

PAPER • OPEN ACCESS

Neural-prior stochastic block model

To cite this article: O Duranthon and L Zdeborová 2023 *Mach. Learn.: Sci. Technol.* **4** 035017

View the [article online](#) for updates and enhancements.

You may also like

- [Detection and localization of change points in temporal networks with the aid of stochastic block models](#)
Simon De Ridder, Benjamin Vandermarliere and Jan Ryckebusch
- [Pitman–Yor process mixture model for community structure exploration considering latent interaction patterns](#)
Jing Wang, , Kan Li et al.
- [Mean-field theory of graph neural networks in graph partitioning](#)
Tatsuro Kawamoto, Masashi Tsubaki and Tomoyuki Obuchi



PAPER

Neural-prior stochastic block model

OPEN ACCESS

RECEIVED
17 April 2023REVISED
3 July 2023ACCEPTED FOR PUBLICATION
10 July 2023PUBLISHED
17 August 2023

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



O Duranthon* and L Zdeborová

Statistical Physics of Computation laboratory (SPOC), EPFL, Lausanne, Switzerland

* Author to whom any correspondence should be addressed.

Keywords: stochastic block model SBM, generative priors, belief propagation BP, approximate message passing AMP, a benchmark for GNN**Abstract**

The stochastic block model (SBM) is widely studied as a benchmark for graph clustering aka community detection. In practice, graph data often come with node attributes that bear additional information about the communities. Previous works modeled such data by considering that the node attributes are generated from the node community memberships. In this work, motivated by a recent surge of works in signal processing using deep neural networks as priors, we propose to model the communities as being determined by the node attributes rather than the opposite. We define the corresponding model; we call it the neural-prior SBM. We propose an algorithm, stemming from statistical physics, based on a combination of belief propagation and approximate message passing. We analyze the performance of the algorithm as well as the Bayes-optimal performance. We identify detectability and exact recovery phase transitions, as well as an algorithmically hard region. The proposed model and algorithm can be used as a benchmark for both theory and algorithms. To illustrate this, we compare the optimal performances to the performance of simple graph neural networks.

1. Introduction

The stochastic block model (SBM) is widely studied as a benchmark for graph clustering aka community detection, see e.g. reviews (Fortunato 2010, Abbe 2017, Peixoto 2019). The standard version of the SBM observed a graph of connections and the goal is to recover the communities from the knowledge of the graph.

However, in practice, graph data often come with node attributes that bear additional information about the communities. In such a case there are several sources of information on communities one can use: the structure of the graph (as in the standard SBM), and the features or attributes of the nodes. Past work developed algorithms and models accounting for such node information. Among the well-known is the CESNA model of Yang *et al* (2013) where the attributes are generated via logistic regression on the community membership. Another model that recently became popular in the context of benchmarking graph neural networks (GNNs) (e.g. Chien *et al* 2021, Tsitsulin *et al* 2021, Fountoulakis *et al* 2022) is the contextual SBM (Binkiewicz *et al* 2017, Deshpande *et al* 2018), where communities determine centroids for a Gaussian mixture model generating the node-features. In both these examples, the node attributes are generated via conditioning on the community label of the node.

In signal processing, another separate line of work, that witnesses a surge of interest, is modeling signals as the output of a deep generative neural network; for recent reviews see e.g. (Ongie *et al* 2020, Shlezinger *et al* 2020). Deep generative neural networks can be trained on data, and due to their expressivity are able to capture generic structural properties of the signal. In community detection the signal can be seen as the community memberships; following the line of work on deep generative priors it is hence of interest to propose a model where the node attributes are an input of a generative neural network and the node community memberships are the output thereof. In this work, motivated by a recent surge of works in signal processing using deep neural networks as priors, we propose to model the communities as being determined by the node attributes rather than the opposite. We define the corresponding model; that we call the neural-prior SBM.

One of the attires of the SBM is that it is amenable to exact statistical analysis of what is the best achievable performance from an information-theoretic and from algorithmic point of view. This has led to a line of work, originating in statistical physics, where statistical and computational thresholds are analyzed; see e.g. Decelle *et al* (2011b), Abbe *et al* (2015), Abbe (2017). It is valuable to have a solvable case for which we know what is statistically and algorithmically achievable; because in the context of modern machine learning, it is rarely known if or how much the observed performance can be further improved. Asymptotically exact analysis of the detectability threshold was also performed for the contextual SBM (Deshpande *et al* 2018, Lu and Sen 2020). The main topic of the present paper is the statistical physics analysis of optimal algorithmic performance for a simplified version of the proposed neural-prior SBM that we call the generalized-linear-model SBM (GLM–SBM).

The GLM–SBM model we propose can be used for benchmarking GNNs. Since the model is analyzable, we can compare the performance of the evaluated GNN to the optimal algorithmic performance in a non-trivial high-dimensional setting. We treat both the unsupervised and the semi-supervised cases and accompany our paper with an implementation that can be readily used for comparison by GNN developers. As far as we found, a model similar to the neural-prior SBM, we propose here, has been used in Cho *et al* (2022). In that work it is used as a building block for a large neural network; it was not analyzed per se.

A large part of this paper is dedicated to the asymptotic analysis of the GLM–SBM model. We identify how the detectability phase transition well known from the SBM changed under the presence of the GLM-prior. We also unveil an exact recovery phase transition that happens when the prior on the latent variables of the GLM is binary, while the average degree of the SBM remains finite. Such an exact recovery phase at a finite average degree came to us as a surprise and we find it rather remarkable in view of the fact that without the GLM prior exact recovery in the standard SBM is only possible for degrees growing logarithmically with the system size (Abbe *et al* 2015, Abbe 2017). The exact recovery transition is discontinuous and makes the problem algorithmically challenging posing a nice set of parameters that can serve as a benchmark in the attempt of improving GNNs.

2. The neural-prior SBM

2.1. Definition

We consider a set V of $|V| = N$ nodes, a graph $G(V, A)$ on those nodes. Nodes have features/attributes $F_\mu \in \mathbb{R}^M$ of dimension M , $\mu = 1, \dots, N$. The features and the graphs are observed. We aim to divide the set of nodes into q communities with labels $s_\mu \in \{1, \dots, q\}$ in such a way that (a) the graph structure correlates with the labels, e.g. nodes being in the same community are more likely to be connected, and (b) the node attributes F_μ are correlated with the labels.

2.1.1. SBM

In the SBM the edges $A_{\mu\nu}$ of the graph G are generated conditioned on the group memberships s_μ ; we consider the following rule:

$$P_{\text{SBM}}(A_{\mu\nu} = 1 | s_\mu, s_\nu) = \begin{cases} c_i/N & \text{if } s_\mu = s_\nu, \\ c_o/N & \text{if } s_\mu \neq s_\nu, \end{cases} \quad (1)$$

and $A_{\mu\nu} = 0$ otherwise. Here c_i and c_o are the affinity coefficients common to the SBM. We define the affinity matrix whose elements are $c_{s,t} = c_i \delta_{s=t} + c_o \delta_{s \neq t}$. We note that the literature often considers a more general SBM where the affinity matrix has arbitrary elements. The model and analysis proposed in this work could be readily generalized to that case. We consider a slightly restricted version of the SBM purely for simplicity. In the SBM the ground truth group memberships s_μ are generated at random from a prior that only accounts for the sizes of the q groups. The node attributes F are simply ignored in the SBM.

2.1.2. Neural-prior SBM

In neural-prior SBM, that we define here, the group memberships s_μ can be a generic function on the attributes F_μ . Such a function can be represented by a deep neural network and learned from ground-truth data. The training data would be pairs $\{F_\mu, s_\mu\}$ where attributes act as the neural network inputs and the group memberships as output labels. For instance, for a L -layer fully connected neural network this reads

$$s_\mu = \varphi^{(L)} \left(W^{(L)} \dots \varphi^{(2)} \left(W^{(2)} \varphi^{(1)} (W^{(1)} F) \right) \dots \right) \quad (2)$$

for the last activation function $\varphi^{(L)}$ chosen as in multi-class classification tasks.

The aim of this paper is to provide a benchmark model where the optimal performance can be analyzed asymptotically exactly. For this we need to (a) define the corresponding asymptotic limit, (b) consider a simple neural network prior that is amenable to asymptotic analysis. We will also limit ourselves to consider community detection with two groups of the same size only, $q = 2$ (this is not a strong limitation, but is considered in the follow-up for simplicity). With this in mind, in the rest of the paper, we will consider the following model generating the group memberships s_μ .

2.1.3. GLM–SBM

In order to make analysis amenable we will consider the features F to be random and drawn independently as $F_{\mu l} \sim \mathcal{N}(0, 1/M)$. We then consider M latent variables $w_l \sim P_w$, $l = 1, \dots, M$ and generate the community memberships as

$$s_\mu = \text{sign} \left(\sum_l^M F_{\mu l} w_l \right). \quad (3)$$

This corresponds to a single-layer neural network with a sign activation function. Such a neural network is also often referred to as the GLM or as the perceptron. We will hence call this variant of the neural-prior SBM the GLM–SBM.

Concerning the asymptotic limit, we work in the challenging sparse case of SBM. We parameterize the SBM by the standard parameterization

$$c_i = c + \sqrt{c}\lambda, \quad c_o = c - \sqrt{c}\lambda. \quad (4)$$

We then consider $N \rightarrow \infty$ with $c = (c_i + c_o)/2 = \mathcal{O}(1)$ is the average degree, and $\lambda = \mathcal{O}(1)$ is the signal-to-noise ratio. We further work in the high-dimensional limit of the GLM where $N/M = \alpha = \mathcal{O}(1)$, with α being the aspect ratio that will play a role of another signal-to-noise ratio. This is because the higher α the more correlation there is between the group memberships and the easier the community detection should be.

The GLM–SBM differs from the SBM because communities are not independent, conditionally on the features. For instance, in the extreme case $M = 1$, all memberships are known, up to a global flip given by w_1 ; that is to say, they are all very strongly correlated. The GLM–SBM tends toward a SBM when $\alpha \rightarrow 0$. Indeed, for large M , $\sum_l^M F_{\mu l} w_l$ tend to independent Gaussian variables.

2.2. Related work

Anticipating the asymptotic analysis that we are aiming at, we note that such an analysis has been done for the standard SBM in Decelle *et al* (2011a, 2011b) using the belief propagation (BP) algorithm and the cavity method from statistical physics for the asymptotic analysis of its behavior. Concerning semi-supervised learning in the SBM, the information coming from the semi-supervision is readily incorporated into the analysis of the above papers as has been done in Zhang *et al* (2014).

The predictions of Decelle *et al* (2011a, 2011b) have then been partially established rigorously see e.g. Mossel *et al* (2015, 2018), Abbe (2017), Coja-Oghlan *et al* (2017). However, the full conjecture of Decelle *et al* (2011a, 2011b) about the asymptotic exactness of their analysis remains an open question from the mathematical point of view. In this paper, we will use the same techniques as Decelle *et al* (2011a, 2011b) anticipating a follow-up work putting the conjectures about optimality on a rigorous basis. For the GLM, which is defined by a dense graphical model, the rigorous analysis is simpler and was carried out in Barbier *et al* (2019).

The analysis of the GLM–SBM requires to *glue* the two graphical models using the GLM as the prior for the SBM, and the SBM as a source of uncertainty of the outputs of the GLM. Such a glueing of two dense exactly solvable graphical models for developed in Manoel *et al* (2017) with rigorous justifications given in Gabrié *et al* (2018), Aubin *et al* (2019), Gerbelot and Berthier (2021). Our work is the first one, as far as we are aware, where a sparse graphical model (the SBM) is glued to a dense graphical model (the GLM). This can be done heuristically and is conjectured asymptotically exact along the lines of the works of Decelle *et al* (2011a, 2011b). A complete rigorous justification would have to be preceded by the proof of the conjecture for the SBM that is still open.

The contextual stochastic block model (CSBM) introduced and studied theoretically in Binkiewicz *et al* (2017), Deshpande *et al* (2018) is another version of the SBM incorporating node information. In the CSBM

the node information is modeled via a Gaussian mixture model with each community having their own centroid. From the analysis point of view, this model takes into account two sources of observation about the latent variables—the community memberships. This is hence different from the GLM–SBM where one model serves as a prior for the other instead of as an independent source of information. Modulo this difference, some of the analysis performed for the CSBM is related to our work. Notably, the detectability threshold and the linearized message passing algorithm presented in Deshpande *et al* (2018), Lu and Sen (2020) are obtained in a similar manner in which we obtain the detectability phase transition and the linearized algorithm. We note that the semi-supervised version of the CSBM has not been analyzed, but this could be done rather straightforwardly using the same methods as in Zhang *et al* (2014).

3. Bayes-optimal estimation of communities

We consider the GLM–SBM as defined above and aim to analyze the Bayes-optimal inference of the community structure. We will consider in general the semi-supervised setting where next to the structure of the graph A and the covariates F we observe the communities for a subset Ξ of the nodes, $\rho = |\Xi|/N$. We denote by s the vector of unobserved nodes and s_Ξ the vector of observed nodes. The unsupervised case is then recovered as the special case where Ξ is an empty set, $\rho = 0$.

The analysis of this paper is set in the so-called Bayes-optimal setting where we know the details of the GLM–SBM model. The only quantity that we do not observe is the ground truth values of the latent variables w that generate the group memberships s . For the group memberships, we only observe a fraction ρ of them in the semi-supervised setting and none of them in the unsupervised setting.

The optimal inference is then done using the posterior distribution over the unobserved communities

$$P(s|A, s_\Xi, F) = \frac{P(A|s, s_\Xi, F)P_{\text{prior}}(s|s_\Xi, F)}{Z(A, s_\Xi, F)} = \frac{P_{\text{prior}}(s|s_\Xi, F)}{Z(A, s_\Xi, F)} \prod_{\mu < \nu} P_{\text{SBM}}(A_{\mu\nu}|s_\mu, s_\nu) \tag{5}$$

where $Z(A, F, s_\Xi)$ is the normalization constant. We used here the definition of the GLM–SBM model that implies $P(A|s, s_\Xi, F) = P(A|s, s_\Xi)$. For GLM–SBM the prior on s is

$$P_{\text{prior}}(s|s_\Xi, F) = \frac{1}{Z(s_\Xi, F)} \int dw P_w(w) \prod_{\mu} \left[P_{s,\mu}(s_\mu) P_0\left(s_\mu \left| \sum_l F_{\mu l} w_l \right. \right) \right] \tag{6}$$

where we define $P_0(t|z) = \delta_{t=\text{sign}(z)}$ the output distribution and $P_{s,\mu}$ the additional prior distribution, which is used to inject information about the membership of node μ :

$$P_{s,\mu}(t) = \begin{cases} \delta_{t=s_\mu} & \text{if } \mu \in \Xi, \\ 1/2 & \text{if } \mu \notin \Xi. \end{cases} \tag{7}$$

In equation (6) we marginalize over the latent variable w . However, since the estimation of the latent variable w is crucial in order to exploit the full power of the prior (6) it will be instrumental to consider the posterior as a joint probability of the unobserved nodes and the latent variable

$$P(s, w|A, s_\Xi, F) = \frac{P_w(w)}{\tilde{Z}(A, s_\Xi, F)} \prod_{\mu} \left[P_{s,\mu}(s_\mu) P_0\left(s_\mu \left| \sum_l F_{\mu l} w_l \right. \right) \right] \prod_{\mu < \nu} P_{\text{SBM}}(A_{\mu\nu}|s_\mu, s_\nu) \tag{8}$$

\tilde{Z} is the Bayesian evidence. We define the free entropy of the problem as its logarithm:

$$\phi(A, s_\Xi, F) = \frac{1}{N} \log \tilde{Z}(A, s_\Xi, F). \tag{9}$$

We seek an estimator \hat{s} that maximizes the overlap with the ground truth. The Bayes-optimal estimator \hat{s} that maximizes it is given by

$$\hat{s}_\mu^{\text{MMO}} = \operatorname{argmax}_t p_\mu(t) \tag{10}$$

where p_μ is the marginal posterior probability of node μ . Using the ground truth values s_μ of the communities the maximal mean overlap is then computed as

$$\text{MMO} = \frac{1}{N} \sum_{\mu=1}^N \delta_{\hat{s}_\mu^{\text{MMO}}, s_\mu}. \tag{11}$$

To estimating the latent variable w , we consider minimizing the mean squared error via the MMSE estimator

$$\hat{w}_l^{\text{MMSE}} = \int ds dw P(s, w | A, s_\Xi, F) w_l \tag{12}$$

i.e. \hat{w}^{MMSE} is the mean of the posterior distribution. Again using the ground truth values w_l of the latent variables the MMSE is then computed as

$$\text{MMSE} = \frac{1}{M} \sum_{l=1}^M (\hat{w}_l^{\text{MMSE}} - w_l)^2. \tag{13}$$

The problem is invariant by a global sign flip of s and w so in practice we measure the following overlaps

$$q_s = \frac{|\hat{s} \cdot s|}{N}, \quad q_w = \frac{|\hat{w} \cdot w|}{\|\hat{w}\|_2 \|w\|_2}. \tag{14}$$

In general, the Bayes-optimal estimation requires the evaluation of the averages over the posterior that is in general exponentially costly in N and M . In the next section, we will derive the AMP–BP algorithm and argue that, in the limit $N \rightarrow \infty$ and $M \rightarrow \infty$ with $N/M = \alpha = \mathcal{O}(1)$ and all other parameters being of $\mathcal{O}(1)$ this algorithm approximates the MMSE and MMO estimators with an error that vanishes. We give more precise statements below.


4. The AMP–BP algorithm

To retrieve the communities for the GLM–SBM, our main results rely on an algorithm that we call AMP–BP. We conjecture that in the large system size, this algorithm cannot be beaten by another polynomial algorithm. We can also extract the so-called hard phases where the randomly initialized algorithm fails, but an exponentially costly algorithm would succeed; we do this using an informed initialization and the free entropy. We then analyze the performance of the algorithm and the associated phase transitions.

4.1. Algorithm

The algorithm is based on BP and approximate message-passing (AMP). BP was used to solve SBM in Decelle *et al* (2011b) and conjectured asymptotically optimal among efficient algorithms in doing so. AMP was used to solve GLM, see e.g. Donoho *et al* (2009), Krzakala *et al* (2012), and again conjectured asymptotically optimal among efficient algorithms in doing so with strong evidence for this being provided by Celentano *et al* (2021). We glue these two algorithms together along the lines of Manoel *et al* (2017), Aubin *et al* (2019) to solve the GLM–SBM; we call the resulting algorithm AMP–BP. Using statistical physics arguments analogous to those in Decelle *et al* (2011b), Krzakala *et al* (2012) we conjecture that it provides asymptotically optimal performance in the considered cases.

We derive the AMP–BP algorithm for the GLM–SBM starting from the factor graph of the problem:



The χ s and ψ s are probability distributions on the variables s and w ; they are called cavity messages. We write the BP equations for these distributions that read:

$$\chi_{w_l}^{l \rightarrow \mu} \propto P_w(w_l) \prod_{\nu \neq \mu} \psi_{w_l}^{\nu \rightarrow l} \quad (15)$$

$$\psi_{w_l}^{\nu \rightarrow l} \propto \sum_{s_\nu} \chi_{s_\nu}^{\nu \rightarrow \nu} \int \prod_{m \neq l} (dw_m \chi_{w_m}^{m \rightarrow \nu}) P_0(s_\nu | F_\nu \cdot w) \quad (16)$$

$$\psi_{s_\nu}^{\nu \rightarrow \nu} \propto \int \prod_m (dw_m \chi_{w_m}^{m \rightarrow \nu}) P_0(s_\nu | F_\nu \cdot w) \quad (17)$$

$$\chi_{s_\mu}^{\mu \rightarrow \mu} \propto P_{s,\mu}(s_\mu) \prod_{\nu \neq \mu} \psi_{s_\mu}^{\nu \rightarrow \mu} \quad (18)$$

$$\chi_{s_\mu}^{\mu \rightarrow \nu} \propto P_{s,\mu}(s_\mu) \psi_{s_\mu}^{\mu \rightarrow \mu} \prod_{\eta \neq \mu, \nu} \psi_{s_\mu}^{\eta \rightarrow \mu} \quad (19)$$

$$\psi_{s_\nu}^{\mu \rightarrow \nu} \propto \sum_{s_\mu} \chi_{s_\mu}^{\mu \rightarrow \nu} P_{\text{SBM}}(A_{\mu\nu} | s_\mu, s_\nu). \quad (20)$$

The proportionality signs \propto denote that all the messages are non-negative numbers summing to one over their lower indices, the corresponding normalization factors being omitted in our notation.

These BP equations still include a high-dimensional integral and hence cannot be implemented efficiently. We simplify them to obtain AMP–BP by using the central limit theorem on the dense side of the graphical model and keeping only the means and variances of the resulting Gaussians. This is standard in the derivation of the AMP algorithm, see e.g. Krzakala *et al* (2012). The details of this derivation are given in appendix A.

In order to state the final algorithm, we introduce the denoising function:

$$g_0(\omega, \chi, V) = \frac{\int dz \sum_s \chi_s P_0(s|z) (z - \omega) e^{-(z-\omega)^2/2V}}{V \int dz \sum_s \chi_s P_0(s|z) e^{-(z-\omega)^2/2V}}. \quad (21)$$

We define the input functions as

$$f_a(\Lambda, \Gamma) = \frac{\int dw P_w(w) w e^{-\Lambda w^2/2 + \Gamma w}}{\int dw P_w(w) e^{-\Lambda w^2/2 + \Gamma w}}, \quad f_v(\Lambda, \Gamma) = \partial_\Gamma f_a(\Lambda, \Gamma). \quad (22)$$

We denote by Z s the normalization factors obtained so that the messages sum to one over their lower indices.

To give some intuition we explain what are the variables AMP–BP employs. a_l is an estimation of the posterior mean of w_l , v_l is an estimation of its variance; ω_μ is an estimation of the mean of $\sum_l F_{\mu l} w_l$ and V an estimation of its variance. $\psi_{s_\mu}^{\mu \rightarrow \mu}$ is a marginal distribution on s_μ , as estimated by the AMP on the GLM side, while $\chi_{s_\mu}^{\mu \rightarrow \mu}$ is the distribution as estimated by the BP on the SBM side. Γ_l is a proxy for estimating the mean of w_l in absence of the prior P_w and Λ is for the variance. h_t can be interpreted as an external field enforcing the nodes not to be in the same group; $\chi_{s_\mu}^{\mu \rightarrow \nu}$ is a marginal distribution on s_μ (these variables are the messages of a sum-product message-passing algorithm) and $\chi_{s_\mu}^\mu$ is the estimated posterior marginal on s_μ , that we are interested in.

The AMP–BP algorithm reads:

AMP–BP

input features $F_{\mu l}$, graph G , affinity matrix $c_{s,t}$, prior information $P_{s,\mu}$.

Initialize $a_l^{(0)} = \epsilon_l$, $v_l^{(0)} = 1$, $g_{o,\mu}^{(0)} = 0$, $\chi_{s_\mu}^{\mu \rightarrow \nu, (0)} = \frac{1}{2} + s_\mu \epsilon^{\mu \rightarrow \nu}$, $\chi_{s_\mu}^{\mu \rightarrow \mu, (0)} = \frac{1}{2}$, $\chi_{s_\mu}^{\mu, (0)} = \frac{1}{2}$, $t = 0$; where ϵ s are zero-mean small random variables.

repeat

AMP update of ω_μ, V_μ

$$V^{(t+1)} \leftarrow \frac{1}{M} \sum_l v_l^{(t)}$$

$$\omega_\mu^{(t+1)} \leftarrow \sum_l F_{\mu l} a_l^{(t)} - V^{(t+1)} g_{o,\mu}^{(t)}$$

AMP update of $\psi^{\mu \rightarrow \mu}, g_{o,\mu}, \Lambda, \Gamma_l$

$$\psi_{s_\mu}^{\mu \rightarrow \mu, (t+1)} \leftarrow \int \frac{dz P_0(s_\mu | z)}{\sqrt{2\pi V_\mu^{(t+1)}}} e^{-\frac{(z - \omega_\mu^{(t+1)})^2}{2V_\mu^{(t+1)}}}$$

$$g_{o,\mu}^{(t+1)} \leftarrow g_o(\omega_\mu^{(t+1)}, \chi^{\mu \rightarrow \mu, (t)}, V^{(t+1)})$$

$$\Lambda^{(t+1)} \leftarrow \frac{1}{M} \sum g_{o,\mu}^{2, (t+1)}$$

$$\Gamma_l^{(t+1)} \leftarrow \Lambda^{(t+1)} a_l^{(t)} + \sum_\mu F_{\mu l} g_{o,\mu}^{(t+1)}$$

AMP update of the estimated marginals a_l, v_l

$$a_l^{(t+1)} \leftarrow f_a(\Lambda^{(t+1)}, \Gamma_l^{(t+1)})$$

$$v_l^{(t+1)} \leftarrow f_v(\Lambda^{(t+1)}, \Gamma_l^{(t+1)})$$

BP update of the field h

$$h_s^{(t+1)} \leftarrow \frac{1}{N} \sum_\mu \sum_{s_\mu} c_{s,s_\mu} \chi_{s_\mu}^{\mu, (t)}$$

BP update of the messages $\chi^{\mu \rightarrow \nu}$ for $(\mu\nu) \in G$ and of the marginals χ^μ

$$\chi_{s_\mu}^{\mu \rightarrow \nu, (t+1)} \leftarrow \frac{P_{s,\mu}(s_\mu)}{Z^{\mu \rightarrow \nu}} e^{-h_{s_\mu}^{(t+1)}} \psi_{s_\mu}^{\mu \rightarrow \mu, (t+1)} \prod_{\eta \in \partial \mu \setminus \nu} \sum_{s_\eta} c_{s_\eta, s_\mu} \chi_{s_\eta}^{\eta \rightarrow \mu, (t)}$$

$$\chi_{s_\mu}^{\mu, (t+1)} \leftarrow \frac{P_{s,\mu}(s_\mu)}{Z^\mu} e^{-h_{s_\mu}^{(t+1)}} \psi_{s_\mu}^{\mu \rightarrow \mu, (t+1)} \prod_{\eta \in \partial \mu} \sum_{s_\eta} c_{s_\eta, s_\mu} \chi_{s_\eta}^{\eta \rightarrow \mu, (t)}$$

BP update of the SBM-to-GLM messages $\chi^{\mu \rightarrow \mu}$

$$\chi_{s_\mu}^{\mu \rightarrow \mu, (t+1)} \leftarrow \frac{P_{s,\mu}(s_\mu)}{Z^{\mu \rightarrow \mu}} e^{-h_{s_\mu}^{(t+1)}} \prod_{\eta \in \partial \mu} \sum_{s_\eta} c_{s_\eta, s_\mu} \chi_{s_\eta}^{\eta \rightarrow \mu, (t)}$$

$t \leftarrow t + 1$

until convergence of a_l, v_l, χ^μ

output estimated mean a_l and variance v_l of w_l and marginal distribution χ^μ of s_μ .

We provide an implementation of AMP–BP in the supplementary material. It is also available from our repository¹.

We draw attention to the output function g_o that covers the difference between AMP for GLM–SBM and AMP for GLM alone. In AMP for GLM alone g_o depends on the observed labels while here we use their estimated marginals. On the other side, the difference between BP for GLM–SBM and BP for SBM alone are the messages $\psi^{\mu \rightarrow \mu}$ in BP update. $\psi^{\mu \rightarrow \mu}$ can be interpreted as the conditional probability of s_μ given w without SBM.

4.1.1. Estimators

The Bayes-optimal estimators of s and w are obtained according to equations (10) and (12). Expressed using the AMP–BP messages they become

$$\hat{s}_\mu^{\text{AMP-BP}} = \text{sign}(2\chi_+^\mu - 1), \quad \hat{w}_l^{\text{AMP-BP}} = a_l \quad (23)$$

where χ_+^μ is the estimated marginal probability of the event $s_\mu = +1$ and a_l is the estimated mean of w_l .

¹ <https://gitlab.epfl.ch/spoc-idephics/glm-sbm>.

4.1.2. Free entropy

We express also the free entropy ϕ in terms of the messages and variables of AMP–BP at the fixed point; it is called the Bethe free entropy ϕ_{Bethe} . The derivation from the factor graph is done in appendix B. Up to a term that diverges with N we obtain that the Bethe free entropy is

$$\phi_{\text{Bethe}} = \phi_{\text{SBM}} + \phi_{\text{GLM}} \tag{24}$$

$$\begin{aligned} \phi_{\text{SBM}} &= \frac{1}{N} \sum_{\mu} \log \sum_{s_{\mu}} P_{s_{\mu}}(s_{\mu}) e^{-h_{s_{\mu}}} \prod_{\eta \in \partial \mu} \sum_{s_{\eta}} c_{s_{\eta}, s_{\mu}} \chi_{s_{\eta}}^{\eta \rightarrow \mu} \\ &\quad - \frac{1}{N} \sum_{(\mu\nu) \in G} \log \sum_{s_{\mu}, s_{\nu}} c_{s_{\mu}, s_{\nu}} \chi_{s_{\mu}}^{\mu \rightarrow \nu} \chi_{s_{\nu}}^{\nu \rightarrow \mu} + \frac{c}{2} \end{aligned} \tag{25}$$

$$\begin{aligned} \phi_{\text{GLM}} &= \frac{1}{N} \sum_{\mu} \log \int \frac{dz_{\mu}}{\sqrt{2\pi V_{\mu}}} \sum_{s_{\mu}} \chi_{s_{\mu}}^{\mu \rightarrow \mu} P_{\theta}(s_{\mu} | z_{\mu}) e^{-(z_{\mu} - \omega_{\mu})^2 / 2V_{\mu}} \\ &\quad + \frac{1}{N} \sum_l \log \int dw_l P_w(w_l) e^{-\Lambda_l w_l^2 / 2 + \Gamma_l w_l} \\ &\quad + \frac{1}{N} \left(\sum_l \frac{\Lambda_l}{2} (a_l^2 + v_l) - \Gamma_l a_l + \sum_{\mu} \frac{(\omega_{\mu} - \sum_l F_{\mu l} a_l)^2}{2V_{\mu}} \right). \end{aligned} \tag{26}$$

If the AMP–BP has more than one fixed point then the free entropy serves to select the fixed point of AMP–BP that corresponds to Bayes-optimal performance. It is the one with the largest free entropy that should be selected.

We compare later the free entropy of the fixed point of AMP–BP to the free entropy of the fully informative point where $q_S = q_W = 1$. We write it ϕ_{info} . At this point the messages are delta functions of the ground truth; we can derive ϕ_{info} directly from the factor graph and it reads

$$\phi_{\text{info}} = \frac{1}{\alpha} \mathbb{E}_{P_w} \log P_w + \frac{1}{N} \sum_{(\mu\nu) \in G} \log c_{s_{\mu}, s_{\nu}} - \frac{c}{2} - (1 - \rho) \log 2. \tag{27}$$

4.2. Asymptotic optimality conjecture

We conjecture that AMP–BP gives the Bayes-optimal estimator for GLM–SBM in the following sense.

We define the two possible initializations: (a) random initialization, where we initialize the messages randomly according to their prior distribution, adding no information, as described in the algorithm above; and (b) informed initialization, where we initialize the estimators to delta functions of the true values of s and w .

We consider the fixed point of AMP–BP that has the largest Bethe free entropy ϕ_{Bethe} . We argue that it suffices to check the random and informed initializations to find all the relevant fixed points.

We conjecture that, asymptotically exactly, the AMP–BP fixed point that has the largest ϕ_{Bethe} provides the Bayes-optimal estimators for the GLM–SBM model. Its overlap q_S is asymptotically equal to the Bayes-optimal MMO overlap and ϕ_{Bethe} is equal to ϕ , with high probability as $N \rightarrow +\infty$. This is aligned with the same conjecture for BP and the standard SBM from Decelle *et al* (2011a) and the proofs of this property for the AMP algorithm and the pure GLM model in Barbier *et al* (2019).

5. Bayes-optimal estimation with AMP–BP and phase transitions

5.1. Gaussian prior, 2nd order transition to partial recovery

In this subsection we consider the GLM prior P_w to be a standard Gaussian. The GLM then produces binary labels, the group memberships, with the same probability of being in each of the groups.

We conjecture that for this prior the fixed point of AMP–BP reached from random initialization always corresponds to the Bayes-optimal estimation and no computationally hard phase is present. We observe that the algorithm converges to the same fixed point for the two possible initializations. In figure 7 in the appendix E we illustrate that the system size we use is close enough to the thermodynamical limit in the sense that the change in the curves is small when the size is changed.

The accuracy AMP–BP achieves is depicted in figures 1 (unsupervised case) and 8 (semi-supervised case, in appendix E). We observe that the larger the snr λ or the aspect ratio α the better the recovery. The recovery is eased when community memberships are explained by a few features i.e. when α is large.

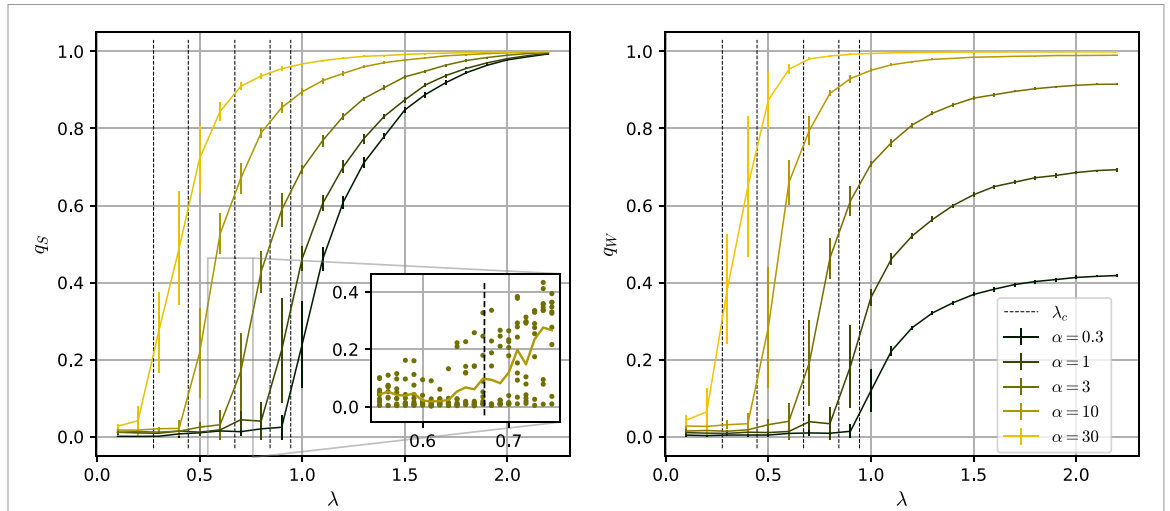


Figure 1. Left and right: overlaps q_S (group membership estimation) and q_W (the GLM latent vector estimation) of the fixed point of AMP–BP, vs λ for a range of compression ratios α . Vertical dashed lines: theoretical thresholds λ_c to partial recovery, equation (29). $N = 10^4$, $c = 5$, P_w Gaussian. We run ten experiments per point. Inset: we plot the ten data points and their mean.

In the unsupervised case, we observe a phase transition from a non-informative fixed point $q_S = q_W = 0$ to an informative fixed point $q_S > 0, q_W > 0$. The transition is located at a particular critical threshold λ_c . This transition is well known for standard SBM, which is recovered here in the $\alpha \rightarrow 0$ limit, where $\lambda_c = 1$ for $q = 2$. The transition is of 2nd order; this means that the overlaps vary continuously with respect to λ . In the semi-supervised case the 2nd order transition disappears.

5.1.1. Linearization, spectral algorithm

λ_c can be computed by a linear stability analysis of the non-informative fixed point of AMP–BP: at a given λ , if the algorithm is not stable it will move away from the non-informative fixed point to the informative fixed point. The linearization of the algorithm is done in appendix C. We obtain the following update equation:

$$x^{\mu \rightarrow \nu, (t+1)} = \frac{\lambda}{\sqrt{c}} \left(\sum_{\eta \in \partial \mu \setminus \nu} x^{\eta \rightarrow \mu, (t)} + \frac{2}{\pi} \sum_{\eta} (FF^T - I_N)_{\mu, \eta} \sum_{\rho \in \partial \eta} x^{\rho \rightarrow \eta, (t-1)} \right) \quad (28)$$

where the x s are real random variables and $(FF^T)_{\mu\nu} = \sum_l F_{\mu l} F_{\nu l}$. Taking the variance of this equation and averaging over the realizations of the graph we obtain the stability criterion

$$1 = \lambda_c^2 \left(1 + \frac{4\alpha}{\pi^2} \right). \quad (29)$$

Equation (28) can be interpreted as a spectral algorithm; we apply iteratively a linear operator to the variables

$$x = \frac{\lambda}{\sqrt{c}} (B + G\tilde{B}) x \quad (30)$$

where

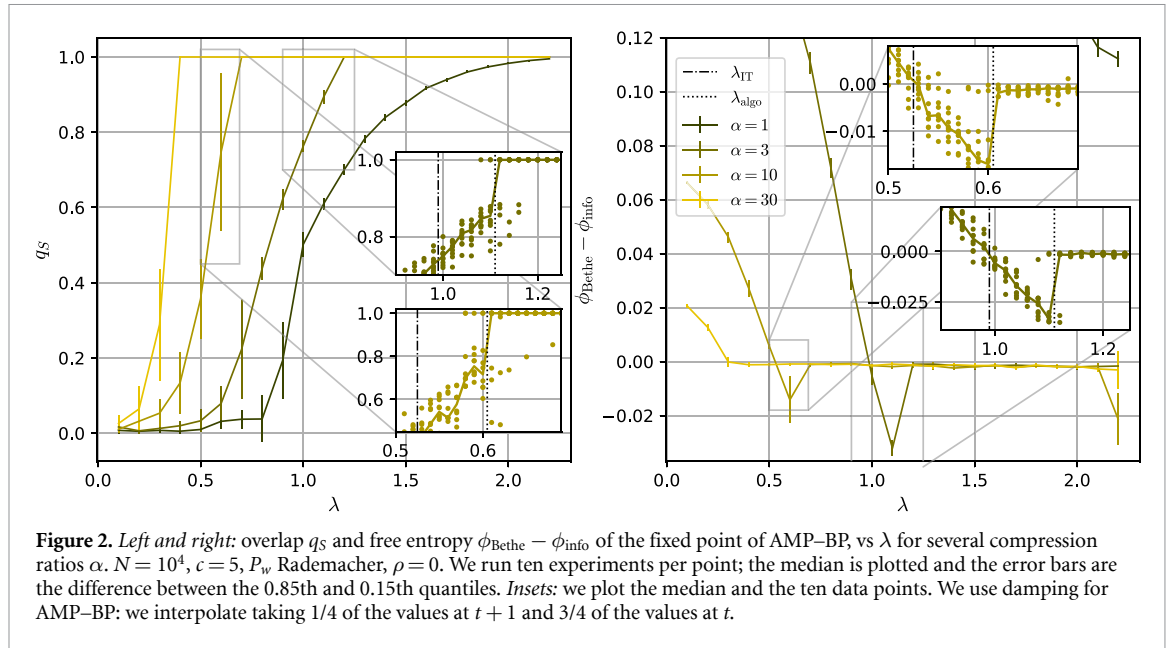
$$B_{\mu \rightarrow \nu, \rho \rightarrow \eta} = \delta_{\mu=\eta}(1 - \delta_{\rho=\nu}), \quad G_{\mu \rightarrow \nu, \eta} = \frac{2}{\pi} (FF^T - I_N)_{\mu, \eta}, \quad \tilde{B}_{\eta, \mu \rightarrow \rho} = \delta_{\rho=\eta}. \quad (31)$$

B is the non-backtracking matrix (Krzakala *et al* 2013). Such a spectral algorithm will share the phase transition at signal-to-noise ratio (snr) given by equation (28). The study of the resulting overlap is also of interest, but we do not consider it in the present article.

5.2. Binary prior, 1st order transition to exact recovery

In this subsection the GLM prior is considered to be $P_w = (\delta_{w=1} + \delta_{w=-1})/2$ Rademacher. This still produces two groups with unbiased sizes.

The fixed point AMP–BP achieved from a random initialization is depicted on figures 2 and 9 (in appendix E). We observe it admits the same transition to partial recovery at λ_c , equation (29), as the Gaussian prior does; this is also predicted by the linearization of the previous part.



For values of $\alpha > \alpha_{\text{algo}}$ (that we determine below), we observe another transition; it is discontinuous, from partial recovery $q_S > 0, q_W > 0$ to exact recovery $q_S = q_W = 1$. There is a value λ_{algo} such that for $\lambda > \lambda_{\text{algo}}$ randomly initialized AMP–BP recovers the group memberships exactly for all nodes. The overlap q_S does not vary continuously at λ_{algo} ; over the many independent trials we observe that there is an interval of overlaps below 1 that cannot be reached by AMP–BP for any λ .

Discontinuous thresholds are related to the existence of several fixed points of AMP–BP and to first-order phase transitions. A 1st order phase transition is located by comparing the free entropies ϕ_{Bethe} of the various fixed points. We notice that next to the AMP–BP fixed point that is reached from a random initialization, the exact recovery point is a fixed point at all values of λ and α (still considering P_w binary). In the region of λ and α where these two fixed points differ we need to compare their free entropies. The fixed point with larger free entropy describes the Bayes-optimal performance that can in general be better than the one of AMP–BP. The difference between the free entropies of the fixed point reached by AMP–BP from random initialization and the informative fixed point is depicted on the rhs of figure 2. We see that for $\lambda < \lambda_{\text{IT}}$ the fixed point reached from random initialization has larger free entropy $\phi_{\text{Bethe}} > \phi_{\text{info}}$ and hence describes the optimal performance. In the region $\lambda_{\text{IT}} < \lambda < \lambda_{\text{algo}}$ the informative fixed point has larger free entropy $\phi_{\text{info}} > \phi_{\text{Bethe}}$, but randomly initialized AMP–BP does not reach it. This is an algorithmically hard phase where exact recovery is statistically possible, but the AMP–BP algorithm is sub-optimal. At the same time the AMP–BP algorithm is conjectured optimal among efficient algorithms (Gamarnik *et al* 2022) and thus the hardness of this phase is believed to be intrinsic. For $\lambda > \lambda_{\text{algo}}$ we only find the exact recovery fixed point.

An exact recovery for the standard SBM is only achievable for graphs of average degrees c diverging logarithmically with the size of the system (Abbe *et al* 2015), where the logarithm comes from a type of coupon collector problem. The existence of an exact recovery phase in graphs of constant degrees is novel as far as we know. It nicely illustrates the power of the GLM prior that is able to induce it. It is well known that a 1st order phase transition appears for GLM alone with binary weights and known labels (Györgyi 1990, Sompolinsky *et al* 1990, Barbier *et al* 2019). We note, however, that in the GLM–SBM the labels are not observed directly but via the graph. It is thus not *a priori* clear that an exact recovery phase can appear. Without our analysis its existence would not be easy to anticipate.

Let us finally derive the values α_{algo} above which the exact recovery phase exists. We consider the limit $\lambda = \sqrt{c}$; then the graph G consists of two disconnected components, one for each community; and AMP–BP performs as AMP for GLM alone, up to a global sign. We take into account the proportion e^{-c} of nodes that are isolated and do not bring information. We obtain that

$$\alpha_{\text{algo}} = \alpha_{\text{algo,perceptron}} (1 - e^{-c})^{-1} \quad (32)$$

where $\alpha_{\text{algo,perceptron}} \approx 1.493$ is the algorithmic critical compression ratio of the binary perceptron (Barbier *et al* 2019). Similarly the λ_{IT} will exist above

$$\alpha_{\text{IT}} = \alpha_{\text{IT,perceptron}} (1 - e^{-c})^{-1} \quad (33)$$

where $\alpha_{\text{algo,perceptron}} \approx 1.249$ is the information-theoretic critical threshold of the binary perceptron (Györgyi 1990, Sompolinsky *et al* 1990, Barbier *et al* 2019).

The 1st order phase transition λ_{IT} and its spinodal λ_{algo} are still present in the semi-supervised case $\rho > 0$, for small values of ρ , see figure 9 in appendix E, contrary to the 2nd order phase transition to partial recovery that vanishes in the semi-supervised case. Moreover, for $\rho > \alpha_{\text{algo,perceptron}}/\alpha$, perfect recovery is achieved at any λ , because one has enough train labels to infer w .

6. Analysis in the dense limit

As in Decelle *et al* (2011a, 2011b) for the sparse SBM, the analysis of the AMP–BP is based on the numerical investigation of the fixed points and their free entropies on systems large enough that the behavior is representative of the large-size limit. This is also at the basis of the mathematical difficulty to establish this prescription rigorously. At the same time, a dense version of the SBM has been proposed and studied fully rigorously in Lesieur *et al* (2017), Miolane (2017). This rigorous analysis has then been extended to include the GLM prior in Aubin *et al* (2019) that studies a generic instance of low-rank matrix factorization problem with a generative prior. We hence study the phenomenology of the AMP–BP algorithm in the limit of large degree c , where it becomes a special case of the framework developed in Aubin *et al* (2019).

The dense limit is defined by taking $p_i, p_o = \mathcal{O}(1)$ and $p_i - p_o = \mathcal{O}(1/\sqrt{N})$. SBM is then a low-rank matrix factorization problem. It is parameterized by its signal-to-noise ratio Δ_I (which is defined as the inverse variance of an equivalent additive Gaussian channel). We need Δ_I as a function of the parameters of the SBM, that is to say to equalize their signal-to-noise ratios. We compute the Fisher information of the channel $P_{\text{SBM}}(A_{\mu\nu} = 1|x_{\mu\nu}) = \frac{1}{N}c_o + \frac{1}{\sqrt{N}}(c_i - c_o)x_{\mu\nu}$, where $x_{\mu\nu} = \frac{\delta_{\mu=\nu}}{\sqrt{N}}$ is taken to zero. The mapping is then

$$\Delta_I = \frac{1}{4} \frac{N(p_i - p_o)^2}{p_o(1 - p_o)} = \frac{c\lambda^2}{c - \sqrt{c}\lambda} + \mathcal{O}\left(\frac{c}{N}\right) \quad (34)$$

where $p_i = c_i/N$ and $p_o = c_o/N$. It is of order one in both sparse case and dense case. Also, we add the factor $1/4$ to obtain a phase transition at $\Delta_I = 1$ in the dense case when $\alpha = 0$. In the following $\rho = 0$.

Aubin *et al* (2019) give the algorithm corresponding to the dense case of AMP–BP algorithm. We reproduce it in appendix D. Its performances can be tracked by a few scalar equations that are named state evolution (SE) equations. For P_w Rademacher they read:

$$q_w^{t+1} = \mathbb{E}_\xi \left[Z_w \left(\sqrt{\hat{q}_w^t} \xi, \hat{q}_w^t \right) f_w \left(\sqrt{\hat{q}_w^t} \xi, \hat{q}_w^t \right)^2 \right] \quad (35)$$

$$\hat{q}_w^t = \alpha \mathbb{E}_{\xi, \eta} \left[Z_o \left(\sqrt{q_s^t} \Delta_I \xi, q_s^t \Delta_I, \sqrt{q_w^t} \eta, 1 - q_w^t \right) f_o \left(\sqrt{q_s^t} \Delta_I \xi, q_s^t \Delta_I, \sqrt{q_w^t} \eta, 1 - q_w^t \right)^2 \right] \quad (36)$$

$$q_s^{t+1} = \mathbb{E}_{\xi, \eta} \left[Z_o \left(\sqrt{q_s^t} \Delta_I \xi, q_s^t \Delta_I, \sqrt{q_w^t} \eta, 1 - q_w^t \right) f_s \left(\sqrt{q_s^t} \Delta_I \xi, q_s^t \Delta_I, \sqrt{q_w^t} \eta, 1 - q_w^t \right)^2 \right] \quad (37)$$

where q_s and q_w are the s- and w-overlaps, Δ_I is the signal-to-noise ratio of the problem, ξ and η are standard Gaussians and

$$Z_w(\gamma, \Lambda) = e^{\Lambda/2} \cosh \gamma, \quad f_w(\gamma, \Lambda) = \partial_\gamma \log Z_w \quad (38)$$

$$Z_o(B, A, \omega, V) = e^{-A/2} (\cosh(B) + \sinh(B) \text{erf}(\omega/\sqrt{2V})) \quad (39)$$

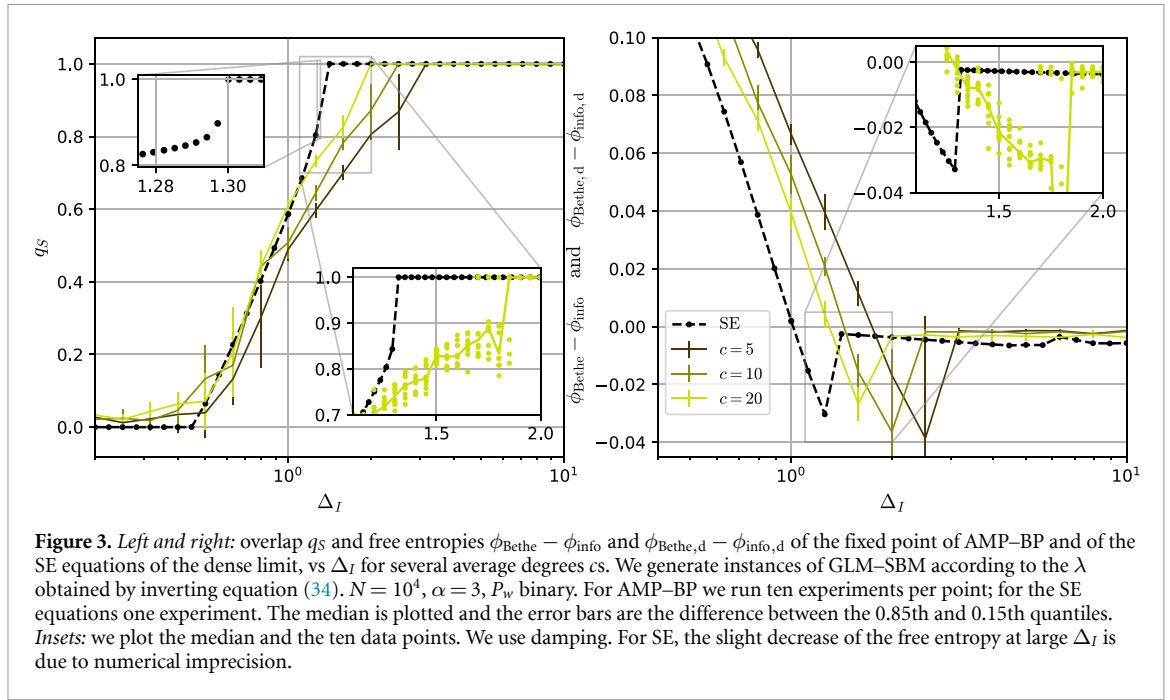
$$f_o(B, A, \omega, V) = \partial_\omega \log Z_o, \quad f_s(B, A, \omega, V) = \partial_B \log Z_o. \quad (40)$$

Aubin *et al* (2019) gives also the free entropy of the fixed point of the algorithm for the dense problem. It reads

$$\phi_{\text{Bethe,d}}(q_s, q_w, \hat{q}_w) = -\frac{\Delta_I}{4} q_s^2 - \frac{1}{2\alpha} \hat{q}_w q_w + \psi_o(\Delta_I q_s, q_w) + \frac{1}{\alpha} \psi_w(\hat{q}_w) \quad (41)$$

where

$$\psi_o(\Delta_I q_s, q_w) = \mathbb{E}_{\xi, \eta} \text{xlogx} Z_o \left(\sqrt{\Delta_I q_s} \xi, \Delta_I q_s, \sqrt{q_w} \eta, 1 - q_w \right) \quad (42)$$



$$\psi_w(\hat{q}_w) = \mathbb{E}_\xi x \log x Z_w \left(\sqrt{\hat{q}_w} \xi, \hat{q}_w \right) \tag{43}$$

$x \log x$ being the function $x \rightarrow x \log x$.

The convergence to the dense limit is quite fast; the large degree results are close to the observed results even for c quite small. Numerically it appears that $c \approx 20$ is enough ($N = 10^4$) to already observe quite small difference, see figure 3.

The fully informative fixed point is $(q_s, \hat{q}_w, q_w) = (1, +\infty, 1)$. Its free entropy is

$$\phi_{\text{info,d}} = -\frac{\log 2}{\alpha} + \frac{\Delta_I}{4}. \tag{44}$$

The analysis of the system of SE equations is done in appendix D; we summarize the four main points: (a) the fully informative fixed point is stable for all Δ_I ; (b) the width of its stability domain shrinks to zero when Δ_I tends to zero; (c) a general necessary condition to observe a fully informative fixed point is that P_w does not admit everywhere a twice differentiable density; (d) the algorithmic critical compression ratio $\alpha_{\text{algo,d}}$ is close to $\alpha_{\text{algo,preceptron}}$.

We also obtain an approximation for the critical point λ_c of the transition to partial recovery. Aubin *et al* (2019) gives us that in the dense limit, the critical snr is

$$\Delta_{I,c} = \left(1 + 4 \frac{\alpha}{\pi^2} \right)^{-1}. \tag{45}$$

The limit $c = \omega(1)$ large gives $\lambda_c = (1 + 4\alpha/\pi^2)^{-1/2} + \mathcal{O}(c/N)$, as predicted by the linearization.

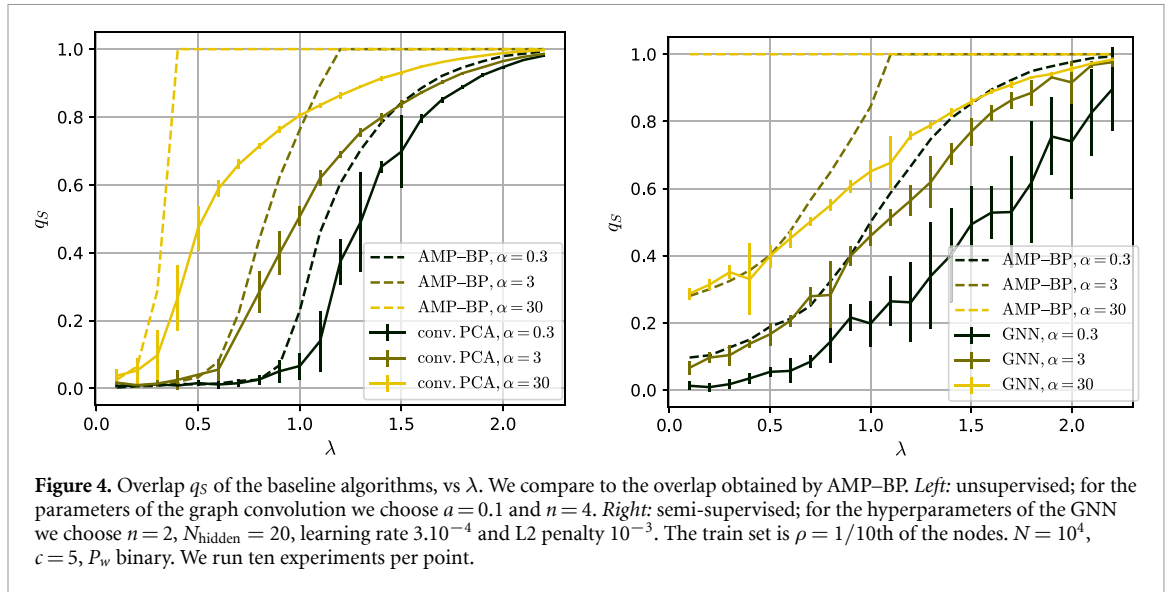
7. Comparison of performance with standard GNNs on GLM–SBM

GLM–SBM can be used as a benchmark for clustering or classification tasks on attributed graphs. We compare two simple baselines with AMP–BP. We show that GLM–SBM is simple to define yet challenging algorithmically, in particular in the case of binary prior close to the first order phase transition.

7.1. An unsupervised baseline

The algorithm is inspired by graph convolution networks; it performs binary clustering. We compare its performances to the optimal ones given by AMP–BP. Its performances are shown on figures 4 left (P_w binary) and 10 left (P_w Gaussian, in appendix E).

Data is generated according to the GLM–SBM. We stack the features $F_{\mu l}$ into vectors $F_\mu^{(0)} \in \mathbb{R}^M$ or a matrix $F \in \mathbb{R}^{N \times M}$. The observed graph G is used for the convolution steps.



We compute n steps of graph convolution on the features; perform PCA on the transformed features and keep the largest component; threshold its left vector to obtain the membership of each node. Formally, we consider the features $F_\mu^{(0)} \in \mathbb{R}^M$; we apply n times

$$F_\mu^{(t+1)} = F_\mu^{(t)} + a \sum_{\nu \in \partial \mu} F_\nu^{(t)} \tag{46}$$

where a is a scalar. We apply PCA on the new matrix \hat{F} whose rows are $F_\mu^{(n)}$. Writing $u \in \mathbb{R}^N$ the left vector of its largest component, the estimator is $\hat{s} = \text{sign}(u)$. We tune n and a empirically to optimize the recovery. We observe that roughly it depends on n and a only by their product an . Also, the optimal a scales like $1/c$.

7.2. A semi-supervised baseline

The algorithm is a simple GNN, trained in a semi-supervised way for node classification.

Again data is generated according to the GLM–SBM, with $\rho = 1/10$. We stack the features $F_{\mu l}$ into vectors $F_\mu^{(0)} \in \mathbb{R}^M$. We use the observed graph G for the message-passing steps.

The GNN is made of a two-layer perceptron and a readout layer for the binary classification. It reads:

$$F_\mu^{(t+1)} = F_\mu^{(t)} + B \text{relu} \left(A \sum_{\nu \in \partial \mu} F_\nu^{(t)} \right), \quad \hat{s}_\mu = w^T F_\mu^{(n)} \tag{47}$$

where A is $N_{\text{hidden}} \times M$ learnable, B is $M \times N_{\text{hidden}}$ learnable, $w \in \mathbb{R}^M$ learnable and n is the number of steps. We train it given the labels of the subset of nodes Ξ . We use gradient descent with logistic loss, momentum and L2 regularization. We do not fine-tune the hyperparameters. Its performances are shown on figures 4 right (P_w binary) and 10 right (P_w Gaussian, in appendix E).

We also performed experiments where the GNN is made of a single-layer perceptron (no relu), as Cheng *et al* (2022) does on CSBM. The performances are similar to the multi-layer perceptron, but it requires much more parameters to be trained (M^2 vs MN_{hidden} , and we take $N_{\text{hidden}} = \mathcal{O}(1)$).

7.3. Conclusion on the comparison

As to the GLM–SBM dataset, figure 4 illustrates that, both in the unsupervised and the semi-supervised settings, the baseline methods have a considerable gap to the optimal performances given by the AMP–BP algorithm. The GLM–SBM setting is hence suitable to develop GNN algorithms that are able to provide higher accuracy.

As to the AMP–BP algorithm, it is very scalable. It has a running time similar to the GNN-based approaches, around a few minutes per point on figure 1 or 2 (including the ten experiments). Its complexity is $\mathcal{O}(NM)$ in time and in memory. This is the smallest any algorithm can do, for reading the input. The number of steps needed for convergence does not depend on N .

8. Conclusions

We propose a model of attributed graphs. It is a sparse SBM where the nodes carry features that determine their community memberships. We solve it, in the sense that we derive an algorithm that is conjectured to perform optimally among polynomial algorithms. We analyze a linearization of the algorithm and the dense limit of the model. The model, yet simple, exhibits a rich phenomenology with detectable and exact recovery phase transitions. It can be used as a challenging benchmark for GNNs.

In the analysis of this paper we only considered two groups. For more than two groups, $q > 2$, the analysis can also be done by writing an AMP–BP algorithm; just, the AMP-side would need to correspond to a single-layer network with multi-class output. The AMP for such a model has been written and studied in Cornacchia *et al* (2022) and one would have to merge it with the BP of Decelle *et al* (2011a). Another generalization that would be possible to analyze is when the attributes F are drawn from a Gaussian with a generic covariance. This can be done along the lines of Loureiro *et al* (2021). On the other hand considering as a prior the multi-layer neural network (2) with learned weights W would be more challenging; a corresponding AMP algorithm that would provide an asymptotically exact solution is not known.

A future direction of work could also be to theoretically analyze the learning of GLM–SBM by a GNN, i.e. to give insights on the generalization performance of the neural network of section 7; as Cheng *et al* (2022) does for a perceptron-based graph convolution network on CSBM. This would be interesting because few theoretical works address the generalization ability of GNNs.

Data availability statement


No new data were created or analysed in this study.

Acknowledgment

We acknowledge funding from the ERC under the European Union’s Horizon 2020 Research and Innovation Program Grant Agreement 714608-SMiLe.

Appendix A. Derivation of the algorithm

We write BP for this problem. We start with the factor graph. It contains six different messages:



These messages satisfy these equations:

$$\chi_{w_l}^{l \rightarrow \mu} \propto P_w(w_l) \prod_{\nu \neq \mu} \psi_{w_l}^{\nu \rightarrow l} \tag{48}$$

$$\psi_{w_l}^{\nu \rightarrow l} \propto \sum_{s_\nu} \chi_{s_\nu}^{\nu \rightarrow \nu} \int \prod_{m \neq l} (dw_m \chi_{w_m}^{m \rightarrow \nu}) P_0(s_\nu | F_\nu, w) \tag{49}$$

$$\psi_{s_\nu}^{\nu \rightarrow \nu} \propto \int \prod_m (dw_m \chi_{w_m}^{m \rightarrow \nu}) P_0(s_\nu | F_\nu, w) \tag{50}$$

$$\chi_{s_\mu}^{\mu \rightarrow \mu} \propto P_{s,\mu}(s_\mu) \prod_{\nu \neq \mu} \psi_{s_\mu}^{\nu \rightarrow \mu} \tag{51}$$

$$\chi_{s_\mu}^{\mu \rightarrow \nu} \propto P_{s,\mu}(s_\mu) \psi_{s_\mu}^{\mu \rightarrow \mu} \prod_{\eta \neq \mu, \nu} \psi_{s_\mu}^{\eta \rightarrow \mu} \tag{52}$$

$$\psi_{s_\nu}^{\mu \rightarrow \nu} \propto \sum_{s_\mu} \chi_{s_\mu}^{\mu \rightarrow \nu} P_{\text{SBM}}(A_{\mu\nu} | s_\mu, s_\nu). \tag{53}$$

We can plug ψ messages into the χ to obtain, for the GLM part:

$$\chi_{w_l}^{l \rightarrow \mu} = \frac{P_w(w_l)}{Z^{l \rightarrow \mu}} \prod_{\nu \neq \mu} \left[\sum_{s_\nu} \chi_{s_\nu}^{\nu \rightarrow \nu} \int \prod_{m \neq l} (dw_m \chi_{w_m}^{m \rightarrow \nu}) P_0(s_\nu | F_\nu, w) \right] \tag{54}$$

the marginals

$$\chi_{w_l}^l = \frac{P_w(w_l)}{Z^l} \prod_{\nu} \left[\sum_{s_\nu} \chi_{s_\nu}^{\nu \rightarrow \nu} \int \prod_{m \neq l} (dw_m \chi_{w_m}^{m \rightarrow \nu}) P_0(s_\nu | F_\nu, w) \right] \tag{55}$$

where the Z s are normalization factors; and for the SBM part:

$$\chi_{s_\mu}^{\mu \rightarrow \nu} \propto P_{s,\mu}(s_\mu) \psi_{s_\mu}^{\mu \rightarrow \mu} \prod_{\eta \neq \mu, \nu} \sum_{s_\eta} \chi_{s_\eta}^{\eta \rightarrow \mu} P_{\text{SBM}}(A_{\mu\eta} | s_\mu, s_\eta) \tag{56}$$

the marginals

$$\chi_{s_\mu}^\mu \propto P_{s,\mu}(s_\mu) \psi_{s_\mu}^{\mu \rightarrow \mu} \prod_{\eta \neq \mu} \sum_{s_\eta} \chi_{s_\eta}^{\eta \rightarrow \mu} P_{\text{SBM}}(A_{\mu\eta} | s_\mu, s_\eta) \tag{57}$$

and

$$\chi_{s_\mu}^{\mu \rightarrow \mu} \propto P_{s,\mu}(s_\mu) \prod_{\nu \neq \mu} \sum_{s_\nu} \chi_{s_\nu}^{\nu \rightarrow \mu} P_{\text{SBM}}(A_{\mu\nu} | s_\mu, s_\nu). \tag{58}$$

A.1. SBM

We can apply the standard simplifications for sparse SBM (Decelle *et al* 2011b, Zdeborová and Krzakala 2016). We consider only messages on G . This gives

$$\chi_{s_\mu}^{\mu \rightarrow \mu} = \frac{1}{Z^{\mu \rightarrow \mu}} P_{s,\mu}(s_\mu) e^{-h_{s_\mu}} \prod_{\eta \in \partial \mu} \sum_{s_\eta} c_{s_\eta, s_\mu} \chi_{s_\eta}^{\eta \rightarrow \mu} \tag{59}$$

$$\chi_{s_\mu}^{\mu \rightarrow \nu} = \frac{1}{Z^{\mu \rightarrow \nu}} P_{s,\mu}(s_\mu) \psi_{s_\mu}^{\mu \rightarrow \mu} e^{-h_{s_\mu}} \prod_{\eta \in \partial \mu \setminus \nu} \sum_{s_\eta} c_{s_\eta, s_\mu} \chi_{s_\eta}^{\eta \rightarrow \mu} \tag{60}$$

and the marginals

$$\chi_{s_\mu}^\mu = \frac{1}{Z^\mu} P_{s,\mu}(s_\mu) \psi_{s_\mu}^{\mu \rightarrow \mu} e^{-h_{s_\mu}} \prod_{\eta \in \partial \mu} \sum_{s_\eta} c_{s_\eta, s_\mu} \chi_{s_\eta}^{\eta \rightarrow \mu} \tag{61}$$

where $h_s = \frac{1}{N} \sum_{\mu} \sum_{s_\mu} c_{s, s_\mu} \chi_{s_\mu}^\mu$.

A.2. GLM

For the GLM, we follow closely (Zdeborová and Krzakala 2016).

A.2.1. r-BP

We apply first the simplifications that lead to r-BP. We define and consider the inner part of the $\chi_{w_l}^{l \rightarrow \mu}$ message:

$$\tilde{\psi}_{w_l}^{\nu \rightarrow l} = \sum_{s_\nu} \chi_{s_\nu}^{\nu \rightarrow \nu} \int \prod_{m \neq l} (dw_m \chi_{w_m}^{m \rightarrow \nu}) P_0(s_\nu | F_\nu, w). \tag{62}$$

We set $z_\nu = F_{\nu l} w_l + \sum_{m \neq l} F_{\nu m} w_m$. By independence of the w the partial sum behaves like a Gaussian with mean and variance

$$\omega_{\nu \rightarrow l} = \sum_{m \neq l} F_{\nu m} a_{m \rightarrow \nu}, \quad V_{\nu \rightarrow l} = \sum_{m \neq l} F_{\nu m}^2 v_{m \rightarrow \nu} \tag{63}$$

with

$$a_{m \rightarrow \nu} = \int dw_m \chi_{w_m}^{m \rightarrow \nu} w_m \quad v_{m \rightarrow \nu} = \int dw_m \chi_{w_m}^{m \rightarrow \nu} w_m^2 - a_{m \rightarrow \nu}^2. \tag{64}$$

We replace the integral over all w s by a Gaussian integral over z_ν ; we obtain

$$\tilde{\psi}_{w_l}^{\nu \rightarrow l} = \sum_{s_\nu} \chi_{s_\nu}^{\nu \rightarrow \nu} \int \frac{dz_\nu}{\sqrt{2\pi V_{\nu \rightarrow l}}} e^{-(z_\nu - F_{\nu l} w_l - \omega_{\nu \rightarrow l})^2 / 2V_{\nu \rightarrow l}} P_0(s_\nu | z_\nu). \tag{65}$$

We can simplify. $F_{\nu l}$ is small, we expand the exponential:

$$e^{-\frac{(z_\nu - F_{\nu l} w_l - \omega_{\nu \rightarrow l})^2}{2V_{\nu \rightarrow l}}} = e^{-\frac{(z_\nu - \omega_{\nu \rightarrow l})^2}{2V_{\nu \rightarrow l}}} \left(1 - \frac{F_{\nu l}^2 w_l^2}{2V_{\nu \rightarrow l}} + \frac{(z_\nu - \omega_{\nu \rightarrow l}) F_{\nu l} w_l}{V_{\nu \rightarrow l}} + \frac{(z_\nu - \omega_{\nu \rightarrow l})^2 F_{\nu l}^2 w_l^2}{2V_{\nu \rightarrow l}^2} \right). \tag{66}$$

We introduce the denoising function; its expression differs from the one of Zdeborová and Krzakala (2016):

$$g_o(\omega, \chi, V) = \frac{\int dz \sum_s \chi_s P_0(s|z) (z - \omega) e^{-(z - \omega)^2 / 2V}}{V \int dz \sum_s \chi_s P_0(s|z) e^{-(z - \omega)^2 / 2V}}. \tag{67}$$

So

$$\tilde{\psi}_{w_l}^{\nu \rightarrow l} \propto \left(1 - \frac{F_{\nu l}^2 w_l^2}{2V_{\nu \rightarrow l}} + g_o F_{\nu l} w_l + \frac{1}{2} \left(\frac{1}{V_{\nu \rightarrow l}} + \partial_\omega g_o + g_o^2 \right) F_{\nu l}^2 w_l^2 \right) \tag{68}$$

where we evaluate g_o in $(\omega_{\nu \rightarrow l}, \chi^{\nu \rightarrow \nu}, V_{\nu \rightarrow l})$. We exponentiate:

$$\tilde{\psi}_{w_l}^{\nu \rightarrow l} \propto e^{g_o F_{\nu l} w_l + \frac{1}{2} \partial_\omega g_o F_{\nu l}^2 w_l^2}. \tag{69}$$

We take the product of the $\tilde{\psi}$ to obtain

$$\chi_{w_l}^{l \rightarrow \mu} \propto P_w(w_l) e^{-\Lambda_{l \rightarrow \mu} w_l^2 / 2 + \Gamma_{l \rightarrow \mu} w_l} \tag{70}$$

where

$$\Lambda_{l \rightarrow \mu} = - \sum_{\nu \neq \mu} \partial_\omega g_o(\omega_{\nu \rightarrow l}, \chi^{\nu \rightarrow \nu}, V_{\nu \rightarrow l}) F_{\nu l}^2, \quad \Gamma_{l \rightarrow \mu} = \sum_{\nu \neq \mu} g_o(\omega_{\nu \rightarrow l}, \chi^{\nu \rightarrow \nu}, V_{\nu \rightarrow l}) F_{\nu l}. \tag{71}$$

We close the loop defining the input functions

$$f_a(\Lambda, \Gamma) = \frac{\int dw P_w(w) w e^{-\Lambda w^2 / 2 + \Gamma w}}{\int dw P_w(w) e^{-\Lambda w^2 / 2 + \Gamma w}}, \quad f_v(\Lambda, \Gamma) = \partial_\Gamma f_a(\Lambda, \Gamma) \tag{72}$$

so

$$a_{l \rightarrow \mu} = f_a(\Gamma_{l \rightarrow \mu}, \Lambda_{l \rightarrow \mu}), \quad v_{l \rightarrow \mu} = f_v(\Gamma_{l \rightarrow \mu}, \Lambda_{l \rightarrow \mu}). \tag{73}$$

The mean and the variance of the marginals are estimated by

$$a_l = f_a(\Gamma_l, \Lambda_l), \quad v_l = f_v(\Gamma_l, \Lambda_l) \tag{74}$$

where

$$\Lambda_l = - \sum_\nu \partial_\omega g_o(\omega_{\nu \rightarrow l}, \chi^{\nu \rightarrow \nu}, V_{\nu \rightarrow l}) F_{\nu l}^2, \quad \Gamma_l = \sum_\nu g_o(\omega_{\nu \rightarrow l}, \chi^{\nu \rightarrow \nu}, V_{\nu \rightarrow l}) F_{\nu l}. \tag{75}$$

We obtain also the expression of the GLM-to-SBM message


$$\psi_{s_\mu}^{\mu \rightarrow \mu} = \frac{1}{\sqrt{2\pi V_\mu}} \int dz P_0(s_\mu | z) e^{-(z - \omega_\mu)^2 / 2V_\mu} \tag{76}$$

where


$$\omega_\mu = \sum_m F_{\mu m} a_{m \rightarrow \mu}, \quad V_\mu = \sum_m F_{\mu m}^2 v_{m \rightarrow \mu}. \tag{77}$$

A.2.2. Time indices

There are two possibilities for mixing the GLM part and the SBM part:



or



We try both; we do not observe any numerical difference.

A.2.3. AMP

Then we go from r-BP to AMP. We remove the dependence of the messages on the target. We keep only the marginals. The derivation is given by Zdeborová and Krzakala (2016). We obtain that

$$V_\mu^{(t+1)} = \sum_l F_{\mu m}^2 v_l^{(t)} \tag{78}$$

$$\omega_\mu^{(t+1)} = \sum_l F_{\mu l} a_l^{(t)} - V_\mu^{(t+1)} g_{o,\mu}^{(t)} \tag{79}$$

$$g_{o,\mu}^{(t+1)} = g_o(\omega_\mu^{(t+1)}, \chi^{\mu \rightarrow \mu, (t)}, V_\mu^{(t+1)}) \tag{80}$$

$$\Lambda_l^{(t+1)} = - \sum_\mu F_{\mu l}^2 \partial_\omega g_o(\omega_\mu^{(t+1)}, \chi^{\mu \rightarrow \mu, (t)}, V_\mu^{(t+1)}) \tag{81}$$

$$\Gamma_l^{(t+1)} = \Lambda_l^{(t+1)} a_l^{(t)} + \sum_\mu F_{\mu l} g_{o,\mu}^{(t+1)}. \tag{82}$$

A.2.4. Further simplifications

$F_{\mu m}^2$ self-averages. We can replace it by its average $1/M$ in equations (78) and (81). So Λ and V become scalars. Also, on average, $-\partial_\omega g_{o,\mu} = g_{o,\mu}^2$. We obtain the algorithm given in the main part.

Appendix B. Free entropy

We start with the factor graph. The Bethe free entropy is the sum of the free entropies of the nodes plus the factors minus the edges i.e.

$$\begin{aligned}
 N\phi_{\text{Bethe}} = & \sum_\mu \phi^\mu + \sum_l \phi^l + \sum_{\mu < \nu} \phi^{\mu\nu} + \sum_\mu \phi^{\mu\mu} \\
 & - \sum_{\mu \neq \nu} \phi^{\mu \rightarrow \nu} - \sum_\mu \phi^{\mu \rightarrow \mu} - \sum_{l, \mu} \phi^{l \rightarrow \mu}
 \end{aligned}$$

where

$$\phi^\mu = \log \sum_{s_\mu} P_{s,\mu}(s_\mu) \psi_{s_\mu}^{\mu \rightarrow \mu} \prod_{\nu \neq \mu} \psi_{s_\mu}^{\nu \rightarrow \mu} = \log Z^\mu \prod_{\nu \neq \mu} \frac{1}{Z_\psi^{\nu \rightarrow \mu}} \tag{83}$$

$$\phi^l = \log \int dw_l P_w(w_l) \prod_{\mu} \psi_{w_l}^{\mu \rightarrow l} = \log Z^l \prod_{\mu} \frac{1}{Z_\psi^{\mu \rightarrow l}} \tag{84}$$

$$\phi^{\mu\nu} = \log \sum_{s_\mu, s_\nu} \chi_{s_\mu}^{\mu \rightarrow \nu} \chi_{s_\nu}^{\nu \rightarrow \mu} P(A_{\mu\nu} | s_\mu, s_\nu) \tag{85}$$

$$\phi^{\mu\mu} = \log \sum_{s_\mu} \int \prod_m (dw_m \chi_{w_m}^{m \rightarrow \mu}) \chi_{s_\mu}^{\mu \rightarrow \mu} P_0(s_\mu | F_\mu, w) \tag{86}$$

$$\phi^{\mu \rightarrow \nu} = \log \sum_{s_\mu} \chi_{s_\mu}^{\mu \rightarrow \nu} \psi_{s_\mu}^{\nu \rightarrow \mu} = \log \frac{1}{Z_\psi^{\nu \rightarrow \mu}} + \phi^{\mu\nu} \tag{87}$$

$$\phi^{\mu \rightarrow \mu} = \log \sum_{s_\mu} \chi_{s_\mu}^{\mu \rightarrow \mu} \psi_{s_\mu}^{\mu \rightarrow \mu} = \log \frac{1}{Z^{\mu \rightarrow \mu}} + \phi^\mu \tag{88}$$

$$\phi^{l \rightarrow \mu} = \log \int dw_l \chi_{w_l}^{l \rightarrow \mu} \psi_{w_l}^{\mu \rightarrow l} = \log \frac{1}{Z_\psi^{\mu \rightarrow l}} + \phi^{\mu\mu}. \tag{89}$$

This simplifies to

$$N\phi_{\text{Bethe}} = \sum_{\mu} \log Z^{\mu \rightarrow \mu} + \sum_l \log Z^l - \sum_{\mu < \nu} \phi^{\mu\nu} + (1 - M) \sum_{\mu} \phi^{\mu\mu}. \tag{90}$$

On the SBM side, we have (Zdeborová and Krzakala 2016)

$$Z^{\mu \rightarrow \mu} = \sum_{s_\mu} P_{s,\mu}(s_\mu) e^{-h_{s_\mu}} \prod_{\eta \in \partial \mu} \sum_{s_\eta} c_{s_\eta, s_\mu} \chi_{s_\eta}^{\eta \rightarrow \mu} \tag{91}$$

$$\sum_{\mu < \nu} \phi^{\mu\nu} = \sum_{(\mu\nu) \in G} \log \sum_{s_\mu, s_\nu} c_{s_\mu, s_\nu} \chi_{s_\mu}^{\mu \rightarrow \nu} \chi_{s_\nu}^{\nu \rightarrow \mu} - N \frac{c}{2}. \tag{92}$$

On the GLM side, we have

$$\phi^{\mu\mu} = \log \int \frac{dz_\mu}{\sqrt{2\pi V_\mu}} \sum_{s_\mu} \chi_{s_\mu}^{\mu \rightarrow \mu} P_o(s_\mu | z_\mu) e^{-(z_\mu - \omega_\mu)^2 / 2V_\mu} \tag{93}$$

$$\log Z^l = \sum_{\mu} \log \hat{Z}^{\mu \rightarrow l} + \log \int dw_l P_w(w_l) e^{-\Lambda_l w_l^2 / 2 + \Gamma_l w_l}. \tag{94}$$

We compute $\log \hat{Z}^{\mu \rightarrow l}$ as a function of the target-free elements (we start using that $\omega_{\mu \rightarrow l} = \omega_\mu - F_{\mu l} a_{l \rightarrow \mu}$ and $V_{\mu \rightarrow l} = V_\mu - F_{\mu l}^2 v_{l \rightarrow \mu}$ and expanding). This gives:

$$\sum_{l, \mu} \log \hat{Z}^{\mu \rightarrow l} = M \sum_{\mu} \phi^{\mu\mu} + \sum_l \frac{\Lambda_l}{2} (a_l^2 + v_l) - \Gamma_l a_l + \sum_{\mu} \frac{(\omega_\mu - \sum_l F_{\mu l} a_l)^2}{2V_\mu} \tag{95}$$

which is what Krzakala *et al* (2014) gives (taking $\Sigma_l = 1/\Lambda_l$ and $R_l = \Gamma_l/\Lambda_l$). Finally, we obtain that the free entropy ϕ is exactly the sum of the free entropies of the two sub-problems:

$$\phi_{\text{Bethe}} = \phi_{\text{SBM}} + \phi_{\text{GLM}} \tag{96}$$

$$\begin{aligned} \phi_{\text{SBM}} &= \frac{1}{N} \sum_{\mu} \log \sum_{s_\mu} P_{s,\mu}(s_\mu) e^{-h_{s_\mu}} \prod_{\eta \in \partial \mu} \sum_{s_\eta} c_{s_\eta, s_\mu} \chi_{s_\eta}^{\eta \rightarrow \mu} \\ &\quad - \frac{1}{N} \sum_{(\mu\nu) \in G} \log \sum_{s_\mu, s_\nu} c_{s_\mu, s_\nu} \chi_{s_\mu}^{\mu \rightarrow \nu} \chi_{s_\nu}^{\nu \rightarrow \mu} + \frac{c}{2} \end{aligned} \tag{97}$$

$$\begin{aligned}
\phi_{\text{GLM}} &= \frac{1}{N} \sum_{\mu} \log \int \frac{dz_{\mu}}{\sqrt{2\pi V_{\mu}}} \sum_{s_{\mu}} \chi_{s_{\mu}}^{\mu \rightarrow \mu} P_{\theta}(s_{\mu} | z_{\mu}) e^{-(z_{\mu} - \omega_{\mu})^2 / 2V_{\mu}} \\
&+ \frac{1}{N} \sum_l \log \int dw_l P_w(w_l) e^{-\Lambda_l w_l^2 / 2 + \Gamma_l w_l} \\
&+ \frac{1}{N} \left(\sum_l \frac{\Lambda_l}{2} (a_l^2 + v_l) - \Gamma_l a_l + \sum_{\mu} \frac{(\omega_{\mu} - \sum_l F_{\mu l} a_l)^2}{2V_{\mu}} \right). \tag{98}
\end{aligned}$$

Appendix C. Linearization and partial recovery threshold

We take $P_{s,\mu}(s) = 1/2$. The non-informative point $q_S = q_W = 0$ is a fixed point of the AMP–BP algorithm. At this point, we have $\chi^{\mu \rightarrow \nu} = \frac{1}{2}$, $\chi^{\mu \rightarrow \mu} = \frac{1}{2}$, $\chi^{\mu \rightarrow \nu} = \frac{1}{2}$, $a_l = 0$, $v_l = 1$, $\omega_{\mu} = 0$, $V = 1$, $\psi^{\mu \rightarrow \mu} = \frac{1}{2}$, $g_{o,\mu} = 0$, $\Lambda = 0$ and $\Gamma_l = 0$.

We linearize the equations of the algorithm around this point. We write $|_*$ the evaluation of functions in this point. We have

$$\delta \chi^{\mu \rightarrow \nu, (t+1)} = \sum_{\eta \in \partial \mu \setminus \nu} \frac{1}{2} \left(\frac{c_{\dots}}{c} - 1 \right) \delta \chi^{\eta \rightarrow \mu, (t)} + \partial_{\omega} \psi^{\mu \rightarrow \mu} |_* \delta \omega_{\mu}^{(t+1)} + \partial_V \psi^{\mu \rightarrow \mu} |_* \delta V^{(t+1)} \tag{99}$$

$$\delta \chi^{\mu \rightarrow \mu, (t+1)} = \sum_{\eta \in \partial \mu} \frac{1}{2} \left(\frac{c_{\dots}}{c} - 1 \right) \delta \chi^{\eta \rightarrow \mu, (t)} \tag{100}$$

$$\delta a_l^{(t+1)} = \partial_{\Lambda} f_a |_* \delta \Lambda^{(t+1)} + \partial_{\Gamma} f_a |_* \delta \Gamma_l^{(t+1)} \tag{101}$$

$$\delta v_l^{(t+1)} = \partial_{\Lambda} \Gamma f_a |_* \delta \Lambda^{(t+1)} + \partial_{\Gamma} \Gamma f_a |_* \delta \Gamma_l^{(t+1)} \tag{102}$$

$$\delta g_{o,\mu}^{(t+1)} = \partial_{\omega} g_o |_* \delta \omega_{\mu}^{(t+1)} + \nabla_{\chi} g_o |_* \delta \chi^{\mu \rightarrow \mu, (t+1)} + \partial_V g_o |_* \delta V^{(t+1)} \tag{103}$$

where we write c_{\dots} for the affinity matrix and where we have used the standard linearization for SBM. We have also

$$\delta \omega_{\mu}^{(t+1)} = \sum_l F_{\mu l} \delta a_l^{(t)} - \delta V^{(t+1)} g_o |_* - V |_* \delta g_{o,\mu}^{(t)} \tag{104}$$

$$\delta V^{(t+1)} = \frac{1}{M} \sum_l \delta v_l^{(t)} \tag{105}$$

$$\delta \Lambda^{(t+1)} = \frac{2}{M} \sum_{\mu} g_o |_* \delta g_{o,\mu}^{(t+1)} \tag{106}$$

$$\delta \Gamma_l^{(t+1)} = \delta \Lambda^{(t+1)} a_l |_* + \Lambda |_* \delta a_l^{(t)} + \sum_{\mu} F_{\mu l} \delta g_{o,\mu}^{(t+1)}. \tag{107}$$

We simplify: $g_o |_* = 0$, $\partial_{\omega} g_o |_* = 0$, $\partial_V g_o |_* = 0$, $\partial_V \psi |_* = 0$ and $\partial_{\Gamma} f_a |_* = 1$ (for P_w both Gaussian or Rademacher). We compute that $\partial_{\omega} \psi |_* = \frac{1}{\sqrt{2\pi}} \begin{pmatrix} 1 \\ -1 \end{pmatrix}$ and $\nabla_{\chi} g_o |_* = \frac{2}{\sqrt{2\pi}} \begin{pmatrix} 1 \\ -1 \end{pmatrix}^T$. We assemble equations together:

$$\delta \chi^{\mu \rightarrow \nu, (t+1)} = \sum_{\eta \in \partial \mu \setminus \nu} \frac{1}{2} \left(\frac{c_{\dots}}{c} - 1 \right) \delta \chi^{\eta \rightarrow \mu, (t)} + \partial_{\omega} \psi |_* \left(\sum_l F_{\mu l} \delta a_l^{(t)} - \nabla_{\chi} g_o |_* \delta \chi^{\mu \rightarrow \mu, (t)} \right) \tag{108}$$

$$\begin{aligned}
\delta \chi^{\mu \rightarrow \mu, (t+1)} &= \sum_{\eta \in \partial \mu} \frac{1}{2} \left(\frac{c_{\dots}}{c} - 1 \right) \delta \chi^{\eta \rightarrow \mu, (t)} \\
&+ \sum_{\eta, l} F_{\mu l} F_{\eta l} (\partial_{\omega} \psi |_* \cdot \nabla_{\chi} g_o |_*) \delta \chi^{\eta \rightarrow \eta, (t)} - (\partial_{\omega} \psi |_* \nabla_{\chi} g_o |_*) \delta \chi^{\mu \rightarrow \mu, (t)} \tag{109}
\end{aligned}$$

$$\delta \chi^{\eta \rightarrow \eta, (t)} = \sum_{\rho \in \partial \eta} \frac{1}{2} \left(\frac{c_{\dots}}{c} - 1 \right) \delta \chi^{\rho \rightarrow \eta, (t-1)}. \tag{110}$$

The matrices $\frac{1}{2} \left(\frac{c_{\nu}}{c} - 1 \right)$ and $\partial_{\omega} \psi |_{*} \nabla_{\chi} g_o |_{*}$ share the same eigenvectors. They have one null eigenvalue and one positive: $\frac{c_{\nu} - c_o}{2c} = \frac{\lambda}{\sqrt{c}}$ and $\frac{2}{\pi}$. We project to obtain

$$x^{\mu \rightarrow \nu, (t+1)} = \frac{\lambda}{\sqrt{c}} \left(\sum_{\eta \in \partial \mu \setminus \nu} x^{\eta \rightarrow \mu, (t)} + \frac{2}{\pi} \sum_{\eta} (FF^T)_{\mu, \eta} \sum_{\rho \in \partial \eta} x^{\rho \rightarrow \eta, (t-1)} - \frac{2}{\pi} \sum_{\eta \in \partial \mu} x^{\eta \rightarrow \mu, (t-1)} \right) \quad (111)$$

$$x^{\mu \rightarrow \nu, (t+1)} = \frac{\lambda}{\sqrt{c}} \left(\sum_{\eta \in \partial \mu \setminus \nu} x^{\eta \rightarrow \mu, (t)} + \frac{2}{\pi} \sum_{\eta} (FF^T - I_N)_{\mu, \eta} \sum_{\rho \in \partial \eta} x^{\rho \rightarrow \eta, (t-1)} \right) \quad (112)$$

where $(FF^T)_{\mu\nu} = \sum_l F_{\mu l} F_{\nu l}$.

We obtain the threshold λ_c of partial recovery taking the variance of the expression 112, discarding the time indices. We use that $(FF^T - I_N)_{\mu, \nu}^2$ averages to $1/M$ if $\mu \neq \nu$ and to $\mathcal{O}(1/M)$ otherwise. We obtain:

$$1 = \lambda_c^2 \left(1 + \frac{4\alpha}{\pi^2} \right). \quad (113)$$

Appendix D. Dense limit

We consider the limit c large. GLM-SBM is equivalent to a low-rank matrix factorization problem with a generative prior. It has been studied in Aubin *et al* (2019). We follow it closely.

We set $p_o = c_o/N$ and $p_i = c_i/N$ of order one and $\mu = \sqrt{N}(p_i - p_o)$ of order one. The effective inverse noise of the SBM is Lesieur *et al* (2017)

$$\Delta_I = \frac{\mu^2}{p_o(1 - p_o)}. \quad (114)$$

D.1. Algorithm

We reproduce here the algorithm given by Aubin *et al* (2019); we simplify it for a binary output channel and P_w Rademacher; and we complete it with the semi-supervised case. In the dense limit BP can be approximated by AMP and our algorithm AMP-BP becomes AMP-AMP.

We set, as in part 6:

$$Z_w(\gamma, \Lambda) = e^{\Lambda/2} \cosh \gamma, \quad f_w(\gamma, \Lambda) = \partial_{\gamma} \log Z_w \quad (115)$$

$$Z_{o, \mu}(B, A, \omega, V) = \begin{cases} e^{-A/2 + s_{\mu} B \frac{1}{2} (1 + s_{\mu} \operatorname{erf}(\omega / \sqrt{2V}))} & \text{if } s_{\mu} \in s_{\Xi} \\ e^{-A/2} (\cosh(B) + \sinh(B) \operatorname{erf}(\omega / \sqrt{2V})) & \text{else} \end{cases} \quad (116)$$

$$f_{o, \mu}(B, A, \omega, V) = \partial_{\omega} \log Z_{o, \mu}, \quad f_{s, \mu}(B, A, \omega, V) = \partial_B \log Z_{o, \mu}. \quad (117)$$

We define the input matrix

$$S_{\eta\nu} = \mu \frac{1}{1 - p_o} \left(\frac{Y_{\eta\nu}}{p_o} - 1 \right) \quad (118)$$

where Y is the observed adjacency matrix; $Y_{\eta\nu} = 1$ if there is an edge between η and ν , 0 otherwise.

In the following a_l and v_l are estimators of the mean and the variance of w_l ; σ_{μ} and Σ_{μ} the mean and the variance of s_{μ} .

AMP-AMP

input features $F_{\mu l}$, input matrix $S_{\eta\nu}$, effective inverse noise Δ_I , prior information $P_{s,\mu}$.

Initialize $a_l^{(0)} = \epsilon_l, v_l^{(0)} = 1, \sigma_\nu^{(0)} = \epsilon_\nu, \Sigma_\nu^{(0)} = 1, g_{\sigma,\mu}^{(0)} = 0, t = 0$; where ϵ s are zero-mean small random variables.

repeat

AMP update of ω_μ, V_μ

$$V^{(t+1)} \leftarrow \frac{1}{M} \sum_l v_l^{(t)}$$

$$\omega_\mu^{(t+1)} \leftarrow \sum_l F_{\mu l} a_l^{(t)} - V^{(t+1)} g_{\sigma,\mu}^{(t)}$$

AMP update of $g_{\sigma,\mu}, \Lambda, \Gamma_l$

$$g_{\sigma,\mu}^{(t+1)} \leftarrow f_{\sigma,\mu}(B_\mu^{(t+1)}, A^{(t+1)}, \omega_\mu^{(t+1)}, V^{(t+1)})$$

$$\Lambda^{(t+1)} \leftarrow \frac{1}{M} \sum_\mu g_{\sigma,\mu}^{2,(t+1)}$$

$$\Gamma_l^{(t+1)} \leftarrow \Lambda^{(t+1)} a_l^{(t)} + \sum_\mu F_{\mu l} g_{\sigma,\mu}^{(t+1)}$$

AMP update of the matrix factorization part

$$A^{(t+1)} \leftarrow \frac{\Delta_I}{N} \sum_\eta \sigma_\eta^{2,(t)}$$

$$B_\mu^{(t+1)} \leftarrow \frac{1}{\sqrt{N}} \sum_\eta S_{\mu\eta} \sigma_\eta^{(t)} - \frac{\Delta_I}{N} \sum_\eta \Sigma_\eta^{(t)} \sigma_\mu^{(t-1)}$$

AMP update of the estimated marginals a_l, v_l

$$a_l^{(t+1)} \leftarrow f_w(\Lambda^{(t+1)}, \Gamma_l^{(t+1)})$$

$$v_l^{(t+1)} \leftarrow \partial_{\Gamma} f_w(\Lambda^{(t+1)}, \Gamma_l^{(t+1)})$$

AMP update of the estimated marginals σ_ν, Σ_ν

$$\sigma_\mu^{(t+1)} \leftarrow f_{s,\mu}(B_\mu^{(t+1)}, A^{(t+1)}, \omega_\mu^{(t+1)}, V^{(t+1)})$$

$$\Sigma_\mu^{(t+1)} \leftarrow \partial_{B} f_{s,\mu}(B_\mu^{(t+1)}, A^{(t+1)}, \omega_\mu^{(t+1)}, V^{(t+1)})$$

$t \leftarrow t + 1$

until convergence of $a_l, v_l, \sigma_\mu, \Sigma_\mu$

output estimated mean a_l and variance v_l of w_l , estimated mean σ_μ and variance Σ_μ of s_μ

D.2. Analysis of the SE equations near the full recovery point

The state evolution equations are given in section 6, equations (35)–(37). We study the conditions of stability for the fully informative fixed point $(q_s, \hat{q}_w, q_w) = (1, +\infty, 1)$.

We use the following notation for the update:

$$(q_s, \hat{q}_w, q_w)^{t+1} = (f_1(r, s), f_2(r, s), f_3(t)), \quad (r, t, s) = (\Delta_I q_s^t, \hat{q}_w^{t+1}, q_w^t) \tag{119}$$

where the f_i are given by the SE update equations:

$$f_1(r, s) = \mathbb{E}_{\xi, \eta} \left[e^{-r/2} \frac{\left(\sinh(\sqrt{r}\xi) + \cosh(\sqrt{r}\xi) \operatorname{erf} \left(\sqrt{\frac{s}{2(1-s)}} \eta \right) \right)^2}{\cosh(\sqrt{r}\xi) + \sinh(\sqrt{r}\xi) \operatorname{erf} \left(\sqrt{\frac{s}{2(1-s)}} \eta \right)} \right] \tag{120}$$

$$f_2(r, s) = \alpha \mathbb{E}_{\xi, \eta} \left[e^{-r/2} \frac{\sinh^2(\sqrt{r}\xi) \frac{2}{\pi(1-s)} e^{-\eta^2 \frac{s}{1-s}}}{\cosh(\sqrt{r}\xi) + \sinh(\sqrt{r}\xi) \operatorname{erf} \left(\sqrt{\frac{s}{2(1-s)}} \eta \right)} \right] \tag{121}$$

$$f_3(t) = \mathbb{E}_\xi \left[e^{-t/2} \sinh(\sqrt{t}\xi) \tanh(\sqrt{t}\xi) \right] \tag{122}$$

where ξ and η are standard Gaussians.

We expand around $(1, +\infty, 1)$; we use the parametrization $(r, t, s) = (\Delta_I + \epsilon_r, 1/\epsilon_t^2, 1 - \epsilon_s^2)$.
 $-f_3$ - We expand the integrand in f_3 around $+\infty$. This is valid only for $\sqrt{t}\xi = \xi/\epsilon_t \gg 1$; so we introduce a cut-off δ such that both $\delta = o(1)$ and $\delta = \omega(\epsilon_t)$. For $\xi > \delta$ we use the asymptotic $\sinh(x) \tanh(x) = \frac{1}{2}e^x - \frac{3}{2}e^{-x} + o(e^{-x})$; for $\xi < \delta$ we develop the Gaussian density to the first (constant) order. Then

$$f_3(1/\epsilon_t^2) = e^{-1/2\epsilon_t^2} 2 \left(\int_0^\infty \frac{d\xi}{\sqrt{2\pi}} e^{-\xi^2/2} \left(\frac{1}{2}e^{\xi/\epsilon_t} - \frac{3}{2}e^{-\xi/\epsilon_t} \right) - \int_0^\delta \frac{d\xi}{\sqrt{2\pi}} \left(\frac{1}{2}e^{\xi/\epsilon_t} - \frac{3}{2}e^{-\xi/\epsilon_t} \right) \right) + e^{-1/2\epsilon_t^2} 2 \int_0^\delta \frac{d\xi}{\sqrt{2\pi}} \sinh(\xi/\epsilon_t) \tanh(\xi/\epsilon_t) \tag{123}$$

$$= \frac{1}{2} \left(1 + \operatorname{erf} \left(\frac{1}{\sqrt{2}\epsilon_t} \right) \right) - \frac{3}{2} \left(1 - \operatorname{erf} \left(\frac{1}{\sqrt{2}\epsilon_t} \right) \right) - 2e^{-1/2\epsilon_t^2} \epsilon_t \int_0^{\delta/\epsilon_t} \frac{d\xi}{\sqrt{2\pi}} \left(\frac{1}{2}e^\xi - \frac{3}{2}e^{-\xi} \right) + 2e^{-1/2\epsilon_t^2} \epsilon_t \int_0^{\delta/\epsilon_t} \frac{d\xi}{\sqrt{2\pi}} \sinh(\xi) \tanh(\xi) \tag{124}$$

$$= 1 - C_3 \epsilon_t e^{-1/2\epsilon_t^2} \tag{125}$$

$$C_3 = 2 \left(\sqrt{\frac{2}{\pi}} + \int_0^\infty \frac{d\xi}{\sqrt{2\pi}} \left(\frac{1}{2}e^\xi - \frac{3}{2}e^{-\xi} - \sinh(\xi) \tanh(\xi) \right) \right) \approx 1.3 \tag{126}$$

where in the last lines we expanded the error function around $+\infty$.

$-f_1$ - We use the shorthand notation

$$g(x, y) = \frac{(\sinh(x) + \cosh(x) \operatorname{erf}(y))^2}{\cosh(x) + \sinh(x) \operatorname{erf}(y)} \tag{127}$$

$$f_1(r, s) = \mathbb{E}_{\xi, \eta} \left[e^{-r/2} g \left(\sqrt{r}\xi, \eta \sqrt{\frac{s}{2(1-s)}} \right) \right]. \tag{128}$$

Expanding the function is difficult: we can obtain quite easily that

$$f_1(\Delta_I + \epsilon_r, 1 - \epsilon_s^2) = 1 - C_1(\Delta_I)\epsilon_s + \mathcal{O}(\epsilon_r\epsilon_s, \dots) \tag{129}$$

but the function $C_1(\Delta_I)$ is harder to obtain. Rather we compute directly the derivative to the constant order:

$$\partial_{\epsilon_s} f_1(\Delta_I + \epsilon_r, 1 - \epsilon_s^2) = -C_1(\Delta_I) \tag{130}$$

$$= \mathbb{E}_{\xi, \eta} \left[e^{-(\Delta_I + \epsilon_r)/2} \frac{-\eta}{\sqrt{2}\epsilon_s^2} \frac{2}{\sqrt{\pi}} e^{-\eta^2/2\epsilon_s^2} \partial_y g \left(\sqrt{\Delta_I + \epsilon_r}\xi, \frac{\eta}{\sqrt{2}\epsilon_s} \right) \right] \tag{131}$$

$$= -\mathbb{E}_{\xi, \eta} \left[e^{-(\Delta_I + \epsilon_r)/2} \eta \sqrt{\frac{2}{\pi}} \partial_y g \left(\sqrt{\Delta_I + \epsilon_r}\xi, \frac{\eta}{\sqrt{2}} \right) \right] + \mathcal{O}(\epsilon_s). \tag{132}$$

So $C_1(\Delta_I) = \sqrt{\frac{2}{\pi}} \mathbb{E}_{\xi, \eta} \left[e^{-\Delta_I/2} \eta \partial_y g \left(\sqrt{\Delta_I}\xi, \eta \right) \right]$.


$-f_2$ - We introduce

$$h(x, y) = \frac{\sinh^2(x)}{\cosh(x) + \sinh(x) \operatorname{erf}(y/\sqrt{2})} \tag{133}$$

$$f_2(r, s) = \alpha \mathbb{E}_{\xi, \eta} \left[e^{-r/2} \frac{2}{\pi(1-s)} e^{-\eta^2 \frac{s}{1-s}} h \left(\sqrt{r}\xi, \eta \sqrt{\frac{s}{1-s}} \right) \right]. \tag{134}$$

The first order is enough since it is not constant; we have:

$$f_1(\Delta_I + \epsilon_r, 1 - \epsilon_s^2) = \frac{\alpha}{\epsilon_s^2} \mathbb{E}_{\xi, \eta} \left[e^{-(\Delta_I + \epsilon_r)/2} \frac{2}{\pi} e^{-\eta^2 \frac{1-\epsilon_s^2}{\epsilon_s^2}} h \left(\sqrt{\Delta_I + \epsilon_r}\xi, \eta \sqrt{\frac{1-\epsilon_s^2}{\epsilon_s^2}} \right) \right] \tag{135}$$



$$= \frac{\alpha}{\epsilon_s} \mathbb{E}_{\xi, \eta} \left[e^{-\Delta_I/2} \frac{\sqrt{2}}{\pi} h \left(\sqrt{\Delta_I} \xi, \frac{\eta}{\sqrt{2}} \right) + \mathcal{O}(\epsilon_r, \epsilon_s) \right]. \tag{136}$$

So

$$f_1(\Delta_I + \epsilon_r, 1 - \epsilon_s^2) = \frac{\alpha}{\epsilon_s} C_2(\Delta_I), \quad C_2(\Delta_I) = \frac{\sqrt{2}}{\pi} \mathbb{E}_{\xi, \eta} \left[e^{-\Delta_I/2} h \left(\sqrt{\Delta_I} \xi, \frac{\eta}{\sqrt{2}} \right) \right] \tag{137}$$

$C_2(\Delta_I)$ is positive for all Δ_I .

D.2.1. Stability

We obtain the following update of the perturbation:

$$\begin{pmatrix} \epsilon_r \\ \epsilon_t \\ \epsilon_s \end{pmatrix} \rightarrow \begin{pmatrix} -\Delta_I C_1(\Delta_I) \epsilon_s \\ \sqrt{\epsilon_s} / \sqrt{\alpha C_2(\Delta_I)} \\ \sqrt{C_3} \sqrt{\epsilon_t} e^{-1/4 \epsilon_t^2} \end{pmatrix}. \tag{138}$$

We consider only the s variable because r does not affect the dynamics and the initialization is done on r and s , t being inferred then. We have

$$\epsilon_s^{t+1} = \sqrt{C_3} \sqrt[4]{\frac{\epsilon_s^t}{\alpha C_2(\Delta_I)}} e^{-\alpha C_2(\Delta_I)/4 \epsilon_s^t} \tag{139}$$

which is stable for all α and Δ_I .

Numerically, however, instability can be detected: for ϵ_s large enough the system diverges from the fully informative fixed point. We compute numerically the limiting ϵ_s^* , such that $\epsilon_s^{t+1} = \epsilon_s^t$; we find that ϵ_s^* tends to zero fast for Δ_I or μ going to zero (see figure 5).

In the Gaussian case we have $f_3(t) = t/(1+t)$ and so $\epsilon_s^{t+2} = \epsilon_t^{t+1} = \sqrt{\epsilon_s^t} / \sqrt{\alpha C_2(\Delta_I)}$; so this fixed point is unconditionally not stable.

D.2.2. Generalization of the prior

We ask for which prior P_w the fully-informative point is stable. We recall that

$$f_3(t) = \mathbb{E}_{\xi} \left[\left(\int dw w P_w(w) e^{-tw^2 + \sqrt{i} \xi w} \right)^2 \left(\int dw P_w(w) e^{-tw^2 + \sqrt{i} \xi w} \right)^{-1} \right]. \tag{140}$$

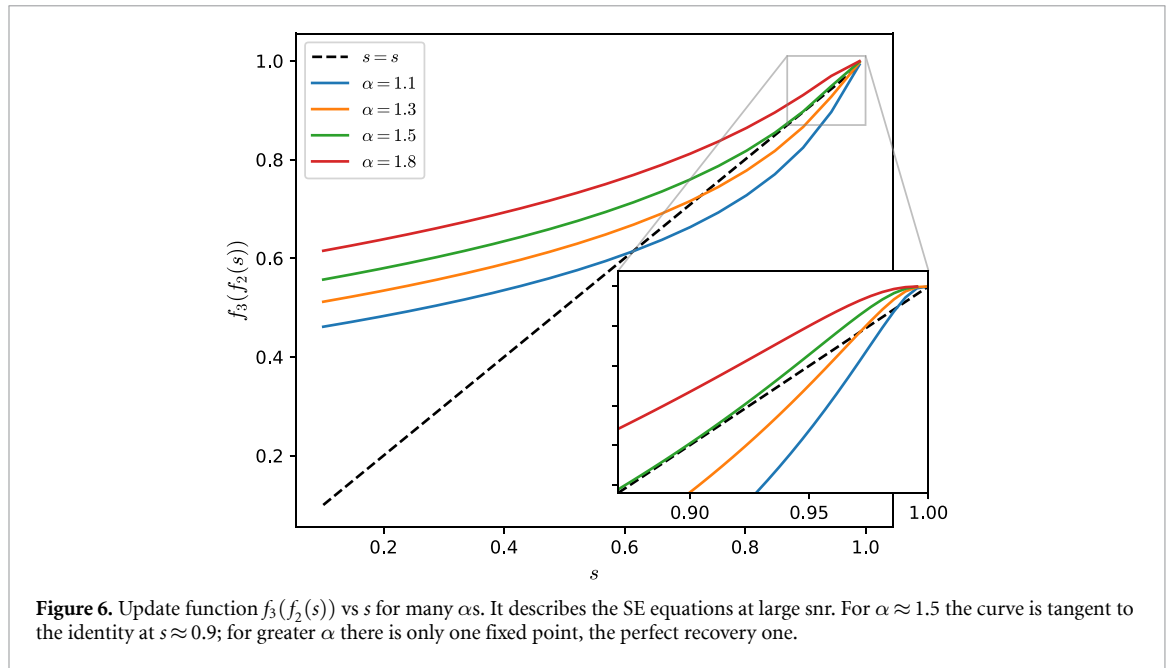


Figure 6. Update function $f_3(f_2(s))$ vs s for many α s. It describes the SE equations at large snr. For $\alpha \approx 1.5$ the curve is tangent to the identity at $s \approx 0.9$; for greater α there is only one fixed point, the perfect recovery one.

We assume that $\mathbb{E}_{P_w} w^2 = \rho_w^2 = 1$.

We show that if P_w admits everywhere a density twice differentiable, then the fully-informative fixed point is unstable. Indeed, at large t we have:

$$f_3(t) = \mathbb{E}_\xi \left[\frac{1}{\sqrt{t}} e^{\frac{1}{2}\xi^2} \left(\int dw \left(\frac{w}{\sqrt{t}} + \frac{\xi}{\sqrt{t}} \right) P_w \left(\frac{w}{\sqrt{t}} + \frac{\xi}{\sqrt{t}} \right) e^{-\frac{1}{2}w^2} \right)^2 \times \left(\int dw P_w \left(\frac{w}{\sqrt{t}} + \frac{\xi}{\sqrt{t}} \right) e^{-\frac{1}{2}w^2} \right)^{-1} \right] \tag{141}$$

$$= \frac{1}{\sqrt{t}} \int d\xi \left(\int \frac{dw}{\sqrt{2\pi}} e^{-\frac{1}{2}w^2} \left(\frac{w}{\sqrt{t}} + \frac{\xi}{\sqrt{t}} \right) \left(P_w \left(\frac{\xi}{\sqrt{t}} \right) + \frac{w}{\sqrt{t}} P_w' \left(\frac{\xi}{\sqrt{t}} \right) \right) \right)^2 \times \left(\int \frac{dw}{\sqrt{2\pi}} e^{-\frac{1}{2}w^2} \left(P_w \left(\frac{\xi}{\sqrt{t}} \right) + \frac{w^2}{2t} P_w'' \left(\frac{\xi}{\sqrt{t}} \right) \right) \right)^{-1} \tag{142}$$

$$= \int dx \left(x P_w(x) + \frac{1}{t} P_w'(x) \right)^2 \left(P_w(x) + \frac{1}{2t} P_w''(x) \right)^{-1} \tag{143}$$

$$= 1 - \mathcal{O}(1/t) \tag{144}$$

and obtain an equation similar to the one of the Gaussian case: $\epsilon_s^{t+2} = C' \sqrt{\epsilon_s^t}$, $C' > 0$, which is unstable.

D.2.3. Large snr

We give an implicit value for the critical compression ratio $\alpha_{\text{algo,d}}$. We take the limit $\Delta_I \gg 1$ and seek whether the SE updates converge to the fully informative point.


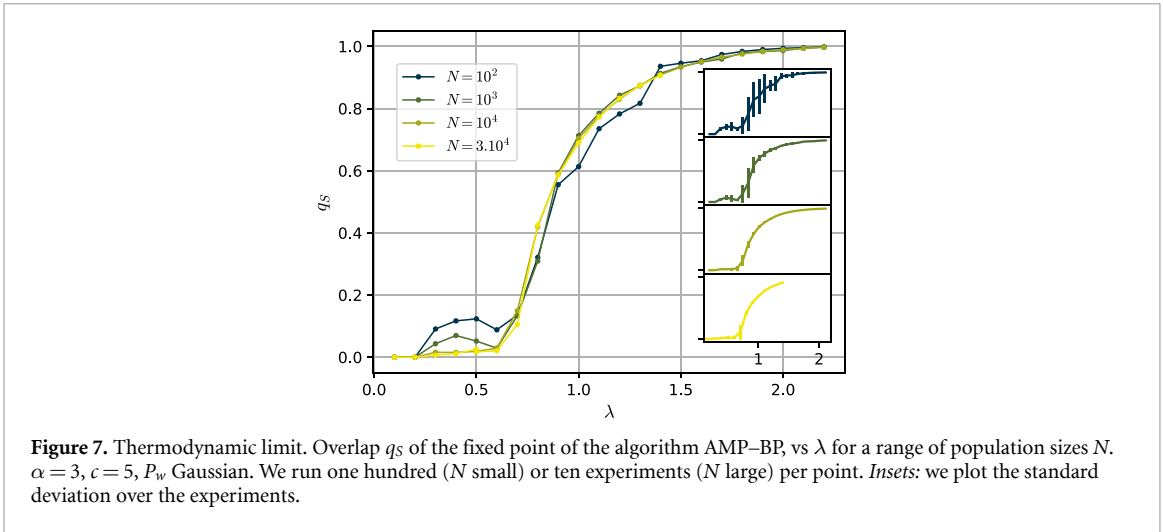
We can simplify the SE equations to one scalar equation. We expand on $r = \Delta_I q_s \gg 1$. We have $q_s = f_1(r, s) \rightarrow 1$ for all s . As to q_z and \hat{q}_z , we have

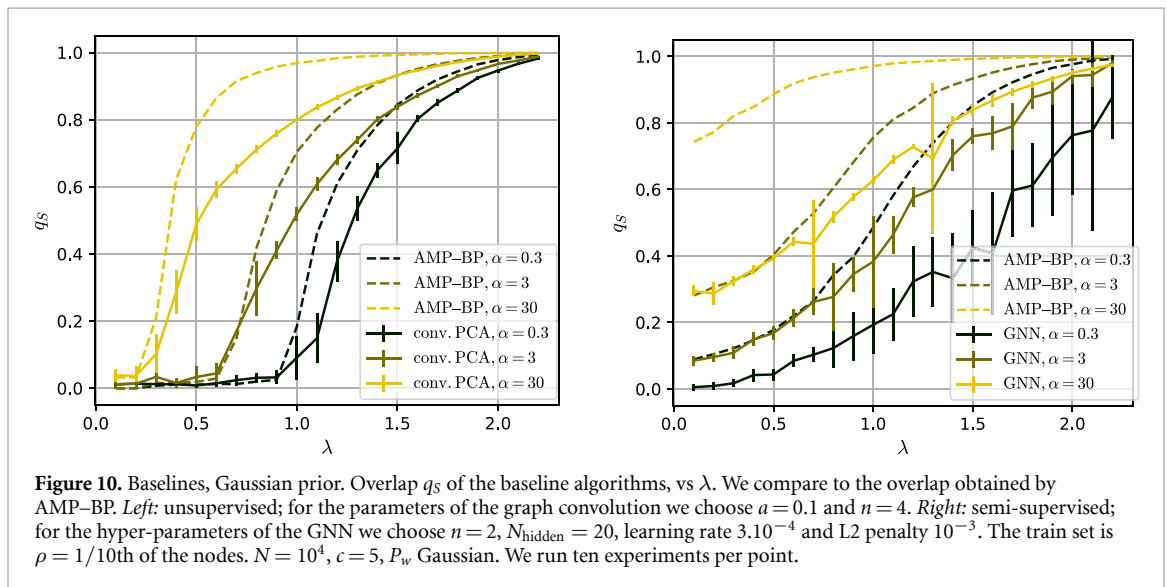
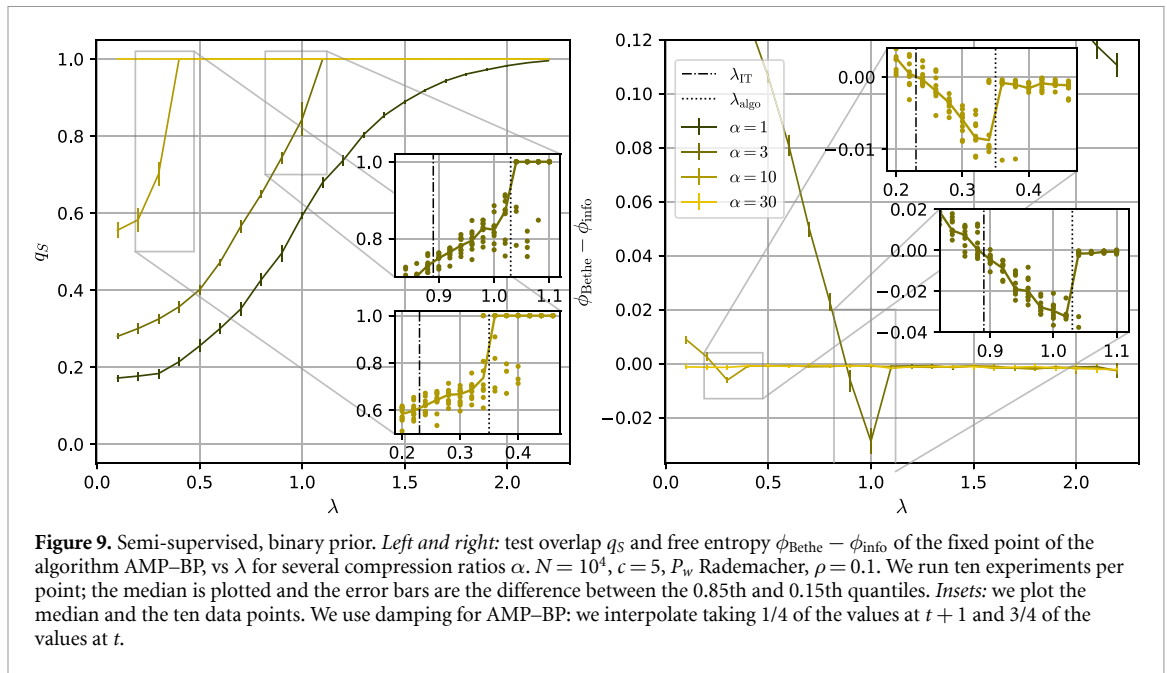
$$f_2(s) = \alpha \frac{2}{\pi} \frac{1}{\sqrt{1-s^2}} \mathbb{E}_\eta \left[1 - \text{erf} \left(\eta \sqrt{\frac{s}{2(1+s)}} \right)^2 \right]^{-1}. \tag{145}$$

We plug f_2 and f_3 together. The fixed points are the s that satisfy the equation $s = f_3(f_2(s))$. The function $f_3 \circ f_2$ is plotted in figure 6.

For $\alpha > \alpha_{\text{algo,d}} \approx 1.5$ the updates lead to perfect recovery $s = 1$, starting from any s . For $\alpha < \alpha_{\text{algo,d}}$, perfect recovery is possible only starting from s close to 1; otherwise the iterations lead to a sub-optimal fixed point.

Appendix E. Supplementary figures





ORCID iD

O Duranthon  <https://orcid.org/0009-0006-2122-6932>

References

- Abbe E 2017 Community detection and stochastic block models: recent developments *J. Mach. Learn. Res.* **18** 6446–531 (arXiv:1703.10146)
- Abbe E, Bandeira A S and Hall G 2015 Exact recovery in the stochastic block model *IEEE Trans. Inf. Theory* **62** 471–87
- Aubin B, Loureiro B, Maillard A, Krzakala F and Zdeborová L 2019 The spiked matrix model with generative priors *Advances in Neural Information Processing Systems* (arXiv:1905.12385)
- Barbier J, Krzakala F, Macris N, Miolane L and Zdeborová L 2019 Optimal errors and phase transitions in high-dimensional generalized linear models *Proc. Natl Acad. Sci.* **116** 5451–60
- Binkiewicz N, Vogelstein J T and Rohe K 2017 Covariate-assisted spectral clustering *Biometrika* **104** 361–77
- Celentano M, Cheng C and Montanari A 2021 The high-dimensional asymptotics of first order methods with random data (arXiv:2112.07572)
- Cheng S, Pan L, Hu H and Dokmanić I 2022 Statistical mechanics of generalization in graph convolution networks (arXiv:2212.13069)
- Chien E, Peng J, Li P and Milenkovic O 2021 Adaptive universal generalized pagerank graph neural network *Int. Conf. on Learning Representations* (arXiv:2006.07988)
- Cho S, Min S, Kim J, Lee M, Lee H and Hong S 2022 Transformers meet stochastic block models: attention with data-adaptive sparsity and cost *Advances in Neural Information Processing Systems* ed A H Oh, A Agarwal, D Belgrave and K Cho (arXiv:2210.15541)

- Coja-Oghlan A, Krzakala F, Perkins W and Zdeborová L 2017 Information-theoretic thresholds from the cavity method *Proc. 49th Annual ACM SIGACT Symp. on Theory of Computing* pp 146–57 (arXiv:1611.00814)
- Cornacchia E, Mignacco F, Veiga R, Gerbelot C, Loureiro B and Zdeborová L 2022 Learning curves for the multi-class teacher-student perceptron *Mach. Learn.: Sci. Technol.* **4** 015019
- Decelle A, Krzakala F, Moore C and Zdeborová L 2011a Inference and phase transitions in the detection of modules in sparse networks *Phys. Rev. Lett.* **107** 065701
- Decelle A, Krzakala F, Moore C and Zdeborová L 2011b Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications *Phys. Rev. E* **84** 066106
- Deshpande Y, Sen S, Montanari A and Mossel E 2018 Contextual stochastic block models *Advances in Neural Information Processing Systems* vol 31, ed S Bengio, H Wallach, H Larochelle, K Grauman, N Cesa-Bianchi and R Garnett (arXiv:1807.09596)
- Donoho D L, Maleki A and Montanari A 2009 Message-passing algorithms for compressed sensing *Proc. Natl Acad. Sci.* **106** 18914–9
- Fortunato S 2010 Community detection in graphs *Phys. Rep.* **486** 75–174
- Fountoulakis K, He D, Lattanzi S, Perozzi B, Tsitsulin A and Yang S 2022 On classification thresholds for graph attention with edge features (arXiv:2210.10014)
- Gabriel M, Manoel A, Luneau C, Macris N, Krzakala F, Zdeborová L *et al* 2018 Entropy and mutual information in models of deep neural networks *Advances in Neural Information Processing Systems* vol 31 (arXiv:1805.09785)
- Gamarnik D, Moore C and Zdeborová L 2022 Disordered systems insights on computational hardness *J. Stat. Mech.* **114015**
- Gerbelot C and Berthier R 2021 Graph-based approximate message passing iterations (arXiv:2109.11905)
- Györfyi G 1990 First-order transition to perfect generalization in a neural network with binary synapses *Phys. Rev. A* **41** 7097
- Krzakala F, Manoel A, Tramel E W and Zdeborová L 2014 Variational free energies for compressed sensing *2014 IEEE Int. Symp. on Information Theory* pp 1499–503 (arXiv:1402.1384)
- Krzakala F, Mézard M, Sausset F, Sun Y and Zdeborová L 2012 Probabilistic reconstruction in compressed sensing: algorithms, phase diagrams and threshold achieving matrices *J. Stat. Mech.* **P08009**
- Krzakala F, Moore C, Mossel E, Neeman J, Sly A, Zdeborová L and Zhang P 2013 Spectral redemption in clustering sparse networks *Proc. Natl Acad. Sci.* **110** 20935–40
- Lesieur T, Krzakala F and Zdeborová L 2017 Constrained low-rank matrix estimation: phase transitions, approximate message passing and applications *J. Stat. Mech.* **073403**
- Loureiro B, Gerbelot C, Cui H, Goldt S, Krzakala F, Mezard M and Zdeborová L 2021 Learning curves of generic features maps for realistic datasets with a teacher-student model *Advances in Neural Information Processing Systems* vol 34 pp 18137–51 (arXiv:2102.08127)
- Lu C and Sen S 2020 Contextual stochastic block model: sharp thresholds and contiguity (arXiv:2011.09841)
- Manoel A, Krzakala F, Mézard M and Zdeborová L 2017 Multi-layer generalized linear estimation *2017 IEEE Int. Symp. on Information Theory (ISIT)* (IEEE) pp 2098–102 (arXiv:1701.06981)
- Miolane L 2017 Fundamental limits of low-rank matrix estimation: the non-symmetric case (arXiv:1702.00473)
- Mossel E, Neeman J and Sly A 2015 Reconstruction and estimation in the planted partition model *Probab. Theory Relat. Fields* **162** 431–61
- Mossel E, Neeman J and Sly A 2018 A proof of the block model threshold conjecture *Combinatorica* **38** 665–708
- Ongie G, Jalal A, Metzler C A, Baraniuk R G, Dimakis A G and Willett R 2020 Deep learning techniques for inverse problems in imaging *IEEE J. Sel. Areas Inf. Theory* **1** 39–56
- Peixoto T P 2019 Bayesian stochastic blockmodeling *Advances in Network Clustering and Blockmodeling* (Wiley) pp 289–332 (arXiv:1705.10225)
- Shlezinger N, Whang J, Eldar Y C and Dimakis A G 2020 Model-based deep learning (arXiv:2012.08405)
- Sompolinsky H, Tishby N and Sebastian Seung H 1990 Learning from examples in large neural networks *Phys. Rev. Lett.* **65** 1683
- Tsitsulin A, Rozemberczki B, Palowitch J and Perozzi B 2021 Synthetic graph generation to benchmark graph learning *Workshop on Graph Learning Benchmarks* (arXiv:2204.01378)
- Yang J, McAuley J and Leskovec J 2013 Community detection in networks with node attributes *2013 IEEE 13th Int. Conf. on Data Mining* (IEEE) pp 1151–6 (arXiv:1401.7267)
- Zdeborová L and Krzakala F 2016 Statistical physics of inference: thresholds and algorithms *Adv. Phys.* **65** 453–552
- Zhang P, Moore C and Zdeborová L 2014 Phase transitions in semisupervised clustering of sparse networks *Phys. Rev. E* **90** 052802