Statistical Science
2021, Vol. 36, No. 1, 2–15
https://doi.org/10.1214/19-STS715

Convex Relaxation Methods for Community Detection

Xiaodong Li, Yudong Chen and Jiaming Xu

Abstract. This paper surveys recent theoretical advances in convex optimization approaches for community detection. We introduce some important theoretical techniques and results for establishing the consistency of convex community detection under various statistical models. In particular, we discuss the basic techniques based on the primal and dual analysis. We also present results that demonstrate several distinctive advantages of convex community detection, including robustness against outlier nodes, consistency under weak assortativity, and adaptivity to heterogeneous degrees.

This survey is not intended to be a complete overview of the vast literature on this fast-growing topic. Instead, we aim to provide a big picture of the remarkable recent development in this area and to make the survey accessible to a broad audience. We hope that this expository article can serve as an introductory guide for readers who are interested in using, designing, and analyzing convex relaxation methods in network analysis.

Key words and phrases: Community detection, semidefinite program, strong consistency, weak consistency, assortativity, degree correction, robustness.

1. INTRODUCTION

Convex relaxation has arisen as a powerful framework for developing computationally tractable and statistically efficient solutions to the community detection problems. Particularly in the last few years, this area has enjoyed remarkable progress: a variety of new methods have been proposed with established strong performance guarantees demonstrating their power and statistical advantages. In this expository article, we give a survey of these convex optimization approaches from the perspective of convexified maximum likelihood and discuss their major theoretical properties as well as relevant analytical tools. We hope this exposition is helpful to the readers who are interested in proposing and analyzing their own convex optimization methods for network analysis.

Convex optimization approaches for community detection can be traced back to the computer science and mathematical programming literature in the study of the planted partition problem; see, for example, Mathieu and Schudy (2010), Oymak and Hassibi (2011), Chen et al. (2014), Ames and Vavasis (2014), Ames (2014). For community detection under statistical models, various theoretical properties of convex optimization methods have been studied in depth recently; a partial list includes strong consistency[1] with a growing number of communities (Chen, Sanghavi and Xu, 2014, Chen and Xu, 2016, Cai and Li, 2015), sharp threshold under sparse networks for strong consistency (Abbe, Bandeira and Hall, 2016, Hajek, Wu and Xu, 2016a, Bandeira, 2018, Hajek, Wu and Xu, 2016b, Agarwal et al., 2015, Perry and Wein, 2015), weak consistency (Guédon and Vershynin, 2016, Fei and Chen, 2019a), nontrivial recovery (Montanari and Sen, 2015, Fan and Montanari, 2017, Mei et al., 2017), robustness against outlier nodes (Cai and Li, 2015, Moitra, Perry and Wein, 2016, Makarychev, Makarychev

Xiaodong Li is Assistant Professor, Department of Statistics, University of California, Davis, 4232 Math Sciences Bldg, 399 Crocker Lane, Davis, California 95616, USA (e-mail: xdgli@ucdavis.edu). Yudong Chen is Assistant Professor, School of Operations Research and Information Engineering, Cornell University, Ithaca, New York 14853, USA (e-mail: yudong.chen@cornell.edu). Jiaming Xu is Assistant Professor, Fuqua School of Business, Duke University, Durham, North Carolina 27708, USA (e-mail: jx77@duke.edu).

[1] An estimator of the community structure is said to achieve *strong consistency*, if it equals the underlying true community partition with probability tending to 1 as network size grows. It is said to achieve *weak consistency*, if the fraction of misclassified nodes (up to a permutation of community labels) goes to 0 with probability tending to 1. It is said to achieve *nontrivial recovery*, if the fraction of misclassified nodes (up to a permutation of community labels) is strictly better than the random guess.

and Vijayaraghavan, 2016), consistency under degree-corrected models (Chen, Li and Xu, 2018), and consistency under weak assortativity (Amini and Levina, 2018, Yan, Sarkar and Cheng, 2018).

1.1 Stochastic Block Models

In the above literature, a standard setting for deriving and analyzing the convex relaxation formulations is the so-called *stochastic block model* (SBM) (Holland et al., 1983), which we shall focus on. The SBM is a generative model for a random graph $G = (V, E)$, where $V = [n]$ is a set of n nodes and and E is the set of edges. Under SBM, the nodes are partitioned into r clusters according to the mapping $\phi : [n] \to [r]$, where the node i belongs to the cluster $\phi(i)$. Each pair of nodes i and j are connected independently with probability $B_{\phi(i)\phi(j)}$. Note that the connectivity probability only depends on the cluster memberships of the nodes. We refer to the symmetric matrix $\boldsymbol{B} := [B_{ab}]_{1\le a,b\le r}$ as the *connectivity probability matrix*. We also denote the size of the ath cluster by $n_a := |\phi^{-1}(a)|$. The general community detection problem under SBM is to estimate the unknown cluster structure ϕ given the observed graph G.

SBM is a powerful and versatile modeling tool that captures several key aspects of the community detection problem in real world networks:

- *Sparsity*: Most large-scale networks are sparse. In SBM, sparsity can be modeled by assuming that the matrix $\boldsymbol{B}$ has small entries.
- *Connectivity probabilities*: Different connectivity patterns can be captured by the matrix $\boldsymbol{B}$. Examples include *strong assortativity*, where $\min_{1\le a\le r} B_{aa} > \max_{1\le a<b\le r} B_{ab}$, and *weak assortativity* $B_{aa} - \max_{b\neq a} B_{ab} > 0, \forall a$ (Amini and Levina, 2018).
- *Number of communities*: The number of clusters r may be large and in particular is allowed to grow with the number of nodes n.
- *Unbalancedness*: SBM allows for significantly unbalanced clusters by imposing high variation in the cluster sizes $n_1, \dots, n_r$.

For expository purposes, we mostly focus on the simple (p, q)-SBM setting, where the diagonal entries of $\boldsymbol{B}$ are all p and the off-diagonal entries are all q. In other words, nodes in the same clusters are connected with probability p, whereas nodes in different clusters are connected with probability q. The results and techniques discussed in this paper can often be extended to the general SBM (or at least to the strong assortativity setting), though in some cases the extension is more difficult.

To prepare for discussion of the convex relaxation approaches, we introduce an alternative way to parameterize the cluster structure through an $n \times n$ *partition matrix* $\boldsymbol{X}$. Here the binary variable X_{ij} indicates whether or not the nodes i and j are assigned to the same cluster, that is, $X_{ij} = 1$ if $\phi(i) = \phi(j)$ and $X_{ij} = 0$ if $\phi(i) \neq \phi(j)$. Correspondingly, the observed network can also be represented by its $n \times n$ adjacent matrix $\boldsymbol{A}$, where $A_{ij} = 1$ if the nodes i and j are connected and $A_{ij} = 0$ otherwise.

1.2 Convex Relaxation Approaches

Convex relaxations for community detection can be derived in various ways, for instance through the relaxation of modularity maximization, k-means, or min k-cut. Such derivations often give rise to convex programs of similar forms. Here we focus on the perspective of convexifying maximum likelihood estimators. To the best of our knowledge, this perspective was first considered in the context of SBM by Chen, Sanghavi and Xu (2014).

Let us start with the (p, q)-SBM, and derive the likelihood function $\ell(\boldsymbol{A}|\boldsymbol{X}, p, q)$ of the community structure given the observed network data. If we denote by $p(A_{ij}|X_{ij}, p, q)$ the probability mass function for the Bernoulli random variable A_{ij}, then by definition of the model we have

$$\log p(A_{ij}|X_{ij}, p, q) = \begin{cases} \log p, & A_{ij} = X_{ij} = 1, \\ \log(1-p), & A_{ij} = 0, X_{ij} = 1, \\ \log q, & A_{ij} = 1, X_{ij} = 0, \\ \log(1-q), & A_{ij} = X_{ij} = 0. \end{cases}$$

This formula can be rewritten as a linear function of X_{ij}:

$$\begin{aligned} &\log p(A_{ij}|X_{ij}, p, q) \\ &= \left[\left(\log\frac{p}{1-p} - \log\frac{q}{1-q}\right)A_{ij} - \log\frac{1-q}{1-p}\right]X_{ij} \\ &\quad + \left[A_{ij}\log q + (1-A_{ij})\log(1-q)\right]. \end{aligned}$$

Since all pairs of nodes are connected independently, that is, the A_{ij}'s are independent across $1 \le i < j \le n$, the log-likelihood function given the observed adjacency matrix $\boldsymbol{A}$ has the summation form

$$\ell(\boldsymbol{A}|\boldsymbol{X}, p, q) = \sum_{1\le i<j\le n} \log p(A_{ij}|X_{ij}, p, q).$$

Our goal is to estimate the true matrix of membership, denoted by $\boldsymbol{X}^\star$, using the maximum likelihood approach. If p and q are fixed and given, then maximizing the log likelihood $\ell(\boldsymbol{A}|\boldsymbol{X}, p, q)$ is equivalent to the maximization of

$$\sum_{1\le i<j\le n}\left[\left(\log\frac{p}{1-p} - \log\frac{q}{1-q}\right)A_{ij} - \log\frac{1-q}{1-p}\right]X_{ij}$$

over all possible cluster matrices $\boldsymbol{X}$ with any number of clusters r. In real world networks, individuals within the same communities are often more likely to be connected

than those across different communities, which corresponds to the setting with $p > q$ and hence $\log \frac{p}{1-p} - \log \frac{q}{1-q} > 0$. In this case, the above optimization problem is equivalent to

$$\max_{\boldsymbol{X}} \sum_{1 \le i < j \le n} (A_{ij} - \lambda) X_{ij}, \tag{1.1}$$

where

$$\begin{aligned} \lambda &\equiv \lambda(p, q) \\ &= \frac{\log(1-q) - \log(1-p)}{\log p - \log q + \log(1-q) - \log(1-p)}. \end{aligned} \tag{1.2}$$

If p and q are unknown, we simply treat λ as a tuning parameter.[2]

To derive a convex relaxation, it is more convenient to write the problem (1.1) in a matrix form. Since both $\boldsymbol{A}$ and $\boldsymbol{X}$ are symmetric and the diagonal entries of $\boldsymbol{A}$ are all 0, the problem (1.1) can be rewritten as

$$\begin{aligned} \max_{\boldsymbol{X}} \quad & \langle \boldsymbol{X}, \boldsymbol{A} - \lambda \boldsymbol{J}_n \rangle \\ \text{subject to } & \boldsymbol{X} \text{ is a partition matrix,} \end{aligned} \tag{1.3}$$

where $\boldsymbol{J}_n$ is the $n \times n$ matrix with all entries equal to 1, and $\langle \boldsymbol{A}, \boldsymbol{B} \rangle := \operatorname{trace}(\boldsymbol{A}^\top \boldsymbol{B})$ denotes the trace inner product. The problem (1.3) is nonconvex due to its constraint. To convexify this program, let us investigate the properties of a partition matrix $\boldsymbol{X}$. By definition we see that $\boldsymbol{X}$ must have form

$$\boldsymbol{X} = \boldsymbol{P} \begin{bmatrix} \boldsymbol{J}_{n_1} & & \\ & \ddots & \\ & & \boldsymbol{J}_{n_r} \end{bmatrix} \boldsymbol{P}^\top, \tag{1.4}$$

where $\boldsymbol{P}$ is some permutation matrix and the number of communities r is unknown. The set of all matrices $\boldsymbol{X}$ of this form is of course nonconvex. The key observation is that any such $\boldsymbol{X}$ satisfies several convex constraints such as (i) all entries of $\boldsymbol{X}$ are nonnegative, (ii) all diagonal entries of $\boldsymbol{X}$ are 1, and (iii) $\boldsymbol{X}$ is positive semidefinite. Various convex relaxations can hence be derived by replacing the constraint in (1.3) with these convex constraints (or variants thereof).

One typical relaxation obtained in this way is

$$\begin{aligned} \max \quad & \langle \boldsymbol{X}, \boldsymbol{A} - \lambda \boldsymbol{J}_n \rangle \\ \text{subject to } & \boldsymbol{X} \succeq \boldsymbol{0}, \boldsymbol{X} \ge \boldsymbol{0}, \\ & X_{ii} = 1 \text{ for } 1 \le i \le n, \end{aligned} \tag{1.5}$$

where $\boldsymbol{X} \ge \boldsymbol{0}$ means $X_{ij} \ge 0$ for $1 \le i, j \le n$. In Section 2 we present results on when (1.5) recovers the ground truth clusters. These results highlight the following attractive properties of the formulation (1.5):

- The communities are allowed to be significantly unbalanced;
- The number of communities r may grow as n increases;
- Although (1.5) is derived from the (p, q)-model, it is applicable to a more general connectivity probability matrix $\boldsymbol{B}$ with strong assortativity;
- The knowledge of r is not required in (1.5);
- There is only one tuning parameter, λ.

Convex relaxation methods of similar forms to (1.5) have been extensively studied in the literature; see, for example, Ames (2014), Chen and Xu (2016), Cai and Li (2015), Amini and Levina (2018), Guédon and Vershynin (2016), Hajek, Wu and Xu (2016b), Agarwal et al. (2015), Perry and Wein (2015), Chen, Li and Xu (2018), Fei and Chen (2019a). Moreover, other convex relaxation formulations, such as those for the simple two-community setting considered by Abbe, Bandeira and Hall (2016), Hajek, Wu and Xu (2016a), Bandeira (2018), Montanari and Sen (2015), Javanmard, Montanari and Ricci-Tersenghi (2015), are also intrinsically related to (1.5).

The subsequent sections are devoted to understanding the statistical performance of the convex relaxation (1.5) and some of its variants. In particular, we present theoretical results on when these relaxations yield, with high probability, a good estimate of the ground truth communities of SBM in terms of strong and weak consistency. Particular focus is given to elucidating several important techniques, including dual and primal analysis, for establishing such results.

In the sequel, we use C_0 etc. to denote a positive and sufficiently large absolute constant. By "with high probability" we mean with probability at least $1 - 10n^{-10}$.

2. STRONG CONSISTENCY VIA DUAL CERTIFICATION ANALYSIS

In this section we introduce a dual certification technique for establishing strong consistency of the convex relaxation (1.5). This technique involves constructing a set of dual variables that *certify* the optimality of the desired primal solution $\boldsymbol{X}^\star$ via the KKT condition. This type of arguments are widely employed in the literature to study exact recovery under SBM using convex optimization; see, for example, Chen, Sanghavi and Xu (2014), Cai and Li (2015), Ames (2014).

For the ease of presentation, we present the dual certification analysis under the two-cluster (p, q)-model with $n_1 = O(n)$ and $n_2 = O(n)$. Without loss of generality we assume that the permutation matrix corresponding to the ground truth communities is the identity, that is, $\boldsymbol{P} = \boldsymbol{I}$. Under this assumption, the adjacency matrix $\boldsymbol{A}$ and the true partition matrix $\boldsymbol{X}^\star$ as defined in (1.4) can be written as

$$\boldsymbol{A} = \begin{bmatrix} \boldsymbol{A}_{11} & \boldsymbol{A}_{12} \\ \boldsymbol{A}_{12}^\top & \boldsymbol{A}_{22} \end{bmatrix} \quad \text{and} \quad \boldsymbol{X}^\star = \begin{bmatrix} \boldsymbol{J}_{n_1} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{J}_{n_2} \end{bmatrix}, \tag{2.1}$$

[2] An alternative approach is profile likelihood maximization, that is, we maximize $\ell(\boldsymbol{A}|\boldsymbol{X}, \widehat{p}, \widehat{q})$ with $(\widehat{p}(\boldsymbol{X}), \widehat{q}(\boldsymbol{X})) := \arg\max_{p,q} \ell(\boldsymbol{A}|\boldsymbol{X}, p, q)$. However, the framework of profile likelihood will result in a highly nonlinear function of $\boldsymbol{X}$ to maximize.

where the entries of A_{11} and A_{22} are Bernoulli random variables with parameter p and those A_{12} are Bernoulli with parameter q.

Our goal is to prove that the convex relaxation (1.5) recovers the true $X^\star$ as an optimal solution. We do so by showing that $X^\star$ satisfies the KKT condition, a sufficient condition for optimality (cf. Boyd and Vandenberghe, 2004, Section 5.5.3). The KKT condition stipulates the existence of some Λ, Φ, and β satisfying:

- Stationarity: $\lambda J_n - A = \Lambda + \Phi + \operatorname{diag}(\beta)$,
- Complementary slackness: $\langle X^\star, \Lambda\rangle = \langle X^\star, \Phi\rangle = 0$,
- Dual feasibility: $\Phi \geq \mathbf{0}$, $\Lambda \succeq \mathbf{0}$,

where $\Lambda, \Phi \in \mathbb{R}^{n\times n}$ and $\beta \in \mathbb{R}^n$ are the dual variables associated with the constraints of (1.5) in order. We write these variables in blocks as in (2.1):

$$\Lambda = \begin{bmatrix} \Lambda_{11} & \Lambda_{12} \\ \Lambda_{12}^\top & \Lambda_{22} \end{bmatrix}, \quad \Phi = \begin{bmatrix} \Phi_{11} & \Phi_{12} \\ \Phi_{12}^\top & \Phi_{22} \end{bmatrix} \quad \text{and} \quad \beta = \begin{bmatrix} \beta_1 \\ \beta_2 \end{bmatrix}.$$

Our strategy for finding Λ, Φ and β involves two steps: we first construct these variables so that the equality constraints in the KKT conditions are satisfied deterministically, and then verify that the inequality constraints in KKT hold with high probability.

Step 1(a): Constructing Λ_{11}, Λ_{22}, Λ_{12} and β

We begin by rewriting the stationarity condition in blocks:

$$\begin{cases} -A_{11} + \lambda J_{n_1} - \Lambda_{11} - \operatorname{diag}(\beta_1) = 0, \\ -A_{22} + \lambda J_{n_2} - \Lambda_{22} - \operatorname{diag}(\beta_2) = 0, \\ -A_{12} + \lambda J_{n_1,n_2} - \Lambda_{12} - \Phi_{12} = 0, \end{cases}$$

where J_{n_1,n_2} is the n_1-by-n_2 all-one matrix. We next observe that the complementary slackness and dual feasibility constraints are implied by the following equalities:

$$\Lambda_{11}\mathbf{1}_{n_1} = \mathbf{0}, \quad \Lambda_{22}\mathbf{1}_{n_2} = \mathbf{0}, \quad \Lambda_{12}\mathbf{1}_{n_2} = \mathbf{0},$$
$$\Lambda_{12}^\top \mathbf{1}_{n_1} = \mathbf{0}, \quad \Phi_{11} = \mathbf{0}, \quad \Phi_{22} = \mathbf{0}.$$

Combining the first two equalities with the stationarity condition, we see that

$$\beta_1 = -A_{11}\mathbf{1}_{n_1} + n_1\lambda J_{n_1}, \quad \beta_2 = -A_{22}\mathbf{1}_{n_2} + n_2\lambda J_{n_2},$$

and that Λ can be expressed in terms of Φ_{12} as

$$\Lambda = \begin{bmatrix} \operatorname{diag}(A_{11}\mathbf{1}_{n_1}) - n_1\lambda I_{n_1} - A_{11} + \lambda J_{n_1} & -\Phi_{12} - A_{12} + \lambda J_{n_1,n_2} \\ -\Phi_{12}^\top - A_{12}^\top + \lambda J_{n_2,n_1} & \operatorname{diag}(A_{22}\mathbf{1}_{n_2}) - n_2\lambda I_{n_2} - A_{22} + \lambda J_{n_2} \end{bmatrix},$$

where I_{n_1} is the $n_1 \times n_1$ identity matrix. Consequently, the equalities $\Lambda_{12}\mathbf{1}_{n_2} = \mathbf{0}$ and $\Lambda_{12}^\top \mathbf{1}_{n_1} = \mathbf{0}$ are equivalent to

$$\begin{aligned} \Phi_{12}\mathbf{1}_{n_2} &= (\lambda J_{n_1,n_2} - A_{12})\mathbf{1}_{n_2}, \\ \mathbf{1}_{n_1}^\top \Phi_{12} &= \mathbf{1}_{n_1}^\top (\lambda J_{n_1,n_2} - A_{12}). \end{aligned} \tag{2.2}$$

Therefore, to establish the KKT condition, it suffices to construct a matrix Φ_{12} that satisfies the equality (2.2) as well as the inequality constraints $\Phi_{12} \geq \mathbf{0}$ and $\Lambda \succeq \mathbf{0}$.

Step 1(b): Constructing Φ_{12}

Such construction is often an art. Here we propose as a candidate using the following rank-2 matrix $\Phi_{12} = y\mathbf{1}_{n_2}^\top + \mathbf{1}_{n_1} z^\top$ for some vectors $y \in \mathbb{R}^{n_1}$ and $z \in \mathbb{R}^{n_2}$. We solve for y and z by plugging this candidate into the equality (2.2), which yields that

$$\Phi_{12} = -\frac{1}{n_2} A_{12} J_{n_2} - \frac{1}{n_1} J_{n_1} A_{12} + \frac{1}{n_1 n_2} J_{n_1} A_{12} J_{n_2} + \lambda J_{n_1,n_2}.$$

It remains to verify that the inequality constraints $\Phi_{12} \geq \mathbf{0}$ and $\Lambda \succeq \mathbf{0}$ hold with high probability under appropriate choices of the tuning parameter λ.

Step 2(a): Verifying $\Phi_{12} > \mathbf{0}$

Note that $\mathbb{E}\Phi_{12} = (\lambda - q) J_{n_1,n_2} > \mathbf{0}$ provided that $\lambda > q$. To translate this expectation bound to a high probability bound, we make use of the Chernoff's inequality and the assumption that $n_1 = O(n)$ and $n_2 = O(n)$, which guarantee that $\Phi_{12} > \mathbf{0}$ with high probability as long as $\lambda > q + O((\log n)/n + \sqrt{(q\log n)/n})$.

Step 2(b): Verifying $\Lambda \succeq \mathbf{0}$

Here we use an important consequence of the rank-2 structure of Φ_{12}: the matrix Λ is the projection of another matrix onto a lower dimensional space. To see this, we define the projective matrix

$$\Pi := \begin{bmatrix} \frac{1}{n_1} J_{n_1} & \mathbf{0} \\ \mathbf{0} & \frac{1}{n_2} J_{n_2} \end{bmatrix}$$

and let

$$\widetilde{\Lambda} := \begin{bmatrix} \operatorname{diag}(A_{11}\mathbf{1}_{n_1}) - n_1\lambda I_{n_1} - A_{11} + pJ_{n_1} & -A_{12} + qJ_{n_1,n_2} \\ -A_{12}^\top + qJ_{n_2,n_1} & \operatorname{diag}(A_{22}\mathbf{1}_{n_2}) - n_2\lambda I_{n_2} - A_{22} + pJ_{n_2} \end{bmatrix}.$$

Using the fact that the rank-2 matrix Φ_{12} satisfies the equality $(I_{n_1} - J_{n_1}/n)\Phi_{12}(I_{n_2} - J_{n_2}/n) = \mathbf{0}$, we obtain the identity $\Lambda = (I_n - \Pi)\widetilde{\Lambda}(I_n - \Pi)$. Therefore, the condition $\Lambda \succeq \mathbf{0}$ is implied by $\widetilde{\Lambda} \succ \mathbf{0}$. To proceed, we decompose $\widetilde{\Lambda}$ into two terms as $\widetilde{\Lambda} = (\mathbb{E}[A] - A) + (A - \mathbb{E}[A] +$

$\widetilde{\boldsymbol{\Lambda}}$). The first term can be interpret as the "noise," that is, deviation of the adjacency matrix $\boldsymbol{A}$ from its expectation. A known result in random matrix theory guarantees that $\|\mathbb{E}[\boldsymbol{A}] - \boldsymbol{A}\| \le C_0\sqrt{np}$ provided $p \ge C_0 \log n/n$, where $\|\cdot\|$ denotes the spectral norm; see, for example, Feige and Ofek (2005), Vu (2018), Hajek, Wu and Xu (2016a) and Bandeira and van Handel (2016). On the other hand, one can verify that the second term is a diagonal matrix of the form

$$\begin{aligned}&\boldsymbol{A} - \mathbb{E}[\boldsymbol{A}] + \widetilde{\boldsymbol{\Lambda}}\\ &\quad= \begin{bmatrix} \operatorname{diag}(\boldsymbol{A}_{11}\mathbf{1}_{n_1}) - (n_1\lambda - p)\boldsymbol{I}_{n_1} & \\ & \mathbf{0} \end{bmatrix}\\ &\qquad\qquad \begin{bmatrix} \mathbf{0} & \\ & \operatorname{diag}(\boldsymbol{A}_{22}\mathbf{1}_{n_2}) - (n_2\lambda - p)\boldsymbol{I}_{n_2} \end{bmatrix}.\end{aligned}$$

Each diagonal entry above is (roughly) equal to the number of neighbors of a node in its own cluster minus λ times the cluster size, and hence can be interpreted as the "signal" strength, which is bounded from below thanks to the Chernoff's inequality. In particular, letting $n_{\min} := \min(n_1, n_2)$, we obtain that

$$\boldsymbol{A} - \mathbb{E}[\boldsymbol{A}] + \widetilde{\boldsymbol{\Lambda}} \succeq (n_{\min}p - 2\sqrt{n_{\min}p\log n} - n_{\min}\lambda)\boldsymbol{I}_n$$

with high probability. We conclude that $\widetilde{\boldsymbol{\Lambda}} \succ \mathbf{0}$ as long as $\lambda < p - O(\sqrt{p\log n/n})$.

The KKT condition in general does not guarantee $\boldsymbol{X}$ to be the *unique* solution to (1.5), establishing which requires additional steps. We refer the readers to Cai and Li (2015) for the details of such steps using Cauchy's interlacing theorem. Summarizing the above arguments, we obtain the following result:

THEOREM 2.1. *Under the two-cluster (p, q)-model with $n_1 = O(n)$ and $n_2 = O(n)$, if $p - q > C_0 \times \sqrt{p\log n/n}$, then there exists a tuning parameter λ such that the ground truth partition matrix $\boldsymbol{X}^\star$ is the unique solution to the convex relaxation* (1.5) *with high probability.*

It is worth highlighting that the above argument is quite generic, and can be extended in a straightforward manner to more general and realistic setups of SBM, such as those with significantly unbalanced clusters, a fast growing number of communities and general connectivity probability matrix satisfying strong assortativity; see, for example, Chen, Sanghavi and Xu (2014).

3. SHARP THRESHOLD FOR STRONG CONSISTENCY

A line of recent work on community detection studies the necessary and sufficient conditions on p and q with *sharp* constants for achieving strong consistency using convex relaxations. This line of work, initiated by Abbe, Bandeira and Hall (2016) and followed by Hajek, Wu and Xu (2016a), Bandeira (2018), Hajek, Wu and Xu (2016b), Agarwal et al. (2015), Perry and Wein (2015), has achieved remarkable progress. Some additional assumptions are usually essential in deriving conditions with sharp constants, such as relatively balanced clusters and fixed (or slowly growing) number of clusters. Nevertheless, these results demonstrate the statistical power of the convex relaxation approaches, and also highlight the theoretical and practical importance of certain constraints (e.g., row-sum constraints) used in convex relaxations. In this section, we briefly outline the main results in this line, with emphasis on the relevant theoretical tools developed therein.

As suggested in Amini and Levina (2018), Hajek, Wu and Xu (2016b), Agarwal et al. (2015), we here consider an important variation of (1.5) which arises when assuming that all communities have equal size n/r and that r is known. In this case, the partition matrix $\boldsymbol{X}$ in (1.4) satisfies another convex constraint: $\boldsymbol{X}\mathbf{1}_n = \frac{n}{r}\mathbf{1}_n$, where $\mathbf{1}_n$ is the all one vector in $\mathbb{R}^n$. Adding this constraint to (1.5) leads to the convex relaxation

$$\begin{aligned}&\max \quad \langle \boldsymbol{X}, \boldsymbol{A}\rangle\\ &\text{subject to } \boldsymbol{X} \succeq \mathbf{0}, \boldsymbol{X} \ge \mathbf{0},\\ &\qquad \boldsymbol{X}\mathbf{1}_n = \frac{n}{r}\mathbf{1}_n, X_{ii} = 1 \text{ for } 1 \le i \le n.\end{aligned} \tag{3.1}$$

The term $\lambda\boldsymbol{J}_n$ in (1.5) disappears because $\langle \boldsymbol{X}, \boldsymbol{J}_n\rangle$ is a constant under the row-sum constraint $\boldsymbol{X}\mathbf{1}_n = \frac{n}{r}\mathbf{1}_n$.

We consider the two-cluster (p, q)-SBM setting as in Section 2. For ease of exposition, we further assume that the two clusters are of equal size ($n_1 = n_2 = n/2$) and that $p = \frac{a\log(n)}{n}$ and $q = \frac{b\log(n)}{n}$, where $a > b > 0$ are two fixed constants. In this setting, the information-theoretically limit (i.e., sufficient and necessary condition) for exact recovery is known to be $\sqrt{a} - \sqrt{b} > \sqrt{2}$ (Abbe, Bandeira and Hall, 2016, Mossel, Neeman and Sly, 2016). It turns out that the convex relaxation (3.1) achieves this information-theoretic limit. In particular, if $\sqrt{a} - \sqrt{b} > \sqrt{2}$, then the convex relaxation (3.1) with row-sum constraint recovers the true partition matrix $\boldsymbol{X}^*$ in (2.1) as the unique optimum with probability $1 - n^{-\Omega(1)}$. Below we present the proof of this result, which makes use of the dual certification analysis argument.

The dual certificate construction is similar to that introduced in the previous section. A key difference is that the extra row-sum constraint induces a new dual variable $\boldsymbol{\mu} \in \mathbb{R}^n$, allowing for extra freedom in the construction. The KKT condition of the relaxation (3.1) with primal solution $\boldsymbol{X}^*$ defined in (2.1) reads:

- Stationarity: $-\boldsymbol{A} - \boldsymbol{\Lambda} - \boldsymbol{\Phi} - \operatorname{diag}(\boldsymbol{\beta}) - \boldsymbol{\mu}\mathbf{1}_n^\top - \mathbf{1}_n\boldsymbol{\mu}^\top = \mathbf{0}$,
- Complementary slackness: $\langle \boldsymbol{X}^\star, \boldsymbol{\Lambda}\rangle = \langle \boldsymbol{X}^\star, \boldsymbol{\Phi}\rangle = 0$,
- Dual feasibility: $\boldsymbol{\Phi} \ge \mathbf{0}$, $\boldsymbol{\Lambda} \succeq \mathbf{0}$.

Under the balanced-cluster assumption $n_1 = n_2 = n/2$, we can simply let $\boldsymbol{\Phi} = \mathbf{0}$. Similarly to the argument in Section 2, we construct the other dual variables as

$$\begin{aligned}
\boldsymbol{\mu}_1 &= -\frac{1}{n_2}\boldsymbol{A}_{12}\mathbf{1}_{n_2} + \frac{\mathbf{1}_{n_1}^\top \boldsymbol{A}_{12}\mathbf{1}_{n_2}}{2n_1 n_2}\mathbf{1}_{n_1},\\
\boldsymbol{\mu}_2 &= -\frac{1}{n_1}\boldsymbol{A}_{12}^\top\mathbf{1}_{n_1} + \frac{\mathbf{1}_{n_1}^\top \boldsymbol{A}_{12}\mathbf{1}_{n_2}}{2n_1 n_2}\mathbf{1}_{n_2},\\
\boldsymbol{\beta_1} &= -\boldsymbol{A}_{11}\mathbf{1}_{n_1} + \boldsymbol{A}_{12}\mathbf{1}_{n_2},\\
\boldsymbol{\beta_2} &= -\boldsymbol{A}_{22}\mathbf{1}_{n_2} + \boldsymbol{A}_{12}^\top\mathbf{1}_{n_1}
\end{aligned}$$

and

$$\begin{aligned}
\boldsymbol{\Lambda} &= -\boldsymbol{A} - \boldsymbol{\Phi} - \operatorname{diag}(\boldsymbol{\beta}) - \boldsymbol{\mu}\mathbf{1}_n^\top - \mathbf{1}_n\boldsymbol{\mu}^\top\\
&= (\boldsymbol{I} - \boldsymbol{\Pi})\big((\mathbb{E}[\boldsymbol{A}] - \boldsymbol{A}) + p\boldsymbol{I} - \operatorname{diag}(\boldsymbol{\beta})\big)\\
&\quad\times(\boldsymbol{I} - \boldsymbol{\Pi}),
\end{aligned}$$

where the projection $\boldsymbol{\Pi}$ is defined in Section 2. To verify that $\boldsymbol{\Lambda} \succeq \mathbf{0}$, it suffices to show that $(\mathbb{E}[\boldsymbol{A}] - \boldsymbol{A}) + p\boldsymbol{I} - \operatorname{diag}(\boldsymbol{\beta}) \succ \mathbf{0}$. This condition can be established by combining the following two facts. First, by a tight Chernoff's inequality (Hajek, Wu and Xu, 2016a, Lemma 1), the inequality $\min_{i\in[n]}(-\beta_i) \geq \frac{\log n}{\log\log n}$ holds with probability at least $1 - n^{1-(\sqrt{a}-\sqrt{b})^2/2+o(1)}$, which is $1 - n^{-\Omega(1)}$ as long as $\sqrt{a} - \sqrt{b} > \sqrt{2}$. Second, as shown in Section 2, $\|\boldsymbol{A} - \mathbb{E}[\boldsymbol{A}]\| \leq c'\sqrt{\log n}$ with high probability for a positive constant c'. With an addition step proving the uniqueness of the optimal solution, which we skip here, we obtain the following result in the line of Abbe, Bandeira and Hall (2016), Hajek, Wu and Xu (2016a), Bandeira (2018), Hajek, Wu and Xu (2016b), Agarwal et al. (2015):

THEOREM 3.1. *Under the two-cluster (p, q)-model with equal cluster sizes, suppose that $p = \frac{a\log(n)}{n}$ and $q = \frac{b\log(n)}{n}$. If $\sqrt{a} - \sqrt{b} > \sqrt{2}$, then with probability $1 - n^{-\Omega(1)}$, the ground truth partition matrix $\boldsymbol{X}^\star$ is the unique solution to the convex relaxation* (3.1).

In the above proof with two clusters, we simply choose $\boldsymbol{\Phi} = \mathbf{0}$. This simple choice not longer works for $r > 2$ equal-sized clusters. In that case, denoting by $\boldsymbol{\Phi}_{k,\ell}$ the block in $\boldsymbol{\Phi}$ corresponding to $\boldsymbol{A}_{k,\ell}$, we may choose $\boldsymbol{\Phi}_{k,\ell} = \boldsymbol{y}_{k,\ell}\mathbf{1}_{n_\ell}^\top + \mathbf{1}_{n_k}\boldsymbol{z}_{k,\ell}^\top$ for some vectors $\boldsymbol{y}_{k,\ell}$ and $\boldsymbol{z}_{k,\ell}$. The detailed argument can be found in Hajek, Wu and Xu (2016b).

We note that the dual variable $-\beta_i$ above corresponds to the number of neighbors node i has in its own cluster minus the number of its neighbors in the other cluster. The variable $-\beta_i$ is closely related to the information-theoretic lower bounds of strong consistency. In particular, it was shown in (Abbe, Bandeira and Hall, 2016) that if $\sqrt{a} - \sqrt{b} < \sqrt{2}$, then with probability $1 - o(1)$, there exists a pair of nodes i and j from different clusters such that $-\beta_i < -1$ and $-\beta_j < -1$, which further implies that the maximum likelihood estimator cannot coincide with the ground truth $\boldsymbol{X}^\star$ and hence the impossibility of achieving strong consistency.

Other convex relaxation formulations

It is noteworthy that besides the formulation (3.1) with row-wise constraints, there exist other convex optimization methods that achieve the sharp thresholds for strong consistency. For example, in the two-community setting, one may consider the following formulation:

$$\begin{aligned}
&\max \qquad \langle \boldsymbol{Y}, \boldsymbol{A} - \lambda\boldsymbol{J}_n\rangle\\
&\text{subject to } \boldsymbol{Y} \succeq \mathbf{0},\ Y_{ii} = 1 \text{ for } 1 \leq i \leq n,
\end{aligned} \tag{3.2}$$

which is a variant of the well-known SDP relaxation for MaxCut (Goemans and Williamson, 1995). With some appropriate λ, the formulation (3.2) is shown to recover the ground truth centered partition matrix $\boldsymbol{Y}^* = 2\boldsymbol{X}^* - \boldsymbol{J}_n$ as the unique optimum with probability $1 - o(1)$, within the information-theoretically feasible range of SBM parameters; see Abbe, Bandeira and Hall (2016), Hajek, Wu and Xu (2016a), Bandeira (2018) for the case of equal cluster sizes and Hajek, Wu and Xu (2016b) for the unbalanced case. Extensions to the more general setting with multiple clusters of unequal sizes can be found in Perry and Wein (2015).

4. PROJECTIVE MATRIX BASED CONVEX OPTIMIZATION AND WEAK ASSORTATIVITY

In Section 2, we present strong consistency results for the relaxation (1.5) under the (p, q)-SBM, which can be extended to more general connectivity probability matrices $\boldsymbol{B}$. However, the strong consistency of (1.5) requires that the tuning parameter λ is between minimum within-community edge density $\min_{1\leq a\leq r} B_{aa}$ and the maximum cross-community edge density $\max_{1\leq a<b\leq r} B_{ab}$. This means that one must have $\min_{1\leq a\leq r} B_{aa} > \max_{1\leq a<b\leq r} B_{ab}$, namely, the strong assortativity of $\boldsymbol{B}$. The notion of strong assortativity, coined in Amini and Levina (2018), is a common assumption in the literature of community detection under the SBM; see, for example, Rohe, Chatterjee and Yu (2011) and Chaudhuri, Chung and Tsiatas (2012).

Strong assortativity is sometimes too restrictive. To address this issue, Amini and Levina (2018) consider the SBM under the weak assortativity assumption $B_{aa} > \max_{b\neq a} B_{ab}$ for any $a = 1, \ldots, r$, and show that the convex relaxation (3.1) achieves strong consistency under certain conditions of the model parameter. The dual analysis used to prove this result is similar to that in Section 3, although they do not focus on identifying sharp thresholds. However, this result relies on the strong assumption that all cluster have equal sizes. Indeed, as discussed in Section 3, this assumption is necessary for the true partition matrix $\boldsymbol{X}^*$ to be feasible to the formulation (3.1).

It is desirable to develop a convex relaxation that does not require this unrealistic assumption of equal-cluster-size while still guaranteeing strong consistency under weak assortativity. To this end, we keep a version of the row-sum constraint $\boldsymbol{X}\mathbf{1}_n = \frac{n}{r}\mathbf{1}_n$, which appears essential for the weak assortativity setting, but try to remove its explicit diagonal constraints in (3.1) with a trace constraint, that is,

$$\max_{\boldsymbol{X}} \quad \langle \boldsymbol{X}, \boldsymbol{A}\rangle$$
$$\text{subject to } \boldsymbol{X} \succeq \mathbf{0},\ \boldsymbol{X} \geq \mathbf{0},\ \boldsymbol{X}\mathbf{1}_n = \frac{n}{r}\mathbf{1}_n,\ \operatorname{trace}(\boldsymbol{X}) = n.$$

With a change of variable $\boldsymbol{Z} = \frac{r}{n}\boldsymbol{X}$, we obtain the following relaxation:

$$\begin{aligned} \max_{\boldsymbol{Z}} \quad & \langle \boldsymbol{Z}, \boldsymbol{A}\rangle \\ \text{subject to } & \boldsymbol{Z} \succeq \mathbf{0},\ \boldsymbol{Z} \geq \mathbf{0},\ \ \boldsymbol{Z}\mathbf{1}_n = \mathbf{1}_n, \\ & \operatorname{trace}(\boldsymbol{Z}) = r. \end{aligned} \tag{4.1}$$

Note that the above relaxation has no explicit dependence on the cluster sizes. The relaxation (4.1) appears to be less studied in the community detection literature compared to (1.5) and (3.1), but it in fact coincides with the well-known Peng-Wei relaxation (Peng and Wei, 2007) for K-means in the literature of clustering in Euclidean space. With this relaxation, the goal is to recover the the ground truth *projective matrix*

$$\boldsymbol{Z}^\star = \boldsymbol{P}\begin{bmatrix} \frac{1}{n_1}\boldsymbol{J}_{n_1} & & \\ & \ddots & \\ & & \frac{1}{n_r}\boldsymbol{J}_{n_r} \end{bmatrix}\boldsymbol{P}^\top, \tag{4.2}$$

which is a symmetric matrix with all eigenvalues equal to either one or zero, and can be viewed as a normalized version of the true partition matrix $\boldsymbol{X}^\star$ given in (1.4). The work in that literature shows that under certain mixture models satisfying appropriate affinity or separation conditions, the $\boldsymbol{Z}^\star$ is the unique solution to the relaxation (4.1); see, for example, Iguchi et al. (2015, 2017), Li et al. (2017) and Fei and Chen (2018).

The relaxation (4.1) was recently analyzed in the context of community detection by Yan, Sarkar and Cheng (2018). They show that this relaxation enjoys strong consistency (in the sense of recovering $\boldsymbol{Z}^\star$ as the unique solution with high probability) under SBM with weak assortativity and unbalanced communities. This result substantially extends the result in Amini and Levina (2018).

The proof of the above result again uses dual certification analysis. The primal feasibility of the projective matrix $\boldsymbol{Z}^\star$ in (4.2) to the relaxation (4.1) can be directly verified. To further establish the optimality of $\boldsymbol{Z}^\star$ to (4.1), it suffices, similarly to Section 2 and Section 3, to construct dual variables $\boldsymbol{\Phi}$, $\boldsymbol{\Lambda}$, $\boldsymbol{\mu}$ and z for which the following KKT conditions hold:

- Stationarity: $-\boldsymbol{A} - \boldsymbol{\Lambda} - \boldsymbol{\Phi} + z\boldsymbol{I} - \boldsymbol{\mu}\mathbf{1}_n^\top - \mathbf{1}_n\boldsymbol{\mu}^\top = \mathbf{0}$,
- Complementary slackness: $\langle \boldsymbol{Z}^\star, \boldsymbol{\Lambda}\rangle = \langle \boldsymbol{Z}^\star, \boldsymbol{\Phi}\rangle = 0$,
- Dual feasibility: $\boldsymbol{\Phi} \succeq \mathbf{0}$, $\boldsymbol{\Lambda} \succeq \mathbf{0}$.

Notice that the above conditions differ from the KKT conditions of the relaxation (3.1) given in Section 3; in particular, the diagonal matrix $-\operatorname{diag}(\boldsymbol{\beta})$ is replaced by the scaled identity matrix $z\boldsymbol{I}$, which reduces the freedom for the dual construction. The detailed construction of these dual variables can be found in Yan, Sarkar and Cheng (2018). Here we cite the result established therein:

THEOREM 4.1. *Denote by $n_{\max}$ and $n_{\min}$ the maximum and minimum cluster sizes in the* SBM, *respectively. If the following separation condition holds*:

$$\begin{aligned} &\min_k \Big(B_{kk} - \max_{l\neq k} B_{kl}\Big) \\ &\quad \geq 2\sqrt{6\log n}\max_k \sqrt{B_{kk}/n_k} \\ &\qquad + 6\max_{1\leq k<l\leq r}\sqrt{B_{kl}\log n/n_{\min}} \\ &\qquad + C_0\sqrt{(n/n_{\min}^2)\Big(\max_k B_{kk}\Big)}, \end{aligned}$$

then with high probability, the matrix $\boldsymbol{Z}^\star$ in (4.2) is the unique solution to (4.1).

5. WEAK CONSISTENCY BY PRIMAL ANALYSIS

In this section, we describe a "primal" approach for analyzing convex relaxation formulations. As opposed to the dual certification analysis presented in the previous sections, this approach directly makes use of the primal optimality and feasibility of the SDP solution $\widehat{\boldsymbol{X}}$ and the ground-truth $\boldsymbol{X}^\star$. Combined with the celebrated Grothendieck's inequality (Grothendieck, 1996), this approach leads to a simple proof of weak consistency (Guédon and Vershynin, 2016). Under the equal-cluster-size assumption, a more refined primal analysis provides unified guarantees covering both weak and strong consistency (Fei and Chen, 2019a, 2019b).

To present the above results, we consider the (p, q)-SBM and the convex relaxation (1.5) with $q < \lambda < p$. Let $\tau := \min(p - \lambda, \lambda - q)$. For expository convenience, we simply choose $\lambda = \frac{p+q}{2}$, which is an approximation of (1.2) used in the exact log likelihood. In this case we have $\tau = \frac{p-q}{2}$. The arguments below can be extended to other values of λ in a straightforward fashion. Recall that our goal is to estimate the ground truth partition matrix $\boldsymbol{X}^\star$. We use the solution $\widehat{\boldsymbol{X}}$ to the convex relaxation (1.5) as an estimator, and aim to characterize the efficiency of (1.5) by bounding the distance between $\widehat{\boldsymbol{X}}$ and $\boldsymbol{X}^\star$.

The primal analysis begins with noting that the optimality of $\widehat{\boldsymbol{X}}$ and the feasibility of $\boldsymbol{X}^\star$ to the convex relaxation

(1.5) imply that $\langle \boldsymbol{X}^*, \boldsymbol{A} - \lambda \boldsymbol{J}_n \rangle \le \langle \widehat{\boldsymbol{X}}, \boldsymbol{A} - \lambda \boldsymbol{J}_n \rangle$. Rearranging the inequality and separating the expectation and deviation terms, we obtain the following inequality:

$$\begin{aligned} 0 &\le \langle \widehat{\boldsymbol{X}} - \boldsymbol{X}^\star, \boldsymbol{A} - \lambda \boldsymbol{J}_n \rangle \\ &= \langle \widehat{\boldsymbol{X}} - \boldsymbol{X}^\star, \mathbb{E}\boldsymbol{A} - \lambda \boldsymbol{J}_n \rangle + \langle \widehat{\boldsymbol{X}} - \boldsymbol{X}^\star, \boldsymbol{A} - \mathbb{E}\boldsymbol{A} \rangle. \end{aligned}$$

To proceed, we make use of a simple observation on the relationship between $\boldsymbol{X}^\star$ and $\mathbb{E}\boldsymbol{A}$. For each $i \neq j$, $\mathbb{E}A_{ij} = p > (p+q)/2$ if and only if $X^*_{ij} = 1$, whereas $\mathbb{E}A_{ij} = q < (p+q)/2$ if and only if $X^*_{ij} = 0$. Due to the feasibility of $\widehat{\boldsymbol{X}}$, we deduce that the entries of the error matrix $\boldsymbol{X}^\star - \widehat{\boldsymbol{X}}$ must have matching signs with those of $\mathbb{E}\boldsymbol{A} - \frac{p+q}{2}\boldsymbol{J}_n$. This observation implies that

$$\langle \boldsymbol{X}^\star - \widehat{\boldsymbol{X}}, \mathbb{E}\boldsymbol{A} - \lambda \boldsymbol{J}_n \rangle \ge \frac{p-q}{2} \|\widehat{\boldsymbol{X}} - \boldsymbol{X}^\star\|_1.$$

The key step in Guédon and Vershynin (2016) involves bounding the "deviation term" $\langle \widehat{\boldsymbol{X}} - \boldsymbol{X}^\star, \boldsymbol{A} - \mathbb{E}\boldsymbol{A} \rangle$. To this end, one applies the triangle inequality to obtain

$$\begin{aligned} &\langle \widehat{\boldsymbol{X}} - \boldsymbol{X}^\star, \boldsymbol{A} - \mathbb{E}\boldsymbol{A} \rangle \\ &\quad \le |\langle \widehat{\boldsymbol{X}}, \boldsymbol{A} - \mathbb{E}\boldsymbol{A} \rangle| + |\langle \boldsymbol{X}^\star, \boldsymbol{A} - \mathbb{E}\boldsymbol{A} \rangle| \\ &\quad \le 2 \sup_{\substack{\boldsymbol{X} \succeq 0 \\ \mathrm{diag}(\boldsymbol{X}) \le \boldsymbol{1}}} |\langle \boldsymbol{X}, \boldsymbol{A} - \mathbb{E}\boldsymbol{A} \rangle|, \end{aligned}$$

where $\mathrm{diag}(\boldsymbol{X}) \le \boldsymbol{1}$ means that $X_{ii} \le 1$ for all $i = 1, \dots, n$, and the second inequality follows from the feasibility of $\widehat{\boldsymbol{X}}$ and $\boldsymbol{X}^\star$. The Grothendieck's inequality (Grothendieck, 1996, Lindenstrauss and Pełczyński, 1968) guarantees that

$$\sup_{\boldsymbol{X} \succeq 0, \mathrm{diag}(\boldsymbol{X}) \le \boldsymbol{1}} |\langle \boldsymbol{X}, \boldsymbol{A} - \mathbb{E}\boldsymbol{A} \rangle| \le K_G \|\boldsymbol{A} - \mathbb{E}\boldsymbol{A}\|_{\infty \to 1},$$

where K_G denotes the Grothendieck's constant ($0 < K_G \le 1.783$) and $\|\boldsymbol{M}\|_{\infty \to 1} := \sup_{\boldsymbol{x}: \|\boldsymbol{x}\|_\infty \le 1} \|\boldsymbol{M}\boldsymbol{x}\|_1$ is the $\ell_\infty \to \ell_1$ operator norm for a matrix $\boldsymbol{M}$. We next make use of the identity

$$\begin{aligned} &\|\boldsymbol{A} - \mathbb{E}\boldsymbol{A}\|_{\infty \to 1} \\ &\quad = \sup_{\substack{\boldsymbol{x}: \|\boldsymbol{x}\|_\infty \le 1 \\ \boldsymbol{y}: \|\boldsymbol{y}\|_\infty \le 1}} \boldsymbol{y}^\top (\boldsymbol{A} - \mathbb{E}\boldsymbol{A}) \boldsymbol{x} \\ &\quad = \sup_{\substack{\boldsymbol{x} \in \{\pm 1\}^n \\ \boldsymbol{y} \in \{\pm 1\}^n}} \boldsymbol{y}^\top (\boldsymbol{A} - \mathbb{E}\boldsymbol{A}) \boldsymbol{x}. \end{aligned}$$

For each pair of fixed vectors $\boldsymbol{x}, \boldsymbol{y} \in \{\pm 1\}^n$, $\boldsymbol{y}^\top (\boldsymbol{A} - \mathbb{E}\boldsymbol{A}) \boldsymbol{x}$ can be written as the sum of independent random variables and hence be bounded using Bernstein's inequality. Taking a union bound over all $\boldsymbol{x}$ and $\boldsymbol{y}$, we obtain that with probability at least $1 - 2(e/2)^{-2n}$,

$$\|\boldsymbol{A} - \mathbb{E}\boldsymbol{A}\|_{\infty \to 1} \le 2\sqrt{2p(n^3 - n^2)} + (8/3)n.$$

Combining pieces, we conclude that with probability at least $1 - 2(e/2)^{-2n}$,

$$\begin{aligned} \|\widehat{\boldsymbol{X}} - \boldsymbol{X}^\star\|_1 &\le 2/(p-q)\big(8\sqrt{2p(n^3 - n^2)} + (32/3)n\big) \\ &\overset{(i)}{\le} 45\sqrt{pn^3}/(p-q), \end{aligned}$$

where step (i) holds provided that $p \ge 1/n$. We have therefore established the following result, which was first proved in Guédon and Vershynin (2016):

THEOREM 5.1. *Under the (p,q)-SBM, if $p \ge 1/n$ and $p > q$, then with high probability, any optimal solution $\widehat{\boldsymbol{X}}$ of the SDP (1.5) satisfies the bound $\|\widehat{\boldsymbol{X}} - \boldsymbol{X}^\star\|_1 \le 45\sqrt{pn^3}/(p-q)$.*

A prominent advantage of this result is that it is applicable even when the networks is very sparse, that is, $p \asymp 1/n$. Moreover, the result does not require any assumption on the number of clusters, the cluster sizes, or the knowledge thereof.

Given the above bound on $\|\widehat{\boldsymbol{X}} - \boldsymbol{X}^\star\|_1$, one can further derive bounds on the clustering errors. For example, in the two-cluster case one can simply estimate the underlying community label ϕ by taking the entry-wise sign of the leading eigenvector of $\widehat{\boldsymbol{Y}} = 2\widehat{\boldsymbol{X}} - \boldsymbol{J}_n$, the centered version of $\widehat{\boldsymbol{X}}$. It is easy to show that the fraction of misclassified vertices is upper bounded by $(1/n^2)\|\widehat{\boldsymbol{X}} - \boldsymbol{X}^\star\|_1$ up to constant factors (Guédon and Vershynin, 2016). Combined with Theorem 5.1, we conclude that the fraction of misclassified vertices is smaller than ϵ with high probability provided that $n(p-q)^2 \gtrsim p\epsilon^{-2}$.

Under the equal-cluster-size assumption, the error bound in Theorem 5.1 has been substantially improved by Fei and Chen (2019a). In particular, by deriving a tighter bound on the term $\langle \widehat{\boldsymbol{X}} - \boldsymbol{X}^\star, \boldsymbol{A} - \mathbb{E}\boldsymbol{A} \rangle$, they show that the estimation error of the SDP (1.5) in fact decays exponentially:

$$\|\widehat{\boldsymbol{X}} - \boldsymbol{X}^\star\|_1 \lesssim n^2 \exp\left[-\Omega\left(\frac{n(p-q)^2}{pr}\right)\right] \tag{5.1}$$

provided that $(p-q)^2/p \gtrsim r^2/n$. Assuming $r = O(1)$, the bound (5.1) guarantees that the fraction of misclassified nodes is at most ϵ as long as $n(p-q)^2 \gtrsim p\log(\epsilon^{-1})$. Consequently, the convex relaxation achieves weak consistency (i.e., $\epsilon \to 0$) if $n(p-q)^2/p \log n \to \infty$, and strong consistency (i.e., $\epsilon < 1/n$, which implies $\epsilon = 0$) if $n(p-q)^2/p \gtrsim \log n$. In fact, in the setting with $r = 2$ equal-sized clusters and using the SDP (3.1), Fei and Chen (2019b) show that the exponent in (5.1) can be taken to be $-(1 - o(1))(\sqrt{p} - \sqrt{q})^2 n/2$, which is minimax-optimal (cf. Zhang and Zhou, 2015) and as a corollary implies the sharp threshold for strong consistency presented in Section 3.

Let us provide an additional remark on the setting with two equal-sized clusters and sparse networks, that is,

$p = a/n$ and $q = b/n$ for two constants a, b. Assuming $a > b$ and $(a-b)^2 \geq C_0(a+b)$, a related question is what is the smallest C_0 to guarantee that the SDP formulation can achieve $\epsilon < 1/2$, that is, producing a nontrivial estimate that is better than random guess. It is known that the information-theoretic *necessary* condition for achieving $\epsilon < 1/2$ fraction of misclassified nodes is $(a-b)^2 > 2(a+b)$ (Mossel, Neeman and Sly, 2015, Massoulié, 2014, Mossel, Neeman and Sly, 2018). In the converse direction, recent work by Montanari and Sen (2015) shows that the SDP (3.2) with appropriate λ achieves $\epsilon < 1/2$ with probability $1 - o_n(1)$ if $(a-b)^2 > (2+\epsilon)(a+b)$ and $a + b > C_1$ for a sufficiently large constant C_1 depending on ϵ. Therefore, the results in Montanari and Sen (2015) show that the convex relaxation approach is nearly optimal in achieving nontrivial estimation. More recent work by Fan and Montanari (2017) gives more precise bounds on C_1 by proving that a simple local algorithm approximately solves the convex program within a factor $1 + O(1/(a+b))$.

6. DEGREE CORRECTION AND CONVEXIFIED MODULARITY MAXIMIZATION

An alternative to the likelihood-based approach (1.1) for community detection is modularity maximization. Proposed by Newman and Girvan (Newman, 2006), this approach involves finding a partition matrix $\boldsymbol{X}$ that maximizes the so-called "modularity":

$$\max_{\boldsymbol{X}} \sum_{1 \leq i < j \leq n} \left(A_{ij} - \frac{d_i d_j}{2L} \right) X_{ij}, \tag{6.1}$$

where $d_i := \sum_{j=1}^n A_{ij}$ is the degree of node i, and $L := \frac{1}{2}\sum_{i=1}^n d_i$ is the total number of edges. Compared to likelihood maximization (1.1), modularity maximization can be viewed as an adaptive variant, where the single tuning parameter λ is replaced by the quantity $\frac{d_i d_j}{2L}$ that is adaptive to each pair of nodes (i, j).

To overcome the so-called "resolution limit" of modularity maximization (Fortunato and Barthelemy, 2007), the work in Reichardt and Bornholdt (2006), Lancichinetti and Fortunato (2011) proposed replacing the number $\frac{1}{2L}$ in (6.1) with a tuning parameter λ. Convexifying the program similarly to (1.5), we obtain the following convex relaxation:

$$\begin{aligned} &\max_{\boldsymbol{X}} \quad \langle \boldsymbol{X}, \boldsymbol{A} - \lambda \boldsymbol{d}\boldsymbol{d}^\top \rangle \\ &\text{subject to } \boldsymbol{X} \succeq \boldsymbol{0}, \boldsymbol{X} \geq \boldsymbol{0}, X_{ii} = 1, \\ &\qquad \text{for } 1 \leq i \leq N. \end{aligned} \tag{6.2}$$

This approach, termed *convexified modularity maximization*, was first proposed by Chen, Li and Xu (2018). Notice that if we replace $\boldsymbol{d}$ with $\boldsymbol{1}_n$, then the formulation (6.2) falls back to convexified maximum likelihood (1.5).

An important advantage of modularity maximization over likelihood maximization lies inits applicability beyond the standard SBM. Indeed, it was shown in Chen, Li and Xu (2018) that the formulation (6.2) is consistent even if the data is generated from the more general *degree corrected* stochastic block model (DCSBM). We now describe this result, beginning with an introduction of DCSBM.

One obvious limitation of standard SBM is that nodes within the same group are statistically equivalent and hence have homogeneous degrees. Therefore, standard SBM fails to model degree heterogeneity that is common in real world networks. This limitation motivates researchers to consider the more general DCSBM. In DCSBM, each pair of nodes i and j are connected with probability $\theta_i \theta_j B_{\phi(i)\phi(j)}$—as opposed to $B_{\phi(i)\phi(j)}$ in SBM—where θ_i is a degree heterogeneity parameter for the node i.

A first step in deriving statistical guarantees under DCSBM is an appropriate generalization of the "density gap" condition $p - q > 0$ used in the (p, q)-SBM. We first note that for each node i in community a, its expected degree is

$$\begin{aligned} \mathbb{E}d_i &= \sum_{j \neq i} \mathbb{E}A_{ij} \\ &= \sum_{j \neq i} \theta_i \theta_j B_{a\phi(j)} \\ &= \theta_i \sum_j \theta_j B_{a\phi(j)} - \theta_i^2 B_{aa} \approx \theta_i \sum_j \theta_j B_{a\phi(j)}. \end{aligned}$$

Define $H_a := \sum_j \theta_j B_{a\phi(j)}$ and note the alternative expression $H_a = \sum_{b=1}^r B_{ab} G_b$, where $G_a := \sum_{i:\phi(i)=a} \theta_i$. Here H_a captures the *average degree* of the nodes in the a-th cluster in the sense that $H_a \approx \mathbb{E}d_i/\theta_i$. Based on these observations, Chen, Li and Xu (2018) proposed the following *degree-corrected density gap condition*:

$$\max_{1 \leq a < b \leq r} \frac{B_{ab}}{H_a H_b} < \min_{1 \leq a \leq r} \frac{B_{aa}}{H_a^2}. \tag{6.3}$$

This degree-corrected density gap condition enjoys several desirable properties. First, it only depends on the community aggregated quantities H_a's instead of individual node parameters θ_i's, so it is robust against abnormal nodes, such as one small θ_i. Moreover, while different pairs of $(\boldsymbol{B}, \boldsymbol{\theta})$ may correspond to the same DCSBM, the condition (6.3) is invariant under equivalent DCSBM's.

Chen, Li and Xu (2018) showed that under appropriate conditions stated in terms of the gap condition (6.3), the convexified modularity maximization (6.2) provides a good estimate of the ground truth partition matrix $\boldsymbol{X}^*$:

THEOREM 6.1. *Under* DCSBM, *suppose that the degree-corrected density gap condition* (6.3) *holds and*

the tuning parameter λ in the formulation (6.2) satisfies

$$\max_{1\le a<b\le r}\frac{B_{ab}+\delta}{H_aH_b}\le\lambda\le\min_{1\le a\le r}\frac{B_{aa}-\delta}{H_a^2} \tag{6.4}$$

for some number $\delta>0$. Then with high probability, any solution $\widehat{\boldsymbol{X}}$ to the convex relaxation (6.2) satisfies the bound

$$\begin{aligned}\sum_{i=1}^n\sum_{j=1}^n\theta_i\theta_j|\widehat{X}_{ij}-X^*_{ij}|&\\ \le\frac{C_0}{\delta}(1+\lambda\sigma)(\sqrt{n\sigma}+n),&\end{aligned} \tag{6.5}$$

where $\sigma:=\sum_{a,b}B_{ab}G_aG_b$.

This theorem is proved via a primal analysis argument similar to that in Guédon and Vershynin (2016) as presented in Section 5.

Note that the error metric in (6.5) is weighted by the degree heterogeneity parameters. This captures the fact that nodes with different θ_i's and hence different degrees have different contributions to the overall clustering quality. Moreover, the error bound (6.5) is insensitive to $\theta_{\min}:=\min_i\theta_i$. Therefore, the presence of a small $\theta_{\min}$ does not hinder recovery of the memberships of nodes with large θ_i.

To better appreciate the results in Theorem 6.1, let us consider a simple sub-class of DCSBM with symmetric and balanced clusters, where $B_{aa}=p$ for all $1\le a\le r$, $B_{ab}=q$ for all $1\le a<b\le r$, and $G_a=g$ for all $1\le a\le r$. In this setting, direct calculation gives $H_a=((r-1)q+p)g$ for all $1\le a\le r$. The constraint (6.4) on δ and λ then simplifies to

$$\begin{aligned}&p-q\ge2\delta\quad\text{and}\\ &\frac{q+\delta}{(p+(r-1)q)^2g^2}\le\lambda\le\frac{p-\delta}{(p+(r-1)q)^2g^2}.\end{aligned} \tag{6.6}$$

Note that the first inequality above is identical to the standard density gap condition imposed in SBM. Furthremore, one has $\sigma=r(p+(r-1)q)g^2\le r^2pg^2$. Substituting these expressions into equation (6.5) yields the error bound

$$\sum_{i=1}^n\sum_{j=1}^n\theta_i\theta_j|\widehat{X}_{ij}-X^*_{ij}|\lesssim\frac{r}{\delta}(rg\sqrt{np}+n). \tag{6.7}$$

As discussed in the previous sections, the solution $\widehat{\boldsymbol{X}}$ of the convex relaxation is not necessarily a partition matrix; Chen, Li and Xu (2018) show that an explicit clustering can be extracted from $\widehat{\boldsymbol{X}}$ using a *weighted* k-medoids algorithm. The number of misclassified nodes, *weighted* by their degree heterogeneity parameters $\{\theta_i\}$, satisfies the bound

$$\sum_{i\in\mathcal{E}}\theta_i\lesssim\frac{r}{\delta}\Big(r\sqrt{np}+\frac{n}{g}\Big), \tag{6.8}$$

where $\mathcal{E}$ is the set of misclassified nodes. We refer the readers to Chen, Li and Xu (2018), Sections 2.4, 3.2, for the details of this result.

We note that the bound in (6.8) gives nontrivial guarantees all the way down to the *sparse graph regime* with bounded average degrees. For example, suppose that $p=a/n$ and $q=b/n$ for two fixed constants $a>b, r=O(1)$, and $g\asymp n$. With $(a-b)/\sqrt{a}$ sufficiently large and the choice $\delta\asymp(a-b)/n$, one may bound the right-hand side of (6.8) by an arbitrarily small constant times n, thereby guaranteeing that an arbitrarily small fraction of nodes are misclassified. In contrast, standard spectral clustering methods are known to be inconsistent in this sparse regime (Krzakala et al., 2013).

Strong consistency: If in addition the minimum degree heterogeneity parameter $\theta_{\min}$ is not too small, Chen, Li and Xu (2018) further show that convexified maximum likelihood perfectly recovers the ground truth communities with high probability. To illustrate this strong consistency result, consider a sub-class of DCSBM with $B_{aa}=p$ for all $a=1,\dots,r$, and $B_{ab}=q$ for all $1\le a<b\le r$. Under this setup, the degree-corrected density gap condition (6.4) becomes

$$\max_{1\le a<b\le r}\frac{q+\delta}{H_aH_b}\le\lambda\le\min_{1\le a\le r}\frac{p-\delta}{H_a^2}. \tag{6.9}$$

Further define $G_{\min}:=\min_{1\le a\le r}G_a$. The following theorem, again extracted from Chen, Li and Xu (2018), provides sufficient conditions for strong consistency.

THEOREM 6.2. *Suppose that the tuning parameter λ satisfies the degree-corrected density gap condition (6.9) for some number $\delta>0$, and that*

$$\delta>C_0\Big(\frac{\sqrt{qn}}{G_{\min}}+\sqrt{\frac{p\log n}{G_{\min}\theta_{\min}}}\Big). \tag{6.10}$$

With high probability, the ground truth partition matrix $\boldsymbol{X}^\star$ is the unique solution to the convex relaxation (6.2).

This theorem is proved using a dual certification analysis argument similar to those presented in Section 2. Note that the condition (6.10) for perfect recovery depends on the minimum values $G_{\min}$ and $\theta_{\min}$. This is expected: it is impossible to recover the membership of a node with an overly small θ_i and G_a, as this node will have too few edges.

7. ROBUSTNESS AGAINST OUTLIER NODES

The results presented in the previous sections highlight some distinguishing statistical advantages of convex relaxation methods for community detection, including consistency under sparse networks and weak assortativity. In this section, we demonstrate another important benefit of convex relaxation, namely, robustness. In Cai

and Li (2015), it is shown, both theoretically and empirically, that a simple variant of (1.5) is robust against outlier/adversarial nodes. That is, convex optimization methods can yield consistent estimates even if a proportion of the nodes have arbitrary connections that do not satisfy the SBM. Below we give an expository introduction of this result and the associated analytical arguments.

For analytical convenience, Cai and Li (2015) consider the following modified version of the convex relaxation (1.5):

$$\begin{aligned} \max_{\boldsymbol{X}} \quad & \langle \boldsymbol{X}, \boldsymbol{A} - \lambda \boldsymbol{J} - (\alpha - \lambda)\boldsymbol{I} \rangle \\ \text{subject to } & \boldsymbol{X} \succeq \boldsymbol{0}, \boldsymbol{X} \geq \boldsymbol{0}, X_{ii} \leq 1, \\ & \text{for } 1 \leq i \leq n. \end{aligned} \tag{7.1}$$

The optimization problem (7.1) differs from (1.5) in that it does not require all diagonal entries of $\boldsymbol{X}$ to equal one. Instead, the entries of $\boldsymbol{X}$ are constrained to fall between 0 and 1, and the trace of $\boldsymbol{X}$ is penalized by introducing a second tuning parameter α.

To study the robustness of the formulation (7.1), let us consider an extension of SBM that allows for abnormal nodes. We assume that the set of nodes V can be decomposed as $V = I \cup O$, where I represents the inliers and O is the outliers. We assume the edges between nodes within I are generated randomly according to the (p, q)-SBM, whereas if one of i and j is an outlier, then i and j are connected in an arbitrary or even adversarial manner.

Under the above model, one can show that the formulation (7.1) can still consistently recover the community memberships of the nodes in I. The argument is based on analyzing the KKT condition for (7.1), but it differs from the purely dual certification analysis presented in Section 2. There when studying the standard SBM and the formulations (1.5), (3.1) and (4.1), it is clear which is the desired solution whose optimality we want to certify—it is either the ground truth partition matrix in (1.4) or the ground truth projective matrix in (4.2). Under the generalized SBM with outlier nodes, however, the desired solution to (7.1) is much more complicated. Consequently, the analysis of (7.1) involves both the construction of a candidate primal solution and that of its dual certificate. This analytic technique, sometimes called *primal-dual witness*, is powerful but relatively uncommon in the convex community detection literature. Below we briefly outline the key ideas of this argument.

Similarly to Section 2, we focus on the (p, q)-model with two inlier clusters. We assume that the $|I| = n$ inliers are partitioned into two clusters of sizes n_1 and n_2, and the number of outliers is $|O| = m$, with $n = n_1 + n_2 + m$ being the total number of nodes. We refer to this model as the (p, q, n_1, n_2, m)-*generalized* SBM. In the presence of outliers, the adjacency matrix (up to some permutation) has the following augmented block structure:

$$\boldsymbol{A} = \begin{bmatrix} \boldsymbol{A}_{11} & \boldsymbol{A}_{12} & \boldsymbol{Z}_1 \\ \boldsymbol{A}_{12}^\top & \boldsymbol{A}_{22} & \boldsymbol{Z}_2 \\ \boldsymbol{Z}_1^\top & \boldsymbol{Z}_2^\top & \boldsymbol{W} \end{bmatrix}. \tag{7.2}$$

Here $\boldsymbol{A}_{11}$, $\boldsymbol{A}_{22}$ and $\boldsymbol{A}_{12}$ correspond to the connectivity within and between the two inlier clusters, and hence have the same statistical properties as in SBM. On the other hand, the blocks $\boldsymbol{Z}_1$ and $\boldsymbol{Z}_2$ represent the connectivity between the outliers and inliers, whereas the $m \times m$ symmetric block $\boldsymbol{W}$ is the adjacency matrix within the outliers. Note that $\boldsymbol{Z}_1$, $\boldsymbol{Z}_2$ and $\boldsymbol{W}$ are all arbitrary.

We say that the formulation (7.1) is robust against outliers if it achieves strong consistency within the inlier nodes in I; that is, its solution $\widehat{\boldsymbol{X}}$ is of the form

$$\widehat{\boldsymbol{X}} = \begin{bmatrix} \boldsymbol{J}_{n_1} & \boldsymbol{0} & \star \\ \boldsymbol{0} & \boldsymbol{J}_{n_2} & \star \\ \star & \star & \star \end{bmatrix}. \tag{7.3}$$

To establish this property, we begin with the KKT condition of (7.1) for the primal solution $\widehat{\boldsymbol{X}}$:

$$\begin{cases} \text{Stationarity:} \\ \quad -\boldsymbol{A} + \lambda \boldsymbol{J} + (\alpha - \lambda)\boldsymbol{I} = \boldsymbol{\Lambda} + \boldsymbol{\Phi} - \operatorname{diag}(\boldsymbol{\gamma}), \\ \text{Primal feasibility:} \\ \quad \widehat{\boldsymbol{X}} \succeq \boldsymbol{0}, \widehat{\boldsymbol{X}} \geq \boldsymbol{0}, \widehat{X}_{ii} \leq 1 \text{ for } 1 \leq i \leq n, \\ \text{Dual feasibility:} \\ \quad \boldsymbol{\Lambda} \succeq \boldsymbol{0}, \boldsymbol{\Phi} \geq \boldsymbol{0}, \boldsymbol{\gamma} \geq \boldsymbol{0}, \\ \text{Complimentary slackness:} \\ \quad \langle \widehat{\boldsymbol{X}}, \boldsymbol{\Lambda} \rangle = 0, \langle \widehat{\boldsymbol{X}}, \boldsymbol{\Phi} \rangle = 0, \gamma_i = 0 \text{ iff } \widehat{X}_{ii} < 1. \end{cases} \tag{7.4}$$

Notice that $\boldsymbol{\Lambda}$ is determined by $\boldsymbol{\Phi}$ and $\boldsymbol{\gamma}$. It remains to construct the primal solution $\widehat{\boldsymbol{X}}$ and dual variables $\boldsymbol{\Phi}$ and $\boldsymbol{\gamma}$ for which the above conditions hold.

Primal solution candidate: Inspired by numerical simulations, Cai and Li (2015) consider a candidate primal solution of the following form:

$$\widehat{\boldsymbol{X}} = \begin{bmatrix} \boldsymbol{J}_{n_1} & \boldsymbol{0} & \boldsymbol{1}_{n_1} \boldsymbol{x}_1^\top \\ \boldsymbol{0} & \boldsymbol{J}_{n_2} & \boldsymbol{1}_{n_2} \boldsymbol{x}_2^\top \\ \boldsymbol{x}_1 \boldsymbol{1}_{n_1}^\top & \boldsymbol{x}_2 \boldsymbol{1}_{n_2}^\top & \boldsymbol{x}_1 \boldsymbol{x}_1^\top + \boldsymbol{x}_2 \boldsymbol{x}_2^\top \end{bmatrix}. \tag{7.5}$$

Assume that

$$-\boldsymbol{A} + \lambda \boldsymbol{J} + (\alpha - \lambda)\boldsymbol{I} := \begin{bmatrix} \star & \star & \widetilde{\boldsymbol{Z}}_1 \\ \star & \star & \widetilde{\boldsymbol{Z}}_2 \\ \widetilde{\boldsymbol{Z}}_1^\top & \widetilde{\boldsymbol{Z}}_2^\top & \widetilde{\boldsymbol{W}} \end{bmatrix}.$$

Plugging the parameterization (7.5) of the primal solution into the convex relaxation (7.1), we obtain the following

quadratic program:

$$\min_{\boldsymbol{x}_1,\boldsymbol{x}_2\in\mathbb{R}^m} \quad \sum_{i=1}^{2}\langle \boldsymbol{x}_i, \widetilde{\boldsymbol{Z}}_i^\top \mathbf{1}_{n_i}\rangle + \frac{1}{2}\sum_{i=1}^{2}\boldsymbol{x}_i^\top \widetilde{\boldsymbol{W}}\boldsymbol{x}_i$$

$$\text{subject to } \boldsymbol{x}_1 \geq \mathbf{0}, \quad \boldsymbol{x}_2 \geq \mathbf{0}, \tag{7.6}$$

$$\sum_{i=1}^{2}\boldsymbol{x}_i^\top(\boldsymbol{e}_j\boldsymbol{e}_j^\top)\boldsymbol{x}_i \leq 1, \text{ for } 1\leq j\leq m.$$

Slightly abusing notation, we use $(\boldsymbol{x}_1, \boldsymbol{x}_2)$ to denote the solution to the above program. Plugging this pair $(\boldsymbol{x}_1, \boldsymbol{x}_2)$ into (7.5) gives our primal candidate $\widehat{\boldsymbol{X}}$, which by definition automatically satisfies primal feasibility in KKT condition.

As an interesting and useful byproduct, we know that the solutions to the dual problem of (7.6), denoted by $\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \boldsymbol{\xi} \geq \mathbf{0} \in \mathbb{R}^m$, satisfy the conditions

$$\begin{cases} \widetilde{\boldsymbol{W}}\boldsymbol{x}_i + \widetilde{\boldsymbol{Z}}_i^\top \mathbf{1}_{l_i} = \boldsymbol{\beta}_i - \operatorname{diag}(\boldsymbol{\xi})\boldsymbol{x}_i, & i = 1, 2, \\ \xi_j\left(1 - \sum_{i=1}^{r}\boldsymbol{x}_i^\top(\boldsymbol{e}_j\boldsymbol{e}_j^\top)\boldsymbol{x}_i\right) = 0, & j = 1, \ldots, m, \\ \langle \boldsymbol{x}_i, \boldsymbol{\beta}_i\rangle = 0, & i = 1, 2. \end{cases}$$

We then use $\boldsymbol{\beta}_1, \boldsymbol{\beta}_2$ and $\boldsymbol{\xi}$ to construct the dual certificates $\boldsymbol{\Phi}$ and $\boldsymbol{\gamma}$ as follows:

$$\boldsymbol{\Phi} = \begin{bmatrix} \mathbf{0} & \boldsymbol{\Phi}_{12} & \frac{1}{n_1}\mathbf{1}_{n_1}\boldsymbol{\beta}_1^\top \\ \boldsymbol{\Phi}_{12}^\top & \mathbf{0} & \frac{1}{n_2}\mathbf{1}_{n_2}\boldsymbol{\beta}_2^\top \\ \frac{1}{n_1}\boldsymbol{\beta}_1\mathbf{1}_{n_1}^\top & \frac{1}{n_2}\boldsymbol{\beta}_2\mathbf{1}_{n_2}^\top & \mathbf{0} \end{bmatrix},$$

$$\boldsymbol{\gamma} = \begin{bmatrix} \boldsymbol{\gamma}_1 \\ \boldsymbol{\gamma}_2 \\ \boldsymbol{\xi} \end{bmatrix};$$

see Cai and Li (2015) for the expressions of $\boldsymbol{\gamma}_1$, $\boldsymbol{\gamma}_2$ and $\boldsymbol{\Phi}_{12}$.

With the primal and dual solutions constructed above, the KKT conditions in (7.4) can be validated with high probability under appropriate conditions on p, q, n_1, n_2, m and the tuning parameters λ and α. Doing so proves the main Theorem 3.1 in Cai and Li (2015). A corollary thereof is given below.

THEOREM 7.1. *Under the above* (p, q, n_1, n_2, m)*-generalized* SBM, *suppose that* $n_1 = O(n)$ *and* $n_2 = O(n)$, *and parametrize* p, q, m *as* $p = a(\log n)/n$, $q = b(\log n)/n$, *and* $m = \theta \log n$. *If*

$$a - b > C_0(1 + \sqrt{\theta})\sqrt{a},$$

then there exist tuning parameters λ *and* α *such that with high probability any solution to* (7.1) *is of the form* (7.3).

This result implies that if $p = O(\log n/n)$ and $q = O(\log n/n)$, the convex relaxation (7.1) is able to recover the true clusters in presence of $O(\log n)$ outlier nodes. The result can be extended to the more general setting with a growing number of clusters and a general connectivity probability matrix satisfying strong assortativity; see Cai and Li (2015) for the details.

Before concluding this section, we note that other notions of robustness have been considered in the literature on convex community detection. For example, one may consider a semirandom model where a monotone adversary is allowed to add edges within communities and delete edges across communities. Weak consistency results have been established under this model (Moitra, Perry and Wein, 2016). A more general semirandom model has later been introduced by Makarychev, Makarychev and Vijayaraghavan (2016); their results on outlier-robustness are comparable to (and sometimes better than) the results in Cai and Li (2015).

8. CONCLUSIONS

The convex relaxation approaches to community detection have attracted much recent attention. The first and foremost advantage of such approaches is that they provide a *computationally* tractable solution to the worst-case hard combinatorial problems that arise in community detection and clustering. In this survey, we have focused on the *statistical* advantages of the convex relaxation approaches, including the power in providing rigorous and strong recovery guarantees, the robustness against data corruption and model misspecification, and the flexibility in handling a broad range of variations of the community detection problems. We have highlighted several key analytical tools for establishing such properties under certain statistical models of the data. As can be seen, such analysis is largely decoupled from the specific algorithms used to solve the convex programs, which is an important benefit of the convex relaxation framework.

Due to space limit, we did not provide a comprehensive coverage of the fast-growing recent literature in the area of convex community detection. Some notable omissions include:

- Solving the convex programs: This can be done using a variety of existing powerful solvers such as interior point methods and ADMM (Nocedal and Wright, 2006). By leveraging specific structures (such as sparsity and low-rankness) that arise in community detection, even more efficient solvers have been developed that scale well to large datasets. For work in this direction we refer to the papers by Chen, Li and Xu (2018), Javanmard, Montanari and Ricci-Tersenghi (2015), Bandeira, Boumal and Voroninski (2016), Montanari (2016), Mei et al. (2017) and the references therein.
- Empirical performance: The convex relaxation approaches have been shown to enjoy competitive performance on various synthetic and real-world benchmarks; see, for example, the work in Cai and Li (2015),

Chen, Li and Xu (2018), Javanmard, Montanari and Ricci-Tersenghi (2015).

- Computational barriers in solving problems that are hard on average: The work in Chen and Xu (2016), Javanmard, Montanari and Ricci-Tersenghi (2015), Hajek, Wu and Xu (2016c) studies the interplay between the statistical and computational limits of SDP relaxations. Based on deep but nonrigorous statistical physics arguments, Decelle et al. (2011) conjectures that there exists a fundamental computational barrier that prevents all polynomial-time procedures from achieving the optimal statistical performance. Interested readers are referred to the recent surveys in Moore (2017), Abbe (2017), Wu and Xu (2018) for detailed discussions.
- Other considerations and extensions: Examples include extracting explicit clustering from the convex relaxation solutions (Guédon and Vershynin, 2016, Chen, Li and Xu, 2018), dealing with an unknown number of clusters and other unknown problem parameters (Yan, Sarkar and Cheng, 2018, Chen, Sanghavi and Xu, 2014), and clustering partially observed or weighted networks (Chen et al., 2014, Lim, Chen and Xu, 2017).

We hope that this survey will stimulate further research exploiting the power of the convex relaxation approaches to network analysis, and in particular provide impetus to the development of efficient implementations of these approaches for practical and large-scale problems.

ACKNOWLEDGMENTS

Y. Chen is partially supported by NSF Grants 1657420 and 1704828.

REFERENCES

ABBE, E. (2017). Community detection and stochastic block models: Recent developments. *J. Mach. Learn. Res.* **18** Paper No. 177, 86. MR3827065

ABBE, E., BANDEIRA, A. S. and HALL, G. (2016). Exact recovery in the stochastic block model. *IEEE Trans. Inform. Theory* **62** 471–487. MR3447993 https://doi.org/10.1109/TIT.2015.2490670

AGARWAL, N., BANDEIRA, A. S., KOILIARIS, K. and KOLLA, A. (2015). Multisection in the Stochastic Block Model using Semidefinite Programming. Available at arXiv:1507.02323.

AMES, B. P. W. (2014). Guaranteed clustering and biclustering via semidefinite programming. *Math. Program.* **147** 429–465. MR3258531 https://doi.org/10.1007/s10107-013-0729-x

AMES, B. P. W. and VAVASIS, S. A. (2014). Convex optimization for the planted k-disjoint-clique problem. *Math. Program.* **143** 299–337. MR3152071 https://doi.org/10.1007/s10107-013-0733-1

AMINI, A. A. and LEVINA, E. (2018). On semidefinite relaxations for the block model. *Ann. Statist.* **46** 149–179. MR3766949 https://doi.org/10.1214/17-AOS1545

BANDEIRA, A. S. (2018). Random Laplacian matrices and convex relaxations. *Found. Comput. Math.* **18** 345–379. MR3777782 https://doi.org/10.1007/s10208-016-9341-9

BANDEIRA, A. S., BOUMAL, N. and VORONINSKI, V. (2016). On the low-rank approach for semidefinite programs arising in synchronization and community detection. In 29*th Annual Conference on Learning Theory* 361–382.

BANDEIRA, A. S. and VAN HANDEL, R. (2016). Sharp nonasymptotic bounds on the norm of random matrices with independent entries. *Ann. Probab.* **44** 2479–2506. MR3531673 https://doi.org/10.1214/15-AOP1025

BOYD, S. and VANDENBERGHE, L. (2004). *Convex Optimization*. Cambridge Univ. Press, Cambridge. MR2061575 https://doi.org/10.1017/CBO9780511804441

CAI, T. T. and LI, X. (2015). Robust and computationally feasible community detection in the presence of arbitrary outlier nodes. *Ann. Statist.* **43** 1027–1059. MR3346696 https://doi.org/10.1214/14-AOS1290

CHAUDHURI, K., CHUNG, F. and TSIATAS, A. (2012). Spectral clustering of graphs with general degrees in the extended planted partition model. **23** *J. Mach. Learn. Res.* 35.1–35.23.

CHEN, Y., LI, X. and XU, J. (2018). Convexified modularity maximization for degree-corrected stochastic block models. *Ann. Statist.* **46** 1573–1602. MR3819110 https://doi.org/10.1214/17-AOS1595

CHEN, Y., SANGHAVI, S. and XU, H. (2014). Improved graph clustering. *IEEE Trans. Inform. Theory* **60** 6440–6455. MR3265033 https://doi.org/10.1109/TIT.2014.2346205

CHEN, Y. and XU, J. (2016). Statistical-computational tradeoffs in planted problems and submatrix localization with a growing number of clusters and submatrices. *J. Mach. Learn. Res.* **17** 882–938. MR3491121

CHEN, Y., JALALI, A., SANGHAVI, S. and XU, H. (2014). Clustering partially observed graphs via convex optimization. *J. Mach. Learn. Res.* **15** 2213–2238. MR3231602

DECELLE, A., KRZAKALA, F., MOORE, C. and ZDEBOROVÁ, L. (2011). Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Phys. Rev. E* **84** 066106.

FAN, Z. and MONTANARI, A. (2017). How well do local algorithms solve semidefinite programs? In *STOC'*17*—Proceedings of the* 49*th Annual ACM SIGACT Symposium on Theory of Computing* 604–614. ACM, New York. MR3678214

FEI, Y. and CHEN, Y. (2018). Hidden integrality of SDP relaxation for sub-Gaussian mixture models. In *Conference on Learning Theory* (*COLT*).

FEI, Y. and CHEN, Y. (2019a). Exponential error rates of SDP for block models: Beyond Grothendieck's inequality. *IEEE Trans. Inform. Theory* **65** 551–571. MR3901009 https://doi.org/10.1109/TIT.2018.2839677

FEI, Y. and CHEN, Y. (2019b). Achieving the Bayes error rate in stochastic block model by SDP, robustly. In *Conference on Learning Theory* (*COLT*).

FEIGE, U. and OFEK, E. (2005). Spectral techniques applied to sparse random graphs. *Random Structures Algorithms* **27** 251–275. MR2155709 https://doi.org/10.1002/rsa.20089

FORTUNATO, S. and BARTHELEMY, M. (2007). Resolution limit in community detection. *Proc. Natl. Acad. Sci. USA* **104** 36–41.

GOEMANS, M. X. and WILLIAMSON, D. P. (1995). Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.* **42** 1115–1145. MR1412228 https://doi.org/10.1145/227683.227684

GROTHENDIECK, A. (1996). Résumé de la théorie métrique des produits tensoriels topologiques. *Resenhas* **2** 401–480. MR1466414

GUÉDON, O. and VERSHYNIN, R. (2016). Community detection in sparse networks via Grothendieck's inequality. *Probab. Theory Related Fields* **165** 1025–1049. MR3520025 https://doi.org/10.1007/s00440-015-0659-z

HAJEK, B., WU, Y. and XU, J. (2016a). Achieving exact cluster recovery threshold via semidefinite programming. *IEEE Trans. Inform. Theory* **62** 2788–2797. MR3493879 https://doi.org/10.1109/TIT.2016.2546280

HAJEK, B., WU, Y. and XU, J. (2016b). Achieving exact cluster recovery threshold via semidefinite programming: Extensions. *IEEE Trans. Inform. Theory* **62** 5918–5937. MR3552431 https://doi.org/10.1109/TIT.2016.2594812

HAJEK, B., WU, Y. and XU, J. (2016c). Semidefinite programs for exact recovery of a hidden community. In *Proceedings of Conference on Learning Theory* (*COLT*). Available at arXiv:1602.06410.

HOLLAND, P. W., LASKEY, K. B. and LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Soc. Netw.* **5** 109–137. MR0718088 https://doi.org/10.1016/0378-8733(83)90021-7

IGUCHI, T., MIXON, D. G., PETERSON, J. and VILLAR, S. (2015). On the tightness of an SDP relaxation of k-means. Preprint. Available at arXiv:1505.04778.

IGUCHI, T., MIXON, D. G., PETERSON, J. and VILLAR, S. (2017). Probably certifiably correct k-means clustering. *Math. Program.* **165** 605–642. MR3707375 https://doi.org/10.1007/s10107-016-1097-0

JAVANMARD, A., MONTANARI, A. and RICCI-TERSENGHI, F. (2015). Phase transitions in semidefinite relaxations. Available at arXiv:1511.08769.

KRZAKALA, F., MOORE, C., MOSSEL, E., NEEMAN, J., SLY, A., ZDEBOROVÁ, L. and ZHANG, P. (2013). Spectral redemption in clustering sparse networks. *Proc. Natl. Acad. Sci. USA* **110** 20935–20940. MR3174850 https://doi.org/10.1073/pnas.1312486110

LANCICHINETTI, A. and FORTUNATO, S. (2011). Limits of modularity maximization in community detection. *Phys. Rev. E* **84**.

LI, X., LI, Y., LING, S., STROHMER, T. and WEI, K. (2017). When do birds of a feather flock together? K-means, proximity, and conic programming. Preprint. Available at arXiv:1710.06008.

LIM, S. H., CHEN, Y. and XU, H. (2017). Clustering from general pairwise observations with applications to time-varying graphs. *J. Mach. Learn. Res.* **18** Paper No. 49, 47. MR3670494

LINDENSTRAUSS, J. and PEŁCZYŃSKI, A. (1968). Absolutely summing operators in L_p-spaces and their applications. *Studia Math.* **29** 275–326. MR0231188 https://doi.org/10.4064/sm-29-3-275-326

MAKARYCHEV, K., MAKARYCHEV, Y. and VIJAYARAGHAVAN, A. (2016). Learning communities in the presence of errors. In *Conference on Learning Theory* 1258–1291.

MASSOULIÉ, L. (2014). Community detection thresholds and the weak Ramanujan property. In *STOC'*14*—Proceedings of the* 2014 *ACM Symposium on Theory of Computing* 694–703. ACM, New York. MR3238997

MATHIEU, C. and SCHUDY, W. (2010). Correlation clustering with noisy input. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms* 712–728. SIAM, Philadelphia, PA. MR2768627

MEI, S., MISIAKIEWICZ, T., MONTANARI, A. and OLIVEIRA, R. I. (2017). Solving SDPs for synchronization and MaxCut problems via the Grothendieck inequality. In *Proceedings of Machine Learning Research* **65** 1–40.

MOITRA, A., PERRY, W. and WEIN, A. S. (2016). How robust are reconstruction thresholds for community detection? In *STOC'*16*—Proceedings of the* 48*th Annual ACM SIGACT Symposium on Theory of Computing* 828–841. ACM, New York. MR3536617

MONTANARI, A. (2016). A Grothendieck-type inequality for local maxima. Preprint. Available at arXiv:1603.04064.

MONTANARI, A. and SEN, S. (2015). Semidefinite programs on sparse random graphs. Available at arXiv:1504.05910.

MOORE, C. (2017). The computer science and physics of community detection: Landscapes, phase transitions, and hardness. *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS* **121** 26–61. MR3699594

MOSSEL, E., NEEMAN, J. and SLY, A. (2015). Reconstruction and estimation in the planted partition model. *Probab. Theory Related Fields* **162** 431–461. MR3383334 https://doi.org/10.1007/s00440-014-0576-6

MOSSEL, E., NEEMAN, J. and SLY, A. (2016). Consistency thresholds for the planted bisection model. *Electron. J. Probab.* **21** Paper No. 21, 24. MR3485363 https://doi.org/10.1214/16-EJP4185

MOSSEL, E., NEEMAN, J. and SLY, A. (2018). A proof of the block model threshold conjecture. *Combinatorica* **38** 665–708. MR3876880 https://doi.org/10.1007/s00493-016-3238-8

NEWMAN, M. E. J. (2006). Modularity and community structure in networks. *Proc. Natl. Acad. Sci. USA* **103** 8577–8582. https://doi.org/10.1073/pnas.0601602103.

NOCEDAL, J. and WRIGHT, S. J. (2006). *Numerical Optimization*, 2nd ed. *Springer Series in Operations Research and Financial Engineering*. Springer, New York. MR2244940

OYMAK, S. and HASSIBI, B. (2011). Finding dense clusters via low rank + sparse decomposition. Available at arXiv:1104.5186.

PENG, J. and WEI, Y. (2007). Approximating K-means-type clustering via semidefinite programming. *SIAM J. Optim.* **18** 186–205. MR2299680 https://doi.org/10.1137/050641983

PERRY, W. and WEIN, A. S. (2015). A semidefinite program for unbalanced multisection in the stochastic block model. Available at arXiv:1507.05605.

REICHARDT, J. and BORNHOLDT, S. (2006). Statistical mechanics of community detection. *Phys. Rev. E* (3) **74** 016110, 14. MR2276596 https://doi.org/10.1103/PhysRevE.74.016110

ROHE, K., CHATTERJEE, S. and YU, B. (2011). Spectral clustering and the high-dimensional stochastic blockmodel. *Ann. Statist.* **39** 1878–1915. MR2893856 https://doi.org/10.1214/11-AOS887

VU, V. (2018). A simple SVD algorithm for finding hidden partitions. *Combin. Probab. Comput.* **27** 124–140. MR3734334 https://doi.org/10.1017/S0963548317000463

WU, Y. and XU, J. (2018). Statistical problems with planted structures: Information-theoretical and computational limits. Preprint. Available at arXiv:1806.00118.

YAN, B., SARKAR, P. and CHENG, X. (2018). Provable estimation of the number of blocks in block models. In *Proceedings of the* 21*st International Conference on Artificial Intelligence and Statistics* (*AISTATS*) 2018, *Lanzarote*, *Spain*.

ZHANG, A. Y. and ZHOU, H. H. (2015). Minimax rates of community detection in stochastic block models. Available at arXiv:1507.05313.