

Large pseudocounts and L_2 -norm penalties are necessary for the mean-field inference of Ising and Potts models

J. P. Barton,^{1,2} S. Cocco,³ E. De Leonardis,^{3,4} and R. Monasson⁵

¹*Department of Chemical Engineering, MIT, Cambridge, Massachusetts 02139, USA*

²*Ragon Institute of MGH, MIT and Harvard, Boston, Massachusetts 02129, USA*

³*Laboratory of Statistical Physics of the Ecole Normale Supérieure, associated to CNRS and University P&M Curie, 24 rue Lhomond, 75005 Paris, France*

⁴*UMR 7238, Computational and Quantitative Biology, UPMC Univ Paris 06, France Sorbonne Universités, 75005 Paris, France*

⁵*Laboratory of Theoretical Physics of the Ecole Normale Supérieure, associated to CNRS and University P&M Curie, 24 rue Lhomond, 75005 Paris, France*

(Received 1 April 2014; revised manuscript received 17 June 2014; published 28 July 2014)

The mean-field (MF) approximation offers a simple, fast way to infer direct interactions between elements in a network of correlated variables, a common, computationally challenging problem with practical applications in fields ranging from physics and biology to the social sciences. However, MF methods achieve their best performance with strong regularization, well beyond Bayesian expectations, an empirical fact that is poorly understood. In this work, we study the influence of pseudocount and L_2 -norm regularization schemes on the quality of inferred Ising or Potts interaction networks from correlation data within the MF approximation. We argue, based on the analysis of small systems, that the optimal value of the regularization strength remains finite even if the sampling noise tends to zero, in order to correct for systematic biases introduced by the MF approximation. Our claim is corroborated by extensive numerical studies of diverse model systems and by the analytical study of the m -component spin model for large but finite m . Additionally, we find that pseudocount regularization is robust against sampling noise and often outperforms L_2 -norm regularization, particularly when the underlying network of interactions is strongly heterogeneous. Much better performances are generally obtained for the Ising model than for the Potts model, for which only couplings incoming onto medium-frequency symbols are reliably inferred.

DOI: [10.1103/PhysRevE.90.012132](https://doi.org/10.1103/PhysRevE.90.012132)

PACS number(s): 02.50.Tt, 05.10.-a, 05.50.+q, 87.10.Mn

I. INTRODUCTION

Inferring effective interaction networks from the measured time series of a population of variables is a problem of increasing importance across multiple scientific disciplines, including biology (for the study of protein residue coevolution [1–7], gene networks [8–10], neuroscience [11–15], and ecology [16–19], among others), sociology [20–22], and finance [23,24]. One popular approach to this problem is to infer a simple graphical model which reproduces the low-order stationary statistics of the data, such as the single-variable frequencies and the pairwise correlations. Inferred model parameters then give clues about the underlying network of interactions between the variables.

When the “true” model is uncertain, in practice one often searches for the maximum entropy, or least constrained [25], model capable of reproducing the data. The Ising model is the maximum entropy model capable of reproducing the one- and two-point statistical constraints between binary variables, e.g., the activity of a population of neurons, which are either silent or emit a spike within a certain time window. When the variables take more than two values, for example, specifying the amino acid present at each site in a protein sequence, the Potts model is a natural extension. In both cases the computational problem consists of finding the set of couplings $J_{ij}(a,b)$, expressing the interactions between the “symbol” a of variable i and the symbol b of variable j , from the measured correlations $c_{ij}(a,b)$. This is referred to as the inverse Ising, or Potts, problem. An exact solution generally requires computational efforts increasing exponentially with the system size (number

of spin variables) N . Efficient and accurate approximation schemes are thus required for the analysis of real data and a host of methods have been developed with this goal in mind [26–32].

Among the algorithms developed for solving the inverse Ising and Potts problems, the mean-field (MF) inference procedure is certainly the simplest. MF is fast as it runs in a time growing polynomially with N . MF simply approximates the coupling matrix with minus the inverse of the correlation matrix, a result which would be exact for Gaussian distributed variables but which is only approximate for the Ising or the Potts model. This method has been shown to give very good results for various biologically motivated problems, such as the study of amino acid covariation in proteins [4–6] and multielectrode recordings of neural activity [33].

Despite its popularity, key ingredients for the success of MF inference remain poorly understood. In particular, an essential ingredient of the inference from real data is the presence of a regularization term ensuring that the inverse problem is always well defined. To this aim L_1 - or L_2 -norm regularization of the couplings, or pseudocount regularization of the correlations, can be used [26]. However, from a Bayesian point of view the optimal strength of these regularization terms is expected to decrease with the level of the noise and should vanish in the limit of perfect sampling. Empirical studies show that this is not the case for MF inference: The best performance of MF inference is achieved only in the presence of very strong regularization. Predictions of contacts between residues on protein folds based on MF inference are optimal

when the regularization terms are very strong, without any apparent dependence on the number of data [4–6,34]. Another related finding which has lacked any explanation so far is why pseudocounts are generally better than other regularization schemes, such as L_2 - or L_1 -norm regularization of couplings, when combined with MF inference.

Here we explore the performance of various regularization schemes for MF inference on Ising and Potts models, and the reasons behind their success, through the analytical analysis of small model spin systems combined with extensive numerical studies of larger systems. First, we show that abnormally strong regularization is necessary to correct for errors produced by the MF approximation itself, which we explain using the simplest case of models with few variables ($N = 2,3$). In addition we show that, in systems with homogeneous interactions, pseudocounts and L_2 -norm regularization are performing similar functions but not L_1 . When the (Ising or Potts) model includes a large number N of spin variables we show based on numerical simulations that the same phenomenon takes place: Large regularization is necessary, but pseudocounts do a better job than L_2 for strongly heterogeneous networks. We explain why this is so using analytical arguments based on the analysis of the $O(m)$ continuous spin model for large but finite m . MF is exact for this model in the $m \rightarrow \infty$ limit, and we show that the optimal pseudocount remains finite in the absence of sampling noise: The optimal penalty is of the order of $\frac{1}{m}$, which estimates the deviation of the model with respect to Gaussianity. Moreover, inference is less affected by sampling noise when using large pseudocount than when using a large L_2 -norm. Finally, we show that inference performances, even with large pseudocount, may be much poorer for the Potts than for the Ising model, especially so when the symbols on each site largely differ in their frequencies. Our study therefore provides a strong basis for the use of large regularization penalties with mean-field inference, which thus far had been totally empirical.

The paper is organized as follows. In Sec. II we present the different regularization schemes studied in the paper and briefly recall how couplings are inferred from correlations within the Gaussian (MF) approximation. In Sec. III we present a detailed analysis of the error of inference due to MF and how those errors are corrected, with varying success, in the presence of regularization. Section IV reports the performances of the MF as a function of the regularization strength based on extensive numerical simulations of Ising and Potts models and with diverse interaction distributions and structures. The statistical mechanics of the $O(m)$ model and results regarding the optimal value of the regularization penalty with MF inference are presented in Sec. V. Conclusions are proposed in Sec. VI.

II. REMINDER ON MEAN-FIELD INFERENCE AND REGULARIZATION

The mean-field approximation consists, as far as inference is concerned, of approximating the Ising (or the Potts) model couplings with the off-diagonal elements of minus the inverse of the correlation matrix. We recall below that this result can be found for the Ising model within the Gaussian approximation (Sec. II A), where the discrete nature of the spin variables

is omitted. Section II B briefly presents the regularization schemes studied here, namely the pseudocount and the L_2 -norm (as well as the L_1 -norm). Specificities of the inference applied to the Potts model are discussed in Sec. II C.

A. Inference of couplings within the Gaussian approximation

In the Gaussian approximation the differences between the Ising spin variables $\sigma_i = 0, 1$ and their empirical average values,

$$\hat{\sigma}_i = \sigma_i - \langle \sigma_i \rangle, \quad (1)$$

are assumed to be drawn from a Gaussian distribution, with zero mean and empirical covariance matrix $c_{ij} = \langle \hat{\sigma}_i \hat{\sigma}_j \rangle = \langle \sigma_i \sigma_j \rangle - \langle \sigma_i \rangle \langle \sigma_j \rangle$. The likelihood of a configuration is

$$P(\hat{\sigma}_1, \dots, \hat{\sigma}_N) = \frac{\sqrt{\det \hat{J}}}{(2\pi)^{N/2}} \exp\left(-\frac{1}{2} \sum_{i,j} \hat{J}_{ij} \hat{\sigma}_i \hat{\sigma}_j\right), \quad (2)$$

where the off-diagonal elements of the \hat{J} matrix coincide with the opposite of the couplings: $\hat{J}_{ij} = -J_{ij}$, for all $i \neq j$. Contrary to the Ising or Potts model the diagonal elements of \hat{J} are important to define the measure (2), and their values will be specified later.

In the formulas above $\langle \cdot \rangle$ denotes the empirical average over the data set, composed of B independently sampled configurations of the model. The log-likelihood of the data within the Gaussian model (2) is a function of its empirical covariance, given by

$$L(c|J) = \frac{B}{2} (-\text{trace}(\hat{J}c) + \log \det \hat{J}). \quad (3)$$

Maximization of L over \hat{J} for a fixed c gives $\hat{J} = c^{-1}$. For the off-diagonal entries we obtain

$$J_{ij} = -\hat{J}_{ij} = -(c^{-1})_{ij}, \quad (4)$$

while the diagonal couplings \hat{J}_{ii} are Lagrange parameters enforcing the N conditions $\langle \hat{\sigma}_i^2 \rangle = \langle \sigma_i \rangle (1 - \langle \sigma_i \rangle)$ for 0,1 spins or $\langle \hat{\sigma}_i^2 \rangle = 1 - \langle \sigma_i \rangle^2$ for ± 1 spins. Hence, inference with the Gaussian model provides the same expression for the couplings as the mean-field approximation [35]. In the following, we will use the subscript MF to refer to the couplings given by expression (4).

B. Regularization schemes

Empirical averages are computed from a finite number B of configurations, and therefore the correlation matrix c is not always invertible. Zero modes are found when some configurations of variables are never sampled, e.g., when variables σ_i and σ_j are equal in all B configurations. The invertibility of the correlation matrix can be ensured by introducing some form of regularization to the model. From a Bayesian point of view, such terms can be thought of as the contribution of prior distributions for the model parameters, and their amplitude should vanish in the limit of perfect sampling ($B \rightarrow \infty$). Below we review three popular regularization schemes: pseudocount and the L_1 - and L_2 -norm regularization of couplings. For the sake of simplicity,

definitions are given for the Ising model, the extension to the Potts case being straightforward.

1. Pseudocount

A very simple regularization scheme consists in adding extra ‘‘pseudo’’ observations to the real data in order to cure singularities caused by strong correlations. For instance, consider the case of a frozen variable, e.g., such that $\sigma_i = 0$ in all B configurations in the real data; σ_i could be given value 1 in a pseudo $(B + 1)$ th configuration. This is the ‘‘pseudocount’’ method, popular in the analysis of protein sequence data in biology, which can be interpreted in terms of a Dirichlet prior distribution for the observation of each state of the σ_i [36]. Typically, one chooses a prior distribution in which each value of σ_i is considered equally likely. In this case the pseudocount changes the frequencies and correlations in the following way:

$$\begin{aligned} \langle \sigma_i \rangle &\rightarrow (1 - \alpha) \langle \sigma_i \rangle + \frac{\alpha}{2}, \\ \langle \sigma_i \sigma_j \rangle &\rightarrow (1 - \alpha) \langle \sigma_i \sigma_j \rangle + \frac{\alpha}{4} \quad (i \neq j), \end{aligned} \quad (5)$$

for 0, 1 spins and

$$\langle \sigma_i \rangle \rightarrow (1 - \alpha) \langle \sigma_i \rangle, \quad \langle \sigma_i \sigma_j \rangle \rightarrow (1 - \alpha) \langle \sigma_i \sigma_j \rangle \quad (i \neq j), \quad (6)$$

for ± 1 spins. Diagonal terms are constrained to $\langle \sigma_i^2 \rangle = \langle \sigma_i \rangle$ for $\sigma_i = 0, 1$, and $\langle \sigma_i^2 \rangle = 1$ for $\sigma_i = \pm 1$. The amplitude α determines the strength of the pseudocount. We expect α to vanish as $B \rightarrow \infty$, as regularization should not be necessary in the case of perfect sampling (assuming all the fields and all the couplings in the underlying model for the data are finite). The frozen-variable example given above amounts to choosing $\alpha = \alpha^B = \frac{2}{B+1}$. Note, however, that pseudocount regularization can be applied to any data, even in the absence of frozen variables.

2. L_1 - and L_2 -norm regularization of couplings

Another possibility to prevent couplings from being infinite is to consider a prior probability distribution for the couplings which discounts large coupling values. The log of the prior distribution then adds to the log-likelihood (3), with the contribution

$$\Delta L(J) = -\frac{\gamma B}{2} \sum_{i < j} J_{ij}^2 \quad (7)$$

in the case of a Gaussian prior with variance $(\gamma B)^{-1}$ and mean zero. We have factored out B in (8) to allow for a direct comparison with the log-likelihood of the data set (3), which is proportional to B . In principle, the prior should depend neither on the data, nor on their number, B . Hence γ is expected to scale as $\frac{1}{B}$ and to vanish in the case of perfect sampling. In the following we use the following L_2 regularization, instead of (7):

$$\Delta L(J) = -\frac{\gamma B}{2} \sum_{i < j} p_i (1 - p_i) p_j (1 - p_j) J_{ij}^2, \quad (8)$$

where $p_i = \langle \sigma_i \rangle$ are the empirical averages of the binary variables. The reason is that the coupling matrix \mathbf{J} maximizing the

log-likelihood $L(c|J) + \Delta L(J)$ can be analytically calculated, see Eqs. (21) and (35) in Ref. [37] with $\mathbf{J} = Id - \mathbf{J}'$.

Another frequently used regularization scheme is L_1 -norm regularization of the couplings, corresponding to a Laplacian prior distribution, which gives to the following additive contribution to the log-likelihood:

$$\Delta L(J) = -\gamma B \sum_{i < j} |J_{ij}|. \quad (9)$$

L_1 -norm regularization favors zero instead of small couplings and produces sparse interaction graphs. In the following analysis we will see that the L_1 -norm regularization is less adequate than the two other schemes presented here for mean-field inference. No analytical expression exists for the optimal J . However, it can be found in a polynomial time using convex optimization techniques [38].

C. Case of the Potts model

The Potts model is a generalization of the Ising model in which each spin can take $q \geq 2$ values, hereafter called symbols, $a = 1, 2, \dots, q$. A mapping can be made onto the Ising model through the introduction of binary spins, $\sigma_i^a = 1$ if spin i carries symbol a , and $\sigma_i^a = 0$ otherwise. Any N Potts-spin configuration $\{a_i\}$ can be uniquely written as an $N \times q$ Ising-spin $\{\sigma_i^a\}$ configuration through this mapping. Reciprocally admissible Ising-spin configurations are such that the constraints

$$\sum_a \sigma_i^a = 1 \quad (10)$$

hold for all sites i . In other words, at each site i , there is one and only one Ising spin equal to 1, with the remaining $q - 1$ spins being equal to zero.

The Hamiltonian of the Potts model may then be recast as an Ising Hamiltonian,

$$\begin{aligned} H[\{a_i\}] &\equiv -\sum_i h_i(a_i) - \frac{1}{2} \sum_{i \neq j} J_{ij}(a_i, a_j) \\ &= -\sum_i h_i(a) \sigma_i^a - \frac{1}{2} \sum_{i \neq j} J_{ij}(a, b) \sigma_i^a \sigma_j^b. \end{aligned} \quad (11)$$

Due to the constraints (10) the local fields and the coupling parameters $h_i(a), J_{ij}(a, b)$ can be concomitantly changed, without affecting the Hamiltonian. One can check that H is invariant under the change $J_{ij}(a, b) \rightarrow J_{ij}(a, b) + K_{ij}(b), h_i(b) \rightarrow h_i(b) - \sum_{j(\neq i)} K_{ij}(b)$ for arbitrary K . This invariance is called gauge invariance. In the following, we will restrict to one particular gauge, called zero-sum gauge, where for every pair of sites (i, j) , the sums of couplings along each column and each row of the $q \times q$ coupling matrix $J_{ij}(a, b)$ vanish.

The MF inference procedure presented in Sec. II A can be readily applied to the Ising representation of the Potts model. We obtain that the inferred coupling matrix, $J_{ij}^{\text{MF}}(a, b)$ is the pseudoinverse of the correlation matrix, $c_{ij}(a, b) = \langle \sigma_i^a \sigma_j^b \rangle - \langle \sigma_i^a \rangle \langle \sigma_j^b \rangle$. The pseudoinverse must be considered here in order to invert the correlation matrix in the $N(q - 1)$ -dimensional subspace orthogonal to the constraints (10).

III. EFFECTS OF REGULARIZATION SCHEMES: A TOY-MODEL ANALYSIS

In this section we consider the effects of regularization through the analysis of very simple models with two or three spins only. For simplicity we will focus on the Ising model when exploring generic properties of regularized MF inference. Results particular to Potts models are discussed in Secs. III E 1 and III E 2.

In the following, we first infer the coupling between the spins using various regularization schemes and compare the outcome to the true value when sampling is perfect. We find that the MF approximation introduces errors into the inferred couplings, which can be corrected by strong regularization terms. In the Potts model case, the optimal regularization strength is found to depend on the number of symbols. We then investigate the effect of poor sampling on the pseudocount performance. Finally, we consider the case of nonuniform couplings. Here we find that uniform L_2 -norm regularization is suboptimal when couplings are strongly heterogeneous, and in this case the pseudocount offers superior results compared to L_2 -norm regularization.

A. Ising models: Case of perfect sampling

1. Mean-field overestimates strong couplings

Let us consider two spins $\sigma_1, \sigma_2 = \pm 1$ which are coupled through an interaction $J_{12} = J$, with no local fields ($h_1 = h_2 = 0$). The magnetizations $\langle \sigma_1 \rangle, \langle \sigma_2 \rangle$ vanish, and the correlation is $c_{12} = \langle \sigma_1 \sigma_2 \rangle = \tanh J$. The 2×2 correlation matrix therefore has elements unity on the diagonal and c_{12} on the off-diagonal. Using (4) we obtain the inferred coupling within MF approximation,

$$J^{\text{MF}}(J) = \frac{c_{12}}{1 - c_{12}^2} = \frac{\tanh J}{1 - \tanh^2 J} = \frac{1}{2} \sinh 2J. \quad (12)$$

The MF prediction is plotted versus the true coupling J in Fig. 1. The inferred coupling is in good agreement with the true value only for small couplings ($|J| < 1$) and diverges very quickly, $|J^{\text{MF}}| \sim \frac{1}{4} e^{2|J|}$, as J increases. The MF approximation drastically overestimates strong couplings.

2. Strong regularizations with pseudocount or L_2 correct for MF errors

We start with the pseudocount regularization of intensity α , see (6). The inferred coupling is given by

$$J^{\text{PC}}(J, \alpha) = \frac{(1 - \alpha) \tanh J}{1 - (1 - \alpha)^2 \tanh^2 J}. \quad (13)$$

The MF prediction with pseudocount (PC) regularization, $J^{\text{PC}}(J, \alpha)$, is plotted versus the true coupling J in Fig. 2(a) for various values of α . Unlike pure MF inference ($\alpha = 0$), J^{PC} saturates for very large couplings J to a finite value, $J^{\text{PC}}(\infty, \alpha) = (1 - \alpha)/(\alpha(2 - \alpha))$. For intermediate couplings $|J| < J^{\text{PC}}(\infty, \alpha)$ we find that the agreement between J^{PC} and J is remarkably good for $\alpha \simeq 0.2$. Hence the pseudocount with a large intensity (compared to the inverse of the number of data, which is infinite here since sampling is perfect) can correct for the dramatic overestimation of large couplings by the mean-field approximation. Couplings weaker in absolute




FIG. 1. Coupling J^{MF} inferred with the mean-field approximation (12) overestimates the true coupling value J (dotted line) in a system of two spins when J is large (in absolute value).

value than the saturation value $J^{\text{PC}}(\infty, 0.2)$ are precisely inferred, while larger couplings cannot be distinguished from $J^{\text{PC}}(\infty, \alpha)$.

A similar correction can be achieved with the L_2 regularization. Adding the penalty term (8) to the mean-field expression for the log-likelihood (3) and maximizing over the 2×2 -coupling matrix J we find that the off-diagonal coupling, $J^{L_2}(J, \gamma)$, is the root of the following implicit equation:

$$\tanh J = J^{L_2}(J, \gamma) \left[\gamma + \frac{2}{1 + \sqrt{1 + 4 J^{L_2}(J, \gamma)^2}} \right]. \quad (14)$$

Figure 2(b) compares the outcome to the true coupling for various values of γ . As in the pseudocount case J^{L_2} saturates for very large couplings J to a finite and γ -dependent value, approximately equal to $1/\sqrt{2\gamma}$ for small γ . For intermediate couplings we find that the agreement between J^{L_2} and J is very good for $\gamma \simeq 0.13$. Hence L_2 regularization with a large intensity (again, compared to the inverse of the number of data, which is infinite here since sampling is perfect) avoids the divergence at large couplings introduced by the mean-field approximation, while being accurate for small coupling values. We observe that the coupling saturation value for the optimal γ with the L_2 -norm ($\simeq 1.8$) is, however, smaller than with the pseudocount regularization ($\simeq 2.2$), compare Figs. 2(a) and 2(b); hence, the pseudocount regularization offers an accurate inference over a slightly wider range of coupling values.

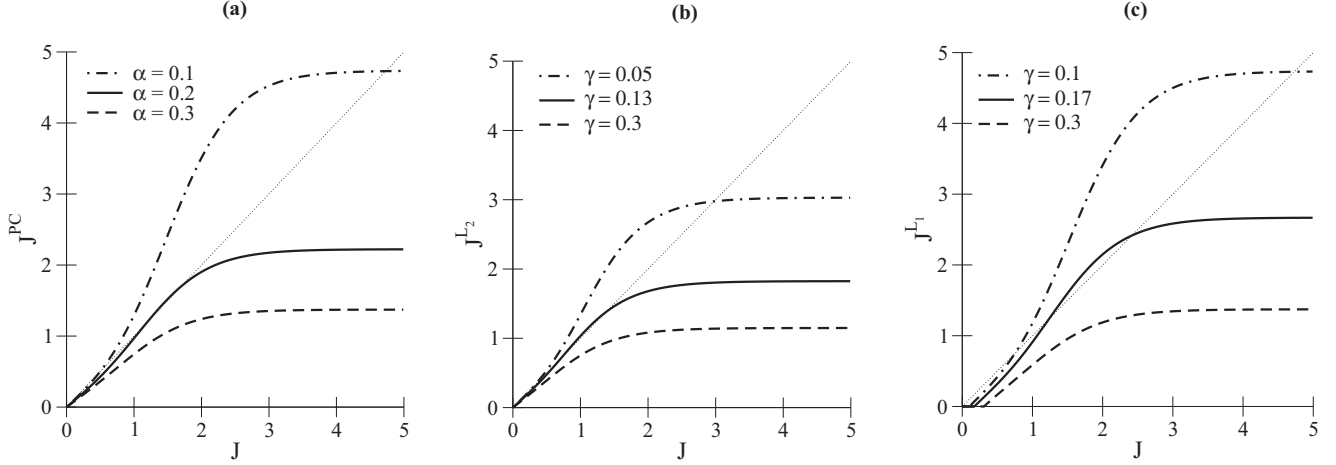


FIG. 2. Coupling inferred with the mean-field approximation with (a) pseudocount of intensity α , J^{PC} ; (b) with L_2 -norm regularization of intensity γ , J^{L_2} ; (c) with the L_1 -norm regularization of intensity γ , J^{L_1} , vs true coupling value J (dotted line) for a system of two spins. Due to the symmetry $J \rightarrow -J$ only the positive quadrant is shown.

3. L_1 -norm regularization is less accurate than pseudocount and L_2

We now consider the L_1 -norm regularization of intensity γ , obtained by adding the penalty term (9) to the log-likelihood (3). An immediate calculation shows that the total log-likelihood is maximized by the coupling value

$$J^{L_1}(J, \gamma) = \begin{cases} \frac{\tanh J - \gamma}{1 - (\tanh J - \gamma)^2} & \text{if } \gamma \leq \tanh J, \\ 0 & \text{if } \gamma \geq \tanh J. \end{cases} \quad (15)$$

The expression above for holds for positive J ; for negative J we have $J^{L_1}(J, \gamma) = -J^{L_1}(|J|, \gamma)$. The inferred coupling $J^{L_1}(J, \gamma)$ is plotted versus the true coupling J in Fig. 2(c) for various values of γ . As with pseudocount and L_2 J^{L_1} saturates for very large couplings J to a finite value, $J^{L_1}(\infty, \alpha) = (1 - \gamma)/(\gamma(2 - \gamma))$. The novelty is that the inferred coupling vanishes for small J . Overall, for intermediate couplings $|J| < J^{L_1}(\infty, \alpha)$, the agreement between J^{L_1} and J is less precise than what can be achieved with the pseudocount and L_2 regularization when, respectively, α and γ are properly chosen. In Ref. [33] the inferred Ising couplings are compared with Gaussian couplings with L_1 - and L_2 -norms on neural data coming from multielectrode recordings. Also on these data sets, in agreement with previous findings, L_2 performs better over L_1 and large regularization strengths are needed to improve the inference.

B. Performance of pseudocount as a function of the sampling quality

We now focus on the pseudocount scheme. We assume that a number, say, B , of configurations of the two spins are drawn at random, from the Ising model measure,

$$P_J(\sigma_1, \sigma_2) = \frac{e^{J\sigma_1\sigma_2}}{2(e^J + e^{-J})}. \quad (16)$$

The magnetizations $p_1 \equiv \langle \sigma_1 \rangle$ and $p_2 \equiv \langle \sigma_2 \rangle$, and the correlation $p_{12} \equiv \langle \sigma_1 \sigma_2 \rangle$ are then computed as empirical averages over the data. The joint probability density for these three

quantities is given by

$$\begin{aligned} \rho(p_1, p_2, p_{12}; B, J) &= \sum_{0 \leq B_{++}, B_{+-}, B_{-+}, B_{--} \leq B} \binom{B}{B_{++}, B_{+-}, B_{-+}, B_{--}} \\ &\times \frac{e^{J(B_{++} - B_{+-} - B_{-+} + B_{--})}}{(2(e^J + e^{-J}))^B} \\ &\times \delta\left(p_1 - \frac{B_{++} + B_{+-} - B_{-+} - B_{--}}{B}\right) \\ &\times \delta\left(p_2 - \frac{B_{++} - B_{+-} + B_{-+} - B_{--}}{B}\right) \\ &\times \delta\left(p_{12} - \frac{B_{++} - B_{+-} - B_{-+} + B_{--}}{B}\right). \end{aligned} \quad (17)$$

We then define the average squared relative error on the inferred coupling as

$$\begin{aligned} \epsilon(B, J, \alpha) &= \int_0^1 dp_1 dp_2 dp_{12} \rho(p_1, p_2, p_{12}; B, J) \\ &\times \left[\frac{J^{\text{PC}}(p_1, p_2, p_{12}, \alpha)}{J} - 1 \right]^2, \end{aligned} \quad (18)$$

where

$$J^{\text{PC}}(p_1, p_2, c_{12}, \alpha) = J^{\text{MF}}[(1 - \alpha)p_1, (1 - \alpha)p_2, (1 - \alpha)c_{12} + \alpha(1 - \alpha)p_1 p_2], \quad (19)$$

$$J^{\text{MF}}(p_1, p_2, c_{12}) = \frac{c_{12}}{(1 - p_1^2)(1 - p_2^2) - c_{12}^2},$$

are the PC and MF predictions for the coupling given the magnetizations p_1, p_2 and the connected correlation $c_{12} = p_{12} - p_1 p_2$. These expressions extend formulas (12) and (13) to the case of nonzero magnetizations.

Peculiar samples for which J^{PC} and J^{MF} diverge, e.g., such that $p_1 = p_2 = 0$ and $p_{12} = -1$, may happen with nonzero (albeit exponentially small in B) probabilities. To get a well-defined and finite expression for the squared error in (18) we

replace the term squared within parenthesis with the minimum of this term and one. The latter constant is arbitrary; any other

choice would lead to $e^{-O(B)}$ changes to the error. With our choice, the relative square error cannot be larger than unity by construction.

We show in Fig. 3(a) the relative square error as a function of B in the absence of pseudocount ($\alpha = 0$). We observe that the error is small for small couplings J if the number B of configurations is large and increases rapidly with J , in agreement with the findings of Fig. 2(a). For the optimal value of the pseudocount strength, $\alpha = 0.2$, the relative error is a decreasing function of B and saturates to a small value for all couplings, see Fig. 3(b); larger couplings produce larger correlations and are easier to recover (require a smaller number of configurations) than smaller interactions. If the pseudocount strength is too large, e.g., $\alpha = 0.4$, the error saturates to a finite and larger value, as shown in Fig. 3(c). We conclude that the presence of a pseudocount with fixed strength compensates the errors due to the mean-field approximation, even for a small number of sampled configurations.

C. Network effects on zero interactions

An important question is to estimate network effects on the performance of mean-field inference. In order to study this point in a simple case we consider a system of three spins, with zero external fields and couplings $J_{12} = J_{13} \equiv \hat{J}_0$ and $J_{23} \equiv \hat{J}_1 \neq \hat{J}_0$. We assume that sampling is perfect and derive the spin-spin correlations $p_{ij} \equiv \langle \sigma_i \sigma_j \rangle$, with $1 \leq i < j \leq 3$. We obtain

$$p_{12} = p_{13} = \frac{e^{2\hat{J}_0} \sinh(2\hat{J}_1)}{e^{2\hat{J}_0} \cosh(2\hat{J}_0) + 1}, \quad (20)$$

$$p_{23} = \frac{e^{2\hat{J}_1} \cosh(2\hat{J}_0) - 1}{e^{2\hat{J}_1} \cosh(2\hat{J}_0) + 1},$$

while all three magnetizations vanish. We infer the three coupling values with the MF approximation by maximizing the log-likelihood $L^{\text{MF}}(c|J)$, where J is given by (22) and the regularized correlation matrix is

$$c = \begin{bmatrix} 1 & (1-\alpha)p_{12} & (1-\alpha)p_{13} \\ (1-\alpha)p_{12} & 1 & (1-\alpha)p_{13} \\ (1-\alpha)p_{12} & (1-\alpha)p_{13} & 1 \end{bmatrix}. \quad (21)$$

We specialize hereafter to the case $\hat{J}_1 = 0$ and ask whether the mean-field inference procedure is able to detect that this coupling vanishes; nonzero values for \hat{J}_1 are considered in Sec. III D. Results are reported in Fig. 4. The inferred coupling is correctly found to vanish for $\alpha \rightarrow 1$ (all connected correlations vanish) and $\alpha \rightarrow 0$ (as sampling is perfect here). However, for intermediate values of α , the inferred coupling shows a nonmonotonic behavior and may reach very large values, depending on the value of J_0 . The maximum is located in $\alpha^* \simeq 4/\sqrt{3} \times \exp(-2\hat{J}_0)$, $J_1^* \simeq (2 - \sqrt{3})/8 \times \exp(2\hat{J}_0)$ at the leading order in \hat{J}_0 . Hence, at small (but nonvanishing) pseudocount, mean field assigns a large value to a coupling whose true value is zero. The origin of the inferred nonzero coupling can be easily understood. In the absence of coupling between sites 1 and 3 the correlation p_{13} is equal to the product of the correlations p_{12} and p_{13} . After the pseudocount is introduced, all correlations are multiplied by a factor $(1 - \alpha)$, and the correlation p_{13} is now larger (by a factor $1 - \alpha$) than the

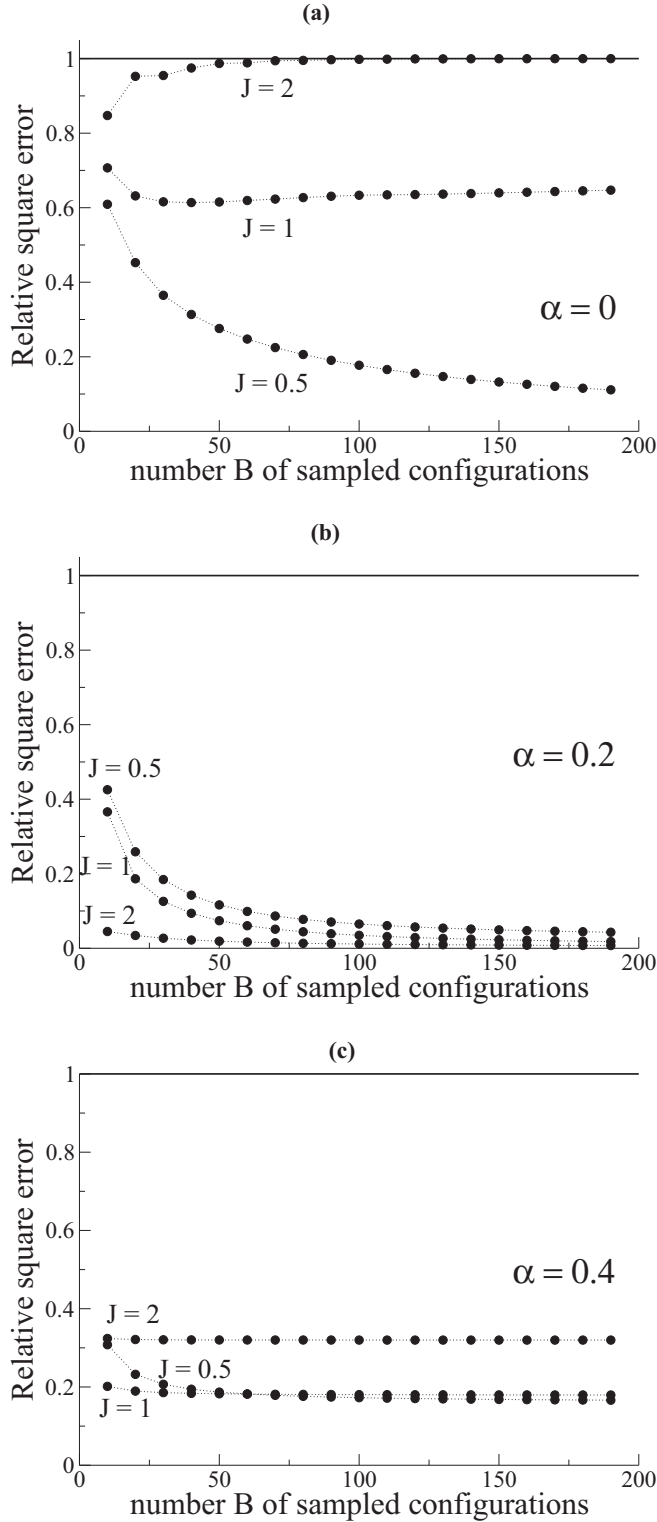


FIG. 3. Relative squared error $\epsilon(B, J, \alpha)$ (18) between the inferred coupling with pseudocount of intensity α and the true coupling for a system of two spins as a function of the number B of sampled configurations. Each panel corresponds to one value of α and each curve to one value of J . Pseudocount decreases the error on inferred couplings even when the couplings are weak, or sampling is large, but regularization that is too large leads to worse performance.

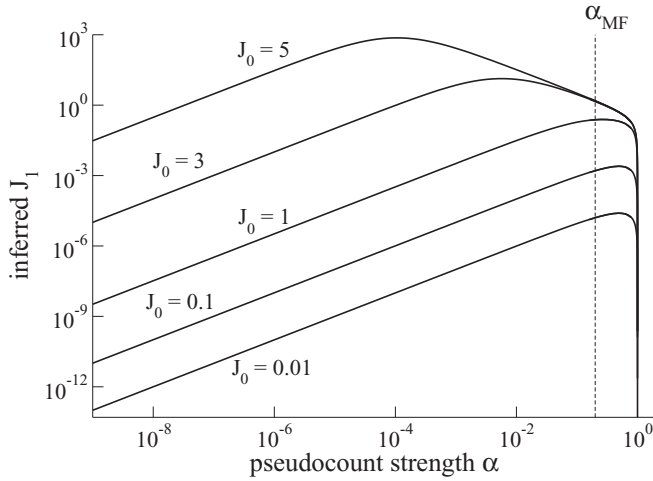


FIG. 4. Mean-field prediction for the coupling \hat{J}_1 as a function of the pseudocount strength α (log-log scale). The true interaction network connects three spins, with two couplings equal to \hat{J}_0 (values shown) and one vanishing coupling \hat{J}_1 . The vertical dashed line locates $\alpha^{\text{MF}} \simeq 0.2$.

product of the two other regularized correlations. As a result, a fictitious positive coupling is inferred to account for this excess in correlation. This artifact is cured by the introduction of a large pseudocount. For $\alpha = \alpha^{\text{MF}}$ the inferred J_1 remains small whatever the value of \hat{J}_0 (Fig. 4).

D. Efficiency of uniform regularization for nonuniform couplings

We now ask whether a regularization scheme with uniform penalties, i.e., equal for all pairs (i, j) is appropriate in the case of a network of interactions with heterogeneous interactions. We start to investigate this issue with the L_2 -norm regularization and the 3-spin model described above. We infer the three coupling values with MF by maximizing the

log-likelihood $L^{\text{MF}}(c|J)$, where

$$J = \begin{pmatrix} K_0 & J_0 & J_0 \\ J_0 & K_1 & J_1 \\ J_0 & J_1 & K_1 \end{pmatrix}, \quad c = \begin{pmatrix} 1 & p_{12} & p_{12} \\ p_{12} & 1 & p_{13} \\ p_{12} & p_{13} & 1 \end{pmatrix}, \quad (22)$$

with an additive L_2 -penalty term given by (8). In the following, rather than fixing the values for \hat{J}_0, \hat{J}_1 and calculating the inferred couplings J_0, J_1 , we do the opposite. The reason is that the maximization equations are complicated implicit equations over J_0, J_1 for given \hat{J}_0, \hat{J}_1 and are simpler to solve for \hat{J}_0, \hat{J}_1 given the values of J_0, J_1 .

We show in Fig. 5(a) the relative squared error between the true and inferred values for the couplings, as a function of J_1 for a fixed $J_0 = 1$. The value of the penalty is chosen to be $\gamma = 0.13$, see Fig. 2(b). We observe that the relative squared errors are small when J_1 and J_0 are close to each other as expected. However, when J_1 departs from the value of the other couplings, J_0 , the relative error on J_1 becomes large. We conclude that imposing uniform L_2 penalties is not optimal for nonuniform couplings and can lead to substantial errors in the inferred couplings.

We now consider the pseudocount regularization. Couplings are computed by inverting the correlation matrix (21). We show in Fig. 5(b) the relative squared error between the true and inferred values for the couplings, as a function of J_1 for a fixed $J_0 = 1$. The PC strength is chosen to be $\alpha = 0.2$, see Fig. 2(a). We observe that for a large range of values of J_1 compared to J_0 the errors on both couplings remains small. The accuracy is better than with the L_2 -norm regularization.

E. Potts model: Case of perfect sampling

1. Homogeneous case

We now consider a two-spin Potts model and denote by $q \geq 2$ the number of spin symbols. The model has Hamiltonian $H(a_1, a_2) = -J_0 \delta_{a_1, a_2}$, where a_1 and a_2 are the symbols corresponding to, respectively, spins 1 and 2. Each spin is equally likely to be in any of the q symbols, and we will hereafter refer to this model as the homogeneous Potts model.




FIG. 5. Relative squared error over the couplings with (a) uniform L_2 -norm penalty and (b) uniform pseudocount for a system of three spins, see text. The dotted lines in (b) reproduce the relative squared errors found for the L_2 -norm regularization in (a). Note the change in the vertical axis scale.

The Ising model is recovered when $q = 2$; note that, as the difference between the energies of the equal-spin and different-spin configurations is J in the Potts model, the corresponding Ising model coupling is $J/2$. In the following we will express the $q \times q$ coupling matrix \mathbf{J} in a specific gauge, in which the sum of couplings over a row or a column of the matrix vanish

$$\mathbf{J} = \begin{pmatrix} J_A & J_B & J_B & \dots & J_B \\ J_B & J_A & J_B & \dots & J_B \\ J_B & J_B & J_A & \dots & J_B \\ \dots & \dots & \dots & \dots & J_B \\ J_B & J_B & J_B & J_B & J_A \end{pmatrix} \quad (23)$$

$$\text{with } J_A = \frac{q-1}{q} J_0, \quad J_B = -\frac{J_0}{q}.$$

The two-spin correlations are

$$\begin{aligned} \langle \delta_{\sigma_1,a} \delta_{\sigma_1,a} \rangle &= \langle \delta_{\sigma_2,a} \delta_{\sigma_2,a} \rangle = \frac{1}{q}, \\ \langle \delta_{\sigma_1,a} \delta_{\sigma_1,b} \rangle &= \langle \delta_{\sigma_2,a} \delta_{\sigma_2,b} \rangle = 0, \\ \langle \delta_{\sigma_1,a} \delta_{\sigma_2,a} \rangle &\equiv p_{aa} = \frac{e^{J_0}}{q(e^{J_0} + q - 1)}, \\ \langle \delta_{\sigma_1,a} \delta_{\sigma_2,b} \rangle &\equiv p_{ab} = \frac{1}{q(e^{J_0} + q - 1)}, \end{aligned} \quad (24)$$

where a and b denote different spin symbols. The $2q \times 2q$ entries of the connected correlation matrix \mathbf{C} can be computed from those values, after subtraction of $1/q^2$.

Note that the sum of the elements of \mathbf{C} over a line or a column is equal to zero. This property of \mathbf{C} reflects the fact that each spin takes one symbol value. To obtain the mean-field prediction for the coupling, \mathbf{J}^{MF} , we consider minus the pseudoinverse of \mathbf{C} . Again, the sum of the elements of \mathbf{J}^{MF} over a line or a column is equal to zero. The inferred Potts couplings J_A^{MF} and J_B^{MF} correspond to, respectively, the diagonal and off-diagonal entries of the off-diagonal $q \times q$ blocks of \mathbf{J}^{MF} . Some simple algebra gives

$$\begin{aligned} J_A^{\text{MF}}(J_0, q) &= \frac{q(q-1)(p_{aa} - p_{ab})}{1 - q^2(p_{aa} - p_{ab})^2}, \\ J_B^{\text{MF}}(J_0, q) &= -\frac{q(p_{aa} - p_{ab})}{1 - q^2(p_{aa} - p_{ab})^2}, \end{aligned} \quad (25)$$

where the correlations p_{aa} and p_{ab} are given in (24). The MF inferred couplings are shown as functions of J_0 for different q in Fig. 6(a). For each value of q there are two branches corresponding to J_A and J_B . The upper branch J_A coincides with the lower branch J_B after rescaling of J and J^{MF} by the factor $-1/(q-1)$. As in the Ising case, the MF prediction is quantitatively correct for weak couplings but strongly overestimates the right coupling value for large J (in absolute value). Contrary to the Ising case there is an asymmetry between the positive and negative values of J (for $q \geq 2$) along each branch J_A or J_B .

In the presence of a pseudocount of strength α the difference of the correlations $p_{aa} - p_{ab}$ is multiplied by $(1 - \alpha)$. We can use again formula (25) to obtain the corresponding couplings, which we denote by J_A^{PC} and J_B^{PC} . Results for $q = 5$ symbols and three values of the pseudocount, ranging from $\alpha = 0.25$ to 0.55, are shown in Fig. 6(b). As in the Ising case we find that the inferred coupling J^{PC} saturates to a finite value




FIG. 6. Homogeneous q -symbol Potts model. Couplings inferred with the mean-field approximation, $J_A^{\text{MF}}(J_0, q)$ (diagonal) and $J_B^{\text{MF}}(J_0, q)$ (off-diagonal), vs true values coupling J_A and J_B in (23). Curves were obtained through a parametric representation with J_0 running from -4 to 4 ; longer stretches would be obtained by increasing the range of values for J_0 . The dashed line represents the $J^{\text{MF}} = J$ curve. (a) No pseudocount. (b) With a pseudocount of strength α (see values in the figure) and for $q = 5$ symbols.

when $J \rightarrow \pm\infty$. There is an optimal value of the pseudocount strength α such that the inferred and true coupling values are close to one another for positive J_A , and negative J_B . We observe, however, that for negative J_A and positive J_B , the pseudocount is not able to correct the errors produced by the MF approximation.

We may define the optimal pseudocount $\alpha^{\text{MF}}(q)$ as the largest value of α such that $J^{\text{PC}} = J$ has a nonzero solution (for positive J_A or, equivalently, for negative J_B). In other words, when $\alpha = \alpha^{\text{MF}}(q)$, the representative curve for the inferred coupling touches the $J^{\text{PC}} = J$ line tangentially [dotted line in Fig. 6(b)]. The value of $\alpha^{\text{MF}}(q)$ is shown as a function of q in Fig. 7 for q ranging between 2 and 20. We observe a monotonic increase of the optimal pseudocount with q , from $\alpha \simeq 0.2$ for $q = 2$ to $\simeq 0.74$ for $q = 20$. Our finding is in good agreement with empirical works on protein covariation, where the




FIG. 7. Optimal pseudocount strength $\alpha^{\text{MF}}(q)$ as a function of the number of Potts symbols, q . The dotted line serves as a guide to the eye.

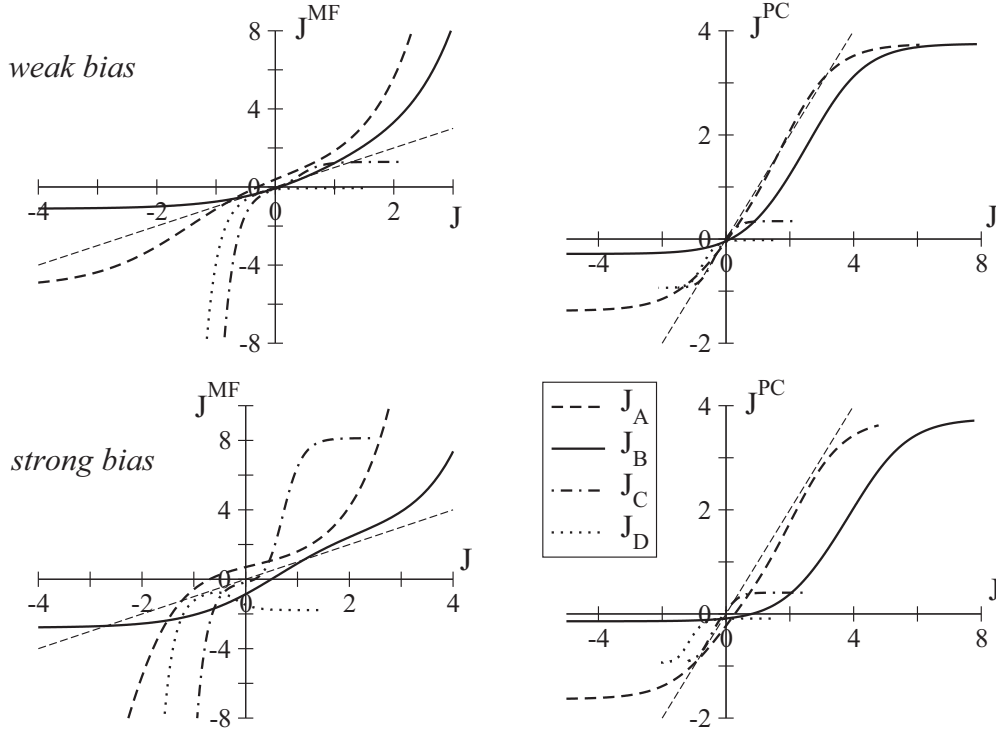


FIG. 8. Heterogeneous q -symbol Potts model. Couplings inferred with the mean-field approximation, $J_L^{\text{MF}}(J_0, J_1, q = 5)$ (left) and with a pseudocount $J_L^{\text{PC}}(J_0, J_1, q = 5, \alpha = 0.4)$ (right), vs true couplings $J_L(J_0, J_1, q)$; $L = A, B, C, D$ labels the four different branches. The value of the pseudocount strength $\alpha = 0.4$ has been chosen according to Fig. 7. Curves were obtained through a parametric representation with J_0 running from -8 to 8 ; longer stretches would be obtained by increasing the range of values for J_0 . Top: “Weak” bias ($J_1 = 3$); bottom: “strong” bias ($J_1 = 5$).

pseudocount is often taken to be 0.5 but may vary between 0.3 and 0.7 depending on the protein family under consideration.

2. Heterogeneous case

We will now study a simple heterogeneous case, in which one of the q symbols, say $a = 1$, has a larger frequency than the other symbols. To do so we consider the Hamiltonian $H(a_1, a_2) = -J_0 \delta_{a_1, a_2} - J_1 \delta_{a_1, 1} (1 - \delta_{a_2, 1})$ for the two spin

symbols on the two sites. The coupling parameter J_1 acts as a field along the $a = 1$ direction. The frequency p_i^1 of the first symbol, $a = 1$, is the same on both sites and is larger than $\frac{1}{q}$ if $J_1 > 0$, and smaller than $\frac{1}{q}$ if $J_1 < 0$. All other symbols $a = 2, \dots, q$ are equally likely with a frequency $\frac{1-p_i^1}{q-1}$.

In the zero-sum gauge, in which the sums of couplings along each row and each column vanish, the coupling matrix reads

$$\mathbf{J} = \begin{pmatrix} J_A & J_B & J_B & J_B & \dots & J_B \\ J_B & J_C & J_D & J_D & \dots & J_D \\ J_B & J_D & J_C & J_D & \dots & J_D \\ J_B & J_D & J_D & J_C & \dots & J_D \\ \dots & \dots & \dots & \dots & \dots & J_D \\ J_B & J_D & J_D & J_D & J_D & J_C \end{pmatrix} \quad \text{with} \quad \begin{cases} J_A = \frac{q-1}{q} J_0 - \frac{(q-1)^2}{q^2} J_1, \\ J_B = -\frac{J_0}{q} + \frac{(q-1)}{q^2} J_1, \\ J_C = \frac{q-1}{q} J_0 - \frac{J_1}{q^2}, \\ J_D = -\frac{J_0}{q} - \frac{J_1}{q^2}. \end{cases} \quad (26)$$

The off-diagonal $q \times q$ block of minus the pseudoinverse of the correlation matrix is the coupling matrix \mathbf{J}^{MF} within the MF approximation, which fulfills the same gauge condition as \mathbf{J} . We obtain four couplings $J_L^{\text{MF}}(J_0, J_1, q)$, which can be compared to the four couplings $J_L(J_0, J_1, q)$ defined in (26), with $L = A, B, C, D$. The homogeneous case, studied in the previous section, is recovered when $J_1 = 0$. In this case we have degenerate couplings: $J_A = J_C$ and $J_B = J_D$.

Figure 8 shows the couplings inferred with the MF approximation against their true values for the $q = 5$ -symbol Potts model, in a “weakly” biased case (top row, corresponding to $J_1 = 3$) and in a “strongly” biased case (bottom row, corresponding to $J_1 = 5$). Note that the terms “weak” and “strong” have no absolute meaning here, as the bias is not constant when J_0 varies. In the weak-bias case, we observe that the degeneracy between the couplings is lifted compared with the $q = 5$ curves in Fig. 6(a). As the bias gets stronger

(Fig. 8, bottom and left), the branches corresponding to J_B and J_D show markedly different behavior; for instance, the discrepancy between the inferred and true values of the couplings J_B may largely exceed the errors on the diagonal couplings J_A and J_C . Varying the value of q does not qualitatively affect the results above, with the exception that the off-diagonal couplings J_B, J_D become smaller as q grows, in agreement with Fig. 6(a).

The introduction of a large pseudocount corrects, to some extent, the errors resulting from the MF approximation, see Fig. 8, right panels. Two remarks can be made. First, the quality of the inference is better for the couplings corresponding to the symbols with high probability, here J_A and J_B , as $\sigma_1 = 1$ is the most frequent symbol for $J_1 > 0$, see (26). Second, when the bias increases, the quality of the inference does not decrease much for the couplings associated to frequent symbols (J_A, J_B) but strongly deteriorates for the other couplings (J_C, J_D). As a consequence, the inferred couplings occupy a larger part of the $J^{\text{PC}} < J, J > 0$ and $J^{\text{PC}} > J, J < 0$ portions of the (J, J^{PC}) plane.

IV. NUMERICAL SIMULATIONS FOR LARGE SYSTEMS

To better understand mean-field inference on larger, more realistic data sets, we have tested the accuracy of the interaction graph recovered by pseudocount and L_2 -regularized MF inference for a variety of Ising (Sec. IV A) and Potts models (Sec. IV B).

A. Results for the Ising model

We have tested Ising models with different network topologies and random distributions of the couplings and using differing numbers of samples to compute the correlations (Fig. 9). All the simulations reported in this section were performed with spins taking the values 0, 1.

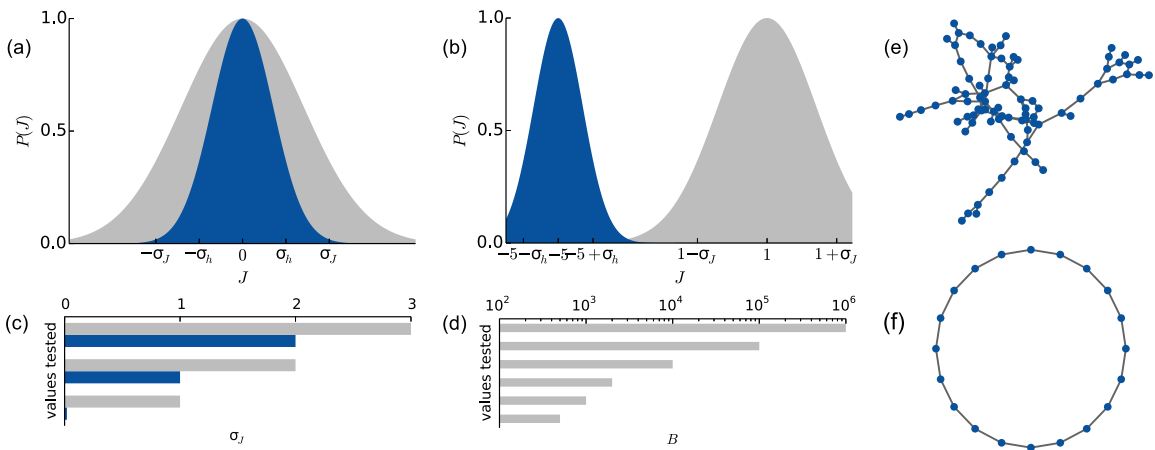


FIG. 9. (Color online) Representation of tested model parameters. Fields and nonzero couplings were selected according to model (a), all fields and couplings normally distributed with means $\bar{h} = 0$, $\bar{J} = 0$ and standard deviations σ_h , σ_J , respectively, or model (b), strong negative fields and couplings normally distributed with means $\bar{h} = -5$, $\bar{J} = 1$ and standard deviations σ_h , σ_J . Distributions for fields are shaded dark, values for couplings are light. (c) For each model a range of σ_h (dark) and σ_J (light) was tested. (d) Correlations used for the MF inference were computed using B samples from a Monte Carlo simulation of the model, with B tested over a range from 500 to 10^6 . All permutations of the above parameters were considered for each choice of the network topology: (e) Erdős-Rényi graph where edges are kept with probability $p = 2/N$ or $p = 4/N$ and (f) 1D lattice with nearest-neighbor couplings. In all cases we take the system size $N = 100$.

The accuracy of the MF inference was quantified in two ways. First, we considered the standard root-mean-square (rms) error between the inferred couplings J^{inf} and those in the true model J^{true} ,

$$\Delta_J = \sqrt{\frac{2}{N(N-1)} \sum_{i < j} (J_{ij}^{\text{inf}} - J_{ij}^{\text{true}})^2}, \quad (27)$$

where N is the system size. The rms error captures the absolute difference between the true and inferred couplings but is unable to clearly distinguish whether the *relative* ordering of the couplings has been correctly inferred. This limitation is problematic since many practical applications, such as the prediction of protein contacts from MF inference on sequence data [4,5,34], rely on proper rank ordering of the inferred couplings rather than their absolute magnitude.

Information about the correct rank ordering can be determined from the rank correlation between the true and inferred couplings. To do this, we assigned each true coupling a rank according to its absolute value, with the largest coupling assigned rank 1, and the smallest rank n_{NZ} , where n_{NZ} is the total number of nonzero couplings. All couplings exactly equal to zero are simply assigned rank $n_{\text{NZ}} + 1$. We then computed the Pearson correlation ρ_J between the rank of the top n_{NZ} inferred couplings and their true counterparts, measuring how well the ordering of the top inferred couplings matches the true ordering in the underlying model. Letting $\{i_k, j_k\}$, with $k = 1, \dots, n_{\text{NZ}}$, denote the pair indices of the largest n_{NZ} inferred couplings, this is

$$\rho_J = \frac{1}{\sigma_{r(J^{\text{true}})} \sigma_{r(J^{\text{inf}})}} \sum_{k=1}^{n_{\text{NZ}}} \left(k - \frac{n_{\text{NZ}} + 1}{2} \right) [r(J_{i_k j_k}^{\text{true}}) - \bar{r}^{\text{true}}]. \quad (28)$$




FIG. 10. (Color online) Typical example trajectory of inference quality as a function of the regularization strength for a single randomly chosen set of couplings and fields. At larger regularization strengths the inferred couplings are similar in magnitude to the true couplings, but at lower regularization strengths the value of the strongest inferred couplings begins to diverge. The true model is a 1D spin chain, with zero fields and couplings normally distributed with mean zero and standard deviation $\sigma_J = 3$. $B = 10^5$ Monte Carlo samples were used to compute the correlations used for the MF inference. (a) rms error Δ_J (27) (circles, left axis), rank correlation ρ_J (28) (squares, right axis), and fraction of nonzero couplings recovered R (triangles, right axis), as a function of the pseudocount α . [(b1), (b2), and (b3)] Comparison of true and inferred couplings for the three pseudocount values $\alpha_1, \alpha_2, \alpha_3$ shown in panel (a) by dashed vertical lines. Dashed lines mark the $J^{\text{inf}} = J^{\text{true}}$ lines. The N largest (in absolute value) inferred couplings are denoted by open circles and others are denoted by squares; the cut-off number N is chosen to match the number of nonzero couplings on the true interaction graph. Note the change of vertical scale between panels. The performance is maximal for $\alpha = \alpha_1 \approx 0.2$. This value of the pseudocount agrees well with the optimal pseudocount for the Ising model $\alpha^{\text{MF}} = 0.2$ shown in Fig. 7. The rank correlation is worse for $\alpha = \alpha_2 \approx 0.002$, as the inference procedure assigns large values to couplings equal to zero on the true interaction network, see Sec. III C. For $\alpha = \alpha_3 = 2 \cdot 10^{-5} \simeq \alpha_B$ those fake couplings have essentially disappeared; the inferred couplings have much larger values than the true couplings but are correctly ordered (large rank correlation).

Here $r(J)$ is the rank of coupling J , \bar{r} the average rank, and $\sigma_{r(J)}$ the standard deviation of the ranks. Note that since the inferred couplings are ranked from 1 to n_{NZ} , their average rank $\bar{r}^{\text{inf}} = (n_{\text{NZ}} + 1)/2$. The fraction R of true nonzero couplings included within the top n_{NZ} inferred couplings was also computed. In this way, we can assess how well the inferred couplings recover real couplings from the underlying model and the degree to which their relative ordering is preserved.

1. Regularization improves the quality of mean-field inference: An example on a 1D lattice

As a typical example, in Fig. 10 we show the performance of the mean-field inference as a function of pseudocount for a model system with nearest-neighbor interactions on a 1D lattice and with a good sampling on $B = 10^5$ configurations. As the pseudocount is lowered from its maximum at $\alpha = 1$,

rms error Δ_J (27) decreases and the rank correlation ρ_J (28) improves until a peak is reached at $\alpha \simeq 0.2$ [Fig. 10(a)], in excellent agreement with the optimal value of the pseudocount strength necessary to correct the MF approximation, $\alpha^{\text{MF}} = 0.2$, theoretically found for the Ising system with two spins (Fig. 7). At $\alpha \simeq \alpha^{\text{MF}}$ the largest true couplings are recovered well, and the inferred couplings are similar to the true ones in magnitude [Fig. 10(b1)]. At lower values of the pseudocount the largest inferred couplings are much larger than their true counterparts, and couplings that are zero in the true model are more likely to be inferred as large [Fig. 10(b2)]. This phenomenon agrees with the results of Sec. III C and considerably lowers the rank correlation. As the pseudocount strength is decreased further to $\alpha_3 \simeq \alpha^B = 1/B$, we observe the rank correlation ρ_J and fraction of nonzero couplings recovered R reach a high value again, as the couplings whose true values are zero are correctly inferred with small values [Fig. 10(b3)]. Though the ordering is good, the inferred couplings have

magnitude much larger than their true counterparts, which produces large rms error Δ_J compared to the $\alpha \simeq \alpha^{\text{MF}}$ case [Fig. 10(a)].

2. Scaling of the optimal regularization strength with sampling depth and coupling strength on random graphs

Analysis presented in Sec. III suggests an optimal value for the regularization strength needed to correct for errors introduced by the MF approximation, which is independent of the amount of data. In contrast, in a Bayesian framework the regularization strength should scale as $\alpha^B \sim 1/B$ as the sampling depth is increased, where B is the number of independent samples, as described in Sec. II B. Our simulation results agree with the former picture: The optimal regularization strength α minimizing the rms error between the true and inferred couplings is nearly independent of the sampling depth B , even when the latter is varied over four orders of magnitude. This is demonstrated in Fig. 11 for a system with underlying interactions given by an Erdős-Rényi graph, but the result is completely general, holding for every model we have considered.

Independent of the value of B we find that the value of the pseudocount α which gives the best performance (smaller rms error Δ_J , largest rank correlation ρ_J , and fraction R of recovered nonzero couplings) is typically of the order of $\alpha^{\text{MF}} = 0.2$, as computed in Sec. III E 1 for Ising spins. The rank correlation (28) [Fig. 11(a), middle] and the fraction of nonzero couplings recovered [Fig. 11(a), bottom] reach similar values with a small regularization strength $\simeq \alpha^B$ for very good sampling. For intermediate sampling depths the performances show a saturation for pseudocount values $\alpha < \alpha^B$, see [Fig. 11(a), middle]. The nonmonotonic behavior of ρ_J and R is similar to the one reported in Fig. 10 and is due to

the large values erroneously assigned to zero couplings by the mean-field inference procedure (Sec. III C).

It is important to stress that for the pseudocount the value of the optimal regularization strength is also largely independent of the strength of the interactions (Fig. 12). In the case of very weak interactions and good sampling, MF inference is almost exact: The nonmonotonic behavior observed in Fig. 10 is not present, and performance with or without pseudocount is comparable.

3. Comparison of pseudocount and L_2 regularization performance on random graphs

We report the performances of L_2 regularization in Fig. 11(b) and Fig. 12(b). In Fig. 13 we compare the L_2 and pseudocount scheme for one example system on a relatively well-sampled Erdős-Rényi graph. Consistent with the analytical arguments presented in Sec. III for small systems, regularization improves the quality of couplings inferred via MF in large systems for a wide variety of underlying models. In particular, the rms error Δ_J (27) is improved by regularization unless couplings are weak or sampling is very good. In the latter case, the rms error may slightly increase with the pseudocount at intermediate values of α [Fig. 10(a)], due to the presence of large inferred couplings, whose true values are equal to zero [Fig. 10(b2)].

The pseudocount can substantially improve rank correlation ρ_J (28) as well as the fraction of true nonzero couplings recovered R , particularly when sampling is poor. L_2 -norm regularization has some effect on the rank correlation and fraction of nonzero couplings recovered but tends to improve them only slightly compared to couplings inferred via MF with no regularization. Moreover in Fig. 11(b) (middle and bottom panels) we show that with the L_2 -norm the value of the rank correlation and the fraction of nonzero couplings depends,

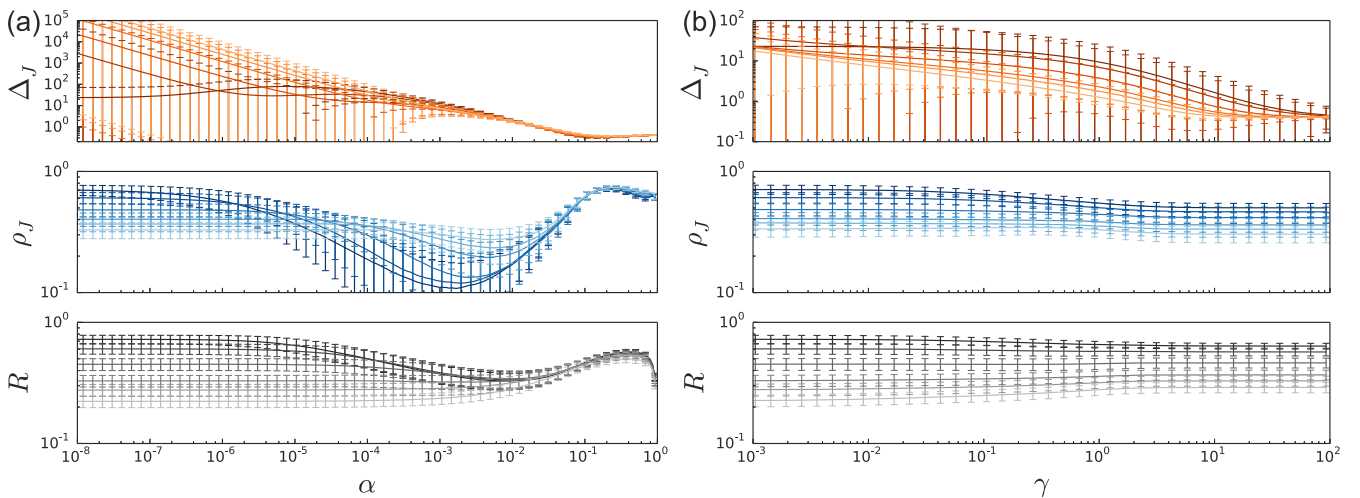


FIG. 11. (Color online) Optimal values of the regularization strength are only weakly affected by sampling depth, even when varied over the full range from $B = 500$ (lightest) to $B = 10^6$ (darkest). Trajectory of the rms error Δ_J (27) (top), rank correlation ρ_J (28) (middle), and fraction of nonzero couplings recovered R (bottom) as the pseudocount α (a) and L_2 -norm regularization strength γ (b) is varied, averaged over 10^3 sets of random couplings. Relevant values of the pseudocount are $\alpha = \alpha^{\text{MF}} = 0.2$ and $\alpha^B = 1/B$, roughly where ρ_J and R begin to plateau for the pseudocount. Each random set of interactions has all fields set to zero. The coupling network is an Erdős-Rényi graph ($N = 100$ nodes) where edges are kept with probability $p = 2/N$. Nonzero couplings are normally distributed with mean zero and standard deviation $\sigma_J = 3$. Bars denote one half standard deviation over the sample.

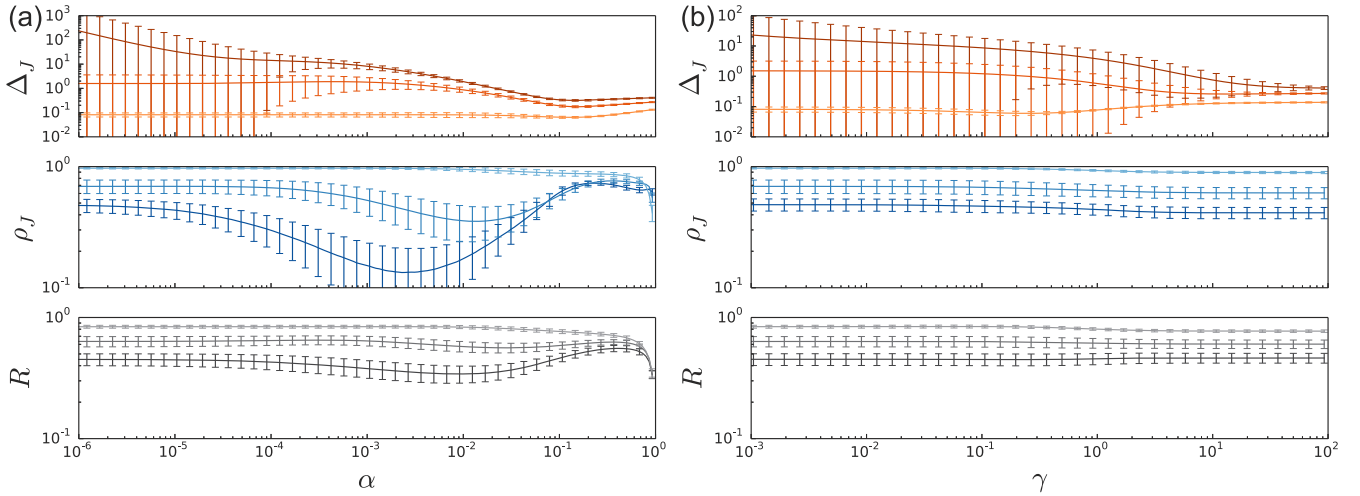


FIG. 12. (Color online) Varying the strength of the underlying interactions can shift the optimal value of the regularization strength. Trajectory of the rms error Δ_J (27) (top), rank correlation ρ_J (28) (middle), and fraction of nonzero couplings recovered R (bottom) as the pseudocount α (a) and L_2 -norm regularization strength γ (b) is varied, averaged over 10^3 sets of random couplings, over a range of coupling distribution widths $\sigma_J = 1$ (light), $\sigma_J = 2$ (medium), and $\sigma_J = 3$ (dark). Relevant values of the pseudocount are $\alpha = \alpha^{\text{MF}} = 0.2$ and $\alpha^B = 1/B = 10^{-4}$, roughly where ρ_J and R plateau for the pseudocount. Each random set of interactions has all fields set to zero. The coupling network is an Erdős-Rényi graph ($N = 100$ nodes) where edges are kept with probability $p = 2/N$. Nonzero couplings are normally distributed with mean zero and standard deviation σ_J . MF couplings were inferred from correlations computed from $B = 10^4$ Monte Carlo samples of the true model. Bars denote one half standard deviation over the sample.

even at large regularization strengths, on the sampling depth. In particular, for poorly sampled systems, a large regularization γ does not improve as much as the one with the pseudocount




FIG. 13. (Color online) Performance of the pseudocount (a) and L_2 -norm regularization (b) differ as a function of the regularization strength, particularly in the behavior of the rank correlation. The rms error Δ_J (27) (circles, left axis), rank correlation ρ_J (28) (squares, right axis), and fraction of nonzero couplings recovered R (triangles, right axis), as a function of pseudocount α and L_2 -norm regularization strength γ , averaged over 10^3 sets of random couplings. Dashed lines in (a) mark $\alpha = \alpha^{\text{MF}} = 0.2$ and $\alpha^B = 1/B = 10^{-3}$, roughly where ρ_J and R plateau for the pseudocount. Each random set of interactions has all fields $h = -5$. The coupling network is an Erdős-Rényi graph where edges are kept with probability $p = 4/N$. Nonzero couplings are normally distributed with mean $\bar{J} = 1$ and standard deviation $\sigma_J = 2$. MF couplings were inferred from correlations computed from $B = 10^3$ Monte Carlo samples of the true model. Bars denote one half standard deviation over the sample.

regularization. The reason is that, even for the large values of γ ($=100$) considered in Figs. 11 and 12, the effective penalty acting on coupling between sites i, j is $\gamma \times [p_i(1 - p_i)p_j(1 - p_j)]$, see Eq. (8), which can be much smaller than γ if p_i or p_j are small. Hence the couplings incoming onto those sites are essentially free of any regularization and may take large values (differing substantially from their true values) when the sampling depth is poor. Those large couplings completely reshuffle the ordering and make poor values for ρ_J and R . This is not the case with large pseudocounts, as p_i and p_j are bounded from below by $\alpha/2$. As a consequence, the rank correlation does not exhibit any local maximum at large γ with L_2 regularization.

Generally, we find that the pseudocount is well suited to situations where the sampling depth is poor and where the true interactions are strong. In such cases ρ_J and R can achieve much larger values than with L_2 regularization, while maintaining similar rms errors Δ_J . This difference between the pseudocount and L_2 -norm regularization schemes can be understood through analysis of the $O(m)$ model, presented in Sec. V. Performance of the pseudocount can be sensitive to changes in α , but the optimal value of α , while varying some with the strength of the true interactions, is generically of the same order as α^{MF} (Fig. 7). Additionally, Δ_J is typically small in the same range of α that maximizes the rank correlation and fraction of nonzero couplings recovered, making the pseudocount particularly attractive in this regime.

L_2 regularization can offer modest advantages compared to the pseudocount when sampling is very good, if the true couplings are not too strong. In these cases L_2 regularization can achieve slightly higher values of ρ_J and R at large values of γ , where the rms error Δ_J is minimized. This method also has the advantage of being much less sensitive to the value of the regularization strength γ .

B. Results for the Potts model

In this section we consider the Potts model, with $q = 5$ and $q = 21$ symbols. We report below results for a one-dimensional interaction network, with $N = 50$ sites; the qualitative conclusions we draw from the study of this model are in agreement with simulations on other interaction network geometries, not shown here. Two variants of this model will be studied, depending on how the Potts interactions vary between the symbols σ_i, σ_{i+1} on neighboring sites along the chain,

(i) *Homogeneous variant*: For each pair of neighbors $i, i + 1$, we draw randomly a number, J_0 , uniformly between $-L$ and L and set all the $q \times q$ couplings of the interaction matrix $\mathbf{J}_{i,i+1}$ through (23). The process is repeated, independently, for all N pairs of neighbors. The model is such that the q Potts symbols have equal frequencies $p_i^a = \frac{1}{q}$.

(ii) *Heterogeneous variants*: Extensions of the above model to nonequal frequencies can be easily obtained. To do so we consider a local field $h_i(a)$ on each site and symbol, which is also, for simplicity, drawn uniformly at random from the $[-L, L]$ range. The explicit introduction of a field allows us to increase the bias between the frequencies of the q symbols. This model will be called the heterogenous-A model in the following. We may, in addition to the introduction of random fields, draw randomly, for each pair of neighbors $i, i + 1$, and for each pair of symbols a, b , a coupling $J_{i,i+1}(a, b)$, uniformly between $-L$ and L . Again, the process is repeated, independently, for all pairs of neighbors. We refer to this model as the heterogeneous-B model.

The values of the one- and two-point correlations are obtained from Monte Carlo simulations with the Hamiltonian (11) in the case of limited sampling (from B configurations) and through a transfer matrix calculation in the case of perfect sampling ($B = \infty$). The following reports results of the MF inference, with and without pseudocount. In the latter case we use the strengths $\alpha = 0.41$ for $q = 5$ and $\alpha = 0.75$ for $q = 21$. Those choices correspond to the “optimal” pseudocount values found in Sec. III E 1 (Fig. 7).

We make sure that the coupling matrices $\mathbf{J}_{i,i+1}$ satisfy the zero-sum gauge: The sum of all couplings along each column

and row of the coupling matrix vanish. The gauge is imposed through

$$J_{i,i+1}(a, b) \rightarrow J_{i,i+1}(a, b) - \frac{1}{q} \sum_{a=1}^q J_{i,i+1}(a, b) - \frac{1}{q} \sum_{b=1}^q J_{i,i+1}(a, b) + \frac{1}{q^2} \sum_{a,b=1}^q J_{i,i+1}(a, b). \quad (29)$$

This choice allows us to compare the original and the inferred couplings.

1. Importance of the bias in frequencies on the quality of inference

In this section we consider the perfect sampling ($B = \infty$) case, in the presence or the absence of a pseudocount regularization. We start with the homogeneous variant. The results of the MF inference are shown in Fig. 14(a). We observe a perfect agreement with the analytical curves, see (23) and (25), if no regularization is present. As a result of the presence of the pseudocount (Sec. III C), many couplings whose true values vanish are inferred with nonzero values, corresponding to the points ($J = 0, J^{\text{PC}} \neq 0$) in the scatter plot of Fig. 14(b). In turn, to account for those extra fictitious couplings, the inferred values for the “existing” (between adjacent sites) couplings are lowered with respect to their true values, see Fig. 14(b). This effect is weaker if the true couplings are chosen from a smaller range, i.e., if L is decreased.

Performance of MF inference for the heterogeneous variants are shown in Figs. 15 and 16. Comparison with the toy-model analysis of Sec. III E 2, see Fig. 8, may be only qualitative here, because the heterogeneous models considered in the simulations have a richer distribution of frequencies p_i^a . In the toy model and in the presence of a pseudocount, where frequencies can take only two values, four branches of couplings values appear in the regions $0 < J^{\text{PC}} < J$ and $J < J^{\text{PC}} < 0$ of the (J, J^{PC}) plane. In the heterogeneous models A and B the frequencies of the q

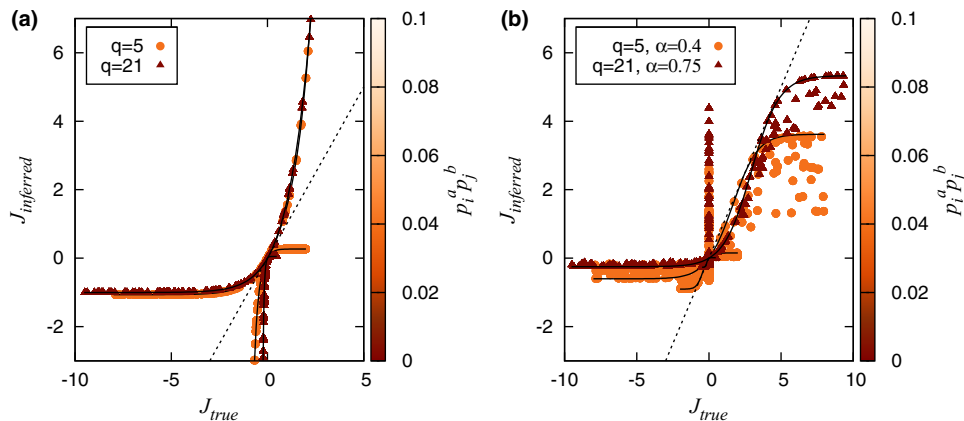


FIG. 14. (Color online) Scatter plot of the couplings $J_{ij}(a, b)$ for the homogenous Potts model for $q = 5$ (filled circles) and for $q = 21$ (triangles) symbols; perfect sampling. (a) No pseudocount. (b) With pseudocount. Each panel shows results from three realizations with different sets of couplings ($L = 10$). Black solid lines correspond to the analytical predictions of Sec. III E 1. Colors show values of $p_i^a p_j^b$, here equal to $1/q^2$ for all interacting sites and for all symbols, see right scale.

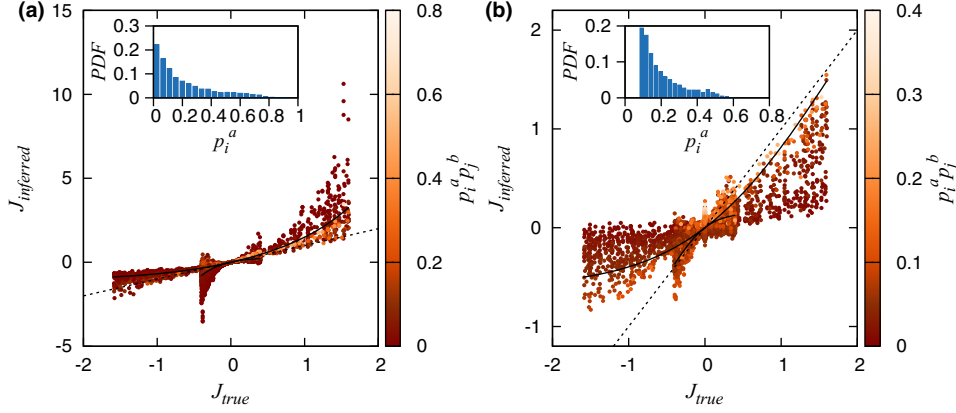


FIG. 15. (Color online) Scatter plot of the couplings $J_{ij}(a,b)$ for the heterogeneous-A Potts model for $q = 5$ symbols, with perfect sampling. (a) No pseudocount. (b) With pseudocount ($\alpha = 0.4$). Each panel shows results from five realizations with different sets of couplings and fields ($L = 2$). Insets: Distributions of the frequencies p_i^a . Black solid lines correspond to the analytical predictions of Sec. III E 1. Colors show values of $p_i^a p_j^a$, see right scale.

symbols are not bimodal, see the inset of Figs. 15 and 16. We anticipate that branches will be less easy to identify but will occupy the same regions of the plane in a dense way.

Results for the heterogeneous-A model are shown in Fig. 15 for $q = 5$ symbols; similar results were obtained for the $q = 21$ -symbol Potts model with and without the pseudocount ($\alpha = 0.75$). The agreement with the real couplings is generally better for couplings $J_{i,i+1}(a,b)$ corresponding to *conserved* sites and symbols, i.e., with large values of the product $p_i^a p_{i+1}^b$. This statement holds also in the presence of a pseudocount, which improves the inference for pairs of sites and symbols with medium and large frequencies [lightest points in Fig. 15(b)]. The behavior of these inferred J^{PC} couplings with a large or medium level of conservation is indeed similar to the branches J^A, J^B from which we have tuned α^{MF} (Fig. 7) in the toy model for the homogeneous case. Moreover, as seen in the homogeneous case [Fig. 14(b)], noninteracting but strongly conserved sites can generate fictitious and strong inferred couplings [Fig. 15(b)]. One consequence of the pseudocount is, indeed, to produce larger correlations between very conserved sites and, in turn, nonzero inferred couplings between those sites.

Performances of the MF inference for the heterogeneous-B model are shown in Fig. 16. The distribution of frequencies p_i^a is more peaked at low values (inset of Fig. 16) than in the heterogeneous-A model (inset of Fig. 15). The global picture is similar to the one of Fig. 15, with an even wider dispersion. Again, we find that couplings corresponding to

medium or strongly conserved sites are generally better inferred than the ones corresponding to nonconserved sites for the mean-field inference. As in the heterogeneous-A model, in the presence of a pseudocount, nonzero couplings appear between nonadjacent and strongly conserved sites.

2. Effects of finite sampling and reconstruction of the network structure

We now study the effect of finite sampling on MF inference for the heterogeneous-B model (Fig. 17). The errors on the inferred MF couplings are strongly affected by the sampling size for weak values of the pseudocount, e.g., the large peak in $J^{\text{MF}} \neq 0, J^{\text{true}} = 0$ corresponding to nonadjacent sites in Fig. 17(a). Remarkably, for strong pseudocount (optimal value defined in Sec. III E 1), limited sampling has little effect on the inference error [Fig. 17(b)], which seems to be due primarily to the poor performance of MF inference in the presence of a wide distribution of the local frequencies p_i^a .

We present in Fig. 18 the scatter plots of the Frobenius norm $F_{ij} = (\sum_{a,b} J_{ij}(a,b)^2)^{1/2}$ for the same heterogeneous-B model as in Fig. 17. The Frobenius norm F_{ij} is a simple way to include contributions from interactions between all symbols at a pair of sites (i,j) and to extract information about the presence or the absence of a link between i and j . Indeed, nonzero links (i,j) in the interaction network generally give rise to larger F_{ij} than zero links. With a weak pseudocount, $\alpha = \frac{1}{B}$, smaller sample sizes result in poorer performance. Many

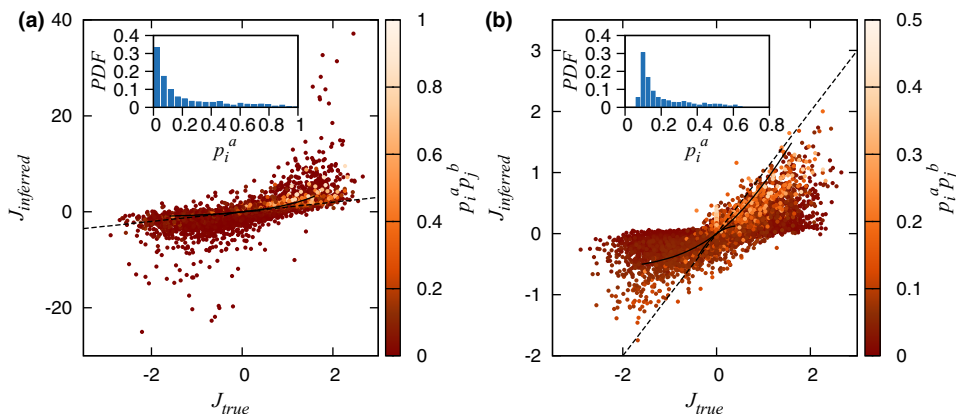


FIG. 16. (Color online) Scatter plot of the couplings $J_{ij}(a,b)$ for the heterogeneous-B Potts model for $q = 5$ symbols, with perfect sampling. (a) No pseudocount. (b) With pseudocount ($\alpha = 0.4$). Each panel shows results from five realizations with different sets of couplings and fields ($L = 2$). Insets: Distributions of the frequencies p_i^a . Black solid lines correspond to the analytical predictions of Sec. III E 1. Colors show values of $p_i^a p_j^a$, see right scale.

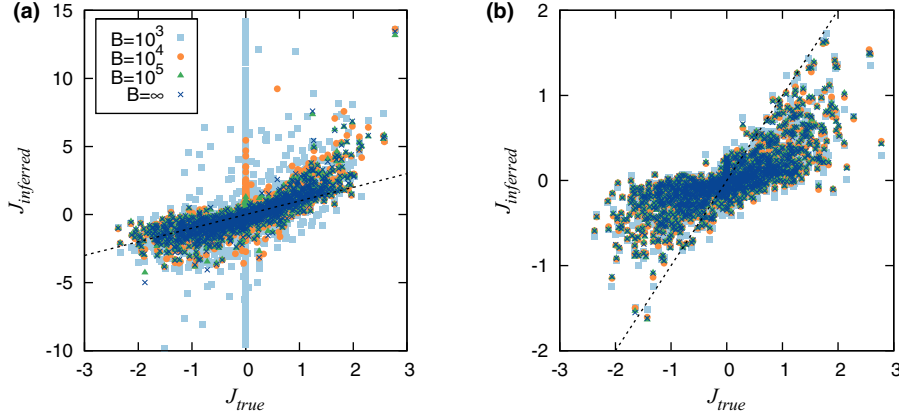


FIG. 17. (Color online) Heterogeneous-B Potts model for $q = 5$ symbols, for various depths of sampling. (a) Small pseudocount $\alpha = \alpha^B = \frac{1}{B}$. (b) Large (optimal) pseudocount $\alpha = \alpha^{\text{MF}} = 0.4$. Each panel shows results from one realization of the Potts model with random couplings and fields ($L = 2$) and three sets of B sampled configurations (for finite B).

pairs (i, j) with zero couplings give rise to Frobenius norms larger than the ones found for pairs of neighbors $(i, i + 1)$, which have real nonzero couplings. This artifact implies that the graph structure cannot be correctly reconstructed, without having an estimate of the statistical error bars on single couplings due to the sampling noise [37]. However, with the optimal strength $\alpha = 0.4$, finite sampling effects are better corrected for, and the correct structure of the graph is recovered. Even though the inferred couplings $J_{ij}^{\text{PC}}(a, b)$ differ from their true values $J_{ij}(a, b)$, the summation over the Potts symbols in the Frobenius norms seem to average out those errors and to allow for a good inference of the underlying graph structure in the presence of a large pseudocount. This result supports the use of large pseudocounts in real applications such as protein contact predictions from covariation data [4].

V. ANALYSIS OF THE $O(m)$ MODEL FOR LARGE BUT FINITE m

In this section we consider a large system of N spins, and we want to estimate the error on the inferred couplings due to the MF approximation and how this error can be corrected for with appropriate regularization. Estimating exactly this error would require that one solves exactly the inverse Ising or Potts model, which is computationally intractable. We therefore concentrate on an extension of the Ising model, the $O(m)$ model, which can be solved analytically for large m . The spin variables in the model are m -dimensional vectors, $\vec{\sigma}_i$, with squared norms constrained to be equal to m : $|\vec{\sigma}_i|^2 = m, \forall i$. The case $m = 1$

corresponds to Ising spins. We first recall how the properties of the $O(m)$ model can be exactly solved in the infinite m limit and how a systematic expansion in powers of $1/m$ can be carried out. We then compute the error done by MF on the inverse $O(m)$ model and study to which extent these errors are compensated by pseudocount and L_2 regularizations.

A. Statistical mechanics of the inverse $O(m)$ model

Given a set of interactions J_{ij} , the likelihood of a configuration $\{\vec{\sigma}_1, \vec{\sigma}_2, \dots, \vec{\sigma}_N\}$ of the model is

$$p(\vec{\sigma}_1, \vec{\sigma}_2, \dots, \vec{\sigma}_N) = \frac{\exp\left(\sum_{i < j} J_{ij} \vec{\sigma}_i \cdot \vec{\sigma}_j\right)}{Z(\{J_{ij}\})}, \quad (30)$$

where the partition function reads

$$Z(\{J_{ij}\}) = \int_{|\vec{\sigma}_1|^2=m} d\vec{\sigma}_1 \dots \int_{|\vec{\sigma}_N|^2=m} d\vec{\sigma}_N \times \exp\left(\sum_{i < j} J_{ij} \vec{\sigma}_i \cdot \vec{\sigma}_j\right). \quad (31)$$

In the above formulas the center dot represents the dot product between two spin vectors. The correlation per spin component is defined through

$$c_{ij} = \frac{1}{m} \int_{|\vec{\sigma}_1|^2=m} d\vec{\sigma}_1 \dots \int_{|\vec{\sigma}_N|^2=m} d\vec{\sigma}_N \times p(\vec{\sigma}_1, \vec{\sigma}_2, \dots, \vec{\sigma}_N) \vec{\sigma}_i \cdot \vec{\sigma}_j. \quad (32)$$

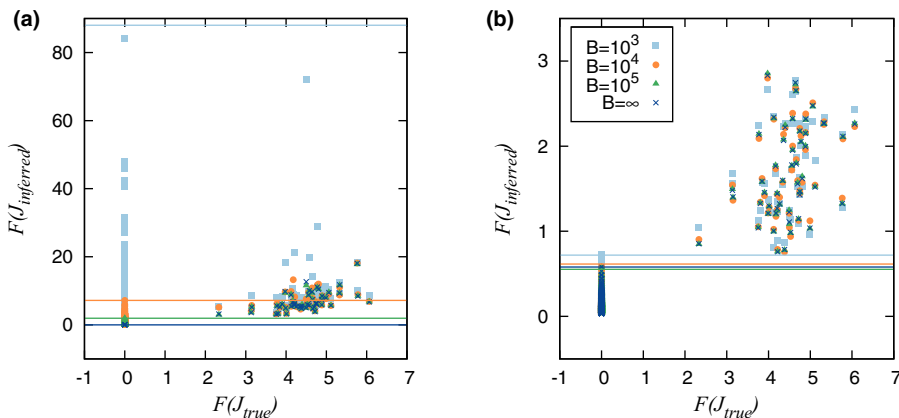


FIG. 18. (Color online) Scatter plot of the Frobenius norms of the inferred couplings vs their true values for the pseudocount strengths $\alpha = \frac{1}{B}$ (a) and $\alpha = 0.4$ (b). Same heterogeneous-B model and same conditions as in Fig. 17. Lines locate the largest Frobenius norm corresponding to a pair of sites (i, j) which are not neighbors on the one-dimensional graph, i.e., which have zero true coupling.

To compute the partition function Z we introduce imaginary-valued Lagrange multipliers λ_i to enforce the constraint over the norm of $\vec{\sigma}_i$ for all $i = 1, \dots, N$. We obtain the following:

$$\begin{aligned} Z(\{J_{ij}\}) &= \int_{i_R} \frac{d\lambda_1}{4\pi} \dots \int_{i_R} \frac{d\lambda_N}{4\pi} \int d\vec{\sigma}_1 \dots \int d\vec{\sigma}_N \\ &\times \exp \left[\sum_{i < j} J_{ij} \vec{\sigma}_i \cdot \vec{\sigma}_j + \sum_j \frac{\lambda_j}{2} (m - \vec{\sigma}_j^2) \right] \\ &= 2^{-N} (2\pi)^{N(m/2-1)} \int_{i_R} d\lambda_1 \dots \int_{i_R} d\lambda_N \\ &\times \exp \left[\frac{m}{2} \sum_j \lambda_j - \frac{m}{2} \log \det A(\lambda, J) \right], \quad (33) \end{aligned}$$

where $A(\lambda, J)$ is a $N \times N$ symmetric matrix, with diagonal elements $A_{ii} = \lambda_i$ and off-diagonal elements $A_{ij} = -J_{ij}$. For large m we estimate the integral according to the saddlepoint method: The values of the Lagrange multipliers λ_i^* are such that the diagonal elements of the inverse matrix of A are all equal to 1: $[A(\lambda^*, J)^{-1}]_{ii} = 1, \forall i$. Gaussian corrections to the saddlepoint are easy to compute with the following expression for the log-likelihood of the data given the coupling matrix:

$$\begin{aligned} L(J|c) &= m \sum_{i < j} J_{ij} c_{ij} - \frac{m}{2} \sum_i \lambda_i^* + \frac{m}{2} \log \det A(\lambda^*, J) \\ &+ \frac{1}{2} \log \det H(J), \\ &= -\frac{m}{2} \text{Trace}[A(\lambda^*, J) c] + \frac{m}{2} \log \det A(\lambda^*, J) \\ &+ \frac{1}{2} \log \det H(J), \quad (34) \end{aligned}$$

where we have omitted an irrelevant J -independent additive constant, and the $N \times N$ symmetric matrix H is defined through

$$H_{ij}(J) = ([A(\lambda^*, J)^{-1}]_{ij})^2, \quad (35)$$

where H is the pointwise square of a positive definite matrix; according to the Schur product theorem it is itself a positive matrix.

We now consider the inverse $O(m)$ problem. We want to determine the coupling matrix $J = \{J_{ij}\}$ fulfilling the constraints (32). To do so we maximize the log-likelihood L with respect to J . In the infinite- m limit the solution is simply $(J_\infty)_{ij} = -(c^{-1})_{ij}$, as expected for the Gaussian model, which is equivalent to the mean-field approximation. Deviations from the MF inference are found when m is large but finite,

$$\delta J_{ij} \equiv J_{ij} - (J_\infty)_{ij} = \frac{1}{m} (H_\infty^{-1})_{ij} c_{ij} + O\left(\frac{1}{m^2}\right), \quad (36)$$

where

$$(H_\infty)_{ij} = (c_{ij})^2. \quad (37)$$

Expression (36) is our ‘‘exact’’ value for the couplings given the correlation matrix. Below we study the accuracy of the MF prediction (in the presence of regularization) compared to this expression.

B. Effect of regularization schemes

We start with the L_2 regularization with link-dependent penalty, γ_{ij} , i.e., we add a penalty term $-\frac{1}{4} \sum_{i,j} \gamma_{ij} J_{ij}^2$ to the log-likelihood L (recall that diagonal couplings J_{ii} coincide here with $-\lambda_i^*$). We then extremize with respect to J and ask for the change in J resulting from the presence of this new L_2 penalty term to compensate exactly δJ given by (36). A straightforward calculation leads to

$$\gamma_{ij} = -\frac{1}{m (c^{-1})_{ij}} \sum_{k,\ell} c_{ik} (H_\infty^{-1})_{k\ell} c_{k\ell} c_{\ell j} + O\left(\frac{1}{m^2}\right). \quad (38)$$

We now repeat the approach with the pseudocount regularization. We consider the general case of a link-dependent pseudocount of strength α_{ij} . The off-diagonal entries c_{ij} of the correlation matrix are now equal to $(1 - \alpha_{ij})c_{ij}$, while the diagonal entries are unchanged: $c_{ii} = 1$. In the $m \rightarrow \infty$ limit the change in the coupling J_{ij} resulting from the presence of the pseudocount is, to the first order in α ,

$$\delta J_{ij}^{\text{PC}} = -\sum_{k \neq \ell} (c^{-1})_{ik} \alpha_{k\ell} c_{k\ell} (c^{-1})_{\ell j}. \quad (39)$$

We want to compute the values of the strengths α_{kl} , with $k \neq l$, such that $\delta J_{ij}^{\text{PC}}$ and δJ_{ij} in (36) sum up to zero for all $i \neq j$. The solutions are given by

$$\begin{aligned} \alpha_{ij} &= \frac{1}{m c_{ij}} \left[\sum_{k,\ell} c_{ik} (H_\infty^{-1})_{k\ell} c_{k\ell} c_{\ell j} - \sum_k c_{ik} d_k c_{kj} \right] \\ &+ O\left(\frac{1}{m^2}\right), \quad (40) \end{aligned}$$

where

$$d_k = \sum_i (H_\infty^{-1})_{ki} \sum_{a,b} c_{ia} (H_\infty^{-1})_{ab} c_{ab} c_{bi}. \quad (41)$$

The presence of the second term in the brackets in (40) ensures that the pseudocount vanishes on the diagonal, i.e., $\alpha_{ii} = 0$.

Two conclusions can be drawn as follows from the previous calculations:

(i) We find that the optimal penalties, with L_2 regularization and pseudocount, do not vanish in the perfect sampling limit considered here. The need for regularization can therefore not be due to poor sampling only. More precisely, the optimal penalties are of the order of $\frac{1}{m}$ for large m . Loosely speaking, they are proportional to the deviation from the Gaussian model (recovered when $m \rightarrow \infty$, for which MF inference is exact).

(ii) We also understand from the formulas above why uniform pseudocount is generally better than uniform L_2 penalty, as found in Sec. III D. In (40) the pseudocount strength scales as the inverse of c_{ij} , which saturates to 1 for very strongly correlated spins. In contrast, in (38), the penalty scales as the inverse of c_{ij}^{-1} , that is, as $1/J_{ij}$. This quantity is not bounded from below when the coupling increases. Hence we expect a much wider range of values for the optimal γ coefficients than for the optimal α coefficients. Uniform L_2 penalties are therefore far away from being optimal.

VI. CONCLUSION

The present paper summarizes our efforts to understand the empirically observed necessity of large regularization terms in the mean-field inference of Ising or Potts interaction networks. In the usual Bayesian interpretation pseudocount and L_2 -norm regularization penalties are required in case of undersampling. As more data become available, the sampling noise becomes smaller, and so do the optimal values of the regularization terms. A combination of analytical and numerical evidence suggests that this interpretation is not correct for MF inference and that the need for large regularization penalties rather comes from the non-Gaussian character of the variable statistics. In other words, large penalties, particularly for the pseudocount, correct for the error in the inferred couplings introduced by MF. The importance of large regularization penalties to correct for errors introduced by the MF approximation is confirmed by analysis of the m -components $O(m)$ spin model: The optimal amplitude for the regularization scales, for large but finite m , as $\frac{1}{m}$, which coincides with the measure of the discrepancy of the model with a Gaussian statistics.

In this work we also explored the performance of MF inference with different regularization schemes for a diverse set of underlying model systems. In general cases both large pseudocount and L_2 -norm regularization can yield couplings which correlate well with the true couplings; the pseudocount is easier to use (requires less fine-tuning) in general, especially when no knowledge of the true couplings exists to guide the choice of an appropriate regularization strength. Moreover, a large pseudocount gives extremely stable performance in the inference in case of limited sampling. For the Ising model case we find that the optimal pseudocount to infer the network topology and the coupling values is $\alpha^{\text{MF}} = 0.2$, independently of the sampling depth, of the values of the true interactions, and of the structure of the interaction network. This value corresponds to what was analytically obtained for a toy model with only two spins.

The toy-model approach also allows us to show that the mean-field inference of the coupling parameters $J_{ij}(a,b)$ in the Potts model case is poorer than in the Ising model case. The quality of the inference generally worsens with the heterogeneity among the frequencies of the q Potts symbols (states). The introduction of a pseudocount helps in reducing the errors on the inferred couplings $J_{ij}(a,b)$, especially for

symbols a and b having medium and large frequencies on the site i and j . For those couplings, an optimal value of the pseudocount can be determined, which increases with q and agrees with what was analytically found with the toy model in the homogeneous case. Couplings attached to nonconserved sites are the ones with the largest inference errors and cannot be reliably inferred even in the presence of a pseudocount. Moreover, the introduction of a pseudocount may lead to the prediction of nonzero couplings between noninteracting but very conserved sites, an artifact known in the context of the application of MF inference techniques to residue covariation in protein families [4]. Even if for heterogeneous Potts models many couplings are poorly inferred, the introduction of a pseudocount is helpful to improve the reconstruction of the interaction network structure in the case of limited sampling. This finding is, again, in good agreement with empirical results in the context of protein covariation.

Our work could be extended in several ways. A potentially important finding of Sec. V is that there exist optimal values for the penalties, which depend on the empirical correlations, see (38) and (40). It would be interesting to pursue this direction and to see whether the introduction of link-dependent penalties could, indeed, improve the quality of MF inference in practical applications. Another issue of interest would be determine if and how the optimal regularization strengths depend on additional and specific constraints on the coupling matrix. A practical example is provided by the inverse Hopfield model [6,39], in which the rank p of the interaction matrix J is small compared to the system size. Last, it would be interesting to understand the observed stability of the inferred couplings against the sampling depth in the presence of a pseudocount, see Fig. 17 and Fig. 12 for, respectively, the Potts and Ising cases. We expect performances to deteriorate when B gets of the order of, or smaller than N [6], but a more quantitative understanding of the minimal sampling depth required would be useful in practical applications.

ACKNOWLEDGMENTS

We are grateful to M. Weigt for very useful discussions. We thank S. Rosay for her contribution to the analytical study of the $O(m)$ model in Sec. V. This work was partly funded by the Agence Nationale de la Recherche Coevstat project (Grant No. ANR-13-BS04-0012-01).

-
- [1] M. Socolich, S. W. Lockless, W. P. Russ, H. Lee, K. H. Gardner, and R. Ranganathan, *Nature* **437**, 512 (2005).
 - [2] M. Weigt, R. A. White, H. Szurmant, J. A. Hoch, and T. Hwa, *Proc. Natl. Acad. Sci. USA* **106**, 67 (2009).
 - [3] L. Burger and E. van Nimwegen, *PLoS Comput. Biol.* **6**, e1000633 (2010).
 - [4] 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. USA* **108**, E1293 (2011).
 - [5] T. A. Hopf, L. J. Colwell, R. Sheridan, B. Rost, C. Sander, and D. S. Marks, *Cell* **149**, 1607 (2012).
 - [6] S. Cocco, R. Monasson, and M. Weigt, *PLoS Comput. Biol.* **9**, e1003176 (2013).
 - [7] A. L. Ferguson, J. K. Mann, S. Omarjee, T. Ndung'u, B. D. Walker, and A. K. Chakraborty, *Immunity* **38**, 606 (2013).
 - [8] E. Segal, M. Shapira, A. Regev, D. Pe'er, D. Botstein, D. Koller, and N. Friedman, *Nat. Genet.* **34**, 166 (2003).
 - [9] T. R. Lezon, J. R. Banavar, M. Cieplak, A. Maritan, and N. V. Fedoroff, *Proc. Natl. Acad. Sci. USA* **103**, 19033 (2006).
 - [10] R. Bonneau, M. T. Facciotti, D. J. Reiss, A. K. Schmid, M. Pan, A. Kaur, V. Thorsson, P. Shannon, M. H. Johnson, J. C. Bare *et al.*, *Cell* **131**, 1354 (2007).
 - [11] E. Schneidman, M. J. Berry, II, R. Segev, and W. Bialek, *Nature* **440**, 1007 (2006).
 - [12] G. D. Field and E. J. Chichilnisky, *Annu. Rev. Neurosci.* **30**, 1 (2007).

- [13] J. Pillow, J. Shlens, L. Paninski, A. Sher, A. Litke, E. Chichilnisky, and E. Simoncelli, *Nature* **454**, 995 (2008).
- [14] S. Cocco, S. Leibler, and R. Monasson, *Proc. Natl. Acad. Sci. USA* **106**, 14058 (2009).
- [15] G. Tkacik, J. Prentice, V. Balasubramanian, and E. Schneidman, *Proc. Natl. Acad. Sci. USA* **107**, 14419 (2010).
- [16] K. Faust and J. Raes, *Nat. Rev. Microbiol.* **10**, 538 (2012).
- [17] J. Friedman and E. J. Alm, *PLoS Comput. Biol.* **8**, e1002687 (2012).
- [18] W. Bialek, A. Cavagna, I. Giardina, T. Mora, E. Silvestri, M. Viale, and A. M. Walczak, *Proc. Natl. Acad. Sci. USA* **109**, 4786 (2012).
- [19] D. R. Hekstra, S. Cocco, R. Monasson, and S. Leibler, *Phys. Rev. E* **88**, 062714 (2013).
- [20] N. A. Christakis and J. H. Fowler, *New Engl. J. Med.* **357**, 370 (2007).
- [21] D. Lazer, A. Pentland, L. Adamic, S. Aral, A.-L. Barabasi, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann *et al.*, *Science* **323**, 721 (2009).
- [22] S. P. Borgatti, A. Mehra, D. J. Brass, and G. Labianca, *Science* **323**, 892 (2009).
- [23] V. Plerou, P. Gopikrishnan, B. Rosenow, L. A. Nunes Amaral, and H. E. Stanley, *Phys. Rev. Lett.* **83**, 1471 (1999).
- [24] E. Moro, J. Vicente, L. G. Moyano, A. Gerig, J. D. Farmer, G. Vaglica, F. Lillo, and R. N. Mantegna, *Phys. Rev. E* **80**, 066102 (2009).
- [25] E. T. Jaynes, *Proc. IEEE* **70**, 939 (1982).
- [26] T. Hastie, R. Tibshirani, and J. H. Friedman, *The Elements of Statistical Learning*, Data Mining, Inference, and Prediction (Springer-Verlag, Berlin, 2009).
- [27] V. Sessak and R. Monasson, *J. Phys. A: Math. Theor.* **42**, 055001 (2009).
- [28] H. C. Nguyen and J. Berg, *J. Stat. Mech.: Theor. Exp.* (2012) P03004.
- [29] J. Sohl-Dickstein, P. B. Battaglino, and M. R. DeWeese, *Phys. Rev. Lett.* **107**, 220601 (2011).
- [30] S. Cocco and R. Monasson, *Phys. Rev. Lett.* **106**, 090601 (2011).
- [31] P. Ravikumar, M. J. Wainwright, and J. D. Lafferty, *Ann. Stat.* **38**, 1287 (2010).
- [32] E. Aurell and M. Ekeberg, *Phys. Rev. Lett.* **108**, 090201 (2012).
- [33] J. Barton and S. Cocco, *J. Stat. Mech.: Theor. Exp.* (2013) P03002.
- [34] D. S. Marks, L. J. Colwell, R. Sheridan, T. A. Hopf, A. Pagnani, R. Zecchina, and C. Sander, *PLoS One* **6**, e28766 (2011).
- [35] H. Huang and Y. Kabashima, *Phys. Rev. E* **87**, 062129 (2013).
- [36] R. Durbin, S. Eddy, A. S. Krogh, and G. Mitchison, *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids* (Cambridge University Press, Cambridge, 1998).
- [37] S. Cocco and R. Monasson, *J. Stat. Phys.* **147**, 252 (2012).
- [38] J. Friedman, T. Hastie, and R. Tibshirani, *Biostatistics* **9**, 432 (2008).
- [39] S. Cocco, R. Monasson, and V. Sessak, *Phys. Rev. E* **83**, 051123 (2011).