Submitted to the Annals of Statistics

OPTIMIZATION VIA LOW-RANK APPROXIMATION WITH APPLICATIONS TO COMMUNITY DETECTION IN NETWORKS

By Can Le, Elizaveta Levina, and Roman Vershynin

University of Michigan

Community detection is one of the fundamental problems of network analysis. A number of methods for community detection have been proposed, including spectral clustering, modularity, and likelihood-based methods. Most of these methods have to solve an optimization problem over a discrete set of labels, which is computationally infeasible. Some fast algorithms have been proposed for specific methods or models have been proposed, but only on a case by case basis. Here we propose a general approach for maximizing a function of a network adjacency matrix over discrete labels by projecting the set of labels onto a subspace spanned by leading eigenvectors of the adjacency matrix. The main idea is that projection onto a low-dimensional space makes the feasible set of labels much smaller and the optimization problem much easier. We prove a general result on this method and show how to apply it to several previously proposed community detection criteria, establishing its consistency for label estimation in each case. Simulations and applications to real-world data are included to demonstrate our method performs well for multiple problems over a wide range of parameters.

1. Introduction. Networks are studied in a wide range of fields, including social psychology, sociology, physics, computer science, probability, and statistics. One of the fundamental problems in network analysis, and one of the most studied, is detecting network community structure. Community detection is the problem of inferring the latent label vector $\boldsymbol{c} \in \{1, \dots, K\}^n$ for the n nodes from the observed $n \times n$ adjacency matrix A, specified by $A_{ij} = 1$ if there is an edge from i to j, and $A_{ij} = 0$ otherwise. While the problem of choosing the number of communities K is an important open problem, for the purposes of this paper we will assume K is fixed, as does most of the existing literature. We will also focus on the undirected network case, where the matrix A is symmetric. Roughly speaking, the large recent literature on community detection has followed one of two tracks: fitting probabilistic models for the adjacency matrix A, or optimizing global cri-

AMS 2000 subject classifications: Primary 62E10; secondary 62G05

Keywords and phrases: Community detection, Spectral clustering, Stochastic block model, Social networks

teria derived from other considerations over label assignments $\boldsymbol{c}$, often via spectral approximations.

The simplest and most popular probabilistic model for fitting community structure is the stochastic block model Holland et al. (1983). The label vector $\boldsymbol{c}$ is assumed to be drawn from a multinomial distribution with parameter $\boldsymbol{\pi} = \{\pi_1, \ldots, \pi_K\}$, where $0 \leq \pi_k \leq 1$ and $\sum_{k=1}^{K} \pi_k = 1$. Edges are then formed independently between every pair of nodes (i, j) with probability $P_{c_i c_j}$, and the $K \times K$ matrix $P = [P_{kl}]$ controls the probability of edges within and between communities. Thus the labels are the only node information affecting edges between nodes, and all the nodes within the same community are stochastically equivalent to each other. This rules out the commonly encountered "hub" nodes, which are nodes of unusually high degrees that are connected to many members of their own community, or simply to many nodes across the network. To address this limitation, a relaxation that allows for arbitrary expected node degrees within communities was proposed by Karrer and Newman (2011): the degree-corrected stochastic block model has $P(A_{ij} = 1) = \theta_i \theta_j P_{c_i c_j}$, where θ_i's are "degree parameters" satisfying some identifiability constraints. In the "null" case of $K = 1$, both the block model and the degree corrected block model correspond to well-studied random graph models, the Erdos-Renyi graph (Erdős and Rényi, 1959) and the configuration model (Chung and Lu, 2002), respectively. Many other network models have been proposed to capture the community structure, for example, the latent space model (Hoff et al., 2002) and the latent position cluster model (Handcock et al., 2007). There has also been work on extensions of the block model which allow nodes to belong to more than one community (Airoldi et al., 2008; Ball et al., 2011). For a more complete review of network models, see Goldenberg et al. (2010).

Fitting models such as the stochastic block model typically involves maximizing a likelihood function over all possible label assignments, which is in principle NP-hard. MCMC-type and variational methods have been proposed, see for example Snijders and Nowicki (1997), Nowicki and Snijders (2001), Mariadassou et al. (2010), as well as maximizing profile likelihoods by some type of greedy label-switching algorithms. The profile likelihood was derived for the block model by Bickel and Chen (2009) and for the degree-corrected block model by Karrer and Newman (2011), but the label-switching greedy search algorithms only scale up to a few thousand nodes. Amini et al. (2013) proposed a much faster pseudo-likelihood algorithm for fitting both the regular and the degree-corrected block model, which is based on compressing A into block sums and modeling them as a Poisson mixture. Another fast algorithm for the block model based on belief propagation has

been proposed by Decelle et al. (2012). Both these algorithms rely heavily on the particular form of the block model likelihood and are not easily generalizable.

The block model likelihood is just one example of a function that can be optimized over all possible node labels in order to perform community detection. Many other functions have been proposed for this purpose, often not tied to a generative network model. One of the best-known such functions is modularity (Newman and Girvan, 2004; Newman, 2006). The key idea of modularity is to compare the observed network to a null model that has no community structure. To define this, let e be an n-dimensional label vector, $n_k(e) = \sum_{i=1}^{n} I\{e_i = k\}$ the number of nodes in community k,

$$O_{kl}(e) = \sum_{i,j=1}^{n} A_{ij} I\{e_i = k, e_j = l\} \tag{1.1}$$

the number of edges between community k and l, and $O_k = \sum_{l=1}^{K} O_{kl}$ the sum of node degrees in community k. Let $d_i = \sum_{j=1}^{n} A_{ij}$ be the degree of node i, and $m = \sum_{i=1}^{n} d_i$ be (twice) the total number of edges in the graph. The Newman–Girvan modularity is derived by comparing the observed number of edges within communities to the number that would be expected under the Chung-Lu model for the entire graph, and can be written in the form

$$Q_{NG}(e) = \frac{1}{2m} \sum_{k} (O_{kk} - \frac{O_k^2}{m}) \tag{1.2}$$

The quantities O_{kl} and O_k turn out to be the key component of many community detection criteria. The profile likelihoods of the stochastic block models (BM) and the degree-corrected block models (DCBM) discussed above can also be expressed in this form, as

$$Q_{BM}(e) = \sum_{k,l=1}^{K} O_{kl} \log \frac{O_{kl}}{n_k n_l} \ , \tag{1.3}$$

$$Q_{DCBM}(e) = \sum_{k,l=1}^{K} O_{kl} \log \frac{O_{kl}}{O_k O_l} \ . \tag{1.4}$$

Another example is the extraction criterion proposed by Zhao et al. (2011) to extract one community at a time, allowing for arbitrary structure in the remainder of the network. The main idea of extraction is to recognize that that some nodes may not belong to any community, and the strength

of a community should depend on ties between its members and ties to the outside world, but not on ties between non-members. This criterion is therefore not symmetric with respect to communities, unlike the criteria discussed, and has the form (using slightly different notation due to lack of symmetry),

$$Q_{EXTR}(V) = |V||V^c| \left(\frac{O(V)}{|V|^2} - \frac{B(V)}{|V||V^c|} \right), \tag{1.5}$$

where V is the set of nodes in the community to be extracted, V^c is the complement of V, $O(V) = \sum_{i,j \in V} A_{ij}$, $B(V) = \sum_{i \in V, j \in V^c} A_{ij}$. The only known method for optimizing this criterion is through greedy label switching, such as the tabu search algorithm (Glover and Lagunas, 1997).

For all these methods, finding the exact solution requires optimizing a function of the adjacency matrix A over all possible label vectors, which is an infeasible optimization problem. In another line of work, spectral decompositions have been used in various ways to obtain approximate solutions that are much faster to compute. A generic algorithm relying on a spectral decomposition is spectral clustering (see, for example, Ng et al. (2001)), a general clustering method which became popular for community detection. In this context, the method has been analyzed by Rohe et al. (2011), Chaudhuri et al. (2012), Riolo and Newman (2012), and Lei and Rinaldo (2013), among others, while Jin (2012) proposed a spectral method specifically for the degree-corrected block model. In spectral clustering, typically one first computes the normalized Laplacian matrix $L = D^{-1/2} A D^{-1/2}$, where D is a diagonal matrix with diagonal entries being node degrees d_i, though other normalizations and no normalization at all are also possible (see Sarkar and Bickel (2013) for an analysis of why normalization is beneficial). Then the K eigenvectors of the Laplacian corresponding to the first K largest eigenvalues are computed, and their rows clustered using K-means into K clusters corresponding to different labels. It has been shown that spectral clustering performs better with further regularization, namely if a small constant is added either to D or to A (Chaudhuri et al., 2012; Amini et al., 2013). This phenomenon was further analyzed by Qin and Rohe (2013) and Joseph and Yu (2013).

The contribution of this paper is a general method of optimizing a function $f(A, e)$ over labels e by projecting the set of labels onto the subspace spanned by leading eigenvectors of A. The main idea of the method is that projecting the feasible set of labels onto a low-dimensional space makes it much smaller and the optimization problem much easier. This approach is distinct from spectral clustering since one can specify any objective function f to be optimized (as long as it satisfies some fairly general conditions),

and thus applicable to a wide range of network problems. We show how our method can be applied to maximize the likelihoods of the stochastic block model and its degree-corrected version, Newman-Girvan modularity, and community extraction, which all solve different network analysis problems. While spectral approximations to some specific criteria that can otherwise be only maximized by a search over labels have been obtained on a case-by-case basis (Newman, 2006; Riolo and Newman, 2012; Newman, 2013), ours is, to the best of our knowledge, the first general method that would apply to any function of the adjacency matrix. In this paper, we focus on the case of two communities ($K = 2$). For methods that are run recursively, such as modularity and community extraction, this is not a restriction. For the block model, the case $K = 2$ is of special interest and has received a lot of attention in the probability literature (see Mossel et al. (2013) for recent advances). For both the regular and the degree-corrected block models, the general case can be treated similarly for the most part in our method but becomes cumbersome, and we leave it for future work.

The rest of the paper is organized as follows. In Section 2, we set up notation and describe our general approach to solving a class of optimization problems over label assignments via projection onto a low-dimensional subspace. In Section 3, we show how the general method can be applied to several community detection criteria. Section 4 compares numerical performance of different methods. The proofs are given in the Appendix.

2. A general method for optimization via low-rank approximation. It is often the case that a community detection method is equivalent to the problem of maximizing an objective function $f(A, e) \equiv f_A(e)$ over a discrete set of labels e, for example, $\{-1, 1\}^n$ for detecting two communities. Since A can be thought of as a noisy realization of $\mathbb{E}[A]$, the "ideal" solution corresponds to maximizing $f_{\mathbb{E}[A]}(e)$ instead of maximizing $f_A(e)$. For a natural class of functions f described below, $f_{\mathbb{E}[A]}(e)$ is essentially a function over the set of projections of labels e onto the subspace spanned by eigenvectors of $\mathbb{E}[A]$ and possibly some other known vectors. In many cases $\mathbb{E}[A]$ is a low-rank matrix, which makes $f_{\mathbb{E}[A]}(e)$ a function of only a few variables. It is then much easier to investigate the behavior of $f_{\mathbb{E}[A]}(e)$, which typically achieves its maximum on the set of extreme points of the convex hull generated by the projection of the label set e. Further, Proposition 1 shows that most of the 2^n possible label assignments e become interior points after the projection, and in fact the number of extreme points is at most polynomial in n. Therefore, we can find the maximum simply by performing an exhaustive search over the labels corresponding to the extreme

points. Section 3.5 provides an alternative method to the exhaustive search, which is faster but approximate.

In reality we do not have access to eigenvectors of $\mathbb{E}[A]$, so we use the eigenvectors of A instead. The algorithm works as follows.

1. Compute leading eigenvectors of A.
2. Find the labels e associated with the extreme points of the projection of the cube $[-1,1]^n$ onto the leading eigenvectors of A.
3. Find maximum of $f_A(e)$ by performing an exhaustive search over the set of labels found in step 2.

Note that the first step of replacing eigenvectors of $\mathbb{E}[A]$ with those of A is very similar to spectral clustering. However, this is where the similarity ends, because instead of following the dimension reduction by an ad-hoc clustering procedure like K-means, we proceed to maximize the original objective function over a particular subset of labels. This subset of labels is found by taking into account the specific behavior of $f_{\mathbb{E}[A]}(e)$, and in turn, the behavior of $f_A(e)$.

While our goal in the context of community detection is to compare $f_A(e)$ to $f_{\mathbb{E}[A]}(e)$, the results and the algorithm in this section apply in a more general setting, where A may be any deterministic symmetric matrix. To emphasize this generality, we write all the results in this section for a generic matrix A and a generic low-rank matrix B, even though we will later apply them to an adjacency matrix A and $B = \mathbb{E}[A]$.

Let A and B be $n \times n$ symmetric matrices with entries bounded by an absolute constant, and assume B has rank $m \ll n$. Assume that $f_A(e)$ has the general form (see Section 3 for specific examples of such functions)

$$f_A(e) = \sum_{j=1}^{\kappa} g_j(h_{A,j}(e)), \tag{2.1}$$

where g_j are scalar functions on $\mathbb{R}$ and $h_j(e)$ are quadratic forms of A and e, namely

$$h_{A,j}(e) = (e + s_{j1})^T A(e + s_{j2}). \tag{2.2}$$

Here κ is a fixed number, s_{j1} and s_{j2} are constant vectors in $\{-1,1\}^n$. We similarly define f_B and $h_{B,j}$, by replacing A with B in (2.1) and (2.2). By allowing e to take values on the cube $[-1,1]^n$, we can treat h and f as functions over $[-1,1]^n$. Note that the algorithm does not use this relaxation.

Let U_A and U_B be $m \times n$ matrices whose rows are the m leading eigenvectors of A and B, respectively. For any $e \in [-1,1]^n$, $U_A(e)$ and $U_B(e)$ are the

coordinates of the projections of e onto the column spaces of U_A and U_B, respectively. Since $h_{B,j}$ are quadratic forms of B and e and B is of rank m, $h_{B,j}$'s depend on e through $U_B(e)$ only, and therefore f_B also depends on e only through $U_B(e)$. In a slight abuse of notation, we also use $h_{B,j}$ and f_B to denote the corresponding induced functions on $U_B[-1,1]^n$.

Let $\mathcal{E}_A$ and $\mathcal{E}_B$ denote the subsets of labels $e \in \{-1,1\}^n$ corresponding to the sets of extreme points of $U_A[-1,1]^n$ and $U_B[-1,1]^n$, respectively. The output of our algorithm is

$$e^* = \operatorname{argmax}\big\{f_A(e), e \in \mathcal{E}_A\big\}. \tag{2.3}$$

Note that every $x \in U_A[-1,1]^n$ has the form $x = s_1v_1 + ... + s_nv_n$, where v_i are column vectors of U_A and $s_i \in [-1,1]$. Proposition 1 shows that the number of points in $\mathcal{E}_A$ and $\mathcal{E}_B$ is at most polynomial in n.

PROPOSITION 1. *Let $x_1, x_2, ..., x_n$ be vectors in $\mathbb{R}^m$ and*

$$\mathcal{M}_{m,n} = \Big\{s_1x_1 + s_2x_2 + \cdots + s_nx_n : s_i \in [-1,1], 1 \le i \le n\Big\}.$$

Then the number of extreme points of $\mathcal{M}_{m,n}$ is bounded by $2n^{m-1}$.

The proof of Proposition 1 by induction on m and n is given in Appendix A.

Our goal is to get a bound on the difference between the maxima of f_A and f_B that can be expressed through some measure of difference between A and B themselves. In order to do this, we make the following assumptions.

(1) Functions g_j are continuously differentiable and there exists $M_1 > 0$ such that $|g_j'(t)| \le M_1 \log(t+2)$ for $t \ge 0$.
(2) Function f_B is convex on $U_B[-1,1]^n$.

Assumption (1) essentially means that Lipschitz constants of g_j do not grow faster than $\log(t+2)$. The convexity of f_B in assumption (2) ensures that f_B achieves its maximum on $U_B\mathcal{E}_B$.

Let $c \in \{-1,1\}^n$ be the maximizer of f_B over the set of label vectors $\{-1,1\}^n$. As a function on $U_B[-1,1]^n$, f_B achieves its maximum at $U_B(c)$, which is an extreme point of $U_B[-1,1]^n$ by assumption (2). Lemma 1 provides a upper bound for $f_A(c) - f_A(e^*)$.

Throughout the paper, we write $\|\cdot\|$ for the l_2 norm (i.e., Euclidean norm on vectors and the spectral norm on matrices), and $\|\cdot\|_F$ for the Frobenius norm on matrices. Note that for label vectors $e, c \in \{-1,1\}^n$, $\|e-c\|^2$ is four times the number of nodes on which e and c differ.

LEMMA 1. *If assumptions (1) and (2) hold then there exists a constant* $M_2 > 0$ *such that*

$$f_T(c) - f_T(e^*) \leq M_2 n \log(n)\big(\|B\| \cdot \|U_A - U_B\| + \|A - B\|\big), \tag{2.4}$$

where T *is either* A *or* B.

Note that in some cases (see Section 3), convexity of f_B can be replaced with a weaker condition. The proof of Lemma 1 is given in Appendix A. To get a bound on $\|c - e^*\|$, we need further assumptions on B and f_B.

(3) There exists $M_3 > 0$ such that for any $e \in \{-1, 1\}^n$,

$$\|c - e\|^2 \leq M_3\sqrt{n}\|U_B(c) - U_B(e)\|.$$

(4) There exists $M_4 > 0$ such that for any $x \in U_B[-1, 1]^n$

$$\frac{f_B(U_B(c)) - f_B(x)}{\|U_B(c) - x\|} \geq \frac{\max f_B - \min f_B}{M_4\sqrt{n}}.$$

Assumption (3) rules out the existence of multiple label vectors with the same projection $U_B(c)$. Assumption (4) implies that the slope of the line connecting two points on the graph of f_B at $U_B(c)$ and at any $x \in U_B[-1, 1]^n$ is bounded from below. Thus, if $f_B(x)$ is close to $f_B(U_B(c))$ then x is also close to $U_B(c)$. These assumptions are satisfied for all functions considered in Section 3.

THEOREM 1. *If assumptions (1)–(4) hold, then there exists a constant* M_5 *such that*

$$\|e^* - c\|^2 \leq \frac{M_5 n^2 \log(n)\big(\|B\| \cdot \|U_A - U_B\| + \|A - B\|\big)}{\max f_B - \min f_B}.$$

Theorem 1 follows directly from Lemma 1 and assumptions (3) and (4). When A is a random matrix and $B = \mathbb{E}[A]$, a standard bound on $\|A - B\|$ can be applied (see Lemma 6), which in turn yields a bound on $\|U_A - U_B\|$ by the Davis-Kahan Theorem. Under certain conditions, the upper bound in Theorem 1 is of order $o(n)$ (see Section 3), which shows consistency of e^* as an estimator of c (i.e., the fraction of mislabeled nodes goes to 0 as $n \to \infty$).

3. Applications to community detection. In this section we apply the general results from Section 2 to a network adjacency matrix A and $B = \mathbb{E}[A]$, for the purpose of maximizing several community detection criteria. Our goal is to show that applying our maximization method to A gets us an estimate close to the true label vector c, which is the maximizer of the corresponding function with B plugged in for A. We focus on the case of two communities and use $m = 2$ for our method.

Note that the quantities O_{11}, O_{22}, and O_{12} defined in (1.1) are quadratic forms of A and e, and can be written as

$$\begin{aligned} O_{11}(e) &= \frac{1}{4}(1+e)^T A(1+e), \quad O_{22}(e) = \frac{1}{4}(1-e)^T A(1-e), \\ O_{12}(e) &= \frac{1}{4}(1+e)^T A(1-e). \end{aligned}$$

3.1. *Maximizing the likelihood of the degree-corrected block model.* When a network has two communities, (1.4) takes the form

$$\begin{aligned} Q_{DCBM}(e) &= O_{11}\log O_{11} + O_{22}\log O_{22} + 2O_{12}\log O_{12} \\ &- 2O_1\log O_1 - 2O_2\log O_2. \end{aligned}$$

Thus, Q_{DCBM} has the form defined by (2.1).

We parameterize the degree corrected block model as follows. For simplicity, instead of drawing c from a multinomial distribution with parameter $\pi = (\pi_1, \pi_2)$, we fix the true label vector by assigning the first $\bar{n}_1 = n\pi_1$ nodes to community 1 and the remaining $\bar{n}_2 = n\pi_2$ nodes to community 2. Let r be the out-in probability ratio, and

$$P = \lambda_n \begin{pmatrix} 1 & r \\ r & \omega \end{pmatrix}$$

be the probability matrix. Denote by θ_i the node degree parameters. We assume that θ_i is a bounded (away from zero and infinity) i.i.d sample from a distribution with $\mathbb{E}[\theta_i] = 1$. The adjacency matrix A is symmetric and for $i > j$ has independent entries generated by $A_{ij} = \text{Bernoulli}(\theta_i\theta_j P_{c_i c_j})$. Throughout the paper, we let λ_n depend on n, and fix r, ω, and π. Since λ_n and the network expected node degree are of the same order, with little abuse of notation, we also denote by λ_n the network expected node degree. The regular block model is a special case when $\theta_i = 1$ for all i.

Theorem 2 shows the consistency of our method in this setting. The upper bound in Theorem 2 depends on the parameters of the network, and can be computed explicitly.

THEOREM 2. *Assume A is generated from a degree corrected block model with expected node degree λ_n growing faster than $\log^3(n)$ as $n \to \infty$. Let $\delta \in (0,1)$ and e^* be defined by (2.3) with $f_A = Q_{DCBM}$. Then with probability at least $1-\delta$, the fraction of mis-clustered nodes goes to zero as $n \to \infty$, namely*

$$\frac{1}{n}\|c - e^*\|^2 = O\left(\frac{\log(n)\sqrt{\log(n/\delta)}}{\sqrt{\lambda_n}}\right).$$

Note that assumption (2) is difficult to check for Q_{DCBM} but a weaker version, namely convexity along a certain direction, is sufficient for proving Theorem 2. The proof of Theorem 2 consists of checking assumptions (1), (3), (4), and a weaker version of assumption (2). For details, see the supplement.

3.2. *Maximizing the likelihood of the block model.* When a network has two communities, (1.3) admits the form

$$Q_{BM}(e) = Q_{DCBM}(e) + 2O_1 \log\frac{O_1}{n_1} + 2O_2 \log\frac{O_2}{n_2},$$

where $n_1 = n_1(e)$ and $n_2 = n_2(e)$ are the numbers of nodes in two communities and can be written as

$$n_1 = \frac{1}{2}(\mathbf{1}+e)^T\mathbf{1} = \frac{1}{2}(n + e^T\mathbf{1}), \quad n_2 = \frac{1}{2}(\mathbf{1}-e)^T\mathbf{1} = \frac{1}{2}(n - e^T\mathbf{1}). \tag{3.1}$$

THEOREM 3. *Assume A is generated from a block model with expected node degree λ_n growing faster than $\log^3(n)$ as $n \to \infty$. Let $\delta \in (0,1)$ and e^* be defined by (2.3) with $f_A = Q_{BM}$. Then with probability at least $1-\delta$, the fraction of mis-clustered nodes goes to zero as $n \to \infty$, namely*

$$\frac{1}{n}\|c - e^*\|^2 = O\left(\frac{\log(n)\sqrt{\log(n/\delta)}}{\sqrt{\lambda_n}}\right).$$

Note that Q_{BM} does not have the exact form of (2.1) but a small modification shows that Lemma 1 still holds for Q_{BM}. Also, assumption (2) is difficult to check for Q_{BM} but again a weaker condition of convexity along a certain direction is sufficient for proving Theorem 3. The proof of Theorem 3 consists of showing the analog of Lemma 1, checking assumptions (3), (4), and a weaker version of assumption (2). For details, see the supplement.

3.3. *Maximizing the Newman–Girvan modularity.* When a network has two communities, up to a factor of $1/2m$, the modularity in (1.2) takes the form

$$\begin{aligned} Q_{NG}(e) &= O_{11} + O_{22} - \frac{O_1^2 + O_2^2}{O_1 + O_2} \\ &= \frac{2O_1O_2}{O_1 + O_2} - 2O_{12}. \end{aligned}$$

Note that Q_{NG} does not have the exact form (2.1), but with a small modification, the argument used for proving Lemma 1 and Theorem 1 still holds for Q_{NG} under regular block models.

THEOREM 4. *Assume A is generated from a degree corrected block model with expected node degree λ_n growing faster than $\log(n)$ as $n \to \infty$. Let $\delta \in (0,1)$ and e^* be defined by (2.3) with $f_A = Q_{NG}$. Then with probability at least $1-\delta$, the fraction of mis-clustered nodes goes to zero as $n \to \infty$, namely*

$$\frac{1}{n}\|c - e^*\|^2 = O\left(\sqrt{\frac{\log(n/\delta)}{\lambda_n}}\right).$$

It is easy to see that Q_{NG} is Lipschitz with respect to O_1, O_2, and O_{12}, which is stronger than assumption (1) and ensures that the argument in the proof of Lemma 1 holds. The proof of Theorem 4 consists of checking assumptions (2), (3), (4), and the Lipschitz condition for Q_{NG}. For details, see the supplement.

3.4. *Maximizing the community extraction criterion.* The community to be extracted V can be identified with the label vector e. The criterion (1.5) can be written as

$$Q_{EXTR}(e) = \frac{n_2}{n_1}O_{11} - O_{12},$$

where n_1, n_2 are defined by (3.1).

We assume that the data are generated from the regular block model described in Section 3.1 with $\omega = r$. Note that Q_{EXTR} also does not have the exact form (2.1), but with a small modification, the argument used for proving Lemma 1 and Theorem 1 still holds for Q_{EXTR}.

THEOREM 5. *Assume A is generated from a block model with $\omega = r$, such that the expected node degree λ_n grows faster than $\log(n)$ as $n \to \infty$. Let $\delta \in (0,1)$ and e^* be defined by (2.3) with $f_A = Q_{EXTR}$. Then with*

probability at least $1-\delta$*, the fraction of mis-clustered nodes goes to zero as* $n \to \infty$*, namely*

$$\frac{1}{n}\|c - e^*\|^2 = O\left(\sqrt{\frac{\log(n/\delta)}{\lambda_n}}\right).$$

The proof of Theorem 5 consists of showing the analog of Lemma 1 and assumptions (2), (3), and (4). For details, see the supplement.

3.5. *An Alternative to Exhaustive Search.* While the projected feasible space is much smaller than the original space, we may still want to avoid the exhaustive search for finding e^* in (2.3). In this section, we derive an alternative estimator of the label vector based on the geometry of the projection of the cube.

Recall that U_A and $U_{\mathbb{E}[A]}$ are $2 \times n$ matrices whose rows are the leading eigenvectors of A and $\mathbb{E}[A]$, respectively. For block models it is easy to see that $U_{\mathbb{E}[A]}[-1,1]^n$ is a parallelogram with vertices $\{\pm U_{\mathbb{E}[A]}(\mathbf{1}), \pm U_{\mathbb{E}[A]}(c)\}$ (see Lemma 8 in the supplement), where $\mathbf{1} \in \mathbb{R}^n$ is a vector of all 1s. Figure 3.1 gives an illustration of these in the projection space. Note that $\pm U_{\mathbb{E}[A]}(c)$ are the farthest points from the line connecting $U_{\mathbb{E}[A]}(\mathbf{1})$ and $-U_{\mathbb{E}[A]}(\mathbf{1})$. Motivated by that observation, we can estimate c by

$$\begin{aligned} \hat{c} &= \arg\max\left\{\langle U_A(e), (U_A\mathbf{1})^\perp)\rangle : e \in \{-1,1\}^n\right\} \\ &= \operatorname{sign}(u_1^T\mathbf{1}u_2 - u_2^T\mathbf{1}u_1), \end{aligned} \tag{3.2}$$

where $U_A = (u_1, u_2)^T$ and $(U_A(\mathbf{1}))^\perp$ is the unit vector perpendicular to $U_A(\mathbf{1})$.

Note that the estimator $\hat{c}$ does not depend on the form of the likelihood. Therefore, when the community detection problem is relatively easy, we expect it to perform well, but when the problem becomes harder, the exhaustive search should provide better results. This intuition is confirmed by simulations in Section 4. Theorem 6 shows that $\hat{c}$ is a consistent estimator. The proof is given in Appendix B.

THEOREM 6. *Assume that data are generated from the degree-corrected block models such that the expected node degree* λ_n *grows faster than* $\log(n)$ *as* $n \to \infty$*. Then the fraction of nodes mis-clustered by* $\hat{c}$ *defined in 3.2 tends to zero as* $n \to \infty$*, namely for* $\delta \in (0,1)$*, with probability at least* $1-\delta$*,*

$$\frac{1}{n}\|\hat{c} - c\|^2 = O\left(\frac{\log(n/\delta)}{\lambda_n}\right).$$

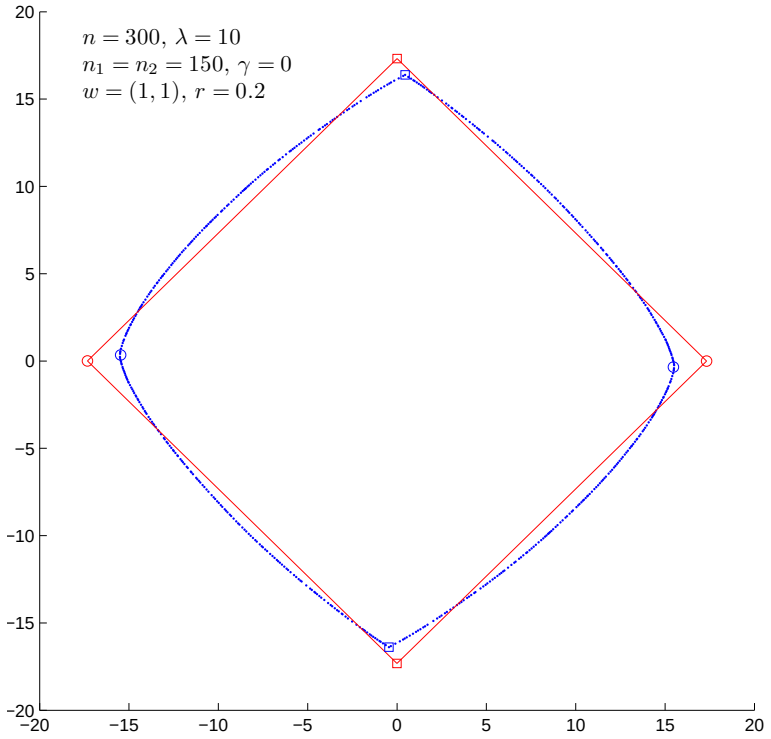


FIG 3.1. *The projection of the cube* $[-1, 1]^n$ *onto two-dimensional subspace. Blue corresponds to the projection onto eigenvectors of* A*, and red onto the eigenvectors of* $\mathbb{E}[A]$*. The red contour is the boundary of* $U_{\mathbb{E}[A]}[-1, 1]^n$*; the blue dots are the extreme points of* $U_A[-1, 1]^n$*. Circles (at the corners) are* $\pm$ *projections of the true label vector; squares are* $\pm$ *projections of the vector of all 1s.*

4. Numerical comparisons. In this section, we perform a brief numerical comparison of our extreme point projection method to several other approximate methods for community detection, both general (spectral clustering) and those designed specifically for optimizing a particular community detection criterion. The methods are compared on simulated networks for the block model fitting, degree-corrected block model fitting, modularity maximization, and community extraction. We also illustrate the different methods on two real networks, the political blogs and the dolphins datasets described in detail in Section 4.5. Since many algorithms have been developed specifically for these community detection approaches already, our goal is to show that our general method does as well as the algorithms tailored to a particular criterion, and thus we are not trading off accuracy for generality.

For the four criteria discussed in Section 3, we compare our method of maximizing the relevant criterion by exhaustive search over the set $\mathcal{E}_A$ of label vectors corresponding to the set of extreme points of $U_A[-1,1]^n$ (EP, for extreme points), the approximate version based on the geometry of the feasible set described in Section 3.5 (AEP, for approximate extreme points), and regularized spectral clustering (SCR) as proposed by Amini et al. (2013), which are all general methods. Further, we include one method specific to the criterion in each comparison. For the stochastic block model, we compare to the (unconditional) pseudo-likelihood (UPL) and for the degree-corrected version, to the conditional pseudo-likelihood (CPL), both currently considered state-of-the-art methods developed by Amini et al. (2013). For Newman-Girvan modularity, we compare to the spectral algorithm of Newman (2006), which assigns labels based on the signs of the entries of the leading eigenvector of the modularity matrix (see details in Section 4.3). Finally, for community extraction we compare to the algorithm proposed in the original paper (Zhao et al., 2011) based on greedy label switching, as there are no faster approximate algorithms available.

The simulated networks are generated using the parametrization of Amini et al. (2013), as follows. Throughout this section, the number of nodes in the network is fixed at $n = 300$, the number of communities $K = 2$, and the true label vector c is fixed. The number of replications for each setting is 100. First, the node degree parameters θ_i are drawn independently from the distribution $\mathbb{P}(\Theta = 0.2) = \gamma$, and $\mathbb{P}(\Theta = 1) = 1 - \gamma$. Setting $\gamma = 0$ gives the standard stochastic block model, and $\gamma > 0$ gives the degree-corrected block model, with $1 - \gamma$ the fraction of hub nodes. The matrix of edge probabilities P is controlled by two parameters: the out-in probability ratio r, which determines how likely edges are formed within and between communities, and the weight vector $\boldsymbol{w} = (w_1, w_2)$, which determines the relative node degrees within communities. Let

$$P_0 = \begin{bmatrix} w_1 & r \\ r & w_2 \end{bmatrix}.$$

The difficulty of the problem, in addition to r, is controlled by the overall expected network degree λ. Thus we rescale P_0 to obtained the desired overall expected degree, setting

$$P = \frac{\lambda P^0}{(n-1)(\pi^T P^0 \pi)(\mathbb{E}[\Theta])^2},$$

where $\pi = n^{-1}(n_1, n_2)$, and n_k is the number of nodes in community k. Finally, edges A_{ij} are drawn independently from a Bernoulli distribution with $\mathbb{P}(A_{ij} = 1) = \theta_i \theta_j P_{c_i c_j}$.

As pointed out by Amini et al. (2013), when the network is sparse, the leading eigenvectors of a small perturbation of A better approximate the eigenvectors of $\mathbb{E}[A]$ than those of A itself. Therefore, for all the methods using eigenvectors of the adjacency matrix (SCR, EP and AEP) we replace the leading eigenvectors of A by the leading eigenvectors of $A+\tau\mathbf{1}\mathbf{1}^T$, where τ is a tuning parameter set to $0.5\lambda/n$ (see Amini et al. (2013) for more details). For the second step of spectral clustering, we use the kmeans function in Matlab with 40 random initial starting points. For CPL and UPL, following Amini et al. (2013), we initialize with the SCR solution, and set the number of outer iterations to 20. We measure the accuracy of all the methods by the normalized mutual information (NMI) between the label vector c and its estimate e (see Strehl and Ghosh (2003)). NMI takes values between 0 (random guessing) and 1 (perfect match).

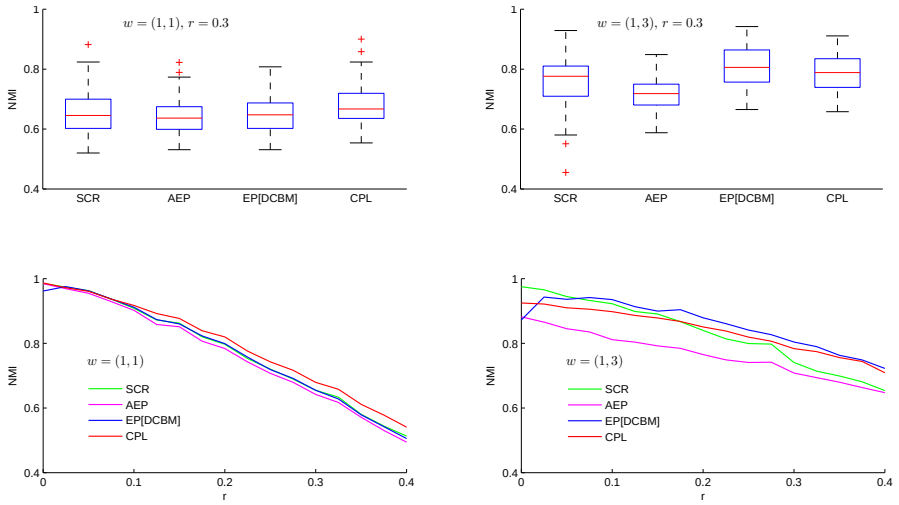


FIG 4.1. *The degree-corrected stochastic block model. Top row: boxplots of NMI between true and estimated labels. Bottom row: average NMI against the out-in probability ratio r. In all plots, $n_1 = n_2 = 150$, $\lambda = 15$, and $\gamma = 0.5$.*

4.1. *The degree-corrected stochastic block model.* Figure 4.1 shows the performance of the four methods for fitting the DCMB under different parameter settings. We use the notation EP[DCBM] to emphasize that EP here maximizes the log-likelihood of the degree corrected block model. In this case, all methods perform similarly, with EP a little better than the benchmark CPL and overall the best method when the community weights are unbalanced ($w = (1,3)$), but a little worse than CPL when $w = (1,1)$. The AEP is always somewhat worse than the exact version, especially when

$w = (1, 3)$, but overall, their results are comparable.

4.2. *The stochastic block model.* Figure 4.2 shows the performance of the four methods for fitting the regular block models ($\gamma = 0$). Over all, four methods provide quite similar results, as we would hope good fitting methods will.

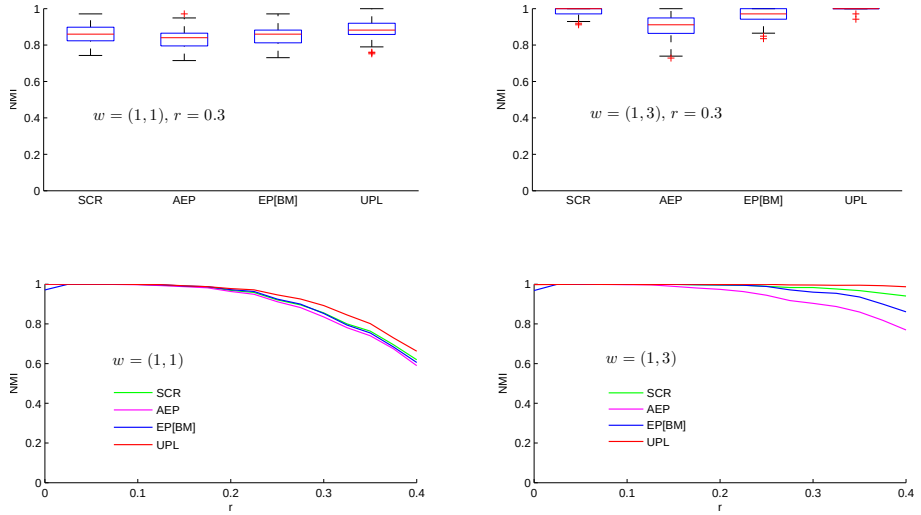


FIG 4.2. *The stochastic block model. Top row: boxplots of NMI between true and estimated labels. Bottom row: average NMI against the out-in probability ratio r. In all plots, $n_1 = n_2 = 150$, $\lambda = 15$, and $\gamma = 0$.*

4.3. *Newman–Girvan Modularity.* Newman (2006) derives an approximate spectral solution to the problem of maximizing $\hat{Q}_{NG}$ when the network has two communities. Let $B = A - P$ where $P_{ij} = d_i d_j / m$, then $\hat{Q}_{NG}$ can be written as $\hat{Q}_{NG}(e) = \frac{1}{2m} e^T B e$. The approximate solution (LES, for leading eigenvector signs) assigns node labels according to the signs of the corresponding entries of the leading eigenvector of B. For a fair comparison to our other methods, we also use $A + \tau \mathbf{1}\mathbf{1}^T$ instead of A here, since empirically we found that adding a small perturbation to A slightly improves the performance of LES. Figure 4.3 shows the performance of AEP, EP[NG], and LES, when data are generated from a regular block model ($\gamma = 0$). Over all, the four methods provide comparable results, with EP[NG] again providing the best solution for the unbalanced case of $w = (1, 3)$.

4.4. *Community Extraction Criterion.* Following the original extraction paper of Zhao et al. (2011), we generate a community with background

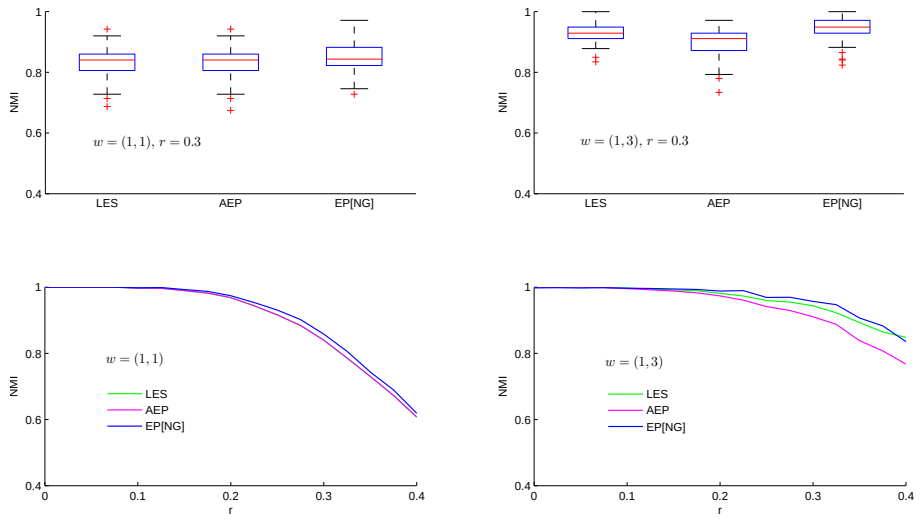


FIG 4.3. *Newman-Girvan modularity. Top row: boxplots of NMI between true and estimated labels. Bottom row: average NMI against the out-in probability ratio r. In all plots, $n_1 = n_2 = 150$, $\lambda = 15$, and $\gamma = 0$.*

from the regular block model with two communities as follows: set $n_1 = 60$, $n_2 = 240$, and take the probability matrix proportional to

$$P_0 = \begin{pmatrix} 0.4 & 0.1 \\ 0.1 & 0.1 \end{pmatrix}.$$

Thus, nodes within the first community are tightly connected, while the rest of the nodes have equally weak links with all other nodes and represent the background. The average expected node degree is varied in the set $[15, 20, 25, 30]$. Figure 4.4 shows that EP[EXTR] performs better than SC, AEP, but slightly worse than the Tabu search in maximizing the community extraction criterion (TS). However, this slight improvement is achieved at a high computational cost, since tabu search is very computationally intensive.

4.5. *Real-world Network Data.* In this section we apply different methods to two real-world networks where node labels that can be viewed as ground truth are available, so we can compare the estimated communities to this ground truth.

The first network, assembled by Adamic and Glance (2005), consists of blogs about US politics and hyperlinks between blogs. Each blog has been manually labeled as either liberal or conservative, which we use as the ground

FIG 4.4. *Community extraction. The boxplots of NMI between true and estimated labels. In all plots,* $n_1 = 60$*,* $n_2 = 240$*, and* $\gamma = 0$*.*

truth. Following Karrer and Newman (2011), and Zhao et al. (2012), we ignore directions of the hyperlinks and only examine the largest connected component of this network, which has 1222 nodes and 16,714 edges, resulting in an average degree of approximately 27. Table 1 and Figure 4.5 show the performance of different methods. While AEP, EP[DCBM], and CPL give reasonable results, SCR, UPL, and EP[BM] clearly miscluster the nodes. This is consistent with previous analyses which showed that the degree correction has to be used for this network to achieve the correct partition.

TABLE 1
The NMI between true and estimated labels for real-world networks.

Method	SCR	AEP	EP[BM]	EP[DCBM]	UPL	CPL
Blogs	0.290	0.660	0.389	0.723	0.001	0.725
Dolphins	0.889	0.753	0.889	0.889	0.889	0.889

The second network, assembled by Lusseau et al. (2003), has as its nodes 62 bottlenose dolphins living in Doubtful Sound, New Zealand. The edges represent the social ties between dolphins. At some point one well-connected dolphin left the group, and the group split into two separate subgroups, which we use as the ground truth in this example. Table 1 and Figure 4.6 show the performance of different methods. In Figure 4.6, shapes of nodes

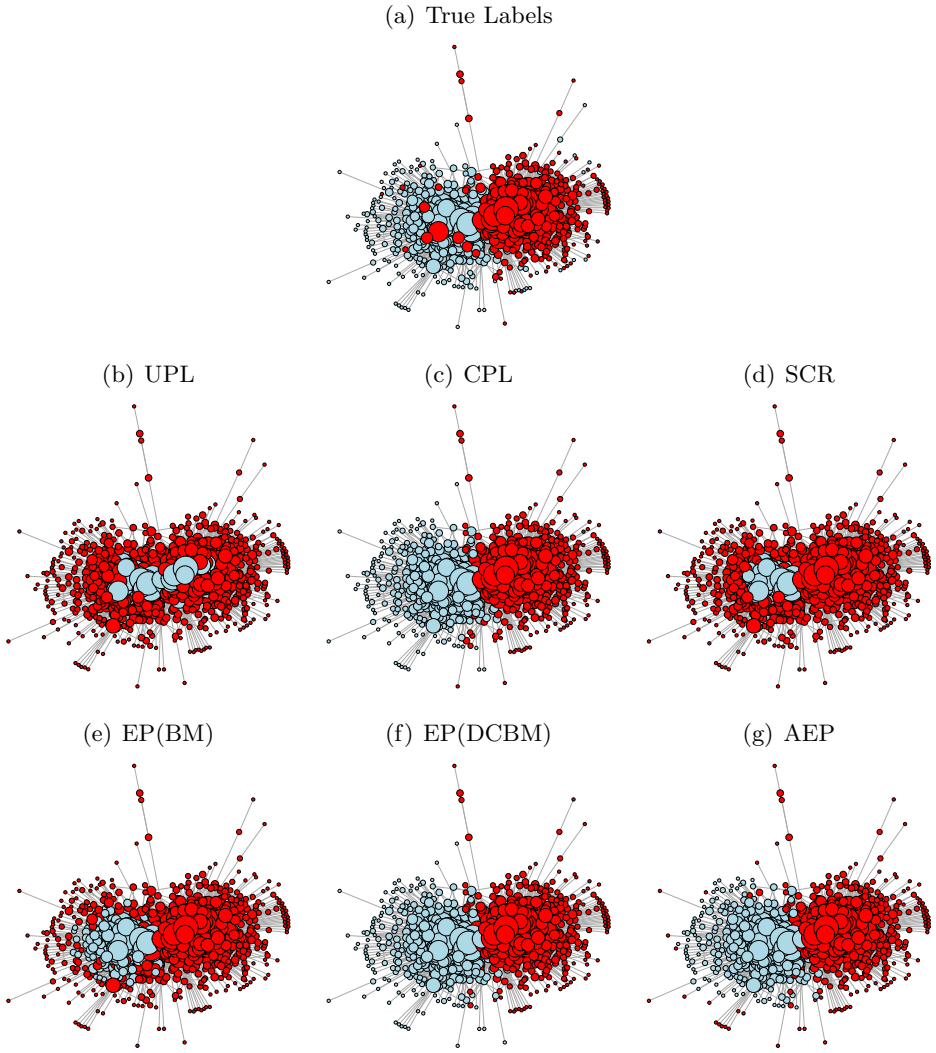


FIG 4.5. *The network of political blogs. Node diameter is proportional to the logarithm of its degree and the colors represent community labels assigned by different methods.*

represent the actual split, while colors represent the estimated labels by different methods. The star-shaped node represents the dolphin that left the group. Excepting the dolphin that left, SCR, EP[BM], EP[DCBM], UPL, and CPL miscluster one node, while AEP misclusters three nodes.

Acknowledgments. We thank Yunpeng Zhao (George Mason University) for sharing his code for the Tabu search, and Arash Amini (University of

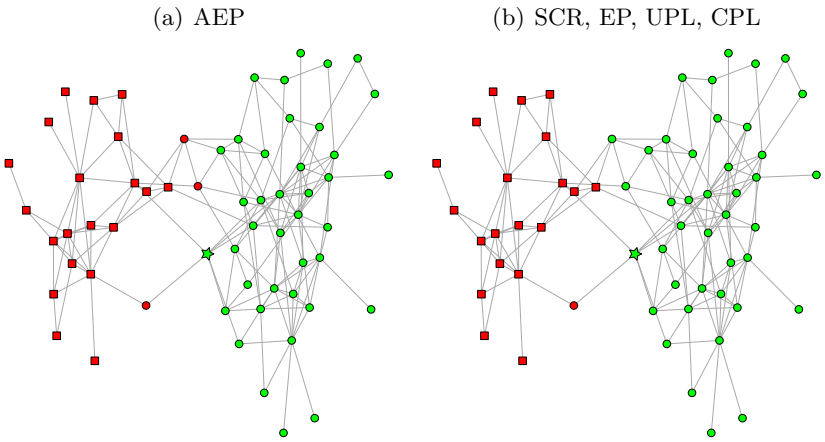


FIG 4.6. *The network of 62 bottlenose dolphins. Shapes of nodes represent the actual split of the group into two smaller subgroups after one dolphin, represented by the star, left the group. Colors represent the estimated labels.*

Michigan) for sharing his code for the pseudo-likelihood methods and helpful discussions. E.L. is partially supported by NSF grants DMS-01106772 and DMS-1159005. R.V. is partially supported by NSF grants DMS 1161372, 1001829, 1265782 and USAF Grant FA9550-14-1-0009.

References.

Adamic, L. A. and Glance, N. (2005). The political blogosphere and the 2004 US election. In *Proceedings of the WWW-2005 Workshop on the Weblogging Ecosystem.*

Airoldi, E. M., Blei, D. M., Fienberg, S. E., and Xing, E. P. (2008). Mixed membership stochastic blockmodels. *J. Machine Learning Research*, 9:1981–2014.

Amini, A., Chen, A., Bickel, P., and Levina, E. (2013). Fitting community models to large sparse networks. *Annals of Statistics*, 41(4):2097–2122.

Ball, B., Karrer, B., and Newman, M. E. J. (2011). An efficient and principled method for detecting communities in networks. *Physical Review E*, 34:036103.

Bhatia, R. (1996). *Matrix Analysis.* Springer-Verlag New York.

Bickel, P. J. and Chen, A. (2009). A nonparametric view of network models and Newman-Girvan and other modularities. *Proc. Natl. Acad. Sci. USA*, 106:21068–21073.

Chaudhuri, K., Chung, F., and Tsiatas, A. (2012). Spectral clustering of graphs with general degrees in the extended planted partition model. *JMLR Workshop and Conference Proceedings*, 23:35.1 – 35.23.

Chung, F. and Lu, L. (2002). Connected components in random graphs with given degree sequences. *Annals of Combinatorics*, 6:125–145.

Decelle, A., Krzakala, F., Moore, C., and Zdeborová, L. (2012). Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84:066106.

Erdős, P. and Rényi, A. (1959). On random graphs. I. *Publ. Math. Debrecen*, 6:290–297.

Glover, F. W. and Lagunas, M. (1997). *Tabu search.* Kluwer Academic.

Goldenberg, A., Zheng, A. X., Fienberg, S. E., and Airoldi, E. M. (2010). A survey of statistical network models. *Foundations and Trends in Machine Learning*, 2:129–233.

Handcock, M. D., Raftery, A. E., and Tantrum, J. M. (2007). Model-based clustering for social networks. *J. R. Statist. Soc. A*, 170:301–354.

Hoff, P. D., Raftery, A. E., and Handcock, M. S. (2002). Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97:1090–1098.

Holland, P. W., Laskey, K. B., and Leinhardt, S. (1983). Stochastic blockmodels: first steps. *Social Networks*, 5(2):109–137.

Jin, J. (2012). Fast network community detection by score. *arXiv:1211.5803*.

Joseph, A. and Yu, B. (2013). Impact of regularization on spectral clustering. *arXiv:1312.1733*.

Karrer, B. and Newman, M. E. J. (2011). Stochastic blockmodels and community structure in networks. *Physical Review E*, 83:016107.

Lei, J. and Rinaldo, A. (2013). Consistency of spectral clustering in sparse stochastic block models. arxiv:1312.2050.

Lusseau, D., Schneider, K., Boisseau, O. J., Haase, P., Slooten, E., and Dawson, S. M. (2003). The bottlenose dolphin community of doubtful sound features a large proportion of long-lasting associations. can geographic isola- tion explain this unique trait? *Behavioral Ecology and Sociobiology*, 54:396–405.

Mariadassou, M., Robin, S., and Vacher, C. (2010). Uncovering latent structure in valued graphs: A variational approach. *The Annals of Applied Statistics*, 4(2):715–742.

Mossel, E., Neeman, J., and Sly, A. (2013). Belief propagation, robust reconstruction, and optimal recovery of block models. arXiv:1309.1380.

Newman, M. E. J. (2006). Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E*, 74(3):036104.

Newman, M. E. J. (2013). Spectral methods for network community detection and graph partitioning. arxiv:1307.7729.

Newman, M. E. J. and Girvan, M. (2004). Finding and evaluating community structure in networks. *Phys. Rev. E*, 69(2):026113.

Ng, A., Jordan, M., and Weiss, Y. (2001). On spectral clustering: Analysis and an algorithm. In Dietterich, T., Becker, S., and Ghahramani, Z., editors, *Neural Information Processing Systems 14*, pages 849–856. MIT Press.

Nowicki, K. and Snijders, T. A. B. (2001). Estimation and prediction for stochastic blockstructures. *Journal of the American Statistical Association*, 96(455):1077–1087.

Oliveira, R. (2010). Concentration of the adjacency matrix and of the laplacian in random graphs with independent edges. *Arxiv preprint*.

Qin, T. and Rohe, K. (2013). Regularized spectral clustering under the degree-corrected stochastic blockmodel. *arXiv:1309.4111*.

Riolo, M. and Newman, M. E. J. (2012). First-principles multiway spectral partitioning of graphs. arXiv:1209.5969.

Rohe, K., Chatterjee, S., and Yu, B. (2011). Spectral clustering and the high-dimensional stochastic block model. *Annals of Statistics*, 39(4):1878–1915.

Sarkar, P. and Bickel, P. (2013). Role of normalization in spectral clustering for stochastic blockmodels. arXiv:1310.1495.

Snijders, T. and Nowicki, K. (1997). Estimation and prediction for stochastic blockstructures for graphs with latent block structure. *Journal of Classification*, 14:75–100.

Strehl, A. and Ghosh, J. (2003). Cluster ensembles–a knowledge reuse framework for combining multiple partitions. *The Journal of Machine Learning Research*, 3:583–617.

Zhao, Y., Levina, E., and Zhu, J. (2011). Community extraction for social networks. *Proc. Natl. Acad. Sci. USA*, 108(18):7321–7326.

Zhao, Y., Levina, E., and Zhu, J. (2012). Consistency of community detection in networks under degree-corrected stochastic block models. *Annals of Statistics*, 40(4):2266–2292.

APPENDIX A: OPTIMIZATION VIA LOW-RANK APPROXIMATION

PROOF OF PROPOSITION 1. Denote by $\text{Ext}(\mathcal{M}_{m,n})$ the set of all extreme points of $\mathcal{M}_{m,n}$. We prove Proposition 1 by induction on m and n. For $m = 1$, it is trivial that $|\text{Ext}(\mathcal{M}_{1,n})| = 2$ for all n. For $m = 2$, each time when we add a new vector to $\mathcal{M}_{m,n}$, the number of extreme points of the set $s_{n+1}x_{n+1} + \mathcal{M}_{m,n}$ is at most $|\text{Ext}(\mathcal{M}_{m,n})| + 2$. Therefore $|\text{Ext}(\mathcal{M}_{m,n})| \leq 2n$.

For $m \geq 3$, suppose that $|\text{Ext}(\mathcal{M}_{m,n})| < 2n^{m-1}$ holds for every n. We show that for every n, the inequality $|\text{Ext}(\mathcal{M}_{(m+1),n})| < 2n^m$ also holds. For each k consider the set

$$\mathcal{M}_{m+1,k+1} = s_{k+1}x_{k+1} + \mathcal{M}_{m+1,k}.$$

Let P be the projector onto the m-dimensional subspace orthogonal to x_{k+1} and $\mathcal{N}_k = P\mathcal{M}_{m+1,k+1}$. Since

$$\mathcal{N}_k = \Big\{ s_1 Px_1 + s_2 Px_2 + \cdots + s_k Px_k : s_i \in [-1, 1], i \leq k \Big\}$$

and $Px_1, Px_2, ..., Px_k$ are vectors in an m-dimensional subspace, $\mathcal{N}_k$ and $\mathcal{M}_{m,k}$ have the same form. Therefore by induction $|\text{Ext}(\mathcal{N}_k)| \leq 2k^{m-1}$ and so

$$|\text{Ext}(\mathcal{M}_{m+1,k+1})| \leq |\text{Ext}(\mathcal{M}_{m+1,k})| + |\text{Ext}(\mathcal{N}_k)| \leq |\text{Ext}(\mathcal{M}_{m+1,k})| + 2k^{m-1}.$$

The last inequality shows that by changing from k to $k + 1$, the number of extreme points increases by at most $2k^{m-1}$. Therefore

$$|\text{Ext}(\mathcal{M}_{(m+1),n})| \leq 2\sum_{k=1}^{n-1} k^{m-1} \leq 2n^m,$$

which completes the proof. □

The following Lemma bounds the Lipschitz constants of $h_{B,j}$ and f_B on $U_B[-1, 1]^n$.

LEMMA 2. *Assume that Assumption (1) hold. For any $j \leq \kappa$ (see 2.1), and $x, y \in U_B[-1, 1]^n$, we have*

$$\begin{aligned} \big|h_{B,j}(x) - h_{B,j}(y)\big| &\leq 4\sqrt{n}\|B\| \cdot \|x - y\|, \\ \big|f_B(x) - f_B(y)\big| &\leq M\sqrt{n}\log(n)\|B\| \cdot \|x - y\|, \end{aligned}$$

where M is a constant independent of n.

PROOF OF LEMMA 2. Let $e, s \in [-1,1]^n$ such that $x = U_B(e), y = U_B(s)$ and denote $L = \left|h_{B,j}(x) - h_{B,j}(y)\right|$. Then

$$\begin{aligned} L &= \left|(e+s_{j1})^T B(e+s_{j2}) - (s+s_{j1})^T B(s+s_{j2})\right| \\ &= \left|e^T B(e-s) + (e-s)^T Bs + (s_{j2}+s_{j1})^T B(e-s)\right| \\ &\le 4\sqrt{n}\|B(e-s)\|. \end{aligned}$$

Let $B = \sum_{i=1}^m \rho_i u_i u_i^T$ be the eigendecomposition of B. Then

$$\begin{aligned} \|B(e-s)\|^2 &= \left\|\sum_{i=1}^m \rho_i u_i u_i^T (e-s)\right\|^2 = \left\|\sum_{i=1}^m \rho_i (x_i - y_i) u_i\right\|^2 \\ &= \sum_{i=1}^m \rho_i^2 (x_i - y_i)^2 \le \|B\|^2 \sum_{i=1}^m (x_i - y_i)^2 = \|B\|^2 \cdot \|x-y\|^2. \end{aligned}$$

Therefore $L \le 4\sqrt{n}\|B\| \cdot \|x-y\|$. Since $h_{B,j}$ are quadratic, they are of order $O(n^2)$. Hence by Assumption (1), the Lipschitz constants of g_j are of order $\log(n)$. Therefore

$$\left|f_B(x) - f_B(y)\right| \le 4\sqrt{n}\log(n)\|B\| \cdot \|x-y\|,$$

which completes the proof. □

In the following proofs we use M to denote a positive constant independent of n, whose value may change from line to line.

PROOF OF LEMMA 1. Since $\|e + s_{j1}\| \le 2\sqrt{n}$ and $\|e + s_{j2}\| \le 2\sqrt{n}$,

$$\begin{aligned} |h_{A,j}(e) - h_{B,j}(e)| &= |(e+s_{j1})^T (A-B)(e+s_{j2})| \\ &\le 4n\|A-B\|. \end{aligned}$$

Since $h_{A,j}$ and $h_{B,j}$ are of order $O(n^2)$, g_j' are bounded by $\log(n)$. Together with assumption (1) it implies that there exists $M > 0$ such that

$$|f_A(e) - f_B(e)| \le Mn\log(n)\|A-B\|. \tag{A.1}$$

Let $\hat{e} = \arg\max\{f_B(e), e \in \mathcal{E}_A\}$. Then $f_A(e^*) \ge f_A(\hat{e})$ and by (A.1) we get

$$\begin{aligned} f_B(\hat{e}) - f_B(e^*) &\le f_B(\hat{e}) - f_A(\hat{e}) + f_A(e^*) - f_B(e^*) \\ &\le Mn\log(n)\|A-B\|. \end{aligned} \tag{A.2}$$

Denote by $\text{conv}(S)$ the convex hull of a set S. Then $U_A(c) \in \text{conv}(U_A\mathcal{E}_A)$ and therefore, there exists $\eta_e \geq 0$, $\sum_{e\in\mathcal{E}_A} \eta_e = 1$ such that

$$U_A(c) = \sum_{e\in\mathcal{E}_A} \eta_e U_A(e) = U_A\left(\sum_{e\in\mathcal{E}_A} \eta_e e\right).$$

Hence

$$\begin{aligned}
\text{(A.3)}\quad \text{dist}\big(U_B(c), \text{conv}(U_B\mathcal{E}_A)\big) &\leq \left\| U_B(c) - U_B\left(\sum_{e\in\mathcal{E}_A} \eta_e e\right)\right\| \\
&= \left\|(U_B - U_A)c + (U_A - U_B)\sum_{e\in\mathcal{E}_A} \eta_e e\right\| \\
&\leq 2\sqrt{n}\ \|U_A - U_B\|.
\end{aligned}$$

Let $y \in \text{conv}(U_B\mathcal{E}_A)$ be the closest point from $\text{conv}(U_B\mathcal{E}_A)$ to $U_B(c)$, i.e.

$$\|U_B(c) - y\| = \text{dist}\big(U_B(c), \text{conv}(U_B\mathcal{E}_A)\big).$$

By A.3 and Lemma 2, we have

$$\text{(A.4)}\qquad f_B(U_B(c)) - f_B(y) \leq Mn\log(n)\|B\| \cdot \|U_A - U_B\|.$$

The convexity of f_B implies that $f_B(y) \leq f_B(U_B(\hat{e}))$, and in turn,

$$\text{(A.5)}\qquad f_B(U_B(c)) - f_B(U_B(\hat{e})) \leq Mn\log(n)\|B\| \cdot \|U_A - U_B\|.$$

Note that $f_B(U_B(e)) = f_B(e)$ for every $e \in [-1,1]^n$. Adding (A.2) and (A.5), we get (2.4) for $T = B$. The case $T = A$ then follows from (A.1) because replacing B with A induces an error which is not greater than the upper bound of (2.4) for $T = B$. □

APPENDIX B: PROOF OF THEOREM 6

We first present the closed form of eigenvalues and eigenvectors of $\mathbb{E}[A]$ under the regular block models.

LEMMA 3. *Under regular block models (described in Section 3.1), the nonzero eigenvalues ρ_i and corresponding eigenvectors $\bar{u}_i$ of $\mathbb{E}[A]$ have the following form. For $i = 1, 2$,*

$$\rho_i = \frac{\lambda_n}{2}\left[(\pi_1 + \pi_2\omega) + (-1)^{i-1}\ \sqrt{(\pi_1 + \pi_2\omega)^2 - 4\pi_1\pi_2(\omega - r^2)}\right],$$

$$\bar{u}_i = \frac{1}{\sqrt{n(\pi_1 r_i^2 + \pi_2)}}(r_i, r_i, ..., r_i, 1, 1, ..., 1)^T, \text{ where}$$

$$r_i = \frac{2\pi_2 r}{(\pi_2\omega - \pi_1) + (-1)^i \sqrt{(\pi_1 + \pi_2\omega)^2 - 4\pi_1\pi_2(\omega - r^2)}}.$$

The first $\bar{n}_1 = n\pi_1$ *entries of* $\bar{u}_i$ *equal* $r_i \left(n(\pi_1 r_i^2 + \pi_2)\right)^{-1/2}$ *and the last* $\bar{n}_2 = n\pi_2$ *entries of* $\bar{u}_i$ *equal* $\left(n(\pi_1 r_i^2 + \pi_2)\right)^{-1/2}$.

Proof of Lemma 3. It is easy to verify directly that $\mathbb{E}[A]\bar{u}_i = \rho_i \bar{u}_i$ for $i = 1, 2$. □

In the degree-corrected model setting, eigenvalues and eigenvectors of $\mathbb{E}[A]$ may not have a closed form. Nevertheless, we can approximate them using ρ_i and $\bar{u}_i$ from Lemma 3. To do so, we need the following lemma.

Lemma 4. *Let* $M = \rho_1 x_1 x_1^T + \rho_2 x_2 x_2^T$, *where* $x_1, x_2 \in \mathbb{R}^n$, $\|x_1\| = \|x_2\| = 1$, $\rho_1 \neq 0$, *and* $\rho_2 \neq 0$. *If* $c = \langle x_1, x_2 \rangle$ *then the eigenvalues* z_i *and corresponding eigenvectors* y_i *of* M *have the following form. For* $i = 1, 2$,

$$\begin{aligned} z_i &= \frac{1}{2}\left[(\rho_1 + \rho_2) + (-1)^{i-1}\sqrt{(\rho_2 - \rho_1)^2 + 4\rho_1\rho_2 c^2}\right], \\ y_i &= (c\rho_1)x_1 + (z_i - \rho_1)x_2. \end{aligned}$$

If ρ_1 *and* ρ_2 *are fixed,* $\rho_1 \geq \rho_2$, *and* $c = o(1)$ *as* $n \to \infty$ *then eigenvalues and eigenvectors of* M *have the form*

$$\begin{aligned} z_1 &= \rho_1 + O(c^2), \quad z_2 = \rho_2 + O(c^2), \\ y_1 &= x_1 + O(c)x_2, \quad y_2 = x_2 + O(c)x_1. \end{aligned}$$

Proof of Lemma 4. It is easy to verify that $My_i = z_i y_i$ for $i = 1, 2$. The asymptotic formulas of z_i and y_i then follow directly from the forms of z_i and y_i. □

The next lemma shows the approximation of eigenvalues and eigenvectors of $\mathbb{E}[A]$.

Lemma 5. *Consider the degree-corrected block models (described in Section 3.1), and let* $\theta = (\theta_1, ..., \theta_n)^T$ *and* $D_\theta = \mathrm{diag}(\theta)$. *Then for any* $\delta \in (0, 1)$, *with probability at least* $1 - \delta$, *the nonzero eigenvalues* ρ_i^θ *and corresponding eigenvectors* $\bar{u}_i^\theta$ *of* $\mathbb{E}[A]$ *have the following form. For* $i = 1, 2$,

$$\rho_i^\theta = \rho_i \|D_\theta \bar{u}_i\|^2 \left(1 + O(1/n)\right),$$

$$\bar{u}_1^\theta = \frac{\tilde{u}_1^\theta}{\|\tilde{u}_1^\theta\|}, \quad \textit{where} \quad \tilde{u}_1^\theta = \frac{D_\theta \bar{u}_1}{\|D_\theta \bar{u}_1\|} + O\left(n^{-1/2}\right) \frac{D_\theta \bar{u}_2}{\|D_\theta \bar{u}_2\|},$$

$$\bar{u}_2^\theta = \frac{\tilde{u}_2^\theta}{\|\tilde{u}_2^\theta\|}, \quad \textit{where} \quad \tilde{u}_2^\theta = \frac{D_\theta \bar{u}_2}{\|D_\theta \bar{u}_2\|} + O\left(n^{-1/2}\right) \frac{D_\theta \bar{u}_1}{\|D_\theta \bar{u}_1\|},$$

where ρ_i, $\bar{u}_i$, and r_i are defined in Lemma 3.

PROOF OF LEMMA 5. Let $M = \rho_1 \bar{u}_1 \bar{u}_1^T + \rho_2 \bar{u}_2 \bar{u}_2^T$ be the expectation of the adjacency matrix in the regular block model setting. In the degree-corrected block model setting, given θ, we have

$$\begin{aligned} \mathbb{E}[A] &= D_\theta M D_\theta = \rho_1 D_\theta \bar{u}_1 (D_\theta \bar{u}_1)^T + \rho_2 D_\theta \bar{u}_2 (D_\theta \bar{u}_2)^T \\ &= \rho_1 \|D_\theta \bar{u}_1\|^2 \frac{D_\theta \bar{u}_1}{\|D_\theta \bar{u}_1\|} \frac{(D_\theta \bar{u}_1)^T}{\|D_\theta \bar{u}_1\|} + \rho_2 \|D_\theta \bar{u}_2\|^2 \frac{D_\theta \bar{u}_2}{\|D_\theta \bar{u}_2\|} \frac{(D_\theta \bar{u}_2)^T}{\|D_\theta \bar{u}_2\|}. \end{aligned}$$

We are now in the setting of Lemma 4 with

$$\begin{aligned} c &= \left(\|D_\theta \bar{u}_1\| \|D_\theta \bar{u}_2\|\right)^{-1} \langle D_\theta \bar{u}_1, D_\theta \bar{u}_2 \rangle \\ &= c_\theta \left[\pi_1 \sqrt{(\pi_1 r_1^2 + \pi_2)(\pi_1 r_2^2 + \pi_2)} \|D_\theta \bar{u}_1\| \|D_\theta \bar{u}_2\|\right]^{-1}, \\ \text{where } c_\theta &= \frac{1}{n} \left[\pi_1 (\theta_{\bar{n}_1+1}^2 + \cdots + \theta_n^2) - \pi_2 (\theta_1^2 + \cdots + \theta_{\bar{n}_1}^2)\right]. \end{aligned}$$

Note that the two sums in the formula of c_θ have the same expectation. It remains to apply Hoeffding's inequality to each sum. □

Lemma 6 estimates the error between eigenvalues and eigenvectors of A and $\mathbb{E}[A]$. It also provides a way to simplify the upper bound of Theorem 1.

LEMMA 6. *Consider the degree-corrected block model such that the expected node degree λ_n grows faster than $\log(n)$. Let U_A and $U_{\mathbb{E}[A]}$ be $2 \times n$ matrices whose rows are leading eigenvectors of A and $\mathbb{E}[A]$ respectively. For any $\xi > 0$ there exists $C = C(\xi) > 0$ such that the following holds. If the expected node degree $\lambda_n > C \log(n)$ and $n^{-\xi} \le \delta \le 1/2$, then there exists $M > 0$ such that with probability at least $1 - \delta$,*

$$\|A - \mathbb{E}[A]\| \le M \sqrt{\lambda_n \log(n/\delta)}, \tag{B.1}$$

$$\|U_A - U_{\mathbb{E}[A]}\| \le M \sqrt{\frac{\log(n/\delta)}{\lambda_n}}. \tag{B.2}$$

In particular, $\|\mathbb{E}[A]\| \cdot \|U_A - U_{\mathbb{E}[A]}\| \le M \sqrt{\lambda_n \log(n/\delta)}$.

PROOF OF LEMMA 6. Inequality (B.1) follows directly from Theorem 3.1 of Oliveira (2010) and the fact that the maximum of the expected node degrees is of order λ_n. Inequality (B.2) is a consequence of (B.1) and Davis-Kahan Theorem (see Theorem VII.3.2 of Bhatia (1996)) as follows. By Lemma 5, the nonzero eigenvalues ρ_1^θ and ρ_2^θ of $\mathbb{E}[A]$ are of order λ_n. Let

$$\mathcal{S} = \left[\rho_2^\theta - M\sqrt{\lambda_n \log(n/\delta)}, \rho_1^\theta + M\sqrt{\lambda_n \log(n/\delta)}\right].$$

Then $\rho_1^\theta, \rho_2^\theta \in \mathcal{S}$ and the gap between $\mathcal{S}$ and zero is of order λ_n. Let $\bar{P}$ be the projector onto the subspace spanned by two leading eigenvectors of $\mathbb{E}[A]$. Since λ_n grows faster than $\|A - \mathbb{E}[A]\|$ by B.1, only two leading eigenvalues of A belong to $\mathcal{S}$. Let P be the projector onto the subspace spanned by two leading eigenvectors of A. By Davis-Kahan Theorem,

$$\|U_A - U_{\mathbb{E}[A]}\| = \|\bar{P} - P\| \leq \frac{M\|A - \mathcal{E}[A]\|}{\lambda_n} \leq M\sqrt{\frac{\log(n/\delta)}{\lambda_n}},$$

which completes the proof. □

Before proving Theorem 6 we need the following Lemma.

LEMMA 7. *Let x, y, $\bar{x}$, and $\bar{y}$ be unit vectors in $\mathbb{R}^n$ such that $\langle x, y\rangle = \langle \bar{x}, \bar{y}\rangle = 0$. Let P and $\bar{P}$ be the orthogonal projections on the subspaces spanned by $\{x, y\}$ and $\{\bar{x}, \bar{y}\}$ respectively. If $\|P - \bar{P}\| \leq \epsilon$ then there exists an orthogonal matrix $\mathcal{K}$ of size 2×2 such that $||(x, y)\mathcal{K} - (\bar{x}, \bar{y})||_F \leq 9\epsilon$.*

PROOF OF LEMMA 7. Let $x_0 = P\bar{x}$ and $y_0 = P\bar{y}$. Since $\|P - \bar{P}\| \leq \epsilon$, it follows that $\|\bar{x} - x_0\| \leq \epsilon$ and $\|\bar{y} - y_0\| \leq \epsilon$. Let $x^\perp = \frac{x_0}{\|x_0\|}$, then

$$\|\bar{x} - x^\perp\| \quad \leq \quad \|\bar{x} - x_0\| + \|x_0 - x^\perp\| \leq \epsilon + |1 - \|x_0\|| \leq 2\epsilon.$$

Also $\langle x^\perp, y_0\rangle = \langle x^\perp, y_0 - \bar{y}\rangle + \langle x^\perp - \bar{x}, \bar{y}\rangle$ implies that $|\langle x^\perp, y_0\rangle| \leq 3\epsilon$. Define $z = y_0 - \langle y_0, x^\perp\rangle x^\perp$. Then $\langle z, x^\perp\rangle = 0$, $\|\bar{y} - z\| \leq \|\bar{y} - y_0\| + \|y_0 - z\| \leq 4\epsilon$, and $|1 - \|z\|| = |\|\bar{y}\| - \|z\|| \leq 4\epsilon$. Let $y^\perp = \frac{1}{\|z\|}z$, then

$$\|\bar{y} - y^\perp\| \quad \leq \quad \|\bar{y} - z\| + \|z - y^\perp\| \leq 4\epsilon + |1 - \|z\|| \leq 8\epsilon.$$

Therefore $\|(\bar{x}, \bar{y}) - (x^\perp, y^\perp)\|_F \leq 9\epsilon$. Finally, let $\mathcal{K} = (x, y)^T(x^\perp, y^\perp)$. □

PROOF OF THEOREM 6. Let ε be the upper bound of (B.2), and denote $U = (u_1, u_2)^T = U_A$, and $\bar{U} = (\bar{u}_1, \bar{u}_2)^T = U_{\mathbb{E}[A]}$. We first show that there exists a constant $M > 0$ such that with probability at least $1 - \delta$,

$$\min \left\|(u_1^T \mathbf{1} u_2 - u_2^T \mathbf{1} u_1) \pm (\bar{u}_1^T \mathbf{1} \bar{u}_2 - \bar{u}_2^T \mathbf{1} \bar{u}_1)\right\| \leq M\varepsilon\sqrt{n}. \tag{B.3}$$

Let $\mathcal{R} = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ be the $\pi/2$-rotation on $\mathbb{R}^2$. Then

$$u_1^T \mathbf{1} u_2 - u_2^T \mathbf{1} u_1 = U^T \mathcal{R} U \mathbf{1}, \quad \bar{u}_1^T \mathbf{1} \bar{u}_2 - \bar{u}_2^T \mathbf{1} \bar{u}_1 = \bar{U}^T \mathcal{R} \bar{U} \mathbf{1}.$$

By Lemma 6 and Lemma 7, there exists an orthogonal matrix $\mathcal{K}$ such that if $E = (E_1, E_2) = U^T - \bar{U}^T \mathcal{K}$ then $||E||_F \leq 9\varepsilon$. By replacing U^T with $E + \bar{U}^T \mathcal{K}$, the left hand side of (B.3) becomes

$$\min \left\| \left(E + \bar{U}^T \mathcal{K}\right) \mathcal{R} \left(E + \bar{U}^T \mathcal{K}\right)^T \mathbf{1} \pm \bar{U}^T \mathcal{R} \bar{U} \mathbf{1} \right\|.$$

Note that $\mathcal{K}^T \mathcal{R} \mathcal{K} = \mathcal{R}$ if $\mathcal{K}$ is a rotation, and $\mathcal{K}^T \mathcal{R} \mathcal{K} = -\mathcal{R}$ if $\mathcal{K}$ is a reflection. Therefore, it is enough to show that

$$\left\| \bar{U}^T \mathcal{K} \mathcal{R} E^T \mathbf{1} + E \mathcal{R} \mathcal{K}^T \bar{U} \mathbf{1} + E \mathcal{R} E^T \mathbf{1} \right\| \leq M \epsilon \sqrt{n}.$$

Note that $|E_i^T \mathbf{1}| \leq \sqrt{n} \|E_i\| \leq 9\varepsilon\sqrt{n}$ and $\|E\|_F \leq 9\varepsilon < 1$ if λ_n is large enough, so

$$\|E \mathcal{R} E^T \mathbf{1}\| = \|E_2^T \mathbf{1} E_1 - E_1^T \mathbf{1} E_2\| \leq 18\varepsilon\sqrt{n}.$$

Since $\bar{U}\mathbf{1} = \sqrt{n}(s_1, s_2)^T$ for some constant s_1 and s_2, it follows that

$$\|E \mathcal{R} \mathcal{K}^T \bar{U} \mathbf{1}\| = \sqrt{n} \|(E_2 - E_1) \mathcal{K}^T (s_1, s_2)^T\| \leq M\varepsilon\sqrt{n}$$

for some $M > 0$. Analogously,

$$\|\bar{U}^T \mathcal{K} \mathcal{R} E^T \mathbf{1}\| = \|\bar{U}^T \mathcal{K} (-E_2^T \mathbf{1}, E_1^T \mathbf{1})^T\| \leq M\varepsilon\sqrt{n},$$

and (B.3) follows. By Lemma 3 and Lemma 5, with probability at least $1 - \delta$, $\bar{U}^T \mathcal{R} \bar{U} \mathbf{1}$ is of order

$$(\pi_2\theta_1, \pi_2\theta_2, ..., \pi_2\theta_{\bar{n}_1}, -\pi_1\theta_{\bar{n}_1+1}, ..., -\pi_1\theta_n)^T.$$

For simplicity, assume that in (B.3) the minimum is when the sign is negative (because $\hat{c}$ is unique up to a factor of -1). If node i is mis-clustered by $\hat{c}$ then

$$|(U^T \mathcal{R} U \mathbf{1})_i - (\bar{U}^T \mathcal{R} \bar{U} \mathbf{1})_i| \geq \min_i |(\bar{U}^T \mathcal{R} \bar{U} \mathbf{1})_i| =: \eta.$$

Let k be the number of mis-clustered nodes, then by (B.3), $\eta\sqrt{k} \leq M\varepsilon\sqrt{n}$. Therefore the fraction of mis-clustered nodes, k/n, is of order ε^2. □

439 West Hall, 1085 S. University Ave.
Ann Arbor, MI 48109-1107
E-mail: canle@umich.edu

439 West Hall, 1085 S. University Ave.
Ann Arbor, MI 48109-1107
E-mail: elevina@umich.edu

2074 East Hall, 530 Church St.
Ann Arbor, MI 48109-1043
E-mail: romanv@umich.edu