

RESEARCH ARTICLE

Encoding in Balanced Networks: Revisiting Spike Patterns and Chaos in Stimulus-Driven Systems

Guillaume Lajoie^{1,2*}, Kevin K. Lin³, Jean-Philippe Thivierge⁴, Eric Shea-Brown^{1,5,6}

1 University of Washington Institute for Neuroengineering, University of Washington, Seattle, Washington, United States of America, **2** Department of Nonlinear Dynamics, Max Planck Institute for Dynamics and Self-Organization, Göttingen, Germany, **3** School of Mathematics, University of Arizona, Tucson, Arizona, United States of America, **4** School of Psychology, University of Ottawa, Ottawa, Ontario, Canada, **5** Department of Applied Mathematics, University of Washington, Seattle, Washington, United States of America, **6** Department of Physiology and Biophysics, University of Washington, Seattle, Washington, United States of America

* glajoie@uw.edu



OPEN ACCESS

Citation: Lajoie G, Lin KK, Thivierge J-P, Shea-Brown E (2016) Encoding in Balanced Networks: Revisiting Spike Patterns and Chaos in Stimulus-Driven Systems. *PLoS Comput Biol* 12(12): e1005258. doi:10.1371/journal.pcbi.1005258

Editor: Bard Ermentrout, University of Pittsburgh, UNITED STATES

Received: April 25, 2016

Accepted: November 21, 2016

Published: December 14, 2016

Copyright: © 2016 Lajoie et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the paper.

Funding: GL was supported by a Bernstein Fellowship from the Bernstein Center for Computational Neuroscience, by a postdoctoral fellowship from the Fonds de Recherche du Quebec, and by an Innovation Fellowship from the Washington Research Foundation. KKL was supported in part by NSF grant DMS-1418775. JPT is supported by a Discovery grant from the Natural Sciences and Engineering Council of Canada (NSERC Grant No. 210977 and No. 210989), operating funds from the Canadian Institutes of Health Research.

Abstract

Highly connected recurrent neural networks often produce chaotic dynamics, meaning their precise activity is sensitive to small perturbations. What are the consequences of chaos for how such networks encode streams of temporal stimuli? On the one hand, chaos is a strong source of randomness, suggesting that small changes in stimuli will be obscured by intrinsically generated variability. On the other hand, recent work shows that the type of chaos that occurs in spiking networks can have a surprisingly low-dimensional structure, suggesting that there may be room for fine stimulus features to be precisely resolved. Here we show that strongly chaotic networks produce patterned spikes that reliably encode time-dependent stimuli: using a decoder sensitive to spike times on timescales of 10's of ms, one can easily distinguish responses to very similar inputs. Moreover, recurrence serves to distribute signals throughout chaotic networks so that small groups of cells can encode substantial information about signals arriving elsewhere. A conclusion is that the presence of strong chaos in recurrent networks need not exclude precise encoding of temporal stimuli via spike patterns.

Author Summary

Recurrently connected populations of excitatory and inhibitory neurons found in cortex are known to produce rich and irregular spiking activity, with complex trial-to-trial variability in response to input stimuli. Many theoretical studies found this firing regime to be associated with chaos, where tiny perturbations explode to impact subsequent neural activity. As a result, the precise spiking patterns produced by such networks would be expected to be too fragile to carry any valuable information about stimuli, since inevitable sources of noise such as synaptic failure or ion channel fluctuations would be amplified by chaotic dynamics on repeated trials. In this article we revisit the implications of chaos in

Health Research (CIHR Grant No. 6105509), and the University of Ottawa Brain and Mind Institute (uOBMI). ESB and GL by NSF CAREER Grant DMS - 1056125 and NIH Training grant 5T90DA032436. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

input-driven networks and directly measure its impact on evoked population spike patterns. We find that chaotic network dynamics can, in fact, produce highly patterned spiking activity which can be used by a simple decoder to perform input-classification tasks. This can be explained by the presence of low-dimensional, input-specific chaotic attractors, leading to a form of trial-to-trial variability that is intermittent, rather than uniformly random. We propose that chaos is a manageable by-product of recurrent connectivity, which serves to efficiently distribute information about stimuli throughout a network.

Introduction

Highly recurrent connectivity occurs throughout the brain. It is believed that recurrent cortical circuits typically operate in a “balanced state” in which every neuron receives a large number of excitatory (E) and inhibitory (I) inputs (see, e.g., [1–3]). This means synaptic currents nearly cancel on average, but feature strong fluctuations, giving rise to sustained irregular spiking [4]. Well-established results show that such strongly recurrent networks operating in a balanced regime can produce *chaotic* dynamics in a range of settings, from abstract firing rate models with random connectivity [5] to networks of spiking units with excitatory and inhibitory cell classes [2, 6–8]. Chaos implies that the network dynamics depend very sensitively on network states, so that tiny perturbations to initial conditions may lead to large effects over time. As a consequence, when the same stimulus is presented to a chaotic network multiple times, it may fail to generate reproducible responses. How can stimuli be encoded in the variable spike trains that result (c.f. [9, 10])? A central issue is the relationship between trial-to-trial variability and input discriminability: since exactly the same sensory input can elicit different neural responses from one trial to the next (as can distinct inputs), how can one decide which stimulus is driving a network based on its response?

To illustrate the variable stimulus responses due to chaos, Fig 1 shows a model balanced network (to be described in detail below) driven by a fixed multi-dimensional stimulus, as well as a raster plot of the spiking output of the network on two trials. In each of these trials, the exact same stimulus is presented, but on the second trial a small perturbation is artificially introduced in the form of an extra spike for neuron 1 (as in [11–13]). This small perturbation quickly reverberates throughout the whole network in a seemingly random fashion [11–13]. Can a network this sensitive to small differences in its internal state produce sufficiently reliable outputs for discrimination? If we were to present two different stimuli to such a network,

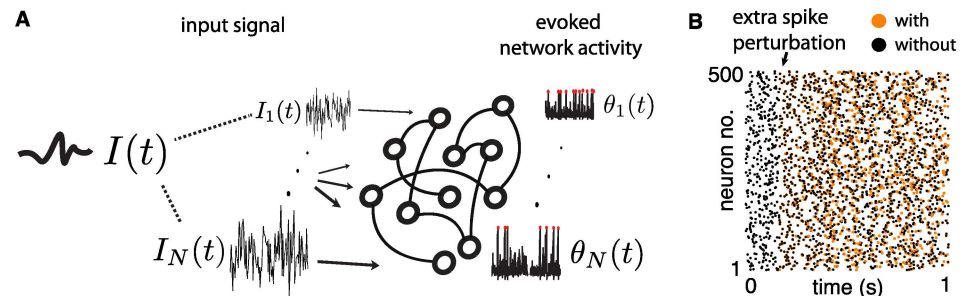


Fig 1. (A) An input $I(t)$, with N independent components, is presented to a chaotic network of N spiking cells. (B) Raster plot of the network response to a fixed input $I(t)$, on two trials. Arrow indicates a perturbation to the network on the second trial, in which an extra spike is “added” to neuron number 1. Different color markers indicate the widely divergent spike rasters that occur with and without the perturbation.

doi:10.1371/journal.pcbi.1005258.g001

could a decoder be trained to discriminate the spikes evoked by the first input from the spikes evoked by the second, despite both sets of output spikes being subject to the type of variability shown here? This is the central question that we investigate here.

We study this stimulus coding question using a strongly chaotic, recurrent spiking network model driven by temporal stimuli. The strength of stimulus inputs is comparable to network interactions, so that dynamics are not dominated by external stimuli alone. We find that, despite chaos, the network's spike patterns encode temporal features of stimuli with sufficient precision so that the responses to close-by stimuli can be accurately discriminated. We relate this coding precision to previous work grounded in the mathematical theory of dynamical systems, which shows that—at the level of multi-neuron spike patterns—chaotic networks do not produce as much variability as one might guess at first glance [14, 15]. This is because in such networks, the trial-to-trial variability of spike trains evoked by time-dependent stimuli leads to the formation of low-dimensional chaotic attractors. Our main findings are:

1. It is possible for strongly chaotic recurrent networks to produce multi-cell spike responses that remain discriminable even for highly similar temporal inputs.
2. The same recurrent connections that produce chaos distribute stimulus information throughout the network, so that stimuli can be discriminated based on only a small subset of “readout” neurons.
3. Input statistics influence the strength of chaotic fluctuations that can obscure stimuli. We quantify this via a chaos-induced “noise floor”; stimuli whose strength exceeds this floor will be easily discriminable.

We show that these findings are consistent with well-understood properties of irregular activity produced in balanced networks [2, 5, 16]. We also demonstrate that classical mean-field theories describe the overall firing statistics in our networks, but that additional tools are needed to describe their response to temporal stimuli.

Our results are based on numerical simulations guided by mathematical theory, but connect to a broad experimental literature: trial-to-trial variability in neural responses to repeated stimuli is often observed empirically [1, 17–19] (though see also [20–24]). Even though there are many likely contributors to this variability, ranging from stochasticity in neurotransmitter release or ion channels ([25], but see [26, 27]) to confounding factors like behavioral state, activity level, and levels of adaptation [28, 29], chaotic interactions may play an inescapable role. Indeed, chaos may represent a mechanism by which other sources of variability are amplified. Importantly, we find that chaos in driven networks produce a type of intermittent variability, with some spiking “events” predictably repeated across trials [14, 15] (cf. [30–32]). This is also reported in a number of experiments (see e.g. [33, 34] and [23]).

The remainder of this paper is organized as follows. In the *Methods* section, we describe our network model, the input discrimination task we use throughout, and outline our implementation of the Tempotron classifier [35], which we train on the spikes produced by our network. We also compare the across-trials statistics of our network to more standard mean-field ensembles, showing that fluctuating time-dependent stimuli can induce complex statistical structure across trials. In the *Results* sections, we first present an analysis of single neuron statistics in our network. This is followed by an overview of dynamical systems concepts useful to describe chaotic dynamics, which we use to interpret the spike outputs our model produces. Finally, we present our main findings based on the performance of trained classifiers, offering explanations of underlying mechanisms at the level of dynamical network interactions.

Methods

We study a recurrent network of excitatory and inhibitory neurons with random, sparse coupling, as in [2, 3, 6, 14]. Every neuron $i = 1, \dots, N$ in our network receives an external input signal $I_i(t)$, which we describe in more detail below. The collection of all these signals is an N -dimensional input that we denote by $I(t) = \{I_i(t)\}_{i=1, \dots, N}$ and simply call the *network's input or stimulus* (see Fig 1 (A)). We emphasize that the inputs have N independently varying components; i.e., they are N -dimensional.

Our goal is to explore the ability of a decoder to discriminate between two distinct network inputs $I^A(t)$ and $I^B(t)$, based on the activity of the network (or of a subset of cells from the network). We now describe the model in detail, as well as the specific discrimination task and decoding paradigm we use.

Model

We consider a network of N neurons separated into excitatory (E) and inhibitory (I) populations of sizes N_E and N_I obeying “Dale’s Law” [36] (i.e. neuron can either have all inhibitory or all excitatory outgoing synapses). We set $N_I = 0.2N$, $N_E = 0.8N$ and couple the network according to the random and sparse, *balanced state* architecture [2, 3, 6, 14, 37] following a standard Erdős-Rényi scheme with mean in-degree $K \ll N$. This means a synaptic connection from a neuron j to a neuron i is drawn randomly and independently with probability $K/N_{E,I}$ where E or I denotes the type of neuron j . Throughout the paper, we report simulations carried out with N in a range from 500 to 5000 and K from 10 to 500 but note that the majority of the results we outline are independent of network size, or scale with N in simple ways while K plays a more subtle role that does not impact the qualitative nature of our results.

Individual neurons are modelled as Quadratic-Integrate-and-Fire (QIF) units [38]. The state of each neuron $i = 1, \dots, N$ is represented by a voltage variable $v_i \in (-\infty, \infty)$. These voltages evolve according to intrinsic voltage-dependent (QIF) dynamics

$$\tau_m \frac{dv_i}{dt} = \frac{(v_i - v_R)(v_i - v_T)}{\Delta v} + I_{syn}(t) \tag{1}$$

where τ_m is the membrane time-constant, v_R and v_T are rest and threshold potentials, $\Delta v = v_T - v_R$ and $I_{syn}(t)$ represents incoming inputs to the neurons, both from the network and from our external stimulus. The dynamics of Eq (1) are discontinuous in time: once the membrane potential v exceeds the threshold, it blows up to infinity in finite time at which point a spike is said to be emitted and v is manually reset to $-\infty$. For convenience, we use $v_T = -v_R = 1$ and apply a smooth change of coordinates $v(\theta) = \tan(2\pi\theta - \pi)/2$ that maps these unbounded values to phase variables $\theta_i \in [0, 1]$. QIF dynamics acting in these coordinates is known as the *θ -neuron* model [38]. Here $\theta_i = 0$ and $\theta_i = 1$ represent the same state of $v_i = \pm\infty$, and a neuron is said to “spike” once it reaches this point. Mathematically, this means that the voltage of each neuron is represented by a point on a circle and the state of the entire network at time t is given by the vector of phases $\theta(t) = (\theta_1(t), \dots, \theta_N(t))$. For the sake of clarity, we report spikes and other temporal observables using milliseconds by fixing the neural time constant $\tau_m = 10$ ms in the QIF coordinates and rescaling unit-less time by $2\pi\tau_m$ (see [38] for more details about this coordinates change).

Thus, the state of the network as a whole can be thought of as a moving point on an N -torus (\mathbb{T}^N). The dynamics of each neuron –representing an axis on the torus– is given by

$$\dot{\theta}_i = F(\theta_i) + Z(\theta_i) \sum_{j=1}^N a_{ij}g(\theta_j) + Z(\theta_i) \underbrace{[\eta + \varepsilon\xi_i(t)]}_{I_i(t)} \tag{2}$$

where $F(\theta_i) = 1 + \cos(2\pi\theta_i)$, $Z(\theta_i) = 1 - \cos(2\pi\theta_i)$ (the canonical phase response curve of a Type I neuron [39]), and $g(\theta_j)$ is a sharp “bump” function, nonzero only near the spiking phase $\theta_j = 1 \sim 0$. As in [14, 15], we set

$$g(\theta) = \begin{cases} d \left(b^2 - \left[\left(\theta + \frac{1}{2} \right) \bmod 1 - \frac{1}{2} \right]^2 \right)^3 & ; \theta \in [-b, b] \\ 0 & ; \text{else} \end{cases}$$

with $b = 1/20$ and $d = 35/32$. This phase coupling function is chosen to model the rapid rise and fall of post-synaptic currents, while being differentiable everywhere so that the vector field defined by Eq (2) remains smooth. Note also that $\int_0^1 d\theta g(\theta) = 1$. The synaptic weight a_{ij} from neuron j to neuron i , if non-zero, is set to a fixed value that depends on both neuron types. Abusing notation slightly, we set: $a_{EE} = a_{IE} = \alpha/\sqrt{K}$, $a_{EI} = -\alpha/\sqrt{K}$ and $a_{II} = -\rho\alpha/\sqrt{K}$. We set $\alpha = 0.35$ and $\rho = 0.75$ so that I-neurons are slightly less inhibited by other I-neurons than E-neurons, as in the original balanced state architecture [2]. The term $I_i(t)$ represents an external input stimulus to neuron i ; here modelled by the sum of a DC current η and independent Gaussian white-noise processes $\xi_i(t)$ scaled by ε . We note that η can take negative values which places neurons in an excitable regime [38]. Both η and ε are global parameters that are fixed across all neurons. Throughout most of the paper, they are set to $\varepsilon = 0.5$, $\eta = -0.5$ so that the network is fluctuation-driven, producing sustained irregular activity characterized by a broad firing rate distribution with a mean roughly between 10 and 20 Hz. Together, the input signals to each neuron form the global input to the network $I(t) = (I_1(t), \dots, I_N(t))$ as depicted in Fig 1 (A). We require that the realizations $I_i(t)$ be statistically independent from each other across neurons i , so that the network’s input $I(t) = \{I_i(t)\}_i$ is free of redundancies and can be thought of as a N -dimensional signal, but we briefly address the implication of such correlations in the Discussion section.

We stress that $I(t)$ models inputs to the system, and not the various molecular and cellular sources of noise associated with neuronal dynamics. We give more details about the statistics of $I(t)$ below, but note that in this regime and all other considered in this paper, we have verified that our network is chaotic by showing the presence of positive Lyapunov exponents, a standard measure for sensitivity to initial conditions (see [14, 15] for more details).

Numerical simulation details. A standard Euler-Maruyama scheme [40] was used to numerically integrate Eq (2), treated as a stochastic differential equation (SDE). Because of the state-dependent factor multiplying white noise, multiple interpretations of the SDE are possible (see, e.g., [41] for an explanation). Following [42], we interpret our equation as a Stratonovich SDE. Accordingly, a correction term $\frac{\varepsilon^2}{2} Z(\theta_i)Z'(\theta_i)\Delta t$ is added to the right hand side of the discretized equation [40]. A time-step of $\Delta t = 0.05$ ms was used for all simulations. We verified that smaller temporal resolution did not change our results. For estimates involving sampling of many trajectories within response-ensembles, initial states of the network were uniformly sampled over state space \mathbb{T}^N . Numerical estimates of Lyapunov exponents were obtained by evolving the associated variational equation of Eq (2); further details can be found in [14, 15]. Large batched simulations were carried out on the NSF XSEDE *Science Gateway* supercomputing platform.

Numerical simulations were implemented in the Python and Cython languages. Computations of statistical quantities such as pairwise trajectory distances and spike-time reliability, as well as Tempotron classifier training and testing, were implemented in MATLAB.

Discrimination task and response ensembles

The stimulus $I(t) = (I_1(t), \dots, I_N(t))$ mimics a collection of highly featured temporal inputs to each neuron in the network. In this framework, the response of the network to a specific input can be modelled by “freezing,” or choosing specific realizations of the stochastic processes $\xi_i(t)$, the fluctuating component of $I_i(t)$; this is sometimes known as a “frozen noise” experiment [26, 27]. We emphasize that even though input components $I_i(t)$ ’s are modelled as random processes, they represent signals impinging on the network, rather than various biological sources of noise. In fact, we do not model any biological noise; our model should be viewed as a deterministic, non-autonomous dynamical system. All trial-to-trial variability is generated by chaos which amplifies discrepancies in initial states. These discrepancies are abstracted as randomly sampled initial states and may, in reality, be produced by noise. We revisit this distinction in greater detail in the *Discussion*.

Response ensembles. Suppose that an input $I(t)$ is presented to the network starting at time $t = 0$, and that the initial state of the network at that time is $\theta^0 = (\theta_1^0, \dots, \theta_N^0)$ (i.e., the “voltages” of all neurons at $t = 0$). The subsequent evoked trajectory $\theta(t)$ is uniquely defined and depends on both $I(t)$ and θ_0 . However, for a chaotic system like our network, even small changes in θ_0 , or even worse, a completely unknown initial state, can lead to big differences in the subsequent trajectories $\theta(t)$. If one of two inputs $I^A(t)$ and $I^B(t)$ is presented to the network starting at $t = 0$ and that the initial state is unknown, could a decoder reading out the evoked network activity be trained to tell which input was shown? How similar can the two inputs be before the decoder fails? To answer these questions, an essential concept is that of “response ensembles.”

To begin, consider an ensemble of many network states θ^0 . To this ensemble corresponds another one, consisting of network responses, i.e., a version of the network’s activity parameterized by $t > 0$ for each choice of initial state θ^0 , in response to the same $I(t)$. Each “response” in this ensemble represents a different “trial,” much like in an experiment where exactly the same stimulus is repeatedly presented to a system. Trial-to-trial variability thus depends on how distributed these ensembles of responses are about the network’s state space.

Formally, we define a *response ensemble* associated with an input $I(t)$ as Θ_I : the collection of network trajectories through state space, in response to $I(t)$, for which initial states were sampled independently from some initial probability distribution. In this paper, we choose this distribution to be the uniform one, meaning that each point θ^0 in state space has equal probability of being picked. We call individual trajectories within a response ensemble *trials*. A response ensemble is thus indexed by 3 components: neuron number i (“space”), time t , and trial number l as depicted in Fig 2 (A). For example, $\Theta_I(i, t, l)$ represents the state of neuron i at time t on trial l , the average across trials $\langle \Theta(i, t) \rangle_l$ represents the peristimulus time-histogram (PSTH) [43] of neuron i , etc. Throughout, we also study the spike patterns associated with the trajectories forming Θ_I and often refer to the collections of these spike times as response ensembles as well (see Fig 2 (A)).

We note that the response ensembles defined here differ fundamentally from the statistical ensembles that underlie the *mean-field* (MF) theories often used in mathematical neuroscience (see e.g. [2, 16]). While response ensembles are defined for a given, specific stimulus history $I(t)$ (so that for each stimulus realization there is, in principle, a different ensemble), the statistical ensembles used in those other theories typically consist of trajectories generated by both random initial conditions *and* independent, random stimuli. As such, those ensembles cannot be used to model the type of discrimination task we are interested in here. In more probabilistic language, equating “ensembles” with stochastic processes, our response ensembles are precisely the conditional probability distributions obtained by conditioning a MF-type statistical

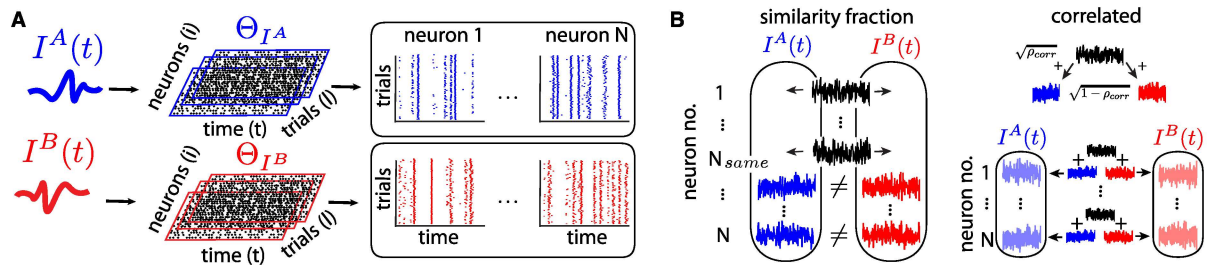


Fig 2. (A) Two ensembles of spike patterns are generated by distinct, N -dimensional inputs I^A and I^B repeatedly presented to the network on multiple trials, where initial states are chosen at random. This results in two network-wide response ensembles Θ_{IA} and Θ_{IB} containing spike patterns across neurons, time and trials. **(B)** Illustration of the two procedures used to control input similarity. Left: ρ_{same} controls the number of neurons $N_{same} = \rho_{same} N$ that receive identical inputs $I_i(t)$ under both stimuli $I^A(t)$ and $I^B(t)$. Right: correlation coefficient ρ_{corr} controls the correlation of all pairs of neural inputs $I_i^A(t)$ and $I_i^B(t)$ (note that $I_i(t)$ and $I_j(t)$ remain uncorrelated if $i \neq j$).

doi:10.1371/journal.pcbi.1005258.g002

ensemble by a particular stimulus history $I(t)$; see, e.g., [44] for a more detailed discussion of response ensembles.

The choice of statistical ensemble can have significant consequences for statistical measures like correlations (between cells and in time): for the same system, correlations computed with MF-type statistical ensembles can differ from those computed with respect to the response ensemble evoked by a specific $I(t)$. We will return to this distinction in the context of our networks in the *Results* section.

Discrimination task. With the concept of response ensembles, we can now give more precise formulations of various questions related to discrimination. First, we ask how much overlap exists between the response ensembles Θ_{IA} and Θ_{IB} . This is controlled by two factors: how much state space is occupied by each ensemble at a given time, and how close they are to each other. Second we train a decoder on the spike patterns associated with the network responses. We describe this decoder in the next section. Importantly, when we compare the network’s response to two stimuli, we also require ε and η to be the same for $I^A(t)$ and $I^B(t)$. This way, the averaged statistics of network responses are identical for any stimulus pair, and discrimination must rely on the differences in random fluctuations between the two specific realizations $I^A(t)$ and $I^B(t)$.

To study our network’s sensitivity to changes in its inputs, we introduce two notions of the similarity, or “distance”, between the stimuli $I^A(t)$ and $I^B(t)$. Fig 2 (B) illustrates the effect of both paradigms. First, we define $\rho_{same} \in [0, 1]$ as the proportion of the network’s neurons that receive identical inputs ($I_i^A(t) = I_i^B(t)$) under both A and B paradigms; the remaining fraction $1 - \rho_{same}$ receive non-identical, independent inputs ($I_i^A(t) \neq I_i^B(t)$). Second, we vary the correlation $\rho_{corr} \in [0, 1]$ between input pairs driving each neuron i , simultaneously for all i . Thus, $I_i^A(t)$ and $I_i^B(t)$ are jointly gaussian processes, each a realization of white noise, with correlation coefficient equal to ρ_{corr} . Note that the inputs to distinct cells remain independent regardless of stimulus choice ($\langle I_i^{A,B}(t) I_j^{A,B}(t) \rangle_t = 0$ for $i \neq j$) and that ρ_{corr} and ρ_{same} only control the similarity of inputs to the same neuron across different stimuli. In both cases, $\rho_{corr} = \rho_{same} = 0$ enforces that all pairs ($I_i^A(t), I_i^B(t)$) are independent, whereas if $\rho_{corr} = \rho_{same} = 1$, they are identical.

The Tempotron and discrimination

Our goal is to discriminate between two statistically identical random stimuli based on the network responses they evoke. In the second part of the paper, we do this by training a classifier on the collections of spike times $t = t_s^i$ evoked by distinct network inputs $I^A(t)$ and $I^B(t)$ on

finite time intervals $[0, T]$. There are many machine-learning techniques that can perform this task, but our main criteria for a preferred approach are: (i) it should be a useful metric to compare the encoding performance of our network under different conditions and (ii) it should isolate important spike features for coding (interpretability). We therefore opt for a simple, linear classification approach. The results serve as a lower bound on the classification capacity of the network.

There are several approaches one can use to find a hyperplane that separates sets of points in a high-dimensional space, such as the Support Vector Machine [45] (see [46] in the context of spiking data) or other regression techniques. Here, we use a classification method called the *Tempotron* [35], a gradient-descent approach acting on linear weights of temporal kernels designed to mimic the post-synaptic potentials induced by individual spikes. Importantly, the resulting classification hyperplane corresponds to the threshold of a spiking Linear Integrate-and-Fire (IF) neuron model.

The Tempotron receives vector-valued filtered spike trains $s(t) = \{s_i(t)\}_{i=1 \dots N}$ where $s_i(t) = \sum_{t_s^i} K(t - t_s^i)$ and $K(t) = V_0[e^{-t/\tau_1} - e^{-t/\tau_2}]$ where V_0 is a normalizing constant. The double-exponential filtering is meant to mimic the rise and fall of synaptic potentials in an IF neuron whose voltage obeys the equation

$$V(t) = \sum_{i=1}^N w_i \sum_{t_s^i} K(t - t_s^i) + V_{rest} \tag{3}$$

with voltage threshold set at $V_{thr} = 1$. Thus, the Tempotron computes the sum of the filtered network spike trains, according to the readout weights w_i and the timescale set by its kernel. We tune the filter’s decay and rise time-constants to $\tau_1 = 20$ and $\tau_2 = 3.75$ ms as in [35], to impose an intrinsic sensitivity to spike timing at that resolution.

The Tempotron’s goal is to fire at least one spike (i.e. cross its threshold V_{thr}) when presented with a network spike output associated to $I^A(t)$ while refraining from firing when the network responds to $I^B(t)$. Following [35], we train the Tempotron to classify spike outputs on finite-length time intervals $[t_0, t_0 + T]$ using a fixed number of trials from Θ_{IA}^t and Θ_{IB}^t respectively, and test the trained classifier on new trials. Thus, beyond discriminating between two “training” ensembles of spikes, we test the ability of the Tempotron to generalize and discriminate new patterns, related to training sets in that they are sampled from the same chaotic response ensemble. The robustness of the Tempotron to Gaussian, random spike-time jitters is well established [35]; here we investigate the effect of chaotic variability of the type generated by our networks.

Fig 3 illustrates the process. Out of 100 trials from each ensemble, we select 50 from Θ_{IA} and 50 from Θ_{IB} to train the readout weights w_i in Eq (3). Fig 3 (A) illustrates the filtered spike output of a randomly selected neuron on the training trials from I^A and I^B . The remaining trials will be used for testing. We employ the algorithm described in [35] to find a set of weights w_i imposing that the voltage V of the Tempotron will exceed the threshold $V_{thr} = 1$ at least once in the pattern time-window when presented with spikes from a trial from Θ_{IA}^t while remaining sub-threshold when receiving spikes from Θ_{IB} . We report results for a time window of $T = 1.25$ seconds. We find that training often fails for $T < 50$ ms, but that the results we present below remain qualitatively unchanged for bigger T . We use a margin of $V_{thr} \pm 0.1$ in the training to ensure a good separation (c.f. [35]). Fig 3 (B) shows the output of the Tempotron during training and testing while panel (C) compares the training and testing outcomes by showing the maximal V within the T -window. In this example, even though inputs are quite similar ($\rho_{same} = 0.9$), only a few test trials are misclassified.

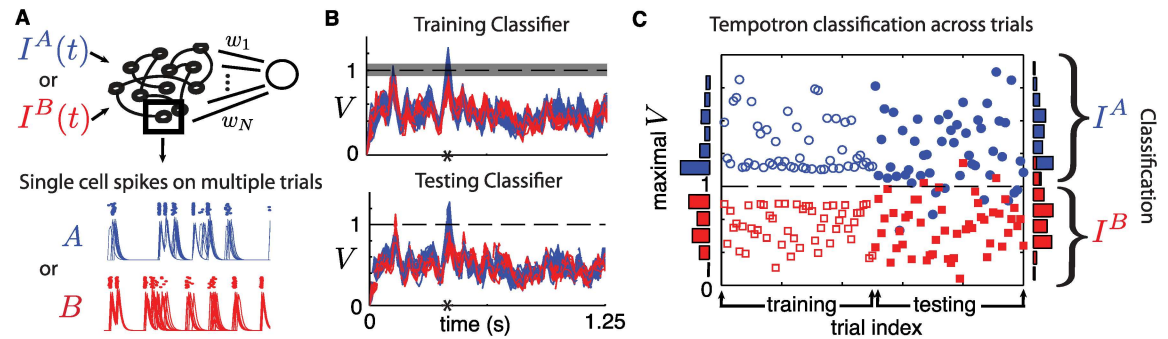


Fig 3. (A) Schematic of a Tempotron with readout weights w_i . For one neuron from the chaotic network, we show spike times (dots) of a randomly chosen neuron across several trials for inputs I^A or I^B , together with the filtered traces for these spike times. **(B)** Tempotron voltage traces $V(t)$ for a time window of $T = 1.25$ seconds. Top: traces used for training. Bottom: traces used for testing. Grey shaded area surrounding threshold V_{thr} shows training margin. Asterisk shows time of maximal output. **(C)** Classification for all training and testing trials. Markers show the Tempotron’s maximal voltage V in time window T . A correct classification corresponds to blue circles (I^A) above threshold and red squares (I^B) below. For all panels, the stimuli I^A or I^B have $\rho_{same} = 0.9$.

doi:10.1371/journal.pcbi.1005258.g003

To quantify the discriminability of spike patterns, we define the *performance* P as the fraction of successful test classifications. Note that P has a maximal value of 1 and a minimal value of 0.5 which corresponds to chance. We average P by retraining the Tempotron 20 times using different training and testing trials from our ensembles. As an example, the performance P of the classification in Fig 3 (C) is about 0.9.

Asynchronous activity and mean-field approximations

Our network has basic statistical features expected from prior work on sparse balanced networks: the activity of neurons is mostly decorrelated in time and between cells, and is statistically stationary (see e.g. [2, 16], but also our note on “correlations” below). These features are useful in deriving simplified expressions for neuronal population temporal statistics. In *mean-field* (MF) theory, one replaces a complex network of interacting units with a reduced model of a single unit driven by a relatively simple independent stochastic process, meant to model the outputs of all other units within the network impinging on the given unit. We now follow this approach to (1) derive an analytic expression for E and I population firing rates and fluctuation amplitudes, and (2) demonstrate that our network operates in an asynchronous balanced regime that is consistent with prior work. The results of this section will help us design surrogate population dynamics—used throughout the paper—to compare chaotic dynamics to basic assumptions of independent noise.

A note about our use of the term “correlation:” as mentioned earlier, even for the same system, correlations depend on the choice of statistical ensemble. Because the term has many potentially different meanings, we have strived to be as explicit as we reasonably can in what follows. In this section, we are concerned with correlations averaged across independent stimulus realizations, as is the custom in MF theory. The nature of correlations across trials in our network is revisited in *Results*, and is treated in detail in [14, 15].

We follow a similar approach to [47, 48] where QIF neurons driven by noise and in a MF setting are studied for exponentially decaying synapses. In contrast to networks of leaky integrate-and-fire (LIF) neurons, QIF dynamics introduce some complications because their firing rate response to fluctuating drives is not straightforward to compute (c.f. [48–50]). Here we follow the general approach of [16] using a number of simplifications to arrive at an estimate for the mean firing rate or neurons in our network.




Fig 4. (A) Top: Typical normalized spike-time cross-correlogram $C_{ij}(u)$: expected proportion of pairs of spikes separated by u ms, for random pairs of neurons within the network. In blue: for 30 individual pairs; in black: average. Bottom: Typical normalized spike-time auto-correlogram $C_{ii}(u)$ for a neuron in the network. The QIF dynamics impose a relative refractory period close to spiking which gives a departure from purely Poisson firing statistics (notice the dip toward time 0). Typical Fano Factors are between 0.77 and 1. For A and B, $N = 500$, $K = 20$; however, the shape of these functions is not affected by N or K . Statistics sampled over network simulations of 300 seconds. **(B)** Mean firing rate of E and I neurons within the network as a function of K . Dotted line shows MF estimate (see Eq 8). Various markers show estimates sampled from all neurons in simulated networks of different sizes N . Error bars show one standard deviation. Inset shows zoom for low- K values.

doi:10.1371/journal.pcbi.1005258.g004

We begin by re-writing Eq (2) as

$$\frac{d\theta_i^\Omega}{dt} = F(\theta_i^\Omega) + Z(\theta_i^\Omega)[I_{net}^\Omega(t) + I_i(t)] \quad (4)$$

with Ω denoting neuron type (E or I) and

$$I_{net}^\Omega(t) = \sum_{j=1}^N a_{ij} g(\theta_j(t))$$

$$I_i(t) = \eta + \varepsilon \zeta_i(t).$$

Since connectivity is sparse and K, N are large, we assume that (i) Pre-synaptic spike trains to a neuron are statistically independent and (ii) Each spike train is Poisson distributed with constant rate v_E or v_I . Assumption (i) is justified by the flat spike-time cross-correlograms observed across pairs of neurons in the network, shown in the top panel of Fig 4 (A). This holds for networks as small as $N = 500$ and in-degree as small as $K = 10$. Assumption (ii) on the other hand is not quite met, as shown by the dip in the typical spike-time auto-correlogram of a neuron in the bottom panel of the same figure. This dip occurs because QIF dynamics produce a relative refractory period, leading to typical Fano Factors lower than one (between 0.77 and 1 in all of our simulations). This refractory period can be regarded as a realistic feature; in any case, the use of assumption (ii) still leads to correct estimates for our dynamical regime, as was also observed in [47].

Since $\frac{d\theta}{dt}|_{\theta=1} = 2$, we make the approximation $\int dt g(\theta_j(t)) = \frac{1}{2} \int d\theta_j g(\theta_j) \simeq \frac{1}{2} \int d\theta_j \delta(\theta_j - 1)$. Thus, the mean and variance of the compounded network input $I_{net}^\Omega(t)$ over a small time interval Δt are approximated by $K(a_{\Omega E} v_E - a_{\Omega I} v_I)\Delta t/2$ and $K(a_{\Omega E}^2 v_E + a_{\Omega I}^2 v_I)\Delta t/4$ respectively. It follows that the dynamics of a typical neuron of type Ω can be represented by the following stochastic differential equation (SDE) [16]

$$\frac{d\theta_i^\Omega}{dt} = F(\theta_i^\Omega) + Z(\theta_i^\Omega) \underbrace{[\mu_\Omega + \sigma_\Omega \zeta(t)]}_{\sim I_{net}^\Omega(t)} + \underbrace{\eta + \varepsilon \zeta(t)}_{I_i(t)} \quad (5)$$

with

$$\begin{aligned}
 \mu_E &= \frac{\alpha}{2} \sqrt{K} (v_E - v_I) \\
 \mu_I &= \frac{\alpha}{2} \sqrt{K} (v_E - \rho v_I) \\
 \sigma_E^2 &= \frac{\alpha^2}{4} (v_E + v_I) \\
 \sigma_I^2 &= \frac{\alpha^2}{4} (v_E + \rho^2 v_I)
 \end{aligned} \tag{6}$$

where $\zeta(t)$ and $\xi(t)$ are independent Gaussian white noise processes with zero mean and unit variance. We can combine both input terms $I_{net}^\Omega(t)$ and $I_i(t)$ in Eq (5) to get an SDE with a single stochastic term:

$$\frac{d\theta_i^\Omega}{dt} = F(\theta_i^\Omega) + Z(\theta_i^\Omega)[(\mu_\Omega + \eta) + \sqrt{\sigma_\Omega^2 + \varepsilon^2} \zeta(t)]. \tag{7}$$

Here, the combination of network interactions and of the external input signal into one noise term is made possible by the fact that mean E and I firing rates do not depend on any specific realization of input $I_i(t)$, as long as its mean and fluctuation strength are given by the fixed parameters η and ε . Thus, just like network interactions, we approximate the effect of any external input by stochastic noise. We show later that for a fixed input stimulus $I_i(t)$ across trials, as for “frozen noise” experiments (c.f. [51, 52]), both stochastic approximations about network and external inputs fail, because important statistical dependencies are introduced across trials.

Using Equation (3.22) from [48] and the standard change of coordinates from θ to v and time rescaling described above, we obtain:

$$\bar{v}_\Omega = \frac{1}{\tau_m} \left[\int_{-\infty}^{\infty} \frac{dx}{\sqrt{\pi}} \exp \left(- \left(\beta_\Omega - \frac{\tau_s}{\tau_m} \pi \gamma_\Omega^2 \right) x^2 - \frac{\pi^2 \gamma_\Omega^4 x^6}{12} \right) \right]^{-1} \tag{8}$$

where $\beta_\Omega = \mu_\Omega + \eta$, $\gamma_\Omega^2 = \sigma_\Omega^2 + \varepsilon^2$, τ_m is a membrane time-constant and τ_s is a synaptic time constant for exponential synapses. We choose $\tau_m = 10$ ms and take $\tau_s \rightarrow 0$ as we already approximated our synaptic interactions with δ -pulse coupling. This gives time units of ms to the self-consistency Eq (8). We solve for v_E and v_I numerically using a function-minimization algorithm (“fminsearch” in MATLAB).

Fig 4 (B) shows that the mean field approximation for the firing rates of neurons in the network matches simulations with great precision for a wide range of N and K , and that the average firing rates are independent of N and change monotonically with K , as expected from previous work [16, 47]. Later, we will make use of the quantities $\mu_{E,I}$ and $\sigma_{E,I}^2$ derived in Eq (6) to outline the difference between network and MF responses for fixed stimuli.

Results

We now discuss stimulus discrimination based on the output of our network. We begin by investigating the role of network interactions in shaping the statistics of trial-to-trial population responses, comparing these to a simplified model of network dynamics that follows a mean-field approach. We then offer a state-space view of the mechanisms responsible for the observed statistics from the perspective of random dynamical systems. From there, we use spike pattern classification to investigate stimulus encoding in the presence of chaos and formulate our main findings.

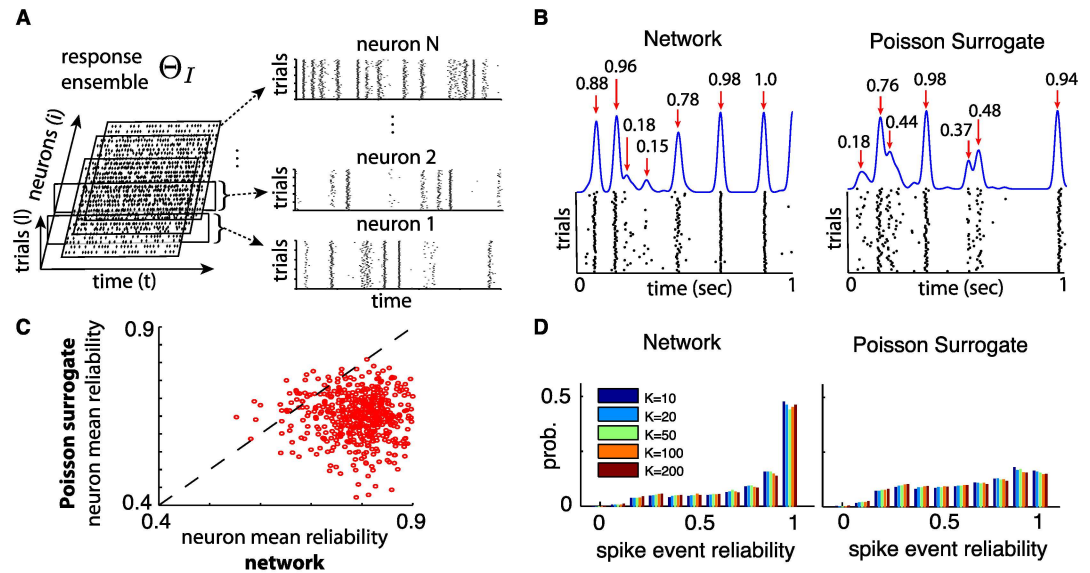


Fig 5. (A) Spike patterns from a response ensemble, Θ_I , containing network responses to input $I(t)$ on many trials. **(B)** Single neuron spike response examples across 50 trials. Left: from a randomly chosen neuron in the network. Right: same neuron as right but replacing network interactions by Poisson spike surrogates following MF assumptions (see text for details). Bottom row shows spike raster plots and top row shows corresponding Gaussian-filtered PSTH ($\sigma = 10$ ms). Red arrows indicate spike events and numbers indicate spike event reliability. See text for details about calculation of reliability. **(C)** Scatter plot showing the mean reliability of each neuron in a network ($N = 500$, $K = 20$) against its counterpart in the Poisson surrogate inputs paradigm. Neuron mean reliability is the average reliability of all spike events produced by the neuron. **(D)** Left: Histogram of spike event reliability sampled over all spike events produced by networks with different values of K ($N = 2000$). Right: Same as left for Poisson surrogate inputs paradigm. For (C) and (D), quantities sampled from simulations lasting 100 seconds over 100 trials.

doi:10.1371/journal.pcbi.1005258.g005

Network interactions are essential for response reliability

We begin by illustrating the statistical structure that emerges across trials when our networks are driven by fluctuating stimuli. Previous work has shown that basic statistical measures such as “noise correlations” may not adequately capture across-trial statistics in driven systems [15]. Here, we follow [14, 15] and take a different approach, comparing the spike patterns produced by our network to those produced by a surrogate spiking model, based on the reduced MF approximation from the *Methods*. Our approach is very similar to that in Fig. 4c of [14], though our comparison here is more extensive and employs quantities derived from MF. Our overall goal is to separate the entraining effect of a fixed stimulus $I(t)$ from the variability caused by network interactions, or a MF approximation of them.

To do so, it is convenient to recall the response ensembles introduced in *Methods* and depicted in Fig 5 (A): $\Theta_I(i, t, l)$. In the MF Eq (8), the mean firing rates v_E and v_I represent averages over all indices of response ensembles, $\langle \Theta_I \rangle_{i,t,b}$ and are used as parameters of a Poisson process. Let us define surrogate response ensembles: $\tilde{\Theta}_I(i, t, l)$, which are generated numerically in the same way we do for our network (see *Methods*), except that we replace the interactions from network coupling with Poisson spike trains with homogeneous firing rates v_E and v_I given by Eq (8). Note that the same input $I(t)$ is presented to the N neurons in this surrogate “disconnected” network as is presented to the original chaotic network. We verified that the firing rates’ mean and variance are conserved, evidence that $\langle \Theta_I \rangle_{i,t,l} = \langle \tilde{\Theta}_I \rangle_{i,t,l}$, as expected.

For our network, it is clear from Fig 5 (A) that some spikes are repeated across all trials (*reliable* spikes) and others are not (*unreliable* spikes). We return to the mechanisms leading to this

intermittent variability in the next section. For now, we ask: how much of this spike time reliability will be conserved in our surrogate response ensemble? We quantify the *reliability* of a spike by estimating the probability of it being repeated on other trials. This is done by convolving the cross-trial spike trains of a neuron with a gaussian filter (standard deviation 10 ms) and adding the resulting waveforms from 100 trials to obtain a filtered peri-stimulus time histogram (PSTH) as illustrated in Fig 5 (B). We call each peak in this histogram a *spike event* and a spike is assigned to an event if it falls within a tolerance time-window defined by the width of the peak at half height [14, 53]. The event's *reliability* is then estimated by the number of member spikes divided by the number of trials considered. If a spike event has a reliability of 1, it means that one can expect a spike from that neuron at that moment on every trial while a lower reliability indicates more variability. Fig 5 (B) shows examples of spike events with their computed *reliability*, for a randomly chosen neuron in the network and in the Poisson surrogate system.

To quantify the reliability of individual neurons over the course of 100 seconds of activity, we compute their mean reliability, taken to be the average *reliability* of all spiking events they produce. Fig 5 (C) compares the mean reliability of neurons during network responses and in the presence of surrogate inputs. It is clear that most neurons are significantly more reliable when network interactions are present, as opposed to independent stochastic inputs. To access the overall reliability of population-wide response ensembles, we estimate the distribution of spiking event *reliability* over all neurons. Fig 5 (D) shows histograms obtained by sampling all spike events in the network, together with surrogate ensembles. Here we take $N = 2000$ and multiple values of K , simulated for 10 seconds over 100 trials. A number of features can be observed from these plots: (i) Despite chaotic dynamics, a majority of spike events in driven networks are reliable (as in [14]). (ii) There are substantially more reliable spikes in networks than in surrogate populations (cf. [14, 15]). (iii) Spike reliability does not depend on the connectivity in-degree K .

Important conclusions can be drawn from these observations. First, as shown in [14, 15], MF-like approximations of network interactions are insufficient to capture the statistical properties of trial-to-trial variability, and actually over-estimate it. Cross-trial statistics in chaotic networks show population-wide dependencies despite independent temporal statistics, a feature that is lost by averaging network interactions.

Second, it shows that reliable spikes observed in chaotic networks are not solely due to strong driving inputs. Such reliability could indeed be generated if inputs $I_i(t)$ had fluctuations much stronger than the ones from the network. This would mean that external inputs drown network interactions, effectively driving neurons into an “uncoupled” regime for which reliable responses are expected [51, 52]. For the parameters chosen here ($\eta = -0.5$, $\varepsilon = 0.5$) however, the input and network fluctuations have comparable strengths with $\varepsilon/\sigma_{E,I} \simeq 1.1$ (see Methods). Additionally, the greater proportion of reliable spikes from the network compared to the surrogate population indicates that some reliable spikes in a given neuron are generated when the network inputs to that neuron are repeated across trials. We verified that spike reliability is unchanged by network size N .

Finally, the fact that spike reliability is robust to moderate changes in connectivity in-degree K (Fig 5 (D)) is consistent with the scaling of fluctuation strengths $\sigma_{E,I}$. Indeed, $\sigma_{E,I}$ does not explicitly depend on K and is independent of N (see Eq (6) in Methods). The ratio $\varepsilon/\sigma_{E,I}$ only changes with overall firing rates $\nu_{E,I}$ which vary slowly and monotonically with K (see Fig 4 (B)). This suggests that N should have little effect on the discrimination properties described in the rest of this paper, and that K 's only influence manifests through its modulation of $\nu_{E,I}$. Thus, we concentrate on networks of fixed size $N = 500$ and in degree $K = 20$. We chose these parameters both for ease of simulation and to demonstrate that even small networks can have rich coding properties while remaining in balanced regimes.

We summarize these findings in the following heuristic description. The response variability of our recurrent network is characterized by intermittent reliability, where some spikes are always reproduced across trials and others are not. Which spikes are reliable depends on the input and its history in non-trivial ways. In some cases, a spike is reliably elicited when the signal $I_i(t)$ has a sufficiently strong upswing, thus directly encoding a feature of a local input. However this is not the only way reliable spikes are produced; others occur when a neuron receives network input from other cells with coincident reliable spikes. Network interactions are themselves partially stimulus locked, but still show some cross-trial variability because of chaos. This has the effect of increasing the “fraction” of all inputs to a neuron that is frozen from trial to trial, compared to the naive mean field assumption, thus increasing reliability. It is this interplay between network and external inputs that create complex response statistics, a signature of which is the intermittent presence of reliable spikes.

Intriguingly, similar types of “intermittent” spiking variability have been reported in *in vivo* experiments (cf. [33, 34]). As we see below, this is best described by attributes of stimulus-dependent chaotic attractors, with low dimensionality, and occupying specific regions of state space. These network interactions are not easily captured by the kind of statistical ensembles usually used to derive MF equations, in which one considers system trajectories with random initial conditions *and* independent stimuli; the assumptions that are valid for such ensembles no longer hold for our response ensembles, which are predicated on a specific stimulus.

Thus, we conclude that *trial-to-trial variability in chaotic networks is more complex, and less severe than that of simplified stochastic models, leading to a great number of reliable spikes repeated across trials*. While it is conceivable that a stochastic model of network interactions can be derived to capture this phenomenon, it is not clear how to implement the various statistical dependencies outlined above. We show below that a state space view based on a dynamical systems approach is better suited to understand the mechanisms underlying this phenomenon.

Chaotic attractors shape statistical dependencies across trials

To better understand the prevalence of reliable spikes in driven networks, and why this reliability appears to be intermittent (see Fig 5 (A)), we turn to a geometric view of network dynamics as captured by response ensembles. Below, we first review important concepts of input-driven chaos that were originally presented in [14, 15, 54]. In turn, we relate features of chaotic attractors to trial-to-trial variability and then to input-dependence, toward our goal of studying discriminability.

Chaos: A state space view of variability. Recall that the variability of a network’s response to $I(t)$ is characterized by the breadth of differences between the trajectories forming its response ensemble Θ_t , i.e., the “size” of Θ_t as measured by, e.g., its diameter. Although difficult to describe by simple mathematical formulae, for specific systems, response ensembles are well-defined mathematical objects with well-understood geometric properties that can be numerically characterized. The theory of random dynamical systems (RDS) provides a framework for studying these properties.

Snapshots of all trials in Θ_t at any time $t > 0$ are ensembles of points, corresponding to a probability distribution that describes all possible network states given that the system has been subjected to the stimulus $I(t)$ up to time t . Taken together for all $t > 0$, these “snapshot distributions” [55] define input- and time-dependent probability distributions describing the network’s evoked activity for all possible initial states at once and they dictate statistical attributes of our network such as trial-to-trial variability. One of the key results of RDS theory is that under very general conditions, the ensembles Θ_t are concentrated on time-evolving

geometric objects known as *random attractors*, so called because almost all initial conditions, when subjected to the same forcing $I(t)$, will converge to the attractor. Unlike “classical attractors” in un-driven systems, which have fixed positions in space (e.g. fixed points, periodic orbits, strange attractors), the position of Θ_I changes in time t and with input choice $I(t)$. This is because our networks are driven by time-dependent stimuli, and are governed by non-autonomous systems of differential equations (see [Methods](#)).

Geometrically, these random attractors can be points, curves, or higher-dimensional surfaces, and can wind around state space in complicated ways. The exact position and shape of an attractor depends on both the system parameters and the specific realization of the stimulus. Incidentally, this is why attractors are poorly modeled by models of stochastic noise that assume independence across dimensions, as demonstrated in the previous section. However, RDS theory also states that certain important properties of the attractor—and thus the response ensembles Θ_I —are independent of specific choices of input $I(t)$, are invariant in time, and depend only on system parameters. An important one is the sensitivity of network responses to small perturbations, as measured by the *Lyapunov exponents* of the system. For an N -dimensional system, these exponents are real numbers $\lambda_1 \geq \dots \geq \lambda_N$ that describe the rate of separation of nearby trajectories in different state space directions. For a system like our network, the Lyapunov exponents do not depend on the choice of input realization $I(t)$ so long as it is a realization of white noise with same parameters ε and η . A criterion for chaos is the presence of positive Lyapunov exponents. Moreover, the number of positive Lyapunov exponents roughly indicates the number of unstable directions in state space, and their magnitude indicates how strong the amplification of small perturbations is in those directions. In a chaotic system, almost any nearby trajectories will diverge from each other exponentially fast, but they do so only along unstable directions of attractors.

Other geometric properties of random attractors can be related to their Lyapunov exponents as well. For example, the *attractor dimension* is a quantity describing how much of state space is occupied by the attractor—the source of trial-to-trial variability for our network—and is (roughly) given by the number of positive Lyapunov exponents. To see that this should be the case, imagine a cloud of initial conditions evolving according to the same stimulus. The state-space expansion associated with positive exponents tends to “stretch” this cloud, leading to the formation of smooth Λ^+ -dimensional surfaces, where Λ^+ is the number of positive exponents (see, e.g., [44] for a general, non-technical introduction and [56, 57] for more details). If all exponents are negative, for example, then the attractor is just a (time-dependent) point, whereas the presence of a single positive exponent suggests that the attractor is curve-like, etc.

In previous work, we showed that in a wide variety of dynamical regimes, the number of positive Lyapunov exponents in driven balanced networks is less than 20% of the network’s dimension N , and often remains below 10% [14, 15]. In most of the paper below, we choose network and input parameters, described in *Methods*, so that about 8% of Lyapunov exponents are positive. Here, the rate of trajectory separation is strong with $\lambda_1 \simeq 3.5$, and the network is by all accounts in a chaotic regime. At the same time, there are plenty of directions in which the attractor is “thin”, leading to trial-to-trial variability that is far from homogeneously random at the population level. We explore different parameter regimes, leading to distinct attractor dimensions, toward the end of this article.

Geometry of response ensembles and state-space separability. Our goal is to describe the geometry of chaotic attractors in state space sufficiently well so to explain the presence of reliable spike events described in the last section. Moreover, we use the same approach to ask about differences that arise between two attractors, generated by distinct inputs $I^A(t)$ and $I^B(t)$.

From the discussion above, we know that for two inputs $I^A(t)$ and $I^B(t)$ with identical statistics, the dimension of their associated response ensembles, and thus their level of trial-to-trial

variability, will be the same. Revisiting our discrimination task, network responses will be discriminable if and only if the corresponding ensembles Θ_{I^A} and Θ_{I^B} do not overlap most of the time. To predict when this is the case, knowing only the dimension of Θ_I and Θ_{I^B} is insufficient; we need to understand how the position of Θ_I in state space depends on the choice of $I(t)$. We will present evidence that for the balanced spiking networks studied here, there exist broad parameter regimes where the relation: “*diameter of a response ensemble* \ll *distance between different ensembles*” generally holds.

Consider the following pairwise distance quantities, the statistics of which we sample using numerically simulated network trajectories (see [Methods](#) for details):

$$\begin{aligned} X(t) &= \|\theta(t, \theta_0, I^A) - \theta(t, \tilde{\theta}_0, I^A)\| \\ Y(t) &= \|\theta(t, \theta_0, I^A) - \theta(t, \tilde{\theta}_0, I^B)\| \end{aligned} \tag{9}$$

where $\theta(t, \theta_0, I)$ denotes the network trajectory in state space, given an initial state θ_0 at $t = 0$ and subject to input stimulus $I(t)$. Initial states θ_0 and $\tilde{\theta}_0$ are independently, randomly chosen and $\|\theta_1 - \theta_2\|$ denotes the (shortest) distance between two states θ_1 and θ_2 . [Fig 6 \(A\)](#) illustrates these measurements. Both $X(t)$ and $Y(t)$ denote the distance between a pair of trajectories at time t : “within-ensemble” for $X(t)$ and “between-ensemble” for $Y(t)$. Formally, they are random variables because they depend on a pair of random initial conditions.

The quantity $X(t)$ represents the “typical size” of an ensemble at time t ; we view it like the diameter in that it measures the distance between two typical points on the corresponding attractor. The quantity $Y(t)$ is the typical distance between two ensembles elicited by stimuli $I^A(t)$ and $I^B(t)$. In the context of the present model, we say that the two stimuli are *separable* (in the state space sense) if the distributions of $X(t)$ and $Y(t)$ do not have a significant overlap (see [Fig 6](#)). The dimension of the underlying attractor is reflected in $X(t)$ in the following way: if the dimension were 0 (so the attractor is a single point), then $X(t) \approx 0$, and if the dimension is positive, then $X(t)$ would also be positive (but can be large or small). In general, $X(t)$ will fluctuate in a complicated fashion as the time-dependent attractor evolves and changes shape. Based on previous work [[14](#), [15](#)], we expect these fluctuations to be relatively small, so that $X(t)$ is nearly constant in time. Furthermore, while the fluctuations in $X(t)$ depend on the specific choice of stimulus, standard results from probability theory tell us that its mean depends only on the system parameters and the statistical distribution of the stimuli. Thus, the mean value of $X(t)$ remains unchanged if we replace $I^A(t)$ by $I^B(t)$ in [Eq \(9\)](#) (since we require $I^A(t)$ and $I^B(t)$ to have identical parameters). Similar reasoning applies to $Y(t)$, which is akin to measuring the rate of separation between two independent realizations of the same stochastic process. Once again, we expect $Y(t)$ to fluctuate around some mean, with the exact fluctuations dependent on the input realizations but with mean and variance depending only on system parameters and stimulus statistics.

From simulations, we produce pairs of trajectories subject to both $I^A(t)$ and $I^B(t)$, all with random initial states. We compute $X(t)$ and $Y(t)$ and plot the result as a function of time in [Fig 6 \(B\)](#) where $\rho_{corr} = 0.5$ (I^A and I^B are 50% correlated). As expected, after a short transient $X(t)$ settles quickly to a positive constant, which is unchanged if the trajectory pair is selected from $\Theta_{I^A}^t$ or from $\Theta_{I^B}^t$. Likewise, $Y(t)$ settles quickly to a steady-state in which it fluctuates around a well-defined mean.

For a more complete view of this phenomenon, we consider the distributions of pairwise distances, sampled over 100 trajectory pairs, starting from uniformly random states. [Fig 6 \(C\)](#) shows the time evolution of these distributions for the first second of elapsed time. As for the single-pair measurements ([Fig 6 \(B\)](#)), both distributions settle into near-constant, steady values after a very fast transient (~ 10 -50 ms). This remains true for any similarity parameter

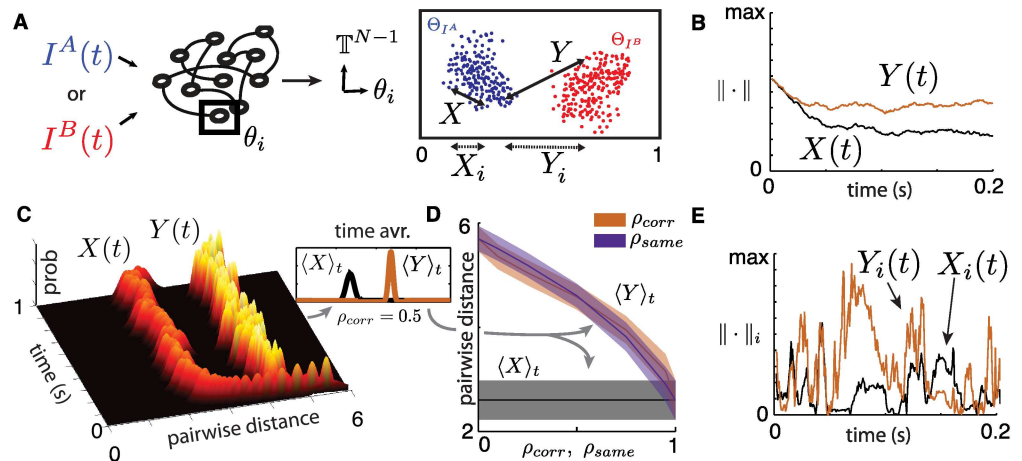


Fig 6. (A) Illustration of pairwise distances between trajectories within an ensemble (X) and across ensembles (Y). Also illustrated are equivalent distances for a single neuron's state variable θ_i (X_i, Y_i). (B) Pairwise distances between two trajectories for the entire network (on the N -dimensional torus \mathbb{T}^N , $N = 500$) as a function of time. (C) Time evolution of pairwise distance distributions for entire network ($N = 500$), sampled with 100 trials from each ensemble. After a very fast transient (~ 10 -50 ms), distributions settle into stable values. For panels (B,C,E), $\rho_{corr} = 0.5$. (D) Mean of stationary pairwise distances for a range of similarity parameters ρ_{corr} or ρ_{same} (independently modulated). Shaded areas show one standard deviation. Both mean and standard deviation were averaged over 1000 time points between $t = 100$ and $t = 12,000$ ms. (E) Pairwise distances between two trajectories in a single neuron's coordinate θ_i as a function of time.

doi:10.1371/journal.pcbi.1005258.g006

value ρ_{same}, ρ_{corr} . We note that this short transient validates our general definition of “trial” which includes trajectories with distinct initial states as well as any trajectories that received some sort of perturbation—such as a synaptic failure or the event of an extra spike—in the recent past. To capture the stationary nature of pairwise distances, we compute time-averaged distributions $\langle X \rangle_t$ and $\langle Y \rangle_t$, which we calculate using 1000 time points between $t = 100$ and $t = 12,000$ ms. While $\langle X \rangle_t$ remains the same regardless of the similarity between $I^A(t)$ and $I^B(t)$, $\langle Y \rangle_t$ can be used to measure the effect of stimulus similarity on the location of response ensembles in state space. Fig 6 (D) shows the means $\langle X \rangle_t$ and $\langle Y \rangle_t$ as well as one averaged standard deviation, for a range of input similarity parameters ρ_{corr} and ρ_{same} between 0 and 1. Here, only one similarity parameter is varied at a time, while the other is kept at zero. As expected, when both inputs are identical ($\rho_{corr} = 1$ or $\rho_{same} = 1$), $\langle Y \rangle_t$ collapses to $\langle X \rangle_t$ since Θ_{IA} and Θ_{IB} describe the same ensemble.

However, we see that $\langle X \rangle_t$ and $\langle Y \rangle_t$ become more than two standard deviations apart as soon as the stimuli become less than 90% similar. This is true for both definitions of stimulus similarity. The conclusion is that the chaotic, balanced networks at hand produce dynamical responses that stay separated in state space even for stimuli that are very similar.

Finally, the geometric attributes described above can be related to the reliability of spike times already discussed in the previous sections (see Fig 5). Previous work has shown that the unstable directions in attractors—i.e., the directions in which chaotic dynamics will spread the response ensemble produced by a single stimulus—generally align with neural coordinates. Moreover, the identity of the corresponding “unreliable” neurons change in time [14]. This leads to intermittent variability in single neurons: at any moment there is a small fraction of neurons in the network that have variable dynamics across trials, while the rest behave in a reliable fashion.

We can observe this phenomenon by restricting the definition of pairwise distance Eq (9) to single-neuron coordinates: $X_i(t)$, and $Y_i(t)$ for the (transformed) voltage variable θ_i . The

value $X_i(t) = 0$ means that different initial network states nevertheless lead to the same state for cell i at time t , i.e., the neuron i has the same voltage across all trials. A similar interpretation applies to $Y_i(t)$. Fig 6 (E) shows these quantities for a randomly selected neuron i . In contrast to pairwise distances in the full state space (Fig 6 (B)), $X_i(t)$ regularly collapses to zero. This is a reflection of the intermittent spike event reliability discussed earlier. $Y_i(t)$ also varies over time although it remains greater on average. This suggests that at any moment in time, for a given input $I(t)$, some network coordinates may offer trial-to-trial reliable responses that can be used to distinguish similar input stimuli. We investigate the use of evoked spike times for discriminability in the next section.

To summarize, pairwise distances between evoked trajectories inform us about geometric properties of response ensembles, and how they organize in state space. In parameter regimes where $X \ll Y$, which occur for stimuli that are less than 90% similar, we expect that a decoder could classify novel trajectories, given information about the distributions $\langle X \rangle_t$ and $\langle Y \rangle_t$. Conversely, in parameter regimes where $X \approx Y$, it may be difficult for a downstream decoder to discriminate response ensembles using simple criteria like linear separability. Moreover, the separation between trajectories within an ensemble remains constant in time but is supported by only a small fraction of neurons, the identity of which changes over time. This leads to intermittent spike-time reliability with statistics that have population-wide dependencies.

Thus, *the geometric attributes of chaotic attractors both enable separability of network responses evoked by distinct stimuli, even when the stimuli have similarities, and explain intermittent spike time reliability within a response ensemble.* We show below how a decoder can exploit these features to classify inputs.

Finding 1: Chaotic spike patterns are linearly discriminable

The state-space separability studied above assumes that one has access to the full state of the network at all times; any biologically realistic decoder would only have access to a network's spiking activity. We next present evidence that the *spikes* generated by balanced, chaotic networks can also be used by a simple linear classifier, the Tempotron (see [Methods](#)), to discriminate between two stimuli with identical statistical properties.

We observed that the Tempotron classification performance P roughly follows the trends found above for the pairwise distance between response ensembles (see Fig 6): there is perfect classification ($P = 1$) until input similarity reaches about 90% (i.e. $\rho_{same}, \rho_{corr} = 0.9$). This means there exists a linear combination of evoked spike patterns that reliably sum to cross a threshold for only one of the two stimuli, regardless of network initial conditions. Thus, when distinct stimuli are presented to chaotic networks, even ones with very similar features, it is not only the network states they produce that are highly discriminable, but also the resulting spike trains. We will return to the question of performance versus input similarity below, but first turn to its mechanism.

We propose that reliable spiking of a few neurons at precise moments in time (i.e., reliable spike patterns) drive successful classification of stimuli by the Tempotron. In this section we demonstrate this by deconstructing the Tempotron's readout and by observing the impact of "jittering" underlying spikes.

We first turn to the readout weights w_i , which are the result of a global optimizing algorithm [35]. Fig 7 (A) shows the values of all w_i 's for a Tempotron trained to distinguish two stimuli, $I^A(t)$ and $I^B(t)$, with perfect performance ($P = 1$). We find that the neurons with the strongest weights spike very reliably when the network is presented with input $I^A(t)$, right before the peak of the Tempotron's output (asterisk, see Fig 7 (A) bottom and Fig 3 (B)). Conversely, the same neurons either do not spike or spike unreliably in response to $I^B(t)$ around

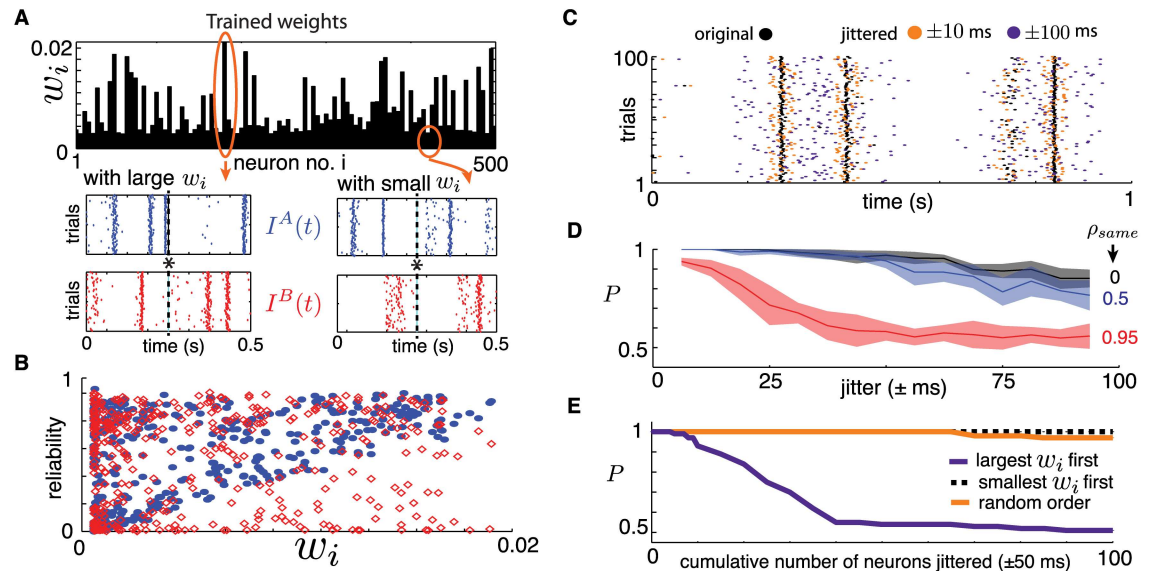


Fig 7. (A) Top: Trained Tempotron weights w_i . Bottom: example cross-trial spiking output for neurons associated with large and small weights w_i . Asterisk shows the time at which the maximal value of the Tempotron variable $V(t)$ is obtained (see also Fig 3 (B)). **(B)** Spike-time reliability for the closest spike to the maximum (asterisk), plotted for each cell against the readout weight w_i . Solid blue markers indicate responses to $I^A(t)$ and hollow red markers to $I^B(t)$. **(C)** Example raster plot for a single neuron across 100 trials, with and without spike-time jittering. Two jitter strengths shown: ± 25 and ± 100 ms. **(D)** The Tempotron's performance P when all neurons are jittered, plotted against jitter strength, for three values of ρ_{same} . The classifier was trained and tested on jittered spike time data. **(E)** Tempotron's performance P against the cumulative inclusion of neurons where the spike times were strongly jittered (± 50 milliseconds). Different lines indicate three orderings of neurons the neurons that were jittered, based on their trained readout weight w_i : (i) decreasing, (ii) increasing and (iii) random. The classifier was trained on original data and tested on jittered data from a network with $N = 500$, $K = 20$. For all panels except (D), $\rho_{corr} = \rho_{same} = 0$.

doi:10.1371/journal.pcbi.1005258.g007

the same moment in time. In Fig 7 (B), we quantify this finding by showing a scatter plot of spiking event reliability for each neuron, plotted against its output weight w_i . This shows that any neuron with a high w_i also shows high spiking reliability under input $I^A(t)$.

Next, we degrade the temporal precision of network responses by randomly “jittering” spike times and study the impact on classification performance by the Tempotron. To “jitter” a spike train, we shift each spike time on all trials by a random amount, uniformly and independently drawn from an interval of $(-r, r)$, where r is the jitter radius (see Fig 7 (C)). This leaves the total number of spikes fired the same, but strongly disrupts their temporal precision, as illustrated in Fig 7 (C). We use this jittering in two ways. First, we train and test the Tempotron on jittered spikes, to probe the dependence of classification performance on the temporal precision of spike times produced by chaotic networks. Second, we train the Tempotron on the original spike data but test using jittered spike times for subsets of neurons, to probe the learned role of these neurons in classifying stimuli.

Training and testing the Tempotron on jittered spike time data shows that classification performance P progressively declines as the jitter radius increases. The rate at which performance drops depends on the similarity between inputs $I^A(t)$ and $I^B(t)$, as illustrated in Fig 7 (D) for three values of ρ_{same} (the fraction of neurons receiving identical direct inputs under both signals). Evidently, the lower ρ_{same} is, the more distinguishing features there will be in the two response ensembles to be classified, the combination of which enables the Tempotron to classify stimuli even with substantial spike time jitter. Overall, when spikes are jittered by 10's of ms, classification performance drops significantly; for similar stimuli, performance drops

halfway to chance when the jitter radius is 25 ms. We conclude that the Tempotron uses quite precisely timed spikes to distinguish the responses of chaotic networks to nearby stimuli.

To probe the question of what spiking features lead to the reliable classification learned by the Tempotron, we set the jitter radius to 50 ms and jitter different subsets of neurons in the testing data. In contrast to the procedure described above, where training and testing spike times were jittered, only testing spike times are now jittered; the Tempotron is trained on the original spikes produced by the network. We apply this jittering procedure to an increasing number of neurons in the network, re-testing the classification performance P as neurons are cumulatively added to this “jittered” pool. We do this for three different orderings in which neurons are added to the jittered pool: (i) neurons with largest trained weights w_i listed first, (ii) neurons with smallest w_i 's listed first and (iii) neurons randomly listed. The results are shown in Fig 7 (E). First, note that perturbing the spike times of neurons with large w_i 's quickly reduces performance down to chance, whereas jittering the spike times of random neurons or those with low readout weights has little to no effect on performance. The details of the observations above are likely to depend on the choice of time constants for the Tempotron's filters, but we expect the overall trends to persist for a range of these constants, based on the spike-time reliability described earlier. From this we conclude that the Tempotron learns to classify stimuli based on the reliable spikes of a few neurons, at precise moments in time.

Finally, we shuffled the spike ensembles across trials by building surrogate spike patterns in which the spike output of each neuron is taken from a randomly chosen trial from the ensemble as was originally done in [14, 15]. If the readout cells that serve as strong inputs to the Tempotron (i.e., the cells with relatively large w_i) were unreliable, this would have the effect of shifting their spike times and changing the spike counts within the test window. Numerical results indicate that shuffling has no appreciable effect on classification performance. This further suggests that repeatable spike patterns are responsible for good classification, rather than statistics like spike counts on longer time scales.

Taken together, these tests show that *despite chaos, the network response preserves a fair degree of spike reliability across trials in key neural coordinates, and a simple decoding scheme is capable of taking advantage of this reliability to accurately classify spike trains*. Note that the identity of these “key” neurons will change for different input stimuli and thus, at least in principle, many readout schemes could be trained in parallel on the same network.

Finding 2: Recurrence enables information distribution within networks and discrimination using few readouts

We have shown that a neural-like readout, the Tempotron, can use the reliable spikes embedded in a chaotic network's response in order to classify input stimuli. Up to now, this classifier had access to spikes from every neuron. A natural question is whether this complete access is necessary. In other words, how does classification performance depend on the number of inputs and outputs to the network? For example, if only a few neurons in the network receive discriminable inputs ($I_i^A(t) \neq I_i^B(t)$) and the decoder only reads out from a few (different) neurons, do network interactions distribute enough stimulus information to enable a successful classification?

In our model, the proportion of the input that is directly discriminable is controlled by the parameter ρ_{same} —the fraction of cells that receive identical inputs under $I^A(t)$ and $I^B(t)$. We adopt an enumeration that lists these cells first for ease of notation: the first N_{same} neurons receive indistinguishable inputs while the remaining N_{diff} receive independent ones for each stimulus; with $N_{same} = \rho_{same} N$ and $N_{diff} = N - N_{same}$ (see Fig 2 (B)). In addition, we introduce a second parameter, N_r , which controls the number of readout neurons providing inputs to the

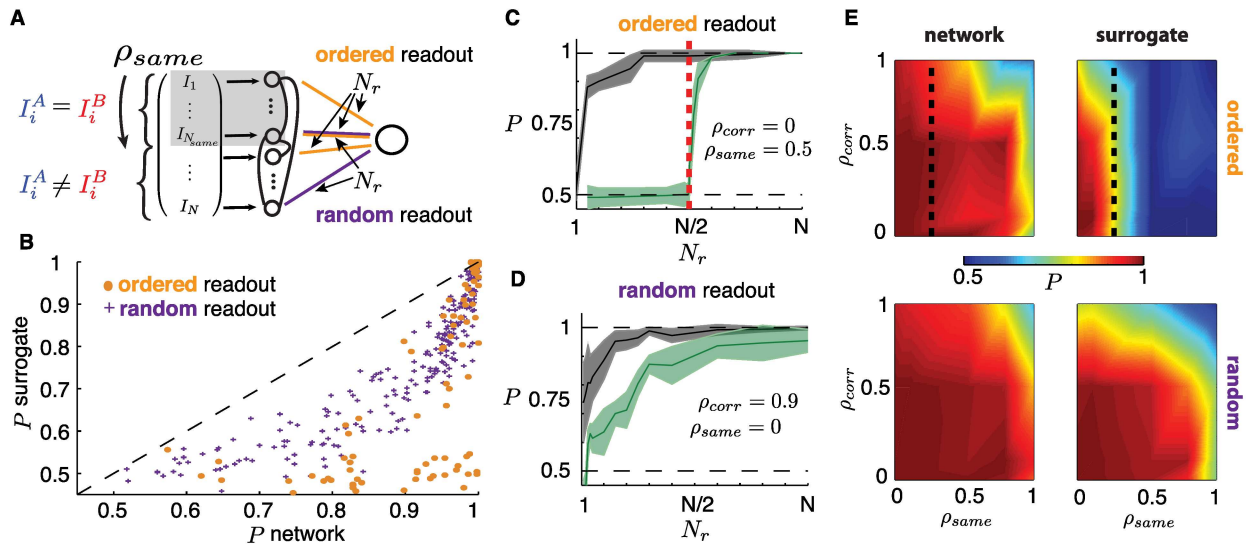


Fig 8. (A) Illustration of input similarity ρ_{same} for which the first N_{same} neurons receive identical direct inputs $I_i^A(t)$ and $I_i^B(t)$. Also depicted is partial readout dimension N_r under two types of readout orderings: *ordered* and *random*, see text for details. (B) Scatter plot of classification performance P of surrogate population vs. network for a range of parameters ρ_{corr} , ρ_{same} and N_r , for both ordered and random readouts (same data as (E)). (C) P for network (black) and surrogate population (green) as a function of N_r with ordered readouts, for $\rho_{corr} = 0$, $\rho_{same} = 0.5$. Red dashed line indicates $N_r = N_{same}$. (D) P for network (black) and surrogate population (green) as a function of N_r with random readout, for $\rho_{corr} = 0.9$, $\rho_{same} = 0$. For (C) and (D), shaded areas show one standard deviation (see sampling details below). (E) Color plot of P as a function of ρ_{same} and ρ_{corr} for $N_r = N/4$. Top row: ordered readouts (dashed line indicates $N_r = N_{same}$). Bottom row: random readouts. First column: network. Second column: surrogate population. For all P values shown, quantities were sampled over 20 repeated simulations where networks, inputs and Tempotron training were reprocessed. $N = 500$, $K = 20$.

doi:10.1371/journal.pcbi.1005258.g008

Tempotron. We select these N_r cells in two ways as depicted in Fig 8 (A): *ordered readouts* where we follow the same ordering as above (i.e. the Tempotron weights w_i are defined for indices $i = 1, \dots, N_r$) and *random readouts* where N_r neurons are selected at random throughout the network.

We investigate the classification performance P of our network with different configurations of input similarity, readout size and type. For example, if $\rho_{same} = 0.5$ (equivalently $N_{same} = N/2$), $\rho_{corr} = 0.75$ and $N_r = N/4$ with ordered readouts, this means that neurons $i = 1, \dots, N/2$ receive identical local inputs $I_i^A(t)$ and $I_i^B(t)$, the remaining neurons $i = N/2 + 1, \dots, N$ receive local inputs that are 75% correlated, and the Tempotron only reads out from neurons $i = 1, \dots, N/4$. In addition, for all parameter configurations, we also produce surrogate population spike patterns where we replace network interactions by independent Poisson-distributed spike trains as was done in Fig 5 (B), following the assumptions from earlier. The comparison is summarized in Fig 8 (B) where a scatter plot of surrogate performances is plotted against network performances for an exhaustive range of parameter choices. Importantly, the chaotic network always performs better or as well as the surrogate population. Furthermore, the only cases in which the two perform equally is when the discrimination task is easy enough to allow perfect performance.

The scatter also shows that the network has a significant advantage over the surrogate population when ordered readouts are used. To better understand this, Fig 8 (C) shows network P (black) and surrogate P (green) for $\rho_{corr} = 0$, $\rho_{same} = 0.5$ as a function of N_r for ordered readouts. Notably, the network shows very high classification performance in many cases where $N_r < N_{same}$. Here the classifier can *only* read out from neurons that receive identical inputs under the two stimuli (i.e. not directly discriminable). This means that discriminable features

of these neurons' spike patterns cannot originate from their direct inputs $I_i^A(t)$ and $I_i^B(t)$, and must instead result from network interactions that communicate activity from other neurons. Not surprisingly, the surrogate population fails completely when $N_r < N_{same}$ since network interactions are replaced by independent stochastic spike trains.

For random readouts, it can be assumed that at least one neuron selected for readout receives a directly discriminable inputs when $\rho_{same} < 1$. This makes the task easier so that network and surrogate population performances are closer when local inputs $I_i(t)$ are sufficiently different. Nevertheless, when the correlation between them grows, i.e. for larger ρ_{corr} , the network does produce significantly more discriminable responses. This is shown in Fig 8 (D) where P is plotted as a function of N_r for $\rho_{same} = 0$, $\rho_{corr} = 0.9$ and random readouts.

The general landscapes of discriminability P as a function of similarity parameters ρ_{corr} and ρ_{same} for both ordered and random readouts, are shown in Fig 8 (E) for $N_r = N/4$. Consistent with the results about the geometric properties of response ensembles shown in Fig 6 (D), we achieve near-perfect performance ($P \approx 1$) up to $\rho_{same} \sim 0.9$, despite reading out from only a quarter of the network's neurons. Crucially, we see that the phenomenon introduced in panel (C) of the same figure—that network can accurately discriminate inputs when reading out from neurons not receiving directly discriminable inputs—is robust to strong input correlations ρ_{corr} .

We conclude from these results that *recurrent interactions of the type found in our network model, despite generating chaotic instabilities, are an effective and robust way to distribute information about local inputs throughout a network so that distant readouts can be used for stimulus discrimination without specifying precise connectivity*. We expect these findings to hold generally in recurrent networks that are well-connected, in the sense that for every pair of vertices in the associated connectivity graph, there is a relatively short directed path between them. In the next section, we show how the relationship between the number of readouts N_r needed for classification and the number of distinct inputs N_{diff} depends on stimulus statistics.

Finding 3: Signal strength modulates a “noise floor” from chaos

So far, our investigation of stimulus encoding has been restricted to a single choice of the stimulus amplitude ε and mean η ($\varepsilon = 0.5$, $\eta = -0.5$). For these parameters, we showed that even highly similar stimuli can be distinguished based on the responses of chaotic networks. How does this depend on the stimulus amplitude (i.e. strength of temporal features)? When this amplitude drops, one might expect that it will eventually fall below a limit when any differences in stimuli will be obscured by variability induced by the chaotic dynamics of a network. We call this limit the chaos-induced *noise floor*. Below, we study this noise floor, and thereby establish how the discriminability of stimuli in chaotic networks depends on their statistics.

To systematically compare how stimulus statistics impact discriminability, care is needed to keep the network activity in a fairly consistent firing regime. We do this by varying parameters (ε, η) together in a way that will produce a fixed firing rate, averaged across the network. This way, we can be certain that classifiers will be trained on the same number of spikes on average, a quantity that could affect the interpretation performance P if left uncontrolled. Specifically, we vary η and ε together along a path that leaves the network-wide average E firing rate fixed at 13 Hz, as was the case for parameter values used above (see Fig 4 (B)). As illustrated in Fig 9 (A), we parameterize this path by a normalized arclength parameter x : for higher values of x , η is larger and ε is smaller (see Fig 9 (B) for illustration of input as x changes).

We begin by visualizing the relative amplitudes of external inputs $I_i(t)$ and network interactions. If the strength of stimulus fluctuations ε was so great as to simply overwhelm interactions between neurons in the network, then the role of chaos in any conclusion about stimulus

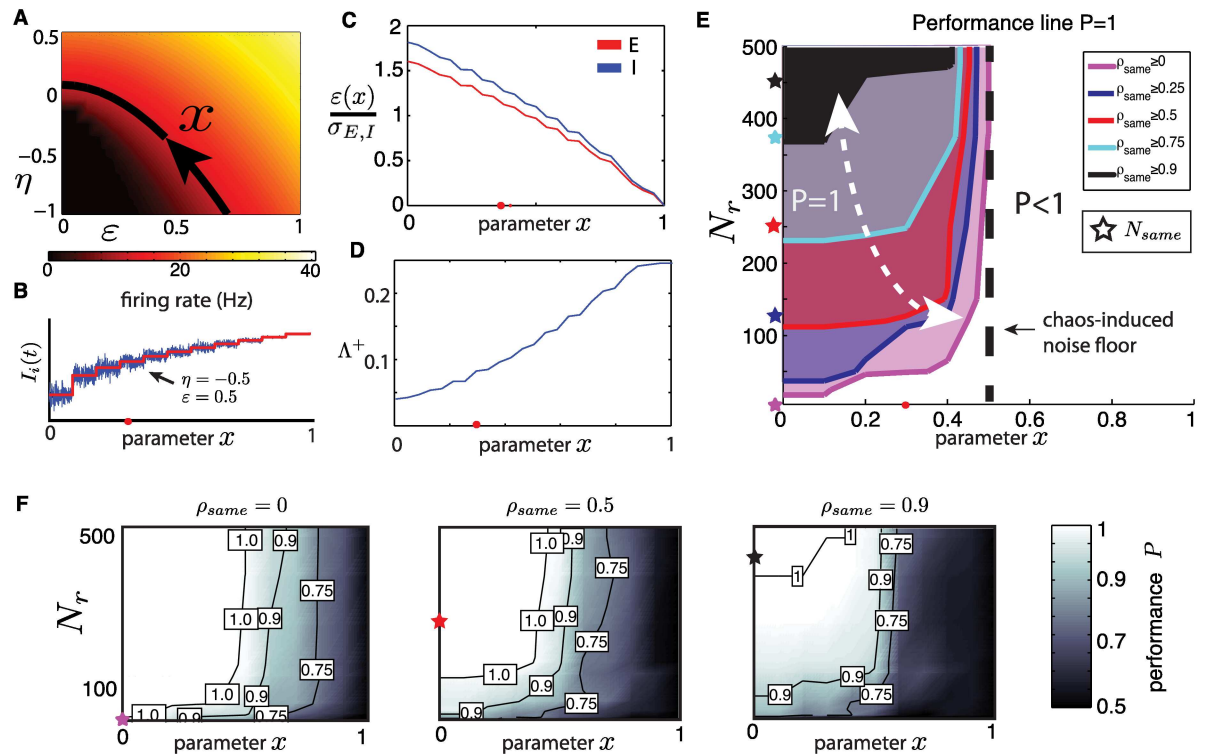


Fig 9. (A) Mean E firing rate of the network as a function of mean input η and signal amplitude ϵ . Black curve shows level set at 13 Hz, and is parametrized by the normalized arclength x in direction of arrow. Parameter set used throughout the rest of paper ($\epsilon = 0.5$, $\eta = -0.5$) corresponds to $x \approx 0.3$. (B) Illustration of input $I_i(t)$ presented to a neuron as a function of the parameter x . The red line shows the mean $\eta(x)$ while the blue line shows fluctuations of amplitude $\epsilon(x)$. For all parameters, the mean excitatory firing rate is constant. Red dot shows the benchmark regime ($\eta = -0.5$, $\epsilon = 0.5$). (C) Ratio of stimulus amplitude ϵ to that of network interaction amplitude $\sigma_{E,I}$, as derived from the mean field equations (see Eq (6)). Red dot shows the benchmark regime ($\eta = -0.5$, $\epsilon = 0.5$). (D) The fraction of positive Lyapunov exponents $\#\{\lambda_i > 0\}/N$ as a function of the parameter x . Red dot shows the benchmark regime ($\eta = -0.5$, $\epsilon = 0.5$). (E) Level curves of classification performance $P = 1$ in the (x, N_r) parameter space for several values of ρ_{same} . For parameter pairs to the left of these curves, the network achieves perfect classification on average. Stars indicate the corresponding $N_{same} = \rho_{same} N$. Red dot shows the benchmark regime ($\eta = -0.5$, $\epsilon = 0.5$). (F) Classification performance P as a function of x and readout dimension N_r for three input similarity fractions $\rho_{same} = 0, 0.5, 0.9$. Stars indicate the corresponding $N_{same} = \rho_{same} N$. For all P values shown, quantities were sampled over 20 repeated simulations where networks, inputs and Tempotron training were re-processed. Ordered readouts, $\rho_{corr} = 0$, $N = 500$, $K = 20$.

doi:10.1371/journal.pcbi.1005258.g009

encoding would be trivial: the stimulus simply overwrites any intrinsic dynamics. Fig 9 (C) shows the ratio of stimulus input amplitude ϵ to that of network interactions $\sigma_{E,I}$ computed from Eq (6). We can see that for the range of our parameter x this ratio goes from 0 to 1.8. Not surprisingly, the highest value of that ratio corresponds to the maximal value of ϵ , when $x = 0$. Even for ratios greater than 1 we argue that our network is not completely overwhelmed by inputs since for all parameters considered, our network is still chaotic with positive Lyapunov exponents (see Fig 9 (D) and text below). This means that on repeated trials, network interactions are strong enough to induce instabilities and create discrepancies between responses. However, these discrepancies will be smaller than in networks where input amplitude ϵ is weaker, impacting discriminability. This is what we investigate below. See also [14] for further investigation of this mechanism using spike-triggered averages.

Even if the stimulus input does not completely overwhelm the network, its statistics play an important role in shaping the level of discriminable features carried by network spike patterns. As we vary input parameters, two key quantities change. One is the “signal strength:” as ϵ becomes smaller, the direct impact of stimulus fluctuations on neurons subsides, and, as a

consequence, so does the magnitude of the differences between stimuli $I^A(t)$ and $I^B(t)$. The second is the “noise:” the chaotic variability in the responses to a single input signal. These two quantities are not independent, and their relationship is not *a priori* obvious. Together, they combine to create the chaos-induced noise floor described above. For stimulus parameters that fall below this noise floor, chaotic variability is too widespread to allow different stimuli to be accurately classified.

As described earlier, the variability across spike patterns from the same response ensemble depends on the dimension of the underlying chaotic attractor. There are many ways to quantify this dimensionality, here we use the number of positive Lyapunov exponents divided by the dimension of the system (see [Methods](#) or [14] for details of their computation):

$$\Lambda^+ \equiv \#\{\lambda_i > 0\}/N.$$

Intuitively, Λ^+ indicates the fraction of unstable directions in state space at any given time. As the geometric properties of our network attractors impose that those directions generally align with neural coordinates θ_i (see earlier text about spike-time reliability), Λ^+ dictates how many neurons are “unreliable” in the system at any given time [14, 15]. [Fig 9 \(D\)](#) shows Λ^+ as a function of x . Notice that the network becomes more chaotic (more positive Lyapunov exponents) as the fluctuation amplitude ε of the inputs $I_i(t)$ shrinks. This can be interpreted as weaker stimulus fluctuations giving less entrainment of neural dynamics by the inputs, and hence allowing intrinsic dynamics, the mechanism by which chaos emerges, to dominate (c.f. [31] for an example of stimulus entrainment effects on chaos in networks of rate units). Thus, as x increases, both signal and noise factors should conspire to make input stimuli less discriminable. We next quantify this effect, and study how it depends on input similarity (quantified by ρ_{same}) and readout dimension N_r .

We numerically estimate the Tempotron’s discrimination performance P along the x -parametrized curve of stimulus parameters, for a range of ρ_{same} and N_r values. We concentrate on regimes where $\rho_{corr} = 0$ and consider only ordered readouts, in order to better investigate the phenomenon of information distribution throughout the network described in the previous section. This determines regions of the parameter space (x, N_r, ρ_{same}) where perfect classification performance, $P = 1$, is achieved. The boundaries of these regions are shown in [Fig 9 \(E\)](#), where every parameter point to the left of the boundaries yields perfect classification. These boundaries have the following interpretation: for a given input parameters specified by x , one needs to readout from a number of neurons N_r greater than a given ρ_{same} -boundary to be able to achieve perfect discrimination of inputs with a similarity level given by ρ_{same} . As expected, the regions of perfect performance shrink as inputs become more similar. This means that more readout dimensions N_r are needed to discriminate more-similar inputs. Importantly, the boundaries are positioned at much lower N_r than the corresponding $N_{same} = \rho_{same} N$ for all cases, indicating that networks across a broad parameter range can classify inputs using neurons that themselves do not receive discriminable inputs (i.e. $I_i^A(t) = I_i^B(t)$) as demonstrated for our benchmark parameter set in the previous section.

Moreover, there is a critical region of stimulus statistics ($x \sim 0.5$) where all classification boundaries aggregate for high N_r . This represents the chaotic noise floor for x , beyond which inputs cannot be perfectly discriminated. [Fig 9 \(F\)](#) illustrates this noise floor in more detail, by showing contour plots of classification performance P in (x, N_r) -space for three values of ρ_{same} . It shows that while P eventually drops to chance (0.5) as $x \rightarrow 1$ —an expected behavior since $I^A(t)$ and $I^B(t)$ become indistinguishable when $\varepsilon = 0$ —this transition becomes sharper for higher input similarity ρ_{same} . This suggests that when two stimuli have many identical components (high ρ_{same}), the network can either classify them very well or not at all, depending on

the stimulus amplitude. When inputs are significantly dissimilar, this transition is more gradual. Importantly, for large enough readout size N_r , the noise floor is almost identical for all input similarities, indicating that for this network, perfect classification becomes impossible for any input similarity once Λ^+ reaches about 11% of N . Lastly, we note that the noise floor boundary is roughly aligned with input parameters yielding a ratio of stimulus-to-network fluctuation amplitude $\varepsilon/\sigma_E \sim 0.9$. This means that input stimuli need to be comparable or stronger than other network interactions to achieve discriminability.

In summary, for the chaotic networks at hand, *perfect stimulus classification can be achieved for a wide range of stimulus statistics and similarities, and classification can be achieved using spikes from relatively few readout neurons. However, once the stimulus amplitude falls below a chaos-induced noise floor, classification performance degrades rapidly.*

Discussion

Summary

Sparse, strongly connected recurrent neural networks used to model cortical activity in the brain often produce a *balanced state*, leading to chaotic dynamics [2]. We showed that despite asynchronous dynamics—where neurons have temporally de-correlated activity—chaos-induced variability in driven networks is not easily approximated by simple stochastic processes (e.g. using a mean-field approach). Instead, we found complex statistical dependencies across network spike outputs conditioned on a given input, i.e. across trials. We studied how this chaos—viewed as an intrinsic source of variability—impacts the capacity of recurrent networks to accurately encode temporal stimuli. With detailed numerical simulations grounded in the theoretical literature, we studied how similar stimuli—modelled by multi-dimensional frozen white noise inputs—can be decoded from chaotic network responses, both at the level of the network state space and output spike trains.

Two factors influence the ability of a decoder to successfully classify stimuli based on network outputs. The first is the strength of the chaos-induced “noise”: the trial-to-trial variability of evoked patterns due to chaos. The second is the “signal”: the sensitivity of evoked patterns to the choice of input. Our analysis of these separate factors leads to three main points: (1) Chaos in recurrent spiking networks does not, in and of itself, preclude the accurate encoding of temporal stimuli; simple decoders read out these stimuli based on reliable multi-spike patterns that chaotic networks produce via low-dimensional attractors. (2) Recurrent connectivity distributes stimulus information throughout chaotic networks, enabling high-dimensional stimuli to be classified with low-dimensional readouts. (3) Stimulus statistics (i.e., the amplitude of stimulus fluctuations relative to their mean) modulate the number of readout neurons necessary to successfully classify them: as their amplitude decreases, more neurons are required to discriminate stimuli, until a “noise floor” is reached where discrimination is no longer possible.

Biological implications of encoding and computing in chaotic networks

Chaotic dynamics appear as an emergent property of recurrent connectivity between neurons [2, 6, 12, 31, 58, 59] that would be otherwise very stable and reliable (see [51, 52] and [42, 60] for reliability of single neurons). It is conceivable that such chaos is a significant contributor to experimentally observed variability as well (e.g. [17, 23, 61]). In this way, chaos amplifies and adds to other stochastic noise sources in biological networks. Moreover, we reiterate that the type of “intermittent spiking reliability” that is produced by chaotic dynamics is also observed experimentally *in vivo* [33, 34]. Although not direct evidence that chaotic dynamics is indeed

present in recurrent neuronal networks in the brain, these observations are consistent with this hypothesis.

Beyond contributing variability and noise, there is a substantial literature addressing the potential advantages of chaotic dynamics for encoding and memory [62–64]. In many cases, chaotic networks act as “reservoirs” and synaptic connections are trained to use their activity to perform a given task. Here, we take a more pragmatic perspective, studying how chaotic networks work as “channels” that receive inputs and produce spike outputs that carry usable information. Furthermore, we show that the same recurrent connectivity that produces chaos also serves to distribute stimulus information throughout the network: discriminability is maintained even if a decoder only has access to small subpopulations of neurons, and even when the inputs to be discriminated do not directly drive the subpopulations. As such, recurrence may serve to simplify the process of reading out stimuli from large populations, eliminating the need for precise wiring—with the resulting chaotic dynamics being a manageable by-product. While this type of stimulus “spread” can also occur in multilayer feed-forward networks with fan-out between layers, a recurrent architecture does the same operation locally, without requiring that decoders be located downstream. We speculate that this mechanism may also be relevant for contextual coding [65, 66], where the response of some neurons to a fixed local input changes if a secondary contextual input to others differs. This role for recurrence complements many other functions that it may serve in neuronal computation (e.g., maintaining working memory, enabling winner-take-all computation, sharpening tuning curves, etc.).

Finally, we argue that the encoding mechanism based on spike patterns we outline in this paper enhances earlier balanced network encoding mechanisms. Classic results point to important properties of balanced population-averaged activity: its response to global external inputs is both rapid—much faster than single neuron time constants—and linear [5]. If the inputs to our network evoke different population firing rates, then population averages carry the necessary information for discrimination. In contrast, when two inputs have similar statistics and differ only in the fine temporal patterns they carry, we show that the same network can rely on spike-time based mechanisms to classify them. It is unclear if such dynamics are present in cortical circuits and if so, in which regime they typically operate. However, there is evidence of different activity states a given cortical network can take (e.g. up and down states) depending on various contextual factors [67]. In light of the results we outlined, it is conceivable that cortical networks encode different aspect of inputs depending on these input features. Under this assumption, the emergent nature of chaos in recurrent networks may act as a natural mechanism to implement adaptive coding schemes, without any changes required to the network or neurons themselves.

Future work

The results presented in this manuscript address a specific class of models, albeit one that is fairly prototypical. Further studies should focus on the effect of single neuron dynamics and connectivity statistics on stimulus encoding. Specifically, while we briefly investigated the effect of modest variations of network size N and connectivity in-degree K , a more substantial study is needed to understand network behavior in large-size limits. Moreover, beyond the amplitude effects studied here, the correlation of stimulus inputs across neurons can also impact the resulting chaotic network responses (see Discussion of [14]). At the same time, these input correlations effectively diminish the dimensionality of the stimulus by introducing redundancies. An interesting area of future work is to better understand the relationship

between input and output dimension with respect to stimulus coding in recurrent, spiking networks.

In experiments, is not an easy task to test whether or not a particular neural circuit is chaotic. Indeed, even for a dynamical system that does not receive input drive, and for which one can observe all degrees of freedom, it is still a hard problem to attribute variability to stochastic or deterministic (chaotic) mechanisms (see e.g. [68, 69]). Therefore, the problem of experimentally verifying the nature of variability in neural circuits found in the brain is not a simple one. Nevertheless, we note that some *in vivo* experiments show stimulus-evoked spikes that appear to have the type of intermittent variability we described in this article [17, 33]. This invites future work to make closer connections between mechanistic models of chaotic dynamics and neural recordings.

Acknowledgments

The authors thank Fred Wolf and Lai-Sang Young for very helpful discussions.

Author Contributions

Conceptualization: GL JPT KKL ESB.

Formal analysis: GL.

Funding acquisition: GL JPT KKL ESB.

Methodology: GL JPT KKL ESB.

Software: GL.

Writing – original draft: GL.

Writing – review & editing: GL JPT KKL ESB.

References

1. Shadlen MN, Newsome WT. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci*. 1998; 18:3870–3896. PMID: [9570816](#)
2. van Vreeswijk C, Sompolinsky H. Chaotic balanced state in a model of cortical circuits. *Neural Comput*. 1998; 10:1321–1371. doi: [10.1162/089976698300017214](#) PMID: [9698348](#)
3. Renart A, de la Rocha J, et al. The Asynchronous State in Cortical Circuits. *Science*. 2010; 327:587–590. doi: [10.1126/science.1179850](#) PMID: [20110507](#)
4. Softky WR, Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *The Journal of neuroscience: the official journal of the Society for Neuroscience*. 1993; 13(1):334–350. PMID: [8423479](#)
5. Sompolinsky H, Crisanti A, Sommers H. Chaos in random neural networks. *Physical Review Letters*. 1988; 61(3):259–262. doi: [10.1103/PhysRevLett.61.259](#) PMID: [10039285](#)
6. Monteforte M, Wolf F. Dynamical Entropy Production in Spiking Neuron Networks in the Balanced State. *Phys Rev Lett*. 2010; 105(26). doi: [10.1103/PhysRevLett.105.268104](#) PMID: [21231716](#)
7. Latham PE, Richmond BJ, Nelson PG, Nirenberg S. Intrinsic Dynamics in Neuronal Networks. I. Theory. *J Neurophysiol*. 2000; 83:808–827. PMID: [10669496](#)
8. Luccioli S, Olmi S, Politi A, Torcini A. Collective dynamics in sparse networks. *Phys Rev Lett*. 2012; 109(13). doi: [10.1103/PhysRevLett.109.138103](#) PMID: [23030123](#)
9. Stanley G. Reading and writing the neural code. *Nature Neuroscience*. 2013;. doi: [10.1038/nn.3330](#)
10. Faisal AA, Selen LPJ, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci*. 2008; 9(4):292–303. doi: [10.1038/nrn2258](#) PMID: [18319728](#)
11. Thivierge J, Cisek P. Nonperiodic synchronization in heterogeneous networks of spiking neurons. *The Journal of Neuroscience*. 2008;. doi: [10.1523/JNEUROSCI.0870-08.2008](#) PMID: [18685022](#)

12. London M, Roth A, Beeren L, Häusser M, Latham PE. Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature*. 2010; 466(7302):123–127. doi: [10.1038/nature09086](https://doi.org/10.1038/nature09086) PMID: [20596024](https://pubmed.ncbi.nlm.nih.gov/20596024/)
13. Rosenbaum R, Doiron B. Balanced Networks of Spiking Neurons with Spatially Dependent Recurrent Connections. *Physical Review X*. 2014;. doi: [10.1103/PhysRevX.4.021039](https://doi.org/10.1103/PhysRevX.4.021039)
14. Lajoie G, Lin KK, Shea-Brown E. Chaos and reliability in balanced spiking networks with temporal drive. *Phys Rev E*. 2013; 87(5):052901. doi: [10.1103/PhysRevE.87.052901](https://doi.org/10.1103/PhysRevE.87.052901) PMID: [23767592](https://pubmed.ncbi.nlm.nih.gov/23767592/)
15. Lajoie G, Thivierge JP, Shea-Brown E. Structured chaos shapes spike-response noise entropy in balanced neural networks. *Front Comput Neurosci*. 2014; 8:1–10. doi: [10.3389/fncom.2014.00123](https://doi.org/10.3389/fncom.2014.00123) PMID: [25324772](https://pubmed.ncbi.nlm.nih.gov/25324772/)
16. Brunel N. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*. 2000; 8(3):183–208. PMID: [10809012](https://pubmed.ncbi.nlm.nih.gov/10809012/)
17. Kara P, Reinagel P, Reid RC. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron*. 2000; 27(3):635–646. doi: [10.1016/S0896-6273\(00\)00072-6](https://doi.org/10.1016/S0896-6273(00)00072-6) PMID: [11055444](https://pubmed.ncbi.nlm.nih.gov/11055444/)
18. Lu T, Liang L, Wang X. Temporal and rate representations of time-varying signals in the auditory cortex of awake primates. *Nature Neuroscience*. 2001;. doi: [10.1038/nn737](https://doi.org/10.1038/nn737) PMID: [11593234](https://pubmed.ncbi.nlm.nih.gov/11593234/)
19. Bair W, Zohary E, Newsome W. Correlated firing in macaque visual area MT: time scales and relationship to behavior. *The Journal of Neuroscience*. 2001;. PMID: [11222658](https://pubmed.ncbi.nlm.nih.gov/11222658/)
20. Lottem E, Gugig E, Azouz R. Parallel Coding Schemes of Whisker Velocity in the Rat's Somatosensory System. *Journal of Neurophysiology*. 2014;. doi: [10.1152/jn.00485.2014](https://doi.org/10.1152/jn.00485.2014) PMID: [25552637](https://pubmed.ncbi.nlm.nih.gov/25552637/)
21. Trussell L. Synaptic mechanisms for coding timing in auditory neurons. *Annual review of physiology*. 1999;. doi: [10.1146/annurev.physiol.61.1.477](https://doi.org/10.1146/annurev.physiol.61.1.477) PMID: [10099698](https://pubmed.ncbi.nlm.nih.gov/10099698/)
22. Scaglione A, Moxon K, Aguilar J. Trial-to-trial variability in the responses of neurons carries information about stimulus location in the rat whisker thalamus. *Proceedings of the ...* 2011;. doi: [10.1073/pnas.1103168108](https://doi.org/10.1073/pnas.1103168108) PMID: [21873241](https://pubmed.ncbi.nlm.nih.gov/21873241/)
23. Tiesinga P, Fellous JM, Sejnowski TJ. Regulation of spike timing in visual cortical circuits. *Nat Rev Neurosci*. 2008; 9(2):97–107. doi: [10.1038/nrn2315](https://doi.org/10.1038/nrn2315) PMID: [18200026](https://pubmed.ncbi.nlm.nih.gov/18200026/)
24. Dan Y, Alonso J, Usrey W, Reid R. Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nature Neuroscience*. 1998;. doi: [10.1038/2217](https://doi.org/10.1038/2217) PMID: [10196548](https://pubmed.ncbi.nlm.nih.gov/10196548/)
25. Faisal A, Selen L, Wolpert D. Noise in the nervous system. *Nature Reviews Neuroscience*. 2008;. doi: [10.1038/nrn2258](https://doi.org/10.1038/nrn2258) PMID: [18319728](https://pubmed.ncbi.nlm.nih.gov/18319728/)
26. Bryant H, Segundo J. Spike initiation by transmembrane current: a white-noise analysis. *J Physiology*. 1976; 260:279–314. doi: [10.1113/jphysiol.1976.sp011516](https://doi.org/10.1113/jphysiol.1976.sp011516) PMID: [978519](https://pubmed.ncbi.nlm.nih.gov/978519/)
27. Mainen ZF, Sejnowski TJ. Reliability of spike timing in neocortical neurons. *Science*. 1995; 268(5216):1503–1506. doi: [10.1126/science.7770778](https://doi.org/10.1126/science.7770778) PMID: [7770778](https://pubmed.ncbi.nlm.nih.gov/7770778/)
28. Niell CM, Stryker MP. Highly Selective Receptive Fields in Mouse Visual Cortex. *Journal of Neuroscience*. 2008; 28(30):7520–7536. doi: [10.1523/JNEUROSCI.0623-08.2008](https://doi.org/10.1523/JNEUROSCI.0623-08.2008) PMID: [18650330](https://pubmed.ncbi.nlm.nih.gov/18650330/)
29. McCormick DA, Mcginley MJ, Salkoff DB. Brain state dependent activity in the cortex and thalamus. *Current Opinion in Neurobiology*. 2015; 31:133–140. doi: [10.1016/j.conb.2014.10.003](https://doi.org/10.1016/j.conb.2014.10.003) PMID: [25460069](https://pubmed.ncbi.nlm.nih.gov/25460069/)
30. Marre O, Yger P, Davison AP, Fregnac Y. Reliable Recall of Spontaneous Activity Patterns in Cortical Networks. *J of Neurosci*. 2009; 29(46):14596–14606. doi: [10.1523/JNEUROSCI.0753-09.2009](https://doi.org/10.1523/JNEUROSCI.0753-09.2009) PMID: [19923292](https://pubmed.ncbi.nlm.nih.gov/19923292/)
31. Rajan K, Abbott L, Sompolinsky H, Proment D. Stimulus-dependent suppression of chaos in recurrent neural networks. *Phys Rev E*. 2010; 82. doi: [10.1103/PhysRevE.82.011903](https://doi.org/10.1103/PhysRevE.82.011903) PMID: [20866644](https://pubmed.ncbi.nlm.nih.gov/20866644/)
32. Molgedey L, Schuchhardt J, Schuster HG. Suppressing chaos in neural networks by noise. *Phys Rev Lett*. 1992; 69. doi: [10.1103/PhysRevLett.69.3717](https://doi.org/10.1103/PhysRevLett.69.3717) PMID: [10046895](https://pubmed.ncbi.nlm.nih.gov/10046895/)
33. Reinagel P, Reid RC. Temporal coding of visual information in the thalamus. *J Neurosci*. 2000; 20(14):5392–5400. PMID: [10884324](https://pubmed.ncbi.nlm.nih.gov/10884324/)
34. Fellous JM, Tiesinga PHE, Thomas PJ, Sejnowski TJ. Discovering Spike Patterns in Neuronal Responses. *The Journal of neuroscience: the official journal of the Society for Neuroscience*. 2004; 24(12):2989–3001. doi: [10.1523/JNEUROSCI.4649-03.2004](https://doi.org/10.1523/JNEUROSCI.4649-03.2004) PMID: [15044538](https://pubmed.ncbi.nlm.nih.gov/15044538/)
35. Gütiğ R, Sompolinsky H. The tempotron: a neuron that learns spike timing–based decisions. *Nature neuroscience*. 2006; 9(3):420–428. doi: [10.1038/nn1643](https://doi.org/10.1038/nn1643) PMID: [16474393](https://pubmed.ncbi.nlm.nih.gov/16474393/)
36. Strata P, Harvey R. Dale's principle. *Brain Research Bulletin*. 1999; 50(5-6):349–350. PMID: [10643431](https://pubmed.ncbi.nlm.nih.gov/10643431/)

37. van Vreeswijk C, Sompolinsky H. Chaos in Neuronal Networks with Balanced Excitatory and Inhibitory Activity. *Science* (New York, NY). 1996; 274(5293):1724–1726. doi: [10.1126/science.274.5293.1724](https://doi.org/10.1126/science.274.5293.1724) PMID: [8939866](https://pubmed.ncbi.nlm.nih.gov/8939866/)
38. Ermentrout B. Type I membranes, phase resetting curves, and synchrony. *Neural Comput*. 1996; 8(5):979–1001. doi: [10.1162/neco.1996.8.5.979](https://doi.org/10.1162/neco.1996.8.5.979) PMID: [8697231](https://pubmed.ncbi.nlm.nih.gov/8697231/)
39. Ermentrout B, Terman D. *Mathematical foundations of neuroscience*. vol. 35. Springer Iam, editor. Springer, Interdisciplinary applied mathematics; 2010. doi: [10.1007/978-0-387-87708-2](https://doi.org/10.1007/978-0-387-87708-2)
40. Kloeden PE, Platen E. *Numerical solution of stochastic differential equations*. vol. 23 of *Applications of Mathematics* (New York). Springer-Verlag, Berlin; 1992. Available from: <http://dx.doi.org/10.1007/978-3-662-12616-5>. doi: [10.1103/PhysRevE.80.026704](https://doi.org/10.1103/PhysRevE.80.026704) PMID: [19792278](https://pubmed.ncbi.nlm.nih.gov/19792278/)
41. Evans LC. *An introduction to stochastic differential equations*. vol. 82. American Mathematical Soc.; 2012. doi: [10.1090/mbk/082](https://doi.org/10.1090/mbk/082)
42. Lindner B, Longtin A, Bulsara A. Analytic expressions for rate and CV of a type I neuron driven by white gaussian noise. *Neural Comput*. 2003; 15(8):1761–1788. doi: [10.1162/08997660360675035](https://doi.org/10.1162/08997660360675035) PMID: [14511512](https://pubmed.ncbi.nlm.nih.gov/14511512/)
43. Rieke F, Warland D, de Ruyter van Steveninck R, Bialek W. *Spikes: Exploring the Neural Code*. Cambridge M, editor. Cambridge, MA; 1996.
44. Lin KK. Stimulus-Response Reliability of Biological Networks. In: *Nonautonomous Dynamical Systems in the Life Sciences*. Cham: Springer International Publishing; 2013. p. 135–161.
45. Cortes C, Vapnik V. Support-vector networks. *Machine learning*. 1995; 20(3):273–297. doi: [10.1007/BF00994018](https://doi.org/10.1007/BF00994018)
46. Ambard M, Rotter S. Support vector machines for spike pattern classification with a leaky integrate-and-fire neuron. *Front Comput Neurosci*. 2012; 6:1–13. doi: [10.3389/fncom.2012.00078](https://doi.org/10.3389/fncom.2012.00078) PMID: [23181017](https://pubmed.ncbi.nlm.nih.gov/23181017/)
47. Grabska-Barwińska A, Latham PE. How well do mean field theories of spiking quadratic-integrate-and-fire networks work in realistic parameter regimes? *Journal of Computational Neuroscience*. 2014; 36(3):469–481. doi: [10.1007/s10827-013-0481-5](https://doi.org/10.1007/s10827-013-0481-5) PMID: [24091644](https://pubmed.ncbi.nlm.nih.gov/24091644/)
48. Brunel N, Latham PE. Firing rate of the noisy quadratic integrate-and-fire neuron. *Neural Computation*. 2003; 15(10):2281–2306. doi: [10.1162/089976603322362365](https://doi.org/10.1162/089976603322362365) PMID: [14511522](https://pubmed.ncbi.nlm.nih.gov/14511522/)
49. Gutkin BS, Ermentrout GB. Dynamics of Membrane Excitability Determine Interspike Interval Variability: A Link Between Spike Generation Mechanisms and Cortical Spike Train Statistics. *Neural Computation*. 1998; 10(5):1047–1065. doi: [10.1162/089976698300017331](https://doi.org/10.1162/089976698300017331) PMID: [9654767](https://pubmed.ncbi.nlm.nih.gov/9654767/)
50. Lindner B, Longtin A, Bulsara A. Analytic expressions for rate and CV of a type I neuron driven by white gaussian noise. *Neural Computation*. 2003; 15(8):1761–1788. doi: [10.1162/08997660360675035](https://doi.org/10.1162/08997660360675035) PMID: [14511512](https://pubmed.ncbi.nlm.nih.gov/14511512/)
51. Mainen Z, Sejnowski T. Reliability of spike timing in neocortical neurons. *Science* (New York, NY). 1995; 268(5216):1503–1506. doi: [10.1126/science.7770778](https://doi.org/10.1126/science.7770778) PMID: [7770778](https://pubmed.ncbi.nlm.nih.gov/7770778/)
52. Bryant H, Segundo J. Spike initiation by transmembrane current: a white-noise analysis. *Journal of Physiology*. 1976; 260:279–314. doi: [10.1113/jphysiol.1976.sp011516](https://doi.org/10.1113/jphysiol.1976.sp011516) PMID: [978519](https://pubmed.ncbi.nlm.nih.gov/978519/)
53. Fellous J, Whitmer D, Tiesinga P, Sejnowski T. A new correlation-based measure of spike timing reliability. *Neurocomputing*. 2003; 52-54. doi: [10.1016/S0925-2312\(02\)00838-X](https://doi.org/10.1016/S0925-2312(02)00838-X) PMID: [20740049](https://pubmed.ncbi.nlm.nih.gov/20740049/)
54. Lin KK, Shea-Brown E, Young LS. Spike-time reliability of layered neural oscillator networks. *J Comput Neuro*. 2009; 27(1):135–160. doi: [10.1007/s10827-008-0133-3](https://doi.org/10.1007/s10827-008-0133-3) PMID: [19156509](https://pubmed.ncbi.nlm.nih.gov/19156509/)
55. Namenson A, Ott E, Antonsen TM. Fractal dimension fluctuations for snapshot attractors of random maps. *Physical Review E Statistical, Nonlinear, and Soft Matter Physics*. 1996; 53(3):2287–2291. doi: [10.1103/PhysRevE.53.2287](https://doi.org/10.1103/PhysRevE.53.2287) PMID: [9964511](https://pubmed.ncbi.nlm.nih.gov/9964511/)
56. Kifer Y. *Ergodic theory of random transformations*. Birkhauser B, editor. Birkhauser, Boston; 1986. doi: [10.2307/2288883](https://doi.org/10.2307/2288883)
57. Arnold L. *Random Dynamical Systems*; 2003. doi: [10.1007/978-3-662-12878-7](https://doi.org/10.1007/978-3-662-12878-7)
58. Banerjee A. On the sensitive dependence on initial conditions of the dynamics of networks of spiking neurons. *J Comp Neuro*. 2006; 20(3):321–348. doi: [10.1007/s10827-006-7188-9](https://doi.org/10.1007/s10827-006-7188-9) PMID: [16683210](https://pubmed.ncbi.nlm.nih.gov/16683210/)
59. Banerjee A, Seriès P, Pouget A. Dynamical constraints on using precise spike timing to compute in recurrent cortical networks. *Neural Comput*. 2008; 20(4):974–993. doi: [10.1162/neco.2008.05-06-206](https://doi.org/10.1162/neco.2008.05-06-206) PMID: [18085984](https://pubmed.ncbi.nlm.nih.gov/18085984/)
60. Ritt J. Evaluation of entrainment of a nonlinear neural oscillator to white noise. *Phys Rev E*. 2003; 68(4). doi: [10.1103/PhysRevE.68.041915](https://doi.org/10.1103/PhysRevE.68.041915) PMID: [14682981](https://pubmed.ncbi.nlm.nih.gov/14682981/)

61. de Ruyter van Steveninck RR, Lewen GD, Strong SP, Koberle R, Bialek W. Reproducibility and variability in neural spike trains. *Science*. 1997; 275(5307):1805. doi: [10.1126/science.275.5307.1805](https://doi.org/10.1126/science.275.5307.1805) PMID: [9065407](https://pubmed.ncbi.nlm.nih.gov/9065407/)
62. Buonomano D, Maass W. State-dependent computations: spatiotemporal processing in cortical networks. *Nature Rev Neurosci*. 2009; 10:113–125. doi: [10.1038/nrn2558](https://doi.org/10.1038/nrn2558) PMID: [19145235](https://pubmed.ncbi.nlm.nih.gov/19145235/)
63. Sussillo D, Abbott LF. Generating Coherent Patterns of Activity from Chaotic Neural Networks. *Neuron*. 2009; 63(4):544–557. doi: [10.1016/j.neuron.2009.07.018](https://doi.org/10.1016/j.neuron.2009.07.018) PMID: [19709635](https://pubmed.ncbi.nlm.nih.gov/19709635/)
64. Laje R, Buonomano D. Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat Neurosci*. 2013; 16(7). doi: [10.1038/nn.3405](https://doi.org/10.1038/nn.3405) PMID: [23708144](https://pubmed.ncbi.nlm.nih.gov/23708144/)
65. Schwartz O, Hsu A, Dayan P. Space and time in visual context. *Nature Reviews Neuroscience*. 2007;. doi: [10.1038/nrn2155](https://doi.org/10.1038/nrn2155) PMID: [17585305](https://pubmed.ncbi.nlm.nih.gov/17585305/)
66. Mante V, Sussillo D, Shenoy KV, Newsome WT. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*. 2013; 503(7474):78. doi: [10.1038/nature12742](https://doi.org/10.1038/nature12742) PMID: [24201281](https://pubmed.ncbi.nlm.nih.gov/24201281/)
67. Mante V, Sussillo D, Shenoy KV, Newsome WT. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*. 2013; 503(7474):78. doi: [10.1038/nature12742](https://doi.org/10.1038/nature12742) PMID: [24201281](https://pubmed.ncbi.nlm.nih.gov/24201281/)
68. Brock WA. Distinguishing random and deterministic systems: Abridged version. *Journal of Economic Theory*. 1986;. doi: [10.1016/0022-0531\(86\)90014-1](https://doi.org/10.1016/0022-0531(86)90014-1)
69. Sugihara G, Mayf R. Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series. *Nature*. 1990; 344(6268):734–741. doi: [10.1038/344734a0](https://doi.org/10.1038/344734a0) PMID: [2330029](https://pubmed.ncbi.nlm.nih.gov/2330029/)