &= \frac{\tau \sqrt{\pi}}{\sigma_b^2} r_b^2 [h(\tilde{V}_{th}^b)(\tilde{V}_{th}^b - \mu_b) - h(\tilde{V}_0^b)(\tilde{V}_0^b - \mu_b)] \end{aligned} \tag{9}$$

where $F(\mu_b, \sigma_b) = r_b$, and \tilde{V}_{th}^b , \tilde{V}_0^b , μ_b and σ_b are the corresponding parameters evaluated at the baseline (for further details on this derivation, see [21]).

We also need to express $\delta\vec{\mu}$ and $\delta\vec{\sigma}$ in terms of the input perturbations. In fact, they can be written in terms of $\delta\vec{s}$ and $\delta\vec{r}$ from Eq. (6) as:

$$\begin{aligned} \delta\vec{\mu} &= \tau [J_s \delta\vec{s} + W \delta\vec{r}], \\ \delta\vec{\sigma} &= \tau \left[\frac{J_s^2}{2\sigma_b} \delta\vec{s} + \frac{V}{2\sigma_b} \delta\vec{r} \right]. \end{aligned} \tag{10}$$

For the total output perturbation, $\vec{r}_m = \delta\vec{r}$, we therefore obtain

$$\vec{r}_m = \alpha\tau[J_s\vec{s}_m + W\vec{r}_m] + \beta\tau\left[\frac{J_s^2}{2\sigma_b}\vec{s}_m + \frac{V}{2\sigma_b}\vec{r}_m\right]. \tag{11}$$

With the simulation parameters used here, our network typically operates in a fluctuation-driven regime of activity with a comparable level of input mean and fluctuations, $\mathbb{O}(\sigma) \approx \mathbb{O}(\mu)$. As a result, the contribution of the mean, $\alpha\vec{\mu}$, to output modulation in Eq. (11) is $\mathbb{O}(\sigma\alpha/\beta)$ larger than the contribution of the variance, $\beta\vec{\sigma}$. In the noise-dominated regime, \tilde{V}_0^b and \tilde{V}_{th}^b are small compared to μ_b in Eq. (9), and hence we can write $\beta \approx \frac{-\tau\sqrt{\pi}}{\sigma_b^2} r_b^2 \mu_b [h(\tilde{V}_{th}^b) - h(\tilde{V}_0^b)]$, yielding $\beta \approx -\frac{\mu_b}{\sigma_b} \alpha$. Thus, with a comparable level of mean and fluctuations, the contribution of the mean to output modulation is $\mathbb{O}(\sigma)$ larger than the contribution of the variance. In fact, the more the network operates in the noise-dominated regime, the more $\alpha\vec{\mu}$ becomes dominant over $\beta\vec{\sigma}$, making the second term on the right hand side of Eq. (11) negligible.

For the network shown in Fig. 1 and Fig. 2, for instance, $r_b \approx 5$ spikes/s. Given the general parameters of our simulation, we obtain $\mu_b = 7$ mV and $\sigma_b = 10$ mV. This yields $\tilde{V}_0^b = -0.7$ and $\tilde{V}_{th}^b = 1.3$, and finally $\alpha = 1.12$ and $\beta = 0.63$. In response to feedforward input perturbations, therefore, the contribution of the mean term ($\alpha\tau J_s \vec{s}_m$) is $\frac{2\alpha\sigma_b}{\beta J_s} \approx 250$ times the contribution of the variance term ($\beta\tau \frac{J_s^2}{2\sigma_b} \vec{s}_m$). In response to recurrent perturbation vectors with zero mean, both the mean term ($\alpha\tau W \vec{r}_m$) and the variance term ($\beta\tau \frac{V}{2\sigma_b} \vec{r}_m$) would respond with zero output, on average. The variance, in contrast, is not zero; a similar computation as in Eq. (3) yields $\tau(\alpha J_r)^2 r N \epsilon [f + g^2(1-f)]$ and $\tau(\frac{\beta J_r^2}{2\sigma_b})^2 r N \epsilon [f + g^4(1-f)]$, the terms resulting from the mean and variance contributions, respectively. That is, the mean contribution is dominant again by a factor of $\frac{4\sigma_b^2 \alpha [f + g^2(1-f)]}{\beta^2 J_r^2 [f + g^4(1-f)]} \approx 300$.

In the rest of our computation we therefore ignore the second part of the right hand side in Eq. (11) and approximate the output modulation as:

$$\vec{r}_m \approx \alpha\tau[J_s\vec{s}_m + W\vec{r}_m]. \tag{12}$$

We call

$$\zeta = \tau\alpha = \frac{\tau^2\sqrt{\pi}}{\sigma_b} r_b^2 [h(\tilde{V}_{th}^b) - h(\tilde{V}_0^b)] \tag{13}$$

the ‘‘linearized gain’’ and write the linearized rate equation of the network in response to small input perturbations as:

$$\vec{r}_m = \zeta W \vec{r}_m + \zeta J_s \vec{s}_m. \tag{14}$$

Linear and Supralinear Gains

The gain ζ is the linearized gain in the firing rate of a single LIF neuron in response to small changes in its mean input, while it is embedded in a recurrent network operating in its baseline AI state. That is, the extra firing rate, δr , of a neuron in response to a perturbation in its input, δs , when all other neurons are receiving the same, untuned input as before, divided by the input modulation weighted by its effect on the postsynaptic membrane $\zeta = \frac{\delta r}{J_s \delta s}$.

Alternative to the analytic derivation we pursued above, this gain can also be evaluated numerically by perturbing the baseline firing rate with an extra input, δs :

$$\zeta = \frac{\delta r}{J_s \delta s} = \frac{\overbrace{F(\mu(s_b + \delta s, r_b), \sigma(s_b + \delta s, r_b))}^{r_b + \delta r} - \overbrace{F(\mu(s_b, r_b), \sigma(s_b, r_b))}^{r_b}}{J_s \delta s}. \quad (15)$$

(Note that, as this is the response gain of an individual neuron to an individual perturbation in its input when all other neurons receive the same baseline input, it is not needed to consider the perturbation in the recurrent firing rate, r , in the baseline state.)

If this procedure is repeated for each δs , a numerical $f-I$ curve is obtained. This is the curve we have plotted in [Fig. 3A](#) as “Numerical perturbation”. If this curve was completely linear, it should not be much different from the results of our analytical perturbation ([Eq. \(13\)](#)), denoted by “Linearized gain” in [Fig. 3A](#). The results of the numerical perturbation, however, show some supralinear behavior, i.e. larger perturbations lead to a higher input-output gain. As a result, if we compute the gain at a perturbation size equal to the input modulation (s_m), a different gain is obtained. We use the term “stimulus gain” to refer to this supralinear gain at the modulation size of input (i.e. when $\delta s = s_m$):

$$\zeta_s = \frac{\delta r}{J_s s_m} = \frac{F(\mu(s_b + s_m, r_b), \sigma(s_b + s_m, r_b)) - F(\mu(s_b, r_b), \sigma(s_b, r_b))}{J_s s_m}. \quad (16)$$

This is shown by the red line in [Fig. 3A](#).

Linear Tuning in Recurrent Networks

Once we obtained the linearized gains at the baseline state of network operation, the linearized rate equation of the network for modulations about the baseline activity is obtained. Each neuron responds to the aggregate perturbation in its input with a gain obtained by the linearization formalism employed. The total perturbation consists of a feedforward component, which is the modulation in the input (stimulus) firing rate of the neuron, and a recurrent component, which is a linear sum of the respective output perturbations of the pre-synaptic neurons in

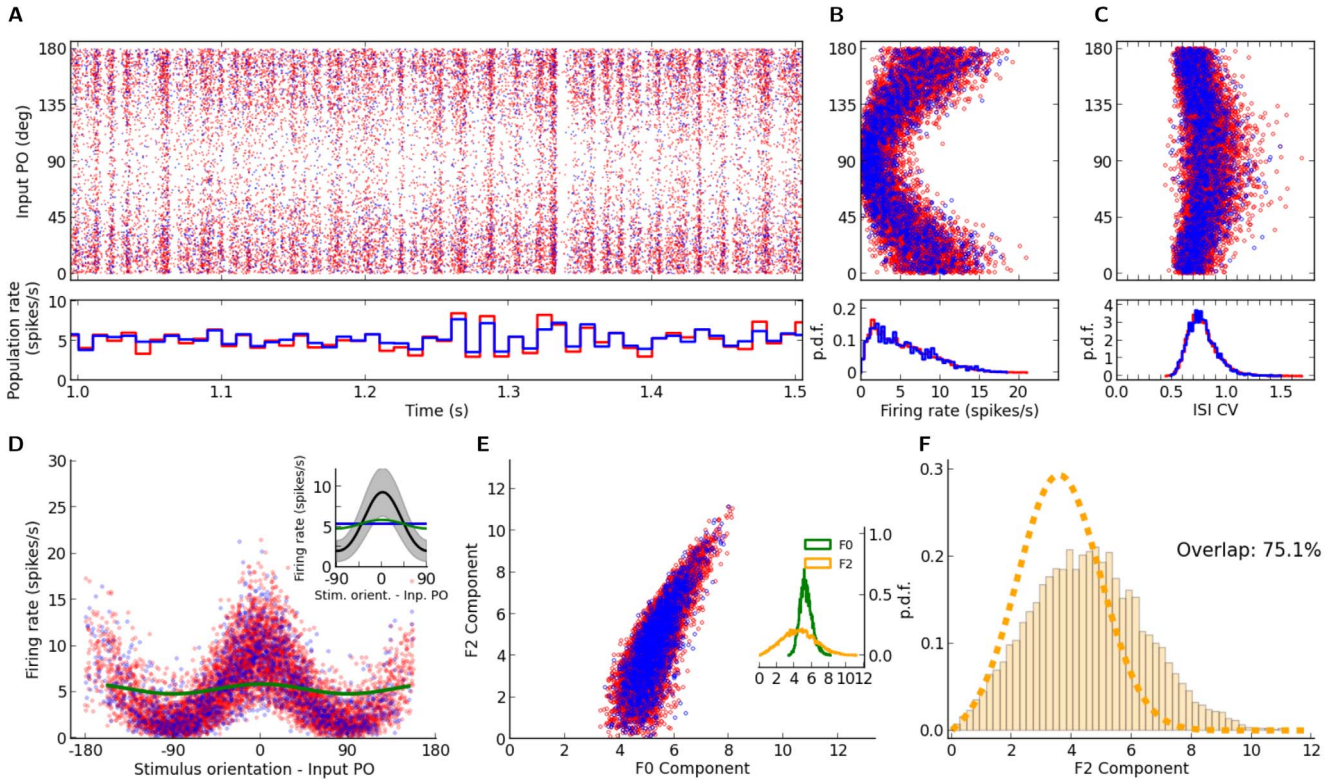


Figure 1. Distribution of orientation selectivity in networks with Erdős-Rényi random connectivity. (A) Raster plot of network activity in response to stimulus with orientation $\theta = 0^\circ$. Neurons are sorted according to their input preferred orientations, θ^* , indicated on the vertical axis. The histogram on the bottom shows the population firing rates, averaged in time bins of 10 ms width. Here, and in all other figures, red and blue colors denote excitatory and inhibitory neurons, or neuronal populations, respectively. (B) Average firing rates, for all neurons in the network, estimated from the spike count over the whole stimulation period ($t_{stim} = 15$ s). The distribution of firing rates over the population is depicted in the histogram at the bottom. (C) Coefficient of Variation (CV) of the inter-spike intervals (ISI), $CV_{ISI} = \text{std}(\text{ISI}) / \text{mean}(\text{ISI})$, computed for all neurons in the network with more than 10 spikes during the stimulation. The distribution of CV_{ISI} is plotted at the bottom. (D) Sample output tuning of 800 excitatory and 200 inhibitory neurons randomly chosen from the network, all aligned at their input preferred orientations. The input tuning (green, same as Eq. (4)) is normalized to the population average of the baseline (mean over all orientations) of output tuning curves. Inset: The mean (across population) of aligned output tunings are shown in black. The gray shading indicates mean \pm std extracted from the population. Linearly interpolated versions of individual tuning curves (generated at a resolution of 1°) have been used to compute mean and std of aligned tuning curves. The population average of the baseline (mean over all orientations) of output tuning curves is shown separately for excitatory and inhibitory populations with a red and a blue line, respectively (the lines highly overlap, since the average activity almost coincide for both populations). The normalized input tuning curve (green) is obtained by the same method as used for the main plot. (E) Scatter plot of F0 and F2 components, extracted from individual output tuning curves in the network. The individual distributions of F0 and F2 components over the population are plotted in the inset. (F) Distribution of single-neuron F2 components from a network simulation (histogram) compared with the prediction of our theory (dashed line, computed from Eq. (25)). To evaluate the goodness of match, the overlap of the empirical and predicted probability density functions (Pr_{emp} and Pr_{prd} , respectively) is computed as $\int_{-\infty}^{\infty} \min(\text{Pr}_{emp}(x'), \text{Pr}_{prd}(x')) dx'$. This returns an overlap index between 0% and 100%, corresponding to no overlap and perfect match of distributions, respectively. Parameters of the network simulation are: $N = 10000$, $\epsilon_{exc} = \epsilon_{inh} = 10\%$, $J_r = 0.25$ mV, $g = 8$, $s_b = 15000$ spikes/s, $J_s = 0.1$ mV, $m = 10\%$.

doi:10.1371/journal.pone.0114237.g001

the recurrent network. This can, therefore, be written, in vector-matrix notation, as:

$$\vec{r}_m = \zeta W \vec{r}_m + \zeta J_s \vec{s}_m. \quad (17)$$

If $\mathbb{1} - \zeta W$ is invertible, the output firing rates can be computed directly as

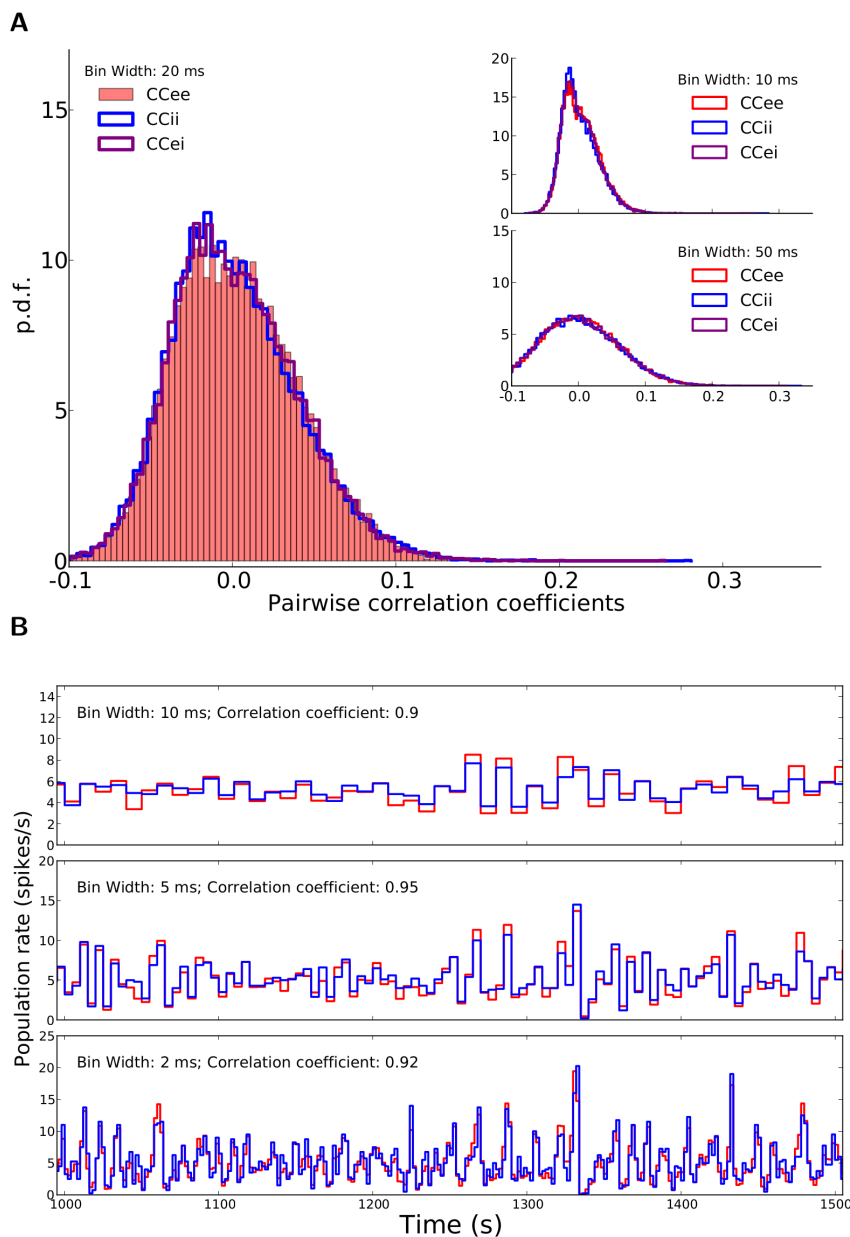


Figure 2. Correlations in the network. (A) Distribution of correlation coefficients for pairs of neurons in the network. For the example network of Fig. 1, the distribution of Pearson correlation coefficients (CC) between spike trains of pairs of neurons is plotted. 200 excitatory and 200 inhibitory neurons are randomly sampled from the network and all pairwise correlations (between pairs of excitatory, CC_{ee} , between pairs of inhibitory, CC_{ii} , and between excitatory and inhibitory, CC_{ei} , samples), based on spike counts in bins of width 20 ms are computed. The corresponding distributions for smaller (10 ms) and larger (50 ms) bins are shown in the inset (top and bottom, respectively). (B) The time series for the excitatory and inhibitory population spike counts indicate a fine balance on the population level. The correlation of activity between excitatory (red) and inhibitory (blue) populations is quite high on different time scales. The similarity of the temporal pattern of population activities is again quantified by the Pearson correlation coefficient.

doi:10.1371/journal.pone.0114237.g002

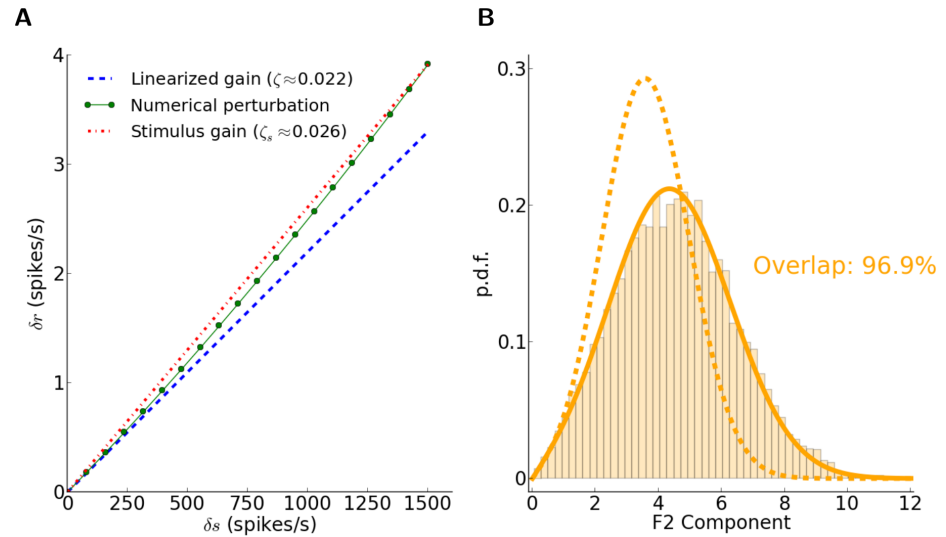


Figure 3. Supralinear neuronal gain affects the linear prediction. (A) Discrepancy of the linearized gain with the gain computed at stronger input modulations. The linearized gain of the neuron obtained analytically from Eq. (13) (dashed blue line) is compared with the numerical solution of Eq. (5) with an input perturbation equal to the modulation in the feedforward input, $\delta s = s_m = m s_b = 1500$ spikes/s (see Eq. (15) in Methods). The red line shows the corresponding linearized gain that would have been computed with this perturbation, ζ_s (Eq. (16)). (B) Comparison of our theoretical prediction of the distribution with ζ and ζ_s (dashed and solid lines, respectively). The overlap index of the improved prediction, i.e. when ζ is replaced by ζ_s in Eq. (30), has greatly increased.

doi:10.1371/journal.pone.0114237.g003

$$\vec{r}_m = (\mathbb{1} - \zeta \mathbf{W})^{-1} \zeta J_s \vec{s}_m = A \zeta J_s \vec{s}_m, \quad (18)$$

which can be further expanded into

$$\vec{r}_m = \sum_{k=0}^{\infty} (\zeta \mathbf{W})^k \zeta J_s \vec{s}_m. \quad (19)$$

Ignoring higher-order contributions $\mathcal{O}((\zeta \mathbf{W})^2)$, Eq. (19) can be approximated as

$$\vec{r}_m \approx (\mathbb{1} + \zeta \mathbf{W}) \zeta J_s \vec{s}_m. \quad (20)$$

Eq. (20) for each stimulus orientation returns the modulation of the output firing rate of all neurons in the network in response to a given input modulation.

We then assume that all inputs s_i are linearly tuned to the stimulus $\vec{\phi}$ according to

$$s_i(\vec{\phi}) = \psi_i^* + \langle \vec{\phi}_i^*, \vec{\phi} \rangle, \quad (21)$$

where ψ_i^* is the baseline rate in absence of stimulation and the vector $\vec{\phi}_i^*$ is the vector of preferred feature for the i -th neuron. The length of the vector that

represents the preferred feature $\|\vec{\phi}_i^*\|$ is the tuning strength. To ensure the linearity of operation, the firing rate $s_i(\vec{\phi})$ should remain always positive

$$0 \leq \min_{\vec{\phi}} s_i(\vec{\phi}) = \psi_i^* - \|\vec{\phi}_i^*\| \|\vec{\phi}\|. \tag{22}$$

If this condition is satisfied, the linearity of the tuning and positivity of firing rates remain compatible. If the condition is violated, partial rectification of the neuronal tuning curve follows and the linear analysis does not fully hold.

To obtain the operation of the network on input preferred feature vectors, we can write [Eq. \(20\)](#) for input tuning curves

$$\vec{r}_m(\vec{\phi}) = A\zeta J_s \vec{s}_m(\vec{\phi}) = A\zeta J_s \vec{\Phi}^* \vec{\phi}. \tag{23}$$

Here $\vec{\Phi}^*$ is a matrix the rows of which are given by the transposed preferred features $(\vec{\phi}_i^*)^T$. Therefore, all neurons in the recurrent network are again linearly tuned, with preferred features encoded by the rows of the matrix $AJ_s \vec{\Phi}^*$. From here we can compute the matrix of output feature vectors, $\vec{\Phi}_{out}^*$, as

$$\vec{\Phi}_{out}^* = A\zeta J_s \vec{\Phi}^* \approx (\mathbb{1} + \zeta W)\zeta J_s \vec{\Phi}^*. \tag{24}$$

The first term on the right-hand side is the weighted tuning vector of the feedforward input each neuron receives, and the second term is the mixture of tuning vectors of corresponding pre-synaptic neurons in the recurrent network.

Distribution of Orientation Selectivity

The length of the output feature vector represents the amplitude of the modulation component of output tuning curves. This is a measure of orientation selectivity, and we compute its distribution here.

Orientation is a two-dimensional feature, and the input feature vector ($\vec{\Phi}^*$ in [Eq.\(24\)](#)) is now a vector of two-dimensional input feature vectors (a vector of vectors). Its each entry, corresponding to the input orientation selectivity vector of each neuron, can, therefore, be determined by a length and a direction. The length of all vectors is $s_m = ms_b$, as all inputs have the same modulation, and the direction is twice the input PO of neurons (see [Eq. \(4\)](#)), which are drawn independently from a uniform distribution on $[0, \pi)$. They are assumed to be independent of the weight matrix W , implying the absence of feature specific connectivity.

The feedforward tuning vector of each neuron is accompanied by a contribution from the recurrent network ([Eq. 24](#)). For each neuron, the recurrent contribution is a vectorial sum of the input tuning vectors of its pre-synaptic neurons. According to the multivariate Central Limit Theorem, the summation of a large number of independent random variables leads to an approximate multivariate normal distribution of the output features. Tuning strength is given by the

length of output tuning vectors, $L = \|\phi\| = \sqrt{\phi_{\text{out},x}^* + \phi_{\text{out},y}^*}$. For a bivariate normal distribution with parameters μ_L and σ_L^2 , we can compute the distribution of this length

$$A_2(L; \mu_L, \sigma_L^2) = \frac{L}{\sigma_L^2} e^{-(L^2 + \mu_L^2)/(2\sigma_L^2)} I_0\left(\frac{L\mu_L}{\sigma_L^2}\right), \tag{25}$$

where

$$I_0(z) = \frac{1}{\pi} \int_0^\pi e^{z \cos(\psi)} d\psi$$

is the modified Bessel function of the first kind and zeroth order. Therefore, we only need to compute the mean and the variance of the resulting distribution.

The mean of the distribution μ_L is equal to the length of feedforward feature vector, $\zeta J_s s_m$. This is because the expected value of the contribution of the recurrent network vanishes in each direction

$$\mathbb{E}[\Phi_{\text{rec},x}^*] = \mathbb{E}[\zeta W \zeta J_s s_m \cos(2\Theta^*)] = \zeta^2 J_s s_m \mathbb{E}[W] \mathbb{E}[\cos(2\Theta^*)] = 0. \tag{26}$$

W and Θ^* denote, respectively, the random variables from which the weights and input POs are drawn. A similar computation yields $\mathbb{E}[\Phi_{\text{rec},y}^*] = 0$. Here we have used the property that the two random variables W and Θ^* are independent, and that all orientations are uniformly represented in the input ($\mathbb{E}[\cos(2\Theta^*)] = 0$). As a result, we obtain

$$\mathbb{E}[L] = \mathbb{E}[\Phi_{\text{ffw}}^* + \Phi_{\text{rec}}^*] = \zeta J_s s_m. \tag{27}$$

The recurrent contribution does not, on average, change the length of output feature vectors. However, it creates a distribution of selectivity, which can be quantified by its variance

$$\begin{aligned} \text{Var}[\Phi_{\text{rec},x}^*] &= (\zeta^2 J_s s_m)^2 \text{Var}[W] \text{Var}[\cos(2\Theta)] \\ &= (\zeta^2 J_s s_m)^2 \text{Var}[W] \int_0^\pi \cos^2(2\Theta) d\Theta \\ &= \frac{1}{2} (\zeta^2 J_s s_m)^2 \text{Var}[W]. \end{aligned} \tag{28}$$

Again, we have exploited the independence of random variables W and Θ , and the uniform representation of input POs ($\mathbb{E}[\cos(2\Theta^*)] = 0$), to factorize the variance, i.e. $\text{Var}[W \cos(2\Theta)] = \text{Var}[W] \text{Var}[\cos(2\Theta)]$. Similar computation yields the same variance for the second dimension.

For our random networks, the weights for each row of the weight matrix are drawn from a binomial distribution, W . The number of non-zero elements is determined by connection probabilities (ϵ_{exc} and ϵ_{inh} for excitation and inhibition respectively), and each non-zero entry is weighted by the synaptic strength (J_r and

$-gJ_r$ for excitation and inhibition respectively). The variance $\text{Var}[W]$ can therefore be computed explicitly:

$$\text{Var}[W] = J_r^2 N \epsilon (1 - \epsilon) [f + g^2(1 - f)]. \quad (29)$$

For more complex connectivities, the variance can be numerically computed from the weight matrix. For our networks here, the mean and the variance of the distribution of output tuning vectors can, therefore, be expressed as

$$\mu_L = \zeta J_s s_M \quad \text{and} \quad \sigma_L^2 = \frac{1}{2} (\zeta^2 J_s s_M)^2 J_r^2 N \epsilon (1 - \epsilon) [f + g^2(1 - f)]. \quad (30)$$

For an output tuning curve with a cosine shape, $R(\theta) = R_b + R_m \cos(\theta - \theta^*)$, the tuning strength we introduced above corresponds to R_m , namely the modulation (F2) component of the tuning curve. R_b is also obtained as the baseline firing rate of the network, r_b , from Eq. (5). To compare the prediction with the result of our simulations, we compute the mean and modulation of individual output tuning curves from the simulated data. Mean and modulation are taken from the zeroth and the second Fourier components of each tuning curve (F0 and F2 components), respectively. The distribution given by Eq. (25) should, therefore, precisely match the distribution of modulation (F2) component of output tuning curves obtained from simulations, if our linear analysis grasps the essential mechanisms of orientation selectivity in model recurrent networks.

Results

Erdős-Rényi Random Networks

We first study excitatory-inhibitory Erdős-Rényi random networks of LIF neurons (Eq. (1)) with a doubly fixed in-degree, namely where both the excitatory in-degree and the inhibitory in-degree is fixed for both excitatory and inhibitory neurons. Figs. 1A–C show the response of a network with $J_r = 0.25$ mV and $\epsilon_{\text{exc}} = \epsilon_{\text{inh}} = 0.1$ to the stimulus of 0° orientation. The network with these parameters operates in the fluctuation-driven regime, which shows asynchronous-irregular (AI) dynamics (Fig. 1A), with low firing rates (Fig. 1B) and high variance of inter-spike intervals (ISI) (Fig. 1C). The network at this regime is capable of amplifying the weak tuning of the input, as it is reflected both in the network tuning curve in response to one orientation (Fig. 1B) and in individual tuning curves in response to different stimulus orientations (Fig. 1D).

The joint distribution of the modulation (F2) component of (individual) output tuning curves and the respective baseline (F0) component (Fig. 1E) shows that the average values of these two components have become comparable after network operation. However, the F2 component has a much broader distribution (Fig. 1E, inset). The distribution predicted by our theory (Eq. (25)) matches partially with the distribution measured in the simulations (Fig. 1F). The degree

of match is quantified by an index, which assesses the overlap area of the two probability distributions.

As our analysis is based on the assumption of linearity of network interactions, the result of our theoretical prediction holds only if the network is operating in the linear regime. Any violation of our linear scheme would, therefore, lead to a deviation of the linear prediction from the measured distribution. The remaining discrepancy should, therefore, be attributed to any factor which invalidates our approximation scheme here. Possible contributing factors of this sort in our networks are partial rectification of tuning curves, correlations and synchrony in the network providing the input, and supralinearity of neuronal gains.

Partial rectification of firing rates is obvious in [Fig. 1B](#). However, this does not seem to be a very prominent effect. Only a small fraction of the population is strictly silent, as is evident in the distribution of firing rates ([Fig. 1B](#), bottom). Correlations, in contrast, seems to be a more important contributor, as is reflected in the raster plot of network activity ([Fig. 1A](#)).

To investigate the possible contribution of correlations in the distribution of orientation selectivity, we plotted the distribution of pairwise correlations in the network ([Fig. 2](#)). Although the distribution of pairwise correlations has a very long tail, on average correlations are very small in the network ([Fig. 2A](#)). This is the case for excitatory-excitatory, excitatory-inhibitory, and inhibitory-inhibitory correlations, and there is the same trend when spike counts are computed for different bin widths ([Fig. 2A](#), insets). Low pairwise correlations in the network are a result of recurrent inhibitory feedback, which actively decorrelates the network activity [22–24]. As illustrated in [Fig. 2B](#), upsurges in the population activity of excitatory neurons are tightly coupled to a corresponding increase in the activity of the inhibitory population. However, the cancellation is not always exact and some residual correlations remain.

Since each neuron receives random inputs from 10% of the population, approximately the same correlation of excitation and inhibition is, on average, also expected in the recurrent input to each neuron. Note that, as our networks are inhibition-dominated, the net recurrent inhibition would be stronger than the net recurrent excitation (indeed twice as strong, given the parameters we have used). Altogether, this implies that inhibition is capable of fast tracking of excitatory upsurges ([Fig. 2B](#)) such that fast fluctuations in the population activity would not be seen in the recurrent input from the network.

Finally, the single-neuron gain that we computed by linearization ([Eq. \(13\)](#)) could be a source of mismatch, as for a highly non-linear system it might only be valid for small perturbations in the input, and not for stronger modulations. This is shown in [Fig. 3A](#), where the linearized gain, ζ , from [Eq. \(13\)](#) is compared with ζ_s , the numerically obtained neuronal gain (see [Eq. \(15\)](#) in Methods) when the perturbation has the size of the input modulation, $s_m = ms_b$. This gain could be approximated analytically by expanding [Eq. \(5\)](#) to higher order terms. Here, however, we have computed this gain numerically ([Eq. \(16\)](#)).

When the prediction of [Eq. \(25\)](#) is repeated with the new gain (ζ_s), a great improvement in the match between the measured and predicted distributions is

indeed observed (Fig. 3B). We therefore concluded that the main source of mismatch in our prediction was our misestimate of the actual neuronal gains. Other sources of nonlinearity, like rectification and correlations, could therefore be responsible for the remaining discrepancy of distributions (less than 5% in the regime considered here). However, given so many possible sources of nonlinearity in our networks, both at the level of spiking neurons and network interactions, it is indeed quite surprising that a linear prediction works so well.

A remark about rectification in our networks should be made at this point. In the type of networks we are considering here, rectification is in fact not a single-neuron property, i.e. only the result of a rectification effect due to the spike threshold in the LIF neuron. This is not the case as the linearized gain of neurons within the network (Eq. (13)) implies a non-zero response even to small perturbations in the input. This is a result of (internally generated) noise within the recurrent network, as a consequence of balance of excitation and inhibition, which smoothens the embedded $f-I$ curve [25, 26]. Rectification could therefore only happen at the level of network, e.g. by increasing the amount of inhibition.

As our networks are inhibition-dominated, increasing the recurrent coupling would be one way to increase the inhibitory feedback within the network. This can be done in two different ways, either by increasing the connection density or by increasing the weights of synaptic connections. The first strategy is tried in Fig. 4A, where the connection probability has been increased (from $\epsilon_{\text{exc}} = \epsilon_{\text{inh}} = 0.1$ to $\epsilon_{\text{exc}} = \epsilon_{\text{inh}} = 0.2$). The second strategy is added to the first in Fig. 4B, where an increase in the connection density is accompanied by an increase in synaptic weights (from $J_r = 0.25\text{mV}$ to $J_r = 0.5\text{mV}$). In both cases, however, a significant rectification of tuning curves did not result, and the prediction of our linear theory still holds.

This unexpected effect can be explained intuitively as follows: An increase in recurrent coupling not only decreases the baseline firing rate of the network, but also changes neuronal gains (ζ and ζ_s). A crucial factor in determining this gain is the average membrane potential of neurons in the network, which in turn sets the mean distance to threshold. The larger the mean distance to threshold is in the network, the less is the neuronal gain. This in turn decreases the mean F2 component of output tuning curves. As a result, with a reduced baseline firing rate, a significant rectification of tuning curves still does not follow, as output modulation components have been scaled down by a comparable factor. This is indeed the case in networks of Fig. 4, where the mean (over neurons) membrane potential (temporally averaged) and the neuronal gains have both been decreased compared to the network of Fig. 1 (results not shown; for a detailed analysis, see [16]).

Networks With Distance-Dependent Connectivity

To extend the scope of the linear analysis, we asked if our theory can also account for networks with different statistically defined topologies. In particular, we considered networks with a more realistic pattern of distance-dependent

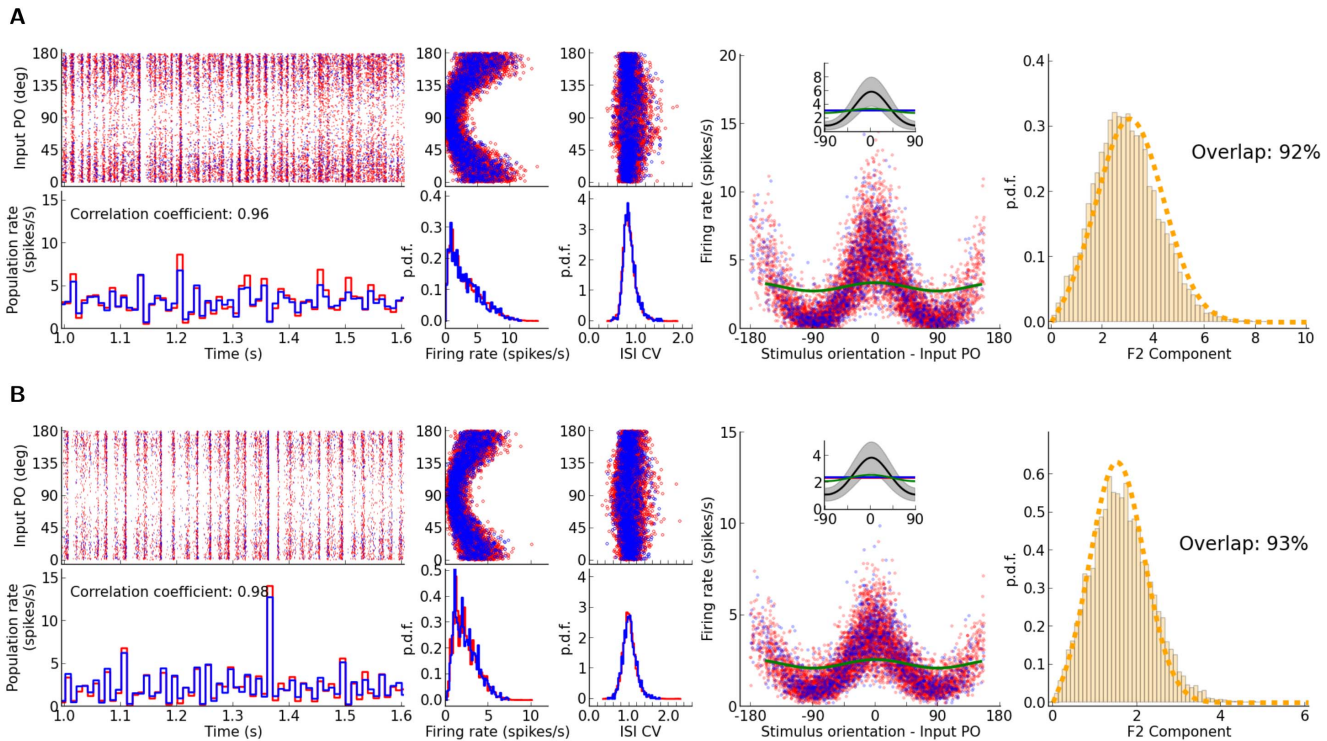


Figure 4. The impact of the strength of recurrent coupling on the distribution of selectivities. The figure layout is similar to Fig. 1 (panel (E) not included), shown are networks with stronger recurrent couplings. In (A), the recurrent coupling is increased by doubling the connection density; in (B), this is further enhanced by doubling all recurrent weights. The parameters of network simulations are: (A) $N = 10000$, $\epsilon_{exc} = \epsilon_{inh} = 20\%$, $J_r = 0.25$ mV, $g = 8$, $s_b = 15000$ spikes/s, $J_s = 0.1$ mV, $m = 10\%$, and (B) $N = 10000$, $\epsilon_{exc} = \epsilon_{inh} = 20\%$, $J_r = 0.5$ mV, $g = 8$, $s_b = 15000$ spikes/s, $J_s = 0.1$ mV, $m = 10\%$. The predicted distributions are computed by considering ζ_s (see Fig. 3).

doi:10.1371/journal.pone.0114237.g004

connectivity: Each neuron is assigned a random position in a two-dimensional rectangle representing a 1 mm \times 1 mm flat sheet of cortex (Fig. 5A). The probability of having a connection between a pre-synaptic excitatory (inhibitory) neuron to a given post-synaptic neuron falls off as a Gaussian function with distance, with parameter σ_{exc} (σ_{inh}). Similar to the Erdős-Rényi random networks considered before, we fix the in-degree, i.e. each neuron receives exactly $\epsilon_{exc}N_{exc}$ excitatory and $\epsilon_{inh}N_{inh}$ inhibitory connections. Multiple synaptic contacts and self-contacts are not allowed.

The connectivity profile is illustrated in Figs. 5B, C. The pre-synaptic sources of a sample neuron are plotted in Fig. 5B, for $\sigma_{exc} = \sigma_{inh} = 0.55$ mm. The resulting distribution of the distances of connected neurons, for the example neuron and for the entire population, is shown in Fig. 5C.

Note that the connectivity depends only on the physical distance. As input preferred orientations are assigned randomly and independently of the actual position of neurons in space, distance-dependent connectivity does not imply any feature-specific connectivity. That is, neither a spatial nor a functional map of orientation selectivity is present here.

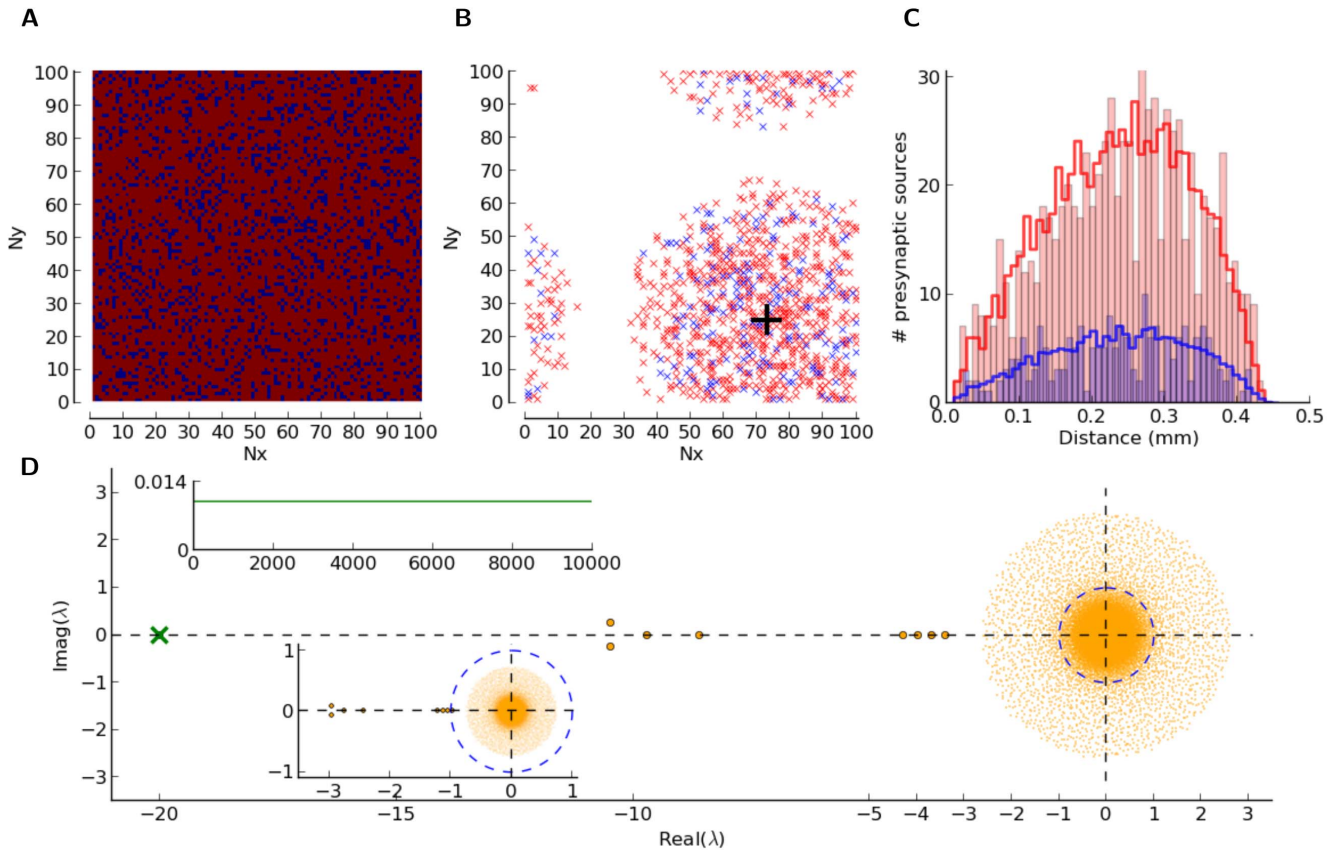


Figure 5. Networks with distance-dependent connectivity. (A) Random positioning of $N_{exc} = 8000$ excitatory (red) and $N_{inh} = 2000$ inhibitory (blue) neurons in a square, representing a flat $1\text{ mm} \times 1\text{ mm}$ sheet of cortex, wrapped to a torus. (B) For a sample (excitatory) neuron (large black cross), positions of excitatory (red) and inhibitory (blue) pre-synaptic neurons are explicitly shown as little crosses. A Gaussian connectivity profile with $\sigma_{exc} = \sigma_{inh} = 0.55\text{ mm}$ was assumed. For each post-synaptic neuron, we fixed the number of randomly drawn pre-synaptic connections of either type, i.e. $C_{exc} = \epsilon_{exc} N_{exc}$ and $C_{inh} = \epsilon_{inh} N_{inh}$ ($\epsilon_{exc} = \epsilon_{inh} = 10\%$). Multiple synapses and self-coupling were not allowed. (C) Histogram of distances to pre-synaptic neurons for the sample neuron (bars) and for the entire population (lines). (D) Eigenvalue spectrum of the weight matrix, W . Weights are normalized by the reset voltage, $V_{reset} = V_{th} - V_0 = V_{th}$, leading to $w_{ij} = J_r/V_{th}$ or $-g_r/V_{th}$, depending on whether the synapse is excitatory or inhibitory, respectively. We used $J_r = 0.5\text{ mV}$. For better visibility, the eigenvalues outside the bulk of the spectrum are shown by larger dots. The green cross marks the eigenvalue corresponding to the uniform eigenmode, which is plotted in the top inset. Re-normalized spectrum, according to the gain ζ_s , is shown in the bottom inset; i.e. $w_{ij} = \zeta_s J_r$ and $-g_{sJ_r}$, for excitatory and inhibitory connections, respectively.

doi:10.1371/journal.pone.0114237.g005

Before discussing the simulations of the spiking networks, it is informative to look at the eigenvalue spectrum of the associated weight matrix, W . It is plotted, for $J_r = 0.5\text{ mV}$ and $\epsilon_{exc} = \epsilon_{inh} = 10\%$, in Fig. 5D. Each entry of the matrix is normalized by the reset voltage, $V_{reset} = V_{th} - V_0 = V_{th}$, for the eigenvalue spectrum shown in the main panel. The effective firing rate equation of the network can then be written as $\tau d\vec{r}/dt = -\vec{r} + \frac{1}{V_{th}} [W\vec{r} + J_s\vec{s}]$. The exceptional eigenvalue (green cross) corresponding to the uniform eigenvector (inset, top) and the bulk of eigenvalues (orange dots) are the structural properties that this network has in common with the previous Erdős-Rényi networks (not shown). There is, however, a small number of additional (in this case, 8) eigenvalues in between, which are the consequence of the specific realization of our

distance-dependent connectivity here. The corresponding eigenmodes will, in principle, affect the response of the network, both in its spontaneous state and in response to stimulation.

All these eigenvalues have, however, negative real parts. They will, therefore, ensure the stability of the linearized network dynamics, as far as these eigenmodes are concerned. The bulk of the spectrum, in contrast, also comprises eigenvalues with real parts larger than 1, which implies an instability. An alternative normalization of the weight matrix according to the neuronal gain ζ_s (Fig. 5, inset, bottom; see also [16]), however, does not render these modes unstable.

Here, we are resorting to a linearized rate equation describing the response of the network to (small) perturbations, $\tau d\vec{r}/dt = -\vec{r} + \zeta_s[W\vec{r} + J_s\vec{s}]$ (see Eq. (17) in Methods). The eigendynamics corresponding to the common-mode (green cross) is faster, and hence it relaxes to the fixed point more rapidly than the other eigenmodes. The common mode effectively leads to the uniform, baseline state of the network (reflected in the baseline firing rate, r_b), about which the network dynamics has indeed been linearized in our linear prediction. The effect of other eigenmodes, in the stationary state, should therefore be computed by considering the linearized gain about that uniform, baseline state.

Simulation results for a network with this connectivity are illustrated in Fig. 6. Inspection of the spiking activity of the network (Fig. 6A) does not suggest a behavior very different from the behavior of random networks shown in Fig. 1. The irregularity of firing is, however, more pronounced, as the variance of interspike intervals is larger (Fig. 6C); the ISI CV has indeed a distribution about 1, which is more similar to the strongly coupled networks described in Fig. 4.

Similar to Erdős-Rényi networks, networks with distance-dependent local connectivity are capable of amplifying the weak tuning of the input signal, and comparable levels of baseline (F0) and modulation (F2) components are emerging (Fig. 6E). When the predicted distribution of F2 components is obtained applying the normalization by the linear gain ζ_s , a very good match to the measured distribution is obtained (Fig. 6F), comparable to predictions in Fig. 4, and only slightly worse than the prediction in Fig. 1.

Although partial rectification of tuning curves seems to be negligible in the example shown (Fig. 6B), correlations in the network could still be responsible for the remaining discrepancy. Moreover, size and structure of correlations in the network might be different here as compared to random networks due to non-homogeneous connectivity. Distance-dependent connectivity implies that connectivity is locally dense, which can lead to more shared input and this way impose strong correlations at the output.

In fact, however, pairwise correlations do not seem to be systematically larger than in random networks Fig. 2A, judged by the distribution of Pearson correlation coefficients (Fig. 7A). In contrast, the fluctuations in the activity of excitatory and inhibitory populations seem to be even less correlated (compare Fig. 7B with Fig. 2B). Occasional partial imbalance of excitatory and inhibitory input may therefore cause systematic distortions of our linear prediction.

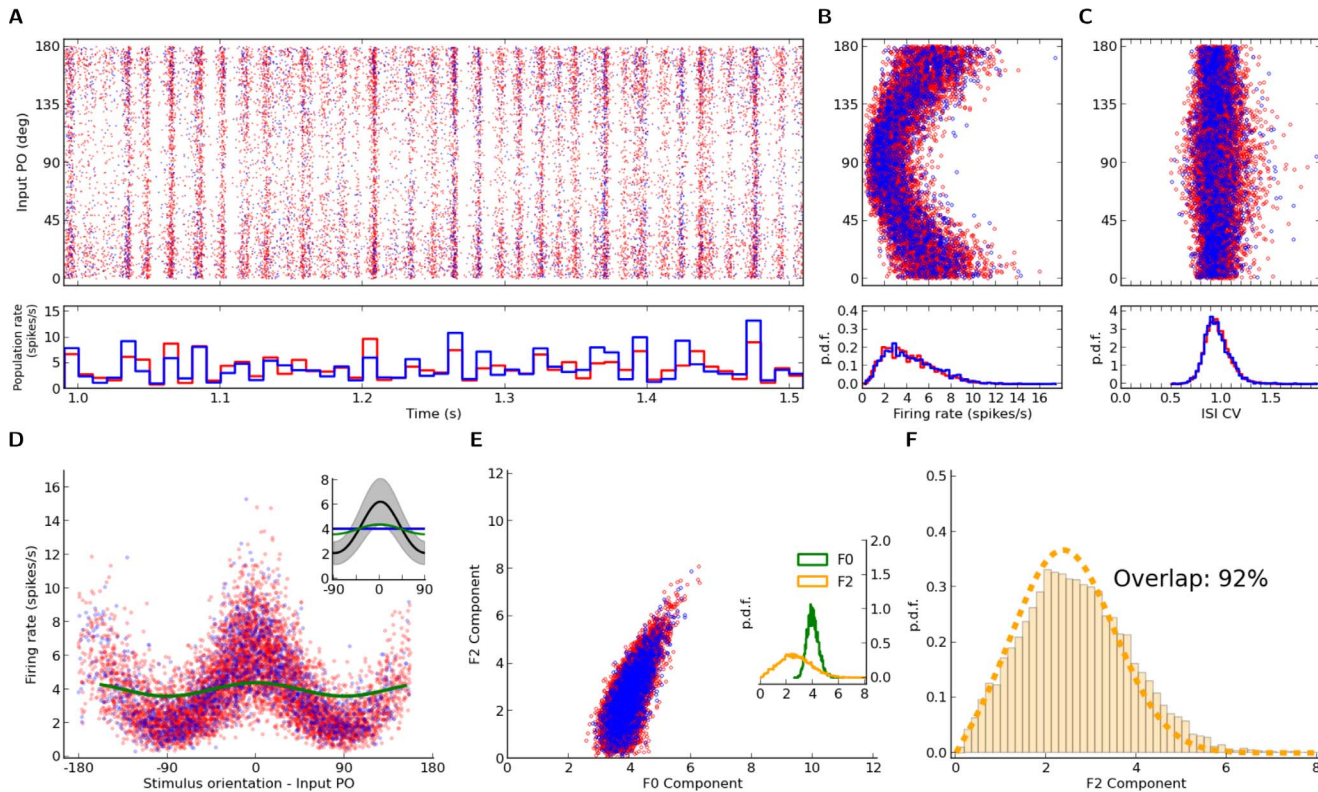


Figure 6. Distribution of orientation selectivity in a network with distance-dependent connectivity. Same figure layout as Fig. 1, for a network with distance-dependent connectivity, similar to Fig. 5. Parameters of the network simulation are: $N = 10000$, $\epsilon_{exc} = \epsilon_{inh} = 10\%$, $\sigma_{exc} = \sigma_{inh} = 0.55$ mm, $J_r = 0.5$ mV, $g = 8$, $s_b = 15000$ spikes/s, $J_s = 0.1$ mV, $m = 10\%$. Note that the distribution of F2 components is computed by using the stimulus gain ζ_s , as in Fig. 3.

doi:10.1371/journal.pone.0114237.g006

Another potential contributor to the discrepancy of predictions are the different structural properties of these networks, reflected among other things in their respective eigenvalue spectrum. It is therefore informative to look more carefully into the eigenvalues which mark the difference to Erdős-Rényi networks, i.e. the ones localized between the bulk spectrum and the exceptional eigenvalue corresponding to the common-mode. To evaluate this, the first ten eigenvectors (corresponding to the ten largest eigenvalues sorted by their magnitude) of the network are plotted (Fig. 8A). The first eigenvector is the uniform vector (common-mode), and the tenth one is hardly distinguishable from noise. (Note that the corresponding eigenvalue is already part of the bulk.) In between, there are eight eigenvectors with non-random spatial structure.

These eigenvectors reflect the specific sample from the network ensemble we are considering here, and they can, in principle, prefer a specific pattern of stimulation in the input. While other patterns of input stimulation would be processed by the network W with a small gain, any input pattern matching these special eigenmodes would experience the highest gain (in absolute terms) from the network. The corresponding eigenvalues λ have, however, a negative real part, therefore these modes would in this case be attenuated: the corresponding

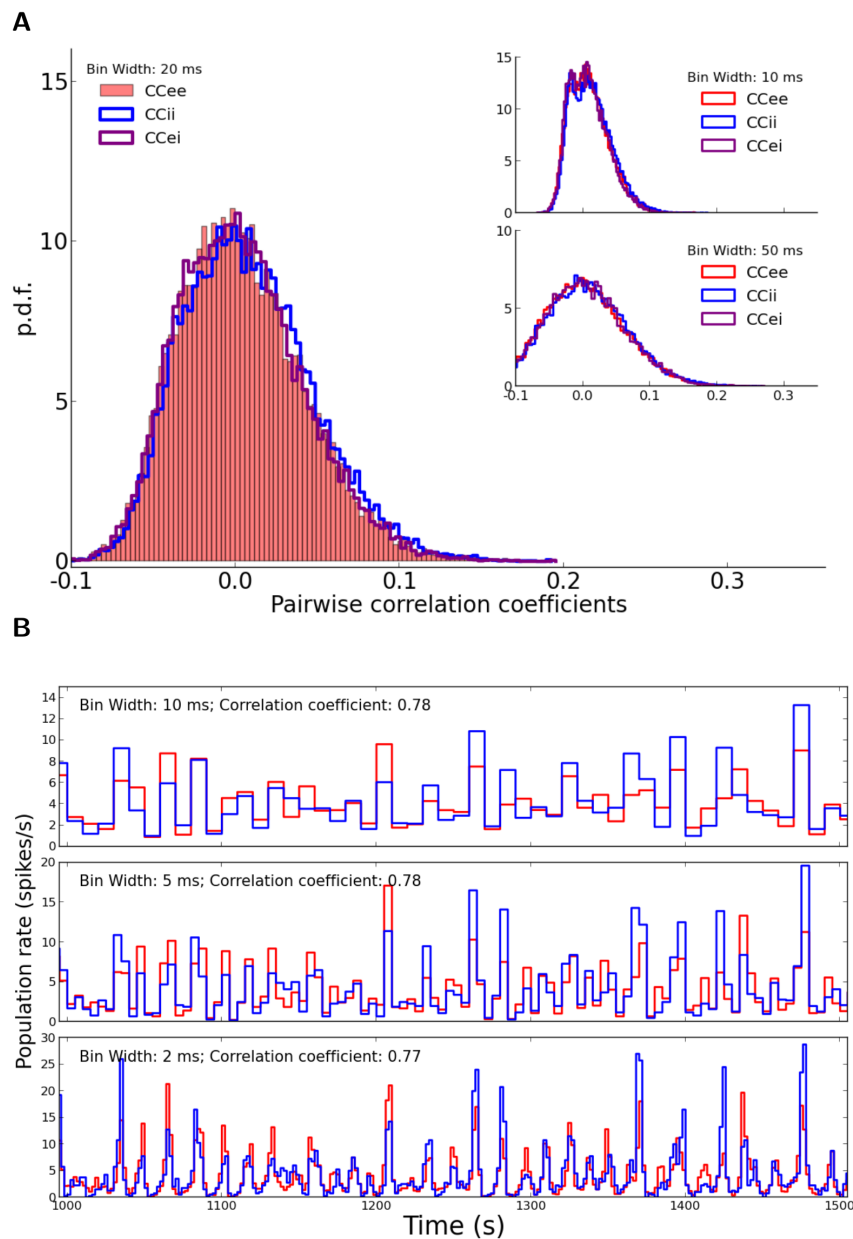


Figure 7. Correlations in a network with distance-dependent connectivity. Distribution of correlation coefficients for pairs of neurons (A) and temporal correlation of population activities (B) in the example network of Fig. 5 with distance-dependent connectivity. Other conventions are similar to Fig. 2.

doi:10.1371/journal.pone.0114237.g007

eigenvalues of the operator $A = (\mathbb{I} - W)^{-1}$ that yields the stationary firing rate vector, namely $\lambda_A = \frac{1}{1 - \lambda_W}$, would then be very small.

We do not, however, explicitly represent any of these patterns in our stimuli. The stimuli considered in this work can be broken down to a linear sum of the common-mode (i.e. the first eigenvector) and the modulation component (i.e. a random pattern, as preferred orientations are assigned randomly and indepen-

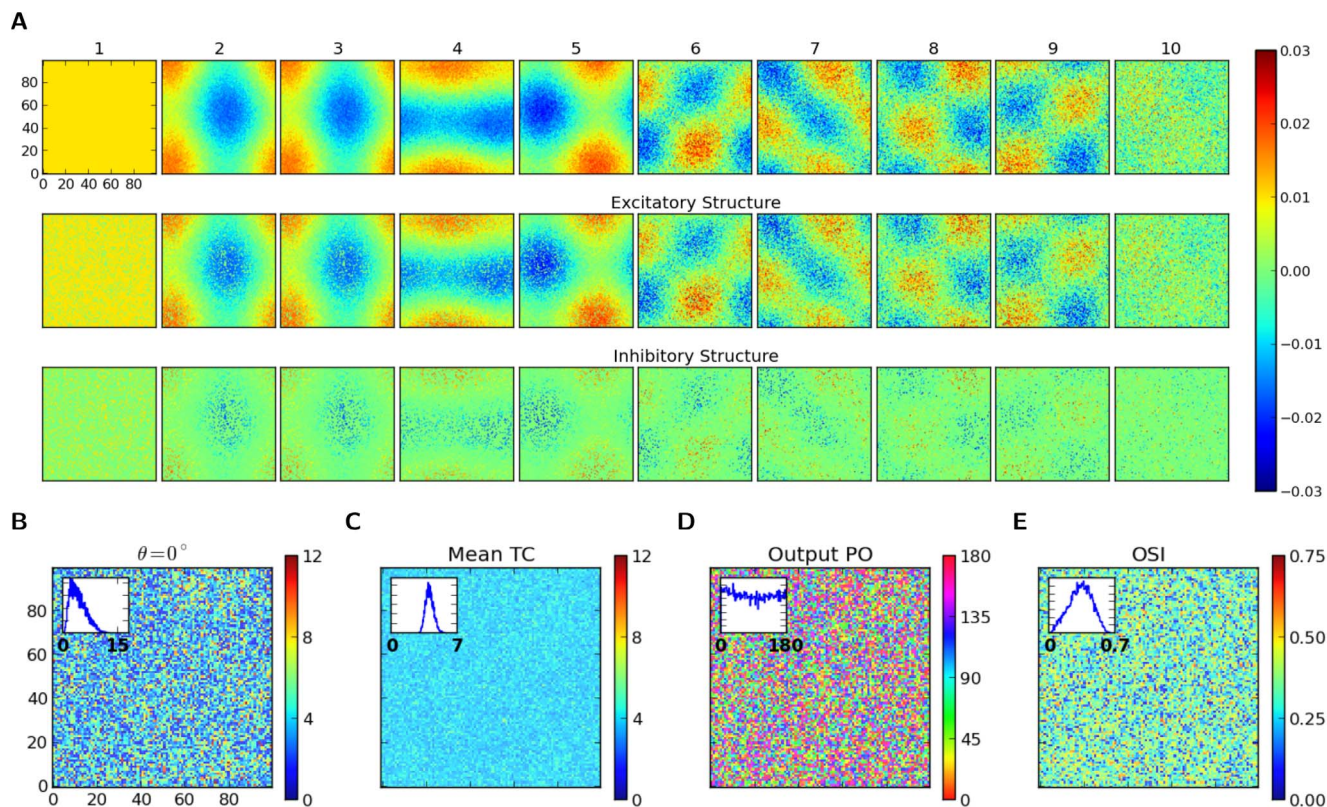


Figure 8. Structure and dynamics of a network with distance-dependent connectivity. (A) First ten eigenvectors, corresponding to the ten eigenvalues of largest magnitude, are plotted for the sample network described and discussed in Figs. 5 and 6. For each eigenvector, the value of the vector corresponding to each neuron is plotted at the respective spatial position of the neuron (as in Fig. 5A). In the first row, this is shown for all neurons, and in the bottom rows, the structure of eigenvectors are separately plotted for excitatory and inhibitory neurons, respectively (with zeros replaced on the positions of the other population, respectively). Only the real part of the components of the eigenvectors are plotted here. Note that the tenth eigenvector already corresponds to an eigenvalue from the bulk of the spectrum in Fig. 5D. (B) Shown is the mean firing rate of neurons in the network, extracted from a 15 s simulation, in response to a stimulus with orientation $\theta = 0^\circ$. (C) For each neuron, the mean tuning curve (Mean TC) is plotted as the average (over different orientations) of the mean firing rate. (C, D) From each tuning curve, $r(\theta)$, the output preferred orientation (Output PO) and output orientation selectivity index (Output OSI) is extracted and plotted, respectively. They are obtained as the angle and length of the orientation selectivity vector, $OSV = \sum_{\theta} r(\theta) \exp(2\pi i \theta / 180^\circ) / \sum_{\theta} r(\theta)$; i.e. $OSI = |OSV|$ and $PO = \arg(OSV)$. Insets show the distributions in each case.

doi:10.1371/journal.pone.0114237.g008

dently to all neurons, irrespective of the position of the neuron in space). The modulation component would therefore only have a very small component in the direction of each any special eigenmode. It is however possible that for non-stationary inputs to the network, transient patterns with a bias for selected eigenmodes resonate more than others.

The question arises, if spatially structured eigenmodes (cf. Fig. 8A) have an impact on the observed pattern of spontaneous and evoked neuronal activity. Plotting the response of the network to a stimulus reflecting one particular orientation, as well as the mean activity of neurons over different orientations, do not reveal any visible structure (Fig. 8B, C). The baseline activity of the network seems to be quite uniform, and the response to a certain orientation does not reveal any structure beyond the random spatial pattern one would expect from the random assignment of preferred orientations of the input. This is further

supported by visual inspection of the map of preferred orientations for the output ([Fig. 8D](#)) and orientation selectivity index ([Fig. 8E](#)) in the network.

In principle, it is conceivable that spatially structured eigenmodes could affect the response of the network by setting the operating point of the network differently at different positions in space, as a result of the selective attenuation of certain eigenmodes. However, we have never observed such phenomena in our simulations. The fact that those structured modes get attenuated (and not amplified) might be one reason; another reason might be the fact that eigenmodes are typically heterogeneous and non-local, which makes the selection of the corresponding overall preferred pattern unlikely. Spatial structure of the network, and of its built-in linear eigenmodes, are therefore not dominant in determining the distribution of orientation selectivity. They could, however, be potential contributors in the small deviation of the predicted distribution from the measured one.

Spatial Imbalance of Excitation and Inhibition

To test the robustness of our predictions, we went beyond the case of spatial balance of excitation and inhibition, and also simulated networks with different extents of connectivity. Roughly the same overall behavior of the network, and accuracy of our predictions, were observed for the case of more localized inhibition and less localized excitation ($\sigma_{\text{inh}} = 0.45$ mm and $\sigma_{\text{exc}} = 0.75$ mm, [Fig. 9A](#)), as well as for the case of more localized excitation and less localized inhibition ($\sigma_{\text{inh}} = 0.75$ mm and $\sigma_{\text{exc}} = 0.45$ mm, [Fig. 9B](#)).

This trend was further corroborated when we systematically scanned the accuracy of our predictions for a large set of different networks, by scanning the parameter space ([Fig. 10A](#)). Indeed, for most of the parameters studied, the predicted distribution of orientation selectivity matched very well with the actual distribution (more than 90% overlap). For the more “extreme” combinations of parameters, however, where the spatial extent of excitation and inhibition were highly out of balance, the quality of the match degraded. The deviation was more significant when excitation was more local and inhibition was more global ([Fig. 10A](#), upper left portion). Note that, even for the most extreme cases of local excitation ($\sigma_{\text{exc}} = 0.25$ mm), the accuracy of our prediction is still fairly good, as long as the inhibition has a similar extent ($\sigma_{\text{inh}} = 0.25 - 0.45$ mm).

To investigate what happens in each extreme case, we chose two examples (marked in [Fig. 10A](#)) for further analysis. The connectivity patterns of these two examples, with $(\sigma_{\text{exc}}, \sigma_{\text{inh}}) = (0.75, 0.25)$ and $(0.25, 0.75)$ (numbers indicated in mm), are illustrated in [Fig. 10B, C](#), respectively. The eigenvalue spectra of the corresponding weight matrices are shown in [Fig. 10D, E](#). When the weights are normalized with respect to the reset voltage (upper panels), both spectra suggest an unstable linearized dynamics, as they both have eigenvalues with a real part larger than one.

The picture changes, however, when a normalization according to the effective gain, ζ_s , is performed. While the network with local excitation still has several

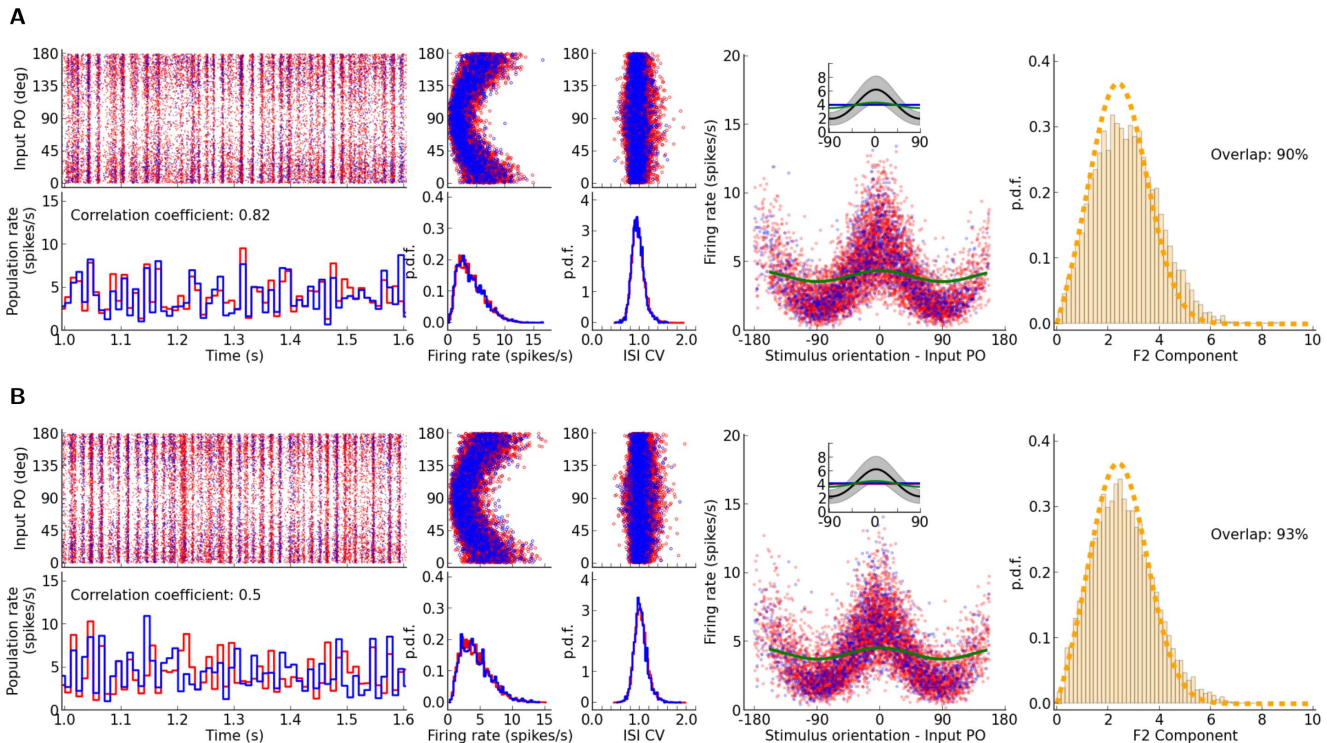


Figure 9. The impact of spatial extent of excitation and inhibition on the distribution of F2 components. Same illustration as in Fig. 4, for simulations with different extents of excitatory and inhibitory connectivity. (A) shows the results for a network with inhibition being more localized than excitation ($\sigma_{inh} = 0.45$ mm and $\sigma_{exc} = 0.75$ mm). In (B) we show the results for excitation being more localized than inhibition ($\sigma_{inh} = 0.75$ mm and $\sigma_{exc} = 0.45$ mm). Other parameters are the same as in Fig. 6. The distribution of F2 components is computed after re-normalization of the connectivity matrix by ζ_s , as explained before.

doi:10.1371/journal.pone.0114237.g009

clearly unstable eigenmodes (Fig. 10E, bottom), the spectrum of the network with local inhibition comprises only one positive eigenvalue which is only slightly larger than one (Fig. 10D, bottom). Some of the eigenvectors corresponding to the largest positive eigenvalues are plotted for both networks in Fig. 10F, G, respectively. From this, it seems therefore possible that the source of deviation from the linear prediction is indeed instability of the linearized dynamics (namely the instability of the uniform asynchronous-irregular state about which we perform the linearization) for these extreme parameter settings. When this instability is more pronounced, i.e. for the network with local excitation, the deviation is highest. When the network is at the edge of instability, i.e. for the network with local inhibition, our predictions show only a modest deviation.

To test this hypothesis further, namely that instability of the linearized dynamics is the source of mismatch between the linear prediction and the actual distribution of orientation selectivity, we need to scrutinize the response behavior of the sample networks. The outcome of this is shown in Fig. 11. While the network with local inhibition does not look very different from other examples considered before (Fig. 11A), the behavior of the network with local excitation very clearly shows deviating behavior (Fig. 11B). First, firing rates are much

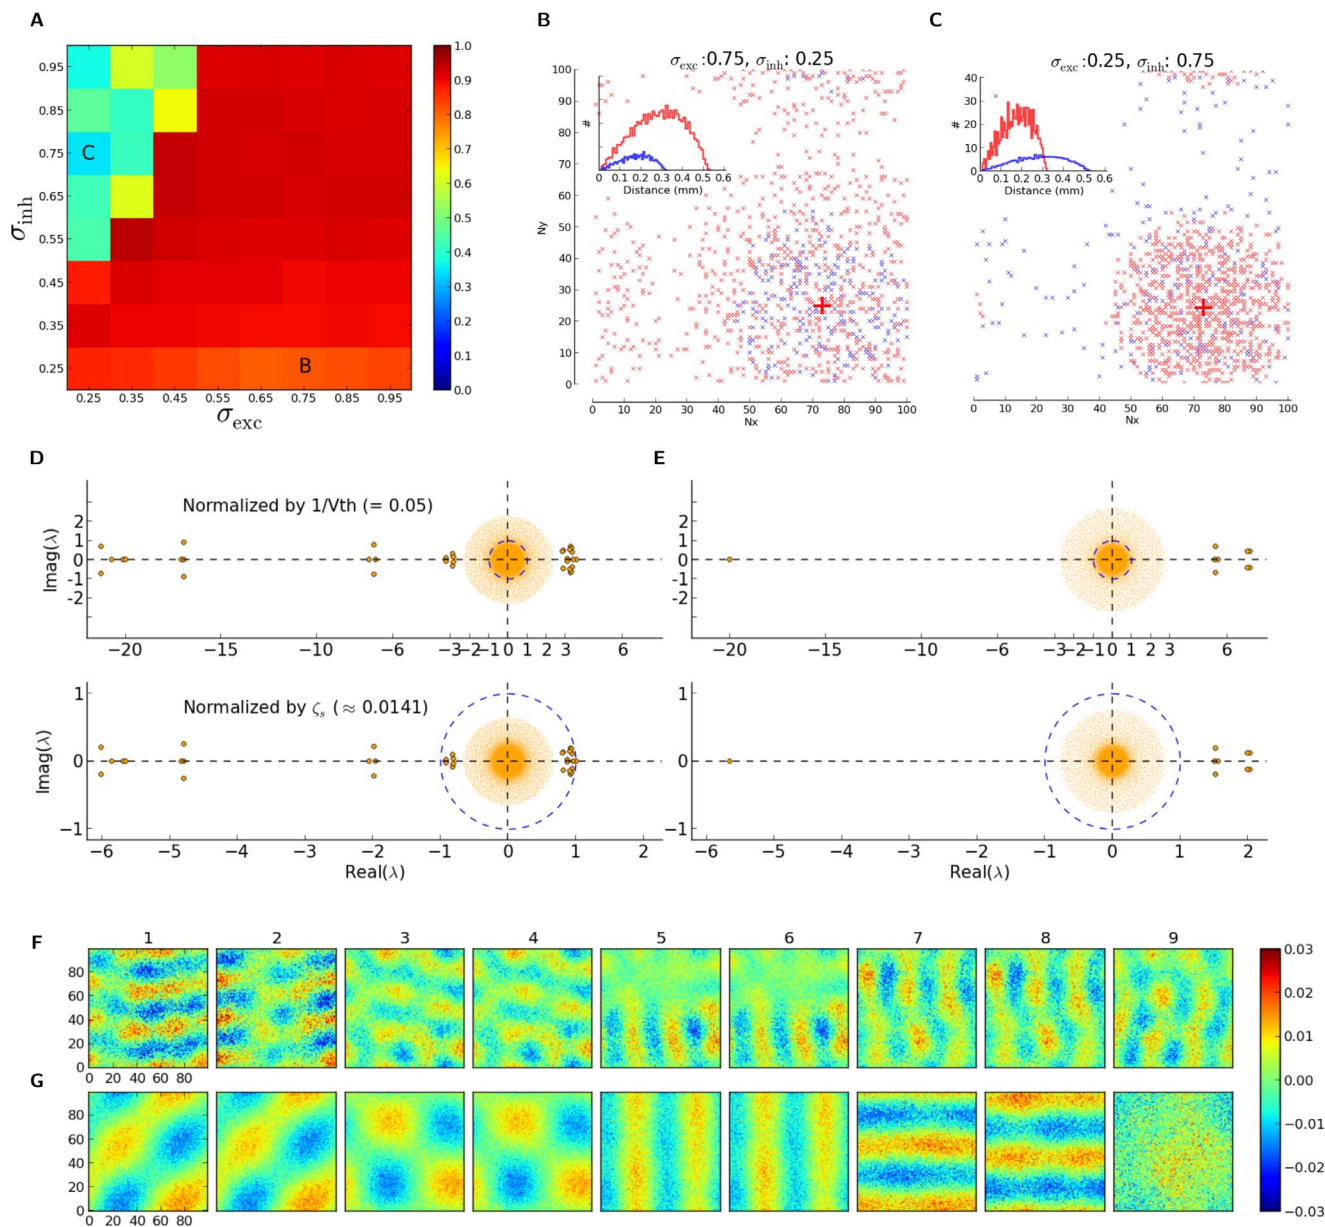


Figure 10. Accuracy of the linear prediction for different spatial extents of excitation and inhibition. (A) The overlap index (using ζ_c) is plotted for networks with different extents of excitation and inhibition. (B, C) Pre-synaptic connections for a sample post-synaptic neuron, along with the histogram of distances to pre-synaptic neurons for the entire population (inset), are shown here for two extreme cases, marked in panel (A). (D, E) Eigenvalue distribution of the example networks in (B) and (C), respectively. Two ways of normalization of the weight matrix are compared in the top and bottom panels. (F, G) First nine eigenmodes, corresponding to the nine largest positive eigenvalues (in terms of their real component), are plotted for the example networks in (A). Panels (F) and (G) correspond to the networks in (B) and (C), respectively. Note that the ninth eigenvector in (G) corresponds to an eigenvalue from the bulk of the spectrum in (E). Only the real part of the components of the eigenmodes are plotted.

doi:10.1371/journal.pone.0114237.g010

higher than in the less extreme cases, for both excitatory and inhibitory populations (Fig. 11B, first column). Moreover, the activity of excitatory and inhibitory neuronal populations are not well correlated in time, as it is the case for the other networks (Fig. 11B, first column, bottom). The firing rate distribution

has a very long tail, and the tail is longer for the excitatory than for the inhibitory population (Fig. 11B, second column). The long tail is accompanied by a peculiar peak at zero firing rate (which is cut for illustration purposes in Fig. 11B, second column, bottom). It reflects the fact that most of the neurons in the network are actually silent, and a small fraction of the population is highly active. The average irregularity of spike trains (the CV of the inter-spike intervals) in the network is reduced compared to our previous examples (Fig. 11B, third column). All these properties are consistent with the presumed instability of the linearized dynamics, as inferred from the eigenvalue spectrum.

In terms of functional properties of the network, the output tuning curves are much more scattered when aligned by the respective preferred orientations of the inputs (Fig. 11B, fourth column, upper panel). In fact, the mean output tuning curve for all neurons of the network does not show any amplification, if it is aligned at the Input PO (Fig. 11B, fourth column, lower panel). The picture changes, however, if tuning curves are aligned according to their Output PO (Fig. 11B, fifth column). Here a clear amplification of the modulation is evident in output tuning curves, although the relation to the feedforward input gets lost. Also, the average output tuning curve is not smooth, i.e. not all orientations are uniformly represented in the distribution of output preferred orientations.

This breaking of the symmetry becomes even more obvious when we look at the response of the two networks to stimuli of different orientations (Fig. 12A, B). While both networks show some degree of inhomogeneity in the spatial pattern of their firing rate responses, the response pattern of the second network is much more clustered (Fig. 12B). In fact, it seems that the internal connectivity structure of the network determines the position of a discrete set of potential activity bumps, and the orientation bias in the input can only choose between these bumps. As the nonlinear dynamics of the unstable network is crucially affecting the activity in response to stimuli, it is not surprising that the distribution of orientation selectivity is not matching the prediction which relies on a linearization about the uniform asynchronous-irregular state (compare Fig. 12C and D, first columns).

In fact, this internal structure is even reflected in the pattern of baseline firing rates (mean of the tuning curves over orientation). While for the network with local inhibition this pattern is covert and ineffective (Fig. 12C, second column), in the network with local excitation clear clusters of activity, resembling the ones in Fig. 12B, are evident (Fig. 12D, second column). One may, therefore, expect that there exists a corresponding pattern in the spatial organization of orientation selectivity. Larger domains of neighboring neurons, who get activated together, also exhibit the same selectivity. This is reflected in the clustering of output preferred orientations (Fig. 12C, third column) and orientation selectivity index (Fig. 12C, fourth column).

Note that a consequence of this clustering of PO is a degenerate representation of orientation selectivity, i.e. not all orientations are represented equally in the network. While the distribution of Output POs is almost uniform in the network with local inhibition (inset in Fig. 12C, third column), clear peaks are present in

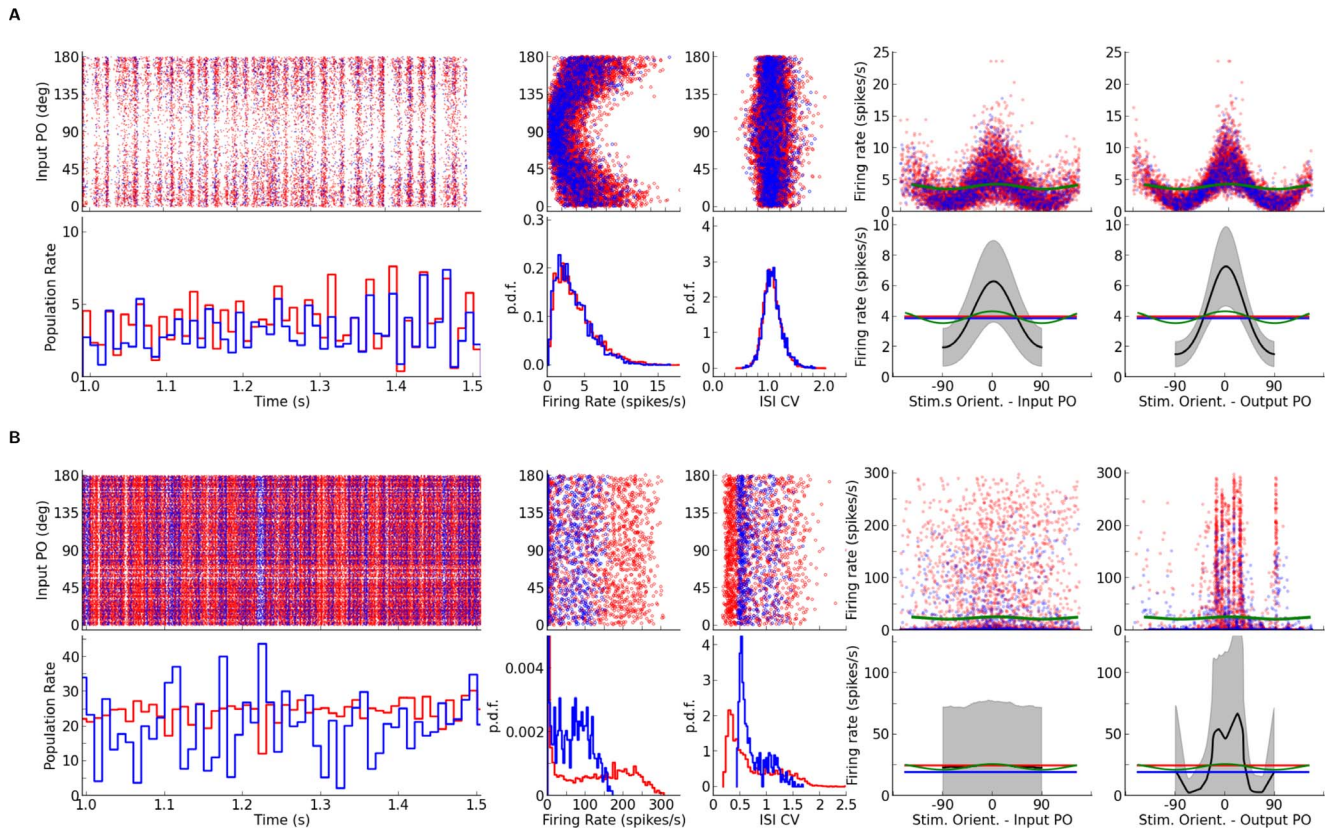


Figure 11. Orientation selectivity in networks with extreme spatial imbalance of excitation and inhibition. (A, B) As extreme examples, networks with highly local inhibition ($\sigma_{inh} = 0.25$ mm and $\sigma_{exc} = 0.75$ mm, Fig. 10B) or highly local excitation ($\sigma_{inh} = 0.75$ mm and $\sigma_{exc} = 0.25$ mm, Fig. 10C), were considered, respectively. The spiking activity of the network (first column), distribution of firing rates (second column) and spike train irregularity index (third column), as well as output tuning curves (fourth and fifth columns). In the fourth column, the tuning curves are aligned according to their Input PO, whereas in the fifth column they are aligned according to their Output PO. Other conventions are the same as Fig. 9.

doi:10.1371/journal.pone.0114237.g011

the distribution of Output POs in the network with local excitation (inset in Fig. 12D, third column). This is in line with our observation of broken symmetry described before, reflected in the pattern of mean output tuning curve in Fig. 11B.

Discussion

We presented a linear analysis, which was capable of predicting the distribution of orientation selectivity in networks with different patterns of random connectivity, including some degree of spatial organization, and for a wide range of parameters. The effective strength of excitation and inhibition in the network (Figs. 1 and Fig. 4), as well as the spatial extent of excitatory and inhibitory connectivity (Fig. 10), did not affect the prediction accuracy very strongly, as long as the linearized dynamics remained stable. We therefore conclude that linear mechanisms are the major network operations that explain amplification and

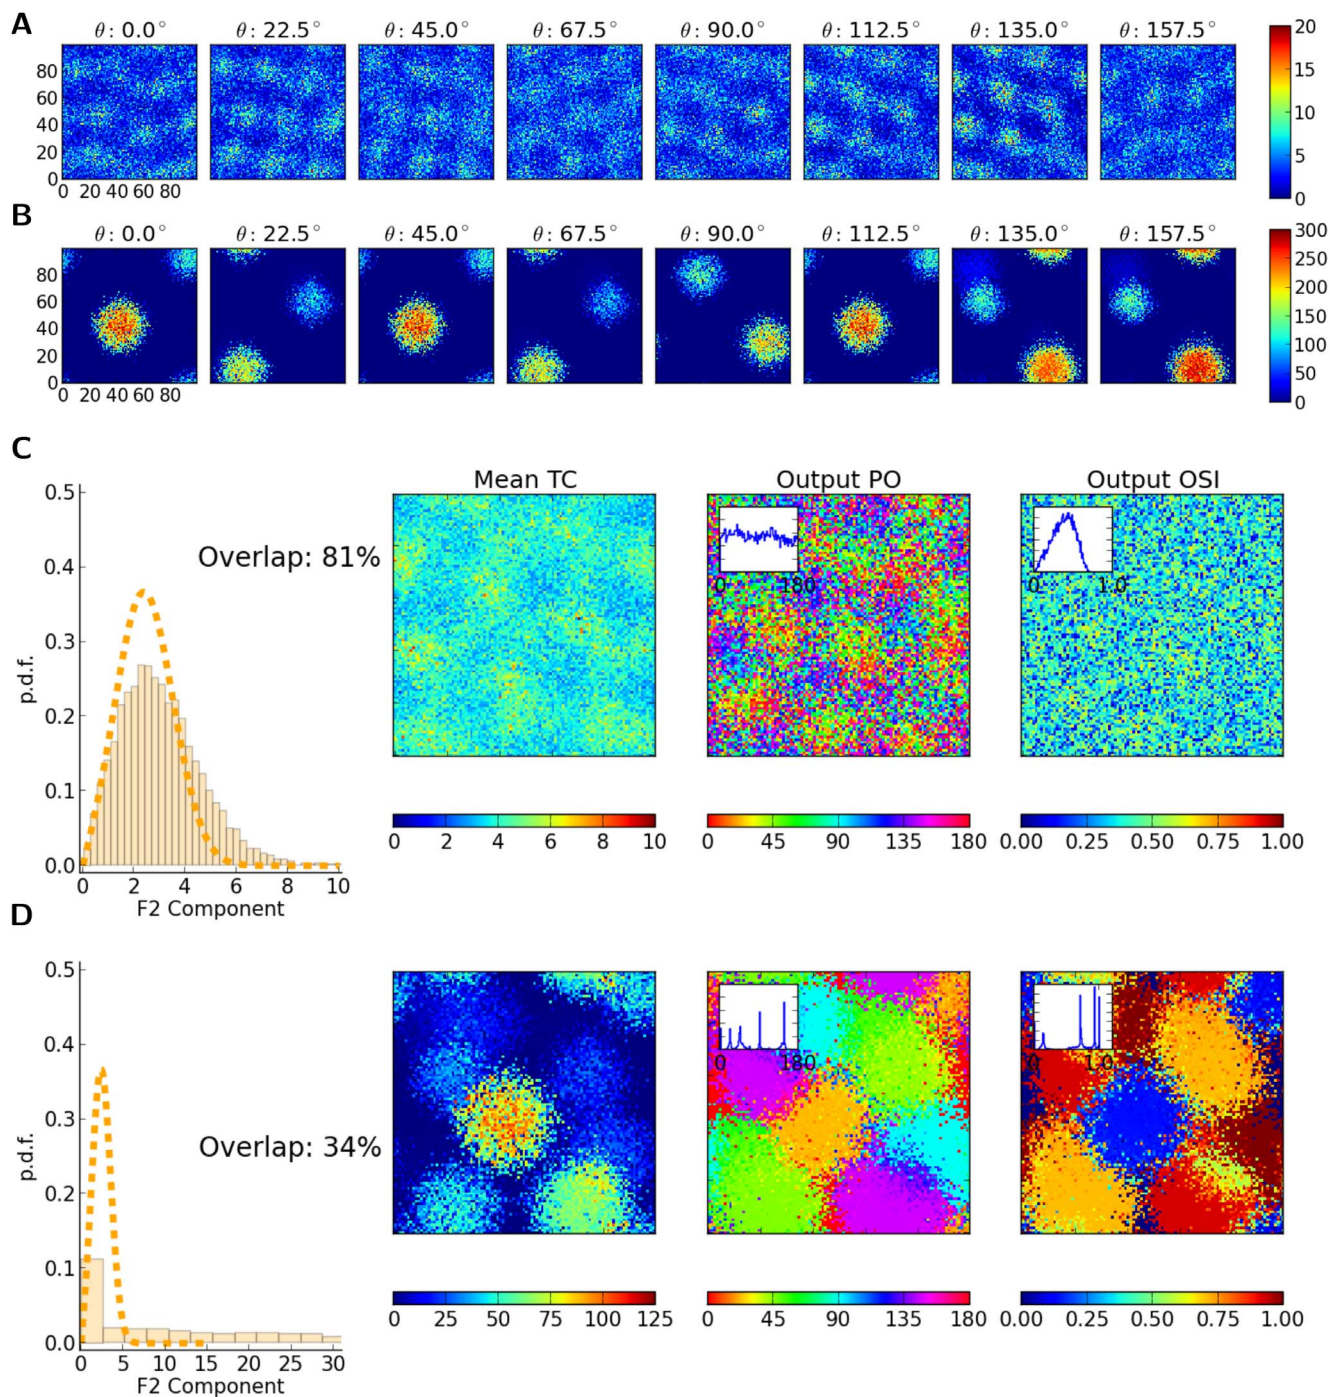


Figure 12. Dynamic instability leads to nonlinear distortions in the processing of orientation selectivity. (A, B) Mean firing rate of neurons in the network (the same networks as in Fig. 11, (A) and (B), respectively) in response to stimuli of different orientations. (C, D) For the networks in (A) and (B), respectively, the distributions of F2 components are compared with the linear prediction (using ζ_2 , dashed line) in the first column. The subsequent columns depict the map of average (over orientation) tuning curves (Mean TC), Output PO and Output OSI. Insets in the last two columns show the distribution of Output PO and Output OSI, respectively.

doi:10.1371/journal.pone.0114237.g012

attenuation, and the distribution of the resulting orientation selectivity in our networks, within their stable regimes of linearized dynamics.

Operating Regime of Orientation Selectivity

Note that even in networks with localized connectivity of excitation and/or inhibition, the linearized dynamics remained stable for a vast set of parameter combinations. Even when excitation was highly local and clustered, as long as inhibition had the same spatial connectivity profile, stability of the network was guaranteed. A similar conclusion has been recently obtained from an analysis of spatially embedded balanced networks [27]. It has also been shown before that networks with distance-dependent connectivity can show the same macroscopic behavior similar to random networks without local connectivity [28].

The asynchronous irregular (AI) state has been argued to best match the activity of cortical networks *in vivo* (see e.g. [19, 29, 30]). The relevance of this regime has only been discussed, however, for cortical networks in response to uniform stimulation. On the other hand, with regard to the processing of a non-uniformly modulated input, it has been claimed that a “marginal state of recurrent dynamics” might be the relevant regime of operation for the processing of weakly tuned inputs [2]. Also, it has recently been suggested that a recurrent regime with “macroscopic chaos” (probably corresponding to our regime of unstable dynamics) might be advantageous for sensory processing, as it may support a better separation of trajectories [31].

In contrast to these proposals, the results of our study suggest that a stable AI state of dynamics might indeed be the relevant regime of operation also for sensory processing in cortical networks in response to tuned inputs. Notably, the dense and local pattern of inhibition in real cortical circuits [8, 32, 33] is in line and consistent with our proposal. It might indeed be a general strategy biological networks of spiking neurons have exploited to ensure their overall stability to modulated inputs. We note again that we are talking about dynamic stability here, where the network dynamics is linearized about the uniform asynchronous-irregular state, and the effective weights of coupling linearized about this baseline state are considered.

Distribution of Orientation Selectivity

A broad distribution of orientation selectivity is reported across all cortical layers in the primary visual cortex of macaque monkeys [34], as well as in mice [35] (for a comparison of the distributions, see panel C in Fig. S2 therein). Although we chose random connectivities by fixing the in-degree of all neurons (which we refer to as “structural homogeneity”), a broad distribution of orientation selectivities also emerged in all our networks. The main contributor to this broad distribution was, therefore, not the structural heterogeneity of synaptic connectivity. In fact, there is no heterogeneity at all, if one only considers the number of connections each neuron receives from pre-synaptic excitatory and inhibitory neurons Nor

were the temporal fluctuations of activity generated by our networks a major source of this variability, although the networks were mostly operated in the fluctuation-driven regime with high amounts of temporal and trial-by-trial variability. As we have generally chosen a homogeneous connectivity pattern, this temporal variance would be essentially the same for all neurons, at least in the baseline state. (This also justifies the mean-field ansatz we have employed for our analysis.) This is again reflected in the narrow distribution of F0 components in all our networks.

The main source of variability in orientation selectivity is rather the “functional heterogeneity” in synaptic connectivity, namely heterogeneous preferred features (here, preferred orientations of inputs) of the pre-synaptic sources within the recurrent network. Receiving input from neurons with different preferred features may be a computational strategy to integrate the information, and help to remove distractive correlations in the activity. The fact that each neuron within the recurrent network receives input from a heterogeneous pool of neurons with a wide range of preferred orientations leads to a random “summation” of pre-synaptic preferred orientations, which eventually changes the output preferred orientation of the post-synaptic neuron [16].

The quenched noise of preferred orientations, and not structural or dynamic fluctuations, is, therefore, the main mechanism responsible for the distribution of orientation selectivity in our networks. We showed that even with this most conservative estimate of neuronal heterogeneity, consistent with recent experiments [7], a broad distribution of neuronal selectivities can be obtained. However, we cannot rule out a possible contribution of other sources of heterogeneity, like heterogeneous connectivity and heterogeneous amounts of excitation and inhibition different neurons may receive in their baseline state (leading to different levels of spontaneous activity, see e.g. [34]), as well as variability in neuron parameters [36] and synaptic noise. Also, heterogeneity in the pattern of feedforward projections to neurons in V1 can be a prominent source of distribution in orientation selectivity. However, if the distribution of orientation selectivity is mainly dominated by feedforward heterogeneity, or/and if single neuron heterogeneities like variability in threshold and synaptic noise are the main source of this distribution, the distribution should not much change when the recurrent network is absent. On the other hand, if functional heterogeneity resulting from recurrent interactions is a major contributor to this distribution, it should get narrowed when the intra-cortical circuitry is deactivated. It therefore awaits further experimental tests which mechanisms are dominant in creating the distribution of feature-selectivity in the cortex.

Future Directions

There are several ways in which the the current study could be expanded. First, sticking to a linear framework of analysis enabled us to analytically compute the distribution of orientation selectivity. In this simplified framework, however, we neglected several nonlinearities, both at the level of neuronal properties and

network interactions. These nonlinearities are deemed to be more prominent in biological networks, for instance in the form of rectification [37, 38], or an expansive-compressive transfer nonlinearity [26, 39, 40]. Such mechanisms might play a major role in sharpening and amplification of orientation selectivity. A more complete theoretical treatment of the problem should therefore consider the contribution of nonlinear mechanisms as well, although this may come at the expense of less rigorous analytical predictions.

One way to embrace additional nonlinear mechanisms that are effective in biological networks, at least at the level of simulations, is to use a more realistic and more detailed neuron model. In our simulations here we used the current-based LIF neuron model. Simulating networks of more realistic neuron models, like conductance-based LIF neurons, may change certain behaviors of the network [41, 42]. For instance, increasing the recurrent coupling in our inhibition-dominated networks can decrease the mean membrane potential of neurons in the network to very negative values, as there is no reversal potential limiting it. This is not the case in a conductance-based neuron model, and therefore a network of that sort might show a different behavior, especially when operated in extreme regimes.

Finally, it would be interesting to see how the predictions of our current theory change when one considers networks with feature-specific connectivity. This scenario might be corresponding to species with orientation maps, where neighboring neurons tend to have a similar preferred orientation [4–6], or to species without spatial map of selectivity, but with feature-specific functional connectivity [8–12]. A linear amplification of feedforward input, for instance, has been recently reported in cortical circuits of mice [43–45]. How this effect could be modeled within our theoretical framework, and how it affects the distribution of orientation selectivity, should therefore be a next step in our research.

Acknowledgments

We thank the developers of the simulation software NEST (see <http://www.nest-initiative.org>) and the maintainers of the BCF computing facilities for their support throughout this study.

Author Contributions

Conceived and designed the experiments: SS SR. Performed the experiments: SS. Analyzed the data: SS SR. Contributed reagents/materials/analysis tools: SS SR. Wrote the paper: SS SR.

References

1. Peters A, Payne BR (1993) Numerical relationships between geniculocortical afferents and pyramidal cell modules in cat primary visual cortex. *Cerebral Cortex* 3: 69–78.

2. **Ben-Yishai R, Bar-Or RL, Sompolinsky H** (1995) Theory of Orientation Tuning in Visual Cortex. *Proceedings of the National Academy of Sciences* 92: 3844–3848.
3. **Somers DC, Nelson SB, Sur M** (1995) An emergent model of orientation selectivity in cat visual cortical simple cells. *The Journal of Neuroscience* 15: 5448–5465.
4. **Blasdel GG, Salama G** (1986) Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature* 321: 579–85.
5. **Bonhoeffer T, Grinvald A** (1991) Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature* 353: 429–31.
6. **Ohki K, Chung S, Kara P, Hübener M, Bonhoeffer T, et al.** (2006) Highly ordered arrangement of single neurons in orientation pinwheels. *Nature* 442: 925–8.
7. **Jia H, Rochefort NL, Chen X, Konnerth A** (2010) Dendritic organization of sensory input to cortical neurons in vivo. *Nature* 464: 1307–12.
8. **Hofer SB, Ko H, Pichler B, Vogelstein J, Ros H, et al.** (2011) Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nature Neuroscience* 14: 1045–52.
9. **Ko H, Hofer SB, Pichler B, Buchanan KA, Sjöström PJ, et al.** (2011) Functional specificity of local synaptic connections in neocortical networks. *Nature* 473: 87–91.
10. **Ko H, Cossell L, Baragli C, Antolik J, Clopath C, et al.** (2013) The emergence of functional microcircuits in visual cortex. *Nature* 496: 96–100.
11. **Ko H, Mrcsic-Flogel TD, Hofer SB** (2014) Emergence of Feature-Specific Connectivity in Cortical Microcircuits in the Absence of Visual Experience. *Journal of Neuroscience* 34: 9812–9816.
12. **Ishikawa AW, Komatsu Y, Yoshimura Y** (2014) Experience-Dependent Emergence of Fine-Scale Networks in Visual Cortex. *Journal of Neuroscience* 34: 12576–12586.
13. **Harris KD, Mrcsic-Flogel TD** (2013) Cortical connectivity and sensory coding. *Nature* 503: 51–8.
14. **Hansel D, van Vreeswijk C** (2012) The Mechanism of Orientation Selectivity in Primary Visual Cortex without a Functional Map. *The Journal of Neuroscience* 32: 4049–4064.
15. **Pehlevan C, Sompolinsky H** (2014) Selectivity and Sparseness in Randomly Connected Balanced Networks. *PLoS ONE* 9: e89992.
16. **Sadeh S, Cardanobile S, Rotter S** (2014) Mean-field analysis of orientation selectivity in inhibition-dominated networks of spiking neurons. *Springer Plus* 3: 148.
17. **Siebert AJF** (1951) On the first passage time probability problem. *Phys Rev* 81: 617–623.
18. **Ricciardi LM** (1977) *Diffusion Processes and Related Topics on Biology*. Berlin: Springer-Verlag.
19. **Amit DJ, Brunel N** (1997) Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex* 7: 237–52.
20. **Brunel N** (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience* 8: 183–208.
21. **Helias M, Deger M, Rotter S, Diesmann M** (2010) Instantaneous non-linear processing by pulse-coupled threshold units. *PLoS computational biology* 6.
22. **Renart A, de la Rocha J, Bartho P, Hollender L, Parga N, et al.** (2010) The asynchronous state in cortical circuits. *Science* 327: 587–90.
23. **Pernice V, Staude B, Cardanobile S, Rotter S** (2011) How structure determines correlations in neuronal networks. *PLoS Computational Biology* 7: e1002059.
24. **Tetzlaff T, Helias M, Einevoll GT, Diesmann M** (2012) Decorrelation of Neural-Network Activity by Inhibitory Feedback. *PLoS Computational Biology* 8: e1002596.
25. **Hansel D, van Vreeswijk C** (2002) How noise contributes to contrast invariance of orientation tuning in cat visual cortex. *The Journal of Neuroscience* 22: 5118–28.
26. **Miller KD, Troyer TW** (2002) Neural noise can explain expansive, power-law nonlinearities in neural response functions. *Journal of Neurophysiology* 87: 653–9.
27. **Rosenbaum R, Doiron B** (2014) Balanced Networks of Spiking Neurons with Spatially Dependent Recurrent Connections. *Physical Review X* 4: 021039.

28. **Yger P, El Boustani S, Destexhe A, Frégnac Y** (2011) Topologically invariant macroscopic statistics in balanced networks of conductance-based integrate-and-fire neurons. *Journal of computational neuroscience* 31: 229–45.
29. **van Vreeswijk C, Sompolinsky H** (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274: 1724–6.
30. **van Vreeswijk C, Sompolinsky H** (1998) Chaotic balanced state in a model of cortical circuits. *Neural computation* 10: 1321–71.
31. **Ostoic S** (2014) Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons. *Nature neuroscience* 17: 594–600.
32. **Fino E, Yuste R** (2011) Dense inhibitory connectivity in neocortex. *Neuron* 69: 1188–203.
33. **Packer AM, Yuste R** (2011) Dense, unspecific connectivity of neocortical parvalbumin-positive interneurons: a canonical microcircuit for inhibition? *The Journal of Neuroscience* 31: 13260–71.
34. **Ringach DL, Shapley RM, Hawken MJ** (2002) Orientation Selectivity in Macaque V1: Diversity and Laminar Dependence. *The Journal of Neuroscience* 22: 5639–5651.
35. **Niell CM, Stryker MP** (2008) Highly selective receptive fields in mouse visual cortex. *The Journal of Neuroscience* 28: 7520–36.
36. **Yim M, Aertsen A, Rotter S** (2013) Impact of intrinsic biophysical diversity on the activity of spiking neurons. *Physical Review E* 87: 032710.
37. **Carandini M, Ferster D** (2000) Membrane potential and firing rate in cat primary visual cortex. *The Journal of Neuroscience* 20: 470–84.
38. **Priebe NJ, Ferster D** (2008) Inhibition, spike threshold, and stimulus selectivity in primary visual cortex. *Neuron* 57: 482–97.
39. **Anderson JS, Lampl I, Gillespie DC, Ferster D** (2000) The Contribution of Noise to Contrast Invariance of Orientation Tuning in Cat Visual Cortex. *Science* 290: 1968–1972.
40. **Priebe NJ, Ferster D** (2012) Mechanisms of Neuronal Computation in Mammalian Visual Cortex. *Neuron* 75: 194–208.
41. **Kuhn A, Aertsen A, Rotter S** (2004) Neuronal integration of synaptic input in the fluctuation-driven regime. *The Journal of Neuroscience* 24: 2345–56.
42. **Kumar A, Rotter S, Aertsen A** (2008) Conditions for propagating synchronous spiking and asynchronous firing rates in a cortical network model. *The Journal of Neuroscience* 28: 5268–80.
43. **Lien AD, Scanziani M** (2013) Tuned thalamic excitation is amplified by visual cortical circuits. *Nature neuroscience* 16: 1315–23.
44. **Li LY, Li YT, Zhou M, Tao HW, Zhang LI** (2013) Intracortical multiplication of thalamocortical signals in mouse auditory cortex. *Nature neuroscience* 16: 1179–81.
45. **Li YT, Ibrahim LA, Liu BH, Zhang LI, Tao HW** (2013) Linear transformation of thalamocortical input by intracortical excitation. *Nature neuroscience* 16: 1324–30.