Statistical inference with correlated disorder: the planted spin glass

Andrea Vincenzo Dell'Abate

Alfredo Braunstein
Louise Budzynski

Master's degree in Physics of Complex Systems
Academic Year: 2024-25

Contents

Abstract

Inference problems defined on random graphs can be studied by using statistical physics methods of disordered systems. In particular, these methods allows us to find the phase transitions characterizing the feasibility of said problems. A largely unexplored setting of Bayesian inference is that of structured signal, as most theoretical work studies models where the signal is made up of i.i.d. components. We study the planted spin glass model in the case of structured signal using the cavity method and compare the results with those obtained by performing inference via Belief Propagation on generated instances of the problem. We present the results in the form of a phase diagram, in which we can distinguish two phases: one in which inference is impossible and one where inference is easy.

I Introduction

An inference problem consists of reconstructing a signal, law or pattern from partial and/or noisy observations (or data). In Bayesian inference, this task is achieved by weighting a prior belief on the signal with the likelihood of data. Inference problems appear ubiquitously in many scientific areas involving data analysis, such as signal processing [1], artificial intelligence [2], computational biology [3] and epidemiology [4, 5]. A central question in Bayesian inference is to assess under which conditions the information contained in the observations is enough to reconstruct the signal. Such an information-theoretic point of view should also be complemented by an algorithmic perspective: understanding what are the most efficient algorithms, and what is the best performance achievable computationally.

Although some recent efforts have focused on the theoretical study of structured datasets [6, 7, 8], most theoretical work on Bayesian Inference model inputs as i.i.d. component-wise draws from some probability distribution. Despite providing valuable insights, these approaches are blind to the structure of real-world data, where one cannot assume that randomness enters in an uncorrelated way. Determining whether the presence of structure in the signal will be beneficial or detrimental to the inference process is a priori an open question and will depend on the particular problem to be considered. While it might make inference easier, as the knowledge of some structure in the signal can help to restrict the search to a smaller space, it might as well make the problem harder from an algorithmic viewpoint, as observed in [5, 9] when an ergodicity breaking phenomenon prevents numerical methods (such as Simulated Annealing, or message-passing approaches) to provide a reasonable approximation of the posterior probability distribution of the signal given the data. In many problems of interest, the signal to reconstruct and the observations are high-dimensional objects, making the theoretical analysis challenging. In such settings, tools from statistical physics, in particular the replica and cavity methods [10, 11], led to a detailed description of the information-theoretical and algorithmic limits in many Bayesian inference problems [12, 13], predicting important properties of inference problems, in particular computational-to-statistical gaps, i.e. regimes where reconstructing the signal is information-theoretically possible although no efficient algorithm exists. Many of these theoretical predictions were confirmed rigorously later on [14, 15].

Of particular interest are minimal - analytically tractable - models, in which the signal's structure can be introduced in a controlled way and its effects on the inference task can be quantified precisely. One such example is the *planted spin glass*, whose theoretical and computational limits are well understood in the case of unstructured signal [16]. This study focuses on this model, by studying its information-theoretic and algorithmic properties in the case in which signal and observations are structured, focusing on the Bayes optimal setting in which the parameters are known. We provide a quantitative analysis of the feasibility of the inference task in the form of a phase diagram depending on the two parameters of the model: κ which quantifies the structure of the signal and β quantifying the signal-to-noise ratio. The analysis is comprised of both a theoretical and a numerical part: in the former we studied the properties of the posterior distribution in typical realizations of the inference problem, focusing on the thermodynamic limit in order to use the Replica Symmetric cavity method, while in the latter we performed numerical resolutions of the cavity equations and compared with the results obtained by running Belief Propagation on finite size instances of the problem. From analytical calculations we were able to find a phase diagram in accordance to the numeric results. In particular, we can distinguish two regions (or phases): easy and impossible. In the former inference is information-theoretically possible, and an algorithm such as Belief Propagation is able to achieve the optimal performance achieved with a perfect knowledge of the posterior probability distribution. In the latter, inference better than random guessing from the prior probability distribution is not possible, even with perfect knowledge of the posterior distribution. We are able to reproduce the critical value β_c separating the easy and impossible phases, respectively $\beta > \beta_c$ and $\beta < \beta_c$, for the unstructured signal, i.e. when $\kappa = 0$, and we show that β_c decreases when the signal is structured. This supports our intuition that introducing structure in this model aids inference. Notable exceptiona are given by the cases of values of κ lower than $-\kappa_{KS}$ and higher than κ_c. In the former the system exhibits an RS to RSB transition, indicating that the inference problem may be hard, i.e. still informational-theoretically feasible but without any algorithm able to find the signal faster than exponential time. In the latter, the critical value β_c increases with κ, eventually becoming larger than its value for $\kappa = 0$, implying that inference is impeded by the structure of the signal.

The thesis is organized as follows. In Section II we define the model under study. In Section III we present the Belief Propagation algorithm which can be applied to single instances of our inference problem, while in Section IV we derive the cavity equations obtained from the replica-symmetric cavity method for a random ensemble of instances. In Section V we compute analytically the phase diagram of the problem by studying the stability of the trivial fixed point of the cavity equations. The main results are presented

in Section VI. By solving numerically the cavity equations with different initial conditions we are able to predict both the information-theoretical optimal performance and the algorithmic performance and we recover the phase diagram computed analytically. We then compare the results with those obtained from the BP algorithm on single instances in order to have numerical confirmations of our predictions. Finally, we discuss the results and present our perspectives in Section VII.

II Planted Spin glass model

A The model

We consider the planted spin glass model defined on a graph $G = (V, E)$. To each node $i \in V = \{1, \dots, N\}$ we assign a spin variable $s_i \in \{\pm 1\}$ in such a way that the planted configuration $\underline{s} = \{s_i\}_{i=1,\dots,N}$ as a whole is sampled from the Ising model with coupling constant κ :

$$P_\kappa(\underline{s}) = \frac{1}{Z(G,\kappa)} \prod_{(i,j)\in E} e^{\kappa s_i s_j} \tag{1}$$

For each edge (i, j) we define the couplings $J_{ij} \in \{\pm 1\}$. These are the observations of our inference problem and are drawn from the probability distribution

$$P(J_{ij}|s_i, s_j) = \rho\delta(J_{ij} - s_i s_j) + (1 - \rho)\delta(J_{ij} + s_i s_j) \tag{2}$$

The couplings tell us whether the spins s_i and s_j are aligned with a probability ρ that the information they give us is truthful. It is easy to see that $\rho = 1$ (and $\rho = 0$ by symmetry) corresponds to the noiseless case in which all observations are truthful, while $\rho = 1/2$ corresponds to the case in which the observations are pure noise and give no information at all about the signal.

Without loss of generality we can define the parameter β:

$$\beta = \frac{1}{2} \ln\left(\frac{\rho}{1-\rho}\right) \tag{3}$$

which allows us to write the probability distribution for sampling the couplings in a way more similar to the Ising model:

$$P_\beta(J_{ij}|s_i, s_j) = \frac{e^{\beta J_{ij} s_i s_j}}{2\cosh\beta} \tag{4}$$

and therefore

$$P_\beta(\underline{J}|\underline{s}) = \prod_{ij\in E} \frac{e^{\beta J_{ij} s_i s_j}}{2\cosh(\beta)} \tag{5}$$

The inference task is that of reconstructing the planted configuration $\underline{s}$, i.e. our signal, starting from the couplings $\underline{J}$, i.e. our observations. The signal is sampled from the probability distribution (1) and will therefore be structured: the spins s_i of the planted configuration $\underline{s}$ are not independent. An example of unstructured signal is the following

$$P(\underline{s}) = \prod_{i=1}^{N} \left(\frac{1}{2}\delta(s_i - 1) + \frac{1}{2}\delta(s_i + 1)\right) \tag{6}$$

which is the one used in [16].

B Bayesian inference and Bayes optimality

To perform this task we utilize the approach of Bayesian inference. Using Bayes' rule we are able to write the posterior distribution of the planted configuration:

$$P_{\beta,\kappa}(\underline{\sigma}|\underline{J}) = \frac{P_\beta(\underline{J}|\underline{\sigma})P_\kappa(\underline{\sigma})}{P(\underline{J})} \tag{7}$$

in which we identify

- the **prior** $P_\kappa(\underline{\sigma})$ which reflects our belief of what the planted configuration should be before receiving any information from the observations;

- the **likelihood** $P_\beta(\underline{J}|\underline{\sigma})$ which gives us the compatibility of the observed couplings $\underline{J}$ with a given configuration $\underline{\sigma}$.

Substituting (1) and (5) into the expression we can write the posterior as

$$P_{\beta,\kappa}(\underline{\sigma}|\underline{J}) = \frac{\prod_{ij\in E}\psi_{ij}(\sigma_i,\sigma_j)}{Z(\kappa,\beta,G,\underline{J})} \tag{8}$$

where

$$\psi_{ij}(\sigma_i,\sigma_j) = e^{(\beta J_{ij}+\kappa)\sigma_i\sigma_j} \tag{9}$$

The parameters The meaning of the two parameters is quite clear: κ quantifies the structure in the signal, while β quantifies the signal-to-noise ratio.

It is important to touch briefly upon the role which κ and the signal's structure play in the inference task. Let us first of all consider the case of $\kappa > 0$, i.e. the ferromagnetic Ising model. It is known [17] that for random c-regular graphs a ferromagnetic phase transition occurs at

$$\kappa_c = \operatorname{atanh}\left(\frac{1}{c-1}\right) \tag{10}$$

In the ferromagnetic phase the planted configuration has non vanishing magnetization and one is able to obtain, even without any information from the observations ($\beta = 0$), a better result than simply guessing randomly each spin with probability one half. This can be done by simply guessing all spins to be equal to each other. One could therefore say that the inference has become easier. However this improved performance comes from the prior (1) and not the likelihood (5). If one defines successful inference as the posterior (8) containing more information than the prior (1), an increase of κ, which amounts to an increase structure of the signal, does not aid inference. On the contrary, one might be led to believe the opposite: most of the information contained in the posterior comes from the prior and a larger signal-to-noise ratio β is needed to balance it with the knowledge coming from the observations J_{ij}. The extreme case corresponds to $\kappa \to \infty$ in which all spins of the planted configuration have the same value. The inference task becomes trivial: the observations give us no additional information and to guess the planted one simply guesses with probability one half what will be the values of all the spins.

Moving our focus to the anti-ferromagnetic case, $\kappa < 0$, an RS to RSB phase transition occurs at

$$-\kappa_{\mathrm{KS}} = -\operatorname{atanh}\left(\frac{1}{\sqrt{c-1}}\right) \tag{11}$$

as discussed in [18]. Not much can be said intuitively on the implications on the inference problem. The consequence of most relevance is the fact that in this regime the replica-symmetric cavity method is not suitable anymore.

Finally, even if we restrict ourselves to the RS paramagnetic phase of the planted configuration, $\kappa \in [-\kappa_{\mathrm{KS}}, \kappa_c]$, we show in Section V that the structure aids the inference process: the critical value of β separating the easy, for $\beta > \beta_c$ and impossible, for $\beta < \beta_c$, phases is at its highest when $\kappa = 0$ as shown in Figure 2a.

Bayes optimality It is important to note that there are two possible settings which we can consider: the case in which the parameters κ and β are known when performing the inference task and the one in which they are not. In the latter case the parameters need to be estimated, while in the former the parameters κ and β used in the posterior probability are the same as the true parameters used to generate the observations. This is called Bayes optimal setting and is the one which we will focus on.

Let f be a generic function, $\underline{\sigma}, \underline{\sigma}_1, \underline{\sigma}_2$ three independent samples generated from the posterior (8) and $\underline{s}$ the planted configuration. In the Bayes optimal setting the following expectation values coincide

$$\mathbb{E}[f(\underline{\sigma}_1, \underline{\sigma}_2)] = \mathbb{E}[f(\underline{\sigma}, \underline{s})]. \tag{12}$$

In other words, when taking expectation values there is no statistical difference between the planted configuration $\underline{s}$ and a configuration sampled from the posterior distribution. This property is called

Nishimori condition. This equality can be easily proven for our model

$$\mathbb{E}[f(\underline{\sigma}_1, \underline{\sigma}_2)] = \sum_{\underline{\sigma}_1, \underline{\sigma}_2, \underline{J}} f(\underline{\sigma}_1, \underline{\sigma}_2) P_{\beta,\kappa}(\underline{\sigma}_1|\underline{J}) P_{\beta,\kappa}(\underline{\sigma}_2|\underline{J}) P(\underline{J}) \tag{13}$$

$$\begin{aligned}\mathbb{E}[f(\underline{\sigma}, \underline{s})] &= \sum_{\underline{\sigma}, \underline{s}, \underline{J}} f(\underline{\sigma}, \underline{s}) P_{\beta,\kappa}(\underline{\sigma}|\underline{J}) P_{\beta}(\underline{J}|\underline{s}) P_{\kappa}(\underline{s}) \\ &= \sum_{\underline{\sigma}, \underline{s}, \underline{J}} f(\underline{\sigma}, \underline{s}) P_{\beta,\kappa}(\underline{\sigma}|\underline{J}) P_{\beta,\kappa}(\underline{s}|\underline{J}) P(\underline{J})\end{aligned} \tag{14}$$

This property is of particular interest because it was conjectured that the Nishimori condition implied the absence of an RSB phase [16]. This equality can be applied for any function $f(\sigma, \sigma')$ of two configurations, in particular it can be applied for any moment of the overlap. When all the moments of two random variables coincide, it holds that their probability distributions coincide. When considering an unstructured signal, the spin glass presents a gauge invariance, which can be used to show that the overlap $\mathrm{O}(\underline{s}, \underline{\sigma})$ between the planted configuration and a configuration sampled from the posterior coincides with the magnetization, a quantity which is argued to be self-averaging. In particular by applying the gauge transformation

$$\sigma_i s_i \to \tilde{\sigma}_i, \qquad J_{ij} s_i s_j \to \tilde{J}_{ij}, \tag{15}$$

the overlap becomes

$$\mathrm{O}(\underline{s}, \underline{\sigma}) = \frac{1}{N} \sum_i \tilde{\sigma}_i = m \tag{16}$$

One can therefore argue that the overlap $\mathrm{O}(\underline{\sigma}, \underline{\sigma}')$ is self-averaging and therefore has a trivial probability distribution, which in turn implies the absence of a replica symmetry breaking phase. A complementary argument in favor of the absence of an RSB (or spin glass) phase was presented by Montanary in [19], in which he proved rigorously that in sparse systems in the Bayes optimal setting two points correlations decay, which does not occur in the spin glass phase. It is important to note that both these arguments have been made when studying systems in which the signal is not structured: it is not clear that in such cases the overlap $\mathrm{O}(\underline{s}, \underline{\sigma})$ is self averaging. Finally, it is important to note that this conjecture only applies to static replica symmetry breaking and does not exclude the presence of the dynamical one-step replica symmetry breaking phase (d1RSB). In this phase, however, the marginals obtained through the belief propagation algorithm, presented in Section III, exactly describe the marginals of the d1RSB phase, which is what is of main interest in Bayesian inference in the optimall setting.

C Ensemble averages

The objective is to estimate how well an inferred configuration computed from the knowledge of the posterior (8), approximates the planted configuration $\underline{s}$ in a typical instance. To do so, we introduce a random ensemble of instances, called the planted ensemble. To create an instance we sample a random graph from the random c-regular ensemble and then we follow the steps in Section II.A to generate the planted configuration and the couplings.

Random c-regular graphs, along with graphs from other ensembles, have the following crucial property: in the thermodynamic limit they converge locally to trees. This is particularly important because it allows us to utilize the cavity method to study the typical properties of the posterior.

D Overlaps and Estimator

To quantify feasibility of the inference task we are now going to introduce the estimator and other quantities used in the results section.

Maximum Mean Overlap estimator To quantify how well an estimator has managed to infer the planted configuration we look at the fraction of variables the two have in agreement, the overlap.

We therefore define the overlap between the planted configuration $\underline{s}$ and an estimator $\hat{\underline{\sigma}}$ as

$$\mathrm{O}(\underline{s}, \hat{\underline{\sigma}}) = \frac{1}{N} \sum_{i=1} \delta(s_i, \hat{\sigma}_i) \tag{17}$$

When considering the case of a single instance of the inference problem, the planted configuration is not known. We consider instead the best Bayesian estimate, the Mean Overlap, obtained by averaging

over the posterior distribution

$$\mathrm{MO}(\hat{\underline{\sigma}}) = \frac{1}{N} \sum_{\underline{\sigma}} P_{\beta,\kappa}(\underline{\sigma}|\underline{J}) \sum_{i=1}^{N} \delta_{\sigma_i,\hat{\sigma}_i} \tag{18}$$

The *Maximum Mean Overlap* estimator is the one maximizing such quantity and is achieved for

$$\hat{\sigma_i}^{\mathrm{MMO}} = \arg\max_{\sigma_i} \mu_i(\sigma_i) \tag{19}$$

where $\mu_i(\sigma_i)$ is the marginal distribution of node i. The overlap between the MMO estimator and the planted configuration provides a quantitative estimation of the accuracy of the estimator. Furthermore, when averaging over the disorder $\mathbb{E}_{\underline{s},\underline{J}}[\mathrm{O}(\underline{s},\hat{\underline{\sigma}})]$ we obtain information regarding the feasibility of the problem.

In the case of unstructured signal, when this quantity is equal to one half inference is impossible: the estimator is effectively guessing each spin randomly. For values different than 0.5 inference is feasible.

In the case of structured signal however, the planted configuration undergoes a ferromagnetic transition. For $\kappa > \kappa_c$, one can obtain values of the overlap (17) different than 0.5 by simply guessing all the spins to be the same. To quantify whether the posterior probability (8) gives us any additional information with respect to the prior probability (1) we need to define a new overlap.

More formally, for $\kappa > \kappa_c$ the marginal probability of the prior probability (1) becomes, for random regular graphs, equal to

$$b_i(s_i) = \frac{1 + s_i\gamma(\kappa)}{2} \tag{20}$$

where $\gamma(\kappa)$ is the non-vanishing magnetization. We define the random estimator as

$$\hat{\sigma}_i^r = \arg\max_{\sigma_i} b_i(\sigma_i) \tag{21}$$

Its overlap with the planted configuration is

$$\mathrm{O}(\underline{s},\hat{\underline{\sigma}}^r) = \frac{1}{N}\sum_{i=1} \delta(s_i,\hat{\sigma}_i^r) = \frac{1+\gamma(\kappa)}{2} \tag{22}$$

We can now define the rescaled overlap between the planted configuration $\underline{s}$ and an estimator $\hat{\underline{\sigma}}$ as

$$\tilde{\mathrm{O}}(\underline{s},\hat{\underline{\sigma}}) = \frac{\mathrm{O}(\underline{s},\hat{\underline{\sigma}}) - \mathrm{O}(\underline{s},\hat{\underline{\sigma}}^r)}{1 - \mathrm{O}(\underline{s},\hat{\underline{\sigma}}^r)} \tag{23}$$

where the denominator is such that the rescaled overlap is equal to zero when $\mathrm{O}(\underline{s},\hat{\underline{\sigma}}) = \mathrm{O}(\underline{s},\hat{\underline{\sigma}}^r)$ and equal to one when $\mathrm{O}(\underline{s},\hat{\underline{\sigma}}) = 1$.

We can use this definition to distinguish when the inference problem is feasible or not. When the average over the disorder $\mathbb{E}_{\underline{s},\underline{J}}[\tilde{\mathrm{O}}(\underline{s},\hat{\underline{\sigma}})]$ is equal to zero inference is not possible: perfect knowledge of the posterior gives us the same result as randomly guessing from the prior. When instead this quantity is greater than zero inference is feasible. It is important to note that the rescaled overlap can, in principle, also be negative: this would correspond to the case in which an estimator performs worse than the random estimator (21) obtained with the knowledge of the prior. We can ignore this case as the posterior (8) always contains the information stored in the prior (1).

It is important to note that for $\kappa < \kappa_c$ overlap and rescaled overlap are equivalent: it can be easily shown that $\mathrm{O}(\underline{s},\hat{\underline{\sigma}}^r) = 1/2$ which implies $\tilde{\mathrm{O}}(\underline{s},\hat{\underline{\sigma}}) = 2\,\mathrm{O}(\underline{s},\hat{\underline{\sigma}}) - 1$.

Nishimori conditions Let us consider the overlap between two arbitrary configurations

$$O(\underline{\sigma},\underline{\sigma}') = \frac{1}{N}\sum_{i=1}^{N} \delta(\sigma_i,\sigma_i') \tag{24}$$

Similarly to (23) we define the rescaled overlap

$$\tilde{\mathrm{O}}(\underline{\sigma},\underline{\sigma}') = \frac{\mathrm{O}(\underline{\sigma},\underline{\sigma}') - \mathrm{O}^r(\underline{\sigma},\underline{\sigma}')}{1 - \mathrm{O}^r(\underline{\sigma},\underline{\sigma}')} \tag{25}$$

with respect to the average over the prior (1):

$$\mathrm{O}^r(\underline{\sigma},\underline{\sigma}') = \frac{1}{N}\sum_{\underline{\sigma}} P_\kappa(\underline{\sigma})P_\kappa(\underline{\sigma}')\sum_{i=1}^{N}\delta_{\sigma_i,\sigma_i'} = \frac{1+\gamma^2(\kappa)}{2} \tag{26}$$

Let us consider the average over the posterior of the rescaled overlap between two configurations sampled from the posterior and the rescaled overlap between a configuration sampled from the posterior and the planted configuration:

$$\mathbb{E}_{\underline{\sigma},\underline{\sigma}'}[\tilde{\mathrm{O}}(\underline{\sigma},\underline{\sigma}')] = \sum_{\underline{\sigma},\underline{\sigma}'} P_{\beta,\kappa}(\underline{\sigma}|\underline{J})P_{\beta,\kappa}(\underline{\sigma}'|\underline{J})\,\tilde{\mathrm{O}}(\underline{\sigma},\underline{\sigma}') \tag{27}$$

$$\mathbb{E}_{\underline{\sigma},\underline{s}}[\tilde{\mathrm{O}}(\underline{\sigma},\underline{s})] = \sum_{\underline{\sigma},\underline{s}} P_{\beta,\kappa}(\underline{\sigma}|\underline{J})P_{\kappa}(\underline{s})\,\tilde{\mathrm{O}}(\underline{\sigma},\underline{s}) \tag{28}$$

We denote the first by "rescaled overlap config" and the second by "rescaled overlap planted".

In the Bayes optimal setting their average over the disorder must coincide (Nishimori condition):

$$\mathbb{E}_{\underline{J},\underline{\sigma},\underline{\sigma}'}[\tilde{\mathrm{O}}(\underline{\sigma},\underline{\sigma}')] = \mathbb{E}_{\underline{J},\underline{\sigma},\underline{s}}[\tilde{\mathrm{O}}(\underline{\sigma},\underline{s})] \tag{29}$$

We will use this property when performing numerical simulations to check whether they satisfy these conditions or not.

III Belief Propagation

In order to utilize the Maximum Mean Overlap estimator (19) we need to compute the marginals of the posterior distribution (8).

Suppose the graph G taken in consideration is a tree, one can calculate exactly both the partition function $Z(\kappa,\beta,G,\underline{J})$ and the marginals of the posterior through Belief Propagation (BP).

We define the message $\nu_{i\to j}(\sigma_i)$ as the marginal probability distribution of σ_i when the edge $i\to j$ is removed. When the graph is a tree, it can be shown that the messages obey the following rule:

$$\nu_{i\to j}(\sigma_i) = \frac{1}{Z_{i\to j}}\prod_{k\in\partial i\setminus j}\sum_{\sigma_k}\psi_{ik}(\sigma_i,\sigma_k)\nu_{k\to i}(\sigma_k) \tag{30}$$

where ∂i denotes the neighboring nodes of i and $Z_{i\to j}$ is the normalization factor

$$Z_{i\to j} = \sum_{\sigma_i}\prod_{k\in\partial i\setminus j}\sum_{\sigma_k}\psi_{ik}(\sigma_i,\sigma_k)\nu_{k\to i}(\sigma_k) \tag{31}$$

Let us define the shorthand notation

$$\nu_{i\to j}(\sigma_i) = f^{\mathrm{BP}}\Big(\{\nu_{k\to i},J_{ik}\}_{k\in\partial i\setminus j}\Big) \tag{32}$$

where we explicited the dependence on the couplings J_{ik}, present in the factors ψ_{ik}.

From these messages we can obtain the marginal probability of spin σ_i in the full graph:

$$\mu_i(\sigma_i) = \frac{1}{Z_i}\prod_{k\in\partial i}\sum_{\sigma_k}\psi_{ik}(\sigma_i,\sigma_k)\nu_{k\to i}(\sigma_k) \tag{33}$$

Finally, let us define the Bethe free entropy

$$\phi_B = \frac{1}{N}\sum_{i\in V}\ln Z_i - \frac{1}{N}\sum_{(i,j)\in E}\ln Z_{ij} \tag{34}$$

where

$$Z_i = \sum_{\sigma_i}\prod_{k\in\partial i}\sum_{\sigma_k}\psi_{ik}(\sigma_i,\sigma_k)\nu_{k\to i}(\sigma_k) \quad\text{and}\quad Z_{ij} = \sum_{\sigma_i,\sigma_k}\psi_{ij}(\sigma_i,\sigma_j)\nu_{i\to j}(\sigma_i)\nu_{j\to i}(\sigma_j) \tag{35}$$

When the graph is a tree the fixed point of the BP equations is unique, the resulting marginals are correct and the Bethe free entropy coincides with the free entropy $\phi = \frac{1}{N}\ln Z$.

In our specific case we assume the graph G to be sampled from the regular random graph ensemble with small enough degree c. In the thermodynamic limit, sparse random graphs are locally tree like. In these conditions one can still run the Belief Propagation algorithm and expect it to converge.

A Generating the planted configuration

One major consequence of the structure of the planted configuration appears when generating one on a given graph G. One naive way of doing it would be the following: randomly choose a node i of the graph, assign to s_i a value taken from the uniform distribution $P(s_i) = 1/2$ and then sampling all other spins following the conditional probability

$$P(s'|s) = \frac{e^{\kappa s s'}}{2\cosh\kappa} \tag{36}$$

However, in the ferromagnetic phase ($\kappa > \kappa_c$) the system presents two equilibrium states with opposite, non-vanishing, magnetization. When following this procedure, the generated planted configuration has vanishing magnetization: by not expliciting in which of the two states the system finds itself in, we are effectively averaging over both them.

We can solve this problem by recurring once again to BP. In this case however, whether we are working with a tree or not matters. Let us start by considering the former.

We introduce the "planted messages" $m_{i\to j}(s_i)$, which satisfy the BP equations

$$m_{i\to j}(s_i) = g^{\mathrm{BP}}(\{m_{k\to i}\}_{k\in\partial i\setminus j}) \tag{37}$$

where

$$g^{\mathrm{BP}}(\{m_{k\to i}\}_{k\in\partial i\setminus j}) = \frac{\prod_{k\in\partial i\setminus j}\sum_{s_k} e^{\kappa s_i s_k} m_{k\to i}(s_k)}{\sum_{s_i}\prod_{k\in\partial i\setminus j}\sum_{s_k} e^{\kappa s_i s_k} m_{k\to i}(s_k)}, \tag{38}$$

Now, when sampling the first spin s_i we do so by using the marginal probability

$$b_i(s_i) = \frac{\prod_{j\in\partial i}\sum_{s_j} e^{\kappa s_i s_j} m_{j\to i}(s_j)}{\sum_{s_i}\prod_{j\in\partial i}\sum_{s_j} e^{\kappa s_i s_j} m_{j\to i}(s_j)} \tag{39}$$

which depends solely on the incoming messages $m_{j\to i}$. We can sequentially sample all other spins by using

$$b_k(s_k|s_i) = \sum_{\underline{s}_{\partial i\setminus k}}\prod_{j\in\partial i}\frac{e^{\kappa s_i s_j} m_{j\to i}(s_j)}{\sum_{s_j} e^{\kappa s_i s_j} m_{j\to i}(s_j)} = \frac{e^{\kappa s_i s_k} m_{k\to i}(s_k)}{\sum_{s_k} e^{\kappa s_i s_k} m_{k\to i}(s_k)} \tag{40}$$

starting with the neighbors of spin s_i, then sampling the spins at distance $n = 2$ and recursively reaching the leaves of the graph. When the system is in the paramagnetic phase, $m_{i\to j}(s_i) = 1/2$, this procedure coincides with the one introduced at the beginning. When the system is instead in the ferromagnetic phase this is not the case and the message tells us whether the system has positive or negative magnetization.

When the graph is not a tree, and thus has loops, the algorithm does not explore all edges of the graph, resulting in wrong values of the correlation between neighboring spins.

In this case we can resort to BP-guided decimation [20]. We randomly choose a node i in the graph and sample the spin s_i using the marginal of the prior probability (1). In particular, we use Belief Propagation's estimate (39). We define the two sets

$$S = \{i \in V : s_i \text{ fixed}\} \quad \text{and} \quad R = V\setminus S \tag{41}$$

We then choose randomly a node j among the remaining nodes R and sample s_j following the marginal probability of $P_\kappa(\{s_k\}_{k\in R})|\{s_k\}_{k\in S})$. To estimate this marginal probability we run BP again, this time adding for each $i \in V$ a factor node enforcing the value s_i. We add j to S and repeat the procedure until $S = V$.

This procedure is correct if one is able to exactly compute the marginal probabilities. In that case the procedure is equivalent to decomposing the prior (1) using the chain rule. In BP-guided decimation we compute these marginals using Belief Propagation, leading to a good approximation.

The pseudo-code is shown in Algorithm 1.

IV Replica-symmetric cavity method

Consider the following thought experiment: suppose to generate a graph G from the random c-regular graph ensemble, sample the planted configuration $\underline{s}$ from the prior probability (1) and the couplings $\underline{J}$ from the likelihood probability (5) and to find the fixed point of the Belief Propagation equation; suppose

Algorithm 1 BP-guided decimation

Input: Graph G
Initialize BP messags $\{\nu_{i\to j}\}$
Initialize list of undecided spins $U = \{i \in V\}$
while $U \neq \emptyset$ **do**
  run BP until it converges
  draw i from U
  compute marginal $b_i(s_i)$
  sample spin $s_i^* \sim b_i(s_i)$
  remove i from U
  fix $m_{i\to j}(s_i) = \delta(s_i - s_i^*)$
end while
Output: planted configuration $\underline{s}^*$

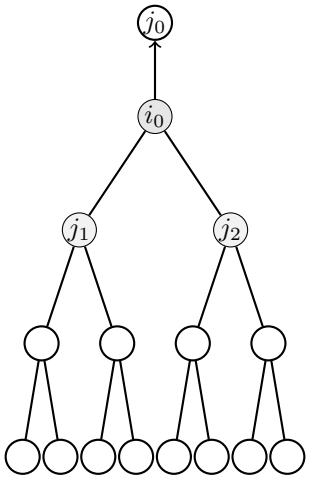

Figure 1: 2-ary regular tree of depth 3, rooted in $i_0 \to j_0$

then to randomly choose an edge $i \to j$. The resulting message $\nu_{i\to j}$ associated to this edge is a random variable with probability distribution $P(\nu_{i\to j})$.

The replica-symmetric cavity method assumes the following: in the thermodynamic limit a solution to the BP equations exists; this solution approximates well the marginals of the model and the messages in this solution are distributed according to a density evolution equation. This entails that the message $\nu_{i\to j}$ converges in distribution to a random variable ν, whose probability distribution $P(\nu)$ satisfies a fixed point equation

$$P(\nu) = \mathcal{F}^{\mathrm{RS}}(P(\nu)) \tag{42}$$

called distributional cavity equation or cavity equation in short.

A Replica-symmetric distributional cavity equations for the prior probability

In this section we illustrate how to obtain the cavity equations for the prior probability (1) by retracing the steps shown in [21].

Let $m^{(t)}_{i_0\to j_0}$ be the message sent by the BP algorithm at iteration t along the edge (i_0, j_0). Let us assume that the initial messages $\nu^{(0)}_{i_0\to j_0}$ are i.i.d. random variables with distributions independent of N.

Let us introduce the directed neighborhood of radius t of a directed edge $i_0 \to j_0$, denoted by $B_{i_0 j_0}(t)$. This is the subgraph of G made up of all the nodes which can be reached from i_0 by a non-reversing path of length at most t, when edge (i_0, j_0) is removed.

Because G is sampled from the random c-regular graph ensemble, in the thermodynamic limit $N \to \infty$, the neighborhood $B_{i_0 j_0}(t)$ converges in distribution to T_t, a c-ary regular tree of depth t (see Figure 1).

The message $m^{(t)}_{i_0\to j_0}$ is a function of the graph G and the initial conditions $\left\{m^{(0)}_{i\to j}\right\}$. However the dependence on G only occurs on through the directed neighborhood $B_{i_0\to j_0}(t+1)$, while the dependence

on the initial condition is only through the messages

$$\underline{m}^{(0)}_{B_{i_0 j_0}(t)} = \left\{ m^{(0)}_{k\to l} : (kl) \in B_{i_0 j_0}(t) \right\} \tag{43}$$

Consider the case in which the neighborhood of (i_0, j_0) really is a regular c-ary random tree T_{t+1}. We define $m^{(t)}$ as the message passed through the root edge of this tree after t BP iterations. Since $B_{i_0\to j_0}(t+1)$ converges in distribution to the tree T_{t+1}, we find that $m^{(t)}_{i_0\to j_0} \xrightarrow{d} m^{(t)}$ as $N \to \infty$. Let us now consider an edge $k \to l$ at a distance d from the root and directed towards it. The directed sub-tree rooted in $k \to l$ will be a c-ary tree of depth $t-d+1$. The message passed through it after $t-d$ iterations is distributed as $m^{(t-d)}$. This implies that, for consistency reasons, the probability distribution $P(m^{(t)})$ satisfies the following distributional equation

$$P(m^{(t)}) = \int \prod_{i=1}^{d} \mathrm{d}m_i \, P(m_i^{(t-1)}) \delta\left[m^{(t)} - g^{\mathrm{BP}}\left(\left\{ m_i^{(t-1)} \right\}_{i=1,\dots,d} \right) \right] \tag{44}$$

where m_i are independent copies of $m^{(t-1)}$ and $d = c - 1$.

The cavity equation is obtained by substituting $P(m^{(t)})$ with the fixed point $P(m)$:

$$P(m) = \int \prod_{i=1}^{d} \mathrm{d}m_i \, P(m_i) \delta\left[m - g^{\mathrm{BP}}\left(\{m_i\}_{i=1,\dots,d} \right) \right] \tag{45}$$

We will refer to it in the following as planted cavity equation.

B Replica-symmetric distributional cavity equations for the posterior probability

Applying the replica-symmetric cavity method directly to our posterior distribution (8) is not as straightforward since the BP equation (30) depends on the couplings variables $\underline{J}$. In particular, these random variable are not independent, but are instead sampled from the probability distribution

$$P(\underline{J}) = \sum_{\underline{s}} P_\beta(\underline{J}|\underline{s}) P_\kappa(\underline{s}) \tag{46}$$

Let us consider the message $\nu_{i\to j}$ of an edge $(i,j) \in G$, where G is a generic graph: it will depend on the incoming messages $\{\nu_{k\to i}\}_{k\in\partial i\setminus j}$ and the couplings $\{J_{ik}\}_{k\in\partial i\setminus j}$. As previously stated, the couplings are correlated. However, this is not true anymore when conditioning on s_i. Furthermore, it is important to introduce the planted messages $m_{i\to j}$ in order to sequentially sample the planted spins, as discussed in Section IV.A. We can therefore write

$$P(\nu_{i\to j}) = \sum_{s_i} \int \mathrm{d}m_{i\to j} \, P(\nu_{i\to j}|s_i, m_{i\to j}) P(s_i|m_{i\to j}) P(m_{i\to j}) \tag{47}$$

and look for a cavity equation of $P(\nu_{i\to j}|s_i, m_{i\to j})P(m_{i\to j})$, since we will show in the following that $P(s_i|m_{i\to j})$ simplifies.

We start by considering the message $\nu^{(t)}_{i_0\to j_0}$ going from node i_0 to j_0 at iteration t of the BP algorithm and by assuming that the initial messages $\underline{\nu}^{(0)}$ are i.i.d. random variables. The message $\nu^{(t)}_{i_0\to j_0}$ will depend on the couplings

$$\underline{J}_{B_{i_0 j_0}(t)} = \{ J_{kl} : (kl) \in B_{i_0 j_0}(t) \} \tag{48}$$

and the initial messages

$$\underline{\nu}^{(0)}_{B_{i_0 j_0}(t)} = \left\{ \nu^{(0)}_{k\to l} : (kl) \in B_{i_0 j_0}(t) \right\} \tag{49}$$

where $B_{i_0 j_0}(t)$ is the neighborhood of $i_0 \to j_0$ of depth t.

Similarly to before, the neighborhood $B_{i_0 j_0}(t)$ converges in distribution to T_t, a c-ary tree of depth t. It follows that $\nu^{(t)}_{i_0\to j_0} \xrightarrow{d} \nu^{(t)}$ where $\nu^{(t)}$ is an N-independent random variable.

The random variable $\nu^{(t)}$ will depend on $\underline{J}_{T_t}, \underline{\nu}^{(0)}_{T_t} = \left\{ \nu^{(0)}_{k\to l} : k \text{ leaf of } T_t \right\}$, in particular:

$$\nu^{(t)} = f(\underline{J}_{T_t}, \underline{\nu}^{(0)}_{T_t}) \tag{50}$$

where f is some function that can be computed by recursively applying (30) starting from the leaves of the tree T_t.

Therefore, the probability distribution $P(\nu^{(t)})$ must satisfy

$$P(\nu^{(t)}_{i_0\to j_0}) = \sum_{\underline{J}_{T_t}} \int \mathrm{d}\underline{\nu}^{(0)}_{T_t}\, P(\underline{J}_{T_t},\underline{\nu}^{(0)}_{T_t})\delta\Big[\nu^{(t)}_{i_0\to j_0} - f(\underline{J}_{T_t},\underline{\nu}^{(0)}_{T_t})\Big]$$
$$= \sum_{\underline{s}_{T_t}}\sum_{\underline{J}_{T_t}} \int \mathrm{d}\underline{\nu}^{(0)}_{T_t}\, P(\underline{J}_{T_t}|\underline{s}_{T_t})P(\underline{s}_{T_t})P(\underline{\nu}^{(0)}_{T_t})\delta\Big[\nu^{(t)}_{i_0\to j_0} - f(\underline{J}_{T_t},\underline{\nu}^{(0)}_{T_t})\Big] \tag{51}$$

$$\tag{52}$$

Because T_t is a tree, we can sample $\underline{s}_{T_t}$ sequentially, as shown in Section IV.A, and introduce the messages of the planted configuration $m_{i\to j}$. Since we considering a random graph instance these messages are now random variables and the probabilities (33) and (40) become

$$b_i(s_i) = P(s_i|\{m_{j\to i}\}_{j\in\partial i}) \quad \text{and} \quad b_k(s_k|s_i) = P(s_k|s_i, m_{k\to i}) \tag{53}$$

We can explicit the dependence of our message $\nu_{i_0\to j_0}$ from the planted spin s_{i_0} and the message $m_{i_0\to j_0}$ by expressing $P(\nu^{(t)}_{i_0\to j_0})$ as a marginal probability, as done in (47).

Introducing the messages of the planted configuration defined on the tree T_t in the r.h.s we can write

$$P(\nu^{(t)}_{i_0\to j_0}|s_{i_0}, m_{i_0\to j_0})P(s_{i_0}|m_{i_0\to j_0})P(m_{i_0\to j_0}) = \iint \mathrm{d}\underline{m}_{T_t}\,\mathrm{d}\underline{\nu}^{(0)}_{T_t} \sum_{\underline{s}_{T_t\setminus i_0}}\sum_{\underline{J}_{T_t}} \times$$
$$\times\, P(\underline{J}_{T_t},\underline{m}_{T_t},\underline{s}_{T_t\setminus i_0}|s_{i_0}, m_{i_0\to j_0})P(s_{i_0}|m_{i_0\to j_0})P(m_{i_0\to j_0})P(\underline{\nu}^{(0)}_{T_t})\delta\Big[\nu^{(t)}_{i_0\to j_0} - f(\underline{J}_{T_t},\underline{\nu}^{(0)}_{T_t})\Big] \tag{54}$$

We can remove $P(s_{i_0}|m_{i_0\to j_0})$ from both sides and obtain

$$P(\nu^{(t)}_{i_0\to j_0}|s_{i_0}, m_{i_0\to j_0})P(m_{i_0\to j_0}) = \iint \mathrm{d}\underline{m}_{T_t}\,\mathrm{d}\underline{\nu}^{(0)}_{T_t} \sum_{\underline{s}_{T_t\setminus i_0}}\sum_{\underline{J}_{T_t}} P(\underline{J}_{T_t}|\underline{s}_{T_t})P(s_{T_t\setminus i_0}|s_{i_0}, m_{i_0\to j_0}, \underline{m}_{T_t})P(\underline{m}_{T_t})\times$$
$$\times\, P(\underline{\nu}^{(0)}_{T_t})\delta\Big[m_{i_0\to j_0} - g^{\mathrm{BP}}(\{m_{j\to i_0}\}_{j\in\partial i_0\setminus j_0})\Big]\delta\Big[\nu^{(t)}_{i_0\to j_0} - f(\underline{J}_{T_t},\underline{\nu}^{(0)}_{T_t})\Big] \tag{55}$$

Introducing the BP equation (30) and exploiting the fact that T_t is a tree we can write

$$\delta\Big(\nu^{(t)}_{i_0\to j_0} - f(\underline{J}_{T_t},\underline{\nu}^{(0)}_{T_t})\Big) =$$
$$\sum_{\left\{\nu^{(t-1)}_{j\to i_0}\right\}_{j\in\partial i_0\setminus j_0}} \delta\left(\nu^{(t)}_{i_0\to j_0} - f^{\mathrm{BP}}\left(\left\{\nu^{(t-1)}_{j\to i_0}, J_{i_0,j}\right\}_{j\in\partial i_0\setminus j_0}\right)\right)\prod_{j\in i_0\setminus j_0}\delta\Big(\nu^{(t-1)}_{j\to i_0} - f(\underline{J}_{T_j},\underline{\nu}^{(0)}_{T_j})\Big) \tag{56}$$

$$P(\underline{J}_{T_t}|\underline{s}_{T_t}) = \prod_{j\in\partial i_0\setminus j_0} P(J_{i_0,j}|s_{i_0}, s_j)P(\underline{J}_{T_j}|\underline{s}_{T_j}) \tag{57}$$

$$P(\underline{s}_{T_t\setminus i_0}|s_{i_0}, m_{i_0\to j_0}, \underline{m}_{T_t}) = \prod_{j\in\partial i_0\setminus j_0} P(\underline{s}_{T_j\setminus j}|s_j, m_{j\to i_0}, \underline{m}_{T_j})P(s_j|m_{j\to i_0}, s_{i_0}) \tag{58}$$

$$P(\underline{m}_{T_t}) = \prod_{j\in\partial i_0\setminus j_0}\delta\Big[m_{j\to i_0} - g^{\mathrm{BP}}(\{m_{k\to j}\}_{k\in\partial j\setminus i_0})\Big]P(\underline{m}_{T_j}) \tag{59}$$

We therefore obtain the RS distributional cavity equation for the random regular graph

$$P(\nu|s,m)P(m) = \int\prod_{i=1}^{d}\mathrm{d}m_i\, P(m_i)\delta\Big[m - g^{\mathrm{BP}}\Big(\{m_i\}_{i=1,\dots,d}\Big)\Big]\times$$
$$\times \sum_{s_1,\dots,s_d}\sum_{J_1,\dots,J_d}\prod_{i=1}^{d}P(J_i|s,s_i)P(s_i|s,m_i)\int\prod_{i=1}^{d}\mathrm{d}\nu_i\, P(\nu_i|s_i,m_i)\delta\Big[\nu - f^{\mathrm{BP}}\Big(\{\nu_i, J_i\}_{i=1,\dots,d}\Big)\Big] \tag{60}$$

where $d = c-1$. We define the shorthand notation

$$P(\nu|s,m)P(m) = \mathcal{F}^{\mathrm{RS}}(P(\nu|s,m)P(m)) \tag{61}$$

It easy to see that by integrating over ν on both sides we obtain the planted cavity equation (45).

As mentioned before, for $\kappa < \kappa_c$ the planted configuration is in the paramagnetic state and all planted messages are $m_{i\to j}(s_i) = 1/2$. When this is the case, the dependence of $P(\nu|s,m)$ on the messages m is trivial and the distributional cavity equation can be simplified to

$$P(\nu|s) = \prod_{i=1}^{d}\left(\sum_{s_i}\sum_{J_i} P(J_i|s,s_i)P(s_i|s)\int \mathrm{d}\nu_i\, P(\nu_i|s_i)\right)\delta\Big[\nu - f^{\mathrm{BP}}\Big(\{\nu_i, J_i\}_{i=1,\dots,d}\Big)\Big] \tag{62}$$

Finally, let us calculate the joint probability $P(\mu, s, b)$ as it will be useful when calculating replica-symmetric estimates of the observables.

Consider the marginals μ_i and b_i of node i. From (30) and (38) we know them to depend, respectively, on $\underline{\nu}_{\partial i} = \{\nu_{j\to i} : j \in \partial i\}, \underline{J}_{\partial i} = \{J_{ij} : j \in \partial i\}$, and $\underline{m}_{\partial i} = \{m_{j\to i} : j \in \partial i\}$. The joint probability will therefore be

$$P(\mu_i, s_i, b_i) = \sum_{\underline{\nu}_{\partial i}, \underline{J}_{\partial i}} \int \mathrm{d}\underline{m}_{\partial i}\, P(\underline{\nu}_{\partial i}, s_i, \underline{m}_{\partial i}, \underline{J}_{\partial i})\delta[\mu_i - h_1(\underline{\nu}_{\partial i}, \underline{J}_{\partial i})]\delta[b_i - h_2(\underline{m}_{\partial i})] \tag{63}$$

where h_1 and h_2 are shorthand notations for the (33) and (39) respectively. Following a similar procedure as before we obtain

$$\begin{aligned} P(\mu_i, s_i, b_i) = \sum_{\underline{s}_{\partial i}, \underline{\nu}_{\partial i}, \underline{J}_{\partial i}} \int \mathrm{d}\underline{m}_{\partial i}\, \delta[\mu_i - h_1(\underline{\nu}_{\partial i})]\delta[b_i - h_2(\underline{m}_{\partial i})]b_i \times \\ \times \prod_{j\in\partial i} P(\nu_{j\to i}|m_{j\to i}, s_j)P(J_{j,i}|s_j, s_i)P(s_j|m_{j\to i}, s_i)P(m_{j\to i}) \end{aligned} \tag{64}$$

and therefore

$$\begin{aligned} P(\mu, s, b) = \sum_{s_1,\dots,s_d}\sum_{J_1,\dots,J_c}\int\left(\prod_{i=1}^{c}\mathrm{d}m_i\,\mathrm{d}\nu_i\, P(\nu_i|m_i,s_i)P(m_i)P(J_i|s,s_i)P(s_i|s,m_i)\right)\times \\ \times\,\delta[\mu - h_1(\{\nu_i, J_i\}_{i=1,\dots,c})]\delta[b - h_2(\{m_i\}_{i=1,\dots,c})]b(s) \end{aligned} \tag{65}$$

C Replica symmetric estimates of the observables

Let us now present how to obtain RS estimates of the observables of interest: the free entropy and the overlaps defined in Section IV.D.

The RS prediction of the free entropy is obtained by averaging the Bethe free entropy (34) over the message distribution and can be shown (see Appendix A) to be equal to

$$\phi^{\mathrm{RS}} = \mathbb{E}^{\mathrm{RS}}[\ln Z_i] + \frac{c}{2}\mathbb{E}^{\mathrm{RS}}[\ln Z_{ij}] \tag{66}$$

where

$$\mathbb{E}^{\mathrm{RS}}[\ln Z_i] = \sum_s\left[\prod_{k=1}^{c}\sum_{s_k}\iint \mathrm{d}\nu_k\,\mathrm{d}m_k\, P(\nu_k|s_k,m_k)P(m_k)P(J_k|s,s_k,m_k)P(s_k|s,m_k)\right]b(s)\ln Z_i \tag{67}$$

$$\mathbb{E}^{\mathrm{RS}}[\ln Z_i] = \prod_{k=i,j}\left(\sum_{s_k}\int \mathrm{d}m_k\,\mathrm{d}\nu_k\, P(\nu_k,|s_k,m_k)P(m_k)\right)P(J_{ij}|s_i,s_j)P(s_j|s_i,m_j)b(s_i)\ln Z_{ij} \tag{68}$$

with Z_i and Z_{ij} defined in (35).

The overlaps defined in Section IV.D are calculated in a similar manner, so we can focus on only one of the three. We consider the overlap between two configurations sampled from the posterior (8). It can easily be shown that the overlap depends on the marginals

$$\mathbb{E}_{\underline{J},\underline{\sigma},\underline{\sigma}'}[O(\underline{\sigma},\underline{\sigma}')] = \frac{1}{N}\sum_{i=1}^{N}\sum_{\sigma_i}\sum_{\underline{J}} P(\underline{J})\mu_i^2(\sigma_i) \tag{69}$$

From (33) we know that the marginal μ_i is a function of the incoming messages $\nu_{j\to i}$, which in turn are functions of the couplings $\underline{J}$ and the initial messages $\underline{\nu}^{(0)}$. Therefore we can write

$$P(\mu_i) = \sum_{\underline{J}} \int \mathrm{d}\underline{\nu}^{(0)}\, P(\underline{\nu}^{(0)}) P(\underline{J}) \delta[\mu_i - \tilde{f}(\underline{\nu}^{(0)}, \underline{J})] \tag{70}$$

Expanding this probability distribution we can write the expectation value in function of (65)

$$\mathbb{E}_{\underline{J},\underline{\sigma},\underline{\sigma}'}[O(\underline{\sigma},\underline{\sigma}')] = \sum_s \iint \mathrm{d}\mu\, \mathrm{d}b\, P(\mu, s, b) \sum_\sigma \mu^2(\sigma) \tag{71}$$

Similarly we obtain

$$\mathbb{E}_{\underline{s},\underline{J},\underline{\sigma}}[O(\underline{\sigma},\underline{s})] = \iint \mathrm{d}b\, \mathrm{d}\mu \sum_s P(\mu, s, b)\mu(s) \tag{72}$$

$$\mathbb{E}_{\underline{s},\underline{J}}[O(\underline{s},\underline{\hat{\sigma}})] = \iint \mathrm{d}b\, \mathrm{d}\mu \sum_s P(\mu, s, b)\delta(s, \hat{\sigma}) \tag{73}$$

D Population dynamics

The distributional cavity equation (60) cannot be solved analytically and therefore one must use numerical methods to estimate the probability distribution of the random variables ν and m. One example is the population dynamics algorithm.

The idea is to approximate the distribution $P(\nu|s,m)P(m)$ with a sample, called population, of N i.i.d. copies of ν and m. As previously stated in Section IV.B, we condition the messages on the value of the planted spins. Therefore, for each element ν_i of the population we associate two elements ν_i^+ and ν_i^-, one for each the two values the planted spin s_i can take.

Our population will therefore consist of N i.i.d. copies ν_i^+, ν_i^-, m_i with $i = 1, \ldots, N$.

In the thermodynamic limit $N \to \infty$ the empirical distributions of the population converges to the actual distribution:

$$P(\nu|s,m)P(m) \simeq \frac{1}{N}\sum_{i=1}^{N} \delta(\nu - \nu_i^{(s)})\delta(m - m_i) \tag{74}$$

The population is then used to iteratively solve the cavity equation (60) by interpreting it as an update equation

$$P^{(t+1)}(\nu|s,m)P^{(t+1)}(m) = \mathcal{F}^{\mathrm{RS}}\left(\left\{P^{(t)}(\nu_i|s_i,m_i)P^{(t)}(m_i)\right\}_{i=1,\ldots,d}\right) \tag{75}$$

The pseudo code of the algorithm is shown in Algorithm 2.

Before solving (60) in such a way, we solved the planted cavity equation (45) by running the population dynamics algorithm only on the population $\{m_i\}$ which we then used to initialize the combined population $\left\{\nu_i^+, \nu_i^-, m_i\right\}$.

E Interpretation of the cavity equation and link with inference

In Section IV.D we defined the feasibility of the inference problem in function of the overlap between the MMO estimator and the planted configuration, distinguishing between two cases: feasible and impossible. Further distinctions can be made for the feasible case. In particular one can distinguish between: easy, hard and hybrid-hard inference[22].

We define the accuracy a of an estimation procedure as the measure of the additional information it exploits in the posterior distribution with respect to the prior. This quantity is non-negative and $a = 0$ corresponds to the case in which the inference problem is not feasible, i.e. the posterior gives no additional information with respect to the prior. The rescaled overlap (23) satisfies such requirements:

$$a(\beta,\kappa) = \tilde{\mathrm{O}}\left(\underline{s}, \underline{\hat{\sigma}}^{\mathrm{MMO}}\right) \tag{76}$$

To define the three regimes, let us consider the BP algorithm, in particular let us define the initial condition for the messages as follows

$$\nu_{i\to j}(\sigma_i) = \begin{cases} \delta(\sigma_i, s_i) & \text{with probability } \varepsilon \\ m_{i\to j}(\sigma_i) & \text{with probability } 1-\varepsilon \end{cases} \tag{77}$$

Algorithm 2 Population dynamics

Input: population size N, max number of iterations T, parameters β, κ
Initialize populations $\left\{\nu_i^{(s,0)}\right\}, \left\{m_i^{(0)}\right\}$
for $t = 1, \dots, T$ **do**
  for $i = 1, \dots, N$ **do**
    for $s = 1, 2$ **do**
      draw $j(1), \dots, j(c) \overset{i.i.d.}{\sim} \text{Unif}\{1, \dots, N\}$
      set $m_i^{(t)} = g^{\text{BP}}\left(\left\{m_{j(k)}^{(t-1)}\right\}_{k=1,\dots,c-1}\right)$
      sample $s_i \sim b_i(s_i)$
      sample $s_{j(k)} \sim P_{\kappa^*}(s_{j(k)}|s, m_{j(k)}) \quad \forall k \in 1, \dots, c$
      sample $J_{j(k)} \sim P(J|s_{j(k)}, s) \quad \forall k \in 1, \dots, c$
      set $\nu_i^{(s,t)} = f^{\text{BP}}\left(\left\{\nu_{j(k)}^{(s_{j(k)},t-1)}, J_{j(k)}\right\}_{k=1,\dots,c-1}\right)$
      calculate $\mathbb{E}^{\text{RS}}[\ln Z_i^{(t)}], \mathbb{E}[\text{O}^{(t)}(\underline{\sigma}, \underline{\sigma}')], \mathbb{E}[\text{O}^{(t)}(\underline{\sigma}, \underline{s})], \mathbb{E}[\text{O}^{(t)}(\hat{\underline{\sigma}}, \underline{s})]$
      sample $j \sim \text{Unif}\{1, \dots, N\}$
      sample $s_j \sim P_{\kappa^*}(s_j|s_i, m_j)$
      sample $J_{ij} \sim P(J|s_i, s_j)$
      calculate $\mathbb{E}^{\text{RS}}[\ln Z_{ij}^{(t)}]$
    end for
  end for
end for
Output: $\left\{\nu_i^{(s,T)}\right\}, \left\{m_i^{(T)}\right\}, \left\{\mathbb{E}^{\text{RS}}[\phi^{(t)}], \mathbb{E}[\text{O}^{(t)}(\underline{\sigma}, \underline{\sigma}')], \mathbb{E}[\text{O}^{(t)}(\underline{\sigma}, \underline{s})], \mathbb{E}[\text{O}^{(t)}(\hat{\underline{\sigma}}, \underline{s})]\right\}_{t=1,\dots,T}$

Two cases are of particular interest: the informed $\varepsilon = 1$ and the uninformed $\varepsilon \to 0$ initial conditions. When both conditions lead to the same solution inference is easy, otherwise we are in one of the two hard cases.

We define $a^*(\beta, \kappa)$ as the information-theoretical optimal accuracy and $a_{\text{alg}}(\beta, \kappa)$ as the accuracy obtained by the most efficient algorithm (BP in our case) starting only from the knowledge of the prior.

We characterize these quantities using the cavity equations. In particular, let us define the initial conditions used to iteratively solve the cavity equations in a similar way to (77):

$$P(\nu|s, m) = \varepsilon\delta(\nu - \delta_{s,\cdot}) + (1 - \varepsilon)\delta(\nu - m) \tag{78}$$

The uninformed initial condition ($\varepsilon \to 0$) allows us to compute the algorithmic accuracy $a_{\text{alg}}(\beta, \kappa)$, while the information-theoretical optimal accuracy $a^*(\beta, \kappa)$ is described by the fixed point of the cavity equation (60). It might be that several non-trivial fixed points exist, in such cases the one describing the optimal accuracy is the one with highest free entropy. In our case such a fixed point is obtained with the informed initial condition ($\varepsilon = 1$). This can be justified intuitively as this initial condition corresponds to the cases in which the messages are already correct with respect to the planted configuration.

Following these definitions we can define the easy, hard, hybrid-hard and impossible phases as follows:

- easy: $a^*(\beta, \kappa) = a_{\text{alg}}(\beta, \kappa) > 0$
- hard: $a^*(\beta, \kappa) > a_{\text{alg}}(\beta, \kappa) = 0$
- hybrid-hard: $a^*(\beta, \kappa) > a_{\text{alg}}(\beta, \kappa) > 0$
- impossible: $a^*(\beta, \kappa) = a_{\text{alg}}(\beta, \kappa) = 0$

V Stability analysis of the paramagnetic solution

Let us focus on the region $\kappa \in [-\kappa_{\text{KS}}, \kappa_c]$. In this region the cavity equation (60) admits the trivial paramagnetic fixed point

$$P(\nu) = \delta[\nu - \bar{\nu}] \quad \text{with} \quad \bar{\nu}(\sigma) = \frac{1}{2} \tag{79}$$

Such a fixed point gives vanishing rescaled overlap (23).

This fixed point is the one found by the BP algorithm when the initial condition (77) is completely uninformed, i.e. $\varepsilon = 0$. We wish to study its stability with respect to a small perturbation: if this is the case, any initial condition close to $\varepsilon = 0$ leads to the trivial fixed point and therefore this will be the only solution the BP algorithm will be able to find.

We can study the stability by considering a distribution $P(\nu|s)$ solution of the RS equations close to the paramagnetic one:

$$\nu_{k\to i}(\sigma_k) = \frac{1}{2} + \varepsilon_{k\to i}(\sigma_k) \tag{80}$$

where $\varepsilon_{k\to i}(\sigma_k)$ is a small perturbation such that $\sum_{\sigma_k} \varepsilon_{k\to i}(\sigma_k) = 0$. Introducing this expression in the BP equation (30) and keeping only linear terms in ε we find

$$\nu_{i\to j}(\sigma_i) = \frac{1}{2} + \sum_{k\in\partial i\setminus j} \sum_{\sigma_k} \frac{1}{2\cosh(\beta J_{ik} + \kappa)} e^{(\beta J_{ik}+\kappa)\sigma_i\sigma_k} \varepsilon_{k\to i}(\sigma_k) \tag{81}$$

Let us now calculate the first moment of the distance between the paramagnetic solution and the perturbed solution:

$$M_s(\sigma) = \int \mathrm{d}\nu\, P(\nu|s) \left(\nu(\sigma) - \frac{1}{2}\right) \tag{82}$$

This quantity quantifies the distance of the RS solution $P(\nu|s)$ from the trivial fixed point. Our goal is to determine how this quantity evolves under the iterations of recursively solving the cavity equation (75). In particular, we want to know in which conditions this distance grows for each iteration.

Introducing the RS equation (62) and utilizing the normalization condition $\sum_\sigma M_s(\sigma) = 0$ we find

$$M_s(+) = d \sum_J \frac{\tanh(\beta J + \kappa)}{4\cosh\kappa\cosh\beta} \sum_{s'} e^{(\beta J+\kappa)s's} M_{s'}(+) \tag{83}$$

where $d = c - 1$.

For simplicity of notation let us define $M_s(+) \equiv M_s$. We can write this in the form of a linear system

$$\underline{M}^{(t+1)} = A\underline{M}^{(t)} \tag{84}$$

where

$$\underline{M} = \begin{pmatrix} M_+ \\ M_- \end{pmatrix}, \quad A = \begin{pmatrix} a & b \\ b & a \end{pmatrix} \tag{85}$$

and

$$a = (c-1)\sum_J \frac{\tanh(\beta J + \kappa)}{4\cosh\kappa\cosh\beta} e^{\beta J+\kappa}, \quad b = (c-1)\sum_J \frac{\tanh(\beta J + \kappa)}{4\cosh\kappa\cosh\beta} e^{-(\beta J+\kappa)}. \tag{86}$$

The paramagnetic fixed point will be stable if and only if both the eigenvalues are smaller than one. It is easily shown that the eigenvalues are

$$\lambda_1 = (c-1)\tanh\kappa \qquad \lambda_2 = (c-1)\frac{\tanh(\beta+\kappa)\sinh(\beta+\kappa) + \tanh(\beta-\kappa)\sinh(\beta-\kappa)}{\cosh\kappa\cosh\beta} \tag{87}$$

which implies

$$\kappa < \operatorname{atanh}\left(\frac{1}{c-1}\right) \quad \text{and} \quad (c-1)\frac{\tanh(\beta+\kappa)\sinh(\beta+\kappa) + \tanh(\beta-\kappa)\sinh(\beta-\kappa)}{\cosh\kappa\cosh\beta} < 1 \tag{88}$$

The resulting phase diagram is shown in Figure 2a. The blue line corresponds to

$$\kappa = \operatorname{atanh}\left(\frac{1}{c-1}\right) \tag{89}$$

while the red one corresponds to

$$(c-1)\frac{\tanh(\beta+\kappa)\sinh(\beta+\kappa) + \tanh(\beta-\kappa)\sinh(\beta-\kappa)}{\cosh\kappa\cosh\beta} = 1 \tag{90}$$

By fixing $\kappa = 0$ we recover the critical value of β in the case of the unstructured signal[16]:

$$\beta_c(\kappa = 0) = \operatorname{atanh}\left(\frac{1}{\sqrt{c-1}}\right) \tag{91}$$

The region delimited by the full lines is the one in which the trivial, paramagnetic, solution is stable. From this diagram it is clear to see that $\beta_c(\kappa)$ decreases as the absolute value of κ increases, showing that the fact that the signal is structured aids inference.

In the first condition we can identify the critical value κ_c previously discussed in Section V.B. For all $\kappa > \kappa_c$, the paramagnetic fixed point is not stable. In this regime, the definition of $M_s(\sigma)$ given in (82) is not optimal as it doesn't tell us whether the probability distribution of the messages ν differs from the trivial paramagnetic fixed point because meaningful inference has been done or simply because we know that the planted configuration has non-vanishing magnetization. A possible solution would be to substitute $1/2$ with $(1 + \sigma\gamma(\kappa))/2$ where $\gamma(\kappa)$ is the magnetization.

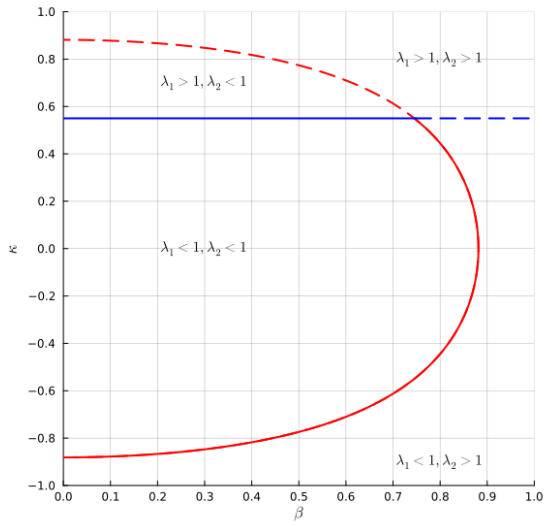

(a) Phase diagram of the stability of the paramagnetic fixed point. The blue line corresponds to (89), while the red line corresponds to (90). The region inside the full lines is where both eigenvalues are smaller than 1 and the paramagnetic fixed point is stable.

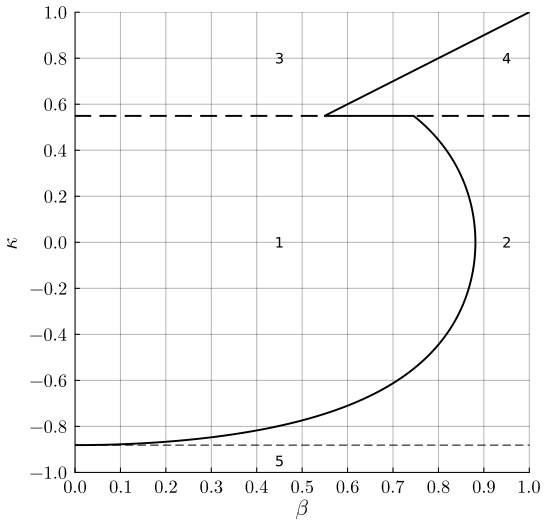

(b) Phase diagram of the feasibility of the inference problem, obtained by numerically solving the cavity equation (60). Regions 1 and 3 correspond to the impossible phase. Regions 2 and 4 correspond to the easy phase. In regions 1 and 2 the planted configuration is in a paramagnetic state. In regions 3 and 4 the planted configuration is in a ferromagnetic phase, while in region 5 it is in a RSB phase.

Figure 2: Analytical and numerical phase diagrams

VI Numerical results

In this section we present the results of the numerical analysis of the feasibility of the inference task for the planted spin glass with structured signal. As stated previously, we focus on the Bayes optimal setting in which the parameters are known and do not need to be inferred. In particular, we solve numerically the replica-symmetric distributional cavity equation (60) and compare the results with those obtained by running the Belief Propagation algorithm on generated instances of the inference problem.

We use the population dynamics algorithm presented in Section VI.D to numerically solve the RS cavity equation. First of all, we initialize a population of $N = 10^5$ planted BP messages and perform a fixed number of iterations *plantedmaxiter*=500. We subsequently initialize the population of BP messages following (78) and perform a fixed number of iterations *maxiter*=1000. Each simulation is performed twice: once with the informed ($\varepsilon = 1$) initial condition, in order to obtain the information-theoretical optimal performance, and once with the uninformed initial condition ($\varepsilon = 0.001$), to obtain the optimal algorithmic performance.

We generate instances of the inference problem by generating a random regular graph of $N = 10^4$ nodes, sampling a planted configuration using BP-guided decimation (see Section VI.A) and then sampling the observations (see Section VI.A). The messages are initialized following (77) and the BP algorithm is iterated until all messages satisfy the same convergence criterion of not differing by more than 10^{-6}

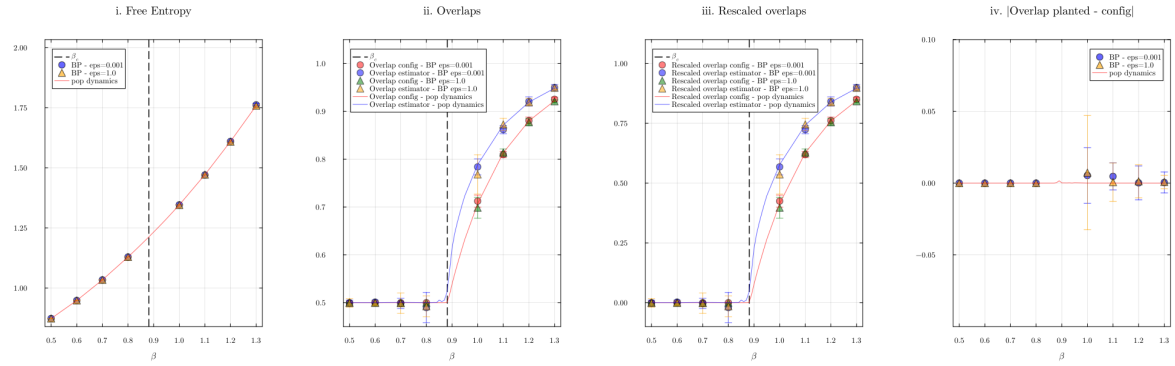

(a) $\kappa = 0$

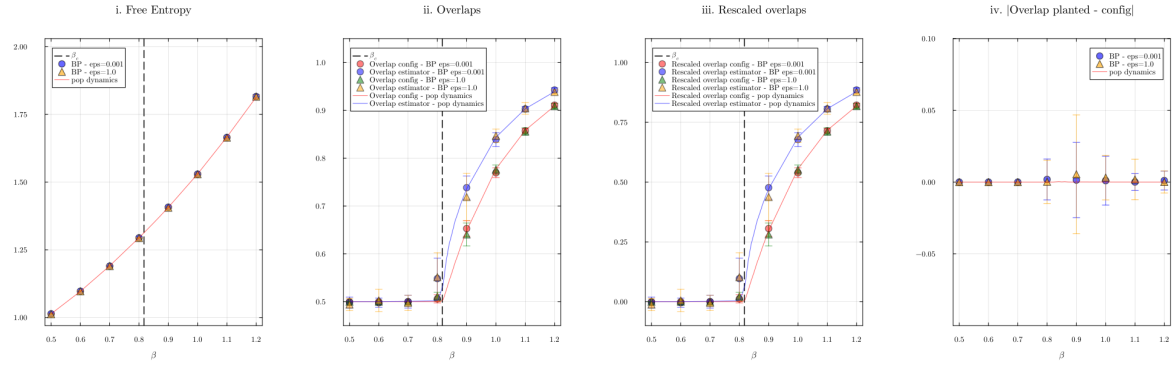

(b) $\kappa = 0.4$

Figure 3: Phase transition ($1 \to 2$) with respect to β for $\kappa = 0$ and $\kappa = 0.4$. i: plot of free entropy. ii: plot of "overlap estimator" (17), and "overlap config", i.e. the overlap (24) between two configurations sampled from the posterior (8). iii: plot of accuracy (76), i.e. "rescaled overlap estimator" (23), and "rescaled overlap config" (27). iv: plot of Nishimori condition, obtained by comparing "overlap config" and "overlap planted", i.e. the overlap (24) between a configuration sampled from the posterior (8) and the planted configuration. The population dynamics result are illustrated by the straight lines, while the circles and triangles correspond to the Belief propagation results obtained with $\varepsilon = 0.001$ and $\varepsilon = 1$ respectively. The vertical dashed line is the critical value β_c obtained with (90).

from the value at the previous iteration. Once again each is simulation is performed twice, once for $\varepsilon = 1$ and once for $\varepsilon = 0.001$.

We calculate the free entropy (34), the overlaps and the rescaled overlaps (23)(27)(28). To obtain an average over different finite size instances, for each set of parameters κ, β, we generate 10 finite-size instances and average the results. The results obtained with Belief Propagation coincide with those obtained with population dynamics, confirming our analytical findings.

Figure 2b shows the phase diagram obtained with the population dynamics algorithm. We can distinguish 5 regions: regions 1 and 2 correspond to $\kappa \in [-\kappa_{\rm KS}, \kappa_c]$ and are divided by the critical line (90), region 3 and 4 correspond to $\kappa > \kappa_c$, obtained with (89), while region 5 corresponds to $\kappa < -\kappa_{\rm KS}$. Comparing with Figure 2a we see that the two coincide in the region $\kappa \in [-\kappa_{\rm KS}, \kappa_c]$, i.e. where the cavity equation (60) admits the paramagnetic trivial fixed point.

In regions 1 and 2, the planted configuration is in the paramagnetic phase and has vanishing magnetization. In both regions the overlaps are equivalent to their rescaled counterparts, as shown in Section VI.D. In region 1 the population dynamics algorithm is only able to find the trivial solution, for both initial conditions: the optimal accuracy $a^*(\beta, \kappa)$ and the algorithmic accuracy $a_{\rm alg}(\beta, \kappa)$ coincide and are both zero. This corresponds to the impossible phase. In region 2 the population dynamics algorithm is able to find a non-trivial solution for both initial conditions: the optimal accuracy $a^*(\beta, \kappa)$ and the algorithmic accuracy $a_{\rm alg}(\beta, \kappa)$ coincide and are both different from zero. This corresponds to the easy phase.

Figure 3, Figure 4, Figure 5 and Figure 6 show the transitions between the different regions. Each figure has four plots displaying, from left to right: the free entropy, the "overlap estimator" (17) and the "overlap config", i.e. the overlap (24) between two configurations sampled from the posterior (8), the

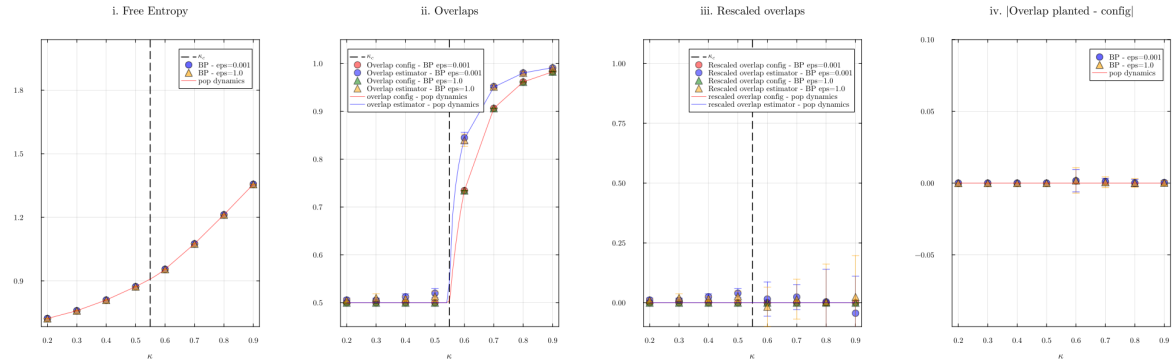

(a) $\beta = 0$

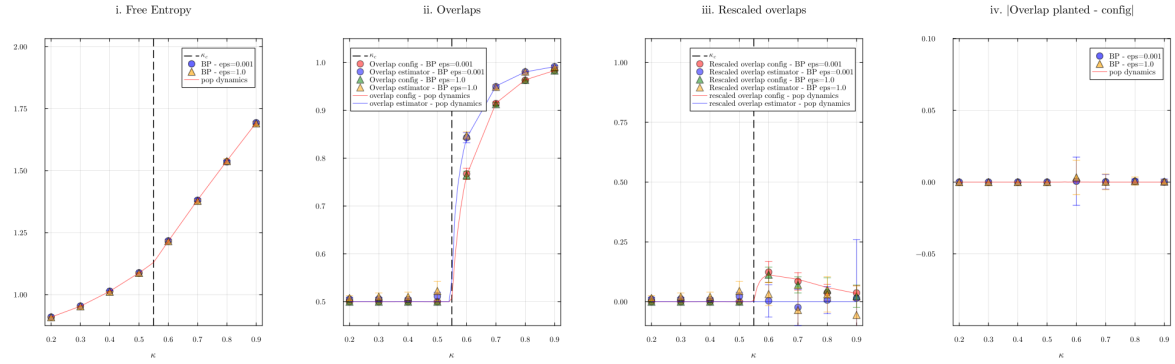

(b) $\beta = 0.5$

Figure 4: Transition between regions 1 and 3 with respect to κ for $\beta = 0$ and $\beta = 0.5$. i: plot of free entropy. ii: plot of "overlap estimator" (17), and "overlap config", i.e. the overlap (24) between two configurations sampled from the posterior (8). iii: plot of accuracy (76), i.e. "rescaled overlap estimator" (23), and "rescaled overlap config" (27). iv: plot of Nishimori condition, obtained by comparing "overlap config" and "overlap planted", i.e. the overlap (24) between a configuration sampled from the posterior (8) and the planted configuration. The population dynamics result are illustrated by the straight lines, while the circles and triangles correspond to the Belief propagation results obtained with $\varepsilon = 0.001$ and $\varepsilon = 1$ respectively. The vertical dashed line is the critical value κ_c (89).

accuracy (76), denoted by "recaled overlap estimator", and the "rescaled overlap config" (27), and the difference between said overlap with the "overlap planted" (28). All plots feature the quantities obtained with the BP algorithm with both initial conditions and population dynamics. In order not to clutter the plots, in the case of population dynamics, only the estimates for $\varepsilon = 1$ are shown, since $\varepsilon = 0.001$ leads to the same fixed point. The BP results are in good agreement with those of population dynamics. In all figures the population dynamics results are independent of the initial condition, therefore the optimal accuracy $a^*(\beta, \kappa)$ and the algorithmic accuracy $a_{\text{alg}}(\beta, \kappa)$ coincide. Furthermore, the Nishimori condition is always satisfied.

Figure 3a and Figure 3b show the phase transition with respect to β for $\kappa = 0$ and $\kappa = 0.4$ respectively. The plots are similar, the only difference being the value of $\beta_c(\kappa)$. The same holds for all plots of the phase transition occurring between regions 1 and 2. Plots for other values of κ are present in Appendix B. Plot iii shows that the rescaled overlap (23) (in the plot denoted as "rescaled overlap estimator") is zero for β up to a critical value β_c after which it becomes finite. This corresponds to the easy phase. The phase transition is between an impossible and an easy phase and is of the second order, as shown in the free energy plot. For $\kappa = 0$, i.e. the case in which the signal is not structured, we are able to recover the known value $\beta_c = \operatorname{atanh}(\frac{1}{\sqrt{c-1}})$ separating the two phases, as shown in Figure 3a. From plot ii and iii we see that the overlaps and their rescaled counterparts are equivalent since the planted configuration is in the paramagnetic state.

For values of κ in regions 3 and 4, i.e. $\kappa > \kappa_c$, the planted configuration is in the ferromagnetic phase and has non-vanishing magnetization: the overlaps (17) (24) and the rescaled overlaps (23) (25) are not equivalent.

Figure 4a and Figure 4b show the transition between regions 1 and 3 with respect to κ for $\beta = 0$ and

$\beta = 0.5$ respectively. The accuracy is vanishing for values of κ smaller and larger than κ_c, for all different values of β: regions 1 and 3 both correspond to the impossible phase. For $\kappa > \kappa_c$ the overlaps become larger than 0.5. However, this does not correspond to the presence of an easy phase, which instead corresponds to the rescaled overlaps being larger than zero. As previously mentioned, the overlaps being larger than 0.5 is caused by the planted configuration being in the ferromagnetic phase, while the rescaled overlaps being vanishing is caused by the posterior distribution not adding any information to what is already known in the prior. Therefore the critical value κ_c does not correspond to a phase transition of the feasibility of the inference problem (inference is impossible in both regions 1 and 3), but only to a phase transition of the planted configuration. For increasing values of β the "overlap config" increases slightly after κ_c, slowly going back to zero as κ increases. Furthermore, the height of the "bump" increases with the value of β. This can be explained intuitively by the marginal of the posterior $\mu_i(\sigma_i)$ being more polarized than the marginal of the prior $b_i(\sigma_i)$ for values of κ close to κ_c due to the effect of $\beta \neq 0$. As κ increases, with fixed β, the effect of the likelihood on the posterior diminishes with respect to that of the prior, and μ_i converges to b_i.

Figure 5a displays the transition region 1 $\rightarrow$ region 4 $\rightarrow$ region 3 with respect to β for $\kappa = 0.7$. Regions 1 and 3 correspond to the impossible phase, while region 4 corresponds to the easy phase. The phase transition between 1 and 4 is of the first order, while the phase transition between 4 and 3 is of second order, as can be seen in the free entropy plot i. In region 1 the overlaps and rescaled overlaps are equivalent, which is not true for regions 3 and 4.

Figure 5b displays the transition region 2 $\rightarrow$ region 4 $\rightarrow$ region 3 with respect to β for $\kappa = 0.9$. The accuracy is non vanishing for all values of κ except for $\kappa = 0.9$, which corresponds to region 3. Regions 2 and 4 both correspond to the easy phase, their only difference being the fact that in the former the overlap and rescaled overlaps are equivalent and in the latter they are not. Similarly to Figure 4, no phase transition occurs between the two regions, except the ferromagnetic phase transition of the planted configuration.

Finally, Figure 6a and Figure 6b display the phase transition with respect to β for $\kappa = 0.6$ and $\kappa = 0.9$ respectively. The plots are similar, the only difference being the value of $\beta_c(\kappa)$. The same holds for all plots of the phase transition occurring between regions 3 and 4. Plots for other values of κ are present in Appendix B. Once again, the two regions correspond respectively to the hard and easy phases and the transition is of second order. Furthermore, Figure 6b shows that the critical value β_c dividing the impossible ($\beta < \beta_c$) and easy ($\beta > \beta_c$) phases is larger than $\beta_c(\kappa = 0) = \text{atanh}(\frac{1}{\sqrt{c-1}})$.

With the definition of the rescaled overlap (23) one is able to make the distinction between feasible (in this case easy) and impossible inference, however for $\kappa = 0.9$, the non-rescaled "overlap estimator" and "overlap config" are both very close to 1. The inference problem is trivial: all the spins of the planted configuration are equal and therefore the signal can be recovered with the prior (1) alone.

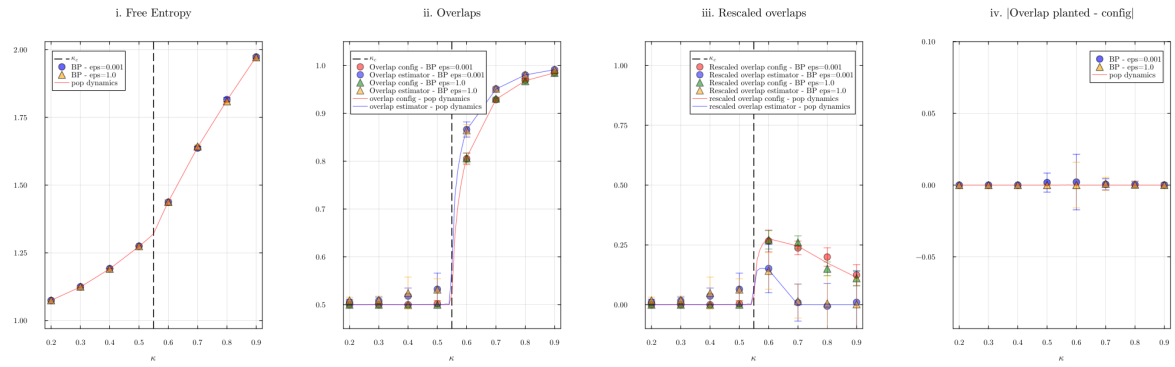

(a) $\beta = 0.7$

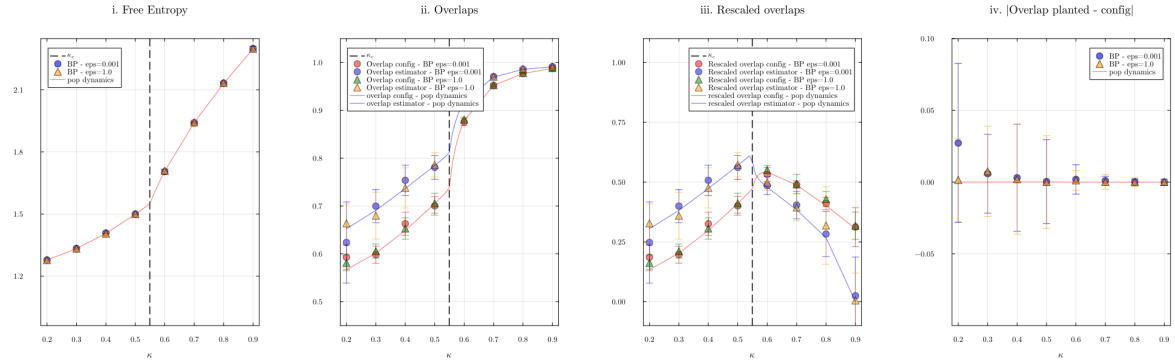

(b) $\beta = 0.9$

Figure 5: Transitions $1 \to 4 \to 3$, and $2 \to 4 \to 3$, with respect to β for $\kappa = 0.7$ and $\kappa = 0.9$ respectively. i: plot of free entropy. ii: plot of "overlap estimator" (17), and "overlap config", i.e. the overlap (24) between two configurations sampled from the posterior (8). iii: plot of accuracy (76), i.e. "rescaled overlap estimator" (23), and "rescaled overlap config" (27). iv: plot of Nishimori condition, obtained by comparing "overlap config" and "overlap planted", i.e. the overlap (24) between a configuration sampled from the posterior (8) and the planted configuration. The population dynamics result are illustrated by the straight lines, while the circles and triangles correspond to the Belief propagation results obtained with $\varepsilon = 0.001$ and $\varepsilon = 1$ respectively. The vertical dashed line is the critical value κ_c (89).

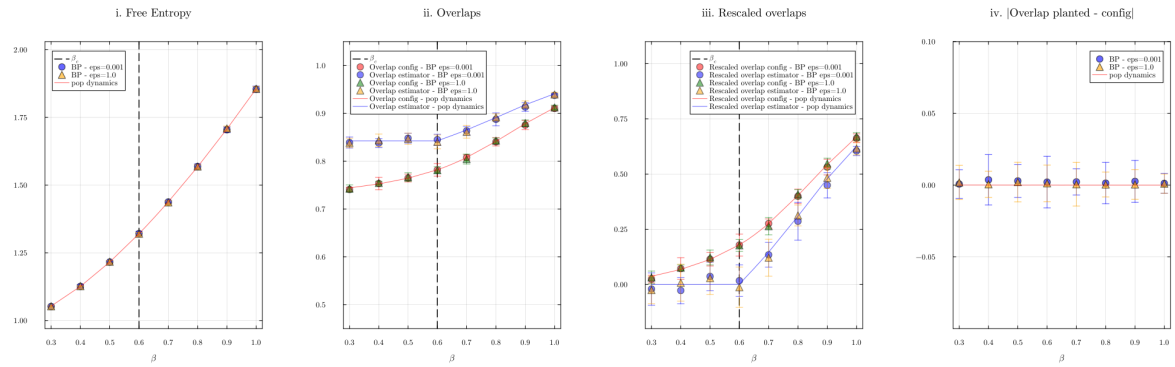

(a) $\kappa = 0.6$

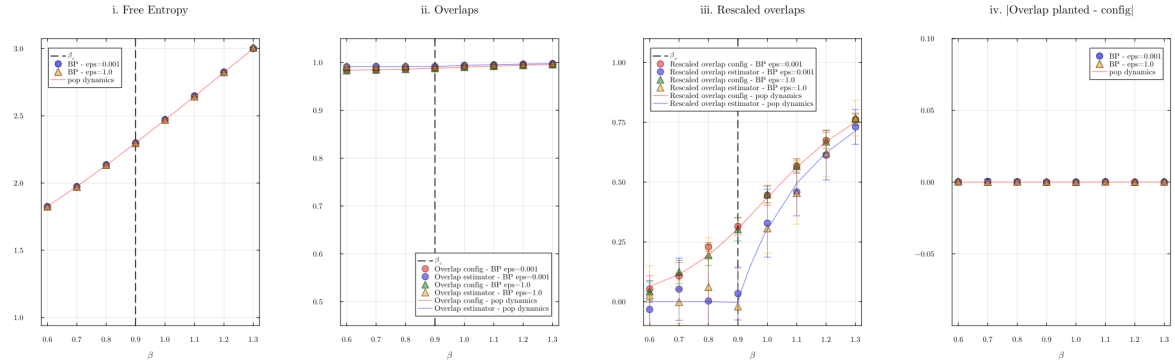

(b) $\kappa = 0.9$

Figure 6: Phase transition with respect to β for $\kappa = 0.7$ and $\kappa = 0.9$. i: plot of free entropy. ii: plot of "overlap estimator" (17), and "overlap config", i.e. the overlap (24) between two configurations sampled from the posterior (8). iii: plot of accuracy (76), i.e. "rescaled overlap estimator" (23), and "rescaled overlap config" (27). iv: plot of Nishimori condition, obtained by comparing "overlap config" and "overlap planted", i.e. the overlap (24) between a configuration sampled from the posterior (8) and the planted configuration. The population dynamics result are illustrated by the straight lines, while the circles and triangles correspond to the Belief propagation results obtained with $\varepsilon = 0.001$ and $\varepsilon = 1$ respectively. The vertical dashed line is the critical value β_c dividing the two phases, which was found numerically.

VII Discussion

In this thesis, we study the feasibility of inferring the signal of the planted spin glass in the case in which said signal is structured. Using the Replica Symmetric cavity method, we predict the phase diagram of the inference problem and we give quantitative predictions of the Maximum Mean Overlap (MMO) estimator for the different regions of the phase diagram. We show that these predictions are in good agreement with the results obtained on finite size instances.

We show that the introduction of the structure in the signal influences the inference problem. In particular, we show that for $\kappa \in [-\kappa_{\mathrm{KS}}, \kappa_c]$ the structure aids inference: for all non-zero values of κ in this range, the associated critical value $\beta_c(\kappa)$, dividing the impossible ($\beta < \beta_c$) and the feasible ($\beta > \beta_c$) phases, is lower than the critical value of the unstructured signal $\beta_c(\kappa = 0)$, value we recovered in accordance with [16]. This implies that when introducing structure in our model, a lower signal-to-noise ratio is needed to be able to successfully perform the inference task, with respect to the unstructured case. We show that the feasible phase corresponds to an easy phase, in which inference is informational-theoretically possible and that an algorithm is able to replicate the performance obtained with a perfect knowledge of the posterior probability distribution. The phase transition is therefore between an impossible phase and an easy phase, and is of second order. We compute analytically $\beta_c(\kappa)$ by studying the stability of the trivial fixed point of the cavity equation, which we find to be in agreement with the numerical simulations.

We show that the structure causes phase transitions in the signal, i.e. the planted configuration. In particular, we recover the critical value κ_c, corresponding to the paramagnetic-to-ferromagnetic phase transition, and $-\kappa_{\mathrm{KS}}$, corresponding to the RS-to-RSB phase transition.

The ferromagnetic transition causes the planted configuration to have non-vanishing magnetization: we introduce a rescaled overlap to take into account this magnetization and compare the information between the prior and the posterior. Furthermore, inference becomes trivial for $\kappa \to \infty$. We show that in this regime a second order phase transition between impossible and easy phase occurs. We show this by solving numerically the cavity equation and calculating the accuracy. The presence of structure in the signal aids inference up to certain value κ^*: for all larger values of κ, the critical value β_c is larger than $\beta_c(\kappa = 0) = \mathrm{atanh}(\frac{1}{\sqrt{c-1}})$. A stability study similar to that presented in this thesis could be performed in order to compute analytically β_c for $\kappa > \kappa_c$. In particular by changing the definition of the mean distance $M_s(\sigma)$.

The implications of the RSB phase transition are, a priori, difficult to state and the RS cavity method cannot be used. To study the inference problem in this regime the 1RSB formalism should be introduced.

The two main technical difficulties addressed in this thesis occur when applying the cavity method framework to the considered model: the correlated nature of the signal and, as a consequence of the observations, is in contradiction with the assumptions of said method. By introducing messages of the planted configuration and by conditioning the messages of the posterior on such messages and on the planted spins, we were able to circumvent these issues and apply the replica-symmetric cavity method. The method presented in this thesis can also be applied to the case in which the coupling constant κ of the planted configuration is not homogeneous or other more complicated models, such as models defined on hypergraphs in which more than 2 spins interact at the same time.

References

[1] D. L. Donoho, A. Maleki, and A. Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, **106**(45), 18914–18919 (2009).

[2] H. Cui, F. Krzakala, and L. Zdeborová. Bayes-optimal Learning of Deep Random Networks of Extensive-width, 2023.

[3] S. Goldt, F. Krzakala, L. Zdeborová, and N. Brunel. Bayesian reconstruction of memories stored in neural networks from their connectivity. *PLOS Computational Biology*, **19**(1), e1010813 (2023).

[4] F. Altarelli, A. Braunstein, L. Dall'Asta, A. Lage-Castellanos, and R. Zecchina. Bayesian Inference of Epidemics on Networks via Belief Propagation. *Phys. Rev. Lett.*, **112**, 118701 (2014).

[5] A. Braunstein, L. Budzynski, and M. Mariani. Statistical mechanics of inference in epidemic spreading. *Physical Review E*, **108**(6) (2023).

[6] O. Duranthon and L. Zdeborová. Neural-prior stochastic block model. *Machine Learning: Science and Technology*, **4**(3), 035017 (2023).

[7] S. Goldt, M. Mézard, F. Krzakala, and L. Zdeborová. Modeling the Influence of Data Structure on Learning in Neural Networks: The Hidden Manifold Model. *Physical Review X*, **10**(4) (2020).

[8] B. Loureiro, C. Gerbelot, H. Cui, S. Goldt, F. Krzakala, M. Mézard, and L. Zdeborová. Learning curves of generic features maps for realistic datasets with a teacher-student model*. *Journal of Statistical Mechanics: Theory and Experiment*, **2022**(11), 114001 (2022).

[9] A. Braunstein, L. Budzynski, M. Mariani, and F. Ricci-Tersenghi. Evidence of Replica Symmetry Breaking under the Nishimori conditions in epidemic inference on graphs, 2025.

[10] G. Parisi. *Spin Glass Theory*, pages 317–329. Springer US, Boston, MA, 1987.

[11] H. Nishimori. *Statistical Physics of Spin Glasses and Information Processing: an Introduction.* Oxford University Press, Oxford; New York, 2001.

[12] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, **84**(6) (2011).

[13] T. Lesieur, F. Krzakala, and L. Zdeborova. Phase transitions in sparse PCA. In *2015 IEEE International Symposium on Information Theory (ISIT)*, page 1635–1639. IEEE, June 2015.

[14] A. Coja-Oghlan, F. Krzakala, W. Perkins, and L. Zdeborová. Information-theoretic thresholds from the cavity method. *Advances in Mathematics*, **333**, 694–795 (2018).

[15] M. Lelarge and L. Miolane. Fundamental limits of symmetric low-rank matrix estimation. In S. Kale and O. Shamir, editors, *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 1297–1301. PMLR, 07–10 Jul 2017.

[16] L. Zdeborová and F. Krzakala. Statistical physics of inference: thresholds and algorithms. *Advances in Physics*, **65**(5), 453–552 (2016).

[17] V. H. Can. Annealed limit theorems for the ising model on random regular graphs, 2017.

[18] A. Coja-Oghlan, P. Loick, B. F. Mezei, and G. B. Sorkin. The Ising antiferromagnet and max cut on random regular graphs, 2020.

[19] A. Montanari. Estimating Random Variables from Random Sparse Observations, 2007.

[20] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian. Solving Constraint Satisfaction Problems through Belief Propagation-guided decimation, 2019.

[21] M. Mézard and A. Montanari. *Information, Physics, and Computation.* Oxford University Press, 01 2009.

[22] F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborová. Typology of phase transitions in Bayesian inference problems. *Physical Review E*, **99**(4) (2019).

A RS free entropy

Let us consider the two terms separately. The first term can be written as

$$\mathbb{E}^{\mathrm{RS}}[\ln Z_i] = \sum_{\underline{J}_c} \int \prod_{k=1}^{c} \mathrm{d}\nu_k \, P(\underline{\nu}_c, \underline{J}_c) \ln Z_i \tag{92}$$

where $\underline{\nu}_c = \{\nu_i : i = 1, \ldots, c\}$, $\underline{J}_c = \{J_i : i = 1, \ldots, c\}$ and c is the degree of the random regular graph. Focusing on the probability:

$$\begin{aligned}
P(\underline{\nu}_c, \underline{J}_c) &= \sum_{s, \underline{s}_c} \int \prod_{k=1}^{c} \mathrm{d}m_k \, P(\underline{\nu}_c, \underline{J}_c | s, \underline{s}_c, \underline{m}_c) P(s, \underline{s}_c, \underline{m}_c) && (93)\\
&= \sum_{s, \underline{s}_c} \int \prod_{k=1}^{c} \mathrm{d}m_k \, P(\underline{\nu}_c | s, \underline{s}_c, \underline{m}_c) P(\underline{J}_c | s, \underline{s}_c, \underline{m}_c) P(s, \underline{s}_c, \underline{m}_c) && (94)\\
&= \sum_{s, \underline{s}_c} \int \prod_{k=1}^{c} \mathrm{d}m_k \, P(\underline{\nu}_c | \underline{s}_c, \underline{m}_c) P(\underline{J}_c | s, \underline{s}_c, \underline{m}_c) P(\underline{s}_c | s, \underline{m}_c) P(s | \underline{m}_c) P(\underline{m}_c) && (95)\\
&= \sum_{s} \left[\prod_{k=1}^{c} \sum_{s_k} \int \mathrm{d}m_k \, P(\nu_k | s_k, m_k) P(m_k) P(J_k | s, s_k, m_k) P(s_k | s, m_k) \right] P(s | \underline{m}_c) && (96)\\
&= \sum_{s} \left[\prod_{k=1}^{c} \sum_{s_k} \int \mathrm{d}m_k \, P(\nu_k | s_k, m_k) P(m_k) P(J_k | s, s_k, m_k) P(s_k | s, m_k) \right] b(s) && (97)
\end{aligned}$$

Therefore

$$\mathbb{E}^{\mathrm{RS}}[\ln Z_i] = \sum_{s} \left[\prod_{k=1}^{c} \sum_{s_k} \iint \mathrm{d}\nu_k \, \mathrm{d}m_k \, P(\nu_k | s_k, m_k) P(m_k) P(J_k | s, s_k, m_k) P(s_k | s, m_k) \right] b(s) \ln Z_i \tag{98}$$

The second term can be written as

$$\mathbb{E}^{\mathrm{RS}}[\ln Z_{ij}] = \sum_{J_{ij}} \int \mathrm{d}\nu_i \, \mathrm{d}\nu_j \, P(\nu_i, \nu_j, J_{ij}) \ln Z_{ij} \tag{99}$$

Focusing on the probability:

$$\begin{aligned}
P(\nu_i, \nu_j, J_{ij}) &= \sum_{si, s_j} \int \mathrm{d}m_i \, \mathrm{d}m_j \, P(\nu_i, \nu_j, J_{ij} | s_i, s_j, m_i, m_j) P(s_i, s_j, m_i, m_j) && (100)\\
&= \sum_{s_i, s_j} \int \mathrm{d}m_i \, \mathrm{d}m_j \, P(\nu_i, \nu_j | s_i, s_j, m_i, m_j) P(J_{ij} | s_i, s_j) P(s_i, s_j | m_i, m_j) P(m_i) P(m_j) && (101)\\
&= \sum_{s_i, s_j} \int \mathrm{d}m_i \, \mathrm{d}m_j \, P(\nu_i, | s_i, m_i) P(m_i) P(\nu_j | s_j, m_j) P(m_j) P(J_{ij} | s_i, s_j) P(s_j | s_i, m_j) b(s_i) && (102)
\end{aligned}$$

Therefore

$$\mathbb{E}^{\mathrm{RS}}[\ln Z_i] = \sum_{s_i, s_j} \int \mathrm{d}m_i \, \mathrm{d}m_j \, \mathrm{d}\nu_i \, \mathrm{d}\nu_j \, P(\nu_i, | s_i, m_i) P(m_i) P(\nu_j | s_j, m_j) P(m_j) P(J_{ij} | s_i, s_j) P(s_j | s_i, m_j) b(s_i) \ln Z_{ij} \tag{103}$$

B Plots

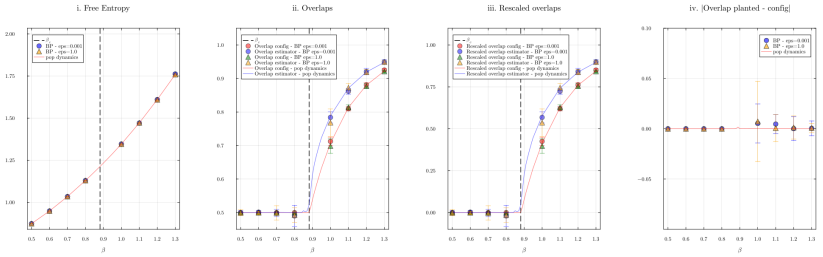

(a) $\kappa = 0$

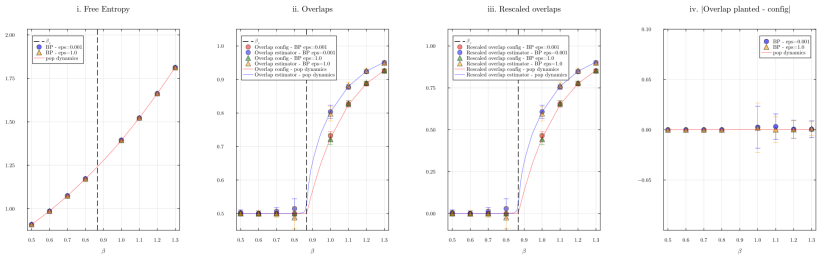

(b) $\kappa = 0.2$

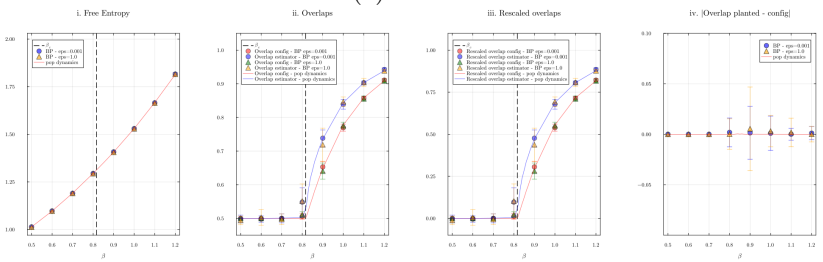

(c) $\kappa = 0.4$

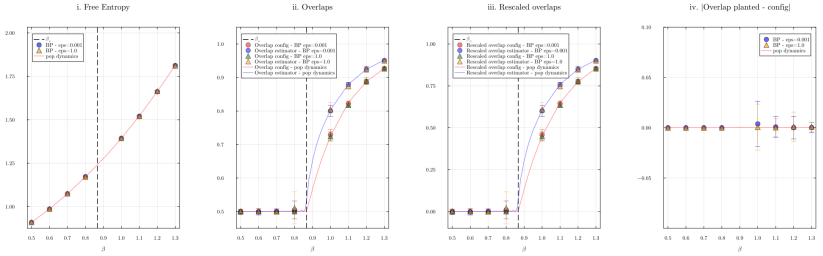

(d) $\kappa = -0.2$

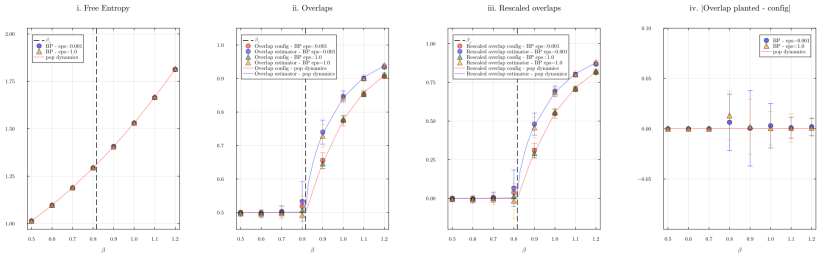

(e) $\kappa = -0.4$

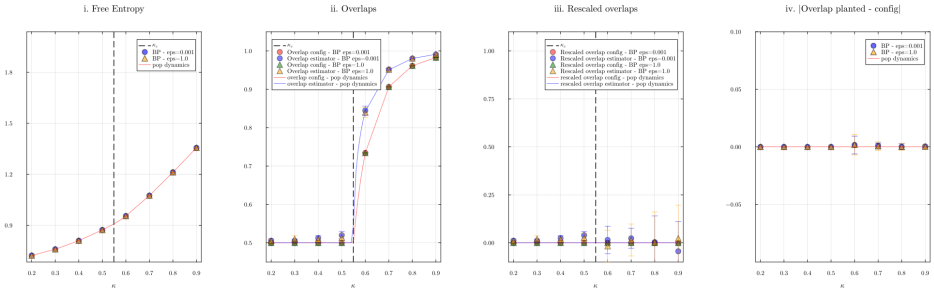

(a) $\beta = 0$

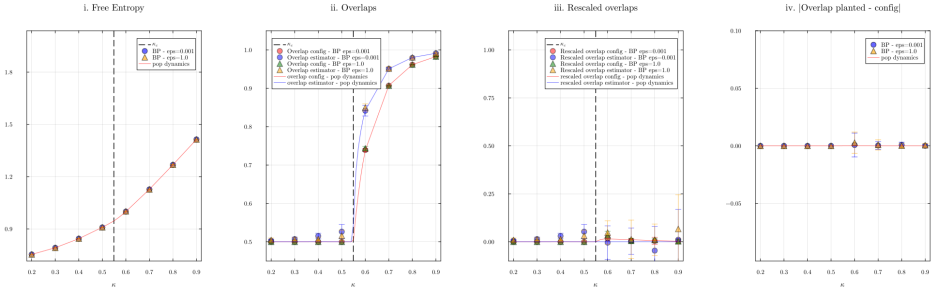

(b) $\beta = 0.2$

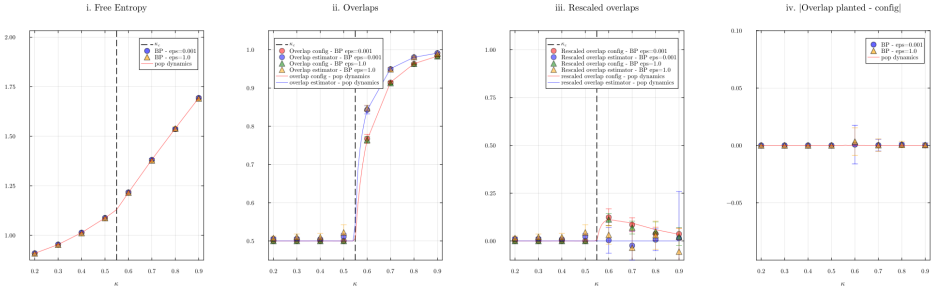

(c) $\beta = 0.5$

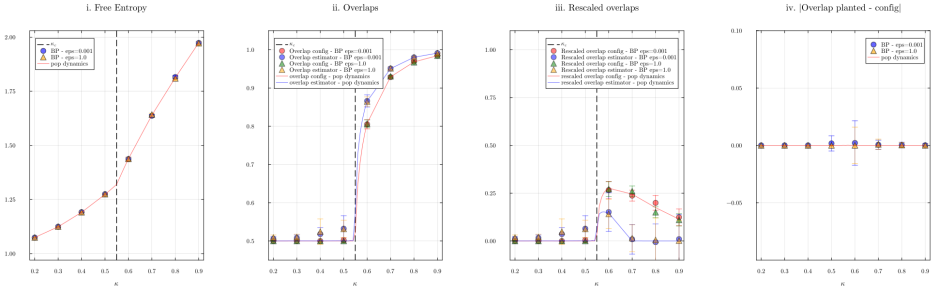

(d) $\beta = 0.7$

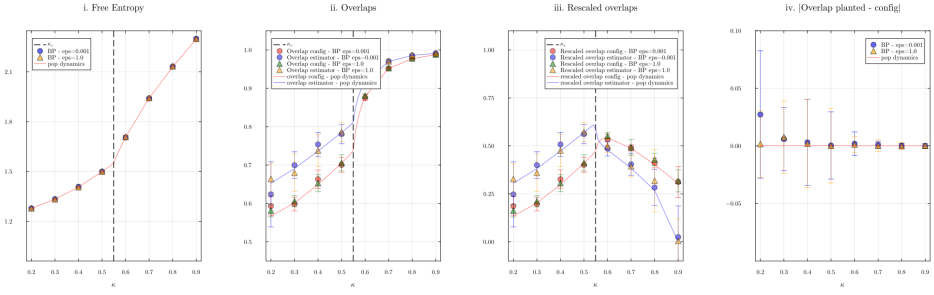

(e) $\beta = 0.9$

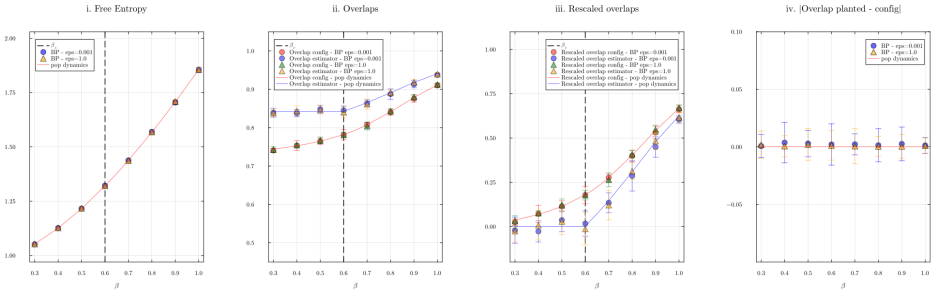

(a) $\kappa = 0.6$

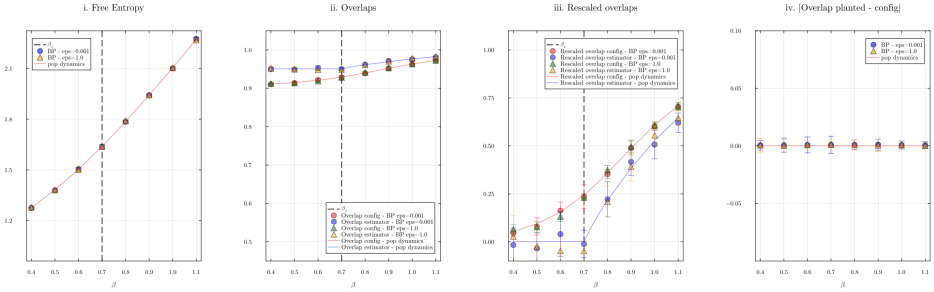

(b) $\kappa = 0.7$

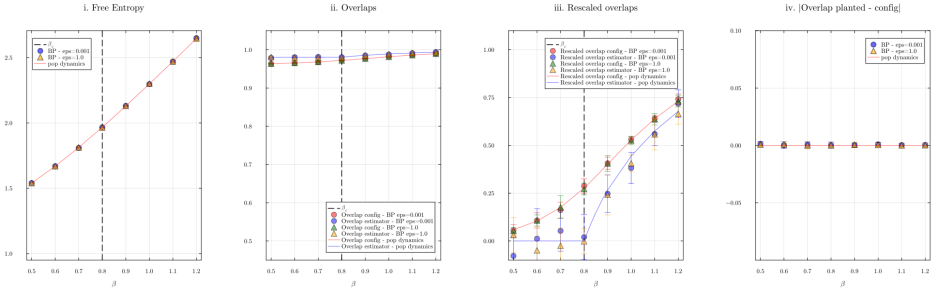

(c) $\kappa = 0.8$

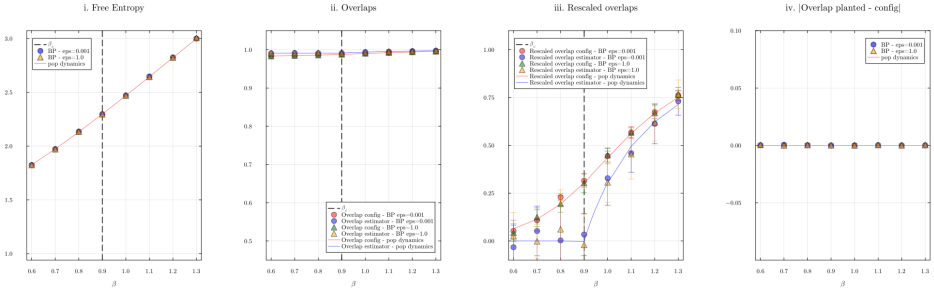

(d) $\kappa = 0.9$

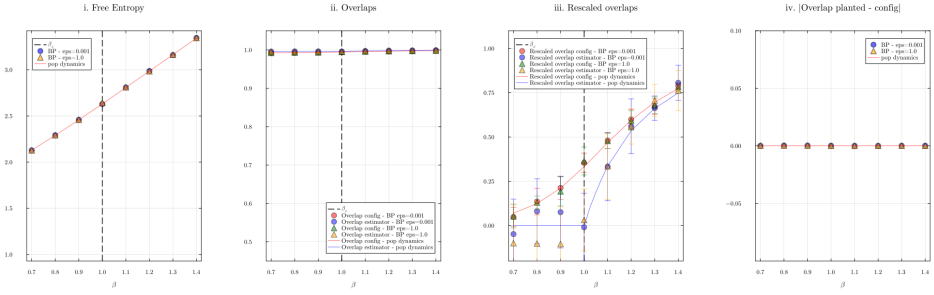

(e) $\kappa = 1.0$