

Common Input Explains Higher-Order Correlations and Entropy in a Simple Model of Neural Population Activity

Jakob H. Macke,^{1,*} Manfred Opper,² and Matthias Bethge³

¹*Gatsby Computational Neuroscience Unit, University College London, London, United Kingdom and University of Tübingen, Tübingen, Germany*

²*Artificial Intelligence Group, Technical University Berlin, Berlin, Germany*

³*Institute for Theoretical Physics, Werner Reichardt Centre, Bernstein Center for Computational Neuroscience, MPI for Biological Cybernetics and University of Tübingen, Tübingen, Germany*

(Received 11 September 2010; published 17 May 2011)

Simultaneously recorded neurons exhibit correlations whose underlying causes are not known. Here, we use a population of threshold neurons receiving correlated inputs to model neural population recordings. We show analytically that small changes in second-order correlations can lead to large changes in higher-order redundancies, and that the resulting interactions have a strong impact on the entropy, sparsity, and statistical heat capacity of the population. Our findings for this simple model may explain some surprising effects recently observed in neural population recordings.

DOI: [10.1103/PhysRevLett.106.208102](https://doi.org/10.1103/PhysRevLett.106.208102)

PACS numbers: 87.19.lj, 87.19.lm, 87.19.1o, 89.70.Cf

Modeling the statistical structure of firing patterns distributed across multiple neurons is a major challenge in neuroscience. Recently, the Ising model [1] has become popular for studying neural population recordings [2–4]. Its use for neural data analysis is motivated by the fact that it is the maximum entropy (MaxEnt) model for given second-order correlations [5,6]. Therefore, deviations from the model indicate that higher-order correlations have to be taken into account for modeling the population statistics. (See [7] for alternative approaches.)

Recent studies have reported that higher-order redundancies (i.e., reductions in entropy due to the presence of higher-order interactions) are weak in small retinal populations [2,3], but that they are substantial in local cortical populations [4]. It has also been shown that they are dynamically modulated by the stimulus [4], and that they affect the coding properties [8] and the sparsity of the population [4]. However, there has arguably been a lack of theoretical models that could explain this diversity of findings. In particular, as most experimental studies have been limited to small populations, current models do not answer fundamental questions about the behavior of the entropy [2,9] and the statistical heat capacity [10] of large populations.

Here, we provide a parsimonious, mathematically tractable model which can account for this multitude of empirical observations. In our model, correlations between neurons arise from Gaussian inputs into threshold neurons. The model is therefore equivalent to the Dichotomized Gaussian distribution (DG) [11,12], which has previously also been investigated by Amari *et al.* [13] to study correlations in neural populations. Although the inputs to the model are Gaussian distributed and therefore have no interactions beyond second order, the nonlinear spiking thresholds give rise to statistical interactions of all orders.

Importantly, the magnitude of the resulting redundancies is strongly dependent on second-order correlations, and in a manner which is quantitatively consistent with neural recordings [4,14]. Using this model, we investigate the scaling of the entropy and redundancies in large neural populations, and show their impact on the sparsity of the population. Finally, we find that our model, because of its higher-order interactions, has a diverging heat capacity and a critical point at unit temperature, as has been found empirically in retinal populations [10] and natural images [15].

The Dichotomized Gaussian is a model of correlated input.—We model a population of n binary neurons X_i , where a neuron is said to spike ($X_i = 1$) if its input is positive, and to be silent ($X_i = 0$) otherwise. The inputs are modeled by a correlated Gaussian with mean γ and covariance Λ , which are chosen such that the outputs X have mean μ and covariance Σ (see [12] for details). Figure 1(a) shows that this model has a characteristic relationship between correlations and firing probabilities which is similar to that found in neural recordings [16]. For analytical tractability, we here focus on homogeneous populations, i. e., $\mu_i = \mu$ and $\Sigma_{ij} = \sigma$, $\Lambda_{ij} = \lambda \mathbb{V}(i \neq j)$ [1,8,13], and define the pairwise correlation coefficient $\rho = \sigma / [\mu(1 - \mu)]$. By symmetry, all activity patterns with the same number of spikes are equally likely, and thus the model is fully specified by the distribution over spike counts $K = \sum_i X_i$.

Higher-order redundancies depend on pairwise correlations.—We want to determine the effect of common input on the entropy of the population. MaxEnt models of different orders define a hierarchy of models which are decreasing in entropy. We define S_{DG} to be the entropy of the full DG model, S_q of the corresponding MaxEnt model with interactions of order q (which matches correlations up to order q of the full model) as well as $\Delta_2 = S_1 - S_2$ and

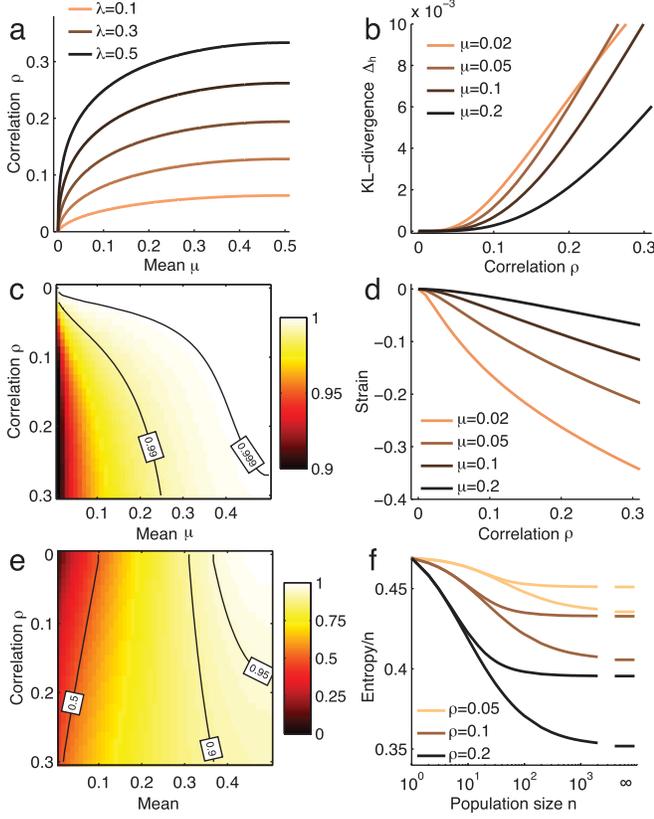


FIG. 1 (color online). Correlations in the DG. (a) The correlation coefficient ρ increases with firing probability μ for constant input correlation λ . (b) The KL divergence Δ_h between the DG and its second-order approximation is modulated by the mean firing rate μ and the correlation ρ in a population of size $n = 5$. (c) In small populations ($n = 5$), the multi-information explained (I_2) by a DG is close to 1. (d) The strain of the homogeneous DG is negative and correlation dependent. (e) For large populations, I_2 between the models is close to 0 for small correlations, i.e., the models are very dissimilar. (f) Scaling of the entropy rate (i.e., entropy per neuron) of the Ising and DG model for mean $\mu = 0.1$, and comparison with asymptotic values (labeled ∞). The entropy rate drops initially before settling to the asymptotic value. For weak correlations, differences between models only become apparent for large n .

$\Delta_h = S_2 - S_{\text{DG}}$ to be the reduction in entropy attributed to second- and higher-order correlations. We refer to Δ_h as the higher-order redundancy. It is equivalent to the Kullback-Leibler (KL) divergence, i.e., the expected log-likelihood ratio per sample between a model and its second-order approximation [17], a popular measure of the magnitude of higher-order redundancies in neural recordings [2,4]. A large Δ_h implies that a pairwise Ising model is insufficient for capturing the population statistics.

Figure 1(b) shows Δ_h for a population model of size $n = 5$ neurons. Notably, small changes in firing probabilities and pairwise correlations can result in large changes in Δ_h . For example, a change of correlation coefficient from 0.05 to 0.1 for $\mu = 0.1$ leads to an increase of Δ_h by a factor of 10.3 [from 6.6 to 68×10^{-5} , corresponding to

0.40 and 4.08 bits per minute at 10 ms bins, which is well inside the range reported by [4] Fig. 2d]. This constitutes a possible quantitative explanation for the phenomenon that higher-order redundancies are much more pronounced amongst nearby cortical neurons [4], for which pairwise correlations are also higher [14]. In addition, our results are also consistent with the finding that the mutual information explained $I_2 = \Delta_2 / (\Delta_2 + \Delta_h)$ is large in retinal recordings with weak correlations driven by visual stimuli. For example, for $\mu = 0.02$, $\rho = 0.04$, $n = 10$, we get $I_2 = 0.92$ [c.f. the values of 0.9 and above in [2] Fig. 2d, moments were fit to their Fig. 1e]. Thus, the DG model makes quantitative predictions for how deviations from pairwise models depend on pairwise correlations.

We also find that the triplet interaction (“strain”) [14] of the DG model, is negative [-0.026 for $\mu = \rho = 0.1$, c.f. [14], Fig. 2d], and decreases monotonically with increasing correlation coefficients [14]. This is quantitatively consistent with experimental observations and surprising, as it has been suggested that common input would lead to a higher occurrence of spike triplets, and thus to a positive strain [14]. Simulations with heterogeneous correlations in the DG show that its strain is negative for weak correlations of the same sign. Thus, these properties of our common input model are consistent with those observed in small neural populations.

For large populations, Δ_2 and Δ_h scale linearly with population size with the correlation-dependent rate.—We are interested in how the entropies of the two models grow with population size. The observation of sublinear entropy growth in small neural populations [2] has raised the question of how entropy scales in large populations. While entropy is extensive for many physical systems [9], it is possible to construct systems with bounded pairwise correlations but nonextensive entropy [18].

For the DG, the asymptotic probability density of the normalized counts $r = k/n$, $r \in (0, 1)$ is given by [19]

$$f_{\text{DG}}(r) = \frac{1}{Z_{\text{DG}}} \exp^{-\frac{1}{2} \frac{\left(\Phi^{-1}(r) - \frac{\gamma\sqrt{1-\lambda}}{(1-2\lambda)}\right)^2}{\lambda/(1-2\lambda)}}, \quad (1)$$

where $\Phi(r)$ denotes the Gaussian cumulative density. We can calculate the asymptotic entropy rate of the DG, $s_{\text{DG}} = \lim_{n \rightarrow \infty} S_{\text{DG}}/n$ by decomposing it into the entropy of the spike count and the entropy conditional on the spike count, $S(X) = S(X|K) + S(K)$. We note that $S(K)$ is bounded above by $\log_2 n$, and that $S(X|K = k) = \log_2 \binom{n}{k}$. Using the identity $\log_2 \binom{n}{nr} / n \approx -[r \log_2(r) + (1-r) \log_2(1-r)] =: \eta_2(r)$, we can see that entropy in this model with all-to-all correlations grows linearly with population size with rate $s_{\text{DG}} = \int_0^1 f_{\text{DG}}(r) \eta_2(r) dr$. By solving a constrained linear optimization problem, we find the distribution $f_{\text{isi}}(r)$ which maximizes the entropy for large n , and find it to be a sum of two delta peaks, $f_{\text{isi}}(r) = p_1 \delta(r - r_1) + p_2 \delta(r - r_2)$ with locations $r_{1,2} = 1/2 \pm \sqrt{1/4 - \mu + \mu^2 + \sigma}$ [18]. Hence, the

asymptotic entropy per neuron of the Ising model is $s_{isi} = \eta_2(r_1)$.

Thus, while the entropies of both models grow sublinearly initially [2] [see Fig. 1(f)], they scale linearly with population size for large n [9]. As a consequence, Δ_h also scales linearly with population size. The entropy rate of the DG for $\mu = 0.1$ and $\rho = 0.05$ is 0.35, and the rate of $\Delta_h = 0.016$. For large populations [depending on the correlations, for n between 10 and 100, Fig. 1(f)], the two models become strongly dissimilar; i.e., I_2 of the DG is asymptotically 0.57 for $\mu = \rho = 0.1$ and 0.24 for $\mu = 0.02, \rho = 0.04$. Figure 1(e) also shows that the close similarity (as measured by I_2) between the Ising model and the DG conjectured by [12] asymptotically holds for firing probabilities near 0.5, but not necessarily otherwise. Our results readily generalize to populations consisting of a finite number of homogeneous pools, in which case the asymptotic entropy is dominated by within-pool correlations. Furthermore, our results could be used to derive lower bounds on the entropy of general MaxEnt models.

The correlations of the DG increase sparsity.—In addition to the entropy, correlations also affect other population statistics. In particular, we are interested in their effect on the sparsity of the population. Sparsity is considered to be an important feature of population coding, and has been shown to be increased by correlations [4]. We quantify sparsity as the probability of the population being quiet [4], i.e. $P(K=0)$, and concentrate on the effects of correlations on sparsity in large populations.

The mode of the asymptotic spike-count distribution $f_{DG}(r)$ [Eq. (1)] is at 0, i.e., quiescence is the most likely population state whenever the input correlation λ exceeds the value $\lambda = 0.5$ [Fig. 2(a)], which is a critical point for $f_{DG}(r)$. Interestingly, this is independent of the parameter γ controlling the mean firing rate (as long as $\gamma < 0$). For small spike probabilities μ , even small correlations ρ correspond to a supercritical λ [Fig. 1(a)].

For the corresponding MaxEnt distribution, the binary infinite-range Ising model with $P(K=k) = Z^{-1} \binom{n}{k} \times \exp(h_n k + J_n k^2)$, we need to identify the parameter scaling yielding the desired means and correlations. It should be noted that this limit is subtly, but critically different from the usual thermodynamic one [1,10,13]: Setting $J_n = J/n$ and $h_n = h$ as assumed in [13] is inconsistent with the moment constraints, as it yields the large- n distribution $\exp\{n[\eta_e(r) + hr + Jr^2]\}$, which collapses to a single delta-peak and has vanishing correlations. However, ensuring $(h + J) = \alpha/n$ with $\alpha = (\log p_2 - \log p_1)/(r_2 - r_1)$ yields constant correlations and the large- n distribution

$$f_{isi}(r) = Z_{isi}^{-1} \exp\{\alpha r + n[\eta_e(r) + J(r^2 - r)]\} \quad (2)$$

with $J = [\log(r_2) - \log(r_1)]/(r_2 - r_1)$ [18].

From Fig. 2, we can see that the two models have very different spike counts, and that the DG has increasing sparsity for supercritical correlation $\rho = 0.25$. The count

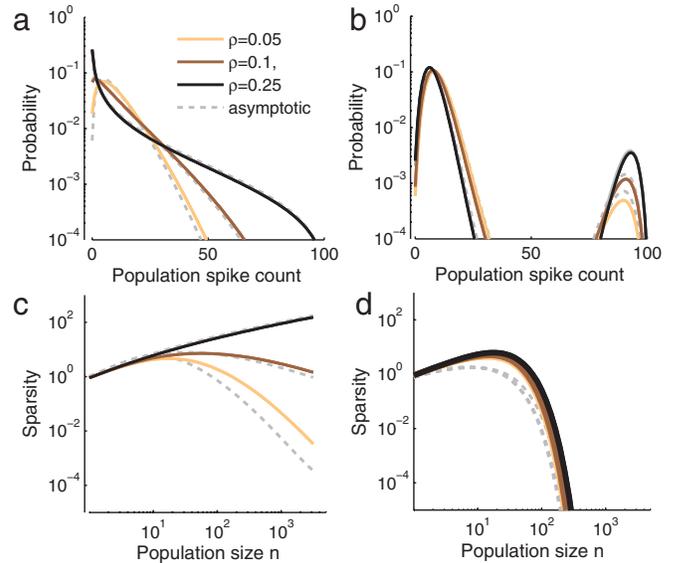


FIG. 2 (color online). Population spike-count distributions and sparsity: (a),(b) The spike-count distributions for the DG (a) and Ising model (b) for population size $n = 100$ and $\mu = 0.1$ are substantially different. (Large- n approximation, dashed grey lines). Note that $\rho = 0.25$ is the critical correlation, and that the Ising model is bimodal, despite the absence of higher-order interactions. (c),(d) Population sparsities for the DG (c) and the corresponding Ising model (d), (parameters as above, sparsity normalized by population size). The DG is much sparser than the Ising model for large population sizes.

distribution of the Ising model with matching pairwise statistics is bimodal (corresponding to a ferromagnetic phase), behaves very much like a mixture of two independent distributions, and has vanishing sparsity. In fact, any model with interactions of finite order will converge to a mixture of independent distributions [13]. Thus, interactions of all orders are necessary for achieving a continuous asymptotic spike-count distribution, and the same sparsity scaling as the DG. While these results were derived assuming a completely homogeneous population, we conjecture that sparsity in large, heterogeneous populations is still dominated by high-order interactions.

Higher-order interactions increase heat capacity.—Finally, we investigate the impact of common input on the heat capacity of the population. A sharply peaked and diverging specific heat (i.e., heat capacity normalized by population size) is evidence for a physical system being at a critical point [1]. This behavior has been observed for neural populations as well as natural images [10,15], and it has been argued that it has desirable consequences for sensory coding. In addition, a diverging specific heat implies that fluctuations around the mean log-probability—the entropy—are nonvanishing in large populations. We here show that our common input model is able to capture this behavior: The distribution of any model $P(x)$ at some temperature $T = 1/\beta$ is given by $P_\beta(x) = P(x)^\beta / Z_\beta$, and the specific heat by $c = \text{Var}[\log_2 P_\beta(x)]/n$.

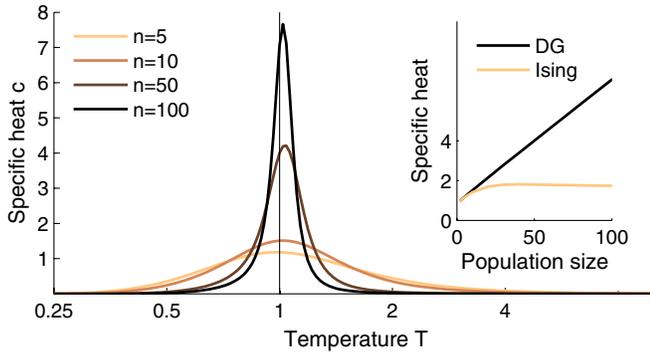


FIG. 3 (color online). Scaling of specific heat: Specific heat of the DG with mean $\mu = 0.1$ and correlation $\rho = 0.1$ diverges at $T = 1$. Inset: The specific heat of the DG at $T = 1$ grows linearly with population size, whereas it saturates for the Ising model.

For large n , the spike-count distribution is $P_\beta(K) = \exp[n(1 - \beta)\eta_e(k/n)]P(K)^\beta/Z_\beta$, and asymptotically this yields $c_\beta = n \int f_\beta(r)[\eta_2(r)^2 - s_\beta^2]dr$, where f_β is the limiting distribution of $P_\beta(K)$.

Therefore, c_β diverges linearly whenever this integral is nonzero, which is the case for the DG and many other models at $\beta = 1$. For $\beta \neq 1$, however, $f_\beta(r)$ is dominated by the exponential, collapses to a delta-peak, and has finite specific heat. Thus, the DG and any other infinite-range model with interactions beyond second order has a critical point at $T = 1$ (Fig. 3). The pairwise Ising model is a notable exception, as it does not have a phase transition, and its specific heat is finite (Fig. 3 inset).

It is therefore informative to calculate the specific heat at unit temperature as a function of the moments μ and ρ . In this case, the specific heat of the Ising model is

$$c_{isi} = \frac{r_1 r_2 J^2 (\sigma + \mu^2 - \mu + 1/4)}{4(1 - 2Jr_1 r_2)} \log_2^2(e). \quad (3)$$

Asymptotically, the heat capacity of the Ising model is maximized for vanishing correlation, whereas the DG attains its maximum at strong correlations, e.g., $\rho = 0.37$ for $\mu = 0.1$. We conclude that higher-order interactions can have a strong impact on the specific heat: They lead to a qualitatively different scaling behavior and criticality, and influence the moments which maximize it.

Conclusions.—We showed that a simple model with common inputs and spiking nonlinearities could qualitatively account for some of the complex structure of higher-order correlations in neural population recordings. The DG captures higher-order redundancies which increase with second-order correlations, has a negative strain, increased sparsity, and a divergent specific heat. Given the ubiquity of common input in sensory systems, the model thus gives quantitative predictions for the conditions under which deviations from second-order models will be strong, and

how these quantities should scale with population size. It is worth remarking that our formulations can readily be generalized to general input distributions or spike generation mechanisms. Further investigations will have to show whether our results would also quantitatively account for observations in heterogeneous populations. Finally, the role of higher-order correlations for stimulus coding has only been investigated experimentally in small populations [4,8], and their impact in large populations remains unclear. As, in our model, a change in the mean input modifies interactions of all orders, it can also facilitate our understanding of the role of higher-order interactions in stimulus coding.

M.B. and J.H.M. were supported by BMBF; FKZ: 01GQ0601, J.H.M. by a European Commission Marie Curie Fellowship and the Gatsby Charitable Foundation. We thank S. Gerwinn, E. Mukamel and P. Latham for discussions.

*jakob@gatsby.ucl.ac.uk

- [1] G. Parisi, *Statistical Field Theory* (Perseus Books, Reading, MA, 1998).
- [2] E. Schneidman, M.J. n. Berry, R. Segev, and W. Bialek, *Nature (London)* **440**, 1007 (2006).
- [3] J. Shlens *et al.*, *J. Neurosci.* **26**, 8254 (2006).
- [4] I.E. Ohiorhenuan, F. Mechler, K.P. Purpura, A.M. Schmid, Q. Hu, and J.D. Victor, *Nature (London)* **466**, 617 (2010).
- [5] E. Jaynes, *Phys. Rev.* **106**, 620 (1957).
- [6] S. Watanabe, *IBM J. Res. Dev.* **4**, 6682 (1960).
- [7] W. Ehm, B. Staude, and S. Rotter, *Electronic Journal of Statistics* **1**, 473 (2007).
- [8] F. Montani *et al.*, *Phil. Trans. R. Soc. A* **367**, 3297 (2009).
- [9] Y. Roudi, S. Nirenberg, and P.E. Latham, *PLoS Comput. Biol.* **5**, e1000380 (2009).
- [10] G. Tkacik, E. Schneidman, M.J. Berry, II, and W. Bialek, *arXiv:q-bio/0611072v2*.
- [11] D.R. Cox and N. Wermuth, *Biometrika* **89**, 462 (2002).
- [12] J.H. Macke, P. Berens, A.S. Ecker, A.S. Tolias, and M. Bethge, *Neural Comput.* **21**, 397 (2009).
- [13] S.-I. Amari, H. Nakahara, S. Wu, and Y. Sakai, *Neural Comput.* **15**, 127 (2003).
- [14] I. Ohiorhenuan and J. Victor, *J. Comput. Neurosci.* **30**, 125 (2010).
- [15] G.J. Stephens, T. Mora, G. Tkacik, and W. Bialek, *arXiv:0806.2694*.
- [16] D.S. Greenberg, A.R. Houweling, and J.N.D. Kerr, *Nat. Neurosci.* **11**, 749 (2008).
- [17] P. Grünwald and A. Dawid, *Ann. Stat.* **32**, 1367 (2004).
- [18] See supplemental material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.106.208102> for further details, derivations, and explanations.
- [19] f_{DG} can be derived using the method of steepest descent as in [13] or by finding the likelihood of an input which has probability r of inducing a spike.