Random versus maximum entropy models of neural population activity

Ulisse Ferrari,^{1,*} Tomoyuki Obuchi,^{2,*} and Thierry Mora^{3,†}

¹*Institut de la Vision, INSERM and UPMC, 75012 Paris, France*

²*Department of Mathematical and Computing Science,*

Tokyo Institute of Technology, Yokohama 226-8502, Japan

³*Laboratoire de physique statistique, École normale supérieure, CNRS and UPMC, 75005 Paris, France*

(Dated: December 8, 2016)

The principle of maximum entropy provides a useful method for inferring statistical mechanics models from observations in correlated systems, and is widely used in a variety of fields where accurate data are available. While the assumptions underlying maximum entropy are intuitive and appealing, its adequacy for describing complex empirical data has been little studied in comparison to alternative approaches. Here data from the collective spiking activity of retinal neurons is reanalysed. The accuracy of the maximum entropy distribution constrained by mean firing rates and pairwise correlations is compared to a random ensemble of distributions constrained by the same observables. In general, maximum entropy approximates the true distribution better than the typical or mean distribution from that ensemble. This advantage improves with population size, with groups as small as 8 being almost always better described by maximum entropy. Failure of maximum entropy to outperform random models is found to be associated with strong correlations in the population.

The principle of maximum entropy was introduced in 1957 by Jaynes [1, 2] to formulate the foundations of statistical mechanics as an inference problem. Its interest has been recently rekindled by its application to a variety of data-rich fields, starting with the correlated activity of populations of retinal neurons [3, 4]. The method has since been used to study correlations in other neural data, such as cortical networks [5, 6] and functional magnetic resonance imaging [7], as well as in other biological and non-biological contexts, including multiple sequence alignments of proteins [8–10] and nucleic acids [11, 12], the collective motion of bird flocks [13], the spelling rules of words [14], and the statistics of decisions by the United States supreme court [15]. In many cases, the close link between maximum entropy and statistical mechanics has led to new insights into the thermodynamics of the system in terms of phase transitions [16–18], or multi-valley energy landscape [19, 20]. In other cases, the method has allowed for predictions of crucial practical relevance, such as residue contacts in proteins [21], or deleterious mutations in HIV [22].

Although the motivations of maximum entropy seem intuitive and can be formalized rigorously [23], the perceived arbitrariness of its assumptions has led to question its validity [24, 25]. The starting point is to consider models that match empirical observations on a few key statistics of the data. Maximum entropy’s crucial—and arguably debatable—assumption is to pick, out of the many models that satisfy that constraint, the one with the largest Gibbs entropy. This choice seems natural, since it ensures that the model is as random as possible. However, it is not clear why it should describe the data better than other models satisfying the same constraints. To address this question directly on empirical data, we reanalyse the original neural data from [3], which contributed to the recent surge of interest in maximum en-

FIG. 1: Random models. The space $\mathcal{C}$ of models P is a simplex of $2^N - M - 1$ dimensions, defined by the intersection of the hyperplanes satisfying the constraints that the mean observables under the model, $\sum_{\sigma} P(\sigma) \mathcal{O}_a(\sigma)$, $a = 1, \dots, M$, equal the empirical means, $\bar{\mathcal{O}}_a$, and by a normalization and positivity constraint. The true distribution to be approximated, $\hat{P}$ (red dot), is not accessible in general. An entropy-dependent measure μ_{Γ} (Eq. 1) is defined on $\mathcal{C}$ (red map). At $\Gamma = 0$ (random ensemble), the measure is uniform over that space. As Γ is increased, the measure concentrates onto the maximum entropy distribution P_{ME} (blue dot), and so does its mean $P_{\Gamma} = \langle P \rangle_{\Gamma}$ (black dot).

trophy. We compare the accuracy of maximum entropy distributions to ensembles of distributions that satisfy the same constraints, using the approach developed in [26, 27].

The collective state of a population of N variables is described by $\sigma = (\sigma_1, \dots, \sigma_N)$. In general, σ_i may denote any degree of freedom, such as the identity of an amino-acid in a protein, the orientation of a bird in a flock, etc. To fix idea, in this paper σ_i will be a binary variable describing the spiking activity of neuron i : $\sigma_i = 1$ if neuron i spikes within a given time window, and 0 otherwise. The joint distribution of the collective activity σ , denoted by $P(\sigma)$, lives in a $2^N - 1$ dimensional space, represented schematically in Fig. 1. Because that space is huge for even moderately large populations, it is often impossible to sample the true distribution, $\hat{P}$, reliably

from the data. Simplifying assumptions are needed.

To restrict the search of models, one can focus on distributions that agree with the data on the average value of a few observables. Calling these observables $\mathcal{O}_a(\sigma)$, $a = 1, \dots, M$, the condition reads $P \cdot \mathcal{O}_a \equiv \sum_{\sigma} P(\sigma) \mathcal{O}_a(\sigma) = \bar{\mathcal{O}}_a$, where $\bar{\mathcal{O}}_a$ is the empirical mean. The observables must be chosen carefully depending on the problem at hand, and may include local or global order parameters, marginal probabilities, correlation functions, etc. Let us denote by $\mathcal{C}$ the subspace of models P that satisfy those constraints, as well as the conditions $P(\sigma) \geq 0$ and $\sum_{\sigma} P(\sigma) = 1$. $\mathcal{C}$ is convex because of the linear nature of the constraints.

A probability law on $\mathcal{C}$ may be defined which weighs models $P \in \mathcal{C}$ according to their Gibbs entropy, $S(P) = -\sum_{\sigma} P(\sigma) \log P(\sigma)$, through the following measure [26]:

$$\mu_{\Gamma}(P) = \frac{e^{\Gamma S(P)}}{\mathcal{Z}}, \quad \mathcal{Z} = \int_{P \geq 0} \mathcal{D}P e^{\Gamma S(P)}, \quad (1)$$

with

$$\mathcal{D}P = \delta(\sum_{\sigma} P(\sigma) - 1) \prod_{a=1}^M \delta(P \cdot \mathcal{O}_a - \bar{\mathcal{O}}_a) \prod_{\sigma} dP(\sigma), \quad (2)$$

where $\delta(\cdot)$ is Dirac's delta function. The parameter Γ is conjugate to the entropy, and sets its average value: $\langle S(P) \rangle_{\Gamma} = \partial \ln \mathcal{Z} / \partial \Gamma$, where we use the brackets $\langle \cdot \rangle_{\Gamma}$ for averages over the measure μ_{Γ} . Γ plays the same role with respect to the entropy as the inverse temperature with respect to the energy in standard statistical mechanics. When $\Gamma = 0$, all distributions satisfying the constraints have the same probability. We call this the unbiased ensemble. As $\Gamma \rightarrow \infty$, the measure becomes increasingly peaked onto a single distribution, P_{ME} , of maximum entropy (or, in the previous analogy, the ground state reached at zero temperature). This distribution defines the classical maximum entropy model, and takes the form [28]:

$$P_{\text{ME}}(\sigma) = \frac{1}{Z} \exp \left[\sum_{a=1}^M \lambda_a \mathcal{O}_a(\sigma) \right], \quad (3)$$

where λ_a are Lagrange multipliers enforcing the constraints on the mean observables, and Z is a normalization constant. We define the average distribution as $P_{\Gamma}(\sigma) \equiv \langle P(\sigma) \rangle_{\Gamma}$, which belongs to $\mathcal{C}$ by convexity. P_{Γ} only coincides with P_{ME} in the limit $\Gamma \rightarrow \infty$. At the other extreme, P_0 is the unbiased, center-of mass distribution that satisfies all the constraints.

We follow the approach of the random ensemble defined by (1) to describe the joint spiking activity of retinal ganglion cells reported in [3]. There, the spiking activities of 40 ganglion cells from the salamander retina were recorded by multielectrode arrays for about an hour, and segmented into $\approx 1.5 \cdot 10^5$ binary spike words σ of 20

FIG. 2: Small networks. Illustration of the random ensemble on 2 and 3 neurons. (a) Pairwise correlation $\langle \sigma_1 \sigma_2 \rangle_c = \langle \sigma_1 \sigma_2 \rangle - \langle \sigma_1 \rangle \langle \sigma_2 \rangle$ predicted by maximum entropy and random models constrained by the mean spiking rates of two neurons, $\langle \sigma_1 \rangle = \langle \sigma_2 \rangle$, as a function of that rate. The mean unbiased model P_0 is the center of mass between the lower and upper allowed limits of the correlation, which delimit the shaded area. (b) Triplet connected correlation, $\langle \sigma_1 \sigma_2 \sigma_3 \rangle_c = \langle \sigma_1 \sigma_2 \sigma_3 \rangle - \langle \sigma_1 \sigma_2 \rangle \langle \sigma_3 \rangle - \langle \sigma_1 \sigma_3 \rangle \langle \sigma_2 \rangle - \langle \sigma_2 \sigma_3 \rangle \langle \sigma_1 \rangle + 2 \langle \sigma_1 \rangle \langle \sigma_2 \rangle \langle \sigma_3 \rangle$, as a function of the pairwise correlation between 3 neurons firing with probability $\langle \sigma_1 \rangle = \langle \sigma_2 \rangle = \langle \sigma_3 \rangle = 0.02$ (mean empirical value). Pairwise correlation in the retinal data range from -10^{-3} to 0.03 with a median of $2 \cdot 10^{-4}$. Key is as in (a).

FIG. 3: Random versus maximum entropy models. (a) The normalized Kullback-Leibler divergence relative to maximum entropy, $D_{\text{KL}}(\hat{P} \| P) / D_{\text{KL}}(\hat{P} \| P_{\text{ME}})$, is represented as a function of the entropy-conjugated variable Γ for (a) a random group of $N = 7$ neurons (out of 40). Values above unity (dashed line) mean that maximum entropy outperforms the random model. The violin plots show the distributions over random models drawn from μ_{Γ} , while the red lines show the value for the average model, $D_{\text{KL}}(\hat{P} \| P_{\Gamma})$. (b) Normalized KL divergence at $\Gamma = 0$ for 20 random subsets of 7 neurons (blue), as well as the group of most correlated neurons (as measured by Pearson's correlation coefficient, green), and the set of neurons with the highest spike rate (yellow).

ms. The collective behavior of small networks (up to 10 neurons) was shown to be well described by maximum entropy distributions constrained by spike rates and pairwise correlations (and later to much larger populations [20]). This choice of constraints corresponds to the observables $\mathcal{O}_a = \sigma_i$ for all neuron i , and $\mathcal{O}_a = \sigma_i \sigma_j$ for all pairs i, j , for which the maximum entropy distribution (3) takes the form of a disordered Ising model, $P_{\text{ME}}(\sigma) = (1/Z) \exp(\sum_i h_i \sigma_i + \sum_{ij} J_{ij} \sigma_i \sigma_j)$.

It is instructive first to consider the unbiased measure

μ_0 over very small networks, for which everything can be calculated analytically. The simplest case of two neurons constrained by just their firing rate is illustrated by Fig. 2a. The maximum entropy distribution factorizes over the two neurons, which are thus independent [29]: $P(\sigma) = p_1(\sigma_1)p_2(\sigma_2)$. By contrast, random models drawn from μ_0 are biased towards a positive correlation $\langle\sigma_1\sigma_2\rangle - \langle\sigma_1\rangle\langle\sigma_2\rangle > 0$ when both firing rates $\langle\sigma_1\rangle, \langle\sigma_2\rangle$ are on the same side of 0.5 (in the retinal data $\langle\sigma_i\rangle \sim 0.02$). A similar bias in the triplet correlation is also found when considering 3 neurons constrained by uniform firing rates and pairwise correlations (Fig. 2b). When pairwise correlations are weak, as is the case in the retina [3], random models predict on average a higher 3-point connected correlation than maximum entropy, although the bias is reversed for large correlations.

Thanks to its exponential form (3), the maximum entropy distribution can be inferred with relative ease for systems of size $N \leq 20$, yet requiring to calculate sums of 2^N terms [3]. Sampling from μ_Γ or calculating P_Γ , on the other hand, is a much harder task, involving the exploration of $\mathcal{C}$ of dimension $2^N - N(N+1)/2 - 1$. To apply the random ensemble to populations of neurons, we sampled from μ_Γ using the Metropolis-Hastings algorithm, for various subgroups of neurons of different sizes. At each step, starting from a distribution P in $\mathcal{C}$, one picks a random direction V in the Fourier basis of the hyperplane orthogonal to all observables $\mathcal{O}_a$ [27]. The new distribution is taken to be $P' = P + \alpha V$, where α is drawn uniformly in the interval $(\alpha_{\min}, \alpha_{\max})$ defined by the lower and upper limits so that $P'(\sigma) \geq 0$ for all σ . P' is accepted with probability $\min(1, e^{\Gamma[S(P') - S(P)]})$. The process is repeated until equilibration is reached. High space dimension limits us to relatively small group sizes, $N \leq 8$. Fortunately for these sizes the true distribution $\hat{P}$ may be accurately estimated from the data, and directly compared to models.

The accuracy of a given model is assessed by the Kullback-Leibler (KL) divergence between the model distribution P and the true one $\hat{P}$, $D_{\text{KL}}(\hat{P}||P) = \sum_{\sigma} \hat{P}(\sigma) \ln[\hat{P}(\sigma)/P(\sigma)]$. Fig. 3 shows, in the form of violin plots, the distribution of KL divergence (normalized relative to maximum entropy) when sampling P from μ_Γ , for groups of $N = 7$ cells. This distribution is plotted in Fig. 3a for a random groups of 7 cells. Maximum entropy is found to have a clear advantage: its accuracy is matched by only a negligible fraction of models drawn from μ_Γ , and it also does better than their mean P_Γ (red line). The advantage of maximum entropy over the unbiased ensemble generalizes to 20 random groups of 7 cells (Fig. 3b), as well as the groups comprising the most correlated (green) and most active (yellow) cells. Interestingly, in all cases the mean distribution P_Γ is more accurate than the typical distribution P sampled from μ_Γ , a consequence of Jensen's inequality which implies $D_{\text{KL}}(\hat{P}||\langle P \rangle_\Gamma) \leq \langle D_{\text{KL}}(\hat{P}||P) \rangle_\Gamma$. In general, $0 < \Gamma < \infty$

FIG. 4: Dependence on populations size. (a) The normalized divergence of the average model, $D_{\text{KL}}(\hat{P}||P_\Gamma)$, is averaged over 20 random subsets, and plotted as a function of $(1/N) \ln \Gamma$. Errors bars show standard error on the mean. (b) Fraction of random groups (out of hundreds) of N neurons that are better described by the mean unbiased distribution P_0 than by the maximum entropy model P_{ME} .

interpolates between the unbiased ensemble and the maximum entropy distribution. For these reasons, in the following the maximum entropy model P_{ME} will only be compared to the mean distribution of the unbiased ensemble, P_0 .

We now investigate the dependence on the population size. Fig. 4 shows the average normalized KL divergence of the mean model P_Γ for random cell groups of varying sizes, as a function of $(1/N) \ln \Gamma$ (the scaling of Γ is assumed to be exponential in N , as suggested by calculations with random observables [26]). The general trend noted before for $N = 7$ generalizes to all sizes: the larger the entropy bias Γ , the better the model (Fig. 4a). However, this average behaviour masks large heterogeneities across different choices of cell groups, especially for small groups, of which a sizeable fraction is better described by the mean distribution P_0 than by P_{ME} . Evaluating this fraction from hundreds of random groups for each N , we find that maximum entropy is more likely to outperform the random ensemble in larger groups (Fig. 4b), and even does so in all of the 200 tested groups of size $N = 8$.

What sets apart groups of cells that are better described by P_0 than by P_{ME} ? Since both share the same 1- and 2-point correlations by construction, we examine their predictions for 3-point correlations in triplets of cells ($N = 3$). Fig. 5a shows that random models typically fail because they overestimate small 3-point correlations. By contrast, maximum entropy is more likely to be outperformed by random models when the triplet correlation is large, in which case maximum entropy overestimates it. Both these findings are in agreement with the results of Fig. 2b. This observation can be generalized to larger groups of neurons ($N > 3$) by considering the total amount of correlations in the network, quantified by the loss of entropy due to correlations, or multi-information [29], $I = S(P_{\text{ind}}) - S(\hat{P})$, where $P_{\text{ind}} = \prod_i p_i(\sigma_i)$ is the model distribution of independent neurons. Groups that are better described by P_0 than by P_{ME} are found to have

FIG. 5: Correlations and maximum likelihood performance. Three-point connected correlation $\langle \sigma_1 \sigma_2 \sigma_3 \rangle_c$ for (a) 100 random triplets whose joint activity is best described by maximum entropy and (b) 100 random triplets whose joint activity is best described by the mean unbiased model, when constraining the values of the pairwise correlations. The error bar shows, for each triplet, the allowed range of values for the 3-point correlation. (c) The multi-information, which measures the overall amount of correlation in the collective activity, is plotted as a function of system size, for groups of neurons that are best described by the maximum entropy model P_{ME} (blue) or by the mean unbiased model P_0 (red). Error bars show standard deviation across groups of cells (the red point at $N = 7$ has no error bar because only one group of that size was better described by P_0)

a higher multi-information on average (Fig. 5c).

Since maximum entropy was proposed as a method for building statistical models from high-dimensional data, its accuracy, relevance, and epistemological validity have been questioned. In this study we have shown that the maximum entropy model describes the spiking activity of populations in the retina better than the mean model satisfying the same constraints, which itself performs better than the vast majority of random models under these constraints. This better performance of maximum entropy gets more marked as the population size N grows, and is essentially always true for $N \geq 8$. The analysis of 3-point and higher-order correlations suggests that the rare instances where the mean model outperforms maximum entropy is when correlations are relatively large. In that case, maximum entropy predicts high triplet correlations within the allowed range compared to the mean unbiased model (Fig. 2b), and may thus overestimate their true value, consistent with previous observations in large populations [20]. In that case, models that take a “middle-of-the-road” value of the correlations may be preferred to maximum entropy.

By providing a test on empirical data, our results complement previous work aimed at explaining or refuting the efficiency of maximum entropy based on theoretical arguments and simulated datasets. Calculations on synthetic learning problems have suggested that maximum entropy is no more accurate than random [26], unless the chosen observables are smooth as a function of configuration space [27]. However, in these studies the choice of observables and true distributions were taken to be completely random, and it is not clear how applicable they are to real distributions and to pairwise constraints.

Other simulation studies have more specifically addressed the role of pairwise interactions. It was suggested that pairwise maximum entropy models should fail for large populations [30]. On the other hand, strongly interacting systems with interactions of arbitrary order have been numerically shown to be well described by pairwise interactions, with an analogy to Hopfield networks [31]. The principle of maximum entropy has also been advocated by contrast to non-additive (or Rényi) entropies, but on purely theoretical grounds [32]. Our results do not preclude that other objective functions than entropy may help better describe empirical data. They suggest, however, that it is better to pick the most random model than to pick a model at random.

We thank Michael Berry for sharing the data from [3], and Olivier Marre for his comments on the manuscript.

* Equal contribution

† Corresponding author: tmora@lps.ens.fr

- [1] E. T. Jaynes, Phys. Rev. **106**, 620 (1957).
- [2] E. T. Jaynes, Phys. Rev. **108**, 171 (1957).
- [3] E. Schneidman, M. J. Berry, R. Segev, and W. Bialek, Nature **440**, 1007 (2006).
- [4] J. Shlens, G. D. Field, J. L. Gauthier, M. I. Grivich, D. Petrusca, A. Sher, A. M. Litke, and E. J. Chichilnisky, J. Neurosci. **26**, 8254 (2006).
- [5] A. Tang, D. Jackson, J. Hobbs, W. Chen, J. L. Smith, H. Patel, A. Prieto, D. Petrusca, M. I. Grivich, A. Sher, et al., J. Neurosci. **28**, 505 (2008).
- [6] G. Tavoni, U. Ferrari, F. P. Battaglia, and S. Cocco (2015).
- [7] T. Watanabe, S. Hirose, H. Wada, Y. Imai, T. Machida, I. Shirouzu, S. Konishi, Y. Miyashita, and N. Masuda,

- Nat. Commun. **4**, 1370 (2013).
- [8] M. Weigt, R. a. White, H. Szurmant, J. a. Hoch, and T. Hwa, Proc. Natl. Acad. Sci. U. S. A. **106**, 67 (2009).
 - [9] T. Mora, A. M. Walczak, W. Bialek, and C. G. Callan, Proc. Natl. Acad. Sci. **107**, 5405 (2010), 0912.5175.
 - [10] M. Figliuzzi, H. Jacquier, A. Schug, O. Tenaillon, and M. Weigt, Mol. Biol. Evol. **33**, 268 (2016), 1510.03224.
 - [11] M. Santolini, T. Mora, and V. Hakim, PLoS One **9**, e99015 (2014).
 - [12] E. De Leonardis, B. Lutz, S. Ratz, S. Cocco, R. Monasson, A. Schug, and M. Weigt, Nucleic Acids Res. **43**, 10444 (2015).
 - [13] W. Bialek, A. Cavagna, I. Giardina, T. Mora, E. Silvestri, M. Viale, and A. M. Walczak, Proc. Natl. Acad. Sci. U. S. A. **109**, 4786 (2012), 1107.0604.
 - [14] G. J. Stephens and W. Bialek, Phys. Rev. E - Stat. Non-linear, Soft Matter Phys. **81**, 3 (2010), 0801.0253.
 - [15] E. D. Lee, C. P. Broedersz, and W. Bialek, J. Stat. Phys. **160**, 275 (2015), 1306.5004.
 - [16] T. Mora and W. Bialek, J. Stat. Phys. **144**, 268 (2011).
 - [17] W. Bialek, A. Cavagna, I. Giardina, T. Mora, O. Pohl, E. Silvestri, M. Viale, and A. M. Walczak, Proc. Natl. Acad. Sci. U. S. A. **111**, 7212 (2014).
 - [18] G. Tkačik, T. Mora, O. Marre, D. Amodei, S. E. Palmer, M. J. Berry, and W. Bialek, Proc. Natl. Acad. Sci. **112**, 11508 (2015).
 - [19] T. Watanabe, N. Masuda, F. Megumi, R. Kanai, and G. Rees, Nat. Commun. **5**, 1 (1), arXiv:1011.1669v3.
 - [20] G. Tkačik, O. Marre, D. Amodei, E. Schneidman, W. Bialek, and M. J. Berry, PLoS Comput. Biol. **10**, e1003408 (2014), /arxiv.org/abs/1306.3061.
 - [21] F. Morcos, A. Pagnani, B. Lunt, A. Bertolino, D. S. Marks, C. Sander, R. Zecchina, J. N. Onuchic, T. Hwa, and M. Weigt, Proc. Natl. Acad. Sci. **108**, E1293 (2011), 1110.5223.
 - [22] A. L. Ferguson, J. K. Mann, S. Omarjee, T. Ndung'u, B. D. Walker, and A. K. Chakraborty, Immunity **38**, 606 (2013).
 - [23] J. E. Shore and R. W. Johnson, IEEE Trans. Inf. Theory **26**, 26 (1980).
 - [24] E. Aurell, PLoS Comput. Biol. **12**, 1 (2016).
 - [25] E. van Nimwegen, PLoS Comput. Biol. **12**, 1 (2016).
 - [26] T. Obuchi, S. Cocco, and R. Monasson, J. Stat. Phys. **161**, 598 (2015), arXiv:1503.02802v1.
 - [27] T. Obuchi and R. Monasson, J. Phys. Conf. Ser. **638**, 012018 (2015).
 - [28] S. Presse, K. Ghosh, J. Lee, and K. A. Dill, Rev. Mod. Phys. **85**, 1115 (2013).
 - [29] E. Schneidman, S. Still, M. J. Berry, and W. Bialek, Phys. Rev. Lett. **91**, 238701 (2003).
 - [30] Y. Roudi, S. Nirenberg, and P. E. Latham, PLoS Comput. Biol. **5**, e1000380 (2009).
 - [31] L. Merchan and I. Nemenman, J. Stat. Phys. **162**, 1294 (2016), 1505.02831.
 - [32] S. Pressé, K. Ghosh, J. Lee, and K. A. Dill, Phys. Rev. Lett. **111**, 180604 (2013), arXiv:1312.1186v1.