CONVEXIFIED MODULARITY MAXIMIZATION FOR DEGREE-CORRECTED STOCHASTIC BLOCK MODELS

BY YUDONG CHEN[1], XIAODONG LI[2] AND JIAMING XU[3]

Cornell University, University of California, Davis and Purdue University

The stochastic block model (SBM), a popular framework for studying community detection in networks, is limited by the assumption that all nodes in the same community are statistically equivalent and have equal expected degrees. The degree-corrected stochastic block model (DCSBM) is a natural extension of SBM that allows for degree heterogeneity within communities. To find the communities under DCSBM, this paper proposes a *convexified modularity maximization* approach, which is based on a convex programming relaxation of the classical (generalized) modularity maximization formulation, followed by a novel doubly-weighted ℓ_1-norm k-medoids procedure. We establish nonasymptotic theoretical guarantees for approximate and perfect clustering, both of which build on a new *degree-corrected density gap condition*. Our approximate clustering results are insensitive to the minimum degree, and hold even in sparse regime with bounded average degrees. In the special case of SBM, our theoretical guarantees match the best-known results of computationally feasible algorithms. Numerically, we provide an efficient implementation of our algorithm, which is applied to both synthetic and real-world networks. Experiment results show that our method enjoys competitive performance compared to the state of the art in the literature.

1. Introduction. Detecting communities/clusters in networks and graphs is an important task in many applications across computer, social and natural sciences and engineering. A standard framework for studying community detection in a statistical setting is the stochastic block model (SBM) proposed in [27]. Also known as the planted partition model in the computer science literature [16], SBM generates a random graph from a set of underlying clusters. The statistical task is to accurately recover the underlying true clusters given a single realization of the random graph.

SBM is arguably the most popular model for studying community detection due to its versatility and analytic tractability. It however falls short of abstracting a key aspect of real-world networks: an unrealistic assumption of SBM is that within each community, the degree distributions of each node are the same and have light tails (Poisson or Gaussian tails). In empirical network data sets, however, the degree distributions are often highly inhomogeneous across nodes, sometimes exhibiting a heavy tail behavior with some nodes having very high degrees (so-called hubs). At the same time, sparsely connected nodes with small degrees are also common in real networks. To overcome this shortcoming of the SBM, the degree-corrected stochastic block model (DCSBM) was introduced in the literature to allow for degree heterogeneity within communities, thereby providing a more flexible and accurate model of real-world networks [17, 29].

A number of community detection methods have been proposed under DCSBM. Several methods are model-based, including profile likelihood maximization and modularity maximization [29, 38]. Although these methods enjoy certain statistical guarantees [48], they often involve optimization over all possible partitions, which is computationally intractable. Recent work in [5, 32] discusses efficient solvers, but theoretical guarantees are only established under restricted settings such as those with two communities. Another popular class of algorithms are spectral methods, which estimate the communities using the graph eigenvectors and are often computationally fast. Statistical guarantees are derived for spectral methods under certain settings (see, e.g., [11, 15, 17, 23, 28, 34, 41]), but numerical validation on synthetic and real data has not been as thorough. One notable exception is the SCORE method in [28], which achieved one of the best known performances on the political blogs dataset from [2]. Spectral methods are also known to suffer from inconsistency in sparse graphs [30] as well as sensitivity to outliers [9]. See Section 5 for more detailed discussion of the literature.

In this paper, we seek for a clustering algorithm that is computationally feasible, has strong statistical performance guarantees under DCSBM, and provides competitive empirical performance. Our approach makes use of the robustness and computational power of convex optimization. Under the standard SBM, methods based on convex optimization have been proven to be statistically efficient under a broad range of model parameters; see, for example, [4, 6, 13, 14, 22, 39]. A recent line of work shows that these methods can in fact achieve the optimal recovery thresholds under SBM [1, 3, 7, 24–26, 37, 40]. In the work [9], convex methods are proved to be robust against arbitrary outlier nodes, and moreover shown to achieve state-of-the-art misclassification rates in the political blogs dataset, in which the node degrees are highly heterogeneous. These observations motivate us to study whether strong theoretical guarantees under DCSBM can be established for convex-optimization-based methods.

Building on the work of [13] and [9], we introduce in Section 2 a new community detection approach called Convexified Modularity Maximization (CMM). CMM is based on convexifying the elegant modularity maximization formulation,

followed by a novel and computationally tractable weighted ℓ_1-norm k-medoids clustering procedure. As we show in Section 3 and Section 4, our approach has strong theoretical guarantees, applicable even in the sparse graph regime with bounded average degree, and at the same time enjoys state-of-the-art empirical performance. In both aspects, our approach is comparable to or improves upon the best-known results in the literature.

Computationally, CMM involves solving semidefinite and linear programs, whose time complexity is polynomial in the number of nodes n [a generous bound is $O(n^6)$]. Empirically, our implementation of CMM (see Section 4) is much faster than the theoretical bound suggests. Though still not as fast as spectral methods and conditional pseudo-likelihood maximization, convex optimization based methods continue to benefit from the advances of fast LP/SDP solvers. We note that the convergence of greedy methods such as profile likelihood maximization [29] and conditional pseudo-likelihood maximization [5] are not established in general, so their theoretical computational complexity is in fact unknown.

2. Problem setup and algorithms.
In this section, we set up the community detection problem under DCSBM, and describe our algorithm based on convexified modularity maximization and weighted k-medoids clustering. Throughout this paper, we use lower-case and upper-case bold letters such as $\boldsymbol{u}$ and $\boldsymbol{U}$ to represent vectors and matrices, respectively, with u_i and U_{ij} denoting their elements. We let $\boldsymbol{U}_{i\bullet}$ denote the ith row of the matrix $\boldsymbol{U}$, and $\boldsymbol{U}_{\bullet j}$ the jth column vector of $\boldsymbol{U}$ If all coordinates of a vector $\boldsymbol{v}$ are nonnegative, we write $\boldsymbol{v} \geq \boldsymbol{0}$. The notation $\boldsymbol{v} > \boldsymbol{0}$, as well as $\boldsymbol{U} \geq 0$ and $\boldsymbol{U} > 0$ for matrices, are defined similarly. For a symmetric matrix $\boldsymbol{U} \in \mathbb{R}^{n \times n}$, we write $\boldsymbol{U} \succ \boldsymbol{0}$ if $\boldsymbol{U}$ is positive definite, and $\boldsymbol{U} \succeq \boldsymbol{0}$ if it is positive semidefinite. For any sequences $\{a_n\}$ and $\{b_n\}$, we write $a_n \lesssim b_n$ if there is an absolute constant $c > 0$ such that $a_n/b_n \leq c$, $\forall n$, and we define $a_n \gtrsim b_n$ similarly.

2.1. *The degree-corrected stochastic block model.* In DCSBM, a graph $\mathcal{G}$ is generated randomly as follows. A total of n nodes, which we identify with the set $[n] := \{1, \ldots, n\}$, are partitioned into r fixed but unknown clusters $C_1^*, C_2^* \ldots, C_r^*$. Each pair of distinct nodes $i \in C_a^*$ and $j \in C_b^*$ are connected by an (undirected) edge with probability $\theta_i \theta_j B_{ab} \in [0, 1]$, independently of all others. Here, the vector $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_n)^\top \in \mathbb{R}_+^n$ is referred to as the *degree heterogeneity parameters* of the nodes, and the symmetric matrix $\boldsymbol{B} \in \mathbb{R}_+^{r \times r}$ is called the *connectivity matrix* of the clusters. Note that if $\theta_i = 1$ for all nodes i, DCSBM reduces to the classical SBM. Given a single realization of the resulting random graph $\mathcal{G} = ([n], E)$, the statistical goal is to estimate the true clusters $\{C_a^*\}_{a=1}^r$.

Before describing our algorithms, let us first introduce some useful notation. Denote by $\boldsymbol{A} \in \{0, 1\}^{n \times n}$ the *adjacency matrix* associated with the graph $\mathcal{G}$, with $A_{ij} = 1$ if and only if nodes i and j are connected. For each candidate partition of n nodes into r clusters, we associate it with a *partition matrix* $\boldsymbol{Y} \in \{0, 1\}^{n \times n}$, such

that $Y_{ij} = 1$ if and only if nodes i and j are assigned to the same cluster, with the convention that $Y_{ii} = 1$, $\forall i$. Let $\mathbb{P}_{n,r}$ be the set of all such partition matrices, and Y^* the true partition matrix associated with the ground-truth clusters $\{C_a^*\}_{a=1}^r$. The notion of partition matrices plays a crucial role in the subsequent discussion.

2.2. *Generalized modularity maximization.* Our clustering algorithm is based on Newman and Girvan's classical notion of *modularity* (see, e.g., [38]). Given the graph adjacency matrix A of n nodes, the modularity of a partition represented by the partition matrix $Y \in \bigcup_r \mathbb{P}_{n,r}$, is defined as

$$(2.1) \qquad Q(Y) := \sum_{1 \le i, j \le n} \left(A_{ij} - \frac{d_i d_j}{2L} \right) Y_{ij},$$

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. The *modularity maximization* approach to community detection is based on finding a partition Y_m that optimizes $Q(Y)$:

$$(2.2) \qquad Y_m \leftarrow \arg \max_{Y \in \bigcup_r \mathbb{P}_{n,r}} Q(Y).$$

This standard form of modularity maximization is known to suffer from a "resolution limit" and cannot detect small clusters [19]. To address this issue, several authors have proposed to replace the normalization factor $\frac{1}{2L}$ by a tuning parameter λ [31, 42], giving rise to the following generalized formulation of modularity maximization:

$$(2.3) \qquad Y_m \leftarrow \arg \max_{Y \in \bigcup_r \mathbb{P}_{n,r}} Q_\lambda(Y) := \sum_{1 \le i, j \le n} (A_{ij} - \lambda d_i d_j) Y_{ij}.$$

While modularity maximization enjoys several desirable statistical properties under SBM and DCSBM [48], the associated optimization problems (2.2) and (2.3) are not computationally feasible due to the combinatorial constraint, which limits the practical applications of these formulations. In practice, modularity maximization is often used as a guidance for designing heuristic algorithms [18], Section VI.

Here, we take a more principled approach to computational feasibility while maintaining provable statistical guarantees: we develop a tractable convex surrogate for the above combinatorial optimization problems, whose solution is then refined by a novel weighted k-medoids algorithm.

2.3. *Convex relaxation.* Introducing the degree vector $d = (d_1, \ldots, d_n)^\top$, we can rewrite the generalized modularity maximization problem (2.3) in matrix form as

$$(2.4) \qquad \begin{aligned} \max_{Y} \quad & \langle Y, A - \lambda d d^\top \rangle \\ \text{subject to} \quad & Y \in \bigcup_r \mathbb{P}_{n,r}, \end{aligned}$$

where $\langle \cdot, \cdot \rangle$ denotes the trace inner product between matrices. The objective function is linear in matrix variable $\boldsymbol{Y}$, so it suffices to convexify the combinatorial constraint $\boldsymbol{Y} \in \bigcup_r \mathbb{P}_{n,r}$.

Recall that each matrix $\boldsymbol{Y}$ in $\mathbb{P}_{n,r}$ corresponds to a unique partition of n nodes into r clusters. There is another representation of such a partition via a *membership matrix* $\boldsymbol{\Psi} \in \{0, 1\}^{n \times r}$, where $\Psi_{ia} = 1$ if and only if node i belongs to cluster a. These two representations are related by the identity

$$(2.5) \qquad \boldsymbol{Y} = \boldsymbol{\Psi} \boldsymbol{\Psi}^\top,$$

which implies that $\boldsymbol{Y} \succeq \boldsymbol{0}$. The membership matrix of a partition is only unique up to permutation of the cluster labels $1, 2, \ldots, r$, so each partition matrix $\boldsymbol{Y}$ corresponds to multiple membership matrices $\boldsymbol{\Psi}$. We use $\mathbb{M}_{n,r}$ to denote the set of all possible membership matrices of r-partitions.

Besides being positive semidefinite, a partition matrix $\boldsymbol{Y}$ also satisfies the linear constraints $0 \leq Y_{ij} \leq 1$ and $Y_{ii} = 1$ for all $i, j \in [n]$. Using these properties of partition matrices, we obtain the following convexification of the modularity optimization problem (2.4):

$$
\begin{aligned}
\widehat{\boldsymbol{Y}} = \arg\max_{\boldsymbol{Y}} \quad & \langle \boldsymbol{Y}, \boldsymbol{A} - \lambda \boldsymbol{d}\boldsymbol{d}^\top \rangle \\
\text{subject to } \quad & \boldsymbol{Y} \succeq \boldsymbol{0}, \\
& \boldsymbol{0} \leq \boldsymbol{Y} \leq \boldsymbol{J}, \\
& Y_{ii} = 1, \qquad \text{for each } i \in [n].
\end{aligned}
$$

(2.6)

Here, $\boldsymbol{J}$ is the $n \times n$ matrix with all entries equal to 1. Implementation of the formulation (2.6) requires choosing an appropriate tuning parameter λ. The theoretical range for λ for consistent clustering is given in Section 3; for all our numerical experiments in Section 4, λ is set to be $\langle \boldsymbol{A}, \boldsymbol{J} \rangle^{-1}$. As our convexification is based on the generalized version (2.3) of modularity maximization, it is capable of detecting small clusters, even when the number of clusters r grows with n, as is shown later.

2.4. *Explicit clustering via weighted k-medoids.* Ideally, the optimal solution $\widehat{\boldsymbol{Y}}$ to the convex relaxation (2.6) is a valid partition matrix in $\mathbb{P}_{n,r}$ and recovers the true partition $\boldsymbol{Y}^*$ perfectly—our theoretical results in Section 3.3 characterize when this happens. In general, the solution $\widehat{\boldsymbol{Y}}$ may not lie in $\mathbb{P}_{n,r}$, but we expect it to be close to $\boldsymbol{Y}^*$. To extract an explicit clustering from $\widehat{\boldsymbol{Y}}$, we introduce a novel and tractable *weighted k-medoids* algorithm.

Recall that by definition, the ith and jth rows of the true partition matrix $\boldsymbol{Y}^*$ are identical if the corresponding nodes i and j belong to the same community, and otherwise orthogonal to each other. If $\widehat{\boldsymbol{Y}}$ is close to $\boldsymbol{Y}^*$, intuitively one can extract a good partition by clustering the row vectors of $\widehat{\boldsymbol{Y}}$ as points in the Euclidean space $\mathbb{R}^n$. While there exist numerous algorithms (e.g., k-means) for such a task, our analysis identifies a particularly viable choice—a k-medoids procedure

appropriately weighted by the node degrees—that is efficient both theoretically and empirically.

Specifically, our weighted k-medoids procedure consists of two steps. First, we multiply the columns of $\widehat{Y}$ by the corresponding degrees to obtain the matrix $\widehat{W} := \widehat{Y}D$, where $D := \operatorname{diag}(d) = \operatorname{diag}(d_1, \dots, d_n)$, which is the diagonal matrix formed by the entries of d. Clustering is performed on the row vectors of $\widehat{W}$ instead of $\widehat{Y}$. Note that if we consider the ith row of $\widehat{Y}$ as a vector of n features for node i, then the rows of $\widehat{W}$ can be thought of as vectors of *weighted* features.

Then we implement a weighted k-medoids clustering on the row vectors of $\widehat{W}$. Denoting by $\widehat{w}_i$ the ith row of $\widehat{W}$, we search for a partition $C_1, \dots, C_r$ of $[n]$ and r cluster medoids $x_1, \dots, x_r \in \{\widehat{w}_i : i \in [n]\}$ in order to minimize the degree-weighted sum of ℓ_1-distances between each w_i and the corresponding center. To be specific, we want to solve the following optimization:

$$
\begin{aligned}
(2.7) \quad & \min_{\{C_a\}_{a=1}^r, \{x_a\}_{a=1}^r} \quad \sum_{1 \le a \le r} \sum_{i \in C_a} d_i \|\widehat{w}_i - x_a\|_1 \\
& \text{s.t.} \quad\quad\quad x_a \in \operatorname{Rows}(\widehat{W}), \quad \forall a = 1, \dots, r,
\end{aligned}
$$

where $\operatorname{Rows}(Z)$ denotes the collection of row vectors of a matrix Z. Representing the partition $\{C_a\}_{a=1}^r$ by a membership matrix $\Psi \in \mathbb{M}_{n,r}$ and the centers $\{x_i\}$ as the rows of a matrix $X \in \mathbb{R}^{r \times n}$, (2.7) is equivalent to

$$
\begin{aligned}
(2.8) \quad & \min_{\Psi, X} \quad \|D(\Psi X - \widehat{W})\|_1 \\
& \text{s.t.} \quad \Psi \in \mathbb{M}_{n,r}, \\
& \quad\quad X \in \mathbb{R}^{r \times n}, \quad \operatorname{Rows}(X) \subseteq \operatorname{Rows}(\widehat{W}),
\end{aligned}
$$

where $\|Z\|_1$ denotes the sum of the absolute values of all entries of Z.

We emphasize that the formulation (2.8) differs from standard clustering algorithms (such as k-means) in several ways. The objective function is the sum of distances rather than that of squared distances, and the distances are in ℓ_1 instead of ℓ_2 norms. Moreover, our formulation has two levels of weighting: each column of $\widehat{Y}$ is weighted to form $\widehat{W}$, and the distance of each row w_i to its cluster center is further weighted by d_i. This doubly-weighted ℓ_1-norm k-medoids formulation is crucial in obtaining strong and robust statistical bounds, and is significantly different from previous approaches, such as those in [23, 34] (which only use the second weighting, and the weights are *inversely* proportional to d_i). Our doubly weighting scheme is motivated by the intuition that nodes with larger degrees tend to be clustered more accurately. This intuition is made concrete in the next section when we study the solution $\widehat{Y}$ of the convex relaxation (2.6)—our analysis naturally leads to a doubly weighted ℓ_1 error bound (Theorem 3.1), with bounds on the unweighted error following as a consequence (Corollary 3.2).

With the constraint $\operatorname{Rows}(X) \subseteq \operatorname{Rows}(\widehat{W})$, the optimization problem (2.8) is precisely the weighted ℓ_1-norm k-medoids problem considered in [10]. Computing

the exact optimizer to (2.8), denoted by $(\overline{\Psi}, \overline{X})$, is NP-hard. Nevertheless, [10] provides a polynomial-time approximation algorithm, which outputs a solution $(\breve{\Psi}, \breve{X})$ feasible to (2.8) and provably satisfying

$$(2.9) \qquad \left\| D(\breve{\Psi}\breve{X} - \widehat{W}) \right\|_1 \le \frac{20}{3} \left\| D(\overline{\Psi X} - \widehat{W}) \right\|_1.$$

As the solution $\widehat{Y}$ to the convex relaxation (2.6) and the approximate solution $\breve{\Psi}$ to the k-medoids problem (2.8) can both be computed in polynomial-time, our algorithm is computationally tractable. In the next section, we turn to the statistical aspect and show that the clustering induced by $\widehat{Y}$ and $\breve{\Psi}$ is close to the true underlying clusters, under some mild and interpretable conditions of DCSBM.

3. Theoretical results. In this section, we provide theoretical results characterizing the statistical properties of our algorithm. We show that under mild conditions of DCSBM, the difference between the convex relaxation solution $\widehat{Y}$ and the true partition matrix Y^*, and the difference between the approximate k-medoids clustering $\breve{\Psi}$ and the true clustering Ψ^*, are well bounded. When additional conditions hold, we further show that $\widehat{Y}$ perfectly recovers the true clusters. Our results are nonasymptotic in nature, valid for any scaling of n, r, θ and B, etc.

In this section, we assume without loss of generality that $\max_{1 \le i \le n} \theta_i = 1$. In general, one can multiply θ by a scalar c and divide B by c^2 without changing the distribution of the graph.

3.1. *Density gap conditions.* In the literature of community detection by convex optimization under standard SBM, it is often assumed that the minimum within-cluster edge density is greater than the maximum cross-cluster edge density, that is,

$$(3.1) \qquad \max_{1 \le a < b \le r} B_{ab} < \min_{1 \le a \le r} B_{aa}.$$

See, for example, [4, 9, 13, 22, 39]. This requirement (3.1) can be directly extended to the DCSBM setting, leading to the condition

$$(3.2) \qquad \max_{1 \le a < b \le r} \max_{i \in C_a^*, j \in C_b^*} B_{ab}\theta_i\theta_j < \min_{1 \le a \le r} \min_{i,j \in C_a^*, i \ne j} B_{aa}\theta_i\theta_j.$$

However, this condition would often be overly restrictive under DCSBM, particularly when the degree parameters $\{\theta_i\}$ are imbalanced with some of them being very small. In particular, this condition is highly sensitive to the minimum value $\theta_{\min} := \min_{1 \le i \le n} \theta_i$, which is unnecessary since the community memberships of nodes with larger θ_i may still be recoverable.

The above observation motivates us to seek a version of the density gap condition that is milder and more appropriate for DCSBM. For each cluster index $1 \le a \le r$, define the quantities

$$(3.3) \qquad G_a := \sum_{i \in C_a^*} \theta_i \quad \text{and} \quad H_a := \sum_{b=1}^r B_{ab}G_b.$$

Simple calculation gives

$$\mathbb{E} d_i = \theta_i H_a - \theta_i^2 B_{aa} \approx \theta_i H_a.$$

Therefore, the quantity H_a controls the *average degree* of the nodes in the ath cluster. With this notation, we consider a particular condition, called the *degree-corrected density gap condition*, as follows:

$$(3.4) \qquad \max_{1 \le a < b \le r} \frac{B_{ab}}{H_a H_b} < \min_{1 \le a \le r} \frac{B_{aa}}{H_a^2}.$$

This condition can be viewed as the "average" version of (3.2), as it depends on the aggregate quantity H_a associated with each cluster a rather than the θ_i's of individual nodes—in particular, the condition (3.4) is robust against small $\theta_{\min}$. Moreover, the condition (3.4) is invariant under equivalent DCSBM s (i.e., those obtained by rescaling $\boldsymbol{B}$ and $\boldsymbol{\theta}$). This degree-corrected density gap condition plays a key role throughout our theoretical analysis, for both approximate and exact cluster recovery under DCSBM.

To gain intuition on the degree-corrected density gap condition, consider the following sub-class of DCSBM with symmetric and balanced clusters.

DEFINITION 3.1. We say that a DCSBM obeys a $\mathcal{F}(n, r, p, q, g)$-model, if $B_{aa} = p$ for all $a = 1, \ldots, r$, $B_{ab} = q$ for all $1 \le a < b \le r$, and $G_1 = G_2 = \cdots = G_r = g$.

In a $\mathcal{F}(n, r, p, q, g)$-model, the true clusters are balanced in terms of the connectivity matrix $\boldsymbol{B}$ and the sum of the degree heterogeneity parameters (rather than the cluster size). Under this model, straightforward calculation gives $H_a = ((r-1)q + p)g$ for all $a = 1, \ldots, r$. The degree-corrected density gap condition (3.4) then reduces to $p > q$, that is, the classical density gap condition (3.1).

3.2. *Theory of approximate clustering.* We now study when the solutions to our convex relaxation (2.6) and weighted k-medoids algorithms (2.8) approximately recover the underlying true clusters. Under DCSBM, nodes with different θ_i's have varying degrees and, therefore, contribute differently to the overall graph and in turn to the clustering quality. Such heterogeneity needs to be taken into account in order to get tight bounds on clustering errors. The following version of ℓ_1 norm, corrected by the degree heterogeneity parameters, is the natural notion of an error metric.

DEFINITION 3.2. For a matrix $\boldsymbol{Z} \in \mathbb{R}^{n \times n}$, its weighted elementwise ℓ_1 norm is defined as

$$\|\boldsymbol{Z}\|_{1,\boldsymbol{\theta}} := \sum_{i=1}^{n} \sum_{j=1}^{n} |\theta_i Z_{ij} \theta_j|.$$

Also recall our definitions of H_a and G_a in equation (3.3). Furthermore, define the vector of (approximate) expected degrees $\boldsymbol{f} \in \mathbb{R}^n$ such that

$$(3.5) \qquad f_i := \theta_i H_a, \qquad \forall 1 \le a \le r, i \in C_a^*.$$

With the notation above, our first theorem shows that the convex relaxation solution $\widehat{\boldsymbol{Y}}$ is close to the true partition matrix $\boldsymbol{Y}^*$ in terms of the weighted ℓ_1 norm.

THEOREM 3.1. *Under* DCSBM, *assume that the degree-corrected density gap condition* (3.4) *holds. Moreover, suppose that the tuning parameter λ in the convex relaxation* (2.6) *satisfies*

$$(3.6) \qquad \max_{1 \le a < b \le r} \frac{B_{ab} + \delta}{H_a H_b} \le \lambda \le \min_{1 \le a \le r} \frac{B_{aa} - \delta}{H_a^2}$$

for some number $\delta > 0$. Then with probability at least $0.99 - 2(e/2)^{-2n}$, the solution $\widehat{\boldsymbol{Y}}$ to the convex relaxation (2.6) *satisfies the bound*

$$(3.7) \qquad \|\boldsymbol{Y}^* - \widehat{\boldsymbol{Y}}\|_{1,\boldsymbol{\theta}} \le \frac{C_0}{\delta}(1 + \lambda \|\boldsymbol{f}\|_1)(\sqrt{n\|\boldsymbol{f}\|_1} + n),$$

where $C_0 > 0$ is an absolute constant and $\|\boldsymbol{f}\|_1 = \sum_{a,b} B_{ab} G_a G_b$.

We prove this claim in Section 7.1. Our analysis is inspired by the seminal work in [22], and makes use of the Grothendieck's inequality [21, 35] to obtain uniform deviation bounds. Notably, the bound (3.7) is insensitive to $\theta_{\min}$ as should be expected, because community memberships of nodes with relatively large θ_i are still recoverable. In contrast, the error bounds of several existing methods, such as that of SCORE method in [28], equations (2.15), (2.16), depend on $\theta_{\min}$ crucially.

Under the $\mathcal{F}(n, r, p, q, g)$-model, recall that $H_a \equiv (p + (r-1)q)g$ and density gap condition (3.4) becomes $p > q$. Moreover, the constraint (3.6) for δ and λ becomes

$$(3.8) \qquad p - q \ge 2\delta \quad \text{and} \quad \frac{q + \delta}{(p + (r-1)q)^2 g^2} \le \lambda \le \frac{p - \delta}{(p + (r-1)q)^2 g^2}.$$

Note that the first inequality above is the same as the standard density gap condition imposed in, for example, [9, 13, 14]. Furthermore, the vector $\boldsymbol{f}$ satisfies $\|\boldsymbol{f}\|_1 = r(p + (r-1)q)g^2 \le r^2 pg^2$. Substituting these expressions into the bound (3.7), we obtain the following corollary for the symmetric DCSBM setting.

COROLLARY 3.1. *Under the $\mathcal{F}(n, r, p, q, g)$-model of* DCSBM, *if the condition* (3.8) *holds for the density gap and tuning parameter, then with probability at least $0.99 - 2(e/2)^{-2n}$, the solution $\widehat{\boldsymbol{Y}}$ to the convex relaxation* (2.6) *satisfies the bound*

$$(3.9) \quad \|\boldsymbol{Y}^* - \widehat{\boldsymbol{Y}}\|_{1,\boldsymbol{\theta}} \lesssim \frac{1}{\delta}\left(1 + \frac{rp}{(p + (r-1)q)}\right)(n + rg\sqrt{np}) \lesssim \frac{1}{\delta}r(n + rg\sqrt{np}).$$

Note that if $\frac{p}{q} = c$ for an absolute constant c, then the first inequality in bound (3.9) takes the simpler form $\|\boldsymbol{Y}^* - \widehat{\boldsymbol{Y}}\|_{1,\theta} \lesssim \frac{n+rg\sqrt{np}}{\delta}$. If $\theta_i = 1$ for all nodes i, the $\mathcal{F}(n, r, p, q, g)$-model reduces to the standard SBM with equal community size. If we assume $r = O(1)$ additionally, and note that $g = n/r$ and let $\delta = \frac{p-q}{4}$, then the error bound (3.9) becomes

$$\|\boldsymbol{Y}^* - \widehat{\boldsymbol{Y}}\|_1 \lesssim \frac{n(1 + \sqrt{np})}{p - q}.$$

This bound matches the error bounds given in [22], Theorem 1.3.

The output $\widehat{\boldsymbol{Y}}$ of the convex relaxation needs not be a partition matrix corresponding to a clustering; a consequence is that the theoretical results in [22] do not provide an explicit guarantee on clustering errors (except for the special case of $r = 2$). We give such a bound below, based on the explicit clustering extracted from $\widehat{\boldsymbol{Y}}$ using the weighted ℓ_1-norm k-medoids algorithm (2.8). Recall that $\breve{\boldsymbol{\Psi}}$ is the membership matrix in the approximate k-medoids solution given in (2.9), and let $\boldsymbol{\Psi}^*$ be the membership matrix corresponding to the true clusters. A membership matrix is unique only up to permutation of its columns (i.e., relabeling the clusters), so counting the misclassified nodes in $\breve{\boldsymbol{\Psi}}$ requires an appropriate minimization over such permutations. This motivates the following definition; here for a matrix $\boldsymbol{M}$, $\boldsymbol{M}_{i\bullet}$ denotes its ith row vector.

DEFINITION 3.3.   Let $\mathcal{S}_r$ denote the set of all $r \times r$ permutation matrices. The set of misclassified nodes with respect to a permutation matrix $\boldsymbol{\Pi} \in \mathcal{S}_r$ is defined as

$$\mathcal{E}(\boldsymbol{\Pi}) := \{i \in [n] : (\breve{\boldsymbol{\Psi}}\boldsymbol{\Pi})_{i\bullet} \neq \boldsymbol{\Psi}^*_{i\bullet}\}.$$

With this definition, we have the following theorem that quantifies the misclassification rate of the weighted k-medoids procedure, whose input is the optimal solution $\widehat{\boldsymbol{Y}}$ of the convex relaxation.

THEOREM 3.2.   *Under the $\mathcal{F}(n, r, p, q, g)$-model, assume that the parameters δ and λ satisfy (3.8). With probability at least $0.99 - 2(e/2)^{-2n}$ and for an absolute constant C_0, the approximate k-medoids solution $\breve{\boldsymbol{\Psi}}$ satisfies*

$$(3.10) \qquad \min_{\boldsymbol{\Pi} \in \mathcal{S}_r} \left\{ \sum_{i \in \mathcal{E}(\boldsymbol{\Pi})} \theta_i \right\} \leq C_0 \frac{r}{\delta} \left(\frac{n}{g} + r\sqrt{np} \right).$$

We prove this claim in Section 7.4. The proof is carried out by establishing two general results of independent interest on the errors of k-medoids clustering (given in Section 7.2 and Section 7.3), and combining them with Theorem 3.1 which bounds the error rate of $\widehat{\boldsymbol{Y}}$, the input to k-medoids.

If we let $\mathbf{\Pi_\theta}$ be a minimizer of the LHS of (3.10) and $\mathcal{E_\theta} := \mathcal{E}(\mathbf{\Pi_\theta})$, then the quantity $\sum_{i \in \mathcal{E_\theta}} \theta_i$ is the number of misclassified nodes *weighted* by their degree heterogeneity parameters $\{\theta_i\}$. Theorem 3.2 controls this weighted quantity. Notably, the bound given in (3.10) is applicable all the way down to the *sparse graph regime* with bounded average degrees, that is, $p, q = O(1/n)$. For example, suppose that $p = a/n$ and $q = b/n$ for two fixed constants $a > b, r = O(1)$ and $g \asymp n$; if $(a - b)/\sqrt{a}$ is sufficiently large, then with the choice $\delta \asymp (a - b)/n$, the right-hand side of (3.10) can be an arbitrarily small constant times n. In comparison, conventional spectral methods are known to be inconsistent in this sparse regime [30]. While this difficulty is alleviated under SBM by the use of regularization or nonbacktracking matrices (e.g., [8, 33]), rigorous justification and numerical validation under DCSBM have not been well explored.

The general weighted bound (3.10) in fact implies direct (unweighted) bounds on the number misclassified nodes, namely $\min_{\mathbf{\Pi} \in \mathcal{S}_r} \{\sum_{i \in \mathcal{E}(\mathbf{\Pi})} 1\} = \min_{\mathbf{\Pi} \in \mathcal{S}_r} |\mathcal{E}(\mathbf{\Pi})|$, as given in the following corollary.

COROLLARY 3.2. *Under the same assumptions as Theorem* 3.2, *the approximate k-medoids solution* $\breve{\mathbf{\Psi}}$ *satisfies the bounds*

$$(3.11) \quad \min_{\mathbf{\Pi} \in \mathcal{S}_r} |\mathcal{E}(\mathbf{\Pi})| \le \sqrt{\frac{C_0}{\delta g} r (n + r g \sqrt{np}) \sum_{i=1}^{n} \frac{1}{\theta_i}}, \quad and$$

$$(3.12) \quad \min_{\mathbf{\Pi} \in \mathcal{S}_r} |\mathcal{E}(\mathbf{\Pi})| \le \frac{C_0}{\delta g \tau} r (n + r g \sqrt{np}) + |\{i : \theta_i < \tau\}|, \qquad \forall \tau > 0.$$

PROOF. By definition $\min_{\mathbf{\Pi} \in \mathcal{S}_r} |\mathcal{E}(\mathbf{\Pi})| \le \mathcal{E}(\mathbf{\Pi_\theta}) = |\mathcal{E_\theta}|$. Furthermore, by the Cauchy–Schwarz inequality, $|\mathcal{E_\theta}| \le \sqrt{\sum_{i \in \mathcal{E_\theta}} \theta_i \sum_{i \in \mathcal{E_\theta}} \frac{1}{\theta_i}}$. Alternatively, for any $\tau > 0$, $|\mathcal{E_\theta}| \le \frac{1}{\tau} \sum_{i \in \mathcal{E_\theta}} \theta_i + |\{i : \theta_i < \tau\}|$. The corollary then readily follows by plugging (3.10) into the last two inequalities. $\square$

Inequality (3.12), in fact a collection of bounds indexed by τ, is insensitive to $\theta_{\min}$. For illustration, consider the following $\mathcal{F}(n, r, p, q, g)$ model of DCSBM, where $\theta_1 = \theta_{\min}, \theta_2 = 1 - \theta_{\min}, \theta_3 = \cdots = \theta_n = \frac{1}{2}$ for some small $\theta_{\min} < \frac{1}{2}$, nodes 1 2 are from the same cluster, and $p \ge 1/n$. In this case, $g = n/(2r)$. By setting $\delta = (p - q)/4$ and $\tau = 1/2$ in (3.12), we get that the number of misclassified nodes satisfies the bound $\min_{\mathbf{\Pi} \in \mathcal{S}_r} |\mathcal{E}(\mathbf{\Pi})| = O(r^2 \sqrt{np}/(p - q)) + 1$, which is not affected if the parameter $\theta_1 = \theta_{\min}$ of node 1 is very small. Intuitively, if a few nodes have small θ_i's, then they cannot be recovered but do not impact the recovery of the other nodes. In comparison, the error bounds in several existing works depend on $\theta_{\min}$ crucially. For example, the error bound of the SCORE method ([28], equations (2.15) and (2.16)), scales with $1/\theta_{\min}$ and approaches infinity when $\theta_{\min} \to 0$.

In the special case of standard SBM with $\theta_i \equiv 1$, $p \ge \frac{1}{n}$ and r equal-sized clusters, setting $\tau = 0.9$ in (3.12) yields that the number of misclassified nodes is

$O(\frac{r^2\sqrt{np}}{p-q})$. This bound is consistent with those in [22] when $r = 2$; it is in fact more general as it applies to $r \geq 3$ clusters.

3.3. *Theory of perfect clustering.* In this section, we show that under an additional condition on the minimum degree heterogeneity parameter $\theta_{\min} = \min_{1 \leq i \leq n} \theta_j$, the solution $\widehat{Y}$ to the convex relaxation perfectly recovers the true partition matrix Y^*. In this case, the true clusters can be extracted easily from $\widehat{Y}$ without using the k-medoids procedure.

For the purpose of studying perfect clustering, we consider a setting of DCSBM with $B_{aa} = p$ for all $a = 1, \ldots, r$, and $B_{ab} = q$ for all $1 \leq a < b \leq r$. Under this setup, the degree-corrected density gap condition (3.6) becomes

$$(3.13) \qquad \max_{1 \leq a < b \leq r} \frac{q + \delta}{H_a H_b} \leq \lambda \leq \min_{1 \leq a \leq r} \frac{p - \delta}{H_a^2}.$$

Recalling the definition of G_a in (3.3), we further define $G_{\min} := \min_{1 \leq a \leq r} G_a$. The following theorem, proved in the Supplementary Material [12], provides a sufficient condition for perfect clustering.

THEOREM 3.3. *Suppose that the degree-corrected density gap condition* (3.13) *is satisfied for some number $\delta > 0$ and tuning parameter λ, and that*

$$(3.14) \qquad \delta > C_0 \left(\frac{\sqrt{qn}}{G_{\min}} + \sqrt{\frac{p \log n}{G_{\min} \theta_{\min}}} \right)$$

for some sufficiently large absolute constant C_0. With probability at least $1 - 10n^{-1}$, the convex relaxation (2.6) *has a unique optimal solution $\widehat{Y} = Y^*$.*

The condition (3.14) depends on the minimum values $G_{\min}$ and $\theta_{\min}$. Such dependence is necessary for perfect clustering, as clusters and nodes with overly small G_a and θ_i will have too few edges and are not recoverable. In comparison, the approximate recovery results in Theorem 3.1 are not sensitive to either $\theta_{\min}$ or $G_{\min}$, as should be expected. Valid for the more general DCSBM, Theorem 3.3 significantly generalizes the existing theory for standard SBM on perfect clustering by SDP in the literature (see, e.g., [9, 13, 14]). Taking $n \to \infty$, Theorem 3.3 guarantees that the probability of perfect clustering converges to one, thereby implying the convex relaxation approach is *strongly consistent* in the sense of [48].

In the special case of standard SBM with $\theta_i = 1, \forall i \in [n]$, the density gap lower bound (3.14) simplifies to

$$\delta \gtrsim \frac{\sqrt{qn}}{\ell_{\min}} + \sqrt{\frac{p \log n}{\ell_{\min}}},$$

where $\ell_{\min} := \min_{1 \leq a \leq r} \ell_a$ is the minimum community size and $\ell_a := |C_a^*|$ is the size of community a. This density gap lower bound is consistent with best existing

results given in [9, 13, 14]—as we discussed earlier, our density condition in (3.8) under the $\mathcal{F}(n, r, p, q, g)$ model (which encompasses SBM with equal-sized clusters) is the same as in these previous papers, with the minor difference that in these papers the term $\frac{\lambda}{d_i d_j}$ in the convex relaxation (2.6) is replaced by a tuning parameter λ' assumed to satisfy the condition $q + \delta \leq \lambda' \leq p - \delta$.

4. Numerical results.
In this section, we provide numerical results on both synthetic and real datasets, which corroborate our theoretical findings.

The convexified modularity maximization problem (2.6) is a semidefinite program (SDP), and can be solved efficiently by a range of general and specialized algorithms. Here, we use the alternating direction method of multipliers (ADMM) suggested in [9]. To specify the ADMM solver, we need some additional notations. For two $n \times n$ matrices X and Y, let $\max\{X, Y\}$ be the matrix whose (i, j)th entry equals $\max\{X_{ij}, Y_{ij}\}$; the matrix $\min\{X, Y\}$ is similarly defined. For a symmetric matrix X with eigenvalue decomposition $X = U \Sigma U^\top$, let $(X)_+ := U \max\{\Sigma, 0\} U^\top$, and let $(X)_I$ be the matrix obtained by setting the diagonal entries of X to 1. Recall that J is the $n \times n$ all-one matrix. Algorithm 1 gives the ADMM procedure for solving (2.6) with the dual update step size equal to 1.

Our choice of the tuning parameter $\lambda = \langle A, J \rangle^{-1}$ is motivated by the following simple observation. By standard concentration inequalities, the number $\langle A, J \rangle$ is close to its expectation $\sum_i \mathbb{E}[d_i] \approx \|f\|_1$. Under the $\mathcal{F}(n, r, p, q, g)$-model, we have $\|f\|_1 = r(p + (r - 1)q)g^2$ and $H_a = ((r - 1)q + p)g$ for all $a \in [r]$. In this case and with the above choice of λ, the density gap assumption (3.13) simplifies to $\frac{q+\delta}{(r-1)q+p} \leq \frac{1}{r} \leq \frac{p-\delta}{(r-1)q+p}$, which holds with $\delta = (p - q)/r$.

After obtaining the solution $\widehat{Y}$ of the convex relaxation, we extract an explicit clustering using the weighted k-medoids procedure in (2.8) with $k = r$, where the number of major clusters r is assumed known. Our complete community detection algorithm, Convexified Modularity Maximization (CMM), is summarized in Al-

---

Algorithm 1 ADMM algorithm for solving the SDP (2.6)

1: Input: A and $\lambda = \langle A, J \rangle^{-1}$.
2: Initialization: $Z^{(0)} = \Lambda^{(0)} = 0$, $k = 0$ and $\texttt{MaxIter} = 100$.
3: while $k < \texttt{MaxIter}$

    1. $Y^{(k+1)} = (Z^{(k)} - \Lambda^{(k)} + A - \lambda dd^\top)_+$
    2. $Z^{(k+1)} = (\min\{\max\{Y^{(k)} + \Lambda^{(k)}, 0\}, J\})_I$
    3. $\Lambda^{(k+1)} = \Lambda^{(k)} + Y^{(k+1)} - Z^{(k+1)}$
    4. $k = k + 1$

    end while
4: Output the final solution $Y^{(k)}$.

Algorithm 2 Convexified Modularity Maximization (CMM)

 1: Input: A, $\lambda = \langle A, J \rangle^{-1}$, and $r \geq 2$.
 2: Solve the convex relaxation (2.6) for $\widehat{Y}$ using Algorithm 1.
 3: Solve the weighted k-medoids problem (2.8) with $\widehat{W} = \widehat{Y} D$ and $k = r$, and output the resulting r-partition of $[n]$.

---

gorithm 2. In our experiments, the weighted k-medoids problem is solved by an iterative greedy procedure that optimizes alternatively over the variables $\boldsymbol{\Psi}$ and X in (2.8), with random initialization.

4.1. *Synthetic data experiments.* In this section, we provide experiment results on synthetic data generated from DCSBM. For each node $i \in [n]$, the degree heterogeneity parameter θ_i is sampled independently from a Pareto(α, β) distribution with the density function $f(x|\alpha, \beta) = \frac{\alpha \beta^\alpha}{x^{\alpha+1}} \mathbf{1}_{\{x \geq \beta\}}$, where α and β are called the *shape* and *scale* parameters, respectively. We consider different values of the shape parameter, and choose the scale parameter accordingly so that the expectation of each θ_i is fixed at 1. Note that the variability of the θ_i's decreases with the shape parameter α. Given the degree heterogeneity parameters $\{\theta_i\}$ and two numbers $0 \leq q < p \leq 1$, a graph is generated from DCSBM, with the edge probability between nodes $i \in C_a^*$ and $j \in C_b^*$ being $\min(1, \theta_i \theta_j B_{ab})$ and $B_{aa} = p$, $B_{ab} = q$, $\forall 1 \leq a \neq b \leq r$.

We applied our CMM approach in Algorithm 2 to the resulting graph, and recorded the misclassification rate [namely, $\min_{\boldsymbol{\Pi} \in \mathcal{S}_r} |\mathcal{E}(\boldsymbol{\Pi})|/n$; cf. the discussion after Theorem 3.2]. For comparison, we also applied several algorithms which are reported to have state-of-the-art empirical performance on DCSBM in the existing literature:

- The SCORE algorithm in [28] is a spectral clustering algorithm that performs k-means on the top-2 to top-r eigenvectors of the adjacency matrix normalized elementwise by the top-1 eigenvector.
- The OCCAM algorithm in [47] is a type of regularized spectral clustering algorithm. It can be instructed to produce nonoverlapping clusters and its regularization parameter is given explicitly in [47].
- The profile likelihood (PL) maximization algorithm in [29] with Tabu search [48] greedily maximizes the profile likelihood function over community partitions, starting with an initial partition generated either randomly or by a regularized spectral clustering algorithm.
- The conditional pseudo-likelihood (CPL) maximization algorithm in [5] maximizes the conditional pseudo-likelihood via EM algorithm, starting with an initial partition provided by regularized spectral clustering.

- To measure the performance gain by incorporating degree correction when the degrees are indeed heterogeneous, we also compared with the SDP1 algorithm in [6], Appendix J, developed under SBM.

For all k-means/medoids procedures used in the experiments, we set $k = r$ and used 100 random initializations. In both PL and CPL, the initial partition was obtained using a standard spectral clustering algorithm applied to the regularized adjacency matrix $A + 0.25(\eta/n)J$, where η is the average degree of the network [5].

We note that some of these methods like SCORE have the advantages of being conceptually simple and computationally fast. On the other hand, methods like PL and CPL employ greedy heuristic procedures, whose convergence properties are not fully established. Here, we mainly focus on their empirical clustering performance in terms of misclassification rates.

The synthetic experiments were performed with different connectivity matrix B, shape parameter α, and numbers and sizes of clusters. Note that as α increases, the degree parameters $\{\theta_i\}$ become less heterogeneous. In the limit $\alpha \to \infty$, we have $\theta_i \equiv 1$ for all i, and thus DCSBM reduces to SBM. In all plots, each point represents the average of 20 independent runs.

We first consider the setting with two equal-sized clusters, with the results given in Figure 1. It shows that the misclassification rate of CMM decreases as the degree parameters $\{\theta_i\}$ become less heterogeneous (larger values of the shape parameter). Moreover, in the dense graph case, CMM consistently has lower misclassification rates than all the other five competing methods; in the sparse graph case, CMM has similar performance as CPL/PL. We have similar observations when cluster sizes are unbalanced, as shown in Figure 2, and when we increase the number of clusters to four, as shown in Figure 3.

To investigate the flexibility and robustness of CMM, we consider a setting where the connectivity matrix B does not satisfy the so-called *strong assortative* condition, that is, $\min_{1 \le a \le r} B_{aa}$ is not necessarily larger than $\max_{1 \le a < b \le r} B_{ab}$. In particular, we consider the following connectivity matrix used in [6] with $p_3 \in$
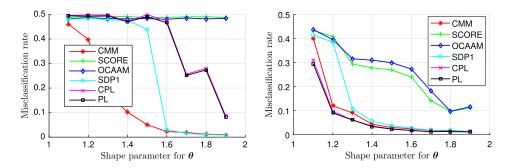

FIG. 1. *Misclassification rate versus variability of θ for $n = 800$ and 2 equal-sized clusters. Left panel*: $p = 0.5, q = 0.35$. *Right panel*: $p = 0.05, q = 0.015$.
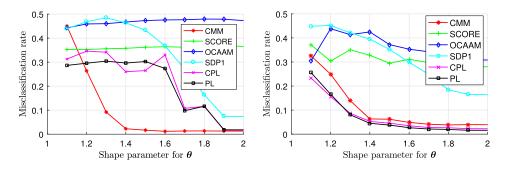
FIG. 2. *Misclassification rate versus variability of θ for* $n = 800$ *and* 2 *clusters of sizes* 600 *and* 200, *respectively. Left panel*: $p = 0.5, q = 0.3$. *Right panel*: $p = 0.05, q = 0.015$.

$[0.05, 0.7]$:

$$
B = \begin{bmatrix} 0.7 & 0.4 & 0.05 & 0.2 \\ 0.4 & 0.6 & 0.05 & 0.2 \\ 0.05 & 0.05 & p_3 & 0.05 \\ 0.2 & 0.2 & 0.05 & 0.4 \end{bmatrix}.
$$

Here, strong assortativity is never satisfied. When $p_3 > 0.05$, this matrix does satisfy the so-called *weak assortativity*, namely, $B_{aa} > \min_{b \neq a} B_{ab}$ for all $a \in [r]$. SDP1 is targeted at this setting; under the additional assumption that all community sizes are equal, it is shown in [6] that SDP1, which has additional row sum constraints, succeeds under weak assortativity.

While CMM does not have such row sum constraints and is oblivious to the equal-size assumption, Figure 4 shows that it still works well under weak assortativity, and in particular is robust to the variation of p_3, similar to SDP1, CPL and PL. In fact, CMM significantly outperforms the other five methods when the degree variation is high (shape parameter $\alpha = 1.2$). When degree variation is low, CMM is still on a par with SDP1, CPL and PL.
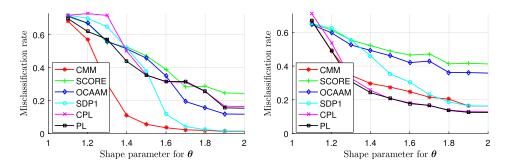

FIG. 3. *Misclassification rate versus variability of θ for* $n = 800$ *and* 4 *equal-sized clusters. Left panel*: $p = 0.5, q = 0.3$. *Right panel*: $p = 0.05, q = 0.015$.
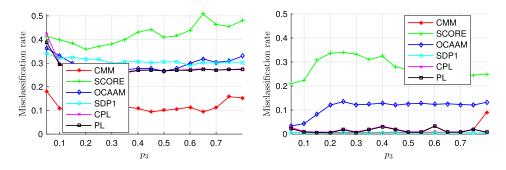
FIG. 4. *Misclassification rate versus p_3 for $n = 800$ and 4 equal-sized clusters. The shape parameter is $\alpha = 1.2$ in the left panel and $\alpha = 2$ in the right.*

We further consider a setting where even weak assortativity fails to hold, as shown in Figure 5. We see that CMM continues to dominate other methods when the network is relatively dense, and has similar performance to CPL/PL in the sparse setting.

In general, the results in this section indicate that CMM is a versatile algorithm for DCSBM, and has competitive performance compared to the state of the art in all the settings we considered. Moreover, CMM outperforms CPL and PL in the relatively dense setting. The comparison with SDP1 also demonstrates the gain of degree correction and the robustness of CMM under unbalanced clusters.

4.2. *Political blog network dataset.* We next test the empirical performance of CMM and the other algorithms on the US political blogs network dataset from [2]. This dataset consists of 19,090 hyperlinks (directed edges) between 1490 political blogs collected in the year 2005. The political leaning (liberal versus conservative) of each blog has been labeled manually by experts. We treat these labels as the true memberships of $r = 2$ communities. We ignore the edge direction, and focus on the largest connected component with $n = 1222$ nodes and 16,714 edges, represented
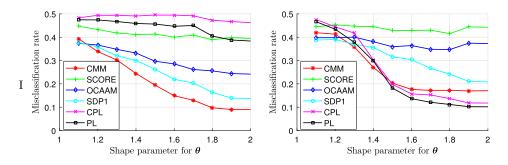

FIG. 5. *$n = 800$ and 2 equal-sized clusters. Left panel: $\boldsymbol{B} = [0.7, 0.3; 0.3, 0.2]$. Right panel: $\boldsymbol{B} = [0.14, 0.06; 0.06, 0.06]$.*
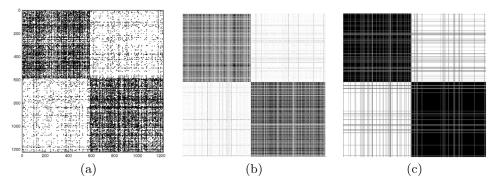
FIG. 6. *Results on largest connected component* (1222 *nodes*) *of the political blogs dataset. The rows and columns are sorted according to the true community labels, with the first* 586 *corresponding to the liberal community and the remaining* 636 *the conservative community. Panel* (a): *The adjacency matrix. Panel* (b): *The output matrix* $\widehat{Y}$ *of the convex relaxation* (2.6), *with the entries truncated to the interval* [0, 1]. *Panel* (c): *The partition matrix corresponding to the output of CMM* (*Algorithm* 2). *Matrix entry values are shown in gray scale with black corresponding to* 1 *and white to* 0.

by the adjacency matrix A. This graph has high degree variation: the maximum degree is 351 while the mean degree is around 27. Panel (a) in Figure 6 shows the adjacency matrix A with rows and columns sorted according to the true community labels. The output of ADMM Algorithm 1 for solving the convex relaxation (2.6) is shown in Figure 6(b). The partition matrix corresponding to the output of the weighted k-medoids step in Algorithm 2 is shown in Figure 6(c).

On the political blogs dataset, SCORE and PL have the best known error rates in the literature [28]. From Table 1, we see that CMM is comparable to the state of the art. In particular, with proper degree correction, CMM significantly improves over SDP1, which is designed for SBM.

4.3. *Facebook dataset.* The Facebook network dataset from [44, 45] consists of 100 US universities and a snapshot of all the "friendship" links between the users within each university in September 2005. The dataset also contains several node attributes such as the gender, dorm, graduation year and academic major of each user. Here, we report results on the friendship network of Simmons College. Results on the friendship network of Caltech can be found in the Supplementary Material [12].

TABLE 1
Number of misclassified nodes in the political blogs dataset

	CMM	SCORE	OCCAM	CPL	PL	SDP1
# of mis. nodes	61	58	65	61	58	181

TABLE 2
Misclassification rates in the Simmons College network

	CMM	SCORE	OCCAM	CPL	PL	SDP1
mis. rate	12.04%	23.57%	22.43%	18.21%	17.85%	25.42%

In the Simmons College network, the subgraph induced by nodes with graduation year between 2006 and 2009 has a largest connected component with 1137 nodes and 24,257 undirected edges, which we shall focus on. It has been observed in [44, 45] that the community structure of the Simmons College network exhibits a strong correlation with the graduation year—students in the same year are more likely to be friends.

We applied CMM (Algorithm 2), SCORE, OCCAM, CPL, PL and SDP1 methods to partition the largest component into $r = 4$ clusters. The error rates of these algorithms are shown in Table 2. In Figure 7, we also provide the confusion matrices of the clustering results against the graduation years; the (i, j)th entry of a confusion matrix represents the number of nodes that are from graduation year $i + 2005$ but assigned to cluster j by the algorithm. We see that our CMM approach produced a partition more correlated with the actual graduation years. In fact, if we treat the graduation years as the ground truth cluster labels, then the fractions of misclassified nodes by various algorithms are listed below.
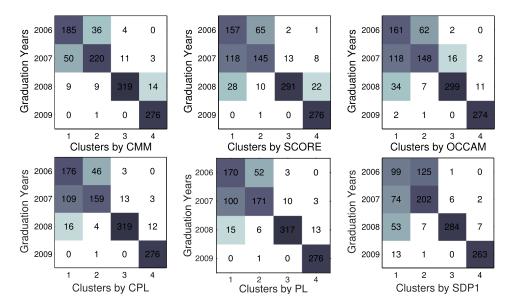

FIG. 7. *The confusion matrices of CMM, SCORE, OCCAM, CPL, PL and SDP1 applied to the largest component of the Simmons College network.*

A closer investigation of Figure 7 indicates that CMM was better in distinguishing between the nodes of year 2006 and 2007. This is consistent with visualization results provided in the Supplementary Material [12].

We emphasize that the partition by graduate years is by no means an absolute ground truth of the network's community structure. Nevertheless, the above results show that our algorithm is able to produce a meaningful, informative clustering without any prior knowledge of the graduation years.

5. Related work. In this section, we discuss prior results that are related to our work. Existing community detection methods for DCSBM include model-based methods and spectral methods. In model-based methods, one fits the model parameters to the observed network by optimizing an objective function derived under the statistical assumptions of SBM or DCSBM. For example, one may consider the profile likelihood function, which is itself obtained by plugging in the maximum likelihood estimates of θ and B. In [29], an estimate of the community structure is obtained by maximizing this profile likelihood using greedy heuristics. No theoretical guarantee is known for this greedy approach, and one usually runs the algorithm with multiple random initial solutions. The work in [48] also considers profile likelihood methods, as well as the closely related modularity maximization approach [38]. With a fixed number of clusters, they prove strong consistency when the average degree is $\Omega(\log n)$, and weak consistency when it is $\Omega(1)$. However, directly solving the associated maximization problems is computationally infeasible, as it involves searching over exponentially many partitions. In practice, these optimization problems are usually solved heuristically using Tabu search and spectral decomposition without theoretical guarantees. The algorithm proposed in [5] involves finding an initial clustering using spectral methods, then iteratively updating the labels via maximizing the conditional pseudo likelihood. The latter is done using the EM algorithm in each iteration. After simplifying the iterations into one E-step, they prove consistency when there are two clusters. The work in [32] proposes to approximate the profile likelihood functions, modularity functions or other criteria using surrogates defined in a 2-dimensional subspace constructed by spectral dimension reduction. Thanks to the convexity of the surrogate functions, the search complexity is polynomial. The algorithm and theory are however only applicable when there are two communities.

Spectral methods for community detection have attracted interest of experts from diverse communities; see, for example, [43] and the references therein for results of spectral clustering under SBM. The seminal work in [17] considers DCSBM (proposed under the name of Extended Planted Partition model) and proposes a spectral method similar to [36]. One major drawback is that the knowledge of θ is required in both the theory and algorithm. In the algorithm proposed in [15], the adjacency matrix is first normalized by the node degrees and then thresholded entrywise, after which spectral clustering is applied. Strong consistency is proved

for the setting with a fixed number of clusters. In [11], a modified spectral clustering method was proposed using a regularized random-walk graph Laplacian, and strong consistency is established under the assumption that the average degree grows at least as $\Omega(\sqrt{n})$. A different spectral clustering approach based on regularized graph Laplacians is considered in [41]. Their theoretical bound on the misclassified rates depends on the eigenvectors of the graph Laplacian, which is still a random object. Spectral clustering based on unmodified adjacency matrices and degree-normalized adjacency matrices are analyzed in [34] and [23], which prove rigorous error rate results but do not provide numerical validation on either synthetic or real data.

It is observed in [28] that spectral clustering based directly the adjacency matrix (or their variants) often result in inconsistent clustering in real data, such as the political blogs dataset [2], a popular benchmark for community detection algorithms. To address this issue, a new spectral clustering algorithm called SCORE is proposed in [28] using self-normalized eigenvectors. In their theoretical results, an implicit assumption is that the number of communities r is bounded by a constant, as implied by the condition (2.14) in [28]. In comparison, our convexified modularity maximization approach works for growing r both theoretically and empirically. As illustrated in Section 4.1, our method exhibited better performance on both the synthetic and real datasets considered there, especially when $r \geq 3$.

6. Discussion and future work. In this paper, we studied community detection in networks with potentially highly skewed degree distributions. We introduced a new computationally efficient methodology, which is based on convexification of the modularity maximization formulation and a novel doubly-weighted ℓ_1 norm k-medoids clustering procedure. Our complete algorithm runs provably in polynomial time and is computationally feasible. Nonasymptotic theoretical performance guarantees were established under DCSBM for both approximate clustering and perfect clustering, which are consistent with the best known rate results in the literature of SBM.

The proposed method also enjoys good empirical performance, as was demonstrated on both synthetic data and real-world networks. On these datasets, our method was observed to have performance comparable to, and sometimes better than, the state-of-the-art spectral clustering methods, particularly when there are more than two communities.

Our work involves several algorithmic and analytical novelties. We provide a tractable solution to the classical modularity maximization formulation via convexification, achieving simultaneously strong theoretical guarantees and competitive empirical performance. The theoretical results are based on an aggregate, degree-corrected version of the density gap condition, which is robust to nodes with small degrees. In our algorithms and error bounds, we made use of some techniques from [22, 28, 34], but departed from these existing works in several important aspects. In particular, we proposed a novel k-medoids formulation using

doubly-weighted ℓ_1 norms, which allows for a tight analysis that produces strong nonasymptotic guarantees on approximate recovery. Furthermore, we developed a nonasymptotic theory on perfect clustering, which is based on a primal-dual analysis and makes crucial use of certain weighted ℓ_1 metrics that exploit the structures of DCSBM.

A future direction important in both theory and practice, is to handle overlapping communities (i.e., a node may belong to multiple communities simultaneously). This setting has been considered under SBM, for example in [47], which proposed a spectral algorithm called OCCAM. As our CMM method is shown to be an attractive alternative to spectral methods for DCSBM, it will be interesting to extend CMM to allow for overlapping communities with heterogeneous degrees. Another direction of interest is to develop a general theory of optimal misclassification rates for DCSBM in the lines of [20, 46].

7. Proofs.
In this section, we prove the theoretical results in Section 3.2. The proof of Theorem 3.3 is deferred to the Supplementary Material [12]. Introducing the convenient shorthand $\boldsymbol{\Theta} := \boldsymbol{\theta}\boldsymbol{\theta}^\top \in \mathbb{R}_+^{n \times n}$, we can write the weighted ℓ_1 norm of a matrix $\boldsymbol{Z}$ in Definition 3.2 as

$$\|\boldsymbol{Z}\|_{1,\boldsymbol{\theta}} = \sum_{1 \le i,j \le n} |\theta_i Z_{ij} \theta_j| = \|\boldsymbol{\Theta} \circ \boldsymbol{Z}\|_1,$$

where $\circ$ denotes the Hadamard (elementwise) product. Several standard matrix norms will also be used: the spectral norm $\|\boldsymbol{Z}\|$ (the largest singular value of $\boldsymbol{Z}$); the nuclear norm $\|\boldsymbol{Z}\|_*$ (the sum of the singular values); the ℓ_1 norm $\|\boldsymbol{Z}\|_1 = \sum_{i,j} |Z_{ij}|$; the ℓ_∞ norm $\|\boldsymbol{Z}\|_\infty = \max_{i,j} |Z_{ij}|$; and the $\ell_\infty \to \ell_1$ operator norm $\|\boldsymbol{Z}\|_{\infty \to 1} = \sup_{\|\boldsymbol{v}\|_\infty \le 1} \|\boldsymbol{Z}\boldsymbol{v}\|_1$.

For any vector $\boldsymbol{v} \in \mathbb{R}^n$, we denote by $\mathrm{diag}(\boldsymbol{v})$ the $n \times n$ diagonal matrix whose diagonal entries are correspondingly the entries of $\boldsymbol{v}$. For any matrix $\boldsymbol{M} \in \mathbb{R}^n$, let $\mathrm{diag}(\boldsymbol{M})$ denote the $n \times n$ diagonal matrix with diagonal entries given by the corresponding diagonal entries of $\boldsymbol{M}$. We denote absolute constants by C, C_0, c_1, etc., whose value may change line by line.

7.1. *Proof of Theorem* 3.1. Recall that the vector $\boldsymbol{f} \in \mathbb{R}^n$ is defined by letting $f_i = \theta_i H_a$ for $i \in C_a^*$, where H_a is defined in (3.3) as $H_a = \sum_{b=1}^r B_{ab} G_b$. It follows from the optimality of $\widehat{\boldsymbol{Y}}$ that

$$0 \le \langle \widehat{\boldsymbol{Y}} - \boldsymbol{Y}^*, \boldsymbol{A} - \lambda \boldsymbol{d}\boldsymbol{d}^\top \rangle$$
$$= \underbrace{\langle \widehat{\boldsymbol{Y}} - \boldsymbol{Y}^*, \mathbb{E}\boldsymbol{A} - \lambda \boldsymbol{f}\boldsymbol{f}^\top \rangle}_{S_1} + \underbrace{\lambda \langle \widehat{\boldsymbol{Y}} - \boldsymbol{Y}^*, \boldsymbol{f}\boldsymbol{f}^\top - \boldsymbol{d}\boldsymbol{d}^\top \rangle}_{S_2} + \underbrace{\langle \widehat{\boldsymbol{Y}} - \boldsymbol{Y}^*, \boldsymbol{A} - \mathbb{E}\boldsymbol{A} \rangle}_{S_3}.$$

We control the terms S_1, S_2 and S_3 separately below.

Upper bound for S_1. For each pair $i, j \in C_a^*$ and $i \neq j$, we have $\widehat{Y}_{ij} - Y_{ij}^* \leq 0$, $\mathbb{E}(A_{ij}) = \theta_i \theta_j B_{aa}$, and $f_i f_j = \theta_i \theta_j H_a^2$. Hence the condition (3.6) implies that $\mathbb{E}(A_{ij}) - \lambda f_i f_j \geq \delta \theta_i \theta_j$, whence

$$\big(\widehat{Y}_{ij} - Y_{ij}^*\big)\big(\mathbb{E}(A_{ij}) - \lambda f_i f_j\big) \leq -\delta \theta_i \theta_j \big|\widehat{Y}_{ij} - Y_{ij}^*\big|.$$

Similarly, for each pair $i \in C_a^*$ and $j \in C_b^*$ with $1 \leq a < b \leq r$, we have $\widehat{Y}_{ij} - Y_{ij}^* \geq 0$, $\mathbb{E}(A_{ij}) = \theta_i \theta_j B_{ab}$, and $f_i f_j = \theta_i \theta_j H_a H_b$. Hence, the condition (3.6) implies that $\mathbb{E}(A_{ij}) - \lambda f_i f_j \leq -\delta \theta_i \theta_j$, whence

$$\big(\widehat{Y}_{ij} - Y_{ij}^*\big)\big(\mathbb{E}(A_{ij}) - \lambda f_i f_j\big) \leq -\delta \theta_i \theta_j \big|\widehat{Y}_{ij} - Y_{ij}^*\big|.$$

Combining the last two inequalities, we obtain the bound

$$S_1 := \big\langle \widehat{Y} - Y^*, \mathbb{E}A - \lambda f f^\top \big\rangle \leq -\delta \big\|Y^* - \widehat{Y}\big\|_{1,\theta}.$$

Upper bound for S_2. By Grothendieck's inequality [21, 35], we have

$$\big\langle \widehat{Y} - Y^*, f f^\top - d d^\top \big\rangle \leq 2 \sup_{Y \succeq 0, \mathrm{diag}(Y)=I} \big|\langle Y, f f^\top - d d^\top \rangle\big|$$

$$\leq 2 K_G \big\| f f^\top - d d^\top \big\|_{\infty \to 1},$$

where K_G is Grothendieck's constant known to satisfy $K_G \leq 1.783$. Applying Lemma 9.1 in the Supplementary Material [12] on $\| f f^\top - d d^\top \|_{\infty \to 1}$ ensures that with probability at least 0.99 and for some absolute constant C,

$$S_2 \leq C \lambda \|f\|_1 \big(\sqrt{n \|f\|_1} + n\big).$$

Upper bound for S_3. Observe that

$$\big\langle \widehat{Y} - Y^*, A - \mathbb{E}A \big\rangle \leq 2 \sup_{Y \succeq 0, \mathrm{diag}\, Y = I} \big|\langle Y, A - \mathbb{E}A \rangle\big|.$$

It follows from Grothendieck's inequality that

$$\sup_{Y \succeq 0, \mathrm{diag}(Y)=I} \big|\langle Y, A - \mathbb{E}A \rangle\big| \leq K_G \|A - \mathbb{E}A\|_{\infty \to 1}.$$

The norm on the last RHS can be expressed as

$$\|A - \mathbb{E}A\|_{\infty \to 1} = \sup_{x : \|x\|_\infty \leq 1} \big\|(A - \mathbb{E}A)x\big\|_1 = \sup_{x, y \in \{\pm 1\}^n} \big|x^\top (A - \mathbb{E}A)y\big|.$$

For each fixed pair of sign vectors $x, y \in \{\pm 1\}^n$, Bernstein's inequality ensures that for each $t > 0$, we have with probability at most $2e^{-t}$,

$$\big|x^\top (A - \mathbb{E}A)y\big| \geq \sqrt{8t\sigma^2} + \frac{4}{3}t,$$

where $\sigma^2 := \sum_{i<j} \text{var} A_{ij} \leq \frac{1}{2} \sum_{a,b=1}^{r} B_{ab} G_a G_b = \frac{1}{2} \|f\|_1$. Setting $t = 2n$ and applying the union bound over all sign vectors, we obtain that with probability at most $2(e/2)^{-2n}$,

$$\|A - \mathbb{E}A\|_{\infty \to 1} \geq \sqrt{8n\|f\|_1} + \frac{8}{3}n.$$

It follows that with probability at least $1 - 2(e/2)^{-2n}$,

$$S_3 \leq 2K_G\sqrt{8n\|f\|_1} + \frac{16K_G}{3}n.$$

Putting together the bounds for S_1, S_2 and S_3, we conclude that with probability at least $0.99 - 2(e/2)^{-2n}$, the bound (3.7) holds.

7.2. *A general error bound for k-medoids clustering.* The remainder of this section is devoted to the proof of Theorem 3.2. Central to the proof are two general results, given in this subsection and in Section 7.3 to follow, on k-medoids clustering error, which may be of independent interest. The first result quantifies the (weighted) misclassification rate of the approximate solution $\breve{\Psi}$ to the (weighted) k-medoids problem (2.7).

LEMMA 7.1. *Suppose $G_a = g, \forall a \in [r]$. For any matrix $Q \in \mathbb{R}^{r \times n}$,*

$$\min_{\Pi \in \mathcal{S}_r} \left\{ \sum_{i \in \mathcal{E}(\Pi)} \theta_i \right\} \leq \frac{2}{g} \|\breve{\Psi} Q - Y^*\|_{1,\theta}.$$

PROOF. Define $\breve{Y} = \breve{\Psi} Q$. For each $a \in [r]$, define the set of node indices

$$S_a := \{i \in C_a^* : \|(\breve{Y}_{i\bullet} - Y_{i\bullet}^*) \text{Diag}(\theta)\|_1 \geq g\},$$

and let $S := \bigcup_{a=1}^{r} S_a$. It follows that

$$(7.1) \qquad \sum_{i \in S} \theta_i \leq \sum_{i=1}^{n} \frac{\theta_i}{g} \|(\breve{Y}_{i\bullet} - Y_{i\bullet}^*) \text{Diag}(\theta)\|_1 = \frac{1}{g}\|\breve{Y} - Y^*\|_{1,\theta}.$$

Consider the set $T_a := C_a^* \setminus S_a$ for each $a = 1, \dots, r$. There are three cases for each T_a. In the first case, $T_a = \varnothing$, and we denote by R_1 the collection of all such indices a. In the second case, $T_a \neq \varnothing$ and $\breve{\Psi}_{i\bullet} = \breve{\Psi}_{j\bullet}$ for all $i, j \in T_a$. We say that these T_a's are pure, and denote by R_2 the collection of all such indices a. Finally, we set $R_3 := \{1, \dots, r\} \setminus (R_1 \cup R_2)$; for each $a \in R_3$, we say that T_a is impure since there exist $i, j \in T_a$ such that $\breve{\Psi}_{i\bullet} \neq \breve{\Psi}_{j\bullet}$.

For each $a \in R_1$, we have $S_a = C_a^*$, which implies that

$$(7.2) \qquad \sum_{i \in S} \theta_i \geq \sum_{i \in \bigcup_{a \in R_1} C_a^*} \theta_i = |R_1|g.$$

For each pair $a, b \in R_2 \cup R_3$ with $a \neq b$, by definition we know that $T_a \neq \varnothing$ and $T_b \neq \varnothing$. Then for each pair $i \in T_a \subseteq C_a^*$, $j \in T_b \subseteq C_b^*$, we have

$$\left\| \breve{Y}_{i\bullet} \operatorname{Diag}(\boldsymbol{\theta}) - \breve{Y}_{j\bullet} \operatorname{Diag}(\boldsymbol{\theta}) \right\|_1$$
$$\geq \left\| (Y_{i\bullet}^* - Y_{j\bullet}^*) \operatorname{Diag}(\boldsymbol{\theta}) \right\|_1 - \left\| (Y_{i\bullet}^* - \breve{Y}_{i\bullet}) \operatorname{Diag}(\boldsymbol{\theta}) \right\|_1$$
$$- \left\| (Y_{j\bullet}^* - \breve{Y}_{j\bullet}) \operatorname{Diag}(\boldsymbol{\theta}) \right\|_1$$
$$> 2g - g - g = 0,$$

whence $\breve{Y}_{i\bullet} \neq \breve{Y}_{j\bullet}$. This implies that $\breve{\boldsymbol{\Psi}}_{i\bullet} \neq \breve{\boldsymbol{\Psi}}_{j\bullet}$.

In conclusion, we have proved that for each pair $a, b \in R_2 \cup R_3$ with $a \neq b$ and each pair $i \in T_a$, $j \in T_b$, we have $\breve{\boldsymbol{\Psi}}_{i\bullet} \neq \breve{\boldsymbol{\Psi}}_{j\bullet}$. Moreover, since for each $a \in R_2$, the set T_a is pure by definition, there exists a permutation matrix $\boldsymbol{\Pi} \in \mathcal{S}_r$ such that for all $i \in \bigcup_{a \in R_2} T_a$, there holds $(\breve{\boldsymbol{\Psi}} \boldsymbol{\Pi})_{i\bullet} = \boldsymbol{\Psi}_{i\bullet}^*$. Recalling Definition 3.3, we conclude that the set $(\bigcup_{a \in R_3} T_a) \cup S$ contains the set of nodes misclassified by $\breve{\boldsymbol{\Psi}}$ with respect to $\boldsymbol{\Pi}$. It follows that

$$(7.3) \qquad \sum_{i \in \mathcal{E}(\boldsymbol{\Pi})} \theta_i \leq \sum_{i \in S} \theta_i + \sum_{i \in \bigcup_{a \in R_3} T_a} \theta_i \leq \sum_{i \in S} \theta_i + |R_3| g.$$

The matrix $\breve{\boldsymbol{\Psi}}$ consists of at most r distinct row vectors. Because R_2 is pure and R_3 is impure by definition, we have the inequality

$$|R_2| + 2|R_3| \leq r = |R_1| + |R_2| + |R_3|,$$

which implies that

$$(7.4) \qquad |R_3| \leq |R_1|.$$

Applying the bounds (7.3), (7.4), (7.2) and (7.1) in order, we obtain

$$\sum_{i \in \mathcal{E}(\boldsymbol{\Pi})} \theta_i \leq 2 \sum_{i \in S} \theta_i \leq \frac{2}{g} \| \breve{Y} - Y^* \|_{1,\boldsymbol{\theta}} = \frac{2}{g} \| \breve{\boldsymbol{\Psi}} \boldsymbol{Q} - Y^* \|_{1,\boldsymbol{\theta}},$$

completing the proof of the lemma. $\square$

7.3. *Comparison between k-medians and k-medoids.* The second general result we need for proving Theorem 3.2, is a comparison of the errors of the weighted k-medoids formulation (2.7) and the following weighted k-medians clustering formulation:

$$(7.5) \qquad \min_{\{C_a\}_{a=1}^r, \{\boldsymbol{x}_a\}_{a=1}^r} \sum_{1 \leq a \leq r} \sum_{i \in C_a} d_i \| \widehat{\boldsymbol{w}}_i - \boldsymbol{x}_a \|_1$$

in which the centers are not necessarily data points and $\{C_a\}_{a=1}^r$ denotes a partition of $[n]$. Similar to (2.7), the formulation (7.5) can be written as

$$(7.6) \qquad \min_{\boldsymbol{\Psi}, X} \| \boldsymbol{D}(\boldsymbol{\Psi} X - \widehat{W}) \|_1, \quad \text{s.t.} \quad \boldsymbol{\Psi} \in \mathbb{M}_{n,r}, \qquad X \in \mathbb{R}^{r \times n}.$$

The following lemma establishes the connection between the clustering errors of k-medoids and k-medians.

LEMMA 7.2.   *Suppose the optimal value of* (2.7) *and* (2.8) *is* m_1 *and the optimal value of* (7.5) *and* (7.6) *is* m_2. *Then we have*

$$m_2 \leq m_1 \leq 2m_2.$$

PROOF.   The optimization problems (7.5) and (2.7) have the same objective function, and the feasible set of (2.7) is contained in that of (7.5), so we have $m_1 \geq m_2$. To prove $m_2 \geq \frac{1}{2}m_1$, it suffices to show that for any partition $\{C_a\}_{a=1}^r$ of $[n]$, and any vectors $\{x_a\}_{a=1}^r$, there holds

$$\sum_{1 \leq a \leq r} \sum_{i \in C_a} d_i \|\widehat{w}_i - x_a\|_1 \geq \frac{m_1}{2}.$$

To this end, we define the vector $\widetilde{x}_a := \arg\min_{x \in \{\widehat{w}_i : i \in C_a\}} \|x - x_a\|_1$ for each $a = 1, \ldots, r$. It follows that

$$\sum_{1 \leq a \leq r} \sum_{i \in C_a} d_i \|\widehat{w}_i - x_a\|_1 \geq \sum_{1 \leq a \leq r} \sum_{i \in C_a} d_i \frac{1}{2} (\|\widehat{w}_i - x_a\|_1 + \|\widetilde{x}_a - x_a\|_1)$$

$$\geq \frac{1}{2} \sum_{1 \leq a \leq r} \sum_{i \in C_a} d_i \|\widehat{w}_i - \widetilde{x}_a\|_1 \geq \frac{m_1}{2},$$

where the last inequality uses the feasibility of $(\{C_a\}, \{\widetilde{x}_a\})$ to (2.7). $\square$

7.4. *Proof of Theorem* 3.2.   We now prove Theorem 3.2 by combining Theorem 3.1, Lemma 7.1 and Lemma 7.2. Our strategy is as follows. We first choose an appropriate Q based on $(\breve{\Psi}, \breve{X})$, so that we can apply Lemma 7.1 to bound the weighted misclassification rate by $\|\breve{\Psi} Q - Y^*\|_{1,\theta}$. We then bound $\|\breve{\Psi} Q - Y^*\|_{1,\theta}$ in terms of the weighted error of the convex relaxation solution, $\|\widehat{Y} - Y^*\|_{1,\theta}$, which is in turn controlled by Theorem 3.1.

To choose a Q to apply Lemma 7.1, we notice that the approximate solution $(\breve{\Psi}, \breve{X})$ is feasible to (2.8), so $\text{Rows}(\breve{X}) \subset \text{Rows}(\widehat{W})$. Since $\widehat{W} = \widehat{Y} D$, there exists an $r \times n$ matrix Q, such that

$$\breve{X} = Q D \quad \text{and} \quad \text{Rows}(Q) \subset \text{Rows}(\widehat{Y}).$$

Define $\breve{Y} = \breve{\Psi} Q$. In view of Lemma 7.1, it suffices to bound $\|\breve{Y} - Y^*\|_{1,\theta}$.

Before bounding $\|\breve{Y} - Y^*\|_{1,\theta}$, let us first control $\|\breve{Y} - Y^*\|_{1,d}$, where $\|Z\|_{1,d} = \sum_{i=1}^n \sum_{j=1}^n |Z_{ij}| d_i d_j$. To this end, we follow the notation in Lemma 7.2 and denote the optimum value of (2.7) and (2.8) as m_1, and that of (7.5) and (7.6) as m_2. First, triangle inequality gives

$$\|\breve{Y} - Y^*\|_{1,d} \leq \|\breve{Y} - \widehat{Y}\|_{1,d} + \|Y^* - \widehat{Y}\|_{1,d}.$$

Since $\breve{Y}D = \breve{\Psi}QD = \breve{\Psi}\breve{X}$ and $\widehat{W} = \widehat{Y}D$, we have

$$\|\breve{Y} - \widehat{Y}\|_{1,d} = \|D(\breve{Y}D - \widehat{W})\|_1 = \|D(\breve{\Psi}\breve{X} - \widehat{W})\|_1 \le \frac{20}{3}m_1,$$

where the last inequality is due to the definition of $(\breve{\Psi}, \breve{X})$. Recall that the true partition matrix admits the decomposition $Y^* = \Psi^*(\Psi^*)^\top$, where $\Psi^* \in \mathbb{M}_{n,r}$ is the true membership matrix. Therefore,

$$\|Y^* - \widehat{Y}\|_{1,d} = \|D(Y^*D - \widehat{W})\|_1 = \|D(\Psi^*((\Psi^*)^\top D) - \widehat{W})\|_1 \ge m_2,$$

where the last inequality is due to the fact that $(\Psi^*, (\Psi^*)^\top D)$ is feasible to (7.6). Then by Lemma 7.2, we have

$$\|\breve{Y} - \widehat{Y}\|_{1,d} \le \frac{20}{3}m_1 \le \frac{40}{3}m_2 \le \frac{40}{3}\|Y^* - \widehat{Y}\|_{1,d}$$

and this implies

(7.7) $$\|\breve{Y} - Y^*\|_{1,d} \le \|\breve{Y} - \widehat{Y}\|_{1,d} + \|Y^* - \widehat{Y}\|_{1,d} \le \frac{43}{3}\|Y^* - \widehat{Y}\|_{1,d}.$$

Next, we translate the inequality between $\|\breve{Y} - Y^*\|_{1,d}$ and $\|Y^* - \widehat{Y}\|_{1,d}$ into an inequality between $\|\breve{Y} - Y^*\|_{1,\theta}$ and $\|Y^* - \widehat{Y}\|_{1,\theta}$. Note that under the $\mathcal{F}(n, r, p, q, g)$-model,

$$H_1 = \cdots = H_r = h := (p + (r-1)q)g,$$

which implies that the vector of (approximate) expected degrees f defined in (3.5) satisfies $f_i = \theta_i h, \forall i \in [n]$ and $\|f\|_1 = r(p + (r-1)q)g^2$.

Recall that $\mathrm{Rows}(Q) \subset \mathrm{Rows}(\widehat{Y})$, which implies that $\mathrm{Rows}(\breve{Y}) = \mathrm{Rows}(\breve{\Psi}Q) \subset \mathrm{Rows}(\widehat{Y})$. Since $\widehat{Y}$ is feasible to the convex relaxation (2.6), we have $0 \le \widehat{Y} \le J$. It follows that $0 \le \breve{Y} \le J$, and hence $\|\breve{Y} - Y^*\|_\infty \le 1$. Setting $M := \|ff^T - dd^T\|_1$, we observe that any matrix Z satisfies the bound

$$\left|\|Z\|_{1,f} - \|Z\|_{1,d}\right| \le \|Z \circ (ff^T - dd^T)\|_1 \le M\|Z\|_\infty,$$

where $\|Z\|_{1,f}$ is defined in the same fashion as $\|Z\|_{1,\theta}$ given in Definition 3.2. Therefore, the bound (7.7) implies that

$$\|\breve{Y} - Y^*\|_{1,\theta} = \frac{1}{h^2}\|\breve{Y} - Y^*\|_{1,f} \le \frac{1}{h^2}(\|\breve{Y} - Y^*\|_{1,d} + M)$$

(7.8) $$\lesssim \frac{1}{h^2}(\|\widehat{Y} - Y^*\|_{1,d} + M) \le \frac{1}{h^2}(\|\widehat{Y} - Y^*\|_{1,f} + 2M)$$

$$\lesssim \|\widehat{Y} - Y^*\|_{1,\theta} + \frac{M}{h^2}.$$

To bound the second term above, we apply Lemma 9.1 in the Supplementary Material [12] to get that with probability at least 0.99,

$$M \le C\|f\|_1(\sqrt{n\|f\|_1} + n),$$

and thus

$$(7.9) \qquad \frac{M}{h^2} \lesssim \frac{r(\sqrt{n}\|f\|_1 + n)}{p + (r-1)q} \leq \frac{r(rg\sqrt{np} + n)}{\delta}.$$

We can control the first term in (7.8) using the bound (3.9) in Corollary 3.1. Putting together, straightforward calculation yields the inequality

$$(7.10) \qquad \|\check{\Psi} Q - Y^*\|_{1,\theta} = \|\check{Y} - Y^*\|_{1,\theta} \lesssim \frac{1}{\delta} r(n + rg\sqrt{np}).$$

We then apply Lemma 7.1 with (7.10) to obtain the desired inequality (3.10).

SUPPLEMENTARY MATERIAL

Additional experiments and remaining proofs. (DOI: 10.1214/17-AOS1595SUPP; .pdf). In this supplement [12], we provide additional numerical results and the remaining proofs of the theoretical results.

REFERENCES

[1] ABBE, E., BANDEIRA, A. S. and HALL, G. (2016). Exact recovery in the stochastic block model. *IEEE Trans. Inform. Theory* **62** 471–487. MR3447993

[2] ADAMIC, A. and GLANCE, N. (2005). The political blogosphere and the 2004 US election: Divided they blog. In *Proceedings of the 3rd International Workshop on Link Discovery* 36–43. ACM, New York.

[3] AGARWAL, N., BANDEIRA, A. S., KOILIARIS, K. and KOLLA, A. (2015). Multisection in the stochastic block model using semidefinite programming. Preprint. Available at arXiv:1507.02323.

[4] AMES, B. P. W. and VAVASIS, S. A. (2014). Convex optimization for the planted k-disjoint-clique problem. *Math. Program.* **143** 299–337. MR3152071

[5] AMINI, A. A., CHEN, A., BICKEL, P. J. and LEVINA, E. (2013). Pseudo-likelihood methods for community detection in large sparse networks. *Ann. Statist.* **41** 2097–2122. MR3127859

[6] AMINI, A. A. and LEVINA, E. (2014). On semidefinite relaxations for the block model. Available at arXiv:1406.5647.

[7] BANDEIRA, A. S. (2015). Random Laplacian matrices and convex relaxations. Preprint. Available at arXiv:1504.03987.

[8] BORDENAVE, C., LELARGE, M. and MASSOULIÉ, L. (2015). Non-backtracking spectrum of random graphs: Community detection and non-regular Ramanujan graphs. Preprint. Available at arXiv:1501.06087.

[9] 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

[10] CHARIKAR, M., GUHA, S., TARDOS, É. and SHMOYS, D. B. (1999). A constant-factor approximation algorithm for the k-median problem (extended abstract). In *Annual ACM Symposium on Theory of Computing (Atlanta, GA, 1999)* 1–10. ACM, New York. MR1797458

[11] CHAUDHURI, K., CHUNG, F. and TSIATAS, A. (2012). Spectral clustering of graphs with general degrees in the extended planted partition model. In *Proceedings of the 25th Annual Conference on Learning Theory (COLT)* 35.1–35.23.

[12] CHEN, Y., LI, X. and XU, J. (2018). Supplement to "Convexified modularity maximization for degree-corrected stochastic block models." DOI:10.1214/17-AOS1595SUPP.

[13] CHEN, Y., SANGHAVI, S. and XU, H. (2012). Clustering sparse graphs. *Adv. Neural Inf. Process. Syst.* 2213–2221.

[14] CHEN, Y. and XU, J. (2014). Statistical-computational phase transitions in planted models: The high-dimensional setting. In *Proceedings of the* 31*st International Conference on Machine Learning* 244–252.

[15] COJA-OGHLAN, A. and LANKA, A. (2009). Finding planted partitions in random graphs with general degree distributions. *SIAM J. Discrete Math.* **23** 1682–1714. MR2570199

[16] CONDON, A. and KARP, R. M. (2001). Algorithms for graph partitioning on the planted partition model. *Random Structures Algorithms* **18** 116–140. MR1809718

[17] DASGUPTA, A., HOPCROFT, J. and MCSHERRY, F. (2004). Spectral analysis of random graphs with skewed degree distributions. In *The* 45*th IEEE FOCS* 602–610.

[18] FORTUNATO, S. (2010). Community detection in graphs. *Phys. Rep.* **486** 75–174. MR2580414

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

[20] GAO, C., MA, Z., ZHANG, A. Y. and ZHOU, H. H. (2015). Achieving optimal misclassification proportion in stochastic block model. Preprint. Available at arXiv:1505.03772.

[21] GROTHENDIECK, A. (1996). Résumé de la théorie métrique des produits tensoriels topologiques. *Resen. Inst. Mat. Estat. Univ. Sao Paulo* **2** 401–480. MR1466414

[22] GUÉDON, O. and VERSHYNIN, R. (2016). Community detection in sparse networks via Grothendieck's inequality. *Probab. Theory Related Fields* **165** 1025–1049. MR3520025

[23] GULIKERS, L., LELARGE, M. and MASSOULIÉ, L. (2017). A spectral method for community detection in moderately sparse degree-corrected stochastic block models. *Adv. in Appl. Probab.* **49** 686–721. MR3694314

[24] HAJEK, B., WU, Y. and XU, J. (2016). Achieving exact cluster recovery threshold via semidefinite programming: Extensions. *IEEE Trans. Inform. Theory* **62** 5918–5937. MR3552431

[25] HAJEK, B., WU, Y. and XU, J. (2016). Achieving exact cluster recovery threshold via semidefinite programming. *IEEE Trans. Inform. Theory* **62** 2788–2797. MR3493879

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

[27] HOLLAND, P. W., LASKEY, K. B. and LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Soc. Netw.* **5** 109–137. MR0718088

[28] JIN, J. (2015). Fast community detection by SCORE. *Ann. Statist.* **43** 57–89. MR3285600

[29] KARRER, B. and NEWMAN, M. E. J. (2011). Stochastic blockmodels and community structure in networks. *Phys. Rev. E* (3) **83** 016107, 10. MR2788206

[30] 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

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

[32] LE, C. M., LEVINA, E. and VERSHYNIN, R. (2016). Optimization via low-rank approximation for community detection in networks. *Ann. Statist.* **44** 373–400. MR3449772

[33] LE, C. M. and VERSHYNIN, R. (2015). Concentration and regularization of random graphs. Preprint. Available at arXiv:1506.00669.

[34] LEI, J. and RINALDO, A. (2015). Consistency of spectral clustering in stochastic block models. *Ann. Statist.* **43** 215–237. MR3285605

[35] LINDENSTRAUSS, J. and PEŁCZYŃSKI, A. (1968). Absolutely summing operators in L_p-spaces and their applications. *Studia Math.* **29** 275–326. MR0231188

[36] McSherry, F. (2001). Spectral partitioning of random graphs. In *42nd IEEE Symposium on Foundations of Computer Science* (*Las Vegas*, *NV*, 2001) 529–537. IEEE Computer Soc., Los Alamitos, CA. MR1948742

[37] Montanari, A. and Sen, S. (2015). Semidefinite programs on sparse random graphs. Preprint. Available at arXiv:1504.05910.

[38] Newman, M. E. J. (2006). Modularity and community structure in networks. *PNAS* **103** 8577–8582.

[39] Oymak, S. and Hassibi, B. (2011). Finding dense clusters via low rank + sparse decomposition. Preprint. Available at arXiv:1104.5186.

[40] Perry, W. and Wein, A. S. (2015). A semidefinite program for unbalanced multisection in the stochastic block model. Preprint. Available at arXiv:1507.05605.

[41] Qin, T. and Rohe, K. (2013). Regularized spectral clustering under the degree-corrected stochastic blockmodel. In *Advances in Neural Information Processing Systems* 3120–3128.

[42] Reichardt, J. and Bornholdt, S. (2006). Statistical mechanics of community detection. *Phys*. *Rev. E* (3) **74** 016110, 14. MR2276596

[43] Rohe, K., Chatterjee, S. and Yu, B. (2011). Spectral clustering and the high-dimensional stochastic blockmodel. *Ann*. *Statist*. **39** 1878–1915. MR2893856

[44] Traud, A. L., Kelsic, E. D., Mucha, P. J. and Porter, M. A. (2011). Comparing community structure to characteristics in online collegiate social networks. *SIAM Rev.* **53** 526–543. MR2834086

[45] Traud, A. L., Mucha, P. J. and Porter, M. A. (2012). Social structure of Facebook networks. *Phys*. *A* **391** 4165–4180.

[46] Zhang, A. Y. and Zhou, H. H. (2015). Minimax rates of community detection in stochastic block models. Preprint. Available at arXiv:1507.05313.

[47] Zhang, Y., Levina, E. and Zhu, J. (2014). Detecting overlapping communities in networks using spectral methods. Preprint. Available at arXiv:1412.3432.

[48] Zhao, Y., Levina, E. and Zhu, J. (2012). Consistency of community detection in networks under degree-corrected stochastic block models. *Ann*. *Statist*. **40** 2266–2292. MR3059083

Y. Chen
Operations Research & Information Engineering
Cornell University
Ithaca, New York
USA
E-mail: yudong.chen@cornell.edu

X. Li
Statistics Department
University of California
Davis, California
USA
E-mail: xdgli@ucdavis.edu

J. Xu
Krannert School of Management
Purdue University
West Lafayette, Indiana
USA
E-mail: xu972@purdue.edu