

# Emergent spike patterns in neuronal populations

Logan Chariker · Lai-Sang Young

Received: 26 March 2014 / Revised: 23 September 2014 / Accepted: 25 September 2014 / Published online: 18 October 2014  
© Springer Science+Business Media New York 2014

**Abstract** This numerical study documents and analyzes emergent spiking behavior in local neuronal populations. Emphasis is given to a phenomenon we call *clustering*, by which we refer to a tendency of random groups of neurons large and small to spontaneously coordinate their spiking activity in some fashion. Using a sparsely connected network of integrate-and-fire neurons, we demonstrate that spike clustering occurs ubiquitously in both high firing and low firing regimes. As a practical tool for quantifying such spike patterns, we propose a simple scheme with two parameters, one setting the temporal scale and the other the amount of deviation from the mean to be regarded as significant. Viewing population activity as a sequence of *events*, meaning relatively brief durations of elevated spiking, separated by *inter-event times*, we observe that background activity tends to give rise to extremely broad distributions of event sizes and inter-event times, while driving a system imposes a certain regularity on its inter-event times, producing a rhythm consistent with broad-band gamma oscillations. We note also that event sizes and inter-event times decorrelate very quickly. Dynamical analyses supported by numerical evidence are offered.

**Keywords** Neuronal networks · Emergent dynamics · Spike clusters · Population activity

---

Action Editor: David Terman

---

This research is supported in part by NSF Grant DMS-1101594.

---

L. Chariker · L.-S. Young (✉)  
Courant Institute of Mathematical Sciences, New York University,  
New York, NY, USA  
e-mail: lsy@cims.nyu.edu

This paper is about emergent dynamics in a network of integrate-and-fire neurons intended as a toy model for a local population somewhere in the brain. Our focus is on population activity rather than the dynamics of individual neurons, and we are interested in spiking behavior both in spontaneous activity and under drive. These patterns are the results of very complex neuron-to-neuron interactions, a rigorous analysis of which is out of reach at the present time. Reported in this paper are results of a numerical study the aim of which is to document the widespread occurrence of a phenomenon we call *spike clustering*, referring to the tendency for neuronal populations to have brief periods of elevated spiking separated by relative lulls. An extreme form of clustering is when the entire population, or a significant fraction of it, synchronizes its activity; that has been studied by many authors and will not be considered here. Another example of clustering – a strong form of it – is the production of rhythms when a system is under drive. But clustering can also occur in more subtle ways; for a sneak preview, see the bottom three panels in Fig. 2. We will show that both in background and under drive, it is quite natural for a randomly connected population of excitatory and inhibitory neurons to produce structured, inhomogeneous dynamics that are nowhere close to synchronization, and we propose a unified dynamical explanation.

First, we acknowledge previous works related to synchronization as an emergent phenomenon: For a small sample of the vast literature, see Kuramoto and Arakai (1975), Mirollo and Strogatz (1990), Tsodyks et al. (1993, 2000), Hansel et al. (1993), Whittington et al. (2000), Börgers and Kopell (2003, 2005), Börgers et al. (2005), Deville and Peskin (2008), Kilpatrick and Ermentrout (2011), and Mark and Tsodyks (2012). Section 5 contains a brief discussion

of these and other results. In most of the papers cited above, either the entire population or a significant fraction of it is synchronized or phase-locked for a good fraction of the time. Structured dynamics that do not involve synchronization or partial synchronization of any kind are studied in the the experimental papers of Plenz (see e.g. Beggs and Plenz 2003; Hahn et al. 2010; Plenz et al. 2011), which coined the term “neuronal avalanches”. Independently, Rangan and Young discovered in their computational modeling of V1 (Rangan and Young 2013a, b) a similar phenomenon which they called “multiple firing events”, or MFEs.

The phenomena studied in the present paper are much closer to those studied in (Beggs and Plenz 2003; Hahn et al. 2010; Plenz et al. 2011; Rangan and Young 2013a, b), but there are differences in emphasis and characterizations. For this reason, we have elected to use a different terminology. It is likely that all are attempts to capture, from one angle or another, the same general phenomena.

Clustering of this kind studied here produces spike patterns that are highly irregular yet clearly structured. Temporally localized elevations in spike rate can be due to the coordinated activity of any number of neurons, ranging from a handful to a sizable fraction of the population, with the set of participating neurons varying from event to event. In particular, the population is not divided into groups that spike together, as is the case in e.g. (Kilpatrick and Ermentrout 2011; Hansel et al. 1993). With regard to time gaps between clusters, not only do they vary nontrivially in duration, the rises and falls in local-in-time spike rates are generally untidy, sometimes ill defined. The dynamics are very far from periodic, a fact that is evident even if one observes the system for only a small fraction of a second.

What exactly, then, constitutes a *cluster*? There are no right or wrong answers to this question, though some answers are more satisfactory than others. Without a formal definition, however, one cannot begin to collect statistics. In this paper we take an operational approach, designed to handle data (experimental or numerical) without *a priori* assumptions. While clusters are easy to identify by visual inspection, it is more challenging to come up with a formal definition that is both meaningful and inclusive. A good definition should be effective not only in obvious situations such as synchronous spiking involving substantial fractions of the population, it should be able to detect oscillations on various timescales, as well as more ambiguous types of clustering such as those discussed in the last paragraph – and it would be good to treat all these situations under an umbrella framework. Multiple tests were conducted on many different regimes of our network, and a definition that performed acceptably is proposed in Section 2.

With a definition of clusters in hand, we propose to view population activity in terms of “events”, i.e., identifiable clusters of spikes, separated by “pauses”. We find that *event*

*sizes* and *inter-event times* are very broadly distributed in background, and that external stimulation fosters a certain regularity in inter-event times. But even these narrower distributions have considerable spread, and are qualitatively consistent with broad-band gamma oscillations (which, in cortex, are known to correspond to elevated frequencies in the range of 25–80 Hz, see e.g. Henrie and Shapley 2005). We propose that the mechanisms discussed here are more realistic than PING (Börger and Kopell 2003; Whittington et al. 2000), which produces essentially time-periodic rhythms. A more detailed comparison is given at the end of Section 4.1.

To recapitulate, the aim of this paper is to shed light on the phenomenon of clustering, and we are especially interested in inhomogeneous, structured dynamics that do not rise to the level of synchronization. In addition to proposing metrics to systematically capture and calibrate spike patterns, we seek to elucidate underlying dynamical mechanisms. Our understanding of these very complex phenomena is far from complete, but whenever we can, we try to offer explanations in terms of neuron-to-neuron interactions, and to support our heuristic arguments and conjectures with further numerical simulations.

Codes used to generate all the figures in this paper can be downloaded from ModelDB.

## 1 Model description

This section contains a detailed description of the models and parameters we use. Very briefly, we consider a homogeneously connected network of leaky integrate-and-fire neurons intended to model a local neuronal population. Such a model involves a nontrivial number of parameters. As it is not feasible to systematically explore a high dimensional parameter space, we will, for the most part, fix all but one or two of these parameters carefully chosen to allow for a wide range of dynamical behaviors, and investigate the low dimensional slices of parameter space so defined.

### 1.1 Model equations and specifications

We consider a network of  $N_E$  excitatory and  $N_I$  inhibitory neurons, with  $N_E + N_I$  in the hundreds, and  $N_E/N_I = 3$ . The simulations shown are for  $N_E = 225$  and  $N_I = 75$ . Each E-neuron is assumed to be presynaptic to 15 % of the E-population and to 30 % of the I-population, while *I-to-E* and *I-to-I* connectivities are taken to be 50 %. For each neuron in the network, its set of postsynaptic neurons is randomly drawn, leading to different realizations of random graphs, which are fixed for the duration of the numerical studies discussed in the sections to follow.

Each neuron is modeled using standard integrate-and-fire equations. Its membrane potential,  $V$ , which we have normalized to have a reset value of  $V_R = 0$  and a spiking threshold of  $V_T = 1$ , evolves according to the equation

$$\dot{V} = -\frac{1}{\tau_{\text{leak}}}V - (V - V_E)g_E - (V - V_I)g_I. \tag{1}$$

Here, time is measured in ms, and  $g_E, g_I \in [0, \infty)$  are the excitatory and inhibitory conductances of the neuron; they are time-dependent variables governed by Eqs. (2) and (3) below;  $V_E = 14/3$  and  $V_I = -2/3$  are excitatory and inhibitory reversal potentials, and  $\tau_{\text{leak}} = 20\text{ms}$  is the leak rate. These numbers follow commonly accepted biophysical parameters, normalized to put  $[V_R, V_T] = [0, 1]$ , see e.g. Koch (1999). The meanings of  $V_T$  and  $V_R$  are as follows: As we will see, most of the time,  $V(t) \in [V_R, V_T]$ . There is an overall upward trend, and when  $V$  reaches  $V_T = 1$ , the neuron fires an action potential, or a spike, after which its  $V$ -value is reset to  $V_R = 0$ , and will remain there for an absolute refractory period of 4ms.

For a neuron  $n$  of type  $Q \in \{E, I\}$ , its conductances satisfy

$$\tau_E \dot{g}_E = -g_E + S^{QE} \sum_{i=1}^{\infty} \delta(t - t_i^{\text{syn},E}) + S^{\text{dr}} \sum_{i=1}^{\infty} \delta(t - t_i^{\text{dr}}) \tag{2}$$

$$\tau_I \dot{g}_I = -g_I + S^{QI} \sum_{i=1}^{\infty} \delta(t - t_i^{\text{syn},I}). \tag{3}$$

These equations can be understood as follows. First there is the *synaptic input* received by neuron  $n$ : We assume a synaptic failure rate of 1/2 for all neurons, so that a spike from a neuron presynaptic to  $n$  will affect  $n$  with probability 1/2, independently from spike to spike. In Eqs. (2) and (3) above,  $\{t_i^{\text{syn},E}, i = 1, 2, \dots\}$  are the times at which a kick from one of the excitatory neurons in the network is (successfully) received by neuron  $n$ . Likewise,  $\{t_i^{\text{syn},I}, i = 1, 2, \dots\}$  are the times at which inhibitory kicks are received. Rise time in conductance upon arrival of an impulse is assumed to be instantaneous, hence the  $\delta$ -function in front of the summands in Eqs. (2) and (3), while decay times for excitatory and inhibitory conductances are taken to be  $\tau_E = 2\text{ms}$  and  $\tau_I = 3\text{ms}$ , respectively. (Convention regarding  $\delta$ -functions: the second term in Eq. (2), for example, should be interpreted to mean  $g_E$  jumps up by an amount equal to  $S^{QE}/\tau_E$  at time  $t_i^{\text{syn},E}$ ). The constants  $S^{QE}$  and  $S^{QI}$  represent the amplitudes of the kicks; they depend only on the types (i.e., E or I) of the pre- and postsynaptic neurons, and  $S^{QQ'}$  means from type  $Q'$  to type  $Q$ .

Additionally, neuron  $n$  receives a drive in the form of small Poisson kicks the arrival times of which are denoted by  $\{t_i^{\text{dr}}, i = 1, 2, \dots\}$ ; this is the third term on the right side of Eq. (2). These kicks are delivered at rate  $\gamma^Q \eta^{\text{dr}}$  and

coupling constant  $S^{\text{dr}}$ ; here  $\eta^{\text{dr}}$  and  $S^{\text{dr}}$  are system constants and  $\gamma^Q \geq 1$  is a constant that can be varied. We assume the arrival times of these kicks are independent for different neurons, and think of this term as representing the combined effect of a background drive and an external stimulation: the background drive has rate  $\eta^{\text{dr}}$ , and the external stimulation is modeled as an increase in the rate of this drive.

### 1.2 Remaining parameters

The parameters we have yet to discuss are

- the coupling strengths  $S^{QQ'}$ ,  $Q, Q' \in \{E, I\}$ ,
- parameters related to background drive, namely  $S^{\text{dr}}$  and  $\eta^{\text{dr}}$ , and
- parameters corresponding to the strengths of external stimulation, i.e.  $\gamma^Q$ ,  $Q \in \{E, I\}$ .

We have tried to use accepted biophysical parameters whenever we can, but there is little experimental guidance for these remaining parameters, and our choices below are not based on compelling biological reasons. Nevertheless, we explain how we arrive at the numbers we use, and introduce the two main parameters to be varied.

First we fix the quantities related to background drive. We assume  $S^{\text{dr}}$  is very small, and would like the firing rate of a neuron in the absence of synaptic input to be very low. In our simulations, we use  $S^{\text{dr}} = 0.0028$  and  $\eta^{\text{dr}} = 3.8 \text{ ms}^{-1}$ , which gives a firing rate of 0.5 spikes/s.

Next we choose  $S^{EE}$ . It is sometimes said that “15-20 excitatory kicks in quick succession should lead to a spike”, though the precise meaning of this statement is up to interpretation. We choose  $S^{EE} = 0.02$ . To give a sense of what this means, observe that the impact of an excitatory kick that arrives at time  $t_0$  depends on (i) the voltage  $V$  of the receiving neuron at time  $t_0$ , (ii) its conductances  $g_E$  and  $g_I$  at time  $t_0$ , (iii) the Poisson input received by this neuron after time  $t_0$ , and (iv) other synaptic inputs received in the several ms before and after  $t_0$ . First we simplify the situation by eliminating (iv) from the picture: in a low-firing regime, it is entirely possible to have no other synaptic event shortly before or after the arrival of the kick in question. Without (iv), it suffices to consider a single neuron. The following numerical experiment was performed: We subject a single neuron to background drive (only), fix a number  $V_0$ , and deliver an  $S^{EE}$ -size excitatory kick to it at various times when its  $V = V_0$ , making sure that these kicks are far enough apart that  $g_E$  and  $g_I$  have ample time to return to their natural values. We then compute the probability that this neuron spikes within 4ms after receiving the kick. For  $S^{EE} = 0.02$ , we find that at  $V_0 = 0.96$ , the probability of eliciting a spike is about 55 %, and for  $V_0 = 0.92$ , it is about 15 %. These are theoretical values. For a neuron embedded in a network, (iv) above generally cannot be ignored,

and taking (iv) into consideration, these spiking probabilities have to be revised upwards or downwards depending on the balance of excitation and inhibition received around the arrival time of the kick in question.

Next we set  $S^{IE} = \frac{1}{2}S^{EE}$ . This is to compensate for the fact that in our system, an I-neuron is postsynaptic to twice as many E-neurons as an E-neuron is (30 vs 15 %). That is to say, other things being equal, we would like the overall excitatory input into an E-neuron to be similar to that into an I-neuron.

With  $S^{EE}$  fixed, we will permit  $S^{EI}$  to vary, but instead of using it as our free parameter, we think it is more intuitive to consider the following: As  $\tau_E, \tau_I \rightarrow 0$ , each synaptic input would cause the voltage of the postsynaptic neuron to “jump”, and the size of this jump, denoted  $\delta V^{QQ'}$ , could be indicative of the synaptic strength from neurons of type  $Q'$  to neurons of type  $Q$  – assuming for definiteness that  $V = 1$  for the postsynaptic neuron when the synapse arrives. Thus for instance,  $\delta V^{EE} = S^{EE} * (14/3 - 1) = 0.073$ , and  $\delta V^{IE} = 0.073/2$ . Using this notation, we now define our first parameter,  $\alpha$ , to be

$$(\delta V^{EE}, \delta V^{IE}, \delta V^{EI}, \delta V^{II}) = 0.073*(1, 1/2, -\alpha, -\alpha/3).$$

Thus  $\alpha = 1$  means that under the idealized conditions above, an I-kick decreases the voltage of an excitatory neuron near threshold by an amount roughly equal to the increase in voltage caused by an E-kick.

We have taken  $S^{II} = \frac{1}{3}S^{EI}$ . This choice was made for reasons of conceptual simplicity: When  $S^{II}$  is comparable to  $S^{EI}$ , increasing  $S^{EI}$  (keeping all else constant) has the effect that each I-spike is more suppressive, but there are fewer I-spikes due to the increase in  $S^{II}$ ; thus the net effect of increasing  $S^{EI}$  on the E-population is unclear. Letting  $S^{EI}$  be somewhat larger than  $S^{II}$  reduces this competition, and enables us to think in the following simplistic way: *increasing  $\alpha$  increases the amount of inhibition in the system.*

The second parameter we consider is  $\gamma = \gamma^E$ . We will allow this constant to vary, but recall that  $\gamma = 1$  corresponds to the absence of external stimulation, and  $(\gamma - 1)\eta^{dr}$  is the drive rate for E-neurons due to external stimuli. Finally we assume  $(\gamma^I - 1) = \frac{1}{3}(\gamma - 1)$ , the  $\frac{1}{3}$ -factor taken in part to offset the ratio in  $S^{II}/S^{EI}$ .

With the exception of Section 4.2, where some additional parameters are explored, *the two parameters that will be varied in this paper are  $\alpha$  and  $\gamma$ :*

- increasing  $\alpha$  means increasing  $S^{EI}$  (and  $S^{II}$ ),
- $\gamma = 1$  and  $> 1$  will be referred to as the “undriven” and “driven” cases respectively;

all other parameters will remain fixed.

### 1.3 Firing rates

To give an idea of the combined effect of the parameter choices we have made, we show in Fig. 1 the firing rates of the networks as  $\alpha$  and  $\gamma$  are varied. As expected, firing rates decrease as  $\alpha$  increases for the undriven network, and they go up with  $\gamma$  for a fixed value of  $\alpha$ .

## 2 Emergent phenomenon of clustering

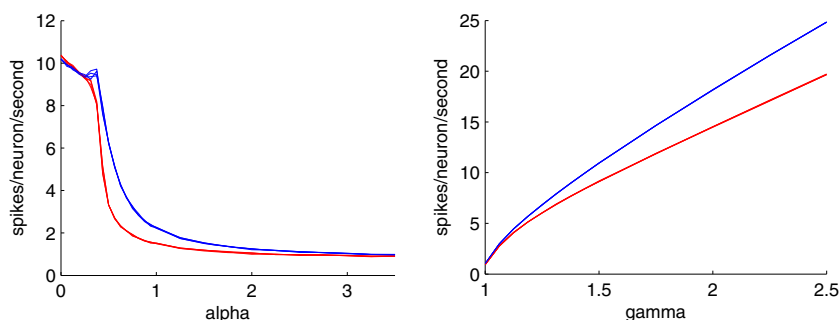
The goal of Section 2.1 is to demonstrate, through the presentation of several raster-plots, that spikes from a population such as ours tend not to occur homogeneously in time but in *clusters*. The goal of Section 2.2 is to formalize suitable definitions of “clusters”, and to devise metrics to aid in their statistical study.

### 2.1 Phenomenology

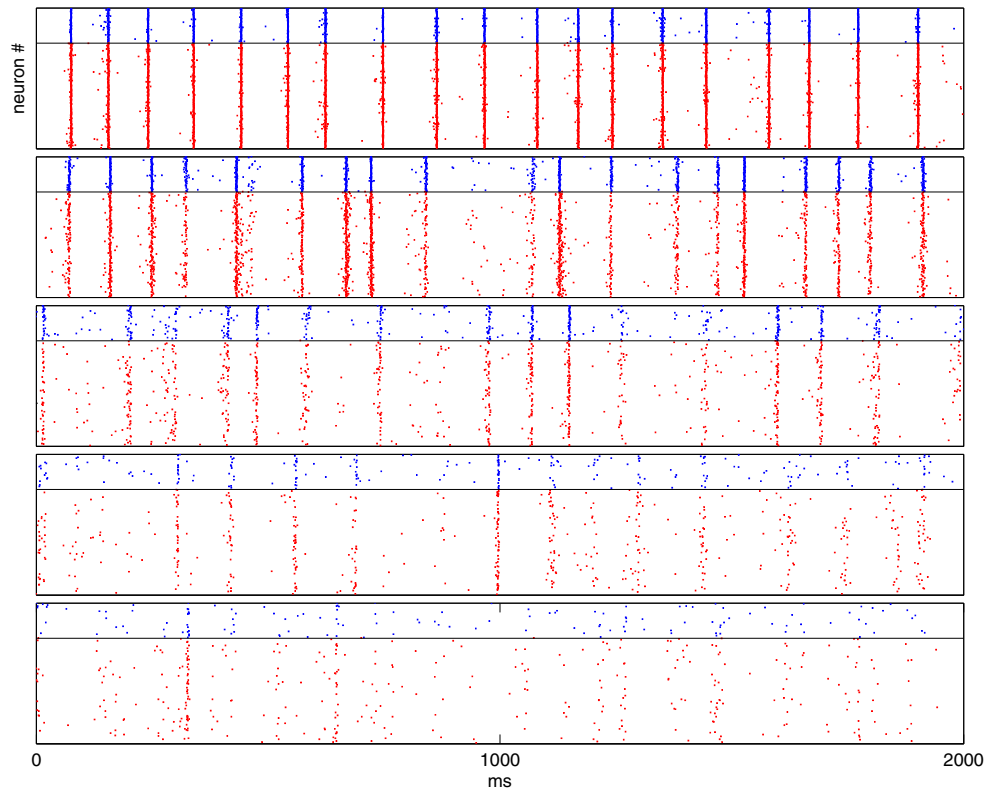
We have examined the spiking patterns of networks corresponding to many values of  $\alpha$  and  $\gamma$ ; a representative sample is presented in Figs. 2 and 3. Figure 2 shows raster-plots of five regimes with increasing  $\alpha$  from top to bottom. The regime in the top panel has very little inhibition and is not intended to be realistic (except possibly in special situations such as epileptic seizures), but we have found it instructive to use its relatively simple dynamics as a departure point to investigate the more complicated spike patterns that develop as the strength of inhibition is increased gradually.

Shown in the top panel of Fig. 2 are *population spikes*, i.e. near-synchronous spikes that involve nearly all of the neurons in the network, occurring semi-regularly at roughly 100 ms apart. This picture is typical of systems with very

**Fig. 1** Firing rates for excitatory (red) and inhibitory (blue) neurons: (a) as functions of  $\alpha$  with  $\gamma = 1$ , and (b) as functions of  $\gamma$ , for  $\alpha = 3$ . Shown in each panel are the superposition of 4 graphs, corresponding to 4 different realizations of our network



**Fig. 2** Raster-plots of the undriven network (i.e.  $\gamma = 1$ ) for various values of  $\alpha$ . From top to bottom:  $\alpha = 0.1, 0.45, 0.75, 1.5$  and  $3$ . Within each panel, the horizontal axis is time, in ms. Rasters for all I-neurons are shown on top (blue), E-neurons at the bottom (red)



low inhibition (e.g.  $\alpha < 0.35$ , see Fig. 1a), and the dynamical mechanisms behind it are relatively simple: When one excitatory neuron spikes, it raises the conductance of a number of postsynaptic E-neurons, making them more likely to spike, and when sufficiently many excitatory neurons cross threshold within a short time, just about the entire population will get pulled over. Here the effects of inhibitory neurons are minimal. Notice that the time lapses between consecutive population spikes are variable but have characteristic lengths. To understand the lengths of these time intervals, consider first the case  $\alpha = 0$ ,  $S^{EE}$  sufficiently large, and  $\tau_E, \tau_I \rightarrow 0$ . In this limit, the crossing of threshold by one excitatory neuron will cause the entire population to spike immediately. This implies that the times between population spikes are *first arrival times* for the population of E-neurons under independent background Poisson drive. More precisely, by “first arrival time”, we mean the time it takes the *first* of the 225 E-neurons to reach threshold when they all start at  $t = 4\text{ms}$  (end of refractory) with  $V = 0$  and a value of  $g_E$  equal to what remains of the rise in conductance caused by a population spike at  $t = 0$ .<sup>1</sup> In the panel shown,

<sup>1</sup>If the refractory period is absent or too short,  $g_E$  may increase with each population spike. Our 4ms refractory period was chosen to quench this build-up in conductance and subsequent runaway spiking. It has virtually no effect on regimes other than those with very small  $\alpha$ .

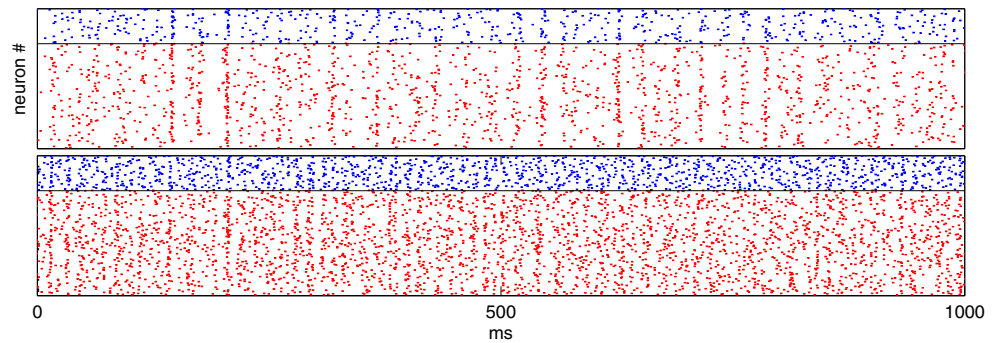
the mean of the first arrival time distribution is about 0.1s, much smaller than the mean arrival time of 2s for a single neuron (see Section 1.2). The situation considered here is an approximation of this limiting regime:  $\tau_{E,I}$  are small relative to other timescales in the system, and  $\alpha$  is small. The parameter  $S^{EE}$ , however, is also quite small; consequently it usually takes the spiking of a number of E-neurons to set off an avalanche.

The next panels show the breaking up and continued degradation of population spikes as  $\alpha$  is increased. At  $\alpha = 0.45$ , while large scale spiking events involving nearly the full population continue to occur, meddling by I neurons leads to a less tidy picture. At  $\alpha = 0.75$ , population spikes are no longer feasible, but instances of temporarily localized, coordinated spiking that involve substantial numbers of both excitatory and inhibitory neurons are still clearly visible. As  $\alpha$  increases, these numbers become smaller, the events become less clearly defined and occur with less regularity, but even at  $\alpha = 3$  (bottom panel), where each neuron fires only once per sec, one sees more coordinated firing than one would as a result of random chance alone.

We remark here on the strong nonlinearity of firing rates and firing statistics as functions of  $\alpha$ . Returning to Fig. 1a, one sees that between  $\alpha = 0.4$  and  $0.6 - 0.7$ , there is a sharp drop in firing rates. Increasing  $\alpha$  further, firing rates level off; as a matter of fact the last 3 panels in Fig. 2 also do not change dramatically. The situation can



**Fig. 3** Raster-plots of the driven network at  $\alpha = 3$  over 1 sec intervals. *Top:*  $\gamma = 1.33$ ; *bottom:*  $\gamma = 2$



be summarized as follows: For very low inhibition, the system fires in near-synchrony; such regimes are quite stable, meaning not very sensitive to small changes in parameters. As inhibition increases, the system settles down to a much lower firing regime which is also very stable qualitatively, and the transition from the first to the second is quick.

Figure 3 shows two sample raster-plots for the driven case. These plots are much “busier” as firing rates are considerably higher (hence we show only a 1 sec stretch), but columns of raster points that are vertically aligned are clearly visible in the weakly driven regime (top panel) and less clearly but still visible in the more strongly driven regime (bottom panel), suggesting temporally localized coordinated spiking events.

*Discussion* By just about any standard, the network regimes depicted in Figs. 2 and 3 exhibit very different dynamical behaviors. What they do have in common are instances of coordinated spiking involving various subpopulations, i.e., the tendency to spike in *clusters*. We will demonstrate in the pages to follow that the dynamical mechanisms behind many seemingly different spike patterns can be seen in a unified light if one is willing to accept a broad enough definition of “clusters”, namely one that permits (i) the involvement of subpopulations of variable sizes, and (ii) variable degrees of regularity and frequencies for inter-event times.

The 2-parameter family of networks defined in Section 1 will be used throughout to illustrate the ideas proposed. This family was not chosen for its clustering properties, but we do not know to what degree our model choice has favored the emergence of clusters. In addition to this family, we have examined many other examples of similarly constructed neuronal networks for different choices of connectivities and  $S^{Q'}$ , and have found them to exhibit spike patterns qualitatively similar to those illustrated in Figs. 2 and 3; in particular, some degree of clustering is seen in virtually all of them. (For examples of larger

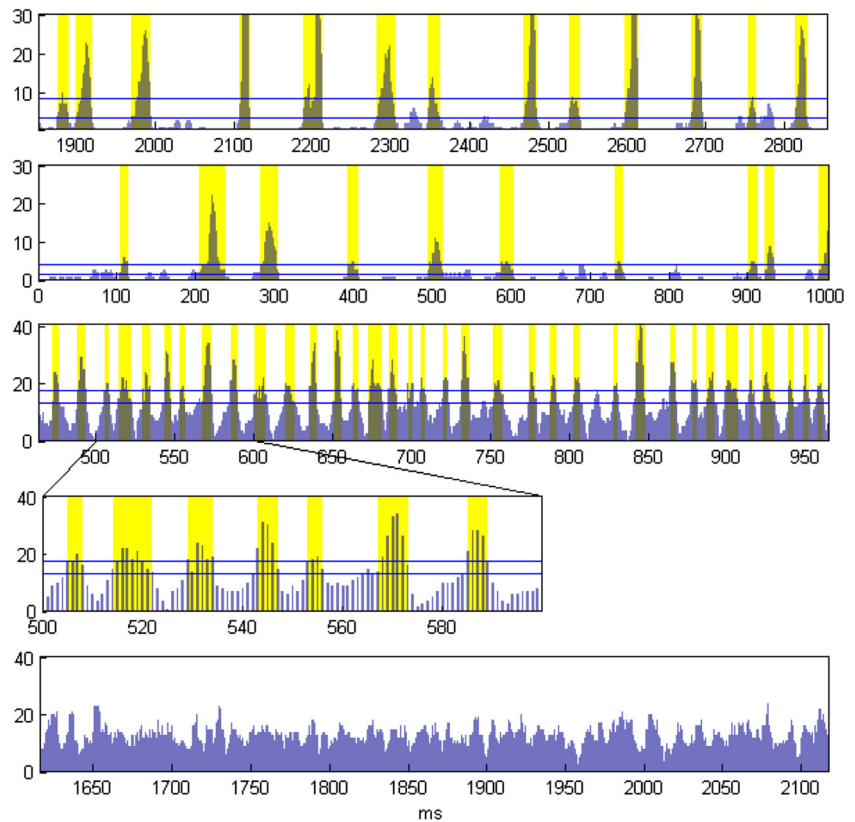
networks, see Rangan and Young 2013a, b). We have not conducted a systematic search in parameter space, as that would not be feasible, though some of the parameters that we have looked at are quite far apart. We also do not claim that spike patterns quantitatively resemble those shown, only that a certain basic phenomenon, namely what we call clustering, appears to be ubiquitous and recognizable.

## 2.2 Defining clusters

Clustered spiking is about the collective behavior in a population of neurons; it has to do not just with how individual neurons behave but how they relate to one another. From close-up examination of the rasters, we know it is not the case that the population in question is divided into distinct subgroups that co-activate; see Fig. 5b below. How, then, would we define a “cluster”? The examples shown tell us what *not* to do: We cannot require the participation of a fixed number of neurons or a fixed fraction of the population, nor can we require periodicity, and we certainly cannot dictate how frequently clusters occur. Requirements of the type above will unnecessarily restrict the scope of our definition.

The following properties, on the other hand, are relevant: First, if our definition is to be effective in detecting rises and falls in spike rates that occur on arbitrary temporal scales, then the investigator must be able to choose a *resolution* at which to view the patterns. Second, if our definition is to capture relevant behavior, then depending on the nature of the study, the investigator must be able to set the *scale*, meaning how large a burst of activity, or how far it has to deviate from the mean, for it to be considered a significant event. Thus our proposed definition will carry two parameters:  $\epsilon$ , which represents the resolution, and  $\delta$ , which represents the amount of deviation from the mean. For definiteness, all discussions of clusters will involve only the E-population, and we assume the time axis is partitioned into 1ms bins.

**Fig. 4** Population activity of several regimes. Shown in the top 3 panels are population activity as defined by the function  $\phi$  in the cluster algorithm above for 3 of the networks considered. Their parameters  $((\alpha, \gamma), (\epsilon, \delta))$  are  $((0.75, 1), (8, 1.5))$ ,  $((3, 1), (8, 1.5))$ , and  $((3, 2), (4, 0.33))$  respectively from top to bottom; inset shows a zoom-in of panel 3. In each of these panels, horizontal lines at height  $\mu$  and  $\mu(1 + \delta)$  are shown; and yellow regions are above bin-intervals defining clusters according to the algorithm. For comparison, we have included an analogous plot for a network of unconnected neurons (bottom panel) with independent background Poisson drive elevated to give a firing rate of 15 (similar to that in panel 3); the function  $\phi$  with  $\epsilon = 4$  is used



*Algorithmic definition of a cluster for given  $(\epsilon, \delta) \in (\mathbb{Z}^+, \mathbb{R}^+)$*

1.  *$\epsilon$ -convolutions:* Let  $\phi_0(n)$  be the number of E-spikes that occur in the time interval corresponding to bin  $n$ . If  $\epsilon$  is odd, we let  $\phi(n) = \sum_{n-(\epsilon-1)/2}^{n+(\epsilon-1)/2} \phi_0(i)$ , i.e., each spike in bin  $n$  contributes a “1” to  $\phi(i)$  for  $n - \frac{\epsilon-1}{2} \leq i \leq n + \frac{\epsilon-1}{2}$ . The even case is dealt with similarly.
2. *Candidate clusters:* Let  $\mu$  be the expected value of  $\phi$ , i.e.  $\mu = rN_E\epsilon/1000$  where  $r$  is the firing rate per sec for individual excitatory neurons and  $N_E$  is the total number of E-neurons in the population. Let  $J$  be an interval of time bins. We say  $J$  defines a candidate cluster if
  - (a)  $\phi(i) \geq \mu(1 + \delta)$  for some  $i \in J$  and
  - (b)  $J$  is the largest interval with the property that  $\phi(i) \geq \mu$  for all  $i \in J$ .
3. *Getting rid of short intervals:* Let  $\mathcal{J}$  denote the set of intervals  $J$  from above. The following two operations are performed, in the order indicated:
  - (a) If  $J_1, J_2 \in \mathcal{J}$  are such that the number of bins between them is  $\leq \epsilon/2$ , then we replace

- (a)  $J_1$  and  $J_2$  by the shortest interval  $J'$  containing  $J_1 \cup J_2$ .
    - (b) If  $J \in \mathcal{J}$  contains  $\leq \epsilon/2$  bins, we remove it from  $\mathcal{J}$ . (Such a  $J$  can, e.g., arise from two volleys, both insignificant, occurring slightly  $< 2\epsilon$  bins apart).
- Let  $\tilde{\mathcal{J}}$  denote the resulting collection of intervals.
4. We declare that each  $J \in \tilde{\mathcal{J}}$  defines a cluster of size  $\frac{1}{\epsilon} \sum_{i \in J} \phi(i)$ .

This is admittedly an *ad hoc* definition, but it is simple conceptually and the algorithm is simple to implement. Nor do we claim that the ideas here are completely novel: algorithms that are not markedly different have been used by experimentalists to capture inhomogeneous spike patterns before. In Hahn et al. (2010), for example, time bins of a few ms are considered; bins containing at least one spike recorded from an electrode array are marked, and contiguous bins separated by blank ones are considered as “events”. The idea of a “multiple firing event” (MFE) in Rangan and Young (2013a, b) is a little different: since the fast synaptic time constants  $\tau_E, \tau_I$  have been taken to zero, it is possible to determine if one spike directly “causes” another, and an MFE is defined to be a maximal sequence of firings that are caused by a single neuron. Without claiming that one

algorithm is better than another (that, needless to say, is dependent on circumstance), our goals are to propose some *general ideas* in the hope of facilitating comparisons and promoting theoretical studies.

We now demonstrate that the algorithm above “works”, meaning for suitably chosen  $(\epsilon, \delta)$  it captures in a reasonable way what we set out to capture, and is flexible enough to incorporate a wide range of circumstances.

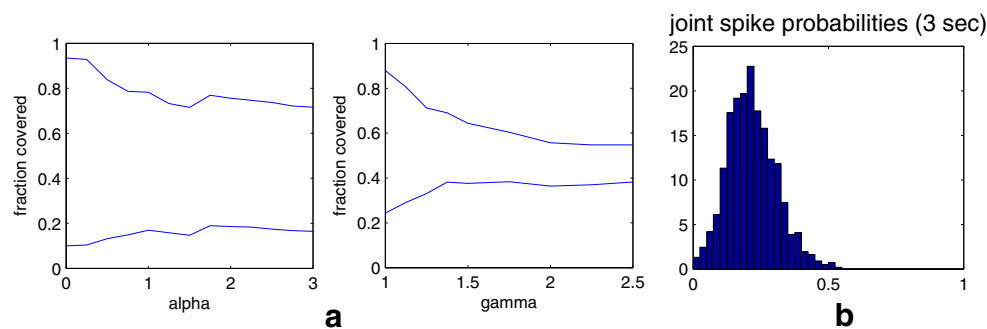
For illustration, we have included in Fig. 4 the outputs of this algorithm for three quite different regimes: at  $\alpha = 0.75$  and 3 (undriven) and for a driven regime with  $\gamma = 2$ . Regimes that produce population spikes are straightforward and omitted. The rationale behind our choices of the parameters  $(\epsilon, \delta)$  in Fig. 4 is as follows: In the top two panels, firing rates are so low that had we used a smaller value of  $\epsilon$ , say  $\epsilon \leq 5$ , the relevant number of spikes used to determine bin intervals would be no more than one or two, and that is not meaningful unless one wishes to consider single spikes as clusters. On the other hand, clusters in these regimes are sufficiently far apart that a larger  $\epsilon$  does no harm. As for  $\delta$ , we chose to consider only moderately large deviations (again because firing rate is low), requiring that  $\phi$  rise to at least 2.5 times the mean before registering the event as a cluster. The third panel together with its zoom-in are for a strongly driven regime. Here we did not wish to use  $\epsilon$  much larger than 4, as the distances between clusters as they naturally occur may not be much larger. Also, in driven regimes, there is a constant stream of E-neurons crossing threshold, so that  $\phi$  is, on average, nontrivial even in the valleys. For such regimes, one would expect the fluctuations to be smaller, and numbers on the order of  $\delta = 1/3$  are likely to capture better the inherent oscillations in the system.

The bottom panel in Fig. 4 is included not as illustration of the cluster capture algorithm but to show how

summed population activity might look in a network in which the neurons are unconnected to one another, i.e. there is no neuron-to-neuron interaction, and each neuron is driven by an independent Poisson drive. The feature that most strongly distinguishes this “null” case from the network regimes shown is that its fluctuations occur in multiple timescales, ranging from very small (a few ms) to quite large (more than 50), whereas in the panels above, consecutive clusters are separated by recognizable characteristic timescales.

To be clear, we do not assert, or even suggest, that for *all* choices of  $\epsilon$  and  $\delta$ , there is significant clustering in our network model (compared to the null model). For each regime, there are choices that will accentuate the clustering property, and others for which it will be less prominent and even statistically insignificant. Thus all statements regarding clustering (as defined in this paper) pertain to specific values of  $\epsilon$  and  $\delta$ .

Figure 5a provides further confirmation that our cluster capture algorithm along with our choices of resolution and scale are meaningful: In the undriven network (left), where clustering is strong, large fractions of spikes are concentrated in small fractions of time bins. In the driven case (right), given that there are constant streams of voltages crossing threshold, one cannot expect the numbers to be as clear-cut, but even at  $(\alpha, \gamma) = (3, 2)$ , for example, more than 55 % of the spikes are contained in a little over 35 % of the bins. Such scales of fluctuation are still fairly visible and not unrealistic. In Fig. 5b, we provide strong evidence to our assertion at the beginning of this subsection that in the clustered spiking observed, neurons are not divided into groups that spike together. In the regime shown, each neuron has, on average, a probability of about 0.18 of appearing in any one cluster. From Fig. 5b, we see that during the 3s following a



**Fig. 5** **a** Coverage: *Upper graph* in each panel shows the fraction of spikes captured in clusters, and *lower graph* shows the fraction of bins they occupy. *Left panel* shows undriven network with  $(\epsilon, \delta) = (8, 1.5)$ ; *right panel* shows the  $\alpha = 3$  network with  $(\epsilon, \delta) = (4, 1/3)$ . **b** Conditional distribution of participation in cluster: For 1000 randomly

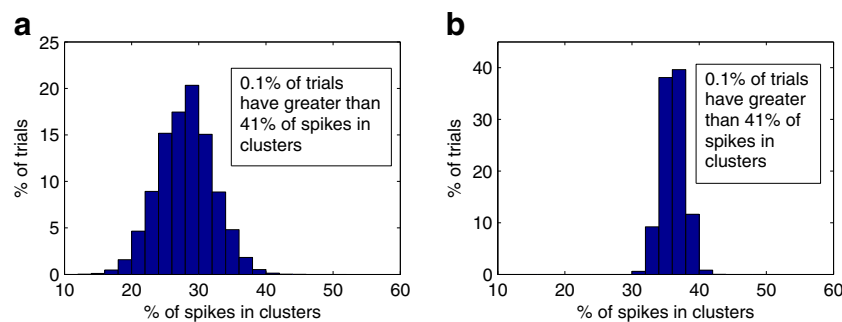
chosen pairs of neurons, named A and B, we compute the conditional probability that neuron B participates in a cluster given that neuron A is a participant, during the 3s following a joint participation. The regime is  $\alpha = 0.75$ ,  $\gamma = 1$ , the *middle panel* in Fig. 2



joint appearance by 2 neurons A and B in a cluster, the conditional probability of B spiking in a cluster given that A spikes in it is, on average, only slightly higher than 0.18. Moreover, there is a nice spread in these conditional probabilities for different pairs, depending on their network connectivities and other circumstances.

The regimes in Fig. 4 exhibit quite obvious clustering. In situations with greater ambiguity, one may wish to carry out an analysis to determine if the amount of clustering seen is statistically significant. There are many ways to do this. One possibility is to compare to a null model as seen through the lens of our cluster capture algorithm. That is to say, given a regime for which we wish to establish  $(\epsilon, \delta)$ -clustering, we proceed as follows. Suppose  $r$  is the mean firing rate per sec of individual E-neurons, and  $N_E$  is the number of excitatory neurons in the system. We divide a time window of  $n$  sec into 1ms bins, and independently for each of the  $rnN_E$  spikes, we pick one of the  $1000n$  bins with uniform distribution and put the spike into this bin. Then for the same  $(\epsilon, \delta)$  as above, we determine the collection of clusters using the algorithm at the beginning of this subsection, and from that we compute the percentages of spikes that are contained in clusters (as was done in the top curves of Fig. 5a). Figure 6 shows histograms of fractions of trials corresponding to these percentages for  $n = 3$ . In (a), the parameters used are those for  $(\alpha, \gamma) = (3, 1)$  and  $(\epsilon, \delta) = (8, 1.5)$ , the same regime and scales as in the second panel of Fig. 4. As can be seen, for a null model with these parameters, typically 20 – 36 % of the spikes occur in clusters, and the probability that clusters account for > 41 % of all spikes is < 0.001, compared to the ~ 70 % of spikes in clusters in our model regime (see Fig. 5a). Similarly, panel (b) corresponds to  $(\alpha, \gamma) = (3, 2)$  and  $(\epsilon, \delta) = (4, 0.33)$ , as in the third panel of Fig. 4.

For related works on statistical significance in relation to firing patterns, see e.g. Okun et al. (2012).



**Fig. 6** Histogram of fraction of spikes captured in null model matching (a) undriven and (b) driven firing rates, over 10,000 trials. For each trial, a number of spikes corresponding to the firing rates of each regime are distributed uniformly over a 3 second interval, and the

### 3 Statistics of clusters

In this section we view population activity in terms of clusters and times between clusters. We will call a cluster an *event*, the size of a cluster as defined in the last section (item 4 under “algorithm”) its *event size* (ES), and the distance between the midpoints of the time intervals corresponding to two consecutive clusters as their *inter-event time* (IET). In the discussion to follow we will view population activity as a sequence of

$$\text{event} \rightarrow \text{pause} \rightarrow \text{event} \rightarrow \text{pause} \rightarrow \text{event} \rightarrow \dots$$

and will examine the probability distributions for ES and IET and relevant correlations.

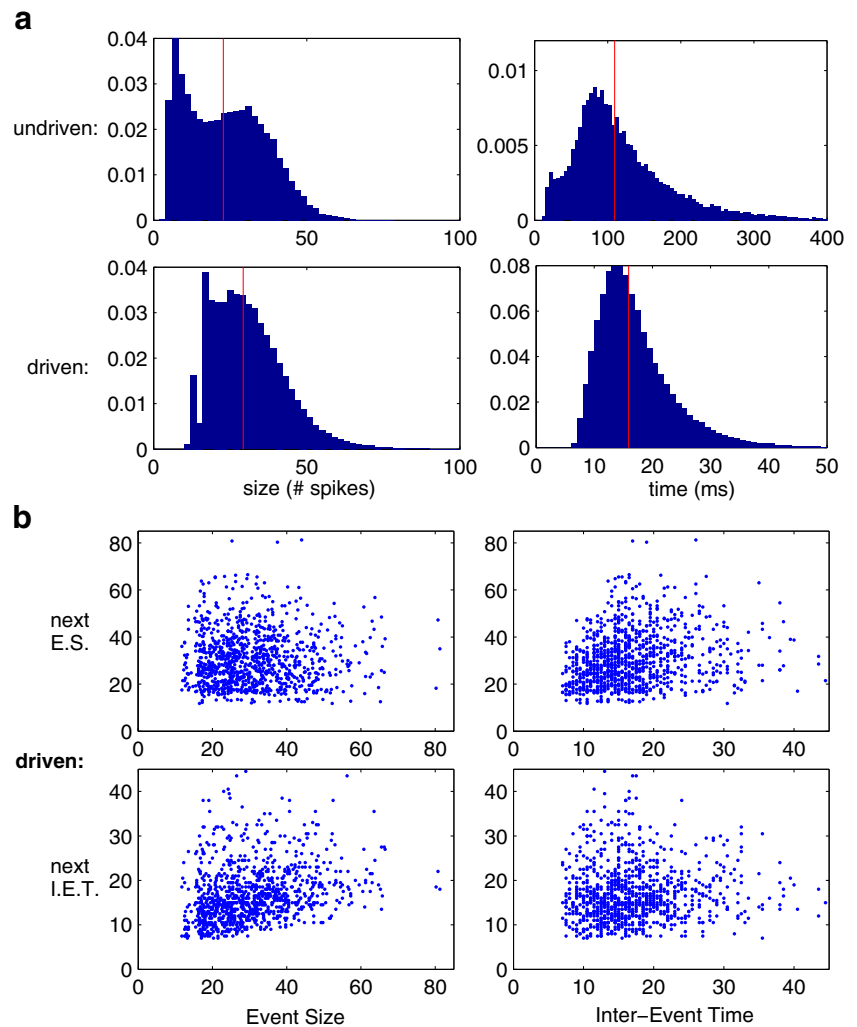
#### 3.1 Event sizes, inter-event times and their correlations

Distributions of event sizes and inter-event times for two sample networks, one driven and one undriven, are shown in Fig. 7a. Corresponding plots for networks with other values of  $\alpha$  and  $\gamma$  are not too different qualitatively—except for the case of undriven networks with very small  $\alpha$ , which we do not discuss (events there involve nearly full populations, and IET distributions are explained in Section 2.1).

We observe from Fig. 7a that the distribution of ES in this undriven network (top left) is *very wide*: cluster sizes can range from a few neurons to a third of the E-population. With such large variations, it makes little sense to speak of a “typical” event size. The IET distribution of this network (top right) is also extremely wide: they run from about 10 to over 300ms even as we discard outliers at the two ends. Similar comments apply to the ES distribution of the driven case (bottom left). The most Gaussian-like of the four is the IET distribution of the driven network, in which external forcing has imposed a certain amount of regularity on its spiking activity. Here, the bulk of the data points

cluster capture algorithm is applied with the same  $(\epsilon, \delta)$  as previously used with each regime. Specifically, (a) corresponds to  $(\alpha, \gamma) = (3, 1)$  with  $(\epsilon, \delta) = (8, 1.5)$ , and (b) corresponds to  $(\alpha, \gamma) = (3, 2)$  with  $(\epsilon, \delta) = (4, 0.33)$

**Fig. 7** Event sizes and inter-event times. The 4 panels in **(a)** show distributions of ES and IET. *Left: ES; right: IET. Top: undriven network with  $\alpha = 1.5$ ; bottom: network with  $\alpha = 3$  and  $\gamma = 2$ .* For ES plots, the *mean* of the distribution is marked by a *red line*; for IET plots, the largest 5 % of inter-event times are discarded (as they may be artifacts of our algorithm) and the *median* is marked by a *red line*. The 4 panels in **(b)** show scatterplots of X followed by Y where  $(X, Y)$  range over all possible pairs of “ES” and “IET”; data shown are from the  $\alpha = 3, \gamma = 2$  network



falling between 10 to 35 ms, a range that is reminiscent of broad-band gamma oscillations observed in real cortex (Henrie and Shapley 2005).

We remark that the wide distributions of ES for undriven networks reported above are consistent with experimental observations in e.g. Hahn et al. (2010) and with the emergent dynamics in the computational model of V1 in Rangan and Young (2013b). Both papers reported power-law distributions, which are in fact part of the definition of “neuronal avalanches” in Hahn et al. (2010). We have not shown our probability density functions in log-log scale or discussed power laws because with our limited population size, it is not clear how meaningful power law distributions for ES will be, nor are we especially interested in very low probability events.

We consider next the correlations between consecutive ES and IET. From the scatterplots shown in Fig. 7b, we observe that some correlations clearly exist, the most obvious one of which being that large events cannot be either

preceded, or followed, by short IET; a proposed explanation will be given in the next section. All in all, however, correlations between current ES and the next IET, or current IET and the next ES, is not strong. Moreover, another half-step later, i.e. from one ES to the next ES, or from one IET to the next IET, biases have subsided considerably, as evidenced by the fact that the joint distributions shown look very much like product distributions.

To document these observations, we compute the variational distances between the joint distributions shown and product measures. More precisely, we view each of the scatterplots shown as representing a probability distribution in the  $x$ - $y$  plane. Discretizing the space by putting  $m$  spikes into one bin for ES and  $n$  ms into one bin for IET, we work with discrete probability distributions supported on a finite grid. The grid sizes used are chosen on a case-by-case basis to ensure that there are enough data points for meaningful comparisons but that local variations, which do not interest us, will not be too dominant. Defining the variational

**Table 1** Distance from product measure (driven)

	Next E.S.	Next I.E.T.
Current E.S.	0.011	0.12
Current I.E.T.	0.11	0.032

distance between two discrete probability distributions  $\mu$  and  $\nu$  by the formula

$$d(\mu, \nu) = \frac{1}{2} \sum_z |\mu(z) - \nu(z)| \tag{4}$$

where the summation runs over all grid points, we compute, for the distribution corresponding to each one of the scatterplots,  $d(\mu, \mu_1 \times \mu_2)$  where  $\mu$  is the joint probability distribution discretized, and  $\mu_1$  and  $\mu_2$  are its two marginals. By definition, then,  $0 \leq d(\mu, \mu_1 \times \mu_2) \leq 1$ , and a distance of 0 means that  $\mu$  is a genuine product measure.

Table 1 below shows  $d(\mu, \mu_1 \times \mu_2)$  for the driven network whose distributions are depicted in Fig. 7b, with IET time bins of size 2ms and ES bins of size 5 spikes. The values shown were computed from a simulation lasting 600 sec during which about 35,000 events occurred. Corresponding variational distances for the undriven network with  $\alpha = 1.5$  using IET time bins of size 20ms and ES bins of size 5 spikes are shown in Table 2.

To summarize, our observations regarding ES and IET distributions in driven and undriven regimes are consistent with experimental and modeling data. As for their time correlations, to our knowledge the rapid decorrelation in ES/IET has not been documented before.

### 3.2 A stochastic model of population activity

The discussion above suggested the following stochastic model for describing the population activity of a large class of networks for integrate-and-fire neurons. It consists of a sequence of random variables

$$X_1, Y_1, X_2, Y_2, X_3, Y_3, \dots$$

where  $X_i$  is the size of the  $i$ th event and  $Y_i$  is the inter-event time between  $X_i$  and  $X_{i+1}$ , these numbers being drawn from computed distributions of ES and IET for the network in question. In a simplest version of this model, one can assume, based on the weak correlations seen in Section 3.1,

**Table 2** Distance from product measure (undriven)

	Next E.S.	Next I.E.T.
Current E.S.	0.026	0.13
Current I.E.T.	0.055	0.042

that these random variables are independent (perhaps requiring that a large value of  $X_i$  cannot be accompanied by small values of  $Y_{i-1}$  and  $Y_i$ ). A slightly more sophisticated version of the model might be to view the process as Markov, with conditional probabilities given by the “ES and next IET” and “IET and next ES” scatterplots. For a more refined model still, one can consider the duration of an event in addition to its size, and so on, adding any amount of detail that one desires.

It is not our aim here to advocate any one specific model. The message we wish to convey is that since clustering is a salient characteristic of population activity in many neuronal systems, and since both event sizes and inter-event times appear quite random with rapidly decaying time correlations, stochastic models of the type above may be a simple yet reasonable way to describe their collective spiking behavior.

Our proposal offers an alternative to the usual power spectra approach, in which one seeks to tie semi-regular behavior to periodicity and preferred frequencies. While spectral densities summarize voltage fluctuations in a tidy way, the stochastic models above are easier to interpret and possibly easier to implement when IETs have very wide distributions. The most significant difference between the two, however, is that the models we have proposed contain information on *event sizes* as well as inter-event times, and these two pieces of information together give a more complete description of neuronal population activity.

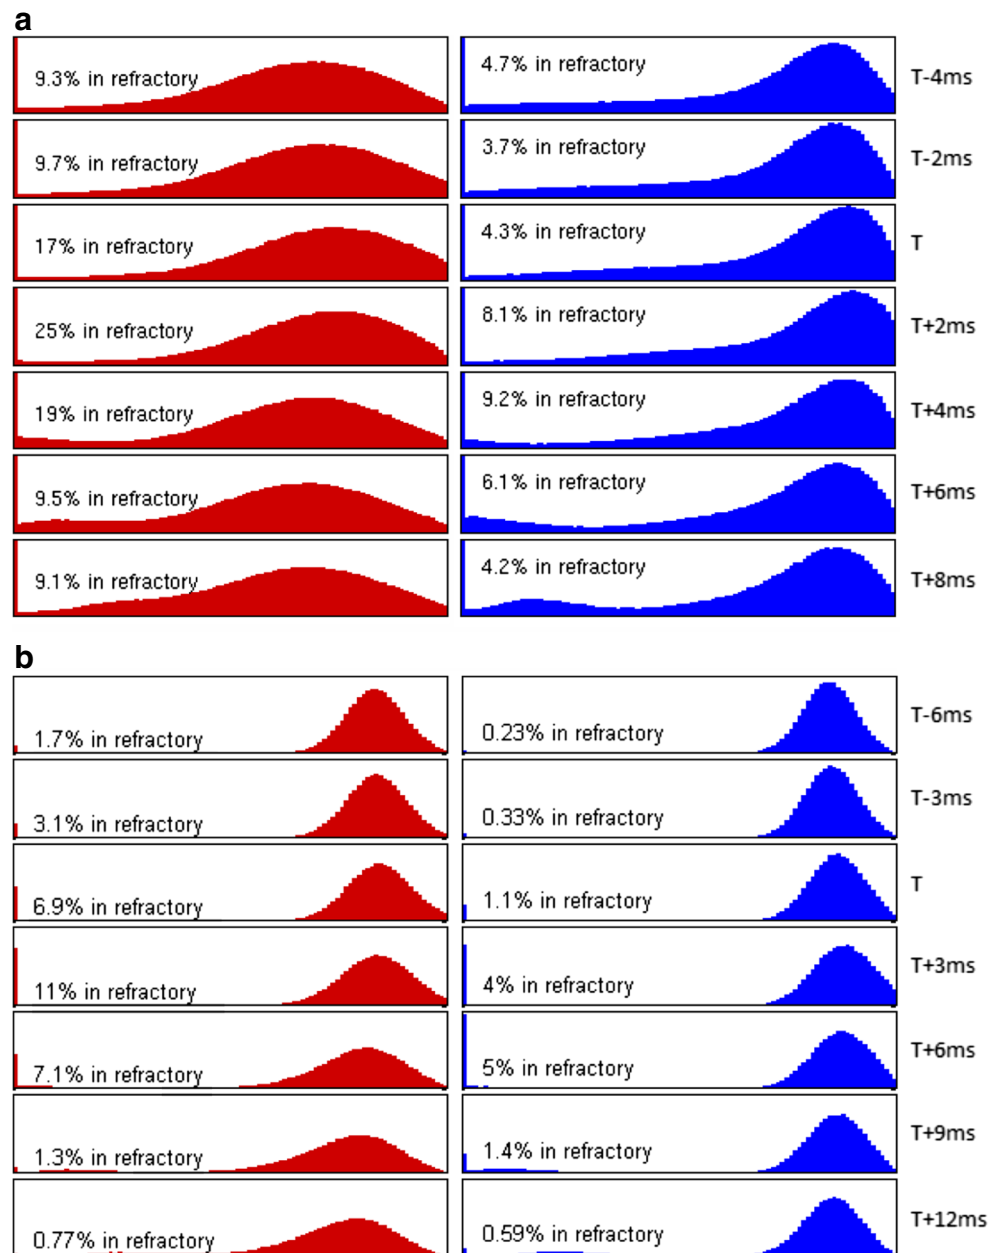
## 4 Dynamical mechanisms

This section contains some dynamical explanation for the phenomena observed. As discussed in the Introduction, the phenomenon of clustering is similar, though not identical, to the “multiple firing events” or MFEs in Rangan and Young (2013a, b). This section will have some small overlap with these two papers. As a rigorous analysis of clustering is currently out of reach, we would like to present here some simulation results that we hope will illuminate the dynamical picture and serve as supporting evidence as we test out some simple hypotheses.

### 4.1 Dynamics of voltage and spike-rate fluctuations

To understand how these fluctuations come about, we consider first the driven case, where the rhythm is more regular. A highly averaged “movie” illustrating the prelude to clustered spiking and the aftermath of such an event is shown in Fig. 8. When interpreting this figure, it is important to remember that the histograms shown are averaged over many events; what happens in each event is much less tidy

**Fig. 8** Time evolution of voltage densities. **a** Driven network with  $\alpha = 3$ ,  $\gamma = 2$ . **b** Undriven network with  $\alpha = 1.5$ . Shown are (*normalized*) voltage densities taken at the 7 points in time indicated,  $T$  being the point in time at which the cluster peaks; E on *left*, I on *right*. Each of the pdfs shown is the average of such snapshots over a large number of clusters



and there is great diversity from event to event; see Fig. 10. Nevertheless, we think this averaged version makes transparent the mechanism of the up-down swings in spike rates seen in rasters and other plots shown earlier:

Looking at the voltage distribution of the E-population in Fig. 8a, we see in the top 3 snapshots an obvious rightward drift, brought about by recurrent excitation triggered by the crossing of threshold by one or more E-neurons. This surge in E-spiking triggers a surge in I-spiking, but that peaks about 2 ms later. The surge in I-spiking then pushes both the E- and I-voltages back, away from threshold. The pushback is completed by the 6th snapshot. In the last 2-3 snapshots, the main action is the rightward march toward the main

bulge of voltages that were set to 0. Since the mean IET for this network is 15 – 20 ms, a new cycle will, on average, begin in several ms. We summarize the action depicted as

E-spike  $\rightarrow$  recurrent excitation  $\rightarrow$  surge in E/I-spikes  
 $\rightarrow$  I-pushback  $\rightarrow$  recharge  $\rightarrow \dots$

This mechanism is reminiscent of the rhythm generation mechanism known as PING. Similarities and differences are discussed in some detail at the end of this subsection.

We now attempt to explain the larger off-diagonal entries in Table 1 (Section 3.1); corresponding scatterplots are shown in Fig. 7b. Looking at the scatterplot for ES vs next IET, we observe that a significant contribution to the

correlation computed appears to be that the lower bound for “next IET” increases monotonically with event size. This is likely due to the fact that the larger an event, the larger is the minimum number of I-spikes needed to bring an end to the E-spiking. Since I-neurons are more densely connected to all other neurons, the I-pushback described above is likely felt by all neurons. With the voltages of all neurons some distance from threshold, there is a minimum recharge time, and we propose that this is reflected in the lower bound seen in Fig. 7b.

The aftermath of a smaller event is more complicated, as not all neurons are uniformly affected by the I-pushback, and what happens in the 5 – 20 ms thereafter depends to a nontrivial degree on “luck”, meaning the configuration of E and I-voltages that remain close, or relatively close, to threshold and how the corresponding neurons happen to be connected to one another. Indeed, even following a medium size event, it is possible for some E and/or I-neurons to be less affected by the I-pushback than others. (We stress that the plots in Fig. 8 are averaged over many events; unaveraged “movies” can be very messy. See Fig. 10). A partial explanation for why large events almost never follow small IET may have to do with this: Since a short IET is likely to be preceded by a not-so-large event, the situation is, as explained above, likely to be “messy”: since smaller events can leave behind neurons near threshold, the smallness of the next event can be due to intervention of I-neurons hovering near threshold, or a group of E-neurons crossing threshold while the bulk is not yet fully recharged. Scenarios of this kind have been observed time and again in the “movies” that we have viewed, though there are likely myriad other possibilities.

The dynamical mechanism described above applies also to undriven networks (Fig. 8b). There are, however, notable differences between the driven and undriven cases, resulting in more regular rhythms in driven regimes and wider distributions of IET for undriven populations: In an undriven network, between events voltages appear to collect near an equilibrium some distance below threshold (at  $V$  slightly  $> 0.8$  for the regime shown), where they remain until, possibly by chance, a large enough E-spiking event occurs to trigger what we define to be a cluster — in contrast to the driven case where the voltages are in a constant march from reset to threshold (as a result of the drive). Thus undriven networks fire volleys at more random times, i.e., their IET distributions are wider, because the occurrence of the next event relies more on random fluctuations than in the driven case. This also explains why, unlike the driven case, correlations between current IET and next ES are smaller; see Table 2.

*Discussion: comparison with PING* We digress here to comment on the relation to PING (pyramidal-interneuronal

network gamma), a previously known mechanism for rhythm generation. To the degree that both involve the interplay between excitatory and inhibitory populations, the ideas above and PING have much in common. Their *outputs*, however, are quite different, and we will argue that our mechanism produces voltage fluctuations that show greater resemblance to gamma band oscillations in the real brain.

PING was first proposed in Whittington et al. (2000), in an attempt to explain EEG oscillations in terms of synchronized activity in neuronal populations. This mechanism was subsequently studied in much greater detail in a series of computational papers of Kopell et al; see e.g. Börgers and Kopell (2003, 2005) and Börgers et al. (2005). In Kopell’s terminology, PING refers to the phenomenon in which near-periodic population spikes are produced in a network of excitatory and inhibitory neurons: E-cells receive identical input in the form of a strong steady current driving them significantly above their spiking threshold, while I-cells are primarily driven by E-cells. A population spike by E-cells triggers a population spike of I-cells, which then inhibits the E-population, leading to a period of low activity. When this inhibition wears off, the E-population spikes again in a synchronized fashion, and the cycle is repeated. Kopell also discussed a phenomenon she called “weak-PING” (Börgers et al. 2005). In weak-PING, E-cells are stochastically driven, and only a fraction of them participate in each event in which the entire I-population spikes (this was also discussed in Whittington et al. (2000)).

The mechanism discussed in this subsection is quite far from PING, as it generally does not involve population spikes. It is closer to weak-PING both in terms of the stochastic drive and the participation of subsets of E-neurons, but differs from weak-PING in that our I-cells do not produce population spikes. For the most part, Kopell’s models are based on the idea that E-cells are driven by external input, I-cells are driven by E-cells, and once aroused they totally suppress all E-cells. In our models, both  $S^{EE}$  and  $S^{II}$  are present, and the competition between the E- and I-populations is more even; neither one is able to overwhelm the other.

We now turn to comparison with experimental data. Gamma oscillations of LFP in cortex are well known to be broad-band, with elevations in a frequency range of roughly 25 Hz to 80 Hz (Henrie and Shapley 2005). PING produces a highly regular rhythm, in the sense that its E-spikes (whether they involve the entire population as in PING or a fraction of the population as in weak-PING) are evenly spaced in time. This corresponds to a power spectrum that is very close to a delta function. Our much wider IET distributions for driven networks are more consistent with broad-band gamma oscillations. For this reason, we propose that our mechanism may be an improvement over PING.



## 4.2 Further illustrative evidence

We provide here two sets of numerical evidence in support of the ideas discussed earlier.

- (1) *The role of recurrent excitation* in the dynamical mechanism described in Section 4.1. To understand the role played by  $S^{EE}$ , we start with the (3,2)-network, i.e., the network corresponding to  $\alpha = 3$  and  $\gamma = 2$ , and investigate the effects of exchanging  $S^{EE}$  for external drive. The second plot in Fig. 9 has the same parameters as the simulation from panel 3 in Fig. 4 without the yellow markers for clusters, i.e., it shows the function  $\phi$  in the cluster algorithm for the (3,2)-network, with  $\epsilon = 4$ .

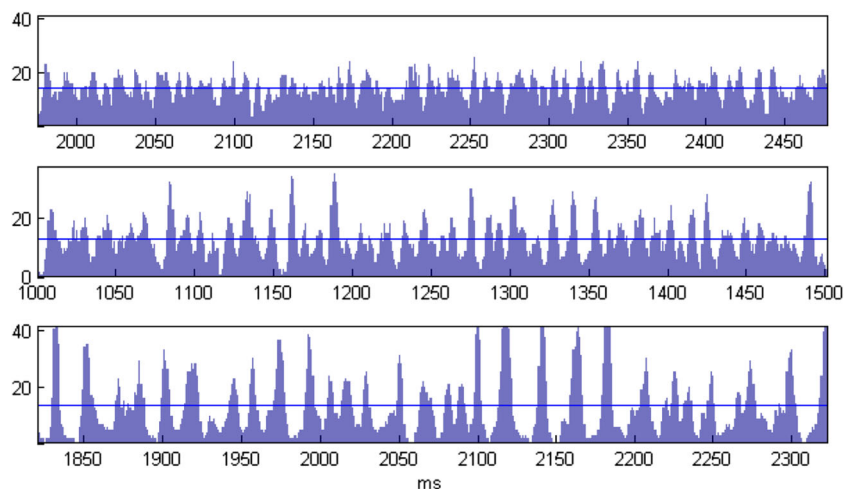
First, we investigate the effect of setting  $S^{EE}$  equal to 0 – while keeping  $S^{IE}$ ,  $S^{EI}$  and  $S^{II}$  exactly as they are in the (3,2)-network, i.e., severing the ties of these parameters to  $S^{EE}$ . In the interest of maintaining the amount of synaptic input to the I-population, we increase the external drive to the E-population (keeping the drive to the I-population unchanged) to bring E-firing rate to roughly the same level as in the (3,2)-network. The  $\phi$  function with  $\epsilon = 4$  for a 500ms stretch of the resulting regime is shown in the top panel of Fig. 9. The dynamical mechanism discussed in Section 4.1 is still at work, producing rises and falls in local spike rates, but we now see a faster rhythm involving smaller clusters, consistent with the fact that in the absence of recurrent excitation, the number of excitatory neurons aroused in each event is smaller (though the external stimulus is still driving a steady stream across threshold). That in turn leads to smaller surges in I-neurons, weaker I-pushbacks than in the (3,2)-network, and shorter recharge times.

The third panel shows a regime in which  $S^{EE}$  is 50 % larger than that in the (3,2)-network, and external drive to the E-population is adjusted to maintain the same E-firing rate, while all other parameters are unchanged. As a consequence, cluster sizes are larger, with stronger pushbacks and slightly larger IET.

- (2) *Competition between the E- and I-populations.* Here we hypothesize that much of what happens hinges on the outcome of the “race” between E- and I-voltages, first to reach threshold and then to act on their postsynaptic neurons. To test this hypothesis, we start with the undriven network with  $\alpha = 1.5$ , change one of the parameters in each of the examples below to put the I-population at either an advantaged or disadvantaged position in this competition, and confirm numerically the predicted outcomes.

- (a) Changing  $\tau_{\text{leak}}^I := \tau_{\text{leak}}$  for I-neurons (while keeping  $\tau_{\text{leak}}^E = 20\text{ms}$ ). Decreasing  $\tau_{\text{leak}}^I$  to 18ms, for example, has the effect of putting the equilibrium voltage distribution in Fig. 8b farther from threshold. One would anticipate that this should give an edge to the E-neurons, which can arouse one another with less interference from I-neurons, whose arrival times have been delayed. If so, larger clusters should result, and *vice versa*. This is indeed the case; see Table 3.
- (b) Changing  $(\tau_E, \tau_I)$ . Another way to alter the effectiveness of the E- or I-population is to change how quickly a neuron can affect its postsynaptic neurons when it spikes. Recall that we have used  $(\tau_E, \tau_I) = (2, 3)$  throughout. Examples of a few other  $(\tau_E, \tau_I)$ -configurations are studied. Setting  $(\tau_E, \tau_I) = (4, 1)$ , for example, will delay the impact of E-spikes, and permit I-neurons to exert their influence considerably sooner. That

**Fig. 9** Summed spikes for three different regimes. *Middle:* the (3,2)-network. *Top:*  $S^{EE}$  in the (3,2)-network is set to zero in exchange for larger external drive. *Bottom:*  $S^{EE}$  in the (3,2)-network is increased by 50 % of the expense of external drive. See the main text for detail



**Table 3** Influence of  $\tau_{leak}^I$  on mean cluster size

$\tau_{leak}^I$	18	19	20	21
mean cluster size	94	53	23	5.1

should discourage cluster formation, a prediction confirmed in Table 4 below. By a “single firing event”, we refer to an E-neuron crossing threshold without another one to follow within 5ms.

### 4.3 Why are cluster sizes so broadly distributed?

We have seen in Fig. 7a that the distributions of cluster sizes in both driven and undriven networks are extremely broad. This is because when an excitatory neuron spikes, the number of possibilities of what can follow is enormous: The spiking neuron has a number of postsynaptic neurons, both E and I, which may or may not be affected due to synaptic failure. Those that are affected may or may not spike depending on their voltages and conductances at that point in time. That determines the first round of action. Those postsynaptic neurons that spike then go on to affect other neurons, which may affect yet others, and so on. In a population of a few hundred neurons such as ours, the number of possible scenarios in the 10ms or so following the spiking of an E-neuron is innumerable, and many of them give rise to clusters of various sizes.

To showcase the complexity of the situation, we present in Fig. 10 three dramatically different pictures following an excitatory spike. In each case, we have included the distribution of E-voltages on (0.8, 1), marked in red those neurons that spike within 10ms or so of the initial spike, and shown (to the degree that is possible) the combinatorial tree of how one spike leads to the next. It must be stressed that this is far from the full picture: we have included neither the action of inhibitory neurons nor the conductances of the excitatory neurons shown. Finally, we remark that not only are the three scenarios shown in Fig. 10 taken from the same undriven network (with parameters  $(\alpha, \gamma) = (1.5, 1)$ ), they in fact occurred within 3 seconds of each other, nor did we have to try hard to locate such scenarios.

## 5 Discussion and conclusions

The goals of this numerical study are to shed light on inhomogeneous spike patterns in local neuronal populations. We

**Table 4** Influence of  $(\tau_E, \tau_I)$

$(\tau_E, \tau_I)$	(4,1)	(3,2)	(2,3)	(2,7)
mean cluster size	2.3	9.2	23	27
single firing event percentage	57 %	30 %	16 %	14 %

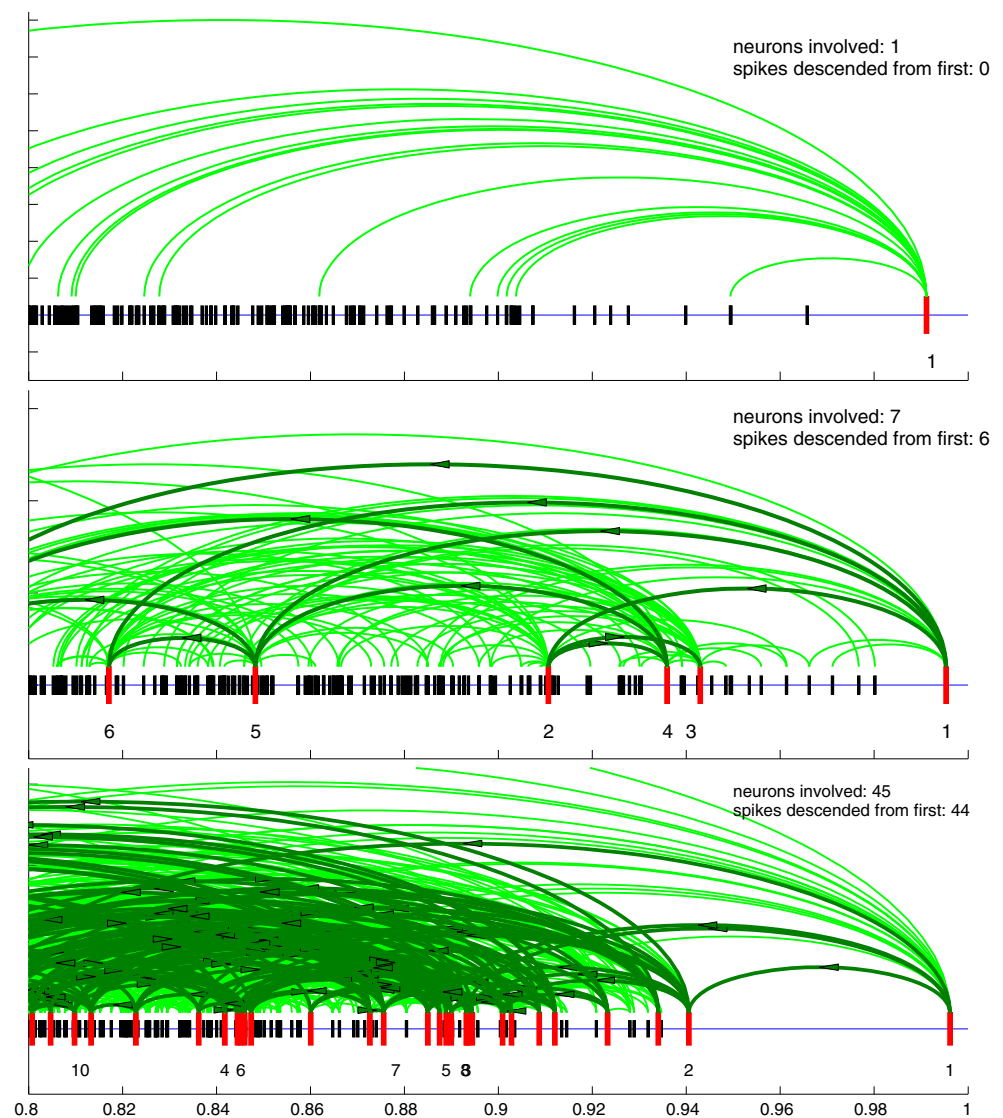
propose practical tools for quantifying a population’s tendency to spike in clusters. Then, using a simple network of integrate-and-fire neurons, we collect statistics on clusters, study them, and propose dynamical explanations whenever we can.

Clustering in this paper refers to temporally localized elevations in spike activity that occur repeatedly, separated by relative lulls in between. Synchronization or partial synchronization is a well recognized form of clustering, on which there is a considerable literature. Notable results include, but are not limited to, the following: First, there is Kuramoto’s model of coupled phase oscillators and his prediction of the dichotomy between phase-locked regimes and “incoherent” regimes (Kuramoto 1975). This is followed by rigorous results, such as Mirollo and Strogatz (1990), as well as extensions and variants of the Kuramoto model, such as Tsodyks et al. (1993) and Hansel et al. (1993). Focusing on various properties of (real) neurons, other researchers have constructed models that also show, as an emergent phenomenon, some form of synchronization. We mention in particular the use of synaptic depression to synchronize a mostly excitatory network by Tsodyks et al. (2000) and Mark and Tsodyks (2012), the production of rhythms via PING by Whittington et al. (2000) and later Kopell and Börges (2003, 2005), the use of synaptic failure in excitation-only populations by DeVill and Peskin to produce synchronous and asynchronous dynamics as well as regimes that can drift back and forth (Deville and Peskin 2008), and more recently, the use of spike frequency adaption by Kilpatrick and Ermentrout (2011) to separate neuronal populations into finite groups of neurons that spike together.

Though synchronized or partially synchronized spiking can be considered a special case of clustering, we have not focused on this topic as it has already received a fair amount of attention. We are primarily interested in structured dynamics in the form of temporally localized barages of spikes that occur at irregular times. These ideas were explored in the experimental work of Plenz et al, see e.g. Beggs and Plenz (2003), Hahn et al. (2010) and Plenz et al. (2011), and were discovered independently in the computational modeling of Rangan and Young (2013a). Plenz et al. emphasized criticality and power law distributions while Rangan and Young stressed the role of dynamical causality. The present work can be seen as a more systematic development of the ideas that originated in Rangan and Young (2013a, b).

There are many other experimental papers reporting on structured dynamics, see e.g. Samonds et al. (2005), Churchland et al. (2010), Leinekugel et al. (2002), Mazzoni et al. (2007) and Luczak et al. (2013), to cite only a few. On the modeling, theoretical and analytical level, papers that investigate various kinds of spike correlations

**Fig. 10** Three different scenarios. *Horizontal axis* shows the membrane potentials of all E-neurons whose voltages are in (0.8, 1) when the neuron marked “1” spikes. *Black markers* depict neurons which do not fire in the next 10ms, and red markers depict those that do, the numbers underneath indicating the order of spiking (*in the 3rd scenario, only the first 10 spikes are labeled*). A green arc is drawn between a spiking neuron and all postsynaptic neurons which received its action potential; the arc is colored dark green if the postsynaptic neuron eventually fires. Note that initial voltage distributions are roughly the same in all three cases, yet the outcomes are very different



include (Litwin-Kumar and Doiron 2012; Yu et al. 2011; Shew et al. 2011; Dehghani et al. 2012; Brunel 2000; Hansel and Sompolinsky 1996; Kriener et al. 2008; Battaglia and Hansel 2011; Trousdale et al. 2013; Staude et al. 2009), again citing only a sample, the literature being much larger.

Our first order of business in this paper was to try to identify a notion of clustering that is both precise and meaningful. A proposal for such a definition is made in Section 2.2. It involves two parameters: one gives the relevant resolution and the other a criterion for what constitutes a notable deviation from the mean. The first parameter permits an investigator to focus on events that occur on any temporal scale of his or her choosing, and the second permits one to give emphasis to small fluctuations or large deviations as the situation warrants. Such a definition is flexible enough that it has the potential

of being applicable to a wide range of circumstances, and we demonstrated that using a simple network of integrate-and-fire neurons.

A precisely formulated notion of clusters opened the door to their statistical analysis, which we carried out using our network models. We have found it productive to view population activity as a sequence of *events*, characterized by clustered spiking, and relative lulls between events called *inter-event times*. Our first observation is that undriven systems have extremely broad distributions of event sizes and inter-event times (consistent with previous works), while driving a system imposes a certain regularity on its inter-event durations, making them more narrowly distributed and fitting the description of broad-band gamma oscillations in real brains (Henrie and Shapley 2005). We have argued, based on comparison with experimental data, that

the dynamical mechanism presented here is more realistic than PING (or weak-PING) as a generator of gamma rhythms in cortex. Our second observation is that event sizes and inter-event times decorrelate very quickly, becoming essentially independent after one full cycle. These observations have led to a probabilistic description of collective neuronal activity in the form of two sequences of *i.i.d.* random variables as explained in Section 3.2.

With regard to dynamics, careful study of our network model has brought home the following two messages: The first is that regimes that exhibit a wide range of behaviors can sometimes be rationalized, up to a point, by a common dynamical mechanism (see Figs. 2, 3 and 8). The second message is the diversity in voltage and conductance configurations that can occur within a regime, leading to endless possibilities each time an excitatory neuron crosses threshold (Fig. 10). The first message gives optimism to theorists. The second highlights the following challenge: how to sort out which aspects of this infinitely complex dynamical landscape matter, in the sense of impacting neural computation downstream?

Finally, the work we have presented is based on a toy model of integrate-and-fire neurons. While there are hints here and there of compatibility with experimental data on cortical dynamics, we do not know the extent to which clustering occurs in the real brain, or what its biological significance may be – and that is precisely why we wish to bring our findings to the neuroscience community.

**Conflict of interest** The authors declare that they have no conflict of interest

## References

- Battaglia, D., & Hansel, D. (2011). Synchronous chaos and broad band gamma rhythm in a minimal multi-layer model of primary visual cortex. *PLoS Computational Biology*, 7(10), e1002176.
- Beggs, J.M., & Plenz, D. (2003). Neuronal avalanches in neocortical circuits. *Journal of Neuroscience*, 23, 11167–11177.
- Börgers, C., & Kopell, N. (2003). Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity. *Neural Computation*, 15(3), 509539.
- Börgers, C., & Kopell, N. (2005). Effects of noisy drive on rhythms in networks of excitatory and inhibitory neurons. *Neural Computation*, 17(3), 557608.
- Börgers, C., Epstein, S., Kopell, N. (2005). Background gamma rhythmicity and attention in cortical local circuits: A computational study. *Proceedings of the National Academy of the Sciences of the United States of America*, 102(19), 7002–7007.
- Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Comparative Neuroscience*, 8, 183–208.
- Churchland, M.M. et al. (2010). Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nature Neuroscience*, 13(3), 369–378.
- Dehghani, N., Hatsopoulos, N.G., Haga, Z.D., Parker, R.A., Greger, B., Halgren, E., Cash, S.S., Destexhe, A. (2012). Avalanche analysis from multi-electrode ensemble recordings in cat, monkey and human cerebral cortex during wakefulness and sleep. arXiv:1203.0738v4[q-bio.NC].
- Deville, L., & Peskin, C.S. (2008). Synchrony and asynchrony in a fully stochastic neural network. *Bulletin of Mathematical Biology*, 70(6), 1608–1633.
- Hahn, G., Petermann, T., Havenith, M.N., Yu, S., Singer, W., Plenz, D., Nikolic, D. (2010). Neuronal avalanches in spontaneous activity *in vivo*. *Journal of Neurophysiology*, 104, 3312–22.
- Hansel, D., & Sompolinsky, H. (1996). Chaos and synchrony in a model of a hypercolumn in visual cortex. *Journal of Comparative Neuroscience*, 3, 7–34.
- Hansel, D., Mato, G., Meunier, C. (1993). Clustering and slow switching in globally coupled phase oscillators. *Physical Review E*, 48(5).
- Henrie, W., & Shapley, R. (2005). LFP power spectra in V1 cortex: the graded effect of stimulus contrast. *Journal of Neurophysiology*, 94, 479–90.
- Kilpatrick, Z.P., & Ermentrout, B. (2011). Sparse gamma rhythms arising through clustering in adapting neuronal networks. *PLoS Computational Biology*, 7(11), e1002281. doi:10.1371/journal.pcbi.1002281.
- Koch, C. (1999). *Biophysics of computation*. London: Oxford University Press.
- Kriener, B., Tetzlaff, T., Aertsen, A., Diesmann, M., Rotter, S. (2008). Correlations and population dynamics in cortical networks. *Neural Computation*, 20, 2185–2226.
- Kuramoto, Y. (1975). In Arakai, H. (Ed.) *International symposium on mathematical problems in theoretical physics, Lecture notes in physics* (Vol. 39, p. 420). New York: Springer.
- Leinekugel, X., Khazipov, R., Cannon, R., Hirase, H., Ben-Ari, Y., Buzsaki, G. (2002). Correlated bursts of activity in the neonatal hippocampus *in vivo*. *Science*, 296(5575), 2049–2052.
- Litwin-Kumar, A., & Doiron, B. (2012). Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience*. doi:10.1038/nn.3220.
- Luczak, A., Bartho, P., Harris, K.D. (2013). Gating of sensory input by spontaneous cortical activity. *Journal of Neuroscience*, 33(4), 1684–1695. doi:10.1523/JNEUROSCI.2928-12.2013.
- Mark, S., & Tsodyks, M. (2012). Population spikes in cortical networks during different functional states. *Frontiers in Computational Neuroscience*, 6, 43. doi:10.3389/fncom.
- Mazzoni, A., et al. (2007). On the dynamics of the spontaneous activity in neuronal networks. *PLoS One*, 5, e539.
- Mirollo, R.E., & Strogatz, S.H. (1990). Synchronization of pulse-coupled biological oscillators. *SIAM Journal on Applied Mathematics*, 50(6), 1645–1662.
- Okun, M., Yger, P., Marguet, S.L., Gerard-Mercier, F., Benucci, A., Katzner, S., et al. (2012). Population rate dynamics and multi-neuron firing patterns in sensory cortex. *Journal of Neuroscience*, 32(48), 17108–17119. doi:10.1523/JNEUROSCI.1831-12.2012.
- Plenz, D. et al. (2011). Multi-electrode array recordings of neuronal avalanches in organotypic cultures. *Journal of Visualized Experiments*, 1:pii, 2949.
- Rangan, A.V., & Young, L.S. (2013a). Emergent dynamics in a model of visual cortex. *Journal of Computational Neuroscience*, 35(2), 155–167.
- Rangan, A.V., & Young, L.S. (2013b). Dynamics of spiking neurons: Between homogeneity and synchrony. *Journal of Computational Neuroscience*, 34(3), 433–460.
- Samonds, J.M., Zhou, Z., Bernard, M.R., Bonds, A.B. (2005). Synchronous activity in cat visual cortex encodes collinear and cocircular contours. *Journal of Neurophysiology*, 95(4), 2602–2616.

- Shew, W.L. et al. (2011). Information capacity and transmission are maximized in balanced cortical networks with neuronal avalanches. *Journal Neuroscience*, *31*, 55–63.
- Stauder, B., Rotter, S., Grün, S. (2009). CuBIC: cumulant based inference of higher-order correlations in massively parallel spike trains. *Journal of Computational Neuroscience*. doi:[10.1007/s10827-009-0195-x](https://doi.org/10.1007/s10827-009-0195-x).
- Trousdale, J., Hu, Y., Shea-Brown, E., Josić, K. (2013). Newblock a generative spike train model with time-structured higher order correlations. *Frontiers in Computational Neuroscience*, *7*.
- Tsodyks, M., Mitkov, I., Sompolinsky, H. (1993). Pattern of synchrony in inhomogeneous networks of oscillators with pulse interactions. *Physical Review Letters*, *71*(8), 1280–1283.
- Tsodyks, M., Uziel, A., Markram, H. (2000). Synchrony generation in recurrent networks with frequency- dependent synapses. *Journal of Neuroscience*, *20*, RC50.
- Whittington, M.A., Traub, R.D., Kopell, N., Ermentrout, B., Buhl, E.H. (2000). Inhibition-based rhythms: experimental and mathematical observations on network dynamics. *International Journal of Psychophysiology*, *38*, 315–336.
- Yu, S. et al. (2011). Higher-order interactions characterized in cortical activity. *Journal Neuroscience*, *31*, 17514–17526.