



Data quality for the inverse Ising problem

Aurélien Decelle, Federico Ricci-Tersenghi, Pan Zhang

► To cite this version:

Aurélien Decelle, Federico Ricci-Tersenghi, Pan Zhang. Data quality for the inverse Ising problem. *Journal of Physics A: Mathematical and Theoretical*, 2016, 49 (38), 10.1088/1751-8113/49/38/384001 . hal-01250822

HAL Id: hal-01250822

<https://hal.science/hal-01250822>

Submitted on 5 Jan 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Data quality for the inverse Ising problem

Aurélien Decelle¹, Federico Ricci-Tersenghi², Pan Zhang³

¹ Laboratoire de Recherche en Informatique, TAO - INRIA, CNRS, Univ. Paris-Sud, Université Paris-Saclay, Bât. 660, 91190 Gif-sur-Yvette, France.

² Dipartimento di Fisica, INFN-Sezione di Roma1 and CNR-Nanotec, Università La Sapienza, Piazzale Aldo Moro 5, I-00185 Roma, Italy.

³ Institute of Theoretical Physics, Chinese Academy of Sciences, Zhong-Guan-Cun-Dong-Lu 55, Beijing 100190, China.

Abstract. There are many methods proposed for inferring parameters of the Ising model from given data, that is a set of configurations generated according to the model itself. However little attention has been paid until now to the data, e.g. how the data is generated, whether the inference error using one set of data could be smaller than using another set of data, etc. In this paper we address the data quality problem in the kinetic inverse Ising problem. We quantify the quality of data using effective rank of the correlation matrix, and show that data gathered in a out-of-equilibrium regime has a better quality than data gathered in equilibrium for coupling reconstruction. We also propose a matrix-perturbation based method for tuning the quality of given data and for removing bad-quality (i.e. redundant) configurations from data.

Submitted to: *J. Phys. A: Math. Theor.*

1. Introduction

In the past few years, considerable attention has been drawn to the inverse Ising problems that study how to infer or reconstruct the parameters of an Ising model from configurations generated according to the model itself. This inference process, also known as “Boltzmann machine learning” in computer science [1], is linked to the maximum entropy principle applied to models of pairwise interacting variables when the first two moments of data are measured [2], and has a capability to model rich behaviors of observed data. Thus it has been used to reconstruct interaction patterns of complex systems, such as the coupling constants in a magnetic alloy, the interactions between firing neurons in neural networks (either in vivo, in vitro or in silico), the way chemical reaction are coupled together in metabolic networks, sociological interactions in social network, etc. The applications of the inverse Ising model can be found in in different fields of science including physics [3, 4, 5, 6, 7], computer science [8], neuroscience [2, 9], social network [10] and biology [11, 12, 13, 14].

The canonical approach to tackle this problem is by the inference methods that maximize the likelihood of the parameters given the data. Most of the study on this subject focused on improving the performance of the inference by increasing accuracy and efficiency of inference methods [3, 4, 5, 15, 6, 7, 16], for instance by applying improved mean-field and cluster variational methods [17], by using regularizations of different forms [8], etc. A similar and very effective approach can be used also in case data comes from a dynamical process [18].

However little attention has been paid to the data side, for instance how to improve the performance of inference by increasing the quality of data. We think that in the recent “age of big-data”, the data-side consideration could become more and more important, as recently we have been observing that the amount of available data in many fields has been growing so quickly that, in some cases taking all of them into account for the inference becomes a difficult task. Then one question arises: do we really need all the data for the inference? To put it differently, does every configuration in the data contains equal amount of information about the system? Obviously, this question concerns the *data quality* problem: how to quantify the quality of the data and how to improve the data quality.

In this article we address directly the questions posed in the last paragraph – the data quality in the inverse Ising problems. We will focus in this work on the dynamical inverse Ising case where the data are generated by a stochastic process where variables update synchronously. Using this model, we will demonstrate that data coming from the out-of-equilibrium regime are much more informative than equilibrium configurations, or configurations gather from a steady state of the system. Our results may suggest new experimental protocols to acquire data reconstructing the interaction network, where the system under study is first perturbed to an out-of-equilibrium state, and then measured.

The paper is organized as follows. Sec. 2 contains descriptions of the model and the stochastic process to update/evolve the system variables. In Sec. 3 we quantify the

data quality using the effective rank of the correlation matrix, and show that out-of-equilibrium data has higher quality than the equilibrium/stationary data. In Sec. 4 we propose a method based on the perturbation analysis of the correlation matrix for tuning the quality of data, i.e. removing configuration from the data in such a way that the data quality keeps improving. Finally we conclude this work in Sec. 5.

2. The dynamical inverse Ising model

The dynamical inverse Ising model is based on pairwise interactions amongst discrete variables and a dynamical rule to update these variables. The general setting considers n variables corresponding to the n nodes of a graph. Each variable (denoted by i) takes values $s_i = \pm 1$. An edge or coupling between node i and node j , J_{ij} , takes a real value to represent an interaction between two nodes. In this work we consider only the case of symmetric interactions, i.e. $J_{ji} = J_{ij}$, which ensures the existence of an equilibrium (i.e. stationary) measure $P_{\text{eq}}(\underline{s})$ at inverse temperature β (the generalization for non-symmetric couplings is straightforward)[‡].

A common and practical choice for the stochastic process that simulates the system evolution is the so-called parallel dynamics, where a new configuration at time $t + 1$ is drawn synchronously from the state at time t : at each time step, each variable is updated according to the local field acting on it, defined as $h_i^{\text{loc}}(\underline{s}) = \sum_j J_{ij}s_j$. Then, the probability of the configuration at time $t + 1$ can be written as

$$P[\underline{s}^{(t+1)}|\underline{s}^{(t)}] = \prod_{i=1}^N \frac{e^{\beta s_i^{(t+1)} h_i^{\text{loc}}(\underline{s}^{(t)})}}{2 \cosh(\beta h_i^{\text{loc}}(\underline{s}^{(t)}))}.$$

It can be shown that in general this dynamics respects the detailed-balance (some oscillations can arise for $\beta \rightarrow \infty$, but it is not of our concern here). However, it is known that, in many cases, the dynamics can be very slow to reach thermal equilibrium. It is typically the case when the system is in a glassy phase (e.g. a Sherrington-Kirkpatrick model [19] with $\beta > 1$), or when it undergoes a rapid quench beyond a second order phase transition.

Let us then consider the following experiment. First, we generate data, i.e. a set of correlated configurations $\underline{s}^{(t)}$ for $t = 0, \dots, T$, using the parallel dynamics described above for given interaction couplings. Second, we try to infer the values of these couplings $\{J_{ij}\}$ using only the data $\{\underline{s}^{(t)}\}_{t=0, \dots, T}$. Since the dynamical updating rule

[‡] As usual, we scale the interactions such as to have unitary variance (in the units that make the energy extensive) and avoid self-interactions ($J_{ii} = 0$).

is known, we can achieve this goal by maximizing the empirical log-likelihood,

$$\begin{aligned}
\mathcal{L} &= \langle \log \left(\prod_{t=1}^T P[\underline{s}^{(t)} | \underline{s}^{(t-1)}] \right) \rangle_{\text{DATA}} \\
&= \left\langle \sum_{t=1}^T \sum_{i=1}^N \beta s_i^{(t)} h_i^{\text{loc}}(\underline{s}^{(t-1)}) - \log(2 \cosh(\beta h_i^{\text{loc}}(\underline{s}^{(t-1)}))) \right\rangle_{\text{DATA}} \\
&= \sum_{t=1}^T \left[\sum_{i,j} \beta J_{ij} \langle s_i^{(t)} s_j^{(t-1)} \rangle_{\text{DATA}} - \sum_i \langle \log(2 \cosh(\beta h_i^{\text{loc}}(\underline{s}^{(t-1)}))) \rangle_{\text{DATA}} \right], \quad (1)
\end{aligned}$$

where $\langle \cdot \rangle_{\text{DATA}}$ denotes the average over the gathered configurations. At variance to the static inverse Ising problem, we are able to maximize directly the likelihood for this dynamical model, since all the terms in Eq. (1) can be computed in polynomial time.

To evaluate the performance of reconstruction we consider a measure to the difference between inferred couplings $\{J_{ij}\}$ and the true couplings $\{J_{ij}^*\}$:

$$\Delta_J = \sqrt{\frac{\sum_{i < j} (J_{ij} - J_{ij}^*)^2}{n(n-1)/2}}. \quad (2)$$

In cases where the system under study is an instance of ensemble of problems, i.e. it is a disordered model, we should also average the reconstruction error over the disorder ensemble. However, we expect the reconstruction error to be self-averaging, so few samples are enough to estimate it. In practice, for each value of the parameters we are going to use (i.e. β and T , see below for definitions) we choose a different sample (i.e. different couplings), such that sample-to-sample fluctuations can be appreciated in the plots reporting the results on Δ_J .

3. Data quality

In many situations we have the freedom to decide how the data is acquired in the experiments and in the real-world inference problems. This may give a way to select high-quality data rather than poor-quality one. In this paper, we consider an experimental setting where we are able to gather configurations either in-equilibrium or out-of-equilibrium. The details of the protocol are described below:

- (i) First, m initial configurations are randomly and uniformly chosen from all possible configurations;
- (ii) For each initial configuration, T steps of parallel dynamics are performed, and the final configuration is recorded. This process generates m configurations that we store in the rows of the matrix $A \in \{-1, 1\}^{m \times n}$;
- (iii) For each configuration stored in A , we do a single step of parallel dynamics, and then record the new configurations as the rows of the matrix B ;
- (iv) Lastly, we infer the couplings by maximizing the likelihood using matrix A and B .

The expression of the log-likelihood in terms of matrices A and B is

$$\mathcal{L} = \sum_{a=1}^m \left[\sum_{i,j} \beta J_{ij} B_{ai} A_{aj} - \sum_i \log(2 \cosh(\beta \sum_j J_{ij} A_{aj})) \right].$$

At the stationary point of \mathcal{L} , the derivative with respect to each coupling J_{ij} must be zero and this leads to the following moment-matching condition

$$\langle A_{ai} B_{aj} \rangle_{\text{DATA}} = \langle A_{ai} \tanh(\beta \sum_k J_{jk} A_{ak}) \rangle_{\text{DATA}}, \quad (3)$$

where the average over the data is given by $\langle \bullet \rangle_{\text{DATA}} = m^{-1} \sum_{a=1}^m \bullet$, and the following updating rule for the couplings

$$J_{ij} \leftarrow J_{ij} + \eta \left[\langle A_{ai} B_{aj} \rangle_{\text{DATA}} - \langle A_{ai} \tanh(\beta \sum_k J_{jk} A_{ak}) \rangle_{\text{DATA}} \right],$$

being η a small learning parameter.

With respect to the standard inverse kinetic Ising problem, our experimental setup has two main differences. (i) The new parameter T allows us to collect configurations both in the stationary equilibrium regime (as usual) for a large T , but also in the early out-of-equilibrium regime, for small T , where configurations are sampled according to a probability distribution different from $P_{\text{eq}}(\underline{s})$. (ii) For each initial configuration we do not save the entire trajectory but only the last two configurations, corresponding to times T and $T + 1$. This choice allows us to better understand how the data quality depends on the “distance from equilibrium”. Moreover a longer trajectory of R steps can always be seen as the union of $R - 1$ of our one-step experiments, with different T parameters.

The reasons why we expect out-of-equilibrium configurations to be of higher quality for the problem of coupling reconstruction are possibly many. First of all, since we start from m random configurations, we have that the configurations in matrix A , measured at time T , are less correlated and spanning a broader region of the configuration space, with respect to equilibrium configurations; in general, we expect correlations between configurations to increase monotonously with T . Moreover in the early out-of-equilibrium regime the dynamics usually has some drift, which is absent at equilibrium: e.g. the energy decreases towards the equilibrium value, and then stays more or less constant. Generally in the out-of-equilibrium early dynamics the system variables get updated more often, and this may lead to a sensible increase in the measured correlations and fluctuations, which are in turn exploited by the moment matching condition in Eq. (3) to infer the couplings.

We are mostly interested in studying what happens in the low temperature regime (large β) because in the low β regime correlations are weak, and many efficient methods exist for inferring the couplings. The low temperature dynamics, starting from a random initial condition, strongly depends on the kind of system under study (whether it is

homogeneous, heterogeneous, with disordered couplings, etc.); however in general the dynamics shows an initial fast relaxation, when most of the variables get updated often. On later times, the dynamics can easily get trapped in a local energy minimum, keeping oscillating around it: it is clear that this asymptotic regime is much less informative for coupling reconstruction, because only a small fraction of variables keep updating and usually in a repetitive way.

Our goal here is not to characterize in detail the behavior of the inference algorithm for different values of T and m . We are rather interested in showing that, intrinsically, out-of-equilibrium configurations contain more information than equilibrium ones. In order to show that, we will use some particular values of T and m . As said earlier, we will focus only on the two last configurations at time T and $T + 1$. We shall then fix the value of m to a particular value and look at different values of T to probe both the out-of-equilibrium regime, for small T values, and the equilibrium or steady-state regime, for large T values.

In order to study an interesting and difficult case we consider the Sherrington-Kirkpatrick (SK) model [19] where couplings J_{ij} are randomly and independently extracted from a Gaussian distribution of zero mean and variance $1/n$. We will consider small systems ($n = 20$) because our results are of a general validity and do not require to take the thermodynamical limit ($n \gg 1$).

In Figure 1 we show the results of experiments performed with $m = 10^4$ random restarts and several values of T . For each value of T and β we use a different sample (i.e. couplings), such that fluctuations in the data points are meaningful for estimating the sample-to-sample variations. In the upper left panel we show how the reconstruction error varies with β for several different values of T . We see basically two regimes where the reconstruction error behave differently. In the regime $\beta \lesssim 1$ the reconstruction error is roughly the same for any T value. This regime corresponds to the paramagnetic phase of the SK model, where ergodicity ensures that configurations remain mostly uncorrelated for any value of T , even approaching equilibrium. The regime $\beta \gtrsim 1$ corresponds to the glassy phase of the SK model, where ergodicity is broken. In this regime the out-of-equilibrium configurations (gathered at small T values) provide a clearly better quality for the inverse Ising problem, resulting to a much smaller reconstruction error. The figure also shows that the reconstruction error grows monotonously with T , thus becoming larger and larger when the dynamics brings the system close to equilibrium.

As we already discussed above, there may be several reasons for the increase of the reconstruction error when configurations are sampled closer to equilibrium. Among these, one possibility is that the m configurations sampled for $\beta > 1$ and $T > 0$ are somehow similar and thus redundant. In practice the $m \times n$ matrix A has correlated entries and we would like to measure how much one can reduce it without losing information. The simplest way to achieve this is to run an efficient algorithm for lossless compression: we use *gzip* to compress each A matrix and we measure the size of the compressed file. In upper right panel of Fig. 1, we plot the size of the compressed




Figure 1. Inference error (*upper left*) and size of zipped files (*upper right*) for several T and varying β . *Lower panel*: relation between the reconstruction error and the size of zipped files, with the black line being a power law fit with exponent 3.7. Experiments were carried out on a network with $n = 20$ spins, using $m = 10^4$ configurations.

files for several values of T and varying β . A comparison with the curves in upper left panel of the same figure suggests that the increase of the reconstruction error in the low temperature phase ($\beta \gtrsim 1$) is mostly related to the loss of information in the m configurations used to infer the couplings. On the contrary the increase of the error in the high temperature limit ($\beta \rightarrow 0$) is due to the lack of correlations among variables, that makes impossible to extract information about the couplings.

Let us call N_{eff} the effective size of the A matrix as measured from the size of the compressed file (the B matrix has approximately the same effective size). A naive expectation would be that the reconstruction error grows like $1/\sqrt{N_{\text{eff}}}$, when the effective size of the A matrix decreases. However the lower panel in Fig. 1 shows that the increase in the reconstruction error by decreasing N_{eff} is much steeper: the power law curve shown in the figure has slope 3.7, thus suggesting that the error increase also depends on other factors, like (i) the fact that equilibrium configurations evolve more slowly, and thus the number of spin flips in a one-time experiment is smaller and (ii) the presence of long-ranged spatial correlations, that grow approaching the equilibrium.

A more formal way to define the effective size of the matrix A is to compute its *effective rank*, i.e. a measure of how much correlated are the entries of the matrix

A. In principle we would like to do the principal component analysis (PCA) of the configuration matrix A , which would tell us whether there are preferred directions along which configurations tend to align. In practice we consider the eigen-decomposition of the correlation matrix C , which is defined as

$$C = \frac{1}{m} A^T A.$$

In our experiments we have $2^n \gg m \gg n$, thus C has n real eigenvalues $\{\lambda_i\}_{i \leq n}$, satisfying

$$\sum_{i=1}^n \lambda_i = \sum_{i=1}^n C_{ii} = n. \quad (4)$$

In the top panel of Fig. 2 these eigenvalues are plotted in a decreasing order for some T values. We can see that for $T = 0$, when configurations are completely random, the distribution of eigenvalues is flat and every eigenvalue is close to 1. It means that there is no particularly preferred direction and the vectors in A span uniformly the configurational space, thus providing m configurations with practically zero redundancy. However with $T = 32$, being configurations closer to equilibrium, they tend to align along a preferred direction and the first eigenvalue, corresponding to this direction, is much larger than the other eigenvalues. Hence, we see that the distribution of the eigenvalues can be used to characterize the data quality of configurations.

From Eq. (4) we know that the sum of all the eigenvalues is a constant for different sets of configurations of the same system. So the most naive way to measure how flat is the distribution of eigenvalues, is to compute the largest eigenvalue (in absolute value): the smaller the leading eigenvalue is, the flatter the distribution of the eigenvalues and therefore better the quality of the data. A more comprehensive approach is to compute the effective rank of the matrix C , defined by [20]

$$\text{rank}_{\text{eff}}(C) = \exp \left(- \sum_{i=1}^n p_i \log p_i \right),$$

where $p_k = \lambda_k / \sum_{i=1}^n \lambda_i = \lambda_k / n$. Again, an almost flat distribution of eigenvalues ($p_i \simeq 1/n$) implies the rank of C is close to n , while the effective rank decreases if the eigenvalues are very different among them.

Actually we are mostly interested in understanding how much redundant are the vectors by which the matrix of empirical correlations C is build, rather than the matrix itself. Being the eigenvalues of C real and positive, we can write $\lambda_i = \sigma_i^2$ and the following decomposition

$$C = \sum_{i=1}^n \lambda_i \underline{v}_i \underline{v}_i^T = \sum_{i=1}^n (\sigma_i \underline{v}_i) (\sigma_i \underline{v}_i)^T, \quad (5)$$

where $\{\underline{v}_i\}$ are the eigenvectors of C , forming an orthonormal basis. Eq. (5) says that the same C matrix could be obtained if the measured configurations were only equal to

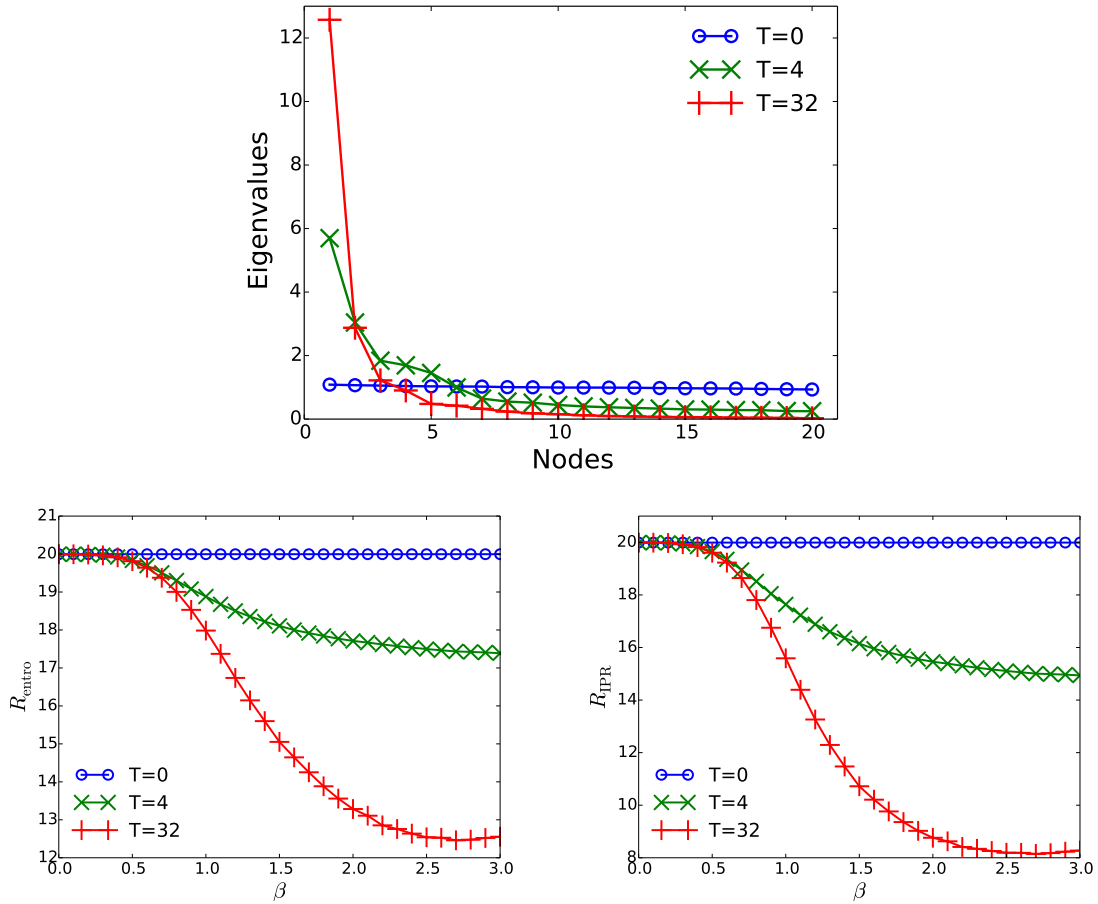


Figure 2. *Top panel:* Eigenvalues of the correlation matrix for $T = 0, 4, 32$ at $\beta = 2.2$. *Effective rank R_{entro} (left panel) and R_{IPR} (right panel) of the correlation matrix for $T = 0, 4, 32$ and varying β . Parameters are the same as in Fig. 1.*

one of the eigenvectors $\{v_i\}_{i \leq n}$, each one chosen with a probability $r_i = \sigma_i / \sum_{k=1}^n \sigma_k$. From these frequencies $\{r_i\}_{i \leq n}$ we can provide two different, but similar, definitions of the effective rank of matrix C

$$R_{\text{entro}} = \exp\left(-\sum_{i=1}^n r_i \log r_i\right),$$

$$R_{\text{IPR}} = \frac{1}{\sum_{i=1}^n r_i^2}. \quad (6)$$

The first definition is simply based on the entropy of the probability law $\{r_i\}_{i \leq n}$, while the second definition is the inverse participation ratio. Both effective ranks R_{entro} and R_{IPR} take values in $[1, n]$: they are equal to n if $r_i = 1/n$, and equal to 1 if the probability concentrates on a single value.

In the lower panels of Fig. 2 we plot the effective ranks R_{entro} and R_{IPR} for the same data used in Fig. 1. We can see that they give similar information as the size of the compressed file: the smaller the effective rank, the worse data quality. The advantage of using the effective rank R over using the size of the compressed file is that the effective

rank is easier to compute. We can thus use it as an objective function to optimize the data quality, as we will show in the next section.

4. Tuning quality of the data

In this section we study how to identify configurations that have relatively bad quality in a given set of configurations. That is those configurations that, being redundant, can be safely removed from the set without losing too much information, and thus actually improving the data quality. The idea is that if we remove a configuration from the dataset, all eigenvalues of the correlation matrix C will shift from $\{\lambda_i\}$ to $\{\lambda_i + \Delta\lambda_i\}$. Thus we can estimate the quality of each configuration in the dataset, according to the shift of the effective rank R in case that configuration is removed. We aim at removing configurations in the direction of increasing R , in order to improve the data quality of set of remaining of configurations.

Since the number of configurations is large, we can treat the effect of removing one configuration, i.e. one row in matrix A , as a perturbation to the correlation matrix C . That is, after removing the configuration \underline{s} , the change of C is

$$\Delta C = \frac{C - \underline{s}\underline{s}^T}{m - 1}.$$

Assuming that after removing the configuration \underline{s} , the i -th eigenvector of C changes from \underline{v}_i to $\underline{v}_i + \Delta\underline{v}_i$, and its associated eigenvalue changes from λ_i to $\lambda_i + \Delta\lambda_i$, then we have

$$(C + \Delta C)(\underline{v}_i + \Delta\underline{v}_i) = (\lambda_i + \Delta\lambda_i)(\underline{v}_i + \Delta\underline{v}_i).$$

Keeping only first-order terms results in

$$\Delta\lambda_i = \underline{v}_i^T \Delta C \underline{v}_i.$$

Then by making use of Eq. (6) and by keeping only the first order of $\Delta\lambda_i$, we can estimate the shift of effective rank as

$$\Delta R_{\text{IPR}} = \sqrt{\frac{R_{\text{IPR}}}{n}} \sum_i \lambda_i^{-\frac{1}{2}} \Delta\lambda_i \quad (7)$$

Then, using Eq. (7), we propose a decimation method to increase the data quality by removing iteratively configurations that provide the largest ΔR . This procedure is similar to the decimation algorithm using marginals of a message passing algorithm in solving constraint satisfaction problems [21], where nodes having most biased marginals are removed (fixed) at each iteration.

In the upper panels of Fig. 3 we plot the evolution of the effective rank as a function of fraction of configurations removed, for the decimation method just described (blue line) and for the process of removing randomly chosen configurations (red line). We clearly see that the choice based on Eq. (7) leads to an increase of the effective rank and

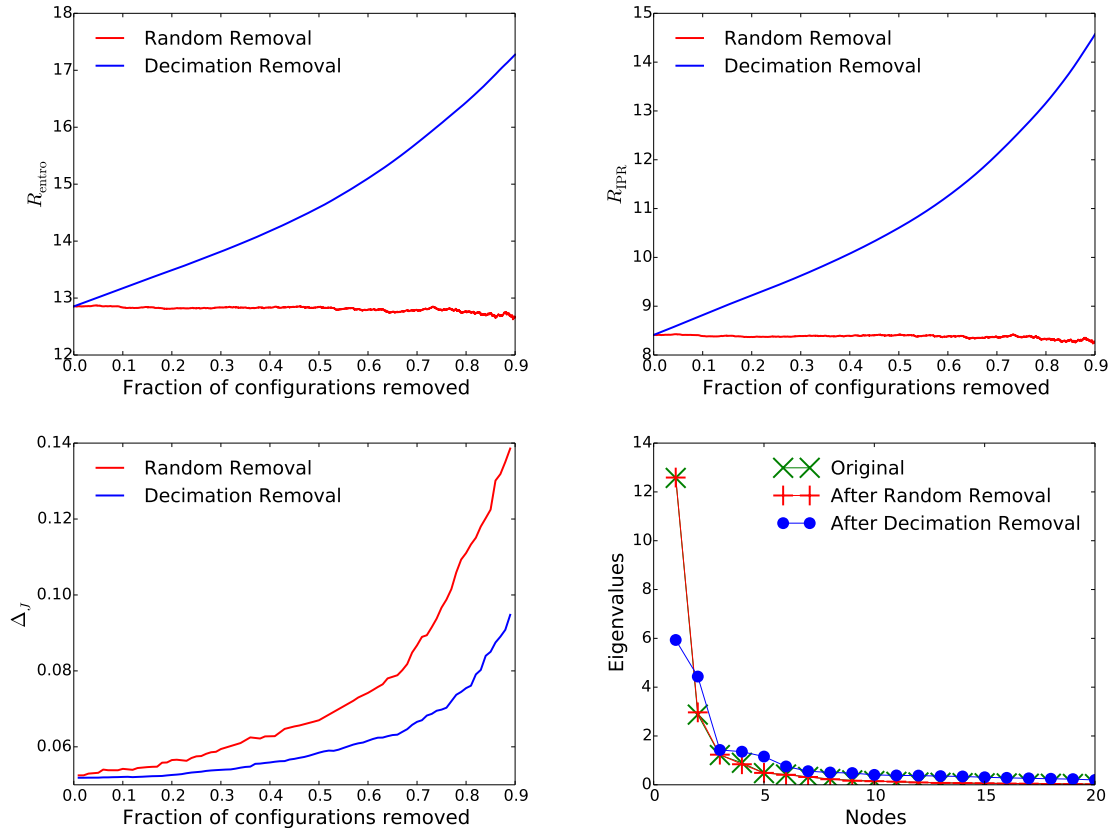


Figure 3. *Top:* Evolution of the effective rank of correlation matrix C , R_{entro} (*left*) and R_{IPR} (*right*), with the fraction of configurations removed randomly and using the decimation algorithm. *Bottom left:* Reconstruction error obtained using as input the decimated dataset. *Bottom right:* Eigenvalues sorted in decreasing order, for the original correlation matrix and for the correlation matrix after random removal and decimation-based-removal of 90% of configurations. In all of the figures, at each step of the decimation the configuration that gives the largest increase in the effective rank among 100 randomly sampled configurations is removed from the dataset.

consequently of the data quality, while the random decimation keeps the effective rank roughly unchanged (actually there is a small, but systematic decrease; notice that ΔR can be negative). As a consequence, if we infer the couplings starting from the decimated dataset, whose size is α times the original one, we obtain a reconstruction error that grows as $1/\sqrt{\alpha}$ for the random decimation (see red line in the lower left panel of Fig. 3). On the contrary, the dataset decimated according to our new rule returns a much smaller reconstruction error. For example, we are able to reduce by a factor 10 the size of the dataset by increasing by less than a factor 2 the reconstruction error (see blue line in the lower left panel of Fig. 3). The lower right panel of Fig. 3 shows the eigenvalues of the C matrix when the dataset is reduced by a factor 10 ($\alpha = 0.1$). The random decimation process keeps the eigenvalues practically unchanged, while our decimation algorithm strongly reduces the largest eigenvalues, thus decreasing the redundancy of the dataset.

5. Conclusion and discussion

We have studied the data quality problem of the kinetic inverse Ising problem. First, we have experimentally shown that data gathered in an out-of-equilibrium way has better quality thus leads to a smaller reconstruction error than data gathered in equilibrium. Then we focused on how to quantify the data quality using the effective rank of the correlation matrix, and how to improve the data quality by a decimation procedure based on a perturbative analysis of the correlation matrix.

Though we only studied the SK model in this paper, we have tested other disordered models such as the Hopfield model and the sequence processing neural networks, where the results are qualitatively similar.

In all of the experiments, the inference is done by maximizing the likelihood, which is usually prone to overfitting, especially when the number of configurations is not large enough. It would be interesting to check our results and the decimation method for Bayesian inference of the model parameters. We leave this for future work.

We believe the results of the present work can be very useful in applications. On the one hand, we have shown that data collected in a strongly out-of-equilibrium regime are much more informative about the interaction network of a set of dynamical variables, and this may suggest new ways to collect the data to infer such an interaction network: for example, perturbing the system out of its equilibrium/stationary regime may allow the system to show up more clearly correlations and fluctuations, that are useful for the reconstruction problem. On the other hand, given that the amount of data available is growing very fast in recent years, the decimation method we have proposed for strongly reducing the size of the input dataset, without losing too much information for the reconstruction problem, may be extremely practical for dealing with huge datasets. Notice that the method can be used also on-the-run, i.e. while data are being generated: in this case one can accept only configurations that bring a substantial improvement in the effective rank, and leaving aside redundant data.

References

- [1] David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. A learning algorithm for boltzmann machines. *Cognitive science*, 9(1):147–169, 1985.
- [2] Elad Schneidman, Michael J. Berry, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440:1007–1012, 2006.
- [3] Erik Aurell and Magnus Ekeberg. Inverse ising inference using all the data. *Physical review letters*, 108(9):90201, 2012.
- [4] H. Chau Nguyen and J. Berg. Bethe-peierls approximation and the inverse ising model. *arXiv*, 1112.3501, 2011.
- [5] H. Chau Nguyen and J. Berg. Mean-field theory for the inverse ising problem at low temperatures. *arXiv*, 1204.5375, 2012.
- [6] Federico Ricci-Tersenghi. The bethe approximation for solving the inverse ising problem: a comparison with other inference methods. *J. Stat. Mech.*, 2012(08):P08015, 2012.
- [7] Aurélien Decelle and Federico Ricci-Tersenghi. Pseudolikelihood decimation algorithm improving

- the inference of the interaction network in a general class of ising models. *Phys. Rev. Lett.*, 112:070603, Feb 2014.
- [8] B. Ravikumar, M.J. Wainwright, and J.D. Lafferty. High-dimensional ising model selection using l1-regularized logistic regression. *The Annals of Statistics*, 38, 2010.
 - [9] Yasser Roudi, Joanna Tyrcha, and John Hertz. Ising model for neural data: Model quality and approximate methods for extracting functional connectivity. *Phys. Rev. E*, 79:051915, May 2009.
 - [10] Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(3):75–174, 2010.
 - [11] Gonzalo A Ruz and Eric Goles. Learning gene regulatory networks with predefined attractors for sequential updating schemes using simulated annealing. In *Machine Learning and Applications (ICMLA), 2010 Ninth International Conference on*, pages 889–894. IEEE, 2010.
 - [12] Martin Weigt, Robert A White, Hendrik Szurmant, James A Hoch, and Terence Hwa. Identification of direct residue contacts in protein–protein interaction by message passing. *Proceedings of the National Academy of Sciences*, 106(1):67–72, 2009.
 - [13] Magnus Ekeberg, Cecilia Lökvist, Yueheng Lan, Martin Weigt, and Erik Aurell. Improved contact prediction in proteins: Using pseudolikelihoods to infer potts models. *Physical Review E*, 87(1):012707, 2013.
 - [14] William Bialek, Andrea Cavagna, Irene Giardina, Thierry Mora, Edmondo Silvestri, Massimiliano Viale, and Aleksandra M Walczak. Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences*, 109(13):4786–4791, 2012.
 - [15] Pan Zhang. Inference of kinetic ising model on sparse graphs. *Journal of Statistical Physics*, 148(3):502–512, 2012.
 - [16] Aurélien Decelle and Pan Zhang. Inference of the sparse kinetic ising model using the decimation method. *Phys. Rev. E*, 91:052136, May 2015.
 - [17] Jack Raymond and Federico Ricci-Tersenghi. Mean-field method with correlations determined by linear response. *Physical Review E*, 87(5):052111, 2013.
 - [18] Yasser Roudi and John Hertz. Mean field theory for nonequilibrium network reconstruction. *Physical review letters*, 106(4):048702, 2011.
 - [19] David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Physical review letters*, 35(26):1792, 1975.
 - [20] Olivier Roy and Martin Vetterli. The effective rank: A measure of effective dimensionality. In *European signal processing conference (EUSIPCO)*, number LCAV-CONF-2007-017, pages 606–610, 2007.
 - [21] M. Mézard, G. Parisi, and R. Zecchina. Analytic and algorithmic solution of random satisfiability problems. *Science*, 297:812, 2002.