

Exact minimax entropy models of large-scale neuronal activityChristopher W. Lynn ^{1,2,3,*} Qiwei Yu ⁴ Rich Pang,⁵ Stephanie E. Palmer ^{6,7} and William Bialek ^{3,4,8}¹*Department of Physics, Quantitative Biology Institute, and Wu Tsai Institute, Yale University, New Haven, Connecticut 06520, USA*²*Initiative for the Theoretical Sciences, The Graduate Center, City University of New York, New York 10016, USA*³*Joseph Henry Laboratories of Physics, Princeton University, Princeton, New Jersey 08544, USA*⁴*Lewis–Sigler Institute for Integrative Genomics, Princeton University, Princeton, New Jersey 08544, USA*⁵*Princeton Neuroscience Institute, Princeton University, Princeton, New Jersey 08544, USA*⁶*Department of Organismal Biology and Anatomy, University of Chicago, Chicago, Illinois 60637, USA*⁷*Department of Physics, University of Chicago, Chicago, Illinois 60637, USA*⁸*Center for Studies in Physics and Biology, Rockefeller University, New York 10065, USA*

(Received 18 December 2023; revised 18 March 2025; accepted 24 March 2025; published 19 May 2025)

In the brain, fine-scale correlations combine to produce macroscopic patterns of activity. However, as experiments record from larger and larger populations, we approach a fundamental bottleneck: the number of correlations one would like to include in a model grows larger than the available data. In this undersampled regime, one must focus on a sparse subset of correlations; the optimal choice contains the maximum information about patterns of activity or, equivalently, minimizes the entropy of the inferred maximum entropy model. Applying this “minimax entropy” principle is generally intractable, but here we present an exact and scalable solution for pairwise correlations that combine to form a tree (a network without loops). Applying our method to over 1000 neurons in the mouse hippocampus, we find that the optimal tree of correlations reduces our uncertainty about the population activity by 14% (over 50 times more than a random tree). Despite containing only 0.1% of all pairwise correlations, this minimax entropy model accurately predicts the observed large-scale synchrony in neural activity and becomes even more accurate as the population grows. The inferred Ising model is almost entirely ferromagnetic (with positive interactions) and exhibits signatures of thermodynamic criticality. Together, these results suggest that a large amount of information may be compressed into a small number of correlations between neurons, and provide the tools for identifying the most important correlations in other complex living systems.

DOI: [10.1103/PhysRevE.111.054411](https://doi.org/10.1103/PhysRevE.111.054411)**I. INTRODUCTION**

Understanding how collective behaviors emerge from webs of fine-scale interactions is a central goal in statistical mechanics approaches to networks of neurons [1–7]. At the same time, exploration of the brain has been revolutionized by experimental methods that monitor, simultaneously, the electrical activity of hundreds or even thousands of neurons [8–14]. One approach to connecting these new data with statistical physics models is maximum entropy, in which we construct the maximally disordered model that is consistent with measured expectation values [15]. In particular, it seems natural to build models that match the mean activity of individual neurons and the correlations between pairs of neurons. These pairwise maximum entropy models have been strikingly successful in describing collective behavior, not only in networks of real neurons, but also in the evolution of protein families, the dynamics of genetic networks, flocks of birds, and social networks [16–27].

But as experiments progress to record from larger and larger numbers of neurons, we face a combinatorial explosion.

Even if we focus on pairwise correlations, the number of correlations approaches the number of independent samples in modern experiments [8–14]. In this undersampled regime, one is forced to select only a sparse subset of the correlations to include in any model. While constructing an accurate model with only a small number of correlations may seem hopeless, one can draw inspiration from statistical physics, where effective descriptions of macroscopic phenomena can often ignore many of the microscopic details.

Here, given restrictions on the number and structure of correlations we can include in a model, we seek to identify those that contain the maximum information about system activity. We demonstrate that the optimal correlations are those that induce the maximum entropy model with minimum entropy [28]. Solving this minimax entropy problem is generally infeasible. But for pairwise correlations that form a tree (a network without loops), the entropy reduction decomposes into a sum over connected pairs; the advantages of tree structure in models of neural activity have been appreciated in other contexts [29]. This decomposition reduces the minimax entropy problem to a minimum spanning tree problem, which can be solved exactly and efficiently [17,30]. The result is a framework for uncovering the maximally informative tree of correlations in very large systems [31].

*Contact author: christopher.lynn@yale.edu

We apply our method to investigate the collective activity of $N \sim 1500$ neurons in the mouse hippocampus [32]. While most pairs of neurons are only weakly correlated, some rare pairs have mutual information orders of magnitude larger than average. By focusing on these exceptionally strong correlations, our minimax entropy model captures 50 times more information than a random tree and, despite containing only 0.1% of all pairwise correlations, produces realistic large-scale synchrony in activity. Moreover, the model becomes even more accurate as the population grows, providing hope for statistical physics descriptions of even larger systems.

The paper is organized as follows. In Sec. II we define the minimax entropy problem and present a solution for trees of pairwise correlations. In Sec. III we review a relatively recent experiment on large-scale recordings of neuronal activity in the mouse hippocampus. Inferring the optimal tree of correlations, in Sec. IV we demonstrate that the minimax entropy model predicts realistic patterns of synchronized activity. In Sec. V we investigate the structural properties of the optimal tree and the functional properties of the induced Ising model. In Sec. VI we show that the minimax entropy model becomes more accurate for larger populations, and then in Sec. VII we investigate the thermodynamic properties of the minimax entropy model, finding that the real system is poised at a special point in its phase diagram. Finally, in Sec. VIII we provide conclusions and outlook.

II. MINIMAX ENTROPY TREES

A. Maximum entropy models

Consider a system of N elements $i = 1, \dots, N$ with states $\mathbf{x} = \{x_i\}$, where x_i is the state of element i . From experiments, we have access to M samples of the system activity $\mathbf{x}^{(m)}$, where $m = 1, \dots, M$. Our knowledge about the system is defined by observables, which can be represented as expectation values,

$$\langle f(\mathbf{x}) \rangle_{\text{exp}} = \frac{1}{M} \sum_{m=1}^M f(\mathbf{x}^{(m)}), \quad (1)$$

where $f(\mathbf{x})$ is an arbitrary function of the state \mathbf{x} . For example, one could measure the average states of individual elements $\langle x_i \rangle_{\text{exp}}$ or the correlations among multiple elements $\langle x_i x_j \rangle_{\text{exp}}$, $\langle x_i x_j x_k \rangle_{\text{exp}}$, and so on. Given a set of K observables $\mathcal{O} = \{f_\nu(\mathbf{x})\}$, where $\nu = 1, \dots, K$, the most unbiased prediction for the distribution over states is the maximum entropy model [15,33]

$$P_{\mathcal{O}}(\mathbf{x}) = \frac{1}{Z} \exp \left[- \sum_{\nu=1}^K \lambda_\nu f_\nu(\mathbf{x}) \right], \quad (2)$$

where Z is the normalizing partition function, and the parameters λ_ν ensure that the model matches the experimental observations, such that

$$\langle f_\nu(\mathbf{x}) \rangle = \langle f_\nu(\mathbf{x}) \rangle_{\text{exp}}. \quad (3)$$

To have control over errors in the K expectation values, we must have $K \ll MN$. But as experiments record from larger systems, one is confronted with an explosion of possible observables. The total number of correlations grows

exponentially with N , and even the $K \propto N^2$ pairwise correlations violate the good sampling condition as N grows large. Thus, to avoid sampling problems, one must focus on a sparse subset of correlations. Here we arrive at the central question: Among a large set of observables, which should we choose to include in a model?

B. Minimax entropy principle

Suppose we want to find the set of observables \mathcal{O} that yields the most accurate description of the system. We can choose \mathcal{O} to maximize the log-likelihood of the model $P_{\mathcal{O}}$ or, equivalently, minimize the KL divergence with respect to the data $D_{\text{KL}}(P_{\text{exp}}||P_{\mathcal{O}})$. Due to the form of $P_{\mathcal{O}}$ in Eq. (2), the KL divergence simplifies to a difference in entropies

$$\begin{aligned} D_{\text{KL}}(P_{\text{exp}}||P_{\mathcal{O}}) &= \left\langle \log \frac{P_{\text{exp}}(\mathbf{x})}{P_{\mathcal{O}}(\mathbf{x})} \right\rangle_{\text{exp}} \\ &= \log Z + \frac{1}{\ln 2} \sum_{\nu} \lambda_\nu \langle f_\nu(\mathbf{x}) \rangle_{\text{exp}} - S_{\text{exp}} \\ &= S_{\mathcal{O}} - S_{\text{exp}}, \end{aligned} \quad (4)$$

where the final equality follows from Eq. (3). We therefore find that the optimal observables \mathcal{O} are the ones that minimize the entropy of the maximum entropy model $S_{\mathcal{O}}$. This is the ‘‘minimax entropy’’ principle, which was proposed 25 years ago but has received relatively little attention [28].

In addition to providing the best description of the data, the optimal observables \mathcal{O} can also be viewed as containing the maximum information about the system. If we begin by observing each element individually, then we only have access to the marginal distributions $P_i(x_i)$; in this case, the maximum entropy model is the independent distribution $P_{\text{ind}}(\mathbf{x}) = \prod_i P_i(x_i)$ with entropy S_{ind} . If, in addition to the marginals, we also observe some of the correlations between elements [34], this knowledge reduces our uncertainty about the system by an amount $I_{\mathcal{O}} = S_{\text{ind}} - S_{\mathcal{O}} \geq 0$. Thus, by minimizing $S_{\mathcal{O}}$, the optimal observables \mathcal{O} also maximize the information $I_{\mathcal{O}}$ contained in the observed correlations.

In practice, applying the minimax entropy principle poses two distinct challenges. First, for each set of observables $\mathcal{O} = \{f_\nu\}$, one must solve the traditional maximum entropy problem, that is, one must compute the parameters λ_ν such that the model $P_{\mathcal{O}}$ matches the expectations $\langle f_\nu(\mathbf{x}) \rangle_{\text{exp}}$ in the data. Second, one must repeat this calculation for all sets of observables \mathcal{O} to find the one that minimizes the entropy $S_{\mathcal{O}}$. This search process is generally intractable. In what follows, we study a class of observables that admits an exact and efficient solution, enabling statistical physics models of very large systems.

C. Trees of pairwise correlations

For simplicity, we focus on binary variables $x_i = 0, 1$, for which the marginals $P_i(x_i)$ are defined by the averages $\langle x_i \rangle$. In the search for sources of order in a system, one might begin with the simplest correlations: those between pairs of elements $\langle x_i x_j \rangle$. In populations of $N \sim 100$ neurons, one often has sufficient data to fit all the pairwise correlations, which can be very effective in capturing key features of the

collective activity [16,18,19]. But this corresponds to $K \propto N^2$ constraints, and at large N we will violate the good sampling condition $K \ll NM$. To avoid undersampling, we are forced to select a sparse subset of pairwise correlations, which can be visualized as a network \mathcal{G} with edges defining the observed correlations between variables. Each network induces a maximum entropy model,

$$P_{\mathcal{G}}(\mathbf{x}) = \frac{1}{Z} \exp \left[\sum_{(ij) \in \mathcal{G}} J_{ij} x_i x_j + \sum_i h_i x_i \right], \quad (5)$$

where the parameters h_i and J_{ij} enforce the constraints on $\langle x_i \rangle$ and $\langle x_i x_j \rangle$ in \mathcal{G} , respectively. The minimax entropy principle tells us that we should find the network \mathcal{G} (within some allowed set) that produces the maximum entropy model $P_{\mathcal{G}}$ with minimum entropy $S_{\mathcal{G}}$.

In statistical physics, calculations are difficult in part due to feedback loops. By eliminating loops, many statistical physics models become tractable, as in one-dimensional systems or on Bethe lattices [35]. In the Ising model—which is equivalent to Eq. (5)—if the interactions J_{ij} lie on a tree \mathcal{T} (or a network without loops), then one can efficiently compute the partition function Z and all statistics of interest (see Appendix A). Inverting this procedure, one can begin with the averages $\langle x_i \rangle$ and the correlations $\langle x_i x_j \rangle$ on a tree \mathcal{T} and analytically derive the maximum entropy parameters [17,30]:

$$J_{ij} = \ln \left[\frac{\langle x_i x_j \rangle (1 - \langle x_i \rangle - \langle x_j \rangle + \langle x_i x_j \rangle)}{(\langle x_i \rangle - \langle x_i x_j \rangle)(\langle x_j \rangle - \langle x_i x_j \rangle)} \right], \quad (6)$$

$$h_i = \ln \frac{\langle x_i \rangle}{1 - \langle x_i \rangle} + \sum_{j \in \mathcal{N}_i} \ln \left[\frac{(1 - \langle x_i \rangle)(\langle x_i \rangle - \langle x_i x_j \rangle)}{\langle x_i \rangle (1 - \langle x_i \rangle - \langle x_j \rangle + \langle x_i x_j \rangle)} \right], \quad (7)$$

where \mathcal{N}_i represents the neighbors of i in \mathcal{T} (see Appendix B). Since each tree contains $N - 1$ correlations, the total number of observables is $K = 2N - 1$, and so we are well sampled if the number of independent samples obeys $M \gg 2$.

Equations (6) and (7) solve the maximum entropy problem for the distribution $P_{\mathcal{T}}$, but we still need to search over all of the N^{N-2} trees to find the one that minimizes the entropy $S_{\mathcal{T}}$. This search simplifies significantly by noticing that the information $I_{\mathcal{T}}$ decomposes into a sum over the connections (ij) in \mathcal{T} ,

$$I_{\mathcal{T}} = S_{\text{ind}} - S_{\mathcal{T}} = \sum_{(ij) \in \mathcal{T}} I_{ij}, \quad (8)$$

where I_{ij} is the mutual information between i and j (see Appendix C) [17,30]. Note that for pairs $(ij) \in \mathcal{T}$, the mutual information I_{ij} is the same in the model and the data, so we can compute the entropy $S_{\mathcal{T}}$ directly from the data without constructing the model itself.

Equation (8) tells us that the tree with the minimum entropy $S_{\mathcal{T}}$ is the one with the largest total mutual information $I_{\mathcal{T}}$. Identifying this optimal tree is a minimum spanning tree problem [17], which can be solved efficiently using a number of different algorithms [36]. To begin, one computes the mutual information I_{ij} between all elements [Fig. 1(a)]. One can then grow the optimal tree by greedily connecting the element i in the tree to the new element j with the largest mutual

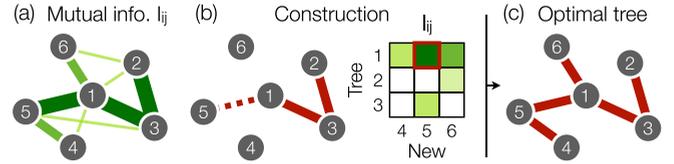


FIG. 1. Constructing the optimal tree. (a) Visualization of the mutual information I_{ij} (edges) between elements in a system (nodes), with darker, thicker edges reflecting larger I_{ij} . (b) Illustration of Prim’s algorithm. At each step, we consider the mutual information I_{ij} between elements in the tree and those not yet connected (matrix). We then connect the two elements with the largest I_{ij} (dashed edge) and repeat until all elements have been added. (c) Optimal tree that maximizes $I_{\mathcal{T}}$ and minimizes $S_{\mathcal{T}}$.

information I_{ij} [Fig. 1(b)]; this is Prim’s algorithm, which runs in $O(N^2)$ time [Fig. 1(c)]. Thus, by restricting to trees of pairwise correlations, we can solve the minimax entropy problem exactly, even at very large N .

To see the strength of this approach, consider data generated by an Ising model on a tree. Recalling Eq. (4), we have $S_{\mathcal{T}} \geq S_{\text{exp}}$, where S_{exp} is the entropy of the data. Since the data is generated by a tree, our minimax entropy approach is guaranteed to find the tree \mathcal{T} that saturates the bound $S_{\mathcal{T}} = S_{\text{exp}}$. Returning to Eq. (4), we see that $D_{\text{KL}}(P_{\text{exp}} || P_{\mathcal{T}}) = S_{\mathcal{T}} - S_{\text{exp}} = 0$, which tells us that we have identified the true underlying tree of interactions and our model is exact.

III. LARGE-SCALE NEURONAL ACTIVITY

We ultimately seek to explain the collective behaviors of very large networks. However, each tree only contains a vanishingly small fraction $2/N$ of all pairwise correlations; and even if we have access to all of the pairwise statistics, there’s still no guarantee of success. Can such a sparse set of observations capture something important about the system as a whole?

To answer this question, we consider patterns of electrical activity in $N = 1485$ neurons in the hippocampus of a mouse, recorded in a recent experiment [32]. Mice are genetically engineered so that their neurons express a protein whose fluorescence is modulated by calcium concentration, which in turn follows the electrical activity of the cells. This fluorescence is recorded using a scanning two-photon microscope as the mouse runs in a virtual environment. The signal from each cell consists of a quiet background punctuated by short bursts of activity [18], providing a natural binarization into active ($x_i = 1$) or silent ($x_i = 0$) within each video frame [Fig. 2(a)]. Capturing images at 30 Hz for 39 min yields $M \sim 7 \times 10^4$ samples of the collective state $\mathbf{x} = \{x_i\}$, but these are not all independent. Nonetheless, we can still estimate the mutual information I_{ij} with small errors after correcting for finite-data effects (see Appendix D).

Among all $\sim 10^6$ pairs of neurons, only 9% exhibit significant mutual information with values shown in Fig. 2(b); see Appendix D. We see that a small number of pairs contain orders of magnitude more information than average ($\bar{I} = 2.9 \times 10^{-4}$ bits). This heavy-tailed distribution provides hope for a tree of correlations that contains much more

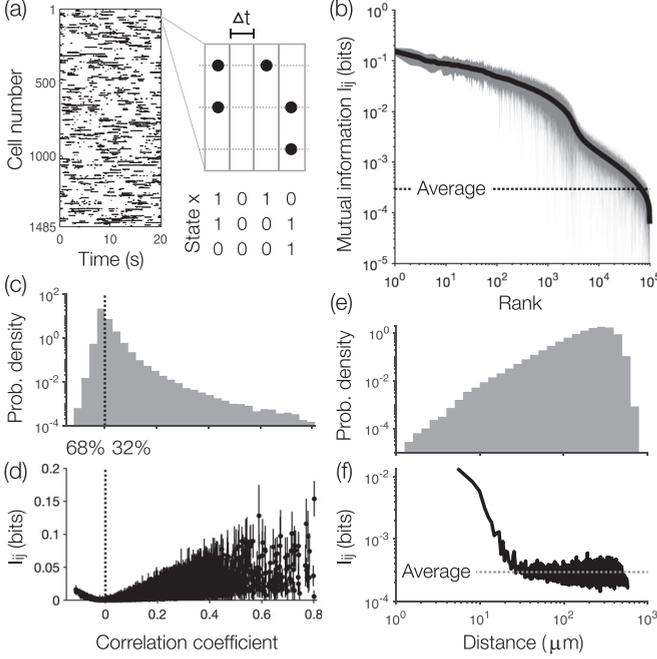


FIG. 2. Collective activity in a large population of neurons. (a) Time series of neuronal activity in the mouse hippocampus, where each dot represents an active neuron (see [32] for experimental details). States $\mathbf{x} = \{x_i\}$ represent the population activity within a single window of width $\Delta t = 1/30$ s. (b) Ranked order of significant mutual information I_{ij} in the population. Solid line and shaded region reflect estimates and errors (two standard deviations) after correcting for finite data (see Appendix D). (c) Distribution of correlation coefficients over neuron pairs, with percentages indicating the fraction of positively and negatively correlated pairs. (d) Mutual information I_{ij} vs correlation coefficient, where each point represents a distinct neuron pair. Estimates and errors are the same in (b). (e) Distribution of physical distances between neurons. (f) Average mutual information I_{ij} as a function of physical distance, computed in bins that contain 500 pairs each; note that individual pairs vary widely around this average.

information than typical $I_{\mathcal{T}} \gg (N-1)\bar{I}$. Additionally, while most pairs of cells are negatively correlated [Fig. 2(c)], the strongest mutual information corresponds to positive correlations [Fig. 2(d)]. And while most neurons are far from one another [Fig. 2(e)], larger values of I_{ij} are concentrated among pairs of cells that are close to one another, as can be seen by plotting the mean mutual information as a function of distance [Fig. 2(f)]. Together, these observations suggest that a backbone of positively correlated and physically proximate neurons may provide a large amount of information about the collective neural activity.

IV. PREDICTIONS OF MINIMAX ENTROPY MODEL

Constructing the minimax entropy tree (Fig. 1), we find that it captures $I_{\mathcal{T}} = 26.2$ bits of information. This reduces our uncertainty about the population activity by $I_{\mathcal{T}}/S_{\text{ind}} = 14.4\%$, which is equivalent to freezing the states of 214 randomly selected neurons. For comparison, we consider two additional networks: (i) a random tree, which represents a

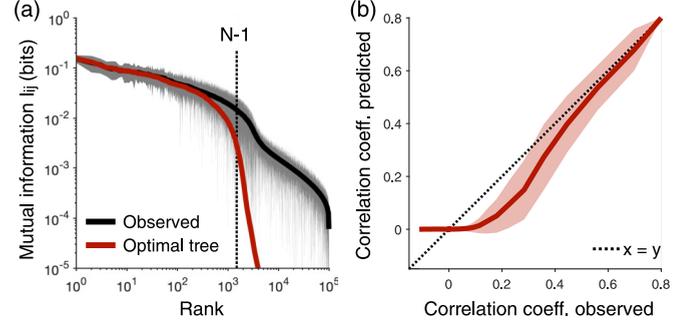


FIG. 3. Predicting pairwise statistics. (a) Ranked order of significant mutual information in the population (black), two-standard-deviation errors (shaded region), and prediction of the minimax entropy model (red). (b) Correlation coefficients predicted in the model vs those in the data, with dashed lines indicating equality. All pairs are divided evenly into bins along the x axis, with solid lines and shaded regions reflecting means and errors (standard deviations) within bins.

typical collection of correlations, and (ii) the tree of minimum physical distances, which reflects the fact that neighboring neurons are more likely to be strongly correlated [Fig. 2(f)]. The optimal tree captures over twice as much information as the minimum distance tree and over 50 times more than random.

While each model $P_{\mathcal{T}}$ is defined to match a sparse subset of the observed correlations $\langle x_i x_j \rangle$ (and thus mutual information I_{ij}) for $(ij) \in \mathcal{T}$, we can ask what $P_{\mathcal{T}}$ predicts for all pairs of neurons (see Appendix E). We note that the optimal tree does not simply match the largest $N-1$ values of I_{ij} ; in general, these will form loops. Yet we find that the minimax entropy model still predicts the distribution of I_{ij} within experimental error for the top $\sim N$ values [Fig. 3(a)]. Indeed, we find that the model captures the strong correlations in the population [Fig. 3(b)]; this accuracy decreases significantly for the minimum distance and random trees (see Appendix F). As expected, the optimal tree underpredicts the strengths of weak and negative correlations [Fig. 3(b)]. Although these correlations may seem unimportant individually, we note that they comprise the vast majority of neuron pairs [Fig. 2(c)].

With knowledge of only $2/N \sim 0.1\%$ of the pairwise correlations, can the optimal tree capture collective behavior in the system? In neuronal populations (and other complex systems), one key collective property is synchronized activity [16,19,20,27], which is characterized by the probability $P(K)$ that K out of the N neurons are simultaneous active. If the neurons were independent, this distribution would be approximately Gaussian at large N (Fig. 4, dashed). But in real populations, the dependencies among neurons leads to a much broader distribution (Fig. 4, black), with moments of extreme synchrony in both activity (large K) and silence (small K). If one builds a model from pairwise correlations chosen at random, then the distribution $P(K)$ is almost indistinguishable from that of an independent system (see Appendix F). By contrast, the optimal tree captures most of this collective behavior [31], correctly predicting $\gtrsim 100$ -fold increases in the probabilities that $K \gtrsim 50$ neurons are active in the same small time bin (Fig. 4, red). Although the detailed patterns of

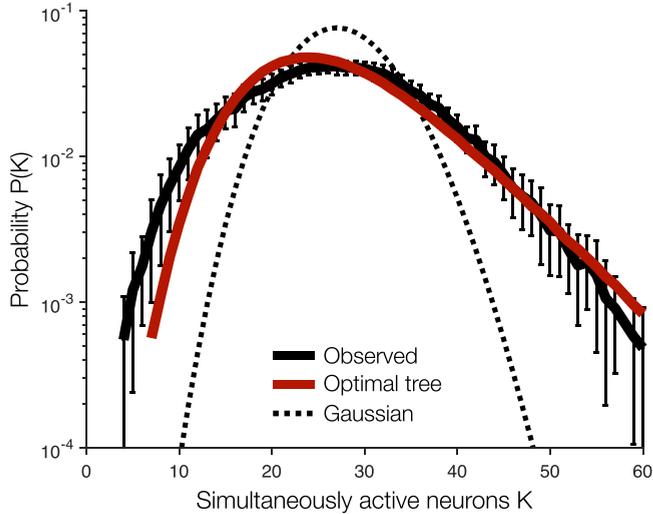


FIG. 4. Predicting synchronized activity. Distribution $P(K)$ of the number of simultaneously active neurons K in the data (black), predicted by the minimax entropy tree (red), and the Gaussian distribution for independent neurons with mean and variance $\langle K \rangle_{\text{exp}}$ (dashed). We confirm that the Gaussian matches the data after shuffling the activity of each neuron in time. To estimate $P(K)$ and error bars (two standard deviations), we first split the data into 1-min blocks to preserve dependencies between consecutive samples. We then select one-third of these blocks at random and repeat 100 times. For each subsample, we compute the optimal tree \mathcal{T} and predict $P(K)$ using a Monte Carlo simulation of the model $P_{\mathcal{T}}$.

activity in the system are shaped by competing interactions that are missing from our optimal tree, this shows that large-scale synchrony can emerge from a sparse network of the strongest correlations.

V. STRUCTURE OF MINIMAX ENTROPY MODEL

A. Ising structure

To understand the nature of the optimal tree, we can study the minimax entropy model $P_{\mathcal{T}}$ itself, which, as discussed above, is equivalent to an Ising model from statistical physics. This mapping is made concrete by considering a system of spins $\sigma_i = 2x_i - 1 \in \pm 1$ with Ising interactions $J_{ij}^1 = J_{ij}/4$ and local fields $h_i^1 = h_i/2 + \sum_j J_{ij}^1$, where J_{ij} and h_i are defined in Eqs. (6) and (7). If the interaction J_{ij}^1 is positive (negative), then activity in neuron j leads to activity (silence) in neuron i , and vice versa. For random trees, the interactions J_{ij}^1 are nearly evenly split between positive and negative [Fig. 5(a)]; this is consistent with previous investigations of fully connected models in populations of $N \sim 100$ neurons [16,19,20]. Meanwhile, we recall that the largest mutual information in the population belongs to positively correlated neurons [Fig. 2(d)]. Accordingly, the optimal tree has interactions that are almost exclusively positive [Fig. 5(a)]. We have arrived, perhaps surprisingly, at a traditional Ising ferromagnet.

While the interactions J_{ij}^1 define effective influences between neurons, the local fields represent individual biases toward activity ($h_i^1 > 0$) or silence ($h_i^1 < 0$). For random trees,

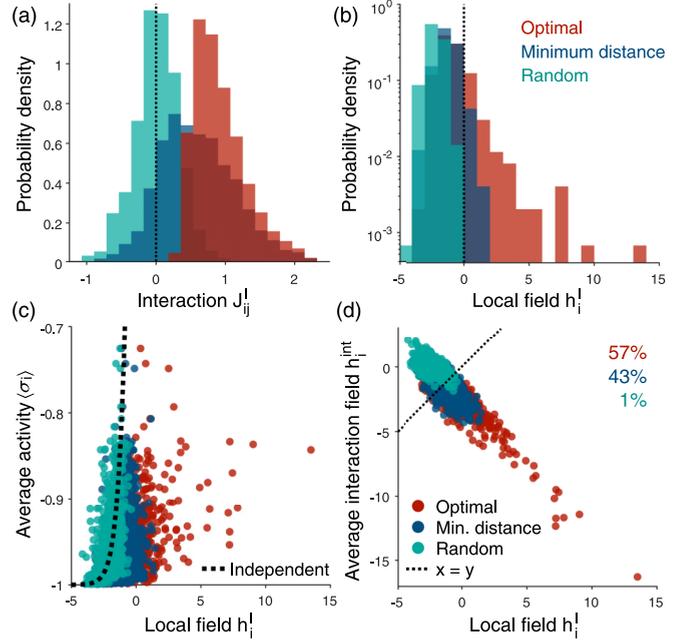


FIG. 5. Maximum entropy models of large-scale activity. (a, b) Distributions of Ising interactions J_{ij}^1 (a) and local fields h_i^1 (b) in the optimal tree (red), the minimum distance tree (blue), and a random tree (cyan). (c) Average activities $\langle \sigma_i \rangle$ vs local fields h_i^1 , where each point represents an individual neuron. Dashed line illustrates the independent prediction $\langle \sigma_i \rangle = \tanh h_i^1$. (d) Average interaction fields $h_i^{\text{int}} = \sum_j J_{ij}^1 \langle \sigma_j \rangle$ vs local fields h_i^1 . Percentages indicate the proportion of neurons for which $h_i^{\text{int}} < h_i^1$ (dashed line indicates equality).

all of the local fields are negative [Fig. 5(b)], reflecting the fact that neurons are more likely to be silent than active. But in the optimal tree, we see that some neurons are counterintuitively biased toward activity with $h_i^1 > 0$ [Fig. 5(b)]. These positive biases stand in competition with the positive interactions in the model, which, because neurons favor silence, tend to induce silence in the population.

To understand the effects of interactions on individual cells, we note that the average activity of an independent neuron i is fully defined by h_i^1 through the relation $\langle \sigma_i \rangle = \tanh h_i^1$. Since random trees contain only weak correlations, the neuronal activity closely tracks this independent prediction [Fig. 5(c)]. As interactions become increasingly positive in the minimum distance and optimal trees, the alignment of neighboring neurons produces average activities that are significantly lower than one would expect from local fields alone [Fig. 5(c)]. For each neuron i , the competition between internal biases and interactions is made clear by comparing the local field h_i^1 to the average influence due to interactions $h_i^{\text{int}} = \sum_j J_{ij}^1 \langle \sigma_j \rangle$. In random trees, only 1% of neurons are dominated by interactions, such that $h_i^{\text{int}} < h_i^1$ [Fig. 5(d)]; this proportion increases to 43% in the minimum distance tree and 57% in the optimal tree [Fig. 5(d)]. So despite the fact that the tree structure constrains each cell to only interact with two others in the entire population (on average), most neurons in the optimal tree are driven more strongly by interactions than internal biases.

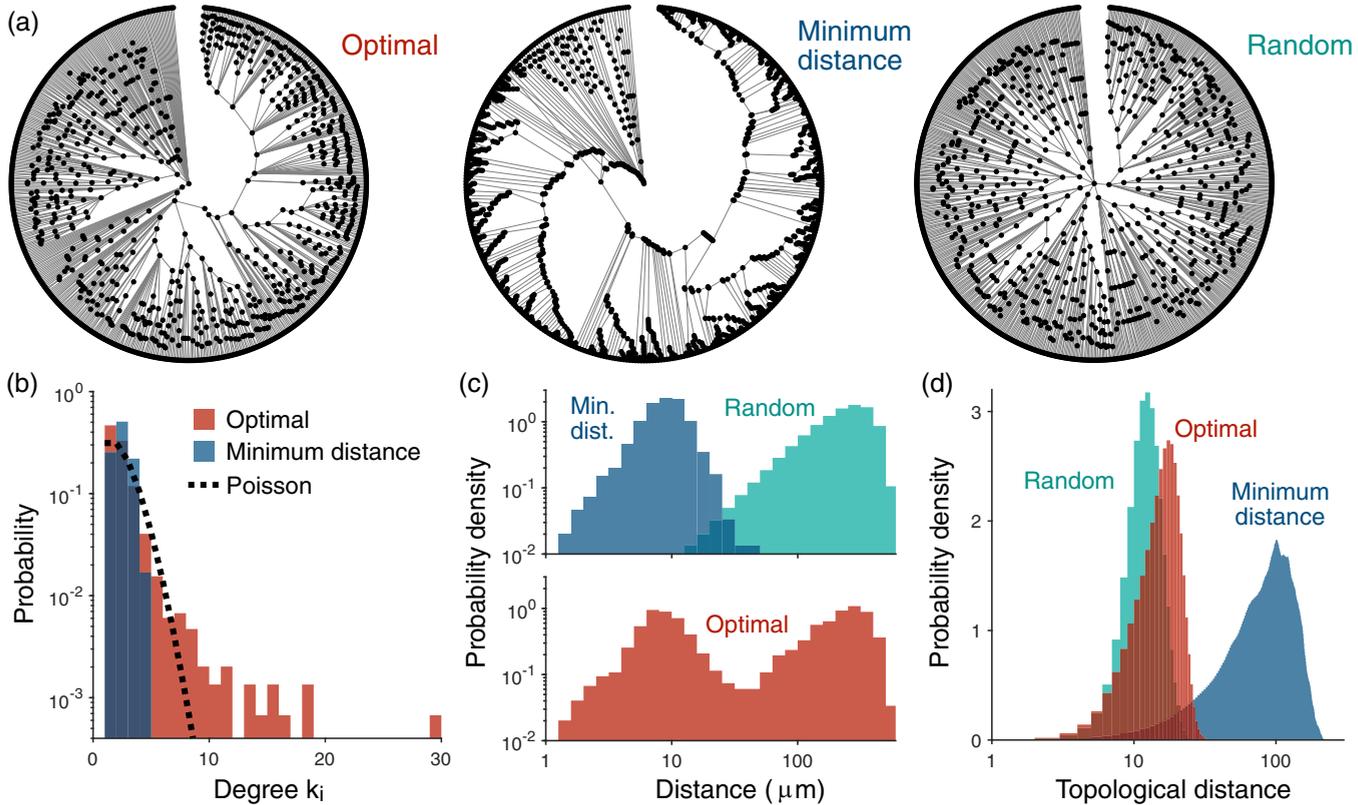


FIG. 6. Network structure of optimal tree. (a) Illustrations of the minimax entropy tree (left), the minimum distance tree (middle), and a random tree (right). In all networks, the central neuron has the largest degree (number of connections), and those with one connection are located on the perimeter with distance from the central neuron decreasing in the clockwise direction. (b) Distributions of degrees in the optimal and minimum distance trees. Dashed line indicates the Poisson distribution of random trees. (c) Distributions of physical distances among the connections in each tree. (d) Distributions of topological distances (or the lengths of shortest paths) across all pairs of neurons in each tree.

B. Network structure

In addition to the functional properties of the model $P_{\mathcal{T}}$, we can also study the graph structure of the optimal tree \mathcal{T} . To visualize each tree, we place the cell with the most connections (or largest degree $k_i = |\mathcal{N}_i|$) at the center and all of the cells with single connections ($k_i = 1$) around the perimeter [Fig. 6(a)]. For random trees, the distribution of degrees is Poisson [Fig. 6(b)], preventing the emergence of high-degree hub nodes. Degrees are even more sharply peaked in the minimum distance tree, such that we do not observe a single neuron with more than four connections [Fig. 6(b)]. By contrast, the optimal tree has a much broader degree distribution, with a central neuron that connects to 29 other cells in the population [Fig. 6(b)]. Such hub nodes are frequently observed in the brain's physical connectivity [37–39] and are thought to play an important role in facilitating communication [40].

By maximizing information about the population, one might hope that the optimal tree captures features of the true interactions between neurons. In the brain, demands on communication are constrained by energetic costs [41]. Networks have evolved to balance efficient communication (minimizing the number of steps between cells in the network, known as topological distance) with energetic efficiency (minimizing the physical lengths of connections) [38,42]. These pressures are in direct competition: Networks with

physically local connections form latticelike structures with long topological distances, and networks with short topological distances (known as the small-world property [43]) require physically long-range connections. Indeed, in the minimum distance tree, which is composed of the physically shortest connections [Fig. 6(c)], communication between two neurons requires ~ 100 intermediate cells on average [Fig. 6(d)]; and random trees, which are known to produce short topological distances [Fig. 6(d)], are mostly composed of long-range connections [Fig. 6(c)]. Meanwhile, the minimax entropy model identifies connections that are much shorter than average [Fig. 6(c)] while simultaneously maintaining small-world structure [Fig. 6(d)], just as observed in real neuronal networks [42].

VI. SCALING WITH POPULATION SIZE

Thus far, we have focused on a single population of $N \sim 1500$ neurons. But as experiments advance to record from even larger populations, how does the minimax entropy model scale with N ? To answer this question, in the spirit of Ref. [19], one can imagine growing a contiguous population centered at a single neuron [Fig. 7(a)] and computing the optimal tree for increasing population sizes. Due to the efficiency of our model, we can repeat this process starting from each of the different neurons and average over the results. As the

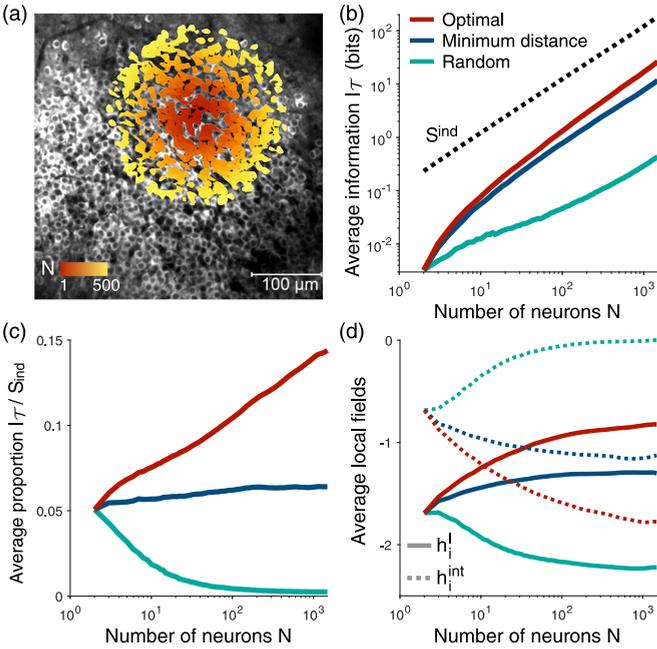


FIG. 7. Scaling of the minimax entropy model. (a) Illustration of our growth process superimposed on a fluorescence image of the $N = 1485$ neurons in the mouse hippocampus. Starting with a single neuron i , we grow the population of N neurons closest to i (red to yellow). We then repeat this process for each neuron and average the results. (b) Information $I_{\mathcal{T}}$ captured by different trees as a function of population size N . Dashed line indicates the independent entropy S_{ind} . (c) Fraction of the independent entropy $I_{\mathcal{T}}/S_{\text{ind}}$ explained by different trees as a function of N . (d) Average Ising local fields h_i^l (solid) and interaction fields h_i^{int} (dashed) for the maximum entropy models on different trees $P_{\mathcal{T}}$.

population grows, the independent entropy S_{ind} must increase extensively (that is, linearly with N) on average [Fig. 7(b)]. Since each tree contains $N - 1$ correlations, one might also expect the information $I_{\mathcal{T}}$ of any tree to scale extensively. However, we find that the scaling of $I_{\mathcal{T}}$ with population size depends critically on which correlations we use in building the tree [Fig. 7(b)].

If the information $I_{\mathcal{T}}$ grows extensively, then the model $P_{\mathcal{T}}$ explains a constant proportion of the independent entropy $I_{\mathcal{T}}/S_{\text{ind}}$ across different population sizes. Indeed, because the properties of the closest neurons do not change as the population grows (on average), the minimum distance tree captures a nearly constant $\sim 6\%$ of the independent entropy [Fig. 7(c)]. By contrast, since the mutual information I_{ij} between neurons tends to decrease with physical distance [Fig. 2(f)], the average mutual information \bar{I} in a spatially contiguous population decreases with N . Thus, the typical information in a random tree grows subextensively with the population size [Fig. 7(b)], and the fractional information $I_{\mathcal{T}}/S_{\text{ind}}$ vanishes [Fig. 7(c)].

But even though the average mutual information \bar{I} decreases, as the population grows we uncover more of the exceptionally large mutual information I_{ij} in the tail of the distribution [Fig. 2(b)]. By identifying these highly informative correlations, the optimal tree accumulates a superextensive amount of information $I_{\mathcal{T}}$ [Fig. 7(b)], thus capturing a greater proportion of the independent entropy as N increases

[Fig. 7(c)]. This increased explanatory power is underpinned by stronger interactions and weaker local fields in the Ising network [Fig. 7(d)]. There is no sign that the trend in Fig. 7(c) is saturating at $N \sim 10^3$, suggesting that our minimax entropy framework may become even more effective for larger populations.

VII. THERMODYNAMICS AND SIGNATURES OF CRITICALITY

As discussed above, each tree of observed correlations \mathcal{T} generates a maximum entropy model $P_{\mathcal{T}}$ [Eq. (5)], which in turn is equivalent to a system of Ising spins. This mapping from experimental observations to statistical physics gives us the opportunity to ask whether the model $P_{\mathcal{T}}$ occupies a special place in the space of possible models. In statistical mechanics, equilibrium systems are described by the Boltzmann distribution,

$$P(\mathbf{x}) = \frac{1}{Z} \exp \left[-\frac{1}{T} E(\mathbf{x}) \right], \quad (9)$$

where T is the temperature of the system and $E(\mathbf{x})$ is the Hamiltonian, which defines the energy of state \mathbf{x} . For a given tree \mathcal{T} , we notice that $P_{\mathcal{T}}$ defines a Boltzmann distribution with temperature $T = 1$ and energy

$$E(\mathbf{x}) = - \sum_{(ij) \in \mathcal{T}} J_{ij} x_i x_j - \sum_i h_i x_i, \quad (10)$$

where J_{ij} and h_i are defined in Eqs. (6) and (7). Note that we do not assume the experimental system itself is in equilibrium; this correspondence is purely mathematical.

By perturbing the temperature away from $T = 1$, we can probe at least one slice through the space of possible networks [20]. For each value of T , we arrive at a hypothetical system $P_{\mathcal{T}}(\mathbf{x}; T)$ with average activities $\langle x_i \rangle_T$ and correlations $\langle x_i x_j \rangle_T$ that are no longer constrained to match experimental observations. Consider the minimax entropy tree, which (as discussed in Sec. V) produces a ferromagnetic Ising model with nearly all positive interactions J_{ij}^l [Fig. 5(a)]. At high temperatures $T \gg 1$, fluctuations destroy the preference for silence over activity, and the system approaches the average activity $\frac{1}{N} \sum_i \langle x_i \rangle_T = 0.5$ [Fig. 8(a)]. Meanwhile, at low temperatures $T \ll 1$, activity vanishes as the network freezes into the all-silent ground state $\mathbf{x} = 0$ [Fig. 8(a)]. In both limits, all of the information contained in correlations is lost.

As the temperature decreases, most systems experience a gradual transition from disorder to order. But for certain combinations of parameters J_{ij} and h_i , a small change in the temperature T can lead to a large change in the behavior of the system, and as N becomes large, this transition becomes sharp [20,35,44]. Such phase transitions mark a critical point in the space of possible systems, with the Ising ferromagnet as the canonical example [45]. In the optimal tree, as the temperature increases just above $T = 1$, the positive interactions lead to a much steeper increase in activity than an independent system, and this transition grows even sharper for larger populations [Fig. 8(a)].

We emphasize that at any finite N there is no true critical point, but $N \sim 1000$ may be large enough that the idealization $N \rightarrow \infty$ is useful. Since we are studying models defined on

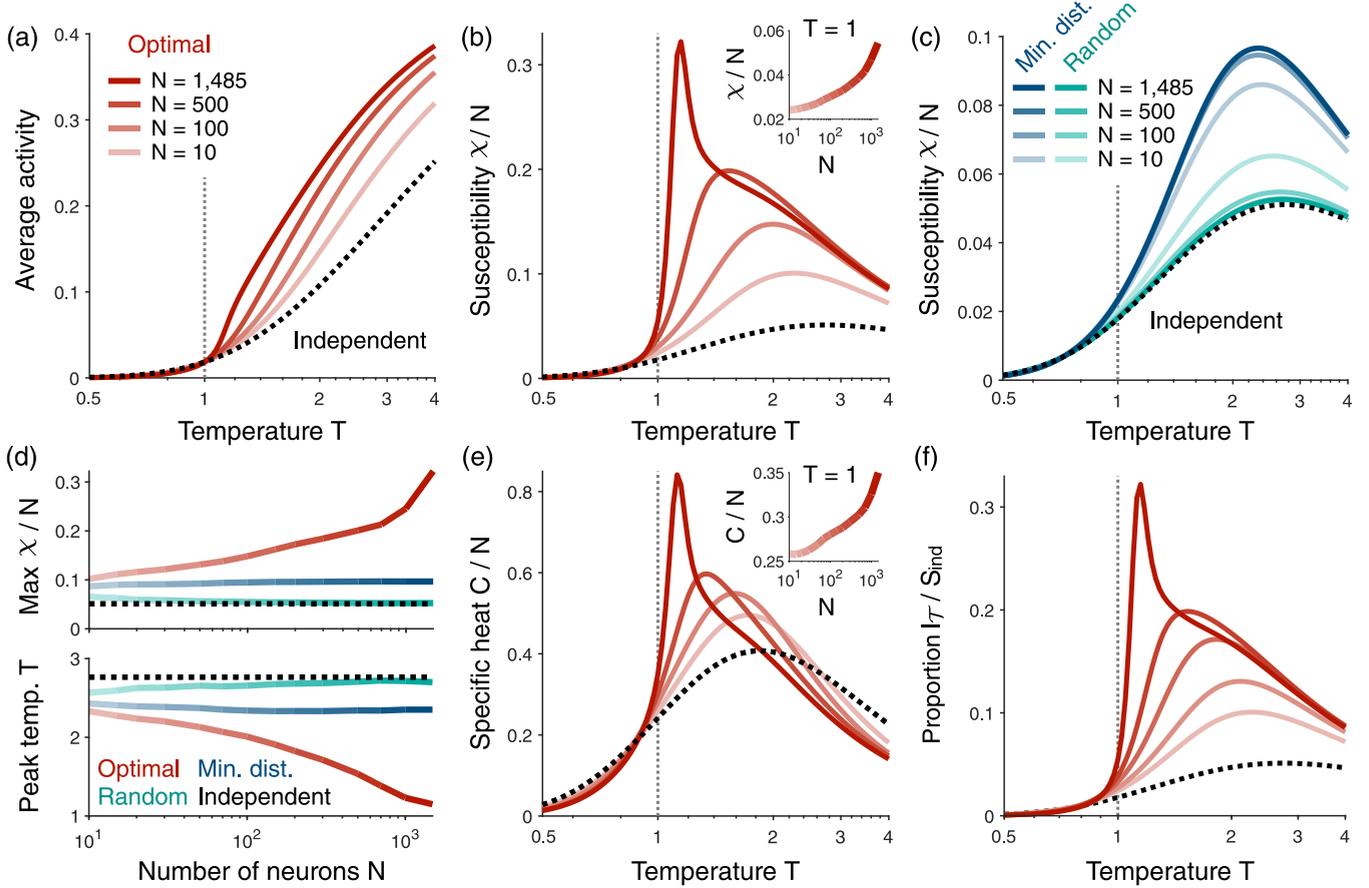


FIG. 8. Thermodynamics of the minimax entropy model. (a) Average activity $\frac{1}{N} \sum_i \langle x_i \rangle_T$ as a function of temperature T in the optimal tree \mathcal{T} for populations of increasing size N and for independent neurons (dashed). (b) Normalized susceptibility χ/N vs temperature T for the same models as (a). Inset shows the increase in χ/N with population size N for the true minimax entropy models ($T = 1$). (c) Normalized susceptibility vs T for the minimum distance tree (blue) and random trees (cyan) across different population sizes. (d) Maximum value of the normalized susceptibility χ/N (top) and peak temperature T (bottom) as functions of the population size N for the optimal tree (red), the minimum distance tree (blue), random trees (cyan), and independent neurons (dashed). (e), (f) Specific heat C/N (e) and information fraction I_T/S_{ind} (f) vs temperature T for the minimax entropy models in (a) and (b). Inset in (e) shows the increase in C/N with population size N for the true minimax entropy models ($T = 1$). In all panels, darker lines reflect populations of increasing size N , constructed using the method in Fig. 7(a) and averaged over 100 random initializations. In (a)–(e), dashed lines represent independent neurons.

trees, there are also subtleties about how one would construct the thermodynamic limit, since such a large fraction of sites are on the boundary [46]. For our purposes, the interesting question is whether real networks of neurons are in any sense at special points in the space of possible networks. One way in which this could happen is if parameters are set so that simple macroscopic quantities have near-extremal values.

One example of a macroscopic quantity that provides a global measure of collective behavior is the total susceptibility of the mean activity to changes in the bias fields,

$$\chi = \sum_{ij} \frac{d\langle x_i \rangle_T}{dh_j} = \frac{1}{T} \sum_{ij} (\langle x_i x_j \rangle_T - \langle x_i \rangle_T \langle x_j \rangle_T), \quad (11)$$

where the rewriting in terms of connected correlations can be derived from the Boltzmann distribution [Eq. (9)]. We recall that at conventional critical points we would see a divergence of χ/N as $N \rightarrow \infty$. At both high and low temperatures, correlations are destroyed and the susceptibility vanishes [Fig. 8(b)]. However, at intermediate temperatures,

the susceptibility exhibits a peak that becomes sharper as the system grows, even after normalizing by the population size N [Fig. 8(b)]. Moreover, as N increases, the peak temperature decreases toward $T = 1$, corresponding to the true minimax entropy model $P_{\mathcal{T}}$. By contrast, the minimum distance and random trees undergo smooth transitions from disorder to order [Fig. 8(c)], with the maximum susceptibility and peak temperatures remaining approximately constant across all population sizes N [Fig. 8(d)].

In addition to the susceptibility χ , we also observe a dramatic peak in the specific heat C/N [Fig. 8(e)], where

$$C = \frac{d\langle E(\mathbf{x}) \rangle_T}{dT} \quad (12)$$

is the heat capacity (see Appendix G). Although there is no meaning to “heat” in this system, because the specific heat is related to the variance in energy, we can think of the peak in C as being a peak in the dynamic range of (log) probabilities across the states of the network. These divergences in the susceptibility and heat capacity also align with a sharp peak

in the information fraction $I_{\mathcal{T}}/S_{\text{ind}}$ [Fig. 8(f)], with larger systems becoming even more strongly correlated. Together, these results indicate that the true minimax entropy model $P_{\mathcal{T}}$ is poised near a special point in the space of models $P_{\mathcal{T}}(\mathbf{x}; T)$, where small changes in parameter values can produce large changes in the collective behavior of the system.

VIII. CONCLUSIONS

The maximum entropy principle provides the most unbiased mapping from experimental observations to statistical physics models. Over the past two decades, this link has proven useful in understanding the emergence of collective behaviors in populations of neurons and other complex living systems [16–27]. Less widely emphasized is the fact that there is not a single maximum entropy model but rather a landscape of possible models depending on what features of the system we choose to constrain. Quite generally, we should choose the features that are most informative—the ones that minimize the entropy of the maximum entropy model—leading to the minimax entropy principle [28]. As experiments record from larger and larger populations of neurons [8–14], we enter an undersampled regime in which selecting a limited number of maximally informative features is not only conceptually appealing, but also practically necessary.

While the minimax entropy problem is generally intractable, here we make progress in two steps. First, we build upon previous work by constraining mean activities and pairwise correlations, resulting in models that are equivalent to systems of Ising spins. Second, taking inspiration from the Bethe lattice, we focus only on trees of correlations, or sparse networks without loops. Under these restrictions, we solve the minimax entropy problem exactly, identifying the optimal tree in quadratic time [17,30]. The result is a nontrivial family of statistical physics models that can be constructed very efficiently for large neuronal populations.

It is far from obvious that these models can capture any of the essential collective behavior in real networks. To answer this question, we study a population of $N \sim 1500$ neurons in the mouse hippocampus [32], identifying the maximally informative tree of pairwise correlations (Figs. 5 and 6). Despite constraining only one correlation per neuron, this minimax entropy model accounts for 14% of the independent entropy (over 50 times more than random trees) and predicts the distribution of large-scale synchrony in activity (Fig. 4). Moreover, the model becomes more effective as the population grows (Fig. 7) and exhibits hints of critical behavior (Fig. 8). The success of such a sparse model hinges on the fact that the distribution of mutual information between neurons is heavy-tailed [Fig. 2(b)] such that a few rare correlations carry much more information than average. In fact, the physical connections between neurons are now understood to be heavy-tailed across a range of animals [47]. Together, these results suggest that our minimax entropy framework may provide at least a starting point for simplified descriptions of the large systems becoming accessible in modern experiments.

While our approach has proven surprisingly effective, it admits a number of key limitations that pave the way for future research. The central feature of our models—their sparsity—is also the main constraint on their accuracy. Indeed, while

our minimax entropy framework captures the sparse network of strong correlations in the population, it underpredicts the strengths of weak correlations (Fig. 3), which (once added together) can have a large impact on the detailed behavior of a system [16]. Although we cannot constrain all $\sim N^2$ correlations in large populations (to avoid undersampling), one might hope to construct a model that includes more than the $N - 1$ in a tree. We hope that in future work it will be possible to build models that extend this (literal) backbone in ways that remain tractable.

ACKNOWLEDGMENTS

We thank L. Meshulam and J. L. Gauthier for guiding us through the data of Ref. [32], and C. M. Holmes and D. J. Schwab for helpful discussions. This work was supported in part by the National Science Foundation through the Center for the Physics of Biological Function (PHY-1734030), by the National Institutes of Health through the BRAIN initiative (R01EB026943), by the James S. McDonnell Foundation through a Postdoctoral Fellowship Award (C.W.L.), and by Fellowships from the Simons Foundation and the John Simon Guggenheim Memorial Foundation (W.B.).

APPENDIX A: ISING CALCULATIONS ON A TREE

To establish notation, we begin by reviewing well-known ideas about statistical mechanics for models without loops. We then proceed, here and in subsequent Appendixes, to technical points needed for the main text.

1. Partition function

Consider a system of N binary variables $x_i \in \{0, 1\}$, $i = 1, 2, \dots, N$, defined by fields h_i and interactions J_{ij} that lie on a tree \mathcal{T} . The Boltzmann distribution [Eq. (5)] takes the form

$$P_{\mathcal{T}}(\mathbf{x}) = \frac{1}{Z} \exp \left[\sum_{(ij) \in \mathcal{T}} J_{ij} x_i x_j + \sum_i h_i x_i - F \right], \quad (\text{A1})$$

where $F = 0$ is the zero-point energy, which will become useful. To begin, we seek to compute the partition function,

$$Z = \sum_{\mathbf{x}} \exp \left[\sum_{(ij) \in \mathcal{T}} J_{ij} x_i x_j + \sum_i h_i x_i - F \right]. \quad (\text{A2})$$

To do so, imagine summing over one variable and finding a new system of $N - 1$ variables with the same partition function Z . If we can repeat this process until no variables remain, then computing Z will be trivial.

We label the nodes i based on the order that they are removed, and we let $h_i^{(i)}$ and $F^{(i)}$ denote the updated parameters at step i while the interactions J_{ij} stay fixed. Consider summing over a variable i with only one connection in the network, say to variable j . We note that such a node is always guaranteed to exist in a tree. To keep the partition function fixed, the new system with i removed must satisfy the equations

$$e^{h_j^{(i)} x_j - F^{(i)}} (e^{J_{ij} x_j + h_i^{(i)}} + 1) = e^{h_j^{(i+1)} x_j - F^{(i+1)}}. \quad (\text{A3})$$

This is a system of two equations (one for each value of x_j), which we can solve for the new parameters:

$$h_j^{(i+1)} = h_j^{(i)} + \ln \left[\frac{e^{J_{ij}+h_j^{(i)}} + 1}{e^{h_j^{(i)}} + 1} \right], \quad (\text{A4})$$

$$F^{(i+1)} = F^{(i)} - \ln[e^{h_j^{(i)}} + 1]. \quad (\text{A5})$$

After removing i , the new system still forms a tree, so we can repeat the above procedure. When all nodes have been removed, we are left with a single parameter $\mathcal{F} = F^{(N+1)}$, which is the free energy of the system, and the partition function is given by

$$Z = e^{-\mathcal{F}}. \quad (\text{A6})$$

2. Average activities and correlations

To compute population statistics, one simply needs to take derivatives of the partition function,

$$\langle x_i \rangle = \frac{d \ln Z}{dh_i} = -\frac{d\mathcal{F}}{dh_i}, \quad (\text{A7})$$

$$\langle x_i x_j \rangle = \frac{d \ln Z}{dJ_{ij}} = -\frac{d\mathcal{F}}{dJ_{ij}}, \quad (\text{A8})$$

where $\frac{d}{dh_i}$ and $\frac{d}{dJ_{ij}}$ represent total derivatives, which account for indirect dependencies via Eqs. (A4) and (A5). Since $\frac{d\mathcal{F}}{dF^{(i+1)}} = 1$ and $\frac{dh_j^{(i)}}{dh_i} = 1$, the above procedure yields

$$\langle x_i \rangle = -\frac{\partial F^{(i+1)}}{\partial h_i^{(i)}} - \frac{d\mathcal{F}}{dh_j^{(i+1)}} \frac{\partial h_j^{(i+1)}}{\partial h_i^{(i)}}. \quad (\text{A9})$$

Noticing that

$$-\frac{d\mathcal{F}}{dh_j^{(i+1)}} = -\frac{d\mathcal{F}}{dh_j} = \langle x_j \rangle, \quad (\text{A10})$$

and taking derivatives of Eqs. (A4) and (A5), we have

$$\langle x_i \rangle = \frac{1}{1 + e^{-h_i^{(i)}}} + \langle x_j \rangle \left(\frac{1}{1 + e^{-J_{ij}-h_i^{(i)}}} - \frac{1}{1 + e^{-h_i^{(i)}}} \right). \quad (\text{A11})$$

The correlation follows analogously,

$$\langle x_i x_j \rangle = -\frac{d\mathcal{F}}{dh_j^{(i+1)}} \frac{\partial h_j^{(i+1)}}{\partial J_{ij}} = \frac{\langle x_j \rangle}{1 + e^{-J_{ij}-h_i^{(i)}}}. \quad (\text{A12})$$

Thus, by proceeding in the opposite order from which the nodes were removed, we can compute the average activities $\langle x_i \rangle$ and correlations $\langle x_i x_j \rangle$ for $(ij) \in \mathcal{T}$. For the correlations $\langle x_i x_j \rangle$ off the tree (that is, for $(ij) \notin \mathcal{T}$), see Appendix E.

APPENDIX B: MAXIMUM ENTROPY ON A TREE

We now solve the inverse problem for the parameters h_i and J_{ij} given the observations $\langle x_i \rangle$ and $\langle x_i x_j \rangle$ on a tree \mathcal{T} . Inverting

Eqs. (A11) and (A12) yields

$$h_i^{(i)} = \ln \left[\frac{\langle x_i \rangle - \langle x_i x_j \rangle}{1 - \langle x_i \rangle - \langle x_j \rangle \langle x_i x_j \rangle} \right], \quad (\text{B1})$$

$$J_{ij} = \ln \left[\frac{\langle x_i x_j \rangle}{\langle x_j \rangle - \langle x_i x_j \rangle} \right] - h_i^{(i)}. \quad (\text{B2})$$

Combining the above equations, we can solve for the interaction J_{ij} in Eq. (6). To compute the local field h_i , we note that we can repeat the procedure in Appendix A ending at any node; this is equivalent to choosing the root of the tree. If we choose i to be the final node, then we have

$$h_i^{(N)} = \ln \frac{\langle x_i \rangle}{1 - \langle x_i \rangle}. \quad (\text{B3})$$

Additionally, for each neighbor $j \in \mathcal{N}_i$, Eq. (A4) tells us that we receive a contribution to $h_i^{(N)}$ of the form

$$\begin{aligned} h_i^{(j+1)} - h_i^{(j)} &= \ln \left[\frac{e^{J_{ij}+h_j^{(j)}} + 1}{e^{h_j^{(j)}} + 1} \right] \\ &= \ln \left[\frac{\langle x_i \rangle (1 - \langle x_i \rangle - \langle x_j \rangle + \langle x_i x_j \rangle)}{(1 - \langle x_i \rangle)(\langle x_i \rangle - \langle x_i x_j \rangle)} \right]. \end{aligned} \quad (\text{B4})$$

Combining these contributions yields

$$\begin{aligned} h_i &= h_i^{(N)} - \sum_{j \in \mathcal{N}_i} (h_i^{(j+1)} - h_i^{(j)}) \\ &= \ln \frac{\langle x_i \rangle}{1 - \langle x_i \rangle} + \sum_{j \in \mathcal{N}_i} \ln \left[\frac{(1 - \langle x_i \rangle)(\langle x_i \rangle - \langle x_i x_j \rangle)}{\langle x_i \rangle (1 - \langle x_i \rangle - \langle x_j \rangle + \langle x_i x_j \rangle)} \right]. \end{aligned} \quad (\text{B5})$$

We have thus arrived at an analytic solution to the maximum entropy problem on a tree.

APPENDIX C: INFORMATION IN A TREE OF CORRELATIONS

Our ability to efficiently construct the optimal tree \mathcal{T} depends critically on the decomposition of the information $I_{\mathcal{T}}$ into the sum of mutual information I_{ij} over pairs $(ij) \in \mathcal{T}$ [Eq. (8)]. To derive this result, we note that for each connection $(ij) \in \mathcal{T}$, the observables $\langle x_i \rangle$, $\langle x_j \rangle$, and $\langle x_i x_j \rangle$ fully define the marginal distribution $P_{ij}(x_i, x_j)$. Now consider a new tree $\mathcal{T}' = \mathcal{T}/(ij)$ with the connection (ij) removed, such that we do not observe $\langle x_i x_j \rangle$. Since \mathcal{T} has no loops, after removing (ij) the two elements i and j become independent. Meanwhile, the dependence of the rest of the system on i and j remains fixed. Thus, observing the correlation $\langle x_i x_j \rangle$ leads to a drop in entropy,

$$S_{\mathcal{T}'} - S_{\mathcal{T}} = S(P_i) + S(P_j) - S(P_{ij}) = I_{ij}, \quad (\text{C1})$$

where I_{ij} is the observed mutual information between i and j . Repeating the above argument for every correlation in \mathcal{T} , we arrive at Eq. (8).

APPENDIX D: ESTIMATING MUTUAL INFORMATION

In order to estimate the mutual information between neurons I_{ij} , one must correct for finite-data effects [48]. To do

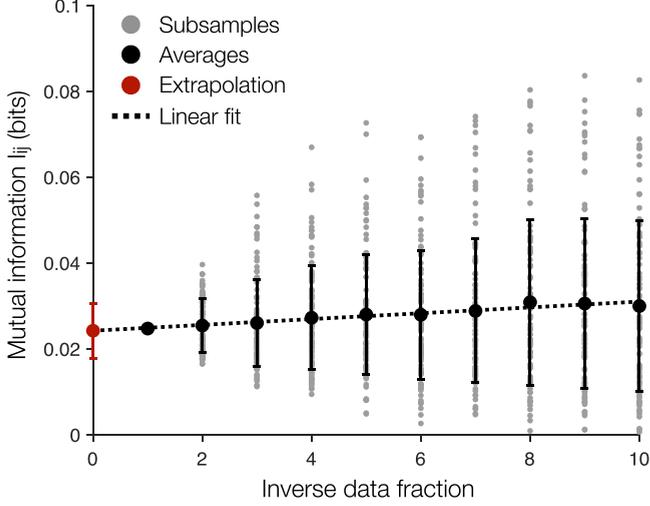


FIG. 9. Correcting for finite-data effects on mutual information. For a given pair of neurons i and j , we plot the estimated mutual information I_{ij} [Eq. (D1)] vs the inverse data fraction for individual subsamples of the data (gray). Repeating 100 times, we plot the average I_{ij} for each data fraction (black), a linear fit (dashed line), and the infinite-data estimate (red). Data points and error bars reflect means and standard deviations over the 100 repetitions.

so, we subsample the data hierarchically for different data fractions $\{1, 0.9, \dots, 0.2, 0.1\}$ such that each subsample is contained within the larger subsamples. Additionally, to preserve the dependencies between consecutive data points $\mathbf{x}^{(m)}$ and $\mathbf{x}^{(m+1)}$, we sample temporally contiguous fractions of the data. To ensure that each point is sampled with equal probability, we allow subsamples that span the beginning and end of the recording.

For each subsample, we estimate the mutual information between neurons i and j using the equation

$$I_{ij} = \sum_{x_i, x_j} \tilde{P}_{ij}(x_i, x_j) \log \frac{\tilde{P}_{ij}(x_i, x_j)}{\tilde{P}_i(x_i) \tilde{P}_j(x_j)}, \quad (\text{D1})$$

where

$$\tilde{P}_{ij}(x_i, x_j) = \frac{1}{M+1} \left(1 + \sum_{m=1}^M \delta_{x_i, x_i^{(m)}} \delta_{x_j, x_j^{(m)}} \right), \quad (\text{D2})$$

$$\tilde{P}_i(x_i) = \sum_{x_j} \tilde{P}_{ij}(x_i, x_j). \quad (\text{D3})$$

The pseudocounts in Eq. (D2) ensure that the mutual information estimates do not diverge. After estimating I_{ij} for each data fraction, following Ref. [48] we extrapolate to the infinite-data limit using a linear fit with respect to the inverse data fraction. Repeating this process 100 times, we arrive at a distribution of infinite-data estimates for I_{ij} from which we can compute a mean and standard deviation (Fig. 9). We consider a mutual information I_{ij} significant if the mean is more than two standard deviations above zero.

To check the above procedure, we note that shuffling the activity of each neuron in time should destroy the mutual information I_{ij} . Indeed, for time-shuffled data, we do not observe a single significant mutual information in the population.

APPENDIX E: COMPUTING ALL CORRELATIONS

Given a maximum entropy model with parameters h_i and J_{ij} on a tree \mathcal{T} , in Appendix A we showed how to compute the averages $\langle x_i \rangle$ and correlations $\langle x_i x_j \rangle$ on the tree. Specifically, we computed the partition function Z by summing over variables x_i in the order $i = 1, 2, \dots, N$ and then computed statistics in the reverse order. Here, we show how to compute the correlations $\langle x_i x_j \rangle$ not on the tree, that is, for $(ij) \notin \mathcal{T}$. To begin, we assume that we have computed the correlations $\langle x_j x_k \rangle$ for all nodes $k > i > j$. Then, if we compute $\langle x_i x_j \rangle$, the procedure will follow by induction.

From the Boltzmann distribution in Eq. (5), we have

$$\frac{d\langle x_i \rangle}{dh_j} = \langle x_i x_j \rangle - \langle x_i \rangle \langle x_j \rangle. \quad (\text{E1})$$

We already know how to compute the averages $\langle x_i \rangle$ and $\langle x_j \rangle$, so all that remains is to calculate the above derivative. Let $p(i)$ denote the parent of i (that is, the final neighbor when i is removed) and likewise for $p(j)$. Differentiating Eq. (A11) with respect to h_j , we have

$$\frac{d\langle x_i \rangle}{dh_j} = \frac{\partial \langle x_i \rangle}{\partial h_i^{(i)}} \frac{dh_i^{(i)}}{dh_j} + \frac{\partial \langle x_i \rangle}{\partial \langle x_{p(i)} \rangle} \frac{d\langle x_{p(i)} \rangle}{dh_j}. \quad (\text{E2})$$

We note that

$$\frac{d\langle x_{p(i)} \rangle}{dh_j} = \langle x_{p(i)} x_j \rangle - \langle x_{p(i)} \rangle \langle x_j \rangle, \quad (\text{E3})$$

which we have already computed by assumption, since $p(i) > i$. From Eq. (A11) we have

$$\begin{aligned} \frac{\partial \langle x_i \rangle}{\partial h_i^{(i)}} &= \frac{e^{-h_i^{(i)}}}{(1 + e^{-h_i^{(i)}})^2} \\ &+ \langle x_{p(i)} \rangle \left(\frac{e^{-J_{ip(i)} - h_i^{(i)}}}{(1 + e^{-J_{ip(i)} - h_i^{(i)}})^2} - \frac{e^{-h_i^{(i)}}}{(1 + e^{-h_i^{(i)}})^2} \right), \end{aligned} \quad (\text{E4})$$

and

$$\frac{\partial \langle x_i \rangle}{\partial \langle x_{p(i)} \rangle} = \frac{1}{1 + e^{-J_{ip(i)} - h_i^{(i)}}} - \frac{1}{1 + e^{-h_i^{(i)}}}. \quad (\text{E5})$$

Finally, we note that the dependence of $h_i^{(i)}$ on h_j runs only through $h_{p(j)}^{(j+1)}$, such that

$$\frac{dh_i^{(i)}}{dh_j} = \frac{dh_i^{(i)}}{dh_{p(j)}^{(j+1)}} \frac{dh_{p(j)}^{(j+1)}}{dh_j}. \quad (\text{E6})$$

Since $p(j) > j$, we can assume that we have already computed $\frac{dh_{p(j)}^{(j+1)}}{dh_{p(j)}^{(j+1)}}$. Finally, Eq. (A4) yields

$$\frac{\partial h_{p(j)}^{(j+1)}}{\partial h_j} = \frac{1}{1 + e^{-J_{jp(j)} - h_j^{(j)}}} - \frac{1}{1 + e^{-h_j^{(j)}}}. \quad (\text{E7})$$

Plugging everything into Eq. (E2), and inducting on $i > j$, one can compute the correlations $\langle x_i x_j \rangle$ between all variables.

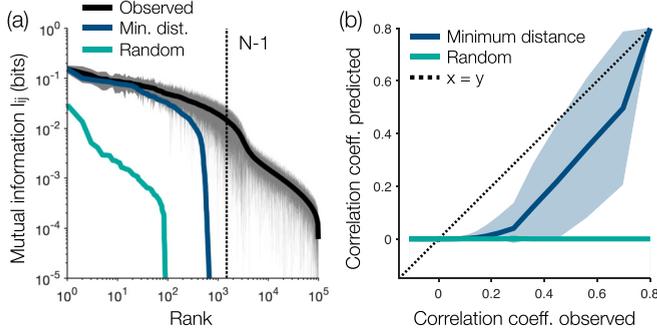


FIG. 10. Pairwise statistics in minimum distance and random trees. (a) Ranked order of significant mutual information in the population (black), two-standard-deviation errors (shaded region), and predictions of the minimum distance tree (blue) and a random tree (cyan). (b) Correlation coefficients predicted in the minimum distance and random tree models vs those in the data, with dashed lines indicating equality. All pairs are divided evenly into bins along the x axis, with solid lines and shaded regions reflecting means and errors (standard deviations) within bins.

APPENDIX F: PREDICTIONS OF RANDOM AND MINIMUM DISTANCE TREES

In Sec. IV we studied the predictions of the minimax entropy model $P_{\mathcal{T}}$ corresponding to the optimal tree \mathcal{T} . For comparison, here we consider the predictions of the minimum distance and random trees. While the optimal tree captures $I_{\mathcal{T}} = 26.2$ bits of information ($I_{\mathcal{T}}/S_{\text{ind}} = 14.4\%$ of the independent entropy), the minimum distance tree only captures 11.9 bits of information (6.5% of the independent entropy), and a typical random tree only captures $(N-1)\bar{I} = 0.4$ bits (0.2% of the independent entropy).

For each tree, we can predict the mutual information I_{ij} and correlation coefficients between all pairs of neurons using the procedure in Appendix E. Since the minimum distance tree includes some of the largest mutual information in the population, it is able to match the distribution of I_{ij} (within errors) out to $N \sim 100$ neurons [Fig. 10(a)]. However, the minimum distance tree fails to predict the observed correlations across most of the dynamic range of the data [Fig. 10(b)]. Meanwhile, random trees typically include only weak mutual information [Fig. 10(a)], such that their predictions are nearly indistinguishable from a population of independent neurons [Fig. 10(b)].

For the minimax entropy model, a backbone of strong positive interactions combine to produce accurate predictions for the distribution $P(K)$ of population-wide synchrony K (Fig. 4). By contrast, random trees predict a Gaussian distribution consistent with independent neurons, and the minimum distance tree only produces a slightly broader distribution (Fig. 11). In both models, large-scale synchrony in activity ($K \gtrsim 50$) and silence ($K \lesssim 10$) occurs significantly less frequently than observed in the data.

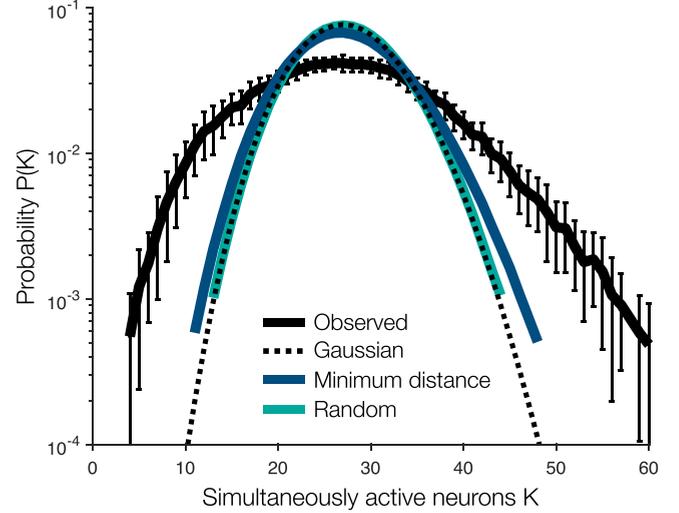


FIG. 11. Synchronized activity in minimum distance and random trees. Distribution $P(K)$ of the number of simultaneously active neurons K in the data (black), the Gaussian distribution for independent neurons with mean and variance $\langle K \rangle_{\text{exp}}$ (dashed), and the predictions of the minimum distance tree (blue) and a random tree (cyan). As in Fig. 4, to estimate $P(K)$ and error bars (two standard deviations), we first split the data into 1-min blocks to preserve dependencies between consecutive samples. We then select one-third of these blocks at random and repeat 100 times. For each subsample of the data, we fit the maximum entropy model $P_{\mathcal{T}}$ for each tree \mathcal{T} and predict $P(K)$ using a Monte Carlo simulation.

APPENDIX G: THERMODYNAMIC QUANTITIES

Consider a system with fields h_i , interactions J_{ij} that lie on a tree \mathcal{T} , and temperature T . The Boltzmann distribution $P(\mathbf{x})$ takes the form in Eq. (9) with energy $E(\mathbf{x})$ defined in Eq. (10). Here, we denote averages over $P(\mathbf{x})$ by $\langle \cdot \rangle$, while dropping the subscript T . The susceptibility χ [Eq. (11)] can be computed using the results of Appendix E.

To compute the heat capacity C [Eq. (12)], we begin with the average energy

$$\langle E(\mathbf{x}) \rangle = T^2 \frac{d \ln Z}{dT} = -T^2 \frac{d(\mathcal{F}/T)}{dT} = \mathcal{F} - T \frac{d\mathcal{F}}{dT}. \quad (\text{G1})$$

To calculate the free energy \mathcal{F} , we proceed as in Appendix A. After including the temperature T , Eqs. (A4) and (A5) take the form

$$h_j^{(i+1)} = h_j^{(i)} + T \ln \left[\frac{e^{\frac{1}{T}(J_{ij} + h_i^{(i)})} + 1}{e^{\frac{1}{T}h_i^{(i)}} + 1} \right], \quad (\text{G2})$$

$$F^{(i+1)} = F^{(i)} - T \ln \left[e^{\frac{1}{T}h_i^{(i)}} + 1 \right]. \quad (\text{G3})$$

Iteratively summing over each variable, we arrive at the free energy $\mathcal{F} = F^{(N+1)}$. To compute $\frac{d\mathcal{F}}{dT}$, we take derivatives of Eqs. (G4) and (G5), yielding

$$\frac{dh_j^{(i+1)}}{dT} = \frac{dh_j^{(i)}}{dT} + \ln \left[\frac{e^{\frac{1}{T}(J_{ij} + h_i^{(i)})} + 1}{e^{\frac{1}{T}h_i^{(i)}} + 1} \right] - \frac{1}{1 + e^{-\frac{1}{T}h_i^{(i)}}} \left(\frac{dh_i^{(i)}}{dT} - \frac{h_i^{(i)}}{T} \right) + \frac{1}{1 + e^{-\frac{1}{T}(J_{ij} + h_i^{(i)})}} \left(\frac{dh_i^{(i)}}{dT} - \frac{J_{ij} + h_i^{(i)}}{T} \right), \quad (\text{G4})$$

$$\frac{dF^{(i+1)}}{dT} = \frac{dF^{(i)}}{dT} - \ln[e^{\frac{1}{T}h_i^{(i)}} + 1] - \frac{1}{1 + e^{-\frac{1}{T}h_i^{(i)}}} \left(\frac{dh_i^{(i)}}{dT} - \frac{h_i^{(i)}}{T} \right). \quad (\text{G5})$$

Iterating the above equations, we arrive at the derivative $\frac{d\mathcal{F}}{dT} = \frac{dF^{(N+1)}}{dT}$, which completes our calculation of the average energy [Eq. (G1)].

The heat capacity is given by

$$C = \frac{d\langle E(\mathbf{x}) \rangle}{dT} = -T \frac{d^2\mathcal{F}}{dT^2}. \quad (\text{G6})$$

Taking derivatives of Eqs. (G4) and (G5), we have

$$\begin{aligned} \frac{d^2h_j^{(i+1)}}{dT^2} &= \frac{d^2h_j^{(i)}}{dT^2} - \frac{1}{1 + e^{-\frac{1}{T}h_i^{(i)}}} \frac{d^2h_i^{(i)}}{dT^2} - \frac{1}{T} \frac{e^{-\frac{1}{T}h_i^{(i)}}}{(1 + e^{-\frac{1}{T}h_i^{(i)}})^2} \left(\frac{dh_i^{(i)}}{dT} - \frac{h_i^{(i)}}{T} \right)^2 \\ &+ \frac{1}{1 + e^{-\frac{1}{T}(J_{ij}+h_i^{(i)})}} \frac{d^2h_i^{(i)}}{dT^2} + \frac{1}{T} \frac{e^{-\frac{1}{T}(J_{ij}+h_i^{(i)})}}{(1 + e^{-\frac{1}{T}(J_{ij}+h_i^{(i)})})^2} \left(\frac{dh_i^{(i)}}{dT} - \frac{J_{ij} + h_i^{(i)}}{T} \right)^2, \end{aligned} \quad (\text{G7})$$

$$\frac{d^2F^{(i+1)}}{dT^2} = \frac{d^2F^{(i)}}{dT^2} - \frac{1}{1 + e^{-\frac{1}{T}h_i^{(i)}}} \frac{d^2h_i^{(i)}}{dT^2} - \frac{1}{T} \frac{e^{-\frac{1}{T}h_i^{(i)}}}{(1 + e^{-\frac{1}{T}h_i^{(i)}})^2} \left(\frac{dh_i^{(i)}}{dT} - \frac{h_i^{(i)}}{T} \right)^2. \quad (\text{G8})$$

Finally, after computing $\frac{d^2\mathcal{F}}{dT^2} = \frac{d^2F^{(N+1)}}{dT^2}$ iteratively, we have arrived at the heat capacity C .

-
- [1] Norbert Wiener, *Nonlinear Problems in Random Theory* (MIT Press, Cambridge, MA, 1958).
- [2] L. N. Cooper, A possible organization of animal memory and learning, in *Collective Properties of Physical Systems: Proceedings of Nobel Symposium 24*, edited by B. Lundqvist and S. Lundqvist (Academic Press, New York, 1973), pp. 252–264.
- [3] W. A. Little, The existence of persistent states in the brain, *Math. Biosci.* **19**, 101 (1974).
- [4] J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. USA* **79**, 2554 (1982).
- [5] D. J. Amit, *Modeling Brain Function: The World of Attractor Neural Networks* (Cambridge University Press, Cambridge, England, 1989).
- [6] J. Hertz, A. Krogh, and R. G. Palmer, *Introduction to the Theory of Neural Computation* (Addison–Wesley, Redwood City, 1991).
- [7] L. Meshulam and W. Bialek, Statistical mechanics for networks of real neurons, [arXiv:2409.00412](https://arxiv.org/abs/2409.00412).
- [8] R. Segev, J. Goodhouse, J. Puchalla, and M. J. Berry II, Recording spikes from a large fraction of the ganglion cells in a retinal patch, *Nat. Neurosci.* **7**, 1155 (2004).
- [9] A. M. Litke, N. Bezayiff, E. J. Chichilnisky, W. Cunningham, W. Dabrowski, A. A. Grillo, M. Grivich, P. Grybos, P. Hottowy, S. Kachiguine *et al.*, What does the eye tell the brain?: Development of a system for the large-scale recording of retinal output activity, *IEEE Trans. Nucl. Sci.* **51**, 1434 (2004).
- [10] J. E. Chung, H. R. Joo, J. L. Fan, D. F. Liu, A. H. Barnett, S. Chen, C. Geaghan-Breiner, M. P. Karlsson, M. Karlsson, K. Y. Lee *et al.*, High-density, long-lasting, and multi-region electrophysiological recordings using polymer electrode arrays, *Neuron* **101**, 21 (2019).
- [11] D. A. Dombeck, C. D. Harvey, L. Tian, L. L. Looger, and D. W. Tank, Functional imaging of hippocampal place cells at cellular resolution during virtual navigation, *Nat. Neurosci.* **13**, 1433 (2010).
- [12] L. Tian, J. Akerboom, E. R. Schreiter, and L. L. Looger, Neural activity imaging with genetically encoded calcium indicators, *Prog. Brain Res.* **196**, 79 (2012).
- [13] J. Demas, J. Manley, F. Tejera, K. Barber, H. Kim, F. M. Traub, B. Chen, and A. Vaziri, High-speed, cortex-wide volumetric recording of neuroactivity at cellular resolution using light beads microscopy, *Nat. Methods* **18**, 1103 (2021).
- [14] N. A. Steinmetz, C. Aydin, A. Lebedeva, M. Okun, M. Pachitariu, M. Bauza, M. Beau, J. Bhagat, C. Böhm, M. Broux *et al.*, Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings, *Science* **372**, eabf4588 (2021).
- [15] E. T. Jaynes, Information theory and statistical mechanics, *Phys. Rev.* **106**, 620 (1957).
- [16] E. Schneidman, M. J. Berry II, R. Segev, and W. Bialek, Weak pairwise correlations imply strongly correlated network states in a neural population, *Nature (London)* **440**, 1007 (2006).
- [17] H. C. Nguyen, R. Zecchina, and J. Berg, Inverse statistical problems: From the inverse Ising problem to data science, *Adv. Phys.* **66**, 197 (2017).
- [18] L. Meshulam, J. L. Gauthier, C. D. Brody, D. W. Tank, and W. Bialek, Collective behavior of place and non-place neurons in the hippocampal network, *Neuron* **96**, 1178 (2017).
- [19] L. Meshulam, J. L. Gauthier, C. D. Brody, D. W. Tank, and W. Bialek, Successes and failures of simplified models for a network of real neurons, [arXiv:2112.14735](https://arxiv.org/abs/2112.14735).
- [20] G. Tkačik, T. Mora, O. Marre, D. Amodèi, S. E. Palmer, M. J. Berry II, and W. Bialek, Thermodynamics and signatures of

- criticality in a network of neurons, *Proc. Natl. Acad. Sci. USA* **112**, 11508 (2015).
- [21] T. R. Lezon, J. R. Banavar, M. Cieplak, A. Maritan, and N. V. Fedoroff, Using the principle of entropy maximization to infer genetic interaction networks from gene expression patterns, *Proc. Natl. Acad. Sci. USA* **103**, 19033 (2006).
- [22] M. Weigt, R. A. White, H. Szurmant, J. A. Hoch, and T. Hwa, Identification of direct residue contacts in protein–protein interaction by message passing, *Proc. Natl. Acad. Sci. USA* **106**, 67 (2009).
- [23] D. S. Marks, L. J. Colwell, R. Sheridan, T. A. Hopf, A. Pagnani, R. Zecchina, and C. Sander, Protein 3D structure computed from evolutionary sequence variation, *PLoS ONE* **6**, e28766 (2011).
- [24] A. Lapedes, B. Giraud, and C. Jarzynski, Using sequence alignments to predict protein structure and stability with high accuracy, [arXiv:1207.2484](https://arxiv.org/abs/1207.2484).
- [25] W. Bialek, A. Cavagna, I. Giardina, T. Mora, E. Silvestri, M. Viale, and A. M. Walczak, Statistical mechanics for natural flocks of birds, *Proc. Natl. Acad. Sci. USA* **109**, 4786 (2012).
- [26] W. P. Russ, M. Figliuzzi, C. Stocker, P. Barrat-Charlaix, M. Socolich, P. Kast, D. Hilvert, R. Monasson, S. Cocco, M. Weigt *et al.*, An evolution-based model for designing chorismate mutase enzymes, *Science* **369**, 440 (2020).
- [27] C. W. Lynn, L. Papadopoulos, D. D. Lee, and D. S. Bassett, Surges of collective human activity emerge from simple pairwise correlations, *Phys. Rev. X* **9**, 011022 (2019).
- [28] S. C. Zhu, Y. N. Wu, and D. Mumford, Minimax entropy principle and its application to texture modeling, *Neural Comput.* **9**, 1627 (1997).
- [29] J. S. Prentice, O. Marre, M. L. Ioffe, A. R. Loback, G. Tkačik, and M. J. Berry, Error-robust modes of the retinal population code, *PLoS Comput. Biol.* **12**, e1005148 (2016).
- [30] C. Chow and C. Liu, Approximating discrete probability distributions with dependence trees, *IEEE Trans. Inf. Theory* **14**, 462 (1968).
- [31] C. W. Lynn, Q. Yu, R. Pang, S. E. Palmer, and W. Bialek, companion paper, Exactly solvable statistical physics models for large neuronal populations, *Phys. Rev. Res.* **7**, L022039 (2025).
- [32] J. L. Gauthier and D. W. Tank, A dedicated population for reward coding in the hippocampus, *Neuron* **99**, 179 (2018).
- [33] C. E. Shannon, A mathematical theory of communication, *Bell Syst. Tech. J.* **27**, 379 (1948).
- [34] E. Schneidman, S. Still, M. J. Berry II, and W. Bialek, Network information and connected correlations, *Phys. Rev. Lett.* **91**, 238701 (2003).
- [35] J. P. Sethna, *Statistical Mechanics: Entropy, Order Parameters, and Complexity* (Oxford University Press, New York, 2021), Vol. 14.
- [36] C. Moore and S. Mertens, *The Nature of Computation* (Oxford University Press, Oxford, 2011).
- [37] S. Song, P. J. Sjöström, M. Reigl, S. Nelson, and D. B. Chklovskii, Highly nonrandom features of synaptic connectivity in local cortical circuits, *PLoS Biol.* **3**, e68 (2005).
- [38] C. W. Lynn and D. S. Bassett, The physics of brain network structure, function and control, *Nat. Rev. Phys.* **1**, 318 (2019).
- [39] A. Lin, R. Yang, S. Dorckenwald, A. Matsliah, A. R. Sterling, P. Schlegel, S.-C. Yu, C. E. McKellar, M. Costa, K. Eichler *et al.*, Network statistics of the whole-brain connectome of *Drosophila* (2023), doi:10.1101/2023.07.29.551086.
- [40] R. Albert and A.-L. Barabási, Statistical mechanics of complex networks, *Rev. Mod. Phys.* **74**, 47 (2002).
- [41] J. J. Harris, R. Jolivet, and D. Attwell, Synaptic energy use and supply, *Neuron* **75**, 762 (2012).
- [42] S. B. Laughlin and T. J. Sejnowski, Communication in neuronal networks, *Science* **301**, 1870 (2003).
- [43] D. J. Watts and S. H. Strogatz, Collective dynamics of “small-world” networks, *Nature (London)* **393**, 440 (1998).
- [44] S. Schnabel, D. T. Seaton, D. P. Landau, and M. Bachmann, Microcanonical entropy inflection points: Key to systematic understanding of transitions in finite systems, *Phys. Rev. E* **84**, 011127 (2011).
- [45] R. Peierls, On Ising’s model of ferromagnetism, in *Mathematical Proceedings of the Cambridge Philosophical Society*, Vol. 32 (Cambridge University Press, Cambridge, England, 1936), pp. 477–481.
- [46] R. J. Baxter, *Exactly Solved Models in Statistical Mechanics* (Elsevier, New York, 2016).
- [47] C. W. Lynn, C. M. Holmes, and S. E. Palmer, Heavy-tailed neuronal connectivity arises from Hebbian self-organization, *Nat. Phys.* **20**, 484 (2024).
- [48] S. P. Strong, R. Koberle, R. R. de Ruyter van Steveninck, and W. Bialek, Entropy and information in neural spike trains, *Phys. Rev. Lett.* **80**, 197 (1998).