

Active Learning for Community Detection in Stochastic Block Models

Akshay Gadde, Eyal En Gad, Salman Avestimehr and Antonio Ortega
University of Southern California, Los Angeles

Email: {agadde,engad}@usc.edu, avestimehr@ee.usc.edu, antonio.ortega@sipi.usc.edu

Abstract—The stochastic block model (SBM) is an important generative model for random graphs in network science and machine learning, useful for benchmarking community detection (or clustering) algorithms. The symmetric SBM generates a graph with $2n$ nodes which cluster into two equally sized communities. Nodes connect with probability p within a community and q across different communities. We consider the case of $p = a \ln(n)/n$ and $q = b \ln(n)/n$. In this case, it was recently shown that recovering the community membership (or label) of every node with high probability (w.h.p.) using only the graph is possible if and only if the Chernoff-Hellinger (CH) divergence $D(a, b) = (\sqrt{a} - \sqrt{b})^2 \geq 1$. In this work, we study if, and by how much, community detection below the clustering threshold (i.e. $D(a, b) < 1$) is possible by querying the labels of a limited number of chosen nodes (i.e., active learning). Our main result is to show that, under certain conditions, sampling the labels of a vanishingly small fraction of nodes (a number sub-linear in n) is sufficient for exact community detection even when $D(a, b) < 1$. Furthermore, we provide an efficient learning algorithm which recovers the community memberships of all nodes w.h.p. as long as the number of sampled points meets the sufficient condition. We also show that recovery is not possible if the number of observed labels is less than $n^{1-D(a,b)}$. The validity of our results is demonstrated through numerical experiments.

I. INTRODUCTION

Community detection (or clustering) is a fundamental problem in the study of networks [1]. In this problem, it is assumed that each node (or vertex) of a network belongs to one out of a few latent groups or communities, and that the topology of the network depends on these *a priori* unknown group memberships (or labels). The goal is to recover the communities by partitioning the nodes into different classes, with denser connections within classes and sparser connections between classes. This problem arises in many areas, such as detecting protein complexes in protein interaction networks [2] and image classification and segmentation [3].

In community detection problems, stochastic block models (SBM) [4] are random graph models that are commonly used for benchmarking the performance of different clustering algorithms and deriving fundamental limitations of graph clustering. In this model, an edge exists between a pair of nodes independently with a probability that depends on the community membership of the two nodes. The simplest example of an SBM, which is the one considered in this paper, contains two communities of equal sizes, such that a pair of nodes from the same community are connected with probability p , and nodes of different communities are connected with probability q . The limits on the performance of

different algorithms in terms of the relative difference between p and q have been established (see [5] and references therein).

In recent years, the fundamental limits of community detection in SBM (using any method) have been studied [6]. In particular, an elegant threshold was discovered for exact recovery of labels (also known as strong consistency), which is defined as finding the correct community memberships of all nodes asymptotically almost surely (a.a.s.). In the case of an SBM with $2n$ nodes and two communities of equal sizes, having within-community edge probability $p = a \ln n/n$ and between-communities edge probability $q = b \ln n/n$ (where a and b are constants greater than 1), exact recovery was shown to be possible if and only if $D(a, b) = (\sqrt{a} - \sqrt{b})^2 \geq 1$ [7]. We call $D(a, b)$ the Chernoff-Hellinger (CH) divergence, in accordance with [5]. The condition $D(a, b) \geq 1$ can be written equivalently as $\Delta \geq \sqrt{M} - 1/4$ by defining

$$M := (a + b)/2 \text{ and } \Delta := (a - b)/2 \quad (1)$$

so that $2M \ln n$ is the average degree of a node and $2\Delta \ln n$ is the average difference between the number of neighbors of a node in the same community and the number of neighbors in the opposite community. In addition, efficient algorithms were proposed that perform exact recovery wherever the above threshold condition holds [5], [8].

This paper is motivated by the following fundamental, and practically important question, to determine *if, and by how much, community detection below the clustering threshold (i.e., $D(a, b) < 1$) is possible by querying the labels of a vanishingly small fraction of carefully selected nodes (i.e., active learning)*. The above question is essential in the context of recent trends of involving human workforce in data analytic problems (e.g., Amazon Mechanical Turk [9]). In such setting, selecting the smallest number of most useful nodes for labeling is important for minimizing human effort. We show that community detection below the threshold is possible by querying the labels of a vanishingly small fraction of nodes (i.e., $o(n)$ number of nodes). We provide a bound on the number of observations (or samples) required as well as an algorithm which specifies which nodes to sample.

We consider an active learning framework in which an algorithm can choose which nodes to sample, up to a limited budget of $B(n)$ nodes. For simplicity, we restrict the budget to take the form $B(n) = n^s$, for a constant s . The main result of this paper is that exact recovery of the community membership is possible if $4\Delta > 1/3 + \sqrt{4M + 1/9}$ and $s \geq 1 - \delta(a, b)$,

where $\delta(a, b)$ is as specified in Theorem 1. Furthermore, we provide an efficient sampling algorithm which recovers the community memberships of all nodes a.a.s. wherever $s \geq 1 - \delta(a, b)$. Additionally, we also show that recovery is not possible if the number of observed labels is less than $n^{1-D(a,b)}$.

The proposed algorithm begins by obtaining an initial clustering using an approach presented in [8], which consists of spectral clustering followed by a replica method. We can characterize the nodes where this initial clustering is more likely to make mistakes. We propose to observe the labels of $B(n)$ nodes selected using this characterization, which allows us to correct the mistakes in the initial clustering. We show that observing $\Omega(n^{1-\delta})$ labels in the proposed way and predicting the labels of the unsampled nodes according to the initial clustering guarantees that we a.a.s. recover all the labels. While most of the techniques used in the proof of our main result are based on previous works [8], [5], to the best of our knowledge, the proposed sampling algorithm is novel.

The problem of selecting a subset of nodes for labeling, which are most helpful in predicting the labels of the rest of the nodes is known as active learning on graphs [10], [11] in the machine learning literature. In a standard active learning setting, the nodes represent data points in a Euclidean space. A k -nearest neighbor graph is used to represent the similarities between them. Since the number of samples required (also called sampling complexity) by the proposed algorithm for exact recovery is sub-linear (and thus, vanishingly small) in the number of nodes in the SBM, we suspect a similar approach to be useful in the above setting.

Prior work. The problem of community detection with random sampling (as opposed to active learning) in stochastic block models has been previously studied in the context of sparse SBMs, with $p = a/n$ and $q = b/n$ [12]. In sparse SBM, only weak recovery is expected, where the goal is to predict labels which have non-trivial correlation with true labels. It was shown in [13], [14] that observing labels of a vanishingly small fraction of nodes *randomly* does not affect the weak recovery asymptotically but can lead to efficient local algorithms (in which the label prediction at each node depends only on its local neighborhood), whenever recovery is possible. The effect of sampling on exact recovery in SBM with log-degree, i.e., with $p = a \ln n/n$ and $q = b \ln n/n$ has not been studied. Our result in this paper demonstrates that active learning can largely impact community detection threshold in the regime of $p = a \ln n/n$ and $q = b \ln n/n$.

II. BACKGROUND AND MAIN RESULT

We begin with some basic definitions and notation.

Definition 1 (Stochastic Block Model [4]). *A stochastic block model is a random graph $G = (V, E) \sim \mathcal{G}(n, p, q)$ with $2n$ nodes constructed as follows: Select a labeling $\sigma: V \rightarrow \{1, -1\}$ for each node $v \in V$ independently and uniformly at random (i.e., $\Pr(\sigma_v = i) = 1/2$, $i \in \{1, -1\}$). Each pair of nodes (v, u) is connected by an edge with probability p if $\sigma_u = \sigma_v$ and with probability q if $\sigma_u \neq \sigma_v$.*

Community detection in SBM is equivalent to reconstruction of labels σ (up to its sign) using G . The correctness of predicted labels is considered in an asymptotic sense, where p and q are allowed to be functions of n and $n \rightarrow \infty$. We focus on the logarithmic degree regime, i.e., the case when $p_n = a \frac{\ln n}{n}$ and $q_n = b \frac{\ln n}{n}$, where a and b are constants. The notion of correctness, which is of most interest in this regime is called strong consistency or exact recovery.

Definition 2 (Strong Consistency [8]). *Given sequences p_n and q_n , the predicted labeling $\tau = \tau(G)$ is said to be strongly consistent if $\Pr(\tau = \sigma \text{ or } \tau = -\sigma) \rightarrow 1$ as $n \rightarrow \infty$.*

Note that strong consistency requires the number of errors to go to 0. A weaker notion of correctness is also considered in the literature, called weak consistency (or partial recovery). This notion only requires that the fraction of errors goes to 0, i.e., the number of errors is $o(n)$.

We next define differential degree, which will be useful for characterizing the fundamental limits on clustering and describing the sampling algorithm proposed in the next section.

Definition 3 (Differential Degree). *Differential degree of a node v w.r.t. a labeling τ is defined as*

$$d_\tau^*(v) = |\{u \in V | u \sim v, \tau_u = \tau_v\}| - |\{u \in V | u \sim v, \tau_u \neq \tau_v\}|.$$

v is said to have a majority if $d_\tau^(v) > 0$. Otherwise, v has a minority.*

Intuitively, any clustering algorithm, which groups the nodes such that there are more edges within communities and fewer edges between communities, can succeed only if each node belongs to the same community as the majority of its neighbors. This necessary local condition also turns out to be sufficient for recovery. Specifically, it has been shown [8] that a strongly consistent estimator for $\mathcal{G}(n, p_n, q_n)$ exists if and only if a.a.s. $d_\sigma^*(v) > 0 \quad \forall v \in V$. This, in turn, holds if and only if $D(a, b) = (\sqrt{a} - \sqrt{b})^2 \geq 1$ or equivalently, $\Delta \geq \sqrt{M} - 1/4$ (hereafter we use D as a shorthand for $D(a, b)$). The necessary part of the condition can be proven by analyzing the probability that a node v has a minority, which can be computed using the fact that the number of neighbors of each node in the same (or different) community follows a binomial distribution with parameter $p = a \ln n/n$ (or $q = b \ln n/n$).

An algorithm to obtain a strongly consistent labeling under the assumption $D \geq 1$ is given in [8]. This algorithm consists of two steps. The first step (spectral clustering followed by replica method) gives a weakly consistent labeling and the final step flips the labels of the minority nodes (with respect to the prediction in the previous step) to ensure that every node is a majority. The first step of the algorithm is *stable* in the sense that it works even when $D \leq 1$ and only requires $a > b$ to guarantee weak consistency.

The fundamental question considered in this paper is whether one can get a strongly consistent estimator for the labels even when $D \leq 1$, by observing labels of a vanishingly small fraction of carefully selected nodes (in addition to G). In




Fig. 1: $1 - \delta$ vs. $\Delta/\sqrt{M - 1/4}$ for different values of M .

order to answer this question, we need to specify the number of samples needed as well as the locations of those samples. To do this, we first define critical differential degree as follows:

Definition 4 (Critical Differential Degree). *Critical differential degree ℓ_{critical} is defined as the smallest number such that a.s. no node v has $d_{\sigma}^*(v) \leq -\ell_{\text{critical}}$, i.e.,*

$$\ell_{\text{critical}} = \inf\{\ell \mid \Pr(d_{\sigma}^*(v) \leq -\ell) = o(n^{-1})\}.$$

Based on this definition, following theorems give the sufficient and necessary conditions on the number of samples needed.

Theorem 1 (Sufficient Condition). *Sampling the labels of a subset of nodes $S := \{v \mid d_{\sigma}^*(v) \leq \ell_{\text{critical}}\}$ is sufficient for exact recovery in $\mathcal{G}(n, a \frac{\ln n}{n}, b \frac{\ln n}{n})$. Moreover, assuming $4\Delta > 1/3 + \sqrt{4M + 1/9}$, $|S| \leq n^s$ if $s \geq 1 - \delta(M, \Delta)$, where*

$$\delta(M, \Delta) = \frac{\left(4\Delta - 1/3 - \sqrt{4M + 1/9}\right)^2}{2\left(2M + 4\Delta - 1/3 - \sqrt{4M + 1/9}\right)}. \quad (2)$$

The above expression shows that the number of samples needed is smaller as Δ increases relative to M .

Theorem 2 (Necessary Condition). *Let $B(n) = n^s$ be the number of observed labels. Then, exact recovery is not possible for $\mathcal{G}(n, a \frac{\ln n}{n}, b \frac{\ln n}{n})$ if $s \leq 1 - D(a, b)$.*

Remark. Note that Theorem 2 is congruous with previous results [5], [8]. If $D \leq 0$ then it is not possible to find even a weakly consistent estimator and the number of samples required for exact recovery is linear in n . On the other hand, if $D \geq 1$, then exact recovery is possible without sampling.

While exact recovery is not possible without sampling when $D < 1$ (or equivalently $\Delta < \sqrt{M - 1/4}$), Figure 1 shows that the number of samples needed (as given by Theorem 1) is sub-linear in n provided $\Delta \gtrsim 0.55\sqrt{M - 1/4}$ (or $D \gtrsim 0.3$). Thus, exact recovery is possible in this regime by observing the labels of a vanishingly small fraction of the nodes. The figure also shows that, for a fixed value of $\Delta/\sqrt{M - 1/4}$, the number of samples needed decreases as M , which specifies the total number of edges, increases.

The next section proves the sufficient condition in Theorem 1 by providing an algorithm for selecting the subset of

nodes to be sampled, and shows that sampling $n^{1-\delta}$ nodes using the proposed algorithm is enough for exact recovery. The proof of the necessary condition is provided in Section IV.

III. SUFFICIENT CONDITION

In this section, we give a constructive proof for the sufficient condition in Theorem 1 by providing an algorithm for sampling. We show that the sampling budget given by Theorem 1 is sufficient for exact recovery using the proposed algorithm. The main idea behind the proposed algorithm is as follows: even when $D < 1$, it is possible to get a weakly consistent estimator τ' for the labels (using Algorithm 1 from [8]), which makes only $o(n)$ mistakes as long as $D > 0$. If we can characterize the nodes which are wrongly labeled by the weakly consistent estimator and then correct the mistakes by sampling their labels, we will have exact recovery. It can be shown that w.h.p. the mistakes in τ' are restricted to nodes which have small differential degrees. Thus, the proposed algorithm samples nodes with smallest differential degrees $d_{\tau'}^*(v)$ (see Algorithm 1).

Algorithm 1 Sampling for exact recovery

Input: Graph $G = (V, E)$, Sampling budget $B(n)$

- 1: Find τ' (an initial guess for the labels) by clustering the nodes using Algorithm 1 from [8].
 - 2: Compute differential degrees $d_{\tau'}^*(v)$ for all $v \in V$.
 - 3: Sample $B(n)$ nodes with the smallest $d_{\tau'}^*(v)$.
 - 4: Update τ' using the sampled labels to get τ .
-

Theorem 3. *Let τ be the labeling given by Algorithm 1 after sampling $B(n) = n^s$ nodes. Then τ is strongly consistent if $4\Delta > 1/3 + \sqrt{4M + 1/9}$ and $s \geq 1 - \delta(M, \Delta)$.*

The proof of Theorem 3 follows by extending and combining several results and techniques from [8]. We begin by characterizing the set of nodes, where the weakly consistent estimator τ makes mistakes.

Definition 5. *For a positive real number ϵ , let V_{ϵ} be the elements of V that either have a differential degree $d_{\sigma}^*(v)$ less than $\epsilon\sqrt{a}\ln n$ or have more than $100np$ neighbors.*

Proposition 1 (Proposition 4.3 in [8]). *For any $\epsilon > 0$, the initial labeling τ' in Algorithm 1 a.s. correctly labels every node in $V \setminus V_{\epsilon}$.*

Proposition 2 (Proposition 4.7 in [8]). *For sufficiently small ϵ , a.s. no two nodes in V_{ϵ} are adjacent.*

The above proposition ensures that a.s. $|d_{\tau'}^*(v)| = |d_{\sigma}^*(v)|$ for $v \in V_{\epsilon}$. Our idea is to find ℓ_{critical} such that a.s. no node has differential degree (w.r.t. σ) less than $-\ell_{\text{critical}}$. Since $|d_{\tau'}^*(v)| = |d_{\sigma}^*(v)|$ for all $v \in V_{\epsilon}$, sampling nodes with $d_{\tau'}^*(v) \leq \ell_{\text{critical}}$ will capture all the nodes in V_{ϵ} with $d_{\sigma}^*(v) \leq \ell_{\text{critical}}$.

Proposition 3. $\ell_{\text{critical}} = \left(\sqrt{4M + 1/9} + 1/3 - 2\Delta\right) \ln n$.

Proof. Let $X_i \sim \text{Bernoulli}(p = a \ln n/n)$ and $Y_i \sim \text{Bernoulli}(q = b \ln n/n)$. Define $Z_i = X_i - Y_i$:

$$Z_i = \begin{cases} +1 & \text{with prob. } p(1-q) \\ 0 & \text{with prob. } pq + (1-p)(1-q) \\ -1 & \text{with prob. } (1-p)q \end{cases}$$

Note that $d_\sigma^*(v) = \sum_{i=1}^n Z_i$. Therefore,

$$\begin{aligned} \Pr(d_\sigma^*(v) \leq -\ell) &= \Pr\left(\sum_{i=1}^n Z_i \leq -\ell\right) \\ &= \Pr\left(\sum_{i=1}^n -\left(Z_i - (a-b)\frac{\ln n}{n}\right) > \ell + (a-b)\ln n\right). \end{aligned}$$

Applying Bernstein inequality to the sum of r.v.'s $\tilde{Z}_i := -(Z_i - (a-b)\frac{\ln n}{n})$, we get $\Pr(d_\sigma^*(v) \leq -\ell) = o(n^{-1})$ if

$$\ell \geq \ell_{\text{critical}} = \left(\sqrt{2(a+b) + 1/9} + 1/3 - (a-b)\right) \ln n.$$

□

Now, $\Pr(d_\sigma^*(v) \leq \ell_{\text{critical}})$ can be similarly bounded using Bernstein inequality to get (3) (note that, in order to apply Bernstein inequality, we need to assume $4\Delta > 1/3 + \sqrt{4M + 1/9}$).

$$\begin{aligned} \Pr(d_\sigma^*(v) \leq \ell_{\text{critical}}) &\leq n^{-\delta}, \quad (3) \\ \text{with } \delta &= \frac{\left(4\Delta - 1/3 - \sqrt{4M + 1/9}\right)^2}{2\left(2M + 4\Delta - 1/3 - \sqrt{4M + 1/9}\right)}. \end{aligned}$$

For sufficiently small ϵ ,¹ $d_\sigma^*(v) \leq \ell_{\text{critical}}$ for all $v \in V_\epsilon$. Therefore, sampling the nodes with $d_{\tau'}^*(v) \leq \ell_{\text{critical}}$ will capture all the nodes in V_ϵ and thus, correct all the mistakes in the initial labeling τ' . A bound on the number of samples required is given by $n \cdot \Pr(d_\sigma^*(v) \leq \ell_{\text{critical}})$. Thus, sampling $n^{1-\delta}$ nodes with the proposed algorithm is sufficient for exact recovery.

IV. NECESSARY CONDITION

To prove the necessary condition of Theorem 2, we assume that $s < 1 - D$, and show that exact recovery is impossible. The proof follows the techniques used in [5] for the necessary condition in clustering. The proof considers the recovery of minority nodes w.r.t. σ . A minority node is a node for which the number of neighbors of the same community is smaller than the number of neighbors of the opposite community. The idea of the proof is to show that if there exists a minority node that was not sampled, then no algorithm can estimate the label of the minority node with vanishing error probability. Then we show, using simple counting, that the number of minority nodes is larger than the budget w.h.p., and therefore there must exist at least one unsampled minority node.

¹Specifically, we need $\epsilon < \frac{\ell_{\text{critical}}}{\sqrt{a} \ln n} = \frac{(\sqrt{2(a+b)+1/9}+1/3-(a-b))}{\sqrt{a}}$. In order to make ϵ smaller, one has to use more partitions in the replica step (see [8] for details.)

Consider estimation of label σ_v of a single node v . If the estimator knows the correct labeling of the rest of the nodes in the graph, then its error probability decreases. Therefore we assume that the estimator possesses this knowledge, to obtain a lower bound on the error probability. Next we notice that the estimator with the lowest error probability is the maximum a posteriori (MAP) estimator. Let X be a random variable that represents the number of neighbors of v of community 1. Similarly let Y be a random variable that represents the number of neighbors of v of community -1 . Given the observation of X and Y , the MAP estimator selects the label

$$\tau_v^* = \arg \max_{\tau_v} \mathbb{P}\{\sigma_v = \tau_v | X = x, Y = y\}.$$

Since the labels are a priori equiprobable, the MAP estimator is equal to the Maximum Likelihood (ML) estimator

$$\tau_v^* = \arg \max_{\tau_v} \mathbb{P}\{X = x, Y = y | \sigma_v = \tau_v\}.$$

By the definition of the graph, we see that the ML estimator always misclassifies a minority node.

It remains to show that w.h.p. the number of minority nodes exceeds the sampling budget. To do this, we use the fact that the probability that a node is a minority is $\Omega(n^{-D}/\ln(n))$ [5], [8]. With probability $1 - o(1)$, the size of each community is within $\sqrt{n} \ln n$ of n . Assume that this holds. We let U be a random set of $2n/\ln^3(n)$ nodes of G . With probability $1 - o(1)$, the number of nodes in U in each community is within \sqrt{n} of $n/\ln^3 n$, and a randomly selected node in U is adjacent to another node in U with probability $o(1)$.

We say that a node in U is ambiguous if it has a minority with respect to the nodes that are not U (i.e., not considering neighbors in U). The number of neighbors of a node in $V \setminus U$ in each community follow a multivariate Poisson distribution approximately. Furthermore, it is shown in [5] that the probability that a node is ambiguous is

$$\Omega(n^{-D}/\ln(n)).$$

By the assumption that $s < 1 - D$, it follows that there exists $\epsilon > 0$ such that a node in U is ambiguous with probability $\Omega(n^{\epsilon+s-1})$. For a fixed community assignment and choice of U , there is no dependence between whether or not any two nodes are ambiguous. Also, an ambiguous node is not adjacent to any other node in U with probability $1 - o(1)$. So, for any choice of n^s nodes to sample, with probability $1 - o(1)$ there is at least one unsampled ambiguous node, which is misclassified by an optimal classifier. Therefore, any classifier misclassifies at least one node with probability $1 - o(1)$.

V. EXPERIMENTS

We generate random graphs using SBM with different values of D . b is set to 2 and D is varied from 0.3 to 0.6. For each value of D , graphs with number of nodes, n , changing from 1000 to 2000 are generated. For each D and n , we compute the following quantities (averaged over 30 trials):

$$\begin{aligned} n_m &= \text{number of minorities} \\ n_s &= \text{number of nodes with } d_{\tau'}^*(v) \leq d_{\tau', \text{err}, \max}^* \end{aligned} \quad (4)$$




Fig. 2: Exponents in (5) as function of D and the theoretical bound on γ given by Theorem 1.

$d_{\tau', \text{err}, \text{max}}^*$ denotes the maximum differential degree w.r.t. the initial label predictions τ' among the nodes, where $\tau'_v \neq \sigma_v$. Note that sampling n_s nodes with the proposed algorithm will capture all the errors in the initial labeling τ' . Number of minorities n_m can be thought of as a necessary lower bound on the number of samples required according to Theorem 2. For each value of D , we express the quantities in (4) as:

$$\begin{aligned} n_m &\approx c_1 n^\alpha \\ n_s &\approx c_2 n^\gamma. \end{aligned} \quad (5)$$

The exponents α and γ are computed by fitting a linear function of $\log n$ to the observed values of $\log n_m$ and $\log n_s$ for each value of D . (Note that the linear regression needs to be done jointly for different values D in order to keep c_1 and c_2 fixed as D changes. If linear regression produces $\gamma > 1$, then it is set to 1, since it implies sampling all nodes.). Figure 2 plots the exponents as a function of D along with the theoretical bound $1 - \delta$ given by Theorem 1. It shows that the exponent in the expression for n_m , $\alpha \approx 1 - D$, which is consistent with the theoretical value given in [8] and thus, acts as a sanity check for the results. Figure 2 also verifies that the exponent γ in the expression for n_s is indeed less than the sufficient theoretical bound given by Theorem 1. The experimentally obtained values of γ seem to be much lower than the theoretical bound $1 - \delta$, suggesting that a potentially tighter bound for n_s can be proven.

We also plot the error rate vs. the fraction of nodes sampled in Figure 3, for graphs of sizes $n = 1000, 2000$ and 4000 keeping D fixed at 0.5. The figure shows that the a smaller fraction of the nodes needs to be sampled to achieve the same error as the graph size increases. The fraction of nodes needed to achieve zero error is well below the theoretical bound and decreases as n increases. This is expected since the number of samples needed for exact recovery is a sub-linear function of n according to Theorem 1.

VI. CONCLUSION

In this paper, we considered the problem of active learning for community detection in SBM. Community detection, using only the graph, is possible if and only if the relative difference between the within-community and between-community




Fig. 3: Error rate vs. fraction of nodes sampled at $D = 0.5$ for different graph sizes n . The theoretical bound on the fraction of samples needed is 0.56, 0.54 and 0.51 for $n = 1000, 2000$ and 4000 respectively.

connection probabilities is larger than certain threshold. We showed that community detection is possible below this threshold by observing the labels of a small fraction of carefully selected nodes in addition to the graph. We gave the necessary and sufficient conditions on the number of samples required and also provided an algorithm that specifies which nodes should be sampled. We also verified our results through numerical experiments. In future, we would like to generalize our result to an SBM with more than two and possibly overlapping communities. Another interesting extension is to consider community detection and active learning in random graph models with geometry.

REFERENCES

- [1] S. Fortunato, "Community detection in graphs," *Physics Reports*, 2010.
- [2] J. Chen and B. Yuan, "Detecting functional modules in the yeast protein-protein interaction network," *Bioinformatics*, 2006.
- [3] J. Shi and J. Malik, "Normalized cuts and image segmentation," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 2000.
- [4] P. W. Holland, K. B. Laskey, and S. Leinhardt, "Stochastic blockmodels: First steps," *Social networks*, 1983.
- [5] E. Abbe and C. Sandon, "Community detection in general stochastic block models: fundamental limits and efficient recovery algorithms," 2015, arXiv preprint arXiv:1503.00609.
- [6] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, "Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications," *Physical Review E*, 2011.
- [7] E. Abbe, A. S. Bandeira, and G. Hall, "Exact recovery in the stochastic block model," *Information Theory, IEEE Transactions on*, vol. 62, no. 1, pp. 471–487, 2016.
- [8] E. Mossel, J. Neeman, and A. Sly, "Consistency thresholds for the planted bisection model," *Electronic Journal of Probability*, 2016, arXiv:1407.1591v4.
- [9] Y. Yan, G. M. Fung, R. Rosales, and J. G. Dy, "Active learning from crowds," in *International conference on machine learning*, 2011.
- [10] B. Settles, "Active learning literature survey," *University of Wisconsin, Madison*, 2010.
- [11] A. Gadde, A. Anis, and A. Ortega, "Active semi-supervised learning using sampling theory for graph signals," in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2014.
- [12] P. Zhang, C. Moore, and L. Zdeborová, "Phase transitions in semi-supervised clustering of sparse networks," *Physical Review E*, 2014.
- [13] V. Kanade, E. Mossel, and T. Schramm, "Global and local information in clustering labeled block models," *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 2014.
- [14] E. Mossel and J. Xu, "Local algorithms for block models with side information," in *ACM Conference on Innovations in Theoretical Computer Science*, 2016.