J Comput Neurosci
DOI 10.1007/s10827-010-0280-1

Spiking neurons that keep the rhythm

Jean-Philippe Thivierge · Paul Cisek

Received: 11 July 2009 / Revised: 16 September 2010 / Accepted: 21 September 2010

Abstract Detecting the temporal relationship among events in the environment is a fundamental goal of the brain. Following pulses of rhythmic stimuli, neurons of the retina and cortex produce activity that closely approximates the timing of an omitted pulse. This omitted stimulus response (OSR) is generally interpreted as a transient response to rhythmic input and is thought to form a basis of short-term perceptual memories. Despite its ubiquity across species and experimental protocols, the mechanisms underlying OSRs remain poorly understood. In particular, the highly transient nature of OSRs, typically limited to a single cycle after stimulation, cannot be explained by a simple mechanism that would remain locked to the frequency of stimulation. Here, we describe a set of realistic simulations that capture OSRs over a range of stimulation frequencies matching experimental work. The model does not require an explicit mechanism for learning temporal sequences. Instead, it relies on spike timing-dependent plasticity (STDP), a form of synaptic modification that is sensitive to the timing of pre- and post-synaptic action potentials. In the model, the transient nature of OSRs is attributed to the heterogeneous nature of neural properties and connections, creating intricate forms of activity that are continuously changing over time. Combined with STDP, neural heterogeneity enabled OSRs to complex rhythmic patterns as well as OSRs following a delay period. These results link the response of neurons to rhythmic patterns with the capacity of heterogeneous circuits to produce transient and highly flexible forms of neural activity.

Keywords Omitted stimulus potential · Synchronization · Neural diversity · Spike timing-dependent plasticity · Integrate-and-fire neurons · Short-term memory

Action Editor: X.-J. Wang

Electronic supplementary material The online version of this article (doi:10.1007/s10827-010-0280-1) contains supplementary material, which is available to authorized users.

J.-P. Thivierge (✉)
Department of Psychological and Brain Sciences,
Indiana University,
1101 East Tenth Street,
Bloomington, IN 47405, USA
e-mail: jthivier@indiana.edu

P. Cisek
Groupe de Recherche sur le Système Nerveux Central,
Département de Physiologie, Université de Montréal,
Montreal, QC, Canada H3T 1J4

1 Introduction

An essential requirement for perception and cognition is the ability of neural circuits to maintain information about sensory events beyond their time of occurrence. Following a rhythmic stimulation, neurons in several locations—including retina (Schwartz et al. 2007) and cortex (Demiralp et al. 1994)—produce action potentials that are closely timed to the frequency of stimulation when the next event should have occurred. This "omitted stimulus response" (OSR) is generally interpreted as a transient response to rhythmic input and is thought to form a basis of short-term perceptual memory for temporal sequences (Friedman et al. 2001; Jaaskelainen et al. 2004). OSRs are found across a wide range of species, including both vertebrates and invertebrates (Ramon and Gronenberg 2005), and a related phenomenon is reported in human scalp recordings (Demiralp and Basar 1992).

Published online: 01 October 2010

Despite the prevalence of OSRs and their role as a carrier of important sensory information, their origin and underlying biophysical mechanisms are unclear. In particular, OSRs are typically limited to a single cycle of activity following stimulation (Schwartz et al. 2007). Afterwards, neurons return to a spontaneous state of largely unpredictable activity, synchronizing in a transient and nonperiodic fashion (Eytan and Marom 2006). This makes OSRs incompatible with models that get "stuck" in a stable attractor state where they reverberate at a fixed frequency (Rabinovich and Abarbanel 1998). A plausible alternative is that OSRs may be generated by recurrent networks with transient dynamics (Ganguli et al. 2008; Jaeger and Haas 2004; Maass et al. 2002; Rabinovich et al. 2008).

Here, we show how a heterogeneous network of spiking neurons (where individual neurons differ in their intrinsic properties) can potentially explain OSRs, producing temporally precise responses to rhythmic signals, and quickly returning to a state of nonperiodic spontaneous activity. The proposed model is based on integrate-and-fire neurons with synaptic delays (see Section 2). No explicit mechanism was required for the model to detect rhythmic patterns; rather, it relied on a combination of reverberating activity and rapid synaptic modification (Hebb 1949; Thivierge and Cisek 2008). A central element of the theory is spike timing-dependent plasticity (STDP), whereby the precise timing of pre- and post-synaptic action potentials determines the degree and direction of change in synaptic efficacy (Abbott and Nelson 2000; Gerstner et al. 1996; Kempter et al. 1999; Markram et al. 1997; Morrison et al. 2007).

Several parameters of the model, reflecting intrinsic cellular properties, were initialized randomly for every cell, thus generating a diverse population of neurons. This aspect of the model was designed to capture the diverse nature of neuronal circuits; for instance, the retina is composed of over fifty-five neuronal types, with different degrees of neuronal excitability (Masland 2001). In cortex, neural heterogeneity in found in the morphology, functional properties, intrinsic biophysical properties, and connectivity of neurons (particularly interneurons, Buzsaki et al. 2004). As a consequence of heterogeneity, neural circuits produce intricate forms of neural activity, giving patterns of spikes an intrinsically "noisy" appearance (Thivierge et al. 2007; Vogels et al. 2005). Here, we suggest that the intricate patterns of activity produced in heterogeneous circuits play a central role in generating transient yet highly precise responses to temporal patterns. Rather than providing a detailed account of retinal or cortical circuitry, our goal is to provide a more general hypothesis of potential mechanisms responsible for generating transient, yet highly precise, responses to rhythmic stimulation.

2 Methods

2.1 Description of the model

The computational model employed here follows similar principles as previous work (Brunel 2000; Brunel and Hakim 1999; Thivierge and Cisek 2008). Briefly, it consists of a population of interconnected neurons (N=100) whose membrane potential (V_i) follows a deterministic integrate-and-fire equation:

$$c_m \frac{dV_i}{dt} = -g_i(V_i - E_i) + \sum_{j=1}^{N} w_{ij} K_j + R(I_{tonic}) \quad (1)$$

where neurons are indexed from i=1,...,N. The model's parameters are as follows: c_m is the membrane capacitance, g is the leak conductance, w_{ij} is the unit-less weight, R is a unit-less scalar gain, E_i is a reversal potential, and I_{tonic} is a tonic current. K_j is the excitatory current produced by incoming spikes (Gerstner and Kistler 2002) (Fig. 1(a)):

$$K_j = V_0 \cdot \sum_{i=1}^{S} \exp\left(\frac{t_i - t}{\tau_{fall}}\right) - \exp\left(\frac{t_i - t}{\tau_{rise}}\right) \quad (2)$$

where t_i is the spike times of the i^{th} afferent and V_0 is a free parameter. Equation (2) considers a maximum of 10 previous spikes (indexed by i= 1,...,S). Because K_j satisfies a second order ODE, an efficient way to update Eq. (2) would be to keep track of fluctuations in $K_j(t)$ over time, and update it every time there is a spike (Matlab code accompanying this article could be modified for that purpose). As in previous work (Izhikevich 2006), each neuron possessed a unique delay of synaptic transmission, drawn from a Gaussian distribution with a mean of 22 ms and a standard deviation of 0.33. This distribution reflects a range of conductance delays, as reported in cortex (Swadlow 1985).

A spike occurs if V_i exceeds the spike threshold (V_{thresh}). At that time, V_i is set to V_{spike} for a duration of T_{spike}. Afterwards, V_i is reset to V_{rest} and held there for a duration of T_{reset}. Membrane potentials are initialized to random values (range: 0–1).

Unless otherwise stated, synaptic efficacies (Δw_{ij}) are adjusted according to spike timing-dependent plasticity (STDP) (Abbott and Nelson 2000):

$$\Delta w_{ij}(t) = \begin{cases} W_+ \cdot \exp\left(-\Delta t_{ij}/\tau_+\right) & \text{if } \Delta t_{ij} > 0 \\ -W_- \cdot \exp\left(\Delta t_{ij}/\tau_-\right) & \text{if } \Delta t_{ij} \leq 0 \end{cases} \quad (3)$$

where $\Delta t_{ij} = t_j - t_i$ reflects the time difference between the last presynaptic (t_i) and postsynaptic (t_j) spikes; W_+ and W_- control the magnitude of change in synaptic efficacy; τ_+ and τ_- control the time-course of plasticity. We explored the consequences of both a symmetrical (LTD ≈ LTP)

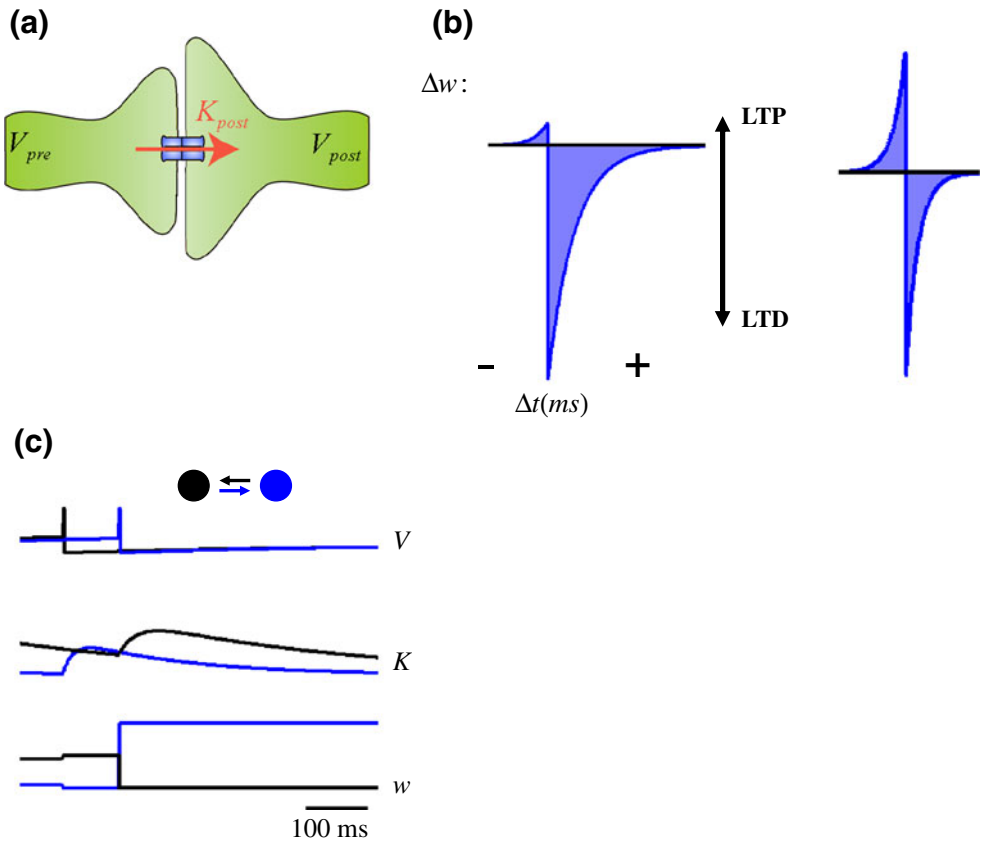

Fig. 1 Physiology of the network. (**a**) Model of synapse with pre- and post-synaptic membrane potentials (V_{pre} and V_{post}), and after-spike excitatory potential (K_{post}). (**b**) Synaptic strengths are altered by spike timing-dependent plasticity (STDP), whereby the difference in pre- and post-synaptic spike arrival times (Δt) determines the change in strength (Δw). Both a symmetrical (LTD=LTP) and an asymmetrical (LTD>LTP) version of the STDP rule are shown. (**c**) Membrane potential (V), after-spike excitatory potential (K), and synaptic strength (w) of two neurons. Changes in synaptic strength were obtained with symmetrical STDP

and asymmetrical (LTD>LTP) version of the STDP rule (Fig. 1(b)). Unless otherwise stated, simulations were performed with an asymmetrical learning rule. Synaptic weights were initialized to random values (range: 0–10) with no self-connections, and were updated as follows:

$$w_{ij}(t) = w_{ij}(t-1) + \eta \Delta w_{ij}(t) \tag{4}$$

where η is a free parameter. Unless otherwise stated, weights were not allowed to take on negative values (i.e., lower bound of zero) and had no upper bound. The proposed model makes no assumption as to the anatomical origins of OSRs; experiments show that STDP is found in several anatomically distinct locations, including both retinotectal projections (Zhang et al. 1998) as well as cortex (Feldman 2000).

A heterogeneous population of neurons—where different neurons have different intrinsic properties—was generated by randomly drawing some of their parametric values from a Gaussian distribution, with means and standard deviations reported in Table 1. This cellular heterogeneity is consistent with experimental evidence showing that cells in retina (Masland 2001) and cortex (Buzsaki et al. 2004) possess an impressive diversity of intrinsic biophysical properties. Unless otherwise noted, different runs of the model re-initialized the intrinsic values of each neuron; these values remained constant for the duration of a given run. The interplay between the membrane potentials, incoming excitatory currents, and synaptic weights of two interconnected neurons is illustrated in Fig. 1(c).

No stochasticity was included in the model, other than for the random initialization of certain parameters (see Table 1). This approach differs from many studies of signal propagation that inject random current into the membrane potential of all neurons (Diesmann et al. 1999; van Rossum et al. 2002). While several models of internal noise have

Table 1 Default values of parameters and distributions

Spike parameters	V_{rest}=1 mV, T_{reset}=3 ms, V_{spike}=100 mV, T_{spike}=1 ms
STDP parameters	Asymmetrical: W_+=10, W_-=100, τ_-=60 ms, τ_+=30 ms. Symmetrical: W_+=5, W_-=8, τ_+=20 ms, τ_-=20 ms.
Means of parameter distributions (σ=0.33)	τ_{fall}=200 ms, E=−65 mV, τ_{rise}=10 ms, g=0.01 pS, V_{thresh}=18 mV
Misc. parameters	I_{tonic}=1 nA, V_0=0.091, η=0.5, τ=0.02, N=100 cells, R=10

been proposed (Brunel 2000; Mehring et al. 2003; van Vreeswijk and Sompolinsky 1996; Vogels et al. 2005), it has not been shown how a simulated neuronal circuit with intrinsic noise can respond in a highly precise, yet transient manner to rhythmic patterns. This problem, which is the primary objective of our study, is by no means trivial. In contrast to models with external noise, internal noise is fixed by the network and cannot be readily adjusted. To argue for the robustness of our model, we also show that the addition of noise in membrane potential does not appreciably change our results (see Section 3).

In one variant of the model, inhibitory neurons were included by randomly selecting a subset of neurons and initializing their outgoing synaptic weights to $w_{ij}=-10$. These weights could not be modified during the simulation. In a second variant, we generated a homogeneous network where all neurons were identical copies of each other. This was achieved by initializing all neurons to the mean parametric values of Table 1. In that version of the model, synapses all had the same conductance delays (fixed to 2 ms). Finally, we also designed a variant of the model where a subpopulation of neurons was homogeneous (i.e., all had the same parametric values), while another subpopulation was heterogeneous.

The robustness of neural dynamics to several factors, including 1) the addition of Gaussian noise to the membrane potential of Eq. (1); 2) the addition of inhibitory neurons; 3) the random deletion of synaptic connections; and 4) delaying the application of changes in the induction of STDP are addressed elsewhere (Thivierge and Cisek 2008). In addition, previous work has canvassed a wide range of simulation parameters (Table 1), showing conditions under which spontaneous and activity-driven dynamics are comparable or different. As described in Section 3 below, the heterogeneous nature of the model suggests that several of its parameters can be altered (within a certain range) without a drastic effect on the behavior of the model. One constraint, as discussed in other work (Kempter et al. 2001), is that the proportion of LTD should be greater than that of LTP in order to prevent runaway excitation in STDP. It has been suggested that $W_{-}\tau_{-}/W_{+}\tau_{+}$ be greater than 1; in our model, the average values of these parameters respect that rule (see Table 1). Unless otherwise stated, all simulations were performed using Matlab software with an integration time-step of 1 ms.

3 Results

3.1 Detecting patterns of rhythmic activity

In a first series of simulations, all neurons in the network received a conditioning stimulation (CS) of 12 pulses with a stimulus period of 100 ms, where each pulse lasted 1 ms and forced all neurons to depolarize (i.e., $V_i>V_{thresh}$). Immediately afterwards, most neurons emitted spikes that were closely timed to the expected occurrence of the next pulse, where the CS was absent (relative response latency averaged over all neurons: 87.63, SEM: 7.99) (Fig. 2(a)). In fact, this response to an omitted stimulus is already foreshadowed by the observation that after several rounds of stimulation, many neurons in the network emit spikes that anticipate the next stimulation cycle (Fig. 2(a), note spikes occurring prior to the stimulation time, shown in red). Following their post-CS synchronization, spikes quickly fall out of phase, as consistent with a rapid decay. Further details of how modeled neurons transition between spontaneous and evoked activity can be found elsewhere (Thivierge and Cisek 2008). Simulations using a smaller integration step (0.1 ms instead of 1 ms) lead to the same qualitative behavior (relative response latency averaged over all neurons: 86.06, SEM: 6.81). For speed of processing all further results were obtained using an integration step of 1 ms.

To quantify the OSR of all neurons, we stored their first post-CS spikes and computed both the mean and standard error across the whole population. Applying these measures to a range of inter-stimulus intervals between 60 and 160 ms—as in recent experiments (Schwartz et al. 2007)—we found that OSRs across the whole population closely approximated the frequency of induced activity (Fig. 2(b)). This result was obtained despite the presence of intrinsic noise in the population dynamics. In complementary analyses, we explored conditions under which changes to the model can alter this result. In these conditions, we manipulated the degree of rhythmicity of the stimulus, the amount of neural heterogeneity, the amount of subthreshold synaptic noise, the addition of spike jittering, the distribution of synaptic strengths, the length of delays in the induction of STDP, and alternative simulation parameters. Taken together, these manipulations suggest that OSRs can occur under a wide range of conditions—with some specific limitations of the model defining the boundaries of these conditions. Details of these manipulations are reported below.

3.2 OSRs depend upon the rhythmicity of the stimulus

First, we examined whether the precision of OSRs is dependent upon the rhythmic nature of the CS. The precision of OSRs was disrupted in simulations where pulses are removed randomly (Fig. 3(a)) or where pulses are delivered at a random time intervals (Fig. 3(b)). These results suggest that, in the model, OSRs were sensitive to the rhythmic nature of the CS.

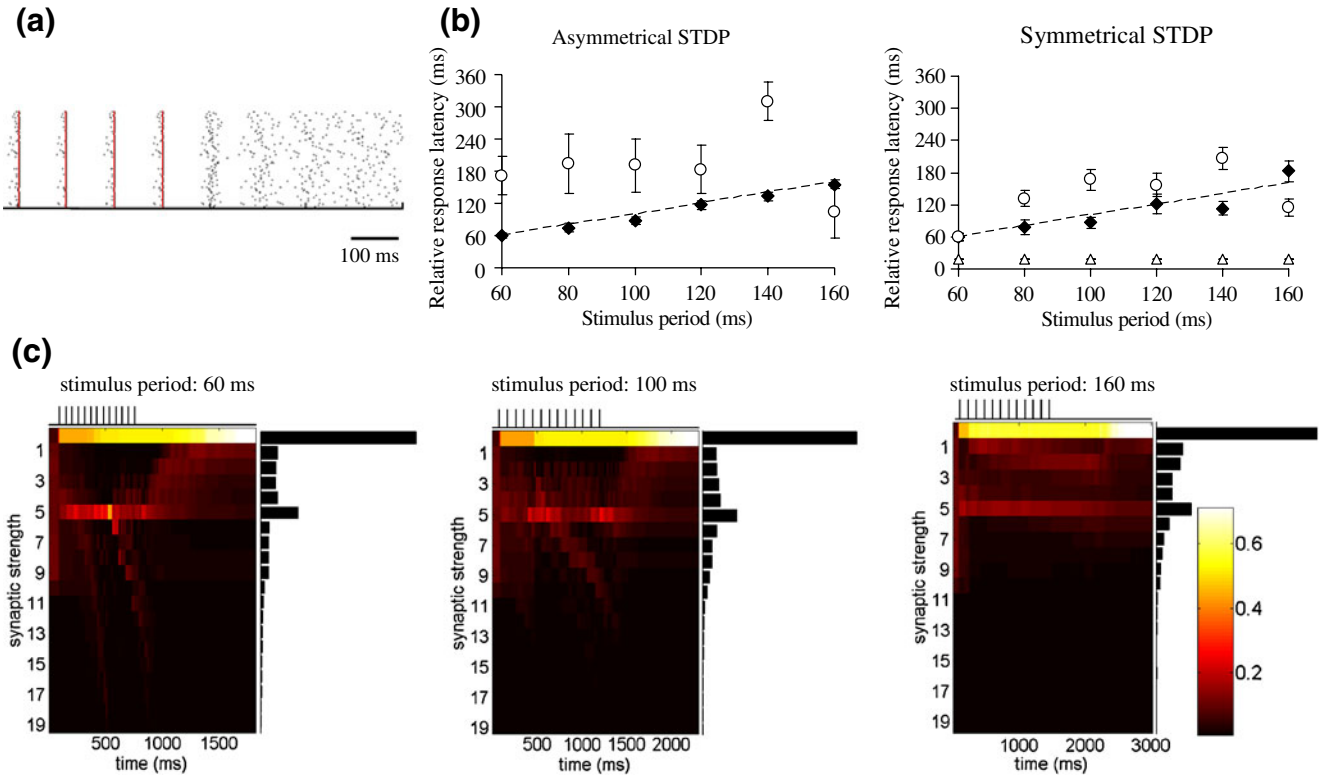

Fig. 2 Omitted stimulus response. (**a**) Spiking response of the model following 12 pulses of induced activity with stimulus period of 100 ms (shown in *red*; only the last 4 pulses are shown). (**b**) Response latency of spikes relative to stimulation offset. Results for different stimulus periods are shown (from 60 ms to 160 ms). *Black diamonds*: latency of first post-CS spikes. *White circles*: latency of second post-CS spikes. *White triangles*: latency of first post-CS spikes in a model with fixed connection strengths. Values are means over 10 independent runs of the model (with different initial conditions) and 100 neurons. *Vertical bars*: SEM. *Dashed line*: Slope of unity. *Left*, asymmetrical STDP. *Right*, symmetrical STDP. (**c**) Histograms showing the distribution of synaptic weights on a millisecond-by-millisecond basis with different stimulation protocols. In each graph, different shades of color represent the proportion of synaptic weights whose values range between [0,1], [1,2], [2,3], etc. (the y-axis represents upper values of those ranges). Results of three different simulations are shown, where the stimulus period is either 60 ms, 100 ms, or 160 ms. The times where pulses were delivered are shown as tick marks above each histogram. The mean proportion of synaptic weights (over the entire simulation) is shown in the *black bars* at the right of each histogram

3.3 The influence of neural heterogeneity on OSRs

Next, we examined the role of neural heterogeneity in generating OSRs. In the model, the intrinsic noise produced by cellular heterogeneity promotes efficient responses to stimulation by preventing the network from getting stuck in a particular state of highly periodic activity. This point is illustrated in a simulation that completely removed cellular heterogeneity by designing a network of identical neurons (see Section 2). Such a homogeneous network produced strictly periodic activity (Fig. 4(a)) (for more elaborate analyses of spontaneous activity, see Thivierge and Cisek 2008). The post-CS spiking of a homogeneous network did not follow the frequency of induced activity (Fig. 4(b), white triangles). While it might be possible to carefully select cell parameters that make a homogeneous network oscillate near the frequency of induced pulses, such parameters would be limited to a single frequency of oscillation. By contrast, a single heterogeneous model can approximate the frequency of induced activity over a range of frequencies. Of course, too much heterogeneity among cells hinders the precision of OSRs. When the distribution of parameters was widened by even a small margin (σ=0.4 compared to the default of σ=0.33), the model lost its ability to respond in a precise fashion following a train of rhythmic stimuli (Fig. 4(b), white squares).

While the above simulations suggest a functional role for neural diversity in generating OSRs, we must acknowledge other influential sources of heterogeneity found in biological circuits, including the intricate patterns of connectivity among neurons (Wang et al. 2006). A question that arises in our model is to what extent random connectivity and synaptic delays—with none of the other sources of heterogeneity of the default model (Table 1)—may provide a basis for precise yet transient OSRs. To address this question, we devised a homogeneous model with random

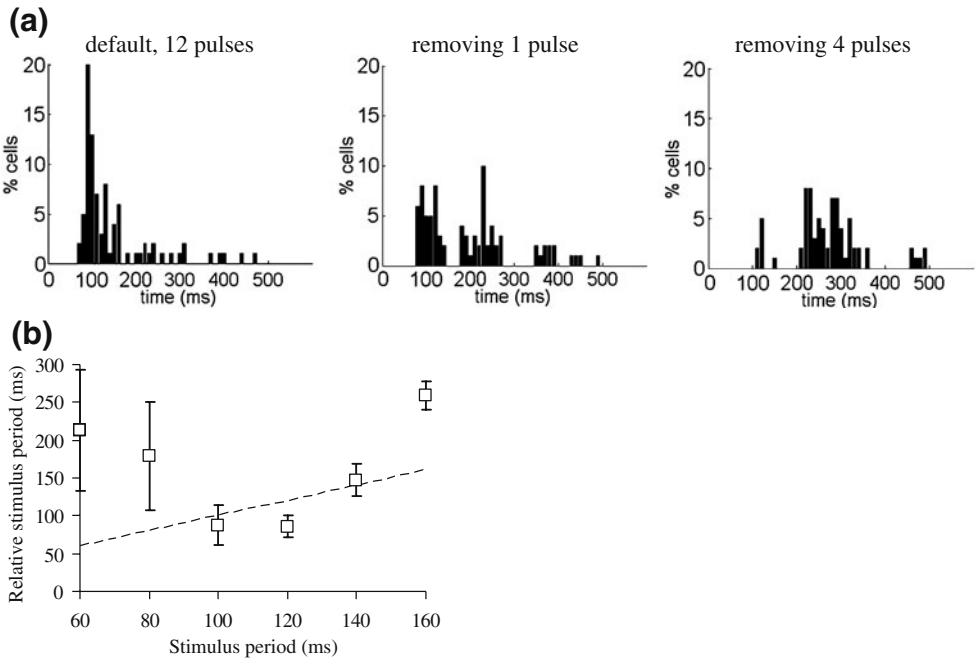

Fig. 3 The precision of omitted stimulus responses depends on the rhythmic nature of the conditioning stimulation. (**a**) Effect of randomly omitting pulses in a stimulation (default: 12 pulses with stimulus period of 100 ms). All three simulations began with the exact same initial conditions. Stimulus offset is at time 0. Only the first post-stimulation spikes of each cell are shown. (**b**) Effect of randomizing the stimulus periods of a conditioning stimulus. Stimulus periods are drawn from a Gaussian distribution with mean shown on the x-axis and standard deviation of σ=0.33 about the mean. Values are means over 10 independent runs of the model; the pulse intervals were randomized for every run of the model. *Dashed line*: slope of unity. *Vertical bars*: SEM

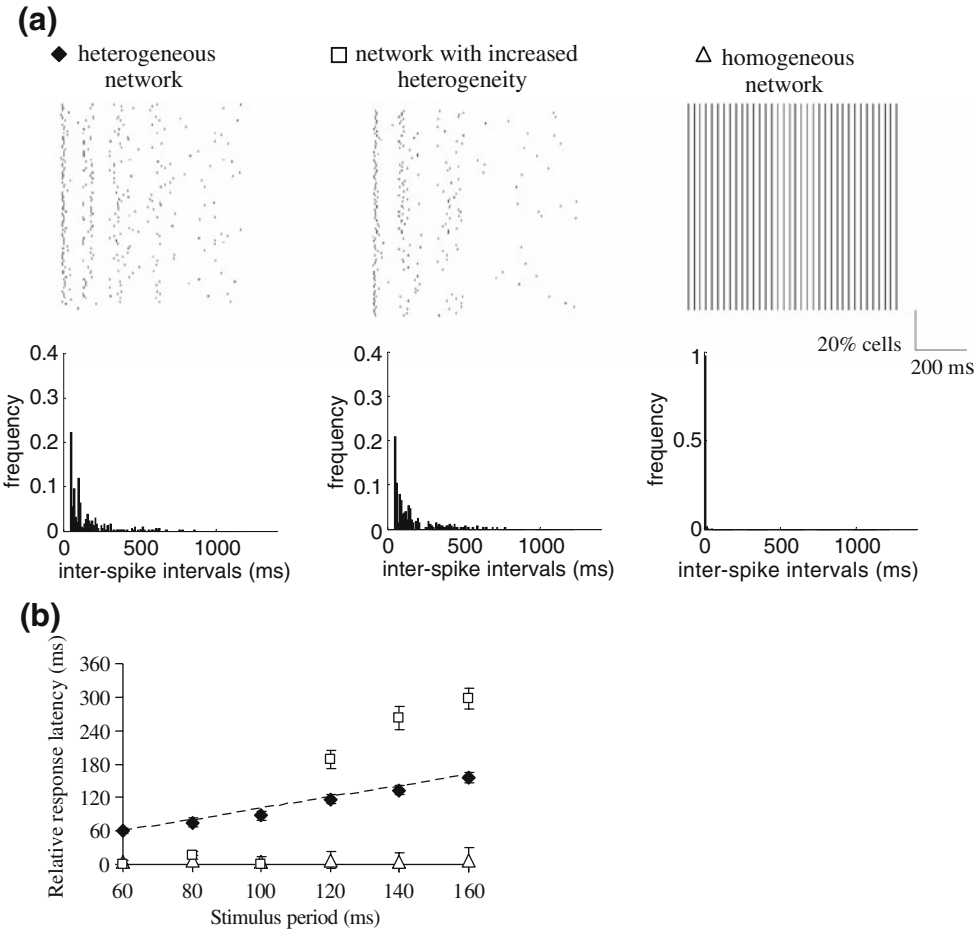

Fig. 4 The precision of OSRs is dependent upon the degree of neural heterogeneity in the model. (**a**) *Top*: representative spike raster for three variations of the model. *Bottom*: distribution of inter-spike intervals taken over 10 s of spontaneous activity (normalized to 1.0). (**b**) Latency of first post-CS spikes in three versions of the model for a range of stimulus periods (from 60 ms to 160 ms). Values are means over 10 independent runs of the model (with different initial conditions) and 100 neurons. *Vertical bars*: SEM. *Dashed line*: Slope of unity. *Black diamonds*: heterogeneous model. *White squares*: model with increased heterogeneity, obtained by increasing the width of the parametric distribution from σ=0.33 (default) to σ=0.4. *White triangles*: homogeneous model. In the heterogeneous network, each cell has a unique configuration of intrinsic properties; in the homogeneous network, all cells have the same properties (see Section 2)

connectivity and synaptic delays drawn from a Gaussian distribution with a mean of 22 ms and a standard deviation of 0.33. In this model, all cells had identical properties, and any two cells had a probability of 0.5 of being connected (no connections could be added or deleted during the simulation). Despite a homogeneity of cell properties, this model produced a reasonable approximation of different stimulus periods (Fig. 5, black diamonds). This result shows that, by itself, the heterogeneity of intrinsic cell properties may not be a necessary condition to generate flexible neural dynamics that can detect a rhythmic stimulation; other factors that prevent symmetries in the interactions among neurons may play a role.

One question that logically arises is how much heterogeneity is required in order to generate precise OSRs. This question is pertinent to cortical circuits, where not all types of neurons have the same degree of diversity. Excitatory cells, forming the principal population, are quite homogeneous; in contrast, inhibitory interneurons are exceptionally diverse in their morphological appearance, functional properties, intrinsic biophysical properties, and connectivity (Buzsaki et al. 2004). Can precise OSRs be produced in a model where inhibitory neurons are heterogeneous but where excitatory neurons are homogeneous? To address this question, we designed a version of the model where excitatory neurons (80% of the population) all had the same parameters, and inhibitory neurons (20% of the population) had a range of values for different intrinsic parameters. Other aspects of the model were as described in Section 2. This model was able to produce relatively precise OSRs for a range of induced rhythms (Fig. 5, white squares). While this simulation represents a highly simplified scenario, it nonetheless illustrates the idea that an entire network need not be heterogeneous in order to produce precise OSRs.

3.4 Addition of synaptic noise

In the default model (see Section 2), the "noisy" activity of neurons is caused by internal dynamics, as opposed to an external source of noise. Can the model produce OSRs that are temporally precise when an external source of noise is added? We addressed this question by adding Gaussian noise (λ) to the subthreshold membrane potential of neurons (Eq. (1)):

$$c_m \frac{dV_i}{dt} = -g_i(V_i - E_i) + \sum_{j=1}^{N} w_{ij}K_j + R(I_{tonic}) + \lambda. \qquad (5)$$

We designed two sets of simulations, where noise had a mean of zero and standard deviation of either 10 or 100 (Fig. 6(a)). Compared to simulations without noise (Fig. 2(b)), the addition of external noise reduced the precision of OSRs. This effect is manifested in two ways. First, the mean response latency of spikes is farther from the slope of unity

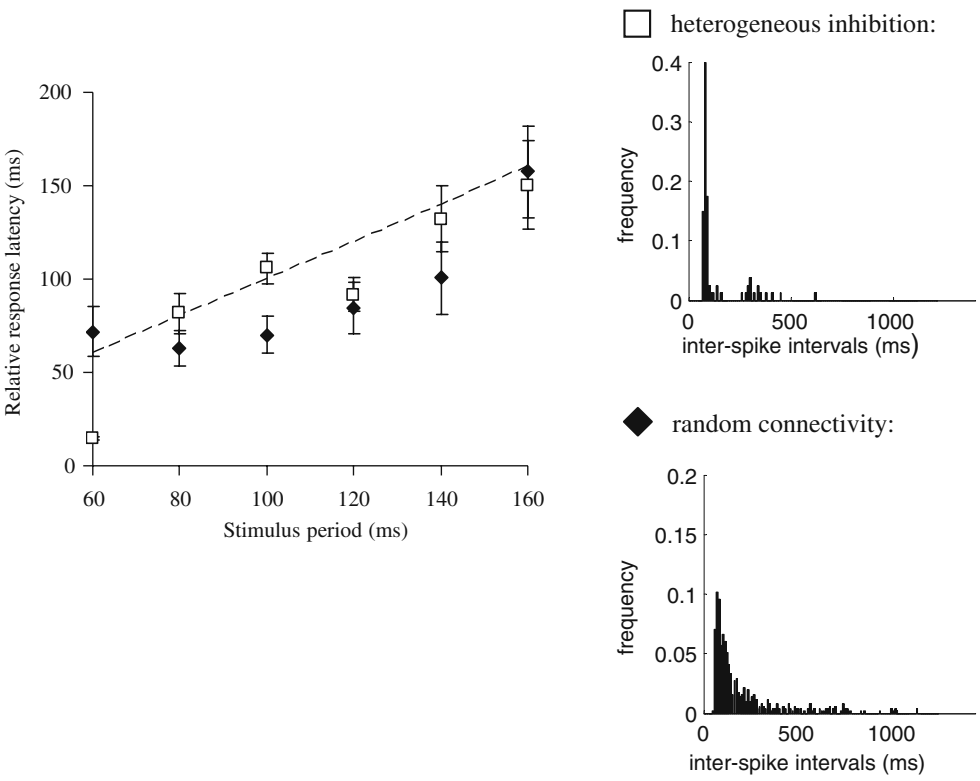

Fig. 5 Different sources of heterogeneity can provide a basis for OSRs. *Left*: relative response latency of first post-CS spikes for two models with different sources of heterogeneity. *White squares*: model where excitatory neurons (80%) are homogeneous and inhibitory neurons (20%) are heterogeneous. *Black diamonds*: model where all intrinsic properties of neurons are homogeneous, and the connectivity amongst neurons is random (with 50% probability of a connection between any two neurons). Neural connectivity is established at the beginning of each simulation; no synapses can be created or deleted during the simulation. *Right*: distribution of inter-spike intervals for 10 s of spontaneous activity (scaled between [0,1])

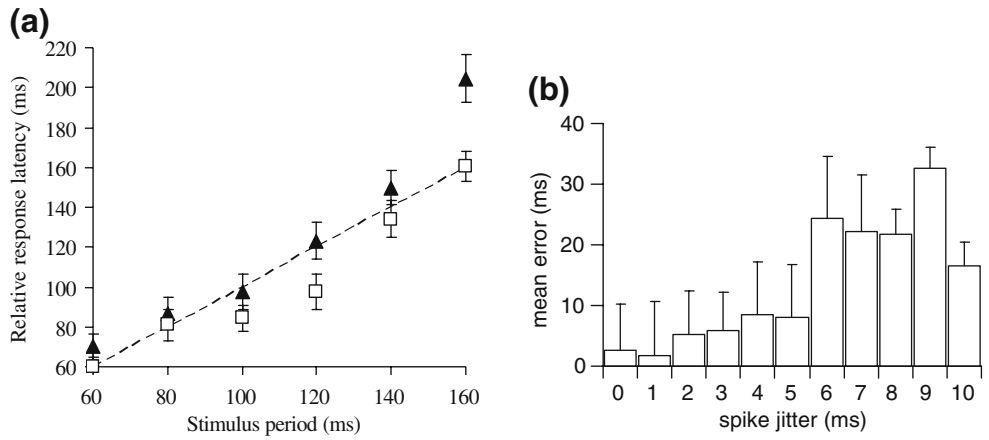

Fig. 6 Robustness of the model to the addition of random noise. (**a**) The addition of subthreshold Gaussian noise reduces the precision of OSRs but maintains a correct order of relative response latencies across different stimulus periods. Two sets of simulations are shown, where Gaussian noise had a mean of zero and standard deviation of either 10 (*white squares*) or 100 (*black triangles*). (**b**) Spike jitters disrupt the precision of OSRs following 12 pulses delivered at 10 Hz. Mean error is the difference between the arrival time of the first post-CS spike of each neuron and the time when the omitted pulse would have occurred (Eq. (6)). Each bar is a mean over all neurons in the network (N=100). *Vertical bars*: SEM. "spike jitter" (x-axis) refers to the range of jitter that was randomly applied to individual spikes during simulation

with external noise than without. Second, the SEM is greater with external noise than without. Despite this reduction in precision, the monotonic relationship between stimulus period and relative response latency was preserved. This result was found in simulations with both lower (standard deviation of 10) and higher (standard deviation of 100) levels of noise. Thus, while some degree of spike precision is lost with the addition of external noise, other, more general properties of OSRs are maintained.

In a complementary series of simulations, we randomly jittered the spike times of neurons both during and after stimulation. To examine the impact of spike jittering, we calculated the precision of OSRs by taking the first post-CS spike of each neuron i, and computing the difference between its arrival time (t_i) and the time when the omitted pulse would have occurred ($t_{omitted}$):

$$\text{precision}_{\text{OSR},i} = |t_i - t_{omitted}|. \quad (6)$$

When spikes were jittered within a restricted range of ~1–2 ms, the impact on the precision of OSRs was marginal (Fig. 6(b)). However, disruptions increased as we applied a gradually broader range of spike jitter. The model's ability to detect temporal sequences may therefore not be dependent upon highly precise (1 ms) activity, but does require some degree of spike accuracy (<10 ms).

3.5 Effect of weight bounding

The model proposed here allows synaptic connections to take on a range of values (i.e., no hard upper bound on weights are imposed). While we are not the first to relax the assumption of a weight bound (Morrison et al. 2008), this nonetheless seems at odds with experimental work suggesting that single synaptic contacts are binary (Petersen et al. 1998; O'Connor et al. 2005). This issue is a matter of debate in the experimental literature, with some work suggesting distributions of synaptic strengths (Sjostrom et al. 2001). While a computational account cannot definitively settle the question one way or the other, it is nevertheless useful to examine the consequences of weight bounds on the computational capacities of neural circuits (Fusi and Abbott 2007).

In a series of simulations, we examined the OSRs of different models with bounds in the range of either [0,1] or [0,10] (Fig. 7(a)). A synaptic bound of [0,1] reduced overall weights by an average of 98% (mean over all stimulus periods from 60–160 ms), and decreased the precision of OSRs. By comparison, a synaptic bound of [0,10] reduced overall weights by 35%, yet had minimal impact on the precision of OSRs (see Fig. 7(b) for weight distributions under different bounds). These results suggest that while stringent weight bounds are detrimental to OSRs, the model can withstand some degree of weight bounding and still produce reasonably precise responses to rhythmic stimuli. Indeed, even in the unbounded model, in which weights are allowed to take on any value, very few ever grew beyond a weight of 10 (Fig. 2(c)).

3.6 Time delays in LTP/LTD induction

Another controversial issue concerns the exact time course involved in the induction of STDP. For instance, pyramidal neurons exhibit an "all-or-none" form of plasticity that is expressed after ~14 s (O'Connor et al. 2005); other work suggests that induction of LTP can take less than a few seconds (Frey and Morris 1997). Few models of synaptic plasticity have addressed this issue directly (Morrison et al.

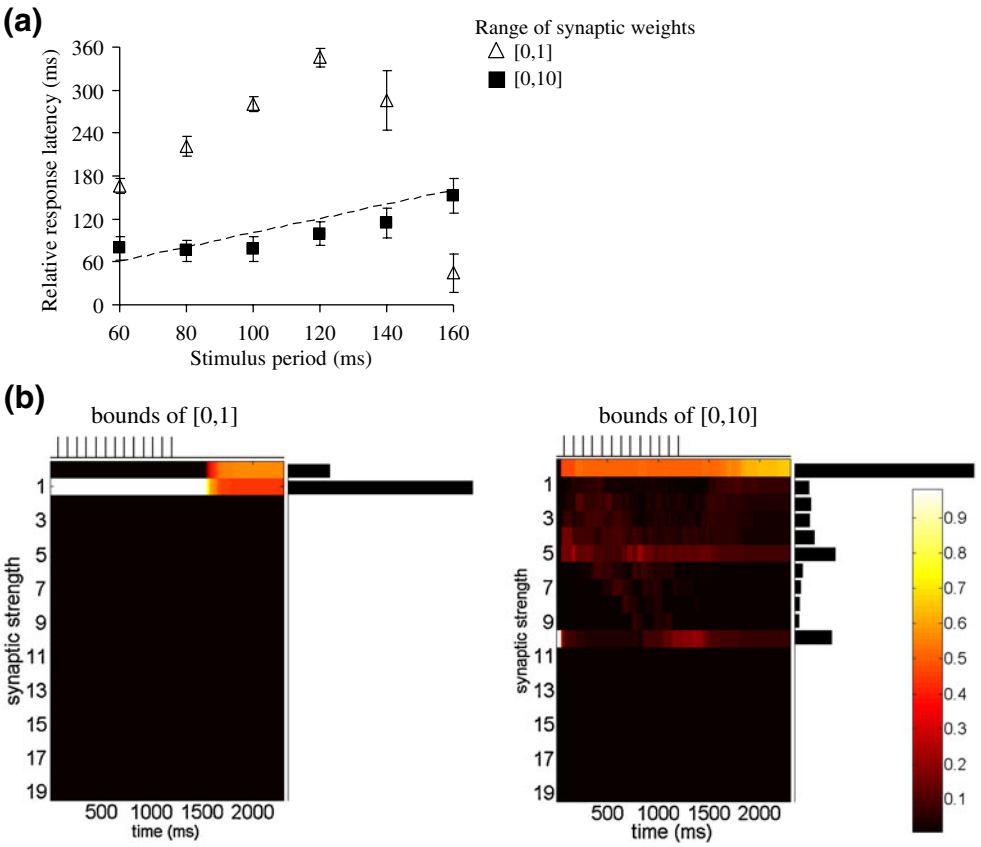

Fig. 7 Imposing hard bounds on synaptic weights influences the precision of OSRs. (**a**) Relative response latency of first post-CS spikes in different simulations that force synaptic weights to remain within a range of either [0,1] or [0,10]. (**b**) Histograms showing the distribution of synaptic weights with hard bounds of either [0,1] or [0,10] (stimulus period: 100 ms). The mean distribution of synaptic weights over the entire simulation is shown on the right of each histogram. The times where pulses were delivered are shown above each histogram

2008)—many models assume a fixed delay in the induction of LTP/LTD (Izhikevich 2006). Based on evidence suggesting that synapses possess different plasticity properties (Abbott and Nelson 2000; Thomson and Lamy 2007), we designed a set of simulations where induction delays varied across synapses, following a Gaussian distribution with mean of either 100 ms or 500 ms and standard deviation of 0.33 (Fig. 8). Results taken over a range of stimulus periods (60–160 ms) suggest that the precision of OSRs is robust to the addition of short delays (mean of 100 ms), but begins to degrade with longer delays (mean of 500 ms). While the exact distributions of delays in the model were not meant to capture precise experimental data, they do suggest that relatively fast induction of synaptic plasticity may promote neural responses to rhythmic sequences. Further experimental work is needed to shed light on the particular time course of plasticity, as well as the distribution of synaptic weights over a complete population of neurons (Song et al. 2005).

3.7 Synaptic plasticity alters population-wide interactions during rhythmic entrainment

The model's ability to detect rhythmic patterns was strongly dependent upon the presence of synaptic plasticity (in particular, STDP). When removed altogether (by fixing synaptic strengths for the entire simulation), activity levels were too high to follow any induced rhythm (Fig. 2(b), white triangles). How does STDP contribute to the post-stimulation response of neurons following a rhythmic pattern? As shown in recent work, STDP synchronizes the activity of heterogeneous neurons that receive a "pacemaker" rhythm (Masuda and Kori 2007). When the imposed rhythm is removed, the heterogeneous nature of neurons forces them to gradually fall out of phase with each other. The synchronization of heterogeneous oscillators, followed by the gradual way in which they fall out of phase after stimulation, have been recently analyzed (Saigusa et al. 2008). As a general rule, both the degree of neural heterogeneity and the synaptic strength between neurons mediate phase synchrony during stimulation (Masuda and Kori 2007). Transient OSRs are produced following stimulation because STDP maintains elevated synaptic strengths for a brief period, during which there is still some structure to the phase of neuronal activity.

The above rule holds in the present model. For each neuron, we calculated the instantaneous frequency as the reciprocal of every inter-spike interval. To obtain the instantaneous phase $\phi(t)$ of a neuron at a given time-step t, we applied the following (Freund et al. 2003):

$$\phi(t) = 2\pi \frac{t - t_i}{t_{i+1} - t_i} + 2\pi n \quad (t_i \le t \le t_{i+1}), \tag{7}$$

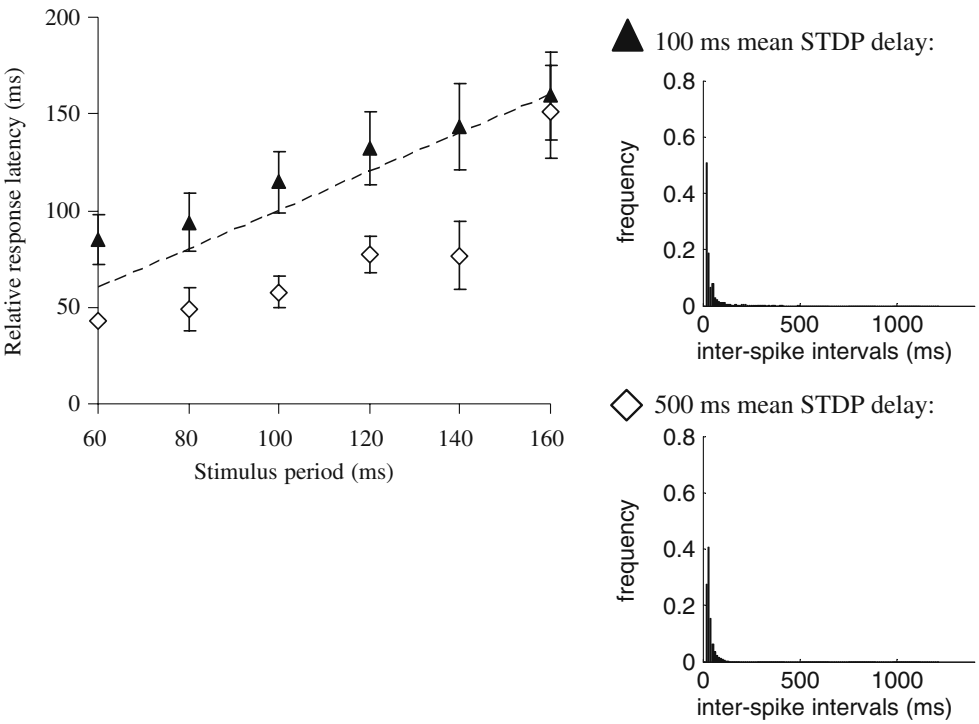

Fig. 8 Relative response latency of first post-CS spikes in a model with a distribution of time delays in the induction of synaptic plasticity. *Left*: relative response latency of first post-CS spikes, showing the mean over 10 independent runs of the model for each stimulus period of induced activity between 60 ms and 160 ms. *Black triangles*: mean induction time delays of 100 ms. *White squares*: mean induction time delays of 500 ms. *Right*: distribution of inter-spike intervals for 10 s of spontaneous activity (scaled between [0,1])

where t_i denotes a time at which the neuron fires a spike. During stimulation, neurons whose firing frequency is close to a 10 Hz rhythmic stimulation (Fig. 9(a), open squares) all respond at the same phase. Here, the contribution of STDP is to align the firing phase of neurons by increasing the strength of synapses between them (Masuda and Kori 2007). When the stimulation halts, these neurons do not immediately go out of phase; rather their decay is gradual (Fig. 9(b)). In a simulation without an imposed rhythm, neurons show no relationship between their phase and firing frequency (Fig. 9(c)). In sum, a heterogeneous network exhibit OSRs because neurons that fire at the induced frequency tend to also fire at the same phase; the gradual decay of firing phase amongst neurons leads to a continuation of the rhythmic response beyond the time of stimulation. Eventual desynchronization occurs when the strength of synapses is no longer sufficient to prevent phases from drifting apart (Masuda and Kori 2007).

3.8 Freezing synaptic weights following stimulation

Following stimulation, neurons gradually fall out of phase because of their intrinsic properties, and it is possible that STDP hastens this process of desynchronization. To explore this idea, we simulated a network where all synaptic weights froze immediately after stimulation—such a network might perform OSRs for more cycles than a network that continued to modify synaptic weights after stimulation. Results confirm this idea and show that a network with frozen weights performs OSRs for more cycles and with more precision (Fig. 10). These results are consistent with the idea that STDP precipitates the desynchronization of heterogeneous neurons following a rhythmic stimulus. In addition, the ability of a network with frozen weights to generate OSRs shows that while STDP enhances neural synchronization to rhythmic stimuli (Masuda and Kori 2007), it does not directly contribute to the post-stimulation response.

3.9 Omitted stimulus responses to more complex stimuli

Going beyond rhythmic sequences composed of a unique stimulus period, simulated neurons also produced OSRs after a sequence composed of two periods, 30 ms (between the 1st and 2nd pulses) and 80 ms (between the 1st and 3rd pulses) (Fig. 11(a), left). To quantify this response, we computed the mean inter-spike interval of each neuron within 500 ms after stimulation (Fig. 11(a), right). This was repeated for 10 independent runs of the model (with different initial conditions), both with and without stimulation. The resulting distribution of inter-spike intervals was significantly different with stimulation ($t(998)=48.61$, $p<0.001$); stimulation resulted in more inter-spike intervals around 20–40 ms as well as 70–90 ms. An ability to produce OSRs for sequences composed of two frequencies was also found for other combinations of stimuli (Fig. 11(b)). These results are compatible with the idea that the detection of complex sequences (e.g., as in retinal ganglion cells) may be driven by the same mechanisms as simpler sequences (Han et al. 2008; Schwartz et al. 2007; Yao et al. 2007). The precision of OSRs diminishes, however, when comparing sequences of

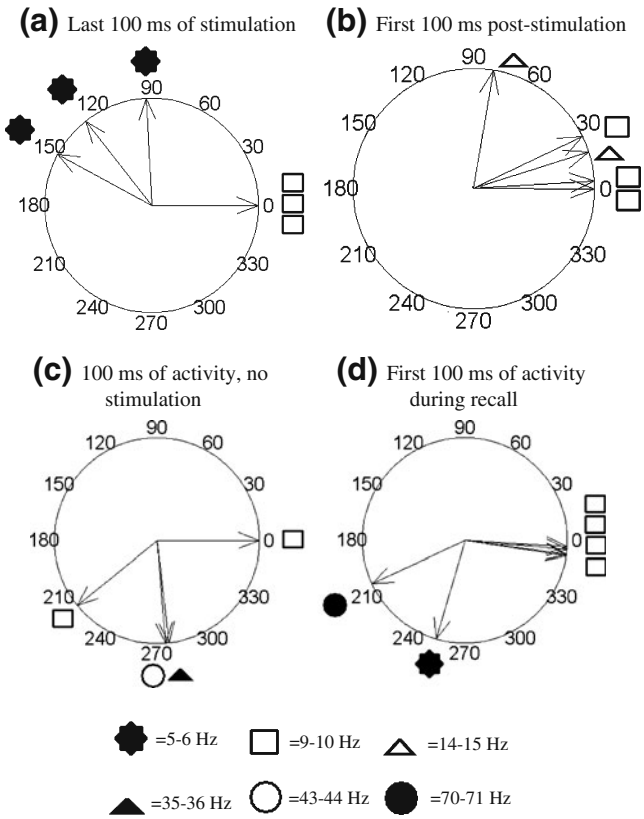

Fig. 9 Phase relationship between neurons during and after rhythmic stimulation. (**a**) Example of six neurons during the last 100 ms of stimulation (12 pulses at 10 Hz frequency). The phase of each neuron is shown as the angle on the ring, and their frequency is denoted by symbols (see legend). Neurons that fire close to the stimulation frequency (9–10 Hz, *open squares*) are aligned in terms of their phase. (**b**) During the first 100 ms after stimulation, neurons that fired close to the stimulation frequency (*open squares*) retain their approximate phase relationship. (**c**) Phase relationship amongst four neurons in a simulation without rhythmic patterns. Neurons with a similar frequency (*open squares*) are not aligned in phase. (**d**) Phase relationship amongst six neurons during the first 100 ms of recall, showing that neurons with a firing frequency close to the induced stimuli (10 Hz) have a similar phase

two frequencies to sequences composed of a single frequency (to display results more clearly, bar graphs in Fig. 11(b) show a sum of spikes over an interval of ±10 ms around the time of omitted pulses).

In a different simulation, we aimed to show that neurons in the model could produce OSRs with a sequence that has both temporal (i.e., occurring at specific times) and spatial (i.e., activating different groups of cells) aspects. We stimulated a network with pulses that alternatively targeted different subgroups of neurons (Fig. 12). Within each subgroup, the stimulus period was 100 ms; a time-lag of 50 ms was introduced between the stimulation of the two subgroups. Following stimulation, each subgroup produced OSRs that approximated the time of its respective omitted pulse.

The above results are related to the capacity of neural circuits to alter the spatiotemporal properties of spontaneous activity in order to reflect sensory experience (Han et al. 2008; Yao et al. 2007). Experiments suggest that some of these effects can be long-lasting (several minutes), and others (in hippocampus) can be delayed to a subsequent period of sleep (Louie and Wilson 2001). While our model emits a precise post-stimulation response, it does not capture the full breadth of those experiments, and does not replicate specific aspects of OSRs in the retina. This includes the ability to detect pattern regularity (e.g., when two short inter-stimulus intervals occur in a sequence of alternating long and short intervals) and to capture the large increase in firing responses during OSRs compared to induced activity (Schwartz et al. 2007). Nonetheless, it is a step forward in capturing the response of a network to complex spatiotemporal patterns of activity, forming a neural substrate of perceptual memory.

3.10 Recall of rhythmic information following a delay period

While experiments and simulations agree that OSRs do not extend much further than a single cycle after stimulation, it may be possible to maintain rhythmic information for a longer time, as illustrated in the following simulation. First, 12 pulses of stimulation were applied as previously described, with the exception that tonic input was low (I_{tonic}=0.1 nA) in order to prevent immediate recall. Then, a delay period lasting 1 s allowed spontaneous activity without stimulation. Finally, a recall period activated neurons by gradually increasing their tonic current (I_{tonic}, see Eq. (1)) by increments of 0.1 nA per 1 ms (Fig. 13(a)). We examined OSRs during the recall period by storing inter-spike intervals between the first and second spikes emitted. During the recall period, these inter-spike intervals approximated the frequency of induced activity (Fig. 13(b), black diamonds), albeit with less precision than OSRs produced immediately after stimulation (Fig. 2(b), black

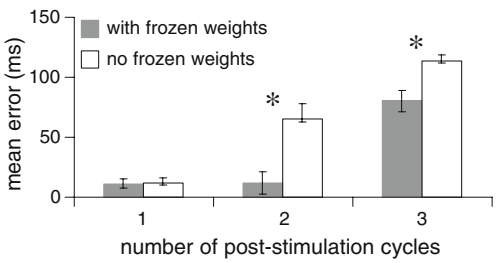

Fig. 10 A network that freezes synaptic weights immediately after stimulation performs OSRs over several cycles. Mean error (in ms) of firing responses for the first three cycles post-stimulation. * = $p<$ 0.001 (Student's *t*-test comparing mean error over 10 independent runs of the model, with vs. without frozen weights). *Vertical bars*: SEM

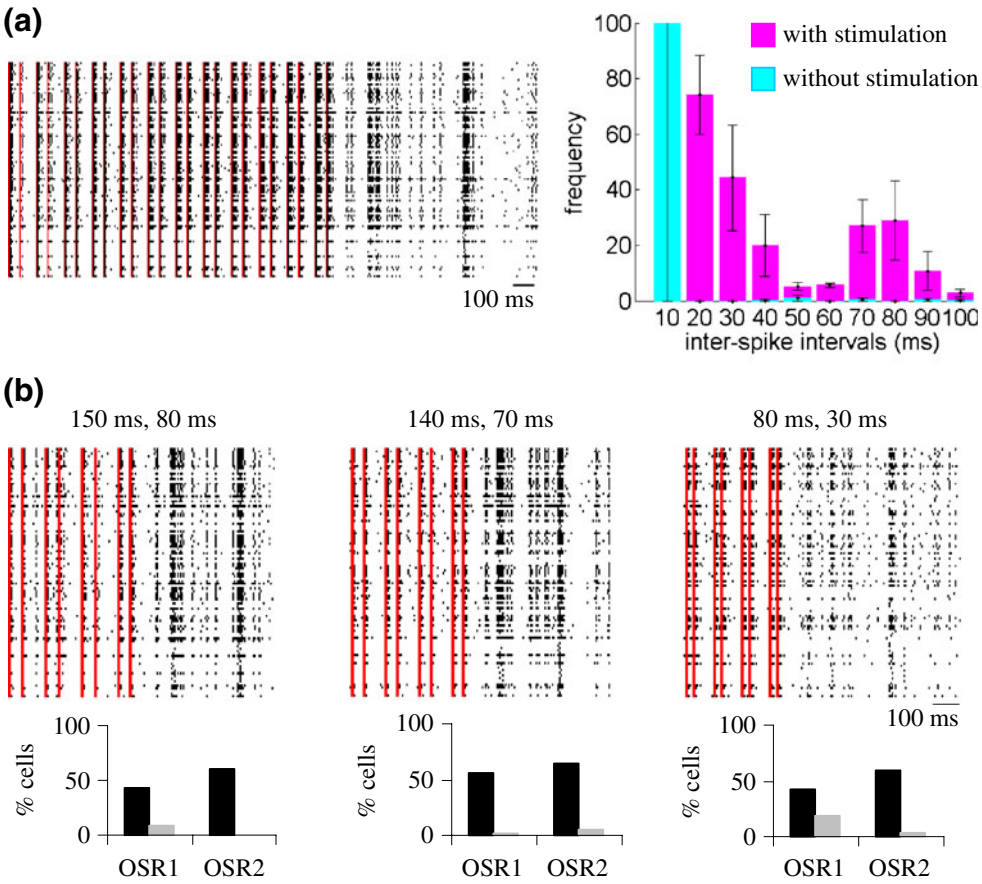

Fig. 11 Omitted stimulus response to more complex conditioning stimulations. (**a**) *Left*, omitted stimulus response for a rhythm composed of two stimulus periods (distance of 30 ms between the 1st and 2nd pulses, and 80 ms between the 1st and 3rd pulses). *Right*, mean distribution of post-CS inter-spike intervals (within 500 ms after the delay period) for all neurons over 10 independent runs of the model, with and without rhythmic stimulation. *Vertical bars*, SEM. (**b**) Each graph shows the response of all cells to different combinations of frequencies. *Top*: spike raster for each cell over time. The times of induced pulses are shown in red. *Bottom*: % of cells active at the time of the first ("OSR1") and second ("OSR2") omitted pulses (within an interval of ±10 ms); *black bars*, with CS; *gray bars*, no CS (i.e., simulation with only spontaneous activity)

diamonds). Extending the delay period to 10 s disrupted the precision of recalled OSRs (Fig. 13(b), white squares).

How does the proposed model recall the frequency of an induced rhythm after a delay period? One possibility is that, after stimulation, the firing rate of neurons decreases back to baseline, but their phase remains synchronized. Indeed, during recall, when firing rates are increased, neurons whose firing frequency is close to the initial 10 Hz rhythmic stimulation respond approximately in phase with one another (Fig. 9(d)). This preserved phase relationship amongst neurons causes them to fire in synchrony at a period that approximates the rhythmic stimulation (Fig. 13(a)). An analysis of heterogeneous oscillators and their ability to recall temporal patterns has recently been provided (Saigusa et al. 2008). One prediction based on the above account is that recall performance will be disrupted if the phase of neuronal activity is altered during the delay period. We tested this hypothesis by randomly jittering spikes, which disrupts the phase of neurons while preserving their frequency. We defined a measure of "recall error" as the difference between 1) the time interval between the first and second spikes emitted during the recall period, and 2) the stimulus period of the CS. Simulations that added a small amount of spike jitter (1 ms range) during the delay period did not lead to a significant increase in recall error when compared to simulations without jitter ($t(18)=0.33$, $p>0.75$) (Fig. 13(c)). The addition of a larger (10 ms range) spike jitter, however, disrupted performance ($t(18)=3.81$, $p<0.001$). The model was therefore not sensitive to small disruptions in spike times, yet some degree of precision was required for adequate recall of a rhythmic pattern.

4 Discussion

This paper describes a hypothesis explaining the omitted stimulus response observed in several species and anatomical systems, and reflecting an ability of neural circuits to respond in a transient fashion to synchronized input. The hypothesis, based on a combination of neural diversity and synaptic plasticity, accounts for the precise yet highly transient nature of OSRs. In addition, it shows that the same mechanisms that detect simple patterns composed of a single stimulus period can also detect more complex patterns, arguing that these mechanisms are not limited to a narrow scope of responses. Finally, the model was able to store rhythmic information and replay a stimulus period during a later phase of recall. These results go beyond previous work that focused on the detection of single pulses (Thivierge and Cisek 2008). In particular, here we show

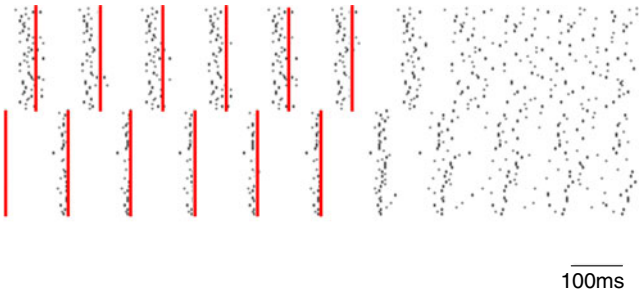

Fig. 12 Omitted stimulus response to spatiotemporal patterns of activity. Rhythm alternating between two subpopulations of cells. *Red*: Induced activity

that heterogeneous spiking networks with STDP can reproduce rhythms composed of multiple frequencies and can recall them after a period of silence. In addition, the present simulations show robustness of the model to a wide range of settings, and discuss potential neurophysiological mechanisms for detecting rhythmic patterns based on a combination of STDP and phase alignment amongst heterogeneous neurons.

4.1 A combination of synaptic plasticity and rich neural dynamics are required to detect temporal patterns

Taken together, our results show that synaptic plasticity based on STDP is a powerful mechanism for learning rhythmic sequences (Abbott and Blum 1996; Bi and Poo 1999). While the present simulations cannot with full certainty settle on STDP as the only possible rule for doing so, we propose, in conjunction with other work showing the limitations of alternative rules based on short-term plasticity (Thivierge and Cisek 2008), that STDP is a plausible candidate that yields precisely timed OSRs to simple as well as more complex patterns.

Our results suggest that the role of STDP is two-fold: 1) to synchronize neuronal activity at a frequency that is close to that of the induced rhythmic pattern (Fig. 2(a)); and 2) to regulate the positive feedback between neurons. In simulations without STDP, positive feedback causes neurons to fire at a very high rate; because this rate is much higher than that of induced activity, they cannot detect temporal patterns, as detailed elsewhere (Thivierge and Cisek 2008).

One prediction derived from the role of STDP in the model is that the down-regulation of activity-dependent synaptic plasticity during stimulation should reduce the probability of OSRs. Consistent with this prediction, whole-cell CA1 recordings have shown diminished post-tetanic responses as a result of a GABA receptor antagonist (Bracci et al. 1999); in addition, other work has shown a decreased cortical response to flash stimuli following NMDA inhibition (Heynen and Bear 2001).

The proposed model is robust to several alterations, including bounded synapses, heterogeneous neurons, random synaptic connectivity, inhibitory synapses, as well as delays in the induction of LTP/LTD. These results are consistent with the notion that, in biological circuits, heterogeneity has many potential sources. Going further, simulation results demonstrate how cellular heterogeneity can play an important role in creating a rich repertoire of neural dynamics. In particular, heterogeneity prevents the network from getting "stuck" in repetitive patterns of activity, instead creating transient dynamics that allow the network to respond to external input. In contrast, a homogeneous network cannot modulate its activity in response to a rhythmic stimulus. Neural heterogeneity also plays a central role in the desynchronization of neurons following stimulation. As a recent analysis makes clear, heterogeneous oscillators fall out of phase after stimulation because of differences in their intrinsic response properties (Saigusa et al. 2008). Our results extend previous work where complex patterns of activity improve associative memories by preventing neurons from getting caught in a particular state of activity (Destexhe and Contreras 2006); in addition, by virtue of their flexibility, these patterns enhance stimulus detection (Collins et al. 1995) and promote fast responses (Silberberg et al. 2004).

4.2 Recalling temporal patterns after a delay period

Our model demonstrates how memory-related information acquired during a stimulation period can be recalled at a later point in time, following a delay period. Despite an apparent lack of structure of the spontaneous spiking activity produced during the delay period (compared to activity produced during stimulation), it played an important role in the subsequent replay of temporal patterns. Jittering activity during the delay period impaired the network's ability to replay patterns. This result suggests that short-term memory traces may be maintained in the precise spontaneous interactions of a population of cells.

The spontaneous neural interactions observed during the delay period suggest that subsequent memory reactivation does not require sustained activity at a rate as high as that generated during the stimulation period. In fact, rates of spiking activity during the delay period were markedly lower than those obtained during stimulation (Fig. 13(a)). This result is consistent with electrophysiological studies that show only modest increases in spike activity during a delay period preceding recall (Rainer and Miller 2002; Shafi et al. 2007). In part because of the high metabolic cost associated with sustained spike activity, neuronal circuits may rely on alternate means of maintaining information in short-term memory; our modeling work, as well as other approaches (Mongillo et al. 2008), begins to shed light on these mechanisms.

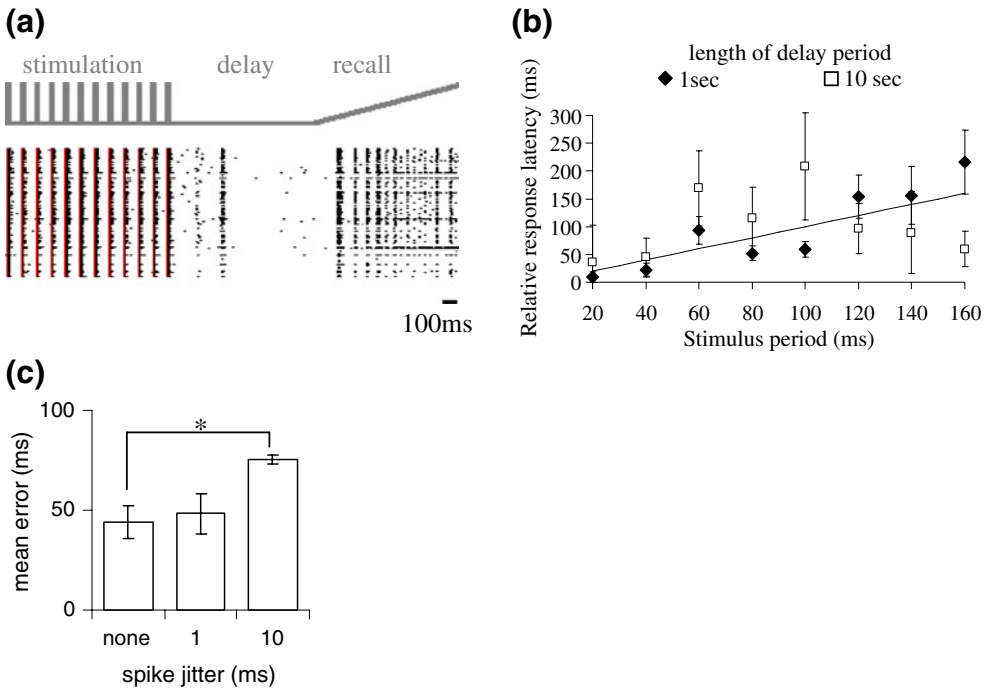

Fig. 13 Recall of a spike rhythm after a delay period. (**a**) Rhythmic activity (100 ms stimulus period) is induced for 2,000 ms (shown in *gray*). Then, following a delay period (lasting 1 s), tonic activity (I_{tonic}, see Eq. (1)) is gradually increased. As a result, neurons produce activity at a period that approximates the induced rhythm. (**b**) Maintenance of spike rhythm for induced activity with various stimulus periods. Values shown are differences between first and second spikes generated during the recall period. (**c**) Random spike jitters with a range of 10 ms but not a range of 1 ms disrupt the recall precision of a rhythmic pattern (12 pulses at 10 Hz). Each bar is mean recall error over 10 independent runs of the model. *Vertical bars*: SEM. $*p<0.01$ (Student's *t*-test), simulations with vs. without spike jitters

A fundamental assumption of our theory is that the replay of rhythmic patterns after a delay period depends upon an increase in nonspecific background activity across a population of neurons (Mongillo et al. 2008). Speculatively, elevated background activity may also play a role in the reactivation of learning-related neural activity during sleep (Ji and Wilson 2007; Louie and Wilson 2001), which promotes memory consolidation. This hypothesis is consistent with a recent computational model, where hippocampal reactivation is driven by persistent input from the entorhinal cortex and dentate gyrus (Koene and Hasselmo 2008). This mechanism, however, does not rule out the possibility of reactivation through intrinsic dynamics. Further modeling work may help explore this question, as our current theory does not aim to capture specific properties of sleep reactivation, including the temporal compression, expansion, or reverse replay of neural patterns (Foster and Wilson 2006; Louie and Wilson 2001).

4.3 Related approaches

Our neural model is preceded by a few related approaches. One of these approaches is the "echo state network" (ESN) (Jaeger and Haas 2004), which consists of a population of randomly connected elements that converge onto an output unit. After training the system with a particular time series, the output unit spontaneously predicts the evolution of the time series for over 2,000 steps into the future. As in our approach, ESNs must be capable of richly varied dynamics; this is achieved by a sparse interconnectivity. Our approach, however, differs from ESNs in two important ways. First, ESNs use sigmoidal units, whereas we used leaky integrate-and-fire neurons that allow us to study spike interactions among neurons, by treating these spikes as a fundamental unit of neuronal communication. Second, ESNs require a massive convergence of connections onto the output unit. By comparison, our model makes no such assumption, treating the entire system as both input and output. Given their impressive capacity for predicting upcoming signals, an interesting prospect might be to combine our present model with ESNs (Maass et al. 2002).

Another related approach is the notion of synfire chains (Abeles 1991; Diesmann et al. 1999), involving signal propagation through synchronous activity. From a general perspective, both our study and the approach of synfire chains are concerned with the broad theme of signal propagation. However, our goal of designing a neural system that can detect rhythmic patterns and emit OSRs in a highly precise—yet transient—manner has not been addressed by synfire chains. In addition, signal propagation

in synfire chains has been studied primarily in feedforward networks (with some exceptions, notably Vogels and Abbott 2005); by contrast, our model does not assume a feedforward structure.

Our approach is also related to "polychronization" (Izhikevich 2006). Accordingly, a model that combines spiking neurons and STDP (similar to our approach) generates patterns of activity extending both in time and across several neurons. These patterns, termed "polychronous groups", occur both spontaneously and as a result of neural stimulation.

While several other models have shown an ability of spiking neurons to synchronize their activity (Levy et al. 2001; Morrison et al. 2007; Suri and Sejnowski 2002) and respond to sequences of patterns (Gerstner et al. 1993; Herz et al. 1989) none displays the OSR behavior that our model shows. This OSR forms a transient yet highly precise synchronization of activity following rhythmic patterns, which is the principal contribution of our study.

4.4 Conclusion

The aim of this study was to propose a hypothesis to explain general aspects of OSRs in several systems including retina and cortex. The main contribution of our work is to show that basic neural properties—including cellular heterogeneity and synaptic plasticity—are sufficient for spike responses to approximate the frequency of induced rhythmic activity. While further work will be required to capture precise quantitative aspects of OSRs in specific neural systems such as retina (Schwartz et al. 2007) and cortex (Demiralp and Basar 1992), as well as to explain responses to longer stimulus periods in the optic tectum (Sumbre et al. 2008), our results nonetheless provide a hypothesis of potentially broad relevance that could be incorporated into more detailed models. Neural mechanisms for detecting temporal patterns likely involve heterogeneity in intrinsic cell properties and connectivity, in addition to the ability of neural circuits to adaptively alter their synaptic strength in response to temporally structured input.

Acknowledgments JPT is supported by postdoctoral fellowships from the Natural Sciences and Engineering Research Council of Canada and the Fonds de Recherche en Santé du Québec. PC is supported by grants from the Natural Sciences and Engineering Research Council of Canada and the Fonds de Recherche en Santé du Québec (infrastructure grant).

References

Abbott, L. F., & Blum, K. I. (1996). Functional significance of long-term potentiation for sequence learning and prediction. *Cerebral Cortex, 6*(3), 406–416.

Abbott, L. F., & Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nature Neuroscience, 3*(Suppl), 1178–1183.

Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex*. Cambridge: Cambridge University Press.

Bi, G., & Poo, M. (1999). Distributed synaptic modification in neural networks induced by patterned stimulation. *Nature, 401*(6755), 792–796.

Bracci, E., Vreugdenhil, M., Hack, S. P., & Jefferys, J. G. (1999). On the synchronizing mechanisms of tetanically induced hippocampal oscillations. *The Journal of Neuroscience, 19*(18), 8104–8113.

Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience, 8*(3), 183–208.

Brunel, N., & Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Computation, 11*(7), 1621–1671.

Buzsaki, G., Geisler, C., Henze, D. A., & Wang, X. J. (2004). Interneuron diversity series: circuit complexity and axon wiring economy of cortical interneurons. *Trends in Neurosciences, 27*(4), 186–193.

Collins, J. J., Chow, C. C., & Imhoff, T. T. (1995). Stochastic resonance without tuning. *Nature, 376*(6537), 236–238.

Demiralp, T., & Basar, E. (1992). Theta rhythmicities following expected visual and auditory targets. *International Journal of Psychophysiology, 13*(2), 147–160.

Demiralp, T., Basar-Eroglu, C., Rahn, E., & Basar, E. (1994). Event-related theta rhythms in cat hippocampus and prefrontal cortex during an omitted stimulus paradigm. *International Journal of Psychophysiology, 18*(1), 35–48.

Destexhe, A., & Contreras, D. (2006). Neuronal computations with stochastic network states. *Science, 314*(5796), 85–90.

Diesmann, M., Gewaltig, M. O., & Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature, 402*(6761), 529–533.

Eytan, D., & Marom, S. (2006). Dynamics and effective topology underlying synchronization in networks of cortical neurons. *The Journal of Neuroscience, 26*(33), 8465–8476.

Feldman, D. E. (2000). Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron, 27*(1), 45–56.

Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature, 440*(7084), 680–683.

Freund, J. A., Schimansky-Geier, L., & Hanggi, P. (2003). Frequency and phase synchronization in stochastic systems. *Chaos, 13*, 225–238.

Frey, U., & Morris, R. G. (1997). Synaptic tagging and long-term potentiation. *Nature, 385*, 533–536.

Friedman, D., Cycowicz, Y. M., & Gaeta, H. (2001). The novelty P3: an event-related brain potential (ERP) sign of the brain's evaluation of novelty. *Neuroscience and Biobehavioral Reviews, 25*(4), 355–373.

Fusi, S., & Abbott, L. F. (2007). Limits on the memory storage capacity of bounded synapses. *Nature Neuroscience, 10*, 485–493.

Ganguli, S., Huh, D., & Sompolinsky, H. (2008). Memory traces in dynamical systems. *Proceedings of the National Academy of Sciences of the United States of America, 105*(48), 18970–18975.

Gerstner, W., & Kistler, W. (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge: Cambridge University Press.

Gerstner, W., Ritz, R., & van Hemmen, J. L. (1993). Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. *Biological Cybernetics, 69*(5–6), 503–515.

Gerstner, W., Kempter, R., van Hemmen, J. L., & Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature, 383*(6595), 76–81.

Han, F., Caporale, N., & Dan, Y. (2008). Reverberation of recent visual experience in spontaneous cortical waves. *Neuron, 60*(2), 321–327.

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

Herz, A., Sulzer, B., Kuhn, R., & van Hemmen, J. L. (1989). Hebbian learning reconsidered: representation of static and dynamic objects in associative neural nets. *Biological Cybernetics, 60*(6), 457–467.

Heynen, A. J., & Bear, M. F. (2001). Long-term potentiation of thalamocortical transmission in the adult visual cortex *in vivo*. *The Journal of Neuroscience, 21*(24), 9801–9813.

Izhikevich, E. M. (2006). Polychronization: computation with spikes. *Neural Computation, 18*(2), 245–282.

Jaaskelainen, I. P., Ahveninen, J., Bonmassar, G., Dale, A. M., Ilmoniemi, R. J., Levanen, S., et al. (2004). Human posterior auditory cortex gates novel sounds to consciousness. *Proceedings of the National Academy of Sciences of the United States of America, 101*(17), 6809–6814.

Jaeger, H., & Haas, H. (2004). Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science, 304*(5667), 78–80.

Ji, D., & Wilson, M. A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience, 10*(1), 100–107.

Kempter, R., Gerstner, W., & van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Physical Review E, 59*(4), 4498–4514.

Kempter, R., Gerstner, W., & van Hemmen, J. L. (2001). Intrinsic stabilization of output rates by spike-based Hebbian learning. *Neural Computation, 13*(12), 2709–2741.

Koene, R. A., & Hasselmo, M. E. (2008). Reversed and forward buffering of behavioral spike sequences enables retrospective and prospective retrieval in hippocampal regions CA3 and CA1. *Neural Networks, 21*(2–3), 276–288.

Levy, N., Horn, D., Meilijson, I., & Ruppin, E. (2001). Distributed synchrony in a cell assembly of spiking neurons. *Neural Networks, 14*(6–7), 815–824.

Louie, K., & Wilson, M. A. (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron, 29*(1), 145–156.

Maass, W., Natschlager, T., & Markram, H. (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Computation, 14*(11), 2531–2560.

Markram, H., Lubke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science, 275*(5297), 213–215.

Masland, R. H. (2001). Neuronal diversity in the retina. *Current Opinion in Neurobiology, 11*(4), 431–436.

Masuda, N., & Kori, H. (2007). Formation of feedforward networks and frequency synchrony by spike-timing-dependent plasticity. *Journal of Computational Neuroscience, 22*, 327–345.

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.

Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic theory of working memory. *Science, 319*(5869), 1543–1546.

Morrison, A., Aertsen, A., & Diesmann, M. (2007). Spike-timing-dependent plasticity in balanced random networks. *Neural Computation, 19*(6), 1437–1467.

Morrison, A., Diesmann, M., & Gerstner, W. (2008). Phenomenological models of synaptic plasticity based on spike timing. *Biological Cybernetics, 98*(6), 459–478.

O'Connor, D. H., Wittenberg, G. M., & Wang, S. S. (2005). Graded bidirectional synaptic plasticity is composed of switch-like unitary events. *Proceedings of the National Academy of Sciences of the United States of America, 102*(27), 9679–9684.

Petersen, C. C., Malenka, R. C., Nicoll, R. A., & Hopfield, J. J. (1998). All-or-none potentiation at CA3-CA1 synapses. *Proceedings of the National Academy of Sciences of the United States of America, 95*, 4732–4737.

Rabinovich, M. I., & Abarbanel, H. D. (1998). The role of chaos in neural systems. *Neuroscience, 87*(1), 5–14.

Rabinovich, M., Huerta, R., & Laurent, G. (2008). Neuroscience. Transient dynamics for neural processing. *Science, 321*(5885), 48–50.

Rainer, G., & Miller, E. K. (2002). Timecourse of object-related neural activity in the primate prefrontal cortex during a short-term memory task. *The European Journal of Neuroscience, 15*(7), 1244–1254.

Ramon, F., & Gronenberg, W. (2005). Electrical potentials indicate stimulus expectancy in the brains of ants and bees. *Cellular and Molecular Neurobiology, 25*(2), 313–327.

Saigusa, T., Tero, A., Nakagaki, T., & Kuramoto, Y. (2008). Amoebae anticipate periodic events. *Physical Review Letters, 100*, 018101.

Schwartz, G., Harris, R., Shrom, D., & Berry, M. J., 2nd. (2007). Detection and prediction of periodic patterns by the retina. *Nature Neuroscience, 10*(5), 552–554.

Shafi, M., Zhou, Y., Quintana, J., Chow, C., Fuster, J., & Bodner, M. (2007). Variability in neuronal activity in primate cortex during working memory tasks. *Neuroscience, 146*(3), 1082–1108.

Silberberg, G., Bethge, M., Markram, H., Pawelzik, K., & Tsodyks, M. (2004). Dynamics of population rate codes in ensembles of neocortical neurons. *Journal of Neurophysiology, 91*(2), 704–709.

Sjostrom, P. J., Turrigiano, G. G., & Nelson, S. B. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron, 32*, 1149–1164.

Song, S., Sjostrom, P. J., Reigl, M., Nelson, S., & Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biology, 3*, e68.

Sumbre, G., Muto, A., Baier, H., & Poo, M. M. (2008). Entrained rhythmic activities of neuronal ensembles as perceptual memory of time interval. *Nature, 456*(7218), 102–106.

Suri, R. E., & Sejnowski, T. J. (2002). Spike propagation synchronized by temporally asymmetric Hebbian learning. *Biological Cybernetics, 87*(5–6), 440–445.

Swadlow, H. A. (1985). Physiological properties of individual cerebral axons studied *in vivo* for as long as one year. *Journal of Neurophysiology, 54*(5), 1346–1362.

Thivierge, J. P., & Cisek, P. (2008). Nonperiodic synchronization in heterogeneous networks of spiking neurons. *The Journal of Neuroscience, 28*(32), 7968–7978.

Thivierge, J. P., Rivest, F., & Monchi, O. (2007). Spiking neurons, dopamine, and plasticity: timing is everything, but concentration also matters. *Synapse, 61*(6), 375–390.

Thomson, A. M., & Lamy, C. (2007). Functional maps of neocortical local circuitry. *Frontiers in Neuroscience, 1*, 19–42.

van Rossum, M. C., Turrigiano, G. G., & Nelson, S. B. (2002). Fast propagation of firing rates through layered networks of noisy neurons. *The Journal of Neuroscience, 22*(5), 1956–1966.

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

Vogels, T. P., & Abbott, L. F. (2005). Signal propagation and logic gating in networks of integrate-and-fire neurons. *The Journal of Neuroscience, 25*(46), 10786–10795.

Vogels, T. P., Rajan, K., & Abbott, L. F. (2005). Neural network dynamics. *Annual Review of Neuroscience, 28*, 357–376.

Wang, Y., Markram, H., Goodman, P. H., Berger, T. K., Ma, J., & Goldman-Rakic, P. S. (2006). Heterogeneity in the pyramidal network of the medial prefrontal cortex. *Nature Neuroscience, 9*, 534–542.

Yao, H., Shi, L., Han, F., Gao, H., & Dan, Y. (2007). Rapid learning in cortical coding of visual scenes. *Nature Neuroscience, 10*(6), 772–778.

Zhang, L. I., Tao, H. W., Holt, C. E., Harris, W. A., & Poo, M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature, 395*(6697), 37–44.