



The Influence of Mexican Hat Recurrent Connectivity on Noise Correlations and Stimulus Encoding

Robert Meyer^{1,2*}, Josef Ladenbauer^{1,2,3} and Klaus Obermayer^{1,2}

¹ Department of Software Engineering and Theoretical Computer Science, Technische Universität Berlin, Berlin, Germany,

² Bernstein Center for Computational Neuroscience, Berlin, Germany, ³ Group for Neural Theory, Laboratoire de Neurosciences Cognitives, École Normale Supérieure, Paris, France

Noise correlations are a common feature of neural responses and have been observed in many cortical areas across different species. These correlations can influence information processing by enhancing or diminishing the quality of the neural code, but the origin of these correlations is still a matter of controversy. In this computational study we explore the hypothesis that noise correlations are the result of local recurrent excitatory and inhibitory connections. We simulated two-dimensional networks of adaptive spiking neurons with local connection patterns following Gaussian kernels. Noise correlations decay with distance between neurons but are only observed if the range of excitatory connections is smaller than the range of inhibitory connections (“Mexican hat” connectivity) and if the connection strengths are sufficiently strong. These correlations arise from a moving blob-like structure of evoked activity, which is absent if inhibitory interactions have a smaller range (“inverse Mexican hat” connectivity). Spatially structured external inputs fixate these blobs to certain locations and thus effectively reduce noise correlations. We further investigated the influence of these network configurations on stimulus encoding. On the one hand, the observed correlations diminish information about a stimulus encoded by a network. On the other hand, correlated activity allows for more precise encoding of stimulus information if the decoder has only access to a limited amount of neurons.

Keywords: noise correlations, neural coding, Mexican hat connectivity, spiking neural networks, pattern formation, primary visual cortex (V1), simulation models, stimulus encoding

OPEN ACCESS

Edited by:

Markus Diesmann,
Forschungszentrum Jülich, Germany

Reviewed by:

Thomas Wennekers,
Plymouth University, UK
Brent Doiron,
University of Pittsburgh, USA

*Correspondence:

Robert Meyer
robert.meyer@ni.tu-berlin.de

Received: 16 January 2017

Accepted: 20 April 2017

Published: 10 May 2017

Citation:

Meyer R, Ladenbauer J and Obermayer K (2017) The Influence of Mexican Hat Recurrent Connectivity on Noise Correlations and Stimulus Encoding.
Front. Comput. Neurosci. 11:34.
doi: 10.3389/fncom.2017.00034

1. INTRODUCTION

One of the fundamental problems in neuroscience is deciphering the neural code and understanding how the brain encodes sensory stimuli. During the past decades the analysis of neural coding has shifted from single cells to investigating population codes due to the development of multi-electrode recordings as well as improved theoretical models. An important aspect regarding population coding is whether neural responses are correlated, especially when driven by the same stimulus. In this case one speaks of *noise correlations* or *shared variability* (Cohen and Kohn, 2011; Hansen et al., 2012). Noise correlations have, for example, been widely observed in the visual cortex (Kohn and Smith, 2005; Martin and Schröder, 2013), where the magnitude of pairwise noise correlations decays with the distance between cell pairs (Smith and Kohn, 2008; Solomon et al., 2014). The underlying mechanisms and the role of noise correlations in information processing are, however, not well understood.

Noise correlations have often been explained by the large amount of shared inputs between neuron pairs (Gawne et al., 1996; Shadlen and Newsome, 1998) or, more recently, by cellular non-linearities (Doiron et al., 2016). Hansen et al. (2012) on the other hand hypothesized recurrent connectivity as one of the potential sources. This is a reasonable assumption considering that the vast majority of connections does not come from the sensory periphery or higher cortical areas, but originates within the visual cortex and projects locally (Markov et al., 2011).

Although work by Renart et al. (2010) showed that a balance between excitation and inhibition among recurrent connectivity can lead to the opposite effect and cause decorrelation of activity, more recent studies indicate that this phenomenon depends on the spatial structure of recurrent connections (Rosenbaum and Doiron, 2014). Rosenbaum and Doiron (2014) demonstrated analytically using a continuous rate model that homogeneous, balanced activity cannot be maintained in networks with Gaussian Mexican hat connectivity profiles, where the spatial reach of recurrent excitatory connections is on average shorter than that of recurrent inhibitory connections. Moreover, Hansel and Sompolinsky (1998) have shown that such connections can lead to the emergence of spatial patterns of activity in form of bumps (also see Roxin et al., 2006, for a more recent summary of the findings). This type of activity is termed *marginal* phase. They further demonstrated that under certain conditions the bumps move across the spatial extent of the network.

In this study we investigate the hypothesis that recurrent connections in the form of a Mexican hat configuration can lead to noise correlations by destabilizing homogeneous, balanced activity, and enabling temporally dynamic, spatially inhomogeneous patterns of activity. Motivated by aforementioned theoretical results we first study the dynamics of two-dimensional networks of spiking model neurons for different lateral interaction kernels and spatially unstructured homogeneous input. We find that a Mexican hat configuration leads to the emergence of dynamic stripe or bump patterns that in turn yield noise correlations. The dynamic spatial heterogeneities produce joint modulations of firing rates on the individual neuron level. Accordingly, we observed oscillatory spatial modulations of positive and negative noise correlations.

We then study the interaction of the neural dynamics with structured input, which is derived from an orientation map model of primary visual cortex. We find interactions among the input dynamics and the marginal phase. In particular, the oscillatory pattern of noise correlations with distance vanishes and we observe a linear decay instead. In addition, we find a suppression of noise correlations with input strengths. Finally we ask, what effects the observed noise correlations have on the quality of the neural code given structured inputs. We investigate whether Mexican hat connectivity with noise correlations would be a better or worse choice than inverse Mexican hat connectivity with decorrelated firing patterns. We find that correlations are detrimental to encoding quality in terms of Fisher information. However, in case readout neurons are sub-sampled, Mexican hat networks show a better encoding performance despite the noise correlations.

2. MATERIALS AND METHODS

2.1. Model Description

Our model is based on a primary visual cortex network developed by Mariño et al. (2005) and Stimberg et al. (2009). Here we provide a brief overview only. For details including the equations describing the underlying dynamics please refer to the Supplementary Material.

The model consisted of $N = N_E + N_I$ excitatory and inhibitory neurons of the adaptive exponential integrate and fire (AEIF) type (Brette and Gerstner, 2005). We assumed a ratio between inhibitory to excitatory neurons of 1–4 (Beaulieu et al., 1992). The AEIF model is frequently applied in computational neuroscience due to its efficiency and good fits to experimental data (Brette and Gerstner, 2005; Jolivet et al., 2008). It can exhibit a very rich set of dynamics including spike frequency adaptation (SFA) (Naud et al., 2008; Touboul and Brette, 2008) which is an essential ingredient for non-stationarity of bump attractors (cf. Hansel and Sompolinsky, 1998; Roxin et al., 2006; and Supplementary Material). We chose moderate sub-threshold and spike-triggered adaptation for excitatory neurons ($a_E = 2.0$ nS, $b_E = 50$ pA). To account for the fact that adaptation is much weaker in inhibitory neurons (Nowak et al., 2003), for inhibitory cells adaptation parameters were set to one tenth of the excitatory values ($a_I = 0.2$ nS, $b_I = 5$ pA).

2.1.1. Topology and Coupling

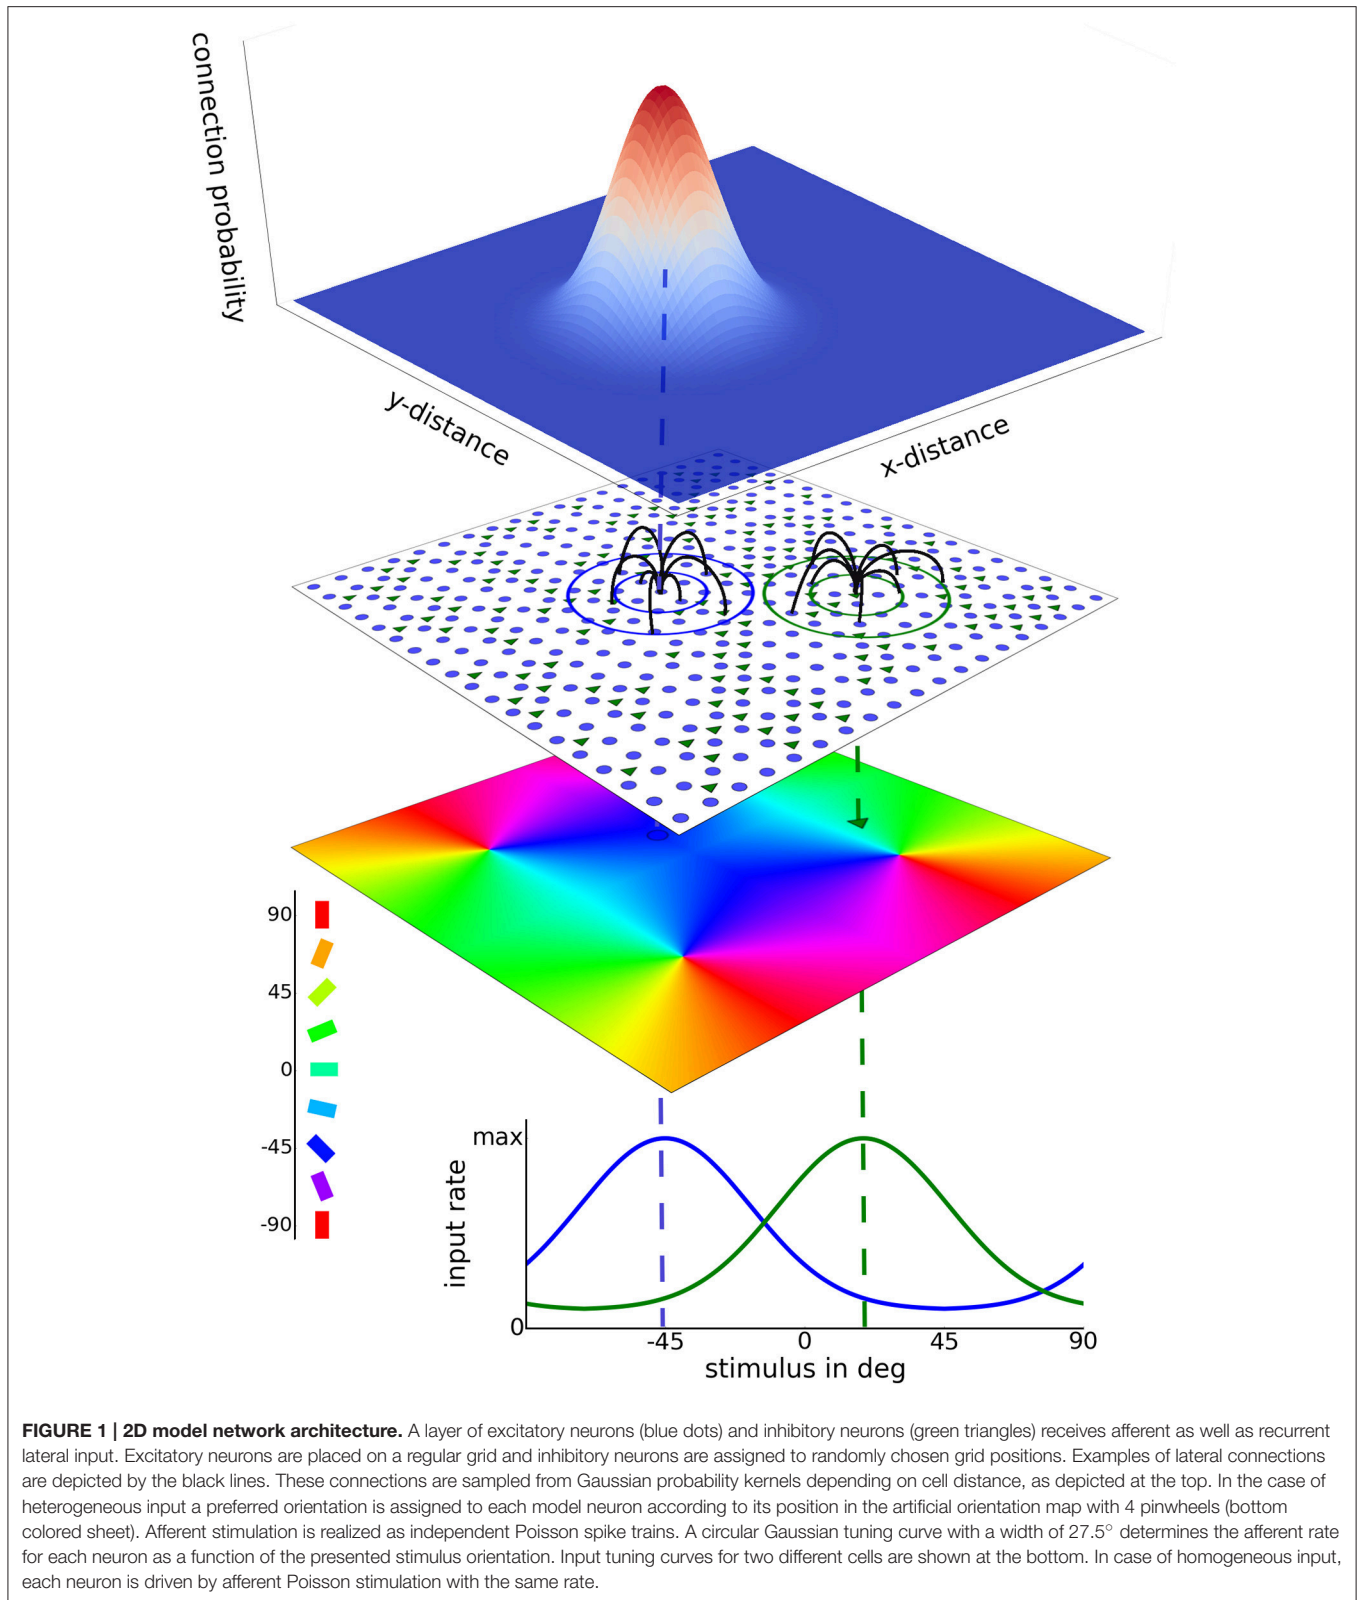
We investigated two-dimensional network models where neurons were placed on a grid. Additional results from one-dimensional networks are presented in the Supplementary Material. Excitatory neurons were regularly and evenly spaced on the grid, while the positions of the inhibitory neurons were chosen randomly from a uniform distribution defined over the network space. Distances in our network are given in pixels (px), with 1 px corresponding to 15 μ m. The conversion factor is calculated from the average pinwheel distance in cat primary visual cortex (Stimberg et al., 2009). Thus, a 100 px \times 100 px network with 4 orientation pinwheels corresponds to a 1.5 mm \times 1.5 mm piece of primary visual cortex. Such a two-dimensional network topology is depicted in **Figure 1**.

Each neuron received $K_E = 400$ excitatory and $K_I = 200$ inhibitory recurrent connections. These synaptic connections were drawn from Gaussian probability distributions defined over the Euclidean distances d_{ij} between two model neurons i and j (periodic boundary conditions):

$$p(i, j) = \begin{cases} 0 & \text{for } i = j \text{ (no self-connections), (1)} \\ 1/(\sqrt{2\pi}\sigma_Y)^D \exp(-d_{ij}^2/2\sigma_Y^2) & \text{otherwise} \end{cases}$$

where $Y \in \{E, I\}$ denotes the population type of pre-synaptic neuron j , D is the dimensionality of the model space (here $D = 2$, for $D = 1$ see Supplementary Material), and σ_Y is the spread of the connections of the corresponding type.

We use the terms *configuration* or *profile* to refer to the connection topology of a network. Important parameters that we investigated are σ_E and σ_I , the connectivity spread of



excitatory and inhibitory connections, respectively. We varied both parameters between 5 and 25 pixels, corresponding to 75 and $375 \mu\text{m}$ in visual cortex scale. These values span a biologically

plausible range. For instance, Hellwig (2000) measured a Gaussian decay of connection probability with spreads between 150 and $350 \mu\text{m}$ for pyramidal neurons in the rat primary

visual cortex. Similarly, Mariño et al. (2005) reported a Gaussian connection spread of about 125 μm for local connections in cat V1.

Throughout this study network topologies with the property of Gaussian kernel widths $\sigma_E < \sigma_I$ are termed **Mexican hat** networks. Likewise, topologies with $\sigma_E > \sigma_I$ are termed **inverse Mexican hat**. The setting $\sigma_E = \sigma_I$ is called a **balanced hat**. This naming scheme for topologies featuring two Gaussian kernels can be found throughout the literature, for example in work by Kang et al. (2003); Blumenfeld et al. (2006); Bressloff (2012).

We considered one type of inhibitory and two types of excitatory synapses. Inhibitory synaptic interactions were modeled using instantaneously rising, exponentially decaying functions with a time constant that describes GABA_A receptor kinetics. Excitatory synaptic interactions were modeled using the same functions with a time constant that describes AMPA receptor kinetics, and, using a bi-exponential function with larger time constants to describe NMDA receptors. We assumed a fixed ratio between both receptor types with a fraction of 70% AMPA receptors (Myme et al., 2003).

When we relate to the connection strengths, we use the term *operating regime*. Two important parameters in this regard are the maximum conductances \bar{g}_{EE} and \bar{g}_{IE} , which quantify the connection strengths between excitatory to excitatory and excitatory to inhibitory synapses. The ratio between maximum AMPA and NMDA was fixed ($\bar{g}_{\text{AMPA},EE} = 0.7\bar{g}_{EE}$ and $\bar{g}_{\text{NMDA},EE} = 0.3\bar{g}_{EE}$, cf. Stimberg et al., 2009). The maximum conductance values were varied in rather small amplitude ranges between 0 and 1.2 nS such that the excitatory post-synaptic potential (EPSP) at resting membrane voltage (-65 mV) had an overall small deflection with a maximum value of up to 1.5 mV. The relation between connection weight and voltage deflection amplitude is visualized in the Supplementary Material. These values are within a biologically plausible range. For instance, Mason et al. (1991) observed single spike-triggered EPSPs ranging from 0.05 up to 2 mV with a mean size of 0.55 mV in rat primary visual cortex. Inhibitory connection strengths were kept fixed ($\bar{g}_{II} = \bar{g}_{EI} = 5.0\text{ nS}$).

Synaptic delays depended linearly on the distance between cells in the model network. We assumed a conduction velocity of 0.2 m/s (about 13,000 px/s) in our networks. Usually values between 0.1 and 0.5 m/s are measured (Bringuier, 1999; González-Burgos et al., 2000).

2.1.2. Network Input

Each neuron received afferent inputs modeled as $K_A = 100$ independent Poisson spike trains targeting AMPA receptors. Hence, inputs to each neuron were decorrelated and all observed correlations could only arise due to recurrent interactions.

We distinguished between a homogeneous, also termed *blank stimulus*, and a heterogeneous tuned input, also termed *orientation stimulus*. In the homogeneous mode, input was not tuned but every neuron received independent Poisson spike trains with the same rate. In the heterogeneous mode the afferent input was tuned according to the preferred orientation of the post-synaptic neuron. The orientation preference of each neuron is determined by its position in

the pinwheel map depicted in **Figure 1**. We assumed moderate orientation tuning. The input firing rate to each cell i was computed by:

$$\nu_{\text{Aff},i}(s) = (\nu_{\text{Aff},\text{max}} - \nu_{\text{Aff},\text{base}}) \exp\left(-\frac{(s - s_{\text{PO}_i})^2}{2\sigma_{\text{Aff}}^2}\right) + \nu_{\text{Aff},\text{base}}, \quad (2)$$

where $\nu_{\text{Aff},\text{max}}$ is the maximum or peak firing rate, $\nu_{\text{Aff},\text{base}}$ is the baseline firing, σ_{Aff} is the tuning width, s is the orientation of the stimulus ($s \in [-90^\circ, 90^\circ]$), and s_{PO_i} denotes the i th neuron's preferred orientation according to the orientation map. Inputs only varied in spatial dimensions and were constant in time.

Network activity was evaluated after the first second of afferent stimulation to avoid artifacts caused by network initializations and onset transients.

2.1.3. Implementation

The network was implemented in Python 2.7 and was partly compiled into C-code for efficiency using the Numba library (The Numba Development Team, 2015). The source code is available online¹. Some numerical experiments were also conducted with smaller networks of Hodgkin-Huxley type neurons (cf. Stimberg et al., 2009) using the BRIAN simulator package (Goodman and Brette, 2008) (data not shown, example run see Supplementary Video). All data and parameter explorations were managed using the simulation toolkit *pypet* (Meyer et al., 2016).

2.2. Data Post-processing

2.2.1. Noise Correlations

Noise correlations were quantified using the linear correlation coefficient among spike counts. Spike counts r were either assessed over repeated fixed length intervals of stimulus presentation or by convolving an observed spike train of a single stimulus presentation with a temporal sliding window kernel K_T (cf. Renart et al., 2010):

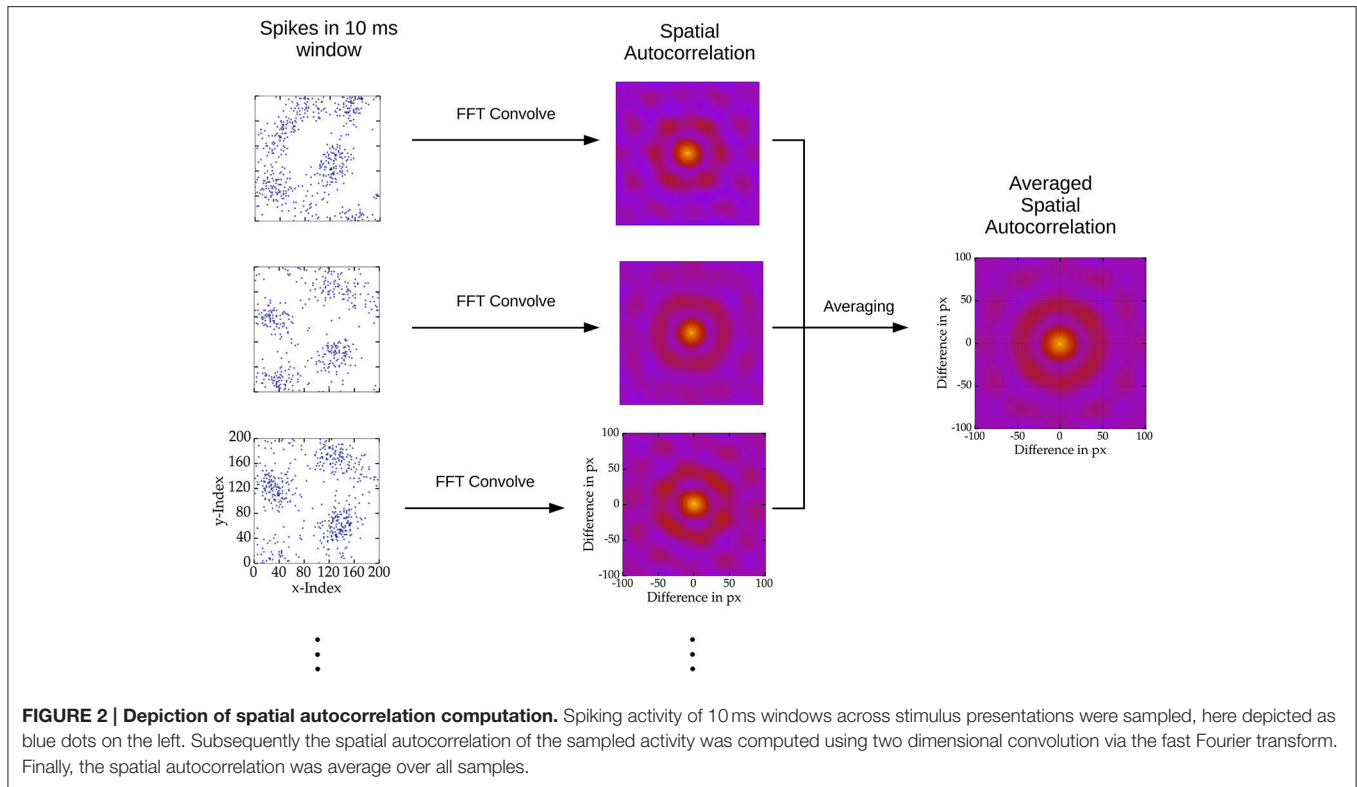
$$r_i(t, T) = \int_{t'} K_T(t' - t) S_i(t') \quad (3)$$

where K_T is a sliding window of length T , and S_i the spike train of neuron i with $S_i(t) = \sum_k \delta(t_i^k - t)$. We used the standard expression for the correlation coefficient r_{SC} of the spike counts:

$$r_{SC} = \frac{\text{Cov}(r_i, r_j)}{\sqrt{\text{Var}(r_i) \text{Var}(r_j)}} \quad (4)$$

where $\text{Cov}(r_i, r_j)$ is the covariance of spike counts between cells i and j and $\text{Var}(r_i)$ the individual spike count variance of one cell.

¹Model available on http://www.ni.tu-berlin.de/menue/software_and_data/visual_cortex_model, direct link to github: <https://github.com/nigroup/visualcortex>.



2.2.2. Spatial Activity Patterns

In order to determine the spatial scale of changing patterns we sampled 100 windows of 10 ms activity for each stimulus presentation and counted the spikes in these intervals for each excitatory cell. Accordingly, for each sample interval of activity of the 2D networks, we computed a spatial grid matrix of spike counts. The autocorrelation was calculated convolving the spike count matrix with itself using the fast Fourier method of the scipy package (Jones et al., 2001). Next, to compute the mean spatial autocorrelation of a network we normalized the autocorrelations and averaged across the sampled time windows and stimulus presentations. This procedure is visualized in **Figure 2**.

2.2.3. Self-sustained Activity

If recurrent excitation becomes too strong, i.e., if the excitatory coupling strengths surpass a threshold determined by the strength of lateral inhibition, persistent spiking activity can be observed even in the absence of afferent stimulation. In this case usually spiking activity diverges and the spiking model neurons fire close to their maximum firing rate. In order to test for self-sustained activity, stimulation was applied for 1 s and 1 s after stimulation was terminated², spikes were counted within another 1 s interval. If spikes were observed the network was classified as *self-sustained*, viz. the network was able to sustain spiking without any external drive. For all analyses we only considered activity that was not self-sustained.

²To verify this method, periods of 9 s were tested as well and the results were similar.

2.2.4. Orientation Selectivity

In order to quantify the sharpening of network tuning we applied a measure called *Orientation Selectivity Index* (OSI) (Swindale, 1998):

$$\text{OSI} = \frac{\sqrt{\left(\sum_{s_i} \langle r(s_i) \rangle \cos(2s_i)\right)^2 + \left(\sum_{s_i} \langle r(s_i) \rangle \sin(2s_i)\right)^2}}{\sum_{s_i} \langle r(s_i) \rangle}, \quad (5)$$

where $\langle r(s_i) \rangle$ is the average spike count of a model neuron for a particular stimulus s_i . We chose twelve evenly spaced orientations $s_i \in \{-90^\circ, -75^\circ, -60^\circ, \dots, 75^\circ\}$. $\langle r(s_i) \rangle$ was averaged over 10 stimulus presentations per orientation. The OSI is a measure of tuning sharpness that ranges from 0 to 1. Values close to 0 correspond to weakly tuned neural responses whereas values close to 1 correspond to well tuned activity. We computed the average OSI then across the whole excitatory neuron population in each network.

2.2.5. Fisher Information

The quality of a population code was quantified using an estimate of Fisher information. The numerical information experiments were based on repeated presentations of oriented stimuli of -1 and 1 degree. After an initial phase of 1 s, 3 s of stimulation were used to estimate Fisher information. The number of stimulus presentations varied between 125 and 7,000 per oriented stimulus depending on the number of neurons we considered for readout.

We chose at least as many stimulus presentations as the number of readout neurons that we sampled. This avoids sets of too few training data items of a number lower than the amount of Fisher estimation parameters (see below). Fisher information was estimated using an approach from Seriès et al. (2004). In most experiments 60% of the simulation data were used to train a locally optimal linear estimator (LOLE) to predict the stimulus orientation from the spiking activity. In the experiments based on 7,000 samples, however, we chose 80% of the stimulus presentations for training. The estimator has the form:

$$\hat{s} = \mathbf{w}^T \mathbf{r} + w_0, \quad (6)$$

where \mathbf{w} is a vector of weights, $\mathbf{r} = (r_1, \dots, r_N)^T$ is a vector containing the spike response of the read out neuron sample in the fixed time window of 3 s, and w_0 a bias weight. \mathbf{w} and w_0 were optimized to reduce decoding error based on the given training data. The parameters were fitted via stochastic gradient descent using the scikit-learn Python library (Pedregosa et al., 2011). Training was stopped in case generalization performance on the validation set—containing 20% (10% for 7,000 presentations) of the data—decreased for 500 consecutive epochs. After training was completed, mean and variance of the estimates were computed for both orientations ($s_1 = -1.0^\circ$ and $s_2 = 1.0^\circ$) based on the remaining 20% (10%) of the data. Fisher information was approximated via:

$$I_{\text{LOLE}} = \frac{\left(\frac{\langle \hat{s}_2 \rangle - \langle \hat{s}_1 \rangle}{s_2 - s_1} \right)^2}{\frac{1}{2}(\text{Var}(\hat{s}_2) + \text{Var}(\hat{s}_1))}, \quad (7)$$

where $\langle \hat{s}_i \rangle$ denotes expectation and $\text{Var}(\hat{s}_i)$ is the variance of the estimates for one particular stimulus across different presentations. I_{LOLE} provides a lower bound on the Fisher information.

We further computed the information available in a data set where all correlations have been artificially removed by shuffling the spike counts of individual neurons observed across different stimulus presentations (I_{shuff}). An example of how data is shuffled across stimulus presentations is given in **Table 1**. A comparison between I_{LOLE} and I_{shuff} shows whether correlations improve ($I_{\text{LOLE}} > I_{\text{shuff}}$), degrade ($I_{\text{LOLE}} < I_{\text{shuff}}$) or have no effect ($I_{\text{LOLE}} = I_{\text{shuff}}$) on stimulus encoding quality. Additionally, we computed the diagonal information I_{diag} to investigate whether the correlations themselves carry information about the stimulus. To estimate I_{diag} the decoder was trained on the shuffled data but applied to the original one. In case I_{diag} is smaller than I_{LOLE} , correlations carry information that would be lost if the decoder ignores the correlation structure. Note that the value of I_{diag} is bounded by I_{LOLE} from above.

Moreover, in order to test if the linear decoder is suitable for information estimation, we used a non-linear support vector regression (SVR) decoder with radial basis function (RBF) kernels for comparison. We transformed the observed spike count data to zero mean and variance one, and optimized the two SVR hyperparameters by grid search. We explored the penalty function parameter of the SVR error ($C_{\text{SVR}} \in$

TABLE 1 | These tables give an example of shuffled data that can be used to compute the shuffled information I_{shuff} .

Neuron	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Total
Neuron I	1	2	3	4	5	15
Neuron II	1	1	0	3	3	8
Neuron III	1	2	0	5	4	12
Neuron	Shuff 1	Shuff 2	Shuff 3	Shuff 4	Shuff 5	Total
Neuron I	3	4	1	5	2	15
Neuron II	3	0	1	1	3	8
Neuron III	4	1	5	2	0	12

The top table lists the original spike counts observed for 3 hypothetical neurons over 5 trials. The bottom table depicts the shuffling of trials. The observed spike counts are shuffled for each neuron. Note that counts are shuffled across trials but not across neurons. Hence, the totally observed number of spikes for each neuron (last column) remains equal between the non-shuffled and shuffled condition.

{0.01, 0.1, 1, 10}) as well as the width of the RBF kernel ($\gamma_{\text{SVR}} \in \{0.0001, 0.001, 0.01, 0.1\}$) using 5-fold cross-validation on the training set before validation on the test data.

For all information processing analyses we only considered excitatory neurons because GABAergic inhibitory interneurons are not known to project out of primary visual cortex (Seriès et al., 2004; Schmolesky, 2007).

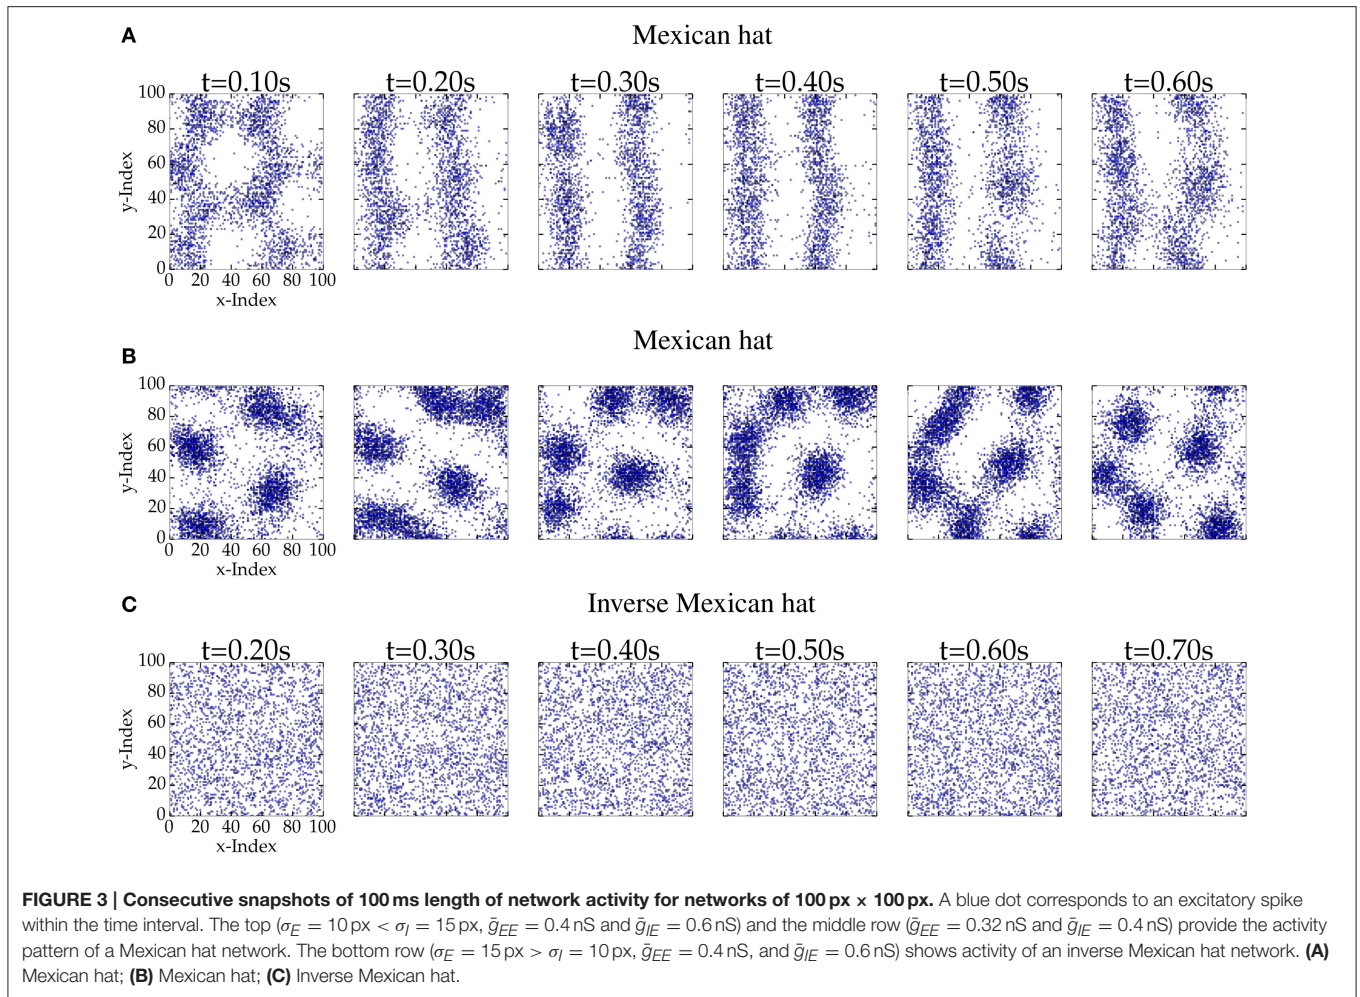
3. RESULTS

We tested network models of different sizes using either 100×100 or 200×200 excitatory neurons with 2,500 or 10,000 inhibitory neurons, respectively. These two sizes correspond to $1.5 \text{ mm} \times 1.5 \text{ mm}$ and $3 \text{ mm} \times 3 \text{ mm}$ of cortical area. For the smaller (100×100) networks we also simulated responses to heterogeneous, orientation tuned input using an orientation map consisting of 4 pinwheels as depicted in **Figure 1**. The smaller networks were also used for the assessment of the quality of stimulus encoding.

3.1. Correlated Variability for Homogeneous External Input

First, we investigated the firing patterns that emerge for different network configurations. **Figure 3** shows consecutive snapshots of spiking activity for different network parameterizations: Two Mexican hat and one inverse Mexican examples. The spiking activity of Mexican hat networks revealed bump or stripe like patterns. Bumps may move and fuse together, but also eventually appear, disappear, pulsate or turn into flickering stripes. These stripes can be stable over a few seconds and may turn back into circular bumps for the same network realization. In contrast, for inverse Mexican hat (**Figure 3C**) and balanced networks (data not shown) we did not observe spatially inhomogeneous patterns.

We next quantified the noise correlations for different network settings. We computed the spike count correlation coefficient among pairs of cells for repeated stimulus presentations. An example distribution of spike counts for



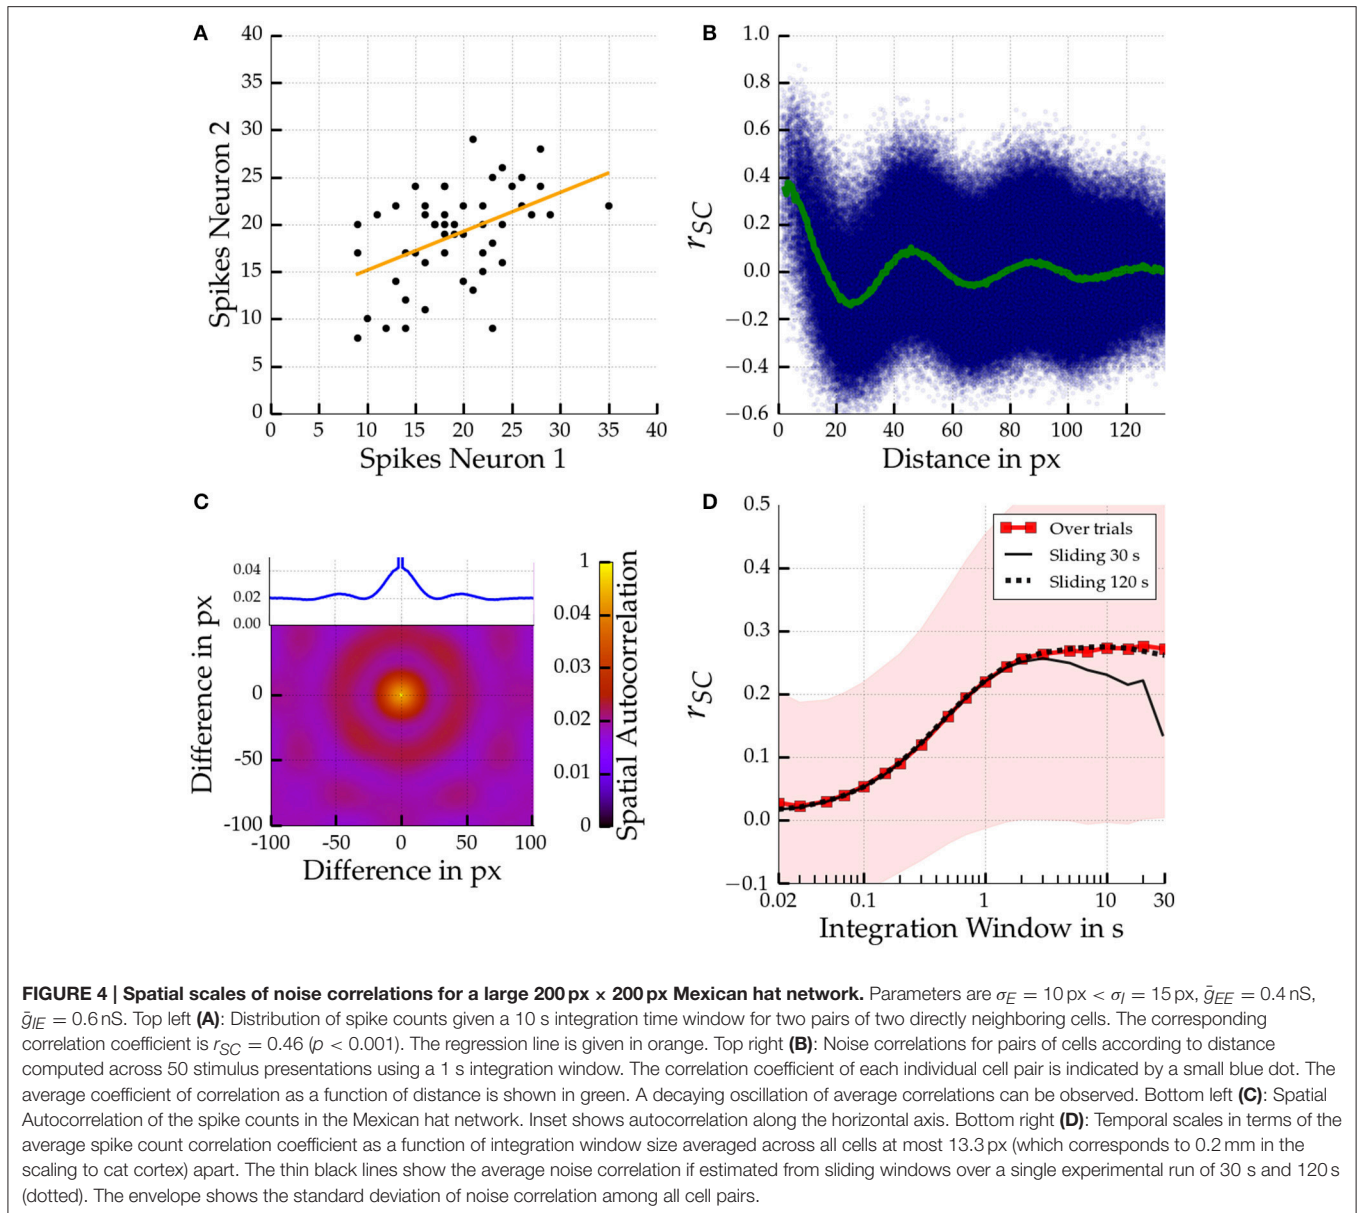
a pair of neighboring cells (1 px apart) is shown in **Figure 4A**. **Figure 4B** displays the distribution of spike count correlation coefficients among different cell pairs for a large 200×200 Mexican hat network as a function of distance. The average correlation as a function of distance is denoted by the green line and follows a damped wave pattern. We did not make such an observation for balanced or inverse networks (data not shown). The decay of the average correlation (green line) happened rather quickly and there was a considerable decrease in the amplitude within the first 100 pixels, which corresponds to 1.5 mm in visual cortex (see Section 2). The oscillation frequency of the average r_{SC} was 2.2 cycles per 100 pixels.

We hypothesized that noise correlations are caused by time-varying, spatially inhomogeneous activity patterns, such as moving bumps. We, therefore, investigated whether the spatial scale of noise correlations corresponds to the spatial scale of the moving bumps. Indeed, the spatial frequency of the noise correlations matches the spatial autocorrelation of the heterogeneous activity profile as depicted in **Figure 4C**. Both are on the order of 2.2 cycles per 100 pixels. We made similar observations for one-dimensional networks where the relation between spatial frequency and bump size becomes even more

obvious because one can easily count the number of bumps emerging in a network (see Supplementary Material).

Experimental studies reported a saturation of spike count noise correlations for increasing time windows (Bair et al., 2001; Reich, 2001; Smith and Kohn, 2008). We were interested if this phenomenon can be observed in our model as well and therefore investigated the temporal scale of the noise correlations. **Figure 4D** shows the average correlation among cell pairs with a maximum distance of 13.3 px, corresponding to 0.2 mm in visual cortex, for the large Mexican hat network. We found a saturation of the magnitude of the distance dependent correlations when increasing the integration time window. The magnitude saturates for time windows of about 1 s and longer. Similar observations could be made for the calculation of noise correlations with a sliding window and a long single stimulus presentation of 120 s (dotted black line). Usage of a shorter stimulus presentation of 30 s yielded similar results but only up to sliding windows with a length of about 3 s (black line).

Rosenbaum and Doiron (2014) demonstrated that neural network activity can be balanced for Mexican hat configurations due to finite size effects. Accordingly, we were interested if we could observe balanced and uncorrelated neural responses for



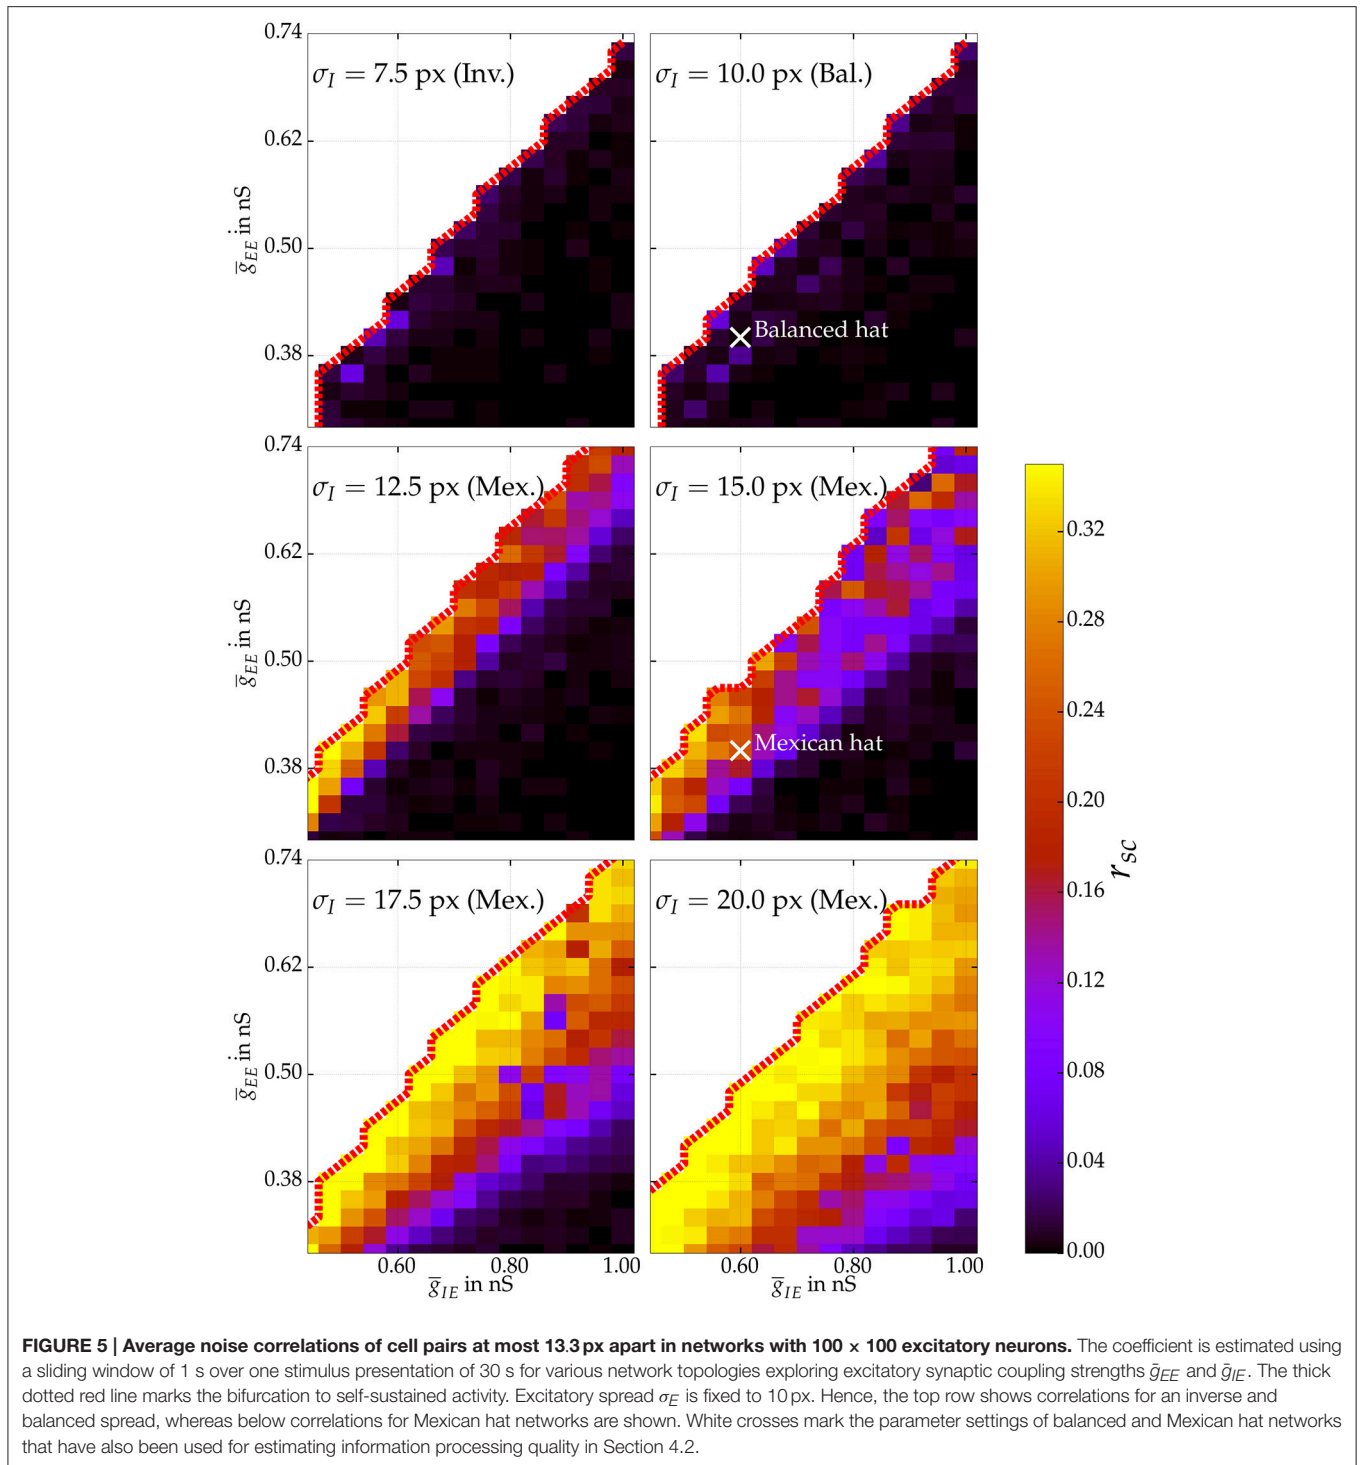
Mexican hat networks and determine the parameter regimes for which correlations emerge. We varied the parameters \bar{g}_{EE} and \bar{g}_{IE} of the connection strengths for several networks with 100×100 excitatory neurons. **Figure 5** depicts the average noise correlation among cell pairs at most 13.3 px apart as a function of different recurrent connection strengths for different network configurations. To reduce simulation time the noise correlations were calculated based on a stimulus presentation of 30 s and using a sliding window of 1 s. Noise correlations were observed for Mexican hat configurations with recurrent connectivity strengths close to the boundary of self-sustained activity (red dotted lines). For weak recurrent connection strengths no correlations were observed, indicating a dominant finite size effect. Increasing the width of the Mexican hat, i.e., scaling σ_I while σ_E was fixed, we measured strong correlations for a much larger range of

parameter values, i.e., these networks are less affected by finite size effects.

Moreover, we investigated the relation between network size and correlations directly in one-dimensional networks (see Supplementary Material). We could show that increasing the number of neurons while keeping all other parameters fixed (including the extent of the network) increased noise correlations in Mexican hat networks.

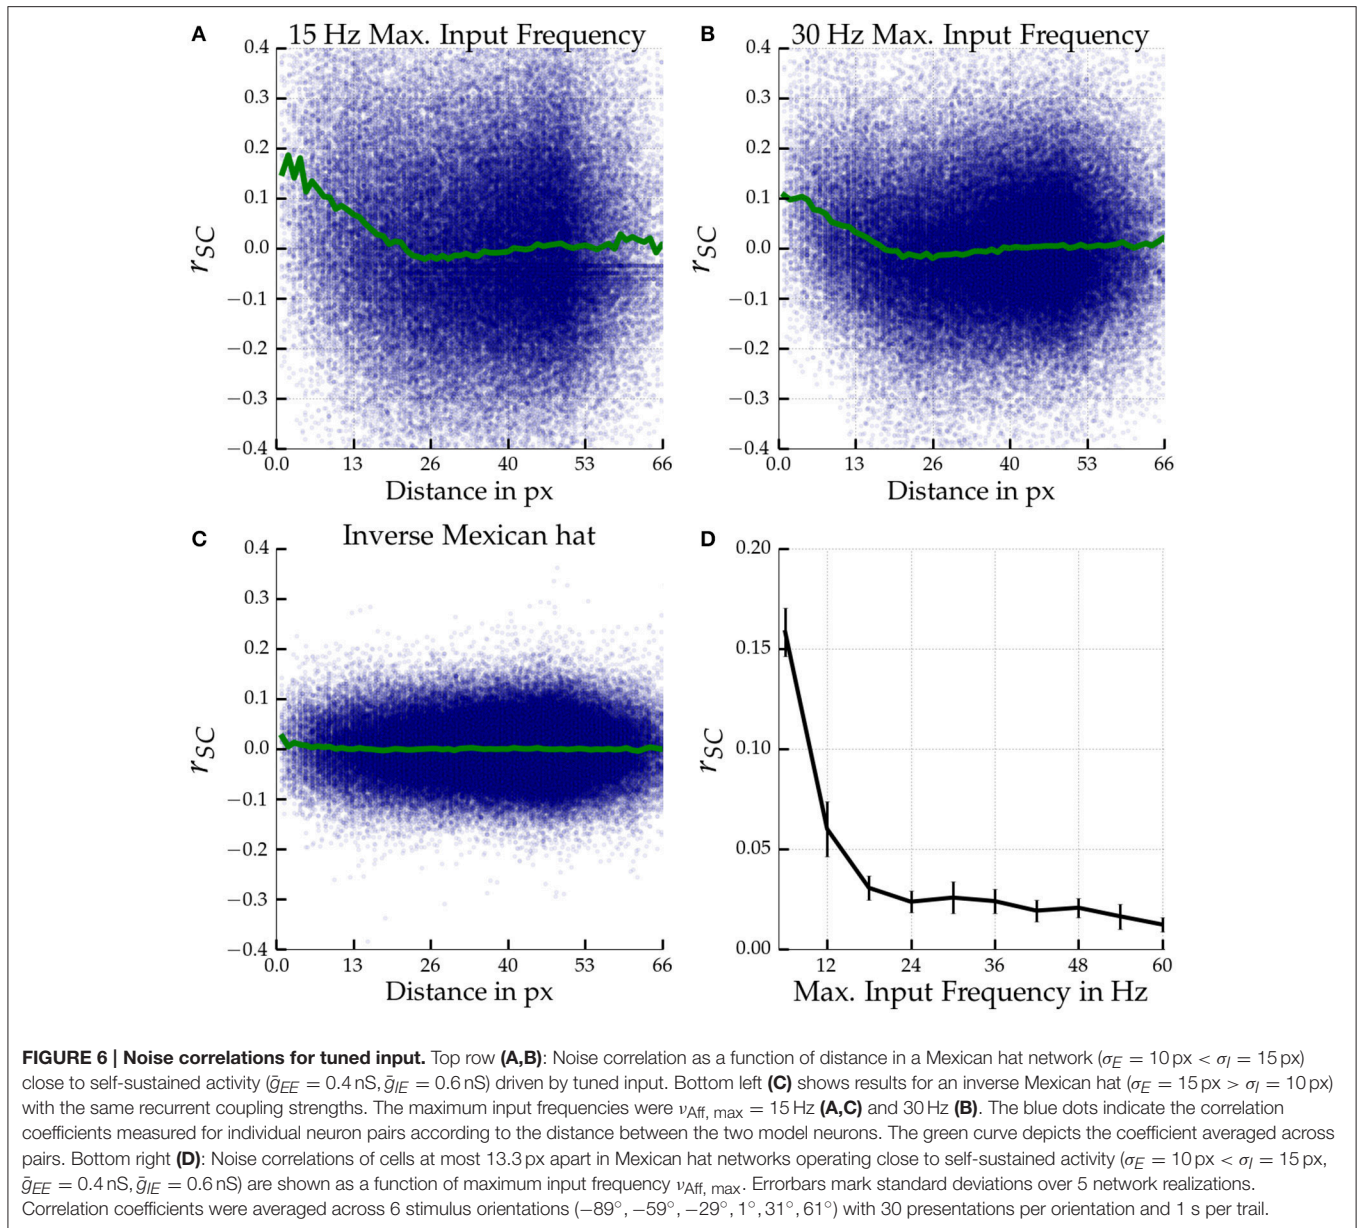
3.2. Correlated Variability for Orientation Stimuli

Figures 6A,B show the distance dependence of noise correlations in a Mexican hat network close to self-sustained activity for orientationally afferent external drive of different maximum strengths. Noise correlations decrease linearly with distance



within the first 25 pixels, corresponding to about 0.4 mm of visual cortex. The strong oscillatory fluctuations of noise correlations that we observed for homogeneous stimuli almost vanished and we observed only a minor dip below zero of the average noise correlation as a function of distance. As expected, no distance dependent correlations were observed

for an inverse Mexican hat network with the same recurrent coupling strengths (**Figure 6C**). Overall the magnitude of noise correlations decreased for Mexican hat networks in comparison to the previous experiments with a blank stimulus. Increasing the maximum input rate $\nu_{Aff, \max}$ from 15 Hz (**Figure 6A**) to 30 Hz (**Figure 6B**) slightly decreased the magnitude of the



noise correlations further. We systematically investigated the influence of the maximum rate $\nu_{\text{Aff, max}}$ of the tuned input on correlated variability. **Figure 6D** shows that the magnitude of noise correlations decreases with an increase in maximum rates.

Next, we were interested if—similar to the results obtained for homogeneous input before—the noise correlations observed for orientationally tuned input can be linked to movement of the bump patterns. For tuned input the network activity was high within four particular regions of the network (see Supplementary Video), where the preferred orientation of the model neurons correspond to the orientation of the driving stimulus in the orientation map with 4 pinwheels. We examined the movement of the centers of these activity clusters over time. In **Figure 7** four trajectories of the activity center (black

lines) at one of the network's pinwheels is depicted for an afferent stimulus with an orientation of 1 degree. For low maximum firing rate of the afferent input we identified a considerable amount of movement in the plane, but increasing the input rate significantly reduced movement. In comparison, the red lines show the same trajectories for an inverse Mexican hat driven by a stimulus of 0 degree. There, the center of activity was stable even for low maximum input rates.

3.3. Information Processing

In this section we demonstrate how the correlations in the network model influence stimulus encoding. We measured and quantified the effects of correlations on coding quality in terms of the three information measures introduced in Section 2:

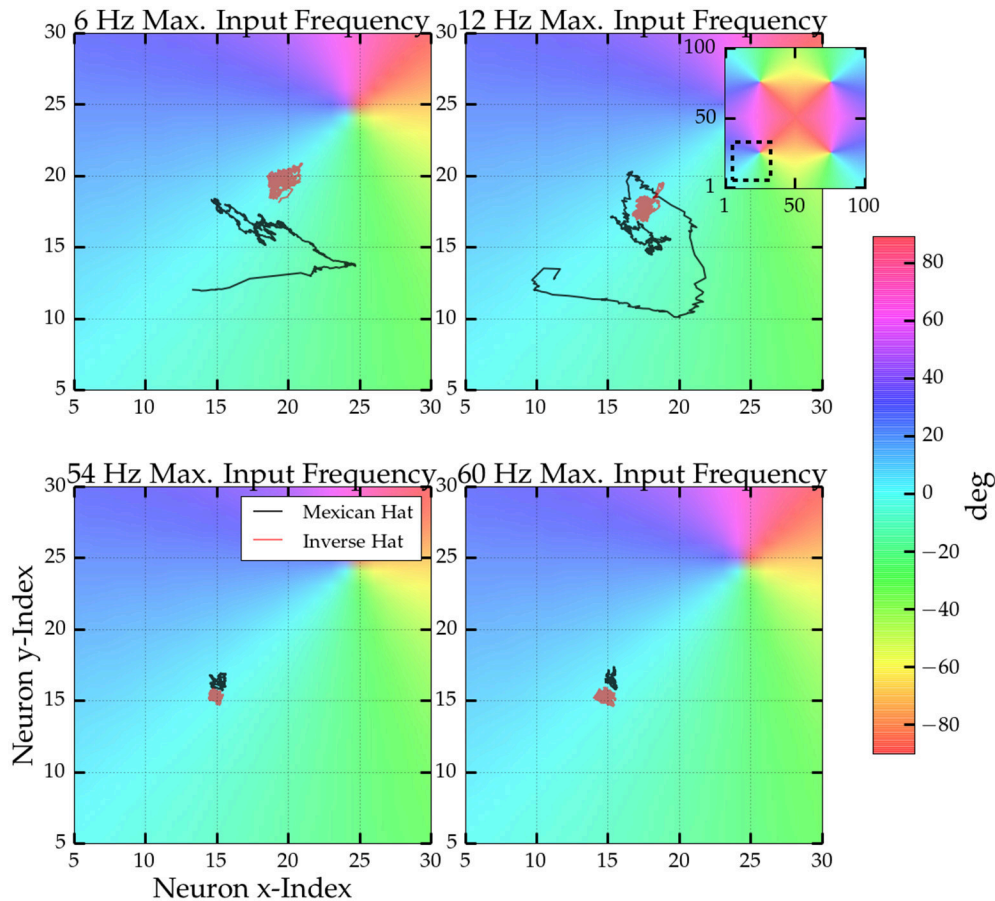


FIGURE 7 | Trajectories of the center of activity at one pinwheel for four different maximum driving frequencies. Colors indicate the orientation preferences of the model neurons. The inset shows which part of the full orientation map is depicted. Centers were computed from the activity within a sliding window of 250 ms length. Every neuron's position was interpreted as a vector in the 2D plane and the trajectory is the average vector weighted by the neurons activity within the sliding window. Black line: Trajectories for a Mexican hat network (parameters as in **Figure 6**) and a stimulus orientation of 1° . Red lines: Trajectories for an inverse Mexican hat network ($\sigma_E = 15 \text{ px} > \sigma_I = 10 \text{ px}$) with the same connections strengths and a stimulus orientation of 0° .

Using Fisher information, shuffled information, and diagonal information.

3.3.1. Fisher Information and Tuning

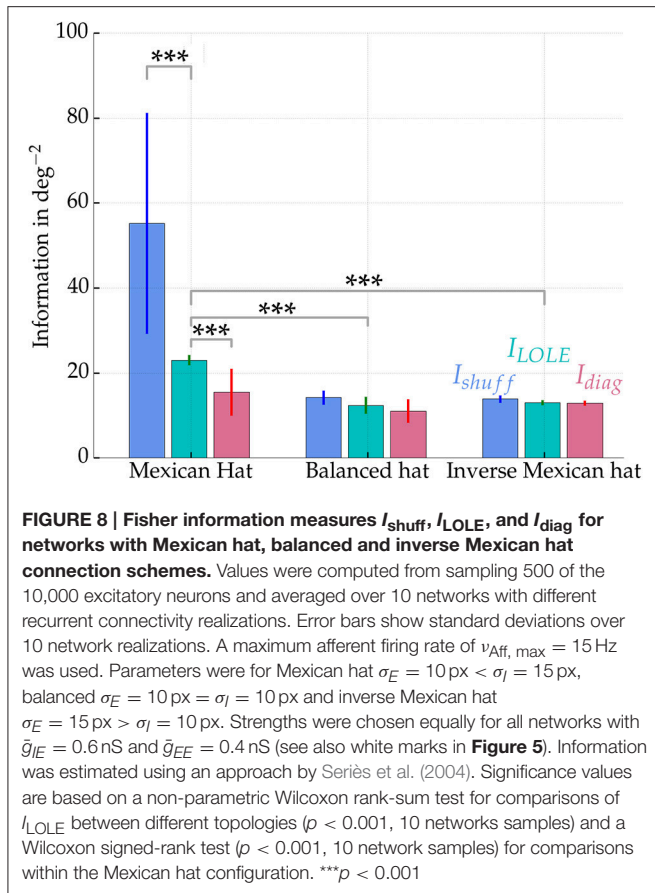
We compared Fisher information between networks with a Mexican hat configuration, a balanced configuration (see also white crosses in **Figure 5**), and an inverse Mexican hat profile.

Figure 8 summarizes the information measures for all three network configurations estimated from 500 excitatory readout neurons. The noise correlations—emerging within the Mexican hat network—significantly reduced the stimulus encoding quality in terms of reducing the Fisher information in comparison to the shuffled data (blue and cyan bars in **Figure 8** on the left, Wilcoxon signed-rank test, $p < 0.001$, 10 network samples). Furthermore, the diagonal information I_{diag} (left magenta bar) is significantly smaller than I_{LOLE} (left cyan bar) (Wilcoxon signed-rank test, $p < 0.001$, 10 network samples). Thus, the correlations themselves carried information, and a decoder cannot safely ignore correlated variability without facing a penalty in performance. Interestingly, the total information

that could be recovered from the Mexican hat network was significantly larger than in the other two networks (I_{LOLE} , cyan bars) despite the presence of noise correlations (Wilcoxon rank sum test, all $p < 0.001$, 10 network samples each).

This increase in Fisher information can be explained by the sharpening of orientation tuning by Mexican hat type connectivity patterns, which enhance stimulus encoding of low dimensional stimuli (Zhang and Sejnowski, 1999; Dayan and Abbott, 2005). Thus, despite the presence of correlations, a Mexican hat profile may better encode the stimulus by sharpening of responses. The response of the individual neurons is more selective to orientation stimuli and thereby facilitate discrimination between stimuli.

In order to quantify the sharpening of tuning we applied the *Orientation Selectivity Index* (OSI) measure (see Equation 5). **Figure 9D** shows the average OSI of the different networks. The Mexican hat networks exhibited the sharpest tuning (0.734 ± 0.003). The inverse configuration showed the second sharpest tuning (0.520 ± 0.003) followed by the balanced network with (0.471 ± 0.003). All these differences are significant (Wilcoxon



rank-sum test, all $p < 0.001$, 10 sample networks each). Notably, the coding was not simply improved due to the availability of more spikes. In **Figure 9A** it can be seen that the average firing rate of the network was smaller for Mexican hat networks. Sharper tuning implies that on average fewer neurons respond to an input stimulus, which in turn decreases the average network firing rate. In summary, despite stronger noise correlations for Mexican hat networks (**Figure 9C**), the coding quality improves alongside a sparser spiking code.

We further tested if the decoding performance could be improved over LOLE by using non-linear SVR. However, there was no increase in performance. On average its mean squared error (MSE) was slightly worse than the error of the linear decoder, as shown in **Figure 9B**. This suggests that the linear decoder is close to optimality and it is unlikely that there is information that could only be obtained by using non-linear methods.

3.3.2. Fisher Information and Lateral Connectivity

We tested whether the performance gain is indeed caused by the Mexican hat configuration or rather an effect of the width of the inhibitory connection spread only. Therefore, we explored different widths of both connection spreads (σ_I and σ_E) and kept recurrent coupling strengths fixed (\bar{g}_{EE} and \bar{g}_{IE}). To reduce simulation time we ran 125 repeated presentations per input stimulus to estimate Fisher information. We sampled

activity from 125 excitatory readout neurons. Clearly, as depicted in **Figure 10**, the phenomenon is related to the Mexican hat configuration. Best stimulus encoding performance was achieved by network topologies with shorter excitatory than inhibitory connection spread ($\sigma_E < \sigma_I$). Still, benefits in information processing are limited to a range of Mexican hats. As the area on the right side of the image shows. A very wide spread of inhibitory connections, however, led to a drop of Fisher information also for Mexican hat topologies.

3.3.3. Fisher Information and Sample Size

We further wanted to know how the difference in Fisher information between the network types depends on sampling of the number of readout neurons for Mexican hat and the inverse Mexican hat networks (cf. Averbeck et al., 2006). We repeated the previous experiments with the same parameter settings but used 7,000 presentations per stimulus and estimated I_{LOLE} again from the spike counts of different numbers of excitatory readout neurons. **Figure 11** summarizes the results. The better performance was observed for Mexican hat networks as long as roughly fewer than one third of the neurons were included. However, if more readout neurons were used to reconstruct the stimulus, I_{LOLE} was higher for the inverse Mexican hat networks. Hence, for larger samples of neurons, the detrimental nature of the noise correlations was more pronounced and eradicated the advantage of sharper tuning in Mexican hat networks.

4. DISCUSSION

4.1. Dynamics and Noise Correlations

4.1.1. Theoretical Considerations

We demonstrated that recurrent connectivity can cause correlated variability. As we hypothesized, in networks of adaptive exponential integrate and fire neurons, Mexican hat coupling with wider inhibitory than excitatory connectivity spread leads to noise correlations.

Our findings are in line with analytical results from Rosenbaum and Doiron (2014). The authors used a rate based model network to investigate how recurrent topologies in 1D and 2D networks affect excitatory and inhibitory recurrent inputs. Inverse Mexican hat configurations and those with an equally wide spread of excitatory and inhibitory connections yield a stable balance of currents, i.e., excitatory and inhibitory recurrent currents were of similar magnitude and canceled each other. This leads to constant stable firing rates of all neurons. We could show in our simulations that this stable balance further leads to a decorrelation of network activity. Rosenbaum and Doiron (2014) derived that Mexican hat coupling cannot maintain a stable balance between recurrent excitation and inhibition, thereby amplifying spatial frequencies in the network activity. The authors further demonstrated that the network size also influences this amplification. If the network size is very small or recurrent connectivity is very weak, Mexican hat networks may still exhibit homogeneous firing activity. Likewise, we observed homogeneous responses for operating regimes far from the region of self-sustained activity (see **Figure 5**). We could also demonstrate (see Supplementary Material) that increasing

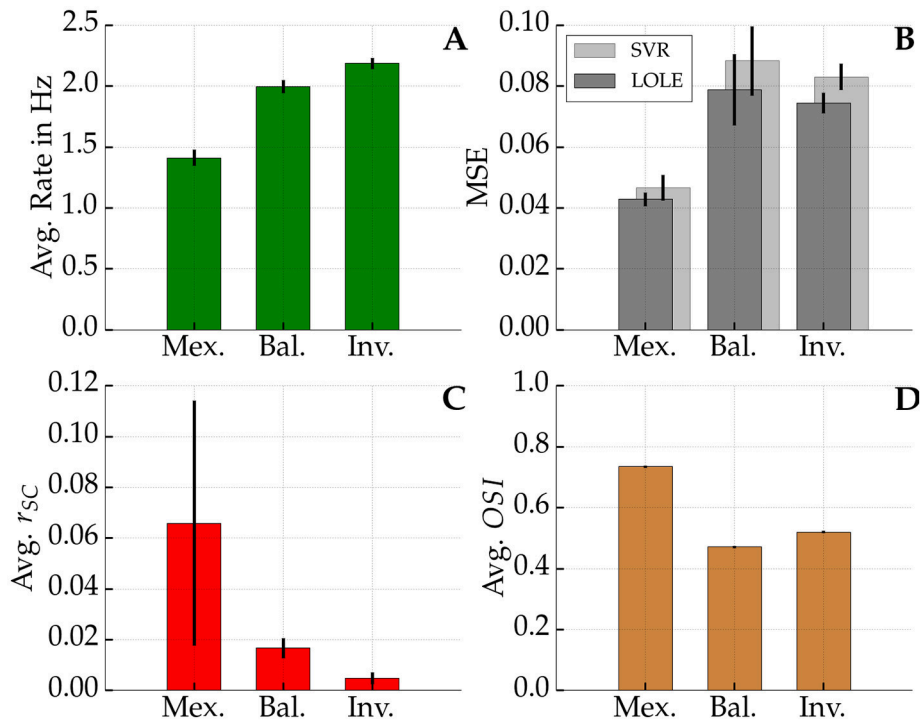


FIGURE 9 | Firing rate, MSE, r_{sc} , and OSI. Top left **(A)**: Average network firing rate averaged over all excitatory neurons. Top right **(B)**: In dark gray the average mean squared error (MSE) of the linear decoder is shown whereas in light gray one sees the average MSE of the best Support Vector Regression (SVR). Bottom left **(C)**: Average noise correlation among pairs of neurons at most 13.3px apart for a 3 s time window. Bottom right **(D)**: Average Orientation Selectivity Index (OSI) among all neurons in the networks. Black bars show the standard deviation across 10 network realizations per profile in each sub-figure. In **(D)** standard deviations are so low that the black bars are not visible. Network parameters as in **Figure 8**.

network size can lead to noise correlations in Mexican hat networks operating in regions afar from self-sustained activity.

An analytical solution of the network state could not be provided for the case that spatial frequencies are amplified by Gaussian Mexican hat connectivity (Rosenbaum and Doiron, 2014). Nevertheless, for simplified connection topologies that also feature wider inhibitory than excitatory connection spread, it has been shown that networks give rise to spatially inhomogeneous responses (Pinto and Ermentrout, 2001; Coombes, 2005; Guo and Chow, 2005; Roxin et al., 2006). For example, Hansel and Sompolinsky (1998) showed that using a sinusoidal spatial connectivity pattern the network undergoes a Turing bifurcation if the recurrent strength of the sinusoidal component reaches a critical value. The resulting network state is given by a single bump solution. Moreover, the authors demonstrated that the spatially inhomogeneous activity starts moving across the spatial extent of the network in case neuronal adaptation is considered in the rate based network model. This movement then leads to joint modulations of firing rates and, consequently, to correlation among rates. This is also sketched in the animation provided in the Supplementary Material. In our simulations of spiking neurons joint modulations in firing rates manifested themselves in joint modulations of spiking activity which appear as dynamic patterns on a network level, i.e., bumps and stripes. Due to the joint changes of spiking activity we measured distance dependent noise correlations. The

pairwise correlations were modulated in a damped oscillatory manner. These phenomena, i.e., the pattern formation and noise correlations, could be observed for a wide range of parameter settings and were most strongly pronounced close to the boundary to self-sustained activity.

The locations of emerging activity bumps can be determined and steered by inhomogeneous inputs Bressloff (2012). We made a similar observation for our networks using an inhomogeneous Poisson input that simulates an oriented bar stimulus. Inhomogeneous input locked bump activity at the locations corresponding to the peaks of the afferent input, but a weak drift or jitter of the activity could be observed as well (see **Figure 7**). Noise correlations persisted for these heterogeneous inputs, but were reduced in comparison to homogeneous inputs. The activity patterns and spatial dependency of correlations were also modulated in comparison to homogeneous inputs. We observed an almost linear decay with distance (**Figures 6A,B**). In contrast to the homogeneous input scenario, the average r_{sc} did not reach values considerably below 0. In addition, noise correlations as well as the movements or jitter around the peak locations were stronger for weak external drive.

4.1.2. Comparison to Other Modeling Studies

Keane and Gong (2015) hypothesized that phenomena such as neuronal spiking variability and noise correlations may be explained by macroscopic network dynamics such as traveling

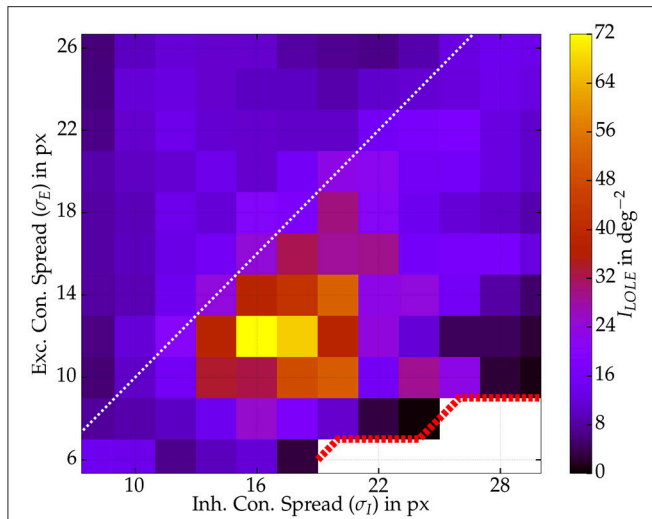


FIGURE 10 | Fisher information I_{LOLE} as a function of both connection widths σ_I and σ_E . One network per spread combination with recurrent strengths of $\bar{g}_{EE} = 0.4$ nS and $\bar{g}_{IE} = 0.6$ nS. Some Mexican hat networks with very narrow σ_E led to self-sustained activity (white area, red dotted border) and were excluded from the analysis. Values were computed from 125 excitatory neurons and 125 repetitions per stimulus ($s \in \{-1^\circ, 1^\circ\}$). Maximum input firing rate was 30 Hz.

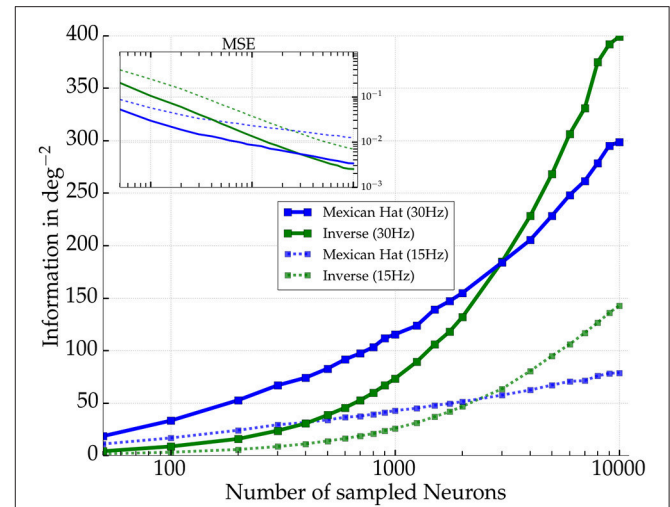


FIGURE 11 | Amount of Fisher information as a function of number of excitatory neurons sampled for two network realizations. Mexican hat ($\sigma_E = 10$ px < $\sigma_I = 15$ px) in blue and inverse Mexican hat ($\sigma_E = 15$ px > $\sigma_I = 10$ px) in green (both with $\bar{g}_{IE} = 0.6$ nS and $\bar{g}_{EE} = 0.4$ nS). Solid (dashed) lines indicate results for a maximum afferent rate of 30 Hz (15 Hz). The inset shows the corresponding mean squared error (MSE) of the LOLE decoder.

waves. The authors simulated a two-dimensional spiking neural network model where the connection probability did not depend on distance between cells but the coupling strengths of excitatory neurons did. The coupling strengths among excitatory neurons exhibited a Gaussian profile similar to our connectivity kernel. Inhibitory neurons, however, were coupled uniformly. Afferent inputs were either noisy or deterministic and homogeneous. The authors observed moving bumps activity as well as traveling wave fronts and noise correlations following a damped sine wave as a function of cell pair distance. They compared their results to topologies with uniform random connections where no pattern formation and noise correlations were observed. In our models connection strengths were constant across all synapses but connection probabilities depended on distance. In comparison to their study, we also applied Gaussian kernels to inhibitory connections, did extensive parameter explorations to uncover the influence of kernel sizes and connection strengths on correlations, and investigated inhomogeneous inputs and the relation between noise correlations and information.

Yger et al. (2011) researched the influence of Gaussian topology in terms of inhibitory and excitatory connection width on noise correlations. Surprisingly, they observed the emergence of bump patterns and noise correlations only for very narrow excitatory spreads ($\sigma_E \ll \sigma_I$). Yet, they kept recurrent weights fixed and used a rather low coupling strength. Therefore, it is likely that the network size was too small for the given recurrent weights and the finite size effect kept network activity stable.

Ponce-Alvarez et al. (2013) used a one-dimensional ring network of non-linear rate model neurons with different preferred stimulus orientations to study phenomena observed in the area MT of awake monkeys. The rate units were connected via

a Mexican hat like sinusoidal coupling. The authors hypothesized that this network topology can explain the experimentally observed directional tuning of noise correlations in MT neurons. For particular parameter regimes with strong weight on the sinusoidal component of the coupling configuration Ponce-Alvarez et al. (2013) observed bump patterns similar to ours.

Moreover, Rosenbaum et al. (2017) provide an alternative hypothesis for the emergence of noise correlations. While we solely focused on recurrent coupling, they investigated the interplay of recurrent and afferent feed-forward connections. In contrast to our model where each neuron simply received independent Poisson input, in Rosenbaum et al. (2017) afferent inputs could be shared among post-synaptic cells, depended on the distance among neurons, and afferent connection probabilities followed a spatial Gaussian kernel (with width σ_A) similar to our recurrent coupling. Using simulations of spiking neurons as well as mean-field approximations, Rosenbaum et al. (2017) showed that if afferent connection spread is narrower than excitatory or inhibitory recurrent connections ($\sigma_A < \sigma_E, \sigma_I$), distance dependent noise correlations similar to our results emerge.

4.1.3. Biological Relevance of the Results

In our study we showed that noise correlations can emerge by moving patterns of higher and lower spiking activity. On the single cell level phases of lower spiking activity alternate with phases of higher activity, depending on whether the neuron is part of a more or a less active network region.

Mochol et al. (2015) showed in the auditory cortex of anesthetized rats that noise correlations are dominated by phases of coinactivation. During spontaneous activity, the authors

observed periods where all neurons in a local network jointly reduced firing. This resembles the occurrence of bump activity states in our networks where firing rates of neighboring model neurons were jointly reduced if they are part of an off region of a bump. Moreover, Mochol et al. (2015) demonstrated that external stimulation led to a reduction of these low activity periods which went hand in hand with a decrease in the magnitude of noise correlations. Similarly, our simulation showed that increasing the strength of an orientation tuned stimulus led to a decrease of correlated variability in a 2D network (see **Figure 6D**).

Moreover, increasing input strength can be a suitable model of an increase of stimulus contrasts in biological experiments (Rubin et al., 2015). Kohn and Smith (2005) reported weaker correlations for strong contrasts of drifting orientation gratings in experiments with Macaque monkeys. These findings match our simulations with increasing afferent input of inhomogeneous stimuli.

In our simulations noise correlations depended on the size of the integration time window (**Figure 4D**). We observed an amplitude saturation for large windows. This is in line with findings from biological experiments (Bair et al., 2001; Reich, 2001; Smith and Kohn, 2008).

Smith and Kohn (2008) and Rosenbaum et al. (2017) measured correlations in visual cortex of monkeys alongside the neural tissue using a micro-electrode array (MEA) and drifting sinusoidal grating stimuli. Both reported a distance dependence of noise correlations. However, the spatial scales are different from what is observed in our model. Both reported a linear decay over a distance of several millimeters in the monkey visual cortex. In our model we observe linear decay for tuned input as well, but correlations vanished after a few ten pixels which corresponds to about 400 μm in cat visual cortex (see **Figure 6**, with 1 px corresponding to 15 μm in our scaling).

In addition, Rosenbaum et al. (2017) also found sinusoidal decay of correlations after removing latent correlations using Gaussian process factor analysis. A full oscillation cycle was observed in about 5 mm in the monkey data as well as in their spiking neuron simulation (see also the last paragraph of Section 4.1.2). In comparison, in our model we measured a full oscillation for a blank stimulus within the first 50 pixels (c.f. **Figure 4B**). This result corresponds to less than a millimeter in cat visual cortex³.

Ch'ng and Reid (2010) observed a sinusoidal modulation of noise correlations with distance for spontaneous activity in the visual cortex of rats using two-photon imaging. In cat visual cortex Ch'ng and Reid (2010) only reported a seemingly exponential decay. However, they could measure noise correlations for a maximum pair distance of only 400 μm . Our model simulations imply that only beyond this distance one

should expect an increase in correlations due to a sinusoidal modulation. Therefore, Ch'ng and Reid (2010) might have measured only the first quarter of a cycle of the spatial oscillation.

The emergence of spatially inhomogeneous patterns of spontaneous activity on the scale of several hundred micrometers has been long known (Arieli et al., 1995). For instance, Kenet et al. (2003) reported bump shaped spontaneous activity in the visual cortex of anesthetized cats using voltage sensitive dyes. More important, the spatial scales of the patterns, i.e., the size of the bumps, matched the distances of pinwheel centers of the preferred orientation map. They measured about one bump per pinwheel. This is also the case in our simulations. In our 2D network model we measured a spatial autocorrelation of about 2.2 cycles per 100 neurons, which agrees with the frequency of 4 pinwheels per 100×100 cells. This means we also observed about one bump per pinwheel.

4.1.4. Limitations

The qualitative comparison of our model with biological data should, however, be taken with care because of our assumption of conditionally independent afferent Poisson input. This afferent input is not consistent with realistic LGN input. For instance, Lin et al. (2012) demonstrated that more realistic LGN models can change response characteristics such as tuning sharpening of post-synaptic V1 neurons. In this manuscript, however, we study the effect of recurrent connectivity in isolation and excluded other sources of correlated variability.

In our model all cells of a population exhibited equal parameter values (except those for coupling). In biology, however, one can observe a large variability among neuronal parameters such as membrane and adaptation time constants (Sanchez-Vives et al., 2000) or for the amplitudes of spike triggered post-synaptic potentials (Mason et al., 1991). A modeling study by Mejias and Longtin (2012) suggests that such heterogeneities can even enhance information processing in neural networks.

Furthermore, in our 2D networks with 100×100 and 200×200 excitatory neurons corresponding to 1.5 mm \times 1.5 mm or 3 mm \times 3 mm of cortical area, respectively, we obtain an average number of about 5,500 neurons per square millimeter (including inhibitory neurons). This is much less than the number of 15,000 neurons per square millimeter estimated for layer 4 (Beaulieu and Colonnier, 1983). Yet, the theory by Rosenbaum and Doiron (2014) states that increasing the network size by keeping the spatial extent constant even fosters inhomogeneous activity and stable Mexican hat networks may yield pattern formation of spiking activity due to a vanishing finite size effect. Accordingly, we could also demonstrate (in the Supplementary Material) that larger neuron numbers per extent yielded correlated variability for even more parameter settings.

We assumed periodic boundary conditions. Yet, we do not believe that our results would change considerably using other boundary conditions. We tested reflecting as well as absorbing boundary conditions in one-dimensional

³Note the scale difference in network extents and parameterization between our and their model is on the order of the difference in the spatial modulation of the correlations. For example, in our model the excitatory recurrent connection had a spread corresponding to $\sigma_E = 150 \mu\text{m}$ whereas Rosenbaum et al. (2017) used $\sigma_E = 500 \mu\text{m}$.

networks (see Supplementary Material) and observed the formation of similar patterns including the moving bump structures.

4.2. Neural Coding

We demonstrated that noise correlations observed for Mexican hat topologies decrease the Fisher information in comparison to shuffled data where correlations have been removed. However, in comparison to other topologies like balanced or inverse Mexican hat, more information about the stimulus could be extracted even if correlations were present (cf. **Figure 8**) as long as a subset of neurons were considered for readout.

4.2.1. Sharpening of Orientation Tuning

The sharpening of the orientation tuning due to the Mexican hat configuration may explain this relative gain in information in comparison to other topologies for neuron sub-sampling. Yet, as number of readout neurons increased, the deteriorating effect of correlation worsened and eliminated the advantage of tuning sharpening. Whether the improvement of coding of a sample readout population in terms of Fisher information carries over to situations where neurons are tuned to more than one feature, however, remains unclear. For stimulus feature spaces with more than two dimensions and uncorrelated neural firing, sharpening of feature tuning yields a reduction in Fisher information (Zhang and Sejnowski, 1999; Dayan and Abbott, 2005).

4.2.2. Comparison to Other Modeling Studies

Similar to our study Serié et al. (2004) investigated the influence of noise correlations and sharpening of tuning curves on stimulus information. They used two different models. In the first model tuning was sharpened by strong recurrent Mexican hat connectivity much like in our model. In the second model the tuning of the afferent input was already sharp and no excitatory recurrent connections were present. The afferent inputs were chosen such that the output tuning curves of the two network models matched. In this setting, the Mexican hat network performed much worse than the model with afferent tuning. Besides a strong reduction in information due to noise correlations, Mexican hat networks showed already a much lower value of Fisher information for shuffled data where correlations were removed. However, Serié et al. (2004) used a Mexican hat defined over orientation space with a very large inhibitory spread of $\sigma_I = 60^\circ$ compared to a very narrow excitatory width of $\sigma_E = 7.5^\circ$. In our simulations we discovered that very wide inhibitory spreads can have devastating effects on the encoding quality. Information was enhanced only for a particular regime of Mexican hats (see **Figure 10**). Thus, the wide inhibitory spread used by Serié et al. (2004) might explain why already shuffled information was low in their Mexican hat network.

Likewise, Hansen et al. (2012) developed a model with a Mexican hat defined over orientation space. However, their ratio between excitatory and inhibitory widths was smaller ($\sigma_E = 15^\circ$ and $\sigma_I = 40^\circ$). Accordingly, in comparison to the work by Serié

et al. (2004), less Fisher information was lost in Mexican hat networks relative to topologies with wider excitatory spread.

Moreno-Bote et al. (2014) demonstrated analytically and numerically with networks of leaky integrate-and-fire neurons that a particular type of correlations is detrimental to stimulus encoding. They termed these *differential correlations*. The authors showed that the noise covariance matrix \mathbf{Q} of neural responses to a stimulus s can be decomposed as:

$$\mathbf{Q}(s) = \mathbf{Q}_0(s) + \varepsilon \mathbf{f}'(s) \mathbf{f}'^T(s), \quad (8)$$

where \mathbf{Q}_0 represents noise that is not harmful to encoding, whereas correlations that are detrimental can take up the form $\varepsilon \mathbf{f}'(s) \mathbf{f}'^T(s)$. ε is a potentially small coefficient. \mathbf{f}' denotes the derivative of the neural tuning curve vector with respect to stimulus s , and \mathbf{f}'^T its transpose. Hence, correlations are limiting encoding quality if they shift joint neural responses tangentially along the stimulus manifold in the neural response space. More simply, assuming a one-dimensional stimulus, like orientation, and a network response in form of a Gaussian curve or a bump profile defined over the stimulus space, the following holds: If noise moves the response curve back and forth across the stimulus space, this yields differential correlations and one cannot discriminate the noise from the actual stimulus. Indeed, we made similar observations in our network model with heterogeneous stimuli. The heterogeneous input locked bump activity to a particular location, but we still observed small jitter around the location of maximum afferent input (cf. **Figure 7**). Consequently, in our 2D networks featuring an orientation map, small movements of the bump responses defined over the two-dimensional neural space simultaneously imply jitter of network responses in the stimulus space because the preferred orientations of neurons smoothly change along the two spatial axes.

Similarly, the spatial profile of the differential correlations found by Moreno-Bote et al. (2014) are reminiscent of a sinusoidal modulation whose amplitude decays with distance in stimulus space akin to our results (cf. **Figure 4A**). The authors argued that differential correlations are small ($\varepsilon \ll 1$) compared to other correlations (\mathbf{Q}_0), which makes them difficult to identify by simply measuring correlation coefficients. In order to detect differential correlations the authors suggested to use the decoder approach of Serié et al. (2004) to estimate Fisher information. Accordingly, in case Fisher information is reduced in comparison to shuffled data, differential correlations are present. This is what we observed in our simulations (cf. **Figure 8**).

Kanitscheider et al. (2015) developed a generative model of differential correlations based on convergent feed-forward projections in a primary visual cortex model. In this model noise correlations emerged due to shared noise from LGN input to V1. The convergence of afferent inputs onto their V1 network model gave rise to shared input among their V1 model neurons. This produced correlations among the input currents which in turn led to correlations among the spiking output of the model neurons. In contrast, one can interpret our findings as a model

of differential correlations originating from recurrent processing instead of afferent input. In our model recurrent connections enable noise and adaptation to cause a drift of the network activity representing an orientation stimulus, which in turn yields differential correlations among neuron pairs.

4.2.3. Biological Relevance of the Results

A modulation of information due to noise correlations has been reported frequently in visual cortex experiments (Gu et al., 2011; Chelaru and Dragoi, 2014). Similar to our observations that noise correlations are less disadvantageous if neurons within a network were sub-sampled (cf. **Figure 11**), Montijn et al. (2014) measured a saturation of decoding performance with sample size using a variety of decoders in mouse visual cortex. Similar observations were made by Freiwald et al. (2002) reconstructing stimuli with a Bayesian decoder from data recorded in rat primary visual cortex. Comparable to our Mexican hat networks (cf. **Figure 8**), Graf et al. (2011) discovered that correlations among neurons in macaque primary visual cortex carry a significant amount of information. They reported that decoding accuracy could drop by more than five percent if a decoder ignored correlations.

Hansen et al. (2012) as well as Smith et al. (2013) found that the magnitude of correlations in monkey primary visual cortex is laminar dependent. For the input layer, often referred to as the granular (Hansen et al., 2012) or middle layer (Smith et al., 2013), the measured average r_{SC} was almost 0. Whereas in the deep or infra-granular (IG) as well as the superficial or supra-granular (SG) layers the experimenters measured significant correlations on the order of 0.1–0.2 for cells up to 300 μm apart. This can give rise to the interpretation that correlations are predominant in layers that are projecting to higher cortical areas. The output layers (IG and SG) use other stimulus encoding strategies that

are less accurate in comparison to the input layer. In contrast, the middle layer, receiving input from the LGN, and projecting only to other layers in V1, provides a very accurate and unmodified representation of a stimulus.

AUTHOR CONTRIBUTIONS

RM and KO designed the study. JL contributed to the final concept. RM implemented the network model and analyzed the simulation results. All authors wrote the manuscript.

FUNDING

This work was funded by the Deutsche Forschungsgemeinschaft (GRK 1589/1, SFB 910).

ACKNOWLEDGMENTS

We thank Timm Lochmann for his helpful comments and remarks. Moreover, we want to express our gratitude to Robert Pröpper and Philipp Meier for their help in Python programming. We further thank Lutz Schimansky-Geier and Benjamin Lindner for their advice, especially for their remarks on the section about Fisher information.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <http://journal.frontiersin.org/article/10.3389/fncom.2017.00034/full#supplementary-material>

REFERENCES

- Arieli, A., Shoham, D., Hildesheim, R., and Grinvald, A. (1995). Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. *J. Neurophysiol.* 73, 2072–2093.
- Averbeck, B. B., Latham, P. E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* 7, 358–366. doi: 10.1038/nrn1888
- Bair, W., Zohary, E., and Newsome, W. T. (2001). Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J. Neurosci.* 21, 1676–1697.
- Beaulieu, C., and Colonnier, M. (1983). The number of neurons in the different laminae of the binocular and monocular regions of area 17 in the cat. *J. Comp. Neurol.* 217, 337–344. doi: 10.1002/cne.902170308
- Beaulieu, C., Kisvárdy, Z. F., Somogyi, P., Cynader, M., and Cowey, A. (1992). Quantitative distribution of GABA-immunopositive and -immunonegative neurons and synapses in the monkey striate cortex (area 17). *Cereb. Cortex* 2(Area 17), 295–309.
- Blumenfeld, B., Bibitchkov, D., and Tsodyks, M. (2006). Neural network model of the primary visual cortex: from functional architecture to lateral connectivity and back. *J. Comput. Neurosci.* 20, 219–241. doi: 10.1007/s10827-006-6307-y
- Bressloff, P. C. (2012). Spatiotemporal dynamics of continuum neural fields. *J. Phys. A Math. Theor.* 45:033001. doi: 10.1088/1751-8113/45/3/033001
- Brette, R., and Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J. Neurophysiol.* 94, 3637–3642. doi: 10.1152/jn.00686.2005
- Bringuier, V. (1999). Horizontal propagation of visual activity in the synaptic integration field of area 17 neurons. *Science* 283, 695–699. doi: 10.1126/science.283.5402.695
- Chelaru, M. I., and Dragoi, V. (2014). Negative correlations in visual cortical networks. *Cereb. Cortex* 26, 246–256. doi: 10.1093/cercor/bhu207
- Ch'ng, Y. H., and Reid, R. C. (2010). Cellular imaging of visual cortex reveals the spatial and functional organization of spontaneous activity. *Front. Integr. Neurosci.* 4:20. doi: 10.3389/fnint.2010.00020
- Cohen, M. R., and Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nat. Neurosci.* 14, 811–819. doi: 10.1038/nn.2842
- Coombes, S. (2005). Waves, bumps, and patterns in neural field theories. *Biol. Cybern.* 93, 91–108. doi: 10.1007/s00422-005-0574-y
- Dayan, P., and Abbott, L. F. (2005). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. Cambridge, MA: The MIT Press.
- Doiron, B., Litwin-Kumar, A., Rosenbaum, R., Ocker, G. K., and Josic, K. (2016). The mechanics of state-dependent neural correlations. *Nat. Neurosci.* 19, 383–393. doi: 10.1038/nn.4242
- Freiwald, W. A., Stemman, H., Wannig, A., Kreiter, A. K., Hofmann, U. G., Hills, M. D., et al. (2002). Stimulus representation in rat primary visual cortex: multi-electrode recordings and estimation theory. *Neurocomputing* 44–46, 407–416. doi: 10.1016/S0925-2312(02)00391-0

- Gawne, T. J., Kjaer, T. W., Hertz, J. A., and Richmond, B. J. (1996). Adjacent visual cortical complex cells share about 20% of their stimulus-related information. *Cereb. Cortex* 6, 482–489. doi: 10.1093/cercor/6.3.482
- González-Burgos, G., Barrionuevo, G., and Lewis, D. A. (2000). Horizontal synaptic connections in monkey prefrontal cortex: an *in vitro* electrophysiological study. *Cereb. Cortex* 10, 82–92. doi: 10.1093/cercor/10.1.82
- Goodman, D., and Brette, R. (2008). BRIAN: a simulator for spiking neural networks in Python. *Front. Neuroinform.* 2:5. doi: 10.3389/neuro.11.005.2008
- Graf, A. B. A., Kohn, A., Jazayeri, M., and Movshon, J. A. (2011). Decoding the activity of neuronal populations in macaque primary visual cortex. *Nat. Neurosci.* 14, 239–245. doi: 10.1038/nn.2733
- Gu, Y., Liu, S., Fetsch, C. R., Yang, Y., Fok, S., Sunkara, A., et al. (2011). Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron* 71, 750–761. doi: 10.1016/j.neuron.2011.06.015
- Guo, Y., and Chow, C. (2005). Existence and stability of standing pulses in neural networks: I. Existence. *SIAM J. Appl. Dyn. Syst.* 4, 217–248. doi: 10.1137/040609471
- Hansel, D., and Sompolinsky, H. (1998). “Modeling feature selectivity in local cortical circuits,” in *Methods Neuronal Model. From Ions to Networks*, eds C. Koch and I. Segev (Cambridge, MA: MIT Press), 499–567.
- Hansen, B. J., Chelaru, M. I., and Dragoi, V. (2012). Correlated variability in laminar cortical circuits. *Neuron* 76, 590–602. doi: 10.1016/j.neuron.2012.08.029
- Hellwig, B. (2000). A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biol. Cybern.* 82, 111–121. doi: 10.1007/PL00007964
- Jolivet, R., Kobayashi, R., Rauch, A., Naud, R., Shinomoto, S., and Gerstner, W. (2008). A benchmark test for a quantitative assessment of simple neuron models. *J. Neurosci. Methods* 169, 417–424. doi: 10.1016/j.jneumeth.2007.11.006
- Jones, E., Oliphant, T., Peterson, P., and Others (2001). *SciPy: Open Source Scientific Tools for Python*. Available online at: <https://www.scipy.org/citing.html>
- Kang, K., Shelley, M., and Sompolinsky, H. (2003). Mexican hats and pinwheels in visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* 100, 2848–2853. doi: 10.1073/pnas.0138051100
- Kanitscheider, I., Coen-Cagli, R., and Pouget, A. (2015). Origin of information-limiting noise correlations. *Proc. Natl. Acad. Sci. U.S.A.* 112, 6973–6982. doi: 10.1073/pnas.1508738112
- Keane, A., and Gong, P. (2015). Propagating waves can explain irregular neural dynamics. *J. Neurosci.* 35, 1591–1605. doi: 10.1523/JNEUROSCI.1669-14.2015
- Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., and Arieli, A. (2003). Spontaneously emerging cortical representations of visual attributes. *Nature* 425, 954–956. doi: 10.1038/nature02078
- Kohn, A., and Smith, M. A. (2005). Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* 25, 3661–3673. doi: 10.1523/JNEUROSCI.5106-04.2005
- Lin, I.-C., Xing, D., and Shapley, R. (2012). Integrate-and-fire vs. Poisson models of LGN input to V1 cortex: Noisier inputs reduce orientation selectivity. *J. Comput. Neurosci.* 33, 559–572. doi: 10.1007/s10827-012-0401-0
- Mariño, J., Schummers, J., Lyon, D. C., Schwabe, L., Beck, O., Wiesing, P., et al. (2005). Invariant computations in local cortical networks with balanced excitation and inhibition. *Nat. Neurosci.* 8, 194–201. doi: 10.1038/nn1391
- Markov, N. T., Misery, P., Falchier, A., Lamy, C., Vezoli, J., Quilodran, R., et al. (2011). Weight consistency specifies regularities of macaque cortical networks. *Cereb. Cortex* 21, 1254–1272. doi: 10.1093/cercor/bhq201
- Martin, K. A., and Schröder, S. (2013). Functional heterogeneity in neighboring neurons of cat primary visual cortex in response to both artificial and natural stimuli. *J. Neurosci.* 33, 7325–7344. doi: 10.1523/JNEUROSCI.4071-12.2013
- Mason, A., Nicoll, A., and Stratford, K. (1991). Synaptic transmission between individual pyramidal neurons of the rat visual cortex *in vitro*. *J. Neurosci.* 11, 72–84.
- Mejias, J. F., and Longtin, A. (2012). Optimal heterogeneity for coding in spiking neural networks. *Phys. Rev. Lett.* 108:228102. doi: 10.1103/physrevlett.108.228102
- Meyer, R., Obermayer, K., and Gerkin, R. C. (2016). pypet : a Python Toolkit for Data Management of Parameter Explorations. *Front. Neuroinform.* 10:38. doi: 10.3389/fninf.2016.00038
- Mochol, G., Hermoso-Mendizabal, A., Sakata, S., Harris, K. D., and de la Rocha, J. (2015). Stochastic transitions into silence cause noise correlations in cortical circuits. *Proc. Natl. Acad. Sci. U.S.A.* 112, 3529–3534. doi: 10.1073/pnas.1410509112
- Montijn, J. S., Vinck, M., and Pennartz, C. M. A. (2014). Population coding in mouse visual cortex: Response reliability and dissociability of stimulus tuning and noise correlation. *Front. Comput. Neurosci.* 8:58. doi: 10.3389/fncom.2014.00058
- Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., and Pouget, A. (2014). Information-limiting correlations. *Nat. Neurosci.* 17, 1410–1417. doi: 10.1038/nn.3807
- Myme, C. I. O., Sugino, K., Turrigiano, G. G., and Nelson, S. B. (2003). The NMDA-to-AMPA ratio at synapses onto layer 2/3 pyramidal neurons is conserved across prefrontal and visual cortices. *J. Neurophysiol.* 90, 771–779. doi: 10.1152/jn.00070.2003
- Naud, R., Marcille, N., Clopath, C., and Gerstner, W. (2008). Firing patterns in the adaptive exponential integrate-and-fire model. *Biol. Cybern.* 99, 335–347. doi: 10.1007/s00422-008-0264-7
- Nowak, L. G., Azouz, R., Sanchez-Vives, M. V., Gray, C. M., and McCormick, D. A. (2003). Electrophysiological classes of cat primary visual cortical neurons *in vivo* as revealed by quantitative analyses. *J. Neurophysiol.* 89, 1541–1566. doi: 10.1152/jn.00580.2002
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830.
- Pinto, D. J., and Ermentrout, G. B. (2001). Spatially structured activity in synaptically coupled neuronal networks: II. Lateral inhibition and standing pulses. *SIAM J. Appl. Math.* 62, 226–243. doi: 10.1137/S0036139900346465
- Ponce-Alvarez, A., Thiele, A., Albright, T. D., Stoner, G. R., and Deco, G. (2013). Stimulus-dependent variability and noise correlations in cortical MT neurons. *Proc. Natl. Acad. Sci. U.S.A.* 110, 13162–13167. doi: 10.1073/pnas.1300098110
- Reich, D. S. (2001). Independent and redundant information in nearby cortical neurons. *Science* 294, 2566–2568. doi: 10.1126/science.1065839
- Renart, A., de la Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., et al. (2010). The asynchronous state in cortical circuits. *Science* 327, 587–590. doi: 10.1126/science.1179850
- Rosenbaum, R., and Doiron, B. (2014). Balanced networks of spiking neurons with spatially dependent recurrent connections. *Phys. Rev. X* 4:021039. doi: 10.1103/PhysRevX.4.021039
- Rosenbaum, R., Smith, M. A., Kohn, A., Rubin, J. E., and Doiron, B. (2017). The spatial structure of correlated neuronal variability. *Nat. Neurosci.* 20, 1–35. doi: 10.1038/nn.4433
- Roxin, A., Brunel, N., and Hansel, D. (2006). Rate models with delays and the dynamics of large networks of spiking neurons. *Prog. Theor. Phys. Suppl.* 161, 68–85. doi: 10.1143/PTPS.161.68
- Rubin, D. B., Van Hooser, S. D., and Miller, K. D. (2015). The stabilized supralinear network: a unifying circuit motif underlying multi-input integration in sensory cortex. *Neuron* 85, 402–417. doi: 10.1016/j.neuron.2014.12.026
- Sanchez-Vives, M. V., Nowak, L. G., and McCormick, D. A. (2000). Cellular mechanisms of long-lasting adaptation in visual cortical neurons *in vitro*. *J. Neurosci.* 20, 4286–4299.
- Schmolesky, M. (2007). “The primary visual cortex,” in *Webvision: The Organization of the Retina and Visual System*, eds H. Kolb, E. Fernandez, R. Nelson (Salt Lake City, UT: University of Utah Health Sciences Center).
- Seriès, P., Latham, P. E., and Pouget, A. (2004). Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. *Nat. Neurosci.* 7, 1129–1135. doi: 10.1038/nn1321
- Shadlen, M. N., and Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* 18, 3870–3896.
- Smith, M. A., Jia, X., Zandvakili, A., and Kohn, A. (2013). Laminar dependence of neuronal correlations in visual cortex. *J. Neurophysiol.* 109, 940–947. doi: 10.1152/jn.00846.2012
- Smith, M. A., and Kohn, A. (2008). Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* 28, 12591–12603. doi: 10.1523/JNEUROSCI.2929-08.2008

- Solomon, S. S., Chen, S. C., Morley, J. W., and Solomon, S. G. (2014). Local and global correlations between neurons in the middle temporal area of primate visual cortex. *Cereb. Cortex* 9, 3182–3196. doi: 10.1093/cercor/bhu111
- Stimberg, M., Wimmer, K., Martin, R., Schwabe, L., Mariño, J., Schummers, J., et al. (2009). The operating regime of local computations in primary visual cortex. *Cereb. Cortex* 19, 2166–2180. doi: 10.1093/cercor/bhn240
- Swindale, N. V. (1998). Orientation tuning curves: empirical description and estimation of parameters. *Biol. Cybern.* 78, 45–56. doi: 10.1007/s004220050411
- The Numba Development Team (2015). *Numba Version 0.22.1*.
- Touboul, J., and Brette, R. (2008). Dynamics and bifurcations of the adaptive exponential integrate-and-fire model. *Biol. Cybern.* 99, 319–334. doi: 10.1007/s00422-008-0267-4
- Yger, P., El Boustani, S., Destexhe, A., and Frégnac, Y. (2011). Topologically invariant macroscopic statistics in balanced networks of conductance-based integrate-and-fire neurons. *J. Comput. Neurosci.* 31, 229–245. doi: 10.1007/s10827-010-0310-z
- Zhang, K., and Sejnowski, T. J. (1999). Neuronal tuning: to sharpen or broaden? *Neural Comput.* 11, 75–84. doi: 10.1162/089976699300016809

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2017 Meyer, Ladenbauer and Obermayer. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.