

UNSUPERVISED ROBUST NONPARAMETRIC LEARNING OF HIDDEN COMMUNITY PROPERTIES

MIKHAIL LANGOVOY*, AKHILESH GOTMARE AND MARTIN JAGGI

Machine Learning and Optimization Laboratory
EPFL
Lausanne, CH-1015, Switzerland

(Communicated by Wan-Quan Liu)

ABSTRACT. We consider learning of fundamental properties of communities in large noisy networks, in the prototypical situation where the nodes or users are split into two classes according to a binary property, e.g., according to their opinions or preferences on a topic. For learning these properties, we propose a nonparametric, unsupervised, and scalable graph scan procedure that is, in addition, robust against a class of powerful adversaries. In our setup, one of the communities can fall under the influence of a knowledgeable adversarial leader, who knows the full network structure, has unlimited computational resources and can completely foresee our planned actions on the network. We prove strong consistency of our results in this setup with minimal assumptions. In particular, the learning procedure estimates the baseline activity of normal users asymptotically correctly with probability 1; the only assumption being the existence of a single implicit community of asymptotically negligible logarithmic size. We provide experiments on real and synthetic data to illustrate the performance of our method, including examples with adversaries.

1. Introduction. We develop robust and scalable methods to uncover global properties of communities hidden in large networks with noisy node attributes. Consider the fundamental situation where the nodes or users in the network are split into two classes according to a binary property, such as their opinion or preferences on a specific topic. We call these two classes the “active” and “inactive” users, respectively. Examples include support of a particular candidate in elections [2], [30], or a level of interest in a particular topic, or a degree of support of certain statement. Pixels in digital images can be viewed as network nodes with attributes, so many applications from image processing, such as road tracking [9], [8], or medical tumor detection [21], [7], can be treated within this framework as well. In disease outbreak detection in epidemiology, people are naturally falling into two categories, healthy or sick [24], [1].

Additionally, motivated by modern real-world settings, we assume that the network of interest is too large to be processed manually, especially for each possible

2010 *Mathematics Subject Classification.* Primary: 62G05, 90B15, 68T05; Secondary: 65C60, 62G35.

Key words and phrases. Nonparametric learning, unsupervised learning, hidden community, scan estimator, community properties, learning for networks, adversarial learning, non-sparse graphs, crawler, scalability.

* Corresponding author: Mikhail Langovoy.

topic of interest. Therefore, activity observations of users are determined and delivered to us by a third-party algorithm called the *crawler*. Naturally, the crawler has its classification and learning errors that are not known to us. Therefore, we treat a general non-parametric case of the crawler error probabilities.

Our general goal is to learn global properties of communities of active and inactive users despite such noise and errors, in an unsupervised way, while additionally being robust to an adversary. The specific goal of this paper is unsupervised learning of the infinite-dimensional vector (a, b, F) , where a and b are baseline activities corresponding to inactive and active users correspondingly. They are fundamental global intrinsic properties of topical communities of users and are needed to be able to detect and recover active and inactive communities in subsequent analysis. An unknown probability distribution function F represents performance of the crawling algorithm for the particular topic, and serves as a natural measure of quality of our knowledge about the network's content.

1.1. Distinctive features of our framework. We infer global characteristic properties of communities, using both the network topology as well as the network's content represented by node attributes (see, e.g., [25] or [31] for related types of setups). Regarding the topology, we assume that the edges can be observed, but do not impose any global assumptions on the graph structure or any edge density assumptions, be it global or local. Naturally, opposing communities in our setup do not overlap.

As for node attributes, we assume that the nodes are corrupted by nonparametric noise of unknown and nonsmooth distribution. Notably, our noise model permits strong and long-distance correlations between observations on active vertices.

We treat the setup where *active* users can fall under the influence of an *adversary* who is capable of directly altering their activity values, pursuing the goal of spoiling our uncovering of community and network properties. The adversary knows the true values on active vertices and which vertices are inactive, knows the full graph structure, and has unlimited computational power and memory. Moreover, the adversary has a special deal with *hackers* and can completely foresee the actions we will be doing on the network, including the outcomes of all our randomized procedures, in case if we choose to use any random or pseudo-random number generator.

In summary, our key **contributions** are:

A robust scalable graph k -NN scan estimator (§3). The k -NN scans presented here generalize sliding windows, moving averages, and scan statistics, extending them to the case of general graphs. Our probabilistic analysis of the estimators is related to extreme order statistics for dependent random variables.

Sufficient conditions for k -NN graph scan estimators to be consistent for learning global properties of communities in an adversarial framework. This is a remarkable property of these local scans, as it can be easily seen that most existing methods can easily be spoiled either by such an adversary or by sheer dimensionality of a very general nonparametric framework that we are considering.

Discussion of aspects of our estimator's computation to allow large scale graphs, in particular via a highly decentralized implementation, allowing scalable distributed and parallel graph processing.

Our estimator is non-parametric, yet scalable method for learning properties of communities in noisy graphs. Moreover, it operates under minimal assumptions about the structure of the network and the communities. Unlike many existing methods, it utilizes the content of the network and thus does not require any global assumptions on the topology of the network or on relations between communities, such as graph sparsity or presence of hidden highly connected communities. The only local condition that we impose on the network is that it has one compact locally connected community of size that grows at least as a logarithm of the total number of vertices; this community can be asymptotically negligible in the large scale setting. Therefore, we believe that these results are foundational to fast consistent algorithms for community detection via percolation on general graphs. For the special case of lattices, the automated detection theorem [16] serves as an example illustration of this approach.

Finally, learning the actual performance of the crawler can also be addressed as a by-product of results of this paper. We note that our results are, of course, valid in their present form in the special case when there is no adversarial action.

1.2. Related work. Scan statistics have been long used for detection of unexpected events and for nonparametric estimation. The initial idea and the first development of the underlying theory, for one-dimensional discrete case, goes back at least to [29], who studied longest runs of successes in Bernoulli sequences. In the two-dimensional case, a surge of interest to continuous scan statistics has been sparked by [12]. This and related types of Euclidean spatial scan statistics found numerous applications in geostatistics, medicine, epidemiology and ecological studies, see examples in [28], [13, 14], [23]. The scan statistic methodology is compatible with Bayesian paradigm as well [22]. Notably, despite substantial efforts, most of the work in this area is based on heuristics and experimentation rather than on rigorous probabilistic analysis.

Recently, a new line of research emerged where discrete versions of multidimensional scan statistics were applied to discrete structures such as pixelized images or lattices [17], [18], [19], [4]. The idea of discrete scans proved useful in application areas like anomaly detection and automated detection of unknown objects in extremely noisy images [16], [15]. Surprisingly, discrete scan statistics were rigorously analyzed by means of random graph and percolation theories [10].

As a natural extension of this idea, several variations of scan statistics for graphs were proposed. Examples include efficient subwindow search for object detection [3], non-parametric scan statistics for event detection and forecasting in heterogeneous social media graphs [6], changepoint detection over graphs with the spectral scan statistic [26], anomaly detection in graphs [27], graph topic scan statistic for spatial event detection [20].

In the present paper, we extend the previous research in this area by constructing a fully nonparametric unsupervised learning graph scan-based procedure that is robust against an extremely malicious adversary. Moreover, we prove consistency of our results in a framework with minimal assumptions.

The outline of this paper is as follows. In Section 2, we introduce the framework for analyzing noisy or indirectly observed graphs with users polarized into two types of communities according to a binary property, and introduce our concept of an adversarial leader influencing communities of one of these types. In Section 3, we introduce the concept of k -NN graph scan estimators and define the main notions related to their construction. These estimators are one of the main ingredients for

robust consistent inference in this paper. In Section 4, we establish consistency and derive important properties of these unsupervised estimators. Proofs of all the results can be found in Appendix. Experimental results for real and synthetic graphs, as well as for the cases with and without adversaries can be found in Section 6. Scalability of the method is established and the algorithm’s distributed realization is discussed in Section 5. Conclusions are stated and future research directions are outlined in Section 7.

2. Formulation: Model and adversaries. Let $G_n = (V_n, E)$ be a graph with n vertices and $|E|$ edges. We do not impose restrictions on G_n such as sparsity, so $|E| \gg n$ is possible. For any vertex $v \in V$ we observe a real-valued random variable $X_v : \Omega \rightarrow \mathbb{R}$, defined on an appropriate probability space Ω .

We call the random variables X_v the *observed activities*. These observations are noisy realizations of the *true activity* A_v (of each vertex v), which we only observe with an additive (*nonparametric*) noise. More precisely,

$$X_v = A_v + \varepsilon_v, \quad (1)$$

where we assume that the noise $\{\varepsilon_v\}$ satisfies

$$\varepsilon_v \sim F, \quad \mathbb{E}[\varepsilon_v] = 0, \quad \text{Var}[\varepsilon_v] = \sigma^2 < +\infty. \quad (2)$$

We *do not* assume knowledge of either the distribution F or the variance σ^2 . Moreover, F need not have any particular parametric form, or be continuous, or be limited to being discrete with finite or countable support.

Given this model, our general aim is to robustly learn, without supervision, as much as possible about the collection of the true underlying activities $\{A_v\}_{v \in V}$ from noisy observations (1); and possibly, to also learn F nonparametrically. Moreover, we admit presence of a powerful adversary who can corrupt the graph to hinder inference (§2.1).

The specific problem solved in this paper can be formulated as unsupervised learning of the infinite-dimensional vector (a, b, F) , where a and b are baseline activities corresponding to inactive and active users correspondingly. They are fundamental global intrinsic properties of topical communities of users and are needed to be able to detect and recover active and inactive communities in subsequent analysis. An unknown probability distribution function F represents performance of the crawling algorithm for the particular topic.

Example 2.1. Consider a large-scale network where manual processing is not possible and the activities of users (nodes) are determined using algorithms that may either be operating on outdated data or have large classification errors. For instance, say in a social network the observed activity X_v measures the level of support of a statement by user v to a specified topic. It might be difficult for an algorithm to realize whether a statement “Topic X is crazy” is supportive of the topic or the opposite.

When estimating baseline activities within communities of a-priori unknown configurations and in the presence of nonparametric noise, it is important to assume that there exists an observability threshold. Thus, we assume that:

$$\text{for each } v \in V : \begin{cases} A_v \geq b > a, & \text{if } v \text{ is active;} \\ A_v = a, & \text{if } v \text{ is inactive.} \end{cases} \quad (3)$$

Vertices with activity level a are normal (inactive). Those with levels above the *unknown threshold* b are active.

This setup with two distinctly different types of communities is typical for applications where users have to be considered in terms of their opinion or preference on an important polarizing topic. Examples include political elections with two candidates, or important issues such as a healthcare reform, or good or bad health status of individuals in epidemiological applications and cells in medical images, etc.

We explicitly remark that we do assume two types of communities, but not the existence of two large communities covering the whole network. In fact, our setup includes the case when there is a large number of small communities of each type. The results of this paper allow strongly consistent inference for the cases when the number of communities of each type can be $O(n \log^{-1} n)$.

We specifically note here that the problem of estimating global baseline activity levels of communities is not the same as the problem of identifying active and inactive vertices in a network. For each particular vertex, the latter problem amounts to classification (or to clustering, in the unsupervised scenario), while the former problem is an infinite-dimensional learning problem. Moreover, as follows from results in this paper, global properties of hidden communities can be learned (in a strongly consistent way) without guessing activity levels of individual vertices. On contrary, it can be easily seen that, in the general setting of nonparametric noisy graphs, it is impossible to guarantee a strong form of consistency for uncovering true activity levels of individual vertices, and it is impossible to even consistently classify any individual vertex as an active or inactive one.

The lower bound for active vertex intensity b is assumed to be unknown, and we propose an unsupervised procedure for learning b as well, even though we do not analyze its performance in this paper. The difficulty of these estimation problems is that locations, shapes and exact sizes of clusters of active users are assumed to be unknown. The number of active users or active clusters is also unknown, and we make no probabilistic assumptions about this number or about the distribution of cluster locations. There can be anywhere between $O(\log n)$ and n inactive users, and between 1 and $O(n \log^{-1} n)$ inactive communities. The contribution of inactive (and active) users and communities can range from negligible to dominant.

It is important to note that we consider the case of a fully nonparametric noise of unknown level and having an unknown distribution; even within the setup with independent identically distributed and bounded noise, this model is far more general than traditional models with normally distributed errors and graphs with simple regular structure or parametric types of degree distributions.

2.1. Adversarial leader. Our model and inference algorithms permit presence of an adversary. We assume that *active* vertices may be under the influence of an *adversarial leader* who is capable of directly *altering their values* A_v , pursuing the goal of spoiling our inference. The adversary knows the true values on active vertices and which vertices are inactive, knows the full graph structure of G_n , and has unlimited computational power and memory. Moreover, the adversary has a special deal with *hackers* and can completely foresee the actions we will be performing on the network, including the outcomes of all our randomized procedures.

It is assumed that the adversary can only command its supporters, and therefore does not influence inactive vertices. This is a natural assumption that is met, for example, by political parties and their leaders, or by communities of bloggers under

the influence of a particular opinion leader. The leader of a political party can command his supporters in a variety of ways and also can gain access to hidden information about the leader’s supporters; meanwhile, all this information will be completely inaccessible for our inference. A general survey of various concepts and approaches to influence analysis can be found in [11].

It is also assumed that the adversary cannot completely tone down the active nodes, so that a separation between (active) supporters and (inactive) non-supporters is maintained and condition (3) still holds, even though the separation rule is not known. Moreover, the adversary does not know and cannot influence the outcomes of the crawling algorithm, so that $\{\varepsilon_v \mid v \in V_n\}$ is a completely independent collection of identically distributed random variables. These assumptions are met in many applications such as, for example, analysis of the blogosphere [2], as the crawler is forming its evaluations using the history of users’ activities; therefore, past activities of the user are accounted for and do not let the user completely mask his activity level.

Without much loss of generality, we assume that the adversary influences the active vertices only once, before we run our inference, and cannot act at later steps. This assumption is not restrictive and is used to have more compact statements of consistency theorems. Indeed, we are considering an effectively static network (even if it is a particular observation of a dynamic network or of a random network); as the adversary knows everything we would do on the network in advance, his single action can be a combined response incorporating series of individual responses to any sequence of our actions. We illustrate this point in our experiments in Section 6, where we allow the adversary that is as strong as above and can act in multiple (even unlimited number of) steps, and we are still able to demonstrate empirical consistency of our method.

Within this adversarial framework, we present below (Sec. 3) new k -nearest neighbors based graph scan estimators and establish sufficient conditions for them to be consistent. This consistency is a strong type of robustness, and a particular strength of our approach, as it can be easily seen that many estimators can be spoiled by such an adversary, as the following example suggests.

Indeed, suppose we are using an ingenious method that allows us to select a “nice” set of local subgraphs on which we run some consistent estimator. The adversary, knowing our method fully, can increase the values on the few active vertices contained in the selected local subgraphs, thus either skewing our averaging on those subgraphs, or knocking off good subgraphs in case model selection is involved.

Of course, all the consistency results of the present paper are still true in the neutral case when there is no adversary spoiling our inference.

3. k -NN graph scan estimators. The true activity level a corresponds to normal (inactive) vertices, while active vertices have the true activity level at least b . Since our estimators will be based on a scan by k -nearest-neighbors, we call these estimators k -NN *scan estimators*, or k -NN graph scan estimators. These estimators can be used for doing nonparametric statistics and unsupervised learning on the graph (network) G_n .

Our sole assumption on the graph structure and active and inactive communities can be formulated in a local form concerning only a negligibly small sub-community of inactive users.

Assumption 3.1. For our search of inactive users, assume that there is an inactive vertex $v \in G$ with a full $k(n)$ -neighborhood of inactive nearest neighbors, and that

$$\lim_{n \rightarrow \infty} \frac{k(n)}{\log n} = \infty. \quad (4)$$

Right below we clarify the terminology used in the Assumption. Let $v \in V$ be a vertex of G and let $m \in \mathbb{N}$ be any number. We define a *nested (multilevel) neighborhoods* of v in G as follows. First define $\mathfrak{N}_0(v) := v$, and then define recursively $\mathfrak{N}_1(v) := \{v_1 \in V \mid (v, v_1) \in E\}$. More generally, for any natural number i ,

$$\begin{aligned} \mathfrak{N}_i(v) := \{v_i \in V \mid & \text{there exists } v_{i-1} \in \mathfrak{N}_{i-1}(v) \\ & \text{such that } (v_{i-1}, v_i) \in E\}. \end{aligned}$$

For natural i , *descendance level sets of i -th order* for v in G , are defined for $i \geq 1$ as

$$\mathfrak{D}_i(v) := \mathfrak{N}_i(v) \setminus \mathfrak{N}_{i-1}(v). \quad (5)$$

For $i = 0$, we have set for convenience

$$\mathfrak{D}_0(v) := v. \quad (6)$$

For any $k \in \mathbb{N}$, a *(full) neighborhood of k -th order* of v in G is defined as the union

$$\Omega_k(v) := \bigcup_{0 \leq i \leq k} \mathfrak{N}_i(v). \quad (7)$$

Lemma 3.2 states a simple relation between \mathfrak{D} and Ω .

Lemma 3.2. Let \mathfrak{D}_i and Ω_k be defined as above. Then,

$$\mathfrak{D}_i(v) = \mathfrak{N}_i(v) \setminus \bigcup_{m=0}^{i-1} \mathfrak{N}_m(v), \quad (8)$$

and

$$\Omega_k(v) := \bigcup_{0 \leq i \leq k} \mathfrak{D}_i(v). \quad (9)$$

For any $m \in \mathbb{N}$, a *full m -neighborhood* $\mathfrak{B}_m^+(v)$ of v in G is $\Omega_r(v)$ such that $|\Omega_{r-1}(v)| < m$, but $|\Omega_r(v)| \geq m$.

For any $m \in \mathbb{N}$, an *exact m -neighborhood* $\mathfrak{B}_m(v)$ of v in G is any subset $\mathfrak{B}_m(v) \subseteq \Omega_r(v)$ such that $\Omega_{r-1}(v) \subseteq \mathfrak{B}_m(v)$ and $|\mathfrak{B}_m(v)| = m$.

In a graph $G = (V, E)$, for a set of vertices $K \subseteq V$, the total sum of values observed over K will be denoted as

$$S_K = \sum_{v \in K} X_v. \quad (10)$$

Definition 3.3 (*k -NN scan estimator*). Let \mathcal{K}_0 be any collection of exact $k(n)$ -neighborhoods of all vertices in G :

$$\mathcal{K}_0 := \{\mathfrak{B}_{k(n)}(v) \mid v \in V\}. \quad (11)$$

Set

$$\hat{K} := \arg \min_{K \in \mathcal{K}_0} \{S_K\} \quad (12)$$

and define a *(sublevel) k -NN scan estimator* as

$$\hat{a} := \frac{1}{k(n)} \sum_{v \in \hat{K}} X_v. \quad (13)$$

Algorithm 1 k -NN scan estimator

- 1: Phase 1 (decentralized). Pick one arbitrary $k(n)$ -neighborhood per vertex in V , and compute the average of X_v over this neighborhood.
 - 2: Phase 2 (collaborative). Identify a node with smallest average, over its neighborhood \hat{K} .
 - 3: Output $\hat{a} := \frac{1}{k(n)} \sum_{v \in \hat{K}} X_v$.
-

The k -NN scan estimator for inactive vertices can be computed via the following Algorithm 1.

Example 3.4. An important special case of graphs is given by lattices. Suppose we have a noisy two-dimensional pixelized image. We are interested in detection of objects that have an unknown color. This color has to be different from the colour of the background. It is also assumed that on each pixel we have random noise that has an unknown nonparametric distribution. This type of model is typical for cryo-electron microscopy. Typically, each cryo-EM picture contains a large number of particles. Particles have unknown, irregular, nonconvex and different shapes and sizes. The only common property for all cryo-EM pictures is that particles are darker than the background and that the noise has a completely unknown irregular distribution. However, both the background intensity and the particle intensity vary from image to image and are not known in advance.

It is possible to view digital images as networks, where individual pixels correspond to vertices. Since the initial noisy image can be naturally viewed as a square lattice graph, where k^2 -nearest neighbors correspond to a $k \times k$ subsquare on the screen, we see that a popular sliding window estimator is a special case of the k -NN scan estimator.

This paradigm was used to solve a number of nonparametric unsupervised learning problems in image analysis and cryo-EM applications (see [15], [16]). Many results on discrete spatial scan estimators from [15] and [16] are special cases of results of the present paper.

In view of the symmetry of the results between sub- and super-level cases, in this paper only the sublevel case is considered in details. For completeness, the super-level scan estimator and the estimator of the crawler's performance are described in the Appendix.

4. Consistency of k -NN scan estimators. In this section, we establish strong consistency of the proposed k -NN scan estimator under the assumption of bounded noise. Suppose, therefore that there is a constant $M > 0$ such that for all $v \in G_n$

$$|\varepsilon_v| \leq M \quad \text{almost surely.} \quad (14)$$

The bounded noise case is not the only case when the k -NN scan estimator is strongly consistent. We mainly use condition (14) to establish tight nonasymptotic performance guarantees for the estimator. Some form of asymptotic consistency can be established for unbounded noise from large nonparametric classes as well. This is illustrated in Section 6, where *all* of our experiments are performed for an unbounded noise.

Let $K \subseteq V_n$ be any collection of vertices. Denote by $S_1(K, n)$ the *number of active vertices* in K . We prove the following statement that provides the foundation of model selection on graph neighborhoods.

Proposition 1. *Let K_0 be any set of inactive vertices with $|K| = k(n)$, and let \mathcal{K} be any collection of exact $k(n)$ -neighborhoods. Define, for any $K \in \mathcal{K}$,*

$$\mathfrak{R}(K, a, b, \{A_v\}) := (b - a) S_1(K, n) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b).$$

Then, we have the following bound:

$$P(\widehat{K} \in \mathcal{K}, S_{\widehat{K}} < S_{K_0}) \leq \sum_{K \in \mathcal{K}} \exp\left(-\frac{3 [\mathfrak{R}(K, a, b, \{A_v\})]^2}{12 \sigma^2 |K \setminus K_0| + 4 M \cdot \mathfrak{R}(K, a, b, \{A_v\})}\right). \tag{15}$$

Notice that the statement of Proposition 1 is actually *non-asymptotic*: even though the bound depends on n , it is valid for small n as well. Additionally, the bound (15) depends on the subgraph K , and automatically tightens for those subgraphs that have more active vertices. This can be used to show that (15) is essentially tight, with or without the adversary, also in finite-sample, non-asymptotic cases.

The bound in Proposition 1 is based on the first Bernstein inequality [5], which substantially improves on the Hoeffding inequality in case if $\text{Var } \varepsilon_v \ll M^2$, which is the case for many distributions of interest. Additionally, this makes the bound (15) extensible to weakly-dependent random variables as well, unlike the corresponding bound that relies on the Hoeffding inequality.

The following is the master theorem governing consistency of graph scan estimators. We understand consistency in its *strong classical form*: as the sample size increases, the algorithm learns the true value of the parameter correctly with probability approaching 1.

Theorem 4.1. *Let $\mathcal{K}_{\mathcal{A}}(n)$ be a collection of exact $k(n)$ -neighborhoods that a k -NN graph scan algorithm \mathcal{A} uses for estimation. Suppose $K_0 \in \mathcal{K}_{\mathcal{A}}(n)$ and*

$$\lim_{n \rightarrow \infty} \frac{k(n)}{\log |\mathcal{K}_{\mathcal{A}}(n)|} = \infty, \tag{16}$$

$$\lim_{n \rightarrow \infty} |\mathcal{K}_{\mathcal{A}}(n)| \exp(-n) = 0. \tag{17}$$

Then the scanning algorithm \mathcal{A} leads to the graph scan estimator $\widehat{a} = \overline{X}_{\widehat{K}}$ that is a consistent estimator of a , regardless of the adversary’s strategy.

As a corollary of this general statement, we derive consistency of the estimator of the present paper. Proofs of all the results can be found in Appendix.

Theorem 4.2. *Suppose that $k(n)$ satisfies Assumption 3.1. Then, the graph scan estimator $\widehat{a} = \overline{X}_{\widehat{K}}$ is a consistent estimator of a , regardless of the adversary’s strategy.*

5. Scalable algorithms and computation. In view of the huge size of many networks of current interest, and also due to the fact that we typically do not get the chance to observe the network in its entirety but rather can only query it in small parts, we need fast and scalable algorithms to implement our estimators. Moreover, these algorithms should be provably statistically reliable and consistent. In this section we discuss some scalability aspects of our estimators; however, to

avoid detracting from our primary modeling focus, we defer a thorough empirical evaluation to the future.

Distributed computation. Note that the computation of our proposed estimator shown in Algorithm 1 is fully decentralized, breaking up into computation of one neighborhood per graph node (Phase 1). However, there is a necessary centralized communication phase at the end that collects all averages (a single number) per neighborhood and only keeps the smallest one out of those (Phase 2). This second phase is easily implemented as a typical map-reduce operation over the nodes, the ‘reduce’ operation being the min operator.

Proposition 2. *Let $k(n)$ be the neighborhood size. One can compute a k -NN scan estimator using $\mathcal{O}(k(n))$ operations per node in the decentralized phase (Phase 1) of Algorithm 1, followed by a single communication round to compute the minimum over the graph.*

Parallelization. The main computational cost is encountered in Phase 1 of Algorithm 1. Fortunately, the computations in this phase decouple completely, enabling our algorithm to be directly applicable to huge networks. In this setting, we assume the nodes of the graph are partitioned over a set of compute agents, each responsible only for the estimators over its own nodes. In the extreme case, in an IoT setting as noted below, every graph node has its own compute agent responsible for it. All independently execute the same BFS style algorithm to build an average over the $k(n)$ sized neighborhood per node (naively needing the communication of $k(n)$ numbers back to node v).

In an IoT framework, we naturally have a processor built-in at each graph node. We therefore readily use those processors (that have immediate access to local observations, and communicate in small neighborhoods) for parallel estimation using Algorithm 1.

In general, it could be productive to combine the scanning method of the present paper with modified algorithms for efficient subwindow search (see [3]).

6. Experimental results.

6.1. Simulated large networks, no adversary. To illustrate performance of our method on a larger scale, we constructed an artificial network that consists of two subgraphs - one of which is made of a million nodes where each of these nodes is connected to 3 nodes chosen uniformly at random within this subgraph. The other subgraph consists of a thousand nodes and, in a similar way, each of these nodes is connected to 3 nodes chosen uniformly at random within this second subgraph. All the 1000 nodes in the smaller subgraph are labeled as inactive, while nodes in the bigger subgraph are labeled as active or inactive randomly with equal probability. This way, we guaranteed that there is an inactive community of size $k(n) = 1000$. Furthermore, we pick 20 random pairs of nodes - one from the bigger subgraph and the other from the smaller one and connect the two. Number of nodes in the graph is 1,001,000 and number of edges is up to 3,003,002.

Activity levels in the network are set to $a = 2$ for the inactive nodes and $A_v = 10$ for the active ones. These weights are corrupted by adding a Gaussian noise of mean 0 and variance $\sigma^2 = 1$ before feeding them to the algorithm.

The estimates returned for the above example using our proposed approach are in accordance with the theory. The algorithm underestimates a when we pick the value of the parameter k that is too small (this value is corresponding to the size

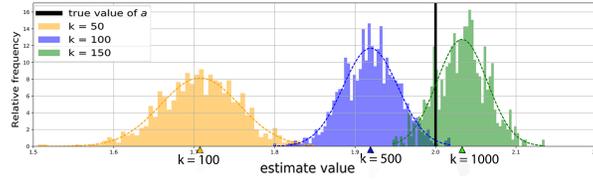


FIGURE 1. Graph scan estimators for the artificial large graph.

of the neighborhood). The estimate’s value increases with an increased value of k . Halfway between the two extreme values we obtain very accurate estimates. Naturally, these values of k that return the optimal estimate ($k = 500$ or $k = 1000$ in this case) can also be considered suggestive of the inactive community sizes inside the network.

If we change the size of the smaller group to 10,000, we observe that the $k = 1000$ setting of the algorithm returns an estimate of $\hat{a} = 1.75$ on average over 400 experiments, implying that we have to scan with bigger neighborhoods to achieve consistent inference. Indeed, in this case $k(n) = 10,000$.

The histogram in Figure 1 illustrates the distribution of the k -NN graph scan estimators for different values of k . In particular, for $k = 500$ the empirical mean of the estimator is 1.91914586 with variance 0.0341272087, while for $k = 1000$ we have estimated mean 2.03303998 with variance 0.031305042. The true value $a = 2$ is, indeed, very close.

6.2. Real network structure, no adversary. In this example, we use the real community structure of political blogs from the 2004 U.S. Presidential Election. Paper [2] studied the linking patterns of political blogs. These blogs can be naturally classified into two classes, the liberal ones and the conservative ones. This is suitable for our framework. The number of libertarian, independent, or moderate blogs was negligible at the time.

In [2], a description of the network of over 1000 blogs is presented, based on a single day snapshot that included blogrolls. Two blogs share an edge if one of the two cites the other one. The blogs were categorized manually. It turned out that neither directory labels relying on self-reported or automated categorizations, nor the manual labels were 100% accurate, with an error probability that has an unknown distribution.

There were 1494 blogs in total, with 759 liberal and 735 conservative. The structure of the underlying linking graph was rather complex, as 91% of the links originating within either the conservative or liberal communities stay within that community, but the number of intercommunity links was non-negligible as well. Figure 2 illustrates the structure of the political blogging graph.

To create a challenge to our method, we add an additional complication to this dataset by corrupting the values attributed to blogs by Gaussian white noise. We assigned the same activity levels as in Section 6.1. Notice that this type of noise is unbounded and so does not satisfy conditions of the consistency theorem. The graph size is also relatively small, so we are far from the asymptotic regime. However, the graph scan estimator produced surprisingly accurate results.

The histogram in Figure 3 illustrates the distribution of the k -NN graph scan estimators for different values of k . It is apparent that in this graph there is an

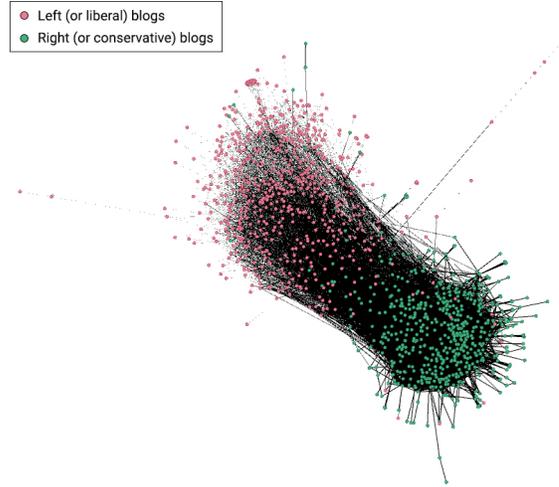


FIGURE 2. Political blogosphere graph for the 2004 Elections.

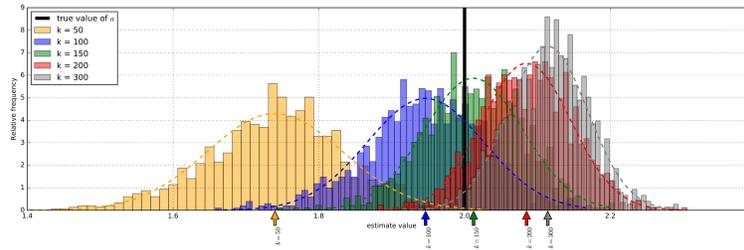


FIGURE 3. Graph scan estimators for the 2004 Elections graph.

inactive neighborhood of size close to 150, but not more. An estimator for $k = 150$ is surprisingly accurate.

6.3. Simulated large network, two types of adversaries. In order to study the impact of an adversary on the performance of the proposed algorithm, the following experiment is designed.

Let us consider the following two scenarios. Scenario 1: the adversary decides to act locally and chooses to influence those active nodes that are present in the best neighborhood \hat{K} of our choice, thus hoping to throw off our specific Algorithm.

Scenario 2: the adversary is strong and decides to use brute force approach, influencing all the active nodes in the network.

For our experiment, we assumed that when the adversary ‘influences’ a node, the activity level for this node is changed to a large number 10^6 . We then run our algorithm on the new set of altered activity levels and generate a new estimate $\hat{a}^{(2)}$ corresponding to a new ‘best neighborhood’ $\hat{K}^{(2)}$.

Over 100 such experiments for the value of $k = 500$, there happen to be only 23 cases where \hat{K} obtained after first running our algorithm contains at least one active node. Hence for the remaining 77 cases, the estimate $\hat{a}^{(2)}$ obtained after running the algorithm over the altered values by either adversary - Scenario 1 or Scenario 2, everything remains the same for the Algorithm since the new best neighborhood

$\hat{K}^{(2)}$ would remain the same as the one obtained previously (\hat{K}). For the 23 times where we have an active node in \hat{K} , the average of the estimate \hat{a} when the algorithm is first run has a mean of 1.9320 and a standard deviation of 0.0339. When we run the algorithm after the action of the weak adversary, the estimate $\hat{a}^{(2)}$ has a mean of 1.9399 and a standard deviation of 0.0333. When the adversary is a strong one, our estimate $\hat{a}^{(2)}$ has a mean of 1.9423 and a standard deviation of 0.0322.

However for the value of $k = 1000$, we reported 97 cases where the neighborhood \hat{K} contains one or more active nodes and for the remaining 3 we have no active nodes within \hat{K} . The mean of the estimates \hat{a} when the algorithm runs for the first time was 2.0293 with a standard deviation of 0.0318, while after the influence of the weak adversary the estimates $\hat{a}^{(2)}$ had a mean of 2.0472 and a standard deviation of 0.0438 and in presence of the strong adversary, the estimates had a mean of 1.9220 and a standard deviation of 0.0354.

This shows that our method exhibits remarkable stability against both crafty and brute force adversaries, even under conditions that are more general than the ones in our consistency theorem.

6.4. Experiments with simulated large networks and multi-step adversaries. In this section, we show that our method can be robust even in those cases where the adversary can act not once, but as many times as he chooses. We build a multi-step multi-stage experiment, where at each step, firstly, we generate an estimate by choosing an optimal neighborhood with our scanning algorithm and, secondly, the adversary acts to spoil our inference by influencing active nodes in the neighborhood that we just selected. This game can be repeated as many times as the adversary chooses. Notice that this is a much more general setup than we permitted in Section 2.1.

For a set of values for $k(n)$ (the size of the scanning neighborhood), we perform 100 experiments for each case. We let the adversary play as long as it made sense, and it happens that we only needed up to 10 steps. In all these cases, we always won the game, by reaching a stage where we select a neighborhood that doesn't contain any active node; therefore, after this step the algorithm is immune to adversary's influence.

The table 1 shows the estimates returned in these experiments, and the number of steps N_w required before winning against the adversary. When $N_w = 0$, the neighborhood chosen by our algorithm in the very first attempt (before adversary's influence) does not contain any active nodes and thus we win at the 0-th step. It can be seen that as k increases, the value of the estimate improves, getting close to the true value of $a = 2$ for the optimal range of values $k = 500$ to $k = 800$, but then we often need more steps to beat the adversary. Thus choosing the parameter k can be considered as facing a *trade-off* between the accuracy of the estimate and the robustness against an adversary.

7. Conclusions and future research. We developed a robust and scalable method to uncover fundamental properties of communities in very large noisy networks, in the situation where the nodes or users are split into two classes according to a binary property, e.g., according to their opinions or preferences on a topic. In our setup, one of the communities can fall under the influence of a knowledgeable adversarial leader, who knows the full network structure, has unlimited computational resources and can completely foresee our planned actions on the network.

k	100	200	500	700	800
$\#(N_w = 0)$	98	98	78	56	55
$\#(N_w = 1)$	2	2	16	22	17
$\#(N_w = 2)$	0	0	3	16	12
$\#(N_w = 3)$	0	0	3	5	9
$\#(N_w \geq 4)$	0	0	0	1	7
$\mu_{\hat{a}}$	1.705	1.815	1.913	1.949	1.961
$\max \hat{a}$	1.806	1.903	1.980	2.030	2.031
$\min \hat{a}$	1.603	1.686	1.824	1.848	1.879
$\sigma_{\hat{a}}$	0.048	0.040	0.033	0.034	0.029

TABLE 1. Multi-step adversary

The adversary is pursuing the goal of spoiling our uncovering of community and network properties.

We proposed a nonparametric, unsupervised, and scalable graph scan procedure that is, in addition, robust against a class of such powerful adversaries. The k -NN scans presented here generalize sliding windows, moving averages, and scan statistics, extending them to the case of general graphs. We proved strong consistency of our results in this nonparametric setup with minimal model or structural assumptions. In particular, the learning procedure estimates the baseline activity of normal users asymptotically correctly with probability 1; the only assumption being the existence of a single implicit community of asymptotically negligible logarithmic size. We provided experiments on real and synthetic data and illustrated the performance of our method, including examples with several types of adversaries.

In addition, activity observations of users are delivered by a third-party algorithm called the crawler. The crawler has its classification and learning errors that are not known to us. Learning the actual performance of the crawler can also be addressed as a by-product of results of this paper.

Baseline activity levels of communities, as well as the crawler’s accuracy, are fundamental global intrinsic properties of topical communities of users and are needed to be able to detect and recover active and inactive communities in subsequent analysis. Therefore, we believe that results of the present paper are foundational to fast consistent algorithms for community detection via percolation on general graphs.

We devised a highly decentralized implementation of our scans, allowing scalable distributed and parallel processing of large scale graphs. However, we deferred a thorough empirical evaluation of demonstrated performance times to the future. In general, it could be productive to combine the scanning method of the present paper with modified algorithms for efficient subwindow search.

Acknowledgments. The authors would like to thank Suvrit Sra for helpful discussions, and the anonymous referee for valuable suggestions that helped to improve the presentation of the paper.

Appendix.

8. Graph scan estimators for the crawler and for active vertices. In view of the symmetry of the results between sub- and super-level cases, in the main body of the paper only the sublevel case was considered in details. For completeness,

the superlevel scan estimator and the estimator of the crawler's performance are described below.

Additionally, we remark here that it is possible to define scan estimators (of each of the three global properties discussed in this paper) that use either *full* $k(n)$ -neighborhoods of all vertices, or *all* exact $k(n)$ -neighborhoods. These estimators are consistent under rather general assumptions as well, even though the corresponding limiting distributions differ from those of the scan estimator of this paper. However, note that the adversary assumption might need to be altered to ensure consistency for each of these modified estimators; we do not discuss the details of this topic in the paper.

8.1. Superlevel graph scans for active vertices. For inference on active clusters we would have to assume that there exists at least one active cluster that contains a full $\varphi(n)$ -neighborhood of vertices, and that

$$\lim_{n \rightarrow \infty} \frac{\varphi(n)}{\log n} = \infty. \quad (18)$$

Construction of superlevel k -NN scan estimators for the active level (b in our notation) requires the following dual assumption (*cf.* (3)):

$$\text{for each } v \in V, \begin{cases} A_v = b > a, & \text{if } v \text{ is active;} \\ A_v \leq a, & \text{if } v \text{ is normal.} \end{cases} \quad (19)$$

Now the above graph scan algorithm can be inverted to get a k -NN scan estimator for active vertices instead. The only modification would be that this time

$$\widehat{K} := \arg \max_{K(v) \subseteq V: |K(v)| = \varphi(n), v \in V} \{S_{K(v)}\} \quad (20)$$

and

$$\widehat{b} := \varphi(n)^{-1} \sum_{v \in \widehat{K}} Y_v. \quad (21)$$

In view of the symmetry of the results between the sublevel and the superlevel cases, in this paper only the sublevel case is considered in details.

8.2. Graph scans for learning the crawler's performance. Our framework involves a third-party network crawling algorithm that provides us the data about the users of the network. We propose here a modified graph scan estimator to estimate the unknown distribution F of the misclassification rate of the crawling algorithm. The difficulty here is again that we do not know which vertices are active or inactive. However, this problem is solved by a combination of the graph scan estimator with the empirical distribution function estimator.

Let \widehat{K} be defined by (12) of Definition 3.3. A *graph scan estimator for the crawler's error* distribution F would be

$$\widehat{F}(t) := \frac{1}{|\widehat{K}|} \sum_{v \in \widehat{K}} \mathbb{1}_{\{X_v \leq t\}}. \quad (22)$$

The following variation gives an *unbiased* consistent estimator for the noise variance σ^2 :

$$\widehat{\sigma}^2 := \frac{1}{|\widehat{K}| - 1} \sum_{v \in \widehat{K}} (X_v - \widehat{a})^2. \quad (23)$$

Both scan estimators $\widehat{\sigma}^2$ and $\widehat{F}(t)$ can be easily calculated once \widehat{a} is calculated. We conjecture that these are consistent estimators of σ^2 and F , respectively. Moreover, there is evidence that \widehat{F} is uniformly $\sqrt{k(n)}$ -consistent in $\|\cdot\|_\infty$ -norm.

9. Proofs. In this section, we start with deriving probabilistic bounds for, and statistical properties of, sublevel graph scan estimators for the case of general unbounded noise of finite variance. Corresponding results for superlevel scan estimators are analogous.

Notice that in the model formulated in Section 2, potentially, $A_v = A_v(\{X_v | v \in V_n\}, G_n, \mathcal{A}(G_n))$ for all active v , where \mathcal{A} is the algorithm we use for inference, and $\mathcal{A}(G_n)$ is the collection of all the steps we would perform together with values of all the random variables that we will generate. Unlike the case of inactive vertices, this is not a collection of numbers, but rather a collection of random variables with a complex and unknown mutual dependence structure. Therefore, $\{X_v | v \in V_n\}$ forms a collection of random variables too. We do not require this collection to be neither completely, nor pairwise independent, nor having identical marginal distributions. In fact, we allow for strong and long-distance correlations (within bounds given by (3)) between observations on active vertices.

We will use a pairing trick for random variables in the proofs below. The trick is based on the following basic lemma.

Lemma 9.1. *Let $\{\varepsilon_v\}_{i=1}^{i=2r}$, for a natural number $r \geq 1$ be a collection of independent identically distributed random variables satisfying (2). Define for $1 \leq i \leq r$*

$$\Delta_i = \varepsilon_i - \varepsilon_{i+r}. \quad (24)$$

*The collection of random variables $\Delta_1, \Delta_2, \dots, \Delta_r$ is completely independent. For all i , it holds that $\mathbb{E} \Delta_i = 0$, $\text{Var} \Delta_i = 2\sigma^2$, and Δ_i is symmetric in distribution around 0. Moreover, $\Delta_i \sim_d F(\cdot) * (1 - F(-\cdot))$.*

Proof. (of Lemma 9.1): The first four assertions are trivial. The last one follows since $\varepsilon_i \sim_d F(\cdot)$ and $-\varepsilon_i \sim_d 1 - F(-\cdot)$, where \sim_d denotes equality in distribution. \square

The following simple geometric lemma would be useful.

Lemma 9.2. *Let $K_0 \subseteq V_n$ be any collection of inactive vertices, with $K_0 = k(n)$, and let $K \subseteq V_n$ be a collection of arbitrary vertices, with $K = k(n)$. Denote by $S_1(K, n)$ a number of active vertices in K , i.e.*

$$S_1(K, n) = |\{v \in K | v \text{ active}\}|. \quad (25)$$

Then $|K \setminus K_0| = |K_0 \setminus K| \geq S_1(K, n)$.

Proof. (of Lemma 9.2): The equality is obvious since $|K| = |K_0|$. The inequality follows since all the $S_1(K, n)$ of active vertices of K must be contained in $K \setminus K_0$ anyways. \square

The following proposition provides the main proxy for application of concentration of measure inequalities to graph scan estimators and to model selection on subgraphs. In essence, the proof is based on the pairing trick on the algebra of random variables.

Proposition 3. *Assume $\{X_v | v \in V_n\}$ satisfies the additive noise model given by (1) and (3), and let $\{\varepsilon_v\}$ be an arbitrary collection of independent identically distributed random variables with mean 0. Let $K_0 \subseteq V_n$ be any collection of inactive vertices, with $K_0 = k(n)$, and let $K \subseteq V_n$ be a collection of arbitrary vertices, with*

$K = k(n)$. Denote by $S_1(K, n)$ a number of active vertices in K . Then, for any $\alpha = \alpha(n) \in \mathbb{R}$

$$P(S_{\widehat{K}} < S_{K_0} - \alpha(n)) = \quad (26)$$

$$\begin{aligned} & P\left(\frac{1}{k(n)}\left(S_1(K, n)(b-a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b) + \alpha(n)\right)\right. \\ & \quad \left. < \frac{1}{k(n)} \sum_{i=1}^{|K \setminus K_0|} \Delta_i\right), \end{aligned}$$

where $\{\Delta_i\}$ is an arbitrary collection of independent identically distributed random variables with $\Delta_i \sim_d F(\cdot) * (1 - F(\cdot))$.

Proof.

$$\begin{aligned} & P(S_{\widehat{K}} < S_{K_0} - \alpha(n)) = \quad (27) \\ & = P\left(\sum_{v \in K \setminus \{v \in K \mid v \text{ active}\}} X_v + \sum_{\{v \in K \mid v \text{ active}\}} X_v < \sum_{v \in K_0} X_v - \alpha(n)\right) \\ & = P\left(\sum_{v \in K \setminus \{v \in K \mid v \text{ active}\}} (a + \varepsilon_v) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v + \varepsilon_v)\right. \\ & \quad \left. < \sum_{v \in K_0} (a + \varepsilon_v) - \alpha(n)\right) \\ & = P\left(\sum_{v \in K \setminus \{v \in K \mid v \text{ active}\}} (a + \varepsilon_v)\right. \\ & \quad \left. + \sum_{\{v \in K \mid v \text{ active}\}} ((A_v - b) + (b - a) + (a + \varepsilon_v)) < \sum_{v \in K_0} (a + \varepsilon_v) - \alpha(n)\right) \\ & = P\left(\sum_{v \in K} (a + \varepsilon_v) + (b - a) S_1(K, n) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b)\right. \\ & \quad \left. < \sum_{v \in K_0} (a + \varepsilon_v) - \alpha(n)\right) \\ & = P\left(\sum_{v \in K} \varepsilon_v + (b - a) S_1(K, n) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b) < \sum_{v \in K_0} \varepsilon_v - \alpha(n)\right) \\ & = P\left((b - a) S_1(K, n) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b) + \alpha(n) < \sum_{v \in K_0} \varepsilon_v - \sum_{v \in K} \varepsilon_v\right) \\ & = P\left((b - a) S_1(K, n) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b) + \alpha(n)\right. \\ & \quad \left. < \sum_{v \in K_0 \cap K} 0 + \sum_{v \in K_0 \setminus K} \varepsilon_v - \sum_{v \in K \setminus K_0} \varepsilon_v\right) \\ & = P\left(\frac{S_1(K, n)(b - a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b) + \alpha(n)}{k(n)}\right) \end{aligned}$$

$$< \frac{1}{k(n)} \sum_{v \in K_0 \setminus K} \varepsilon_v - \frac{1}{k(n)} \sum_{v \in K \setminus K_0} \varepsilon_v \Bigg).$$

Now we perform the following pairing trick. Let us arbitrarily enumerate all pixels in $K \setminus K_0$ by numbers from 1 to $|K \setminus K_0|$. We also perform an arbitrary enumeration on $K_0 \setminus K$. Now in the last equation of the chain (27) we can regroup the difference of two sums into a single sum that consists of differences of pairs of random variables, and each pair is formed from the two noise variables with the same number. We would have $|K \setminus K_0|$ of those pairs. Call those random differences $\Delta_1, \Delta_2, \dots, \Delta_{|K \setminus K_0|}$.

Applying Lemma 9.1 with $r = |K \setminus K_0|$, we continue to transform the above equations:

$$\begin{aligned} &= P \left(\frac{S_1(K, n)(b-a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b) + \alpha(n)}{k(n)} \right. \\ &\quad \left. < \frac{1}{k(n)} \sum_{v \in K_0 \setminus K} \varepsilon_v - \frac{1}{k(n)} \sum_{v \in K \setminus K_0} \varepsilon_v \right) \\ &= P \left(\frac{1}{k(n)} \left(S_1(K, n)(b-a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b) + \alpha(n) \right) \right. \\ &\quad \left. < \frac{1}{k(n)} \sum_{i=1}^{|K \setminus K_0|} \Delta_i \right). \end{aligned} \tag{28}$$

□

In the sequel, we shall use the following general form of the Bernstein inequality (see [5]).

Proposition 4. *Let X_1, \dots, X_n be independent zero-mean random variables. Suppose that $|X_i| \leq M$ almost surely, for all i . Then for all $t > 0$*

$$P \left(\sum_{i=1}^n X_i > t \right) \leq \exp \left(- \frac{t^2/2}{\sum_{i=1}^n \mathbb{E} X_i^2 + Mt/3} \right). \tag{29}$$

We are ready to prove the basic cornerstone probabilistic inequality of this paper.

Proposition 5. *Let K_0 be any set of inactive vertices with $|K| = k(n)$, and let K be any collection of $k(n)$ vertices. Define*

$$\mathfrak{R}(K, a, b, \{Im\}) := (b-a) S_1(K, n) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b).$$

Then, if assumptions of Proposition 3 are satisfied and (14) holds,

$$P(\widehat{K} = K, S_{\widehat{K}} < S_{K_0}) \leq \exp \left(- \frac{3 [\mathfrak{R}(K, a, b, \{Im\})]^2}{12 \sigma^2 |K \setminus K_0| + 4 M \cdot \mathfrak{R}(K, a, b, \{Im\})} \right).$$

Proof. (of Proposition 5): Notice that under the assumption (14) we have $|\Delta_i| \leq 2M$ and $\mathbb{E} \Delta_i = 0$. Setting $\alpha(n) = 0$ for all n and applying Proposition 3, Lemma 9.1, and the Bernstein inequality in the form given by Proposition 4, we derive

$$\begin{aligned}
P(\widehat{K} = K, S_{\widehat{K}} < S_{K_0}) &\leq P(S_{\widehat{K}} < S_{K_0}) = & (30) \\
&= P\left(\frac{1}{k(n)}\left(S_1(K, n)(b-a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b)\right) < \frac{1}{k(n)} \sum_{i=1}^{|K \setminus K_0|} \Delta_i\right) \\
&= \exp\left(-\frac{\left(S_1(K, n)(b-a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b)\right)^2}{2}\right) \\
&\quad : \left(\sum_{i=1}^{|K \setminus K_0|} \mathbb{E} \Delta_i^2 + \frac{2M\left(S_1(K, n)(b-a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b)\right)}{3}\right) \\
&\exp\left(-\frac{\left(S_1(K, n)(b-a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b)\right)^2}{2}\right) \\
&\quad : \left(\sum_{i=1}^{|K \setminus K_0|} \text{Var} \Delta_i + \frac{2M\left(S_1(K, n)(b-a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b)\right)}{3}\right) \\
&= \exp\left(-\frac{\left(S_1(K, n)(b-a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b)\right)^2}{2}\right) \\
&\quad : \left(|K \setminus K_0| \cdot 2\sigma^2 + \frac{2M\left(S_1(K, n)(b-a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b)\right)}{3}\right) \\
&= \exp\left(\frac{-3\left(S_1(K, n)(b-a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b)\right)^2}{12\sigma^2 |K \setminus K_0| + 4M \cdot \left(S_1(K, n)(b-a) + \sum_{\{v \in K \mid v \text{ active}\}} (A_v - b)\right)}\right).
\end{aligned}$$

□

Proof. (Proposition 1) Follows from Proposition 5 by bounding the probability of the union of individual model selection events. □

Proof. (Theorems 4.1 and 4.2) Applying Propositions 3, 1, 5 and 4 for $\alpha(n) = \delta k(n)$ and M replaced by $2M$ directly leads to the consistency statement of Theorem 4.1. Noticing that in the statement of Theorem 4.2 it holds that $|\mathcal{K}_{\mathcal{A}}(n)| = n$, we derive Theorem 4.2 from Theorem 4.1. □

REFERENCES

- [1] F. M. Aarestrup and M. G. Koopmans, [Sharing data for global infectious disease surveillance and outbreak detection](#), *Trends in Microbiology*, **24** (2016), 241–245.
- [2] L. A. Adamic and N. Glance, The political blogosphere and the 2004 us election: Divided they blog, in *Proceedings of the 3rd International Workshop on Link Discovery*, ACM, 2005, 36–43.
- [3] S. An, P. Peursum, W. Liu and S. Venkatesh, Efficient algorithms for subwindow search in object detection and localization, in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, 2009, 264–271.
- [4] E. Arias-Castro and G. Grimmett, [Cluster detection in networks using percolation](#), *Bernoulli*, **19** (2013), 676–719.
- [5] S. N. Bernstein, On certain modifications of chebyshev’s inequality, *Doklady Akademii Nauk SSSR*, **17** (1937), 275–277.

- [6] F. Chen and D. B. Neill, [Non-parametric scan statistics for event detection and forecasting in heterogeneous social media graphs](#), in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2014, 1166–1175.
- [7] V. Cheplygina, M. de Bruijne and J. P. Pluim, [Not-so-supervised: A survey of semi-supervised, multi-instance, and transfer learning in medical image analysis](#), *Medical Image Analysis*, **54** (2019), 280–296.
- [8] C. Deng, Y. Han and B. Zhao, [High-performance visual tracking with extreme learning machine framework](#), *IEEE Transactions on Cybernetics (Early Access)*, 2019, 1–12.
- [9] D. Geman and B. Jedynak, [An active testing model for tracking roads in satellite images](#), *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **18** (1996), 1–14.
- [10] G. Grimmett, [Percolation](#), vol. 321 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, 2nd edition, Springer-Verlag, Berlin, 1999.
- [11] M. Han and Y. Li, [Influence analysis: A survey of the state-of-the-art](#), *Mathematical Foundations of Computing*, **1** (2018), 201–253.
- [12] M. Kulldorff, [A spatial scan statistic](#), *Communications in Statistics-Theory and Methods*, **26** (1997), 1481–1496.
- [13] M. Kulldorff, [Spatial scan statistics: Models, calculations, and applications](#), in *Scan Statistics and Applications*, Springer, 1999, 303–322.
- [14] M. Kulldorff, L. Huang, L. Pickle and L. Duczmal, [An elliptic spatial scan statistic](#), *Statistics in medicine*, **25** (2006), 3929–3943.
- [15] M. Langovoy, M. Habeck and B. Schoelkopf, [Adaptive nonparametric detection in cryo-electron microscopy](#), in *Proceedings of the 58-th World Statistical Congress, Session: High Dimensional Data*, 2011, 4456–4461.
- [16] M. Langovoy, M. Habeck and B. Schoelkopf, [Spatial statistics, image analysis and percolation theory](#), in *The Joint Statistical Meetings Proceedings, Time Series and Network Section*, American Statistical Association, Alexandria, VA, 2011, 5571–5581.
- [17] M. Langovoy and O. Wittich, [Detection of Objects in Noisy Images and Site Percolation on Square Lattices](#), EURANDOM Report No. 2009-035, EURANDOM, Eindhoven, 2009.
- [18] M. Langovoy and O. Wittich, [Randomized algorithms for statistical image analysis and site percolation on square lattices](#), *Statistica Neerlandica*, **67** (2013), 337–353.
- [19] M. Langovoy and O. Wittich, [Robust nonparametric detection of objects in noisy images](#), *Journal of Nonparametric Statistics*, **25** (2013), 409–426.
- [20] Y. Liu, B. Zhou, F. Chen and D. W. Cheung, [Graph topic scan statistic for spatial event detection](#), in *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, 2016, 489–498.
- [21] T. McInerney and D. Terzopoulos, [Deformable models in medical image analysis: A survey](#), *Medical Image Analysis*, **1** (1996), 91–108.
- [22] D. B. Neill, A. W. Moore and G. F. Cooper, [A bayesian spatial scan statistic](#), *Advances in Neural Information Processing Systems (NIPS)*, **18** (2006), 1003.
- [23] G. P. Patil and C. Taillie, [Upper level set scan statistic for detecting arbitrarily shaped hotspots](#), *Environmental and Ecological statistics*, **11** (2004), 183–197.
- [24] L. D. Rotz and J. M. Hughes, [Advances in detecting and responding to threats from bioterrorism and emerging infectious disease](#), *Nature Medicine*, **10** (2004), S130–S136.
- [25] Y. Ruan, D. Fuhry and S. Parthasarathy, [Efficient community detection in large networks using content and links](#), in *Proceedings of the 22nd International Conference on World Wide Web*, ACM, 2013, 1089–1098.
- [26] J. Sharpnack, A. Singh and A. Rinaldo, [Change-point detection over graphs with the spectral scan statistic](#), in *AISTATS*, **13** (2013), 545–553.
- [27] J. L. Sharpnack, A. Krishnamurthy and A. Singh, [Near-optimal anomaly detection in graphs using Lovasz extended scan statistic](#), in *Advances in Neural Information Processing Systems (NIPS)*, 2013, 1959–1967.
- [28] T. Tango and K. Takahashi, [A flexibly shaped spatial scan statistic for detecting clusters](#), *International Journal of Health Geographics*, **4** (2005), 11.
- [29] J. V. Uspensky, *Introduction to Mathematical Probability*, McGraw-Hill, 1937.
- [30] Y. Wang, Y. Feng, J. Luo and X. Zhang, [Voting with feet: Who are leaving hillary clinton and donald trump](#), in *2016 IEEE International Symposium on Multimedia (ISM)*, IEEE, 2016, 71–76.

- [31] J. Yang, J. McAuley and J. Leskovec, [Community detection in networks with node attributes](#), in *Data Mining (ICDM), 2013 IEEE 13th International Conference on*, IEEE, 2013, 1151–1156.

Received for publication April 2019.

E-mail address: mikhail@langovoy.com

E-mail address: akhilesh.gotmare@alumni.epfl.ch

E-mail address: martin.jaggi@epfl.ch