

A Multiscale Scan Statistic for Adaptive Submatrix Localization

Yuchao Liu
Microsoft Corporation
Redmond, Washington
yuchao.liu@microsoft.com

Ery Arias-Castro
University of California, San Diego
La Jolla, California
eariasca@ucsd.edu

ABSTRACT

We consider the problem of localizing a submatrix with larger-than-usual entry values inside a data matrix, without the prior knowledge of the submatrix size. We establish an optimization framework based on a multiscale scan statistic, and develop algorithms in order to approach the optimizer. We also show that our estimator only requires a signal strength of the same order as the minimax estimator with oracle knowledge of the submatrix size, to exactly recover the anomaly with high probability. We perform some simulations that show that our estimator has superior performance compared to other estimators which do not require prior submatrix knowledge, while being comparatively faster to compute.

CCS CONCEPTS

• **Mathematics of computing** → **Multivariate statistics**; *Randomized local search.*

KEYWORDS

Submatrix Localization, Biclustering, Minimax, Scan Statistic, Consistency, Hill-Climbing Algorithm, Golden Section Search

ACM Reference Format:

Yuchao Liu and Ery Arias-Castro. 2019. A Multiscale Scan Statistic for Adaptive Submatrix Localization. In *The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19)*, August 4–8, 2019, Anchorage, AK, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3292500.3330916>

1 INTRODUCTION

Observing a data matrix \mathbf{X} , the problem of submatrix localization, a.k.a., biclustering or coclustering, consists in localizing one or several submatrices whose entries are ‘unusually large’, or significant in some other prescribed way. Such submatrices may be of special interest, since the unusual entries may indicate some potential association or relationship between the corresponding variables. One important application is in the analysis of gene expression data [13, 29], where the row and column indices stand for the genes and conditions, respectively. See [10, 23] for a survey and [26] for a comparison of existing methods.

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

KDD '19, August 4–8, 2019, Anchorage, AK, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6201-6/19/08...\$15.00

<https://doi.org/10.1145/3292500.3330916>

1.1 Submatrix localization

We denote by M and N the number of rows and columns of the data matrix $\mathbf{X} = (X_{ij})$. We assume for simplicity that there is only one submatrix, of size $m^* \times n^*$, to be localized. We further assume that the entries are independent and normally distributed, and with same (known) variance, as well as homogeneous both inside and outside the anomalous submatrix, or more formally,

$$X_{ij} = \theta 1_{\{(i,j) \in (I^* \times J^*)\}} + \varepsilon_{ij}, \quad (1)$$

where the ε_{ij} are IID standard normal distribution, while $I^* \subset [M]$, $J^* \subset [N]$ are the index sets defining anomalous submatrix. Note that $|I^*| = m^*$ and $|J^*| = n^*$. Here we use the symbol $[M]$ to represent integer set $\{1, \dots, M\}$, while $|I|$ is the cardinality of discrete set I . The parameter θ , which is assumed to be strictly positive in this case, quantifies the per-element signal strength within the submatrix.

This parametric setup is used, for example, in [6, 7, 19]. [6] focuses on the simpler problem of detecting the presence of an anomalous submatrix and considers the case where the submatrix size is unknown, while [7, 19] focus on localizing the submatrix with full knowledge of its size (m^*, n^*) . In all these papers, submatrices are ‘scanned’ for anomaly. In detail, the scan statistic, with (m^*, n^*) known, is defined as

$$\text{SCAN}_{m^*, n^*}(\mathbf{X}) = \max_{I \subset [M], |I|=m^*, J \subset [N], |J|=n^*} \sum_{(i,j) \in I \times J} X_{ij}.$$

Localizing the submatrix by the scan statistic is simply returning the row and column index sets achieving the maximum, namely:

$$\Phi_{\text{SCAN}}(\mathbf{X}) = \arg \max_{I \subset [M], |I|=m^*, J \subset [N], |J|=n^*} \sum_{(i,j) \in I \times J} X_{ij}. \quad (2)$$

This estimator has the ‘minimax’ property under the above assumptions.

THEOREM 1 ([7] THEOREMS 2.1, 2.2). *Consider the model (1) as described above in an asymptotic regime where*

$$M, N, m^*, n^* \rightarrow \infty, \quad \frac{\max(m^*, n^*)}{\min(M, N)} \rightarrow 0. \quad (3)$$

Denote

$$\theta_0 = \max \left\{ \frac{\sqrt{2 \log n^*} + \sqrt{2 \log (N - n^*)}}{\sqrt{m^*}}, \frac{\sqrt{2 \log m^*} + \sqrt{2 \log (M - m^*)}}{\sqrt{n^*}}, \frac{\sqrt{2 n^* \log (N/n^*)} + 2 m^* \log (M/m^*)}{\sqrt{m^* n^*}} \right\}.$$

Then if

$$\liminf \theta / \theta_0 > 1,$$

the estimator (2) is strongly consistent in the sense that

$$\mathbb{P}(\Phi_{\text{SCAN}}(\mathbf{X}) \neq I^* \times J^*) \rightarrow 0.$$

Moreover, if

$$\limsup \theta/\theta_0 < 1, \quad (4)$$

there does not exist a strongly consistent estimator.

Condition (4) provides a lower bound of the signal strength for the existence of a strongly consistent estimator with *prior knowledge* of the anomaly size (m^*, n^*) .

We consider here the case when the submatrix size (m^*, n^*) is *not known*. Our strategy is straightforward: we consider a variant of the scan statistic, different from the one defined in (2), which scans submatrices of different sizes.

Contribution 1 (Multiscale scan statistic and its property). We propose a multiscale scan statistic as an estimator under the parametric model defined in Section 1.1. We prove that under some regularity conditions, this statistic exactly recovers the submatrix with high probability when the signal strength is of the same order of magnitude as that implied by Theorem 1 under the setting where the submatrix size is known a priori.

Scanning, even when the submatrix size is known, is computationally difficult. It is in fact known to be NP-hard [13, Thm 1]. To tackle the problem of computing the scan statistic defined in (2), [27] develops an alternate maximization algorithm called LAS (for Large Average Submatrix). Based on the LAS algorithm, we develop two algorithms in order to approach our proposed statistic.

Contribution 2. We develop iterative algorithms to solve the problem of finding the proposed multiscale scan statistic. One algorithm adapts the idea of [27] and is a hill-climbing type optimization algorithm. The other is a variant based on the golden section search.

1.2 More related work

There is an active research line focusing on different aspects of submatrix localization. [3, 6] concentrate on the detection of the existence of such an anomalous submatrix. Besides [7, 19], discussed above, [16] also considers the minimal signal strength needed for accurate localization under symmetric data structure, while [16, 17] develop upper and lower bounds for the existence of weakly consistent estimators. A convex optimization framework for biclustering is proposed and associated algorithms are developed in [14]. A selective inference framework for quantifying the information contained within a selected submatrix is proposed by [20].

The computation issue is attracting an increasing amount of attention from researchers in the field, addressing the tradeoff between statistical power and computational tractability [8, 12, 22]. In this context, a variety of computationally tractable (i.e., running in polynomial time) methods have been proposed and analyzed, such as methods relying on a semidefinite relaxation [12], spectral methods [8], methods relying on message passing heuristics [17], and more [5, 28].

Another closely related research area is network analysis, for example, stochastic blockmodels [18]. Similar to submatrix localization, there are works that study the problem of detection [4, 30, 31], some that study the task of partitioning the graph [25], for which spectral and SDP methods have been proposed [1, 11, 12, 24].

1.3 Content

The remaining of this paper is organized as follows. In Section 2 we introduce the multiscale scan statistic. Section 3 introduces the main algorithms for approximating the proposed statistic: a hill-climbing search algorithm and a golden section search algorithm. Some theory is established. The theoretical properties of the (exact) scan statistic are stated in Section 4. Section 5 presents the result of some numerical simulations.

2 THE MULTISCALE SCAN STATISTIC

The scan statistic as defined in (2) requires knowledge of the size of the submatrix, which we have denoted (m^*, n^*) . When this is unknown, we use a criterium for comparing sums over submatrices of different sizes. For example, if two potential submatrix sizes (m_1, n_1) and (m_2, n_2) are investigated, we need to determine which one of $\text{SCAN}_{m_1, n_1}(\mathbf{X})$ and $\text{SCAN}_{m_2, n_2}(\mathbf{X})$ is more ‘significant’.

When comparing two normalized scan statistics, especially when the sizes of the scan statistics differ significantly, we need to consider the effect of background noise into the comparison. Inspired by the multiscale testing procedure proposed in [6], we define our multiscale scan statistic as

$$\text{MSCAN}(\mathbf{X}) = \max_{I \subset [M], J \subset [N]} \left\{ \frac{\sum_{(i,j) \in I \times J} X_{ij}}{\sqrt{|I||J|}} - \lambda_{|I|, |J|} \right\}, \quad (5)$$

where

$$\lambda_{m,n} = \sqrt{2 \log \left[MN \binom{M}{m} \binom{N}{n} \right]}. \quad (6)$$

Naturally, the corresponding estimator for localizing the anomalous submatrix will be a maximizer of (5), namely

$$\Phi_{\text{MSCAN}}(\mathbf{X}) = \arg \max_{I \subset [M], J \subset [N]} \left\{ \frac{\sum_{(i,j) \in I \times J} X_{ij}}{\sqrt{|I||J|}} - \lambda_{|I|, |J|} \right\}. \quad (7)$$

3 TWO ALGORITHMS

3.1 The LAS algorithm

The computation of (7) seems to require the scanning of all, or most, submatrices, and there are too many submatrices for this to be possible. (The total number of submatrices is equal to 2^{M+N} .) This fact makes the direct computation of the multiscale scan statistic computationally intractable. We also note that the computation of (7) is harder than that of (2), and the latter is proved to be NP-hard by [13]. It seems therefore necessary to resort to approximations.

When the true anomaly size (m^*, n^*) is known, [27] proposed an iterative hill-climbing algorithm, called LAS, to approximate the statistic in (5); see Algorithm 1. As a hill-climbing algorithm, it may be trapped in local maxima, so that in practice the algorithm is run on several (random) initializations [3, 6].

3.2 The adaptive LAS algorithm

We modify Algorithm 1 in order to approximate the multiscale scan statistic. In principle, we could run Algorithm 1 on each submatrix size $(m, n) \in [M] \times [N]$, but although computationally feasible, it remains computationally demanding. Instead, we adopt a more greedy approach. We still rely on alternatively maximizing the objective, over the column index set, then the row index set, etc,

Algorithm 1: Large Average Submatrix (LAS) [27]

Input: submatrix size (m^*, n^*) , initial row index set \hat{I} with $|\hat{I}| = m^*$

Output: \hat{I}, \hat{J}

- 1 Compute the sum $S_j = \sum_{i \in \hat{I}} X_{ij}$, and let \hat{J} be the column index set corresponding to the largest n^* items in $\{S_j\}_{j=1}^N$;
- 2 Compute the sum $T_i = \sum_{j \in \hat{J}} X_{ij}$, and let \hat{I} be the row index set corresponding to the largest m^* items in $\{T_i\}_{i=1}^M$;
- 3 Repeat *Steps 1 and 2* until convergence.

until convergence, but the objective is that of (5). See Algorithm 2 for details. In practice, the algorithm is run on multiple random initializations as well.

Algorithm 2: Adaptive LAS

Input: initial submatrix size (m, n) , initial row index set \hat{I} with $|\hat{I}| = m$

Output: \hat{I}, \hat{J}

- 1 Run Algorithm 1 based on the initial input, obtaining (\hat{I}, \hat{J})
- 2 Compute the sums $S_j = \sum_{i \in \hat{I}} X_{ij}$ and $S_j^\dagger = \frac{1}{\sqrt{j}} \sum_{k=1}^j S_{(k)} - \lambda_{|\hat{I}|, j}$ where $S_{(1)} \geq S_{(2)} \geq \dots \geq S_{(N)}$ are the ordered $\{S_j\}_{j=1}^N$;
- 3 Let n maximize S_j^\dagger over $j = 1, \dots, N$, and let \hat{J} be the column index set corresponding to the largest n items in $\{S_j\}_{j=1}^N$;
- 4 Compute the sums $T_i = \sum_{j \in \hat{J}} X_{ij}$ and $T_i^\dagger = \frac{1}{\sqrt{i}} \sum_{l=1}^i T_{(l)} - \lambda_{i, |\hat{J}|}$ where $T_{(1)} \geq T_{(2)} \geq \dots \geq T_{(M)}$ are the ordered $\{T_i\}_{i=1}^M$;
- 5 Let m maximize T_i^\dagger over $i = 1, \dots, M$, and let \hat{I} be the row index set corresponding to the largest m items in $\{T_i\}_{i=1}^M$;
- 6 Repeat *Steps 2-5* until convergence.

Note that the initial submatrix size input (m, n) in Algorithm 2 does not necessarily represent any prior knowledge of the anomaly size, although it can be informed by such prior knowledge.

3.3 A golden section search variant

We propose a variant based on the (two-dimensional) golden section search, a well-known method for finding the maximum of a unimodal function (see [9] for a general survey of the algorithm in N dimensions). Define

$$f_X(m, n) = \frac{\text{SCAN}_{m,n}(X)}{\sqrt{mn}} - \lambda_{m,n}, \quad (8)$$

so that the multiscale scan statistic is the maximum value of f_X :

$$\text{MSCAN}(X) = \max_{m,n} f_X(m, n).$$

We simply apply the two-dimensional golden section search (with some modifications for discreteness) to f_X . Details are provided in Algorithm 3.

The hope is that, if the signal strength is reasonably large and our initial bound on the submatrix size, (\bar{m}, \bar{n}) below, is sufficiently

accurate, f_X is (with high probability) unimodal on $[\bar{m}] \times [\bar{n}]$ with maximizer (m^*, n^*) . We are not able to support this with a theoretical result, but have performed some numerical experiments, some of them reported in Section 5.3, that seem to support this.

In the algorithm, instead of computing $f_X(m, n)$, we instead apply Algorithm 1 to provide an approximation. In the description, ϕ is the golden ratio constant, namely, $\phi = 0.5(\sqrt{5} - 1)$.

Algorithm 3: Golden Section Search Algorithm

Input: initial submatrix size (\bar{m}, \bar{n})

Output: \hat{I}, \hat{J}

- 1 Denote $(m_{\min}, n_{\min}) = (1, 1)$, $(m_{\max}, n_{\max}) = (\bar{m}, \bar{n})$;
- 2 Compute $m_1 = \lceil m_{\max} + (m_{\min} - m_{\max})\phi \rceil$, $m_2 = \lfloor m_{\min} + (m_{\max} - m_{\min})\phi \rfloor$, $n_1 = \lceil n_{\max} + (n_{\min} - n_{\max})\phi \rceil$, $n_2 = \lfloor n_{\min} + (n_{\max} - n_{\min})\phi \rfloor$;
- 3 Compute $f_X(m_1, n_1)$, $f_X(m_2, n_1)$, $f_X(m_1, n_2)$, $f_X(m_2, n_2)$, denote $(m^\#, n^\#) = \arg \max_{(i,j) \in \{1,2\}^2} f_X(m_i, n_j)$;
- 4 Following the standard Golden Section Search algorithm for continuous functions, renew (m_{\max}, n_{\max}) , (m_{\min}, n_{\min}) based on $(m^\#, n^\#)$;
- 5 Repeat *Steps 2-4*, till $\max(m_{\max} - m_{\min}, n_{\max} - n_{\min}) \leq 3$;
- 6 Compute $f_X(s, t)$ for all $(s, t) \in [m_{\min}, m_{\max}] \times [n_{\min}, n_{\max}] \cap \mathbb{N}^2$. Denote the maximal among the results as $f_X(\hat{m}, \hat{n})$. Finalize (\hat{I}, \hat{J}) as the output of Algorithm 1 with input (\hat{m}, \hat{n}) .

We note that the algorithm needs to stop the loop when the search frame is 3×3 or smaller, for otherwise the loop would not break because of the discrete nature in the index sets (the search frame will not shrink). When this happens, the algorithm switches to an exhaustive search. The initial submatrix size (\bar{m}, \bar{n}) is ideally an upper bound on the actual submatrix size. Its choice can be informed by prior knowledge.

4 THEORETICAL PROPERTY**4.1 Gaussian entries**

We establish a performance result for our multiscale scan statistic (7). Recall from Section 2 that the true anomaly $I^* \times J^*$ has size (m^*, n^*) . As before, we allow θ to change with (M, N, m^*, n^*) .

THEOREM 2 (EXACT RECOVERY). *Consider the model (1) as described in Section 1.1 in an asymptotic regime where (3) holds. Set*

$$\theta_1 = C \max \left\{ \sqrt{\frac{\log(M - m^*) + \log m^*}{n^*}}, \sqrt{\frac{\log(N - n^*) + \log n^*}{m^*}}, \frac{\lambda_{m^*, n^*}}{\sqrt{m^* n^*}} \right\}, \quad (9)$$

where C is the positive real solution to the equation

$$C = 2([C/(C-1)]^{1.5} + [C/(C-1)]^{1.25}). \quad (10)$$

If

$$\liminf \frac{\theta}{\theta_1} > 1, \quad (11)$$

the estimator defined by (7) is equal to (I^*, J^*) with high probability as $M, N, m^*, n^* \rightarrow \infty$.

Remark 1. We note that $C \approx 4.32$. By comparing the terms defining θ_1 with the terms defining θ_0 (first defined in Theorem 1), it is easy to see that $\theta_1/\theta_0 = O(1)$.

Remark 2. The first two terms in the maximum (9) are associated with the row and column structures of the anomaly and data matrix, while the last part is associated with the existence of weak consistent estimators. In fact, [16, 17] show that when (m^*, n^*) are known, the scan statistic (2) is weakly consistent if

$$\liminf \theta \left\{ \frac{\lambda_{m^*, n^*}}{\sqrt{m^* n^*}} \right\}^{-1} > 1.$$

(There is weak consistence if $\mathbb{E}(|(I^* \times J^*) \Delta \Phi_{\text{SCAN}}|/m^* n^*) \rightarrow 0$, where Δ denotes the set symmetric difference.)

PROOF STRATEGY. Denote $S^* = I^* \times J^*$, which identifies the submatrix. Key to the proof is to bound the probability that, for some other submatrix S of size (m, n) ,

$$\frac{\sum_{(i,j) \in S} X_{ij}}{\sqrt{mn}} - \lambda_{m,n} > \frac{\sum_{(i,j) \in S^*} X_{ij}}{\sqrt{m^* n^*}} - \lambda_{m^*, n^*}. \quad (12)$$

The union of such events, as S runs through all matrices, is exactly the event where the multiscale scan fails at exact recovery.

We will start by defining some quantities which quantify the closeness between the candidate matrix S and true anomaly S^* . For the submatrices in $\mathcal{S}_{m,n}$, which is the collection of all submatrices with dimensions (m, n) , define their affinity in size as

$$\rho_{m,n} = \max_{S \in \mathcal{S}_{m,n}} \frac{|S \cap S^*|}{\sqrt{|S||S^*|}} = \frac{(m \wedge m^*)}{\sqrt{mm^*}} \frac{(n \wedge n^*)}{\sqrt{nn^*}}.$$

For a specific submatrix S with size (m, n) , such that $S \cap S^*$ is a $s \times t$ submatrix, define the affinity of coverage as

$$v_{m,n,s,t} = \frac{st}{\sqrt{mm^* nn^*}}.$$

The proof strategy is that, with probability converging to 1, we can sequentially rule out the submatrices with low affinity in size and low affinity in coverage. We then analyze those submatrices left behind, which have size and coverage comparable with the true anomaly.

In fact, the introduction of the regularizer $\lambda_{m,n}$ in the expression of the multiscale scan statistic is there to eliminate the submatrices with low affinity in size or in coverage. While when the affinities are not too small (bounded from zero by a positive number), the number of candidate matrices is controlled. We measure the probability of (12) when S and S^* only differ by one column or one row, and show that this is the largest failure probability among all candidates that have affinities bounded from below. \square

4.2 An extension to an exponential family

In practice the Gaussian assumption is not satisfied in many situations. For example the entries may only take integer values (as with counting data) or binary data (as with presence-absence matrices in Ecology [15]). Therefore it is of importance to extend the result to a distribution family that covers several common data types. Following [2, 3, 6], we consider a one parameter exponential family. Let ν denote a distribution on the real line with mean zero and variance 1. Denote $\varphi(\theta)$ as the moment generating function of ν ,

meaning $\varphi(\theta) = \int e^{\theta t} \nu(dt)$, and suppose that $\varphi(\theta) < \infty$ for all θ in $[0, \theta_*)$. Here θ_* is defined as $\sup\{\theta : \varphi(\theta) < \infty\}$ and could be equal to infinity. For $\theta \in [0, \theta_*)$, define

$$f_\theta(x) = \exp\{\theta x - \log \varphi(\theta)\}, \quad (13)$$

which is a density with respect to ν .

By selecting ν appropriately, the distribution family becomes the normal location family, the Poisson family, or the Rademacher family. Note that with ν fixed, the distribution family is stochastically monotone in θ ([21], Lemma 3.4.2), which is to say, for $X_1 \sim f_{\theta_1}$ and $X_2 \sim f_{\theta_2}$ with $0 \leq \theta_2 \leq \theta_1 \leq \theta_*$ and fixed ν ,

$$\mathbb{P}(X_1 \geq t) \geq \mathbb{P}(X_2 \geq t), \forall t \in \mathbb{R}. \quad (14)$$

This fact enables us to model the submatrix localization problem with θ controlling the signal strength. With θ increasing, the anomaly is more 'anomalous', making the localization problem relatively easier to solve. The localization problem can now be formalized with ν as the role of noise, and entries in the anomaly are distributed as f_θ . Formally,

$$X_{ij} \sim \begin{cases} f_\theta & (i,j) \in I^* \times J^* \\ \nu & (i,j) \notin I^* \times J^* \end{cases} \quad (15)$$

THEOREM 3 (LOWER BOUND UNDER EXPONENTIAL FAMILY ASSUMPTION). *Under the assumption of (13) and suppose (3) holds. Additionally, we assume*

$$\frac{\max(\log M, \log N)}{\min(m^*, n^*)} \rightarrow 0. \quad (16)$$

Then if

$$\theta \leq C \max \left(\sqrt{\frac{\log M}{n^*}}, \sqrt{\frac{\log N}{m^*}} \right),$$

with $C < 1/\sqrt{8}$, even with full knowledge of (m^*, n^*) , there does not exist a strongly consistent estimator of (I^*, J^*) .

PROOF SKETCH OF THEOREM 3. The proof is following most of the arguments in the proof of Theorem 1 in [19]. The result focuses on the normal location family, but extends to other one-parameter exponential families. We aim to prove the following:

$$\mathcal{D}(\mathcal{P}_0 || \mathcal{P}_j) \geq m^* \theta^2 (1 + o(1)), \quad (17)$$

where $\mathcal{D}(\cdot || \cdot)$ is the Kullback Leibler divergence between two distribution, and $\mathcal{P}_0, \mathcal{P}_j$ are the distributions of data \mathbf{X} when $S^* = [m^*] \times [n^*]$ and $S^* = [m^*] \times \{[n^* - 1] \cup \{j\}\}$, with $j \in \mathbb{N} \cap [n^* + 1, N]$, respectively. Replace (15) in the proof of Theorem 1 in [19] with the above inequality, and the result follows. \square

This result also partially answers the open problem raised by [7] asking for a lower bound under a general exponential family.

Before we head into evaluating the performance of our proposed estimator (7) applied under the exponential family assumption, we need to re-define the adjusting quantity in (6). Fixing a constant $\delta > 0$, we re-define $\lambda_{m,n}$ as

$$\lambda_{m,n} = \sqrt{(2 + \delta) \log \left[MN \binom{M}{m} \binom{N}{n} \right]}. \quad (18)$$

The multiscale scan statistic, and its associated estimator of the anomalous submatrix, is defined again according to (5) and (7).

THEOREM 4 (EXACT RECOVERY UNDER EXPONENTIAL FAMILY). *Assume the model (13), and suppose that (3) and (16) hold. Take \underline{m} and \underline{n} such that $\log(N) = o(\underline{m})$ and $\log(M) = o(\underline{n})$. Suppose we scan over submatrices of size (m, n) such that $m \geq \underline{m}$ and $n \geq \underline{n}$, and that the anomalous submatrix is among these, meaning that $m^* \geq \underline{m}$ and $n^* \geq \underline{n}$. Suppose θ_1 is as in Theorem 2. If*

$$\liminf \frac{\theta}{\theta_1} > 1,$$

with high probability as $M, N, m, n \rightarrow \infty$, this scan returns the anomalous submatrix.

Remark 3. There is a normal approximation underneath which drives the asymptotic behavior of the statistic to be close to that under the normal model. This explains why the result is similar under a general exponential family as it is under the normal family. In fact, the proof of Theorem 4 is otherwise essentially the same as that of Theorem 2.

5 NUMERICAL EXPERIMENTS

5.1 Performance of Proposed Algorithms

We performed some simulation experiments to evaluate the performance of our proposed algorithms. We generate a matrix with independent entries according to (1) as well as (15), and perform the proposed algorithms (Adaptive LAS and Golden Section Search LAS) on the generated data. By permutation invariance, we simply choose $I^* = [m^*]$ and $J^* = [n^*]$. We then measure the accuracy of an estimator (\hat{I}, \hat{J}) as follows

$$\text{Err}(\hat{I}, \hat{J}) = \log(|\hat{I} \triangle [m^*]| + |\hat{J} \triangle [n^*]| + 1). \quad (19)$$

We compare the proposed algorithms with two computationally tractable algorithms with proved consistency.

- Spectral algorithm of [8]. The algorithm computes the singular value decomposition of the data matrix \mathbf{X} , then applies k -means algorithm on the first left and right singular vectors with number of clusters set to $k = 2$, and clusters the row and column indices according to corresponding singular vector's clustering results.
- Greatest Marginal Gap method of [5]. The algorithm calculates row and column sums of the data matrix \mathbf{X} , denoted as $\{RS_i\}$ and $\{CS_i\}$, then sorts the row and column sums, obtaining $\{RS_{(i)}\}$ and $\{CS_{(i)}\}$. The algorithm clusters rows based on $i_R = \arg \max_i RS_{(i+1)} - RS_{(i)}$, yielding two sets $\{RS_{(1)}, \dots, RS_{(i_R)}\}$ and $\{RS_{(i_R+1)}, \dots, RS_{(M)}\}$. The clustering of columns is analogous.

We consider a balanced setting where (M, N, m^*, n^*) is taken to be (1000, 1200, 170, 140), and also an unbalanced setting where (M, N, m^*, n^*) is taken to be (4000, 500, 70, 250). For each fixed signal strength, the data is independently generated 30 times and the error counts defined by (19) for all four the algorithms are recorded. Three one-parameter exponential models (normal, Poisson, Rademacher) are investigated to illustrate the performance of the algorithms under different distributions. See Figure 1 and Figure 2.

5.1.1 Signal strength. The signal strength is quantified by the parameter θ of the entries inside the submatrix. Denote the following quantity θ_{crit} :

$$\theta_{\text{crit}} = \max \left(\sqrt{\frac{\log M}{n^*}}, \sqrt{\frac{\log N}{m^*}}, \sqrt{\frac{\log M + \log N}{m^* + n^*}} \right).$$

This is part of the quantity on the left hand-side of (11), inside Theorem 2. We zoom in to the interval $[1.0 \times \theta_{\text{crit}}, 4.0 \times \theta_{\text{crit}}]$. The signal strength is increased in the process of simulation, each time by $0.1 \times \theta_{\text{crit}}$.

5.1.2 Simulation result. The two proposed algorithms perform similarly under the balanced design, across all three data types. See Figure 1. Adaptive Hill-Climbing has a slightly weaker performance when the signal is weak (around $1 \times \theta_{\text{crit}}$ to $1.5 \times \theta_{\text{crit}}$), but the two algorithms' error rates shrink to zero at the same signal strength, showing that the two approximate algorithms return the same result when signal strength goes beyond some threshold, and the result would exactly recover the planted anomaly.

As Figure 2 shows, the two algorithms perform differently under the imbalanced design when the signal is weak, with the Golden Section Search performing better. This is due to the fact that the search space of GSS is smaller, as well as the existence of local maximums when the design is imbalanced (see Figure 4 for an example). However, when considering the successful rate of exact recovery, the signal strength, above which the algorithm's error rate shrinks to zero is still similar for both algorithms.

It is worth mentioning that in both cases, our proposed methods outperform the other two computationally tractable methods. In the balanced case, both Adaptive Hill-Climbing and Golden Section Search beat the spectral method by a small margin in error rate, while the error rate of Greatest Marginal Gap method is much worse in the examined signal strength region. In the imbalanced case, Adaptive Hill-Climbing has similar performance with spectral method in the weak signal region, while the GSS again outperforms both, showing that the multiscale scan statistic has better discovery power compared to spectral method. Again, the error rate of Greatest Marginal Gap is much larger.

5.2 Computing time

Here we present the computing time for executing the programs described in the previous section. The simulated data is from model (1), with corresponding sizes $(M, N, m^*, n^*) = (1000, 1000, 100, 100)$ and signal strength $\theta = 2.5 \times \theta_{\text{crit}}$. We generate data under this model 100 times, and each time record the computation time for each of the competing methods.

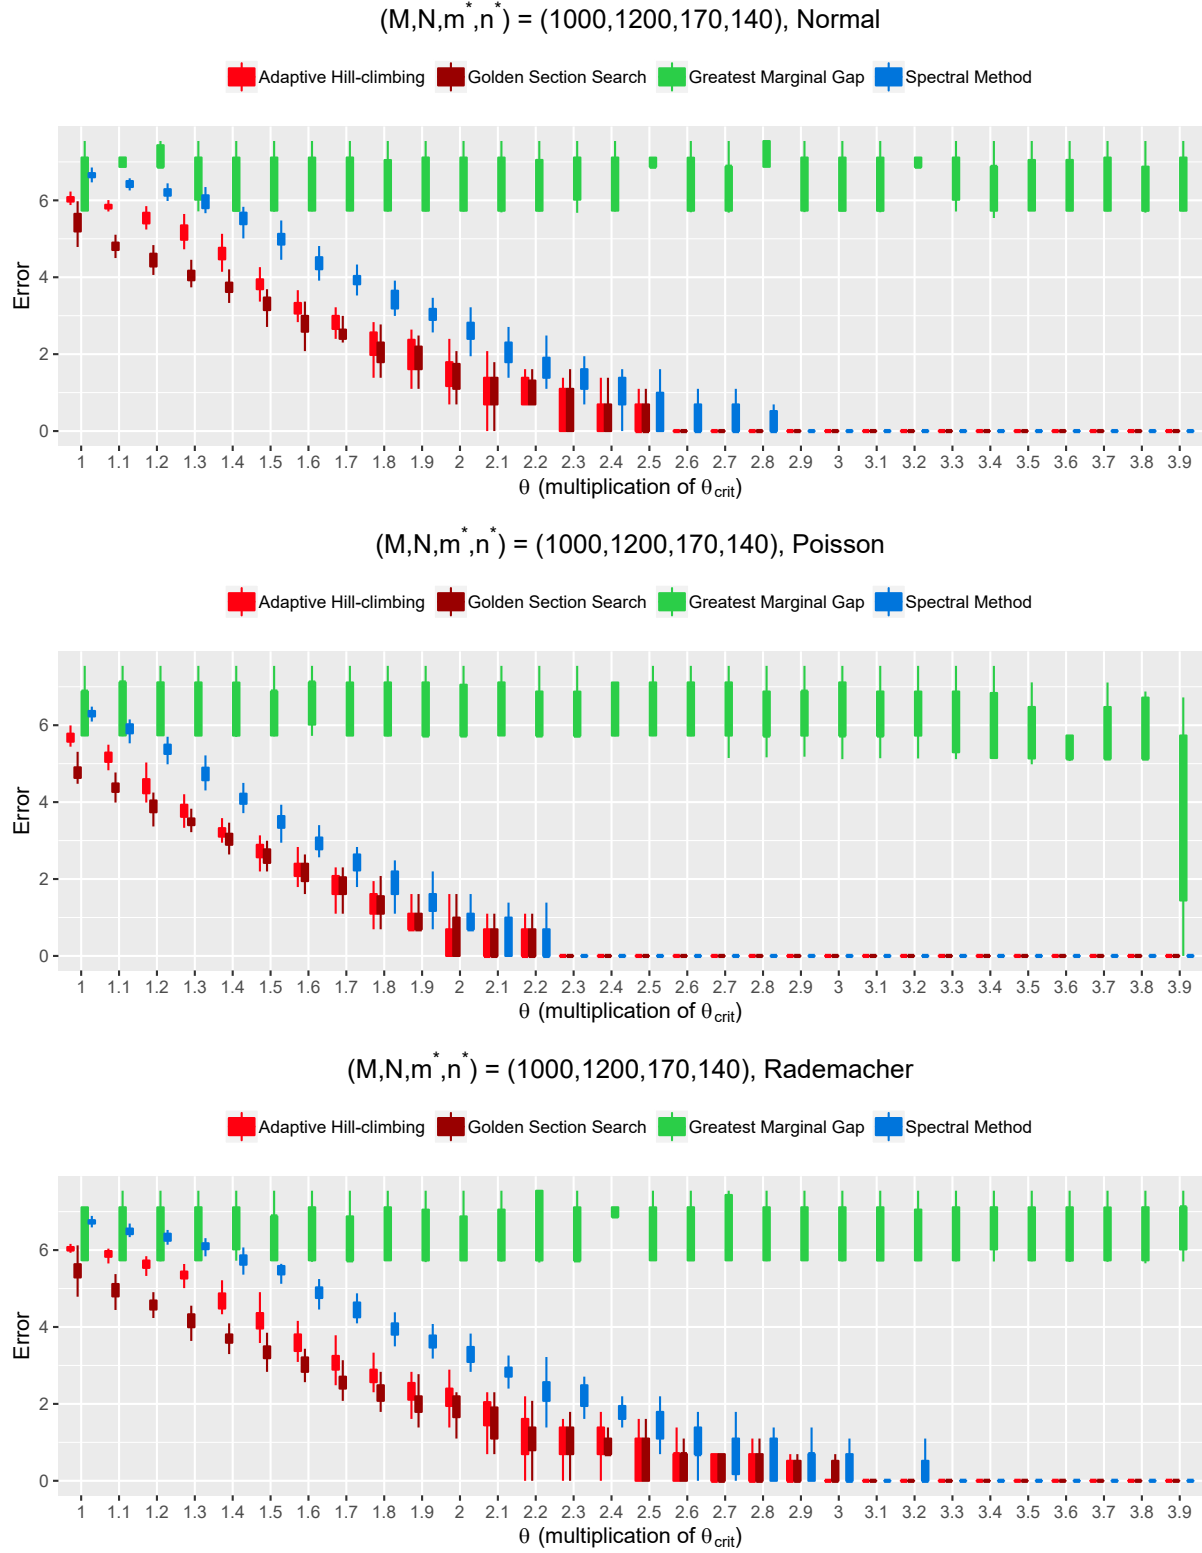


Figure 1: Error counts for a balanced design

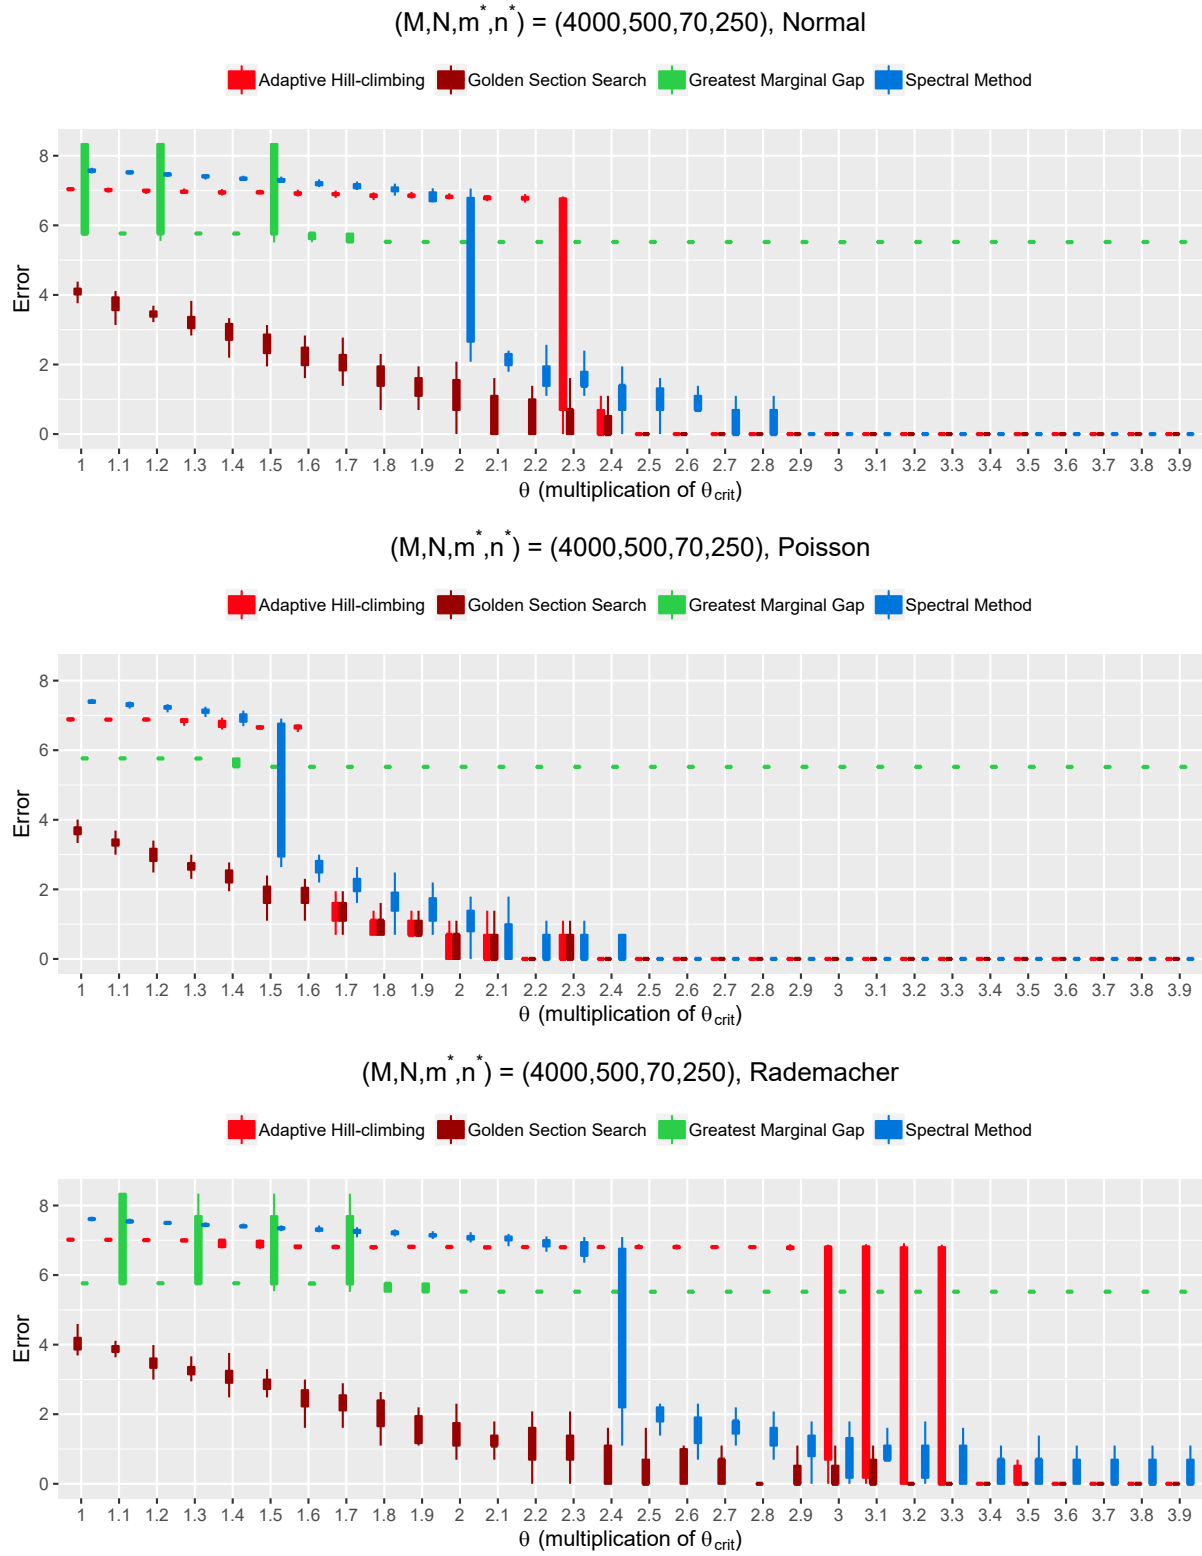


Figure 2: Error counts for an imbalanced design




Figure 3: Computing times for different algorithms. (The computing time is calculated by function ‘proc.time’ in programming language R. Category ‘user’ represents the time for executing the program codes, ‘system’ is the CPU time charged for execution by the system on behalf of the calling process, and ‘elapsed’ is the sum of the other two.)




Figure 4: Levelplots for illustrating unimodality

The spectral method was computationally much more demanding than the other methods. As expected, the clustering algorithm based on the largest marginal gap is the fastest, however our algorithms are not too far (while much more accurate).

5.3 Unimodality issue in Algorithm 3

The golden section search presented in Algorithm 3 requires the function defined in (8) to be unimodal, in order to successfully discover the global maximal. In general, checking the unimodality of a function is often hard and here we present some numerical experiment illustrating the unimodality of the target function (8).

We use two simulated Gaussian datasets. For each combination of (M, N, m^*, n^*) , we set the level of signal to $\theta = 2 \times \theta_{\text{crit}}$, which is just above the signal level such that the search algorithm makes few to zero mistakes. Then for every pair of (m, n) such that $m \leq \bar{m}, n \leq \bar{n}$, we calculate the function value according to (8). The scan statistic $\text{SCAN}_{m,n}(X)$ is calculated by Algorithm 1, and $\lambda_{m,n}$ is approximated by $(2(\log M + \log N + m \log(M/m) + n \log(N/n)))^{1/2}$. Here we use the fact that $\log \binom{K}{k} \approx k \log(K/k)$ when $k \ll K$. The simulation is otherwise set as in the previous section, in that we examine both a balanced and an imbalanced setting. In the balanced setup, $(M, N, m, n) = (300, 360, 40, 60)$ and $(\bar{m}, \bar{n}) = (100, 120)$. In the imbalanced setup, $(M, N, m, n) = (500, 50, 10, 25)$ and $(\bar{m}, \bar{n}) = (100, 50)$.

To better illustrate the unimodality, the simulated values of f is raised to its fourth moment to enlarge the difference around the mode. We can see from Figure 4 that except for a few points around the edge, the target function (8) has a clear unimodal structure on the majority of the search field $[\bar{m}] \times [\bar{n}]$.

6 CONCLUSION AND DISCUSSION

In this paper we propose a new multiscale scan statistic for localizing the anomalous submatrix inside a large noisy matrix, which does not require prior knowledge on the anomalous submatrix size. We show the signal strength needed for strong consistency, and design two algorithms with good approximating accuracy and computing speed. There are, however, some problems for future work and discussions.

Minimaxity of the multiscale scan statistic: [7] showed a sharp minimax signal bound, as we described in Theorem 1. Can our estimator based on the multiscale scan statistic reach the same minimax bound (which relies on knowledge of the submatrix size)? While we are unable to reduce the constant in Theorem 2 to the bound, we conjecture that our estimator is essentially minimax. Another measurement of accuracy of estimators is weak consistency. [16, 17] shows that (2) is minimax in the weak consistency sense. When is the multiscale scan statistic weak consistent and is it minimax?

Proof of unimodality of (8): What is the relationship of $(M, N, m^*, n^*, \bar{m}, \bar{n})$ and θ , such that (8) is unimodal on $[\bar{m}] \times [\bar{n}]$ with high probability? Solution to this problem will directly lead to the success guarantee of Algorithm 3.

Computationally tractable algorithms: There is a growing literature about using computationally tractable methods such as semidefinite programming to approximate NP-hard problems. (for example, [12] on using SDP on finding the maximal likelihood estimator). Can such an algorithm be designed to approximate with guaranteed accuracy the multiscale scan?

Tight minimax bound for exponential family case: As [7] mentioned, minimax theory on the case of exponential family is still open. We give a bound based on [19], but this bound is not

tight. Also, although the scan statistic under the exponential family setup is proved to be minimax in testing by [6], the analysis of its localization performance is still open.

ACKNOWLEDGMENTS

The authors thank the reviewers for the helpful feedback, and Jiaqi Guo for help in simulation studies.

REFERENCES

- [1] Emmanuel Abbe, Afonso S Bandeira, and Georgina Hall. 2016. Exact recovery in the stochastic block model. *IEEE Transactions on Information Theory* 62, 1 (2016), 471–487.
- [2] Ery Arias-Castro, Rui M Castro, Ervin Tanczos, and Meng Wang. 2018. Distribution-free detection of structured anomalies: Permutation and rank-based scans. *J. Amer. Statist. Assoc.* 113, 522 (2018), 789–801.
- [3] Ery Arias-Castro and Yuchao Liu. 2017. Distribution-free detection of a submatrix. *Journal of Multivariate Analysis* 156 (2017), 29–38.
- [4] Ery Arias-Castro and Nicolas Verzelen. 2014. Community detection in dense random networks. *The Annals of Statistics* 42, 3 (2014), 940–969.
- [5] Vincent Brault and Antoine Channarond. 2016. Fast and Consistent Algorithm for the Latent Block Model. *arXiv preprint arXiv:1610.09005* (2016).
- [6] Cristina Butucea and Yuri I Ingster. 2013. Detection of a sparse submatrix of a high-dimensional noisy matrix. *Bernoulli* 19, 5B (2013), 2652–2688.
- [7] Cristina Butucea, Yuri I Ingster, and Irina A Suslina. 2015. Sharp variable selection of a sparse submatrix in a high-dimensional noisy matrix. *ESAIM: Probability and Statistics* 19 (2015), 115–134.
- [8] T Tony Cai, Tengyuan Liang, Alexander Rakhlin, et al. 2017. Computational and statistical boundaries for submatrix localization in a large noisy matrix. *The Annals of Statistics* 45, 4 (2017), 1403–1430.
- [9] Yen-Ching Chang. 2009. N-dimension golden section search: Its variants and limitations. In *2009 2nd International Conference on Biomedical Engineering and Informatics*. IEEE, 1–6.
- [10] Malika Charrad and Mohamed Ben Ahmed. 2011. Simultaneous clustering: A survey. In *International Conference on Pattern Recognition and Machine Intelligence*. Springer, 370–375.
- [11] Kamalika Chaudhuri, Fan Chung, and Alexander Tsiatas. 2012. Spectral clustering of graphs with general degrees in the extended planted partition model. In *Conference on Learning Theory*. 1–23.
- [12] Yudong Chen and Jiaming Xu. 2016. Statistical-computational tradeoffs in planted problems and submatrix localization with a growing number of clusters and submatrices. *Journal of Machine Learning Research* 17, 27 (2016), 1–57.
- [13] Yizong Cheng and George M Church. 2000. Bicustering of expression data. In *Proceedings of the Eighth International Conference on Intelligent Systems for Molecular Biology*. AAAI Press, 93–103.
- [14] Eric C Chi, Genevera I Allen, and Richard G Baraniuk. 2016. Convex bicustering. *Biometrics* (2016).
- [15] Nicholas J Gotelli. 2000. Null model analysis of species co-occurrence patterns. *Ecology* 81, 9 (2000), 2606–2621.
- [16] Bruce Hajek, Yihong Wu, and Jiaming Xu. 2017. Information limits for recovering a hidden community. *IEEE Transactions on Information Theory* 63, 8 (2017), 4729–4745.
- [17] Bruce Hajek, Yihong Wu, and Jiaming Xu. 2017. Submatrix localization via message passing. *The Journal of Machine Learning Research* 18, 1 (2017), 6817–6868.
- [18] Paul W Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. 1983. Stochastic blockmodels: First steps. *Social networks* 5, 2 (1983), 109–137.
- [19] Mladen Kolar, Sivaraman Balakrishnan, Alessandro Rinaldo, and Aarti Singh. 2011. Minimax localization of structural information in large noisy matrices. In *Advances in Neural Information Processing Systems*. 909–917.
- [20] Jason D Lee, Yuekai Sun, and Jonathan E Taylor. 2015. Evaluating the statistical significance of biclusters. In *Advances in Neural Information Processing Systems*. 1324–1332.
- [21] Erich L. Lehmann and Joseph P. Romano. 2005. *Testing Statistical Hypotheses* (3rd ed.). Springer, New York.
- [22] Zongming Ma and Yihong Wu. 2015. Computational barriers in minimax submatrix detection. *The Annals of Statistics* 43, 3 (2015), 1089–1116.
- [23] Sara C Madeira and Arlindo L Oliveira. 2004. Bicustering algorithms for biological data analysis: a survey. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)* 1, 1 (2004), 24–45.
- [24] Frank McSherry. 2001. Spectral Partitioning of Random Graphs. In *Proceedings of the 42nd IEEE symposium on Foundations of Computer Science*. IEEE Computer Society, 529.
- [25] Elchanan Mossel, Joe Neeman, Allan Sly, et al. 2016. Consistency thresholds for the planted bisection model. *Electronic Journal of Probability* 21 (2016).
- [26] Amela Prelić, Stefan Bleuler, Philip Zimmermann, Anja Wille, Peter Bühlmann, Wilhelm Gruissem, Lars Hennig, Lothar Thiele, and Eckart Zitzler. 2006. A systematic comparison and evaluation of bicustering methods for gene expression data. *Bioinformatics* 22, 9 (2006), 1122–1129.
- [27] Andrey A Shabalin, Victor J Weigman, Charles M Perou, and Andrew B Nobel. 2009. Finding large average submatrices in high dimensional data. *The Annals of Applied Statistics* 3, 3 (2009), 985–1012.
- [28] Kean Ming Tan and Daniela M Witten. 2014. Sparse bicustering of transposable data. *Journal of Computational and Graphical Statistics* 23, 4 (2014), 985–1008.
- [29] Amos Tanay, Roded Sharan, and Ron Shamir. 2002. Discovering statistically significant biclusters in gene expression data. *Bioinformatics* 18, suppl 1 (2002), S136–S144.
- [30] Nicolas Verzelen and Ery Arias-Castro. 2015. Community detection in sparse random networks. *The Annals of Applied Probability* 25, 6 (2015), 3465–3510.
- [31] Anderson Y Zhang and Harrison H Zhou. 2016. Minimax rates of community detection in stochastic block models. *The Annals of Statistics* 44, 5 (2016), 2252–2280.

A SIMULATION DETAILS FOR REPRODUCIBILITY

In the appendix we list some detail setups of our simulation in order to provide essential information for reproducing the simulation results in the paper. All the simulation are done with statistical software R version 3.5.2. The graph is generated by graphing package ggplot2. The random number generator are set with seed 100 by the set.seed function for each single graph. For the further reproduction purposes, all the simulation codes are available at <https://github.com/nozoeli/adaptiveBiclustering>.

A.1 Simulation Setup for Figure 1 and Figure 2

We list setups used during our simulation, mainly regarding the inputs of Algorithm 2 and Algorithm 3.

For Algorithm 2, the initial input $(m, n) = (25, 25)$. The initial row index set \hat{I} is uniformly randomly chosen from $[M]$. The algorithm is repeated 50 times with I.I.D. chosen \hat{I} , and the result producing the largest multiscale scan statistic is returned.

For Algorithm 3, the input $(\bar{m}, \bar{n}) = (500, 500)$. When using Algorithm 1 to calculate f , the initial row index is chosen with rows

with the largest row sums across all columns (e.g. input $m = 20$, the initial row index set is made of the indices with top 20 largest row sums), and Algorithm 1 is run once to calculate the scan statistic inside f .

These setups are used across the calculation of all figures in Figure 1 and Figure 2.

A.2 Simulation Setup for Figure 3

The setup of Algorithm 2 during this part of simulation is $(m, n) = (5, 5)$, with the rest setup the same as the previous section.

The setup of Algorithm 3 is the same as the previous section.

The function used in spectral method is as follows. The singular vector decomposition is using the R function svd, and the k -means algorithm is calling R function kmeans with default setup and number of clusters 2.

A.3 Simulation Setup for Figure 4

The calculation of scan statistic during this part of simulation is repeating Algorithm 1 100 times with I.I.D. initial row indexes uniformly randomly chosen from $[M]$, and return the result with the largest entry sum.