

Semi-Supervised Clustering of Sparse Graphs: Crossing the Information-Theoretic Threshold

Junda Sheng^{* 1} and Thomas Strohmer^{† 2}

¹Department of Mathematics, University of California, Davis, CA 95616-5270, USA.

²Department of Mathematics and Center of Data Science and Artificial Intelligence Research, University of California, Davis, CA 95616-5270, USA.

Abstract. The stochastic block model is a canonical random graph model for clustering and community detection on network-structured data. Decades of extensive study on the problem have established many profound results, among which the phase transition at the Kesten-Stigum threshold is particularly interesting both from a mathematical and an applied standpoint. It states that no estimator based on the network topology can perform substantially better than chance on sparse graphs if the model parameter is below a certain threshold. Nevertheless, if we slightly extend the horizon to the ubiquitous semi-supervised setting, such a fundamental limitation will disappear completely. We prove that with an arbitrary fraction of the labels revealed, the detection problem is feasible throughout the parameter domain. Moreover, we introduce two efficient algorithms, one combinatorial and one based on optimization, to integrate label information with graph structures. Our work brings a new perspective to the stochastic model of networks and semidefinite program research.

Keywords:

Clustering,
Semi-supervised learning,
Stochastic block model,
Kesten-Stigum threshold,
Semidefinite programming.

Article Info.:

Volume: 3
Number: 1
Pages: 64 - 106
Date: March/2024
doi.org/10.4208/jml.230624

Article History:

Received: 24/06/2023
Accepted: 29/02/2024

Communicated by:

Qianxiao Li

1 Introduction

Clustering has long been an essential subject of many research fields, such as machine learning, pattern recognition, data science, and artificial intelligence. In this section, we include some background information on its general setting and the semi-supervised approach.

1.1 Clustering on graphs

The basic task of clustering or community detection in its general form is, given a (possibly weighted) graph, to partition its vertices into several densely connected groups with relatively weak external connectivity. This property is sometimes also called assortativity. Clustering and community detection are central problems in machine learning and data science with various applications in scientific research and industrial development. A considerable amount of data sets can be represented in the form of a network that consists of

^{*}Corresponding author. sheng@math.ucdavis.edu

[†]strohmer@math.ucdavis.edu

interacting nodes, and one of the first features of interest in such a situation is to understand which nodes are similar, as an end or as a preliminary step towards other learning tasks. Clustering is used to find genetically similar sub-populations [60], to segment images [65], to study sociological behavior [59], to improve recommendation systems [47], to help with natural language processing [30], etc. Since the 1970s, in different communities like social science, statistical physics, and machine learning, a large diversity of algorithms have been developed such as:

- Hierarchical clustering algorithms [38] build a hierarchy of progressive communities, by either recursive aggregation or division.
- Model-based statistical methods, including the celebrated expectation-maximization (EM) clustering algorithm proposed in [23], fit the data with cluster-exhibiting statistical models.
- Optimization approaches identify the best cluster structures regarding carefully designed cost functions, for instance, minimizing the cut [34] and maximizing the Girvan-Newman modularity [58].

Multiple lines of research intersect at a simple random graph model, which appears under many different names. In the machine learning and statistics literature around social networks, it is called stochastic block model (SBM) [36], while it is known as the planted partition model [17] in theoretical computer science and referred to as inhomogeneous random graph model [14] in the mathematics literature. Moreover, it can also be interpreted as a spin-glass model [22], a sparse-graph code [5], a low-rank random matrix model [50], and more.

The essence of SBM can be summarized as follows: Conditioned on the vertex labels, edges are generated independently and the probability only depends on which clusters the pairs of vertices belong to. We consider its simplest form, namely the symmetric SBM consisting of two blocks, also known as the planted bisection model.

Definition 1.1 (Planted Bisection Model). *For $n \in \mathbb{N}$ and $p, q \in (0, 1)$, let $\mathcal{G}(n, p, q)$ denote the distribution over graphs with n vertices defined as follows. The vertex set is partitioned uniformly at random into two subsets S_1, S_2 with $|S_i| = n/2$. Let E denote the edge set. Conditional on this partition, edges are included independently with probability*

$$P((i, j) \in E \mid S_1, S_2) = \begin{cases} p, & \text{if } \{i, j\} \subseteq S_1 \text{ or } \{i, j\} \subseteq S_2, \\ q, & \text{if } i \in S_1, j \in S_2 \text{ or } i \in S_2, j \in S_1. \end{cases} \quad (1.1)$$

Note that if $p = q$, the planted bisection model is reduced to the so-called Erdős-Rényi random graph where all edges are generated independently with the same probability. Hence, there exists no cluster structure. But if $p \gg q$, a typical graph will have two well-defined clusters. The scale of p and q also plays a significant role in the resulting graph, which will be discussed in detail later. They govern the amount of signal and noise in the graph's generating process. As the key parameters that researchers work with, they depict various regimes and thresholds.

The SBMs generate labels for vertices before the graph. The ground truth allows us to formally discuss the presence of community structures and measure the performance of algorithms in a meaningful way. It also supplies a natural basis to rigorously define the semi-supervised clustering problem. But as a parametrized statistical model, one can only hope that it serves as a good fit for the real data. Although not necessarily a realistic model, SBM provides us with an insightful abstraction and captures some of the key phenomena [2, 6, 11, 19, 54].

Given a single realization of a graph G , our goal is to recover the labels x , up to a certain level of accuracy. Formally, the ground truth of the underlying community structure is encoded using the vector $x \in \{+1, -1\}^n$ with $x_i = +1$, if $i \in S_1$ and $x_i = -1$, if $i \in S_2$. An estimator is a map $\hat{x} : G_n \rightarrow \{+1, -1\}^n$, where G_n is the space of graphs over n vertices. We define the Overlap between an estimator and the ground truth as

$$\text{Overlap}(x, \hat{x}(G)) = \frac{1}{n} |\langle x, \hat{x}(G) \rangle|. \quad (1.2)$$

Overlap induces a measure on the same probability space as the model, which represents how well an (unsupervised) estimator performs on the recovery task. To intuitively interpret the result, we put requirements on its asymptotic behavior, which takes place with high probability as $n \rightarrow \infty$.

Definition 1.2. *Let $G \sim \mathcal{G}(n, p, q)$. The following recovery requirements are solved if there exists an algorithm that takes G as an input and outputs $\hat{x} = \hat{x}(G)$ such that:*

- *Exact recovery:* $\mathbb{P}\{\text{Overlap}(x, \hat{x}(G)) = 1\} = 1 - o(1)$,
- *Weak recovery:* $\mathbb{P}\{\text{Overlap}(x, \hat{x}(G)) \geq \Omega(1)\} = 1 - o(1)$.

In other words, exact recovery requires the entire partition to be correctly identified. Weak recovery only asks for substantially better performance than chance. In some literature, exact recovery is simply called recovery. Weak recovery is also called detection since as long as one can weakly recover the ground truth, there must exist a community structure.

Note that if G is an Erdős-Rényi random graph ($p = q$) then the overlap will be $o_p(1)$ for all estimators. This can be seen by noticing that x and G are independent in this setting and then applying Markov's inequality. This has led to two additional natural questions about SBMs. On the one hand, we are interested in the distinguishability (or testing): Is there a hypothesis test to distinguish a random graph generated by the Erdős-Rényi model (ERM) from a random graph generated by the SBM, which succeeds with high probability? On the other hand, we can ask about the model learnability (or parameter estimation): Assuming that G is drawn from an SBM ensemble, is it possible to obtain a consistent estimator for the parameters (p, q) ? Although each of these questions is of independent interest, for symmetric SBMs with two symmetric communities (planted bisection model) the following holds [1]:

$$\text{learnability} \iff \text{weak recovery} \iff \text{distinguishability}. \quad (1.3)$$

Such equivalence benefits our understanding of the model in turn. For example, direct analysis of weak recovery leads to the converse of phase transition theory [54]. The achievability of the phase transition threshold [49] is proved by counting non-backtracking walks on the graph which gives consistent estimators of parameters. In the recent work [52], hypothesis testing formulation is studied.

SBMs demonstrate the fundamental limits of clustering and community detection as some necessary and sufficient conditions for the feasibility of recovery, information-theoretically or computationally. Moreover, they are usually expressed in the form of phase transition. Sharp transitions exist in the parameter regimes between phases where the task is resolvable or not. For example, when the average degree grows as $\log n$, if the structure is sufficiently obvious then the underlying communities can be exactly recovered [12], and the threshold at which this becomes possible has also been determined [2]. Above this threshold, efficient algorithms exist [5, 7, 24, 61] that recover the communities exactly, labeling every vertex correctly with high probability, below this threshold, exact recovery is information-theoretically impossible. [28, 74] studied a slightly weaker requirement that allows a vanishing proportion of misclassified nodes and proposed algorithms that achieve the optimal rates for various models. In contrast, we study a rather different regime where the reasonable question to ask is whether one can guarantee a misclassification proportion strictly less than $1/2$.

We acknowledge that SBMs in general are theoretical constructs that can not capture all the complicated situations emerging with real-life networks. However, it succeeds in modeling the simplest situation where a community structure can be quantitatively described. In practice, instead of applying the model directly, it is advisable to initially conduct some heuristic and hierarchical analyses on the graph and see if the problem can be effectively reduced to those simple cases. For the graphs violating the model assumptions, e.g. a sparse graph dominated by cliques, alternative and tailored modeling may be necessary.

1.2 Sparse regime and Kesten-Stigum threshold

In the sparse case where the average degree of the graph is $\mathcal{O}(1)$, it is more difficult to find the clusters and the best we can hope for is to label the vertices with nonzero correlation or mutual information with the ground truth, i.e. weak recovery. Intuitively, we only have access to a constant amount of connections about each vertex. The intrinsic difficulty can be understood from the topological properties of the graphs in this regime. The following basic results are derived from [25]:

- For $a, b > 0$, the planted bisection model $\mathcal{G}(n, (a \log n)/n, (b \log n)/n)$ is connected with high probability if and only if $(a + b)/2 > 1$.
- $\mathcal{G}(n, a/n, b/n)$ has a giant component (i.e. a component of size linear in n) with high probability if and only if $d := (a + b)/2 > 1$.

The graph will only have vanishing components if the average degree is too small. Therefore, it is not possible to even weakly recover the labels. But we will see in the next section

that semi-supervised approaches amazingly piece the components together with consistent labeling.

Although it is mathematically challenging to work in the sparse regime, real-world data are likely to have bounded average degrees. [46] and [67] studied a large collection of benchmark data sets, including power transmission networks, website link networks, and complex biological systems, which had millions of nodes with an average degree of no more than 20. For instance, the LinkedIn network they studied had approximately seven million nodes, but only 30 million edges.

The phase transition for weak recovery or detection in the sparse regime was first conjectured in the paper by Decelle *et al.* [22], which sparked the modern study of clustering and SBMs. Their work is based on deep but non-rigorous insights from statistical physics, derived with the cavity method (also known as belief propagation). Since then, extensive excellent research has been conducted to understand this fundamental limit, e.g. [3, 49, 54, 55]. A key result is the following theorem.

Theorem 1.1 (Kesten-Stigum Threshold). *Let $\mathcal{G}(n, a/n, b/n)$ be a symmetric SBM with two balanced clusters and $a, b = \mathcal{O}(1)$. The weak recovery problem is solvable and efficiently so, if and only if $(a - b)^2 > 2(a + b)$.*

In particular, if we denote the probability measures induced by the ERM $\mathcal{G}(n, (a+b)/(2n), (a+b)/(2n))$ and the SBM $\mathcal{G}(n, a/n, b/n)$ by $P_n^{(0)}$ and $P_n^{(1)}$ correspondingly, they are mutually contiguous, that is for any sequence of events $\{E_n\}$'s, $P_n^{(0)}(E_n) \rightarrow 0$ if, and only if, $P_n^{(1)}(E_n) \rightarrow 0$.

Conventionally, signal-to-noise ratio (SNR) is defined as the following:

$$\text{SNR} := \frac{(a - b)^2}{2(a + b)}. \tag{1.4}$$

It is worth noting that we only quoted the KS threshold for the two communities case ($k=2$). For sufficiently large k , namely $k \geq 5$, there is a hard but detectable area where the weak recovery is information-theoretically possible, but computationally hard [6, 11]. This gap between the KS threshold and the information-theoretic (IT) threshold only shows up in the constant degree regime, making it a fertile ground for studying the fundamental tradeoffs in community detection. We focus on the cardinal case, symmetric SBM with two balanced clusters, where two thresholds coincide and a semi-supervised approach crosses them together. Fig. 1.1 provides an example of the SBM.

The terminology ‘‘KS threshold’’ can be traced back to the work of Kesten and Stigum [42] concerning the reconstruction of infinite rooted trees in 1966. The problem consists of broadcasting the root label of a tree with a fixed degree c down to its leaves and trying to recover it from the leaves at a large depth. We start with drawing the root label uniformly in $\{0, 1\}$. Then, in a top-down manner, we independently label every child the same as its parent with probability $1 - \epsilon$ and the opposite as its parent otherwise. Let $x^{(t)}$ denote the labels at depth t in this tree with $t = 0$ being the root. We say the reconstruction is solvable if

$$\lim_{t \rightarrow \infty} \mathbb{E} \left| \mathbb{E} (x^{(0)} | x^{(t)}) - \frac{1}{2} \right| > 0$$

or, equivalently,

$$\lim_{t \rightarrow \infty} I(x^{(0)}; x^{(t)}) > 0,$$

where I is the mutual information. Although it was shown in the original paper, reconstruction is solvable when $c(1 - 2\epsilon)^2 > 1$, the non-reconstruction was proved 30 years later, namely, it is not solvable if $c(1 - 2\epsilon)^2 \leq 1$ [13, 26]. Based on that finding, Mossel *et al.* [54] proved the converse part of Theorem 1.1 by coupling the local neighborhood of an SBM vertex with a Galton-Watson tree with a Markov process. Inspired by this elegant approach, we propose our census method to solve the semi-supervised clustering problem, and we will see in Section 3 how it works by amplifying revealed information with tree-like neighborhoods.




Figure 1.1: The left image represents the adjacency matrix of one realization of $\mathcal{G}(100, 0.12, 0.05)$, where the detection is theoretically possible. Yet the data is given non-colored (middle) and also non-ordered (right).

1.3 Basic algorithms

Information-theoretic bounds can provide the impossibility side of phase transitions, but we still need specific efficient algorithms for the achievability side. One straightforward approach is the spectral method. Under Definition 1.1, let A be the adjacency matrix of the graph $G \sim \mathcal{G}(n, a/n, b/n)$, $a > b$. Up to reordering indices, its expectation is a block matrix except for the diagonal,

$$\mathbb{E} A \approx \frac{1}{n} \begin{pmatrix} a & b \\ b & a \end{pmatrix} \otimes I_{\frac{n}{2} \times \frac{n}{2}}, \tag{1.5}$$

which has three eigenvalues, $(a + b)/n > (a - b)/n > 0$. 0 has multiplicity $n - 2$ and the eigenvector associated with the second largest eigenvalue is $\begin{pmatrix} \mathbf{1}_{n/2} \\ -\mathbf{1}_{n/2} \end{pmatrix}$ which is consistent with the ground truth of the labels. However, we do not observe the expected adjacency matrix. Instead, we only have access to one realization of the model. In modern terms, community detection is a one-shot learning task. But one can still hope that $A - \mathbb{E} A$ is small and the second eigenvector of A gives a reasonable estimator. For example, denoting the ordered eigenvalues of $\mathbb{E} A$ and A as $\{\lambda_i\}$'s and $\{\hat{\lambda}_i\}$'s respectively, the Courant-Fischer-Weyl min-max principle implies

$$|\hat{\lambda}_i - \lambda_i| \leq \|A - \mathbb{E} A\|_{\text{op}}, \quad i = 1, \dots, n. \tag{1.6}$$

Recall that the operator norm of a symmetric matrix M is

$$\|M\|_{\text{op}} = \max(\xi_1(M), -\xi_n(M)),$$

where $\xi_i(M)$ denotes the i -th largest eigenvalue of M . If one can bound $\|A - \mathbb{E} A\|_{\text{op}}$ by half of the least gap between the three eigenvalues mentioned above, the order will be preserved. Then the Davis-Kahan theorem guarantees the eigenvectors are correlated. Namely, if θ denotes the angle between the second eigenvectors (spectral estimator and ground truth), we have

$$\sin \theta \leq \frac{\|A - \mathbb{E} A\|_{\text{op}}}{\min\{\|\lambda_i - \lambda_2\|/2 : i \neq 2\}}. \tag{1.7}$$

Thus, the key is to control the norm of the perturbation. Many deep results from random matrix theory come into play here [4, 57, 70].

This nice and simple approach stops working as we step into the sparse regime [21, 22, 27, 41, 44]. The main reason is that leading eigenvalues of A are about the order of the square root of the maximum degree. High-degree vertices mess up the desired order of eigenvalues. In particular, for Erdős-Rényi random graphs ($\mathcal{G}(n, d/n)$), we have

$$\hat{\lambda}_1 = (1 + o(1)) \sqrt{\frac{\log n}{\log \log n}}$$

almost surely [43]. Furthermore, the leading eigenvectors are concentrated at these outliers with high degree and contain no structural information of the underlying model.

Take the star graph for example, where we assume that only the first node is connected to k neighbors. It is easy to see that the corresponding adjacency matrix has eigenvalue \sqrt{k} and eigenvector $(\sqrt{k}, 1, \dots, 1)$. Various interesting spectrum-based methods are proposed to overcome this challenge [15, 49, 55]. The key idea is to replace adjacency with some combinatorically constructed matrices. However, they typically rely on model statistics and underlying probabilistic assumptions, which leads to the problem of adversarial robustness. For example, they are non-robust to helpful perturbations. Namely, if we allow an adversary to perform the following changes on the graph: (1) adding edges within communities and/or (2) removing edges across communities, spectral approaches are going to fail. It is surprising since, intuitively, these changes help to emphasize community structures.

Meanwhile, semidefinite programming (SDP) sheds light on how we may be able to overcome the limitations of spectral algorithms, which are shown to be robust when SNR is sufficiently large [51]. It is another major line of work on clustering and community detection concerning the performance of SDPs on SBMs. While a clear picture of the unbounded degree case has been figured out by [2, 7, 9, 10, 33, 61], the results for sparse networks are more complicated. [32] proved a sub-optimal condition, $\text{SNR} \geq 10^4$, using Grothendieck inequality. Then, with a Lindeberg interpolation process [68], Montanari *et al.* proved that an SDP algorithm as proposed by [52] is nearly optimal for the case of large bounded average degree by transferring analysis of the original SDPs to the analysis of SDPs of Gaussian random matrices.

Theorem 1.2 ([52]). *Assume $G \sim G(n, a/n, b/n)$. If for some $\epsilon > 0$, $\text{SNR} \geq 1 + \epsilon$ and $d > d^*(\epsilon)$ then the SDP estimator solves the weak recovery.*

The necessary degree d^* depends on the choice of ϵ and goes to infinity as $\epsilon \rightarrow 0$. If we fix d and view ϵ as its function, the condition becomes $\text{SNR} \geq 1 + o_d(1)$. Numerical estimation and non-rigorous statistical mechanism approximation suggest that it is at most 2% sub-optimal. This result seems to be the ceiling of SDP according to the preliminary calculation from [37]. Moreover, they address the irregularity of high-degree nodes by showing SDPs return similar results for Erdős-Rényi random graphs and random regular graphs, which appear to be sensitive only to the average degree. See Section 4 for more discussion on the estimation. Inspired by their work, we propose a natural modification of SDP to incorporate revealed labels in the semi-supervised setting and show that it not only achieves but even crosses, the KS threshold. In turn, our result brings a new perspective to study the (non-)achievability and robustness of (unsupervised) SDPs.

1.4 Semi-supervised learning

Within machine learning, there are three basic approaches: supervised learning, unsupervised learning, and the combination of both, semi-supervised learning. The main difference lies in the availability of labeled data. While unsupervised learning (e.g. clustering, association, and dimension reduction) operates without any domain-specific guidance or preexisting knowledge, supervised learning (e.g. classification and regression) relies on all training samples being associated with labels. However, it is often the case where existing knowledge for a problem domain does not fit either of these extremes.

In real-world applications, unlabeled data comes with a much lower cost not requiring expensive human annotation and laboratory experiments. For example, documents crawled from the Web, images obtained from surveillance cameras, and speech collected from the broadcast are relatively more accessible compared to their labels which are required for prediction tasks, such as sentiment orientation, intrusion detection, and phonetic transcript. Motivated by this labeling bottleneck, the semi-supervised approach utilizes both labeled and unlabeled data to perform learning tasks faster, better, and cheaper. Since the 1990s, semi-supervised learning research has enjoyed an explosion of interest with applications like natural language processing [19, 63] and computer vision [45, 72].

This paper is closely related to the subtopic called constrained clustering, where one has some must-links (i.e. two nodes belong to the same cluster) and cannot-links (i.e. two nodes are in different clusters) as extra information. Although constrained versions of classic algorithms have been studied empirically, such as expectation-maximization [64], k-means [71], and spectral method [39], our methods take different approaches than hard-coding these pairwise constraints into the algorithms and provide theoretically insightful guarantees under SBM.

There also has been some excellent work considering introducing node information into SBMs. One interesting direction is overlaying a Gaussian mixture model with SBM, namely at each node of the graph assuming there is a vector of Gaussian covariates, which are correlated with the community structure. [73] proposed an SDP with k-means regularization and showed that such node information indeed improves the clustering accuracy,

while [48] formally established the information-theoretic threshold for this model when the average degree exceeds one. The node information is depicted as noisy observations for all nodes in [56]. In this setting, random guessing is no longer a meaningful baseline. Hence, the authors refer to the maximum a posteriori (MAP) estimator instead and show a local belief propagation algorithm's performance converges to the MAP accuracy in various regimes of the model. They also conjectured that those regimes can be extended to the entire domain of $a > b$ with arbitrary, but non-vanishing, strength of node information. We will see in the next section that our result establishes this conjecture in the sense that with arbitrary, but non-vanishing, knowledge of the labels, we can beat the meaningful baseline for all $a > b$.

All of these models require input on every node, which does not fall within the scope of semi-supervised learning. Whereas, we consider a realistic and intuitive generalization of SBM where a small random fraction of the labels is given accurately. [40] studied the same model as ours. They demonstrated that a vanishing fraction of labels improves the power of local algorithms. In particular, when the number of clusters diverges, it helps their local algorithm to go below the conjectured algorithmic threshold of the unsupervised case. Elegantly, they also proved the following result, which is closely related to this paper.

Theorem 1.3 ([40]). *When the fraction of revealed node labels is vanishingly small, the (unsupervised) weak recovery problem on the planted bisection model is not solvable if SNR is under the Kesten-Stigum threshold.*

2 Our results

The main goal of this paper is to answer the long-standing open question regarding semi-supervised learning on probabilistic graph models. We would like to quote the version from [1]:

“How do the fundamental limits change in a semi-supervised setting, i.e. when some of the vertex labels are revealed, exactly or probabilistically?”

In the previous section, we discussed deep research related to the clustering/community detection problem on the SBM. Establishing the phase transition phenomenon at the KS threshold is a major focal point. However, such a sharp and intrinsic limit completely disappears when an arbitrarily small fraction of the labels is revealed. This astonishing change is first observed in [75] where the authors provide non-trivial conjectures based on calculations of the belief propagation approach.

The theory of semi-supervised clustering contains some fascinating and fundamental algorithmic challenges arising from both the sparse random graph model itself and the semi-supervised learning perspective. To address them rigorously, we first define the semi-supervised SBM so that it captures the essence of realistic semi-supervised learning scenarios and is a natural and simple generalization of unsupervised models.

Definition 2.1 (Semi-Supervised Planted Bisection Model). *For $n \in \mathbb{N}$, $p, q \in (0, 1)$ and $\rho \geq 0$, let $\mathcal{G}(n, p, q, \rho)$ denote the distribution over graphs with n vertices and n -dimensional*

vectors defined as follows. The vertex set is partitioned uniformly at random into two subsets S_1, S_2 under the balance constraint $|S_1| = |S_2| = n/2$. Then, conditioned on the partition, two processes are undertaken independently:

- Let E denote the edge set of the graph G . Edges are included independently with probability defined as follows:

$$P((i, j) \in E | S_1, S_2) = \begin{cases} p, & \text{if } \{i, j\} \subseteq S_1 \text{ or } \{i, j\} \subseteq S_2, \\ q, & \text{if } i \in S_1, j \in S_2 \text{ or } i \in S_2, j \in S_1. \end{cases} \quad (2.1)$$

- An index set \mathcal{R} of size $m := 2\lfloor \rho \cdot n/2 \rfloor$ is chosen uniformly at random such that $|\mathcal{R} \cap S_1| = |\mathcal{R} \cap S_2| = m/2$. The revealed labels are given as

$$\tilde{x}_i = \begin{cases} 1, & i \in \mathcal{R} \cap S_1, \\ -1, & i \in \mathcal{R} \cap S_2, \\ 0, & \text{otherwise.} \end{cases} \quad (2.2)$$

Remark 2.1. The revealing process is independent of edge formation, i.e. $G \perp \tilde{x} | S_1, S_2$. Moreover, if we set $\rho = 0$ or simply ignore the revealed labels, the model is exactly the unsupervised SBM. In other words, the marginal distribution of the random graph is indeed $\mathcal{G}(n, p, q)$ from Definition 1.1.

Remark 2.2. One can also consider revealing uniformly at random over the index set independent of $\mathcal{G}(n, p, q)$ (instead of requiring revealed communities to have the same size), but this modification makes almost no difference in the context of this work. In practice, one can always achieve the balance requirement by either sampling a few more or dropping the uneven part.

Remark 2.3. The definition is versatile in the sense that it keeps the unsupervised setting as a special case (and with it all the interesting phase transitions). On the other hand, it can be easily generalized to the multiple and/or asymmetric communities case.

Under the semi-supervised setting, we naturally extend community detection problems in a non-trivial way that includes the unsupervised situation as a special case and captures its essence. We will discuss these items in detail when it comes to the corresponding section.

Definition 2.2. *Semi-supervised weak recovery: Finding an estimator to perform substantially better than chance on the unrevealed vertices. Semi-supervised distinguishability: Finding a test that, with high probability, distinguishes $\mathcal{G}(n, d/n, d/n, \rho)$ from $\mathcal{G}(n, a/n, b/n, \rho)$, where $d = (a+b)/2, a > b$.*

Based on the fact that a $\ln(n)$ -neighborhood in $(G, x) \sim \mathcal{G}(n, a/n, b/n)$ asymptotically has the identical distribution as a Galton-Watson tree with Markov process, we propose our first semi-supervised clustering algorithm, called census method. Namely, we decide the label estimation of a certain node according to the majority of its revealed neighbors,

$$\hat{x}_v = \text{sgn} \left(\sum_{i \in \{u \in \mathcal{R}: d(u,v)=t\}} x_i \right), \quad (2.3)$$

where $d(u, v)$ is the length of the shortest path connecting u and v . We conclude that when $\text{SNR} \leq 1$, the optimal choice of t is indeed 1.

Theorem 2.1. *The 1-neighbors census method solves the semi-supervised weak recovery problem with any reveal ratio $\rho > 0$ for arbitrary $\text{SNR} > 0$.*

Furthermore, we derive an explicit constant bound for the overlap with the corresponding tail probability. Namely, under $\mathcal{G}(n, a/n, b/n, \rho)$, if we denote Overlap on the unrevealed nodes, between 1-neighbors census estimator and the ground truth, as Θ and let

$$\delta = \frac{\rho(a-b)}{2e^{\rho(a+b)}},$$

then

$$\mathbb{P}\left(\Theta \geq \frac{\delta}{2}\right) \geq 1 - e^{-\frac{\delta^2(1-\rho)n}{8}}.$$

A detailed discussion is provided in the next section.

Note that if $\rho \rightarrow 0$, semi-supervised weak recovery is equivalent to unsupervised weak recovery. Therefore, Theorem 1.3 implies that our result is also sharp in the sense of minimum requirement on the fraction of revealed labels.

Although it successfully solves the weak recovery problem in the largest possible regime, some limitations are hindering the census method's utility in practice. Its performance depends on a sufficient amount of revealed labels, hence requiring n to be quite large. Besides, without an unsupervised counterpart, it is not applicable when revealing is unreliable.

To address these challenges, we propose our second semi-supervised clustering algorithm which performs well in practice and covers the unsupervised setting as a special case. As discussed in the previous section, SDPs enjoy many nice properties, among which the monotone-robustness is particularly interesting to us. In the semi-supervised setting, the revealed labels are supposed to enhance the community structure. However, the work from [51] suggests such enhancement may not help with, but to the contrary can hurt the performance of many algorithms, which makes SDP an ideal starting point for us. We define the constrained semidefinite program (CSDP) as

$$\text{CSDP}(M) = \max_{\substack{X \succeq 0 \\ X_{ii}=1, \forall i \in [n]}} \{ \langle M, X \rangle : X_{ij} = x_i \cdot x_j, \forall i, j \in \mathcal{R} \}, \quad (2.4)$$

and show that it solves the semi-supervised community detection problem in the form of hypothesis testing.

Theorem 2.2. *Let $(G, x_{\mathcal{R}}) \sim \mathcal{G}(n, a/n, b/n, \rho)$ and A be the adjacency matrix associated with G . For any $a > b$, there exists $\rho_0 < 1$ such that if $\rho \geq \rho_0$, the CSDP-based test*

$$T(G, x_{\mathcal{R}}; \Delta) = \mathbb{1}_{\{\text{CSDP}(A - \frac{a}{n}\mathbf{1}\mathbf{1}^\top) \geq n[\frac{a-b}{2} - \Delta]\}}$$

will succeed with a high probability for some $\Delta > 0$.

2.1 Proof techniques

The technical challenges of establishing Theorem 2.1 root in the fact that the advantage created by revealed labels can be easily blurred out by various approximations of the limit distribution. Instead of the central limit theorem, one needs a Berry-Esseen-type inequality to derive a more quantitative result of the convergence rate. Moreover, since the distribution of each underlying component also depends on n , the conventional sample mean formulation does not apply here. We overcome the difficulty above with a direct analysis of non-asymptotic distributions, which leads to a detailed comparison between two binomial variable sequences with constant expectations.

It is quite surprising that this calculation can be carried out in a rather elegant manner since many other applications of this method are much more technically involved. For example, to establish independence among estimators, one may need to consider the leave-one-out trick. But in our case, it comes in a very natural way.

Regarding CSDP, we first show it can be coupled to an SDP with the surrogate input matrices. Moreover, its optimal value lies between two unsupervised SDPs associated with the same random graph model (different parameters). It means that all the analytical results from SDP research can be transferred into the CSDP study. However, we notice that it is common to make assumptions on the average degree d in the relevant literature. It is quite reasonable in the unsupervised setting since the graph topology is a strong indicator of the possibility of weak recovery. E.g. when $d \leq 1$, there will not exist a giant component that is of size linear in n .

To establish our result without such extra assumptions, we derive a probabilistic bound on the cut norm of the centered adjacency matrix and then use Grothendieck's inequality to bound the SDP on ERMs from above. This idea follows from [32], we give a slightly different analysis fitting for our purpose. A generalized weak law of large numbers is also derived to address the issue that distributions of the entries change as $n \rightarrow \infty$. Then we conclude the proof with a lower bound of the CSDP on SBMs considering a witness that consists of the ground truth of labels.

2.2 Outline

The rest of the paper is organized in the following way. In Section 3, we formally derive the census method and prove that it can solve the weak recovery problem throughout the entire parameter domain. In Section 4, we introduce the constrained SDP and the associated hypothesis test, through which we show that even under the KS threshold (also the information-theoretic threshold), the ERMs and the SBMs become distinguishable in the semi-supervised setting. Section 5 includes some numerical simulation results. We end the paper with concluding remarks in Section 6.

2.3 Notation

For any $n \in \mathbb{N}$, we denote the first n integers by $[n] = \{1, 2, \dots, n\}$. For a set S , its cardinality is denoted by $|S|$. We use lowercase letters for vectors (e.g. $v = (v_1, v_2, \dots, v_n)$) and uppercase letters for matrices (e.g. $M = [M_{ij}]_{i,j \in [n]}$). In particular, for adjacency

matrices, we omit their dependency on underlying graphs. Instead of A_G , we simply write A . $\mathbf{1}_n = (1, \dots, 1) \in \mathbb{R}^n$ stands for the all-ones vector, and I_n is the $n \times n$ identity matrix. $\mathbf{e}_i \in \mathbb{R}_n$ represents the i 's standard basis vector. For two real-valued matrices A and B with the same dimensions, we define the Frobenius inner product as

$$\langle A, B \rangle = \sum_{i,j} A_{ij} \cdot B_{ij} = \text{Tr}(A^\top B).$$

Vector inner product is viewed as a special case of $n \times 1$ matrices. Let

$$\|v\|_p = \left(\sum_{i=1}^p \|v_i\|^p \right)^{\frac{1}{p}}$$

be the ℓ_p norm of vectors with standard extension to $p = \infty$. Let

$$\|M\|_{p \rightarrow q} = \sup_{\|v\|_p \leq 1} \|Mv\|_q$$

be the ℓ_p -to- ℓ_q operator norm and $\|M\|_{\text{op}} := \|M\|_2 := \|M\|_{2 \rightarrow 2}$. Random graphs induce measures on the product space of label, edge, and revealed node assignments over n vertices. For any $n \in \mathbb{N}$, it is implicitly understood that one such measure is specified with that graph size. The terminology ‘‘with high probability’’ means with probability converging to 1 as $n \rightarrow \infty$. Also, we follow the conventional Big-Oh notation for asymptotic analysis. $o_p(1)$ stands for convergence to 0 in probability.

3 Census method

Analysis of the model from Definition 1.1 is a challenging task since conditioned on the graph, it is neither an Ising model nor a Markov random field. This is mainly due to the following facts:

- (1) The balance requirement puts a global condition on the size of each cluster.
- (2) Even if conditioned on sizes, there is a slight repulsion between unconnected nodes.

Namely, if two nodes do not form an edge, the probability of them being in the same community is different from the probability of them being in opposite communities.

Recent years have witnessed a series of excellent contributions to the study of phase transitions in the sparse regime. Our census method for semi-supervised clustering is mainly inspired by the natural connection between community detection on SBMs and reconstruction problems on trees, which was formally established by [54]. Intuitively, for a vertex v in $\mathcal{G}(n, a/n, b/n)$, it is not likely that a node from its small neighborhood has an edge leading back to v . Therefore, the neighborhood looks like a randomly labeled tree with high probability. Furthermore, the labeling on the vertices behaves like broadcasting a bit from the root of a tree down to its leaves (see the survey [53] for a detailed discussion).

In this section, we will first look into the census method of t -neighbors, i.e. deciding the label of a node by the majority on its neighbors at depth t . After defining the general framework, we will show that when $\text{SNR} \leq 1$, 1-neighbors voting is optimal in terms of recovering the cluster structure via informal calculation. Then, we rigorously prove that census on 1-neighbors solves the semi-supervised weak recovery problem for any $\text{SNR} > 0$ with an arbitrarily small fraction of the labels revealed.

3.1 Majority of t -neighbors

Let (G, x) obey the planted bisection model $\mathcal{G}(n, a/n, b/n)$. We denote the set of all vertices by $V(G)$. For a fixed vertex v and $t \in \mathbb{N}$, let $N_t(v)$ denote the number of vertices that are t edges away from v . $\Delta_t(v)$ is defined as the difference between the numbers of t -neighbors in each community. Namely,

$$N_t(v) = |K_t(v)|, \quad (3.1)$$

$$\Delta_t(v) = \sum_{u \in K_t(v)} x_u, \quad (3.2)$$

where

$$K_t(v) := \{u \in V(G) : d(u, v) = t\}$$

denotes the t -neighbors of v .

If one assumes that the subgraph of G induced by the vertices within t edges of v is a tree, the expected value of $N_t(v)$ is approximately $[(a+b)/2]^t$, and the expected value of $x_v \cdot \Delta_t(v)$, i.e. the expected number of these vertices in the same community as v minus the expected number of these vertices in the other community, is approximately $[(a-b)/2]^t$. So, if one can somehow independently determine which community a vertex is in with an accuracy of $1/2 + \alpha$ for some $\alpha > 0$, one will be able to predict the label of each vertex with an accuracy of roughly $1/2 + [(a-b)^2/(2(a+b))]^{t/2} \cdot \alpha$, by guessing it as the majority of v 's t -neighbors. Under the unsupervised learning setting, one can get a small advantage, $\alpha \sim \Theta(1/\sqrt{n})$, by randomly initializing labels. It is guaranteed by the central limit theorem that such a fraction exists in either an agreement or disagreement form. See Fig. 3.1 for a visual illustration.

To amplify this lucky guess, we need t to be sufficiently large so that $[(a-b)^2/(2(a+b))]^{t/2} > \sqrt{n}$, which implies $[(a+b)/2]^t > n$. Note that $d = (a+b)/2$ is the average degree. This means before the signal is strong enough for our purpose, not only our tree approximation will break down, but vertices will also be exhausted. However, if we have access to some of the true labels, i.e. in the semi-supervised setting, we can leverage the tree structure to get a non-vanishing advantage over random guessing.

Let A be the adjacency matrix associated with G . Consider the random variables Y_u representing votes of directly connected neighbors

$$Y_u = \begin{cases} x_u, & \text{if } A_{uv} = 1, \\ 0, & \text{otherwise.} \end{cases} \quad (3.3)$$




Figure 3.1: Neighborhood of node v with a tree structure. The ground truth of clusters is coded in black and white. The shaded area indicates those nodes randomly guessed to be in the same community or the opposite community as v . The annulus represents the collection of its t -neighbors.

We have

$$N_1(v) = \sum_{u \in V(G)} |Y_u|, \tag{3.4}$$

$$\Delta_1(v) = \sum_{u \in V(G)} Y_u. \tag{3.5}$$

By definition of the planted bisection model,

$$P(Y_u = 1 | x_v = 1) = \frac{P(Y_u = 1, x_v = 1)}{P(x_v = 1)} \approx \frac{a}{2n}. \tag{3.6}$$

Similarly,

$$P(Y_u = -1 | x_v = 1) \approx \frac{b}{2n}. \tag{3.7}$$

It is not exact due to the balanced community constraint. But when n is large, such an effect is negligible. Furthermore, if we consider the definition of the planted bisection model without balance constraint, the equation will be exact.

Without loss of generality, we only consider the case, where $x_v = 1$ and omit the condition on it. We have

$$\Delta_1(v) = \sum_{u \in V(G)} Y_u, \quad Y_u = \begin{cases} 1 & \text{with probability } \frac{a}{2n}, \\ -1 & \text{with probability } \frac{b}{2n}, \\ 0 & \text{with probability } 1 - \frac{a+b}{2n}, \end{cases} \tag{3.8}$$

where the Y_u 's are independent. Note that

$$\mathbb{E}(Y_u) = \frac{a-b}{2n}, \quad \mathbb{E}(Y_u^2) = \frac{a+b}{2n}.$$

Recall that $\rho \in [0, 1]$ is the ratio of revealed labels. For the sake of simplicity, we assume the total number of revealed vertices $m = \rho n \in 2\mathbb{N}$ to be an even integer. The revealed vertices are chosen arbitrarily, denoted as

$$\mathcal{R} := \{u_{n-m+1}, u_{n-m+2}, \dots, u_n\}.$$

The model also provides that the number of revealed vertices in each community is $\rho n/2$. Then the majority of revealed vertices among 1-neighborhood of v can be written as

$$\tilde{\Delta}_1(v) = \sum_{u \in \mathcal{R}} Y_u. \tag{3.9}$$

Therefore,

$$\mathbb{E}(\tilde{\Delta}_1(v)) = \sum_{u \in \mathcal{R}} \mathbb{E}(Y_u) = \rho \frac{a-b}{2}, \tag{3.10}$$

$$\text{Var}(\tilde{\Delta}_1(v)) = \sum_{u \in \mathcal{R}} \text{Var}(Y_u) = \rho \frac{a+b}{2} + o(1). \tag{3.11}$$

3.2 Locally tree-like structure

Proceeding to the t -neighbors, we need to first understand the structure of a small neighborhood in the SBM. The neighborhoods in a sparse network locally have no loops. So they have a nice tree-like structure. Moreover, the labels also obey some random broadcasting processes on trees.

A broadcasting process transmits the information from the root of a tree to all the nodes. At each level, nodes inherit the information from its parent. Meanwhile, errors could happen with a certain amount of probability. Usually, the edges are assumed to be included according to the same rule and work independently. It was first considered in genetics [18] since it perfectly describes the propagation of a gene from ancestor to descendants. It can also be interpreted as a communication network that passes out the information from the root. Such processes were intensively studied in information theory and statistical physics [13, 35, 66]. In particular, we are interested in the following Markov process since it can be identified with the labeling process of a small neighborhood in SBM.

Definition 3.1 (Galton-Watson Tree with Markov Process). *Let T be an infinite rooted tree with root v . Given a number $0 \leq \epsilon < 1$ and the offspring rate $d > 0$, we define a random labeling $\tau \in \{1, -1\}^T$. First, draw τ_v uniformly in $\{1, -1\}$. Then, recursively construct the labeling as follows:*

- *Generate children of each parent node according to a Poisson distribution with mean d .*
- *Conditionally independently given τ_v , for every child u of v , set $\tau_u = \tau_v$ with probability $1 - \epsilon$ and $\tau_u = -\tau_v$ otherwise.*

The following lemma shows that a $\ln(n)$ -neighborhood in (G, x) looks like a Galton-Watson tree with Markov process. For any $v \in G$, let G_R be the induced subgraph on $\{u \in G : d(u, v) \leq R\}$.

Lemma 3.1 ([54]). *Let*

$$R = R(n) = \frac{\ln n}{10 \ln(2(a+b))}.$$

There exists a coupling between (G, x) and (T, τ) such that $(G_R, x_{G_R}) = (T_R, \tau_{T_R})$ a.a.s.

Hence, for fixed $t \in \mathbb{N}$, $t \leq R$, and any $v \notin \mathcal{R}$, we can denote the label of a vertex in v 's t -neighborhood as

$$Y_i^{(t)} := \prod_{k=1}^t Y_{u_k},$$

where $\{Y_{u_k}\}_{k=1}^t$ are independent copies of Y_u . Then we have

$$\mathbb{E}(Y_i^{(t)}) = \left(\frac{a-b}{2n}\right)^t, \quad \mathbb{E}\left((Y_i^{(t)})^2\right) = \left(\frac{a+b}{2n}\right)^t.$$

Moreover, $\{Y_i^{(t)}\}$'s are essentially independent. Therefore, the census of v 's revealed t -neighbors can be written as

$$\tilde{\Delta}_t(v) = \sum_{i \in [\rho \cdot n^t]} Y_i^{(t)} \quad (a.a.s.). \tag{3.12}$$

The central limit theorem suggests

$$\tilde{\Delta}_t(v) \rightarrow \mathcal{N}\left(\rho \left(\frac{a-b}{2}\right)^t, \rho \left(\frac{a+b}{2}\right)^t\right) \quad \text{as } n \rightarrow \infty. \tag{3.13}$$

Hence,

$$\begin{aligned} \mathbb{P}(\tilde{\Delta}_t(v) > 0 \mid x_v = 1) &= \frac{1}{2} \left[1 + \operatorname{erf}\left(\frac{\rho[(a-b)/2]^t}{\sqrt{\rho[(a+b)/2]^t} \sqrt{2}}\right) \right] + o(1) \\ &= \frac{1}{2} + \frac{1}{2} \operatorname{erf}\left(\sqrt{\frac{\rho \operatorname{SNR}^t}{2}}\right) + o(1), \end{aligned} \tag{3.14}$$

where

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(-t^2) dt$$

is the Gauss error function.

So one can see that once SNR is less than or equal to 1, it is not beneficial to look into t -neighbors. The optimal choice of t is 1 in this situation. Since we also know that weak recovery is solvable when $\operatorname{SNR} > 1$, it makes the majority of 1-neighbors particularly interesting.

Suppose $\text{SNR} \leq 1$ and include the symmetric part of $x_v = -1$, we have

$$P(\text{sgn}(\tilde{\Delta}_1(v)) = x_v) > \frac{1}{2} + \frac{1}{3}\sqrt{\rho \text{SNR}}. \quad (3.15)$$

Consider the estimator of unrevealed labels

$$\hat{x}_{\mathcal{R}^c} := \text{sgn} \left([\tilde{\Delta}_1(u_1), \tilde{\Delta}_1(u_2), \dots, \tilde{\Delta}_1(u_{n-m})]^\top \right), \quad (3.16)$$

and the ground truth

$$x_{\mathcal{R}^c} = [x_{u_1}, x_{u_2}, \dots, x_{u_{n-m}}]^\top.$$

Recall that

$$\text{Overlap}(x_{\mathcal{R}^c}, \hat{x}_{\mathcal{R}^c}) = \frac{1}{n-m} |\langle x_{\mathcal{R}^c}, \hat{x}_{\mathcal{R}^c} \rangle|. \quad (3.17)$$

We can conclude that

$$\begin{aligned} E[\text{Overlap}(x_{\mathcal{R}^c}, \hat{x}_{\mathcal{R}^c})] &= E \left[\frac{1}{n-m} \left| \sum_{i \in [n-m]} \text{sgn}(\tilde{\Delta}_1(u_i)) x_{u_i} \right| \right] \\ &\geq \frac{1}{n-m} \left| \sum_{i \in [n-m]} E[\text{sgn}(\tilde{\Delta}_1(u_i)) x_{u_i}] \right| \\ &> \frac{2}{3} \sqrt{\rho \text{SNR}}. \end{aligned} \quad (3.18)$$

The expected overlap is not vanishing which suggests that weak recovery is solvable for any SNR. But it is technically impractical to rigorously describe the limit distribution of our census estimator without blurring this edge out. From Fig. 3.2, we can see that our calculation is close to the expectation. But the convergence rate depends on ρ . In particular, when both SNR and ρ are small, the asymptotic behavior of our algorithm remains unclear. Hence, we go through a direct analysis to establish the desired result.

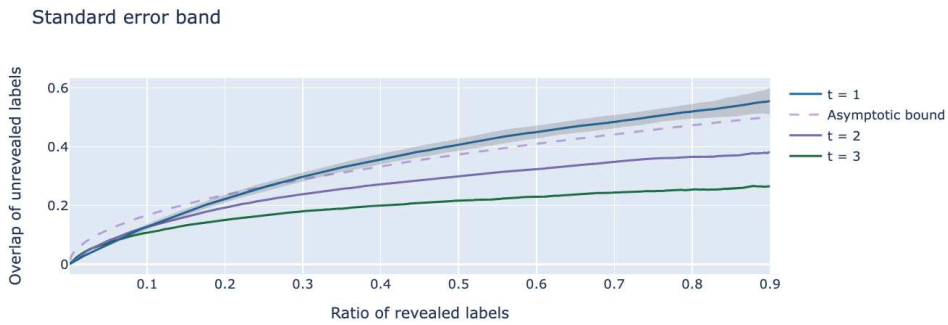


Figure 3.2: The simulation result of $\mathcal{G}(3000, 5/3000, 2/3000)$, $\text{SNR} \approx 0.64$. Solid curves stand for the average overlaps of the t -neighbors census method ($t = 1, 2, 3$) on 60 independent realizations of the random graph. The shaded area represents the standard error band of the 1-neighbors census. The dashed curve stands for the asymptotic lower bound we conclude from our calculation.

3.3 Majority of 1-neighbors

Since the algorithm is invariant under index reordering, without loss of generality, we let the adjacency matrix A be a symmetric matrix with diagonal entries $A_{ii} = 0, i = 1, 2, \dots, n$. For $1 \leq i < j \leq n$, $\{A_{ij}\}$'s are independent,

$$A_{ij} \sim \text{Bernoulli} \left(\frac{a}{n} \right), \quad \left(i \leq \frac{n}{2} \text{ and } j \leq \frac{n}{2} \right) \quad \text{or} \quad \left(i \geq \frac{n}{2} \text{ and } j \geq \frac{n}{2} \right), \quad (3.19)$$

$$A_{ij} \sim \text{Bernoulli} \left(\frac{b}{n} \right), \quad i \geq \frac{n}{2} \text{ and } j \leq \frac{n}{2}. \quad (3.20)$$

The true label x and revealed label \tilde{x} are, respectively,

$$x_i = \begin{cases} 1, & i = 1, 2, \dots, \frac{n}{2}, \\ -1, & i = \frac{n}{2} + 1, \dots, n, \end{cases} \quad (3.21)$$

$$\tilde{x}_i = \begin{cases} 1, & i = 1, 2, \dots, \frac{m}{2}, \\ -1, & i = \frac{n}{2} + 1, \dots, \frac{n+m}{2}, \\ 0, & \text{otherwise.} \end{cases}$$

For an unrevealed vertex, we consider the majority of its 1-neighbors,

$$\tilde{\Delta}_1(i) = \langle A[i, :], \tilde{x} \rangle = \sum_{j: \tilde{x}_j \neq 0} A_{ij} \tilde{x}_j = \sum_{j: \tilde{x}_j \neq 0} A_{ji} \tilde{x}_j. \quad (3.22)$$

Therefore, $\{\tilde{\Delta}_1(i)\}$'s are independent for all $i : \tilde{x}_i = 0$ since they have no common term. Note that this is not the case for all $i \in [n]$. However, we only need to predict the unrevealed labels, hence the independence holds.

The estimator given by majority voting of 1-neighbors is

$$\hat{x}_i = \begin{cases} \tilde{x}_i, & \text{if } \tilde{x}_i \neq 0, \\ \text{sgn}^*(\tilde{\Delta}_1(i)), & \text{if } \tilde{x}_i = 0. \end{cases} \quad (3.23)$$

We toss a fair coin when $\tilde{\Delta}_1(i) = 0$ to break the tie, i.e.

$$P(\text{sgn}^*(\tilde{\Delta}_1(i)) = 1 \mid \tilde{\Delta}_1(i) = 0) = P(\text{sgn}^*(\tilde{\Delta}_1(i)) = -1 \mid \tilde{\Delta}_1(i) = 0) = \frac{1}{2}. \quad (3.24)$$

Formally, for a finite sequence of random variables W_i ,

$$\text{sgn}^*(W_i) = \text{sgn}(W_i) + \mathbb{1}_{\{W_i=0\}}(2B_i - 1), \quad (3.25)$$

where $\{B_i\}$'s are i.i.d. Bernoulli(1/2) random variables, independent of all other random objects. Note that it is only introduced for analysis purposes and is equivalent to the conventional sign function in practice.

Suppose (G, x) is an Erdős-Rényi random graph with revealed label \tilde{x} , any estimator can only have a vanishing correlation with the true label among the unrevealed vertices. So the semi-supervised weak recovery problem on SBM requires finding an estimator such that the correlation restricted on the unrevealed part is non-vanishing. Formally, we want to show that

$$P(\text{Overlap}(x|_{\tilde{x}_i=0}, \hat{x}|_{\tilde{x}_i=0}) \geq \Omega(1)) = 1 - o(1). \quad (3.26)$$

As discussed in Section 2.1, we start with a critical result scrutinizing binomial variable sequences. It gives us an edge over direct analysis via a Berry-Esseen-type inequality, which usually assumes the distribution of individual random variables in the sequence independent of n .

Lemma 3.2. *Let X and Y be independent binomial random variables with $X \sim \text{Binomial}(n, a/n)$ and $Y \sim \text{Binomial}(n, b/n)$, $a > b$. Denote*

$$\delta = \delta(a, b) := \frac{a - b}{2 \exp(a + b)}.$$

Then, for sufficiently large n ,

$$P(X > Y) - P(X < Y) \geq \delta. \quad (3.27)$$

Remark 3.1. By symmetry, we always have $P(X > Y) - P(X < Y) > 0$. This lemma guarantees the difference will not vanish as $n \rightarrow \infty$.

Proof. By the law of total probability and independence, we have

$$\begin{aligned} P(X > Y) &= \sum_{x=1}^n P(Y < x) P(X = x) = \sum_{x=1}^n \sum_{y=0}^{x-1} P(Y = y) P(X = x) \\ &= \sum_{x=1}^n \sum_{y=0}^{x-1} \left[\binom{n}{x} \left(\frac{a}{n}\right)^x \left(1 - \frac{a}{n}\right)^{n-x} \binom{n}{y} \left(\frac{b}{n}\right)^y \left(1 - \frac{b}{n}\right)^{n-y} \right]. \end{aligned} \quad (3.28)$$

Let $\Delta := P(X > Y) - P(X < Y)$, then

$$\begin{aligned} \Delta &= \sum_{x=1}^n \binom{n}{x} \left(\frac{a}{n}\right)^x \left(1 - \frac{a}{n}\right)^{n-x} \binom{n}{x} \left(\frac{b}{n}\right)^x \left(1 - \frac{b}{n}\right)^{n-x} \\ &\quad \times \left\{ \sum_{y=0}^{x-1} \binom{n}{y} \left[\left(\frac{b}{n}\right)^{y-x} \left(1 - \frac{b}{n}\right)^{x-y} - \left(\frac{a}{n}\right)^{y-x} \left(1 - \frac{a}{n}\right)^{x-y} \right] \right\} \\ &= \sum_{x=1}^n \binom{n}{x} \left(\frac{ab}{n}\right)^x \left(1 - \frac{a+b}{n} + \frac{ab}{n^2}\right)^{n-x} \\ &\quad \times \left\{ \sum_{y=0}^{x-1} \binom{n}{y} \frac{1}{n^y} \left[\left(\frac{1}{b} - \frac{1}{n}\right)^{x-y} - \left(\frac{1}{a} - \frac{1}{n}\right)^{x-y} \right] \right\}. \end{aligned}$$

Let $f(x) = \alpha^x - \beta^x, \alpha > \beta > 0$. Since

$$f'(x) = \alpha^x \ln \alpha - \beta^x \ln \beta > 0,$$

we have

$$f(m) \geq f(1) = \alpha - \beta, \quad \forall m \in \mathbb{N}.$$

So

$$\left(\frac{1}{b} - \frac{1}{n}\right)^{x-y} - \left(\frac{1}{a} - \frac{1}{n}\right)^{x-y} \geq \frac{a-b}{ab}.$$

Also notice that

$$\binom{n}{m} = \prod_{i=0}^{m-1} \frac{n-i}{m-i} \geq \left(\frac{n}{m}\right)^m, \quad \forall 1 \leq m \leq n.$$

We have

$$\begin{aligned} \Delta &\geq \sum_{x=1}^n \left(\frac{ab}{x}\right)^x \left(1 - \frac{a+b}{n}\right)^{n-x} \left(\sum_{y=0}^{x-1} \frac{1}{y^y} \cdot \frac{a-b}{ab}\right) \\ &\geq (a-b) \left(1 - \frac{a+b}{n}\right)^n \\ &\geq \frac{a-b}{2 \exp(a+b)}, \end{aligned} \tag{3.29}$$

for sufficiently large n , where we follow the convention that $0^0 = 1$. □

Then, we can simply resort to a classical concentration inequality to bound the overlap.

Lemma 3.3 ([20, Chernoff-Hoeffding Theorem]). *Suppose X_1, \dots, X_n are i.i.d. random variables, taking values in $\{0, 1\}$. Let $p = E(X)$ and $\epsilon > 0$. Then*

$$P\left(\frac{1}{n} \sum X_i \leq p - \epsilon\right) \leq \left(\left(\frac{p}{p-\epsilon}\right)^{p-\epsilon} \left(\frac{1-p}{1-p+\epsilon}\right)^{1-p+\epsilon}\right)^n = e^{-D(p-\epsilon||p)n}, \tag{3.30}$$

where

$$D(x||y) = x \ln \frac{x}{y} + (1-x) \ln \left(\frac{1-x}{1-y}\right)$$

is the Kullback-Leibler-divergence between Bernoulli distributed random variables with parameters x and y .

We now convert the KL divergence to the total variation distance, which is easier to work with. Let P_1 and P_2 be two probability measures defined on the same sample space Ω and sigma-algebra \mathcal{F} . The total variation distance between them is defined as

$$d_{TV}(P_1, P_2) = \sup_{E \in \mathcal{F}} |P_1(E) - P_2(E)|.$$

Moreover, in the discrete case, we have the following identity:

$$d_{TV}(P_1, P_2) = \frac{1}{2} \|P_1 - P_2\|_1 = \sum_{\omega \in \Omega} \frac{1}{2} \|P_1(\omega) - P_2(\omega)\|.$$

It is related to the KL divergence through Pinsker's inequality (see, e.g. [69, Chapter 3]). For completeness, we include an elementary proof of the Bernoulli special case that is sufficient for our usage later.

Lemma 3.4. *Let P_1 and P_2 be two Bernoulli distributions, where $P_1(1) = x$ and $P_2(1) = y$. We have*

$$2(d_{TV}(P_1, P_2))^2 \leq D(x||y). \tag{3.31}$$

Proof. We can manipulate both sides of the inequality as

$$D(x||y) = x \ln \frac{x}{y} + (1-x) \ln \left(\frac{1-x}{1-y} \right), \tag{3.32}$$

$$2(d_{TV}(P_1, P_2))^2 = \frac{1}{2} \|P_1 - P_2\|_1^2 = 2(x-y)^2. \tag{3.33}$$

We denote

$$f(x, y) = x \ln \frac{x}{y} + (1-x) \ln \frac{1-x}{1-y} - 2(x-y)^2.$$

Therefore,

$$\frac{\partial f}{\partial y} = (x-y) \left[4 - \frac{1}{y(1-y)} \right]. \tag{3.34}$$

Notice that since $0 \leq y \leq 1$, $y(1-y) \leq 1/4$. So $4 - 1/(y(1-y))$ is always negative. Thus, for fixed x , $f(x, y) \geq f(x, x) = 0, \forall y$. Hence,

$$D(x||y) - 2(d_{TV}(P_1, P_2))^2 \geq 0. \tag{3.35}$$

The proof is complete. □

Now we prove the main result for the census method.

Proof of Theorem 2.1. Recall that for any i such that $\tilde{x}_i = 0$, our estimator is defined as $\hat{x}_i = \text{sgn}^*(\tilde{\Delta}_1(i))$ and

$$\tilde{\Delta}_1(i) = \sum_{j: \tilde{x}_j \neq 0} A_{ij} \tilde{x}_j = \left(\sum_{\rho \frac{n}{2} < j \leq \frac{n}{2}} A_{ij} \right) - \left(\sum_{(1+\rho) \frac{n}{2} < j \leq n} A_{ij} \right). \tag{3.36}$$

It is indeed the difference between two independent binomial variables with parameters $(\rho n, \rho a / \rho n)$ and $(\rho n, \rho b / \rho n)$. By Lemma 3.2, we have

$$\text{P}(\text{sgn}(\tilde{\Delta}_1(i)) = x_i) - \text{P}(\text{sgn}(\tilde{\Delta}_1(i)) = -x_i) \geq \delta = \frac{\rho(a-b)}{2e^{\rho(a+b)}} \tag{3.37}$$

for sufficiently large n . Also, notice that

$$\mathbb{P}(\text{sgn}(\tilde{\Delta}_1(i)) = -x_i) = 1 - \mathbb{P}(\text{sgn}(\tilde{\Delta}_1(i)) = x_i) - \mathbb{P}(\tilde{\Delta}_1(i) = 0). \quad (3.38)$$

Therefore,

$$\mathbb{P}(\text{sgn}(\tilde{\Delta}_1(i)) = x_i) \geq \frac{1 + \delta}{2} - \frac{1}{2} \mathbb{P}(\tilde{\Delta}_1(i) = 0). \quad (3.39)$$

Then, by the law of total probability, we have

$$\begin{aligned} \mathbb{P}(\hat{x}_i = x_i) &= \mathbb{P}(\text{sgn}^*(\tilde{\Delta}_1(i)) = x_i) \\ &= \mathbb{P}(\text{sgn}(\tilde{\Delta}_1(i)) = x_i) + \frac{1}{2} \mathbb{P}(\tilde{\Delta}_1(i) = 0) \\ &\geq \frac{1}{2} + \frac{\delta}{2}. \end{aligned} \quad (3.40)$$

Since the $\{\hat{x}_i\}$'s are independent for all unrevealed vertices as $\{\tilde{\Delta}_1(i)\}$'s and

$$\mathbb{E} \left[\frac{\hat{x}_i x_i + 1}{2} \right] = \mathbb{P}(\hat{x}_i = x_i).$$

Lemmas 3.3 and 3.4 give us that

$$\mathbb{P} \left(\frac{1}{(1-\rho)n} \sum_{i:\hat{x}_i=0} \frac{\hat{x}_i x_i + 1}{2} \leq \frac{1}{2} + \frac{\delta}{2} - \epsilon \right) \leq e^{-2\epsilon^2(1-\rho)n}. \quad (3.41)$$

Taking $\epsilon = \delta/4$, we have

$$\mathbb{P} \left(\text{Overlap}(x|_{\hat{x}_i=0}, \hat{x}|_{\hat{x}_i=0}) \geq \frac{\delta}{2} \right) \geq 1 - e^{-\frac{\delta^2(1-\rho)n}{8}}. \quad (3.42)$$

As long as $a > b$, we have $\delta > 0$, which concludes the proof. \square

Corollary 3.1. *The semi-supervised SBM and ERM are not mutually contiguous for any given $a > b \geq 0$ and $\rho > 0$.*

Proof. Let

$$P_n^{(0)} = \mathcal{G} \left(n, \frac{a+b}{2n}, \frac{a+b}{2n}, \rho \right), \quad P_n^{(1)} = \mathcal{G} \left(n, \frac{a}{n}, \frac{b}{n}, \rho \right).$$

Then consider the same constant $\delta > 0$ from the proof of Theorem 2.1 and denote the event sequence

$$E_n = \left\{ \text{Overlap}(x|_{\hat{x}_i=0}, \hat{x}|_{\hat{x}_i=0}) \geq \frac{\delta}{2} \right\},$$

where \hat{x} is our semi-supervised census estimator. We have

$$P_n^{(0)}(E_n) \rightarrow 0 \quad (\text{Law of large number}), \quad (3.43)$$

$$P_n^{(1)}(E_n) \not\rightarrow 0 \quad (\text{Bounded from below}). \quad (3.44)$$

The proof is complete. \square

Recall that distinguishability and weak recovery are equivalent in the sense that they share the same phase transition threshold. In the semi-supervised setting, it is straightforward that the weak recovery implies distinguishability. So, Theorem 2.1 suggests that in this case, SBM and ERM are always distinguishable, which is also equivalent to Corollary 3.1. Hence, the two problems share the same solvable region as well.

4 Semi-supervised SDP

We have seen that the census method solves the semi-supervised community detection problem. However, the algorithm is desirable in practice only when the amount of revealed labels is sufficient to support a reasonable performance. In other words, it has no unsupervised fallback built-in. Meanwhile, SDPs enjoy nice properties like optimality and robustness as mentioned earlier. It is also well known that approximate information about the extremal cuts of graphs can be obtained by computing the optimizer for the SDP of their adjacency matrix, for example [29]. From both a practical and a mathematical point of view, we are interested in developing an SDP-based semi-supervised clustering approach, through which we shall be able to see the models, algorithms, and phase transitions with a fresh perspective.

In this section, we will focus on the hypothesis testing formulation of the community detection problem. We have discussed the equivalency between it and the non-vanishing overlap formulation under the unsupervised setting. In the semi-supervised scenario, it is still an interesting question to ask whether there exists a test that can distinguish SBMs from ERMs. Here we understand ERM as the special case of SBM with $a = b$. It also has ground truth of labels, which is uniformly random under the balance constraint. Given that they are originally contiguous when $\text{SNR} \leq 1$, we want to show that revealed labels together with random graphs can separate them.

4.1 SDP for community detection

Under the planted bisection model specified in Definition 1.1, the MAP estimator is equivalent to the maximum likelihood estimator, which is given by min-bisection, i.e. a balanced partition with the least number of crossing edges. Formally, it can be written the following optimization problem:

$$\max_{\substack{x \in \{1, -1\}^n \\ x^\top \mathbf{1} = 0}} x^\top A x. \tag{4.1}$$

By lifting the variable $X := x x^\top$, we can rewrite it as

$$\hat{X}_{\text{MAP}}(G) = \arg \max_{\substack{X \succeq 0 \\ X_{ii} = 1, \forall i \in [n] \\ \text{rank}(X) = 1 \\ X \mathbf{1} = 0}} \langle A, X \rangle. \tag{4.2}$$

Although min-bisection of G is optimal (in the MAP sense) for exact recovery, finding it is NP-hard. Various relaxations have been proposed for the MAP estimator. Since the rank constraint makes the optimization difficult, we can remove it to make the problem convex. One can also get rid of the balance constraint by centralizing the adjacency matrix,

$$\tilde{A} := A - \frac{d}{n}\mathbf{1}\mathbf{1}^\top$$

with $d = (a + b)/2$ the average degree. This can also be justified using Lagrangian multipliers. The resulting semidefinite relaxation is given by

$$\hat{X}_{\text{SDP}}(G) = \underset{\substack{X \succeq 0 \\ X_{ii} = 1, \forall i \in [n]}}{\text{arg max}} \langle \tilde{A}, X \rangle. \quad (4.3)$$

The feasible region $\{X \in \mathbb{R}^{n \times n} : X \succeq 0, X_{ii} = 1, \forall i \in [n]\}$ is indeed the space of correlation matrices, which defines a subset of the unit hypercube and is also called the elliptope. Although it is derived from the relaxation of MAP, one can define the SDP for general symmetric matrices as

$$\text{SDP}(M_{n \times n}) = \max\{\langle M, X \rangle : X \in \text{elliptope}_n\}. \quad (4.4)$$

Proposition 4.1. *For any $n \times n$ symmetric matrix M , if we denote its leading eigenvalue as λ_1 , then*

$$\frac{1}{n}\text{SDP}(M) \leq \lambda_1. \quad (4.5)$$

Proof. For any feasible $X \succeq 0$ and $X_{ii} = 1$, we have $\text{Tr}(X) = n$, moreover,

$$\langle X, M = U\Lambda U^\top \rangle = \text{Tr}(U^\top XU\Lambda) = \langle Y := U^\top XU, \Lambda \rangle = \sum Y_{ii}\lambda_i \leq n\lambda_1. \quad (4.6)$$

The last inequality follows from $\text{Tr}(Y) = n$ and $Y \succeq 0$ so that $Y_{ii} \geq 0$. \square

This proposition relates SDPs to spectra of the underlying matrices, which suffer from those high-degree nodes as we mentioned in the introduction. In contrast, SDPs behave similarly on SBMs and random regular graphs. The optimal values of the SDPs for both are approximately $2n\sqrt{d}$, see [52]. Random regular graphs obey the uniform distribution over graphs with n vertices and uniform degree d , which provide a simple example to illustrate the regularity property of SDPs. We cite an intermediate result from the original proof as Lemma 4.4.

An important way to understand SDPs is by considering the Burer-Monteiro factorization of X , which characterizes the constraints. We have

$$X = \Sigma\Sigma^\top, \quad (4.7)$$

where $\Sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)^\top$ and $\|\sigma_i\|_2 = 1$, for any $i \in [n]$. Therefore, the i -th node of the graph is associated with the vector σ_i that lies on the unit sphere. $X_{ij} = \langle \sigma_i, \sigma_j \rangle$ can be interpreted as the affinity metric between nodes i and j . SDP maximizes the likelihood score of this affinity matrix concerning the given centralized adjacency matrix. The optimizer X^* is a better representation of the structure information than the vanilla adjacency matrix. Then we can identify the labels by simply running the k -means method on it or compute the eigenvector corresponding to the largest eigenvalue.

4.2 Constrained SDP and semi-supervised testing

In this section, we will introduce our SDP modification and prove that it solves the semi-supervised community detection problem with the hypothesis testing formulation. Let x denote labels of $G(n, a/n, b/n)$. And m of them are revealed uniformly at random in a balanced manner. Conditioned on the ground truth of clusters, indices of revealed nodes \mathcal{R} and edges are independent. So, without loss of generality, we denote revealed labels \tilde{x} as follows:

$$x_i = \begin{cases} 1, & i = 1, 2, \dots, \frac{n}{2}, \\ -1, & i = \frac{n}{2}, \frac{n}{2} + 1, \dots, n, \end{cases} \quad (4.8a)$$

$$\tilde{x}_i = \begin{cases} 1, & i = 1, 2, \dots, \frac{m}{2}, \\ -1, & i = \frac{n}{2}, \frac{n}{2} + 1, \dots, \frac{n+m}{2}, \\ 0, & \text{otherwise.} \end{cases} \quad (4.8b)$$

We have shown that the entry value of the optimizer X can be interpreted as an affinity metric among nodes. Moreover, we have $X_{ij} \in [-1, 1]$, for any i, j . It is natural to force the optimizer to have large entry values for those vertex pairs in which we have high confidence to be in the same community and vice versa. Therefore, we propose the CSDP approach to integrate the information provided by the semi-supervised approach. If node i and node j are revealed to have the same label, we add the constraint $X_{ij} = 1$ to the optimization model. If they are revealed to have the opposite labels, we add $X_{ij} = -1$. Formally, the CSDP is defined as

$$\text{CSDP}(M_{n \times n}) = \max \{ \langle M, X \rangle : X \in \text{elliptope}_n, X_{ij} = x_i \cdot x_j, \forall i, j \in \mathcal{R} \}, \quad (4.9)$$

where \mathcal{R} denotes the collection of revealed nodes. After reordering the indices, we can assume it as $\{1, 2, \dots, m/2\} \cup \{n/2, n/2 + 1, \dots, (n+m)/2\}$. It is worth noting that the optimization remains a positive semidefinite programming problem, which can be solved efficiently, for example by interior point methods [8].

Then let

$$\mathcal{S}^{n-1} := \{v \in \mathbb{R}^n : \|v\|_2 = 1\}$$

be the unit $(n-1)$ -sphere and

$$\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n) \in (\mathcal{S}^{n-1})^n.$$

Consider the CSDP in the form of Burer-Monteiro factorization. We have the following identities:

$$\text{SDP}(M) = \max \left\{ \sum_{i,j=1}^n M_{ij} \langle \sigma_i, \sigma_j \rangle : \sigma_i \in \mathcal{S}^{n-1}, \forall i \in [n] \right\}, \quad (4.10)$$

$$\text{CSDP}(M) = \max_{\sigma \in (\mathcal{S}^{n-1})^n} \left\{ \sum_{i,j=1}^n M_{ij} \langle \sigma_i, \sigma_j \rangle : \sigma_i^\top \sigma_j = x_i x_j, \forall i, j \in \mathcal{R} \right\}$$

$$= \max_{\sigma \in (\mathcal{S}^{n-1})^n} \left\{ \sum_{i,j \in [n] \setminus \mathcal{R}} M_{ij} \sigma_i^\top \sigma_j + \sum_{i,j \in \mathcal{R}} M_{ij} x_i x_j + 2 \sum_{i \in \mathcal{R}} \sum_{j \in [n] \setminus \mathcal{R}} x_i M_{ij} \sigma_0^\top \sigma_j \right\}, \quad (4.11)$$

where $\sigma_0 \equiv x_i \sigma_i$, for any $i \in \mathcal{R}$. Now one can consider an alternative matrix with a special margin denoting the algebraic sum of the blocks from M that are associated with \mathcal{R} . We define M^{agg} to be the $(n - m + 1) \times (n - m + 1)$ symmetric matrix indexed from 0 that

$$M_{00}^{\text{agg}} = \sum_{i,j \in \mathcal{R}} M_{ij} x_i x_j, \quad (4.12)$$

$$M_{0j}^{\text{agg}} = \sum_{i \in \mathcal{R}} x_i M_{i,j+\frac{m}{2}}, \quad \forall j \in \left[\frac{n}{2} - \frac{m}{2} \right], \quad (4.13)$$

$$M_{0j}^{\text{agg}} = \sum_{i \in \mathcal{R}} x_i M_{i,j+m}, \quad \forall j \in [n - m] \setminus \left[\frac{n}{2} - \frac{m}{2} \right], \quad (4.14)$$

$$M_{ij}^{\text{agg}} = M_{i+\frac{m}{2},j+\frac{m}{2}}, \quad \forall i, j \in \left[\frac{n}{2} - \frac{m}{2} \right], \quad (4.15)$$

$$M_{ij}^{\text{agg}} = M_{i+m,j+m}, \quad \forall i, j \in [n - m] \setminus \left[\frac{n}{2} - \frac{m}{2} \right]. \quad (4.16)$$

Essentially, we aggregate the rows and columns related to revealed vertices according to their communities into the 0-th row and column and reindex the matrix. It introduces spikiness to the underlying matrix

$$M^{\text{agg}} = \left(\begin{array}{c|cccc} \sum_{i,j \in \mathcal{R}} M_{ij} x_i x_j & M_{01}^{\text{agg}} & M_{02}^{\text{agg}} & \dots & M_{0,n-m}^{\text{agg}} \\ \hline M_{01}^{\text{agg}} & & & & \\ M_{02}^{\text{agg}} & & & & \\ \vdots & & & & \\ M_{0,n-m}^{\text{agg}} & & & & \\ \hline & & M_{\mathcal{R}\mathcal{C}} & & \end{array} \right). \quad (4.17)$$

Although [52] employed a rather different technique to study SDPs, they also noticed that the critical change comes with such built-in structures, where the authors state “We expect the phase transition in $\text{SDP}(\lambda v v^\top + W)/n$ to depend-in general-on the vector v , and in particular on how ‘spiky’ this is”.

Combining the transformed input matrix with Eq. (4.11), we conclude that CSDP is indeed an SDP regarding M^{agg}

$$\begin{aligned} \text{CSDP}(M) &= \max_{\substack{\sigma_i \in \mathcal{S}^{n-m} \\ i=0,1,\dots,n-m}} \left\{ \sum_{i,j \in [n-m]} M_{ij}^{\text{agg}} \sigma_i^\top \sigma_j + M_{00}^{\text{agg}} + 2 \sum_{j \in [n-m]} M_{0j}^{\text{agg}} \sigma_0^\top \sigma_j \right\} \\ &= \text{SDP}(M^{\text{agg}}). \end{aligned} \quad (4.18)$$

Lemma 4.1. *Let $M_{\mathcal{R}\mathcal{C}}$ be the principle submatrix of M obtained by removing the rows and columns associated with \mathcal{R} . The following inequalities hold:*

$$\text{SDP}(M_{\mathcal{R}\mathcal{C}}) \leq \text{CSDP}(M) - M_{00}^{\text{agg}}. \quad (4.19)$$

Proof. Let X^* be the optimizer of $\text{SDP}(M_{\mathcal{R}\mathcal{C}})$. Define its $(n - m + 1) \times (n - m + 1)$ extension \hat{X}^* as

$$\hat{X}_{ij}^* = \begin{cases} 1, & i = j = 0, \\ 0, & i \in [n - m], \quad j = 0, \\ 0, & j \in [n - m], \quad i = 0, \\ X_{ij}^*, & \text{otherwise.} \end{cases} \quad (4.20)$$

Due to the identity from above and the fact that $\hat{X}^* \in \text{elliptope}_{n-m+1}$ is feasible, we can conclude that

$$\text{CSDP}(M) = \text{SDP}(M^{\text{agg}}) \geq \langle \hat{X}^*, M^{\text{agg}} \rangle = \text{SDP}(M_{\mathcal{R}\mathcal{C}}) + M_{00}^{\text{agg}}. \quad (4.21)$$

The proof is complete. \square

So far, all the results are deterministic, M can be an arbitrary symmetric matrix, and \mathcal{R} can be any balanced index set. Next, we will consider

$$M = \tilde{A} := A - \frac{d}{n} \mathbf{1}\mathbf{1}^\top$$

to study CSDPs on probabilistic models.

Remark 4.1. As shown in the Lemma 4.2, $\tilde{A}_{00}^{\text{agg}} \geq m \cdot (a - b) / 2 \geq 0$ with high probability. By definition, we have $\text{CSDP}(\tilde{A}) \leq \text{SDP}(\tilde{A})$. So, with high probability,

$$\text{SDP}(\tilde{A}_{\mathcal{R}\mathcal{C}}) \leq \text{CSDP}(\tilde{A}) \leq \text{SDP}(\tilde{A}). \quad (4.22)$$

The CSDP always lies in between the SDPs of the original adjacency matrix and the submatrix of unrevealed vertices. Moreover, if $\tilde{A} \sim \mathcal{G}(n, a/n, b/n)$, we have

$$\tilde{A}_{\mathcal{R}\mathcal{C}} \sim \mathcal{G}\left(n - m, \frac{a(1 - \rho)}{n - m}, \frac{b(1 - \rho)}{n - m}\right).$$

It is worth mentioning that although $\tilde{A}_{\mathcal{R}\mathcal{C}}$ is just a submatrix of the original centered adjacency matrix, its probabilistic distribution as a random matrix is not simply changed from n nodes to $n - m$ nodes. The edge probability parameters are also changed by a factor of $(1 - \rho)$. It leads to some technical challenges, which we are going to handle later. But intuitively, from the asymptotic behavior of SDP, we can derive a rough understanding of CSDP as $n \rightarrow \infty$. Recall that the phase transition theory tells us that when $\text{SNR} \leq 1$, the optimal value of SDP for SBM will not be large enough to distinguish from the optimal value of SDP for ERM. Therefore, the order of the above quantities from inequality (4.22) suggests that semi-supervised SDP can not help to increase the statistics associated with SBM. The best one can hope for is that it will make the statistics associated with ERM smaller by a factor depending on ρ . This turns out to be enough for community detection.

Recall the community detection problem can be formalized as a binary hypothesis testing problem, whereby we want to determine, with a high probability of success, whether

the random graph under consideration has a community structure or not. As discussed in Section 2, we introduce semi-supervised learning to the problem by revealing a part of the labels involved in the random graph's generating process. Namely, if the labels associated with a graph G over n vertices are denoted as x , we choose m of them uniformly at random and denote the index set by \mathcal{R} such that $\sum_{i \in \mathcal{R}} x_i = 0$.

Given a realization of the random graph G and the revealed labels $x_{\mathcal{R}}$, we want to decide which of the following holds:

Hypothesis 0: $(G, x_{\mathcal{R}}) \sim \mathcal{G}(n, d/n, \rho)$ is an Erdős-Rényi random graph with edge probability d/n , $d = (a + b)/2$ and reveal ratio ρ . We denote the corresponding distribution over graphs by P_0 .

Hypothesis 1: $(G, x_{\mathcal{R}}) \sim \mathcal{G}(n, a/n, b/n, \rho)$ is a planted bisection random graph with edge probabilities $(a/n, b/n)$ and reveal ratio ρ . We denote the corresponding distribution over graphs by P_1 .

A statistical test T is a function defined on the graphs and revealed labels with range $\{0, 1\}$. It succeeds with a high probability if

$$P_0(T(G, x_{\mathcal{R}}) = 1) + P_1(T(G, x_{\mathcal{R}}) = 0) \rightarrow 0, \quad n \rightarrow \infty. \quad (4.23)$$

Note that this is indeed a generalization of the unsupervised community detection. Simply looking at the labels, the two models are indistinguishable. What characterizes their difference is the probabilistic law of how edges are generated, i.e. whether there is a cluster structure. The revealed labels serve as an enhancement of the graph observed. The phase transition theory says that under the unsupervised setting (the special case when $\rho = 0$), no test can succeed with high probability when $\text{SNR} \leq 1$, or equivalently, $a - b \leq \sqrt{2(a + b)}$. While if $\text{SNR} > 1$, several polynomially computable tests are developed. The SDP-based test is nearly optimal, in the sense that it requires

$$\frac{a - b}{\sqrt{2(a + b)}} \geq 1 + \epsilon(d), \quad (4.24)$$

where $\epsilon(d) \rightarrow 0$ as $d \rightarrow \infty$. It is believed to be the best that SDPs can reach. As the monotone-robustness study suggests [51], this gap may be necessary, since SDP is indeed solving a harder problem where no algorithm can approach the threshold. However, we are going to see that when ρ is sufficiently large, SDPs can not only reach but cross the threshold.

4.3 Semi-supervised detectability

With the problem and algorithm defined clearly, we are ready to prove that SBM and ERM can be consistently distinguished in the semi-supervised setting. We take a divide and conquer approach to establish an upper bound of CSDP on ERM, while we bound the CSDP on SBM from below with a witness that consists of the ground truth of labels, $X = xx^T$.

Lemma 4.2. Let (A, x) obey the planted bisection model $\mathcal{G}(n, a/n, b/n)$ and denote $\langle xx^\top, \tilde{A} \rangle$ as Y . Then, for any $\epsilon > 0$, we have $Y/n \in [(a-b)/2 - \epsilon, (a-b)/2 + \epsilon]$ with probability converging to one as $n \rightarrow \infty$.

Proof.

$$\begin{aligned} Y &= \langle xx^\top, \tilde{A} \rangle = \langle xx^\top, A \rangle - \frac{d}{n} \langle xx^\top, \mathbf{1}\mathbf{1}^\top \rangle \\ &\stackrel{d}{=} 2 \cdot \left[\text{Bin} \left(\left(\frac{n}{2} \right)^2 - \frac{n}{2}, \frac{a}{n} \right) - \text{Bin} \left(\left(\frac{n}{2} \right)^2, \frac{b}{n} \right) \right]. \end{aligned} \quad (4.25)$$

We have $EY = n(a-b)/2 - a$ and

$$\text{Var} Y = 4 \left(a \left(\frac{n}{4} - \frac{1}{2} \right) \left(1 - \frac{a}{n} \right) + b \frac{n}{4} \left(1 - \frac{b}{n} \right) \right) \leq n(a+b). \quad (4.26)$$

Then Chebyshev's inequality implies that for any $\delta \in (0, 1)$,

$$\begin{aligned} &P \left(\left| Y - \frac{n}{2}(a-b) + a \right| \geq \sqrt{n(a+b)} \cdot n^{\frac{1-\delta}{2}} \right) \leq \frac{1}{n^{1-\delta}} \\ \implies &P \left(\left| \frac{Y}{n} - \frac{a-b}{2} + \frac{a}{n} \right| \geq \frac{\sqrt{a+b}}{n^{\frac{\delta}{2}}} \right) \leq \frac{1}{n^{1-\delta}}. \end{aligned}$$

Hence, for sufficiently large n , we have

$$P \left(\frac{Y}{n} \geq \frac{a-b}{2} + \epsilon \right) + P \left(\frac{Y}{n} \leq \frac{a-b}{2} - \epsilon \right) \leq \frac{1}{n^{1-\delta}}. \quad (4.27)$$

Therefore,

$$P \left(\frac{Y}{n} \in \left[\frac{a-b}{2} - \epsilon, \frac{a-b}{2} + \epsilon \right] \right) \geq 1 - \frac{1}{n^{1-\delta}}. \quad (4.28)$$

The proof is complete. \square

Besides bounding the outcomes on the SBM from below, this lemma can also be applied to the all revealed blocks to estimate $\tilde{A}_{00}^{\text{agg}}$, which is used several times throughout our proofs.

Lemma 4.3. Let $G \sim \mathcal{G}(n, a/n, b/n)$, $d = (a+b)/2$ and

$$\tilde{A} = A - \frac{d}{n} \mathbf{1}\mathbf{1}^\top$$

be its centered adjacency matrix. Then for any $\epsilon > 0$ and $\gamma > 0$, with probability at least $1 - 1/(n^{1-\gamma})$ for all $n \geq n_0(a, b, \epsilon, \gamma)$, we have

$$\text{CSDP}(\tilde{A}) \geq n \left(\frac{a-b}{2} - \epsilon \right). \quad (4.29)$$

Proof. We prove the lower bound by considering a witness of the constrained optimization problem. Notice that xx^\top is feasible for both SDP and CSDP, where x is the label vector associated with G . Therefore,

$$\text{CSDP}(\tilde{A}) \geq \langle xx^\top, \tilde{A} \rangle. \quad (4.30)$$

Then, we can apply Lemma 4.2 to get the result. \square

This result holds for any $\text{SNR} > 0$ and suggests the following test for the semi-supervised community detection problem:

$$T(G, x_{\mathcal{R}}; \Delta) = \begin{cases} 1, & \text{if } \text{CSDP}(\tilde{A}) \geq n \left[\frac{a-b}{2} - \Delta \right], \\ 0, & \text{otherwise.} \end{cases} \quad (4.31)$$

The following lemma bounds the CSDP of ERM from above. Intuitively, the contribution of blocks of the adjacency matrix, where columns or rows are associated with revealed nodes, concentrates well around zero. So the effective dimension of the SDP is reduced, hence the optimal value.

Lemma 4.4 ([52, Theorem 1]. Reformulated). *Let $G \sim \mathcal{G}(n, d/n)$ and*

$$\tilde{A} = A - \frac{d}{n} \mathbf{1}\mathbf{1}^\top$$

be its centered adjacency matrix. There exists absolute constants C and $d_0 > 1$ such that if $d \geq d_0$, then with high probability,

$$\frac{1}{n\sqrt{d}} \text{SDP}(\tilde{A}) \leq 2 + \frac{C \log d}{d^{\frac{1}{10}}}. \quad (4.32)$$

This result is rigorously derived with profound insights from mathematical physics. However, there is an implicit condition on the average degree d in the proof. It is common to assume at least $d > 1$ in the literature concerning unsupervised clustering because otherwise, the graph has no giant component, not to mention reconstruction, as discussed in Section 1.2. But our approach leads to the subgraph with a possibly small effective average degree. Moreover, we do not want to be limited by the topology structure, although it is indeed a fundamental limit in the unsupervised setting. Theorem 2.2 shows that semi-supervised SDPs are capable of integrating those sublinear components. To achieve that we resort to Grothendieck’s inequality and carry out the analysis without assumption on d .

Theorem 4.1 ([31, Grothendieck’s Inequality]). *Let M be a $n \times n$ real matrix. If for any $s, t \in \{-1, 1\}^n$,*

$$\left| \sum_{i,j} M_{ij} s_i t_j \right| \leq 1. \quad (4.33)$$

Then for all vectors $X_i, Y_i \in \{x \in \mathbb{R}^n : \|x\|_2 \leq 1\}, i = 1, 2, \dots, n$, we have

$$\left| \sum_{i,j} M_{ij} \langle X_i, Y_j \rangle \right| \leq K_G. \quad (4.34)$$

K_G is an absolute constant called Grothendieck’s constant. The inequality was initially proved as a fundamental tool in functional analysis. In this paper, we focus exclusively on the above matrix version in \mathbb{R}^n and consider the following suboptimal estimate derived by [16]:

$$K_G < \frac{\pi}{2\log(1 + \sqrt{2})} \leq 1.78. \tag{4.35}$$

Note that if we restrict the vectors X_i ’s and Y_i ’s to the unit sphere S^{n-1} , the inequality is still true as any inequality valid for a set also holds for its subset. Since s and t are arbitrary, the left-hand side of inequality (4.33) is indeed the $\ell_\infty \rightarrow \ell_1$ norm of matrix M , which is

$$\|M\|_{\infty \rightarrow 1} = \max_{\|x\|_\infty \leq 1} \|Mx\|_1 = \max_{s,t \in \{-1,1\}^n} s^\top M t = \max_{s,t \in \{-1,1\}^n} \left| \sum_{i,j} M_{ij} s_i t_j \right|. \tag{4.36}$$

This norm is also known as the cut norm, whose importance in algorithmic problems is well understood in the theoretical computer science community. With the ellipsope definition of SDP from Eq. (4.4) and the consequential factorization of X in Eq. (4.7), we can rewrite the theorem in the following matrix form.

Lemma 4.5. *For arbitrary matrix $M \in \mathbb{R}^{n \times n}$, we have*

$$\text{SDP}(M) \leq \max_{X \in \text{elliptope}_n} |\langle M, X \rangle| \leq K_G \|M\|_{\infty \rightarrow 1}. \tag{4.37}$$

Next, we use Bernstein’s inequality to establish a probabilistic bound on the cut norm of $A - \mathbb{E} A$ where A is the adjacency matrix of $\mathcal{G}(n, d/n)$.

Theorem 4.2 ([62, Bernstein’s Inequality]). *Let $\{X_i\}_{i=1}^n$ be independent random variables such that $\mathbb{E} X_i = 0$ and $|X_i| \leq M$ for any $i \in [n]$. Denote the average variance as*

$$\sigma^2 = \frac{1}{n} \sum_{i=1}^n \text{Var}(X_i).$$

Then for any $t \geq 0$,

$$\mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n X_i > t \right) \leq \exp \left(-\frac{nt^2/2}{\sigma^2 + Mt/3} \right). \tag{4.38}$$

Lemma 4.6. *Let A be the adjacency matrix of an ERM, $\mathcal{G}(n, d/n)$. Then, with probability at least $1 - 5^{-n+2}$,*

$$\|A - \mathbb{E} A\|_{\infty \rightarrow 1} \leq 6(1 + d)n. \tag{4.39}$$

Proof. According to the identity from inequality (4.33), we want to bound

$$\begin{aligned} \|A - \mathbb{E} A\|_{\infty \rightarrow 1} &= \max_{s,t \in \{-1,1\}^n} \sum_{i,j} (A - \mathbb{E} A)_{ij} s_i t_j \\ &= \max_{s,t \in \{-1,1\}^n} \sum_{i < j} (A - \mathbb{E} A)_{ij} (s_i t_j + s_j t_i). \end{aligned} \tag{4.40}$$

For fixed $s, t \in \{-1, 1\}^n$, denote

$$X_{ij} = (A - \mathbb{E} A)_{ij}(s_i t_j + s_j t_i), \quad 1 \leq i < j \leq n. \quad (4.41)$$

Then we have $\mathbb{E} X_{ij} = 0$, $|X_{ij}| \leq 2$ and $\text{Var}(X_{ij}) \leq 4d/n$ for any $i < j$. There are totally $n(n-1)/2$ of $\{X_{ij}\}$'s. And they are independent by the definition of ERM. So Bernstein's inequality implies

$$\mathbb{P} \left(\frac{2}{n(n-1)} \sum_{i < j} X_{ij} > t \right) \leq \exp \left(-\frac{n(n-1)t^2/4}{4d/n + 2t/3} \right). \quad (4.42)$$

Let $t = 12(1+d)/n$, which guarantees $4d/n + 2t/3 < t$. Hence,

$$\mathbb{P} \left(\sum_{i < j} X_{ij} > 6(1+d)n \right) \leq \exp(-3(n-1)). \quad (4.43)$$

Apply the union bound to all 2^{2n} possible (s, t) , we have

$$\mathbb{P} \left(\max_{s, t \in \{-1, 1\}^n} \sum_{i < j} (A - \mathbb{E} A)_{ij} (s_i t_j + s_j t_i) > 6(1+d)n \right) \leq 2^{2n} \cdot e^{-3(n-1)}. \quad (4.44)$$

We conclude the proof with the identity of $\ell_\infty \rightarrow \ell_1$ norm and the fact that the right-hand side of the above inequality is less than 5^{-n+2} . \square

Since the distribution of each entry in the matrix changes as $n \rightarrow \infty$, we now develop a slightly generalized version of the weak law of large numbers fitting for our purpose. We use the superscripts to explicitly denote dependence on n .

Lemma 4.7. *For any n , let $\{X_i^{(n)}\}_{i=1}^n$ be a collection of independent random variables. Assume there exist universal constants μ and σ such that $\mathbb{E} X_i^{(n)} \leq \mu < \infty$ and $\text{Var}(X_i^{(n)}) \leq \sigma^2 < \infty$ for any $n \in \mathbb{N}$ and $i \leq n$. If we denote the sample mean as*

$$\bar{X}^{(n)} = \frac{1}{n} \left(X_1^{(n)} + X_2^{(n)} + \dots + X_n^{(n)} \right), \quad (4.45)$$

then for any $\epsilon > 0$,

$$\mathbb{P}(\bar{X}^{(n)} \geq \mu + \epsilon) \rightarrow 0 \text{ as } n \rightarrow \infty. \quad (4.46)$$

Proof. For any $n \in \mathbb{N}$, we have

$$\begin{aligned} \text{Var}(\bar{X}^{(n)}) &= \frac{1}{n^2} \text{Var} \left(X_1^{(n)} + X_2^{(n)} + \dots + X_n^{(n)} \right) \\ &= \frac{1}{n^2} \sum_{i=1}^n \text{Var} \left(X_i^{(n)} \right) \quad (\text{by independence}) \\ &\leq \frac{\sigma^2}{n}. \quad (\text{by uniform boundedness}) \end{aligned} \quad (4.47)$$

Then Chebyshev's inequality ensures

$$\begin{aligned} \mathbb{P} (|\bar{X}^{(n)} - \mathbb{E} \bar{X}^{(n)}| \geq \epsilon) &\leq \frac{\sigma^2}{n\epsilon^2} \\ \implies \mathbb{P} \left(\bar{X}^{(n)} \geq \frac{1}{n} \sum_{i=1}^n \mathbb{E} X_i^{(n)} + \epsilon \right) &\leq \frac{\sigma^2}{n\epsilon^2} \\ \implies \mathbb{P} (\bar{X}^{(n)} \geq \mu + \epsilon) &\leq \frac{\sigma^2}{n\epsilon^2}. \end{aligned}$$

The proof is complete. \square

Remark 4.2. Compared with a standard large deviation theory, this result allows the random variables to not be identically distributed. And more importantly, the distributions can depend on n . Furthermore, the random variables associated with different n are not necessary to be independent.

Lemma 4.8. Let $G \sim \mathcal{G}(n, d/n)$, x be the labels, \mathcal{R} be the revealed indices and

$$\tilde{A} = A - \frac{d}{n} \mathbf{1}\mathbf{1}^\top$$

be its centered adjacency matrix. Define

$$B_{ij} = \begin{cases} \sum_{i,j \in \mathcal{R}} \tilde{A}_{ij} x_i x_j, & i = j = 0, \\ \sum_{k \in \mathcal{R}} x_k \tilde{A}_{kj}, & i = 0, \quad j \in [n] \setminus \mathcal{R}, \\ \sum_{k \in \mathcal{R}} x_k \tilde{A}_{ik}, & j = 0, \quad i \in [n] \setminus \mathcal{R}, \\ 0, & \text{otherwise.} \end{cases} \quad (4.48)$$

Then for any $\epsilon > 0$, with high probability,

$$\text{SDP}(B) \leq 2dm \left(1 - \frac{m}{n}\right) + (2n - m)\epsilon. \quad (4.49)$$

Proof. Notice that for any feasible X of the above optimization problem, we have $X \succeq 0$, $X_{ii} = 1$, for any $i \in [n + 1]$. So, for any $i, j \in [n + 1]$,

$$(\mathbf{e}_i \pm \mathbf{e}_j)^\top X (\mathbf{e}_i \pm \mathbf{e}_j) = 2 \pm 2X_{ij} \geq 0 \implies |X_{ij}| \leq 1. \quad (4.50)$$

Therefore,

$$\begin{aligned} \text{SDP}(B) &= \max\{\langle B, X \rangle : X \in \text{elliptope}_{n+1}\} \\ &= B_{00} + 2 \max \left\{ \sum_{j \in [n] \setminus \mathcal{R}} B_{0j} X_{0j} : X \in \text{elliptope}_{n+1} \right\} \end{aligned}$$

$$\leq B_{00} + 2 \sum_{j \in [n] \setminus \mathcal{R}} |B_{0j}|. \quad (4.51)$$

Note that $\{B_{0j} : j \in [n] \setminus \mathcal{R}\}$'s are independent random variables. Moreover, if we let B_1, B_2 be two independent binomial random variables with the same parameter $(m/2, d/n)$ and denote their difference as $Z := B_1 - B_2$, we have $B_{0j} \stackrel{d}{=} Z$ for any $j \in [n] \setminus \mathcal{R}$ with $E Z = 0$ and $\text{Var } Z \leq dm/n$.

Since $Z^2 \geq |Z|$, we have

$$E |Z| \leq E(Z^2) = \text{Var } Z \leq d \frac{m}{n}, \quad (4.52)$$

$$\text{Var } |Z| = E(Z^2) - (E |Z|)^2 \leq \text{Var } Z \leq d \frac{m}{n}. \quad (4.53)$$

Then Lemma 4.7 can be applied to

$$\bar{X}^{(n)} := \frac{\sum_{j \in [n] \setminus \mathcal{R}} |B_{0j}|}{n - m}. \quad (4.54)$$

So, for any $\epsilon > 0$, we have

$$\lim_{n \rightarrow \infty} P \left(\frac{1}{n - m} \sum_{j \in [n] \setminus \mathcal{R}} |B_{0j}| > d \frac{m}{n} + \epsilon \right) = 0. \quad (4.55)$$

Hence,

$$\sum_{j \in [n] \setminus \mathcal{R}} |B_{0j}| \leq (n - m) \left(d \frac{m}{n} + \epsilon \right)$$

with high probability.

Lemma 4.2 implies, with high probability,

$$B_{00} \leq \epsilon m. \quad (4.56)$$

Combining the above results with the union bound completes the proof. \square

Returning to the semi-supervised SDP, based on the notions from Section 4.2, we consider the following decomposition of the transformed input matrix M^{agg} with the unrevealed part and revealed part as

$$M^{\text{agg}} = M^{(\mathcal{R}^c)} + M^{(\mathcal{R})}, \quad (4.57)$$

where we define

$$M_{ij}^{(\mathcal{R})} = \begin{cases} M_{ij}^{\text{agg}}, & i = 0 \text{ or } j = 0, \\ 0, & \text{otherwise.} \end{cases} \quad (4.58)$$

To prove the main result of semi-supervised SDP, we control the $M^{(\mathcal{R}^c)}$ part by Grothendieck's inequality and bound the contribution of $M^{(\mathcal{R})}$ with the generalized law of large numbers shown above.

Proof of Theorem 2.2. Notice that Lemma 4.3 guarantees the test to succeed under the SBM. We only need to show, under ERM,

$$\text{CSDP}(\tilde{A}) < n \left[\frac{a-b}{2} - \Delta \right] \quad (\text{with high probability}). \quad (4.59)$$

According to the identity from Eq. (4.18), we have

$$\begin{aligned} \text{CSDP}(\tilde{A}) &= \text{SDP}(\tilde{A}^{\text{agg}}) \\ &= \max \{ \langle \tilde{A}^{\text{agg}}, X \rangle : X \in \text{elliptope}_n \} \\ &= \max \{ \langle \tilde{A}^{(\mathcal{R}^c)} + \tilde{A}^{(\mathcal{R})}, X \rangle : X \in \text{elliptope}_n \} \\ &\leq \text{SDP}(\tilde{A}^{(\mathcal{R}^c)}) + \text{SDP}(\tilde{A}^{(\mathcal{R})}). \end{aligned} \quad (4.60)$$

Recall that $\tilde{A}_{\mathcal{R}^c}$ is the principal submatrix of \tilde{A} obtained by removing the rows and columns associated with \mathcal{R} . By definition, we have $\text{SDP}(\tilde{A}_{\mathcal{R}^c}) = \text{SDP}(\tilde{A}^{(\mathcal{R}^c)})$. Under the null hypothesis, $\tilde{A}_{\mathcal{R}^c}$ has the same distribution as the centered adjacency matrix associated with $\mathcal{G}(n-m, (1-\rho)d/(n-m))$. Also,

$$\begin{aligned} \text{SDP}(\tilde{A}^{(\mathcal{R}^c)}) &= \text{SDP} \left(A_{\mathcal{R}^c} - \mathbb{E} A_{\mathcal{R}^c} - \frac{(1-\rho)d}{n-m} I_{n-m} \right) \\ &= \text{SDP}(A_{\mathcal{R}^c} - \mathbb{E} A_{\mathcal{R}^c}) - (1-\rho)d. \end{aligned} \quad (4.61)$$

Now we apply Grothendieck's inequality and Lemma 4.6. With probability at least $1 - 5^{-(1-\rho)n+2}$,

$$\text{SDP}(\tilde{A}^{(\mathcal{R}^c)}) \leq 6K_G [1 + (1-\rho)d](n-m) < 12(1+d)(1-\rho)n. \quad (4.62)$$

Combining above estimations and the result from Lemma 4.8 with $\epsilon = d(1-\rho)^2$, we have

$$\frac{1}{n} \text{CSDP}(\tilde{A}) \leq 14(1-\rho)(1+d) \quad (\text{with high probability}). \quad (4.63)$$

Hence, Eq. (4.59) holds with $\Delta = (a-b)/40$ and

$$\rho \geq \rho_0 = 1 - \frac{a-b}{30(1+d)}.$$

We conclude, if $m/n \geq \rho_0$,

$$\mathbb{P}_0(T(G, x_{\mathcal{R}}) = 1) \rightarrow 0, \quad n \rightarrow \infty. \quad (4.64)$$

The proof is complete. \square

Remark 4.3. If $\rho = 0$, CSDP is naturally reduced to SDP. Hence, it shares the same capability to solve the (unsupervised) community detection problem when $\text{SNR} > 1$ as stated in Theorem 1.2. Although the analysis above cannot be directly generalized to a vanishing $\rho \rightarrow 0$ situation, CSDP provides a new perspective for further study on the optimality of SDP.

Remark 4.4. We believe that it is possible to reduce the requirement on ρ through more involved analysis. Empirically, it does not need to be close to 1 for a clear improvement in the result. For instance, our simulation in the next section shows that phase transition disappears with 20% of the nodes revealed.

5 Numerical experiments

We include some simulation results below. $\rho \in [0, 1]$ is the ratio of revealed labels. Results associated with unsupervised SDPs are identified as $\rho = 0$. As discussed in Section 3, to make the comparison fair and keep the problem meaningful, all overlaps are restricted to the unrevealed labels.

Each point in Fig. 5.1 represents one realization of a SBM with $n = 1000$. The dashed line stands for the KS and information-theoretic threshold. The graphs are shared by both the unsupervised and the semi-supervised SDPs. Labels are identified by applying the k-means method to corresponding optimizers. Overlaps of unsupervised SDP essentially drop down to zero on the left-hand side. In contrast, with 20% of the labels revealed, the outcome of our constraint SDP algorithm goes down gradually as the SNR decreases and remains substantially greater than zero even when $\text{SNR} \leq 1$.

Theorem 1.1 guarantees that the upper left corner of the left image will be completely dark as $n \rightarrow \infty$. But we see semi-supervised SDPs successfully light up the area between the two reference lines, see Fig. 5.2. Moreover, when n is sufficiently large, there will be no pixel with a value of 0.

Fig. 5.3 shows color-coded entry values of optimizer X^* in different settings and suggests that the representing of underlying community structure is significantly enhanced by the semi-supervised approach, while no such structure is introduced if there should not be one.




Figure 5.1: Disappearance of the phase transition.




Figure 5.2: Overlap heatmaps of the unsupervised (left) and the semi-supervised (right) SDPs. The coordinates correspond to the model parameters a and b . The solid line represents the KS and information-theoretic threshold. The dashed line corresponds to $a = b$.

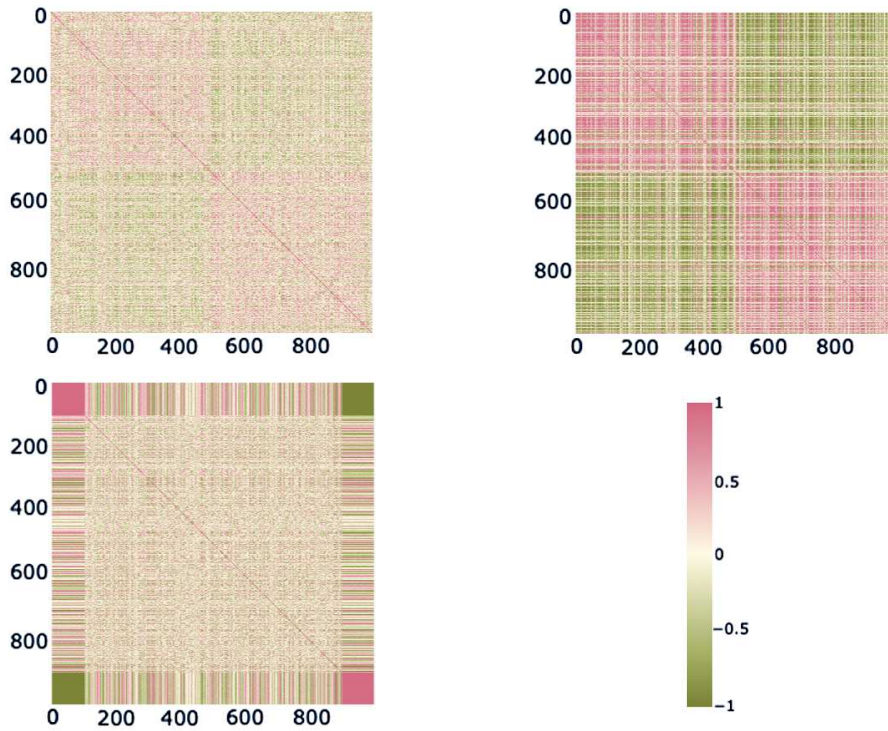


Figure 5.3: Visualization of the optimizer X^* . The upper row is concerned with one realization of the SBM $\mathcal{G}(1000, 12/1000, 5/1000)$, where the left image shows the value of optimizer for the unsupervised SDP and the right image is associated with the semi-supervised SDP with $\rho = 0.2$. The lower left image is the optimizer for one realization of the ERM of the same size with the associated average degree $d = 8.5$, indices of which are reordered such that the entries related to revealed labels are gathered in four corners. It could be understood as the situation of the null hypothesis we defined in Section 4.2.

To see how such a better representation leads to a successful test that is originally impossible, we consider the following simulations. We generate 50 independent realizations of underlying random graphs ($n = 200$) and compute their SDP values with and without the semi-supervised constraints ($\rho = 0.25$). Particularly, the parameters in Fig. 5.4 are chosen to have $\text{SNR} > 1$. The left two boxes imply that we can tell the difference between SBM and the ERM with the same average degree $d = (a + b)/2$. However, as in Fig. 5.5, the vanilla SDP gives essentially the same result since the two models become contiguous if $\text{SNR} \leq 1$. As we have proved in Theorem 2.2, our semi-supervised SDP algorithm still manages to distinguish them by bringing down the optimal value of ERM more significantly when compared to its effect on SBM, which is confirmed by the right two boxes.

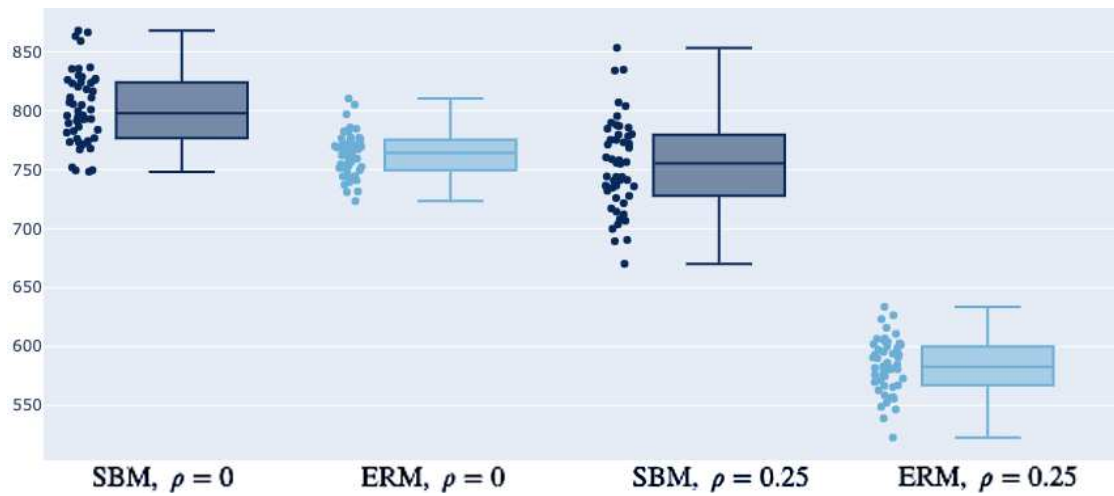


Figure 5.4: SDP optimal value, when SNR is above KS/IT, $a = 9, b = 2$ ($d = 5.5, \text{SNR} \approx 2.23$).




Figure 5.5: SDP optimal value, when SNR is below KS/IT, $a = 5, b = 2$ ($d = 3.5, \text{SNR} \approx 0.64$).

6 Conclusion

The census method comes from the combinatorial perspective, while the CSDP is inspired by convex optimization research. Both algorithms are computationally efficient. The former has no requirement on the reveal ratio. The latter one is more practical and backward compatible with the unsupervised setting. By carefully integrating the revealed information with the observed graph structure, we can not only improve the performance of clustering algorithms but also resolve initially unsolvable problems. The fundamental changes brought by the semi-supervised approach let us cross the KS threshold, information-theoretical threshold, and even the topological limitation.

Our work provides a different angle to study stochastic models of networks and semi-definite programs. In real-world situations, it is almost always the case that we will have a certain fraction of samples being understood fairly well. So, an abstract model should be able to capture the existence of such knowledge instead of being blindly restricted to the unsupervised setting. Combining the universality of revealed information and the insight derived from our census method, it is arguable that the phase transitions, although mathematically beautiful, will never be an issue in practice. Our results on CSDPs, in turn, could be used to study SDPs, e.g. prove or disprove that it can reach the phase transition threshold or the monotone-robustness threshold by a limiting process of $\rho \rightarrow 0$.

Acknowledgments

The authors acknowledge support from the National Science Foundation via Grants NSF DMS-2027248, NSF CCF-1934568, NIH Grants P41EB032840, R01HL16351, DE-SC0023490.

References

- [1] E. Abbe, Community detection and stochastic block models: Recent developments, *J. Mach. Learn. Res.*, 18(177):1–86, 2018.
- [2] E. Abbe, A. S. Bandeira, and G. Hall, Exact recovery in the stochastic block model, *IEEE Trans. Inf. Theory*, 62:471–487, 2016.
- [3] E. Abbe, E. Boix, P. Ralli, and C. Sandon, Graph powering and spectral robustness, *SIAM J. Math. Data Sci.*, 2:132–157, 2020.
- [4] E. Abbe, J. Fan, K. Wang, and Y. Zhong, Entrywise eigenvector analysis of random matrices with low expected rank, *Ann. Statist.*, 48(3):1452–1474, 2020.
- [5] E. Abbe and C. Sandon, Community detection in general stochastic block models: Fundamental limits and efficient algorithms for recovery, in: *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, IEEE, 670–688, 2015.
- [6] E. Abbe and C. Sandon, Proof of the achievability conjectures for the general stochastic block model, *Comm. Pure Appl. Math.*, 71(7):1334–1406, 2018.
- [7] N. Agarwal, A. S. Bandeira, K. Koiliaris, and A. Kolla, *Multisection in the stochastic block model using semidefinite programming*, in: *Compressed Sensing and its Applications. Applied and Numerical Harmonic Analysis*, Birkhäuser, 125–162, 2017.
- [8] F. Alizadeh, Interior point methods in semidefinite programming with applications to combinatorial optimization, *SIAM J. Optim.*, 5:13–51, 1995.

- [9] A. A. Amini and E. Levina, On semidefinite relaxations for the block model, *Ann. Statist.*, 46(1):149–179, 2018.
- [10] A. S. Bandeira, Random Laplacian matrices and convex relaxations, *Found. Comput. Math.*, 18:345–379, 2018.
- [11] J. Banks, C. Moore, J. Neeman, and P. Netrapalli, Information-theoretic thresholds for community detection in sparse networks, in: *29th Annual Conference on Learning Theory*, PMLR, 49:383–416, 2016.
- [12] P. J. Bickel and A. Chen, A nonparametric view of network models and Newman-Girvan and other modularities, *Proc. Natl. Acad. Sci. USA*, 106(50):21068–21073, 2009.
- [13] P. Bleher, J. Ruiz, and V. A. Zagrebnov, On the purity of the limiting gibbs state for the Ising model on the Bethe lattice, *J. Stat. Phys.*, 79:473–482, 1995.
- [14] B. Bollobás, S. Janson, and O. Riordan, The phase transition in inhomogeneous random graphs, *Random Structures Algorithms*, 31(1):3–122, 2007.
- [15] C. Bordenave, M. Lelarge, and L. Massoulié, Non-backtracking spectrum of random graphs: Community detection and non-regular Ramanujan graphs, in: *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, IEEE, 1347–1357, 2015.
- [16] M. Braverman, K. Makarychev, Y. Makarychev, and A. Naor, The Grothendieck constant is strictly smaller than Krivine’s bound, in: *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, IEEE, 453–462, 2011.
- [17] T. Bui, S. Chaudhuri, T. Leighton, and M. Sipser, Graph bisection algorithms with good average case behavior, in: *25th Annual Symposium on Foundations of Computer Science*, IEEE, 181–192, 1984.
- [18] J. A. Cavender, Taxonomy with confidence, *Math. Biosci.*, 40(3):271–280, 1978.
- [19] Y. Chen and J. Xu, Statistical-computational tradeoffs in planted problems and submatrix localization with a growing number of clusters and submatrices, *J. Mach. Learn. Res.*, 17(27):1–57, 2016.
- [20] H. Chernoff, A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations, *Ann. Math. Statist.*, 23:493–507, 1952.
- [21] A. Coja-Oghlan, Graph partitioning via adaptive spectral techniques, *Combin. Probab. Comput.*, 19:227–284, 2009.
- [22] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications, *Phys. Rev. E*, 84(6):066106, 2011.
- [23] A. P. Dempster, N. M. Laird, and D. B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 39(1):1–22, 1977.
- [24] S. Deng, S. Ling, and T. Strohmer, Strong consistency, graph Laplacians, and the stochastic block model, *J. Mach. Learn. Res.*, 22:117:1–117:44, 2021.
- [25] P. L. Erdős and A. Rényi, On the evolution of random graphs, *Trans. Amer. Math. Soc.*, 286:257–257, 1984.
- [26] W. Evans, C. Kenyon, Y. Peres, and L. J. Schulman, Broadcasting on trees and the Ising model, *Ann. Appl. Probab.*, 10(2):410–433, 2000.
- [27] U. Feige and E. O. Ofek, Spectral techniques applied to sparse random graphs, *Random Structures Algorithms*, 27:251–275, 2005.
- [28] C. Gao, Z. Ma, A. Y. Zhang, and H. H. Zhou, Achieving optimal misclassification proportion in stochastic block models, *J. Mach. Learn. Res.*, 18(60):1–45, 2017.
- [29] M. X. Goemans and D. P. Williamson, Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, *J. ACM*, 42(6):1115–1145, 1995.
- [30] P. K. Gopalan and D. M. Blei, Efficient discovery of overlapping communities in massive networks. *Proc. Natl. Acad. Sci. USA*, 110(36):14534–14539, 2013.
- [31] A. Grothendieck, Résumé des résultats essentiels dans la théorie des produits tensoriels topologiques et des espaces nucléaires, *Ann. Inst. Fourier*, 4:73–112, 1952.
- [32] O. Guédon and R. Vershynin, Community detection in sparse networks via Grothendieck’s inequality, *Probab. Theory Related Fields*, 165:1025–1049, 2014.
- [33] B. E. Hajek, Y. Wu, and J. Xu, Achieving exact cluster recovery threshold via semidefinite programming, *IEEE Trans. Inf. Theory*, 62:2788–2797, 2016.
- [34] E. Hartuv and R. Shamir, A clustering algorithm based on graph connectivity, *Inform. Process. Lett.*, 76(4):175–181, 2000.

- [35] Y. Higuchi, Remarks on the limiting gibbs states on a $(d+1)$ -tree, *Publ. Res. Inst. Math. Sci.*, 13:335–348, 1977.
- [36] P. Holland, K. B. Laskey, and S. Leinhardt, Stochastic blockmodels: First steps, *Soc. Networks*, 5:109–137, 1983.
- [37] A. Javanmard, A. Montanari, and F. Ricci-Tersenghi, Phase transitions in semidefinite relaxations, *Proc. Natl. Acad. Sci. USA*, 113:E2218–E2223, 2016.
- [38] S. C. Johnson, Hierarchical clustering schemes, *Psychometrika*, 32(3):241–254, 1967.
- [39] S. D. Kamvar, D. Klein, and C. D. Manning, Spectral learning, in: *Proceedings of the 18th international joint conference on Artificial intelligence*, Morgan Kaufmann Publishers Inc., 561–566, 2003.
- [40] V. Kanade, E. Mossel, and T. Schramm, Global and local information in clustering labeled block models, *IEEE Trans. Inf. Theory*, 62:5906–5917, 2014.
- [41] R. H. Keshavan, A. Montanari, and S. Oh, Matrix completion from noisy entries, *J. Mach. Learn. Res.*, 11:2057–2078, 2010.
- [42] H. Kesten and B. P. Stigum, A limit theorem for multidimensional Galton-Watson processes, *Ann. Math. Statist.*, 37(5):1211–1223, 1966.
- [43] M. Krivelevich and B. Sudakov, The largest eigenvalue of sparse random graphs, *Combin. Probab. Comput.*, 12:61–72, 2003.
- [44] F. Krzakala, C. Moore, E. Mossel, J. Neeman, A. Sly, L. Zdeborová, and P. Zhang, Spectral redemption in clustering sparse networks, *Proc. Natl. Acad. Sci. USA*, 110:20935–20940, 2013.
- [45] D.-H. Lee, Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks, in: *Workshop on challenges in representation learning, International Conference on Machine Learning*, Vol. 3, 896, 2013.
- [46] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney, Statistical properties of community structure in large social and information networks, in: *Proceedings of the 17th International Conference on World Wide Web*, ACM, 695–704, 2008.
- [47] G. Linden, B. Smith, and J. York, Amazon.com recommendations: Item-to-item collaborative filtering, *IEEE Internet Comput.*, 7(1):76–80, 2003.
- [48] C. Lu and S. Sen, Contextual stochastic block model: Sharp thresholds and contiguity, *J. Mach. Learn. Res.*, 24(54):1–34, 2020.
- [49] L. Massoulié, Community detection thresholds and the weak ramanujan property, in: *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, ACM, 694–703, 2014.
- [50] F. McSherry, Spectral partitioning of random graphs, in: *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, IEEE, 529–537, 2001.
- [51] A. Moitra, W. Perry, and A. S. Wein, How robust are reconstruction thresholds for community detection? In: *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, ACM, 828–841, 2016.
- [52] A. Montanari and S. Sen, Semidefinite programs on sparse random graphs and their application to community detection, in: *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, ACM, 814–827, 2016.
- [53] E. Mossel, Survey: Information flow on trees, *Graphs, Morphisms and Statistical Physics*, 155–170, 2001.
- [54] E. Mossel, J. Neeman, and A. Sly, Reconstruction and estimation in the planted partition model, *Probab. Theory Related Fields*, 162(3):431–461, 2015.
- [55] E. Mossel, J. Neeman, and A. Sly, A proof of the block model threshold conjecture, *Combinatorica*, 38:665–708, 2018.
- [56] E. Mossel and J. Xu, Local algorithms for block models with side information, in: *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, ACM, 71–80, 2015.
- [57] R. R. Nadakuditi and M. E. J. Newman, Graph spectra and the detectability of community structure in networks, *Phys. Rev. Lett.*, 108(18):188701, 2012.
- [58] M. E. J. Newman and M. Girvan, Finding and evaluating community structure in networks, *Phys. Rev. E*, 69(2):026113, 2004.
- [59] M. E. J. Newman, D. J. Watts, and S. H. Strogatz, Random graph models of social networks, *Proc. Natl. Acad. Sci. USA*, 99:2566–2572, 2002.
- [60] B. Padhukasahasram, Inferring ancestry from population genomic data and its applications, *Frontiers in*

- Genetics*, 5:204, 2014.
- [61] W. Perry and A. S. Wein, A semidefinite program for unbalanced multisection in the stochastic block model, in: *2017 International Conference on Sampling Theory and Applications*, IEEE, 64–67, 2017.
 - [62] H. T. H. Piaggio, Review of Introduction to Mathematical Probability, by J. V. Uspensky, *The Mathematical Gazette*, 22(249):202–204, 1938.
 - [63] Z. Qiu, E. Cho, X. Ma, and W. M. Campbell, Graph-based semi-supervised learning for natural language understanding, in: *Proceedings of the Thirteenth Workshop on Graph-Based Methods for Natural Language Processing*, Association for Computational Linguistics, 151–158, 2019.
 - [64] N. Shental, A. Bar-Hillel, T. Hertz, and D. Weinshall, Computing Gaussian mixture models with EM using equivalence constraints, in: *Proceedings of the 16th International Conference on Neural Information Processing Systems*, MIT Press, 465–472, 2003.
 - [65] J. Shi and J. Malik, Normalized cuts and image segmentation, *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(8):888–905, 2000.
 - [66] F. Spitzer, Markov random fields on an infinite tree, *Ann. Probab.*, 3(3):387–398, 1975.
 - [67] S. H. Strogatz, Exploring complex networks, *Nature*, 410(6825):268–276, 2001.
 - [68] T. Tao, *Topics in Random Matrix Theory*, AMS, 2012.
 - [69] A. B. Tsybakov, *Introduction to Nonparametric Estimation*, in: *Springer Series in Statistics*, Springer, 2009.
 - [70] V. H. Vu, Spectral norm of random matrices, *Combinatorica*, 27:721–736, 2007.
 - [71] K. L. Wagstaff, C. Cardie, S. Rogers, and S. Schrödl, Constrained k-means clustering with background knowledge, in: *Proceedings of the Eighteenth International Conference on Machine Learning*, Morgan Kaufmann Publishers Inc., 577–584, 2001.
 - [72] Q. Xie, E. H. Hovy, M.-T. Luong, and Q. V. Le, Self-training with noisy student improves ImageNet classification, in: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, 10684–10695, 2020.
 - [73] B. Yan and P. Sarkar, Covariate regularized community detection in sparse graphs, *J. Amer. Statist. Assoc.*, 116:734–745, 2016.
 - [74] S.-Y. Yun and A. Proutiere, Optimal cluster recovery in the labeled stochastic block model, in: *Proceedings of the 30th International Conference on Neural Information Processing Systems*, Curran Associates Inc., 973–981, 2016.
 - [75] P. Zhang, C. Moore, and L. Zdeborová, Phase transitions in semisupervised clustering of sparse networks, *Phys. Rev. E*, 90:052802, 2014.