



Distribution-Free, Size Adaptive Submatrix Detection with Acceleration

Yuchao Liu and Jiaqi Guo

Microsoft Corporation
One Microsoft Way
Redmond, Washington, 98052, USA.
E-mail address: yuchao.liu@microsoft.com

Amazon Inc.
410 Terry Ave. N
Seattle, Washington, 98109, USA.
E-mail address: jiaqig@amazon.com

Abstract. Given a large matrix containing independent data entries, we consider the problem of detecting a submatrix inside the data matrix that contains larger-than-usual values. Different from previous literature, we do not have exact information about the dimension of the potential elevated submatrix. We propose a Bonferroni type testing procedure based on permutation tests, and show that our proposed test loses no first-order asymptotic power compared to tests with full knowledge of potential elevated submatrix. In order to speed up the calculation during the test, an approximation net is constructed and we show that Bonferroni type permutation test on the approximation net loses no power on the first order asymptotically.

1. Introduction

Matrix type data are common in contemporary data analysis and have wide applications in biology, social sciences and other fields. In many situations, the row and column indexes represent individuals or units that could interact with each other, and the entries / data points evaluate the level of the interaction between row and column units. For example, in DNA chips analysis, the rows represent genes and columns represent situations. The entries inside the data matrix can

Received by the editors November 15th, 2018; accepted February 29th, 2020.

2010 Mathematics Subject Classification. 60K35, 82B43.

Key words and phrases. Submatrix detection, Permutation test, Scan statistics, Approximation net, Distribution-free test.

Majority of the work is done when the authors are PhD students at the Mathematics Department of University of California, San Diego.

be expression levels of some genes under some situations. See [Cheng and Church \(2000\)](#) for a detailed introduction.

Bi-clustering, or co-clustering, or sometimes referred as simultaneous clustering, aims to find subsets of row and column indexes, such that entries inside the submatrix indexed by those units, are 'special'. This hidden structure is usually of interest to researchers and can be used for specific references. See [Tanay et al. \(2005\)](#); [Charrad and Ben Ahmed \(2011\)](#); [Kriegel et al. \(2009\)](#) for surveys on this topic.

Before making inferences based on the bi-clustering result, one important question is that, does the result contain any information, or is it just a product of pure noise? This leads to the submatrix detection problem, which is a hypothesis testing procedure distinguishing data from pure noise and data containing some hidden structures.

1.1. Submatrix detection. We consider the simplest case where there is one potential 'special' submatrix to discover. Furthermore, assume that entries are 'larger than usual' inside the submatrix. The observed data is denoted by $\mathbf{X} = \{X_{ij}\}$ with M rows and N columns. The hypothesis testing problem is formulated as follows:

$$H_0: X_{ij} \text{ are IID for all } i \in [M] \text{ and } j \in [N]$$

vs

$$H_1: \text{There exists } \mathcal{I} \subset [M] \text{ and } \mathcal{J} \subset [N] \text{ such that } X_{ij} \text{ are stochastically larger when } (i, j) \in \mathcal{I} \times \mathcal{J}.$$

Here we make a notation that $[M] = \{1, 2, \dots, M\}$. To make the alternative hypothesis mathematically concrete, in H_1 we mean that for any $(i, j) \in \mathcal{I} \times \mathcal{J}$ and $(k, l) \notin \mathcal{I} \times \mathcal{J}$, and for any $c \in \mathbb{R}$, we have $\mathbb{P}(X_{ij} \geq c) \geq \mathbb{P}(X_{kl} \geq c)$. And there exists $c = c_0 \in \mathbb{R}$ that the equality does not hold. This problem is considered in [Butucea and Ingster \(2013\)](#) with test statistics as

$$\text{SUM}(\mathbf{X}) = \sum_{i \in [M]} \sum_{j \in [N]} X_{ij}, \quad (1.1)$$

and

$$\text{SCAN}_{m,n}(\mathbf{X}) = \max_{\mathcal{I} \subset [M], |\mathcal{I}|=m} \max_{\mathcal{J} \subset [N], |\mathcal{J}|=n} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} X_{ij}. \quad (1.2)$$

One rejects H_0 when either of the two values go beyond some corresponding predefined thresholds. Note that in order to make the scan statistic (1.2) work, exact knowledge of the size of \mathcal{I} and \mathcal{J} is required, which here is m and n . When (m, n) are unknown, a Bonferroni testing procedure is used - that is, test all combinations of (m, n) of interest, and reject H_0 when one of the tests indicates to reject.

The method of [Butucea and Ingster \(2013\)](#) relies on parametric assumptions, therefore [Arias-Castro and Liu \(2017\)](#) considered calibrate the p-value by permutation. Since the permutation test is based on (1.2), full knowledge of (m, n) is necessary. That is, it can only distinguish between H_0 and

$$H_1(m, n): \text{There exists } \mathcal{I} \subset [M] \text{ and } \mathcal{J} \subset [N] \text{ such that } X_{ij} \text{ are larger than usual when } (i, j) \in \mathcal{I} \times \mathcal{J}, \text{ where } |\mathcal{I}| = m \text{ and } |\mathcal{J}| = n.$$

In this paper we adapt the permutation test framework, and develop a Bonferroni testing procedure in order to deal with the case when (m, n) are unknown.

Contribution 1.1. We develop a Bonferroni testing procedure based on the permutation test by [Arias-Castro and Liu \(2017\)](#). We show that the testing procedure

(at the first order) is asymptotically as powerful as the test by [Arias-Castro and Liu \(2017\)](#) and [Butucea and Ingster \(2013\)](#). This is analyzed and proved under some standard exponential family parameter assumptions.

As permutation test is computationally hard to calibrate (usually the total number of permutations are increasing exponentially with the sample size), in practice a Monte Carlo calibration method is executed. This requires independent sampling from the group of permutation patterns for a decent number of times. However still, due to the computational complexity of (1.2), performing permutation tests on all combinations of (m, n) could be extremely consuming in time and computational power. We construct a subset of $[M] \times [N]$, such that by performing permutation test on all (m, n) inside this subset, we can still detect the existence of the elevated submatrix with no sacrifice of the power on the first order.

Contribution 1.2. We propose a power-preserving fast test based on the Bonferroni permutation test, with the Bonferroni procedures working on a proper approximate net of $[M] \times [N]$. We show that this test is as powerful as the Bonferroni permutation test on all pairs of (m, n) in $[M] \times [N]$. This is also analyzed and proved under some standard exponential family parameter assumptions.

1.2. *More Related work.* There are works focusing on the localization theory side of this problem, studying the existence of consistent estimators of the elevated matrix under some parametric setup ([Kolar et al., 2011](#); [Chen and Xu, 2016](#); [Hajek et al., 2017a,b](#)). With the realization of the importance of computational efficiency, there is a stream of work line considering the trade-off between statistical and computational power ([Balakrishnan et al., 2011](#); [Cai et al., 2017](#); [Chen and Xu, 2016](#); [Ma and Wu, 2015](#)). Several computationally efficient (with polynomial computational complexity with respect to data size) submatrix localization algorithm are proposed, including convex optimization ([Chen and Xu, 2016](#); [Chi et al., 2017](#)) and spectral method ([Cai et al., 2017](#)).

Another active research line worth mentioning here is on the stochastic block model (SBM). In the setup of SBM the observation is a graph, with edges independently connected. See [Holland et al. \(1983\)](#) for a detailed introduction. The detection problem is to distinguish between an Erdős-Rényi graph and a graph with groups of nodes which nodes are more likely to connect within groups compared to across groups. The localization problem is to cluster the nodes by the closeness of their connection.

If the adjacency matrix of the graph is considered, the problem shares many properties with submatrix detection and localization (the adjacency matrix is symmetric with independent upper triangle nodes, compared with submatrix localization problem). There are works considering the existence of consistency detectors ([Zhang and Zhou, 2016](#)), the existence of consistency clustering methods ([Mossel et al., 2015](#)), semi-definite programming ([Chen and Xu, 2016](#); [Abbe et al., 2016](#)), and spectral methods ([Chaudhuri et al., 2012](#); [McSherry, 2001](#)).

1.3. *Content.* The paper is arranged as follows. Section 2 introduces the parametric setup, the detection boundaries set up by [Butucea and Ingster \(2013\)](#), as well as the permutation test by [Arias-Castro and Liu \(2017\)](#). Section 3 describes the Bonferroni-type testing procedure as well as its theoretical property. Section 4 shows the construction of an approximation net in order to speed up the testing

process, and the associated theoretical results. The numerical experiments are in Section 5.

2. The detection boundaries

We establish the minimax framework corresponding to the hypothesis testing problem raised in Section 1. First we introduce the one-parameter natural exponential family considered throughout this paper, which is also the parametric family used in the analysis of Arias-Castro et al. (2018); Butucea and Ingster (2013); Arias-Castro and Liu (2017). To define such a distribution family, first consider a distribution ν with mean zero and variance 1, and assume that its moment generating function $\phi(\theta) < \infty$ for some $\theta > 0$. Let $\theta_* = \sup\{\theta : \phi(\theta) < \infty\}$, and the family is parameterized by $\theta \in [0, \theta_*)$ with density function

$$f_\theta(x) = \exp\{x\theta - \log \phi(\theta)\}. \quad (2.1)$$

The density function is with respect to ν . Note that when $\theta = 0$, $f_0(x) = 1$ which corresponds to ν . By varying the choice of ν , this parametric model includes several common parametric models such as normal family ($\nu = \mathcal{N}(0, 1)$), Poisson family ($\nu = \text{Pois}(1) - 1$) and Rademacher family ($\nu = 2\text{Rade}(0.5) - 1$).

An important property of f_θ is the stochastic monotonicity with respect to θ . This enables us to model the previous hypothesis testing problem with ν acting as noise distribution and f_θ as the distribution of those unusually large entries. In details, we consider

$$\begin{aligned} H_0: X_{ij} \text{ are IID following } \nu \text{ for all } i \in [M] \text{ and } j \in [N] \\ \text{vs} \\ H_1: \text{There exists } \mathcal{I} \subset [M] \text{ and } \mathcal{J} \subset [N] \text{ such that } X_{ij} \sim f_{\theta_{ij}}, \theta_{ij} > \theta_{\ddagger} > 0 \text{ when} \\ (i, j) \in \mathcal{I} \times \mathcal{J}. \end{aligned}$$

Here the parameter θ_{\ddagger} is the lower bound for all the θ_{ij} inside the raised submatrix, and is acting as the role of signal-noise ratio. In this context, if the potential submatrix's size (m, n) is known, the corresponding $H_1(m, n)$ is described as follows.

$$\begin{aligned} H_1(m, n): \text{There exists } \mathcal{I} \subset [M] \text{ and } \mathcal{J} \subset [N] \text{ such that } X_{ij} \sim f_{\theta_{ij}}, \theta_{ij} > \theta_{\ddagger} > 0 \\ \text{when } (i, j) \in \mathcal{I} \times \mathcal{J}, \text{ where } |\mathcal{I}| = m \text{ and } |\mathcal{J}| = n. \end{aligned}$$

In order to perform the hypothesis testing task between H_0 and $H_1(m, n)$ under this parameterization framework, Butucea and Ingster (2013) developed the following.

Theorem 2.1 (Butucea and Ingster, 2013). *Consider an exponential model as described in (2.1), with ν having finite fourth moment. Assume that*

$$M, N, m, n \rightarrow \infty, \quad m = o(M), n = o(N), \quad \frac{\log(M \vee N)}{m \wedge n} \rightarrow 0. \quad (2.2)$$

For any $\alpha > 0$ fixed, the sum test based on (1.1) is asymptotically powerful if

$$\frac{\theta_{\ddagger} mn}{\sqrt{MN}} \rightarrow \infty, \quad (2.3)$$

and the scan test based on (1.2), is asymptotically powerful if

$$\liminf \frac{\theta_{\ddagger} \sqrt{mn}}{\sqrt{2(m \log(M/m) + n \log(N/n))}} > 1. \quad (2.4)$$

Conversely, if $m = O(n)$ and $\log M = O(\log N)$, and

$$\frac{\theta_{\dagger}mn}{\sqrt{MN}} \rightarrow 0, \quad \liminf \frac{\theta_{\dagger}\sqrt{mn}}{\sqrt{2(m \log(M/m) + n \log(N/n))}} < 1, \quad (2.5)$$

any test with level α has limiting power at most α .

The test rejects H_0 when the sum test statistic (1.1) or scan test statistic (1.2) is larger than some pre-defined threshold. While these tests heavily depend on parameterization, distribution-free methods such as permutation test are proposed to tackle this problem. A permutation test based on the scan statistic (1.2) is analyzed by [Arias-Castro and Liu \(2017\)](#). Here a permutation pattern on set $[M] \times [N]$ is denoted by π , and define

- Π_1 as the collection of permutations that entries are permuted within their row;
- Π_2 as the collection of all permutations.

To illustrate the difference of Π_1 and Π_2 , let $\mathbf{X}_{\pi} = \{X_{\pi(i,j)}\}$, consider

$$\mathbf{X} = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 5 & 6 & 7 & 8 \\ 9 & 10 & 11 & 12 \end{pmatrix}$$

and we will have

$$\mathbf{X}_{\pi_1} = \begin{pmatrix} 1 & 4 & 2 & 3 \\ 6 & 8 & 7 & 5 \\ 10 & 12 & 9 & 11 \end{pmatrix}, \mathbf{X}_{\pi_2} = \begin{pmatrix} 10 & 8 & 2 & 9 \\ 6 & 3 & 11 & 1 \\ 12 & 4 & 5 & 7 \end{pmatrix}$$

for some $\pi_1 \in \Pi_1, \pi_2 \in \Pi_2$. We calculate the scan statistic on the permuted data \mathbf{X}_{π} , and compare to the one from the original data \mathbf{X} . The permutation p -value is defined as

$$\mathfrak{P}_{m,n}(\mathbf{X}) = \frac{|\{\pi \in \Pi : \text{SCAN}_{m,n}(\mathbf{X}_{\pi}) \geq \text{SCAN}_{m,n}(\mathbf{X})\}|}{|\Pi|}. \quad (2.6)$$

Here $\Pi = \Pi_1$ or Π_2 . The permutation test rejects the null hypothesis when $\mathfrak{P}_{m,n}(\mathbf{X})$ is smaller than the pre-determined level. This test has the following property.

Theorem 2.2 ([Arias-Castro and Liu, 2017](#), Theorem 2). *Consider an exponential model as described in (2.1), assume (2.2) and*

$$\log^3(M \vee N)/(m \wedge n) \rightarrow 0 \quad (2.7)$$

and that either (i) ν has support bounded from above, or (ii) $\max_{i,j} \theta_{ij} \leq \bar{\theta}$ for some $\bar{\theta} < \theta_*$ fixed. Additionally if $\Pi_i = \Pi_1$, we require that $\phi(\theta) < \infty$ for some $\theta < 0$. Then the permutation scan test based on (2.6), at any fixed level $\alpha > 0$, has limiting power 1 when (2.4) holds.

This theorem shows that permutation test has the same first-order asymptotic power compared to the parametric test described in Theorem 2.1, under some extra mild conditions.

3. Bonferroni permutation test

We now consider the case where the size of the submatrix to be detected is not specified. Recall the definition of H_1 and $H_1(m, n)$ in Section 2. By realizing the fact that H_1 is true if and only if there exists some (m, n) such that $H_1(m, n)$ is true, we can perform test on $H_1(m, n)$ for all pairs of (m, n) , and use Bonferroni correction in order to control the type I error.

We adapt the distribution-free permutation test of Arias-Castro and Liu (2017) here. For each pair of (m, n) , calculate $\mathfrak{P}_{m,n}(\mathbf{X})$, and calculate the final Bonferroni corrected p -value as

$$\mathfrak{P}(\mathbf{X}) = \min(MN \min_{m,n} \mathfrak{P}_{m,n}(\mathbf{X}), 1). \quad (3.1)$$

One rejects H_0 when $\mathfrak{P}(\mathbf{X})$ is less than some pre-determined level α . Due to the property of Bonferroni type of tests, this test has level α if each test concerning H_0 and $H_1(m, n)$ has level α/MN , which is a proved fact in Arias-Castro and Liu (2017), regardless of the dependencies between tests.

Being a conservative method in multiple testing, Bonferroni method usually loses statistical power in exchange for controlling the family wise error rate. However in some cases (for example, Arias-Castro and Verzelen, 2014; Butucea and Ingster, 2013), the Bonferroni procedure achieves the same first order asymptotic power as the scan test without knowledge of the submatrix size. We illustrate the same phenomenon in this case.

Theorem 3.1. *Consider an exponential model as described in (2.1). Assume that there exists a pair of (m, n) such that $H_1(m, n)$ is true, and all the assumptions in Theorem 2.2 are satisfied, which are:*

$$\begin{aligned} M, N, m, n &\rightarrow \infty, \quad m = o(M), \quad n = o(N), \\ \frac{\log(M \vee N)}{m \wedge n} &\rightarrow 0, \quad \log^3(M \vee N)/(m \wedge n) \rightarrow 0 \end{aligned} \quad (3.2)$$

and that either (i) ν has support bounded from above, or (ii) $\max_{i,j} \theta_{ij} \leq \bar{\theta}$ for some $\bar{\theta} < \theta_*$ fixed. Additionally if $\Pi_i = \Pi_1$, we require that $\phi(\theta) < \infty$ for some $\theta < 0$. Then

$$\mathfrak{P}(\mathbf{X}) \rightarrow 0 \quad (3.3)$$

in probability.

The interpretation of this theorem is as follows. Assume we are under $H_1(m, n)$. With full knowledge of (m, n) , if the permutation test by (2.6) can successfully detect the submatrix with high probability, the test by (3.1) will reject H_0 with high probability, without any knowledge of (m, n) .

Note that if (2.5) and its associated assumptions hold true, any test trying to distinguish H_0 and H_1 with level α will have limiting power at most α , due to the fact that any test distinguishing H_0 versus H_1 can be used to distinguish H_0 and $H_1(m, n)$, thus having limiting power at most α thanks to Theorem 2.1.

4. A power-preserving fast test

We build our testing framework on the permutation test, which is by nature a computationally intensive method. Calculation of scan statistic (1.2) is NP-hard, and the total number of permutations $|\Pi|$ will skyrocket when the number of data

increases. In practice the scan statistic is calculated by LAS algorithm proposed by [Shabalin et al. \(2009\)](#), as did in previous literature ([Butucea and Ingster, 2013](#); [Arias-Castro and Liu, 2017](#)), and the permutation test is done by Monte-Carlo sampling. In detail, a large number B is fixed, and permutations $\{\pi_1, \dots, \pi_B\}$ is the IID sample from uniform distribution on $|\Pi_i|$. Then $\mathfrak{P}_{m,n}$ is approximated by

$$\hat{\mathfrak{P}}_{m,n}(\mathbf{X}) = \frac{|\{i : \text{SCAN}_{m,n}(\mathbf{X}_{\pi_i}) \geq \text{SCAN}_{m,n}(\mathbf{X})\}| + 1}{B + 1}. \quad (4.1)$$

As illustrated in [Butucea and Ingster \(2013\)](#); [Arias-Castro and Liu \(2017\)](#), the calculation of the scan statistic (1.2) is already difficult, even if the size (m, n) is known. Now with the submatrix size unknown, we have added the difficulty since the scan statistic under all possible combinations of (m, n) , will be calculated during the Bonferroni process. In principle, the Bonferroni method requires going over all submatrix sizes, but we only scan a carefully chosen subset to lighten up the computational burden. Inspired by [Arias-Castro et al. \(2005, 2018\)](#), in which the authors scan for anomalous data interval within all possible intervals using a dyadic representation, we illustrate the construction of such subset on $[M] \times [N]$, and show that the first-order statistical power is preserved.

4.1. An approximation net. The subset of $[M] \times [N]$ we are going to construct in order to approximate the elements in $[M] \times [N]$ is called an approximate net. We first construct one-dimensional approximate net on $[M]$.

We start by the following definition.

Definition 4.1. A binary expansion of an integer c is a sequence $\{a_i(c)\}$ with $a_i(c) \in \{0, 1\}$ and $i \leq \lfloor \log_2 c \rfloor$, such that

$$c = \sum_{i=0}^{\lfloor \log_2 c \rfloor} a_i(c) 2^i. \quad (4.2)$$

After representing an integer in the binary numeral system, one may approximate this integer by keeping the first k digits of its binary expansion.

Definition 4.2. A k -binary approximation of an integer c is the integer c' such that

$$a_i(c') = a_i(c) \mathbb{1}\{i \geq \lfloor \log_2 c \rfloor - k + 1\}. \quad (4.3)$$

To find the k -binary approximation of c , represent c in its binary expansion, keep the first k digits, and shrink the rest to zero. Finally calculate c' by the formula in Definition 4.1. It captures the main part of the integer and the difference could be controlled by k . The method is closed related to the binary tree representation of integers, where an integer is represented to the root at distance k from the first non-zero node. The following lemma gives an upper bound of the difference rate.

Lemma 4.3. *If c' is a k -binary approximation of c , then*

$$\frac{c - c'}{c} \leq 2^{1-k}. \quad (4.4)$$

Note that the difference rate is only associated with k , but not with the value of c . Therefore if we apply k -binary approximation to a collection of integers, the difference rate will be controlled *uniformly* among all the integers in the collection by the choice of k .

TABLE 4.1. Approximation net for 1024

Binary	Decimal	Binary	Decimal	Binary	Decimal
1000000000	1024	0001000000	128	00000010000	16
0111000000	896	00001110000	112	00000001110	14
0110000000	768	00001100000	96	00000001100	12
0101000000	640	00001010000	80	00000001010	10
0100000000	512	00001000000	64	00000001000	8
0011100000	448	00000111000	56	00000000111	7
0011000000	384	00000110000	48	00000000110	6
0010100000	320	00000101000	40	00000000101	5
0010000000	256	00000100000	32	00000000100	4
0001110000	224	00000011100	28	00000000011	3
0001100000	192	00000011000	24	00000000010	2
0001010000	160	00000010100	20	00000000001	1

Now we construct the approximation net of $[M]$ based on k -binary approximation.

Definition 4.4. An approximation net $S_k(M)$, based on k -binary approximation, of set $[M]$, is defined as

$$S_k(M) = \{c' : c' \text{ is a } k\text{-binary approximation of some } c \in [M]\} \quad (4.5)$$

To better illustrate the approximation net, Table 4.1 shows the approximation net of 1024 with $k = 3$.

The cardinality of $S_k(M)$ can be much less than M if k is chosen properly. For example, set $k = \log_2 \log_2 M + 1$ and it can be shown that $|S_k(M)| = O((\log_2 M)^2)$. Note that in this case $k \rightarrow \infty$ when $M \rightarrow \infty$, and by Lemma 4.3 we know that for every $c \in [M]$, there exists some $c' \in S_k(M)$ such that $c' = (1 + o(1))c$, and $o(1)$ is uniform among $[M]$.

4.2. *Test power under approximation nets.* Based on the one-dimensional approximation net defined in Definition 4.4, we can similarly extend the idea to sets of two-dimensional integer pairs. We perform the Bonferroni-type testing procedure on $S_{k_M}(M) \times S_{k_N}(N)$, instead of $[M] \times [N]$. In detail, we use the following Bonferroni corrected p -value:

$$\mathfrak{P}_{k_M, k_N}(\mathbf{X}) = \min(|S_{k_M}(M)| |S_{k_N}(N)| \min_{(s,t) \in S_{k_M}(M) \times S_{k_N}(N)} \mathfrak{P}_{s,t}(\mathbf{X}), 1). \quad (4.6)$$

The idea is to use the property of approximation net to eliminate a significant portion of calculation by reducing the scanning region of the Bonferroni process, while keeping the accuracy through choosing a proper pair of (k_M, k_N) . Assume we are under $H_1(m, n)$. When setting $k_M, k_N \rightarrow \infty$, there is a pair of $(m', n') \in S_{k_M}(M) \times S_{k_N}(N)$ that is close enough to (m, n) , and $\mathfrak{P}_{m', n'}(\mathbf{X})$ will converge to zero fast enough such that brings the Bonferroni corrected p -value to zero as well.

The following theorem describes the asymptotic power of the Bonferroni test on the approximate net.

Theorem 4.5. *Assume all the assumptions in Theorem 3.1 hold (specifically, (3.2) and (i) and (ii) following it). Further set $k_M, k_N \rightarrow \infty$. Then*

$$\mathfrak{P}_{k_M, k_N}(\mathbf{X}) \rightarrow 0 \quad (4.7)$$

in probability.

5. Numerical experiments

We use simulations to verify our theoretical findings.¹ The main purpose of the simulation is to illustrate the proposed test's behavior under the alternative hypothesis. We adapt the simulation setup in Arias-Castro and Liu (2017) due to the similarity of the research problem and simulation purpose.

The data \mathbf{X} is generated with IID samples from ν for all entries except $[m] \times [n]$, where the entries are IID from f_θ for some $\theta > 0$. Then the data goes through the proposed test progress and the p -value is recorded. The above process is repeated 100 times for fixed θ in order to see the stochastic behavior of the p -values.

To see how the theory works, we define θ_{crit} as follows.

$$\theta_{\text{crit}} = \sqrt{\frac{2(m \log(M/m) + n \log(N/n))}{mn}}. \quad (5.1)$$

And we slowly increase θ from $0.625 \times \theta_{\text{crit}}$ to $1.5 \times \theta_{\text{crit}}$, with step size $0.125 \times \theta_{\text{crit}}$. This is aim to examine the behavior of the test around $\theta = \theta_{\text{crit}}$, which is claimed to be the convergence threshold from our theory.

The permutation test is calibrated by Monte-Carlo which is described in (4.1), with $B = 500$. As mentioned in Section 4, the calculation of $\text{SCAN}_{m,n}(\mathbf{X})$ is NP-hard in theory. LAS algorithm from Shabalín et al. (2009) is an approximating algorithm to calculate the scan statistic, however due to its hill-climbing optimization process, it suffers from being stuck inside local minimums. We re-initiate the LAS algorithm several times with random initialization, and return the result with the largest output, in order to prevent being stuck at local minimums. Both permutation methods are examined in the analysis. The data size is set as $(M, N) = (200, 100)$ and we examine 2 anomaly sizes, namely $(m, n) = (10, 15)$ and $(30, 10)$. For the approximate net, we set $(k_M, k_N) = (\lceil \log_2(\log_2(M)) \rceil, \lceil \log_2(\log_2(N)) \rceil)$

We choose two representative distributions as ν , which is standard normal and centralized $Pois(1)$. The corresponding f_θ is $\mathcal{N}(\theta, 1)$ and $Pois(e^\theta) - 1$. From Figure 5.1 and 5.2 we see that when $\theta \leq \theta_{\text{crit}}$, the p -values are generally close to 1 due to the conservative property of the Bonferroni methods. However, the p -value shrinks to simulation lower bound at

$$|S_k(M)| \times |S_k(N)| / (B + 1) = 14 \times 12 / (500 + 1) \approx 0.34 \quad (5.2)$$

quickly once θ passes θ_{crit} , in both permutation patterns and distribution setups. This is because of the fact that all the permutation p -values are larger than or equal to $1/(B + 1)$, which makes the Bonferroni corrections no less than $|S_k(M)| \times |S_k(N)| / (B + 1)$. Although reaching the lower bound indicates that no permutation during the test generates the test statistic larger than the one from the original data, it is not very persuasive in the sense to show that the p -values are converging

¹The code used in this section is available in <https://github.com/nozoeli/bonferroniSubmatrix>.

to 0. In order to verify this, we need a larger number of Monte Carlo simulations, namely a larger B , such that $|S_k(M)| \times |S_k(N)| / (B + 1)$ is close to 0.

We perform an additional experiment by setting $B = 5000$, and focus on the pair of $(m', n') \in S_{k_M}(M) \times S_{k_N}(N)$, such that $m' = \min\{p \in S_{k_M}(M) | p > m\}$ and $n' = \min\{q \in S_{k_N}(N) | q > n\}$. This provides us an upper bound of the Bonferroni permutation p -value, since by (4.6),

$$\mathfrak{P}_{k_M, k_N}(\mathbf{X}) \leq \min(|S_{k_M}(M)| |S_{k_N}(N)| \mathfrak{P}_{m', n'}(\mathbf{X}), 1). \tag{5.3}$$

We illustrate the relationship between this upper bound of the Bonferroni permutation p -value and the signal level. Figure 5.3 and Figure 5.4 shows that the upper bound converges to 0 once the signal level θ goes beyond θ_{crit} , and this confirms our theoretical findings.

6. Technical proofs

In this section we present the proofs to the theorems in the previous sections. Unless separately declared, the inequalities hold with high probability (or with probability $1 - o(1)$ as data matrix size $M, N \rightarrow \infty$).

The following lemma is at the center of our argument.

Lemma 6.1 (Arias-Castro et al., 2018, Lemma 2, Bernstein’s inequality for sampling without replacement). *Let (Z_1, \dots, Z_m) be obtained by sampling without replacement from a given a set of real numbers $\{z_1, \dots, z_J\} \subset \mathbb{R}$. Define $z_{\max} = \max_j z_j$, $\bar{z} = \frac{1}{J} \sum_j z_j$, and $\sigma_z^2 = \frac{1}{J} \sum_j (z_j - \bar{z})^2$. Then the sample mean $\bar{Z} = \frac{1}{m} \sum_i Z_i$ satisfies*

$$\mathbb{P}(\bar{Z} \geq \bar{z} + t) \leq \exp \left[- \frac{mt^2}{2\sigma_z^2 + \frac{2}{3}(z_{\max} - \bar{z})t} \right], \quad \forall t \geq 0.$$

The lemma is a result from Serfling (1974). We also refer the reader to Bardenet and Maillard (2015); Boucheron et al. (2013) for further details of the lemma.

6.1. *Proof of Theorem 3.1.* We first show that the test has level α . To show this, we fix a pair of (m, n) and show that $\mathbb{P}(\mathfrak{P}_{m, n}(X) \leq \alpha) \leq \alpha$. By the standard argument on the level of Bonferroni test, we will finish the proof on the level. Assuming the null is true, $\mathbf{scan}_{m, n}(\pi(X))$ has the same distribution with $\mathbf{scan}_{m, n}(X)$ under either permutation methods. Therefore we define

$$T_k = \mathbf{scan}_{m, n}(\pi_k(X)), k \in [(MN)!],$$

and assume $T_{k_0} = \mathbf{scan}_{m, n}(X)$, then $\text{rank}(T_{k_0})$ is uniformly distributed on $[(MN)!]$ (if the ties are broken randomly). We have

$$\mathbb{P}(\mathfrak{P}_{m, n}(X) \leq \alpha) \leq \mathbb{P}(\text{rank}(T_{i_0}) \leq \lfloor \alpha(MN)! \rfloor) \leq \frac{\lfloor \alpha(MN)! \rfloor}{MN!} \leq \alpha.$$

Therefore all we need to show is that the p -value tends to zero under the alternative.

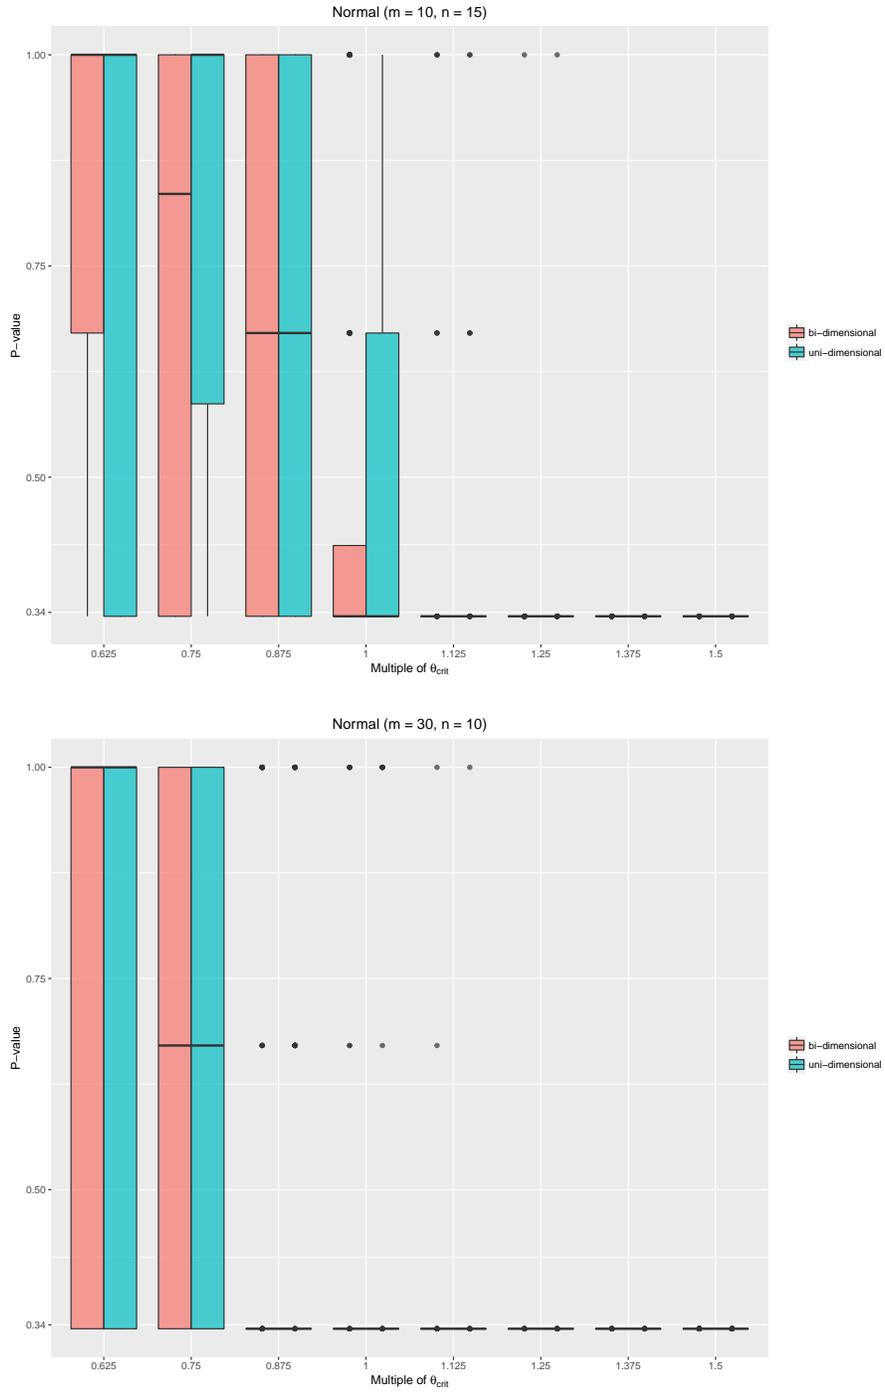
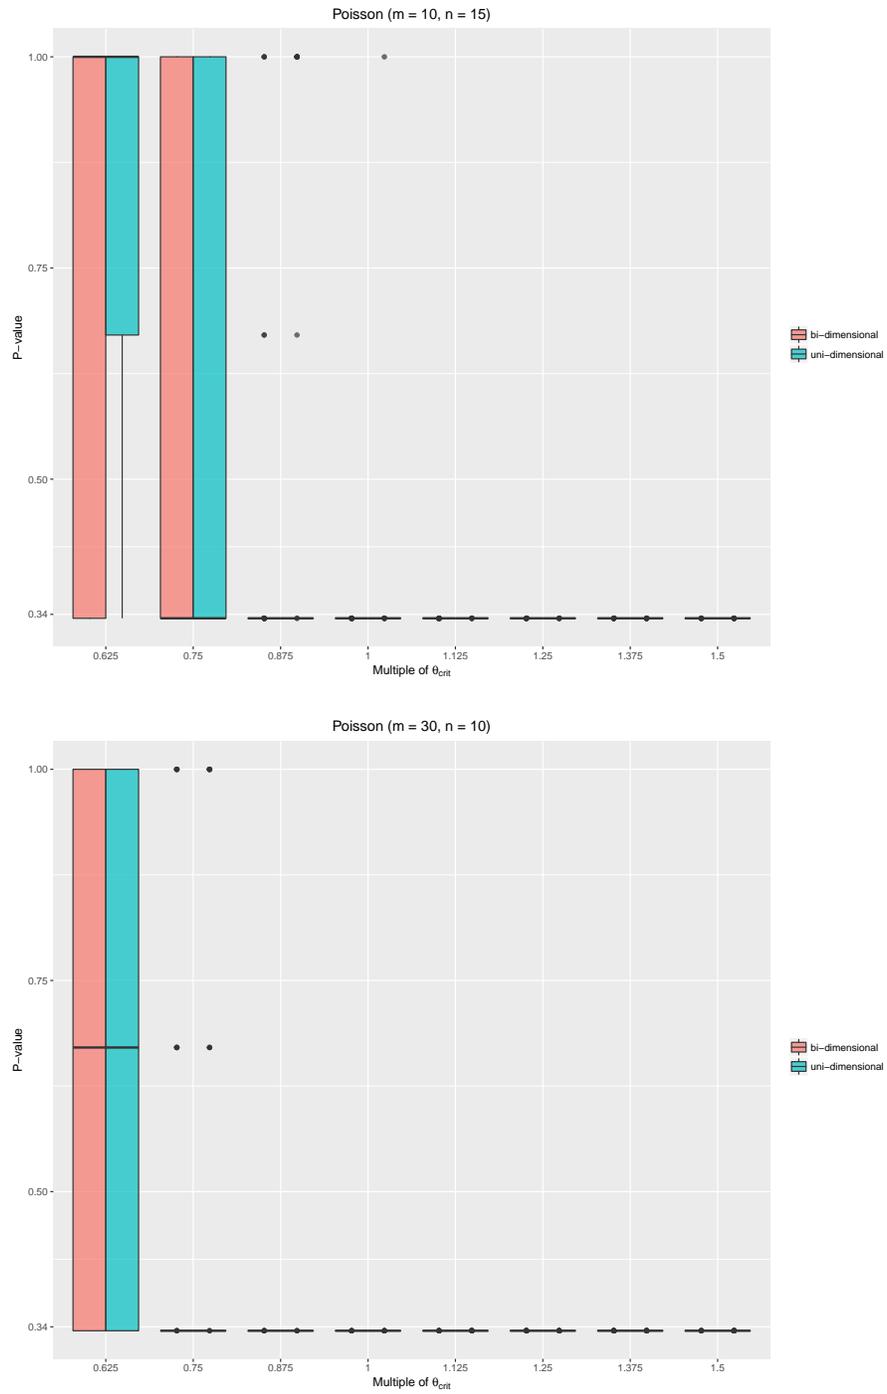


FIGURE 5.1. p -values of proposed test in the normal model

FIGURE 5.2. p -values of proposed test in the Poisson model

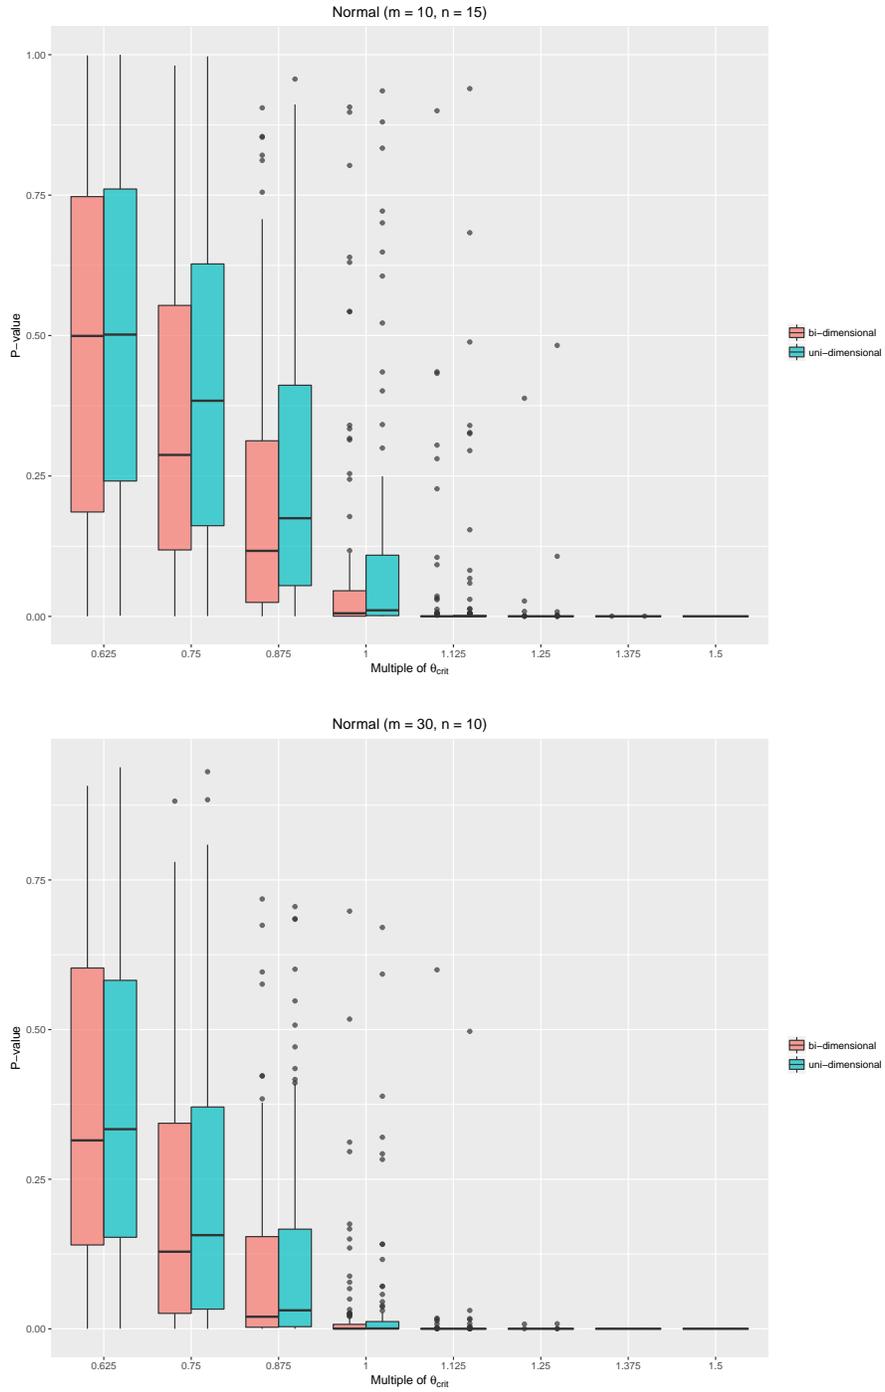


FIGURE 5.3. Upper bounds of p -values of proposed test in the normal model

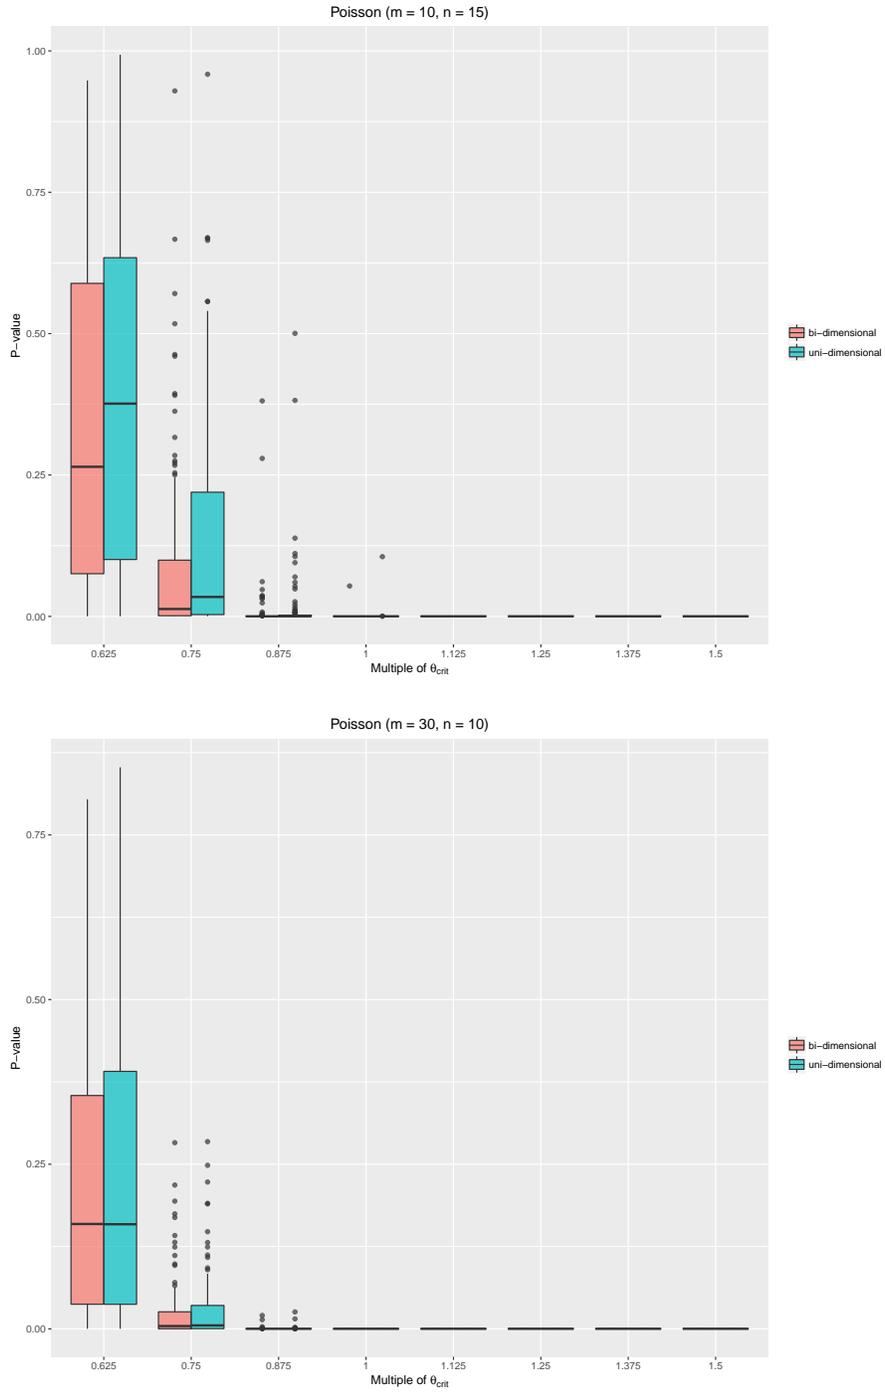


FIGURE 5.4. Upper bounds of p -values of proposed test in the Poisson model

6.1.1. *Bidimensional permutation.* We start with considering bidimensional permutation. Starting from here we let S^* be the true anomaly submatrix, and (m, n) as its row and column size. For a submatrix index set S , let $Y_S(X) = \sum_{(i,j) \in S} X_{ij}$. Now fix some $S \in \mathbb{S}_{mn}$ and uniformly sample a permutation π . Here \mathbb{S}_{mn} is the collection of row/column indexes of all submatrices of size $m \times n$. Also, if we conditional a realization of \mathbf{X} , say \mathbf{x} , by Lemma 6.1, we have

$$\mathbb{P}\left(\frac{Y_S(\pi(\mathbf{x}))}{mn} - \bar{\mathbf{x}} > t\right) \leq \exp\left(-\frac{t^2 mn}{2\sigma_{\mathbf{x}}^2 + \frac{2}{3}(\mathbf{x}_{\max} - \bar{\mathbf{x}})t}\right),$$

where \mathbf{x}_{\max} is the largest value in the data matrix \mathbf{x} , and $\sigma_{\mathbf{x}}^2$ is the sample variance of the data if \mathbf{x} is treated as a one dimensional data vector. Notice that the probability is on the permutation process. Apply a union bound, and set $t = \mathbf{scan}(\mathbf{x})/mn - \bar{\mathbf{x}}$, then we have

$$\begin{aligned} \mathfrak{P}(\mathbf{x}) &\leq MN \mathfrak{P}_{m,n}(\mathbf{x}) \\ &\leq MN |\mathbb{S}_{mn}| \exp\left(-\frac{mn(\mathbf{scan}_{m,n}(\mathbf{x})/mn - \bar{\mathbf{x}})^2}{2\sigma_{\mathbf{x}}^2 + \frac{2}{3}(\mathbf{x}_{\max} - \bar{\mathbf{x}})(\mathbf{scan}_{m,n}(\mathbf{x})/mn - \bar{\mathbf{x}})}\right). \end{aligned} \quad (6.1)$$

Note that this inequality holds for all the realizations of \mathbf{X} , thus we allow \mathbf{X} to change and focusing on the quantity

$$MN |\mathbb{S}_{mn}| \exp\left(-\frac{mn(\mathbf{scan}_{m,n}(\mathbf{X})/mn - \bar{\mathbf{X}})^2}{2\sigma_{\mathbf{X}}^2 + \frac{2}{3}(\mathbf{X}_{\max} - \bar{\mathbf{X}})(\mathbf{scan}_{m,n}(\mathbf{X})/mn - \bar{\mathbf{X}})}\right). \quad (6.2)$$

We show that this quantity is $o_P(1)$.

We start with bounding $\bar{\mathbf{X}}$ by rewriting $\bar{\mathbf{X}}$ as

$$\bar{\mathbf{X}} = \frac{\sum_{i,j} X_{ij}}{MN} = \frac{mn}{MN} \cdot \frac{\sum_{(i,j) \in S^*} X_{ij}}{mn} + \frac{MN - mn}{MN} \cdot \frac{\sum_{(i,j) \notin S^*} X_{ij}}{MN - mn}. \quad (6.3)$$

With θ in the anomalous submatrix bounded from above, or the support of f_θ being bounded, the first term is $o_P(1)$ since $mn = o(MN)$. By Law of Large Numbers, the second term is $o_P(1)$ given that distribution f_0 has mean zero. So

$$\bar{\mathbf{X}} = o_P(1). \quad (6.4)$$

Following the proof of Theorem 1 in Arias-Castro et al. (2018), we can bound \mathbf{X}_{\max} with

$$\mathbb{P}(\mathbf{X}_{\max} - \bar{\mathbf{X}} < \frac{3}{c} \log(MN)) \rightarrow 1. \quad (6.5)$$

with $c \in (0, \theta_* - \bar{\theta})$. Define the event $\mathcal{A} = \{\mathbf{X}_{\max} - \bar{\mathbf{X}} < \frac{3}{c} \log(MN)\}$. All the following arguments are conditional on \mathcal{A} . But since \mathcal{A} happens with probability tending to 1, all the conditional high probability events will happen unconditionally with high probability as well.

We do similar operations to bound $\sigma_{\mathbf{X}}^2$ as follows.

$$\sigma_{\mathbf{X}}^2 = \frac{1}{MN} \sum_{i,j} (X_{ij} - \bar{\mathbf{X}})^2 \leq \frac{1}{MN} \sum_{i,j} X_{ij}^2 \quad (6.6)$$

$$= \frac{mn}{MN} \cdot \frac{\sum_{X_{ij} \in S^*} X_{ij}^2}{mn} + \frac{MN - mn}{MN} \cdot \frac{\sum_{X_{ij} \notin S^*} X_{ij}^2}{MN - mn}. \quad (6.7)$$

The first term is $o_P(1)$ if distribution f_θ has finite second moment, which can be derived from the assumption. The second term is $1 + o_P(1)$ by Law of Large Numbers and Slutsky's Lemma. Therefore

$$\sigma_{\mathbf{X}}^2 = 1 + o_P(1) \quad (6.8)$$

Finally we bound $\text{scan}_{m,n}(\mathbf{X})/mn - \bar{\mathbf{X}}$ as a whole. By the definition of the scan statistic,

$$\frac{\text{scan}_{m,n}(\mathbf{X})}{mn} - \bar{\mathbf{X}} \geq \frac{Y_{S^*}(\mathbf{X})}{mn} - \bar{\mathbf{X}} = \bar{\mathbf{X}}_{S^*} - \bar{\mathbf{X}}, \quad (6.9)$$

here $\bar{\mathbf{X}}_{S^*} = \sum_{(i,j) \in S^*} X_{ij}/mn$ represents the average in the submatrix indexed by S^* . Rewrite $X_{ij} = E(X_{ij}) + Z_{ij}$ for $X_{ij} \in S^*$, where Z_{ij} has mean zero and bounded second moment. By Law of Large numbers, as well as (7) of [Arias-Castro and Liu \(2017\)](#) (which is, $\mathbb{E} X_{ij} \geq \theta_{\dagger}$ for $X_{ij} \in S^*$),

$$\bar{\mathbf{X}}_{S^*} = \frac{1}{mn} \sum_{X_{ij} \in S^*} \mathbb{E} X_{ij} + O_P\left(\frac{1}{\sqrt{mn}}\right) \geq \theta_{\dagger} + O_P\left(\frac{1}{\sqrt{mn}}\right). \quad (6.10)$$

Note that $\bar{\mathbf{X}} = O_P(1/\sqrt{MN})$ and $mn = o(MN)$,

$$\frac{\text{scan}_{m,n}(\mathbf{X})}{mn} - \bar{\mathbf{X}} \geq \theta_{\dagger} + O_P\left(\frac{1}{\sqrt{mn}}\right) \quad (6.11)$$

By (2.4) we know that $\sqrt{mn}\theta_{\dagger} \rightarrow \infty$, or $\theta_{\dagger} \gg 1/\sqrt{mn}$, so we can rewrite the equation above as

$$\frac{\text{scan}_{m,n}(\mathbf{X})}{mn} - \bar{\mathbf{X}} \geq \theta_{\dagger}(1 + o_P(1)). \quad (6.12)$$

Plug in (6.5), (6.8) and (6.12) into (6.2), we have

$$\mathfrak{P}(\mathbf{X}) \leq MN|\mathbb{S}_{mn}| \exp\left(-\frac{(1 + o_P(1))\theta_{\dagger}^2 mn}{2(1 + o_P(1)) + \frac{2}{c}(\log MN)\theta_{\dagger}(1 + o_P(1))}\right). \quad (6.13)$$

Assume in (2.4),

$$\liminf \frac{\theta_{\dagger}^2 mn}{2(m \log \frac{M}{m} + n \log \frac{N}{n})} \geq 1 + \epsilon \quad (6.14)$$

with some constant $\epsilon > 0$. Then eventually with high probability we can bound the exponential part of (6.13) as

$$-\frac{(1 + o_P(1))\theta_{\dagger}^2 mn}{2(1 + o_P(1)) + \frac{2}{c}(\log MN)\theta_{\dagger}(1 + o_P(1))} \leq -\frac{(1 - \epsilon/2)\theta_{\dagger}^2 mn}{2(1 + \epsilon/8)}.$$

Here we used the fact that the second term in the denominator in the exponent component is $o_P(1)$ from (2.7). Combined with (6.14), we have the upper bound of $\log(\mathfrak{P}(\mathbf{X}))$ as

$$\begin{aligned} \log(\mathfrak{P}(\mathbf{X})) &\leq \log(MN|\mathbb{S}_{mn}|) - \frac{(1 - \epsilon/2)\theta_{\dagger}^2 mn}{2(1 + \epsilon/8)} \\ &\leq \log(MN|\mathbb{S}_{mn}|) - \left(1 + \frac{\epsilon}{4 + \epsilon/2}\right)(m \log \frac{M}{m} + n \log \frac{N}{n}). \end{aligned}$$

By the fact that $MN|\mathbb{S}_{mn}(X)| = (1 + o(1))\binom{N}{n}\binom{M}{m}$ and $\log(\binom{N}{n}) = (1 + o(1))n \log(N/n)$, we have eventually

$$\log(MN|\mathbb{S}_{mn}|) \leq \left(1 + \frac{\epsilon}{8 + \epsilon}\right)(m \log \frac{M}{m} + n \log \frac{N}{n}). \quad (6.15)$$

So the log Bonferroni-corrected empirical p-value is bounded from above by

$$\log(\mathfrak{P}(\mathbf{X})) \leq -\left(\frac{\epsilon}{8+\epsilon}\right)\left(m \log \frac{M}{m} + n \log \frac{N}{n}\right) \quad (6.16)$$

which finishes the proof.

6.1.2. Unidimensional permutation. We refer to the proof of Theorem 2 in [Arias-Castro and Liu \(2017\)](#). Under the unidimensional permutation, from (17) and its following arguments of [Arias-Castro and Liu \(2017\)](#), for the permutation p -value with (m, n) known, we have

$$\log \mathfrak{P}_{m,n}(\mathbf{X}) \leq -(1 + o_P(1))\delta \log(|\mathbb{S}_{mn}|). \quad (6.17)$$

Here δ is a positive constant that is less than the following quantity

$$2 \left[\liminf \frac{\theta_{\ddagger} \sqrt{mn}}{2(m \log(M/m) + n \log(N/n))} - 1 \right] \quad (6.18)$$

Note that $\log(MN) = o(1) \log(|\mathbb{S}_{mn}|)$ by (2.7), therefore directly,

$$\begin{aligned} \log \mathfrak{P}(\mathbf{X}) &\leq \log(MN) + \log \mathfrak{P}_{m,n}(\mathbf{X}) \\ &\leq -(1 + o_P(1))\delta \log |\mathbb{S}_{mn}| + \log(MN) = -(1 + o_P(1))\delta \log(|\mathbb{S}_{mn}|). \end{aligned} \quad (6.19)$$

6.2. Proof of Lemma 4.3. From Definition 4.2,

$$c' - c = \sum_{i=0}^{\lfloor \log_2 c \rfloor - k} a_i(c) 2^i \leq \sum_{i=0}^{\lfloor \log_2 c \rfloor - k} 2^i \leq 2^{\lfloor \log_2 c \rfloor - k + 1}. \quad (6.20)$$

Observing that $c = 2^{\log_2 c}$, we have

$$\frac{c - c'}{c} \leq 2^i \leq 2^{\lfloor \log_2 c \rfloor - k + 1 - \log_2 c} \leq 2^{1-k}. \quad (6.21)$$

6.3. Proof of Theorem 4.5. We illustrate the bidimensional case here, since the unidimensional case is following the same proof strategy and basically a rework of the existing proof in Section 6.1.

From Lemma 4.3, we may find $(m', n') \in S_{k_M} \times S_{k_N}$ such that $m' \leq m, n' \leq n$ and $m' = (1 + o(1))m, n' = (1 + o(1))n$. Now we consider performing permutation test on (m', n') and bound $\mathfrak{P}_{m',n'}$ by the same way in Section 6.1.

We rewrite (6.2) as follows,

$$|S_{k_M}| |S_{k_N}| |\mathbb{S}_{m'n'}| \exp \left(- \frac{m'n'(\mathbf{scan}_{m',n'}(\mathbf{X})/m'n' - \bar{\mathbf{X}})^2}{2\sigma_{\mathbf{X}}^2 + \frac{2}{3}(\mathbf{X}_{\max} - \bar{\mathbf{X}})(\mathbf{scan}_{m',n'}(\mathbf{X})/m'n' - \bar{\mathbf{X}})} \right). \quad (6.22)$$

Note that $|S_{k_M}| |S_{k_N}| \leq MN$. Combined with $m' = (1 + o(1))m, n' = (1 + o(1))n$, all we need to verify is a new version of (6.12), which in details, we just need to show that

$$\frac{\mathbf{scan}_{m',n'}(\mathbf{X})}{m'n'} - \bar{\mathbf{X}} \geq \theta_{\ddagger}(1 + o_P(1)). \quad (6.23)$$

This is done by realizing that $\mathbf{scan}_{m',n'}(\mathbf{X}) \geq Y_{S'}(\mathbf{X})$, where $S' \subset S^*$ and have m' rows and n' columns. Since $m', n' \rightarrow \infty$, the same argument yields

$$\frac{\mathbf{scan}_{m',n'}(\mathbf{X})}{m'n'} \geq \theta_{\ddagger} + O_P\left(\frac{1}{\sqrt{m'n'}}\right) \geq \theta_{\ddagger} + O_P\left(\frac{1}{\sqrt{mn}}\right). \quad (6.24)$$

Everything else follows the previous argument, and (6.23) will be verified, which concludes the proof.

Acknowledgements. The authors thank the AE and the reviewer for the comments and suggestions, which lead to great improvement of the paper. The authors are grateful to Ery Arias-Castro for introducing this topic to our attention, and helpful discussions.

References

- Abbe, E., Bandeira, A. S., and Hall, G. Exact recovery in the stochastic block model. *IEEE Trans. Inform. Theory*, **62** (1), 471–487 (2016). [MR3447993](#).
- Arias-Castro, E., Castro, R. M., Tanczos, E., and Wang, M. Distribution-free detection of structured anomalies: permutation and rank-based scans. *J. Amer. Statist. Assoc.*, **113** (522), 789–801 (2018). [MR3832227](#).
- Arias-Castro, E., Donoho, D. L., and Huo, X. Near-optimal detection of geometric objects by fast multiscale methods. *IEEE Trans. Inform. Theory*, **51** (7), 2402–2425 (2005). [MR2246369](#).
- Arias-Castro, E. and Liu, Y. Distribution-free detection of a submatrix. *J. Multivariate Anal.*, **156**, 29–38 (2017). [MR3624683](#).
- Arias-Castro, E. and Verzelen, N. Community detection in dense random networks. *Ann. Statist.*, **42** (3), 940–969 (2014). [MR3210992](#).
- Balakrishnan, S., Kolar, M., Rinaldo, A., Singh, A., and Wasserman, L. Statistical and computational tradeoffs in biclustering. In *NIPS 2011 workshop on computational trade-offs in statistical learning*, volume 4 (2011).
- Bardenet, R. and Maillard, O.-A. Concentration inequalities for sampling without replacement. *Bernoulli*, **21** (3), 1361–1385 (2015). [MR3352047](#).
- Boucheron, S., Lugosi, G., and Massart, P. *Concentration inequalities. A nonasymptotic theory of independence*. Oxford University Press, Oxford (2013). ISBN 978-0-19-953525-5. [MR3185193](#).
- Butucea, C. and Ingster, Y. I. Detection of a sparse submatrix of a high-dimensional noisy matrix. *Bernoulli*, **19** (5B), 2652–2688 (2013). [MR3160567](#).
- Cai, T. T., Liang, T., and Rakhlin, A. Computational and statistical boundaries for submatrix localization in a large noisy matrix. *Ann. Statist.*, **45** (4), 1403–1430 (2017). [MR3670183](#).
- Charrad, M. and Ben Ahmed, M. Simultaneous Clustering: A Survey. In Kuznetsov, S. O., Mandal, D. P., Kundu, M. K., and Pal, S. K., editors, *Pattern Recognition and Machine Intelligence*, pp. 370–375. Springer Berlin Heidelberg, Berlin, Heidelberg (2011). DOI: [10.1007/978-3-642-21786-9_60](#).
- Chaudhuri, K., Chung, F., and Tsiatas, A. Spectral clustering of graphs with general degrees in the extended planted partition model. In *Conference on Learning Theory*, pp. 1–23 (2012).
- Chen, Y. and Xu, J. Statistical-computational tradeoffs in planted problems and submatrix localization with a growing number of clusters and submatrices. *J. Mach. Learn. Res.*, **17**, Paper No. 27, 57 (2016). [MR3491121](#).
- Cheng, Y. and Church, G. M. Biclustering of Expression Data. In *Proceedings of the Eighth International Conference on Intelligent Systems for Molecular Biology*, pp. 93–103. AAAI Press (2000).

- Chi, E. C., Allen, G. I., and Baraniuk, R. G. Convex biclustering. *Biometrics*, **73** (1), 10–19 (2017). [MR3632347](#).
- Hajek, B., Wu, Y., and Xu, J. Information limits for recovering a hidden community. *IEEE Trans. Inform. Theory*, **63** (8), 4729–4745 (2017a). [MR3683533](#).
- Hajek, B., Wu, Y., and Xu, J. Submatrix localization via message passing. *J. Mach. Learn. Res.*, **18**, Paper No. 186, 52 (2017b). [MR3827074](#).
- Holland, P. W., Laskey, K. B., and Leinhardt, S. Stochastic blockmodels: first steps. *Social Networks*, **5** (2), 109–137 (1983). [MR718088](#).
- Kolar, M., Balakrishnan, S., Rinaldo, A., and Singh, A. Minimax localization of structural information in large noisy matrices. In *Advances in Neural Information Processing Systems*, pp. 909–917 (2011).
- Kriegel, H.-P., Kröger, P., and Zimek, A. Clustering high-dimensional data: A survey on subspace clustering, pattern-based clustering, and correlation clustering. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, **3** (1), 1 (2009).
- Ma, Z. and Wu, Y. Computational barriers in minimax submatrix detection. *Ann. Statist.*, **43** (3), 1089–1116 (2015). [MR3346698](#).
- McSherry, F. Spectral partitioning of random graphs. In *42nd IEEE Symposium on Foundations of Computer Science (Las Vegas, NV, 2001)*, pp. 529–537. IEEE Computer Soc., Los Alamitos, CA (2001). [MR1948742](#).
- Mossel, E., Neeman, J., and Sly, A. Consistency thresholds for the planted bisection model [extended abstract]. In *STOC'15—Proceedings of the 2015 ACM Symposium on Theory of Computing*, pp. 69–75. ACM, New York (2015). [MR3388184](#).
- Serfling, R. J. Probability inequalities for the sum in sampling without replacement. *Ann. Statist.*, **2**, 39–48 (1974). [MR0420967](#).
- Shabalin, A. A., Weigman, V. J., Perou, C. M., and Nobel, A. B. Finding large average submatrices in high dimensional data. *Ann. Appl. Stat.*, **3** (3), 985–1012 (2009). [MR2750383](#).
- Tanay, A., Sharan, R., and Shamir, R. Biclustering algorithms: A survey. *Handbook of computational molecular biology*, **9** (1-20), 122–124 (2005).
- Zhang, A. Y. and Zhou, H. H. Minimax rates of community detection in stochastic block models. *Ann. Statist.*, **44** (5), 2252–2280 (2016). [MR3546450](#).