

Optimal Detection of a Localized Perturbation in Random Networks of Integrate-and-Fire Neurons

Davide Bernardi^{1,2} and Benjamin Lindner^{1,2}

¹*Bernstein Center for Computational Neuroscience Berlin, Philippstraße 13, Haus 2, 10115 Berlin, Germany*

²*Physics Department of Humboldt University Berlin, Newtonstraße 15, 12489 Berlin, Germany*

(Received 19 December 2016; published 29 June 2017)

Experimental and theoretical studies suggest that cortical networks are chaotic and coding relies on averages over large populations. However, there is evidence that rats can respond to the short stimulation of a single cortical cell, a theoretically unexplained fact. We study effects of single-cell stimulation on a large recurrent network of integrate-and-fire neurons and propose a simple way to detect the perturbation. Detection rates obtained from simulations and analytical estimates are similar to experimental response rates if the readout is slightly biased towards specific neurons. Near-optimal detection is attained for a broad range of intermediate values of the mean coupling between neurons.

DOI: 10.1103/PhysRevLett.118.268301

How can a local perturbation in a huge network become macroscopically noticeable, as reported in recent experiments [1–5]? One particularly striking example is that awake rats can be trained to report the transient stimulation of a single cortical cell [2]. This finding that one neuron out of millions has a weak but significant behavioral effect seems to be at odds with experimental [6] and theoretical [7,8] studies suggesting that cortical dynamics are chaotic and therefore only population averages matter. A theoretical explanation of this hypersensitivity is missing and many factors make it a challenge: The system is overwhelmingly complex, it is unclear how the learning process modifies the system during the training of the experimental subjects, and it is unknown how the activity of the neural network is linked to the behavioral response.

A widely used class of models in the theoretical investigation of neural dynamics are networks of leaky integrate-and-fire (LIF) neurons [9–15]. Such models represent an enormous simplification of actual cortical networks [16]. Nevertheless, even an unstructured LIF network can successfully reproduce features such as the asynchronous irregular firing pattern [17] or long-range temporal correlations [18–20]. Can such a simple network model explain the detectability of the single-cell stimulation as well?

In the present study, we apply a perturbation to one randomly selected neuron from a large random network of excitatory and inhibitory LIF neurons and consider the activity of a subset of the network as a basis for the perturbation readout. We show that for a random selection of readout neurons the detectability is smaller than what was experimentally observed (reflecting the fact that untrained rats cannot detect the single-cell stimulation). However, biasing the readout towards neurons receiving direct input from the stimulated cell (a proxy for the training) increases the detection performance to values similar to the experimentally measured ones. This finding is

valid over a broad range of synaptic coupling amplitudes. For strong coupling, fluctuations of the spontaneous activity impair the detection of the perturbation, whereas the detectability stays high even for weak coupling. In the presence of external input noise (mimicking the fact that real neural nets are never isolated), the detectability drops for unrealistically small values of the coupling strength, thus leading to an optimal range for the synaptic connection strength.

Model.—We consider a random network of $N_E = 8 \times 10^4$ excitatory and $N_I = 2 \times 10^4$ inhibitory LIF neurons [21]. The membrane voltage of the k th neuron evolves according to

$$\tau_m \dot{v}_k = -v_k + R[I_{\text{ext}} + I_k(t)], \quad (1)$$

where $\tau_m = 20$ ms is the membrane time constant and the constant external input is $RI_{\text{ext}} = 22$ mV. Whenever the voltage (measured with respect to the resting potential) $v_k(t)$ reaches $v_T = 20$ mV, the neuron fires and $v_k(t)$ is reset to $v_R = 10$ mV after a refractory period $\tau_{\text{ref}} = 2$ ms. Delta functions centered on the time of each threshold crossing, $t_{k,i}$, define the output spike train $x_k(t) = \sum_i \delta(t - t_{k,i})$ of the k th neuron. Neurons are coupled by current-based delayed instantaneous synapses [22]

$$I_k(t) = \frac{\tau_m}{R} \left(\sum_{j \in \mathcal{P}_e(k)} J_{kj} x_j(t - D_{kj}) - g \sum_{\ell \in \mathcal{P}_i(k)} J_{k\ell} x_\ell(t - D_{k\ell}) \right). \quad (2)$$

The connectivity is sparse and with fixed in degree as in Ref. [9], but with a larger number of input connections per neuron $C = 5000$ and with couplings J_{kj} drawn from an exponential distribution with mean J , an approximate description of the long-tailed histograms of synaptic

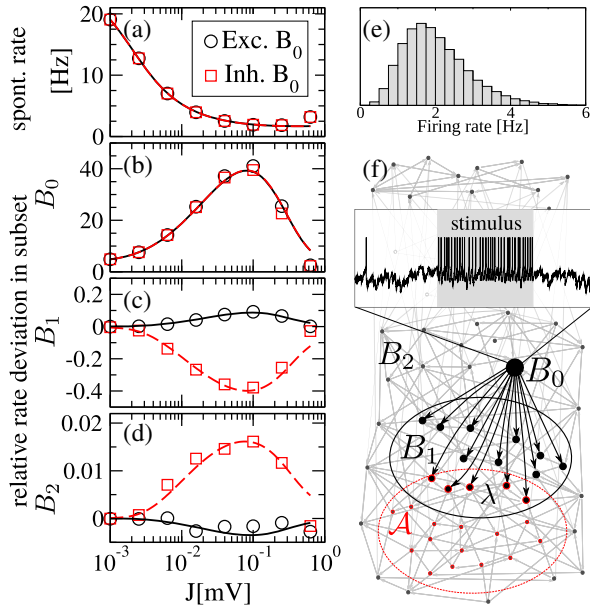


FIG. 1. (a) Spontaneous network firing rate r_{sp} as a function of the average coupling strength. (b),(c),(d) Relative deviation of the firing rate of B_0 , B_1 , and B_2 during the stimulus, respectively. Lines indicate MFT predictions and data points are simulation results. (e) Firing rate distribution for $J = 0.1$ mV. (f) Network model. B_0 : perturbation site; B_1 : neurons receiving direct input from B_0 ; B_2 : all other neurons. \mathcal{A} : readout population. λ : overlap between \mathcal{A} and B_1 ($0 \leq \lambda \leq 1$). Note that there is no structure in the network, subsets are grouped for illustration purposes.

amplitudes measured in cortex [23]. We set the strength of inhibition relative to excitation $g = 7$ to make the net recurrent input inhibitory and achieve low spontaneous firing rates. Figure 1(a) shows the spontaneous firing rate r_{sp} as a function of the mean coupling J . Note that the random weights lead to distributed firing rates despite the fixed in degree [Fig. 1(e)]. Data points are simulation results and lines are predictions of the mean-field theory (MFT) based on the analytical results for a LIF neuron driven by white shot-noise input developed in Ref. [24]: Requesting self-consistency for input and output firing rates leads to mean field equations that can be solved numerically. For $J \leq 0.1$ mV, the spontaneous rate r_{sp} decreases because the recurrent inhibitory input increases with J . Somewhere above $J = 0.1$ mV, fluctuations in the network become stronger and slower [19,20], so that the assumption of temporally uncorrelated input is less and less accurate, causing discrepancies between theory and simulations.

Perturbation and its effect on firing rates.—On each trial, the network is initialized randomly and simulated for a time window of 3 s centered on $t = 0$ [25]. A randomly selected neuron is the perturbation site and labeled as B_0 . For $0 < t < 400$ ms the constant input to B_0 , RI_{ext} , is increased by 23 mV to bring its firing rate from the spontaneous value r_{sp} to a new value r_0 and mimic the single-cell stimulation. Figure 1(b) shows the relative increase of the firing rate during the stimulation, that is $(r_0 - r_{sp})/r_{sp}$.

To study the effects of the increased firing rate of B_0 , one needs to distinguish between two subsets of the network: neurons receiving direct input from B_0 , labeled as B_1 , and all other neurons (that do not receive direct input from B_0), labeled as B_2 [Fig. 1(f)]. Intuitively, the firing rate of B_1 neurons, r_1 , is increased (decreased) when B_0 is an excitatory (inhibitory) neuron: Figure 1(c) shows the relative deviation from the spontaneous firing rate $(r_1 - r_{sp})/r_{sp}$ as a function of J for the case of excitatory (inhibitory) B_0 with a continuous (dashed) line for theory and circles (squares) for simulation results. As B_2 does not receive direct input from B_0 , the effect of the stimulus on r_2 , the firing rate of B_2 , is weaker [Fig. 1(d)].

Detection.—Neurons forming a second circuit reading out the activity of the network can receive input only from a subset of the stimulated network. Therefore, a natural choice is to consider a subpopulation \mathcal{A} of size C as input for a neural “perturbation detector.” The readout population \mathcal{A} can be constructed by choosing at random λC neurons in B_1 and $(1 - \lambda)C$ neurons in B_2 , so that $0 \leq \lambda \leq 1$ parametrizes the overlap between \mathcal{A} and B_1 ; to be conservative, B_0 is excluded from \mathcal{A} . Each realization of \mathcal{A} can be thought of as one sample from an ensemble of populations serving as input to many readout neurons. These neurons would integrate their input over some time scale, which we set at $\tau_f = 100$ ms (comparable to the time scale of NMDA synapses [26]), and react to variations of their input, the readout population activity $A(t) = C^{-1} \sum_{x_i \in \mathcal{A}} x_i \star \mathcal{F}_{\tau_f}(t)$, where \star indicates convolution and \mathcal{F}_{τ_f} is a causal filter standing for an abstract information integration [27]. Two trajectories of $A(t)$ for different values of λ are plotted together with the respective trial average $\langle A(t) \rangle$ in Fig. 2. Here, $\langle A(t) \rangle$ reaches a plateau $r_\lambda = r_1 \lambda + r_2 (1 - \lambda)$ after approximately $3\tau_f$.

We consider the deviation of $A(t)$ from its time average before stimulation $\Delta A(t) = |A(t) - \int_{-T_w}^0 dt \langle A(t) \rangle|$, where $T_w = 1300$ ms is the time window for the perturbation detection. An event is detected when $\Delta A(t)$ crosses a threshold θ_d ; i.e., $A(t)$ leaves the interval $[r_{sp} - \theta_d, r_{sp} + \theta_d]$ (see Fig. 2). If $\Delta A(t)$ exceeds θ_d at least once before the




FIG. 2. Schematics of perturbation readout. Two example trajectories of $A(t)$ for each λ are plotted together with the trial average $\langle A(t) \rangle$ (dashed lines). Parameters: inhibitory B_0 , $J = 0.1$ mV.

stimulus onset (for $-T_w < t < 0$), this is counted as a *false positive* event. The fraction of trials in which this happens defines the false positive rate, $\mathcal{FP}(\theta_d)$, as a function of the detection threshold. The correct detection rate $\mathcal{CD}(\theta_d)$ is defined analogously but for a threshold crossing in $0 < t < T_w$. We note that fluctuations of $A(t)$ that limit the detectability are mainly due to weak cross-correlations in the network.

Estimation of detection rates.—Because $A(t)$ is the filtered sum of many weakly correlated spike trains, its stationary distribution is nearly Gaussian (not shown). The time-dependent mean of $A(t)$ can be well approximated by exploiting the MFT [28]. However, even under the assumption that its autocorrelation function is unchanged by the stimulus, exact detection rates are related to the first-passage time distribution for a Gaussian process with time-dependent mean and a generic autocorrelation function, which has no known analytical solution. Instead, we use a simple approximation: Given the variance of $A(t)$ from the measured autocorrelation function, the probability $p_{\Delta A}(\theta_d, t)$ that $\Delta A(t)$ is below θ_d at any given point in time is straightforward to obtain [29]. We then divide the time axis in T_w/τ_c intervals, where τ_c is a suitably defined correlation time τ_c [30], and approximate the probability of $\Delta A(t)$ to be below θ_d in the whole T_w as the product of the average probability of being below the threshold once in each interval:

$$\mathcal{CD}(\theta_d) \approx 1 - \prod_{k=0}^{T_w/\tau_c} \frac{1}{\tau_c} \int_{k\tau_c}^{(k+1)\tau_c} dt p_{\Delta A}(\theta_d, t), \quad (3)$$

and analogously for $\mathcal{FP}(\theta_d)$, but for negative times.

Plotting $\mathcal{FP}(\theta_d)$ on the x axis and $\mathcal{CD}(\theta_d)$ on the y axis upon variation of θ_d defines the receiver operating characteristic (ROC) curve, a standard way of quantifying the performance of a detector. ROC curves for excitatory B_0 , $J = 0.1$ mV and four values of λ are shown in Fig. 3: continuous lines are analytical estimates, while symbols are simulation results. For each value of λ four realizations of \mathcal{A} are shown. For $\lambda \approx 1$ the ROC curve is well separated from




FIG. 3. ROC curves for excitatory B_0 (maximum distance from the diagonal defines the effect size \bar{Y}).

the diagonal (chance level), and approaches the diagonal on decreasing λ .

Optimal detection.—Following Ref. [2], we define the effect size as the difference between correct detections and false positives. To compare our results to the experiments, we choose the maximum distance of the ROC curve from the diagonal $\bar{Y} = \max_{\theta_d} \{\mathcal{CD}(\theta_d) - \mathcal{FP}(\theta_d)\}$. Selecting the optimal threshold can be thought of as one effect of the training phase that is part of the experimental procedure. One practical issue with this choice is that any finite-size fluctuation of the ROC curve produces a systematic upward bias on \bar{Y} . Increasing the number of trials decreases the bias. Still, a way to quantify the reliability of our measure is desirable. To this end, we compute the p value from Fisher's exact test (with significance level $p < 0.05$, as in Ref. [2]) for each data point. We investigate the effect of the recurrent coupling parameter, by plotting \bar{Y} as a function of J . In Fig. 4(a) the case of excitatory B_0 is depicted for many values of λ . Solid lines are theoretical estimates while data points are simulation results. Closed symbols indicate significant data points. It can be seen that almost all curves display a very broad maximum for $10^{-2} \leq J \leq 10^{-1}$. For larger values of J , i.e., in the range where strong slow fluctuations appear [19,20], the detectability drops. For $J < 10^{-2}$ the effect size decreases less rapidly upon decreasing coupling strength. Nonsignificant data points are those for which large deviations from the theoretical predictions are observed, with the exception of the two smallest values of J : for these data points deviations of firing rates from the spontaneous value are so small that the relative error of the MFT prediction is large, although the absolute error on each firing rate is small. For small values of λ the theoretically predicted effect size is small

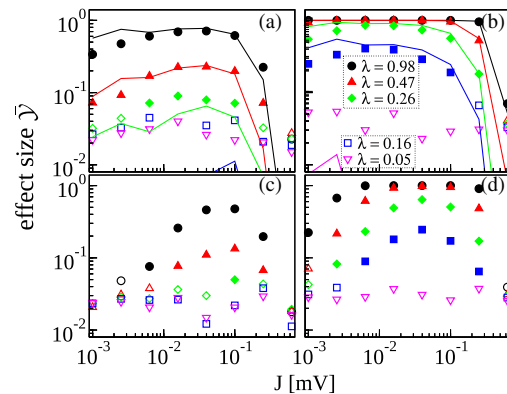


FIG. 4. Detection statistics as a function of the average coupling strength for different overlaps λ . Closed (open) symbols indicate significant (nonsignificant) data points (significance level $p < 0.05$). Each point is the average over 24 different realizations of \mathcal{A} . Lines indicate theoretical estimates. Results for excitatory (a,c) and inhibitory (b,d) B_0 in the absence (a,b) or presence (c,d) of external shot noise. Number of trials: 800 (a), 400 (b), 900 (c), 600 (d).

and not significant. The minimum λ required to obtain a $\bar{\mathcal{Y}}$ similar to the experimental value of ≈ 0.015 is $\lambda \approx 0.25$, while a completely random selection of readout neurons corresponds to $\lambda_0 = C/N = 0.05$, which is not significant for all values of J . In the case of inhibitory B_0 [plotted in Fig. 4(b)] the effect size is generally much stronger, which is in line with experiments. In our model this is caused by inhibitory synapses having a higher weight than excitatory ones. For $\lambda > 0.4$, there is almost perfect detection for small to intermediate values of J , and $\bar{\mathcal{Y}}$ is larger than in experiments (≈ 0.07) even for $\lambda = 0.15$. Also, in this case of inhibitory B_0 , for $\lambda = 0.05$ there is no significant data point.

The effect size $\bar{\mathcal{Y}}$ stays rather high even for extremely small values of J because the fluctuations of $A(t)$ become extremely small, which makes the detection of a very weak effect in the firing rates possible. However, cortical circuits are not isolated systems because they receive both bottom-up and top-down external input. If such a feature is mimicked in the model by adding external excitatory shot noise input to each neuron [31], the detectability drops especially for small values of J : in Figs. 4(c)/4(d) we plot simulation results for excitatory/inhibitory B_0 in the presence of external random input (θ_d was optimized for each data point in Fig. 4 separately). For strong coupling, the reduction in $\bar{\mathcal{Y}}$ is only moderate because the detection is limited by the autonomous fluctuations of the network.

Scaling and thermodynamic limit.—Although cortical networks are indeed finite, a theoretically intriguing question is whether detection is possible in an arbitrarily large network (briefly discussed here, more details in Ref. [33]).

If the network size $N = N_E + N_I$ is increased with constant C , the network becomes sparser, leading to decreased cross-correlations between neurons. Because cross-correlations are the main source of fluctuations in $A(t)$, the effect size $\bar{\mathcal{Y}}$ increases for a given λ . However, λ_0 , the overlap expected from an unbiased selection, becomes smaller with increasing N , making the necessary bias on the readout prohibitively strong.

In contrast, if C is scaled along with N to keep the ratio C/N constant, J needs to be scaled down to prevent trivial limiting behaviors of the network dynamics. In the range of C that can be numerically explored, we find that for *weak* synapses (i.e., $J, J_{\text{ext}} \sim C^{-1}$) [37] the effect size $\bar{\mathcal{Y}}$ increases moderately for increasing N , whereas for *strong* synapses (i.e., $J, J_{\text{ext}} \sim C^{-1/2}$) [7,37] $\bar{\mathcal{Y}}$ decreases slightly. These observations can be related to how cross-correlations and the neuron's linear response to slow signals depend on C [33]. For network sizes beyond our computational possibilities ($NC > 10^{10}$), theoretical arguments suggest that the effect size eventually goes to zero for $N \rightarrow \infty$ for both scalings [33].

Conclusions.—How the transient firing of a single neuron in a large network can be detected is a novel and still unsolved problem. In this study, we showed how the

activity of a single cell can be detected in a simple macroscopic statistic: the instantaneous firing rate of a readout subset \mathcal{A} , a proxy for the response of an ensemble of readout neurons and for the rat's response. Our results show that choosing \mathcal{A} at random is not sufficient; i.e., the readout needs to be trained, as experiments suggest. Note that the experimental training [2] does not target single cells but an area. Therefore, the formation of hubs, essential to the phenomena in Refs. [3,4], seems less likely to be relevant for the considered problem.

For a large range of coupling values J , a biased selection of \mathcal{A} permits the detection; i.e., there is a minimum overlap λ with B_1 for which the effect is significant and of a size similar to experiments. If readout sets \mathcal{A} are chosen at random, the probability of finding a readout set \mathcal{A} with λ above the needed value is low. However, it is conceivable that in neural systems synaptic input populations are not completely random but spatially organized and that during the training phase most readout sets with the necessary λ gain weight, thus enabling the perturbation detection.

We found that strong coupling is detrimental for detecting the perturbation. In the other limit of extremely weak coupling, the detectability stays high even for unrealistic values. In the presence of external noise, the detectability shows a more definite but still broad maximum located somewhat below mean synaptic amplitudes measured in the cortex. The optimum may shift if conductance-based synapses or heterogeneities are included in the model [38].

Concerning the robustness of our findings, replacing the fixed in-degree connectivity with an Erdős-Rényi graph or increasing the network size up to $N = 4 \times 10^5$ does not change our results qualitatively [33]. Although a theoretical analysis reveals that the perturbation detection must be regarded, strictly speaking, as a finite-size effect, we remark that the perturbation stays detectable for any network size in a biologically plausible range.

This study was conducted in the spirit of assessing how short a null hypothesis would fall. Although many aspects need further investigation, it is surprising that a comparatively simple model, which allows a partial analytical treatment, can reproduce some features of the experimental results, thus posing a first promising step in further exploration of this challenging theoretical problem.

The authors would like to thank R. Scheid, G. Vallone, and S. Voronenko for useful comments on the manuscript and BMBF (FKZ:01GQ1001A) and DFG (GRK 1589/2) for financial support.

-
- [1] M. Brecht, M. Schneider, B. Sakmann, and T. W. Margrie, *Nature (London)* **427**, 704 (2004).
 - [2] A. R. Houweling and M. Brecht, *Nature (London)* **451**, 65 (2008).
 - [3] C.-y. T. Li, M.-m. Poo, and Y. Dan, *Science* **324**, 643 (2009).

- [4] P. Bonifazi, M. Goldin, M. A. Picardo, I. Jorquera, A. Cattani, G. Bianconi, A. Represa, Y. Ben-Ari, and R. Cossart, *Science* **326**, 1419 (2009).
- [5] G. Doron, M. von Heimendahl, P. Schlattmann, A. R. Houweling, and M. Brecht, *Neuron* **81**, 653 (2014).
- [6] M. London, A. Roth, L. Beeren, M. Häusser, and P. E. Latham, *Nature (London)* **466**, 123 (2010).
- [7] C. van Vreeswijk and H. Sompolinsky, *Science* **274**, 1724 (1996).
- [8] M. Monteforte and F. Wolf, *Phys. Rev. Lett.* **105**, 268104 (2010).
- [9] N. Brunel, *J. Comput. Neurosci.* **8**, 183 (2000).
- [10] D. Golomb and G. B. Ermentrout, *Phys. Rev. Lett.* **86**, 4179 (2001).
- [11] A. Roxin, H. Riecke, and S. A. Solla, *Phys. Rev. Lett.* **92**, 198101 (2004).
- [12] T. Vogels and L. Abbott, *Nat. Neurosci.* **12**, 483 (2009).
- [13] G. Mongillo, D. Hansel, and C. van Vreeswijk, *Phys. Rev. Lett.* **108**, 158101 (2012).
- [14] J. F. Mejias and A. Longtin, *Phys. Rev. Lett.* **108**, 228102 (2012).
- [15] S. Luccioli, E. Ben-Jacob, A. Barzilai, P. Bonifazi, and A. Torcini, *PLoS Comput. Biol.* **10**, e1003823 (2014).
- [16] D. Feldmeyer, M. Brecht, F. Helmchen, C. C. H. Petersen, J. F. A. Poulet, J. F. Staiger, H. J. Luhmann, and C. Schwarz, *Prog. Neurobiol.* **103**, 3 (2013).
- [17] A. Renart, J. de la Rocha, P. Bartho, L. Hollender, N. Parga, A. Reyes, and K. D. Harris, *Science* **327**, 587 (2010).
- [18] B. Kriener, H. Enger, T. Tetzlaff, H. E. Plesser, M.-O. Gewaltig, and G. T. Einevoll, *Front. Comput. Neurosci.* **8**, 136 (2014).
- [19] S. Ostojic, *Nat. Neurosci.* **17**, 594 (2014).
- [20] S. Wieland, D. Bernardi, T. Schwalger, and B. Lindner, *Phys. Rev. E* **92**, 040901 (2015).
- [21] W. Gerstner and W. M. Kistler, *Spiking Neuron Models: Single Neurons, Populations, Plasticity* (Cambridge University Press, Cambridge, England, 2002).
- [22] Excitatory input spike trains $x_j(t)$ are generated by neurons belonging to $\mathcal{P}_e(k)$, a random selection of $C_E = 4000$ out of N_E excitatory neurons of the network. Analogously, $\mathcal{P}_i(k)$ are $C_I = 1000$ randomly selected inhibitory neurons. The connection probability is, therefore, 5% for any pair of neurons. Self-connections are excluded. Transmission delays D_{kj} are uniformly distributed between 0.5 and 2 ms.
- [23] S. Lefort, C. Tómm, J.-C. F. Sarria, and C. C. H. Petersen, *Neuron* **61**, 301 (2009).
- [24] M. J. E. Richardson and R. Swarbrick, *Phys. Rev. Lett.* **105**, 178102 (2010).
- [25] The network connectivity, random weights, and delays are drawn once and unchanged across trials (playing the role of frozen disorder). To forget random initial conditions, simulations were run for $400 < t_i < 1200$ ms (depending on the coupling strength) before starting data acquisition. Neurons were integrated with a Euler method and time step $\Delta t = 0.1$ ms.
- [26] A. C. Flint, U. S. Maisch, J. H. Weishaupt, A. R. Kriegstein, and H. Monyer, *J. Neurosci.* **17**, 2469 (1997).
- [27] We use a shifted and truncated Gaussian $\mathcal{F}_{\tau_f}(t) = (\pi\tau_f^2/2)^{-1/2} e^{-(t-3\tau_f/2)^2/(\tau_f^2/2)} \theta(t)\theta(3\tau_f - t)$. Changing the time constant in the range 50–200 ms does not affect our results qualitatively (not shown).
- [28] $\langle \Delta A(t) \rangle = \Delta r_\lambda [\Delta A_+(t) - \Delta A_+(t - t_1)]$, where $\Delta A_+(t) = 1/2 \{ \text{erfc}[-(-2t + 3\tau_f)/\sqrt{2}\tau_f] - \exp[(\tau_f^2 - 8t\tau_m + 12\tau_f\tau_m)/8\tau_m^2] \text{erfc}[(\tau_f^2 - 4t\tau_m + 6\tau_f\tau_m)/(2\sqrt{2}\tau_f\tau_m)] \}$.
- [29] $p_{\Delta A}(\theta_d, t) = \frac{1}{2} [\text{erf}(\theta_d - \langle \Delta A(t) \rangle / \sqrt{2}\sigma_A) + \text{erf}(\theta_d + \langle \Delta A(t) \rangle / \sqrt{2}\sigma_A)]$.
- [30] $\tau_c = \int_0^\infty dt |C_{AA}(t)| / C_{AA}(0)$. $C_{AA}(t)$ is the autocorrelation function measured from $A(t)$ for $t < 0$.
- [31] Each receives $C_{\text{ext}} = 700$ Poisson inputs with firing rate $r_{\text{ext}} = 12$ Hz and exponentially distributed amplitudes with mean $J_{\text{ext}} = 0.1$ mV [32]. The input current is adjusted so that the total mean input to each neuron is unchanged $RI_{\text{ext}} + \tau_m J_{\text{ext}} C_{\text{ext}} r_{\text{ext}} = 22$ mV.
- [32] C. E. Schoonover, J.-C. Tapia, V. C. Schilling, V. Wimmer, R. Blazeski, W. Zhang, C. A. Mason, and R. M. Bruno, *J. Neurosci.* **34**, 6746 (2014).
- [33] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.118.268301> for plots and additional information, which includes Refs. [34–36].
- [34] C. van Vreeswijk and H. Sompolinsky, *Neural Comput.* **10**, 1321 (1998).
- [35] S. Ostojic, N. Brunel, and V. Hakim, *J. Neurosci.* **29**, 10234 (2009).
- [36] R. D. Vilela and B. Lindner, *Phys. Rev. E* **80**, 031909 (2009).
- [37] C. Pehlevan and H. Sompolinsky, *PLoS One* **9**, e89992 (2014).
- [38] A. Litwin-Kumar and B. Doiron, *Nat. Neurosci.* **15**, 1498 (2012).