

Functional consequences of correlated excitatory and inhibitory conductances in cortical networks

Jens Kremkow · Laurent U. Perrinet ·
Guillaume S. Masson · Ad Aertsen

Received: 19 March 2009 / Revised: 8 April 2010 / Accepted: 20 April 2010 / Published online: 19 May 2010
© Springer Science+Business Media, LLC 2010

Abstract Neurons in the neocortex receive a large number of excitatory and inhibitory synaptic inputs. Excitation and inhibition dynamically balance each other, with inhibition lagging excitation by only few milliseconds. To characterize the functional consequences of such correlated excitation and inhibition, we studied models in which this correlation structure is induced by feedforward inhibition (FFI). Simple circuits show that an effective FFI changes the integrative behavior of neurons such that only synchronous inputs can elicit spikes, causing the responses to be sparse and precise. Further, effective FFI increases the selectivity for propagation of synchrony through a feedforward network, thereby increasing the stability to background activity. Last, we show that recurrent random networks with effective inhibition are more likely to exhibit dynamical network activity states as have been observed *in vivo*. Thus, when a feedforward signal path is embedded in such recurrent network, the stabilizing effect of effective inhibition creates a suitable substrate for

signal propagation. In conclusion, correlated excitation and inhibition support the notion that synchronous spiking may be important for cortical processing.

Keywords Correlated conductances · Synaptic integration · Sparse coding · Signal propagation

1 Introduction

Cortical neurons receive large numbers of excitatory and inhibitory synaptic inputs whose temporal interplay determines the neuron's spiking behavior. On average, excitation and inhibition balance each other, such that spikes are elicited by membrane potential fluctuations (Gerstein and Mandelbrot 1964; van Vreeswijk and Sompolinsky 1996; Shadlen and Newsome 1998; Kumar et al. 2008b). In addition, it has been shown *in vivo* that excitation and inhibition are correlated, with inhibition lagging excitation by only few milliseconds, creating a small temporal integration window (Okun and Lampl 2008; Hasenstaub et al. 2005; Atallah and Scanziani 2009). This correlation structure could be induced by feedforward inhibition (FFI). Here, an excitatory projection directly synapses onto a neuron, whereas inhibition is provided disynaptically by local inhibitory neurons which also receive excitatory inputs from the same projection (Fig. 1(a)). This type of connectivity pattern has been shown to be ubiquitous in the central nervous system. Its importance in the hippocampal formation was reviewed by Buzsáki (1984), this review discussed the potential differences between feedforward and feedback inhibition in shaping the responses to the afferent input. For the thalamocortical connection it has been studied in detail

Action Editor: X.-J. Wang

J. Kremkow · L. U. Perrinet · G. S. Masson
Institut de Neurosciences Cognitives de la Méditerranée,
UMR6193 CNRS—Aix-Marseille Université,
31 chemin Joseph Aiguier, 13402 Marseille Cedex 20, France

J. Kremkow (✉) · A. Aertsen
Neurobiology and Biophysics, Faculty of Biology,
Albert-Ludwig University, Schänzlestrasse 1,
79104 Freiburg, Germany
e-mail: kremkow@biologie.uni-freiburg.de

J. Kremkow · A. Aertsen
Bernstein Center for Computational Neuroscience,
Hansastrasse 9A, 79104 Freiburg, Germany

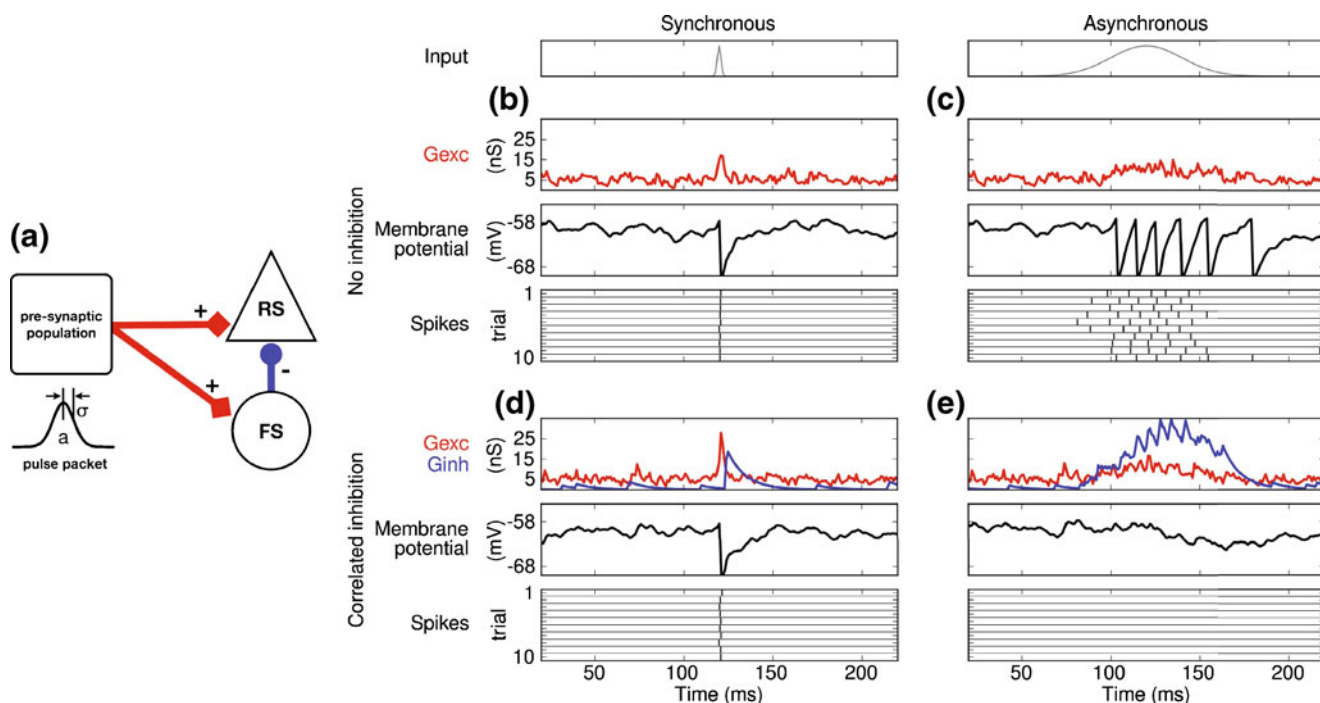


Fig. 1 Correlated inhibition induced by FFI. **(a)** Minimal FFI circuit. **(b)** Integration of a synchronous pulse packet without inhibition. Shown are the excitatory conductance (Gexc), the resulting membrane potential and the spiking responses. **(c)** Integration of an asynchronous pulse packet without inhibition. **(d)** Response to synchronous input when Gexc is accompa-

nied by correlated inhibitory conductance (Ginh). **(e)** Responses to asynchronous input with correlated Ginh. In each case, the spike responses are shown for ten consecutive trials, whereas the excitatory conductance (Gexc) and the resulting membrane potential are shown for a single (the 10th) trial

(Swadlow 2003; Cruikshank et al. 2007). Within the cortical micro-circuitry, this connection scheme seems to be functional as recently demonstrated in slices of rat somatosensory cortex (Silberberg and Markram 2007; Kapfer et al. 2007), where disynaptic inhibition was found in the very local neighborhood, that is within a lateral distance of around 100 μm . In addition, studies of the synaptic physiology of horizontal connections in the visual cortex of the cat (Hirsch and Gilbert 1991) provide evidence that horizontal projections that synapse over some hundreds of micrometers can implement far reaching feedforward inhibition. Thus, electrically stimulating the cortical surface some hundreds of micrometers away from an intracellular electrode may result in excitation, closely followed by inhibition. Since inhibitory connections are generally only local, the source of this inhibition is most likely the activation of inhibitory neurons in the local neighborhood of the recording electrode. These findings have been confirmed by a combination of optical imaging and intracellular recordings of the layer 2/3 network in ferrets (Tucker and Katz 2003a, b).

In view of the many experimental reports documenting the existence of feedforward inhibition, it is important to improve our understanding of its effects on the

dynamics of cortical processing. So far, this question has been addressed either in single neuron models, mainly to explain experimental data (Wehr and Zador 2003), in models of cortical integration of thalamic inputs to show differences between cortical amplification (Somers et al. 1995) and cortical dampening hypothesis (Pinto et al. 2003), or in a complex network model of the olfactory system (Assisi et al. 2007).

Nevertheless, a comprehensive study to understand in details the effects of FFI at different spatial scales, from single neurons to feedforward networks and recurrent networks, is still lacking. Here we study the functional properties of FFI in different scenarios with increasing complexity. First, we will use simple stimuli such as pulse packets (PP) (Aertsen et al. 1996; Diesmann et al. 1999), i.e. spike volleys with gaussian distributed spike times, to investigate the integration behaviors of simple circuits in response to stimuli with a differing degree of synchrony. Next, we assess the signal propagation through a feedforward network (i.e. a synfire chain) made of a sequence of such elementary FFI circuits. Furthermore, we characterize the dynamical states of recurrent random networks with effective inhibition. Lastly, we document the state dependent signal-to-noise level in a feedforward network which is

embedded in a random recurrent network. Preliminary results have been presented in abstract form (Kremkow et al. 2008a, 2009).

2 Materials and methods

2.1 Biophysical prerequisites for effective feedforward inhibition

One important requirement for FFI to be effective is that the inhibitory neurons are very reliably activated to provide efficient inhibition. In other words, they need to have a lower effective spiking threshold than excitatory neurons. This could be achieved in several different ways: by changing the passive cell properties of the inhibitory neurons, by connecting more synapses onto the inhibitory neurons, or by increasing the synaptic weight onto the inhibitory neurons. Although it is known that different neuron types have different passive properties (Nowak et al. 2003) that contribute in facilitating the response of inhibitory neurons, it is less clear if, in addition, a stronger drive is achieved by having more and/or stronger synapses. For the thalamocortical connection, it was shown that the unitary excitatory conductances onto inhibitory neurons are larger and more reliable than those onto excitatory neurons (Cruikshank et al. 2007), thus resulting in larger postsynaptic potentials (Cruikshank et al. 2007; Gibson et al. 1999). Similar findings for the recurrent cortical network were reported by (Povysheva et al. 2006) and were recently confirmed in human cortex (Molnár et al. 2008). However, other studies of cortical synaptic weight distributions could not confirm these findings (Thomson et al. 2002; Brémaud et al. 2007).

Although, *in vivo* it is more likely that a mixture of the above mentioned factors will ensure effective FFI, in the present study we will simplify that problem by having only one parameter controlling the difference in spike probability between the two populations of neurons: the excitatory weights onto the inhibitory neuron ($g_{exc-inh}$). Adopting any of the other mentioned mechanisms as a control parameter would give rise to equivalent results.

2.2 Models

The neurons in the different network models under study were all modeled as leaky-integrate-and-fire neurons, with the subthreshold dynamics of the membrane

potential $V^i(t)$ in neuron i described by the following equation:

$$C \frac{d}{dt} V^i(t) + G_{rest}[V^i(t) - V_{rest}] = I_{syn}^i \quad (1)$$

where I_{syn}^i is the total synaptic input current into neuron i , and C and G_{rest} denote the passive electrical properties of its membrane at rest (V_{rest}). When the membrane potential reached a fixed spike threshold V_{thresh} above rest, a spike was emitted, the membrane potential was reset to its resting value, and synaptic integration was halted for 2 ms to mimic the refractory period observed in real neurons. The parameters used in the simulations were: $C = 290$ pF, $G_{rest} = 29$ nS, $V_{rest} = -70$ mV and $V_{thresh} = -57$ mV. Synaptic input was modeled as transient conductance changes (Kuhn et al. 2004), using exponential functions with $\tau_{exc} = 1.5$ ms and $\tau_{inh} = 10$ ms. Parameters of the passive cell properties as well as synaptic time constants were taken from Muller et al. (2007) and Destexhe et al. (1998). In addition to the synaptic input, each neuron received an independent Gaussian noise current to introduce trial-by-trial variability. The standard deviation of the noise was adjusted to induce membrane potential fluctuations as observed *in vivo* (Destexhe et al. 2003; Rudolph et al. 2007). In all network models studied here, except for the minimal FFI circuit, the mean noise current was set to induce realistic low background firing rates of few spikes per second (Abeles 1991; Diesmann et al. 1999; Gewaltig et al. 2001).

2.2.1 Network architectures

FFI circuit The minimal cortical circuit we studied contained only two types of cells: excitatory (regular-spiking, RS) and inhibitory (fast-spiking, FS) neurons. In the circuit we included one RS neuron and a pool of FS neurons (Fig. 1(a)). The size of the inhibitory pool was set to $n_{FS} = 25$ when investigating the effect of FFI onto the spiking response of the RS neuron (Inoue and Imoto 2006). For the control condition, in which excitation was not accompanied by inhibition, we set $n_{FS} = 0$. The input to the circuit was a spike volley from a population of 100 pre-synaptic neurons, the spike times of which were drawn from a gaussian distribution, creating a pulse packet (Aertsen et al. 1996; Diesmann et al. 1999) with strength (number of spikes/pre-synaptic neuron) ' a ' and temporal dispersion ' σ '. Each cortical neuron received 60 randomly chosen weak excitatory synapses ($g_{exc-exc}$, $g_{exc-inh} = 1$ nS) from the pre-synaptic population (Bruno and Sakmann 2006), except in the cases where we increased the effectivity of FFI by strengthening the excitatory

synapses onto FS neurons to 3.5 nS (Cruikshank et al. 2007). Due to the limited size of the pre-synaptic population, the cortical neurons received highly similar synaptic inputs ('shared input', (Lampl et al. 1999)). The inhibitory synaptic conductance was set to 2 nS to balance the incoming excitation. The synaptic delay from FS to RS neurons ($\text{delay}_{inh-exc}$) was set to 2 ms (Cruikshank et al. 2007). Since in this minimal FFI circuit we were primarily interested in the stimulus induced spike-timing precision, we prevented any background spiking by adjusting the mean of the noise such that without synaptic input the membrane potential stayed below the spiking threshold. Furthermore, to avoid trial-by-trial variability within the stimulating pulse packets, for each (a, σ) -combination we only once drew a pulse packet realization from the respective gaussian distribution and presented this same pulse packet repeatedly in each trial.

Feedforward network The feedforward network was inspired by the synfire chain architecture (Abeles 1991; Aertsen et al. 1996; Diesmann et al. 1999; Gewaltig et al. 2001) in which a group of neurons projects in a divergent-convergent way onto another group of neurons. Repeating this scheme for ten subsequent groups creates a feedforward network (Fig. 3(a)). In the classical synfire chain, these groups contain only RS neurons. Here we extend this architecture by including FS neurons into each group. The FS neurons from one group receive excitatory inputs from the preceding group and make local inhibitory connections within their own group, thus providing correlated and balanced inhibition. To have a sufficiently large number of neurons in the projecting population, we included 100 RS neurons in each group. In this network, as well as in the further network models in this study, the two classes of neurons appeared in the ratio 4:1, following numbers from the classical neuroanatomical literature on neocortex (Braitenberg and Schüz 1991). Each group had 100 RS and 25 FS neurons with each neuron receiving 60 excitatory synapses from the preceding group. Otherwise, synapses and the synaptic conductances were identical to those in the minimal model. Note that the synfire chains used in (Litvak et al. 2003; Aviel et al. 2003; Kumar et al. 2008a) also contained balanced inhibition within the chain (referred to as 'double balance' or 'shadow inhibition'), but it was implemented in a different manner. There, FS neurons projected to the next group, but not within their own group. As a consequence, there was no delay between excitation and inhibition arriving at the target neurons. Because inhibitory neurons only connect to neurons in the local neighborhood, Litvak et al. (2003) concluded

that their model would not be applicable to long-range patchy connections. Our model is different, in that it generalizes also to long-range connections (Hirsch and Gilbert 1991) and that it includes more biophysical details, in particular more realistic time delays and synaptic weights.

Recurrent network The recurrent cortical network contained 10,000 neurons (80% RS and 20% FS, Braitenberg and Schüz (1991)). The connectivity probability was low ($c = 0.1$), creating a sparsely connected network (Fig. 5(a); Kumar et al. 2008b). The synaptic weights for all four connection types ($g_{exc-exc}, g_{exc-inh}, g_{inh-inh}, g_{inh-exc}$) were set to 1 nS. When characterizing the dynamical network states, all inhibitory synaptic conductances ($g_{inh-inh}, g_{inh-exc}$) were scaled by a factor G (similar to Brunel 2000; Kumar et al. 2008b). For example, $G = 2$ would mean that with a $g_{exc-exc}$ of 1 nS, the inhibition $g_{inh-inh}, g_{inh-exc}$ would be 2 nS. For the experiments in which we increased the effective inhibition, the unitary excitatory conductance onto FS neurons ($g_{exc-inh}$) was increased to 3.5 times the conductance onto RS neurons ($g_{exc-exc}$) (Povysheva et al. 2006; Cruikshank et al. 2007). A uniform transmission delay of 2 ms was imposed for all synapses. Next to the synapses from the network, neurons also received an external poissonian input with rate ' v_{ext} ', mimicking non-local cortico-cortical inputs to the local network.

Recurrent network with embedded signal path To study signal propagation in a more realistic scenario, we embedded the feedforward network into a laminar patch of cortical neurons (in a similar manner as in Kumar et al. (2008a) and Vogels and Abbott (2009)). We used 28125 leaky integrate-and-fire neurons, 22,500 excitatory and 5,625 inhibitory, to represent a small cortical sheet of 1×1 mm (Fig. 6(a)). The network was folded as a torus to avoid boundary effects. As in Kumar et al. (2008a) we assumed a sparsely connected network ($c = 0.05$) and described the distance dependent connection probability with a Gaussian profile (Hellwig 2000; Mehring et al. 2003). The spatial widths were set to: $\sigma_{exc-exc} = 0.4$ mm, $\sigma_{exc-inh} = 0.2$ mm, $\sigma_{inh-inh}, \sigma_{inh-exc} = 0.1$ mm, taking into account differences in the spatial extent of connection types (Stepanyants et al. 2008). As in the other models throughout this paper, all synaptic weights were set to 1 nS. Again the factor ' G ' scaled the overall inhibition while $g_{exc-inh}$ controlled the effectiveness of the inhibition. A uniform transmission delay of 2 ms was imposed for all synapses. Next to the synapses from the network, neurons also received an external poissonian input with

rate ' ν_{ext} ', mimicking non-local cortico-cortical inputs to the local network.

To investigate signal propagation we embedded a feedforward network, composed of three synfire groups, into the cortical sheet (Fig. 6(a)). The number of neurons within a group as well as the number of synapses between two groups were the same as in Section 2.2.1. To maintain a balanced connectivity in the overall network, neurons in the feedforward network made accordingly fewer synapses with the embedding recurrent network, such that all neurons, irrespective of whether they were in the feedforward or in the recurrent network, received and projected equal numbers of synapses (Kumar et al. 2008a). The delay from one group in the feedforward network to the next was set to 40 ms for illustrative purposes.

2.3 Characterization of the spike response

Single neuron integration The response strength, that is the number of spikes per pulse packet stimulus recorded in the RS neuron, was measured by dividing the total spike count by the number of stimulus presentations. To quantify the trial-by-trial precision of the response, we simulated multiple trials and estimated the correlation coefficient of the binned (binwidth: 1 ms) spike trials, excluding the correlation of each trial with itself.

Feedforward network To characterize the spiking activity of each group and, thus, the 'a' and ' σ ' of the pulse packet sent to the next group, we measured 'events', i.e. distinct peaks in the peri stimulus time histogram (PSTH) of all neurons from the RS group (similar to Kumbhani et al. 2007). The PSTH was calculated by binning (binwidth 1 ms) the spiking activity and convolving it with a triangular kernel of 2 ms. The beginning and end of an event was set by thresholding the PSTH, to avoid background activity from being detected as events, and searching for gaps within the PSTH which were separated by at least 4 ms. All spike times occurring within the region between the gaps were collected and thus defined an event. The mean spike count of all RS neurons in such an event described the response strength and defined the next 'a'. Likewise, the temporal spread of the spike times, estimated by the standard deviation of the spike times within the event, constituted the new ' σ ' (as in Gewaltig et al. 2001).

We could thereby determine whether the signal could propagate from one group to the next and, if so, whether it synchronized while propagating along the chain. An activation of the synfire attractor would

express itself in a small ' σ ' and a value of 'a' close to 1 in the last group. To calculate the separatrix, we fitted the function: $f = a + (b * x^c)$ as a smooth approximation to the gridwise-determined border separating successful activation of the synfire attractor from propagation failures in the parameter space of ' σ ' and 'a'. In other words, the border between successful propagation and failure was the right-most grid position on the ' σ '-axis (the horizontal axis in the panels in Fig. 4) in which stimulus related activity was detectable in the last group of the feedforward network.

Recurrent network We characterized the spiking activity using the methods in (Kumar et al. 2008b). In short, the mean activity in the network was assessed by measuring the mean firing rate of all active neurons. The degree of synchrony among neurons in the population was measured by the pairwise correlation, defined as the correlation coefficient of the binned (binwidth: 2 ms) spiking activities of 500 randomly selected neuron pairs (400 pairs in the RS and 100 pairs in the FS population). The CV of the inter-spike-interval (ISI) distribution of individual neurons was used to quantify the irregularity of spiking. The CV was given by:

$$CV_{ISI}^2 = \text{Var}[ISI]/E[ISI]^2. \quad (2)$$

Low values reflect more regular spiking and high value irregular spiking.

Since estimates of spike train irregularity from finite length experimental observations suffered from a tendency to under-estimate its value, we adopted the criterion established in (Nawrot et al. 2008) to use 'long enough' observation windows (30 s) that comprised at least 10 spikes.

Recurrent network with embedded signal path To characterize the state dependent quality of signal propagation, we analyzed the spiking activity in the last group of the embedded feedforward network and compared it to the spiking activity of the embedding recurrent network. For this, we estimated the spike count in the 100 RS neurons of the last group in a 4 ms window around the expected arrival time of the signal (Kumar et al. 2008a). The background activity was characterized in the same 4 ms window in 100 randomly chosen RS neurons from the embedding recurrent network.

2.4 Simulation and data analysis tools

Simulations were written in python using PyNN (Davison et al. 2009) as interface to the simulation environment NEST (Gewaltig and Diesmann 2007; Morrison et al. 2005; Eppler et al. 2009). Simulation

management and data analysis were performed using the python package NeuroTools (Yger et al. 2009).

3 Results

3.1 Single neuron integration

To characterize the effect of feedforward inhibition on single neuron integration, we first constructed a minimal cortical circuit containing the principal elements (Fig. 1(a) and see Section 2). As stimuli we applied pulse packets (Diesmann et al. 1999) with strength (number of spikes/pre-synaptic neuron) ‘ a ’ and temporal width ‘ σ ’ (see Section 2). Using the pulse packet framework we were then in the position to compare the interaction of stimuli with differing degrees of synchrony with the model circuit (Diesmann et al. 1999; Kumar et al. 2008a).

3.1.1 Example of single neuron integration with and without correlated inhibition

To illustrate the influence of the correlated inhibition on the integration of incoming pulse packets, we recorded the total impinging conductance, the resulting membrane potential, and the response spikes of the RS neuron. Because we were interested in the stimulus-induced variability of the neuronal response, we adjusted the mean noise level such that, without a stimulus, the membrane potential stayed below spiking threshold (see Section 2). Thus, all spikes were induced by synaptic input from the stimulus. When the neuron could integrate a synchronous pulse packet without receiving inhibition ($\sigma = 1$ ms, $a = 1$), it typically responded with one precise spike per trial (Fig. 1(b)). However, when the pulse packet was composed of many asynchronously arriving spikes ($\sigma = 20$ ms, $a = 5$), the neuron responded with multiple, variable spikes (Fig. 1(c)). These spikes were caused by the long lasting depolarization of the membrane potential, due to the interaction of the impinging excitatory conductances (Gexc) with the membrane time constant. This firing behavior changed when, in addition, correlated inhibition (Ginh) was provided by the feedforward circuit. Whereas the synchronous pulse packets could still induce one precise spike per trial (Fig. 1(d)), the asynchronous pulse packet could hardly elicit a spike response, if at all (Fig. 1(e)). The Ginh induced by the FFI was now balancing the impinging Gexc and, as a result, the mean membrane potential remained almost stationary and clearly below threshold, thereby diminishing the response to at most a single spike in

only few trials (none in the ones shown in Fig. 1(e)), with highly variable spike timing. The fact that feedforward inhibition may increase the temporal fidelity of the spiking response is in agreement with experimental results obtained with hippocampal neurons (Pouille and Scanziani 2001). In fact, this differential behavior of the effective temporal integration in the receiving neuron, depending on the temporal properties of the input stimulus and the presence of correlated inhibition, should be measurable in spike correlation functions, measured extra-cellularly between the pre- and post-synaptic excitatory neurons in Fig. 1(a). In case of no inhibition (Fig. 1(b), (c)), we would expect a delayed peak in the cross-correlogram, narrow for synchronous input (Fig. 1(b)), possibly followed by a small negativity reflecting refractoriness, and broader for asynchronous input (Fig. 1(c)), reflecting the increased jitter in spike response in the latter case. In the case of correlated inhibition (Fig. 1(d), (e)), we would expect a delayed peak in the cross-correlogram, narrow for synchronous input (Fig. 1(d)), followed by a clear negativity reflecting the absence of later spiking due to correlated inhibition, and a strongly reduced positivity for asynchronous input (Fig. 1(e)), reflecting the response reduction due to correlated inhibition. Indeed, some such observations have been reported in the literature, both for *in vitro* recordings (Hirsch and Gilbert 1991) and in network models (Fig. 4 in Kumar et al. 2008a).

3.1.2 Pulse packet state space

To systematically study the circuit properties we applied a range of differently shaped pulse packets. We varied ‘ σ ’ from 1 ms (synchronous) to 30 ms (asynchronous) and ‘ a ’ from 1 spike/pre-synaptic neuron (low signal-to-noise ratio) to 14 spikes/pre-synaptic neuron (high signal-to-noise ratio). For very small ‘ σ ’, ‘ a ’ might become unrealistically high: i.e. 10 spikes per pre-synaptic neuron within 5 ms is not physiologically realistic. However, since our aim was to obtain a complete overview of all possible combinations of ‘ σ ’ and ‘ a ’, we chose not to exclude such unrealistic combinations—they should, however, be interpreted with great care. As the neuronal membrane acts as an integrator, we expected that with increasing temporal width and number of spikes, the response to the incoming pulse packet should increase. Indeed, when the RS neuron integrated the impinging pulse packet without being constrained by inhibition, the response strength showed the expected trend (Diesmann et al. 1999; Kumar et al. 2008a). Thus, for small ‘ a ’, the pulse packet needed to be very synchronous ($\sigma < 5$ ms) to be able to elicit a response spike reliably (Fig. 2(a)). Pulse packets that

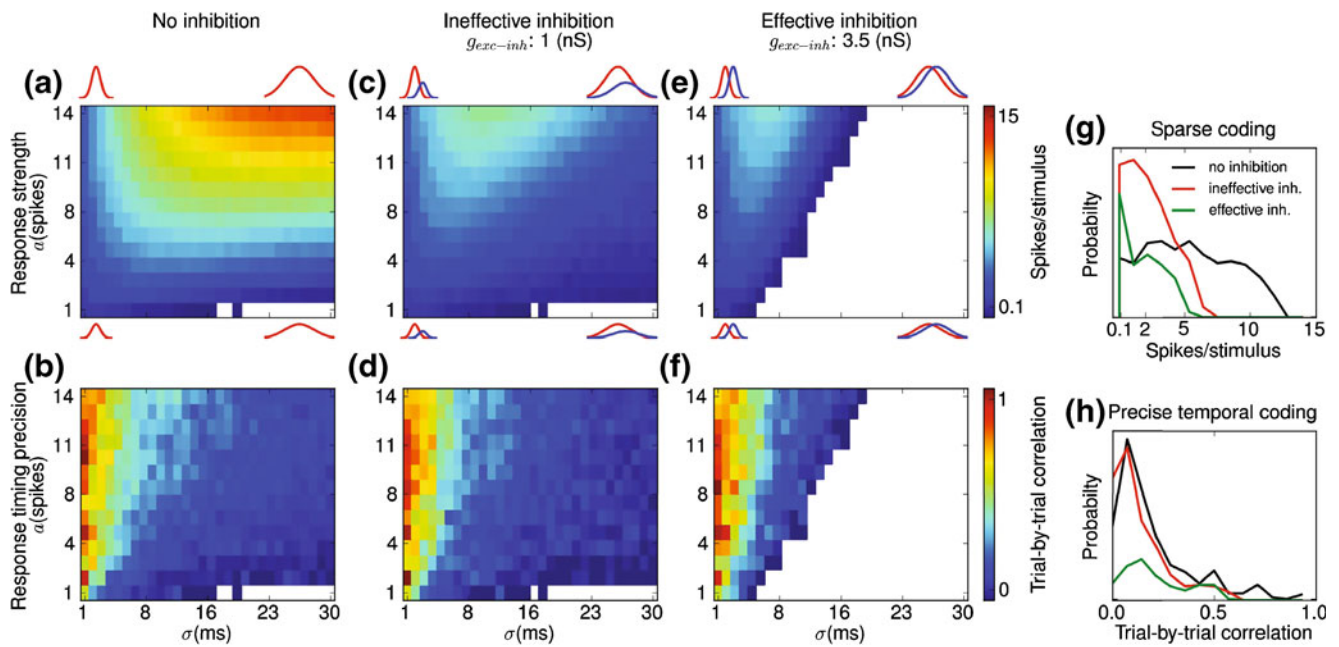


Fig. 2 Sparse and precise coding by effective FFI. **(a)** Response strength when pulse packets of various parameters are integrated in the absence of inhibition. Responses < 0.1 spikes/stimulus are shown as white entries in **(a)–(f)**. **(b)** Response timing precision of the responses in the absence of inhibition. **(c, d)** With FFI, with the excitatory synaptic weight onto the FS neurons and RS neurons both 1 nS, the FS induced ineffective inhibition.

(e, f) With the excitatory weight onto FS 3.5 nS, inhibition was more effective and suppressed the responses to asynchronous inputs. **(g)** Probability of response strengths. Correlated inhibition suppressed dense responses. **(h)** Probability of response timing precision. Effective inhibition increased the probability of precise responses, whereas ineffective inhibition failed to do so

were either not synchronous or strong enough failed to elicit response spikes altogether (responses < 0.1 spikes/stimulus are shown as white entries in Fig. 2(a)–(f)). With increasing ‘ a ’, even larger ‘ σ ’ could induce a response, often in the form of multiple spikes. The temporal precision was highest for synchronous pulse packets, i.e. with small ‘ σ ’ (Fig. 2(b); Mainen and Sejnowski 1995; Diesmann et al. 1999). For asynchronous inputs (large σ), the timing of response spikes was more variable, as the background noise had more time to interact with the less transient input.

When the RS neuron received correlated inhibition together with the excitation (FFI), the state space changed considerably, with the change stronger for less synchronous input (i.e. going towards the right in the panels in Fig. 2), especially in the case of effective inhibition. As mentioned in Section 2.1, an effective FFI requires that the FS neuron already spikes for smaller stimulus input than the RS neuron. To test whether this is indeed the case, we simulated the same state space for two different values of $g_{exc-inh}$ (1 and 3.5 nS).

First, we studied a case where RS and FS neurons received the same amount of conductances from the pre-synaptic population ($g_{exc-inh} = 1$ nS). Under this

condition, the response rate was considerably reduced (Fig. 2(c)). Nevertheless, FFI could not effectively and fully prevent the spiking response to asynchronous pulse packets (large σ), because RS and FS neurons reached threshold at about the same input strength. The remaining RS spike times were again variable (Fig. 2(d)).

Increasing $g_{exc-inh}$ to 3.5 nS (Cruikshank et al. 2007) induced a clearly more effective FFI. The responses to asynchronous pulse packets were now suppressed (large white area in Fig. 2(e) for large σ), whereas a synchronous input could still elicit a spike response (Fig. 2(e), small σ). However, the response magnitude was also reduced for more synchronous inputs, as the RS neuron could not fully integrate the excitatory input, except for the most synchronous case ($\sigma = 1$). The temporal precision showed a similar trend as before, however, by effectively suppressing the responses to less transient inputs, the FFI increased the overall temporal precision of the response (Fig. 2(f)).

In summary, correlated excitation and inhibition, such as induced by FFI, changes dynamically the effective integration behavior of the neuron. The small time lag between the incoming excitation and inhibition

results in high-pass filtering properties and, thus, lowers the effective membrane time-constant. Therefore, only synchronous inputs can reach threshold. These synchronous inputs induce only few (i.e. sparse) but highly precise response spikes. In addition, by reducing multiple spiking, effective FFI renders the responses to become even more sparse, such that the response often contains only one spike (Fig. 2(g)). By contrast, without correlated inhibition, asynchronous inputs reached threshold and induced multiple threshold crossings (Fig. 2(g)), resulting in dense and temporally imprecise spike responses (Fig. 2(h)).

3.2 Signal propagation through feedforward networks (synfire chains)

Next, we studied how the FFI affects the signal propagation through a feedforward network architecture, often called a synfire chain. Synfire chains are known to have a stable fix-point (Diesmann et al. 1999) at small ' σ ', with the basin of attraction (BOA) spanning a wide range of pulse packet parameters. Whereas ' σ ' has to be small at small ' a ', it can become wider with increasing ' a ' (Diesmann et al. 1999), introducing serious problems for the stability of the ground state of the feedforward network: synchronous activity can be spontaneously induced by random fluctuations of background activity or by a stimulus induced rate increase (Tetzlaff et al. 2002).

In the previous section we argued that effective FFI induces an increased selectivity for synchronous inputs. We therefore predicted that embedding FFI into feedforward networks could stabilize the ground state of such networks. To study the effect of FFI onto the signal propagation, we constructed a synfire chain embedding inhibitory populations within each group to generate local FFI connections within each of them (Fig. 3(a); see Section 2).

3.2.1 FFI shapes the selectivity of synfire chains

The FFI had differential effects on the synfire dynamics, depending on the parameters of the pulse packet. When the pulse packet presented to the first group was synchronous (small σ , 1 ms; a , 1 spike), the synfire chain propagated it both with and without correlated Ginh (Figs. 3(b, d)). In either case, the pulse packet was strong and precise enough to elicit spikes in enough neurons of the first group, such that, in turn, it could activate the next group, and reach the BOA, rapidly converging into the fix-point. However, when the pulse packet was asynchronous (large σ , 20 ms) and stronger (larger a , 5 spikes), resembling asynchronous rate input,

the FFI had a strong effect. Without correlated Ginh, the rate input resulted in an elevated firing rate in the neurons of the first group (Fig. 3(c)), which was high enough to fall into the BOA and, hence, to synchronize over subsequent groups. By contrast, the FFI prevented strong spiking responses in the first group (Fig. 3(e)), thereby prohibiting asynchronous inputs from inducing synchronous activity to propagate in subsequent groups.

To assess the change of the BOA upon incorporating FFI into the chain, we again scanned a parameter range of the pulse packets similar to the one used in the Section 3.1. Without FFI, the ' σ ' for which one can detect a synfire chain propagation in the last group depended on the number of input spikes (Fig. 4(a)). The direction and length of the arrows at any particular position represent the direction and strength of the local activity gradient, i.e. how, on average, a pulse packet starting at this position is transformed into a pulse packet in the subsequent neuron group in the feedforward chain (as in Diesmann et al. 1999). The blue line shows the separatrix between failure and successful propagation (see Section 2). Already with one spike in each neuron of the pre-synaptic population, the stimulus could propagate through the network. This low activation threshold is not surprising as the membrane potentials are close to threshold. Embedding FFI into the synfire chain resulted in a change of the BOA. Pulse packets that could elicit synfire activity without FFI now failed to do so (Fig. 4(b)). Thus, the synfire chain with FFI was more selective than without FFI, as it was activated by fewer combinations of pulse packet parameters. To study the effect of the different temporal delays between excitation and inhibition (delay_{inh-exc}), we varied the synaptic delay from 2 ms to 8 ms (2 ms was the default value, see Section 2). Since the delay regulates the integration window of the RS neurons, it also controls the selectivity of the synfire chain. To test this, we calculated again the separatrix between failure and successful propagation for each delay value (Fig. 4(c); see Section 2). At 2 ms delay, activation of the synfire chain was only possible for sufficiently large ' a ' at very small ' σ '. Increasing the delay_{inh-exc} resulted in a widening of the integration window and, thereby, in a reduction of the selectivity (Fig. 4(c)). The input selectivity of the synfire chain is critical when one wants to embed multiple chains in a network. Without FFI, the probability of falling into the BOA is higher, and the synchronization process over subsequent groups will ensure stable propagation, even for asynchronous input, see also (Kumar et al. 2008a). By contrast, with FFI, the propagation of spiking activity along the feedforward chain is restricted to

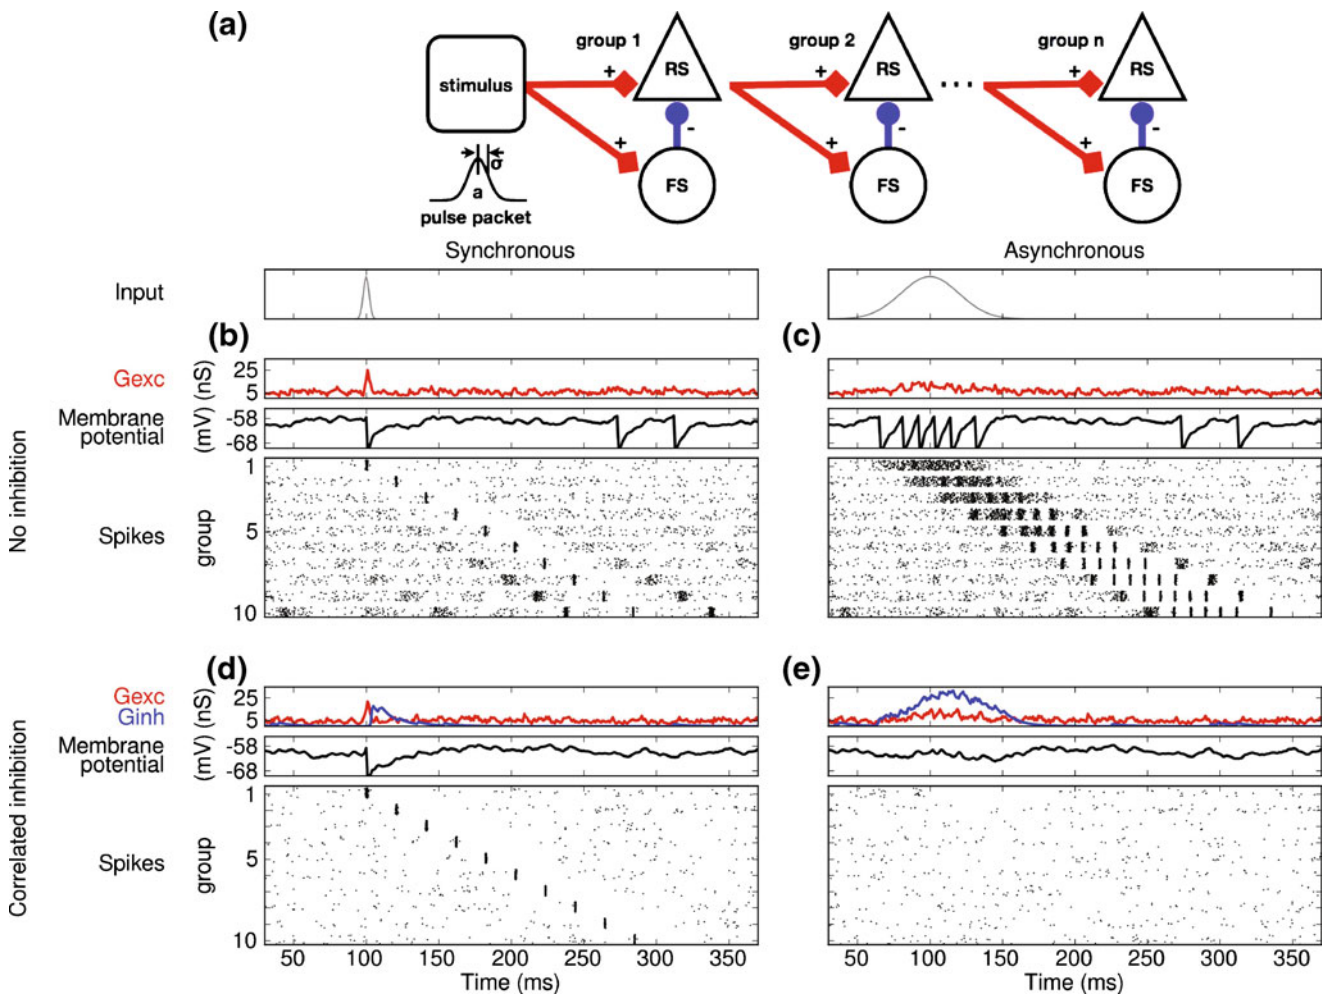


Fig. 3 Conditions for signal propagation through a feedforward network with correlated inhibition. **(a)** Model of a feedforward network with correlated inhibition induced by FFI. **(b)** Signal propagation of a synchronous input through the network when each group projects only onto the RS population of the following group. Due to the instability of the ground-state of purely excitatory feedforward networks (Tetzlaff et al. 2002), transient random fluctuations in the asynchronous background activity may occasionally induce spontaneously propagating synchrony,

as can be observed here some 50 ms before the stimulus onset. **(c)** Propagation of an asynchronous input through the same network. The asynchronous input induced elevated firing rates in the first groups. However, the activity rapidly synchronized over subsequent groups. **(d)** Propagation of synchronous input was hardly affected by correlated inhibition, induced by including the FS neurons in the target population of the successive group. **(e)** FFI in the feedforward network prevented asynchronous inputs from inducing synchronous activity in subsequent groups

synchronous pulse packets, with the required degree of synchrony being determined by the delay between excitation and inhibition in the FFI.

3.3 Stability of random recurrent networks

A common model of the cortical architecture is the randomly connected recurrent network (Fig. 5(a)). Its dynamical states have been well studied (Brunel 2000; Kumar et al. 2008b) and can be categorized, depending on the regularity of the single neuron spike trains (regular ‘R’ or irregular ‘I’) and on the synchrony

across neurons (synchronous ‘S’ or asynchronous ‘A’). Spontaneous *in vivo* activity is best described by the asynchronous irregular state (AI), even though high correlations between spatially close neurons have been reported both in *in vivo* recordings from cortex (Lampl et al. 1999; Okun and Lampl 2008; Smith and Kohn 2008) and in theoretical studies on network models with distance-dependent connectivity (e.g. Mehring et al. 2003). In random networks, the AI state is only possible in a restricted parameter regime, such that the network evolves towards a synchronous state for most parameter values considered. Also, stable propagation

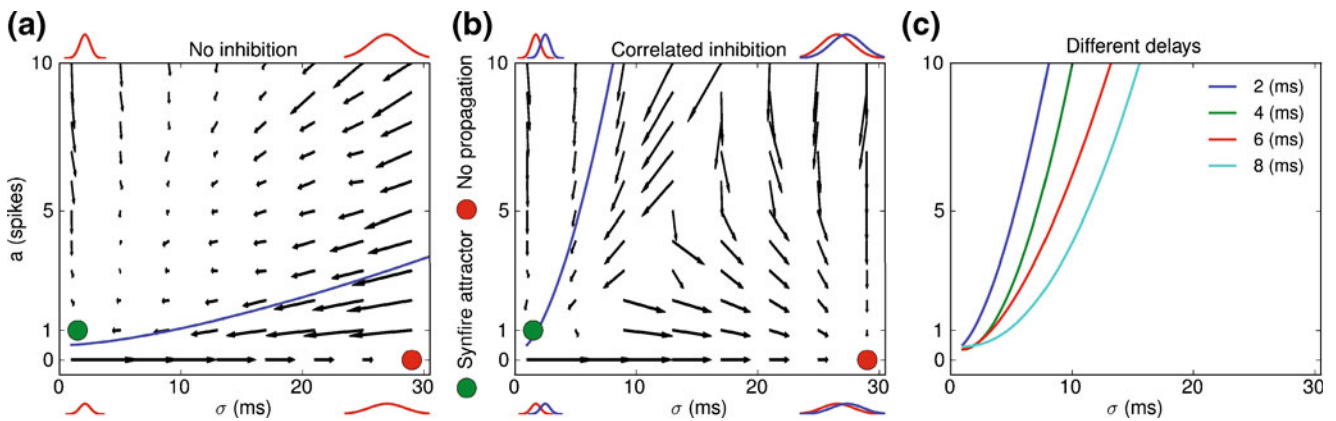


Fig. 4 Correlated inhibition shapes selectivity of feedforward networks. **(a)** Basin of attraction (BAO) when excitation can be integrated, without being constrained by inhibition. *Arrows* represent the local evolution of the pulse packet parameters ‘a’ and ‘σ’, i.e. how, on average, a pulse packet starting at this position is transformed into a pulse packet in the subsequent neuron group in the feedforward chain (as in Diesmann et al. 1999). The green

marker represents the synfire attractor, the red marker indicates the case when no event could be detected in the last group. The BAO for the feedforward network without inhibition is large. **(b)** Including FFI in the feedforward network changes the BAO. Only synchronous inputs can propagate through the network. **(c)** The temporal delay between G_{exc} and G_{inh} controls the selectivity of the feedforward network

of synfire activity is easier in AI type background activity than in synchronous network states (Kumar et al. 2008a).

In blocking experiments using bicuculline to prevent GABA-ergic inhibition, (Hirsch and Gilbert 1991) showed that responses to electrical stimuli became

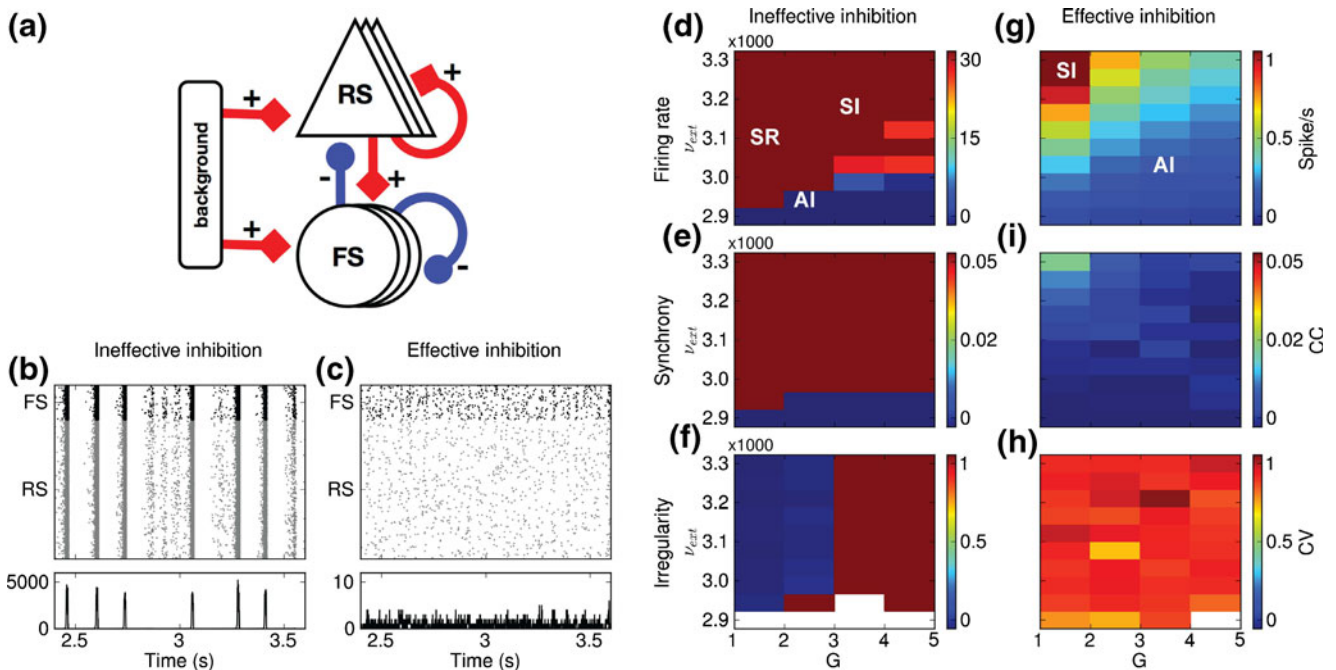
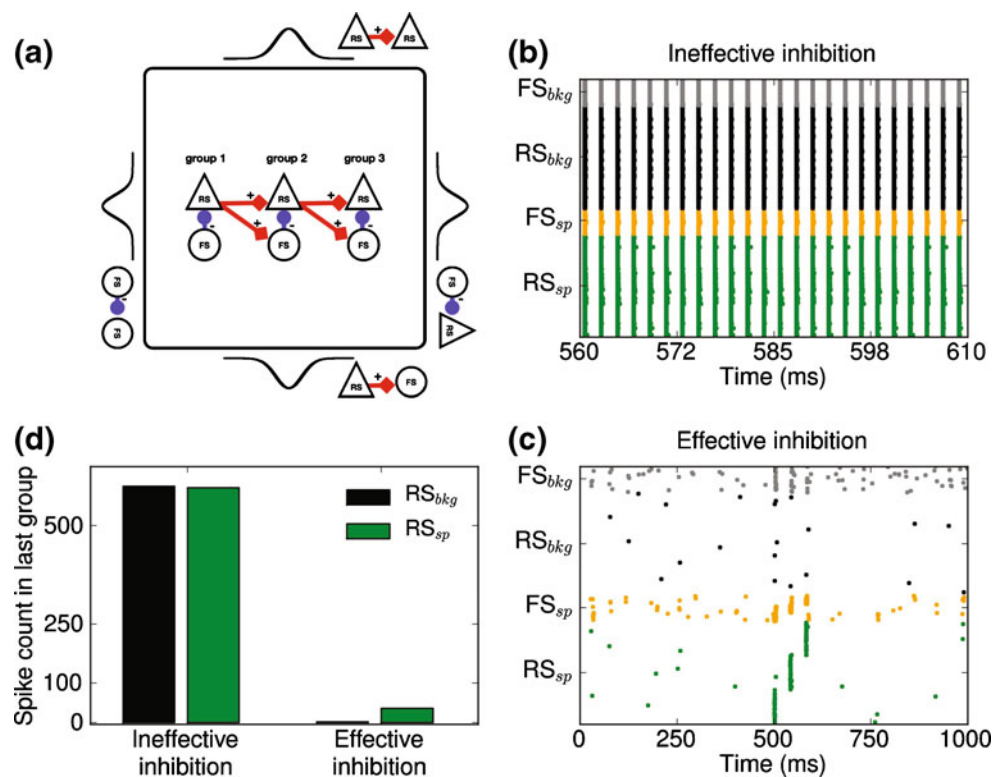


Fig. 5 Stability of recurrent random networks by effective inhibition. **(a)** Model of recurrent cortical random network. **(b)** Spiking activity of the network with ineffective inhibition ($g_{exc-inh} = 1$ nS). The network shows population wide synchronous events. **(c)** With effective inhibition ($g_{exc-inh} = 3.5$ nS), the network is stable and shows asynchronous irregular (AI) activity. **(d–f)** Characterization of the network states with ineffective inhibition

as a function of G and v_{ext} . **(g–h)** Likewise, characterization of the network states with effective inhibition. **(d, g)** Firing rate. **(e, i)** Synchrony. **(f, h)** Irregularity. *White labels* in panels **(d)** and **(g)** exemplarily indicate prototypical network states—further explanation in the main text. Note that for the CV measurement we set a minimum of 10 spikes per trial; data that remained below that number are shown as *white entries* in **(f)** and **(h)**

Fig. 6 Signal-to-noise level in embedded feedforward networks depends on efficacy of inhibition. **(a)** Model of locally connected random network with embedded signal path. **(b)** Spiking activity of recurrent network with embedded signal path with ineffective inhibition ($g_{exc-inh} = 1$ nS, $v_{ext} = 3,000$). **(c)** Spiking activity with effective inhibition ($g_{exc-inh} = 3.5$ nS, $v_{ext} = 3,000$). **(d)** Spiking activity in the background RS neurons (RS_{bkg}) and in the RS neurons of the last group of the signal path (RS_{sp}) in both network configurations



unstable and resulted in a burst of activity. This suggests that inhibition is crucial in preventing runaway excitatory processes and, hence, can increase the stability of the cortical network.

To show that effective FFI not only results in sparse and precise responses in simple circuits, but can also help stabilizing a recurrent network, we constructed a random recurrent network (see Section 2) and studied its state space. Usually the parameters that are used to study the dynamical states of random networks are the rate (v_{ext}) of the external inputs (typically modeled by poisson processes) and the relative strength of the inhibition (G). The excitatory synaptic strength is assumed identical for both RS and FS neurons. To test the stabilizing effect of effective FFI in recurrent networks, we increased the synaptic strength from RS onto FS neurons ($g_{exc-inh}$) by a factor of 3.5 (Povysheva et al. 2006; Cruikshank et al. 2007), in the same manner as in the previous section. We found that, indeed, increasing $g_{exc-inh}$ had a stabilizing effect on the recurrent network dynamics. Thus, with a parameter combination of ' v_{ext} ' and ' G ' at which a classical network ($g_{exc-inh} = 1$ nS) exhibits population wide synchronizations (Fig. 5(b)), a network with effective inhibition ($g_{exc-inh} = 3.5$ nS) showed a stable AI state (Fig. 5(c)). To further characterize the effect of the effective inhibition over a wide parameter space, we then studied the v_{ext} - G state-space

for both a classical random network ($g_{exc-inh} = 1$) and an enhanced FS network ($g_{exc-inh} = 3.5$). The classical random network exhibited the expected dynamical states (Brunel 2000; Kumar et al. 2008b), indicated here by white labels in Fig. 5(d): with increasing v_{ext} , the mean firing rate increased (Fig. 5(d)) and, depending on G , the network exhibited the various states regarding population synchrony (Fig. 5(e)) and spiking irregularity (Fig. 5(f)). Note that for the CV measurement we set a minimum of 10 spikes per trial (Nawrot et al. 2008) (see Section 2); data that remained below that number are shown as white entries in Fig. 5(f) and (h).

For small G , increasing v_{ext} resulted in synchronous and regular (SR) spiking activity. With increasing G , the responses became more irregular (Fig. 5(f)), however, for most of the tested v_{ext} they remained synchronous (SI). Only for small external input v_{ext} , the network showed AI type activity.

By contrast, the state space of the enhanced FS network changed drastically. Although the mean firing rate increased with v_{ext} (Fig. 5(g)), it remained quite low compared to the classical network (Fig. 5(d)). The recurrent network now was much more stable and did not elicit a population wide synchronization at parameter combinations which did induce those in the classical network (compare Fig. 5(e), (i)). A similar effect was observed for the irregularity of the single neuron spike

trains (compare Fig. 5(f), (h)). Thus, due to this large increase in stability, the parameter regime for which the network could sustain AI type activity was vastly increased, now covering most of the state space (compare positions of white labels in Figs. 5(d, g)).

3.4 High signal-to-noise level in embedded feedforward network with effective inhibition

It has been previously shown that propagation of synchrony in embedded feedforward networks is more robust when the embedding recurrent network is in a low rate, low synchrony AI-regime (Kumar et al. 2008a). To test whether effective inhibition, both at the global network level and as FFI in the feedforward network, results in a substrate that supports the propagation of synchrony, we embedded a feedforward network composed of three groups as a signal path into the recurrent background network. As recurrent background network we used a locally connected random network (Kumar et al. 2008a; Mehring et al. 2003) which represents a cortical sheet of about 1×1 mm (Fig. 6(a); see Section 2). We then compared the quality of signal propagation in the network with and in one without effective inhibition. For this we choose a v_{ext} value which resulted in a synchronous state of the background network in the case of ineffective inhibition ($v_{ext} = 3,000$). G was set to 2, such that the inhibitory synapses had the same value as in the first two sections. We then applied a pulse packet stimulus ($\sigma = 1$ ms, $a = 1$ spike) to the first group of the embedded feedforward network. The quality of signal propagation was assessed by analyzing the spiking activity of the 100 RS neurons in the last group (signal, RS_{sp}) and 100 RS neurons randomly chosen from the background network (noise, RS_{bkg}). We made sure that none of the 100 RS neurons of the last group were among the 100 randomly chosen RS background neurons. We then estimated the spike count in a small time window (4 ms) around the expected time of arrival of the propagating pulse packet in the last group. We expected that the signal-to-noise level would depend on the state of the background network (Kumar et al. 2008a). With ineffective inhibition, the recurrent network could sustain a high rate and high synchrony regime (Fig. 6(b), note that for illustrative purposes only a subset of RS and FS neurons in the background (RS_{bkg} , black; FS_{bkg} , gray) is shown. The number of neurons was set to be the same as in the signal path (RS_{sp} , green; FS_{sp} , orange). Activating the signal pathway in such a state gave rise to a signal that was very difficult to distinguish from the dominant background activity (Fig. 6(b)). Figure 6(b) shows a 50 ms time-window around the expected arrival

time (583 ms) of the pulse packet in the last group. By contrast, with effective inhibition the same combination of v_{ext} and G induced only a low rate regime, with only occasionally a very partial and short-lasting synchronization among a small fraction of the neurons (Fig. 6(c)). Here, the start of the signal and its propagation through the groups is easily discernible (please note that the time axis was changed from Fig. 6(b) to Fig. 6(c) to show the successful propagation from group 1 to group 3). The activity in the FS_{bkg} during the signal propagation is the so called after-activity or halo (Kumar et al. 2008a). Here the traveling pulse packet excites the neighborhood of a group after one synaptic delay. Due to the rather low number of FS neurons in this study many FS neurons were affected by the halo. Comparing the spike counts in the background and in the last group measures the state dependent quality of signal propagation. In the case of ineffective inhibition, the spiking activity in the last group was controlled by the background network (Fig. 6(b), (d)). In this situation, it was not evident whether or not a signal arrived in the last group. This changed clearly with effective inhibition. Now the recurrent network was much more stable, providing a substrate in which synchronous spiking activity propagates more robustly. The background activity was low and the signal in the last group could be read out easily (Fig. 6(c), (d)).

4 Discussion

We characterized the consequences of correlated and lagged inhibition, as implemented by the feedforward inhibition scheme, on different network scales.

At the single neuron level, we found that feedforward inhibition changes the effective integration behavior of the neuron, such that only transient inputs can induce spiking. The main critical parameter to induce effective and correlated inhibition was the stronger unitary synaptic weight onto FS neurons, consistent with biological findings (Cruikshank et al. 2007). It is known that neurons respond more precisely to transient inputs (Mainen and Sejnowski 1995). The imprecision is related to the membrane potential variability just prior to the spike (Tiesinga et al. 2008). The FFI architecture used here ensures that the amount of inhibition balances the amount of excitation (Wehr and Zador 2003). As a result, the mean membrane potential is below threshold, such that spikes are mainly driven by transient fluctuations (the fluctuation-driven regime: Shadlen and Newsome 1994, 1998; Kuhn et al. 2004). In the present study all synapses were static. However,

it has been shown that this precise balance is also maintained when the synapses show depressing dynamics (Higley and Contreras 2006). That such gain control can indeed result in sparse responses over a wide range of input intensities has been already demonstrated in a model of the olfactory system (Assisi et al. 2007). In the visual system, balancing feedforward inhibition was proposed as a candidate mechanism to explain the invariant orientation tuning observed with drifting gratings (Troyer et al. 1998). Interestingly, it was shown *in vivo* (Marre et al. 2005; Haider et al. 2010) and in a model of the early visual system (Kremkow et al. 2008b) that conductances are indeed correlated during natural viewing. In the model, the correlation of the conductances was related to the correlation of the contrast polarity in the receptive fields of cortical layer 4 neurons. Thus, FFI not only controls the sparseness, but also induces basic functional properties in thalamic input receiving neurons in the cortex.

In recurrent networks, an effective synaptic connection from excitatory for inhibitory neurons drastically increased the network stability. The same network, which exhibited unstable synchronous states, showed stable AI type activity when effective inhibition was present. It was recently shown that the conditions for signal propagation through feedforward networks embedded in a recurrent network, depend on the activity state of the embedding network (Kumar et al. 2008a). In that study, synchrony could optimally propagate in an AI activity state, although successful propagation was also possible in more synchronous states. By contrast, asynchronous rate inputs could not propagate, because of the synchronizing effect of the feedforward network. This property introduces a problem for the selectivity of the feedforward network. High background activity can de-stabilize the ground state by spontaneously activating the feedforward network by random activity transients and, thereby, create a false-positive response (Tetzlaff et al. 2002). We have shown here how incorporating inhibition into the feedforward network can solve this problem. The correlated inhibition drastically changed the conditions for propagating signals through the feedforward network. Synchronous inputs to the first group could propagate with or without correlated inhibition. However, strong asynchronous inputs such as those inducing propagation dynamics in a synfire chain without FFI, failed to do so with FFI, thereby increasing the selectivity of the chain. Similar findings were recently reported by Vogels and Abbott (2009). These authors studied rate propagation from one sender group through a receiver group. Similar to the architecture used here, the

receiving group consisted of excitatory and inhibitory neurons, thus excitation and inhibition were balanced in the receiving group. The authors showed that in this configuration, rates could not propagate through the receiving group. Only when the gain of the inhibitory neurons was reduced, such that excitation was no longer balanced by inhibition, the firing rate in the receiving neurons could be elevated. Only at stimulus onset, when responses were more synchronized, the receiving group responded transiently, thus confirming our results.

Interestingly, the integration time window, that is shaped by the time lag between excitation and inhibition, controlled the selectivity of the synfire chains. Thereby, this parameter gives the possibility to dynamically regulate the signal flow in cortical networks as it can be controlled by delaying or advancing the spiking of the inhibitory neurons. We have shown that recurrent networks with effective inhibition provide a substrate in which synchronous spiking activity propagates more robustly. The stabilizing effect of the effective inhibition results in low rate background regimes in which synchrony can propagate through an embedded signal path and be easily detected in later groups (Kumar et al. 2008a). Thus, effective inhibition, both on the global scale and in the signal path, facilitates the propagation of synchrony. The two observations, that effective inhibition in the random network results in asynchronous low rate activity, while effective FFI in the signal path induces selectivity for synchronous inputs, may seem contradictory at first sight. However, the difference is that the response of the embedding network reflects the overall dynamical behavior of the system, whereas the feedforward network rather demonstrates signal propagation through a structured network. In the case of the random network, effective inhibition prevents run-away excitatory processes and, thus, stabilizes the system. In the feedforward network, FFI induces the short integration time window and, thereby, makes the transmission more selective.

Furthermore, it was shown that classical feedforward networks, which share a minimum number of group members by coexisting in the same cortical network space, may interact and activate each other (Kumar et al. 2008a). This mutual activation is an interesting concept for compositionality (Abeles et al. 2004; Schrader et al. 2007). However, it necessitates that the mutual activation is selective, such that many feedforward networks can coexist in the same cortical network, in a stable manner and without necessarily activating each other by undesired cross-talks. In this context, embedded FFI could increase the memory capacity, i.e. the number of possible chains that can be embedded in

the network. More work, however, is needed to clarify how, or indeed whether, such synfire dynamics can support computation in cortical networks.

Our goal was to investigate how a balanced sequence of excitation and inhibition can shape spike-pattern output in sensory cortical neurons and help in transmitting this information in a reliable manner. We investigated this at different scales (single neurons and different network configurations) and showed that balanced sequences are helpful in propagating information along multiple processing stages. Moreover, it has been suggested that such balanced sequences may help in processing sensory information and improving detection and discrimination performance. For instance, it was shown that such excitation/inhibition sequences may produce a computationally efficient neural code in the visual system, latency rank order coding (Delorme 2003). This opens the possibility for future investigations into its functional role when the incoming flow is modulated by local oscillatory activity and top-down inputs (see Tiesinga et al. 2008).

5 Conclusion

In summary, the reported results add evidence that correlated excitation and inhibition found *in vivo* are important for computational processes in cortical circuits. In particular, our study suggests that the brain takes advantages of a simple architecture to filter, process and propagate synchronous events in complex networks. By investigating how these properties are important in a more functional context, such as processing in the visual system (Kremkow et al. 2008b), future work will be able to address the functional impact of such architecture onto the dynamics of information processing.

Acknowledgements For helpful discussions we thank Yves Fregnac and Arvind Kumar, the latter also for his careful reading of the manuscript. We thank the reviewers for helpful suggestions. This work was supported by the German Federal Ministry of Education and Research (BMBF grant 01GQ0420 to BCCN Freiburg), by the German Research Council (DFG SFB-780), by the CNRS and the 6th RFP of the EU (grant no. 15879-FACETS).

References

- Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex*. Cambridge University Press.
- Abeles, M., Hayon, G., & Lehmann, D. (2004). Modeling compositionality by dynamic binding of synfire chains. *Journal of Computational Neuroscience*, 17(2), 179–201.
- Aertsen, A., Diesmann, M., & Gewaltig, M.-O. (1996). Propagation of synchronous spiking activity in feedforward neural networks. *Journal of Physiology (Paris)*, 90(3–4), 243–247.
- Assisi, C. G., Stopfer, M., Laurent, G., & Bazhenov, M. (2007). Adaptive regulation of sparseness by feedforward inhibition. *Nature Neuroscience*, 10(9), 1176–1184.
- Atallah, B. V., & Scanziani, M. (2009). Instantaneous modulation of gamma oscillation frequency by balancing excitation with inhibition. *Neuron*, 62(4), 566–577.
- Aviel, Y., Mehring, C., Abeles, M., & Horn, D. (2003). On embedding synfire chains in a balanced network. *Neural Computation*, 15(6), 1321–1340.
- Brémaud, A., West, D. C., & Thomson, A. M. (2007). Binomial parameters differ across neocortical layers and with different classes of connections in adult rat and cat neocortex. *Proceedings of the National Academy of Sciences of the United States of America*, 104(35), 14134–14139.
- Braitenberg, V., & Schüz, A. (1991). *Cortex: Anatomy of the cortex: Statistics and geometry*. Springer.
- Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8(3), 183–208.
- Bruno, R. M., & Sakmann, B. (2006). Cortex is driven by weak but synchronously active thalamocortical synapses. *Science*, 312(5780), 1622–1627.
- Buzsáki, G. (1984). Feed-forward inhibition in the hippocampal formation. *Progress in Neurobiology*, 22(2), 131–153.
- Cruikshank, S. J., Lewis, T. J., & Connors, B. W. (2007). Synaptic basis for intense thalamocortical activation of feedforward inhibitory cells in neocortex. *Nature Neuroscience*, 10(4), 462–468.
- Davison, A. P., Brüderle, D., Eppler, J. M., Kremkow, J., Müller, E., Pecevski, D., et al. (2009). PyNN: A common interface for neuronal network simulators. *Frontiers in Neuroinformatics*, 2, 11. doi:10.3389/neuro.11/011.2008.
- Delorme, A. (2003). Early cortical orientation selectivity: How fast inhibition decodes the order of spike latencies. *Journal of Computational Neuroscience*, 15(3), 357–365.
- Destexhe, A., Contreras, D., & Steriade, M. (1998). Mechanisms underlying the synchronizing action of corticothalamic feedback through inhibition of thalamic relay cells. *Journal of Neurophysiology*, 79(2), 999–1016.
- Destexhe, A., Rudolph, M., & Paré, D. (2003). The high-conductance state of neocortical neurons *in vivo*. *Nature Reviews Neuroscience*, 4(9), 739–751.
- Diesmann, M., Gewaltig, M.-O., & Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402(6761), 529–533.
- Eppler, J. M., Helias, M., Müller, E., Diesmann, M., & Gewaltig, M.-O. (2009). PyNEST: A convenient interface to the nest simulator. *Frontiers in Neuroinformatics*, 2, 12. doi:10.3389/neuro.11/012.2008.
- Gerstein, G., & Mandelbrot, B. (1964). Random walk Models for the spike activity of a single neuron. *Biophysical Journal*, 4, 41–68.
- Gewaltig, M.-O., & Diesmann, M. (2007). NEST (neural simulation tool). *Scholarpedia*, 2(4), 1430.
- Gewaltig, M.-O., Diesmann, M., & Aertsen, A. (2001). Propagation of cortical synfire activity: Survival probability in single trials and stability in the mean. *Neural Networks*, 14(6–7), 657–673.
- Gibson, J. R., Beierlein, M., & Connors, B. W. (1999). Two networks of electrically coupled inhibitory neurons in neocortex. *Nature*, 402(6757), 75–79.
- Haider, B., Krause, M. R., Duque, A., Yu, Y., Touryan, J., Mazer, J. A., et al. (2010). Synaptic and network mechanisms of

- sparse and reliable visual cortical activity during nonclassical receptive field stimulation. *Neuron*, 65(1), 107–121.
- Hasenstaub, A. R., Shu, Y., Haider, B., Kraushaar, U., Duque, A., & McCormick, D. (2005). Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron*, 47(3), 423–435.
- Higley, M. J., & Contreras, D. (2006). Balanced excitation and inhibition determine spike timing during frequency adaptation. *Journal of Neuroscience*, 26(2), 448–457.
- Hirsch, J. A., & Gilbert, C. D. (1991). Synaptic physiology of horizontal connections in the cat's visual cortex. *Journal of Neuroscience*, 11(6), 1800–1809.
- Hellwig, B. (2000). A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biological Cybernetics*, 82(2), 111–121.
- Inoue, T., & Imoto, K. (2006). Feedforward inhibitory connections from multiple thalamic cells to multiple regular-spiking cells in layer 4 of the somatosensory cortex. *Journal of Neurophysiology*, 96(4), 1746–1754.
- Kapfer, C., Glickfeld, L. L., Atallah, B. V., & Scanziani, M. (2007). Supralinear increase of recurrent inhibition during sparse activity in the somatosensory cortex. *Nature Neuroscience*, 10(6), 743–753.
- Kremkow, J., Perrinet, L., Aertsen, A., Masson, G. S. (2008a). Functional properties of feed-forward inhibition. *Proc NeuroComp 2008*
- Kremkow, J., Perrinet, L., Baudot, P., Levy, M., Marre, O., Monier, C. et al. (2008b). Control of the temporal interplay between excitation and inhibition by the statistics of visual input: A V1 network modelling study. Vol. Soc. Neurosci. Abstr. (p. 366.5/II10).
- Kremkow, J., Perrinet, L., Masson, G. S., & Aertsen, A. (2009). Functional consequences of correlated excitation and inhibition on single neuron integration and signal propagation through synfire chains. *Proceedings of the 32nd Göttingen Neurobiology Conference T26-6B*.
- Kuhn, A., Aertsen, A., & Rotter S. (2004). Neuronal integration of synaptic input in the fluctuation-driven regime. *Journal of Neuroscience*, 24(10), 2345–2356.
- Kumar, A., Rotter, S., & Aertsen, A. (2008a). Conditions for propagating synchronous spiking and asynchronous firing rates in a cortical network model. *Journal of Neuroscience*, 28(20), 5268–5280.
- Kumar, A., Schrader, S., Aertsen, A., & Rotter, S. (2008b). The high-conductance state of cortical networks. *Neural Computation*, 20(1), 1–43.
- Kumbhani, R. D., Nolt, M. J., & Palmer, S. E., (2007). Precision, reliability, and information-theoretic analysis of visual thalamocortical neurons. *Journal of Neurophysiology*, 98(5), 2647–2663.
- Lampl, I., Reichova, I., & Ferster, D. (1999). Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron*, 22(2), 361–374.
- Litvak, V., Sompolinsky, H., Segev, I., & Abeles, M. (2003). On the transmission of rate code in long feedforward networks with excitatory-inhibitory balance. *Journal of Neuroscience*, 23(7), 3006–3015.
- Mainen, Z. F., & Sejnowski, T. J., (1995). Reliability of spike timing in neocortical neurons. *Science*, 268(5216), 1503–1506.
- Marre, O., Baudot, P., Levy, M., & Frégnac, Y. (2005). High timing precision and reliability, low redundancy and low entropy code in V1 neurons during visual processing of natural scenes. *Society for Neuroscience Abstracts*, 31, 285.5.
- Mehring, C., Hehl, U., Kubo, M., Diesmann, M., & Aertsen, A. (2003). Activity dynamics and propagation of synchronous spiking in locally connected random networks. *Biological Cybernetics*, 88(5), 395–408.
- Molnár, G., Oláh, S., Komlósi, G., Füle, M., Szabadics, J., Varga, C., et al. (2008). Complex events initiated by individual spikes in the human cerebral cortex. *PLoS Biol*, 6(9), e222.
- Morrison, A., Mehring, C., Geisel, T., Aertsen, A., & Diesmann, M. (2005). Advancing the boundaries of high-connectivity network simulation with distributed computing. *Neural Computation*, 17(8), 1776–1801.
- Muller, E., Buesing, L., Schemmel, J., & Meier, K. (2007). Spike-frequency adapting neural ensembles: Beyond mean adaptation and renewal theories. *Neural Computation*, 19(11), 2958–3010.
- Nawrot, M. P., Boucsein, C., Molina, V. R., Riehle, A., Aertsen, A., & Rotter, S. (2008). Measurement of variability dynamics in cortical spike trains. *Journal of Neuroscience Methods*, 169(2), 374–390.
- Nowak, L. G., Azouz, R., Sanchez-Vives, M. V., Gray, C. M., & McCormick, D. A. (2003). Electrophysiological classes of cat primary visual cortical neurons *in vivo* as revealed by quantitative analyses. *Journal of Neurophysiology*, 89(3), 1541–1566.
- Okun, M., & Lampl, I. (2008). Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nature Neuroscience*, 11(5), 535–537.
- Pinto, D. J., Hartings, J. A., Brumberg, J. C., & Simons, D. J. (2003). Cortical damping: Analysis of thalamocortical response transformations in rodent barrel cortex. *Cerebral Cortex*, 13(1), 33–44.
- Pouille, F., & Scanziani, M. (2001). Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science*, 293(5532), 1159–1163.
- Povysheva, N. V., Gonzalez-Burgos, G., Zaitsev, A. V., Kröner, S., Barrionuevo, G., Lewis, D. A., & Krimer, L. S. (2006). Properties of excitatory synaptic responses in fast-spiking interneurons and pyramidal cells from monkey and rat prefrontal cortex. *Cerebral Cortex*, 16(4), 541–552.
- Rudolph, M., Pospischil, M., Timofeev, I., & Destexhe, A. (2007). Inhibition determines membrane potential dynamics and controls action potential generation in awake and sleeping cat cortex. *Journal of Neuroscience*, 27(20), 5280–5290.
- Schrader, S., Morrison, A., & Diesmann, M. (2007). A composition machine for complex movements. *Proceedings of the 31st Göttingen Neurobiology Conference TS18-1C*.
- Shadlen, M. N., & Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Current Opinion in Neurobiology*, 4(4), 569–579.
- Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *Journal of Neuroscience*, 18(10), 3870–3896.
- Silberberg, G., & Markram, H. (2007). Disynaptic inhibition between neocortical pyramidal cells mediated by martinotti cells. *Neuron*, 53(5), 735–746.
- Smith, M., & Kohn, A. (2008). Spatial and temporal scales of neuronal correlation in primary visual cortex. *Journal of Neuroscience*, 28(48), 12591–12603.
- Somers, D. C., Nelson, S. B., & Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. *Journal of Neuroscience*, 15(8), 5448–5465.
- Stepanyants, A., Hirsch, J., Martinez, L. M., Kisvárdy, Z. F., Ferecskó, A. S., & Chklovskii, D-B. (2008). Local potential connectivity in cat primary visual cortex. *Cerebral Cortex*, 18(1), 13–28.

- Swadlow, H. A. (2003). Fast-spike interneurons and feedforward inhibition in awake sensory neocortex. *Cerebral Cortex*, *13*(1), 25–32.
- Tetzlaff, T., Geisel, T., & Diesmann, M. (2002). The ground state of cortical feed-forward networks. *Neurocomputing*, *44–46*, 673–678.
- Thomson, A. M., West, D. C., Wang, Y., & Bannister, A. P. (2002). Synaptic connections and small circuits involving excitatory and inhibitory neurons in layers 2–5 of adult rat and cat neocortex: Triple intracellular recordings and biocytin labelling *in vitro*. *Cerebral Cortex*, *12*(9), 936–953.
- Tiesinga, P., Fellous, J.-M., & Sejnowski, T. J. (2008). Regulation of spike timing in visual cortical circuits. *Nature Reviews Neuroscience*, *9*(2), 97–107.
- Troyer, T. W., Krukowski, A. E., Priebe, N. J., & Miller, K. D. (1998). Contrast-invariant orientation tuning in cat visual cortex: Thalamocortical input tuning and correlation-based intracortical connectivity. *Journal of Neuroscience*, *18*(15), 5908–5927.
- Tucker, T. R., & Katz, L. C. (2003a). Recruitment of local inhibitory networks by horizontal connections in layer 2/3 of ferret visual cortex. *Journal of Neurophysiology*, *89*(1), 501–512.
- Tucker, T. R., & Katz, L. C. (2003b). Spatiotemporal patterns of excitation and inhibition evoked by the horizontal network in layer 2/3 of ferret visual cortex. *Journal of Neurophysiology*, *89*(1), 488–500.
- van Vreeswijk, C., & Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, *274*(5293), 1724–1746.
- Vogels, T. P., & Abbott, L. F. (2009). Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nature Neuroscience*, *12*(4), 483–491.
- Wehr, M. S., & Zador, A. M. (2003). Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature*, *426*(6965), 442–446.
- Yger, P., Bruderle, D., Eppler, J., Kremkow J., Pevovski, D., Perrinet, L., et al. (2009). NeuralEnsemble: Towards a meta-environment for network modeling and data analysis. *Eight Göttingen Meeting of the German neuroscience society* (pp. T26–4C). <http://www.incm.cnrs-mrs.fr/LaurentPerrinet/Publications/Yger09gns>.