

RESEARCH ARTICLE

Approximate Inference for Time-Varying Interactions and Macroscopic Dynamics of Neural Populations

Christian Donner^{1,2,3}, Klaus Obermayer^{1,2}, Hideaki Shimazaki^{4*}

1 Bernstein Center for Computational Neuroscience, Berlin, Germany, **2** Neural Information Processing Group, Department of Electrical Engineering and Computer Science, Technische Universität Berlin, Berlin, Germany, **3** Group for Methods of Artificial Intelligence, Department of Electrical Engineering and Computer Science, Technische Universität Berlin, Berlin, Germany, **4** RIKEN Brain Science Institute, Wako-shi, Saitama, Japan

* shimazaki@brain.riken.jp



OPEN ACCESS

Citation: Donner C, Obermayer K, Shimazaki H (2017) Approximate Inference for Time-Varying Interactions and Macroscopic Dynamics of Neural Populations. *PLoS Comput Biol* 13(1): e1005309. doi:10.1371/journal.pcbi.1005309

Editor: Matthias Bethge, University of Tübingen and Max Planck Institute for Biological Cybernetics, GERMANY

Received: July 29, 2016

Accepted: December 12, 2016

Published: January 17, 2017

Copyright: © 2017 Donner 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: Data are taken from Synder et. al, *Nature Neurosci.*, Vol 18, pp 736-743 (doi:10.1038/nn.3979) with permissions of the authors, who may be contacted at smithma@pitt.edu.

Funding: This work was in part supported by the Deutsche Forschungsgemeinschaft GRK1589/2 (CD and KO) <http://www.dfg.de/en/>. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Abstract

The models in statistical physics such as an Ising model offer a convenient way to characterize stationary activity of neural populations. Such stationary activity of neurons may be expected for recordings from *in vitro* slices or anesthetized animals. However, modeling activity of cortical circuitries of awake animals has been more challenging because both spike-rates and interactions can change according to sensory stimulation, behavior, or an internal state of the brain. Previous approaches modeling the dynamics of neural interactions suffer from computational cost; therefore, its application was limited to only a dozen neurons. Here by introducing multiple analytic approximation methods to a state-space model of neural population activity, we make it possible to estimate dynamic pairwise interactions of up to 60 neurons. More specifically, we applied the pseudolikelihood approximation to the state-space model, and combined it with the Bethe or TAP mean-field approximation to make the sequential Bayesian estimation of the model parameters possible. The large-scale analysis allows us to investigate dynamics of macroscopic properties of neural circuitries underlying stimulus processing and behavior. We show that the model accurately estimates dynamics of network properties such as sparseness, entropy, and heat capacity by simulated data, and demonstrate utilities of these measures by analyzing activity of monkey V4 neurons as well as a simulated balanced network of spiking neurons.

Author Summary

Simultaneous analysis of large-scale neural populations is necessary to understand coding principles of neurons because they concertedly process information. Methods of thermodynamics and statistical mechanics are useful to understand collective phenomena of the interacting elements, and they have been successfully used to understand diverse activity of neurons. However, most analysis methods assume stationary data, in which activity rates of neurons and their correlations are constant over time. This assumption is easily

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

violated in the data recorded from awake animals. Neural correlations likely organize dynamically during behavior and cognition, and this may be independent from the modulated activity rates of individual neurons. Recently several methods were proposed to simultaneously estimate dynamics of neural interactions. However, these methods are applicable to up to about 10 neurons. Here by combining multiple analytic approximation methods, we made it possible to estimate time-varying interactions of much larger neural populations. The method allows us to trace dynamic macroscopic properties of neural circuitries such as sparseness, entropy, and sensitivity. Using these statistics, researchers can now quantify to what extent neurons are correlated or de-correlated, and test if neural systems are susceptible within a specific behavioral period.

Introduction

Activity patterns of neuronal populations are constrained by biological mechanisms such as biophysical properties of each neuron (e.g., synaptic integration and spike generation [1, 2]) and their anatomical connections [3]. The characteristic correlations among neurons imposed by the biological mechanisms interplay with statistics of sensory inputs, and influence how the sensory information is represented in the population activity [4–6]. Thus accurate assessment of the neural correlations in ongoing and evoked activities is a key to understand the underlying biological mechanisms and their coding principles.

The number of possible activity patterns increases combinatorially with the number of neurons analyzed. The maximum entropy (ME) principle and derived ME models—known as the pairwise ME model or the Ising model—have been used to explain neural population activities using fewer activity features such as event rates or correlations between pairs of neurons [7, 8]. This approach has been employed to explain not only the activity of neuronal networks but also other types of biological networks [9–11]. For large networks, however, exact inference of these models becomes computationally infeasible. Thus researchers have employed approximation methods [12–18]. While they successfully extended the number of neurons that could be analyzed, it was pointed out that the pairwise ME model might fail to explain large neural populations because the effect of higher-order interactions may become prominent [19–21]. Another fundamental problem of the conventional ME models is that these models assume temporarily constant spike rates for individual neurons. The assumption of stationary spike rates is invalid, e.g., when *in vivo* activity is recorded while an animal performs a behavioral task. Ignoring such dynamics might result in erroneous model estimates and misleading interpretations on their correlations [22–26]. Moreover neural correlations themselves likely organize dynamically during behavior and cognition, which can be independent from changes in the spike rates of individual neurons [27–29]. The time-dependence of neural activity may be explained by including stimulus signals in the model, e.g., for analyses of early sensory cells [30]. However, the approach may become impractical when analyzing neurons in higher brain areas in which receptive fields of neurons are not easily characterized. Thus it remains to be examined how much the pairwise ME model can explain the data if the inappropriate stationary assumption is removed.

The state-space analysis [31] offers a general framework to model time-series data as observations driven by an unobserved latent state process. The underlying state changes are uncovered by a sequential estimation method from the noisy measurements. While observations of neuronal activity are often characterized by point events (spikes), a series of studies have established the nonlinear recursive Bayesian estimation of the underlying state that drives the event

activity [32–34]. The method successfully estimated an animal’s position from population activity of hippocampal place cells [32], or estimate arm trajectories from neurons in the monkey motor cortex [35, 36]. Recently, this framework has been extended to the analysis of population activity [37–39]. In addition to the point estimates of interaction parameters suggested by earlier studies [40–42], the state-space analysis provides credible intervals of those estimates through the recursive Bayesian fitting algorithm.

Nevertheless, as previously mentioned, the state-space model of a neural population was restricted by its computational cost. Therefore, it could be utilized to analyze only small populations ($N \leq 15$). Recent advances in electrophysiological and optical recording techniques from a large number of neurons *in vivo* under free moving or virtual reality settings challenge these analysis methods. Thus the challenge is to make it possible to fit the exponentially complex state-space model to such large-scale data. For this goal, we need to incorporate approximation methods into the sequential Bayesian algorithm. More specifically, we need good approximations of mean and variance of the model parameters required in the approximate Bayesian scheme. These approximation methods must be analytical to avoid impractical computation time. By doing so we will be able to directly estimate all time-varying interactions of a large neural population. Such a model will serve as benchmark for alternative unsupervised methods that aim to capture low-dimensional, time-dependent latent structure of the pairwise interactions [43–45] (see also [46–48] for other dimension reduction methods for neuroscience data).

Here by combining the state-space model proposed in [37–39] with analytic approximation methods, we provide a framework for estimating interactions of neuronal populations consisting of up to 60 neurons. To find the mean we used the pseudolikelihood approximation method. To approximate the variance, we provide two alternative methods: the Bethe or the mean-field approximation. The Bayesian analysis methods for larger networks of neurons allow us to better understand macroscopic states of a neural population, such as entropy, free energy and sensitivity, all in a time-resolved manner and with credible intervals. Thus the model provides a new way to investigate effects of stimuli and behavior on activity of neuronal populations. It is expected to provide observations that give us insights into the underlying circuitry and its computation.

Materials and Methods

To clarify the problem of large-scale analysis on dynamic population activity, we first formulate the state-space model and its estimation method originally investigated in [37, 38] in the next subsection. Then we describe how to introduce approximation methods to the state-space model in order to overcome the limitation of the model and make the large-scale analysis possible. The custom-made Python programs are provided on GitHub (https://github.com/christiando/ssll_lib).

The state-space analysis of neural population activity

Spike data. To investigate how neuronal activities realize perception, cognition, and behavior, neurophysiologists record timing of neuronal spiking activity over the course of a behavioral paradigm designed to test specific hypotheses. Typically, these experiments are repeated multiple times under the same experimental conditions to uncover common neuronal dynamics related to the behavioral paradigm from stochastic spiking activities. We assume that neural data is composed of repeated measurements (R times) of spike timing recorded from N neurons simultaneously. Hereafter repetition is termed trial. To analyze activity patterns of neurons, we discretize the parallel spike sequences into T time bins with

bin size Δ , and represent the population activity by a set of binary variables. For neurons $n = 1, \dots, N$, time bins $t = 1, \dots, T$, and trials $r = 1, \dots, R$, the neural activity is represented by a binary variable $X_n^{r,t}$, where $X_n^{r,t} = 1$ when neuron n spiked in time bin t and trial r ; and $X_n^{r,t} = 0$ otherwise. Hence, we describe the whole data as a $N \times R \times T$ dimensional binary matrix. The activity pattern of N neurons at time bin t and trial r is a vector, $\mathbf{X}^{r,t} = (X_1^{r,t}, \dots, X_N^{r,t})'$. Similarly, $\mathbf{X}^t = (\mathbf{X}^{1,t}, \dots, \mathbf{X}^{R,t})$ summarizes observations for all neurons $1, \dots, N$ and all trials $1, \dots, R$ at time bin t . Finally, $\mathbf{X}^{t_1:t_2} = (\mathbf{X}^{t_1}, \dots, \mathbf{X}^{t_2})$ denotes the observations from time bin t_1 to t_2 .

State-space model of neural population activity. We assume a state-space model of dynamic population activity composed of two submodels; an observation model and a state model. First, the observation model specifies the probability distribution of population activity patterns using state variables, whereas the latter dictates how those state variables change. Here we construct the observation model using the exponential family distribution considering up to pairwise interactions of neurons' activities,

$$p(\mathbf{x}|\boldsymbol{\theta}_t) = \exp \left[\sum_{i=1}^N \theta_i^t x_i + \sum_{j>i} \theta_{ij}^t x_j x_i - \psi_t(\boldsymbol{\theta}_t) \right], \tag{1}$$

where $\psi_t(\boldsymbol{\theta}_t)$ is a log normalization term (a.k.a. log partition function). The model contains $d = N + N(N - 1)/2$ parameters $\{\theta_i^t\}, \{\theta_{ij}^t\}$ known as natural or canonical parameters of an exponential family distribution. In statistical mechanics, this model is named "Ising model", where the vector \mathbf{x} represents a spin configuration (up or down). There, the natural parameters $\{\theta_i^t\}, \{\theta_{ij}^t\}$ represent external magnetic field and interactions among the spins, and may be denoted as $\{h_i\}, \{J_{ij}\}$ conventionally. Here we consider these parameters to be time-dependent, and refer to them as state variables of the state-space model. By introducing the d -dimensional state vector $\boldsymbol{\theta}_t = (\theta_1^t, \dots, \theta_N^t, \theta_{1,2}^t, \dots, \theta_{N-1,N}^t)'$, and the feature vector $\mathbf{F}(\mathbf{x}) = (x_1, \dots, x_N, x_1 x_2, \dots, x_{N-1} x_N)'$, the model of Eq 1 is written concisely as $p(\mathbf{x}|\boldsymbol{\theta}_t) = \exp [\boldsymbol{\theta}_t' \mathbf{F}(\mathbf{x}) - \psi_t(\boldsymbol{\theta}_t)]$. The resulting log partition function is then given by

$$\psi_t(\boldsymbol{\theta}_t) = \log \sum_{\mathbf{x}} \exp [\boldsymbol{\theta}_t' \mathbf{F}(\mathbf{x})]. \tag{2}$$

In statistical mechanics, ψ_t is known as the free energy. Note that it specifies the probability that all neurons are simultaneously silent because $p(\mathbf{0}|\boldsymbol{\theta}_t) = \exp[-\psi_t(\boldsymbol{\theta}_t)]$. This model considers individual and pairwise activity of neurons. Hence, we will refer to it as the *pairwise observation model* in the following.

Next, the state model considers that dynamics of the latent state $\boldsymbol{\theta}_t$ is described by a random walk

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \boldsymbol{\xi}_t(\lambda), \tag{3}$$

where $\boldsymbol{\xi}_t$ is a random vector drawn from a multivariate normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{Q})$, and \mathbf{Q} is a diagonal covariance matrix. Here we assume that entries of the diagonal of the inverse matrix \mathbf{Q}^{-1} are given by a scalar λ that determines precision of the noise for all elements. For the initial time bin we set the density to $p(\boldsymbol{\theta}_1) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

It should be noted that here we model the neural dynamics as a *quasistatic* process, similarly to the classical analysis on dynamics of a thermodynamic system, e.g., a heat engine (see also [49]): At each time t , we presume that neural activity is sampled from the *equilibrium* distribution (Eq 1), which is the same across the trials (across-trial stationarity). The free energy (Eq 2) is also defined in the same manner as in the classical thermodynamics. We emphasize that the

quasistatic process is a simplified view of the neural dynamics. See [Discussion](#) for possible extensions of the model.

Estimating the state-space model. Given the data $\mathbf{X}^{1:T}$, our goal is to jointly estimate the posterior density of the latent states and the optimal noise precision λ . By denoting hyperparameters of the model as $\mathbf{w} = (\lambda, \boldsymbol{\mu}, \boldsymbol{\Sigma})$, the posterior density of the state process writes as

$$p(\boldsymbol{\theta}_{1:T} | \mathbf{X}^{1:T}, \mathbf{w}) = \frac{p(\mathbf{X}^{1:T} | \boldsymbol{\theta}_{1:T}) p(\boldsymbol{\theta}_{1:T}; \mathbf{w})}{p(\mathbf{X}^{1:T}; \mathbf{w})}, \tag{4}$$

where the first component in the numerator is constructed from the observation model, and the second component from the state model. In the next section, we provide the iterative method to construct this posterior density by approximating it by a Gaussian distribution (the Laplace approximation). The posterior density depends on the choice of the parameters \mathbf{w} . The optimal \mathbf{w} maximizes the marginal likelihood, a.k.a. evidence, that appears in the denominator in [Eq 4](#), given by

$$l(\mathbf{X}^{1:T} | \mathbf{w}) = p(\mathbf{X}^1 | \boldsymbol{\mu}, \boldsymbol{\Sigma}) \prod_{t=2}^T p(\mathbf{X}^t | \mathbf{X}^{1:t-1}, \lambda). \tag{5}$$

This approach is called the empirical Bayes method. In this study, we optimize noise precision λ and mean $\boldsymbol{\mu}$ of the initial distribution as described below while values for the covariance $\boldsymbol{\Sigma}$ are fixed. For fitting in the subsequent analyses, we set initial values as $\lambda = 100$ and $\boldsymbol{\Sigma} = 10\mathbf{I}$. For initial value of $\boldsymbol{\mu}$ we computed the vector $\boldsymbol{\theta}$ from time and trial averaged data, assuming $\{\theta_{ij}^t\} = \mathbf{0}$.

The optimization is achieved by an EM-algorithm combined with recursive Bayesian filtering/smoothing algorithms [[33](#), [50](#)]. In this approach, we alternately perform construction of the posterior density ([Eq 4](#), E-step) and optimization of the hyperparameters (M-step) until the marginal likelihood ([Eq 5](#)) saturates. In order to update the hyperparameters to new values \mathbf{w}^* from old values \mathbf{w} in the M-step, a lower bound of the marginal likelihood is maximized. This lower bound is obtained by applying the Jensen’s inequality to the marginal likelihood:

$$\begin{aligned} l(\mathbf{X}^{1:T} | \mathbf{w}^*) &= \log \int p(\mathbf{X}^{1:T}, \boldsymbol{\theta}_{1:T} | \mathbf{w}^*) d\boldsymbol{\theta}_{1:T} \\ &= \log \left\langle \frac{p(\mathbf{X}^{1:T}, \boldsymbol{\theta}_{1:T} | \mathbf{w}^*)}{p(\boldsymbol{\theta}_{1:T} | \mathbf{X}^{1:T}, \mathbf{w})} \right\rangle_{\boldsymbol{\theta}_{1:T} | \mathbf{X}^{1:T}, \mathbf{w}} \\ &\geq \left\langle \log \frac{p(\mathbf{X}^{1:T}, \boldsymbol{\theta}_{1:T} | \mathbf{w}^*)}{p(\boldsymbol{\theta}_{1:T} | \mathbf{X}^{1:T}, \mathbf{w})} \right\rangle_{\boldsymbol{\theta}_{1:T} | \mathbf{X}^{1:T}, \mathbf{w}} \\ &= \langle \log p(\mathbf{X}^{1:T}, \boldsymbol{\theta}_{1:T} | \mathbf{w}^*) \rangle_{\boldsymbol{\theta}_{1:T} | \mathbf{X}^{1:T}, \mathbf{w}} - \langle \log p(\boldsymbol{\theta}_{1:T} | \mathbf{X}^{1:T}, \mathbf{w}) \rangle_{\boldsymbol{\theta}_{1:T} | \mathbf{X}^{1:T}, \mathbf{w}} \end{aligned} \tag{6}$$

Here $\langle \cdot \rangle_{\boldsymbol{\theta}_{1:T} | \mathbf{X}^{1:T}, \mathbf{w}}$ is expectation by the posterior density of the state variables ([Eq 4](#)). In order to maximize the lower bound w.r.t. the new hyperparameters \mathbf{w}^* , we only need to maximize the first term, $q(\mathbf{w}^* | \mathbf{w}) \equiv \langle \log p(\mathbf{X}^{1:T}, \boldsymbol{\theta}_{1:T} | \mathbf{w}^*) \rangle_{\boldsymbol{\theta}_{1:T} | \mathbf{X}^{1:T}, \mathbf{w}}$. This term is called expected complete data log-likelihood, where the expectation is taken by the posterior density with the old

w. It is computed as

$$\begin{aligned}
 q(\mathbf{w}^*|\mathbf{w}) &= \sum_{t=1}^T \sum_{r=1}^R \langle \boldsymbol{\theta}_t^r \mathbf{F}(\mathbf{X}^{t,r}) - \psi(\boldsymbol{\theta}_t) \rangle_{\boldsymbol{\theta}_{1:T}|\mathbf{X}^{1:T},\mathbf{w}} \\
 &\quad - \frac{1}{2} \log |2\pi\boldsymbol{\Sigma}^*| - \frac{1}{2} \langle (\boldsymbol{\theta}_1 - \boldsymbol{\mu}^*)' \boldsymbol{\Sigma}^{*-1} (\boldsymbol{\theta}_1 - \boldsymbol{\mu}^*) \rangle_{\boldsymbol{\theta}_{1:T}|\mathbf{X}^{1:T},\mathbf{w}} \\
 &\quad - \frac{T-1}{2} \log |2\pi\mathbf{Q}^*| - \frac{1}{2} \sum_{t=2}^T \langle (\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-1})' \mathbf{Q}^{*-1} (\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-1}) \rangle_{\boldsymbol{\theta}_{1:T}|\mathbf{X}^{1:T},\mathbf{w}}.
 \end{aligned} \tag{7}$$

By considering derivatives of this equation w.r.t. the hyperparameters, we obtain their update rules. The precision $\lambda^* \mathbf{I} (= \mathbf{Q}^{*-1})$ is updated as

$$\lambda^* = \frac{1}{(T-1)d} \text{tr} \left[\sum_{t=2}^T \langle (\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-1})(\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-1})' \rangle_{\boldsymbol{\theta}_{1:T}|\mathbf{X}^{1:T},\mathbf{w}} \right], \tag{8}$$

where d is the dimension of vector $\boldsymbol{\theta}_t$. The initial mean is optimized by $\boldsymbol{\mu}^* = \langle \boldsymbol{\theta}_1 \rangle_{\boldsymbol{\theta}_{1:T}|\mathbf{X}^{1:T},\mathbf{w}}$. Here the key step is to develop an algorithm that constructs the posterior density of Eq 4. This is done by the forward and backward recursive Bayesian algorithms. Below we review this method followed by introduction of the approximations that make the method applicable to larger number of neurons.

Recursive estimation of dynamic neural interactions. The estimation of the latent process is achieved by forward filtering and then backward smoothing algorithms. In the filtering algorithm, we sequentially estimate the state of population activity at time bin t given the data up to time t . This estimate is given by the recursive Bayesian formula

$$p(\boldsymbol{\theta}_t|\mathbf{X}^{1:t}, \mathbf{w}) = \frac{p(\mathbf{X}^t|\boldsymbol{\theta}_t)p(\boldsymbol{\theta}_t|\mathbf{X}^{1:t-1}, \mathbf{w})}{p(\mathbf{X}^t|\mathbf{X}^{1:t-1}, \mathbf{w})}. \tag{9}$$

where $p(\mathbf{X}^t|\boldsymbol{\theta}_t)$ is obtained from the observation model. The second term in the numerator $p(\boldsymbol{\theta}_t|\mathbf{X}^{1:t-1}, \mathbf{w})$ is called the one-step prediction density. It is computed using the state model and the filter density at the previous time bin via the Chapman-Kolmogorov equation,

$$p(\boldsymbol{\theta}_t|\mathbf{X}^{1:t-1}, \mathbf{w}) = \int p(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}, \mathbf{w})p(\boldsymbol{\theta}_{t-1}|\mathbf{X}^{1:t-1}, \mathbf{w})d\boldsymbol{\theta}_{t-1}. \tag{10}$$

Thus the filter density (Eq 9) can be recursively computed for $t = 2, \dots, T$ using Eq 10, given observation and state models as well as an initial distribution of the one-step prediction density at time $t = 1$. Note that the initial one-step prediction density was specified as $p(\boldsymbol{\theta}_1) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. This distribution dictates the density of the state at the initial time step without observing neural activity.

The approximate nonlinear recursive formulae were developed by approximating the posterior density (Eq 9) with a Gaussian distribution [32, 51]. Let us assume that the filter density at time $t - 1$ is given by a Gaussian distribution with mean $\boldsymbol{\theta}_{t-1|t-1}$ and the covariance matrix $\mathbf{W}_{t-1|t-1}$. The subscript $t - 1|t - 1$ means the estimate at time $t - 1$ (left) given the data up to time bin $t - 1$ (right). Because the state model (Eq 3) is also Gaussian, the Chapman-Kolmogorov equation yields the one-step prediction density that is a Gaussian distribution with mean $\boldsymbol{\theta}_{t|t-1} = \boldsymbol{\theta}_{t-1|t-1}$ and covariance $\mathbf{W}_{t|t-1} = \mathbf{W}_{t-1|t-1} + \mathbf{Q}$. We then obtain the following

log posterior density (Eq 9),

$$\log p(\boldsymbol{\theta}_t | \mathbf{X}^{1:t}, \mathbf{w}) = \sum_{r=1}^R [\boldsymbol{\theta}_t' \mathbf{F}(\mathbf{X}^{t:r}) - \psi(\boldsymbol{\theta}_t)] - \frac{1}{2} (\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t|t-1})' \mathbf{W}_{t|t-1}^{-1} (\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t|t-1}) + \text{const.} \quad (11)$$

Here we approximate the posterior density by a Gaussian distribution (the Laplace approximation). We identify the mean of this distribution with the MAP estimate:

$$\boldsymbol{\theta}_{t|t} = \text{argmax}_{\boldsymbol{\theta}_t} \log p(\boldsymbol{\theta}_t | \mathbf{X}^{1:t}, \mathbf{w}). \quad (12)$$

This solution is called a filter mean. It may be obtained by gradient ascent algorithms such as the conjugate gradient algorithm and the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm. These algorithms use the gradient

$$\frac{\partial \log p(\boldsymbol{\theta}_t | \mathbf{X}^{1:t}, \mathbf{w})}{\partial \boldsymbol{\theta}_t} = \sum_{r=1}^R [\mathbf{F}(\mathbf{X}^{t:r}) - \boldsymbol{\eta}_t] - \mathbf{W}_{t|t-1}^{-1} (\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t|t-1}). \quad (13)$$

Here we define the expectation parameters $\boldsymbol{\eta}_t$ as

$$\boldsymbol{\eta}_t \equiv \frac{\partial \psi(\boldsymbol{\theta}_t)}{\partial \boldsymbol{\theta}_t} = \langle \mathbf{F}(\mathbf{x}) \rangle_{\boldsymbol{\theta}_t}, \quad (14)$$

where $\langle \mathbf{x} \rangle_{\boldsymbol{\theta}_t}$ is the expectation of \mathbf{x} with respect to $p(\mathbf{x} | \boldsymbol{\theta}_t)$. This expectation needs to be computed repeatedly in the gradient algorithms. The covariance matrix of the approximated Gaussian distribution is computed from the Hessian of the log posterior evaluated at the MAP estimate:

$$\begin{aligned} \mathbf{W}_{t|t}^{-1} &= - \left. \frac{\partial^2 \log p(\boldsymbol{\theta}_t | \mathbf{X}^{1:t}, \mathbf{w})}{\partial \boldsymbol{\theta}_t \partial \boldsymbol{\theta}_t'} \right|_{\boldsymbol{\theta}_{t|t}} \\ &= \mathbf{R} \mathbf{G}_t + \mathbf{W}_{t|t-1}^{-1}. \end{aligned} \quad (15)$$

\mathbf{G}_t is the Fisher-information matrix:

$$\mathbf{G}_t \equiv \left. \frac{\partial \psi(\boldsymbol{\theta}_t)}{\partial \boldsymbol{\theta}_t \partial \boldsymbol{\theta}_t'} \right|_{\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t|t}} = \langle \mathbf{F}(\mathbf{x}) \mathbf{F}(\mathbf{x})' \rangle_{\boldsymbol{\theta}_{t|t}} - \langle \mathbf{F}(\mathbf{x}) \rangle_{\boldsymbol{\theta}_{t|t}} \langle \mathbf{F}(\mathbf{x}) \rangle_{\boldsymbol{\theta}_{t|t}}' \quad (16)$$

The expectations are taken by $p(\mathbf{x} | \boldsymbol{\theta}_{t|t})$. Note that we initially assumed that the filter density at previous time step is a Gaussian distribution when computing the Chapman-Kolmogorov equation. By the Laplace approximation, this assumption is fulfilled in the next time step. Additionally we assumed that the initial distribution of the state variables is Gaussian. Thus we obtain an approximate nonlinear recursive filter that is consistent across the iterations.

Once the approximate filter density is constructed for $t = 1, \dots, T$, the backward smoothing algorithm is applied to obtain the smoothed posterior density of the state variable at time t [32, 52],

$$p(\boldsymbol{\theta}_t | \mathbf{X}^{1:T}, \mathbf{w}) = p(\boldsymbol{\theta}_t | \mathbf{X}^{1:t}, \mathbf{w}) \int \frac{p(\boldsymbol{\theta}_{t+1} | \mathbf{X}^{1:T}, \mathbf{w}) p(\boldsymbol{\theta}_{t+1} | \boldsymbol{\theta}_t, \mathbf{w})}{p(\boldsymbol{\theta}_{t+1} | \mathbf{X}^{1:t}, \mathbf{w})} d\boldsymbol{\theta}_{t+1}. \quad (17)$$

for $t = T, \dots, 1$. In practice, the following fixed interval smoothing algorithm [32] provides the

smoothed MAP estimate $\theta_{t|T}$ and smoothed covariance $\mathbf{W}_{t|T}$ of the posterior distribution

$$\theta_{t|T} = \theta_{t|t} + \mathbf{A}_t(\theta_{t+1|T} - \theta_{t+1|t}), \tag{18}$$

$$\mathbf{W}_{t|T} = \mathbf{W}_{t|t} + \mathbf{A}_t(\mathbf{W}_{t+1|T} - \mathbf{W}_{t+1|t})\mathbf{A}'_t, \tag{19}$$

where $\mathbf{A}_t = \mathbf{W}_{t|t}\mathbf{W}_{t+1|t}^{-1}$. In addition, the posterior covariance matrix between state variables at time t and $t - 1$ is obtained as $\mathbf{W}_{t-1, t|T} = \mathbf{A}_{t-1}\mathbf{W}_{t|T}$ [53]. This procedure constructs the smoother posterior density of the latent process (Eq 4) by approximating it as a Gaussian process of length $N(N + 1)/2 \times T$ with mean $(\theta'_{1|T}, \theta'_{2|T}, \dots, \theta'_{T|T})$ and a block tridiagonal covariance matrix whose block diagonal is given by $\mathbf{W}_{t|T}$ (for $t = 1, \dots, T$), and block off-diagonals are given by $\mathbf{W}_{t-1, t|T}$ (for $t = 2, \dots, T$).

Approximation methods for large-scale analysis

Approximate estimate of filter mean by pseudolikelihood method. To obtain the filter estimate using iterative gradient ascent methods, the gradient (Eq 13) needs to be evaluated at each iteration. This requires computation of the expectations (Eq 14) by summing over all 2^N states the network can realize. This is infeasible for a large network size N . Thus the method introduced in the previous subsection was limited to $N \leq 15$. However, the *pseudolikelihood* method [40, 54, 55] has been shown to estimate with reasonable accuracy the interactions without requiring evaluation of the expectations. Here we incorporate it into the sequential Bayesian estimation framework.

The pseudolikelihood approximates the likelihood of the joint activity of neurons by a product of conditional likelihoods of each neuron given the activity of the others. Let the activity of neurons except neuron n be $\mathbf{x}_{\setminus n} = (x_1, \dots, x_{n-1}, x_{n+1}, \dots, x_N)$; and $f_t^n(\mathbf{x}_{\setminus n}) = \theta'_t \mathbf{F}(x_n = 1, \mathbf{x}_{\setminus n})$. Then the pseudolikelihood is given by

$$\prod_{r=1}^R \tilde{p}(\mathbf{X}^{t,r} | \theta_t) = \prod_{r=1}^R \prod_{n=1}^N p(X_n^{t,r} | \mathbf{X}_{\setminus n}^{t,r}, \theta_t) = \prod_{r=1}^R \prod_{n=1}^N \frac{\exp(X_n^{t,r} f_t^n(\mathbf{X}_{\setminus n}^{t,r}))}{1 + \exp(f_t^n(\mathbf{X}_{\setminus n}^{t,r}))}. \tag{20}$$

Note that the log partition function does not appear in Eq 20. Replacing the likelihood in Eq 9 with Eq 20 yields

$$\begin{aligned} \log p(\theta_t | \mathbf{X}^{1:T}, \mathbf{w}) &\approx \sum_{r=1}^R \sum_{n=1}^N \left[X_n^{t,r} f_t^n(\mathbf{X}_{\setminus n}^{t,r}) - \log \left(1 + \exp \left(f_t^n(\mathbf{X}_{\setminus n}^{t,r}) \right) \right) \right] \\ &\quad - \frac{1}{2} (\theta_t - \theta_{t|t-1})' \mathbf{W}_{t|t-1}^{-1} (\theta_t - \theta_{t|t-1}) + \text{const.} \end{aligned} \tag{21}$$

The derivative of this approximated filter density results in

$$\frac{\partial \log p(\theta_t | \mathbf{X}^{1:T}, \mathbf{w})}{\partial \theta_t} \approx \sum_{r=1}^R \sum_{n=1}^N \left[(X_n^{t,r} - \tilde{\eta}_n^{t,r}) \frac{\partial f_t^n(\mathbf{X}_{\setminus n}^{t,r})}{\partial \theta_t} \right] - \mathbf{W}_{t|t-1}^{-1} (\theta_t - \theta_{t|t-1}), \tag{22}$$

where $\tilde{\eta}_n^{t,r} = \langle x_n^t | \mathbf{X}_{\setminus n}^{t,r} \rangle_{\theta_t}$, i.e., the expectation of x_n^t being 1 given the activity of the other neurons. Using this gradient in the same gradient ascent algorithms as before we obtain the approximate mean $\theta_{t|t}$ of the filter density.

Approximation of the filter covariance. The pseudolikelihood can provide the approximate mode of the filter density (Eq 12). However, to perform the sequential estimation, we need in addition the filter covariance matrix (Eq 15). This requires to compute the Fisher

information matrix (Eq 16, i.e., the Hessian of the observation model at the filter mean $\theta_{t|t}$). To compute the Fisher information matrix, not only the first and second order but also the third and fourth order expectation parameters need to be evaluated at the filter mean parameters. In order to avoid computing the higher-order expectation parameters and to reduce the computational cost of the matrix inversion, we approximate it by a diagonal matrix. The diagonal is composed of the first and second order expectation parameters $\{\eta_i^{t|t}\}, \{\eta_{ij}^{t|t}\}$, where the expectations parameters are defined as $\eta_i^{t|t} \equiv \langle x_i \rangle_{\theta_{t|t}}$ and $\eta_{ij}^{t|t} \equiv \langle x_i x_j \rangle_{\theta_{t|t}}$. Here we test two different approximation methods to obtain these marginals. One is the *Bethe approximation* [56] and the other the mean-field *Thouless-Anderson-Palmer (TAP)* approach [57].

Bethe approximation. The Bethe approach approximates a probability distribution by assuming that it factorizes into its pairwise marginals. Hence, the approximated joint distribution writes as

$$p(\mathbf{x}|\theta_{t|t}) \approx \frac{\prod_{i,j>i} q_t(x_i, x_j)}{\prod_i q_t(x_i)^{(N-1)-1}} := q_t(\mathbf{x}), \tag{23}$$

where q are so-called *beliefs* [58] that approximate the marginals of the underlying distribution p . Note that for any acyclic graph this yields the true joint distribution. However, here the observation model (Eq 1) is a fully connected graph and hence the Bethe approximation ignores all cycles. Realizing that the beliefs have to fulfill constraints ($\sum_{x_j} q_t(x_i, x_j) = q_t(x_i)$ and $\sum_{x_i} q_t(x_i) = 1$) one can write the problem as a Lagrangian that has to be minimized. This allows to derive a dual representation of the marginals (in terms of the Lagrangian multipliers), which in turn allows to derive messages that are sent from one belief to another. Propagating this beliefs through the Markov field yields the belief propagation algorithm (BP) [56]. While BP is relatively fast in obtaining the expectation values, it is not guaranteed to converge to an unique solution. This guarantee is provided by the alternative concave-convex procedure (CCCP) [59]. CCCP also starts from the same Lagrangian, but updates the beliefs and Lagrangian multipliers in an alternating manner. This more strict procedure comes with the disadvantage that it is much slower than BP. Therefore, here the two algorithms are combined to a *hybrid method*, where BP is utilized primarily and the algorithm falls back to CCCP, when BP does not converge. For more details on the Bethe approximation, see [S1 Text](#).

The estimation of the log partition function for the Bethe approximation is simply computed by the negative logarithm of the approximated probability (Eq 23) that all neurons are silent, i.e.,

$$\psi_t \approx -\log q_t(\mathbf{0}). \tag{24}$$

TAP approximation. The TAP approximation of the expectation parameters $\eta_{t|t}$ given the natural parameters $\theta_{t|t}$ (*forward-problem*) can be derived in multiple ways [13, 60], but here we follow [61, 62] that use the so-called “Plefka expansion”. The following formulae and their derivation are revised for binary variables $x_i \in \{0, 1\}$ instead of $\{-1, 1\}$. See [S2 Text](#) for more details. The method constructs a new free energy as a function of the mixture coordinates ($\{\eta_i^{t|t}\}, \{\theta_{ij}^{t|t}\}$) by the Legendre transformation of the log partition function ψ_t as $\sum_{i=1}^N \theta_i^{t|t} \eta_i^{t|t} - \psi_t$. Then this function is approximated by a second-order expansion around the independent model assuming weak pairwise interactions. This results in the approximate log

partition function,

$$\begin{aligned} \psi_t \approx & \sum_{i=1}^N \theta_i^{t|t} \eta_i^{t|t} - \sum_{i=1}^N (\eta_i^{t|t} \log \eta_i^{t|t} + (1 - \eta_i^{t|t}) \log (1 - \eta_i^{t|t})) + \frac{1}{2} \sum_{j \neq i} \theta_{ij}^{t|t} \eta_i^{t|t} \eta_j^{t|t} \\ & + \frac{1}{8} \sum_{j \neq i} (\theta_{ij}^{t|t})^2 (\eta_i^{t|t} - (\eta_i^{t|t})^2) (\eta_j^{t|t} - (\eta_j^{t|t})^2). \end{aligned} \tag{25}$$

Here we extended the definition of interaction parameters as $\theta_{ii}^{t|t} = 0$ and $\theta_{ij}^{t|t} = \theta_{ji}^{t|t}$. At the independent model, the values for the expectations can be computed and the expansion yields correction terms for the non-zero $\theta_{ij}^{t|t}$. Since derivatives of the new free energy based on the mixture coordinates w.r.t. $\{\eta_i^{t|t}\}$ yield the first order parameters $\{\theta_i^{t|t}\}$, we obtain the following self-consistent equations:

$$\theta_i^{t|t} = \log \left(\frac{\eta_i^{t|t}}{1 - \eta_i^{t|t}} \right) - \sum_{j \neq i} \theta_{ij}^{t|t} \eta_j - \frac{1}{2} \sum_{j \neq i} (\theta_{ij}^{t|t})^2 \left(\frac{1}{2} - \eta_i^{t|t} \right) (\eta_j^{t|t} - (\eta_j^{t|t})^2), \tag{26}$$

for $i, j = 1, \dots, N$. Solving this equations yields the first order expectations which can be used to estimate the log partition function (Eq 25).

Furthermore, from the relation $\frac{\partial \theta_i^{t|t}}{\partial \eta_j^{t|t}} = [\mathbf{G}_t^{-1}]_{ij}$ we obtain

$$[\mathbf{G}_t^{-1}]_{ij} = \frac{1}{\eta_i^{t|t} (1 - \eta_i^{t|t})} \delta_{ij} - \theta_{ij}^{t|t} - (\theta_{ij}^{t|t})^2 \left(\frac{1}{2} - \eta_i^{t|t} \right) \left(\frac{1}{2} - \eta_j^{t|t} \right). \tag{27}$$

Here δ_{ij} is the Kronecker delta function, which is 1 for $i = j$ and 0 otherwise. To obtain the second order expectation parameters, we calculate and then invert the $N \times N$ matrix obtained by Eq 27, and approximate it as the Fisher information matrix for $\{\theta_i\}$ given in Eq 16 to obtain the second order expectation parameters by $\eta_{ij}^{t|t} = [\mathbf{G}_t]_{ij} + \eta_i^{t|t} \eta_j^{t|t}$ [61].

Approximate marginal likelihood. Because the TAP and Bethe approximations provide estimates of the log partition function ψ_t , we are able to evaluate the approximation of the marginal likelihood (Eq 7), and the EM-algorithm for the state-space model can be run until it converges. The approximate marginal likelihood is obtained as (see also [38])

$$\begin{aligned} l(\mathbf{X}^{1:T} | \mathbf{w}) &= \sum_{t=1}^T \log \int p(\mathbf{X}^t | \boldsymbol{\theta}_t) p(\boldsymbol{\theta}_t | \mathbf{X}^{1:t-1}, \mathbf{w}) d\boldsymbol{\theta}_t \\ &\approx \sum_{t=1}^T \sum_{r=1}^R \left[\boldsymbol{\theta}_{t|t}^r \mathbf{F}(\mathbf{X}^{t,r}) - \psi_t(\boldsymbol{\theta}_{t|t}) \right] \\ &\quad - \frac{1}{2} \sum_{t=1}^T (\boldsymbol{\theta}_{t|t} - \boldsymbol{\theta}_{t|t-1})' \mathbf{W}_{t|t-1}^{-1} (\boldsymbol{\theta}_{t|t} - \boldsymbol{\theta}_{t|t-1}) \\ &\quad + \frac{1}{2} \sum_{t=1}^T (\log \det W_{t|t} - \log \det W_{t|t-1}), \end{aligned} \tag{28}$$

where $p(\boldsymbol{\theta}_t | \mathbf{X}^{1:0}, \mathbf{w})$ indicates a prior of the initial distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Similarly, we use $\boldsymbol{\theta}_{1|0} = \boldsymbol{\mu}$ and $\mathbf{W}_{1|0} = \boldsymbol{\Sigma}$. Here the integral with respect to $\boldsymbol{\theta}_t$ at the first equality is approximated as an integral of a Gaussian function, using up to the quadratic information around its mode (the Laplace approximation). From Eqs 11 and 12, it turns out that the mean and covariance of the filter density provide this information.

Results

Model fit to simulated data

In the following subsections, we demonstrate the fit of the state-space model of neural population activity to artificially generated data of 40 neurons with dynamic couplings for $T = 500$ time bins. To be able to compare it to the ground truth we construct 4 populations each consisting of 10 neurons. Individual parameters $\theta_{1:T}$ of the underlying submodels are generated as smooth independent Gaussian processes, where the mean for the first order parameters θ_i^t increases at $t = 100$ and then decreases more slowly shortly after that. The interaction parameters θ_{ij}^t are generated as Gaussian processes whose mean is fixed at 0. In total, 500 trials of spike data are sampled from this generative model. Note that the sampled individual parameters differ and vary over time although we use homogeneous means. The increase of the mean for θ_i^t increases spiking probability followed by a decrease back to baseline (Fig 1A). In the resulting data neurons spike with time averaged probabilities ranging from 0.10 up to 0.21. Supposing bin width $\Delta = 10$ ms these are in a physiologically reasonable range. This exemplary scenario may mimic a population that independently receives an external input elicited by e.g., a sensory stimulus. For details of the generation of the data see S3 Text.

Next we fit the state-space model of neural population activity to the generated data with the combination of pseudolikelihood and Bethe approximation. This combination is chosen for the demonstration because it provides the best estimates of the underlying model as we will assess later in this section. Top panel of Fig 1B shows snapshots of the smoothed estimates of the inferred network at different time points ($t = 50, 150, 300$). The color of the nodes indicate the smoothed estimates of the first order parameters θ_i^{tT} and the one of the edges interactions θ_{ij}^{tT} . Visual inspection of the fitted network suffices to identify that there are 4 independent subpopulations of correlated neurons (one in each quadrant). To check whether the inferred changes over time match those of the underlying generative model, credible intervals of three fitted couplings are compared with their underlying values (Fig 1B Bottom). The fit follows the dynamics, and correctly identifies the parameter that is constantly 0 (the lowest panel).

Estimating macroscopic properties of the network

One of the main motives to model joint activities of a large population of neurons is to assess macroscopic properties of the network in a time-dependent manner with credible intervals. The macroscopic measures obtained for this example are shown in Fig 1C, and in the following we introduce them one by one.

The first and simplest macroscopic property shown in the top left panel of Fig 1C is the probability of spiking in a network (population spike rate). We define it as

$$P_{\text{spike}}(t) = \frac{1}{N} \sum_{i=1}^N \eta_i^t, \tag{29}$$

where η_i^t is the spike rate of i th neuron at time t . Considering the smoothed estimate $\eta_i^t = \eta_i^{tT}$, the method recovers correctly the empirical rate obtained from the data (Fig 1A Bottom). The shaded area in the panel indicates the 98% credible interval of the population spike rate obtained by resampling the natural parameters from the smoothed posterior density 100 times at each bin. The underlying spike probability for $N = 40$ neurons is obtained by calculating the marginals η_i^t independently for each subpopulation and averaging over all neurons.

Next from the state-space model of neural population activity one can estimate the probability of simultaneous silence (i.e., the probability that no neuron elicits a spike, Fig 1C

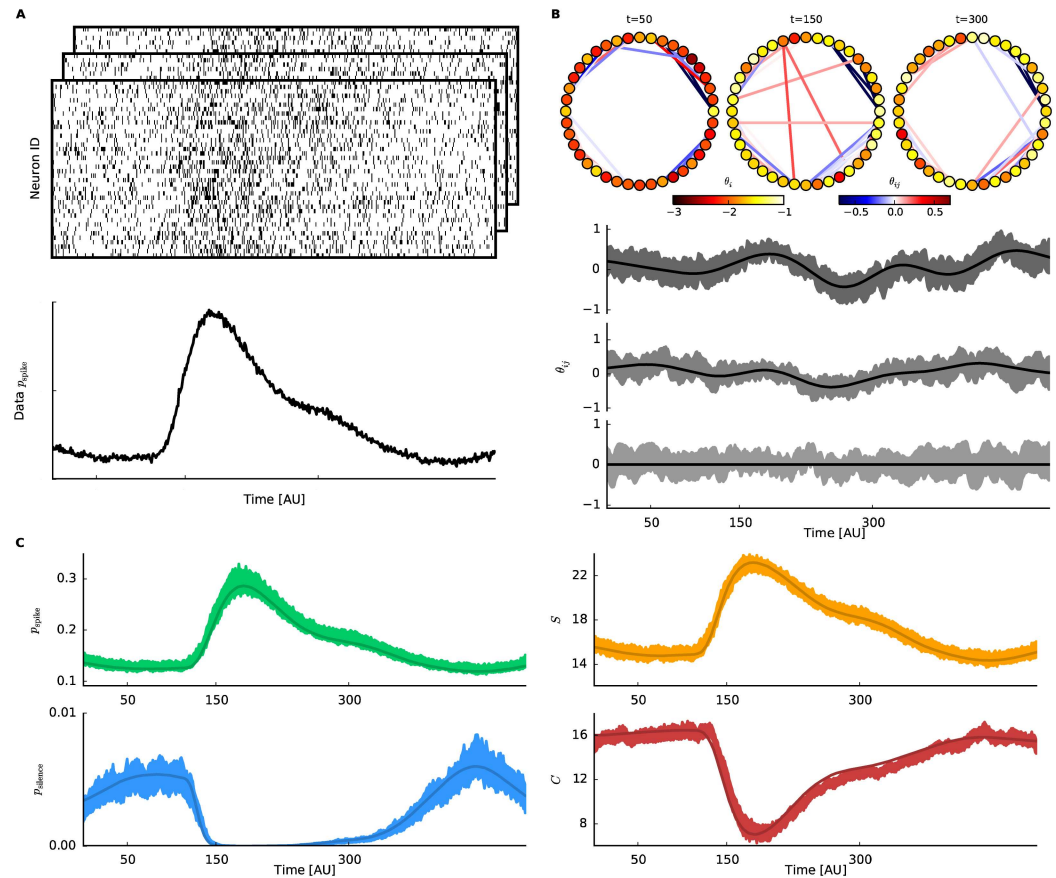


Fig 1. Approximate inference of dynamic neural interactions and macroscopic network properties. Analysis on simulated spike data of 40 neurons. **A** Top: Simultaneous spiking activity of 40 neurons that are repeatedly simulated 500 times (here only 3 trials are visualized). The data is sampled from a time-dependent model of a neural population (Eq 1). The time-varying parameters are chosen such that neurons' spike probability resembles evoked activity in response to stimulus presentation to an animal. The neural interactions are assumed to smoothly change irrespective of the firing rates. See the main text for details. Bottom: Empirical spike probability over time, averaged over trials and neurons. **B** Top: Estimated network states at $t = 50, 150, 300$ by the pseudolikelihood-Bethe approximation method. Neurons are represented by nodes whose colors respectively indicate a value of the smoothed estimate of θ_i^t (for $i = 1, \dots, 40$). Links are color-coded according to estimated strength of the interaction θ_{ij}^t between connected nodes (positive or negative interactions are marked in red or blue, respectively). Only significant edges are displayed, where the corresponding θ_{ij}^t has a 98% credible interval that does not include 0. Bottom: Dynamics of 3 exemplary interaction parameters, θ_{ij}^t . The lines denote the ground truth from which the binary data are sampled. The shaded areas are 98% credible intervals. **C** Estimated population rate (top left). Probability that all neurons are silent (bottom left). Entropy (top right) and heat capacity (bottom right) of the neural population. In all panels, shaded areas indicate 1% and 99% quantiles obtained by resampling the natural parameters from the fitted smoothed distribution. Solid lines represent ground truth computed from the underlying network model.

doi:10.1371/journal.pcbi.1005309.g001

bottom left)

$$P_{\text{silence}}(t) = \exp(-\psi_t). \tag{30}$$

The approximation methods allow us to evaluate the log partition function ψ_t (Eqs 24 and 25). Here we use smoothed estimates to compute the log partition function. Thus we immediately obtain the probability of simultaneous silence. The expected simultaneous silence for $N = 40$ neurons is obtained as multiplication of the silence probabilities of the 4 subpopulations.

The entropy of the network (i.e., expectation of the information content, $\langle -\log p(\mathbf{x}|\boldsymbol{\theta}_t) \rangle_{\boldsymbol{\theta}_t}$) can be also calculated from the model as

$$S(t) = -\boldsymbol{\theta}'_t \mathbf{n}_t + \psi_t. \tag{31}$$

Estimation of this information theoretic measure allows us to quantify the amount of interactions in the network by comparing the pairwise model to the independent one (see following analyses and Eq 36). Since it is an extensive quantity, the entropy of $N = 40$ neurons is obtained by addition of the entropies from the 4 independent subpopulations. The entropy increases while the individual activity rates of neurons also increases (Fig 1C top right).

The last measure shown in the bottom right panel of Fig 1C is the heat capacity, or sensitivity, of the system. It is the variance of information content: $C(t) = \langle \{-\log p(\mathbf{x}|\boldsymbol{\theta}_t)\}^2 \rangle_{\boldsymbol{\theta}_t} - \{\langle -\log p(\mathbf{x}|\boldsymbol{\theta}_t) \rangle_{\boldsymbol{\theta}_t}\}^2$, where the brackets indicate expectation by $p(\mathbf{x}|\boldsymbol{\theta}_t)$. It is also the variance of the Hamiltonian $-\boldsymbol{\theta}'_t \mathbf{F}(\mathbf{x})$. Thus we can obtain it by introducing a nominal dual parameter β to the Hamiltonian in the model, assuming that it is 1 for real data. The log partition function of the augmented model is

$$\psi_t(\beta) = \log \sum_{\mathbf{x}} \exp(\beta \boldsymbol{\theta}'_t \mathbf{F}(\mathbf{x})). \tag{32}$$

The variance of Hamiltonian is given as the Fisher information w.r.t. β , i.e., the second derivative of the log partition function. This allows us to use the approximate ψ_t to assess the heat capacity. Then we further approximate the second derivative by its discrete version

$$C(t) = \left. \frac{\partial^2 \psi_t}{\partial \beta^2} \right|_{\beta=1} \approx \frac{\psi_t(1 + \epsilon) - 2\psi_t(1) + \psi_t(1 - \epsilon)}{\epsilon^2}, \tag{33}$$

and ϵ is chosen to be 10^{-3} . The heat capacity measures sensitivity of the network, namely how much the network activity changes due to subtle changes in its network configuration (i.e., to changes of the $\boldsymbol{\theta}_t$ parameters). Networks with higher sensitivity are more responsive to changes than those with lower sensitivity. Similarly to the entropy, the heat capacity is an extensive quantity. For the simulated data, the heat capacity decreases while activity rates of neurons are increased (Fig 1C bottom right).

Assessment of fitting error with different network sizes and amount of data

Next we examine the goodness-of-fit of the model fitted by the pseudolikelihood and Bethe approximation methods. In particular, we ask how the fitting performance changes with increasing network size. For this reason we generated 6 dynamic models for populations of 10 neurons as described previously (500 time bins, 500 trials). Then we construct smaller or larger populations by concatenating the independent groups. The model is fitted by the pseudolikelihood and Bethe approximation methods to the first subnetwork, then two subnetworks, and so on, until we fit the model to a network containing 60 neurons composed of 6 independent groups. We obtain estimates of the macroscopic measures from the smoothed estimates of the model parameters at each time bin. Fig 2A shows values of these measures averaged over time. The results show extensive properties of macroscopic measures (except for the population spike rate), and that the estimates may slightly deviate for larger number of neurons.

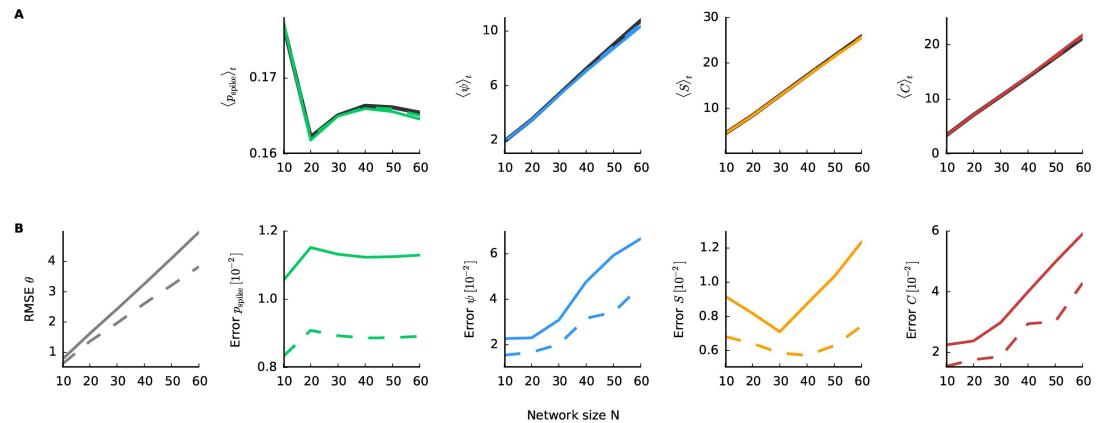


Fig 2. Approximation error and network size. Error analysis on networks consisting of subpopulations with 10 neurons, constructed by the same procedure as in Fig 1. **A:** The average value of the macroscopic properties over time as a function of network size. Black line is the true value, while colored lines show the estimated ones (solid line fit with 500 trials and dashed with 1000 trials) **B:** The corresponding errors (only for θ_t the RMSE is shown) for 500 trials (solid) and 1000 trials (dashed).

doi:10.1371/journal.pcbi.1005309.g002

To assess quality of the fit, first the root mean squared error (RMSE) for the natural parameters averaged across time bins is calculated

$$\text{RMSE}(\theta_{t|T}) = \sqrt{\frac{1}{T} \sum_{t=1}^T \|\theta_{t|T} - \theta_t\|^2}, \quad (34)$$

where $\theta_{t|T}$ is the smoothed estimate of the underlying model θ_t . $\|v\|$ denotes the $L2$ -norm of vector v . For the data sets with 500 trials, the RMSE increases linearly with network size (Fig 2B Left). Furthermore, the error for the macroscopic measures is assessed by

$$\text{Error}[f(\theta_{t|T})] = \frac{\text{RMSE}(f(\theta_{t|T}))}{\frac{1}{T} \sum_{t=1}^T f(\theta_t)}, \quad (35)$$

where $f(\theta_{t|T})$ is any function of the macroscopic measures. The RMSE is defined similarly to Eq 34 while substituting the parameters $\theta_{t|T}$ by the function $f(\theta_{t|T})$. Besides the population rate these errors also increase as the network size increases (Fig 2B). We observe non-monotonic behavior in some of the macroscopic properties (e.g., average spike rate and the entropy's error), which can be explained by fluctuations from the data generation process.

To understand whether these errors increase primarily due to the approximation methods used for the fit or because of the finite amount of data, the fit is repeated but now to spiking data with 1000 trials. The error of the fit is reduced particularly for larger network size (Fig 2B dashed lines), suggesting that the limited amount of data is mainly responsible for the estimation error.

In general, the estimation error is largest at time points where the parameters θ_t change rapidly. This is a general problem of smoothing algorithms, including spike rate estimation, which depend on fixed smoothness parameter(s) (i.e., here λ) optimized for an entire observation period (see e.g., [63] for optimizing a variable smoothness parameter to cope with such abrupt changes).

Comparison between Bethe and TAP approximation

To this end, only the Bethe approximation was used in combination with the pseudolikelihood to fit the model approximately. However, as discussed previously, the TAP approximation constitutes a potential alternative. To assess the quality of both approximations, we investigated a small network (15 neurons, 500 time bins, 1000 trials). The data was generated as described for Fig 1. The smaller network is considered because it allows to fit the model by an exact method without the Bethe or TAP approximations. Here the exact method refers to the method in which the expectation parameters are calculated exactly at the gradient search for the MAP estimates of model parameters (Eq 13). It should be noted that we approximate the posterior density by the Gaussian distribution even for the “exact method” in the recursive Bayesian algorithm. Comparison of the approximation methods with the exact method determines the error that is caused by the approximation methods and not by the finite amount of data.

First, investigation of three exemplary time points (Fig 3A) reveals that both the pseudolikelihood-Bethe and the pseudolikelihood-TAP approximation recover the underlying parameters. We examine the error across time bins by the RMSE. Comparing RMSE of the approximation results with the exact fit (Fig 3B) demonstrates that the both approximations perform worse in the same range. To examine the approximations also for large networks ($N = 60$) we sampled 1000 trials (as for Fig 2). In Fig 3C we observe that errors of the approximations are comparable. Furthermore, we compare running times required for fitting the network of the two methods (Fig 3D). The pseudolikelihood-TAP approximation turns out to be faster than Bethe. We observed that the EM algorithm required more iterations for the Bethe approximation. Furthermore, the occasional use of the CCCP contributed to the long fitting time of the pseudolikelihood-Bethe procedure.

Since both, Bethe and TAP, provide an approximation for the log partition function ψ_t (Eqs 25 and 24), we assess their performance for the same data as in Fig 3. The time evolution of simultaneous silence (directly linked to ψ by Eq 30) is recovered by exact, Bethe, and TAP (Fig 4A). The results show that the TAP approximation slightly overestimated the probability in

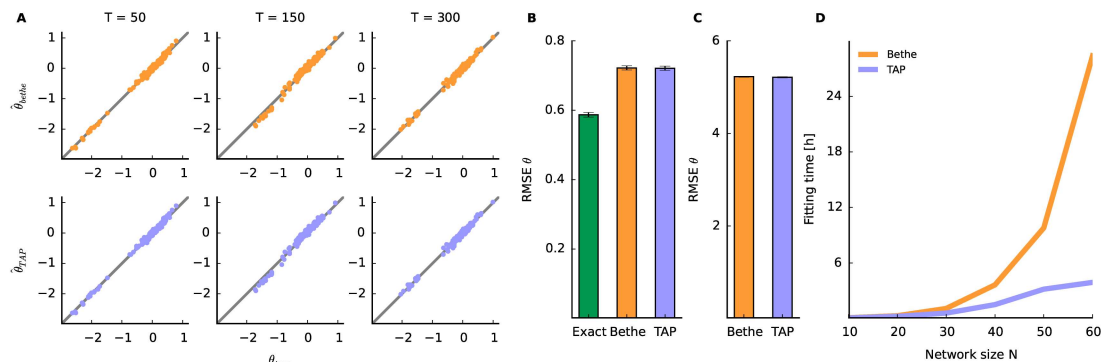


Fig 3. Comparison of the Bethe and TAP approximation. Simulated neural activity composed of 500 time bins, and 1000 trials are used to compare the two approximation methods. The underlying model parameters follow Fig 1. **A** Top: Ground truth θ_t of a network of 15 neurons vs. its smoothed estimate by pseudolikelihood-Bethe approximation at three different time points ($t = 50, 150, 300$). Bottom: The same as above obtained with pseudolikelihood-TAP approximation. **B** The RMSE between the true model parameter θ_t and its smoothed estimate by the exact inference, pseudolikelihood-Bethe, or pseudolikelihood-TAP approximation. The bar height and error bars indicate the mean and standard deviation from 10 realizations of data, each sampled from the same underlying parameters (generated as in Fig 1). **C** As in B the RMSE of the estimated model parameters for a network of 60 neurons, composed of 6 equally sized subnetworks. **D** Running time as function of network size for the two different approximation methods.

doi:10.1371/journal.pcbi.1005309.g003




Fig 4. Time-varying probability of simultaneous silence. Results of different approximation methods. The underlying model parameters are the same as in Fig 3. **A** The probability of simultaneous silence ($\rho_{\text{silence}}(t) = \exp(-\psi_t)$) for a network of 15 neurons as a function of time. The pseudolikelihood-Bethe (orange) and pseudolikelihood-TAP (lavender) method estimate the underlying value with sufficient accuracy (dashed black). For comparison, an estimate by the exact method (green) is shown. **B** The error between the approximate and true free energy ψ_t . **C** The error of free energy ψ_t for large networks ($N = 60$, data same as in Fig 3C).

doi:10.1371/journal.pcbi.1005309.g004

this example. This is also reflected in the $\text{Error}[\psi(\{\hat{\theta}_{i|T}\}_t)]$ (Fig 4B), where the Bethe approximation performs better than the TAP method. However, the error for the Bethe approximation increases compared to the exact method. The relation between the two approximation methods persists also for large networks (Fig 4C). Another disadvantage of the TAP approximation is that the system of non-linear equations occasionally could not be solved. This happens more frequently when fitting larger networks and/or networks with stronger interactions. Therefore, it seems that the pseudolikelihood-Bethe approximation exhibits more accurate estimates; hence we will use it again for the following analysis. However the faster fitting of pseudolikelihood-TAP can be advantageous elsewhere.

Dynamic network inference from V4 spiking data of behaving monkey

We now apply the approximate inference method to analyze activity of monkey V4 neurons recorded while the animal performed repeatedly (1004 trials) the following behavioral task. Each trial began when the monkey fixated its gaze within 1 degree of a centrally-positioned dot on a computer screen. After 150 ms, a drifting sinusoidal grating was presented for 2 s in the receptive field area of the neuronal population that was recorded, at which time the grating stimulus disappeared and the fixation point moved to a new, randomly chosen location on the screen, and the animal made an eye movement to fixate on the new location. Data epochs from 500 ms prior to grating stimulus onset until 500 ms after stimulus offset were extracted from the continuous recording for analysis. The spiking data obtained by micro-electrode recordings includes 112 single and multi units identified by their distinct wave forms. The experiment was performed at the University of Pittsburgh. All experimental procedures were approved by the University of Pittsburgh Institutional Animal Care and Use Committee, and were performed in accordance with the United States' National Institutes of Health (NIH) *Guide for the Care and Use of Laboratory Animals*. For details on experimental setup, recording and unit identification see [64]. The recorded units are tested for across-trial stationarity (which is the assumption of the model): The mean firing rates for each trial are standardized and if more than 5% of the trials were outside the 95% confidence interval the unit is excluded. After this preprocessing 45 units remained. To obtain the binary data, the spike trains are discretized into time bins with $\Delta = 10$ ms resulting into 300 time bins over the course of the trial.

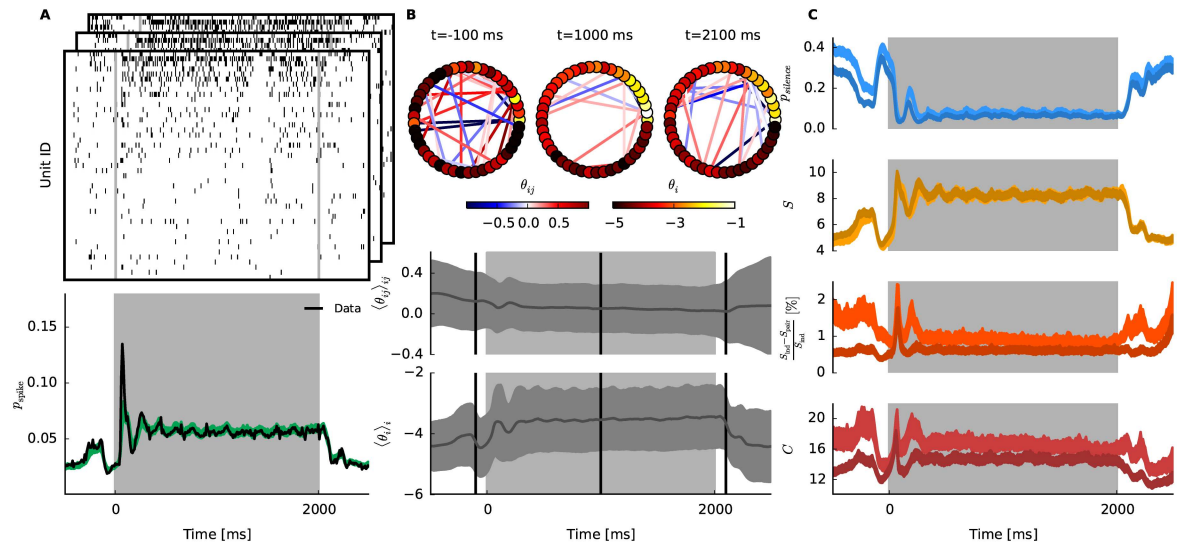


Fig 5. Dynamic network inference from monkey V4 data. In this experiment, a 90° grating on a screen was presented to the monkey for 2s (light gray shaded areas). 1004 trials were recorded, and binary spike trains were constructed with bin width of 10 ms. **A** Top: Exemplary spiking data ($N = 45$). Bottom: Empirical probability (black) of observing a spike over time and spike probability of the fitted model (green). **B** Top: The fitted network at three different time points, before, during, and after stimulation. Edges with significantly non-zero θ_{ij} are displayed (as in Fig 1). Bottom: The mean of smoothed MAP estimates for θ_{ij} and θ_i (dark gray line). The shaded area is the mean \pm standard deviation. **C** Credible intervals of macroscopic measures of the network over time obtained from the smoothed estimates of the model (light color). Dark shaded area corresponds to the credible intervals of the estimates for trial shuffled data.

doi:10.1371/journal.pcbi.1005309.g005

Exemplary data are displayed in Fig 5A Top. We note that the following conclusions of this analysis do not change even if we use smaller and larger bin size ($\Delta = 5$ and 20 ms).

After the data are preprocessed, we analyze the network dynamics of the 45 units during the task period by the state-space model for the neural population activity. Inference is done by using the pseudolikelihood-Bethe approximation. The results of fitting the state-space model are displayed in Fig 5B. Before presenting detailed results, we note that considering dynamics in activity rates and neural correlations better explains the population activity while avoiding overfitting, compared to assuming that they are stationary. To assess this, we compared the predictive ability of the state-space model with that of the stationary model, using the Aikake (Bayesian) Information Criterion (AIC) [65] defined as $-2\ln(\mathbf{X}^{1:T}|\mathbf{w}) + 2k$, where k is the number of free parameters in \mathbf{w} . To obtain the latter, we fitted the state-space model once more but now fixing $\lambda^{-1} = 0$, which results in a stationary model since the state model in Eq 3 no longer contains variability. The result confirms that the dynamic model better predicts the data ($AIC_{\text{dyn}} = 4467026$ for the dynamic model and $AIC_{\text{stat}} = 4576544$ for the stationary model).

We observe stimulus locked oscillations in the population firing rate that are also captured by the model (Fig 5A Bottom). The average of the estimated natural parameters (Fig 5B Bottom) show that these oscillations are explained by the first order parameters $\theta_i^{(t)}$. We note that these oscillations are mainly caused by two units with high firing rates and they should not be considered as a homogeneous property of the network. Investigation of the network states before, during, and after the stimulus (Fig 5B Top) reveals that the interactions $\theta_{ij}^{(t)}$ are altered over time. This is also reflected in an average over the all pairwise interactions (Fig 5B Center), where the mean decreases during the stimulus presentation as well as the standard deviation.

Thus neurons are likely to decorrelate during the stimulus presentation whereas the population rate increases and oscillates at the same time.

Similarly to the analysis of artificial data (Fig 1), we measure the macroscopic properties of the fitted model over the task period (see Fig 5C for credible intervals). To test the contribution of interactions in the recorded data, the model is once again fitted to trial shuffled data [23], which should destroy all correlations among units that do not occur due to chance. Comparison of the macroscopic measures between the models fitted to the original data and to the trial shuffled data shows how interactions among units alter the results. In the following, we will refer to the two models as “actual” and “trial shuffled” model.

The probability of simultaneous silence shows again the stimulus locked oscillations, and decreases during the stimulus period. The difference between the actual and trial shuffled model before the stimulus is larger than during and after the stimulus, suggesting that the observed positive interactions contributed to increasing the silence probability in particular before and after the stimulus period. The entropy reflects the oscillations and shows a strong increase ($\sim 1/3$) during the stimulus period. This is reasonable because we observe an increase in activity rates and a decrease in correlations—both effects should result in an increase in entropy. Next, we examine how much of the entropy is explained by the interactions among the neurons. To do so, at each time point we calculate the corresponding independent model by projecting the fitted interaction model to the independent model (i.e., the model with the same individual firing rates η_i^t but with all $\theta_{ij}^t = 0$). The entropy of the independent model S_{ind} should always be larger than S_{pair} , the entropy of the model with interactions. Hence, a fraction of entropy explained by the interactions can be calculated as

$$\frac{S_{\text{ind}} - S_{\text{pair}}}{S_{\text{ind}}}. \tag{36}$$

In general, contribution of interactions to the entropy is small for these data ($\leq 2\%$). However, the contribution is less during stimulus presentation, compared to the period before the stimulus. Only in the beginning of the stimulus presentation, two peaks of correlated activity can be observed. The observed reduction of the fractional entropy for interactions could be caused by the increase of the first order parameters θ_i^t and/or by the decrease of the interactions θ_{ij}^t during the stimulus period. The decorrelation observed during the stimulus period is successfully dissociated from the oscillatory activity: Previously observed oscillations are absent in this measure of interactions. This result is important because ignoring such firing rate dynamics often leads to erroneous detection of positive correlations among neurons. A clear exception is the first peak appeared during the stimulus presentation, which was also observed in the trial-shuffled model. Indeed, the first sharp increase of the spike rates was not faithfully captured by the models, which caused spurious interactions in the trial-shuffled model. Last, the sensitivity (heat capacity) of the network over time is obtained. While for the artificial data in Fig 1 the sensitivity showed a drastic decrease, such reduction is not observed in the V4 data. The sensitivity of the network is maintained at approximately the same value before and during the stimulus period. This is interesting since we already observed that before and during the stimulus the network seems to be in two qualitatively different states (low vs. high firing rate and strong vs. weak interactions). After stimulus presentation the sensitivity drops. Overall, neural interactions contribute to have higher sensitivity (see light vs. dark credible intervals).

Dynamic network inference from simulated balanced network data

Networks with balanced excitation and inhibition have been used to describe cortical activity [66, 67]. To see whether the balanced network model can reproduce the findings from the

recorded V4, we simulate spiking data using the balanced spiking network following [24], and analyze these data with the state-space model. The network consists of 1000 leaky integrate-and-fire neurons (800 excitatory, 200 inhibitory) (For details see S4 Text). Connection probability is 20%, between all neurons. The network receives input from 800 Poisson neurons. Each input neuron has a Gaussian tuning curve, where the preferred direction is randomly assigned. We choose an experimental paradigm which resembles one of the V4 data. 1000 trials of 3 s duration are simulated. Before each trial, the simulation runs for 500 ms under random Poisson inputs such that the network state at the beginning of each trial is independent. Then the trial starts at -500 ms. At 0 ms a 90° is shown for 2 s followed again by a 500 ms period of stimulus absence. The activity of 140 neurons are recorded for investigation. From the recorded subpopulation, we further selected 40 excitatory and 20 inhibitory neurons with the highest firing rates for the following analysis. Binary spike trains were obtained by binning with $\Delta = 10$ ms. Exemplary data are shown in Fig 6A (top spike trains are from excitatory, and bottom spike trains from inhibitory neurons). We then fitted the state-space model to these data.

As for the V4 data, we show in Fig 6B 3 snapshots of the network ($N = 60$) (Top), as well as mean and standard deviation of θ_i^{tT} and θ_{ij}^{tT} (Bottom). In contrast to the V4 network there are numerous significant non-zero couplings. However, similarly to the monkey data, we observe an increase for θ_i^t and a decrease of θ_{ij}^t during the stimulus period. We also assess the macroscopic states for the balanced network (Fig 6C). As in the V4 data the probability of silence decreases during the stimulus period. Furthermore, compared to the trial shuffled result, the difference is larger before and after the stimulus than during the stimulus, suggesting a larger contribution of the couplings to silence when no stimulus is present. The entropy increases during the stimulus period. The credible interval for the trial shuffled data is narrower than for actual model and the entropy tends to be larger. Up to this point we did not find, in the macroscopic properties, significant qualitative differences between the V4 data and the simulated data from the balanced network. However, the entropy that is explained by the couplings

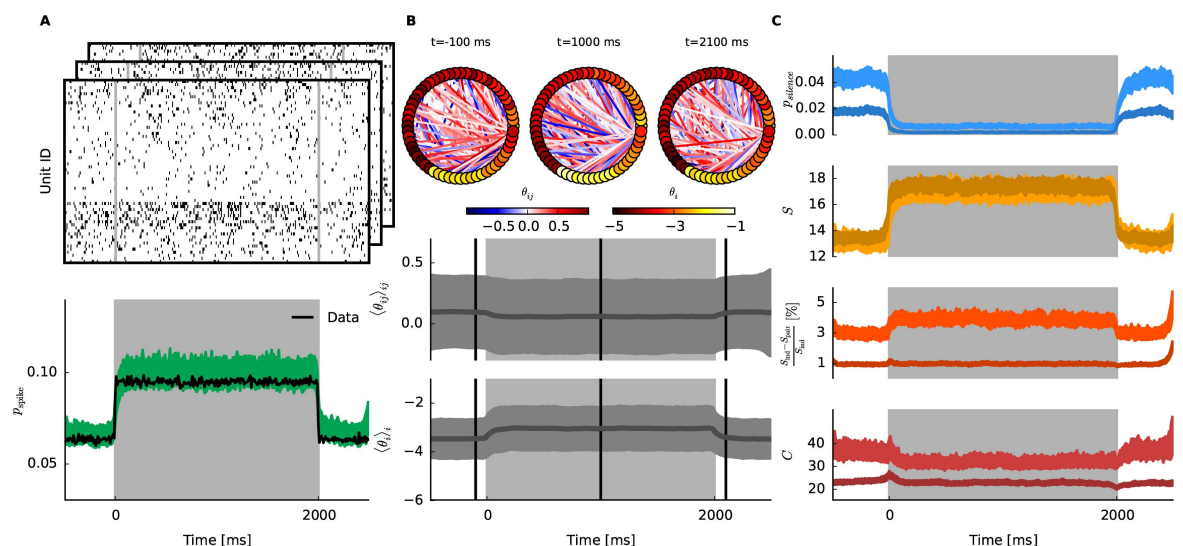


Fig 6. Dynamic network inference from simulated balanced network data. 60 neurons (40 excitatory, 20 inhibitory) are recorded from a simulated balanced network of 1000 leaky integrate-and-fire neurons that receive inputs from 800 excitatory orientation selective Poisson neurons (mean firing rate 7.5 Hz when no stimulus present). See main text for the details. Stimulus was presented for 2 s, and 1000 trials are generated. Bin width is 10 ms. The structure of this figure is the same as in Fig 5.

doi:10.1371/journal.pcbi.1005309.g006




Fig 7. Comparison of model interactions with synapses in the balanced network. The synaptic structure is reflected in the inferred interactions. **A** Histograms of the interactions θ_{ij}^{tT} for all pairs (gray), pairs that are connected by *at least* one excitatory synapse (red), and those that are connected by *at least* one inhibitory synapse (blue) at three different time points. **B** Averages of the couplings θ_{ij}^{tT} across time and pairs as a function of a network size (always consisting of two thirds of excitatory and one third of inhibitory neurons). Colors as in A, and error bars denote standard deviations.

doi:10.1371/journal.pcbi.1005309.g007

increases during the stimulus, while in the V4 data a decrease is observed (Fig 6C, third panel). Hence, the interactions in the balanced network become stronger during the stimulus, even though the mean of the couplings θ_{ij}^{tT} decreases for this period. This can be explained by more negative values in estimated couplings during the stimulus period. The sensitivity slightly decreases when the stimulus is shown and, as for the V4 data, couplings contribute to higher sensitivity.

Observing the dynamics in the model parameters poses the question how the actual synaptic connectivity structure of the network is reflected in the inferred interactions. Do positive values correspond to excitatory synapses, and negative to inhibitory ones? While for the V4 data this is impossible to assess, we compare the values of θ_{ij}^{tT} of pairs, that are at least connected by one excitatory synapse and those that are connected by at least one inhibitory synapse (Fig 7A, red and blue histograms respectively). In general, excitatory connected pairs show more positive values, while inhibiting ones tend to be negative. The most negative values are almost exclusively explained by inhibiting pairs. However, compared to all θ_{ij}^{tT} (gray histogram) many positive couplings θ_{ij}^{tT} do not represent excitatory connected pairs. Thus it is difficult to identify excitatory synapses from the inferred couplings. The result that inhibitory pairs showed stronger negative couplings, while excitatory pairs were mostly represented by weak positive couplings, can be explained by on average much stronger conductance of inhibitory synapses.

Finally we compare the mean values of couplings between different network sizes (Fig 7B). To do so networks of size $N = 15, 30, 60$ are fitted, where the network always consisted of one third inhibitory and two thirds excitatory neurons. However, neither for excitatory, inhibitory or all couplings we could identify dependency on the network sizes that can be analyzed by our model.

Discussion

This study provides approximate inference methods for simultaneously estimating neural interactions of a large number of neurons, and quantifying macroscopic properties of the network in a time-resolved manner. We assessed performance of these methods by using simulated parallel spike sequences, and demonstrated the utility of the proposed approach by revealing dynamic decorrelation of V4 neurons and maintained susceptibility during stimulus presentations. Furthermore we compared those findings with data from a simple balanced

network of LIF neurons, which suggested that further refinements were necessary to reproduce the observed network activity.

Accurate assessment of correlated population activity in ongoing and evoked activity is a key to understand the underlying biological mechanisms and their coding principles. It is critical to model time-dependent firing rates to correctly assess neural interactions. If we apply a stationary model of neural interactions to independent neurons with varying firing rates, we may erroneously observe excess of correlations [22–24, 26, 68]. Such an apparent issue of a stationary model can introduce considerable confusion in search of fundamental coding principles of neurons. Several related studies accounted for the nonstationary activity by modeling time-dependent external fields (c.f., $\{\theta_i^t\}$ in Eq 1) while fixing pairwise interactions [26, 30]. In addition to the external fields, however, we consider that modeling dynamics of correlations are important particularly for analyses of neurons recorded from awake animals because neural correlations are known to appear dynamically in relation to behavioral demand to the animals [27–29, 38, 69]. Indeed, we found dynamic decorrelation of V4 neurons during stimulus presentation (Fig 5C 3rd panel), which may reflect asynchronous neural activities under stimulus processing of an alert animal [70, 71]. In general, it is important to compare the result with that of surrogate data in which one destroys correlations to examine potentially short-lasting time-varying interactions in relation to behavioral paradigms.

The current state-space model presumes that the neural dynamic follows a *quasistatic* process. At each time t , we assumed that population activity is sampled from the *equilibrium* joint distribution given by Eq 1 across trials while the state of population activity smoothly changes within a trial. This is of course a simplified view of neuronal dynamics. Most notably, dependency of the neurons' activity on their past activity makes the system a nonequilibrium one. Such activity is captured by models via the history effect, e.g., using the kinetic Ising model [25, 26, 72, 73] or generalized linear models (GLM) of point and Bernoulli processes [35, 74–76]. Given the past activities, these models construct the joint activity assuming their conditional independence. The equilibrium and non-equilibrium models thus assume different generative processes, even though the pseudo-likelihood approximation for our equilibrium Ising model used similar conditional independence given the activity of other neurons at the same time. It is an important topic to include both modeling frameworks in the sequential Bayes estimation to better account for dynamic and nonequilibrium properties of neural activity [39]. The model goodness-of-fit may be additionally improved by including sparseness constraints on the couplings as was done in the stationary models [40, 77, 78].

In this study, we employed the classical pseudolikelihood method to perform MAP estimation of interactions (i.e., natural parameters) without computing the partition function. For the inverse problem without the prior, we may use alternative approximation methods such as Bethe and TAP approximations, and further state-of-the-art methods such as the Sessak-Monasson [12], minimum-probability-flow [15], and adaptive-cluster expansion [17] method. However, here we chose the pseudolikelihood method because it was not trivial to apply the other methods to the Bayesian estimation. Alternatively, the Bethe and TAP approximation methods may be used to approximate the expectation parameters during the iterative procedure of the exact MAP estimation (Eq 13) because these methods allow us to estimate the expectation parameter from the natural parameters (the forward problem). However, as we found in the estimation of the Fisher information, TAP may occasionally fail and Bethe approximation by BP may not converge. Thus we rather used these methods after the MAP estimation was found by the pseudolikelihood method. The framework, however, is not limited to these approximation methods, and new methods may be incorporated into the state-space model to further increase the number of neurons that can be analyzed.

It should be noted that the current model does not include higher-order interactions to explain the population dynamics. While neural higher-order interactions are ubiquitously observed *in vivo* [38, 79–81] as well as *in vitro* [20, 21, 82, 83] conditions, it remains to be elucidated how they contribute to characterizing evoked activities. It is an important step to include higher-order interactions in the large-scale time-dependent model. However, the proposed method that includes up to pairwise interactions can be used as a null model for testing activity features involving higher-order interactions. For example, both experimental and modeling studies showed that simultaneous silence of neurons constitutes a major feature of higher-order interactions of stationary neural activities [83, 84]. It remains to be tested, though, if silence probability of all neurons recorded from behaving animals exceed prediction by the pairwise model. Such sparse population activity may be expected when animals process natural scenes, compared to artificial stimuli [85].

The limiting factor for the current model on the network size is rather the lack of data than the performance of the approximation methods (Fig 2). Hence, the state-space or other time-resolved methods that include dimension reduction techniques will be important approaches to explain activity of much larger populations than analyzed here. While there is still room for improvement, the currently proposed method already allows researchers to start testing hypotheses of network responses under distinct task conditions or brain states. These observations will serve to construct biophysical models of neural networks by constraining them, therefore revealing their coding principles.

Supporting Information

S1 Text. Bethe approximation.

(PDF)

S2 Text. TAP approximation.

(PDF)

S3 Text. Generation of simulated data.

(PDF)

S4 Text. Simulated experiment with a balanced network.

(PDF)

Acknowledgments

The authors thank Thomas Sharp for originally translating Matlab code written by HS to Python code, and Adam Snyder and Matthew A. Smith for kindly providing the V4 spiking data. CD and HS acknowledge Taro Toyozumi for hosting CD's stay in RIKEN Brain Science Institute, and Timm Lochmann for valuable ideas and discussions.

Author Contributions

Conceptualization: CD KO HS.

Formal analysis: CD HS.

Investigation: CD HS.

Methodology: CD HS.

Project administration: CD KO HS.

Resources: CD HS.

Software: CD HS.

Supervision: KO HS.

Validation: CD HS.

Visualization: CD HS.

Writing – original draft: CD HS.

Writing – review & editing: CD KO HS.

References

1. London M, Häusser M. Dendritic computation. *Annual Review on Neuroscience*. 2005; 28:503–532. doi: [10.1146/annurev.neuro.28.061604.135703](https://doi.org/10.1146/annurev.neuro.28.061604.135703)
2. De La Rocha J, Doiron B, Shea-Brown E, Josić K, Reyes A. Correlation between neural spike trains increases with firing rate. *Nature*. 2007; 448(7155):802–806. doi: [10.1038/nature06028](https://doi.org/10.1038/nature06028) PMID: [17700699](https://pubmed.ncbi.nlm.nih.gov/17700699/)
3. Reyes AD. Synchrony-dependent propagation of firing rate in iteratively constructed networks in vitro. *Nature Neuroscience*. 2003; 6(6):593–599. doi: [10.1038/nm1056](https://doi.org/10.1038/nm1056) PMID: [12730700](https://pubmed.ncbi.nlm.nih.gov/12730700/)
4. Pitkow X, Meister M. Decorrelation and efficient coding by retinal ganglion cells. *Nature Neuroscience*. 2012; 15(4):628–635. doi: [10.1038/nn.3064](https://doi.org/10.1038/nn.3064) PMID: [22406548](https://pubmed.ncbi.nlm.nih.gov/22406548/)
5. Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A. Spontaneously emerging cortical representations of visual attributes. *Nature*. 2003; 425(6961):954–956. doi: [10.1038/nature02078](https://doi.org/10.1038/nature02078) PMID: [14586468](https://pubmed.ncbi.nlm.nih.gov/14586468/)
6. Luczak A, Barthó P, Harris KD. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*. 2009; 62(3):413–425. doi: [10.1016/j.neuron.2009.03.014](https://doi.org/10.1016/j.neuron.2009.03.014) PMID: [19447096](https://pubmed.ncbi.nlm.nih.gov/19447096/)
7. Shlens J, Field GD, Gauthier JL, Grivich MI, Petrusca D, Sher A, et al. The structure of multi-neuron firing patterns in primate retina. *The Journal of Neuroscience*. 2006; 26(32):8254–8266. doi: [10.1523/JNEUROSCI.1282-06.2006](https://doi.org/10.1523/JNEUROSCI.1282-06.2006) PMID: [16899720](https://pubmed.ncbi.nlm.nih.gov/16899720/)
8. Schneidman E, Berry MJ, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*. 2006; 440(7087):1007–1012. doi: [10.1038/nature04701](https://doi.org/10.1038/nature04701) PMID: [16625187](https://pubmed.ncbi.nlm.nih.gov/16625187/)
9. Lezon TR, Banavar JR, Cieplak M, Maritan A, Fedoroff NV. Using the principle of entropy maximization to infer genetic interaction networks from gene expression patterns. *Proceedings of the National Academy of Sciences*. 2006; 103(50):19033–19038. doi: [10.1073/pnas.0609152103](https://doi.org/10.1073/pnas.0609152103) PMID: [17138668](https://pubmed.ncbi.nlm.nih.gov/17138668/)
10. Mora T, Walczak AM, Bialek W, Callan CG. Maximum entropy models for antibody diversity. *Proceedings of the National Academy of Sciences*. 2010; 107(12):5405–5410. doi: [10.1073/pnas.1001705107](https://doi.org/10.1073/pnas.1001705107) PMID: [20212159](https://pubmed.ncbi.nlm.nih.gov/20212159/)
11. Bialek W, Cavagna A, Giardina I, Mora T, Silvestri E, Viale M, et al. Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences*. 2012; 109(13):4786–4791. doi: [10.1073/pnas.1118633109](https://doi.org/10.1073/pnas.1118633109) PMID: [22427355](https://pubmed.ncbi.nlm.nih.gov/22427355/)
12. Sessak V, Monasson R. Small-correlation expansions for the inverse Ising problem. *Journal of Physics A: Mathematical and Theoretical*. 2009; 42(5):055001. doi: [10.1088/1751-8113/42/5/055001](https://doi.org/10.1088/1751-8113/42/5/055001)
13. Roudi Y, Aurell E, Hertz J. Statistical physics of pairwise probability models. *Frontiers in Computational Neuroscience*. 2009; 3(22). doi: [10.3389/neuro.10.022.2009](https://doi.org/10.3389/neuro.10.022.2009) PMID: [19949460](https://pubmed.ncbi.nlm.nih.gov/19949460/)
14. Roudi Y, Tyrcha J, Hertz J. Ising model for neural data: model quality and approximate methods for extracting functional connectivity. *Physical Review E*. 2009; 79(5):051915. doi: [10.1103/PhysRevE.79.051915](https://doi.org/10.1103/PhysRevE.79.051915) PMID: [19518488](https://pubmed.ncbi.nlm.nih.gov/19518488/)
15. Sohl-Dickstein J, Battaglino PB, DeWeese MR. New method for parameter estimation in probabilistic models: minimum probability flow. *Physical Review Letters*. 2011; 107(22):220601. doi: [10.1103/PhysRevLett.107.220601](https://doi.org/10.1103/PhysRevLett.107.220601) PMID: [22182019](https://pubmed.ncbi.nlm.nih.gov/22182019/)
16. Schaub MT, Schultz SR. The Ising decoder: reading out the activity of large neural ensembles. *Journal of Computational Neuroscience*. 2012; 32(1):101–118. doi: [10.1007/s10827-011-0342-z](https://doi.org/10.1007/s10827-011-0342-z) PMID: [21667155](https://pubmed.ncbi.nlm.nih.gov/21667155/)

17. Cocco S, Monasson R. Adaptive cluster expansion for the inverse Ising problem: convergence, algorithm and tests. *Journal of Statistical Physics*. 2012; 147(2):252–314. doi: [10.1007/s10955-012-0463-4](https://doi.org/10.1007/s10955-012-0463-4)
18. Haslinger R, Ba D, Galuske R, Williams Z, Pipa G. Missing mass approximations for the partition function of stimulus driven Ising models. *Frontiers in Computational Neuroscience*. 2013; 7. doi: [10.3389/fncom.2013.00096](https://doi.org/10.3389/fncom.2013.00096) PMID: [23898262](https://pubmed.ncbi.nlm.nih.gov/23898262/)
19. Roudi Y, Nirenberg S, Latham PE. Pairwise maximum entropy models for studying large biological systems: when they can work and when they can't. *PLoS Computational Biology*. 2009; 5(5):e1000380. doi: [10.1371/journal.pcbi.1000380](https://doi.org/10.1371/journal.pcbi.1000380) PMID: [19424487](https://pubmed.ncbi.nlm.nih.gov/19424487/)
20. Ganmor E, Segev R, Schneidman E. Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *Proceedings of the National Academy of Sciences*. 2011; 108(23):9679–9684. doi: [10.1073/pnas.1019641108](https://doi.org/10.1073/pnas.1019641108) PMID: [21602497](https://pubmed.ncbi.nlm.nih.gov/21602497/)
21. Tkačik G, Marre O, Amodei D, Schneidman E, Bialek W, Berry MJ II. Searching for collective behavior in a large network of sensory neurons. *PLoS Computational Biology*. 2014; 10(1):e1003408. doi: [10.1371/journal.pcbi.1003408](https://doi.org/10.1371/journal.pcbi.1003408) PMID: [24391485](https://pubmed.ncbi.nlm.nih.gov/24391485/)
22. Brody CD. Correlations without synchrony. *Neural computation*. 1999; 11(7):1537–1551. doi: [10.1162/089976699300016133](https://doi.org/10.1162/089976699300016133) PMID: [10490937](https://pubmed.ncbi.nlm.nih.gov/10490937/)
23. Grün S. Data-driven significance estimation for precise spike correlation. *Journal of Neurophysiology*. 2009; 101(3):1126–1140. doi: [10.1152/jn.00093.2008](https://doi.org/10.1152/jn.00093.2008) PMID: [19129298](https://pubmed.ncbi.nlm.nih.gov/19129298/)
24. Renart A, de la Rocha J, Bartho P, Hollender L, Parga N, Reyes A, et al. The asynchronous state in cortical circuits. *Science*. 2010; 327(5965):587–590. doi: [10.1126/science.1179850](https://doi.org/10.1126/science.1179850) PMID: [20110507](https://pubmed.ncbi.nlm.nih.gov/20110507/)
25. Roudi Y, Hertz J. Mean field theory for nonequilibrium network reconstruction. *Physical Review Letters*. 2011; 106(4):048702. doi: [10.1103/PhysRevLett.106.048702](https://doi.org/10.1103/PhysRevLett.106.048702) PMID: [21405370](https://pubmed.ncbi.nlm.nih.gov/21405370/)
26. Tyrcha J, Roudi Y, Marsili M, Hertz J. The effect of nonstationarity on models inferred from neural data. *Journal of Statistical Mechanics: Theory and Experiment*. 2013; 2013(03):P03005. doi: [10.1088/1742-5468/2013/03/P03005](https://doi.org/10.1088/1742-5468/2013/03/P03005)
27. Vaadia E, Haalman I, Abeles M, Bergman H, Prut Y, Slovin H, et al. Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature*. 1995; 373(6514):515–518. doi: [10.1038/373515a0](https://doi.org/10.1038/373515a0) PMID: [7845462](https://pubmed.ncbi.nlm.nih.gov/7845462/)
28. Riehle A, Grün S, Diesmann M, Aertsen A. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*. 1997; 278(5345):1950–1953. doi: [10.1126/science.278.5345.1950](https://doi.org/10.1126/science.278.5345.1950) PMID: [9395398](https://pubmed.ncbi.nlm.nih.gov/9395398/)
29. Sakurai Y, Takahashi S. Dynamic synchrony of firing in the monkey prefrontal cortex during working-memory tasks. *The Journal of Neuroscience*. 2006; 26(40):10141–10153. doi: [10.1523/JNEUROSCI.2423-06.2006](https://doi.org/10.1523/JNEUROSCI.2423-06.2006) PMID: [17021170](https://pubmed.ncbi.nlm.nih.gov/17021170/)
30. Granot-Atedgi E, Tkacik G, Segev R, Schneidman E. Stimulus-dependent maximum entropy models of neural population codes. *PLoS Computational Biology*. 2013; 9(3):e1002922. doi: [10.1371/journal.pcbi.1002922](https://doi.org/10.1371/journal.pcbi.1002922) PMID: [23516339](https://pubmed.ncbi.nlm.nih.gov/23516339/)
31. Chen Z, Brown EN. State space model. *Scholarpedia*. 2013; 8(3):30868. doi: [10.4249/scholarpedia.30868](https://doi.org/10.4249/scholarpedia.30868)
32. Brown EN, Frank LM, Tang D, Quirk MC, Wilson MA. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *The Journal of Neuroscience*. 1998; 18(18):7411–7425. PMID: [9736661](https://pubmed.ncbi.nlm.nih.gov/9736661/)
33. Smith AC, Brown EN. Estimating a state-space model from point process observations. *Neural Computation*. 2003; 15(5):965–991. doi: [10.1162/089976603765202622](https://doi.org/10.1162/089976603765202622) PMID: [12803953](https://pubmed.ncbi.nlm.nih.gov/12803953/)
34. Eden UT, Frank LM, Barbieri R, Solo V, Brown EN. Dynamic analysis of neural encoding by point process adaptive filtering. *Neural Computation*. 2004; 16(5):971–998. doi: [10.1162/089976604773135069](https://doi.org/10.1162/089976604773135069) PMID: [15070506](https://pubmed.ncbi.nlm.nih.gov/15070506/)
35. Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology*. 2005; 93(2):1074–1089. doi: [10.1152/jn.00697.2004](https://doi.org/10.1152/jn.00697.2004) PMID: [15356183](https://pubmed.ncbi.nlm.nih.gov/15356183/)
36. Srinivasan L, Eden UT, Willsky AS, Brown EN. A state-space analysis for reconstruction of goal-directed movements using neural signals. *Neural Computation*. 2006; 18(10):2465–2494. doi: [10.1162/neco.2006.18.10.2465](https://doi.org/10.1162/neco.2006.18.10.2465) PMID: [16907633](https://pubmed.ncbi.nlm.nih.gov/16907633/)
37. Shimazaki H, Amari Si, Brown EN, Grün S. State-space analysis on time-varying correlations in parallel spike sequences. In: *Acoustics, Speech and Signal Processing (ICASSP), IEEE International Conference on*. IEEE; 2009. p. 3501–3504. doi: [10.1109/ICASSP.2009.4960380](https://doi.org/10.1109/ICASSP.2009.4960380)

38. Shimazaki H, Amari Si, Brown EN, Grün S. State-space analysis of time-varying higher-order spike correlation for multiple neural spike train data. *PLoS Computational Biology*. 2012; 8(3):e1002385. doi: [10.1371/journal.pcbi.1002385](https://doi.org/10.1371/journal.pcbi.1002385) PMID: [22412358](https://pubmed.ncbi.nlm.nih.gov/22412358/)
39. Shimazaki H. Single-trial estimation of stimulus and spike-history effects on time-varying ensemble spiking activity of multiple neurons: a simulation study. In: *Journal of Physics: Conference Series*. vol. 473. IOP Publishing; 2013. 012009. doi: [10.1088/1742-6596/473/1/012009](https://doi.org/10.1088/1742-6596/473/1/012009)
40. Kolar M, Song L, Ahmed A, Xing EP. Estimating time-varying networks. *The Annals of Applied Statistics*. 2010; p. 94–123. doi: [10.1214/09-AOAS308](https://doi.org/10.1214/09-AOAS308)
41. Long JDI, Carmena JM. A statistical description of neural ensemble dynamics. *Frontiers in Computational Neuroscience*. 2011; 5:52. doi: [10.3389/fncom.2011.00052](https://doi.org/10.3389/fncom.2011.00052) PMID: [22319486](https://pubmed.ncbi.nlm.nih.gov/22319486/)
42. Kass RE, Kelly RC, Loh WL. Assessment of synchrony in multiple neural spike trains using loglinear point process models. *The Annals of Applied Statistics*. 2011; 5(2B):1262. doi: [10.1214/10-AOAS429](https://doi.org/10.1214/10-AOAS429) PMID: [21837263](https://pubmed.ncbi.nlm.nih.gov/21837263/)
43. Hayashi K, Hirayama JI, Ishii S. Dynamic exponential family matrix factorization. In: *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer; 2009. p. 452–462. doi: [10.1007/978-3-642-01307-2_41](https://doi.org/10.1007/978-3-642-01307-2_41)
44. Effenberger F, Hillar C. Discovery of Salient Low-Dimensional Dynamical Structure in Neuronal Population Activity Using Hopfield Networks. In: *International Workshop on Similarity-Based Pattern Recognition*. Springer; 2015. p. 199–208. doi: [10.1007/978-3-319-24261-3_16](https://doi.org/10.1007/978-3-319-24261-3_16)
45. Hirayama Ji, Hyvärinen A, Ishii S. Sparse and low-rank matrix regularization for learning time-varying Markov networks. *Machine Learning*. 2016; p. 1–32. doi: [10.1007/s10994-016-5568-6](https://doi.org/10.1007/s10994-016-5568-6)
46. Byron MY, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. In: *Advances in Neural Information Processing Systems*; 2009. p. 1881–1888. doi: [10.1152/jn.90941.2008](https://doi.org/10.1152/jn.90941.2008) PMID: [19357332](https://pubmed.ncbi.nlm.nih.gov/19357332/)
47. Cunningham JP, Byron MY. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*. 2014; 17(11):1500–1509. doi: [10.1038/nn.3776](https://doi.org/10.1038/nn.3776) PMID: [25151264](https://pubmed.ncbi.nlm.nih.gov/25151264/)
48. Okun M, Steinmetz NA, Cossell L, Iacaruso MF, Ko H, Barthó P, et al. Diverse coupling of neurons to populations in sensory cortex. *Nature*. 2015; 521(7553):511–515. doi: [10.1038/nature14273](https://doi.org/10.1038/nature14273) PMID: [25849776](https://pubmed.ncbi.nlm.nih.gov/25849776/)
49. Shimazaki H. Neurons as an Information-theoretic Engine. 2015;arXiv:1512.07855.
50. Shumway RH, Stoffer DS. An approach to time series smoothing and forecasting using the EM algorithm. *Journal of Time Series Analysis*. 1982; 3(4):253–264. doi: [10.1111/j.1467-9892.1982.tb00349.x](https://doi.org/10.1111/j.1467-9892.1982.tb00349.x)
51. Fahrmeir L. Posterior mode estimation by extended Kalman filtering for multivariate dynamic generalized linear models. *Journal of the American Statistical Association*. 1992; 87(418):501–509. doi: [10.1080/01621459.1992.10475232](https://doi.org/10.1080/01621459.1992.10475232)
52. Kitagawa G. Non-Gaussian state-space modeling of nonstationary time series. *Journal of the American Statistical Association*. 1987; 82(400):1032–1041. doi: [10.2307/2289381](https://doi.org/10.2307/2289381)
53. De Jong P, Mackinnon MJ. Covariances for smoothed estimates in state space models. *Biometrika*. 1988; 75(3):601–602. doi: [10.1093/biomet/75.3.601](https://doi.org/10.1093/biomet/75.3.601)
54. Besag J. Statistical analysis of non-lattice data. *The Statistician*. 1975; p. 179–195. doi: [10.2307/2987782](https://doi.org/10.2307/2987782)
55. Höfling H, Tibshirani R. Estimation of sparse binary pairwise markov networks using pseudo-likelihoods. *The Journal of Machine Learning Research*. 2009; 10:883–906. PMID: [21857799](https://pubmed.ncbi.nlm.nih.gov/21857799/)
56. Yedidia JS, Freeman WT, Weiss Y. Understanding belief propagation and its generalizations. *Exploring Artificial Intelligence in the New Millennium*. 2003; 8:236–239.
57. Thouless DJ, Anderson PW, Palmer RG. Solution of 'solvable model of a spin glass'. *Philosophical Magazine*. 1977; 35(3):593–601.
58. Yedidia J. An idiosyncratic journey beyond mean field theory. *Advanced mean field methods: Theory and practice*. 2001; p. 21–36.
59. Yuille AL. CCCP algorithms to minimize the Bethe and Kikuchi free energies: Convergent alternatives to belief propagation. *Neural Computation*. 2002; 14(7):1691–1722. doi: [10.1162/08997660260028674](https://doi.org/10.1162/08997660260028674) PMID: [12079552](https://pubmed.ncbi.nlm.nih.gov/12079552/)
60. Oppor M, Saad D. *Advanced mean field methods: Theory and practice*. MIT Press; 2001.
61. Tanaka T. Mean-field theory of Boltzmann machine learning. *Physical Review E*. 1998; 58(2):2302.
62. Tanaka T. A theory of mean field approximation. *Advances in Neural Information Processing Systems*. 1999; p. 351–360.

63. Shimazaki H, Shinomoto S. Kernel bandwidth optimization in spike rate estimation. *Journal of Computational Neuroscience*. 2010; 29(1–2):171–182. doi: [10.1007/s10827-009-0180-4](https://doi.org/10.1007/s10827-009-0180-4) PMID: [19655238](https://pubmed.ncbi.nlm.nih.gov/19655238/)
64. Snyder AC, Morais MJ, Willis CM, Smith MA. Global network influences on local functional connectivity. *Nature Neuroscience*. 2015; 18(5):736–743. doi: [10.1038/nn.3979](https://doi.org/10.1038/nn.3979) PMID: [25799040](https://pubmed.ncbi.nlm.nih.gov/25799040/)
65. Akaike H. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*. 1974; 19(6):716–723.
66. Amit DJ, Brunel N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex*. 1997; 7(3):237–252. doi: [10.1093/cercor/7.3.237](https://doi.org/10.1093/cercor/7.3.237) PMID: [9143444](https://pubmed.ncbi.nlm.nih.gov/9143444/)
67. van Vreeswijk C, Sompolinsky H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*. 1996; 274(5293):1724. 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/)
68. Mochol G, Hermoso-Mendizabal A, Sakata S, Harris KD, de la Rocha J. Stochastic transitions into silence cause noise correlations in cortical circuits. *Proceedings of the National Academy of Sciences*. 2015; 112(11):3529–3534. doi: [10.1073/pnas.1410509112](https://doi.org/10.1073/pnas.1410509112) PMID: [25739962](https://pubmed.ncbi.nlm.nih.gov/25739962/)
69. Steinmetz PN, Roy A, Fitzgerald P, Hsiao S, Johnson K, Niebur E. Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature*. 2000; 404(6774):187–190. doi: [10.1038/35004588](https://doi.org/10.1038/35004588) PMID: [10724171](https://pubmed.ncbi.nlm.nih.gov/10724171/)
70. Poulet JF, Petersen CC. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature*. 2008; 454(7206):881–885. doi: [10.1038/nature07150](https://doi.org/10.1038/nature07150) PMID: [18633351](https://pubmed.ncbi.nlm.nih.gov/18633351/)
71. Tan AY, Chen Y, Scholl B, Seidemann E, Priebe NJ. Sensory stimulation shifts visual cortex from synchronous to asynchronous states. *Nature*. 2014; 509(7499):226. doi: [10.1038/nature13159](https://doi.org/10.1038/nature13159) PMID: [24695217](https://pubmed.ncbi.nlm.nih.gov/24695217/)
72. Zeng HL, Alava M, Aurell E, Hertz J, Roudi Y. Maximum likelihood reconstruction for Ising models with asynchronous updates. *Physical Review Letters*. 2013; 110(21):210601. doi: [10.1103/PhysRevLett.110.210601](https://doi.org/10.1103/PhysRevLett.110.210601) PMID: [23745850](https://pubmed.ncbi.nlm.nih.gov/23745850/)
73. Dunn B, Roudi Y. Learning and inference in a nonequilibrium Ising model with hidden nodes. *Physical Review E*. 2013; 87(2):022127. doi: [10.1103/PhysRevE.87.022127](https://doi.org/10.1103/PhysRevE.87.022127) PMID: [23496479](https://pubmed.ncbi.nlm.nih.gov/23496479/)
74. Brillinger DR. Maximum likelihood analysis of spike trains of interacting nerve cells. *Biological Cybernetics*. 1988; 59(3):189–200. doi: [10.1007/BF00318010](https://doi.org/10.1007/BF00318010) PMID: [3179344](https://pubmed.ncbi.nlm.nih.gov/3179344/)
75. Chornoboy E, Schramm L, Karr A. Maximum likelihood identification of neural point process systems. *Biological Cybernetics*. 1988; 59(4–5):265–275. doi: [10.1007/BF00332915](https://doi.org/10.1007/BF00332915) PMID: [3196770](https://pubmed.ncbi.nlm.nih.gov/3196770/)
76. Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky E, et al. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*. 2008; 454(7207):995–999. doi: [10.1038/nature07140](https://doi.org/10.1038/nature07140) PMID: [18650810](https://pubmed.ncbi.nlm.nih.gov/18650810/)
77. Stevenson IH, Rebesco JM, Hatsopoulos NG, Haga Z, Miller LE, Kording KP. Bayesian inference of functional connectivity and network structure from spikes. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*. 2009; 17(3):203–213. doi: [10.1109/TNSRE.2008.2010471](https://doi.org/10.1109/TNSRE.2008.2010471) PMID: [19273038](https://pubmed.ncbi.nlm.nih.gov/19273038/)
78. Köster U, Sohl-Dickstein J, Gray CM, Olshausen BA. Modeling higher-order correlations within cortical microcolumns. *PLoS Computational Biology*. 2014; 10(7):e1003684. doi: [10.1371/journal.pcbi.1003684](https://doi.org/10.1371/journal.pcbi.1003684) PMID: [24991969](https://pubmed.ncbi.nlm.nih.gov/24991969/)
79. Montani F, Ince RA, Senatore R, Arabzadeh E, Diamond ME, Panzeri S. The impact of high-order interactions on the rate of synchronous discharge and information transmission in somatosensory cortex. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*. 2009; 367(1901):3297–3310. doi: [10.1098/rsta.2009.0082](https://doi.org/10.1098/rsta.2009.0082) PMID: [19620125](https://pubmed.ncbi.nlm.nih.gov/19620125/)
80. Ohiorhenuan IE, Mechler F, Purpura KP, Schmid AM, Hu Q, Victor JD. Sparse coding and high-order correlations in fine-scale cortical networks. *Nature*. 2010; 466(7306):617–621. doi: [10.1038/nature09178](https://doi.org/10.1038/nature09178) PMID: [20601940](https://pubmed.ncbi.nlm.nih.gov/20601940/)
81. Yu S, Yang H, Nakahara H, Santos GS, Nikolić D, Plenz D. Higher-order interactions characterized in cortical activity. *The Journal of Neuroscience*. 2011; 31(48):17514–17526. doi: [10.1523/JNEUROSCI.3127-11.2011](https://doi.org/10.1523/JNEUROSCI.3127-11.2011) PMID: [22131413](https://pubmed.ncbi.nlm.nih.gov/22131413/)
82. Tkačik G, Marre O, Mora T, Amodei D, Berry MJ II, Bialek W. The simplest maximum entropy model for collective behavior in a neural network. *Journal of Statistical Mechanics: Theory and Experiment*. 2013; 2013(03):P03011.
83. Shimazaki H, Sadeghi K, Ishikawa T, Ikegaya Y, Toyozumi T. Simultaneous silence organizes structured higher-order interactions in neural populations. *Scientific Reports*. 2015; 5:9821. doi: [10.1038/srep09821](https://doi.org/10.1038/srep09821) PMID: [25919985](https://pubmed.ncbi.nlm.nih.gov/25919985/)

84. Macke JH, Berens P, Ecker AS, Tolias AS, Bethge M. Generating spike trains with specified correlation coefficients. *Neural Computation*. 2009; 21(2):397–423. doi: [10.1162/neco.2008.02-08-713](https://doi.org/10.1162/neco.2008.02-08-713) PMID: [19196233](https://pubmed.ncbi.nlm.nih.gov/19196233/)
85. Froudarakis E, Berens P, Ecker AS, Cotton RJ, Sinz FH, Yatsenko D, et al. Population code in mouse V1 facilitates readout of natural scenes through increased sparseness. *Nature Neuroscience*. 2014; 17(6):851–857. doi: [10.1038/nn.3707](https://doi.org/10.1038/nn.3707) PMID: [24747577](https://pubmed.ncbi.nlm.nih.gov/24747577/)