Searching for a Single Community in a Graph

AVIK RAY, SUJAY SANGHAVI, and SANJAY SHAKKOTTAI, University of Texas at Austin

In standard graph clustering/community detection, one is interested in partitioning the graph into more densely connected subsets of nodes. In contrast, the *search* problem of this article aims to only find the nodes in a *single* such community, the target, out of the many communities that may exist. To do so, we are given suitable side information about the target; for example, a very small number of nodes from the target are labeled as such.

We consider a general yet simple notion of side information: all nodes are assumed to have random weights, with nodes in the target having higher weights on average. Given these weights and the graph, we develop a variant of the method of moments that identifies nodes in the target more reliably, and with lower computation, than generic community detection methods that do not use side information and partition the entire graph. Our empirical results show significant gains in runtime, and also gains in accuracy over other graph clustering algorithms.

CCS Concepts: • **Theory of computation** → **Semi-supervised learning**; *Unsupervised learning and clustering*;

Additional Key Words and Phrases: Community detection, graph clustering, search, side information, semi-supervised

ACM Reference format:

Avik Ray, Sujay Sanghavi, and Sanjay Shakkottai. 2018. Searching for a Single Community in a Graph. *ACM Trans. Model. Perform. Eval. Comput. Syst.* 3, 3, Article 13 (June 2018), 17 pages.
<https://doi.org/10.1145/3200863>

1 INTRODUCTION

Community detection, or graph clustering, is the classic problem of finding subsets of nodes such that each subset has higher connectivity within itself, as compared to the average connectivity of the graph as a whole. Typically, when graphs represent similarity or affinity relationships between nodes, these subsets represent communities of similar nodes. Also typically, this problem has primarily been considered in the unsupervised setting, where the only input is the graph itself and the objective is to partition all or most of the nodes.

In this article, we look at a different, but related, community detection task, which we will refer to as the *search problem*. Our objective is to use the graph to find a single community of

An earlier version of this work has been published as a short poster paper in the Proceedings of ACM Sigmetrics 2016 (Ray et al. 2016).

We would like to acknowledge support from NSF grants CNS-1320175, 0954059, ARO grants W911NF-15-1-0227, W911NF-14-1-0387, W911NF-16-1-0377, and the US DoT supported D- STOP Tier 1 University Transportation Center.

Authors' addresses: A. Ray, S. Sanghavi, and S. Shakkottai, Department of Electrical and Computer Engineering, The University of Texas at Austin, 2501 Speedway, C0803, EER 6.804, Austin, TX 78712; emails: avik@utexas.edu, sanghavi@mail.utexas.edu, shakkott@austin.utexas.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 ACM 2376-3639/2018/06-ART13 \$15.00

<https://doi.org/10.1145/3200863>

nodes – which we will call the *target community* – for which we have been given some relevant but quite noisy side information. We would like to do so more reliably, and with lower computation, than existing methods that do not use side information.

Our motivations are two-fold: (i) it is often the case that the network analyst is looking for nodes with *a priori* specified characteristics, and (ii) it is rare that we are faced with a “pure” graph analysis problem; typically, there is extra non-graphical side information that, if used properly, could make the inference task easier.

As an example setting, consider the case where we have some nodes from the target community explicitly marked as such, and our task is to recover the remaining nodes. This is a situation that frequently arises in military/intelligence settings, and also in analysis of regular consumer social networks, internet/web graphs, and so on. In military intelligence, it can be useful to recover a single community, which a known suspect is part of. Besides explicit node labels, side information could also come from meta-information one may have about the nodes, e.g., from text analysis if the graph is a web graph, or from browse/activity history of users in a social network. In recommendation system or targeted advertising, it is useful to learn a community of users with a specific interest (e.g., sports) using the knowledge of how users interact with relevant contents (e.g., sports news and images). Our aim is to find a principled way to use such side information *and* the graph itself.

Our contributions are as follows.

- (i) We develop a simple yet generic framework for how side information is to be specified: each node is given a (possibly random) weight, with nodes in the target community having higher weight on average than nodes not in the target—we call these *biased weights*. This setting would thus split an overall data + graph analysis objective into two: the analyst needs to devise a (application-dependent) procedure to convert her side information into biased node weights; these are then used by our algorithm.
- (ii) Given such biased weights, we develop a new spectral-like algorithm—specifically, a variant of the second-order method of moments—to find the nodes in the target community. We call this *Community Search* below. In the following, we first provide the basic intuition behind it by considering the case where we have access to the population statistics of a graph coming from a stochastic block model, and then formally describe the algorithm.
- (iii) Our main results characterize the effectiveness of this algorithm in finding the target community; we study this in the standard stochastic block model setting with many communities. Analytically, we show that it matches (potentially up to log factors) the analytical guarantees of the state-of-the-art unsupervised community detection methods; empirically, we show that the method outperforms these methods even with very noisy side information (e.g., very small number of labeled nodes), and has significantly lower computational complexity.
- (iv) We also specialize our results to the case where the side information is in the form of a small number of labeled nodes; for this case, we show how one can effectively convert this to node weights, even for sparse graphs. Our experiments on a real-world network further corroborate the practical applicability of this method.

1.1 Related Work

While no other work has considered the problem of searching a single community in a graph, there has been a lot of research in three closely related fields; that of unsupervised and semi-supervised graph clustering, method of moments, and learning with side information. Each of these threads have a rich history—here, we cover the ones most relevant to this article.

Unsupervised graph clustering: Graph clustering or community detection has been widely studied mainly in the unsupervised setting where nodes do not have any associated labels. There is a vast literature of graph clustering algorithms both in the setting where clusters are non-overlapping (Fortunato 2010) and overlapping (Xie et al. 2013). The most widely studied generative model for non-overlapping clusters in a graph is the planted partition or stochastic block model (Condon and Karp 2001). Assuming this model, many algorithms have been proposed, which provide statistical guarantees of recovery of all hidden clusters. These algorithms can be broadly divided into three categories—(i) spectral clustering (McSherry 2001; Ng et al. 2002; Rohe et al. 2011; Chaudhuri et al. 2012; Amini et al. 2013; Yun and Proutiere 2014), (ii) convex optimization (Chen et al. 2012; Ailon et al. 2013; Abbe et al. 2016), and more recently, (iii) tensor decomposition (Anandkumar et al. 2013; Huang et al. 2013).

Semi-supervised graph clustering: The graph clustering problem has also been explored in a semi-supervised setting, where some of the nodes and/or edges are explicitly labeled. Many optimization and kernel-based algorithms have been proposed (Zhu 2005; Kulis et al. 2009) to solve this problem. The popular label propagation-based clustering algorithms (Zhou et al. 2004; Fujiwara and Irie 2014) are also essentially semi-supervised graph clustering algorithms with labeled nodes. Another related line of work also studies the graph clustering problem where the nodes have additional node features/attributes (McAuley and Leskovec 2012; Xu et al. 2012; Yang et al. 2013; Zhang et al. 2016).

Method of Moments: This is a classical parameter estimation technique, where the parameters to be estimated are described in terms of the moments from the true distribution. Empirical moments are now used to replace the true moments, leading to parameter estimates (Bowman and Shenton 2004). There has been much recent interest in these methods for many statistical learning problems. These include learning Gaussian mixture models (Hsu and Kakade 2013; Anandkumar et al. 2014), latent Dirichlet allocation (LDA) topic models (Anandkumar et al. 2012), hidden Markov models (Chang 1996), and the like.

Others: There is a broader machine learning literature that incorporates the availability of extra side information into existing models and algorithms. In the context of LDA topic models, side information maybe available in the form of extra response variables for each document (Mcauliffe and Blei 2008), or additional text review information of products (Lu and Zhai 2008). In collaborative filtering, side information can be of the form of item or user graph (Rao et al. 2015).

In this article, we consider the community search problem with side information either in the form of biased node weights or a small set of labeled nodes.

2 SETTINGS AND ALGORITHM

Stochastic Block Model: Consider a graph $G = (V, E)$ with n nodes and k non-overlapping communities that partition the vertex set as $V = \cup_{i=1}^k V_i$. Let $\alpha_i = |V_i|/n$ be the fraction of nodes in the i -th community. In a stochastic block model, the edge set E is generated as follows. Let $0 < q < p < 1$. Then, for any two nodes in the same community $r, s \in V_i$, we have $P((r, s) \in E) = p$, and when r, s are in different communities then $P((r, s) \in E) = q$. We define this as the (n, k, p, q) stochastic block model.

Target community and side information: In the search problem, we are interested in the recovery of *one* target community. In this article, without loss of generality, consider V_1 to be this target community. We are also provided with some side information on this target community V_1 . The side information is in the form of *biased node weights*. Suppose for each node $j \in V$ we are given a biased weight $w_j > 0$. These weights are generated by a random process satisfying the condition that for any node $j \in V$, we have $E[w_j | j \in V_1] > E[w_j | j \in V_i]$ for all $i \neq 1$.

These biased weights may be computed using a set of *labeled nodes* from the target community $\mathcal{L} \subset V_1$ (see Section 3.2). These weights can also arise from other available sources of side information. For example, consider a social network graph where the target community consists of users who are sports enthusiasts. Then, we can observe the amount of interaction (e.g., “likes” and “shares” in Facebook) of the users with known sports related contents. Since users in a sports community are more likely to interact with such contents, these will have the above biased node weight property. The main goal is to solve this search problem faster than the time required to recover all k communities, and without any loss in estimation accuracy.

2.1 Algorithm

In this section, we describe our main algorithm called **Community Search**. Let X be the adjacency matrix of the graph G . Also, define community membership vectors $\mu_1, \dots, \mu_k$, where $\mu_i \in \mathbb{R}^n$, as follows. Let $\mu_{j,i}$ be the j -th coordinate of vector μ_i . Then,

$$\mu_{j,i} = \begin{cases} p & \text{if } j \in V_i \\ q & \text{otherwise} \end{cases}$$

Note that these μ_i -s are linearly independent and the community memberships of the nodes can be obtained from these membership vectors via thresholding. The main purpose of our algorithm is to estimate the membership vector of the first community μ_1 (which can then be used to recover nodes in V_1).

Intuition behind our method: To understand the core of our technique, let us suppose here—just for intuition—that we actually had access to the “average” adjacency matrix $E[X]$ (recall that X is the actual adjacency matrix of the stochastic block model), and let $E[X_j]$ be the average of the j^{th} column. Then it is easy to see that $E[X_j] = \mu_{c_j}$, where $c_j \in 1, \dots, k$ is the community that node j belongs to. This means that the following holds for the matrix A defined below:

$$A := \frac{1}{n} \sum_{j=1}^n E[X_j]E[X_j]^T = \sum_{i=1}^k \alpha_i \mu_i \mu_i^T$$

Similarly, let us now also suppose that we see the “average” node weights $\tilde{w}_j = E[w_j]$ for every node j . Then, the following holds for the matrix B defined below:

$$B := \frac{1}{n} \sum_{j=1}^n \tilde{w}_j E[X_j]E[X_j]^T = \sum_{i=1}^k \alpha_i \omega_i \mu_i \mu_i^T,$$

where in the above, for each cluster i , we have defined ω_i be the averaged weights of all nodes in that cluster. By the bias condition, we have that $\omega_1 > \omega_i$ for all $i \neq 1$.

Note that both the A and B as defined above are symmetric positive definite rank- k matrices, with the column space of each spanned by the μ_i ’s. However, note also that our desired vector μ_1 may not be an eigenvector of either A or B ; indeed, if the target community V_1 is small, it may be quite far from the leading eigenvector of either matrix.

The main idea is that we can still recover μ_1 by “whitening” B using A , a process we describe in the proto-algorithm below. The description also provides the (simple) reason why it works—in this idealized case where average X and w are available.

Proto-algorithm (and explanation):

- (1) Compute matrices A and B as described above,
- (2) Perform rank- k svd of A as $A = UDU^T$, and let $W := UD^{-1/2}$. Also note that,

$$W^T A W = I_k = \sum_{i=1}^k \tilde{\mu}_i \tilde{\mu}_i^T,$$

where we define $\tilde{\mu}_i := \sqrt{\alpha_i} W^T \mu_i$. Now, we see that the addition of k terms of the type $\tilde{\mu}_i \tilde{\mu}_i^T$ results in I_k ; this can *only* happen if the corresponding $\tilde{\mu}_i$ are *orthonormal* vectors in $\mathbb{R}^k$. The vectors $\tilde{\mu}$ are thus “whitened” versions of the original μ vectors.

- (3) Next, we compute the following matrix.

$$R := W^T B W = \sum_{i=1}^k \omega_i \tilde{\mu}_i \tilde{\mu}_i^T$$

Now, since $\tilde{\mu}_i$ are orthonormal, the above equation represents an eigenvalue decomposition of the $k \times k$ size matrix R , with eigenvectors $\tilde{\mu}_i$ and corresponding eigenvalues ω_i . Thus, $\tilde{\mu}_1$ —the whitened vector corresponding to the target community—is now the *leading eigenvector* of R , because $\omega_1 > \omega_i$.

- (4) Find $\tilde{\mu}_1$ by setting it to be the leading eigenvector of R . Finally, we can recover μ_1 from $\tilde{\mu}_1$ in two steps. First, compute $z := U D^{1/2} \tilde{\mu}_1 = \sqrt{\alpha_1} \mu_1$. Next, compute vector

$$m_1 := \frac{1}{n} \sum_{j=1}^n E[X_j] = \sum_{i=1}^k \alpha_i \mu_i$$

We can recover $\sqrt{\alpha_1} = \tilde{\mu}_1^T W^T m_1$. Then, simply divide the z defined above by this to find μ_1 .

An issue: Although simple, it is not straightforward to convert this intuition to an algorithm because due to interdependencies, it becomes hard to estimate these A and B matrices. In particular, note that in the actual problem we are given the adjacency matrix X , and a natural impulse is to approximate A using the matrix

$$\frac{1}{n} \sum_{j=1}^n X_j X_j^T$$

Unfortunately, this is a good approximation to $E[XX^T]$, but $E[XX^T] \neq E[X]E[X]^T$ —and we require the latter. However, we can get around these dependencies by first partitioning the graph. This is outlined in Algorithm 1 below.

For any two subsets $P, Q \subset [n]$, let $X_{P,Q}$ denote the submatrix of X corresponding to the rows and columns in set P and Q , respectively. The input parameters to Algorithm 1 are the adjacency matrix X , number of communities k , the set of biased node weights $(w_1, \dots, w_n)$, and a threshold τ . The output is the community estimate $\hat{V}_1$.

3 MAIN RESULT

In this section, we present our main theoretical results. First we show that when the set of biased weights $(w_1, \dots, w_n)$ satisfy certain mild sufficient conditions, then Algorithm 1 is guaranteed to recover the target community V_1 . Later, we show how such weights can be obtained even with a set of labeled nodes from the target community.

3.1 Recovery Using Biased Weights

When side information is available in the form of biased weights w_j for each node $j \in V$, these weights need to be *informative* about the target community V_1 so that it could be recovered. Clearly, *good* side information will lead to a better performance of any search algorithm. We quantify this

ALGORITHM 1: Community Search**Input:** Adjacency matrix X , k , biased weights $(w_1, \dots, w_n)$, threshold τ **Output:** $\hat{V}_1$

- 1 Partition nodes into four sets P_1, P_2, P_3, P_4 at random;
- 2 Compute matrices $\hat{A}_1 = \frac{1}{\sqrt{|P_3|}} X_{P_1, P_3}$, $\hat{A}_2 = \frac{1}{\sqrt{|P_3|}} X_{P_2, P_3}$;
- 3 Compute vector $\hat{m}_1 = \frac{1}{|P_1|} \sum_{j \in P_1} X_{P_1, j}$;
- 4 Compute matrix $\hat{B} = \frac{1}{|P_4|} \sum_{j \in P_4} w_j X_{P_1, j} X_{P_2, j}^T$;
- 5 $\hat{\mu}_{P_1}, \hat{\alpha}_1 \leftarrow \text{SearchSubroutine}(\hat{A}_1, \hat{A}_2, \hat{B}, \hat{m}_1, k)$;
- 6 Compute $V_{P_1} = \{j \in P_1 : \hat{\mu}_{P_1, j} > \tau\}$;
- 7 Repeat steps 2–5 with P_i 's rotated in order to estimate $\hat{\mu}_{P_2}, \hat{\mu}_{P_3}, \hat{\mu}_{P_4}$. Use them to compute $V_{P_2}, V_{P_3}, V_{P_4}$ as in step 6;
- 8 Return community $\hat{V}_1 = V_{P_1} \cup V_{P_2} \cup V_{P_3} \cup V_{P_4}$;

ALGORITHM 2: SearchSubroutine**Input:** $\hat{A}_1, \hat{A}_2, \hat{B}, \hat{m}_1, k$ **Output:** $\hat{\mu}_1, \hat{\alpha}_1$

- 1 Compute rank k -svd of matrices $\hat{A}_1, \hat{A}_2 : \hat{A}_1 = U_1 D_1 V_1^T, \hat{A}_2 = U_2 D_2 V_2^T$;
- 2 Compute matrices $W_1 = U_1 D_1^{-1}, W_2 = U_2 D_2^{-1}$;
- 3 Let u_1 be the largest left singular vector of $W_1^T \hat{B} W_2$;
- 4 Compute $z = U_1 D_1 u_1$;
- 5 Compute $a = u_1^T W_1^T \hat{m}_1$;
- 6 Return $\hat{\mu}_1 \leftarrow z/a$ and $\hat{\alpha}_1 \leftarrow a^2$;

quality of information in the following set of assumptions (Condition (A1) and (A2)) on the biased weights. The third Condition (A3) is a more fundamental condition that determines when the community structure itself is identifiable in a stochastic block model.

- (A1) *Average weight bias*: Under this condition, the expected weight of a node in community V_1 is greater than the expected weight of a node in any other community V_i . Precisely, the weights satisfy:

$$E[w_j | j \in V_1] > E[w_j | j \in V_i], \forall i \neq 1$$

This weight bias allows us to determine that community V_1 is being searched/preferred over the remaining communities. However, we only require this to hold in expectation and the actual weights themselves may vary significantly. Clearly, any algorithm that only uses the weight bias to determine community membership by simple thresholding will perform very poorly.

- (A2) *Weight concentration*: Let $\alpha_{\max} = \max_{i \in [k]} \alpha_i$ and $\alpha_{\min} = \min_{i \in [k]} \alpha_i$. Define $\sigma_1(R) := E[w_j | j \in V_1]$, $\sigma_2(R) := \max_{i \neq 1} E[w_j | j \in V_i]$, $\gamma_2 := \max_{i \in [k], j \in V} |w_j - E[w_j | j \in V_i]|$, and $\xi(n) = o(\sqrt{\log n})$ be any slowly growing function. Then, with high probability, the maximum deviation of the weights are bounded as,

$$\frac{\gamma_2}{(\sigma_1(R) - \sigma_2(R))} = O\left(\min\left\{\frac{\alpha_{\min}^4 (p - q)^4}{\alpha_{\max}^4 p^4 \xi(n)}, \frac{\alpha_{\min}^5 \sqrt{n} (p - q)^5}{\alpha_{\max}^4 p^{4.5} \xi(n)} - 1\right\}\right)$$

This condition dictates that the maximum variation of the weights γ_2 is also small compared to the difference between the largest and second largest expected weights $\sigma_1(R) - \sigma_2(R)$.

Since the weights are used primarily to construct the matrix B in Algorithm 1, this condition ensures that the matrix B can be estimated up to a tolerable error.

– (A3) p, q separation: Let p, q, n satisfy

$$\frac{(p - q)^2}{p\sqrt{p}} = \tilde{\Omega}\left(\frac{\alpha_{\max}}{\alpha_{\min}^2 \sqrt{n}}\right)$$

This condition fundamentally determines when communities are identifiable in a stochastic block model and similar conditions are required for other community detection algorithms (Anandkumar et al. 2013; Chaudhuri et al. 2012; Chen et al. 2012). The more the gap $p - q$, the easier it is to identify communities. Hence, this condition gives a lower bound on $p - q$, which is required for community identifiability.

Theorem 3.1 shows that under the above assumptions on the biased weights, Algorithm 1 can reconstruct community V_1 with high accuracy.

THEOREM 3.1. *Consider a (n, k, p, q) stochastic block model satisfying condition (A3). Given biased weights $(w_1, \dots, w_n)$ satisfying conditions (A1), (A2), then Algorithm 1 recovers community V_1 with a fraction of error nodes $o(1)$ with high probability.*

Remark 3.2. For a stochastic block model with equal community sizes n/k , Condition (A3) reduces to $\frac{(p-q)^2}{p\sqrt{p}} = \tilde{\Omega}(\frac{k}{\sqrt{n}})$. When $p = \Theta(p - q)$, this has the same scaling as other community detection algorithms (Chen et al. 2012; Anandkumar et al. 2013; Chaudhuri et al. 2012). Therefore, even in sparse graphs where $p, q = \Theta(\frac{\log n}{n})$ or for small community sizes up to $\Omega(\sqrt{n})$ nodes Algorithm 1 can recover the community.

In Theorem 3.1, the $o(1)$ fraction error can be easily converted to a zero error guarantee using an additional post-processing step. Instead of estimating community 1 nodes inside partition P_1 , we can estimate those in partition P_2 , first by observing for each node $j \in P_2$ the number of edges shared with the estimated set V_{P_1} , followed by thresholding. Since V_{P_1} estimates $V_1 \cap P_1$ up to only $o(\alpha_1 n)$ error nodes, this does not cause any errors in thresholding, with high probability. This post-processing step is also independent of the previous steps in the algorithm since the edges between partitions P_1 and P_2 are not utilized in Algorithm 1. The following theorem formalizes this idea.

THEOREM 3.3 (EXACT RECOVERY). *In a (n, k, p, q) stochastic block model, under assumptions (A1)–(A3), Algorithm 1 with an additional degree thresholding step can recover community V_1 completely with high probability.*

We prove Theorems 3.1 and 3.3 in Appendix A.

3.2 Recovery Using Labeled Nodes

Biased weights, as required in Theorem 3.1, can be obtained from a small set of labeled nodes $\mathcal{L}$ as follows:

- Choose a radius r .
- Weight w_i is the number of edges between nodes in $\mathcal{L}$ and nodes at a distance of r hops from node i .

Note that the weight can also be viewed as the number of neighbors of the set $\mathcal{L}$, which are at a distance r from node i . Larger choice of radius r means less variance in the weights, but also potentially less bias if it becomes too large. For example, $r = 1$ means only neighbors of labeled nodes get weights; this is very high bias but also high variance.

The theorem below provides the correct way to choose the radius r such that the weights w_i can be made to satisfy Conditions (A1) and (A2). This means that even with such weights computed via labeled nodes, we can efficiently find community V_1 using Algorithm 1. Note that when $p \geq \frac{1}{\sqrt{n}}$, then, with high probability, the labeled nodes in $\mathcal{L}$ have neighbors with any other node $i \in V_1 \setminus \mathcal{L}$; hence, the number of common neighbors between i and nodes $l \in \mathcal{L}$ can be taken as weights w_i , which will satisfy Conditions (A1) and (A2). However, this does not work for sparse graphs when $p < \frac{1}{\sqrt{n}}$. In the following theorem, we show that even for $p = \Theta(\frac{\log n}{n^\epsilon})$, $\frac{1}{2} \leq \epsilon \leq 1$, the weights chosen by the above procedure and a correct r will work.

THEOREM 3.4. *Consider a (n, k, p, q) stochastic block model satisfying condition (A3) where $p = \Theta(\frac{\log n}{n^\epsilon})$, $q = \Theta(\frac{\log n}{n^\epsilon})$, $p - q = \Theta(\frac{\log n}{n^\epsilon})$, and all equal-sized communities, where $\frac{1}{2} \leq \epsilon \leq 1$. Given $L = \tilde{\Omega}(n^{\epsilon/2})$ labeled nodes, the biased weights computed with $r = \frac{2 \log(n^\epsilon/L)}{\log np}$, satisfy Conditions (A1) and (A2) with high probability.*

We prove this in Appendix A.3. For simplicity in Theorem 3.4, we assume equal community sizes; however, this can be extended to unequal but comparable community sizes.

3.3 Parallel Semi-supervised Graph Clustering

Our algorithm naturally provides a method for the standard semi-supervised graph clustering problem. This is the setting where we are given a small number of labeled nodes from every community, and we are interested in recovering all communities. In such a scenario, we can apply the community search algorithm to search for each individual community using the labeled nodes in that target community. Moreover, this search can be performed in parallel. Therefore, Algorithm 1 can also be used as a parallel graph clustering algorithm. Note that the vector m_1 and matrices A_1, A_2 remain the same for individual searches; only matrix B should be computed separately for every target community. Section 4 shows some numerical results evaluating the performance of Algorithm 1 in this semi-supervised graph clustering setting.

3.4 Comparison

In this section, we compare the theoretical performance of our algorithm with other unsupervised graph clustering algorithms.

For graphs with equal communities of size n/k , convex optimization based algorithms by Chen et al. (2012), Ailon et al. (2013), and Agarwal et al. (2015) can achieve the performance bound $\frac{(p-q)}{\sqrt{p}} = \tilde{\Omega}(\frac{k}{\sqrt{n}})$. In comparison, our algorithm achieves a slightly higher bound $\frac{(p-q)^2}{p\sqrt{p}} = \tilde{\Omega}(\frac{k}{\sqrt{n}})$. However, when $p = \Theta(p - q)^1$, both bounds are equivalent (up to log factors), implying our algorithm can recover communities even in sparse graphs with $p, q = \Theta(\frac{\log n}{n})$ and for growing numbers of communities $k = O(\sqrt{n})$. In terms of runtime, our algorithm runs in $O(n^2k)$ time faster than $\Omega(n^3)$ time required by convex optimization-based algorithms.

The Community Search by Whitening algorithm is also faster than tensor decomposition-based graph clustering algorithm by Anandkumar et al. (2013). Note that the first step of this tensor algorithm is to compute a whitening matrix using rank- k svd, which is identical to the search algorithm. In the remaining steps, for the tensor algorithm, the bulk of the computation is a rank- k tensor decomposition requiring $O(k^5)$ computation, which is slower than rank-1 svd computed in $O(k^2)$ time by the search algorithm. This is corroborated by our experiments in Section 4.

¹This is the case in most real sparse networks when $p = \Theta(\log n/n)$; if not, then it becomes impossible for any algorithm to recover communities in this regime as shown by Abbe et al. (2016).

Recently, a quasi-linear time graph clustering algorithm was presented by Abbe and Sandon (2015) for the case when the number of communities $k = O(1)$. In comparison, our algorithm can be applied even when the number of communities scale as $k = O(\sqrt{n})$, and it requires much lesser knowledge of model parameters than the former.

4 EXPERIMENTS

In this section, we present our numerical results showing the performance of the Community Search algorithm on synthetic and real datasets. We compare our algorithm with the Spectral clustering algorithm by Ng et al. (2002) and the Tensor decomposition-based clustering algorithm by Anandkumar et al. (2013). We generate synthetic datasets according to the stochastic block model (see Section 2) with $n = 1,000$ nodes, $k \in \{5, 8\}$ communities, and different values of p and q . The real-world network we consider is the US political blogosphere network first introduced in Adamic and Glance (2005). The Spectral clustering algorithm (Ng et al. 2002) requires clustering of the rows corresponding to bottom k eigenvectors of the normalized Laplacian. Although k-means may be used for this, it tends to converge to local minima resulting in poor performance. To prevent this, we perform clustering of the rows via the hierarchical single-linkage clustering (SLINK) algorithm (Sibson 1973). We refer to this Spectral+SLINK algorithm simply as Spectral clustering in the remaining section. Our algorithm implementations are all in Matlab. We consider two types of side information: *labeled nodes* and *synthetic weights*.

Labeled Nodes: As discussed earlier, this is a natural means of providing side information to the algorithm. A set of m labeled nodes are randomly chosen from the target community V_1 . The corresponding weights are then computed as described in Section 3.2 with $r \in \{1, 2\}$.

Synthetic Weights: We synthetically generate three sets of weights, each of which are (on average) larger over the target community. These weights are generated as follows. A pair of weights (w_1, w_2) are first chosen to be one of $\{(5, 8), (5, 10), (5, 12)\}$. For each node in community V_1 , we set the node's weight to be w_2 with probability 0.8, and w_1 otherwise. For all other nodes in $V \setminus V_1$, we swap the probabilities, i.e., we set a node's weight to be w_2 with probability 0.2 and w_1 otherwise. This process generates three possible values of the expected node weights in the target community, $\sigma_1(R) \in \{7.4, 9, 10.6\}$.

Performance Metrics: Note that our algorithm directly uses labeled nodes/biased node weights, and the graph to infer the target community. The baseline algorithms, however, first estimate all communities in the graph; then, it computes the average node weight in each community. Finally, it outputs as target the community that has the highest average node weight. The estimation error for the i -th community is given as $e_i = |\{j \in V : j \in V_i, j \notin \hat{V}_i \text{ or } j \in \hat{V}_i, j \notin V_i\}|$. We compute the error for searching each community and plot either the overall average, or average over a subset of clusters. Let $T_{\mathcal{A}_1}, T_{\mathcal{A}_2}$ be the runtimes of algorithms $\mathcal{A}_1, \mathcal{A}_2$, respectively. Then, we define speedup s of algorithm $\mathcal{A}_1$ over algorithm $\mathcal{A}_2$ as $s = T_{\mathcal{A}_2}/T_{\mathcal{A}_1}$. $s > 1$ implies algorithm $\mathcal{A}_1$ is faster than $\mathcal{A}_2$.

4.1 Performance and Speedup with Labeled Nodes

First, we compare the error performance of the Community Search algorithm with Spectral clustering and Tensor decomposition algorithms in the setting where side information is given in the form of m labeled nodes from the target community. We then compute biased weights w_j using the tree method of Section 3.2 with a radius $r = 2$. Note that this tree method may assign weights in violation of condition (A1) for small target communities, since for small target clusters, the number of nodes in the tree from a large cluster may exceed those from the target community. In such cases, Algorithm 1 cannot be expected to recover the communities. Therefore, we consider a subset of larger communities that assign the correct weights satisfying Condition (A1) and evaluate

Fig. 1. Labeled Nodes: Comparing the average error performance of Community Search algorithm with Spectral clustering (Ng et al. 2002) and Tensor decomposition (Anandkumar et al. 2013) algorithms in a stochastic block model with (a) $n = 1,000, k = 5, \alpha_{min} = .1$, (b) $n = 1,000, k = 8, \alpha_{min} = .08$. The algorithms use m labeled nodes from target clusters as side information and compute biased weights. The Community Search algorithm outperforms both Spectral clustering and Tensor decomposition.

Fig. 2. Labeled Nodes: The average speedup performance of the Community Search and Spectral clustering (Ng et al. 2002) algorithms with respect to Tensor decomposition (Anandkumar et al. 2013) in a stochastic block model with (a) $n = 1,000, k = 8, \alpha_{min} = .08$; (b) $n = 1,000, k = 5, \alpha_{min} = .1$, and labeled nodes as side information. The Community Search algorithm is faster than both Spectral clustering and Tensor decomposition.

our algorithm over these communities. Figure 1(a) plots the average error over three largest cluster in a stochastic block model (SBM) with $n = 1,000$ and $k = 5$ unequal sized communities. The Community Search shows significantly less error than Tensor decomposition and Spectral clustering. In Figure 1(b) we plot the average over five larger clusters in an SBM with $n = 1,000, k = 8$ unequal communities. Again, Community Search shows better error than Spectral clustering and comparable error to Tensor decomposition.

Figure 2 shows the speedup performance of the Community Search and Spectral clustering algorithms over Tensor decomposition in this setting with labeled nodes. As indicated earlier, all three algorithms were implemented in Matlab. We observe that the Community Search has a much lower runtime than both Spectral clustering and Tensor decomposition.

Fig. 3. Synthetic Weights: Comparing the average error performance of Community Search algorithm with Spectral clustering (Ng et al. 2002) and Tensor decomposition (Anandkumar et al. 2013) algorithms in a stochastic block model with (a) $n = 1,000, k = 8, \alpha_{min} = .08$, (b) $n = 1,000, k = 5, \alpha_{min} = .1$. The algorithms use synthetic weights as side information to search for the target community. Community Search algorithm shows a lower error.

Fig. 4. Synthetic Weights: The average speedup performance of the Community Search and Spectral clustering (Ng et al. 2002) algorithms with respect to Tensor decomposition (Anandkumar et al. 2013) in a stochastic block model with (a) $n = 1,000, k = 8, \alpha_{min} = .08$, (b) $n = 1,000, k = 5, \alpha_{min} = .1$, and synthetic weights as side information. The Community Search algorithm is faster than both Spectral clustering and Tensor decomposition.

4.2 Performance and Speedup with Synthetic Weights

Next, we compare the error and runtime performance of all three algorithms in a setting where side information is available in the form of synthetically generated biased weights (as discussed earlier, three different choices of parameters). Figure 3(a) plots the average error over all communities in an SBM with $n = 1,000, k = 8$. The Community Search algorithm has a better performance over Spectral clustering and comparable performance with Tensor decomposition. Figure 3(b) plots the average error in an SBM with $n = 1,000, k = 5$. In this case, Community Search outperforms Tensor decomposition and has comparable performance to Spectral clustering.

In Figure 4, we plot the average speedup of Community Search and Spectral clustering over Tensor decomposition. Again, the Community Search algorithm is significantly faster than both

Fig. 5. Sensitivity to side information: The average percentage error of Community Search Algorithm in a stochastic block model with (a) $n = 1,000$, $k = 8$, $\alpha_{min} = .08$, (b) $n = 1,000$, $k = 5$, $\alpha_{min} = .1$, with increasing singular value gap $\sigma_1(R) - \sigma_2(R)$. As shown in our analysis, the performance improves with increase in the singular value gap.

Spectral clustering and Tensor decomposition. We also observe that the speedup increases with increasing $p - q$.

4.3 Sensitivity

In order to determine the sensitivity of our algorithm with respect to the quality of side information, and the number of communities k , we perform the following two experiments.

First, to see how the quality of side information effects the performance of our algorithm, we plot the average error with increasing singular value gap $\sigma_1(R) - \sigma_2(R)$ (or the difference between the largest and second largest expected node weight) in Figure 5. In this experiment we fix the synthetic weights ($w_1 = 5$, $w_2 = 10$) and vary $\sigma_1(R) - \sigma_2(R)$ by changing the probabilities with which the weights appear in each community. Note that the singular value gap increases when one weight appears with greater chance than the other. Therefore, $\sigma_1(R) - \sigma_2(R)$ can also be viewed as a measure of quality of side information. As predicted from our analysis, we observe that the error improves with an increase in the gap $\sigma_1(R) - \sigma_2(R)$.

Often, in real applications, one does not have perfect knowledge of the number of communities k in a graph, which is required to run any community search or detection algorithm. Thankfully, there are several methods to estimate the parameter k , e.g., from the spectral properties of the graph (Chen et al. 2012; Newman and Reinert 2016). Another approach is to compute a suitable community quality score metric like modularity (Newman 2006; Yang and Leskovec 2012) after running the algorithm with different values of k , and choosing the k that produces a community with the best score. However, such estimation may not always be accurate. Therefore, it is crucial that any community search algorithm perform robustly with respect to the input parameter k in the algorithm. In our next experiment, we compare the sensitivity of our Community Search algorithm with Spectral clustering and Tensor decomposition when provided with imperfect parameter k . In Figure 6, we plot the average percentage error of all three algorithms on two different SBM. We observe that, even with imperfect knowledge of k , the Community Search algorithm has lower error than Spectral clustering and Tensor decomposition. Interestingly, the Community Search algorithm shows much less sensitivity to higher k values than a lower k (with respect to the ground truth k).

Fig. 6. Sensitivity to k : The average percentage error of Community Search algorithm in a stochastic block model with (a) $n = 1,000, k = 8, \alpha_{\min} = .08$, (b) $n = 1,000, k = 5, \alpha_{\min} = .1, p = .09, q = .01$ in both models, synthetic weights as side information, and with imperfect knowledge of the number of communities k . Our Community Search algorithm exhibits lower sensitivity and error than Spectral clustering (Ng et al. 2002) and Tensor decomposition (Anandkumar et al. 2013) algorithms for both models.

Fig. 7. Labeled Nodes + Parallel Clustering: Comparing the average error performance of Community Search algorithm with Spectral clustering (Ng et al. 2002) and Tensor decomposition (Anandkumar et al. 2013) algorithms in a stochastic block model with $n = 1,000, k = 8, \alpha_{\min} = .08$; (a) $q = .01$; (b) $p = .14$. We consider the semi-supervised graph clustering setting when side information is available in the form of the $m \in \{2, 4, 6, 8\}$ labeled node from each community. The Community Search algorithm has a better performance over both Spectral clustering and Tensor decomposition.

4.4 Parallel Clustering

Finally, we consider the semi-supervised graph clustering setting described in Section 3.3 where we are provided with m labeled nodes from each community, and we want to recover all communities. Recall that the Community Search algorithm can also be used as a semi-supervised parallel graph clustering algorithm. Figure 7 plots the cumulative error over all communities with increasing $p - q$ in an SBM with $n = 1,000, k = 8$, and using different numbers of labeled nodes. The weights in this case are computed using the tree method and with radius $r = 1$. The Community Search algorithm outperforms both Spectral clustering and Tensor decomposition algorithms in both the experiments.

Table 1. Average and Best Classification Error (Number of Nodes that Are Misclassified in Each Estimated Community Compared to Ground-truth)

	W ($m = 2$)	W ($m = 4$)	W ($m = 6$)	W ($m = 8$)	W ($m = 10$)	T	S
Mean	56	55.64	55.32	55.30	54.98	60	70
Best	55	54	54	53	53	60	70

Obtained by Community Search Algorithm (W) with Spectral Clustering (S) (Ng et al. 2002) and Tensor Decomposition (T) (Anandkumar et al. 2013) Algorithms on US Political Blogosphere Network (Adamic and Glance 2005). The Community Search Algorithm Achieves the Best Classification Error of 53 and Better Average Error Over the Competing Algorithms.

4.5 Results on Real Datasets

In this section, we evaluate the performance of the Community Search algorithm on two real-world networks.

In the first experiment, we consider the *US political blogosphere network* first introduced by Adamic and Glance (2005), where nodes correspond to political blogs classified as either liberal or conservative during the 2004 US election, and edges represent hyperlinks between them. We consider the largest connected component of the network having 1,222 nodes and 16,716 edges. This dataset provides the ground-truth labels (liberal or conservative) for each node; these labels were manually generated by authors in Adamic and Glance (2005) according to their content. The largest component has two communities of sizes 586 and 636 according to this ground truth.

In this semi-supervised graph clustering setting, we randomly choose $m \in \{2, 4, 6, 8, 10\}$ labeled nodes from each ground-truth community as side information. Our performance metric is the classification error, namely, the number of nodes wrongly classified in each estimated community compared to the ground-truth communities² ($e = |\{j \in V : j \in V_1, j \notin \hat{V}_1 \text{ or } j \in \hat{V}_1, j \notin V_1\}|$).

For each m , we observe the overall performance of the Community Search algorithm over 50 different random choices of labeled nodes. As before, we compare the performance with Tensor decomposition and Spectral clustering algorithms. For the Community Search algorithm, we compute weights using the tree method with radius $r = 1$. In Table 1, we show the best and average classification error obtained by the clustering algorithms. With $r = 1$, the Community Search algorithm shows a better classification error than both Tensor and Spectral algorithms. In fact, our algorithm achieves the best classification error of 53, which is better than other state-of-the-art algorithms (Jin 2015; Gao et al. 2017), which achieved errors in the range 58–60 on this dataset. We also perform an in-depth analysis of the error cases in this dataset. We observed that 50 nodes in the graph do not satisfy the community definition since they share fewer neighbors in their ground truth community (in degree) than the second community (out degree). Since the best error in our algorithm is 53, it appears that our algorithm performs close to the best achievable error in this dataset.

In our next experiment, we consider the *Facebook-ego network* dataset from Leskovec and Krevl (2014), first introduced in Leskovec and McAuley (2012). The network corresponds to a Facebook graph with 4,039 user nodes and 88,234 edges. The dataset also contains 193 ground-truth ego circles, where each circle is a group of users sharing a particular interest, e.g., circle of college friends, family, and so on. Hence, each circle here corresponds to a community. We consider the 10 largest circles with more than 100 nodes each as the ground-truth communities. Note that in this dataset, some users can belong to multiple circles/communities, unlike a typical SBM. We

²Since there are only two communities, the estimation error in the first community is equal to that in the second; thus, we can count any one of them.

Table 2. Average Percentage Error Obtained by Community Search Algorithm

	W ($m = 5$)	W ($m = 10$)	W ($m = 15$)	T	S
Average Error	4.204%	2.712%	2.708%	4.927%	4.462%

(W) with Spectral Clustering (S) (Ng et al. 2002) and Tensor Decomposition (T) (Anandkumar et al. 2013) Algorithms to Estimate 9 Largest Ego Circles in Facebook–Ego Network (Leskovec and McAuley 2012). The Community Search Algorithm Achieves Lower Average Error than the Competing Algorithms.

remove such nodes from the ground-truth communities, but not from the graph. Even after this pruning, step 9 out of 10 ground-truth communities have more than 100 users each. Next, we try to recover these nine largest ground-truth communities by randomly choosing $m \in \{5, 10, 15\}$ nodes as labeled nodes. We run our Community Search algorithm by computing weights using the tree method with radius $r = 1$, and we compare the average percentage estimation error (according to ground-truth) with that of Spectral clustering and Tensor decomposition algorithms. From the results presented in Table 2, we observe that the Community Search algorithm has lower average error than the baselines even with just $m = 5$ labeled nodes. Also, as expected, the error reduces with increasing number of labeled nodes.

5 CONCLUSION AND DISCUSSION

In this article, we defined the search problem in community detection, provided a simple generic framework for incorporating side information, and a corresponding algorithm to solve the problem. Our algorithm analytically matches the state-of-the-art performance of existing algorithms that do not use side information, and empirically outperforms them on reliability and speed.

More generally, we believe that incorporating side information into graph analysis is a fertile and important area of research, as no real-world problem is a “pure” graph problem (i.e., where the only input is a graph) of the kind studied in, e.g., the vast majority of clustering literature.

There are several possible future directions: (A) Understanding fundamental limits of community detection (Mossel et al. 2014; Montanari 2015) when there is non-trivial side information (e.g., $\Theta(\log n)$ of labeled nodes in a community). (B) Richer notions of side information and corresponding problem definitions beyond search. (C) From a more practical viewpoint, we show in our experimental results (Section 4) that even this simple form of side information can dramatically reduce the computation time for searching communities, and also improve error performance. As discussed in the previous section, this work also provides a new method to parallelize graph clustering, an inherently difficult task. Adapting even faster algorithms, e.g., those based on belief-propagation, to this new semi-supervised setting is also an important prospect.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library.

ACKNOWLEDGMENTS

The authors also thank the Texas Advanced Computing Center (TACC 2018) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this article.

REFERENCES

Emmanuel Abbe, Afonso S. Bandeira, and Georgina Hall. 2016. Exact recovery in the stochastic block model. *IEEE Trans. Information Theory* 62, 1 (2016), 471–487.

- E. Abbe and C. Sandon. 2015. Community detection in general stochastic block models: Fundamental limits and efficient algorithms for recovery. In *Proceedings of the 2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 670–688.
- L. A. Adamic and N. Glance. 2005. The political blogosphere and the 2004 US election: Divided they blog. In *Proceedings of the 3rd International Workshop on Link Discovery*. ACM, 36–43.
- N. Agarwal, A. S. Bandeira, K. Koiliaris, and A. Kolla. 2015. Multisection in the stochastic block model using semidefinite programming. *arXiv Preprint arXiv:1507.02323* (2015).
- Nir Ailon, Yudong Chen, and Huan Xu. 2013. Breaking the small cluster barrier of graph clustering. In *Proceedings of the 30th International Conference on Machine Learning*. 995–1003.
- A. A. Amini, A. Chen, P. J. Bickel, and E. Levina. 2013. Pseudo-likelihood methods for community detection in large sparse networks. *The Annals of Statistics* 41, 4 (2013), 2097–2122.
- Animashree Anandkumar, Rong Ge, Daniel Hsu, and Sham Kakade. 2013. A tensor spectral approach to learning mixed membership community models. In *Proceedings of the Conference on Learning Theory*. 867–881.
- A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky. 2014. Tensor decompositions for learning latent variable models. *The Journal of Machine Learning Research* 15, 1 (2014), 2773–2832.
- A. Anandkumar, Y. Liu, D. J. Hsu, D. P. Foster, and S. M. Kakade. 2012. A spectral algorithm for latent dirichlet allocation. In *Advances in Neural Information Processing Systems*. 917–925.
- K. O. Bowman and L. R. Shenton. 2004. Estimation: Method of moments. In *Encyclopedia of Statistical Sciences*.
- J. T. Chang. 1996. Full reconstruction of Markov models on evolutionary trees: Identifiability and consistency. *Mathematical Biosciences* 137, 1 (1996), 51–73.
- K. Chaudhuri, F. Chung, and A. Tsiatas. 2012. Spectral clustering of graphs with general degrees in the extended planted partition model. *Journal of Machine Learning Research* 12 (2012), 1–23.
- Y. Chen, S. Sanghavi, and H. Xu. 2012. Clustering sparse graphs. In *Advances in Neural Information Processing Systems*, Vol. 25.
- A. Condon and R. M. Karp. 2001. Algorithms for graph partitioning on the planted partition model. *Random Structures and Algorithms* 18, 2 (2001), 116–140.
- S. Fortunato. 2010. Community detection in graphs. *Physics Reports* 486, 3 (2010), 75–174.
- Y. Fujiwara and G. Irie. 2014. Efficient label propagation. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*. 784–792.
- Chao Gao, Zongming Ma, Anderson Y. Zhang, and Harrison H. Zhou. 2017. Achieving optimal misclassification proportion in stochastic block models. *Journal of Machine Learning Research* 18 (2017), 60:1–60:45.
- D. Hsu and S. M. Kakade. 2013. Learning mixtures of spherical gaussians: Moment methods and spectral decompositions. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*. ACM, 11–20.
- F. Huang, UN. Niranjan, M. Hakeem, and A. Anandkumar. 2013. Fast Detection of Overlapping Communities via Online Tensor Methods. (2013). <http://arxiv.org/abs/1309.0787>.
- J. Jin. 2015. Fast community detection by SCORE. *The Annals of Statistics* 43, 1 (2015), 57–89.
- B. Kulis, S. Basu, I. Dhillon, and R. Mooney. 2009. Semi-supervised graph clustering: A kernel approach. *Machine Learning* 74, 1 (2009), 1–22.
- J. Leskovec and A. Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. Retrieved from <http://snap.stanford.edu/data>.
- J. Leskovec and J. McAuley. 2012. Learning to discover social circles in ego networks. In *Advances in Neural Information Processing Systems*. 539–547.
- Y. Lu and C. Zhai. 2008. Opinion integration through semi-supervised topic modeling. In *Proceedings of the 17th International Conference on World Wide Web*. ACM, 121–130.
- J. McAuley and J. Leskovec. 2012. Learning to discover social circles in ego networks. In *NIPS*, Vol. 2012. 548–56.
- J. D. McAuliffe and D. M. Blei. 2008. Supervised topic models. In *Advances in Neural Information Processing Systems*. 121–128.
- F. McSherry. 2001. Spectral partitioning of random graphs. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science, 2001*. IEEE, 529–537.
- Andrea Montanari. 2015. Finding one community in a sparse graph. *Journal of Statistical Physics* 161, 2 (2015), 273–299.
- E. Mossel, J. Neeman, and A. Sly. 2014. Reconstruction and estimation in the planted partition model. *Probability Theory and Related Fields* (2014), 1–31.
- Mark EJ Newman. 2006. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences* 103, 23 (2006), 8577–8582.
- Mark EJ Newman and Gesine Reinert. 2016. Estimating the number of communities in a network. *Physical Review Letters* 117, 7 (2016), 078301.
- A. Y. Ng, M. I. Jordan, and Y. Weiss. 2002. On spectral clustering: Analysis and an algorithm. *Advances in Neural Information Processing Systems* 2 (2002), 849–856.

- N. Rao, H. F. Yu, P. Ravikumar, and I. Dhillon. 2015. Collaborative filtering with graph information: Consistency and scalable methods. In *Advances in Neural Information Processing Systems*.
- A. Ray, S. Sanghavi, and S. Shakkottai. 2016. Searching for a single community in a graph. In *Proceedings of the ACM Sigmetrics 2016 (Poster Paper)*. ACM.
- K. Rohe, S. Chatterjee, and B. Yu. 2011. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics* 39, 4 (2011), 1878–1915.
- R. Sibson. 1973. SLINK: An optimally efficient algorithm for the single-link cluster method. *Comput. J.* 16, 1 (1973), 30–34.
- TACC. 2018. Texas Advanced Computing Center. Retrieved from <http://www.tacc.utexas.edu>.
- J. Tropp. 2015. An introduction to matrix concentration inequalities. *arXiv Preprint arXiv:1501.01571* (2015).
- P. Wedin. 1972. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics* 12, 1 (1972), 99–111.
- J. Xie, S. Kelley, and B. K. Szymanski. 2013. Overlapping community detection in networks: The state-of-the-art and comparative study. *ACM Computing Surveys (CSUR)* 45, 4 (2013), 43.
- Z. Xu, Y. Ke, Y. Wang, H. Cheng, and J. Cheng. 2012. A model-based approach to attributed graph clustering. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*. ACM, 505–516.
- J. Yang and J. Leskovec. 2012. Defining and evaluating network communities based on ground-truth. In *Proc. of the ACM SIGKDD Workshop on Mining Data Semantics*. ACM, 3.
- J. Yang, J. McAuley, and J. Leskovec. 2013. Community detection in networks with node attributes. In *2013 IEEE 13th International Conference on Data Mining (ICDM)*. IEEE, 1151–1156.
- Yi Yu, Tengyao Wang, and Richard J. Samworth. 2015. A useful variant of the Davis-Kahan theorem for statisticians. *Biometrika* 102, 2 (2015), 315–323.
- S. Yun and A. Proutiere. 2014. Accurate community detection in the stochastic block model via spectral algorithms. *arXiv Preprint arXiv:1412.7335* (2014).
- Yuan Zhang, Elizaveta Levina, Ji Zhu, and others. 2016. Community detection in networks with node features. *Electronic Journal of Statistics* 10, 2 (2016), 3153–3178.
- D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf. 2004. Learning with local and global consistency. *Advances in Neural Information Processing Systems* 16, 16 (2004), 321–328.
- X. Zhu. 2005. Semi-supervised learning literature survey. Retrieved from http://pages.cs.wisc.edu/jerryzhu/pub/ssl_survey.pdf.

Received November 2016; revised March 2018; accepted March 2018