SPIKING NEURAL NETWORK MODEL OF MOTOR CORTEX WITH JOINT EXCITATORY AND INHIBITORY CLUSTERS REFLECTS TASK UNCERTAINTY, REACTION TIMES, AND VARIABILITY DYNAMICS

A PREPRINT

Vahid Rostami*[1], Thomas Rost*[1], Alexa Riehle[2,3], Sacha J. van Albada[3,4], Martin P. Nawrot[1]

1 Computational Systems Neuroscience, Institute of Zoology, University of Cologne, Cologne, Germany
2 UMR7289 Institut de Neurosciences de la Timone (INT), Centre National de la Recherche Scientifique (CNRS)—Aix-Marseille Université (AMU), Marseille, France
3 Institute of Neuroscience and Medicine (INM-6) and Institute for Advanced Simulation (IAS-6) and JARA Institute Brain Structure-Function Relationships (INM-10), Jülich Research Center, Jülich, Germany
4 Computational Neuroanatomy, Institute of Zoology, University of Cologne, Cologne, Germany

corresponding: vhdrostami@gmail.com

ABSTRACT

Both neural activity and behavior of highly trained animals are strikingly variable across repetition of behavioral trials. The neural variability consistently decreases during behavioral tasks, in both sensory and motor cortices. The behavioral variability, on the other hand, changes depending on the difficulty of the task and animal performance.

Here we study a mechanism for such variability in spiking neural network models with cluster topologies that enable multistability and attractor dynamics, features subserving functional roles such as decision-making, (working) memory and learning. Multistable attractors have been studied in spiking neural networks through clusters of strongly interconnected excitatory neurons. However, we show that this network topology results in the loss of excitation/inhibition balance and does not confer robustness against modulation of network activity. Moreover, it leads to widely separated firing rate states of single neurons, inconsistent with experimental observations.

To overcome these problems we propose that a combination of excitatory and inhibitory clustering restores local excitation/inhibition balance. This network architecture is inspired by recent anatomical and physiological studies which point to increased local inhibitory connectivity and possible inhibitory clustering through connection strengths.

We find that inhibitory clustering supports realistic spiking activity in terms of a biologically realistic firing rate, spiking irregularity, and trial-to-trial spike count variability. Furthermore, with the appropriate stimulation of network clusters, this network topology enabled us to qualitatively and quantitatively reproduce in vivo firing rate, variability dynamics and behavioral reaction times for different task conditions as observed in recordings from the motor cortex of behaving monkeys.

1 Introduction

The performance of a behavioral task in a human or monkey requires precise processing of sensory cues, rapid decision-making, and the accurate control of movement. At the level of cortical processing, this involves the highly dynamic and successive activation of neuronal populations across different cortical areas. Despite the accuracy of behavior in such a task, it has been observed that neural activity and behavioral performance are highly variable across the repetition of experimental trials with the same conditions (Renart & Machens, 2014). Over the last few decades, attractor dynamics, which is the mathematical instantiation of the hypothesis of cell assemblies (Hebb, 1949), has been established as the most viable computational model for this type of behavioral task. Yet, despite the fact that the basic theory of attractor dynamics is well understood (Hopfield, 1982; Amit & Brunel, 1997; Roudi & Latham, 2007) and there has been considerable experimental evidence to support it (Sakai & Miyashita, 1991; Wills et al., 2005; Knierim & Zhang, 2012; Miconi et al., 2016; Pereira & Brunel, 2018; Inagaki et al., 2019), the implementation of this type of framework in realistic spiking networks, and the link between the precise computational mechanisms and behavioral and neural trial-to-trial variability remain less clear.

The classical cortical model is a balanced network of excitatory and inhibitory neurons (van Vreeswijk & Sompolinsky, 1998; Brunel, 2000) with random connectivity. This model successfully captures certain aspects of cortical spiking dynamics, including low firing rates and irregular spiking statistics. However, it fails to explain the trial-to-trial neural variability in terms of spike counts, behavioral variability in terms of reaction times, and switching between metastable states observed in cortical activity across various tasks and species. Over recent years, balanced random networks have been extended to capture neural trial-to-trial variability and multistability by accommodating cluster topology in the network architecture (Litwin-Kumar & Doiron, 2012; Deco & Hugues, 2012; Mazzucato et al., 2015, 2019; La Camera et al., 2019). These studies use purely excitatory clustering and neglect a possible structure in the topology of local inhibitory networks.

Despite the vital role of inhibitory neurons in cortical dynamics, the circuit connectivity of inhibitory neurons has remained poorly understood until recent years. A series of works found that interneurons connect non-specifically to all the surrounding pyramidal cells (Packer & Yuste, 2011; Fino & Yuste, 2011), inspiring the term “blanket of inhibition”. This understanding of the circuit connectivity motivated early studies of clustered networks (Deco & Hugues, 2012; Litwin-Kumar & Doiron, 2012) to keep the inhibitory connectivity homogeneous and non-specific while forming clusters only between excitatory neurons. Recent growing interest in inhibitory neurons, however, has brought a more complete understanding of the role of inhibitory neurons and their circuit connectivity, suggesting a higher degree of specificity and possible clustering of inhibitory neurons (Xue et al., 2014; Lee et al., 2014; Morishima et al., 2017; Arkhipov et al., 2018; Khan et al., 2018; Znamenskiy et al., 2018; Shin et al., 2019; Najafi et al., 2020). In particular, it has been argued based on anatomical and physiological evidence that inhibitory networks can be strongly interconnected locally. Moreover, excitatory neurons that receive strong excitatory input typically receive strong inhibitory input, supporting local balance at the level of individual cells (Xue et al., 2014; Okun & Lampl, 2008). In addition, theoretical studies have corroborated the importance of inhibition in attractor type dynamics and found, based on the analytical treatment of rate and binary models, that inhibitory clustering strongly improves the robustness of the metastable dynamics (Rost et al., 2017; Najafi et al., 2020).

In the present study, we propose a novel network architecture for spiking cortical attractor networks using combined excitatory and inhibitory clustering. We show that inhibitory clustering helps to retain the local balance of excitation and inhibition and generates the desirable multistability robustly over a wide range of network sizes and parameters. We utilize our model to mechanistically explain in vivo task-related dynamics of neurons recorded from motor cortex during a delayed reaching task, where the monkey receives cues that provide different degrees of uncertainty about the reaching target. We find that our model qualitatively and quantitatively captures in vivo firing rate, variability dynamics, and reaction times for different task conditions as observed in multiple single-unit recordings from the motor cortex of the behaving monkey.

2 Results

We start out with analyzing the temporal dynamics of spiking variability across and within trials in single unit recordings from the motor cortex of the macaque monkey during a delayed center-out reach task. The monkey was instructed to press one of six target buttons at the end of a 1 s delay period previous to which a varying degree of target certainty was cued according to one of three experimental conditions ((Bastian et al., 2003; Rickert et al., 2009); see section "*Experimental paradigm and data analysis*"). In the present section we consider only the simplest task condition (condition 1) in which complete target information was provided to the monkey at the onset of the preparatory period (preparatory signal, PS, indicated by the green circle in Fig. 1a). After the fixed delay period of one second, the monkey

was prompted to perform the movement by the response signal (RS, indicated by the red circle in Fig. 1a). Correct execution was rewarded with fruit juice. Extracellular recordings were taken from the primary motor areas.

2.1 Variability dynamics in monkey motor cortex and clustered spiking network models

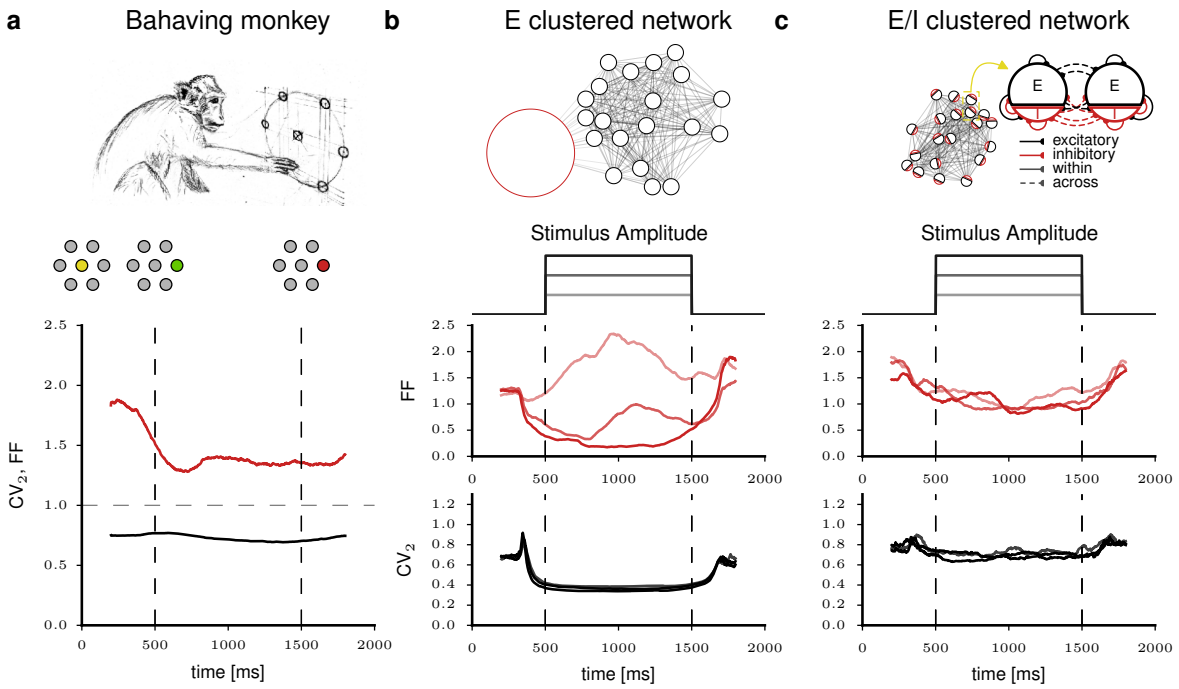

Figure 1: Variability dynamics observed experimentally and in clustered spiking network models. **a)** Experimental data recorded from the motor cortex of a macaque monkey during a delayed center-out reach task (Rickert et al., 2009). Time-resolved Fano factor (FF) and coefficient of variation (CV_2) are shown in red and black, respectively. Gray horizontal dashed line shows FF and CV_2 for a Poisson process and the vertical dashed lines indicate PS (when green LED is lit) and RS (when red LED is lit). **b)** Network model with purely excitatory clusters and global inhibition (E clustered network). Amplitude of the step-wise external input is shown in shades of gray, time-resolved Fano factors for different amplitudes are shown in shades of red, and CV_2 in shades of gray; the lighter the color, the weaker the stimulus amplitude. **c)** Similar to b but for the network model with excitatory and inhibitory clusters (E/I clustered network). In all three panels, Fano factor (FF) and coefficient of variation (CV_2) are computed in a $400\,\mathrm{ms}$ sliding window across 50 trials.

We first quantified spiking variability in our experimental data. Common statistical measures for quantifying the variability dynamics of spiking neurons are the Fano factor (FF) and the local coefficient of variation of inter-spike intervals (CV_2) (Holt et al., 1996; Nawrot et al., 2008). The FF determines the response variability across the repetition of the same experimental task (across trials) while the CV_2 measures the irregularity of spike occurrences within trials (for details see section "*Quantifying neural variability*"). Fig. 1a shows the temporal evolution of FF and CV_2 computed on the experimental data in a sliding window of $400\,\mathrm{ms}$. During the spontaneous activity before cue-onset (PS), FF is high (FF~ 1.8) and after the cue-onset, FF decreases (red curve in Fig. 1a). This behavior of FF has been observed previously in other cortical areas and different experimental tasks (Rickert et al., 2009; Churchland et al., 2010; Riehle et al., 2018). The irregularity of inter-spike intervals, CV_2, on the other hand, is fairly constant over time (black curve in Fig. 1a) and does not show any dependency on the experimental epochs (note that for Poisson variability, we expect FF$= 1$ and $CV_2 = 1$; gray dashed line in Fig. 1a).

Next, we study a biologically plausible spiking network model of cortex which can generate such variability. The reduction in count variability (FF) is commonly modeled using balanced networks of excitatory and inhibitory units where only the excitatory population is divided into clusters of stronger internal connectivity (Deco & Hugues, 2012; Litwin-Kumar & Doiron, 2012; Mazzucato et al., 2015, 2019). We first investigate the variability dynamics in this network model (which we refer to as E clustered network). The model is composed of 4,000 excitatory (E) and 1,000

inhibitory (I) exponential integrate-and-fire neurons. The E neurons are organized into Q recurrent clusters (Fig. 1b top panel). Synaptic connections within each cluster are potentiated by a factor J_+ and to maintain overall balance, connections between E neurons belonging to different clusters are depressed by a factor J_-. Inhibitory neurons are non-specifically connected to themselves and E neurons (see section "*Spiking network model*" for a detailed explanation of the model).

We simulate $2000\,\mathrm{ms}$ activity of the E clustered network in which $\sim 10\%$ of clusters are stimulated for $1000\,\mathrm{ms}$ with different stimulus amplitudes. We quantify the variability of these simulated data with FF and CV_2 in a $400\,\mathrm{ms}$ window, similar to the analysis of the experimental data. The middle panel of Fig. 1b shows the temporal modulation of FF for different stimulus amplitudes indicated by different shades of red (the darker the color, the stronger the stimulus). The consistent reduction in FF upon stimulation can be only achieved with a strong stimulus: the E clustered network shows an increase in variability for stimuli of low and intermediate strength, inconsistent with experimental observations. As for the regularity dynamics, we observe a reduction in CV_2 during stimulation. This reduction tends to be stronger for stronger stimulation and indicates that spiking activity becomes much more regular in comparison with spontaneous activity. Such stimulus dependency of regularity is inconsistent with our experimental observation (see Fig. 1a) and to our knowledge has not been reported in other studies. In summary, the E clustered network suffers from the following unrealistic dynamics: 1) Reduction in count variability (FF) is achieved only with a strong stimulus, and a weak stimulus leads to an increase in FF. 2) In the presence of a stimulus, the spiking activity becomes unrealistically regular (low CV_2).

To tackle the unrealistic dynamics observed in the E clustered network model, we suggest a more realistic connectivity structure. Recent anatomical and physiological studies point to increased local inhibitory connectivity and possible inhibitory clustering through connection strengths (Xue et al., 2014; Lee et al., 2014; Morishima et al., 2017; Khan et al., 2018; Znamenskiy et al., 2018; Shin et al., 2019; Najafi et al., 2020). We therefore combine excitatory and inhibitory clustering as we proposed previously for binary networks (Rost et al., 2017) (Fig. 1c top panel). We simulate this network model with combined excitatory and inhibitory clusters (which we refer to as E/I clustered network) using the same parameters as for the E clustered network (see section "*Spiking network model*"). Fig. 1c shows the variability dynamics in the E/I clustered network. FF decreases in the presence of the stimulus, even for weak stimulation, and CV_2 does not show stimulus dependency, similar to the experimental observations (Fig. 1a). So far we have shown that joint excitatory-inhibitory clustering enables the network to display realistic variability dynamics. In the following sections we investigate the network mechanisms which lead to such realistic dynamics by exploring the E and E/I clustered networks in detail.

2.2 Variability statistics in E versus E/I clustered networks

In this section we examine the effect of clustering parameters, namely the clustering strength J_+ and the total number of clusters Q, on the emergent variability in both E and E/I clustered networks.

Fig. 2a and c show the FF in relation to the excitatory clustering strength J_{E+} in E (panel a) and E/I (panel c) clustered networks for $Q = 50$ clusters. The statistics are computed over 20 trials of $400\,\mathrm{ms}$ duration and averaged over 20 random network realizations. For low values of J_{E+}, the FF remains unchanged on a level similar to that of the balanced random network ($J_{E+} = 1$). At $J_{E+} \approx 3.2$ an increase in FF occurs, spanning a relatively narrow range of connection strengths. It can be seen in the corresponding inset in Fig. 2a that the activity for intermediate values of J_{E+} visits different clusters but for larger J_{E+} the network can get stuck in a state with a single active cluster. Thus, the E clustered network is able to exhibit winnerless competition, but it is rather fragile. As J_{E+} is increased further, states where one cluster wins become dominant and the FF decreases below that of the unstructured networks. The reason is that the highly active clusters fire near the saturation rate where firing is very regular. This also increases the regularity of the input to the remaining clusters, which become progressively more silent as cluster strength increases.

In the case of the E/I clustered network (Fig. 2c) the cluster switching dynamics become more moderate and are spread over a much wider range of J_{E+}. With increasing cluster strength, the duration of individual up-states becomes longer and their amplitude increases. As long as the up-state duration is not much longer than the estimation window, this increases the observed rate variance and thereby the FF. The behavior is shown for the whole range of cluster strengths until the excitatory populations become completely decoupled at $J_{E+} = Q$. For $J_{E+} > 15$ or so, the average FF becomes quite noisy despite averaging over 20 random realizations. For those large values, the variance over repeated simulations is high, as some cases exhibit extensive cycling between clusters while in other cases one or more clusters win, suppressing the winnerless competition dynamics. At $J_{E+} = Q = 50$ coupling between the populations exists only through the inhibitory connections. This coupling is however still relatively strong, so that most populations are quieted by a few winners. The whole range is shown for completeness, but the interesting dynamics take place at relatively low values of J_{E+}.

So far we showed the variability dynamics in relation to clustering strength J_{E+} only for $Q = 50$ clusters. To further investigate the role of the number of clusters Q, we repeat the analysis from Fig. 2a and c but for different values of Q. To change Q we keep the total number of neurons N fixed and change the size of each cluster, i.e. larger Q means a smaller number of neurons per cluster. Fig. 2b shows FF for different values of J_{E+} and Q. There is only a very small parameter region in which the E clustered network shows high variability ($FF > 1$). This makes the E clustered network highly sensitive to the choice of parameter values, so that stable switching dynamics can be achieved only by extensive parameter tuning. In the case of the E/I clustered network, as shown in Fig. 2d, the dynamics are much more robust and a large FF is achieved for a wide range of Q and J_{E+}.

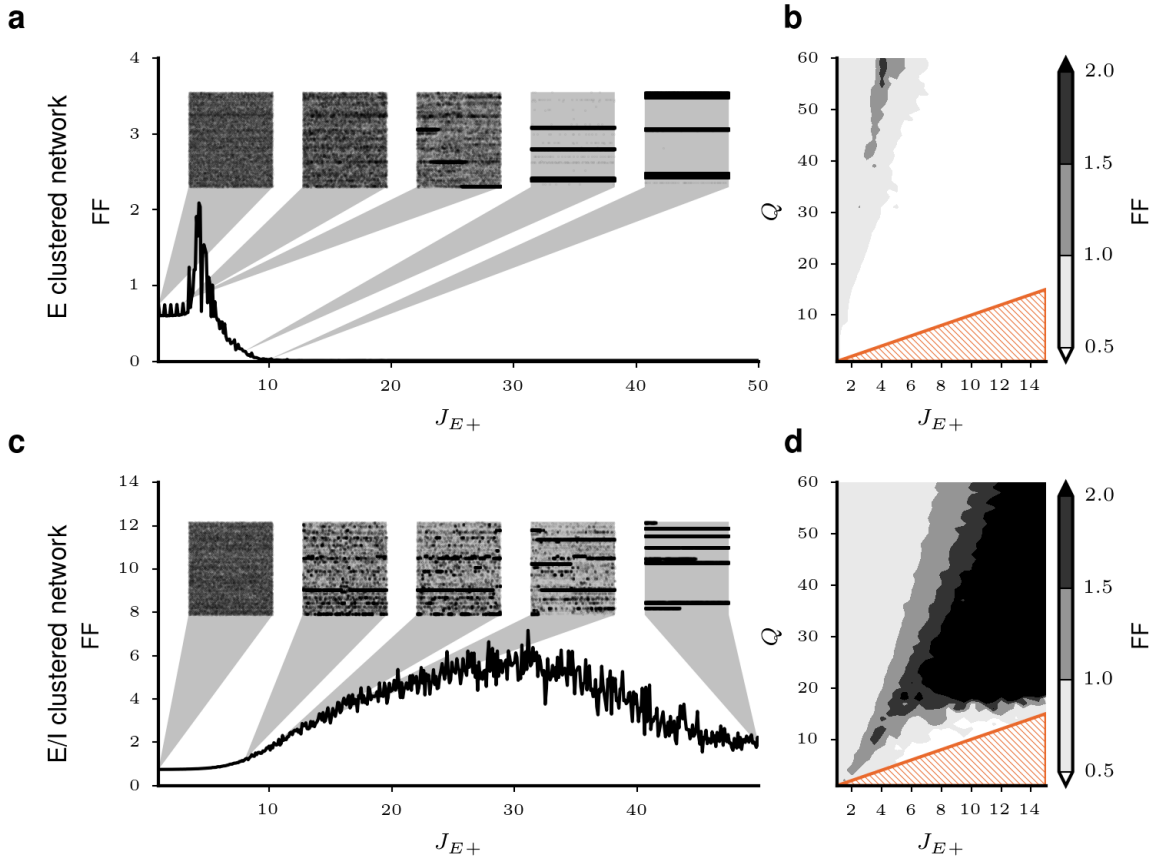

Figure 2: Effect of cluster strength J_{E+} and number of clusters Q on trial-to-trial variability FF. **a, c)** FF versus J_{E+} for networks with $Q = 50$ computed on 20 trials of 400 ms and averaged over 50 network realizations. Other parameters as in Table 2. Insets show two seconds of spiking activity of the excitatory populations. **b, d)** Pseudocolor representation of FF for different values of Q and J_{E+}. Orange shaded triangles indicate the zone below $J_{E+} = Q$, where the clusters are completely decoupled. Panels **a, b** show results for the E clustered network; **c, d** for the E/I clustered network.

2.2.1 Local balance of excitation and inhibition in the E/I network topology facilitates metastable states and spiking irregularity

Synaptic excitation and inhibition are opposing effects that together determine the activity and maintain the excitability of cortical networks (see (Isaacson & Scanziani, 2011) for a review). It has been shown in many studies that these two synaptic influences are proportional during both spontaneous and evoked cortical activity. This mechanism, known as "balance" of excitation and inhibition, is crucial for the functionality of cortical networks and disruptions of this balance can lead to extremes such as seizures (Van Vreeswijk & Sompolinsky, 1996; Anderson et al., 2000; Zhang et al., 2003; Okun & Lampl, 2008; Isaacson & Scanziani, 2011; Dehghani et al., 2016; Rubin et al., 2017).

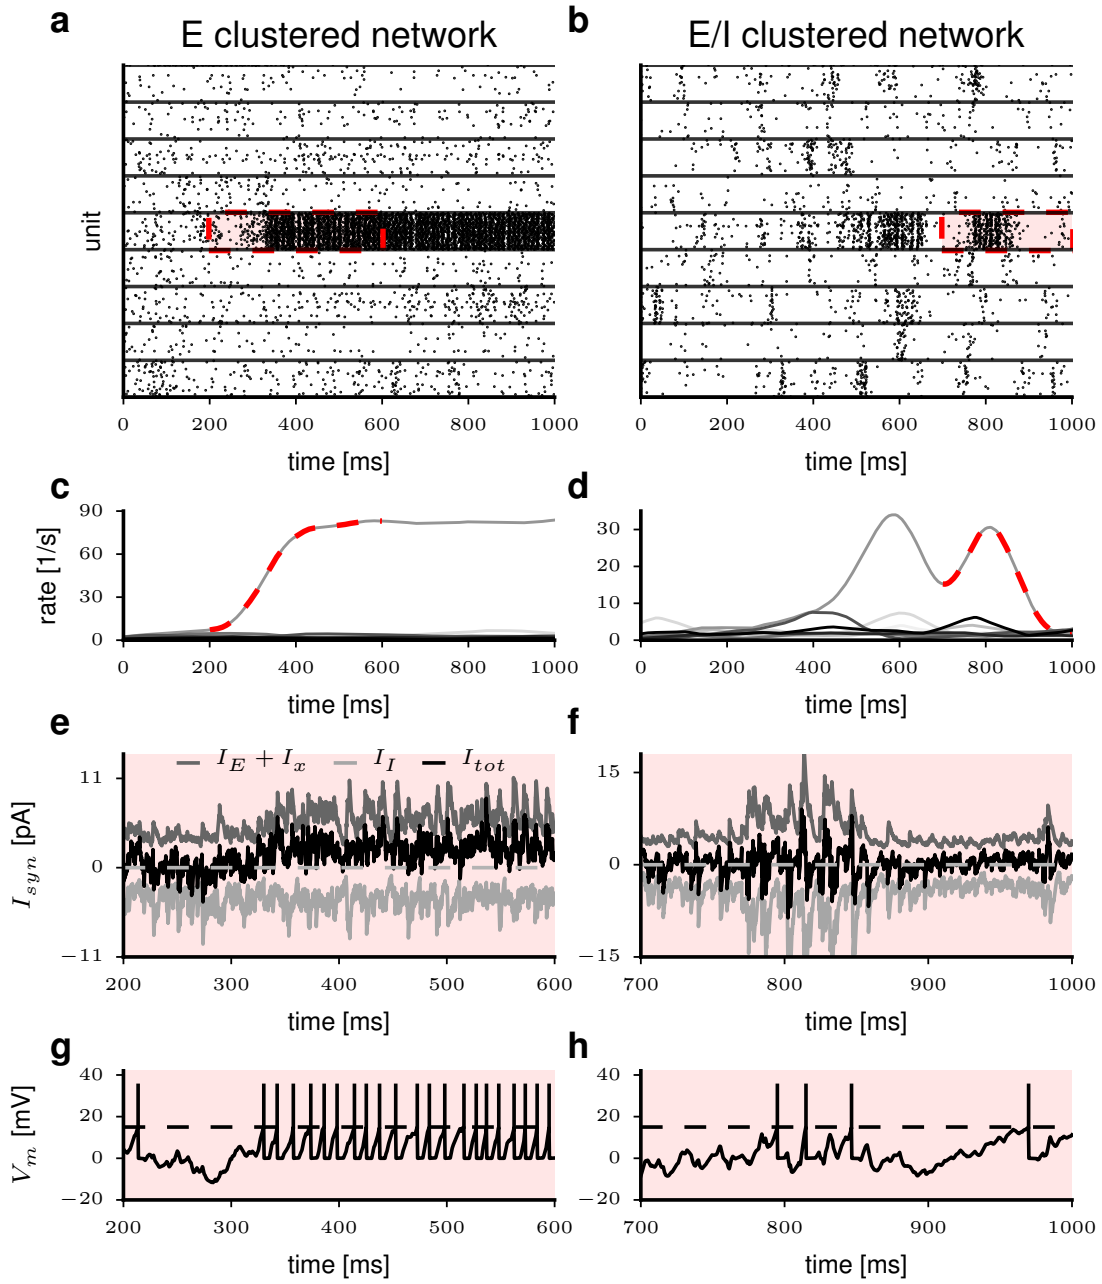

Figure 3: Comparison of attractor dynamics in E clustered (left) and E/I clustered (right) networks. **a, b)** Raster plots of nine excitatory clusters for $1000\,\mathrm{ms}$ of spontaneous activity. Horizontal lines indicate cluster membership. Red shaded area shows the region of interest around the jump from down to up state in one of the clusters. **c, d)** Average firing rate of the above 9 clusters estimated with a $50\,\mathrm{ms}$ triangular kernel. Dashed lines correspond to the red shaded region of interest in the upper panels. **e, f)** Synaptic currents I_{syn} for a randomly chosen unit during the change from down to up state for the region of interest indicated above. **g, h)** Membrane potential with superimposed spikes for the unit shown above in the same time interval.

In this section, we study the balance of excitation and inhibition in E and E/I clustered networks. Using the same parameters as in the previous sections we simulate $1000\,\mathrm{ms}$ of spontaneous activity of E and E/I clustered networks with $Q = 50$ clusters. Fig. 3 illustrates the difference between E clustered (left panels) and E/I clustered (right panels) networks on the level of individual units. The raster plot in Fig. 3a shows the activity of nine sample excitatory clusters in an E clustered network. All clusters fire at low rates until at about $t = 350\,\mathrm{ms}$ one cluster switches to the up state.

The cluster-averaged instantaneous firing rates for each of 9 clusters, estimated with a $50\,\mathrm{ms}$ triangular kernel (Nawrot et al., 1999), are shown in Fig. 3c. The firing rate of the active cluster increases dramatically (~ 90 spikes/s), while the rates of the other clusters are suppressed. Fig. 3e shows the synaptic input currents to a single neuron from the cluster which jumps to the higher activity state around the onset of cluster activation (red area). It can be seen that, while the excitatory input current increases (I_E upper trace) due to self-excitation within the cluster, the net inhibitory input (I_I, lower trace) remains constant. Hence, the total input current (I_{tot}, middle trace) increases and spiking becomes mean-driven rather than fluctuation-driven. This means that a large portion of the units in the active cluster synchronize and fire extremely regularly at the saturation rate dictated by the refractory period, as can be seen in Fig. 3g, which depicts the membrane potential of the same unit.

The right-hand side of Fig. 3 shows the equivalent scenario for an E/I clustered network. The raster plot, Fig. 3b, shows that at this mild stage of winnerless competition, the clusters cycle between short bursts of higher activity and occasional longer excursions to the up state of one or more clusters. Fig. 3d shows that this cycling takes place at much lower firing rates than in the E clustered network. Fig. 3f shows the synaptic currents of a sample neuron at the transition from low to higher activity of its cluster. The effect of inhibitory clustering is that the inhibitory currents also increase during up states, thereby increasing the variance but not the mean of the synaptic currents, keeping inputs balanced throughout. Hence the units remain in the fluctuation-driven balanced state with irregular spike trains, as illustrated in Fig. 3h. Although synchrony arises also in this case as units in the same cluster tend to have correlated rates, the tight locking seen in Fig. 3a is not observed in the fluctuation-driven regime. This synchrony comes along with regular spike trains of individual neurons, which explains the drastic decrease of the CV_2 shown in Fig. 1b for the E clustered network.

2.2.2 Stimulus dependency of variability dynamics

Having shown that inhibitory clustering is a crucial component to achieve stable attractor dynamics in the spontaneous activity of spiking networks, we now investigate the role of inhibitory clusters during evoked activity. In Fig. 1 we have seen that the E clustered network fails to capture the reduction in trial-to-trial variability during stimulation for a weak stimulus. Here we ask what exactly constitutes a weak stimulus when considered in terms of the resulting neuronal firing rates.

We simulate the same E and E/I clustered networks as before, where we stimulate 10% of clusters to mimic the evoked activity. The stimulation consists of a step increase in external current injected into the units making up the chosen clusters (Fig. 4a, top panel). Next, we choose $1000\,\mathrm{ms}$ spiking activity during spontaneous (no stimulation) and evoked (with stimulation) states and calculate the changes in firing rate (Δrate) and Fano factor (ΔFF). Positive Δ means an increase in rate or FF from the spontaneous to the evoked state while negative Δ indicates a decrease in these measures.

Fig. 4b shows the firing rate changes (in spikes/s) caused by different stimulus amplitudes (in pA). Different colors indicate Δrate versus stimulus amplitude for units belonging to stimulated clusters and non-stimulated clusters in both E and E/I clustered networks. Units belonging to the stimulated clusters in E clustered networks increase their firing rate much faster than those of the E/I clustered networks. The non-stimulated clusters only show a small decrease in their average firing rate during stimulation. In stimulated clusters, as shown in Fig. 4c, the trial-to-trial variability in the E clustered network increases (positive Δ) drastically for stimulus amplitude $< 0.4\,\mathrm{pA}$. A stimulus amplitude of $0.4\,\mathrm{pA}$ gives a Δrate of 40 spikes/s (vertical dashed line in Fig. 4b). This means that the E clustered network fails to reproduce the reduction in trial-to-trial variability for weak stimuli that evoke firing rate increases $\lesssim 40$ spikes/s. For both E clustered and E/I clustered networks, trial-to-trial variability decreases when increasing the stimulus amplitude (see Fig. 4c and d). The range of changes matches that observed in experimental data (cf. Fig. 1a). Similar behavior can be seen for the non-stimulated clusters (orange curve in Fig. 4c; purple curve in Fig. 4d: While the E clustered network sometimes increases its variability for weak stimulation, the variability in the E/I clustered network consistently decreases. In summary, the E/I clustered network robustly captures the reduction in trial-to-trial variability during stimulation, while E clustered networks fail to reproduce the reduction in trial-to-trial variability during stimulation for weak stimuli.

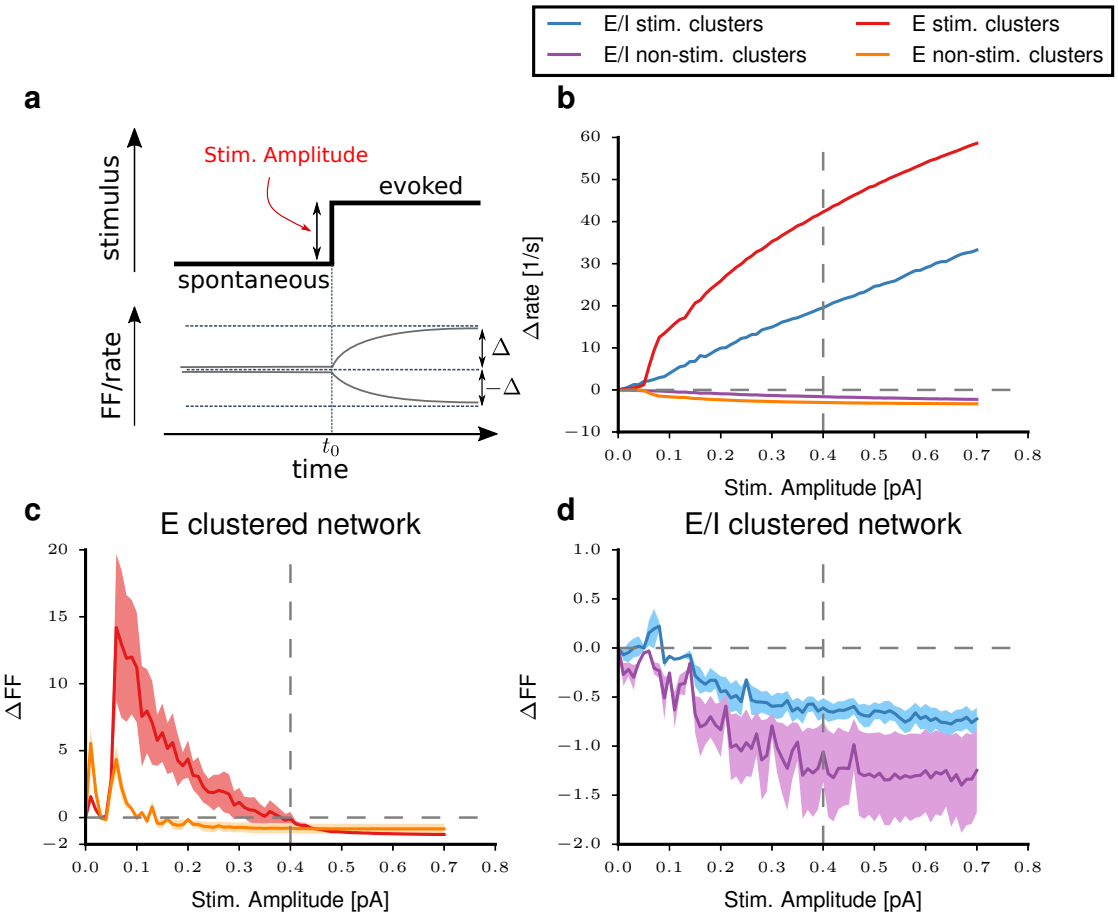

Figure 4: Stimulus dependency of firing rate and FF. **a)** Cartoon depicting the transition from spontaneous to evoked activity. Top panel shows the step change in the external input at time t_0, mimicking the stimulus. Bottom panel shows the changes in the measured FF or rate from spontaneous to evoked activity quantified by Δ. **b)** Difference in the firing rate between spontaneous and evoked states Δrate versus stimulus amplitude. Δrate is defined as the firing rate in the evoked minus that in the spontaneous state, i.e. $\Delta\text{rate} = \text{rate}_{\text{evoked}} - \text{rate}_{\text{spontaneous}}$, and is computed from $1000\,\text{ms}$ spiking data averaged over 100 trials. The network parameters are as before (see section "*Spiking network model*"). Stimulation is applied to 10% of the clusters ($Q = 50$), leading to 5 stimulated clusters and 45 non-stimulated ones. Different colors indicate the behaviors of stimulated and non-stimulated clusters in the E and E/I clustered networks (see the legend). **c)** ΔFF ($= FF_{\text{evoked}} - FF_{\text{spontaneous}}$) versus stimulus amplitude averaged over units in stimulated clusters of E clustered networks. **d)** The same for E/I clustered network. Shaded areas show the standard error calculated over the given clusters (stimulated or non-stimulated) for each case.

2.3 Functional E/I clustered network as a model of monkey motor cortex

In the previous sections, by comparing E and E/I clustered networks we have shown that including inhibitory clustering makes the activity of the clustered networks more realistic and robust. In this section, we further analyze the variability and information content of the motor cortical spike trains during the delay period in our experimental data set. Then, we examine if the E/I clustered networks can reproduce those observations.

2.3.1 Experimental paradigm and data analysis

The experiment (conducted in Alexa Riehle's lab at the CNRS Marseille, France) involved three conditions in which varying amounts of target information were given at the beginning of the trials. The monkey was seated in front of a panel featuring a hexagonal array of touch-sensitive LEDs. The monkey initiated the trials by touching a seventh LED at the center. Five hundred ms after trial start (TS), a preparatory signal (PS) was given. The PS consisted of either one,

two or three of the LEDs lighting up in green depending on the condition, as illustrated in Fig. 5a. After a delay of $1000\,\mathrm{ms}$ a single LED lit up in red representing the response signal (RS) which informed the subject where to move his hand. The times of movement onset (MO) and movement end (ME) were recorded and if the monkey touched the correct LED, the trial was registered as successful and a drop of juice was given as a reward. Only successful trials were used for the present study.

The conditions of one, two or three possible targets presented at PS were executed in blocks. In each block, 150 trials with randomized target directions were carried out so that each of the directions appeared on average 25 times per condition. Note that in order to obtain the same number of possible trial types in all conditions, not all possible combinations of directions for the preparatory stimulus were used in conditions 2 and 3. Since six combinations are possible for condition one, only the pairs 1-2, 3-4 and 5-6 were used in condition 2 and for condition 3, only two cases occurred (6-1-2, 3-4-5).

Extracellular recordings were obtained with a multielectrode microdrive (Reitböck system; Thomas Recording) to insert transdurally seven independently movable electrodes in motor (M1 close to the central sulcus (Bastian et al., 2003)). Online spike sorting resulted in up to seven simultaneously recorded single-unit spike trains. On each recording day, all three conditions were measured so that the responses of individual neurons can be compared across conditions. Some of the available units were excluded from the analyses because of suspected artifacts.

To allow a fair comparison of variability statistics across conditions, additional precautions were taken. Fano factors were computed for each unit and direction separately and we required that units had at least ten spikes in the $2\,\mathrm{s}$ interval after trial start and that at least ten trials were recorded per direction. To enable the comparison across conditions, we only included units and directions where those criteria were met for all conditions.

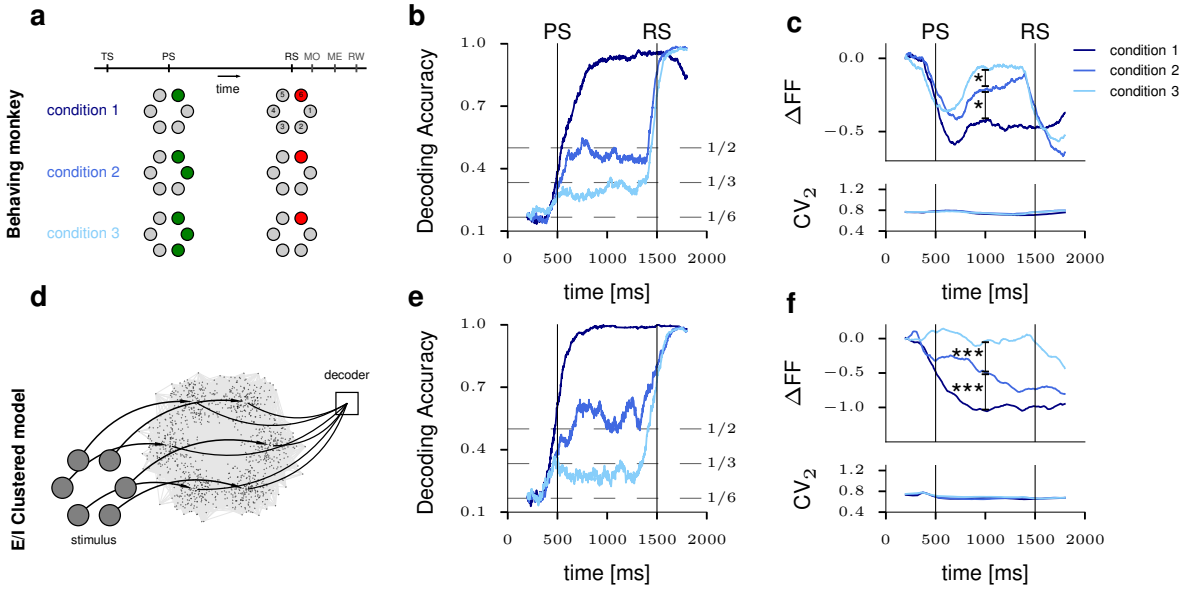

Figure 5: Functional E/I clustered model captures context dependency of information encoding and variability dynamics. **a)** Experimental protocol of the delayed center-out reach task. Five hundred ms after trial start (TS) a preparatory stimulus (PS) is presented, indicating either one, two or three adjacent targets, depending on the condition. The response signal (RS) follows a fixed delay of $1\,\mathrm{s}$ after which the monkey is required to execute a movement to the indicated target. Movement onset (MO), movement end (ME) and the time of reward (RW) depend on behavior and are not fixed in time. **b)** Decoding accuracy of direction classification using a pseudo-population constructed from all available neurons in windows of $400\,\mathrm{ms}$. **c)** Top, relative Fano factors to the Fano factor at TS; bottom, CV_2, computed in sliding windows of $400\,\mathrm{ms}$. $*$ in top panel indicates significance from Wilcoxon signed rank test (conditions 1-2: $p = 0.007$, conditions 2-3: $p = 0.03$). **d)** Architecture of the E/I clustered model where each cluster represents one direction, similar to the experiment, with the parameters given in Table 1. Each cluster is assigned a direction for which it receives its external input. The decoder integrates the average cluster rates and generates a decision. **e, f)** Similar to b-c but for the spiking network model.

Parameter	Unit	Value
N	-	$1200(E), 300(I)$
E_L	mV	0
V_{th}	mV	20
V_R	mV	0
C_m	pF	1
τ_m	ms	$20(E), 10(I)$
τ_{syn}	ms	$3(E), 2(I)$
τ_r	ms	5
p_{EE}	-	0.2
p_{EI}, p_{IE}, p_{II}	-	0.5
g	-	1.2
J_{EE}	pA	0.60
J_{EI}	pA	-1.60
J_{IE}	pA	0.46
J_{II}	pA	-2.44
I_x	pA	$1.25I_{th}(E), 0.78I_{th}(I)$
I_{stim}	pA	0.15
Q	-	6
J_{E+}	-	3.3
R_J	-	$3/4$

Table 1: Parameters used in the spiking network model

To assess how much directional information is contained in the population activity we reproduced the approach of (Rickert et al., 2009) and constructed pseudo-populations of all available units. This was achieved by combining spike counts from randomly chosen trials for the same directions from all units and treating them as if they were simultaneously recorded. Then we computed the classification performance of a logistic regression classifier using spike counts in windows of $400\,\mathrm{ms}$ width. At each point in time the performance was calculated using five-fold cross-validation as follows: the data set was divided into five groups and each group in turn served as the test set while the model was trained on the remaining groups, and the average score of the five models was calculated. We applied the five-fold cross validation to ten randomly assembled populations and computed the decoding accuracy of the predictions to assess classifier performance. Decoding accuracy is simply the fraction of correctly predicted samples averaged over classes.

$$\text{Decoding accuracy} = \frac{1}{C}\sum_{c=1}^{C}\frac{N^c_{correct}}{N^c_{total}} \tag{1}$$

All classification was performed using the `scikit-learn` package (Pedregosa et al., 2011).

The resulting decoding accuracies for the three conditions are shown in Fig. 5b. Using these pseudo-populations, the decoding accuracies approach those theoretically possible given the available information: in condition 1 where full information is available, the classification accuracy is close to unity, when two possible targets are given, the accuracy is close to $1/2$, and in condition 3 it approaches $1/3$. At RS the performance approaches unity for all conditions.

Fig. 5c, top, shows the average relative Fano factor to the Fano factor at TS (ΔFF) across all units for the three conditions. For all conditions, the count variability is *quenched* at the presentation of the PS. The Fano Factor (FF) then recovers to different levels depending on the condition. When the complete information is given at the PS, i.e. condition 1, the FF remains at a low value, while for condition 2 it is higher, and it is still higher when three possible targets are presented. This difference is significant according to a Wilcoxon signed rank test throughout most of the preparatory period (conditions 1-2: $p = 0.007$, conditions 2-3: $p = 0.03$ at $1000\,\mathrm{ms}$ after TS). After the response signal, the FF drops for conditions 2 and 3 but remains more or less constant for condition 1. This may be in part due to the fact that the movement onset times for condition 1 were more broadly distributed than for the more difficult conditions 2 and 3 where the monkey had to wait for the RS before it could form a decision. Hence for conditions 2 and 3 the trials are closer to being aligned to MO (see Fig. 6c). As for the spiking irregularity (CV_2), shown in Fig. 5c, bottom, all 3 conditions display similar behavior, i.e. CV_2 remains constant at a value around 0.8 during the whole experiment.

2.3.2 An Attractor Model of Movement Preparation

Here we propose a mechanistic model employing the E/I clustered network that is based on winnerless competition and can explain the differences in count variance and information content of single units in the preparatory period as

well as the difference in reaction times between the three conditions. Movement preparation is often associated with bringing the firing rates in motor cortex to a certain *optimal subspace* (Churchland et al., 2006), or *initial condition* (Afshar et al., 2011) suitable for the movement to be executed. Alternatively, the neuronal *activation* associated with a movement is brought *close to threshold*. Both types of models assume that a certain combination of firing rates needs to be achieved before a movement can be executed and that this takes some time. Our findings that average firing rates are similar while movement direction can still be decoded optimally from populations during the preparatory period (Fig. 5b) is consistent with this view. Achieving a suitable *position* in some high-dimensional firing rate space prior to movement execution could explain the reduction in FF during the preparatory period. Naturally, if the exact movement direction is known (condition 1), a more *optimal* initial condition for the movement can be prepared than when the target is ambiguous (conditions 2, 3), and shorter reaction times can be achieved.

The aim of our model is to show how winnerless competition can lead to the observed temporal evolution of condition-dependent count variability and information content. We therefore simplify its construction as much as possible. The architecture is schematically shown in Fig. 5d.

The core of the model consists of a balanced network of integrate-and-fire units with $Q = 6$ clusters — one for each direction. Each cluster comprises 200 excitatory units and is assigned a corresponding inhibitory cluster with 50 units. The clustering parameters J_{E+} are adjusted so that transitions between clusters occurred and the Fano factors in the spontaneous state resembled those measured in the physiological data. Due to the small size of the network, the strength of the external input currents had to be adjusted to obtain firing rates of approximately 3 and 5 spikes per second for the excitatory and inhibitory populations, respectively. No parameter tuning was otherwise performed. The parameters used are summarized in Table 1.

Real neurons display various types of tuning curves which can also change over time (Bastian et al., 2003; Rickert et al., 2009). We do not specifically model tuning or how it arises. We simply stimulate one cluster per direction by applying a constant current of $0.15\ \mathrm{pA}$ to all its units. This means that each unit in our model is sharply selective to a single direction. The same stimulation protocol as described above for the experiment was applied to our model. For each condition 150 trials were executed. In each trial one of the six trial types was randomly chosen. The preparatory signal was applied to the model by direct stimulation of 1, 2, or 3 clusters depending on the condition. The amplitude of the stimulation current was identical for all conditions. After $1000\,\mathrm{ms}$, i.e. at the end of the preparatory period, the PS currents were turned off except for the one corresponding to the correct target direction. The target current was applied for an additional $400\,\mathrm{ms}$ to mimic the response signal in the experiment. Trials were cut from one long continuous simulation. To allow the network to relax to its spontaneous state, inter-trial periods randomly varied between 1500 and $1700\,\mathrm{ms}$. The variance in this relaxation period was intended to avoid any effects of periodicity.

Inspired by previous models of perceptual decision making (Gold & Shadlen, 2007; Meckenhäuser et al., 2014), the decoder (depicted in Fig. 5d) has a decision variable associated with each cluster. For each direction d, a leaky integrator governed by an equation of the form

$$\frac{dI_d(t)}{dt} = -\frac{I_d(t)}{\tau_I} + C_d \tag{2}$$

sums up the instantaneous spike counts C_d (in $1\,\mathrm{ms}$ bins) of the corresponding population and forgets with time constant τ_I. Due to the random nature of the balanced network, the firing rates of the individual clusters may vary considerably. To enable a decision by applying a single threshold θ to all directions, a decision variable $DV_d(t)$ is formed for all directions through normalizing by the sum over all directions:

$$DV_d(t) = \frac{I_d(t)}{\sum_{j=1}^{6} I_j(t)} \tag{3}$$

This is similar to multi-class classification, and the decision variable $DV_d(t)$ expresses the probability that direction d is the correct choice at time t. A decision was reached when one of the decision variables crossed the threshold in the interval where the RS was presented ($1000\,\mathrm{ms} < t <= 1400\,\mathrm{ms}$). Threshold crossings after that period were counted as unsuccessful trials. If a decision variable was already above threshold at the beginning of RS, the decision was counted. The threshold θ was adjusted so that the performance of the model was maximized. The time constant of integration τ_I was set to $50\,\mathrm{ms}$ which represents an intermediate value between very fast reactions directly when the threshold is reached at RS and very slow integration where the threshold was not reached during the RS-interval.

Fig. 5e-f shows the same analyses as shown for the physiological data (Fig. 5b-c) computed on the spike trains generated by the model. To allow a fair comparison the same selection conditions were applied as for the experimental data. Since the number of trials was the same for all units, only the minimum count rate criterion had an effect (10 spikes in the inspected time window).

The probabilistic nature of switching between prepared directions does produce average decoding accuracies for single neurons which resemble those seen in the data (Fig. 5b). By following the same classification as in Fig. 5b, performed

on the entire population of analyzed units simultaneously, we recover the optimal decoding scores of 1, $1/2$ and $1/3$ for the three conditions, respectively (Fig. 5e).

Besides showing similar decoding accuracy as in the experimental data, the model captures the condition-dependent modulation of the average relative Fano factor (ΔFF) and spike time irregularity (CV_2) (compare Fig. 5c and D). In the spontaneous state before the PS is turned on, the count variance is high in all conditions. During the preparatory period the Fano factors are differentially quenched for the three target conditions, as in the physiological data. The differences are also highly significant with respect to a Wilcoxon signed rank test (conditions 1-2: $p = 3.6 \times 10^{-13}$, conditions 2-3: $p = 3.7 \times 10^{-30}$). At the RS onset, a single stimulus is applied in all conditions and the ΔFF converge to the same value. The Fano factors for condition 2 and 3 exhibit a second strong decrease after RS in the physiological data. This is again not captured by our model. A possible reason is that the distributions of reaction times are much narrower for conditions 2 and 3 (see Fig. 6c).

Finally, the operation of the decoder is illustrated in Fig. 6. The threshold is adjusted so that the balanced accuracy over all sessions is maximized. Fig. 6a shows a raster plot of the excitatory activity for a sample trial of condition 3. The PS current is applied to clusters 2, 3 and 4 (counting from the bottom). After stimulation, competition arises between these three clusters. Since three clusters now have to *share* the higher probability of being active, the number of occurring rate configurations increases and thereby the reduction in rate variance is weaker. The amount of rate variance in this model increases with the number of stimulated clusters. If all clusters are stimulated at the same time, their chances during winnerless competition are again equal and the rate variance level of the spontaneous state is recovered. Hence for condition 3 the variance is again higher than for condition 2.

The respective decision variables $DV(t)$ are superimposed on each direction. Due to competition between populations, all decision variables remain below the threshold in this case as a longer period of being in the active state is required due to the time scale of integration. When the RS is given to the target cluster 3, the *wrong* cluster is currently in the active state. It takes a short moment until the stimulus takes effect an the network activity switches to the target cluster. An even longer time is taken for the decision variable to reach the threshold. In this particular trial a *reaction time* of $\sim 400\,\mathrm{ms}$ is thus produced.

Fig. 6b-c show the distributions of reaction times produced by the model and experimental data, respectively, for each condition. It can be seen that as in the experimental data the average reaction times in condition 1 were much quicker than in conditions 2 and 3. In contrast to the experiment, anticipated responses were not penalized in the model. If the decision variable of the correct direction was already above threshold at RS, the trial was counted as successful. In condition 1 this was frequently the case. The shape of the reaction time histogram for condition 1 (dark blue in Fig. 6c) suggests that the monkey displayed a similar behaviour. The difference in response times between conditions 2 and 3 is not discernible in the model as in the data. The chance of having *prepared* for the wrong direction in the model hence explains the difference in reaction times between the full information and the ambiguous conditions.

Our model presents a mechanistic explanation for the models proposed by (Churchland et al., 2006) and (Afshar et al., 2011). To our knowledge, it is the first model that links attractor dynamics to context dependent variability modulation in motor cortex and behavioural variability in the form of reaction times.

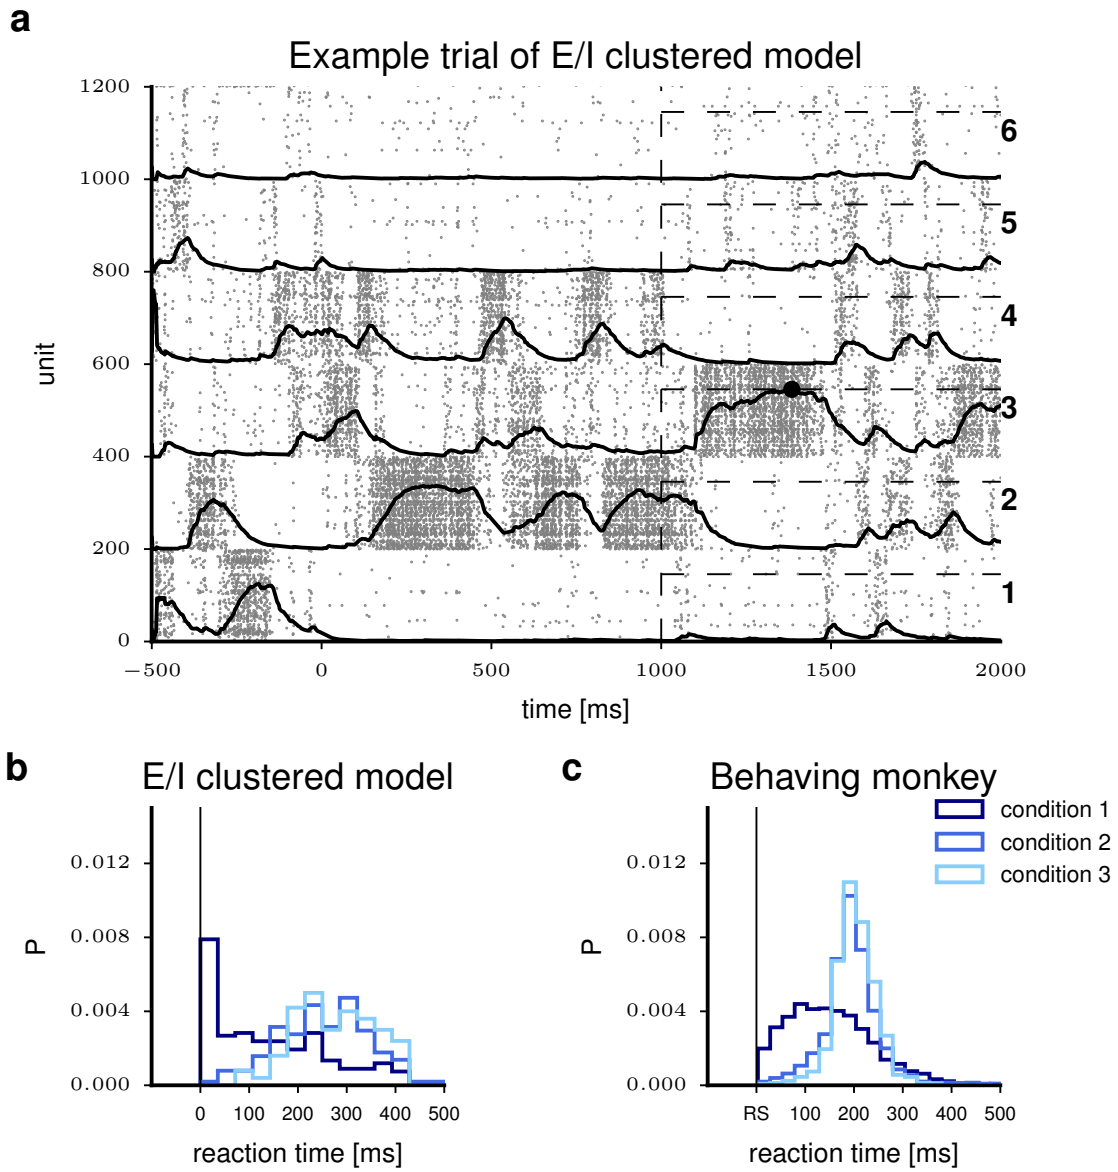

Figure 6: E/I decision model generates task dependent reaction times. **a**) Raster plot of excitatory cluster activity for an example trial of condition 3. Solid lines represent the associated decision variable (DV) for each cluster. At $t = 1000$ ms the decoder starts integrating as indicated by the vertical dashed line. Horizontal dashed lines show the level of the decision threshold θ. To the right, direction numbers are indicated. During the preparatory phase, clusters 2, 3 and 4 are stimulated. Target is direction 3. The black dot indicates a successful threshold crossing. **b-c**) Histograms of *reaction times* of the model (b) and the experimental data (c) for the three conditions.

3 Discussion

Conclusion We have proposed a robust mechanistic model of cortical and behavioral variability dynamics and showed that it accounts for task-related dynamics and coding in motor cortex. The model extends previous spiking attractor network models featuring excitatory clustering (Amit & Brunel, 1997; Litwin-Kumar & Doiron, 2012; Deco & Hugues, 2012; Mazzucato et al., 2015, 2019; La Camera et al., 2019) by adding clustered inhibitory connectivity

motivated by recent experimental (Xue et al., 2014; Lee et al., 2014; Arkhipov et al., 2018; Khan et al., 2018; Shin et al., 2019; Najafi et al., 2020) and theoretical (Rost et al., 2017; Najafi et al., 2020) studies. We showed that the clustered inhibition ensures the local balance of excitation and inhibition within each attractor, and increases the robustness of winnerless competition, able to capture cortical dynamics during decision-making. The model replicates both neural and behavioral variability of monkey performing a delayed reach task with various degrees of prior information about the target direction. At the neuronal level, we showed that our model reproduces the spiking irregularity and trial-to-trial variability of spike counts, and at the behavioral level, our model captures the reaction times during the task. These findings suggest joint excitatory and inhibitory clustering as a powerful mechanism for decision-related activity in cortex.

Joint excitatory and inhibitory clustering vs. purely excitatory clustering Previous studies have considered purely excitatory clustering with a 'blanket of inhibition' to explain decision-related activity in cortex and the associated cortical variability dynamics (Amit & Brunel, 1997; Litwin-Kumar & Doiron, 2012; Deco & Hugues, 2012; Mazzucato et al., 2015, 2019). We here showed that joint excitatory and inhibitory clustering better accounts for cortical variability dynamics, thereby more plausibly reflecting the neural mechanisms underlying decision-making. Furthermore, recent evidence in the literature supports the notion that not only excitatory neurons but also inhibitory neurons in cortex are clustered (Xue et al., 2014; Lee et al., 2014; Morishima et al., 2017; Arkhipov et al., 2018; Khan et al., 2018; Znamenskiy et al., 2018; Shin et al., 2019; Najafi et al., 2020).

Cortical neuronal activity is variable on both the single-trial level, in terms of the irregularity of interspike intervals, and across trials, measured for instance by the Fano factor. Our experimental data show that spike timing irregularity is comparable during spontaneous and evoked activity, while trial-to-trial variability is lower for evoked activity. Switching between clusters in a process of winnerless competition is able to account for the high Fano factor during spontaneous activity, while the variability is quenched during evoked, decision-related activity by selection of a single cluster or subset of clusters.

Spiking attractor network models with clustering only among excitatory neurons have been successful in explaining the trial-to-trial variability observed in in vivo recordings. However, this type of model has unrealistic spiking irregularity and excessive spike rates during cluster activation and exhibits biologically plausible multistability (winnerless competition) only in a narrow parameter regime (see Fig. 2 and discussion of (Litwin-Kumar & Doiron, 2012; Deco & Hugues, 2012); also (Rost et al., 2017)).

By analyzing the stable fixed points of the mean-field equations for networks with excitatory clusters, Rost et al. 2017 have shown that switching is hampered by the high rates attained in active clusters. Our proposed remedy of assigning an inhibitory population to each cluster, by increasing the corresponding $E \to I$, $I \to E$ and $I \to I$ synaptic strengths, solves both of the above problems. Since inhibition is now also selective, each excitatory population is held in check by its inhibitory counterpart and hence the fixed points of the active and passive clusters move closer together. This facilitates switching, and also helps to maintain interval variability because the active clusters remain in the fluctuation-driven balanced state rather than being forced into saturation (see Fig. 3). Thus, inhibitory clustering ensures realistic spike timing variability and greatly increases the robustness of the multistability, extending the range of cluster strengths and sizes over which winnerless competition occurs.

Functional role The balanced random network model with joint excitatory and inhibitory clustering not only accounts for realistic cortical variability dynamics but also reproduces functional and behavioral aspects of movement-preperation and decision-making: task-related encoding, variability, and reaction times match the experimental observations. In our experiment when the monkey has incomplete information during the preparatory period it can only resolve the ambiguity of multiple targets when the GO signal provides full information and the reaction times are increased. We observe the same phenomena in our model. The mechanistic explanation is that in the 2 target condition two clusters compete and the activity switches between them during the preparation period. When the GO signal resolves the ambiguity, only one of the clusters retains input stimulation. If this cluster has been active at this point in time, activity level reaches the threshold faster (short reaction times) than in the case where a switch is required, leading to longer reaction times (Fig. 6). The same mechanistic explanation (is likely to) underlies short and long reaction times in a recent attractor network model to describe behavioral reaction times in anticipatory versus unexpected cues depending on pre-stimulus cluster activation ((Mazzucato et al., 2019)). In our model we find a small but significant increase in reaction times from the 2 target to 3 target condition. The same effect was observed in the monkey's reaction time. Our model makes a mechanistic prediction for this effect where, in the case of three competing clusters, the chance is lower that the one cluster associated with the final goal had been active at time GO.

Multistability and timescales of variability dynamics The timescales on which neural activity varies constitute a recurring theme in the context of cortical variability. The present work is centered around the hypothesis that rate

variance across trials is a consequence of slow switching between clusters of neurons. In our model, we have observed that switching tends to occur less frequently when the clustering is stronger, thereby increasing the time that assemblies spend in the active state. A number of physiological results support the hypothesis that spontaneous activity is made up of sequences of structured activity patterns that emerge seemingly at random. The average or typical amount of time spent in each pattern is still a matter of investigation.

Kenet et al. 2003 used voltage-sensitive dye imaging in the visual cortex of anesthetized cats to show that spontaneous activity appears to cycle through patterns very similar to those evoked by visual stimuli and reported that transitions between patterns occurred on a timescale of around $80\,\mathrm{ms}$. Berkes et al. 2011 recorded extracellular activity using multi-electrode arrays in the visual cortex of developing ferrets. They found that over the course of development, the spontaneous activity became increasingly similar to evoked responses. The study revealed strong correlations between activity patterns for delays up to $20\,\mathrm{ms}$. Ponce-Alvarez et al. 2012 fitted a hidden Markov model (HMM) to simultaneously recorded spike trains from somatosensory and motor cortices in monkeys. During the delay period in a perceptual decision task, they observed switching through sequences of states with transition times on the order of $50\,\mathrm{ms}$. Mazzucato et al. 2015 applied a similar analysis to spontaneous recordings from the gustatory cortex in rats. Their HMM analysis yielded exponentially distributed state durations with a mean of $\sim 700\,\mathrm{ms}$. Luczak et al. 2009 examined evoked and spontaneous activity in the auditory and somatosensory cortices of rats and found not only spatial patterns but temporally stereotyped profiles with durations of $\sim 100\,\mathrm{ms}$.

These results point to timescales on the order of tens to hundreds of milliseconds. Teich et al. 1997 on the other hand have found that Fano factors of retinal ganglion cells in cats increase with counting window width for observation intervals of several minutes. For stationary point processes, the FF depends only weakly on window size (Nawrot et al., 2008). When rate variance is added to the equation, the Fano factor will increase with counting window size until the window spans several periods of the periodicity of the underlying fluctuations. Such long timescale fluctuations do not necessarily have to originate from the cluster switching mechanisms described here. It is likely that winnerless competition dynamics in spontaneous cortical firing will happen on timescales related to the stimulus modalities or movements a particular region codes for.

The time scales of switching in our model depend on the setting of the cluster parameter J_{E+}. We have adjusted this parameter so that the networks' spontaneous activities would exhibit Fano factors similar to those observed in our data set. We have not measured the durations that attractors spend in the active states, although an HMM analysis similar to those by Ponce-Alvarez et al. 2012 and Mazzucato et al. 2015 could yield estimates. Inspection of the raster plots does however suggest that the up states in our model have similar time scales as those reported in the literature ($\sim 20 - 200\,\mathrm{ms}$).

It would in principle be desirable to have a theoretical prediction for the time scales of cluster switching in relation to the model parameters. Lagzi and Rotter (Lagzi & Rotter, 2015) have described the winnerless competition between two populations using a rate model governed by Lokte-Volterra type equations. If noise is introduced, switching can also occur in rate models. They found that the survival times of the active states could be well approximated by an exponential distribution and that the average time between switches grows faster than exponentially with cluster strength. Rost et al. 2017 used the mean-field description of our network configuration to find the stable rate configurations. Switching between these stable attractors is a finite size effect due to chaotic fluctuations in the firing of individual units in the populations which is by definition not captured by the mean field approach. The mean field theory for networks of binary units also predicts the distribution of activity rates within populations (van Vreeswijk & Sompolinsky, 1998). From this, it is theoretically possible to compute the 2Q-dimensional joint distribution of cluster rates. Making some assumptions about the noise caused by the rate fluctuations within clusters may then be possible to make predictions about the switching dynamics. This analysis was however beyond the scope of the current study.

Outlook In our study, we varied cluster size and found more robust winnerless competition for smaller clusters, below about 200 neurons per cluster. Switching of activity states is triggered by fluctuations in cluster firing rate. Hence the probability of switching decreases when clusters become larger, because fluctuations tend to average out (Doiron & Litwin-Kumar, 2014). While fine-scale clusters have been reported to consist of tens rather than thousands of neurons (Perin et al., 2011), it is unclear whether the cortical algorithm in reality already breaks down with clusters of a few hundred units. Future investigations into this issue can include anatomical and physiological estimates of cluster size, as well as computational modeling attempting to increase the robustness of winnerless competition for large clusters.

Our experimental data is recorded after an extensive training period so that the monkey performs the task with a high proportion of correct trials. We, therefore, assume in our model that the connectivity has reached a fixed structure and is no longer plastic. Future work may investigate how the clustered connectivity is learned during training, for instance through spike-timing-dependent plasticity (STDP) combined with selective stimulation (Ocker et al., 2015; Zenke et al., 2015; Litwin-Kumar & Doiron, 2014; Wu et al., 2019). To form and recall these clusters in a stable manner over a long

time some form of homeostatic mechanism is crucial. Zenke et al. 2015 shows that multiple time scales of homeostatic regulation are necessary to form robust and stable clusters that are functionally relevant. Litwin-Kumar & Doiron 2014 investigate homeostatic mechanisms that act on I→E synapses together with E→E STDP rules to form clusters that reflect previously experienced stimuli. The inhibitory plasticity in these studies is globally modulated while excitatory neurons form local clusters and are responsible for functional representations. It will be interesting to investigate how joint excitatory and inhibitory clustering can stably emerge in neural networks through plasticity, and what role is played by homeostatic mechanisms in this context.

4 Materials and Methods

4.1 Spiking network model

Our spiking network model is composed of linear units with exponential synaptic currents where the sub-threshold evolution of the membrane potential V is described by the differential equation

$$\frac{dV}{dt} = \frac{-(V - E_L)}{\tau_m} + \frac{I_{syn} + I_x}{C_m}. \tag{4}$$

In the absence of input, the membrane potential decays exponentially to the resting potential E_L with time constant τ_m. The current I_{syn} represents the synaptic input, I_x is an externally injected current and C_m is the membrane capacitance. If the potential reaches the threshold V_{th} a spike is emitted and V is clamped to a reset voltage V_r for an absolute refractory period τ_r. The synaptic current to a neuron i evolves according to the equations

$$\tau_{syn}\frac{dI^i_{syn}}{dt} = -I^i_{syn} + \sum_j J_{ij} \sum_k \delta\left(t - t^j_k\right) \tag{5}$$

where t^j_k is the time of the arrival of the kth spike from presynaptic neuron j and δ is the Dirac delta function.

To allow comparison to other clustering results, we use parameters similar to those given in (Litwin-Kumar & Doiron, 2012; Mazzucato et al., 2015) (see table 2). For clarity, we will briefly explain how the parameters used here were derived in the following.

Parameter	Unit	Value
N	-	$4000(E), 1000(I)$
E_L	mV	0
V_{th}	mV	20
V_R	mV	0
C_m	pF	1
τ_m	ms	$20(E), 10(I)$
τ_{syn}	ms	$3(E), 2(I)$
τ_r	ms	5
p_{EE}	-	0.2
p_{EI}, p_{IE}, p_{II}	-	0.5
g	-	1.2
J_{EE}	pA	0.33
J_{EI}	pA	-0.89
J_{IE}	pA	0.25
J_{II}	pA	-1.34
I_x	pA	$2.13I_{th}(E), 1.24I_{th}(I)$

Table 2: Summary of parameters used in the spiking network simulations

Calibration of the Balanced State We follow the same approach as in (Rost et al., 2017) for the binary networks by requiring that $\sqrt{K}$ excitatory action potentials arriving within a short time suffice to drive the membrane potential form E_L to V_{th} and hence elicit a spike. For that purpose we need to compute the deflection in the membrane potential caused by a presynaptic spike.

According to equation 5, a spike arriving at $t = 0$ leads to a postsynaptic current of the form

$$I_{psc}(t) = Je^{-t/\tau_{syn}}\,\Theta\,(t) \tag{6}$$

where J and Θ are the synaptic efficacy and step function, respectively. Inserting this into equation 4 and integrating with $V = 0$ at $t = 0$ the postsynaptic potential is obtained:

$$PSP(t) = J\frac{\tau_m\tau_{syn}}{\tau_m - \tau_{syn}}\left(e^{-t/\tau_m} - e^{-t/\tau_{syn}}\right)\Theta\,(t) \tag{7}$$

The maximal deflection of the PSP, PSP_{max} occurs at $t = \frac{log\frac{\tau_{syn}}{\tau_m}}{(1/\tau_m - 1/\tau_{syn})}$.

Note that the PSP amplitude depends on the synaptic as well as the membrane time constants and is therefore different for each synapse type (PSP_{max}^{EE}, PSP_{max}^{EI},...). The scale-free weights are then constructed in the same way as for the binary networks (equations 3 to 8 in (Rost et al., 2017)) but weighted by the respective PSP amplitudes:

$$j_{EE} = \frac{V_{th} - E_L}{\sqrt{p_{EE} n_E}} \frac{1}{PSP_{max}^{EE}} \tag{8}$$

$$j_{EI} = -g j_{EE} \frac{p_{EE} n_E}{p_{EI} n_I} \frac{PSP_{max}^{EE}}{PSP_{max}^{EI}} \tag{9}$$

$$j_{IE} = \frac{V_{th} - E_L}{\sqrt{p_{IE} n_E}} \frac{1}{PSP_{max}^{IE}} \tag{10}$$

$$j_{II} = -j_{IE} \frac{p_{IE} n_E}{p_{II} n_I} \frac{PSP_{max}^{IE}}{PSP_{max}^{II}} \tag{11}$$

where g is the relative strength of inhibition. The final weights $J_{\alpha\beta}$ are obtained by dividing by $\sqrt{N}$.

Since we consider variability dynamics the external inputs are modeled as constant currents to ensure all variability arises deterministically inside the network rather than stemming from externally generated Poisson input. In analogy to the *threshold rate* of (Brunel, 2000) the external current I_x is expressed in terms of the current required to reach the threshold in the absence of synaptic input:

$$I_{th} = \frac{V_{th} - E_L}{\tau_m} C_m \tag{12}$$

A complex interplay exists between the E and I firing rates and the magnitude of the external currents to the populations. The tuning of the injected currents required to obtain the desired firing rates for 3 and 5 spikes per second for the E and I populations respectively was therefore achieved by modeling single units with Poissonian inputs mimicking the network input at the target firing rates. The external inputs could then be increased until the modeled units fired on average at the required rates.

Before introducing structured connectivity we first ensured that the network configuration was operating in the asynchronous-irregular (AI) regime. Irregularity was measured using the squared Coefficient of Variation (CV^2) as explained in section "*Quantifying neural variability*". Synchrony of measures such as the instantaneous firing rate or the membrane potential in neural networks can be quantified according to (Golomb & Hansel, 2000) as:

$$\chi = \sqrt{\frac{\sigma_{pop}^2}{\langle \sigma_i^2 \rangle}} \tag{13}$$

here σ_{pop}^2 is the variance of the the population average and $\langle \sigma_i^2 \rangle$ is the average over the individual units' variances. The measure gives unity for totally synchronized activity and for asynchronous activity in networks of size N, one expects $\chi \sim \mathcal{O}\left(\frac{1}{\sqrt{N}}\right)$. Since recording all membrane potentials in simulations is computationally expensive, we computed χ on spike counts measured in bins of $20ms$.

It can be seen in Fig. 7 that the networks show the usual characteristics of the balanced state. When excitation dominates, synchronous-regular firing near the saturation rate $1/\tau_r$ is observed. The AI state occurs when g is sufficiently large for inhibition to dominate (Fig. 7a). As in the binary network (Rost et al., 2017), we choose $g = 1.2$, where $\chi = 0.02 \sim 1/\sqrt{N}$ and $CV^2 = 0.73$ (Fig. 7b). The raster plot shows no discernible structure (Fig. 7c) and the average firing rate is low and constant over time (Fig. 7d). The synaptic currents from excitatory and inhibitory inputs and the external current Ix cancel so that the net input fluctuates around zero (Fig. 7e). Hence, the membrane potentials fluctuate at low values and only occasionally reach the threshold to produce a spike (Fig. 7f). The parameters used for all simulations in this chapter are summarized in Table 2.

It is worth noting that the synaptic delay had to be set equal to the simulation time step (0.1 ms) in order to avoid synchronous oscillations. This is not biologically realistic and usually synaptic delays between 0.5 and 2 ms are used in network simulations (e.g. (Kriener et al., 2014; Ostojic, 2014; Pernice et al., 2011; Morrison et al., 2007)). Synaptic delays are known to play a role in synchronous oscillations (Brunel, 2000; Bose & Kunec, 2001). We did not conclusively investigate the origin of these oscillations. It is, however, likely that they are due to the dense connectivity employed in our model, which we chose to enable comparisons with (Litwin-Kumar & Doiron, 2012) and (Mazzucato et al., 2015).

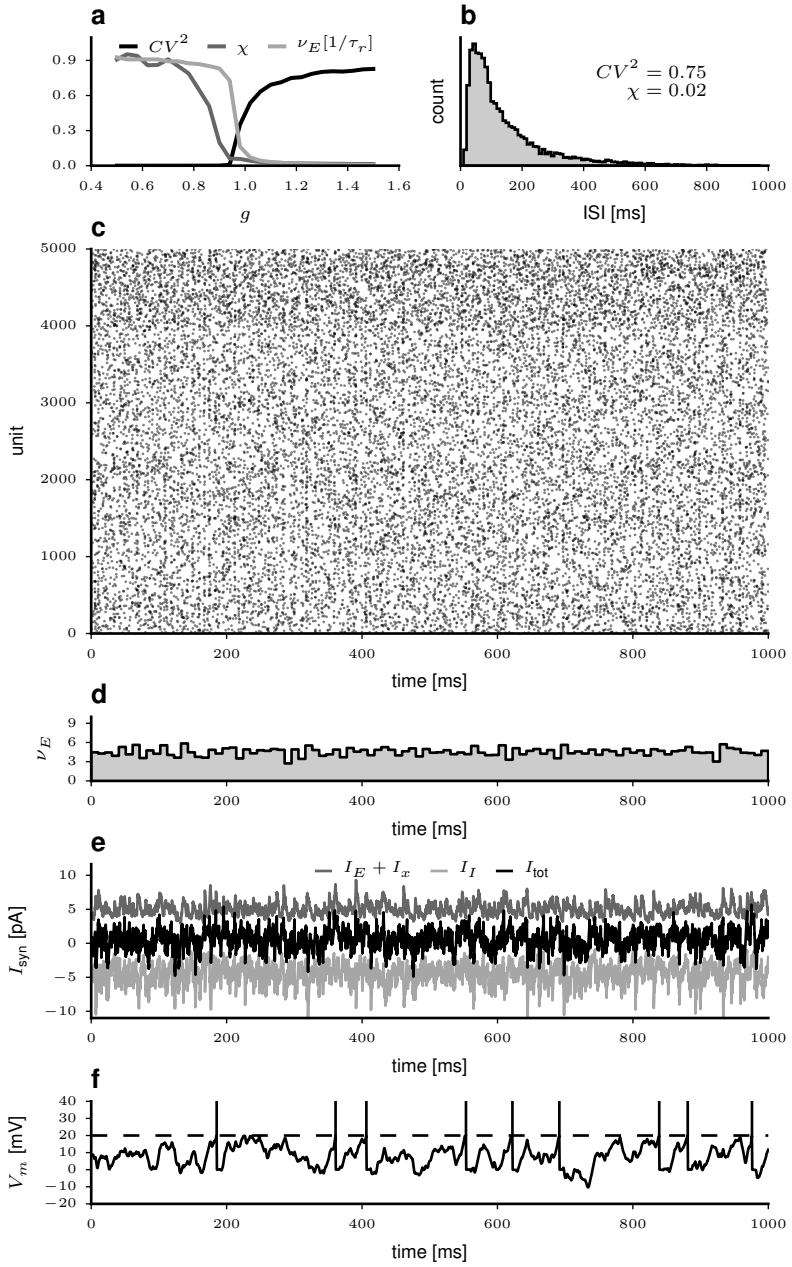

Supplemental Material, Figure 7: Spiking network in the balanced state. Parameters as in Table 2. **a**) Irregularity CV^2, synchrony χ and normalized excitatory firing rate versus relative inhibitory strength g. **b**) Pooled ISI distribution for the E population. **c**) Raster plot of one second of spiking activity for 4000 excitatory (from 0 to 4000) and 1000 (from 4000 to 5000) inhibitory neurons. **d**) E population rate histogram computed in $10\,\text{ms}$ bins. **e**) Synaptic currents of a randomly selected E unit. **f**) Membrane potential for same unit as in **e**. Vertical bars above the threshold (dashed line) represent action potentials.

E and E/I clustered networks We follow the same connectivity scheme that we introduced for binary networks in our previous work (see (Rost et al., 2017) for a detailed explanation). Briefly, for the E clustered networks, we first divide the excitatory population into Q equally sized clusters with uniform connection probability. Then, we potentiate the synaptic connection within each cluster by a factor J_+, which we refer to as cluster strength. $J_+ = 1$ represents the random balanced network and the larger J_+, the stronger the weights within the formed clusters. To maintain the

overall balance, we decrease the weights among units belonging to different clusters by a factor $J_- = \frac{Q-J_+}{Q-1}$. For the E/I clustered networks, we divide not only the excitatory population but also the inhibitory population into Q clusters. Then we require that each excitatory cluster selectively influences its corresponding inhibitory cluster and vice versa by increasing the corresponding *EI*, *IE* and *II* weights. We have shown in (Rost et al., 2017) that the inhibitory clustering needs to be weaker than the excitatory clustering to obtain realistic firing rates. Therefore we introduce separate excitatory and inhibitory clustering strengths, J_{E+} and J_{E-}. The relation between J_{E+} and J_{E-} is defined as follows:

$$J_{I+} = 1 + R_J(J_{E+} - 1),$$

where R_J is a proportionality factor, i.e. $R_J = 1$ implies the same cluster strength for inhibitory and excitatory clusters ($J_{E+} = J_{I+}$) and $R_J = 0$ makes the inhibitory population un-clustered ($J_{I+} = 1$, which represents the E clustered networks). Throughout the current study, we use $R_J = 3/4$ based on our previous results (Rost et al., 2017) where we showed that R_J can prevent firing rate saturation in up states.

4.2 Quantifying neural variability

The Fano factor (FF) quantifies the dispersion of spike counts for single neurons across repeated observations of the same experimental condition or trial.

$$FF = \frac{\sigma_c^2}{\mu_c}, \tag{14}$$

where σ_c^2 and μ_c are the variance and mean count over trials. The estimation of the FF is biased towards unity for small estimation windows. However, this bias quickly becomes negligible when the estimation window is several multiples of the mean inter spike interval (ISI) (Nawrot, 2010).

Interval statistics are usually characterized by the Coefficient of Variation (CV) of the ISI distribution,

$$CV^2 = \frac{\sigma_{ISI}^2}{\mu_{ISI}^2}. \tag{15}$$

Here, σ_{ISI}^2 and μ_{ISI} are the variance and mean of the intervals between action potentials. Estimating the CV^2 requires some caution, as modulations in firing rate increase the interval variability. Another problem with estimating the CV^2 follows from finite-size estimation windows. In an estimation window of width T, only $ISIs < T$ can be observed. If the underlying process has non-zero probabilities for larger intervals, the CV^2 will be under-estimated. This effect is known as *right-censoring* (Wiener, 2003; Nawrot et al., 2008). However, there are other measures of interval variability which to some degree overcome the problem of rate-change sensitivity by considering only pairs of consecutive intervals. CV_2 was developed to compare the irregularity of firing of visual neurons in the anaesthetised cat with that in vitro (Holt et al., 1996):

$$CV_2 = 2\left\langle \frac{|\tau - \tau'|}{\tau + \tau'} \right\rangle \tag{16}$$

Here, $\langle ... \rangle$ denotes averaging and τ and τ' are consecutive ISIs.

5 Acknowledgments

This work is supported by the German Science Foundation under the Institutional Strategy of the University of Cologne within the German Excellence Initiative (DFG-ZUK 81/1). We thank Dr. Moritz Deger for all fruitful discussions at the early phase of the project.

References

Afshar, A., Santhanam, G., Yu, B. M., Ryu, S. I., Sahani, M., & Shenoy, K. V. (2011). Single-Trial Neural Correlates of Arm Movement Preparation. *Neuron 71*(3), 555–564.

Amit, D., & Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex 7*(3), 237–252.

Anderson, J. S., Carandini, M., & Ferster, D. (2000). Orientation Tuning of Input Conductance, Excitation, and Inhibition in Cat Primary Visual Cortex. *Journal of Neurophysiology 84*(2), 909–926.

Arkhipov, A., Gouwens, N. W., Billeh, Y. N., Gratiy, S., Iyer, R., Wei, Z., Xu, Z., Abbasi-Asl, R., Berg, J., Buice, M., Cain, N., da Costa, N., de Vries, S., Denman, D., Durand, S., Feng, D., Jarsky, T., Lecoq, J., Lee, B., Li, L., Mihalas, S., Ocker, G. K., Olsen, S. R., Reid, R. C., Soler-Llavina, G., Sorensen, S. A., Wang, Q., Waters, J., Scanziani, M., & Koch, C. (2018). Visual physiology of the layer 4 cortical circuit in silico. *PLoS Computational Biology 14*(11), e1006535.

Bastian, A., Schoner, G., & Riehle, A. (2003). Preshaping and continuous evolution of motor cortical representations during movement preparation. *European Journal of Neuroscience 18*(7), 2047–2058.

Berkes, P., Orbán, G., Lengyel, M., & Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science 331*(6013), 83–87.

Bose, A., & Kunec, S. (2001). Synchrony and frequency regulation by synaptic delay in networks of self-inhibiting neurons. *Neurocomputing 38-40*, 505–513.

Brunel, N. (2000). Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons. *Journal of Computational Neuroscience 8*(3), 183–208.

Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., Newsome, W. T., Clark, A. M., Hosseini, P., Scott, B. B., Bradley, D. C., Smith, M. A., Kohn, A., Movshon, J. A., Armstrong, K. M., Moore, T., Chang, S. W., Snyder, L. H., Lisberger, S. G., Priebe, N. J., Finn, I. M., Ferster, D., Ryu, S. I., Santhanam, G., Sahani, M., & Shenoy, K. V. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience 13*(3), 369–378.

Churchland, M. M., Yu, B. M., Ryu, S. I., Santhanam, G., & Shenoy, K. V. (2006). Neural variability in premotor cortex provides a signature of motor preparation. *The Journal of neuroscience : the official journal of the Society for Neuroscience 26*(14), 3697–712.

Deco, G., & Hugues, E. (2012). Neural Network Mechanisms Underlying Stimulus Driven Variability Reduction. *PLoS Computational Biology 8*(3), e1002395.

Dehghani, N., Peyrache, A., Telenczuk, B., Le Van Quyen, M., Halgren, E., Cash, S. S., Hatsopoulos, N. G., & Destexhe, A. (2016). Dynamic Balance of Excitation and Inhibition in Human and Monkey Neocortex. *Scientific Reports 6*(1), 23176.

Doiron, B., & Litwin-Kumar, A. (2014). Balanced neural architecture and the idling brain. *Frontiers in Computational Neuroscience 8*(MAY), 56.

Fino, E., & Yuste, R. (2011). Dense inhibitory connectivity in neocortex. *Neuron 69*(6), 1188–1203.

Gold, J. I., & Shadlen, M. N. (2007). The Neural Basis of Decision Making. *Annual Review of Neuroscience 30*(1), 535–574.

Golomb, D., & Hansel, D. (2000). The number of synaptic inputs and the synchrony of large, sparse neuronal networks. *Neural Computation 12*(5), 1095–1139.

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. New York: Wiley.

Holt, G. R., Softky, W. R., Koch, C., & Douglas, J. (1996). Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. *Journal of Neurophysilogy 75*(5), 1806–1814.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America 79*(8), 2554–2558.

Inagaki, H. K., Fontolan, L., Romani, S., & Svoboda, K. (2019). Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature 566*(7743), 212–217.

Isaacson, J., & Scanziani, M. (2011). How Inhibition Shapes Cortical Activity. *Neuron 72*(2), 231–243.

Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., & Arieli, A. (2003). Spontaneously emerging cortical representations of visual attributes. *Nature 425*(6961), 954–956.

Khan, A. G., Poort, J., Chadwick, A., Blot, A., Sahani, M., Mrsic-Flogel, T. D., & Hofer, S. B. (2018). Distinct learning-induced changes in stimulus selectivity and interactions of GABAergic interneuron classes in visual cortex. *Nature Neuroscience 21*(6), 851–859.

Knierim, J. J., & Zhang, K. (2012). Attractor Dynamics of Spatially Correlated Neural Activity in the Limbic System. *Annual Review of Neuroscience 35*(1), 267–285.

Kriener, B., Enger, H., Tetzlaff, T., Plesser, H. E., Gewaltig, M.-O., & Einevoll, G. T. (2014). Dynamics of self-sustained asynchronous-irregular activity in random networks of spiking neurons with strong synapses. *Frontiers in Computational Neuroscience 8*, 136.

La Camera, G., Fontanini, A., & Mazzucato, L. (2019). Cortical computations via metastable activity.

Lagzi, F., & Rotter, S. (2015). Dynamics of competition between subnetworks of spiking neuronal networks in the balanced state. *PLoS ONE 10*(9).

Lee, S.-H., Marchionni, I., Bezaire, M., Varga, C., Danielson, N., Lovett-Barron, M., Losonczy, A., & Soltesz, I. (2014). Parvalbumin-Positive Basket Cells Differentiate among Hippocampal Pyramidal Cells. *Neuron 82*(5), 1129–1144.

Litwin-Kumar, A., & Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience 15*(11), 1498–1505.

Litwin-Kumar, A., & Doiron, B. (2014). Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nature Communications 5*(1), 5319.

Luczak, A., Barthó, P., & Harris, K. D. (2009). Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations. *Neuron 62*(3), 413–425.

Mazzucato, L., Fontanini, A., & La Camera, G. (2015). Dynamics of multistable states during ongoing and evoked cortical activity. *The Journal of neuroscience : the official journal of the Society for Neuroscience 35*(21), 8214–31.

Mazzucato, L., La Camera, G., & Fontanini, A. (2019). Expectation-induced modulation of metastable activity underlies faster coding of sensory stimuli. *Nature Neuroscience 22*, 787–796.

Meckenhäuser, G., Krämer, S., Farkhooi, F., Ronacher, B., & Nawrot, M. P. (2014). Neural representation of calling songs and their behavioral relevance in the grasshopper auditory system. *Frontiers in Systems Neuroscience 8*, 183.

Miconi, T., McKinstry, J. L., & Edelman, G. M. (2016). Spontaneous emergence of fast attractor dynamics in a model of developing primary visual cortex. *Nature Communications 7*(1), 1–10.

Morishima, M., Kobayashi, K., Kato, S., Kobayashi, K., & Kawaguchi, Y. (2017). Segregated excitatoryinhibitory recurrent subnetworks in layer 5 of the rat frontal cortex. *Cerebral Cortex 27*(12), 5846–5857.

Morrison, A., Aertsen, A., & Diesmann, M. (2007). Spike-Timing-Dependent Plasticity in Balanced Random Networks. *Neural Computation 19*(6), 1437–1467.

Najafi, F., Elsayed, G. F., Cao, R., Pnevmatikakis, E., Latham, P. E., Cunningham, J. P., & Churchland, A. K. (2020). Excitatory and Inhibitory Subnetworks Are Equally Selective during Decision-Making and Emerge Simultaneously during Learning. *Neuron 105*(1), 165 – 179.e8.

Nawrot, M. P. (2010). Analysis and interpretation of interval and count variability in neural spikes trains. In S. Grün & S. Rotter (Eds.), *Analysis of Parallel Spike Trains*, pp. 34–58. Springer Verlag, New York, Berlin.

Nawrot, M. P., Aertsen, A., & Rotter, S. (1999). Single-trial estimation of neuronal firing rates: From single-neuron spike trains to population activity. *Journal of Neuroscience Methods 94*(1), 81–92.

Nawrot, M. P., Boucsein, C., Rodriguez Molina, V., Riehle, A., Aertsen, A., & Rotter, S. (2008). Measurement of variability dynamics in cortical spike trains. *Journal of neuroscience methods 169*(2), 374–90.

Ocker, G. K., Litwin-Kumar, A., & Doiron, B. (2015). Self-Organization of Microcircuits in Networks of Spiking Neurons with Plastic Synapses. *PLoS Computational Biology 11*(8).

Okun, M., & Lampl, I. (2008). Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nature Neuroscience 11*(5), 535–537.

Ostojic, S. (2014). Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons. *Nature Neuroscience 17*(4), 594–600.

Packer, A. M., & Yuste, R. (2011). Dense, unspecific connectivity of neocortical parvalbumin-positive interneurons: a canonical microcircuit for inhibition? *The Journal of neuroscience : the official journal of the Society for Neuroscience 31*(37), 13260–71.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, E. (2011). Scikit-learn: Machine learning in python. *J. Mach. Learn. Res. 12*, 2825–2830.

Pereira, U., & Brunel, N. (2018). Attractor Dynamics in Networks with Learning Rules Inferred from In Vivo Data. *Neuron 99*(1), 227–238.e4.

Perin, R., Berger, T. K., & Markram, H. (2011). A synaptic organizing principle for cortical neuronal groups. *Proceedings of the National Academy of Sciences of the United States of America 108*(13), 5419–5424.

Pernice, V., Staude, B., Cardanobile, S., & Rotter, S. (2011). How Structure Determines Correlations in Neuronal Networks. *PLoS Computational Biology 7*(5), e1002059.

Ponce-Alvarez, A., Nácher, V., Luna, R., Riehle, A., & Romo, R. (2012). Dynamics of cortical neuronal ensembles transit from decision making to storage for later report. *Journal of Neuroscience 32*(35), 11956–11969.

Renart, A., & Machens, C. K. (2014). Variability in neural activity and behavior. *Current opinion in neurobiology 25*, 211–20.

Rickert, J., Riehle, A., Aertsen, A., Rotter, S., & Nawrot, M. P. (2009). Dynamic Encoding of Movement Direction in Motor Cortical Neurons. *Journal of Neuroscience 29*(44), 13870–13882.

Riehle, A., Brochier, T., Nawrot, M. P., & Grün, S. (2018). Behavioral Context Determines Network State and Variability Dynamics in Monkey Motor Cortex. *Frontiers in Neural Circuits 12*, 52.

Rost, T., Deger, M., & Nawrot, M. P. (2017). Winnerless competition in clustered balanced networks: inhibitory assemblies do the trick. *Biological Cybernetics*, 1–18.

Roudi, Y., & Latham, P. E. (2007). A Balanced Memory Network. *PLoS Computational Biology 3*(9), e141.

Rubin, R., Abbott, L. F., & Sompolinsky, H. (2017). Balanced excitation and inhibition are required for high-capacity, noise-robust neuronal selectivity. *Proceedings of the National Academy of Sciences of the United States of America 114*(44), E9366–E9375.

Sakai, K., & Miyashita, Y. (1991). Neural organization for the long-term memory of paired associates. *Nature 354*(6349), 152–155.

Shin, M., Kitazawa, A., Yoshinaga, S., Hayashi, K., Hirata, Y., Dehay, C., ichiro Kubo, K., & Nakajima, K. (2019). Both excitatory and inhibitory neurons transiently form clusters at the outermost region of the developing mammalian cerebral neocortex. *Journal of Comparative Neurology 527*(10), 1577–1597.

Teich, M. C., Heneghan, C., Lowen, S. B., Ozaki, T., & Kaplan, E. (1997). Fractal character of the neural spike train in the visual system of the cat. *Journal of the Optical Society of America A 14*(3), 529.

Van Vreeswijk, C., & Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science 274*(5293), 1724–1726.

van Vreeswijk, C., & Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural computation 10*(6), 1321–71.

Wiener, M. C. (2003). An adjustment to the time-rescaling method for application to short-trial spike train data. *Neural computation 15*(11), 2565–76.

Wills, T. J., Lever, C., Cacucci, F., Burgess, N., & O'Keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. *Science 308*(5723), 873–876.

Wu, Y., Hengen, K. B., Turrigiano, G. G., & Gjorgjieva, J. (2019). Homeostatic mechanisms regulate distinct aspects of cortical circuit dynamics. *bioRxiv*, 790410.

Xue, M., Atallah, B. V., & Scanziani, M. (2014). Equalizing excitation–inhibition ratios across visual cortical neurons. *Nature 511*(7511), 596–600.

Zenke, F., Agnes, E. J., & Gerstner, W. (2015). Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nature Communications 6*(1), 6922.

Zhang, L. I., Tan, A. Y. Y., Schreiner, C. E., & Merzenich, M. M. (2003). Topography and synaptic shaping of direction selectivity in primary auditory cortex. *Nature 424*(6945), 201–205.

Znamenskiy, P., Kim, M.-H., Muir, D. R., Iacaruso, M. F., Hofer, S. B., & Mrsic-Flogel, T. D. (2018). Functional selectivity and specific connectivity of inhibitory neurons in primary visual cortex. *bioRxiv*, 294835.