

Probabilistic Models, Compressible Interactions, and Neural Coding

Luisa Ramirez^{1,2,3,*} William Bialek^{3,4,5,†} Stephanie E. Palmer^{1,4,6,‡} and David J. Schwab^{4,‡}

¹*Institute of Developmental Biology and Neurobiology, Johannes-Gutenberg University Mainz, Mainz, Germany*

²*Departamento de Física, Universidade Federal de Minas Gerais, 31270-901 Belo Horizonte, Minas Gerais, Brasil*

³*Joseph Henry Laboratories of Physics and Lewis–Sigler Institute for Integrative Genomics, Princeton University, Princeton, New Jersey 08544, USA*

⁴*Initiative for the Theoretical Sciences, The CUNY Graduate Center, City University of New York, 365 Fifth Avenue, New York, New York 10016, USA*

⁵*Center for Studies in Physics and Biology, Rockefeller University, 1230 York Avenue, New York, New York 10021, USA*

⁶*Department of Organismal Biology and Anatomy and Department of Physics, The University of Chicago, Chicago, Illinois 60637, USA*



(Received 11 April 2022; accepted 24 March 2026; published 23 April 2026)

In physics we often use very simple models to describe systems with many degrees of freedom, but it is not clear why or how this success can be transferred to the more complex biological context. We consider models for the joint distribution of many variables, as with the combinations of spiking and silence in large networks of neurons. In this probabilistic framework, we argue that simple models are possible if the mutual information between two halves of the system is consistently subextensive, and if this shared information is compressible. These conditions are not met generically, but they are met by real-world data, such as natural images and the activity in a population of retinal output neurons. We introduce compression strategies that combine the information bottleneck with an iteration scheme inspired by the renormalization group, and we find that the number of parameters needed to describe the distribution of joint activity scales with the square of the number of neurons, even though the interactions are not well approximated as pairwise. Our results also show that this shared information is essentially equal to the information that individual neurons carry about natural visual inputs, which has surprising implications for the neural code.

DOI: [10.1103/pj4s-g4jc](https://doi.org/10.1103/pj4s-g4jc)

I. INTRODUCTION

In statistical mechanics, we routinely analyze the joint probability distribution of very large numbers of variables; in field theory this number is infinite, at least formally [1,2]. There is considerable interest in giving a similar probabilistic description outside the traditional domains of physics, spurred in part by the availability of “big data” from a wider variety of complex systems. But the models we consider in most physics problems are highly constrained, and without these constraints we must learn the underlying distribution from the data. If what we observe are discrete states or events, then the probability distribution is a list of numbers, one for each possible outcome, and this number is beyond astronomical: in an image with just $N = 100$ pixels, where each pixel can be black or white, the number of possible images ($2^N \sim 10^{30}$) is larger

than the age of the universe in seconds. Under these conditions it is physically impossible to “measure” the underlying probability distribution from data alone, and it will continue to be impossible no matter how our technology evolves.

The problem of inferring large probabilistic models has been made more urgent by enormous growth in our ability to monitor, simultaneously, the functional activity of many degrees of freedom in living systems. Examples range from the expression levels of many genes in a single cell [3–6] to the electrical activity of many neurons in the brain [7–13] and the movements of all the individual organisms in a flock or swarm [14]; we emphasize that these are illustrative rather than exhaustive. To understand these experiments we need a theoretical framework that tames the combinatorial explosion of potential complexity.

We can identify several different reactions to the increased dimensionality of the available experimental data. One view, inspired by the success of modern AI, embraces complex models such as deep neural networks, emphasizing that successful predictions are possible even when the number of parameters in our models far exceeds the number of data points [15–17]. The opposite view is that high-dimensional data may lie on lower-dimensional manifolds, so that the search for these manifolds becomes the central problem of data analysis [18,19]. An intermediate approach focuses on the fact that complex models often have a characteristic

*Contact author: luramire@uni-mainz.de

†These three senior authors (W.B., S.E.P., and D.J.S.) contributed to all aspects of the work.

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

geometry in parameter space, where some combinations of parameters are essential for successful prediction and others are not [20,21]. Other approaches are more explicitly connected to ideas from statistical physics, including the construction of maximum entropy models that are consistent with low-order correlations [22–28] and the search for scaling behaviors that might point toward models described by a fixed point of the renormalization group [29–32]. The maximum entropy approach has been extended to match global features of the data [26], subsets of higher-order correlations [33], nonlinear transformations of the effective energy [34], or the expectation values of thresholded projections of the data as might be computed by real neurons [35]. For a recent review of statistical physics approaches to networks of real neurons, see [36].

In practice, the search for simplification usually is done by hypothesizing a particular family of models and exploring how far these can take us in the description of real data. We would like to go beyond the exploration of particular models to have more general criteria for the learnability of distributions, in the spirit of theories for the learnability of functions or rules [37,38]. To be concrete, suppose that what we observe is a collection of N binary variables, $\sigma \equiv \{\sigma_1, \sigma_2, \dots, \sigma_N\}$, so that there are 2^N possible states, and in general learning the distribution would require many more than 2^N observations. Can we state conditions on the distribution that are sufficient to guarantee effective learning from a much smaller number of examples, perhaps linear or polynomial in N ?¹ Importantly, are these conditions satisfied in natural data?

Here we suggest that two conditions are sufficient for learnability: the consistent subextensivity of the mutual information between parts of a system, and the compressibility of interactions into an efficient representation. We give general arguments, and test our ideas against statistical physics models, the statistical structure of natural images, and the patterns of electrical activity in the retina as it responds to naturalistic inputs. The implications for the analysis of neural data seem especially rich, and so we explore these in more detail. This paper combines and extends work presented in preliminary form [40,41].

Before proceeding, we admit that our focus on simplified models might seem anachronistic in the era of deep networks. These models, which drive the current revolution in artificial intelligence, are far from simple, in some cases being described by trillions of parameters [42]. In truth, we do not understand the success of these models [43,44]. More relevant for our discussion, these models are still small compared with the number of possible “states” taken on by the

relevant variables. For language models, for example, with a five thousand word vocabulary, there are $\sim 10^{37}$ possible 10-word sequences. While most of these are forbidden by grammatical rules, human level performance requires capturing dependencies across ~ 100 words [45]. In this context, trillion-parameter models *are* simplified models.

II. SUBEXTENSIVITY

If we knew that the N binary variables could be broken down into two independent halves, then the full probability distribution could be written in terms of $2 \times 2^{N/2}$ parameters, vastly less than 2^N . More generally, imagine that we can place a bound on the mutual information between the two halves, $I_{1/2}(N)$. If this information is unconstrained, then we need $\sim 2^N$ parameters to describe the system, while if $I_{1/2}(N) \rightarrow 0$ then we can use only $2 \times 2^{N/2}$ and still give an exact description. It seems plausible that if $I_{1/2}(N)$ is sufficiently small, there should be a good approximation that has roughly $2 \times 2^{N/2}$ parameters.

Let us call the two halves of our system right and left,

$$\sigma_R \equiv \{\sigma_1, \sigma_2, \dots, \sigma_{N/2}\}, \quad (1)$$

$$\sigma_L \equiv \{\sigma_{N/2+1}, \sigma_{N/2+2}, \dots, \sigma_N\}. \quad (2)$$

We recall that the shortest possible code which represents the states σ is based on exact knowledge of the probability distribution, where each state σ is represented by a code word of length $L(\sigma) \sim -\ln P(\sigma)$, so that the mean code length is the entropy of the distribution [46,47]. Codes built from approximate models of the distribution will be longer, on average, by an amount $\langle \Delta L \rangle$ equal to the Kullback-Leibler divergence between the model and the true distribution,

$$\begin{aligned} \langle \Delta L \rangle &= \sum_{\sigma} P(\sigma) ([-\log_2 P_{\text{approx}}(\sigma)] - [-\log_2 P(\sigma)]) \\ &= \sum_{\sigma} P(\sigma) \log_2 \left[\frac{P(\sigma)}{P_{\text{approx}}(\sigma)} \right], \end{aligned} \quad (3)$$

and this provides a measure of model quality. If our approximate model is the one in which the two halves of the system are independent,

$$P_{\text{approx}}(\sigma) = P_R(\sigma_R)P_L(\sigma_L), \quad (4)$$

then this coding cost becomes

$$\langle \Delta L \rangle = \sum_{\sigma} P(\sigma_R, \sigma_L) \log_2 \left[\frac{P(\sigma_R, \sigma_L)}{P_R(\sigma_R)P_L(\sigma_L)} \right] \quad (5)$$

$$= I_{1/2}(N), \quad (6)$$

the mutual information between the two halves.

If the variables $\{\sigma_i\}$ are arranged in real space such that there is a finite correlation length ξ , then the division into right and left halves can be taken literally, and the mutual information between the halves arises from correlations among spins within ξ of the boundary. As a result the mutual information must be related to the area of the boundary, not the volume of the system, and hence is subextensive: if the system is of linear dimension l in d dimensions, we have $N \sim l^d$ and $I_{1/2} \sim l^{d-1}$,

¹Our notion of learnability here is a bit different from the usual. In traditional learning theory, one asks for the number of samples required to identify, within some error bounds, the best model in a given class, and the class is defined by the number of parameters or some other complexity measure [39]. Here we are interested in the absolute performance of models, and ask about the number of parameters needed to reach a fixed criterion as the number of degrees of freedom in the system changes. We take as given that the number of needed samples will scale with the complexity of models, and sometimes interchange these quantities, colloquially; our focus is on their scaling with system size.

hence $I_{1/2}(N) = cN^\alpha$ with $\alpha = 1 - 1/d$. In the quantum case this becomes the “area laws” for entanglement [48].

We can ask more generally about systems which, when divided in half, exhibit a mutual information between the halves that behaves as $I_{1/2}(N) = cN^\alpha$ with $\alpha < 1$. Then the approximation of the system as two independent halves has a cost that per degree of freedom

$$\frac{\langle \Delta L \rangle}{N} = cN^{\alpha-1}, \quad (7)$$

which vanishes as N becomes large. Thus, subextensive behavior of the mutual information is sufficient to ensure that, for large systems, the reduction in number of parameters from 2^N down to $2 \times 2^{N/2}$ will result in a model that makes only small errors per degree of freedom.

We can now think about cases where the mutual information is *consistently subextensive*, that is when we look at properly chosen pieces of the system with n variables, and cut these pieces in half, we always find a mutual information between the halves $I_{1/2}(n) \leq cn^\alpha$. This means that we can keep cutting the variables in half, approximating the distribution as being composed of independent halves, and in the process we make errors that are small when measured as the cost of coding per degree of freedom.

If we make b cuts, then we have

$$\langle \Delta L \rangle = cN^\alpha + 2c\left(\frac{N}{2}\right)^\alpha + 4c\left(\frac{N}{4}\right)^\alpha + \dots + 2^{b-1}c\left(\frac{N}{2^{b-1}}\right)^\alpha \quad (8)$$

$$= cN^\alpha \frac{2^{b(1-\alpha)} - 1}{2^{1-\alpha} - 1} \leq \tilde{c}N^\alpha \left(\frac{N}{n_0}\right)^{1-\alpha}, \quad (9)$$

where $\tilde{c} = c/(2^{1-\alpha} - 1)$ and $n_0 = 2^{-b}N$, so that

$$\frac{\langle \Delta L \rangle}{N} \leq \frac{\tilde{c}}{n_0^{1-\alpha}}. \quad (10)$$

This means that we can guarantee a cost $\langle \Delta L \rangle/N \leq \ell$ if we stop cutting once the pieces are of size

$$n_0 = (\tilde{c}/\ell)^{1/(1-\alpha)}. \quad (11)$$

The distribution of n_0 binary variables requires at most 2^{n_0} parameters, and this is independent of N . We need one such model for each of the N/n_0 pieces.

Thus, when the mutual information is consistently subextensive we can make an approximate model that has $\mathbf{P} \sim (N/n_0)2^{n_0}$ parameters, and the error that we make corresponds to an excess coding cost of ℓ bits per degree of freedom, with n_0 and ℓ connected through Eq. (11). This number of parameters is linear in the number of degrees of freedom, and hence we expect that the model can be learned from a number of examples which is also linear in the system size.

Thus, we say that a distribution is *learnable* if there is an approximate model whose cost per variable, $\langle \Delta L \rangle/N$, can be made small ($\leq \ell$) with a number of parameters that scales at most polynomially in N . The resulting number of parameters is far fewer than the number of possible states. For the connection between learnability and subextensivity to be useful, we require two things: First, it must be that typical probability

distributions do *not* have consistently subextensive mutual information. Second, data in interesting systems should exhibit this property.

For any partition of a system into two halves R and L , the mutual information between the two halves is

$$I_{1/2} \equiv I(R; L) = S[P_R(\sigma_R)] + S[P_L(\sigma_L)] - S[P(\sigma_R, \sigma_L)]. \quad (12)$$

If the variables are binary, and each half contains $N/2$ degrees of freedom, then the entropy of each half is bounded, $S[P_R(\sigma_R)] \leq N/2$ bits, and similarly for $S[P_L(\sigma_L)]$. Consequently, we can write

$$I_{1/2} \leq N - S[P(\sigma_R, \sigma_L)] = N - S(N). \quad (13)$$

This global upper bound is depicted by the right side of the triangle in Fig. 1.

The mutual information $I_{1/2}$ also cannot be larger than the entropy of either half system, and neither of these entropies can be larger than $S(N)$ itself. This gives the upper bound on the left side of the triangle in Fig. 1. Together with $I_{1/2} \leq N - S(N)$, this gives an allowed triangle in the plot of $I_{1/2}$ versus $S(N)$, as shown in Fig. 1. There are analogous bounds in quantum systems [49].

We define a typical probability distribution to be one in which the probabilities $P(\sigma_R, \sigma_L)$ are nearly independent random numbers, constrained only by normalization. But then the probability of each state in one half of the system,

$$P_R(\sigma_R) = \sum_{\sigma_L} P(\sigma_R, \sigma_L), \quad (14)$$

is the sum of a large number of nearly independent random variables, and from the central limit theorem this should approach its expectation value. This implies that

$$S[P_R^{\text{typical}}(\sigma_R)] \approx S[P_L^{\text{typical}}(\sigma_L)] \approx \frac{N}{2}. \quad (15)$$

Substituting into Eq. (12), we see the *typical* distribution saturates the bound, $I_{1/2} \approx N - S(N)$.

To illustrate this argument, we consider the random energy model (REM), in which each of 2^N states has an energy drawn at random from a Gaussian distribution with variance $\langle E^2 \rangle = N$, with probabilities given by the Boltzmann distribution at temperature T [50]. The states can be labeled by binary numbers and the digits assigned arbitrarily as left and right halves of a spin system. In Fig. 1(a) we show $I_{1/2}(N)$ versus $S(N)$ for these models, with varying T and N , and compare with the bounds derived above. We see that as N increases the information per spin *increases* to approach the bounds, indicating that $I_{1/2}(N)$ is extensive everywhere above the freezing transition. In contrast, models with subextensive mutual information would approach the x-axis in this plot. The peak in the mutual information versus T shows that, while the REM is unlearnable everywhere in the high-temperature regime, it is most unlearnable in an intermediate regime above T_c , while the entropy is monotonically increasing as a function of T [Fig. 1(b)].

The failure of subextensivity means literally that the influence of one half of this system on the other carries $\mathcal{O}(N)$ bits of information. Intuitively this suggests that we cannot know how one half influences the other unless we specify the




FIG. 1. Mutual information between halves of the system for the random energy model. (a) Along each curve at fixed N , we vary T , and compare with the bounds (dashed lines). (b) For $N = 22$, the mutual information between halves of the system versus T . The infinite-size system (solid line) has a cusp in the mutual information, while the entropy (inset) is monotonically increasing with T .

state completely. This idea that information is available only once we have access to all the bits in the system reminds us of cryptography, and this connection can be made precise in the context of the random energy model [51].

As a first example of real-world data, we consider ensembles of images extracted from a large database of natural movies [52]. We discretize to black/white binary pixels with a threshold such that black and white are equally likely. We then analyze contiguous patches of N pixels, and choose halves as the left/right partitions of these patches. With 1200 frames and roughly 200,000 image samples from within these frames, we are able to make reliable estimates of entropy and mutual information out to $N \sim 16$ pixels; for details see Appendix A. In Fig. 2 (inset) we see that $I_{1/2}(N)$ versus $S(N)$ moves away from the bounds with increasing N , and in the main figure we see explicitly that $I_{1/2}(N) \propto N^\alpha$ is strongly subextensive, with $\alpha = 0.1 \pm 0.03$, consistently across different natural contexts. This is a striking departure from the behavior of the random energy model, and demonstrates that real-world distributions display subextensivity. Adding noise to these image patches would degrade this feature of the distributions and they would flow toward the $I/N = 0, S/N = 1$ lower right corner of the Fig. 2 inset, also losing the scaling we observe with system size.

III. COMPRESSIBILITY

It is perhaps surprising that real-world data meet the conditions for being well approximated by a model of independent pieces. Still, this is unsatisfying, and we would like to do better. Can we build a model in which the *total* cost ΔL is finite, even as the number of degrees of freedom N becomes large? We will see that this is possible if shared information is compressible.

Let us break the N spins into two groups,

$$\vec{\sigma}_K \equiv \{\sigma_1, \sigma_2, \dots, \sigma_K\}, \quad (16)$$

$$\vec{\sigma}_{N-K} \equiv \{\sigma_{K+1}, \sigma_{K+2}, \dots, \sigma_N\}, \quad (17)$$

with $K \ll N$. The smaller group of K spins could be one of the blocks of size n_0 from above, but this is not essential. The mutual information

$$I_0(N, K) \equiv I(\vec{\sigma}_K; \vec{\sigma}_{N-K}) \quad (18)$$

is bounded by $S(K)$ and hence finite, even as $N \rightarrow \infty$, making it plausible that we do not need to specify all the details of the




FIG. 2. Mutual information between halves of image patches vs patch size in pixels. Data from snapshots out of the Chicago motion database, with different natural environments analyzed separately [52]; error bars at the largest N are a few percent, smaller than the symbols. Inset shows $I_{1/2}(N)$ vs $S(N)$, moving farther away from the bounds as N increases.




FIG. 3. Schematic of the information bottleneck. Solid line divides the forbidden (gray) region from the allowed region. If we solve the bottleneck problem with X as a discrete variable, then at fixed cardinality we vary T in Eq. (19) to trace out the dashed lines, each ending at $I(X; \vec{\sigma}_{N-K}) = \log_2(|X|)$. As $|X| \rightarrow \infty$ we approach saturation, $I(X; \vec{\sigma}_K) = I_0$, but at finite $|X|$ we miss by ϵ .

$N - K$ spins to capture their influence on the K spins. The general idea is to compress our description of $\vec{\sigma}_{N-K}$ while maintaining as much information as possible about $\vec{\sigma}_K$, and this is the information bottleneck problem [53]. Concretely, we map $\vec{\sigma}_{N-K} \rightarrow X$, maximizing

$$-\mathcal{F} = I(X; \vec{\sigma}_K) - T I(X; \vec{\sigma}_{N-K}). \quad (19)$$

We can solve this problem with X being a discrete variable of cardinality $|X|$. As $T \rightarrow 0$ we recover a deterministic mapping $\vec{\sigma}_{N-K} \rightarrow X$, and this mapping captures a fraction of the available information,

$$I_{T=0}(X; \vec{\sigma}_K) = [1 - \epsilon_N(|X|)] I_0(N, K), \quad (20)$$

where the notation reminds us that the efficiency of capturing information may depend on N .

The intuition of compressibility is that with only I_0 bits available, we should be able to express the interaction between $\vec{\sigma}_K$ and $\vec{\sigma}_{N-K}$ in roughly I_0 bits, or in a compressed variable X with $\log_2 |X| \sim I_0$. To be more precise, let us define a function $F_N(\epsilon)$, such that if we compress to within a factor F we capture information to within a factor ϵ ,

$$\log_2 |X| = F_N(\epsilon) I_0 \Rightarrow \epsilon_N(|X|) = \epsilon. \quad (21)$$

This is illustrated in Fig. 3.

Compression means that we are approximating

$$P(\vec{\sigma}_K | \vec{\sigma}_{N-K}) \approx P(\vec{\sigma}_K | X). \quad (22)$$

This approximate model has $2^K |X|$ states, and hence this many parameters. To describe the whole system we need N/K of these models, so the total number of parameters $\mathbf{P}$ is given (somewhat generously) by

$$\log_2 \mathbf{P} = K + \log_2 |X| + \log_2(N/K). \quad (23)$$

The cost of coding in this approximate model is the total mutual information we are missing,

$$\Delta L = (N/K) \epsilon_N(|X|) I_0. \quad (24)$$

So to achieve a fixed ΔL at large N , we need to have

$$\epsilon = \frac{K \Delta L}{N I_0}, \quad (25)$$

which means

$$\log_2 \mathbf{P} = K + F_N\left(\epsilon = \frac{K \Delta L}{N I_0}\right) I_0 + \log_2(N/K). \quad (26)$$

Thus, the number of parameters is set by the behavior of $F_N(\epsilon = K \Delta L / N I_0)$ at large N .

The most favorable possibility is that

$$\lim_{N \rightarrow \infty} F_N(\epsilon = 0) = f(K) \quad (27)$$

$$\Rightarrow \log_2 \mathbf{P} = K + f(K) I_0 + \log_2(N/K), \quad (28)$$

and hence $\mathbf{P} \sim N$. This is what happens in physics problems with local interactions: the impact of the $N - K$ spins on the small region of K spins can be captured by enumerating a fixed number of variables even as $N \rightarrow \infty$.

The next case is where there is a logarithmic divergence at small ϵ , so that

$$\lim_{N \rightarrow \infty} F_N\left(\epsilon = \frac{K \Delta L}{N I_0}\right) = g(K) \log_2\left(\frac{N I_0}{K \Delta L}\right) + \text{const.}, \quad (29)$$

which implies

$$\log_2 \mathbf{P} \sim [1 + g(K) I_0] \log_2 N + \text{const.} \quad (30)$$

Thus, the number of parameters is polynomial in the number of spins, although possibly with a large power.

The logarithmic behavior of $F_N(\epsilon = K \Delta L / N I_0)$ as $N \rightarrow \infty$ is realized in certain models with long-ranged interactions, including mean-field models. This is easiest to see at $K = 1$, where the impact of all $N - 1$ spins on the one spin of interest can always be summarized by an effective field h . As $N \rightarrow \infty$, this field becomes a continuous variable, chosen from a distribution $P(h)$ which could be different at every spin. Compressing the state of the $N - 1$ spins is equivalent to representing the continuous h by the discrete X ; information is lost because there is some range of h values that are assigned to the same X . If $|X|$ is large and this information loss is small, then we will have $\epsilon \sim \langle (\delta h)^2 \rangle_X$, the variance of h at fixed X . Crudely speaking, compression takes the full dynamic range H_N of the effective field, which may depend on N , and divides it into $|X|$ bins, so that

$$\langle (\delta h)^2 \rangle_X \sim \frac{H_N^2}{|X|^2}, \quad (31)$$

and hence $F_N(\epsilon) \sim \log_2(H_N^2 / \epsilon)$, so that

$$\log_2 \mathbf{P} \sim \log_2\left(\frac{N H_N^2}{\Delta L}\right) + \log_2(N), \quad (32)$$

where we drop N -independent constants.

As an example, in the disordered phase of a mean-field ferromagnet, we have $H_N \sim 1/\sqrt{N}$, which gives a number of parameters again linear in the number of spins. This still

seems like too many parameters for a mean-field model, but we have not assumed that all spins are identical, which would take $\mathbf{P} \rightarrow \mathbf{P}/N$. In contrast, if $H_N \sim 1$ at large N , we have $\mathbf{P} \propto N^2$. The last case we might worry about is if the typical field H_N grows with N , but then the entropy per spin will vanish as $N \rightarrow \infty$. Notice that these results, perhaps surprisingly, do not depend on the usual assumption of pairwise interactions, although they show that the number of parameters we need is of the same order as in a pairwise model.

While a logarithmic divergence in $F_N(\epsilon)$ leaves us with a polynomial number of parameters, a linear divergence implies that the code words needed to describe the effect of $N - K$ spins on the small cluster of K spins have $\sim N$ bits, and no compression is possible. In this case we are back to a number of parameters that is exponential in N . Together, these three examples illustrate three different regimes, from strongly compressible to fully incompressible, showing how compressibility is sufficient to bound the number of parameters $\mathbf{P}(N, \epsilon)$.

IV. COMPRESSIBILITY OF NEURAL INTERACTIONS

To see how this works in practice, we look at experiments on the activity of $N = 160$ ganglion cells in the salamander retina as it responds to naturalistic movies [26]. In these data, $\sigma_i = 1(0)$ corresponds to the presence (absence) of an action potential from neuron i in a window of duration $\Delta\tau = 20$ ms; we note at the outset that these data are not low dimensional [26,54].

We first test for subextensivity of the mutual information. In contrast to systems with local interactions, however, there is no unique way of considering groups of different size. The test will be more compelling if we have groups that share large amounts of information, so we start with one neuron and then add greedily, each time choosing the cell $N + 1$ so that $I(\sigma_{N+1}; \{\sigma_{i=1, \dots, N}\})$ is maximized. We can then estimate the mutual information between halves of the group, and these estimates are reliable up to $N \sim 10$, following the methods of Appendix A; results are shown in Fig. 4. While there is considerable variation across different groups, the mean behavior is $I_{1/2}(N)/N \sim N^{-1.39}$. For all groups of neurons, the trajectory of $I_{1/2}(N)$ versus $S(N)$ moves away from the bounding triangle as we increase N . Together these results suggest that mutual information is strongly subextensive in this system.

To analyze compressibility, we choose one neuron in the population as σ_0 , and then order the remaining neurons by their mutual information $I(\sigma_0; \sigma_i)$. To be sure that we can calibrate the fraction of information that we capture, we focus on smaller groups of K neurons. Choosing K involves a trade-off between statistical reliability and compression significance: at larger K it is more significant to find a successful compression, but it is more difficult to make reliable statistical inferences. As we will see, $K = 8$ provides an effective compromise (Appendix A). Within the bounds of our ability to make reliable estimates, choosing a different K does not change our main conclusions on the learnability and compressibility of neuronal populations (Appendix B).

The effective interactions between σ_0 and the activity of the other K neurons, $\{\sigma_{j=1, \dots, K}\}$, are described by the conditional




FIG. 4. Mutual information between halves of neuron groups vs group size. Gray traces correspond to different groups from the whole population and the black line shows the linear fit over $N = 160$ groups. As explained in the text, groups are formed greedily from the most highly correlated neurons. Inset shows $I_{1/2}(N)$ vs $S(N)$ as in Fig. 2.

distribution $P(\sigma_0|\{\sigma_j\})$, and we can always write this as an effective field $h_{\text{eff}}(\{\sigma_j\})$ acting on σ_0 ,

$$P(\sigma_0|\{\sigma_j\}) = \frac{1}{Z(\{\sigma_j\})} \exp[\sigma_0 h_{\text{eff}}(\{\sigma_j\})]. \quad (33)$$

If we can understand this distribution for each possible choice of σ_0 , then we will have understood the whole network.

With K neurons in the set $\{\sigma_j\}$, then in principle we need 2^K different values of the effective field

$$h_{\text{eff}}(\{\sigma_j\}) = \ln \left[\frac{P(\sigma_0 = 1|\{\sigma_j\})}{P(\sigma_0 = 0|\{\sigma_j\})} \right]. \quad (34)$$

A conventional simplification is to expand h_{eff} in a series,

$$h_{\text{eff}}(\{\sigma_j\}) = h_0 + \sum_j J_{0j}^{(2)} \sigma_j + \frac{1}{2} \sum_{j,k} J_{0jk}^{(3)} \sigma_j \sigma_k + \dots \quad (35)$$

Stopping with the second term corresponds to allowing only pairwise interactions in the effective Hamiltonian $H = -\ln P$, and gives a description of $P(\sigma_0|\{\sigma_j\})$ with $K + 1$ rather than 2^K parameters.

Consider the $K = 8$ neurons that share the most information with some particular σ_0 . Figure 5 shows the effective field h_{eff} as function of the state $S = \{\sigma_j\}$ for this example. We note that 218 of the 256 possible states S are visible in the data. If we try to describe these data through Eq. (35), then even including terms up to $J^{(3)}$ leaves small but significant scatter beyond the measurement errors. This is a very explicit way of seeing that stopping with $J^{(2)}$, and using a pairwise Ising model, misses significant parts of the underlying correlation structure in this system [26,36]. However, we hope not to need all the possible terms out to $J^{(8)}$. Can we show that interactions are compressible, and this captures the dependences with fewer parameters?

Compressibility means that we do not need to keep every detail of the network state $\{\sigma_j\}$ to make reliable predictions




FIG. 5. Effective field $h_{\text{eff}}(\{\sigma_j\})$ as a function of the states $\{\sigma_j\}$, in rank order. Mean, with error bars estimated from the standard deviation across random halves of data (black), and best least squares fit to Eq. (35) truncated at third order (orange). Ranked states are represented in the figure below the trace, with black indicating that a neuron is “on,” $\sigma_j = 1$. States at far right are not observed in data.

of the effective field. Concretely, this means compressing $\{\sigma_j\} \rightarrow \tilde{\sigma}$, where $\tilde{\sigma}$ takes on M states, with $M \ll 2^K$; we have changed notation from X to $\tilde{\sigma}$ to emphasize that we’re constructing coarse-grained or compressed descriptions of the variables σ . As before we are interested in the information that these compressed variables share with σ_0 , so we want to choose the mapping $\{\sigma_j\} \rightarrow \tilde{\sigma}$ that maximizes $I(\tilde{\sigma}; \sigma_0)$, and we write this maximum as $I_{\text{max}}(M)$. If we can achieve $\text{FI} = I_{\text{max}}(M)/I(\sigma_0; \{\sigma_j\}) \approx 1$ for small M even at large K , then we have tamed the combinatorial explosion.

We consider here only deterministic mappings $\{\sigma_j\} \rightarrow \tilde{\sigma}$, which means that we are solving the zero temperature or hard clustering limit of the information bottleneck problem Eq. (19) [53,55],

$$\max_{\{\sigma_j\} \rightarrow \tilde{\sigma}} I(\tilde{\sigma}; \sigma_0), \quad ||\tilde{\sigma}|| = M. \quad (36)$$

A simple algorithm for solving this problem is to start with some random assignment $\{\sigma_j\} \rightarrow \tilde{\sigma}$, then compute

$$P(\sigma_0; \tilde{\sigma}) = \sum_{\{\sigma_j\} \in \tilde{\sigma}} P(\sigma_0 | \{\sigma_j\}) P(\{\sigma_j\}), \quad (37)$$

$$P(\tilde{\sigma}) = \sum_{\{\sigma_j\} \in \tilde{\sigma}} P(\{\sigma_j\}), \quad (38)$$

with $P(\sigma_0 | \tilde{\sigma}) = P(\sigma_0; \tilde{\sigma})/P(\tilde{\sigma})$ as usual. We then reassign each particular state $\{\sigma_j\}$ to the compressed variable by minimizing the Kullback-Leibler divergence,

$$\{\sigma_j\} \rightarrow \arg \min_{\tilde{\sigma}} \sum P(\sigma_0 | \{\sigma_j\}) \log_2 \left[\frac{P(\sigma_0 | \{\sigma_j\})}{P(\sigma_0 | \tilde{\sigma})} \right]. \quad (39)$$

Iterating, we arrive at a mapping $\{\sigma_j\} \rightarrow \tilde{\sigma}$ that maximizes $I(\sigma_0; \tilde{\sigma})$.

To compare the performance of compression methods with the performance of series expansions, we use once more the idea that approximations to the true probability distribution define suboptimal codes, and the excess code length measures the cost of the approximation. In this case we are building a code for the binary variable σ_0 conditional on the state of the other K neurons; the optimal code length is L_{min} . If we have an

approximate model for $h_{\text{eff}}(\{\sigma_j\}) \approx \hat{h}_{\text{eff}}(\{\sigma_j\})$, then the mean code length is

$$L_{\text{approx}} = -\frac{1}{\ln 2} \langle \sigma_0 \hat{h}_{\text{eff}}(\{\sigma_j\}) \rangle + \langle \log_2 [1 + e^{\hat{h}_{\text{eff}}(\{\sigma_j\})}] \rangle, \quad (40)$$

where $\langle \dots \rangle$ is an average over the observed states of the system. The most naive approximation ignores interactions, assigning the same effective field to all states, and this “independent” code has length L_{ind} . A natural measure of coding cost is then

$$C = (L_{\text{approx}} - L_{\text{min}})/(L_{\text{ind}} - L_{\text{min}}), \quad (41)$$

which ranges between zero and one. For the case where we do a proper compression $\{\sigma_j\} \rightarrow \tilde{\sigma}$, then $C = 1 - \text{FI}$, where FI is the fraction of the mutual information that we capture (see above), but the coding cost is defined more generally, e.g., in the truncated series expansions of Eq. (35). In Figs. 6(a) and 6(b) we compare the optimal compressions with the series expansion, and find that compression into the best choice of $M \sim 10$ states performs as well as including 163 parameters to describe 5th order interactions.

To check that these results do not depend on our choice of the central neuron σ_0 , in Figs. 6(c) and 6(d) we show the distribution of coding cost and fractional information across these choices. We see in Fig. 6(c) that even getting within ten percent of the optimum across the majority of neurons requires extending the series expansion to fifth order, that is including terms up to $J^{(5)}$ in Eq. (35). While for specific choices of σ_0 , making truncations at lower orders might lead to reasonable descriptions of the system, we find that this is not the rule at the population level. In contrast, Fig. 6(d) shows that by compressing into $M \sim 11$ –15 states we achieve a code that captures all but a few percent of the available mutual information, for all neurons in the population. To summarize, in this network we can describe the influence of $K = 8$ neurons on one neuron using just $M = 11$ –15 parameters, but this most efficient description does not correspond to a simple choice of pairwise or other low-order interactions.

Estimates of mutual information come with errors, and so statements about the number of states needed to capture a given fraction of the information also have uncertainty (Appendix A). For each choice of σ_0 and $\{\sigma_j\}$ out of the network, estimates of FI are accompanied by an error $\Delta_{\text{FI}}(\sigma_0, \{\sigma_j\})$, and as a global measure Δ_{FI} we take the median of these errors. If we choose a fixed number of states M for the compression, then across all choices of σ_0 and $\{\sigma_j\}$ we will find a fraction D_{FI} for which the estimate of FI is larger than $1 - \Delta_{\text{FI}}$, i.e., the information captured is within errors of the information available. Figure 7(a) shows the dependence of D_{FI} on the number of states M , and we see that in 90% of all the relevant groups we achieve essentially perfect compression with $M^* \sim 11$ states. We can do this analysis not just for interactions between a single cell σ_0 and its K most informative partners, $S_1 = \{\sigma_1, \dots, \sigma_8\}$, but also for interactions with successively less informative groups $S_n = \{\sigma_{k(n-1)+1}, \dots, \sigma_{kn}\}$, and the result are the same up to $n = 8$. Note that with $n = 8$ we are covering $0.4N$ of the cells in the entire population, and that $I(\sigma_0; S_n)$ is within error bars of zero for $n > 8$.




FIG. 6. Series expansions vs compression. (a) Coding cost [Eq. (41)] as a function of the number of parameters for the series expansion in Eq. (35). Black points from analysis with all data, error bars are standard deviation across random choices of learning from 60% of the data and testing on the remaining 40%. (b) Fraction of mutual information captured as a function of the number of states M in the compressed representation $\tilde{\sigma}_i$. Error bars from analyses of random subsets of the data. (c) Coding cost probability density over all possible choices of σ_0 . Each curve correspond to a different order truncation, $J_{0j}^{(i)}$, of the series expansion [Eq. (35)]. (d) Probability density of the fractional information over all possible choices of σ_0 . Each curve correspond to a different value of M . We used a weighted-KDE method for the inference of the probability densities, considering the measured error bars of each choice of σ_0 .

V. ITERATED COMPRESSION

The compression $\{\sigma_j\} \rightarrow \tilde{\sigma}$ is reminiscent of the block spin construction in the renormalization group (RG) [56,57]. We recall that block spins are coarse-grained variables that replace groups of spins. In the present context, it is important to remember that coarse-graining can be thought of as data compression, and vice versa. By analogy with the RG, then, we would like to do iterative compression.

Concretely, we are focused on a variable σ_0 and have ordered the remaining variables σ_j by their mutual information with σ_0 . Our first coarse-graining step has been to take these variables in groups of $K = 8$, and compress according to the solution of the optimization problem in Eq. (36), which gives us

$$\begin{aligned} \{\sigma_1, \sigma_2, \dots, \sigma_8\} &\rightarrow \tilde{\sigma}_1^{(1)} \\ \{\sigma_9, \sigma_{10}, \dots, \sigma_{16}\} &\rightarrow \tilde{\sigma}_2^{(1)} \\ &\dots, \end{aligned} \quad (42)$$

where each of the variables $\tilde{\sigma}_n^{(1)}$ has M states and the superscript reminds us that this is only the first step of coarse-graining. To iterate, we take pairs of these variables and compress again, e.g.,

$$(\tilde{\sigma}_1^{(1)}, \tilde{\sigma}_2^{(1)}) \rightarrow \tilde{\sigma}_1^{(2)}, \quad (43)$$

where again the mapping is chosen to maximize the mutual information $I(\sigma_0; \tilde{\sigma}_1^{(2)})$. We can keep iterating,

$$(\tilde{\sigma}_1^{(2)}, \tilde{\sigma}_2^{(2)}) \rightarrow \tilde{\sigma}_1^{(3)}, \quad (44)$$




FIG. 7. Capturing mutual information with a limited number of states. (a) Fraction of cells σ_0 and groups $\{\sigma_j\}$ such that compression into M states captures the available mutual information, within error bars. Successive coarse-graining steps as described in the text. We define M^* as the minimum number of states needed to achieve complete compression in 90% of the cases. The black curve corresponds to the naive expectation for for the third step of coarse-graining if the number of states we need grows exponentially with the number of neurons. (b) Minimum number of states M^* as a function of the coarse-graining step. Dashed curves (gray) correspond to different compression iterations, which vary because of noise in our estimates. For comparison, a linear relation is shown in orange.

always with the same principle of choosing the compression that maximizes the mutual information with σ_0 . Note that in our initial compression $\{\sigma_j\} \rightarrow \tilde{\sigma}^{(1)}$, we chose $K = 8$ cells and hence 256 states. In the second step, Eq. (43), we start with $M^2 \sim 256$ states again, which means that we have the same high level of control over sampling problems. At the third step, Eq. (44), we start with somewhat more states, but still sampling is under control, and successive coarse-graining steps are entropically comparable.

It is not surprising that successive stages of compression or coarse-graining require more states to capture all the available mutual information [Fig. 7(a)]. What is surprising is that the minimal number of states M^* seems to grow linearly rather than exponentially as we proceed through multiple stages, as seen in Fig. 7(b). After three stages, we are describing the interactions of σ_0 with 32 other cells using only $M_3^* = 32$ states.

If we did not know about subextensivity we might naively expect that to capture information from a group of N neurons requires $M \sim 2^{\gamma N}$ states, or more generally that our success in capturing information should depend on M and N as $D_{\text{FI}}(M, N) = f(M/e^{\gamma N})$. An estimate of γ comes from $M_2^*/M_1^* = e^{\gamma(N_2-N_1)}$, where $N_1 = 8$ and $N_2 = 16$ are the sizes of the group at the first steps of coarse-graining, so that $\gamma = 0.1026$. We then can use this to predict D_{FI} at the third stage of coarse-graining as a stretched version of the curve at the second stage (black curve in Fig. 7), where the stretching factor along the M axis is $e^{\gamma(N_3-N_2)}$, with $N_3 = 32$. We see that the real data are dramatically different.

The linear growth of M^* with the number of neurons is explicit evidence that we have tamed the combinatorial explosion, combining the compressibility of interactions with an RG-inspired iteration scheme. The scaling of M^* is what we might expect in a model with pairwise interactions, or if single neurons coupled only to the total activity of other neurons, but neither of these simplifications is correct.

As a further test of these ideas we have looked at experiments on a very different network of neurons, in the mouse hippocampus [27,30]. The results, described in Appendix B, are very much the same, but perhaps less surprising since maximum entropy models with only pairwise interactions already provide an excellent description of these data, matching the higher order correlations within experimental error [27,36]. In contrast, as emphasized in Ref. [26], for the population of cells in the retina the pairwise models show small but significant deviations from the data, and this has led to the exploration of several alternatives [26,33,34].

Where previous work has focused on simple forms of the interactions, taking intuition both from physics and from neurobiology, the point of our discussion is not to identify the correct model, but to understand why *any* simple model can succeed. Indeed, the results of this approach in two very different neuronal populations suggest that compressibility is a more general and intrinsic property of large neuronal networks.

VI. IMPLICATIONS FOR THE NEURAL CODE

We have emphasized that compressibility allows us to give a simpler description of correlation structure in the patterns

of activity seen in populations of retinal ganglion cells. It is important that compressibility also has implications for the functional behavior of the network as it encodes the visual world.

The fact that we can compress the interactions means that we have a good estimate of the information that each neuron shares with the rest of the network, $I(\sigma_0; \tilde{\sigma})$. But what is this information? A natural guess is that information shared among visual neurons is information about the visual world. To test this, we note that the experiments in Ref. [26] include many repeated presentations of the same movie. This repetition means we can estimate the distribution of neural activity conditional on the visual stimulus $\mathbf{s}$, and consequently the information that neurons carry about the stimulus, $I(\sigma_0; \mathbf{s})$.

Information estimates would be difficult if we were to ask about the patterns of activity across many cells. But if we ask about just one cell there is more than enough data in these experiments to reach reliable conclusions about the information that this cell's response carries about the visual input [58,59]. Importantly, these information estimates $I(\sigma_0; \mathbf{s})$ are independent of model assumptions about which features of the visual stimulus are being represented and independent of what is happening in the other neurons of the network. Thus, while we can think $I(\sigma_0; \tilde{\sigma})$ as a measure of population-level correlations, $I(\sigma_0; \mathbf{s})$ is an independent measure of visual information.

Figure 8 shows the information that each single cell carries about the visual stimulus, $I(\sigma_0; \mathbf{s})$, compared with the information that this single neuron shares with the network, $I(\sigma_0; \tilde{\sigma})$. We see that the intuition connecting shared information with visual information is surprisingly accurate. In fact, as we consider larger groups of neurons, the shared information approaches the visual information more and more closely, within (small) error bars. This equality has a surprising consequence.

Suppose that we ask not about the visual information carried by a single neuron, but rather about the *extra* information that this neuron carries beyond what all the other neurons tell the brain about the visual inputs. Formally, this extra information is

$$\Delta I(\sigma_0; \mathbf{s}) = I(\{\sigma_j\}, \sigma_0; \mathbf{s}) - I(\{\sigma_j\}; \mathbf{s}). \quad (45)$$

Following arguments about synergy and redundancy among different components of the neural response [59], we can rewrite the extra information in terms of shared information (see Appendix C for details):

$$\Delta I(\sigma_0; \mathbf{s}) = I(\sigma_0; \mathbf{s}) + I(\sigma_0; \{\sigma_j\}|\mathbf{s}) - I(\sigma_0; \{\sigma_j\}). \quad (46)$$

In this expression, $I(\sigma_0; \mathbf{s})$ is the information that a single neuron carries about the visual stimulus and $I(\sigma_0; \{\sigma_j\})$ is the information that it shares with the network, as before, while $I(\sigma_0; \{\sigma_j\}|\mathbf{s})$ is the (average) mutual information between σ_0 and the network given that visual stimulus is known.

The fact that we achieve effective compression of the shared information means we can replace $I(\sigma_0; \{\sigma_j\})$ by $I(\sigma_0; \tilde{\sigma})$. But then Fig. 8 tells us that $I(\sigma_0; \tilde{\sigma}) \approx I(\sigma_0; \mathbf{s})$, so that the first and last terms in Eq. (46) cancel, and we are left with

$$\Delta I(\sigma_0; \mathbf{s}) \approx I(\sigma_0; \{\sigma_j\}|\mathbf{s}). \quad (47)$$




FIG. 8. Shared information vs visual information. a) Information that a single neuron shares with the network vs information that this neuron conveys about the stimulus. Plot corresponds to the first (blue, 8 neurons), second (green, 16 neurons) and third (red, 32 neurons) coarse-graining steps. Estimates and errors as in Ref. [58].) Kernel density estimator (KDE) of the difference between the information shared with the stimulus vs information about the network for each step. Dashed lines correspond to the medians.

This tells us that the retina is operating in a regime where the extra information provided by a single neuron is nearly equal to the information that this neuron shares with the network given that we know the visual stimulus.

A popular model for retinal coding is that individual neurons respond independently to the visual inputs, so that all correlations are inherited from the stimulus;² formally this conditionally independent model is defined by

$$P(\{\sigma_j\}|\mathbf{s}) = \prod_j Q_j(\sigma_j|\mathbf{s}). \quad (48)$$

If this is true, then $I(\sigma_0; \{\sigma_j\}|\mathbf{s}) = 0$ and the neuron at the center of our analysis would be completely redundant with the other K neurons, $\Delta I(\sigma_0; \mathbf{s}) \approx 0$. Stated in a more positive way, the global correlation structure of the retinal population is such that the extra information carried by individual neurons depends entirely on their departure from conditional independence.

Correlations between neurons that persist even when the stimulus is fixed are sometimes called “noise correlations” [61–68]. There is ample experimental evidence for these correlations among retinal ganglion cells, as reviewed for example in Ref. [54], and there are clear mechanisms that could generate such correlations, including the flow of noise from the receptor cells through the retinal circuitry [69]. Nonetheless noise correlations continue to be seen as second order effects, and the conditions under which these correlations can enhance the information content or efficiency of the neural code seem exotic. It thus comes as surprise that the retina is operating in a limit where the information carried by noise correlations is essentially *equal* to the incremental information contributed by a single neuron.

VII. DISCUSSION

To summarize, the consistent subextensivity of mutual information is a determinant of learnability that makes possible approximate models with a number of parameters that is linear in the number of degrees of freedom while suffering a cost

per degree of freedom that vanishes in the thermodynamic limit, and compressibility of the mutual information makes it possible to have only finite total cost in this limit. These results suggest, strongly, that complexity can be tamed without making assumptions about the nature of interactions, generalizing our intuition from physics problems that we understand. Perhaps this also provides new perspectives on why simple models work in the traditional problems of statistical physics.

It is important that the ideas of consistent subextensivity and compressibility apply to real data. We are especially struck by the fact that neural activity in the retina obeys these conditions, providing a basis to understand why the neural code is learnable [33]. As a consequence of learnability, we can perform an iterative, RG-like compression of the effective interactions between neurons that leads to a description in which the influence of N neurons on one central neuron is described by $\sim N$ parameters, even though these interactions are not pairwise. This tells us that the collective nature of the retinal code [70] does not require us to learn higher-order interactions order by order. More work is required to exploit this observation in constructing a full model for the joint distribution of activity across a large network, but our results show that such simplified models are possible with essentially zero information loss.

This compressibility has another intriguing consequence in the context of population coding in the brain. Almost from the start of simultaneous recordings from multiple neurons there has been an attempt to separate “signal” and “noise” correlations, often in the hope that noise correlations are small and could be neglected [60,63,64]. While there are examples where noise correlations may have a structure that enhances neural coding [67,68,71,72], we are not aware of an example in which these correlations almost perfectly cancel the redundancy between individual neurons and the larger population, which is what we find here.

The near perfect compressibility of the effective interactions found here tames the combinatorial explosion naïvely expected for collective codes. Exploiting this result we find that the visual information that one cell’s response adds to the information carried by N other cells in the network is essentially equal to the natural measure of noise correlation in bits. In this sense, the system operates in a regime where noise correlations are essential for coding.

²This idea has many origins; it has been used as a simplifying hypothesis but also as a conjectured principle. A particularly strong form of the idea is presented in Ref. [60].

We have restricted ourselves to samples from visual scenes and neural responses at one moment in time, but we could generalize to sequences of states over time. The natural division of a trajectory into halves would be to look at the past and future relative to some midpoint, and thus the analog of $I_{1/2}$ is the mutual information between past and future or, more simply, the predictive information in the time series. Because interactions do not extend over arbitrarily long temporal scales, subextensivity of the predictive information is automatic [73], and this subextensivity is what guarantees the convergence to a well-defined entropy rate for the trajectory. There has been interest in analyzing predictive information both for movies [74] and in the responses of the retinal neurons to such inputs (as considered here) [75–78]. A discussion of learnability for models of these time series would add another layer to the broader effort to understand the learnability of natural distributions, and should connect to current excitement about large language models.

We have phrased the problem of understanding a network of neurons as being able to write a good approximation for the joint distribution of activity across the whole population, essentially being able to predict the likelihood of seeing any of the 2^N combinations of spiking and silence in the network [36]. The motivation is the analogy to equilibrium statistical mechanics, where being able to write the Boltzmann distribution gives us a starting point for calculating all static physical properties of the system. Our exploration of the compressibility of interactions in the network gives us a clear answer to the question of why such simplified models exist, and a surprisingly direct path to conclusions about the function of the retina as an encoder of the visual world.

ACKNOWLEDGMENTS

We thank M. Bauer, R. Dickman, I. Nemenman, V. Ngampruetikorn, and A. Tan for helpful discussions, and we thank our experimental colleagues D. Amodi, M. J. Berry II, C. D. Brody, J. L. Gauthier, O. Marre, and D. W. Tank for sharing the data of Refs. [26,27,30]. This work was supported in part by the US National Science Foundation, through the Center for the Physics of Biological Function (Grant No. PHY-1734030), the Center for the Science of Information (Grants No. CCF-0939370 and No. PHY-1607612; by the National Institutes of Health BRAIN initiative (Grant No. R01EB026943-01); by the Simons Foundation; by the John Simon Guggenheim Memorial Foundation; by the NSF-Simons National Institute for Theory and Mathematics in Biology, Award No. NSF DMS-2235451 and Simons Foundation Grant No. MP-TMPS-00005320; by the Conselho Nacional de Desenvolvimento Científico e Tecnológico, CNPq, Brazil; and by FAPESP Grant No. 2016/01343-7.

DATA AVAILABILITY

The data that support the findings of this article are openly available [79].

APPENDIX A: ESTIMATION OF MUTUAL AND FRACTIONAL INFORMATION FOR FINITE DATA

Our strategy for estimating information theoretic quantities follows that used in Ref. [58]: we vary the number of sam-

ples that we use in making our estimates, verify that we are in the regime where sample size dependence is as expected from perturbation theory, and extrapolate to infinite data. This approach has a long history, and is reasonably well known; a review can be found in Sec. A.8 of Ref. [80]. We give some details here, in the hope of making the discussion more accessible.

Estimating the mutual information between a single neuron, σ_0 , and a group of K neurons, $\sigma \equiv \{\sigma_{j=1,\dots,K}\}$, requires inferring the corresponding 2^{K+1} state probabilities. To begin, as shown in Fig. 9(a), the choice of $K = 8$ allows the observation of >90% of the states for almost all possible choices of σ_0 . We then choose random fractions of the data, $f = (50, 60, 70, 80, 90)\%$, and calculate the mutual information from these limited samples. If we see the expected simple dependence on the (inverse) number of samples, then we extrapolate to infinite data, and error bars are calculated as the standard deviation from random halves of the data.

We test our estimation procedure in a model with the effective fields $h_{\text{eff}}(\{\sigma_j\})$ are drawn from a normal distribution with zero mean and unit variance, and the distribution over states $P(\{\sigma_j\})$ is Zipf. In Fig. 9(b) we see the expected linear dependence of the information estimate on the inverse number of samples, and the extrapolation matches the exact answer for this model example. We propagate the error to our estimates of the fractional information, Δ_{FI} . Figure 9(c) shows that the estimate of the FI for each choice of σ_0 comes with a different error. Consequently, we use the median standard deviation, over all choices of σ_0 , as the reference to obtain the number of states at the population level (see Fig. 6).

APPENDIX B: COMPRESSION FOR A POPULATION OF NEURONS IN THE HIPPOCAMPUS

We also have tested compressibility in a population of $N = 1485$ neurons in the mouse hippocampus [27,30]. As in the retinal population, neurons are described with the two states $\sigma_i = 1(0)$ corresponding to the presence (absence) of activity. In these data the activity is monitored by fluorescent proteins that are sensitive to the intracellular calcium concentration, which provides a slower and coarser readout than direct electrical measurements, but again we can discretize into binary variables.

Neuronal activity in this population is more sparse than in the retinal population, leading to fewer but still a large number of observable states. Describing this population of neurons with a pairwise approximation of the series expansion in Eq. (35) leads to a scatter similar to that observed in the retina with a third order approximation [Fig. 10(a)], although this pairwise approximation is known to capture many collective properties of the system quite accurately [36]. Following the same formalism as before, we calculate the coding cost and the fractional information at the population level, showing that compressibility also is feasible in this population of neurons [Fig. 10(b)]. Our compression approach outperforms the series expansion of h_{eff} when we compare models with same number of parameters [Fig. 10(c)]. Finally, we implement our iterative coarse-graining algorithm to describe larger populations and find that, as in the retina case, a compression approach exhibits linear growth of number of states that we need (M) as a function of the number of neurons, showing that we can




FIG. 9. (a) Probability density for the fraction of possible states $\{\sigma_j\}$ that we observe in the data, for different choices of the neighborhood size K . The state ratio is defined as the number of states S_{exp} that we find relative to the possible number 2^K . The distribution is across many randomly chosen groups of K neurons from the full population of $N = 160$ cells. (b) Estimates of mutual information as a function of the (inverse) number of samples for the Zipf-like model described in the text. Examples (black), means (red circles), linear fit (red line) and extrapolation with errors; exact result shown for comparison with expected estimation errors (blue). (c) Probability distribution of the fractional information error across all σ_0 .

tame the exponential growth [Figs. 10(d) and 10(e)]. We have done our analysis for groups of $K = 7$ and $K = 8$ neurons, finding a similar (and very slow) linear growth [Fig. 10(e)].

APPENDIX C: DERIVATION OF EQ. (46)

We would like to rewrite the extra visual information carried by a single neuron, Eq. (45), in terms of information shared between that neuron and the network. To do this we follow Ref. [59], and make the various information terms

more explicit. We start with the definition:

$$\begin{aligned} \Delta I(\sigma_0; \mathbf{s}) &\equiv I(\{\sigma_j\}, \sigma_0; \mathbf{s}) - I(\{\sigma_j\}; \mathbf{s}) \\ &= \left\langle \log_2 \left[\frac{P(\{\sigma_j\}, \sigma_0; \mathbf{s})}{P(\{\sigma_j\}, \sigma_0)P(\mathbf{s})} \right] \right\rangle \\ &\quad - \left\langle \log_2 \left[\frac{P(\{\sigma_j\}; \mathbf{s})}{P(\{\sigma_j\})P(\mathbf{s})} \right] \right\rangle. \end{aligned} \quad (\text{C1})$$




FIG. 10. Compressibility in a hippocampal population. (a) The effective field $h_{\text{eff}}(\sigma)$ as a function of the states σ , in rank order for $K = 8$. Mean, with error bars estimated from the standard deviation across random halves of the data (black) and best least squares fit to Eq. (35) truncated at third order (orange). Ranked states are represented in the figure below the trace, with black indicating that a neuron is active, $\sigma_j = 1$. Only experimentally observed states are shown. (b) Probability density of the fractional information over all possible choices of σ_0 . Each curve correspond to a different value of M . We used a weighted-KDE method for the inference of the probability densities, considering the measured error bars at each choice of σ_0 . (c) Probability distribution of coding costs across all choices of σ_0 . Each curve correspond to a different order truncation, $J_{0j}^{(l)}$, of the series expansion. (d) Fraction of cells σ_0 and groups $\sigma^{(l)}$ such that compression into M states captures the available mutual information, within error bars for $K = 8$. Colors represent successive coarse-graining steps as described in the text. We define M^* as the minimum number of states needed to achieve complete compression in 90% of the cases. (e) Number of states as a function of the number of neurons for $K = 7$ (orange markers) and $K = 8$ (black markers). The dashed line shows the possible number of states 2^K .

We then group the terms and insert factors of unity:

$$\Delta I(\sigma_0; \mathbf{s}) = \left\langle \log_2 \left[\frac{P(\{\sigma_j\}, \sigma_0; \mathbf{s}) P(\{\sigma_j\})}{P(\{\sigma_j\}; \mathbf{s}) P(\{\sigma_j\}, \sigma_0)} \right] \right\rangle \quad (\text{C2})$$

$$= \left\langle \log_2 \left[\frac{P(\{\sigma_j\}, \sigma_0; \mathbf{s}) P(\{\sigma_j\})}{P(\{\sigma_j\}; \mathbf{s}) P(\{\sigma_j\}, \sigma_0)} \right] \right\rangle + \left\langle \log_2 \left[\frac{P(\sigma_0)}{P(\sigma_0)} \cdot \frac{P(\mathbf{s})}{P(\mathbf{s})} \cdot \frac{P(\sigma_0 | \mathbf{s})}{P(\sigma_0 | \mathbf{s})} \right] \right\rangle. \quad (\text{C3})$$

Now, we rearrange:

$$\Delta I(\sigma_0; \mathbf{s}) = \left\langle \log_2 \left[\frac{P(\sigma_0 | \mathbf{s}) P(\mathbf{s})}{P(\sigma_0) P(\mathbf{s})} \right] \right\rangle + \left\langle \log_2 \left[\frac{P(\{\sigma_j\}, \sigma_0 | \mathbf{s})}{P(\{\sigma_j\} | \mathbf{s}) P(\sigma_0 | \mathbf{s})} \right] \right\rangle + \left\langle \log_2 \left[\frac{P(\{\sigma_j\}) P(\sigma_0)}{P(\{\sigma_j\}, \sigma_0)} \right] \right\rangle. \quad (\text{C4})$$

Finally, we recognize each of these terms as a mutual information, so that

$$\Delta I(\sigma_0; \mathbf{s}) = I(\sigma_0; \mathbf{s}) + I(\sigma_0; \{\sigma_j\} | \mathbf{s}) - I(\sigma_0; \{\sigma_j\}). \quad (\text{C5})$$

-
- [1] J. Sethna, *Statistical Mechanics: Entropy, Order Parameters, and Complexity* (Oxford University Press, Oxford, UK, 2006).
- [2] C. Itzykson and J. M. Drouffe, *Statistical Field Theory* (Cambridge University Press, Cambridge, UK, 1991).
- [3] E. Lubeck and L. Cai, Single-cell systems biology by super-resolution imaging and combinatorial labeling, *Nat. Methods* **9**, 743 (2012).
- [4] A. M. Klein *et al.*, Droplet barcoding for single cell transcriptomics applied to embryonic stem cells, *Cell* **161**, 1187 (2015).
- [5] E. Z. Macosko *et al.*, Highly parallel genome-wide expression profiling of individual cells using nanoliter droplets, *Cell* **161**, 1202 (2015).
- [6] K. H. Chen, A. N. Boettiger, J. R. Moffitt, S. S. Wang, and X. Zhuang, Spatially resolved, highly multiplexed RNA profiling in single cells, *Science* **348**, aaa6090 (2015).
- [7] 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).
- [8] 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).
- [9] O. Marre, D. Amodèi, N. Deshmukh, K. Sadeghi, F. Soo, T. E. Holy, and M. J. Berry II, Mapping a complete neural population in the retina, *J. Neurosci.* **32**, 14859 (2012).
- [10] M. B. Ahrens, M. B. Orger, D. N. Robson, J. M. Li, and P. J. Keller, Whole-brain functional imaging at cellular resolution using light-sheet microscopy, *Nat. Methods* **10**, 413 (2013).
- [11] J. J. Jun *et al.*, Fully integrated silicon probes for high-density recording of neural activity, *Nature (London)* **551**, 232 (2017).
- [12] J. E. Chung *et al.*, High-density, long-lasting, and multi-region electrophysiological recordings using polymer electrode arrays, *Neuron* **101**, 21 (2019).
- [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] A. Cavagna and I. Giardina, Bird flocks as condensed matter, *Annu. Rev. Condens. Matter Phys.* **5**, 183 (2014).
- [15] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature (London)* **521**, 436 (2015).
- [16] P. Mehta, M. Bukov, C.-H. Wang, A. G. R. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, A high-bias, low-variance introduction to machine learning for physicists, *Phys. Rep.* **810**, 1 (2019).
- [17] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [18] B. M. Yu, J. P. Cunningham, G. Santhanam, S. I. Ryu, K. V. Shenoy, and M. Sahani, Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity, *J. Neurophysiol.* **102**, 614 (2009).
- [19] J. A. Gallego, M. G. Perich, L. E. Miller, and S. A. Solla, Neural manifolds for the control of movement, *Neuron* **94**, 978 (2017).
- [20] M. K. Transtrum, B. B. Machta, K. S. Brown, B. C. Daniels, C. R. Myers, and J. P. Sethna, Perspective: Sloppiness and emergent theories in physics, biology, and beyond, *J. Chem. Phys.* **143**, 010901 (2015).
- [21] K. N. Quinn, H. Wilber, A. Townsend, and J. P. Sethna, Chebyshev approximation and the global geometry of model predictions, *Phys. Rev. Lett.* **122**, 158302 (2019).
- [22] 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).
- [23] W. Bialek and R. Ranganathan, Rediscovering the power of pairwise interactions, *arXiv:0712.4397*.
- [24] 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).
- [25] 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).
- [26] G. Tkačik, O. Marre, D. Amodèi, E. Schneidman, W. Bialek, and M. J. Berry II, Searching for collective behavior in a large network of sensory neurons, *PLoS Comput. Biol.* **10**, e1003408 (2014).
- [27] 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).
- [28] H. C. Nguyen, R. Zecchina, and J. Berg, Inverse statistical problems: From the inverse Ising problem to data science, *Adv. Phys.* **66**, 197 (2017).
- [29] A. Cavagna, I. Giardina, and T. Grigera, The physics of flocking: Correlation as a compass from experiments to theory, *Phys. Rep.* **728**, 1 (2018).
- [30] L. Meshulam, J. L. Gauthier, C. D. Brody, D. W. Tank, and W. Bialek, Coarse-graining, fixed points, and scaling in a large population of neurons, *Phys. Rev. Lett.* **123**, 178103 (2019).

- [31] G. B. Morales, S. Di Santo, and M. A. Muñoz, Quasi-universal scaling in mouse brain neuronal activity stems from edge-of-instability critical dynamics, *Proc. Natl. Acad. Sci. USA* **120**, e2208998120 (2023).
- [32] B. R. Munn, E. Müller, I. Favre-Bulle, E. Scott, J. T. Lizier, M. Breakspear, and J. M. Shine, Multiscale organization of neuronal activity unifies scale-dependent theories of brain function, *Cell* **187**, 7303 (2024).
- [33] E. Ganmor, R. Segev, and E. Schneidman, Sparse low-order interaction network underlies a highly correlated and learnable neural population code, *Proc. Natl. Acad. Sci. USA* **108**, 9679 (2011).
- [34] J. Humplik and G. Tkačik, Probabilistic models for neural populations that naturally capture global coupling and criticality, *PLoS Comput. Biol.* **13**, e1005763 (2017).
- [35] O. Maoz, G. Tkačik, M. S. Esteki, R. Kiani, and E. Schneidman, Learning probabilistic neural representations with randomly connected circuits, *Proc. Natl. Acad. Sci. USA* **117**, 25066 (2020).
- [36] L. Meshulam and W. Bialek, Statistical mechanics for networks of real neurons, *Rev. Mod. Phys.* **97**, 045002 (2025).
- [37] L. G. Valiant, A theory of the learnable, *Commun. ACM* **27**, 1134 (1984).
- [38] T. L. H. Watkin, A. Rau, and M. Biehl, The statistical mechanics of learning a rule, *Rev. Mod. Phys.* **65**, 499 (1993).
- [39] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms* (Cambridge University Press, Cambridge, UK, 2014).
- [40] W. Bialek, S. E. Palmer, and D. J. Schwab, What makes it possible to learn probability distributions in the natural world? [arXiv:2008.12279](https://arxiv.org/abs/2008.12279).
- [41] L. Ramirez and W. Bialek, Compression as a path to simplification: Models of collective neural activity, [arXiv:2112.14334](https://arxiv.org/abs/2112.14334).
- [42] A. Grattafiori *et al.*, The Llama 3 herd of models, [arXiv:2407.21783](https://arxiv.org/abs/2407.21783).
- [43] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, Understanding deep learning (still) requires rethinking generalization, *Commun. ACM* **64**, 107 (2021).
- [44] Y. Zhang, P. Tiño, A. Leonardis, and K. Tang, Survey on neural network interpretability, *IEEE Trans. Emerging Top. Comput. Intell.* **5**, 726 (2021).
- [45] C. E. Shannon, Prediction and entropy of written English, *Bell Sys. Tech. J.* **30**, 50 (1951).
- [46] C. E. Shannon, A mathematical theory of communication, *Bell Sys. Tech. J.* **27**, 379 (1948).
- [47] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley, New York, NY, 1991).
- [48] J. Eisert, M. Cramer, and M. B. Plenio, *Colloquium: Area laws for the entanglement entropy*, *Rev. Mod. Phys.* **82**, 277 (2010).
- [49] D. N. Page, Average entropy of a subsystem, *Phys. Rev. Lett.* **71**, 1291 (1993).
- [50] B. Derrida, Random-energy model: An exactly solvable model of disordered systems, *Phys. Rev. B* **24**, 2613 (1981).
- [51] V. Ngampruetikorn and D. J. Schwab, Random-energy secret sharing via extreme synergy, *Phys. Rev. Lett.* **133**, 257401 (2024).
- [52] <https://cmd.rcc.uchicago.edu>.
- [53] N. Tishby, F. C. Pereira, and W. Bialek, The information bottleneck method, in *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing*, edited by B. Hajek and R. S. Sreenivas (University of Illinois, Urbana, IL, 1999), Vol. 37, pp. 368–377.
- [54] 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).
- [55] D. J. Strouse and D. J. Schwab, The deterministic information bottleneck, *Neural Comput.* **29**, 1611 (2017).
- [56] L. P. Kadanoff, Scaling laws for Ising models near T_c , *Physics Physique Fizika* **2**, 263 (1966).
- [57] J. Cardy, *Scaling and Renormalization in Statistical Physics* (Cambridge University Press, Cambridge, UK, 1996).
- [58] 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).
- [59] N. Brenner, S. P. Strong, R. Koberle, W. Bialek, and R. R. de Ruyter van Steveninck, Synergy in a neural code, *Neural Comput.* **12**, 1531 (2000).
- [60] S. Nirenberg, S. M. Carcieri, A. L. Jacobs, and P. E. Latham, Retinal ganglion cells act largely as independent encoders, *Nature (London)* **411**, 698 (2001).
- [61] H. G. Eyherabide and I. Samengo, When and why noise correlations are important in neural decoding, *J. Neurosci.* **33**, 17921 (2013).
- [62] H. Sompolinsky, H. Yoon, K. Kang, and M. Shamir, Population coding in neuronal systems with correlated noise, *Phys. Rev. E* **64**, 051904 (2001).
- [63] A. Ecker, P. Berens, A. Tolias, and M. Bethge, The effect of noise correlations in populations of diversely tuned neurons, *Journ. Neuro.* **31**, 14272 (2011).
- [64] B. B. Averbeck, P. E. Latham, and A. Pouget, Neural correlations, population coding and computation, *Nat. Rev. Neurosci.* **7**, 358 (2006).
- [65] I. Kanitscheider, R. Coen-Cagli, and A. Pouget, Origin of information-limiting noise correlations, *Proc. Natl. Acad. Sci. USA* **112**, E6973 (2015).
- [66] Y. Hu, J. Zylberberg, and E. Shea-Brown, The sign rule and beyond: Boundary effects, flexibility, and noise correlations in neural population codes, *PLoS Comput. Biol.* **10**, e1003469 (2014).
- [67] J. Cafaro and F. Rieke, Noise correlations improve response fidelity and stimulus encoding, *Nature (London)* **468**, 964 (2010).
- [68] J. Zylberberg, J. Cafaro, M. H. Turner, E. Shea-Brown, and F. Rieke, Direction-selective circuits shape noise to ensure a precise population code, *Neuron* **89**, 369 (2016).
- [69] B. D. Hoshal, C. M. Holmes, K. Bojanek, J. Salisbury, M. J. Berry, O. Marre, and S. E. Palmer, Stimulus-invariant aspects of the retinal code drive discriminability of natural scenes, *Proc. Nat. Acad. Sci.* **121**, e2313676121 (2024).
- [70] 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).
- [71] X. Ding, D. Lee, J. B. Melander, S. Ganguli, and S. A. Baccus, Adaptation of retinal discriminability to natural scenes, *bioRxiv* (2024).
- [72] M. Manu, L. T. McIntosh, D. B. Kastner, B. N. Naecker, and S. A. Baccus, Synchronous inhibitory pathways create both

- efficiency and diversity in the retina, *Proc. Natl. Acad. Sci. USA* **119**, e2116589119 (2022).
- [73] W. Bialek, I. Nemenman, and N. Tishby, Predictability, complexity, and learning, *Neural Comput.* **13**, 2409 (2001).
- [74] J. M. Salisbury and S. E. Palmer, A dynamic scale-mixture model of motion in natural scenes, *eLife* **14**, e104054 (2025).
- [75] S. E. Palmer, O. Marre, M. J. Berry, and W. Bialek, Predictive information in a sensory population, *Proc. Natl. Acad. Sci. USA* **112**, 6908 (2015).
- [76] A. J. Sederberg, J. N. MacLean, and S. E. Palmer, Learning to make external sensory stimulus predictions using internal correlations in populations of neurons, *Proc. Natl. Acad. Sci. USA* **115**, 1105 (2018).
- [77] M. Chalk, O. Marre, and G. Tkačik, Toward a unified theory of efficient, predictive, and sparse coding, *Proc. Natl. Acad. Sci. USA* **115**, 186 (2018).
- [78] B. Liu, A. Hong, F. Rieke, and M. B. Manookin, Predictive encoding of motion begins in the primate retina, *Nat. Neurosci.* **24**, 1280 (2021).
- [79] GitHub repository, https://github.com/luf Ramirezoc/Ramirez_et.al._2026_probabilistic.
- [80] W. Bialek, *Biophysics: Searching for Principles* (Princeton University Press, Princeton, NJ, 2012).