

On Maximum Entropy and Inference

Luigi Gresele ^{1,2} and Matteo Marsili ^{3,*}

¹ Empirical Inference Department, Max Planck Institute for Intelligent Systems, Max-Planck-Ring 4, 72076 Tübingen, Germany; luigi.gresele@tuebingen.mpg.de

² High-field Magnetic Resonance Department, Max Planck Institute for Biological Cybernetics, Max-Planck-Ring 11, 72076 Tübingen, Germany

³ Quantitative Life Sciences Section, The Abdus Salam International Center for Theoretical Physics, Strada Costiera 11, 34151 Trieste, Italy

* Correspondence: marsili@ictp.it; Tel.: +39-040-224-0461

Received: 9 October 2017; Accepted: 20 November 2017; Published: 28 November 2017

Abstract: Maximum entropy is a powerful concept that entails a sharp separation between relevant and irrelevant variables. It is typically invoked in inference, once an assumption is made on what the relevant variables are, in order to estimate a model from data, that affords predictions on all other (dependent) variables. Conversely, maximum entropy can be invoked to retrieve the relevant variables (sufficient statistics) directly from the data, once a model is identified by Bayesian model selection. We explore this approach in the case of spin models with interactions of arbitrary order, and we discuss how relevant interactions can be inferred. In this perspective, the dimensionality of the inference problem is not set by the number of parameters in the model, but by the frequency distribution of the data. We illustrate the method showing its ability to recover the correct model in a few prototype cases and discuss its application on a real dataset.

Keywords: maximum entropy; model selection; spin models; singular value decomposition; high order interactions

1. Introduction

Statistical mechanics stems from classical (or quantum) mechanics. The latter prescribes which are the relevant quantities (i.e., the conserved ones). The former brings this further, and it predicts that the probability to observe a system in a microscopic state s , in thermal equilibrium, is given by:

$$P(s) = \frac{1}{Z} e^{-\beta H[s]} \quad (1)$$

where $H[s]$ is the energy of configuration s . The inverse temperature β is the only relevant parameter that needs to be adjusted, so that the ensemble average $\langle H \rangle$ matches the observed energy U . It has been argued [1] that the recipe that leads from $H[s]$ to the distribution $P(s)$ is maximum entropy: among all distributions that satisfy $\langle H \rangle = U$, the one maximizing the entropy $S = -\sum_s P(s) \log P(s)$ should be chosen. Information theory clarifies that the distribution Equation (1) is the one that assumes nothing else but $\langle H \rangle = U$, or equivalently, that all other observables can be predicted from the knowledge of $H[s]$.

This idea carries through more generally to inference problems: given a dataset of N observations $\hat{s} = \{s^{(1)}, \dots, s^{(N)}\}$ of a system, one may invoke maximum entropy to infer the underlying distribution $P(s)$ that reproduces the empirical averages of a set \mathcal{M} of observables $\phi^\mu(s)$ ($\mu \in \mathcal{M}$). This leads to Equation (1) with:

$$-\beta H[s] = \sum_{\mu \in \mathcal{M}} g^\mu \phi^\mu(s) \quad (2)$$

where the parameters g^μ should be fixed by solving the convex optimization problems:

$$\langle \phi^\mu \rangle = \overline{\phi^\mu} \equiv \frac{1}{N} \sum_{i=1}^N \phi^\mu(s^{(i)}) \quad (3)$$

that result from entropy maximization and are also known to coincide with maximum likelihood estimation (see [2–4]).

For example, in the case of spin variables $s \in \{\pm 1\}^n$, the distribution that reproduces empirical averages $\overline{s_i}$ and correlations $\overline{s_i s_j}$ is the pairwise model:

$$H[s] = - \sum_i h_i s_i - \sum_{i < j} J_{ij} s_i s_j, \quad (4)$$

which in the case $J_{ij} \equiv J \forall i, j$ and $h_i \equiv h \forall i$ is the celebrated Ising model. The literature on inference of Ising models, stemming from the original paper on Boltzmann learning [5] to early applications to neural data [6] has grown considerably (see [7] for a recent review), to the point that some suggested [8] that a purely data-based statistical mechanics is possible.

Research has mostly focused on the estimate of the parameters $g = \{h_i, J_{ij}\}$, which itself is a computationally challenging issue when $n \gg 1$ [7], or in recovering sparse models, distinguishing true interactions ($J_{ij} \neq 0$) from spurious ones ($J_{ij} = 0$; see, e.g., [9]). Little has been done to go beyond pairwise interactions (yet, see [10–13]). This is partly because pairwise interactions offer a convenient graphical representation of statistical dependences; partly because ℓ -th order interactions require $\sim n^\ell$ parameters and the available data hardly ever allow one to go beyond $\ell = 2$ [14]. Yet, strictly speaking, there may be no reason to believe that interactions among variables are only pairwise. The choice of form (4) for the Hamiltonian represents an assumption on the intrinsic laws of motion, which reflects an a priori belief of the observer on the system. Conversely, one would like to have inference schemes that certify that pairwise interactions are really the relevant ones, i.e., those that need to be included in H in order to reproduce correlations $\overline{s_{i_1} s_{i_2} \cdots s_{i_\ell}}$ of arbitrary order ℓ [15].

We contrast a view of inference as parameter estimation of a preassigned (pairwise) model, where maximum entropy serves merely an ancillary purpose, with the one where the ultimate goal of statistical inference is precisely to identify the minimal set \mathcal{M} of sufficient statistics that, for a given dataset \hat{s} , accurately reproduces all empirical averages. In this latter perspective, maximum entropy plays a key role in that it affords a sharp distinction between relevant variables ($\phi^\mu(s)$, $\mu \in \mathcal{M}$) (which are the sufficient statistics) and irrelevant ones, i.e., all other operators that are not a linear combination of the relevant ones, but whose values can be predicted through theirs. To some extent, understanding amounts precisely to distinguishing the relevant variables from the “dependent” ones: those whose values can be predicted.

Bayesian model selection provides a general recipe for identifying the best model \mathcal{M} ; yet, as we shall see, the procedure is computationally unfeasible for spin models with interactions of arbitrary order, even for moderate dimensions ($n = 5$). Our strategy will then be to perform model selection within the class of mixture models, where it is straightforward [16], and then, to project the result on spin models. The most likely models in this setting are those that enforce a symmetry among configurations that occur with the same frequency in the dataset. These symmetries entail a decomposition of the log-likelihood with a flavor that is similar to principal component analysis (yet of a different kind than that described in [17]). This directly predicts the sufficient statistics $\psi_\lambda(s)$ that need to be considered in maximum entropy inference. Interestingly, we find that the number of sufficient statistics depends on the frequency distribution of observations in the data. This implies that the dimensionality of the inference problem is not determined by the number of parameters in the model, but rather by the richness of the data.

The resulting model features interactions of arbitrary order, in general, but is able to recover sparse models in simple cases. An application to real data shows that the proposed approach is able to

spot the prevalence of two-body interactions, while suggesting that some specific higher order terms may also be important.

2. Background

2.1. Spin Models with Interactions of Arbitrary Order

Consider a system of n spin variables $s_i = \pm 1$, a state of which is defined by a configuration $\mathbf{s} = (s_1, \dots, s_n)$. The number of all possible states is 2^n . A generic model is written as:

$$P(\mathbf{s}|\mathbf{g}, \mathcal{M}) = \frac{1}{\mathcal{Z}} e^{\sum_{\mu \in \mathcal{M}} g^\mu \phi^\mu(\mathbf{s})}, \quad (5)$$

where $\phi^\mu(\mathbf{s}) = \prod_{i \in \mu} s_i$ is the product of all spins involved in the corresponding interaction g^μ , the sum on μ runs on a subset of operators ϕ^μ and \mathcal{Z} ensures normalization. We follow the same notation as in [18]: There are 2^n possible such operators, which can be indexed by an integer $\mu = 0, \dots, 2^n - 1$ whose binary representation indicates those spins that occur in the operator ϕ^μ . Therefore, $\mu = 0$ corresponds to the constant operator $\phi^0(\mathbf{s}) = 1$, and for $\mu = 11$, the notation $i \in \mu$ is equivalent to $i \in \{1, 2, 4\}$, i.e., $\phi^{11}(\mathbf{s}) = s_1 s_2 s_4$.

Given a dataset $\hat{\mathbf{s}}$ of N observations of \mathbf{s} , assuming them to be i.i.d. draws from $P(\mathbf{s}|\mathbf{g}, \mathcal{M})$, the parameters g^μ are determined by solving Equation (3). Bayesian inference maintains that different models should be compared on the basis of the posterior $P\{\mathcal{M}|\hat{\mathbf{s}}\}$ that can be computed integrating the likelihood over the parameters \mathbf{g} (we refer to [18] for a discussion of Bayesian model selection within this setup). This can already be a daunting task if n is large. Note that each operator ϕ^μ ($\mu > 0$) can either be present or not in \mathcal{M} ; this implies that the number of possible models is $2^{2^n - 1}$. Finding the most likely model is impossible, in practice, even for moderately large n .

2.2. Bayesian Model Selection on Mixture Models

Let us consider mixture models in which the probability of state \mathbf{s} is of the form:

$$P(\mathbf{s}|\mathbf{q}, \mathcal{Q}) = \sum_{j=1}^q \rho_j \mathbb{1}_{Q_j}(\mathbf{s}) \quad (6)$$

where $\mathbb{1}_A(\mathbf{s})$ is the indicator function $\mathbb{1}_A(\mathbf{s}) = 1$ if $\mathbf{s} \in A$, and $\mathbb{1}_A(\mathbf{s}) = 0$ otherwise. The prior (and posterior) distributions of \mathbf{q} take a Dirichlet form; see Appendix B. In other words, model \mathcal{Q} assigns the same probability ρ_j to all configurations in the same set Q_j . Formally, $\mathcal{Q} = \{Q_i\}$ is a partition of the set of configurations \mathbf{s} , i.e., a collection of subsets such that $\{\pm 1\}^n = \bigcup_j Q_j$ and $Q_j \cap Q_{j'} = \emptyset \forall j \neq j'$. The model's parameters ρ_j are subject to the normalization constraint $\sum_j |Q_j| \rho_j = 1$, where $|Q|$ stands for the number of elements within Q . We denote by $q = |\mathcal{Q}|$ the number of subsets in \mathcal{Q} . The number of independent parameters in model \mathcal{Q} is then $q - 1$.

The number of possible models \mathcal{Q} is the number of partitions of a set of 2^n elements, which is the Bell number B_{2^n} . This grows even faster than the number of spin models \mathcal{M} . Yet, Bayesian model selection can be easily carried out, as shown in [16], assuming Dirichlet's prior. In brief, the most likely partition \mathcal{Q}^* depends on the assumed prior, but it is such that if two states \mathbf{s} and \mathbf{s}' are observed a similar number of times $k_s \simeq k_{s'}$, then the most likely model places them in the same set Q_q and assigns them the same probability ρ_q . In other words, considering the frequency partition:

$$\mathcal{K} = \{Q_k\}, \quad Q_k = \{\mathbf{s} : k_s = k\}, \quad (7)$$

that groups in the same subset all states \mathbf{s} that are observed the same number k_s of times, the optimal partition \mathcal{Q}^* is always a coarse graining of \mathcal{K} , likely to merge together subsets corresponding to similar empirical frequencies.

For example, in the $n = 1$ case, the possible partitions are $\mathcal{Q}^{(1)} = \{\{\pm 1\}\}$ and $\mathcal{Q}^{(2)} = \{\{-1\}, \{+1\}\}$. The first corresponds to a model with no parameters where $P(s|\mathcal{Q}^{(1)}) = 1/2$ for $s = \pm 1$, whereas the second assigns probability $P(s = +1|\mathcal{Q}^{(2)}) = \rho_+$ and $P(s = -1|\mathcal{Q}^{(2)}) = 1 - \rho_+$. Detailed calculation [16] shows that model $\mathcal{Q}^{(1)}$ should be preferred unless \bar{s} is sufficiently different from zero (i.e., the frequency with which the states $s = \pm 1$ are observed is sufficiently different).

We refer the interested reader to Appendix B and [16] for more details, as well as for a heuristic for finding the \mathcal{Q}^* model.

3. Mapping Mixture Models into Spin Models

Model \mathcal{Q} allows for a representation in terms of the variables g^μ , thanks to the relation:

$$g_{\mathcal{Q}}^\mu = \frac{1}{2^n} \sum_s \phi^\mu(s) \log P(s|\mathbf{q}, \mathcal{Q}) = \frac{1}{2^n} \sum_j \chi_j^\mu \log \rho_j, \quad \chi_j^\mu = \sum_{s \in Q_j} \phi^\mu(s), \quad (8)$$

which is of the same nature of the one discussed in [11] and whose proof is deferred to Appendix A. The index in $g_{\mathcal{Q}}^\mu$ indicates that the coupling refers to model \mathcal{Q} and merely corresponds to a change of variables $\mathbf{q} \rightarrow \mathbf{g}$; we shall drop it in what follows, if it causes no confusion.

In Bayesian inference, \mathbf{q} should be considered as a random variable, whose posterior distribution for a given dataset \hat{s} can be derived (see [16] and Appendix B). Then, Equation (8) implies that also \mathbf{g} is a random variable, whose distribution can also be derived from that of \mathbf{q} .

Notice, however, that Equation (8) spans only a $q - 1$ -dimensional manifold in the $2^n - 1$ -dimensional space \mathbf{g} , because there are only $q - 1$ independent variables \mathbf{q} . This fact is made more evident by the following argument: Let \mathbf{v} be a $2^n - 1$ component vector such that:

$$\sum_{\mu=1}^{2^n-1} v^\mu \chi_j^\mu = 0, \quad \forall j = 1, \dots, q. \quad (9)$$

Then, we find that:

$$\sum_{\mu} v^\mu g^\mu = 0. \quad (10)$$

In other words, the linear combination of the random variables g^μ with coefficients v^μ that satisfy Equation (9) is not random at all. There are (generically) $2^n - 1 - q$ vectors \mathbf{v} that satisfy Equation (9) each of which imposes a linear constraint of the form of Equation (10) on the possible values of \mathbf{g} .

In addition, there are q orthogonal directions \mathbf{u}_λ that can be derived from the singular value decomposition of χ_j^μ :

$$\chi_j^\mu = \sum_{\lambda=1}^q \Lambda_\lambda u_\lambda^\mu w_{\lambda,j}, \quad \sum_{\mu} u_\lambda^\mu u_{\lambda'}^\mu = \delta_{\lambda,\lambda'}, \quad \sum_j w_{\lambda,j} w_{\lambda',j} = \delta_{\lambda,\lambda'}. \quad (11)$$

This in turn implies that model \mathcal{Q} can be written in the exponential form (see Appendix D for details):

$$P(\mathbf{s}|\mathbf{g}, \mathcal{Q}) = \frac{1}{Z} e^{\sum_{\lambda=1}^q g_\lambda \psi_\lambda(\mathbf{s})}, \quad (12)$$

where:

$$\psi_\lambda(\mathbf{s}) = \sum_{\mu} u_\lambda^\mu \phi^\mu(\mathbf{s}); \quad (13)$$

$$g_\lambda = \sum_{\mu} u_\lambda^\mu g^\mu. \quad (14)$$

The exponential form of Equation (12) identifies the variables $\psi_\lambda(s)$ with the sufficient statistics of the model. The maximum likelihood parameters \hat{g}^λ can be determined using the knowledge of empirical averages of $\psi_\lambda(s)$ alone, solving the equations $\langle \psi_\lambda \rangle = \overline{\psi_\lambda}$ for all $\lambda = 1, \dots, q$. The resulting distribution is the maximum entropy distribution that reproduces the empirical averages of $\psi_\lambda(s)$. In this precise sense, $\psi_\lambda(s)$ are the relevant variables. Notice that, the variables $\psi_\lambda(s)$ are themselves an orthonormal set:

$$\sum_s \psi_\lambda(s) = 0, \quad \frac{1}{2^n} \sum_s \psi_\lambda(s) \psi_{\lambda'}(s) = \delta_{\lambda, \lambda'}. \quad (15)$$

In particular, if we focus on the \mathcal{K} partition of the set of states, the one assigning the same probability ρ_k to all states s that are observed k times, we find that $P(s|\hat{g}, \mathcal{Q}) = k_s/N$ exactly reproduces the empirical distribution. This is a consequence of the fact that the variables \hat{g}_λ that maximize the likelihood must correspond to the maximum likelihood estimates $\hat{\rho}_k = k/N$, via Equation (8). This implies that the maximum entropy distribution Equation (12) reproduces not only the empirical averages $\psi_\lambda(s)$, but also that of the operators $\phi^\mu(s)$ for all μ . A direct application of Equation (8) shows that the maximum entropy parameters are given by the formula:

$$\hat{g}_\mathcal{K}^\mu = \frac{1}{2^n} \sum_s \phi^\mu(s) \log \frac{k_s}{N} = \frac{1}{2^n} \sum_k \chi_k^\mu \log \frac{k}{N}. \quad (16)$$

Similarly, the maximum likelihood parameters \hat{g}_λ are given by:

$$\hat{g}_\lambda = \frac{1}{2^n} \sum_s \psi_\lambda(s) \log \frac{k_s}{N} = \frac{\Lambda_\lambda}{2^n} \sum_k w_{\lambda,k} \log \frac{k}{N}. \quad (17)$$

Notice that, when the set $Q_0 = \{s : k_s = 0\}$ of states that are not observed is not empty, all couplings \hat{g}^μ with $\chi_0^\mu \neq 0$ diverge. Similarly, all \hat{g}_λ with $w_{\lambda,0} \neq 0$ also diverge. We shall discuss later how to regularize these divergences that are expected to occur in the under-sampling regime (i.e., when $N \leq 2^n$).

It has to be noted that, of the q parameters ρ_q , only $q - 1$ are independent. Indeed, we find that one of the q singular values Λ_λ in Equation (11) is practically zero. It is interesting to inspect the covariance matrix $C^{\mu,\nu} = E[\delta g^\mu \delta g^\nu]$ of the deviations $\delta g^\mu = g^\mu - E[g^\mu]$ from the expected values computed on the posterior distribution. We find (see Appendices C and D) that $C^{\mu,\nu}$ has eigenvalues Λ_λ^2 along the eigenvectors \mathbf{u}_λ and zero eigenvalues along the directions \mathbf{v} . The components λ with the largest singular value Λ_λ are those with the largest statistical error, so one would be tempted to consider them as “sloppy” directions, as in [19]. Yet, by Equation (17), the value of \hat{g}_λ itself is proportional to Λ_λ , so the relative fluctuations are independent of Λ_λ . Indeed “sloppy” modes appear in models that overfit the data, whereas in our case, model selection on mixtures ensures that the model \mathcal{Q}^* does not overfit. This is why relative errors on the parameters g_λ are of comparable magnitude. Actually, variables ψ_λ that correspond to the largest eigenvalues Λ_λ are the most relevant ones, since they identify the directions along which the maximum likelihood distribution Equation (12) tilts most away from the unconstrained maximal entropy distribution $P_0(s) = 1/2^n$. A further hint in this direction is that Equation (17) implies that variables $\psi_\lambda(s)$ with the largest Λ_λ are those whose variation across states s typically correlates mostly with the variation of $\log k_s$ in the sample.

Notice that the procedure outlined above produces a model that is sparse in the q variables, i.e., it depends only on $q - 1$ parameters, where q is, in the case of the \mathcal{K} partition, the number of different values that k_s takes in the sample. Yet, it is not sparse in the g^μ variables. Many of the results that we have derived carry through with obvious modifications if the sums over μ are restricted to a subset \mathcal{M} of putatively relevant interactions. Alternatively, the results discussed above can be the starting point for the approximate scheme to find sparse models in the spin representation.

4. Illustrative Examples

In the following, we present simple examples clarifying the effects of the procedure outlined above.

4.1. Recovering the Generating Hamiltonian from Symmetries: Two and Four Spins

As a simple example, consider a system of two spins. The most general Hamiltonian that should be considered in the inference procedure is:

$$H^{\text{Inf}}[s] = g^1 \cdot s_1 + g^2 \cdot s_2 + g^3 \cdot s_1 s_2. \quad (18)$$

Imagine the data are generated from the Hamiltonian:

$$H^{\text{Gen}}[s] = J \cdot s_1 s_2. \quad (19)$$

and let us assume that the number of samples is large enough, so that the optimal partition \mathcal{Q}^* groups configurations of aligned spins $Q_+ = \{s : s_1 = s_2\}$ distinguishing them from the configuration of unaligned ones $Q_- = \{s : s_1 = -s_2\}$.

Following the strategy explained in Section 2.1, we observe that $\chi_{\pm}^{\mu} = \chi_{\mp}^{\mu} = 0$ for both $\mu = 1$ and 2. Therefore, Equation (8) implies $g^1 = g^2 = 0$. Therefore, the \mathcal{Q}^* model only allows for g^3 to be nonzero. In this simple case, symmetries induced by the \mathcal{Q}^* model (i.e., $(s_1, s_2) \rightarrow (-s_1, -s_2)$) directly produce a sparse model where all interactions that are not consistent with them are set to zero.

Consider now a four-spin system. Suppose that the generating Hamiltonian is that of a pairwise fully-connected model as in Figure 1 (left), with the same couplings $g_3 = g_5 = g_6 = g_9 = g_{10} = g_{12} = J$. With enough data, we can expect that the optimal model is based on the partition \mathcal{Q}^* that distinguishes three sets of configurations:

$$Q_j = \{s : s_1 + s_2 + s_3 + s_4 = \pm 2j\}, \quad j = 0, 1, 2 \quad (20)$$

depending on the absolute value of the total magnetization. The \mathcal{Q}^* model assigns the same probability ρ_j to configurations s in the same set Q_j . Along similar lines to those in the previous example, it can be shown that any interaction of order one is put to zero ($g^1 = g^2 = g^4 = g^8 = 0$), as well as any interaction of order three ($g^7 = g^{11} = g^{13} = g^{14} = 0$), because the corresponding interactions are not invariant under the symmetry $s \rightarrow -s$ that leaves \mathcal{Q}^* invariant. The interactions of order two will on the other hand correctly be nonzero and take on the same value $g_3 = g_5 = g_6 = g_9 = g_{10} = g_{12}$. The value of the four-body interaction is:

$$g^{15} = \frac{1}{2^4} \left[2 \cdot \log \rho_2 - 2 \cdot \binom{4}{1} \log \rho_1 + 2 \cdot \binom{4}{2} \log \rho_0 \right]. \quad (21)$$

This, in general, is different from zero. Indeed, a model with two- and four-body interactions shares the same partition \mathcal{Q}^* in Equation (20). Therefore, unlike the example of two spins, symmetries of the \mathcal{Q}^* model do not allow one to recover uniquely the generative model (Figure 1, left). Rather, the inferred model has a fourth order interaction (Figure 1, right) that cannot be excluded on the basis of symmetries alone. Note that there are $2^{2^4-1} = 32,768$ possible models of four spins. In this case, symmetries allow us to reduce the set of possible models to just two.



Figure 1. A representation of the two models consistent with the Q^* partition in Equation (20). Blue links represent pairwise interactions and the shaded area represents a four body interaction.

4.2. Exchangeable Spin Models

Consider models where $P(\mathbf{s})$ is invariant under any permutation π of the spins, i.e., $P(s_1, \dots, s_n) = P(s_{\pi_1}, \dots, s_{\pi_n})$. For these models, $P(\mathbf{s})$ only depends on the total magnetization $\sum_i s_i$. For example, the fully-connected Ising model:

$$P(\mathbf{s}) = \frac{1}{Z} e^{\frac{\beta}{n} \sum_{i < j} s_i s_j} \quad (22)$$

belongs to this class. It is natural to consider the partition where Q_q contain all configurations with q spins $s_i = -1$ and $n - q$ spins $s_j = +1$ ($q = 0, 1, \dots, n$). Therefore, when computing χ_q^μ , one has to consider $|Q_q| = \binom{n}{q}$ configurations. If μ involves m spins, then $\binom{m}{j} \binom{n-m}{q-j}$ of them will involve j spins $s_i = -1$, and the operator $\phi^\mu(\mathbf{s})$ takes the value $(-1)^j$ on these configurations. Therefore, χ_q^μ only depends on the number $m = |\mu|$ of spins involved and:

$$\chi_q^\mu = \chi_{q,m} \equiv \frac{1}{2^n} \sum_{j=0}^q \binom{m}{j} \binom{n-m}{q-j} (-1)^j, \quad m = |\mu|.$$

This implies that the coefficients g^μ of terms that involve m spins must all be equal. Indeed, for any two operators $\mu \neq \mu'$ with $|\mu| = |\mu'|$:

$$g^\mu = \sum_{q=0}^n \chi_{q,m} \log \mu_q = g^{\mu'}. \quad (23)$$

Therefore, the proposed scheme is able, in this case, to reduce the dimensionality of the inference problem dramatically, to models where interactions g^μ only depend on the number $m = |\mu|$ of spins involved in ϕ^μ .

Note also that any non-null vector v^μ such that $\sum_\mu v^\mu = 0$ and $v^\mu = 0$ if $|\mu| \neq m > 0$ satisfies:

$$\sum_\mu v^\mu \chi_q^\mu = 0.$$

The vectors u_λ corresponding to the non-zero singular values of $\hat{\chi}$ need to be orthogonal to each of these vectors, so they need to be constant for all μ that involve the same number of spins. In other words, $u_\lambda^\mu = u_\lambda(|\mu|)$ only depend on the number $|\mu|$ of spins involved. A suitable choice of a set of n independent eigenvectors in this space is given by $u_\lambda(m) = a \delta_{\lambda,m}$ that correspond to vectors that are constant within the sectors of μ with $|\mu| = \lambda$ and are zero outside. In such a case, the sufficient statistics for models of this type are:

$$\psi_\lambda(\mathbf{s}) = \delta_{\sum_i s_i, n-2\lambda},$$

as it should indeed be. We note in passing that terms of this form have been used in [20].

Inference can also be carried out directly. We first observe that the g_λ are defined up to a constant. This allows us to fix one of them arbitrarily, so we will take $g_0 = 0$. If K_λ is the number of observed configurations with λ spins $s_i = -1$, then the equation $\langle \phi_\lambda \rangle = \bar{\phi}_\lambda$ (for $\lambda > 0$) reads:

$$\binom{n}{\lambda} \frac{e^{g_\lambda}}{\mathcal{Z}} = \frac{K_\lambda}{N}, \quad \mathcal{Z} = 1 + \sum_{\lambda=1}^n \binom{n}{\lambda} e^{g_\lambda}$$

so that, after some algebra,

$$g_\lambda = \log \left(\frac{K_\lambda}{\binom{n}{\lambda} K_0} \right).$$

From this, one can go back to the couplings of operators g^μ using:

$$g^\mu = \sum_{\lambda=0}^n \log \left(\frac{K_\lambda}{\binom{n}{\lambda} K_0} \right) \frac{1}{2^n} \sum_{j=0}^{\lambda} \binom{|\mu|}{j} \binom{n-|\mu|}{\lambda-j} (-1)^j$$

Figure 2 illustrates this procedure for the case of the mean field (pairwise) Ising model Equation (22). As this shows, the procedure outlined above identifies the right model when the number of samples is large enough. If N is not large enough, large deviations from theoretical results start arising in couplings of highest order, especially if β is large.




Figure 2. Inferred parameters of a mean field Ising model Equation (22) with $n = 10$ (bottom) and 20 (top) spins (only interactions up to $m = 10$ spins are shown for $n = 20$). (left) Couplings g^m of m -th order interactions as a function of β ($N = 10^3$ samples for $n = 10$ and $N = 10^4$ for $n = 20$). (right) g^m as a function of N for $\beta = 1$. The correct value $g^2 = \beta/n$ is shown as a full line.

4.3. The Deep Under-Sampling Limit

The case where the number N of sampled configurations is so small that some of the configurations are never observed deserves some comments. As we have seen, taking the frequency partition \mathcal{K} , where $Q_k = \{s : k_s = k\}$, if $Q_0 \neq \emptyset$, then divergences can manifest in those couplings where $\chi_0^\mu \neq 0$.

It is instructive to consider the deep under-sampling regime where the number N of visited configurations is so small that configurations are observed at most once in the sample. This occurs when $N \ll 2^n$. In this case, the most likely partitions are (i) the one where all states have the same probability Q_0 and (ii) the one where states observed once have probability ρ/N and states not yet observed have probability $(1 - \rho)/(2^n - N)$, i.e., $Q_1 = \{Q_0, Q_1\}$ with $Q_k = \{s : k_s = k\}$, $k = 0, 1$. Following [16], it is easy to see that that generically, the probability of model Q_0 overweights the one of model Q_1 , because $P\{\hat{s}|Q_1\} \ll P\{\hat{s}|Q_0\}$. Under Q_0 , it is easy to see that $\chi_0^\mu = 0$ for all $\mu > 0$. This, in turn, implies that $g^\mu = 0$ exactly for all $\mu > 0$. We reach the conclusion that no interaction can be inferred in this case [21].

Taking instead the partition Q_1 , a straightforward calculation shows that Equation (8) leads to $g^\mu = a\bar{\phi}^\mu$. Here, a should be fixed in order to solve Equation (3). It is not hard to see that this leads to $a \rightarrow \infty$. This is necessary in order to recover empirical averages, which are computed assuming that unobserved states $s \in Q_0$ have zero probability.

This example suggests that the divergences that occur when assuming the \mathcal{K} partition, because of unobserved states ($k_s = 0$) can be removed by considering partitions where unobserved states are clamped together with states that are observed once.

These singularities arise because, when all the singular values are considered, the maximum entropy distribution exactly reproduces the empirical distribution. This suggests that a further method to remove these singularities is to consider only the first ℓ singular values (those with largest Λ_λ) and to neglect the others, i.e., to set $g_\lambda = 0$ for all other λ 's. It is easy to see that this solves the problem in the case of the deep under-sampling regime considered above. There, only one singular value exists, and when this is neglected, one derives the result $\hat{g}^\mu = 0$, $\forall \mu$ again. In order to illustrate this procedure in a more general setting, we turn to the specific case of the U.S. Supreme Court data [8].

4.4. A Real-World Example

We have applied the inference scheme to the data of [8] that refer to the decisions of the U.S. Supreme Court on 895 cases. The U.S. Supreme Court is composed of nine judges, each of whom casts a vote against ($s_i = -1$) or in favor ($s_i = +1$) of a given case. Therefore, this is a $n = 9$ spin system for which we have $N = 895$ observations. The work in [8] has fitted this dataset with a fully-connected pairwise spin model. We refer to [8] for details on the dataset and on the analysis. The question we wish to address here is whether the statistical dependence between judges of the U.S. Supreme Court can really be described as a pairwise interaction, which hints at the direct influence of one judge on another one, or whether higher order interactions are also present.

In order to address this issue, we also studied a dataset $n = 9$ spins generated from a pairwise interacting model, Equation (22), from which we generated $N = 895$ independent samples. The value of $\beta = 2.28$ was chosen so as to match the average value of two-body interactions fitted in the true dataset. This allows us to test the ability of our method to recover the correct model when no assumption on the model is made.

As discussed above, the procedure discussed in the previous section yields estimates \hat{g}^μ that allow us to recover empirical averages of all the operators. These, for a finite sample size N , are likely to be affected by considerable noise that is expected to render the estimated \hat{g}^μ extremely unstable. In particular, since the sample contains unobserved states, i.e., states with $k_s = 0$, we expect some of the parameters g^μ to diverge or, with a finite numerical precision, to attain large values.

Therefore, we also performed inference considering only the components with largest Λ_λ in the singular value decomposition. Table 1 reports the values of the estimated parameters \hat{g}_λ obtained for the U.S. Supreme Court considering only the top $\ell = 2$ to 7 singular values, and it compares them to those obtained when all singular values are considered. We observe that when enough singular values are considered, the estimated couplings converge to stable values, which are very different from those obtained when all 18 singular values are considered. This signals that the instability due to unobserved states can be cured by neglecting small singular values $\Lambda_\lambda \ll 1$.

Table 1. Singular values and estimated parameters for the U.S. Supreme Court data. The parameters $\hat{g}_\lambda^{(\ell)}$ refer to maximum entropy estimates of the model that considers only the top ℓ singular values (i.e., $\lambda \leq \ell$), whereas \hat{g}_λ in the last column refers to estimated parameters using all singular values.

λ	Λ_λ	$\hat{g}_\lambda^{(2)}$	$\hat{g}_\lambda^{(3)}$	$\hat{g}_\lambda^{(4)}$	$\hat{g}_\lambda^{(5)}$	$\hat{g}_\lambda^{(7)}$	\hat{g}_λ
1	0.528	0.946	1.023	1.347	1.512	1.510	3.680
2	0.250	−0.506	−0.573	−0.688	−0.722	−0.722	−1.213
3	0.159	0	0.256	0.358	0.378	0.377	0.519
4	0.102	0	0	−0.436	−0.492	−0.491	−0.601
5	0.073	0	0	0	−0.178	−0.131	−0.152
6	0.062	0	0	0	0	0.018	0.087
7	0.062	0	0	0	0	−0.010	−0.041
8	0.055	0	0	0	0	0	−0.222

This is confirmed by Figure 3, which shows that estimates of \hat{g}^μ are much more stable when few singular values are considered (top right panel). The top left panel, which refer to synthetic data generated from Equation (22), confirms this conclusion. The estimates \hat{g}^μ are significantly larger for a two-body interaction than for higher order and one-body interactions, as expected. Yet, when all singular values are considered, the estimated values of a two-body interaction fluctuate around values that are much larger than the theoretical one ($\beta/n \simeq 0.2533 \dots$) and the ones estimated from fewer singular values.




Figure 3. Inference of a system of $n = 9$ spins from a dataset of $N = 895$ samples. (left) Data generated from a pairwise Ising model Equation (22) with $\beta = 2.28$. (right) Data from the U.S. Supreme Court [8]. The upper panels report the estimated values of the parameters \hat{g}^μ as a function of the order m of the interaction. Different colors refer to inference limited to the largest $\ell = 2, 3, 5, 7$ singular values or to the case when all singular values are considered. The lower panels report the change in log likelihood (per sample point) when a single parameter g^μ is set to zero, as a function of the order $m = |\mu|$ of the interaction.

In order to test the performance of the inferred couplings, we measure for each operator μ the change:

$$\Delta_\mu = \max_{\mathbf{g}: g^\mu=0} \sum_\mu g^\mu \overline{\phi^\mu} - \sum_\mu \hat{g}^\mu \overline{\phi^\mu} \quad (24)$$

in log-likelihood when g^μ is set to zero. If Δ_μ is positive or is small and negative, the coupling g^μ can be set to zero without affecting much the ability of the model to describe the data. A large and negative Δ_μ instead signals a relevant interaction g^μ .

Clearly, $\Delta_\mu \leq 0$ for all μ when \hat{g}^μ is computed using all the q components. This is because in that case, the log-likelihood reaches the maximal value it can possibly achieve. When not all singular values are used, Δ_μ can also attain positive values.

Figure 3 confirms our conclusions that inference using all the components is unstable. Indeed for the synthetic data, the loss in likelihood is spread out on operators of all orders, when all singular values are considered. When few singular values are considered, instead, the loss in likelihood is heavily concentrated on two body terms (Figure 3, bottom left). Pairwise interactions stick out prominently because $\Delta_\mu < 0$ for all two-body operators μ . Still, we see that some of the higher order interactions, with even order, also generate significant likelihood losses.

With this insight, we can now turn to the U.S. Supreme Court data, focusing on inference with few singular values. Pairwise interactions stick out as having both sizable \hat{g}^μ (Figure 3, top right) and significant likelihood loss (Figure 3, bottom right). Indeed, the top interactions (those with minimal Δ_μ) are prevalently pairwise ones. Figure 4 shows the hypergraph obtained by considering the top 15 interactions [22], which are two- or four-body terms (see the caption for details). Comparing this with synthetic data, where we find that the top 19 interactions are all pairwise, we conjecture that four-body interactions may not be spurious. The resulting network clearly reflects the orientation of individual judges across an ideological spectrum going from liberal to conservative positions (as defined in [8]). Interestingly, while two-body interactions describe a network polarized across this spectrum with two clear groups, four-body terms appear to mediate the interactions between the two groups. The prevalence of two-body interactions suggests that direct interaction between the judges is clearly important, yet higher order interactions seem to play a relevant role in shaping their collective behavior.




Figure 4. Hypergraph of the top 15 interactions between the nine judges of the second Rehnquist Court. Judges are represented as nodes with labels referring to the initials (as in [8]). Two-body interactions are represented by (red) links of a width that increases with $|\Delta_\mu|$, whereas four-body interactions as (green) shapes joining the four nodes. The shade of the nodes represents the ideological orientation, as reported in [8], from liberal (black) to conservative (white).

As in the analysis in [8], single-body terms, representing a priori biases of individual judges, are not very relevant [23].

5. Conclusions

The present work represents a first step towards a Bayesian model selection procedure for spin models with interactions of arbitrary order. Rather than tackling the problem directly, which would imply comparing an astronomical number of models even for moderate n , we show that model selection can be performed first on mixture models, and then, the result can be projected in the space of spin models. This approach spots symmetries between states that occur with a similar frequency, which impose constraints between the parameters g^μ . As we have seen, in simple cases, these symmetries are enough to recover the correct sparse model, imposing that $g^\mu = 0$ for all those interactions ϕ^μ that are not consistent with the symmetries. These symmetries allow us to derive a set of sufficient statistics $\psi_\lambda(s)$ (the relevant variables) whose empirical values allow one to derive the maximum likelihood parameters \hat{g}_λ . The number q of sufficient statistics is given by the number of sets in the optimal partition \mathcal{Q}^* of states. Therefore, the dimensionality of the inference problem is not related to the number of different interaction terms $\phi^\mu(s)$ (or equivalently, of parameters g^μ), but it is rather controlled by the number of different frequencies that are observed in the data. As the number N of samples increases, q increases and so does the dimensionality of the inference problem, until one reaches the well-sampled regime ($N \gg 2^n$) when all states s are well resolved in frequency.

It has been observed [11] that the family of probability distributions of the form (5) is endowed with a hierarchical structure that implies that high-order and low-order interactions are entangled in a nontrivial way. For example, we observe a non-trivial dependence between two- and four-body interactions. On the other hand, Ref. [18] shows that the structure of interdependence between operators in a model is not simply related to the order of the interactions and is invariant with respect to gauge transformations that do not conserve the order of operators. This, combined with the fact that our approach does not introduce any explicit bias to favor an interaction of any particular order, suggests that the approach generates a genuine prediction on the relevance of interactions of a particular order (e.g., pairwise). Yet, it would be interesting to explore these issues further, combining the quasi-orthogonal decomposition introduced in [11] with our approach.

It is interesting to contrast our approach with the growing literature on sloppy models (see, e.g., [19]). Transtrum et al. [19] have observed that inference of a given model is often plagued by overfitting that causes large errors in particular combinations of the estimated parameters.

Our approach is markedly different in that we stem right from the beginning from Bayesian model selection, and hence, we rule out overfitting from the outset. Our decomposition in singular values identifies those directions in the space of parameters that allow one to match the empirical distribution while preserving the symmetries between configurations observed with a similar frequency.

The approach discussed in this paper is only feasible when the number of variables n is small. Yet, the generalization to a case where the set \mathcal{M} of interactions is only a subset of the possible interactions is straightforward. This entails setting to zero all couplings g^μ relative to interactions $\mu \notin \mathcal{M}$. Devising decimation schemes for doing this in a systematic manner, as well as combining our approach with regularization schemes (e.g., LASSO) to recover sparse models comprise a promising avenue of research for exploring the space of models.

Acknowledgments: We gratefully acknowledge Edward D. Lee and William Bialek for sharing the data of [8]. We are grateful to Iacopo Mastromatteo, Vijay Balasubramanian, Yasser Roudi, Clélia de Mulatier and Paolo Pietro Mazza for insightful discussions.

Author Contributions: Luigi Gresele and Matteo Marsili conceived the research, performed the analysis and wrote the paper. Both authors have read and approved the final manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. The Completeness Relation (8)

We notice that the set of operators satisfies the following orthogonality relations:

$$\sum_{\{s\}} \phi^\mu(s) \phi^\nu(s) = \delta_{\mu,\nu} \cdot 2^n; \quad (\text{A1})$$

$$\sum_{\{s\}} \phi^\mu(s) = 0, \quad \forall \mu > 0. \quad (\text{A2})$$

Then, taking the logarithm of Equation (5), multiplying by $\phi^\nu(s)/2^n$ with $\nu > 0$ and summing over s , one finds that:

$$\begin{aligned} \frac{1}{2^n} \sum_{\{s\}} \phi^\nu(s) \cdot \left(\sum_{\mu} \phi^\mu(s) g^\mu - \log \mathcal{Z} \right) \\ = \sum_{\mu} \delta_{\mu,\nu} g^\mu = g^\nu, \end{aligned} \quad (\text{A3})$$

Combining the above identity with the expression with Equation (6) finally yields Equation (8).

Appendix B. The Posterior Distribution of ϱ

Following [16], we assume a Dirichlet prior for the parameter vector ϱ , i.e.,

$$P_0^Q(\varrho) = \frac{\Gamma(a \cdot 2^n)}{[\Gamma(a)]^{2^n}} \prod_i \rho_i^{a-1} \delta \left(\sum_i \rho_i - 1 \right). \quad (\text{A4})$$

This is a natural choice due to the fact that it is a conjugate prior for the parameters of the multinomial distribution [24]. This means that the posterior distribution has the same functional form as the prior, and the a parameters can be interpreted as pseudocounts. In other words, the posterior probability of $\vec{\rho}$ is still a Dirichlet distribution, i.e.,

$$P_1^Q(\varrho|\hat{s}) = \Gamma \left(\sum_i k_i + a \right) \prod_i \frac{\rho_i^{k_i+a-1}}{\Gamma(k_i+a)} \delta \left(\sum_i \rho_i - 1 \right), \quad (\text{A5})$$

where k_s is the number of times state s_j was observed in the sample. We remind that the choice $a = 0.5$ corresponds to the least informative (Jeffrey's) prior [25], whereas with $a = 0$, the expected values of the parameters ϱ_q coincide with the maximal likelihood estimates.

The likelihood of the sample \hat{s} under model Q is given by:

$$P\{\hat{s}|Q\} = \sum_j \left[\log \frac{\Gamma(K_j + a)}{\Gamma(a)} - K_j \log m_j \right] - \log \frac{\Gamma(N + aq)}{\Gamma(aq)}$$

where:

$$K_j = \sum_{s \in Q_j} k_s,$$

is the number of sample points in partition Q_j and $m_j = |Q_j|$ is the number of states in partition Q_j . The work in [16] shows that the partition that maximizes the likelihood is the one where states with similar frequencies are grouped together, which is a coarse-graining of the \mathcal{K} partition.

Appendix C. The Covariance Matrix of the g^μ Parameters

The covariance matrix \hat{C} can be written as:

$$C^{\mu\nu} = \text{Cov}[g^\mu, g^\nu] = \frac{1}{2^{2n}} \sum_{q,q'} \chi_q^\mu \chi_{q'}^\mu \text{Cov}[\log \rho_q, \log \rho_{q'}]. \quad (\text{A6})$$

In order to compute the covariance of the g parameters, $\text{Cov}[\log \rho_q, \log \rho_{q'}]$ must be computed first. For the following considerations, it is useful to define:

$$\mathcal{Z}(\vec{\lambda}) = \mathbf{E} \left[\prod_q \rho_q^{\lambda_q} \right], \quad (\text{A7})$$

where the expectation is taken over the posterior distribution of $\vec{\rho}$:

$$P_1^Q(\vec{\rho}|\hat{s}) = \Gamma \left(\sum_i k_i + a \right) \prod_i \frac{\rho_i^{k_i+a-1}}{\Gamma(k_i+a)} \delta \left(\sum_i \rho_i - 1 \right), \quad (\text{A8})$$

as in [16]. First, one can see that:

$$\begin{aligned} \partial_{\lambda_q} \partial_{\lambda_{q'}} \mathcal{Z}(\vec{\lambda}) &= \mathbf{E} \left[\delta_{q,q'} \cdot \rho_q^{\lambda_q} (\log \rho_q)^2 \prod_{q' \neq q''} \rho_{q''}^{\lambda_{q''}} \right] + \\ &+ \mathbf{E} \left[(1 - \delta_{q,q'}) \cdot \rho_q^{\lambda_q} \log \rho_q \cdot \rho_{q'}^{\lambda_{q'}} \log \rho_{q'} \prod_{q' \neq q'', q'''} \rho_{q'''}^{\lambda_{q'''}} \right]. \end{aligned}$$

This relation, for $\vec{\lambda} = 0$, yields:

$$\partial_{\lambda_q} \partial_{\lambda_{q'}} \mathcal{Z}(\vec{\lambda}) \Big|_{\vec{\lambda}=0} = \mathbf{E} \left[\delta_{q,q'} \cdot (\log \rho_q)^2 + (1 - \delta_{q,q'}) \cdot \log \rho_q \cdot \log \rho_{q'} \right] \quad (\text{A9})$$

$$= \mathbf{E} \left[\log \rho_q, \log \rho_{q'} \right], \quad (\text{A10})$$

and we find that:

$$\text{Cov}[\log \rho_q, \log \rho_{q'}] = \left[\partial_{\lambda_q} \partial_{\lambda_{q'}} \mathcal{Z}(\vec{\lambda}) - \partial_{\lambda_q} \mathcal{Z}(\vec{\lambda}) \partial_{\lambda_{q'}} \mathcal{Z}(\vec{\lambda}) \right] \Big|_{\vec{\lambda}=0} \quad (\text{A11})$$

$$= \frac{\partial^2}{\partial \lambda \partial \lambda'} \left[\log \mathcal{Z}(\vec{\lambda}) \right] \Big|_{\vec{\lambda}=0} \quad (\text{A12})$$

Given the distribution of $\vec{\rho}$, the expression in Equation (A7) reads:

$$\mathcal{Z}(\vec{\lambda}) = \frac{1}{\prod_q m_q^\lambda} \cdot \frac{\prod_q \Gamma(K_q + \lambda_q + a)}{\Gamma(\sum_q (K_q + \lambda_q + a))}. \quad (\text{A13})$$

then:

$$\begin{aligned} \frac{\partial^2}{\partial \lambda \partial \lambda'} \left[\log \mathcal{Z}(\vec{\lambda}) \right] &= \frac{\partial^2}{\partial \lambda \partial \lambda'} \left[\sum_q \log \Gamma(K_q + \lambda_q + a) - \log \Gamma(\sum_q K_q + \sum_q \lambda_q + \sum_q a) \right] \\ &= \frac{\partial}{\partial \lambda'} \left[\frac{\Gamma'(K_q + \lambda_q + a)}{\Gamma(K_q + \lambda_q + a)} - \frac{\Gamma'(M + \sum_q \lambda_q + Na)}{\Gamma(M + \sum_q \lambda_q + Na)} \right] \\ &= \delta_{q,q'} \cdot \psi^{(1)}(K_q + \lambda_q + a) - \psi^{(1)}(M + \sum_q \lambda_q + Na) \end{aligned}$$

where $\psi^{(n)}(x) = \frac{d^{n+1}}{dx^{n+1}} \log \Gamma(x)$ is the polygamma function.

Evaluated in $\vec{\lambda} = 0$, this yields:

$$\text{Cov}[\log \rho_q, \log \rho_{q'}] = C_{q,q'} = \delta_{q,q'} \cdot \psi^{(1)}(K_q + a) - \psi^{(1)}(M + Na). \quad (\text{A14})$$

Therefore, the covariance matrix among the elements of $\vec{\rho}$ is composed of a diagonal part plus a part proportional to the identity matrix.

Inserting (A14) into (A6),

$$C^{\mu\nu} = \text{Cov}[g^\mu, g^\nu] = \frac{1}{2^{2n}} \sum_q \chi_q^\mu \chi_q^\nu \psi^{(1)}(K_q + a) - \frac{1}{2^{2n}} \cdot \psi^{(1)}(M + Na) \cdot \left(\sum_q \chi_q^\mu \right) \left(\sum_{q'} \chi_{q'}^\nu \right) \quad (\text{A15})$$

$$= \frac{1}{2^{2n}} \sum_q \chi_q^\mu \chi_q^\nu \psi^{(1)}(K_q + a), \quad (\text{A16})$$

due to the fact that, by Equation (A2), $\sum_{q'} \chi_{q'}^\nu = 0$ for all $\nu > 0$.

Appendix D. Sufficient Statistics

Property (10) can be used to give a proof of Equation (12).

Let $v_j^\mu, j = 1, \dots, 2^n - 1 - q$ be the vectors such that $\sum_\mu v_j^\mu g^\mu = 0$. With the u_λ^μ , these constitute an orthonormal basis for the $2^n - 1$ dimensional space with coordinates g^μ . This implies the identity:

$$\delta_{\mu,\nu} = \sum_\lambda u_\lambda^\mu u_\lambda^\nu + \sum_j v_j^\mu v_j^\nu$$

that can be inserted in the expression:

$$\sum_\mu g^\mu \phi^\mu(s) = \sum_{\mu,\nu} g^\mu \phi^\nu(s) \delta_{\mu,\nu} \quad (\text{A17})$$

$$= \sum_\lambda \left(\sum_\mu u_\lambda^\mu g^\mu \right) \left(\sum_\nu u_\lambda^\nu \phi^\nu(s) \right) + \sum_j \left(\sum_\mu v_j^\mu g^\mu \right) \left(\sum_\nu v_j^\nu \phi^\nu(s) \right) \quad (\text{A18})$$

that yields the desired result, because $\sum_\mu v_j^\mu g^\mu = 0$ for all j .

References and Notes

1. Jaynes, E.T. Information Theory and Statistical Mechanics. *Phys. Rev.* **1957**, *106*, 620–630.
2. Pitman, E.J.G. Sufficient statistics and intrinsic accuracy. In *Mathematical Proceedings of the Cambridge Philosophical Society*; Cambridge University Press: Cambridge, UK, 1936; Volume 32, pp. 567–579.
3. Darmon, G. Sur les lois de probabilité à estimation exhaustive. *C. R. Acad. Sci. Paris* **1935**, *200*, 1265–1266. (In French)
4. Koopman, B.O. On distributions admitting a sufficient statistic. *Trans. Am. Math. Soc.* **1936**, *39*, 399–409.
5. Ackley, D.H.; Hinton, G.E.; Sejnowski, T.J. A Learning Algorithm for Boltzmann Machines. *Cogn. Sci.* **1985**, *9*, 147–169.
6. Schneidman, E.; Berry, M.J., II; Segev, R.; Bialek, W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* **2006**, *440*, 1007–1012.
7. Nguyen, H.C.; Zecchina, R.; Berg, J. Inverse statistical problems: From the inverse Ising problem to data science. *arXiv* **2017**, arXiv:1702.01522.
8. Lee, E.; Broedersz, C.; Bialek, W. Statistical mechanics of the US Supreme Court. *J. Stat. Phys.* **2015**, *160*, 275–301.

9. Wainwright, M.J.; Jordan, M.I. Variational inference in graphical models: The view from the marginal polytope. In Proceedings of the Annual Allerton Conference on Communication Control and Computing, Allerton, IL, USA, 23–25 September 1998; Volume 41, pp. 961–971.
10. Sejnowski, T.J. Higher-order Boltzmann machines. *AIP Conf. Proc.* **1986**, *151*, 398–403.
11. Amari, S. *Information Geometry on Hierarchy of Probability Distributions*; IEEE: Hoboken, NJ, USA, 2001; Volume 47, pp. 1701–1711.
12. Margolin, A.; Wang, K.; Califano, A.; Nemenman, I. Multivariate dependence and genetic networks inference. *IET Syst. Biol.* **2010**, *4*, 428–440.
13. Merchan, L.; Nemenman, I. On the Sufficiency of Pairwise Interactions in Maximum Entropy Models of Networks. *J. Stat. Phys.* **2016**, *162*, 1294–1308.
14. Limiting inference schemes to pairwise interactions is non-trivial when variables take more than two values (e.g., Potts spins). A notable example is that of the inference of protein contacts from amino acid sequences. There, each variable can take 20 possible values; hence, there are 200 parameters for each pair of positions. Sequences are typically $n \sim 100$ amino acids long, so a pairwise model contains $200 n^2 / 2 \sim 10^6$ parameters. In spite of the fact that the number of available sequences is much less than that (i.e., $N \sim 10^3 \nabla \cdot 10^4$), learning Potts model parameters has been found to be an effective means to predict structural properties of proteins [7]. However, we will not enter into details related to the Potts model in the present work.
15. As already pointed out in [5], any higher order interaction can be reduced to pairwise interaction, introducing hidden variables. Conversely, higher order interactions may signal the presence of hidden variables.
16. Haimovici, A.; Marsili, M. Criticality of mostly informative samples: A Bayesian model selection approach. *J. Stat. Mech. Theory Exp.* **2015**, *2015*, P10013.
17. Collins, M.; Dasgupta, S.; Schapire, R.E. A Generalization of Principal Component Analysis to the Exponential Family. In *Advances in Neural Information Processing Systems*; The MIT Press: Cambridge, MA, USA, 2001; pp. 617–624.
18. Beretta, A.; Battistin, C.; Mulatier, C.; Mastromatteo, I.; Marsili, M. The Stochastic complexity of spin models: How simple are simple spin models? *arXiv* **2017**, arXiv:1702.07549.
19. Transtrum, M.K.; Machta, B.B.; Brown, K.S.; Daniels, B.C.; Myers, C.R.; Sethna, J.P. Perspective: Slowness and emergent theories in physics, biology, and beyond. *J. Chem. Phys.* **2015**, *143*, 010901.
20. Tkačik, G.; Marre, O.; Mora, T.; Amodei, D.; Berry, M.J., II; Bialek, W. The simplest maximum entropy model for collective behavior in a neural network. *J. Stat. Mech. Theory Exp.* **2013**, *2013*, P03011.
21. Notice that other inference methods may infer non-zero interactions in this case [7]. Note also that the statistics of the frequencies can be very different if one takes a subset of $n' < n$ spin, so the present approach may predict $g'' \neq 0$ when the same dataset is restricted to a subset of spins.
22. A conservative estimate of the number of significant interactions is given by the number of independent parameters g_λ in our data. These are 18 in the U.S. Supreme Court data and 12 in the synthetic data.
23. Reference [8] remarks that the definitions of “yes” and “no” are somewhat arbitrary and do not carry any information on the political orientation associated with a given vote, since they are decided in lower courts; it also shows that, even when a “left-wing/right-wing” label is attached to the “yes/no” votes, the fields alone do not explain the data well.
24. Gelman, A.; Carlin, J.B.; Stern, H.S.; Dunson, D.B.; Vehtari, A.; Rubin, D.B. *Bayesian Data Analysis*; Chapman and Hall/CRC Press: Boca Raton, FL, USA, 2014; Volume 2.
25. Box, G.E.P.; Tiao, G.C. *Bayesian Inference in Statistical Analysis*; Addison-Wesley Publishing Company: Boston, MA, USA, 1973.

